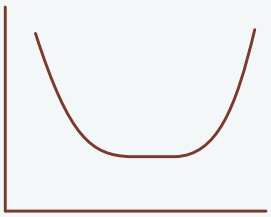
PROBABILISTIC GRAPH MODELING
FOR ENGINEERING DISCIPLINES · VOLUME I



Reliability Engineering by Derivation

From engineering questions to probabilistic inference and decisions

672 pages

Ricardo Marquez, Ph.D.

ricardo-marquez.com

Reliability Engineering by Derivation

Volume I

From engineering questions to probabilistic inference and decisions

A lesson sequence, written up as a book

Ricardo Marquez

October 31, 2026 (work in progress: one chapter per lesson, added as the lessons proceed;
Volume II holds Parts V–VII, Volume III Parts VIII–IX)

How to read this book

This book teaches one skill: to take a reliability question stated in ordinary engineering language, write it as a probability or an expectation over the unknowns of a physical model, derive what that expression requires, compute it, check it, and act on it. It is a companion to the two-volume *Reliability Engineering*, which states the discipline's concepts and vocabulary; this book goes deeper into the formulation and derivation of the probabilistic questions that discipline raises, and it does not repeat that book's introductory material on failure modes, maintenance organization or standards.

Every chapter follows the same progression, and the progression is the method.

Step 0: the plant graph. The physical system as it is, before any question is asked: its nodes, paths and sources; which quantities are pinned and which are free; the physical relations that must hold among them. The free quantities form the state vector $\mathbf{z}$. The graph is built once and serves every question put to it.

Step 1: the question. A probability or an expectation of a function of the state, with the unknown on the left of the conditioning bar and everything treated as known on the right.

Step 2: the graph extended for the question. Scenario nodes, a second drawing when two structures compete, and the evidence attached as rows: readings, counts, handbook values as priors, each with the modeling assumption it rests on, what that assumption buys and how it can fail.

Step 3: the posterior of the state. The rows combined into one objective; its minimum, the MAP state $\hat{\mathbf{z}}$; its curvature, the information $\mathbf{\Lambda} = \mathbf{J}^{\mathsf{T}}\mathbf{J}$; and the covariance $\mathbf{\Sigma} = \mathbf{\Lambda}^{-1}$, derived and then computed for the plant, with the exact posterior beside them where the model admits one. It is built once and serves every question.

Step 4: the answers. Each question read from that posterior, from the engineering question through its general and specialized forms down to its number: an entry of the state or the gradient of a derived quantity, the quantity at a time asked included; when drawings compete, a ratio of evidences; when actions compete, an expected consequence.

Step 5: evaluation and checks. The closed forms evaluated, and their approximations tested against the full calculation.

Step 6: the report. The probability, the Bayes factor or the expected consequence with its band, or a crossing time with its band by inversion.

Step 7: the next measurement. If the band is too wide, the candidate measurements ranked by the width each would remove, and a return to step 3 with the new row.

Five kinds of question cover everything the book asks, and they are distinguished by the mathematical object that answers them rather than by the domain word in the question: *estimation* (a marginal of the posterior: a mean and a spread), *prediction under a threshold* (a tail probability of that marginal, or of the trend when the quantity moves in time), *hypothesis comparison* (a ratio of evidences between two complete models), *diagnosis* (hypothesis comparison

run over a family of candidate faults generated from the graph, returned as a ranking), and *design* (a preposterior: the expected posterior after a measurement not yet taken; or an expected consequence minimized over the actions). Five computations sit underneath: the posterior of the state; the trend, the posterior over a state indexed by time, read as the sign of a trend, a crossing time or a forecast; the evidence of a model and its ratios; the preposterior of a measurement not yet taken; and the expected consequence of an action, minimized over the actions. The notation chapter that follows fixes the symbols, and the appendix collects the questions in an atlas that maps each to its mathematical form and its derivation.

The register is that of a textbook written for a reader who builds or uses physical models of engineered systems and wants every formula derived rather than quoted. Every new quantity is introduced with its name, meaning, units and equation; every derivation proceeds one relation per line from the probability rules; every numerical example is worked with units. Where the field uses a shorthand term, a *Jargon* box expands it once. Three systems recur with growing complexity: an equipment hall with one heat-removal path, a set of valves and pumps whose test and failure records are the evidence, and a redundant cooling plant that later grows into a data center.

The framework the book teaches is physics-constrained probabilistic modeling: a plant written as a graph, its unknowns as a state vector, every piece of evidence as a row, and every question as a probability over that state. TERRA is the author's own implementation of this framework; the book names the framework and the computations, not the software, and every number in it is reproducible from the derivations as written.

Each chapter closes with a statement of what to retain, a table of its terms with their meaning, units and the section that defines each, and exercises at three levels, formulation, derivation and application, each followed by its answer.

Contents

How to read this book **ii**

Notation and vocabulary **viii**

I Thinking in Probabilities **1**

1 Engineering systems as graphs **2**
1.1 One system and several questions . . . 2
1.2 The posterior of the state (step 3) . . . 15
1.3 The answers (step 4) . . . 21
1.4 The checks (step 5) . . . 26
1.5 The report (step 6) . . . 28
1.6 The next measurement (step 7) . . . 29
1.7 What to retain . . . 32
1.8 Exercises . . . 34

2 The language of probabilistic questions **37**
2.1 Engineering questions and their mathematical form . . . 37
2.2 The grammar of probability statements . . . 40
2.3 The general forms and the five computations . . . 43
2.4 Estimation, prediction and joint events . . . 51
2.5 Model comparison, diagnosis and decisions . . . 62
2.6 Measurements, classification and workflow . . . 66
2.7 What to retain . . . 72
2.8 Exercises . . . 74

3 Uncertainty and conditional probability **76**
3.1 Uncertainty and evidence in reliability questions . . . 76
3.2 Conditional probability and Bayes' theorem . . . 80
3.3 Bayesian diagnosis of a cooling fault . . . 82
3.4 Marginalization and dependence . . . 85
3.5 Conditioning in a physical model . . . 89
3.6 Predictive variance and the posterior covariance . . . 97
3.7 What to retain . . . 103
3.8 Exercises . . . 105

4 The anatomy of probabilistic inference **107**
4.1 From a plant graph to an engineering probability . . . 107
4.2 The posterior of the state (step 3) . . . 115
4.3 The answers (step 4) . . . 124

4.4 The checks (step 5) . . . 130
4.5 The report (step 6) . . . 135
4.6 The next measurement (step 7) . . . 136
4.7 What to retain . . . 140
4.8 Exercises . . . 144

II Estimating Reliability from Data 146

5 Estimating failure probability from data 147
5.1 Failure probability from a record of tests . . . 147
5.2 The posterior of the state (step 3) . . . 155
5.3 The answers (step 4) . . . 159
5.4 The checks (step 5) . . . 173
5.5 The report (step 6) . . . 176
5.6 The next measurement (step 7) . . . 178
5.7 What to retain . . . 181
5.8 Exercises . . . 183

6 Estimating failure rates from records 185
6.1 Failure rates from operating records . . . 185
6.2 The posterior of the state (step 3) . . . 197
6.3 The answers (step 4) . . . 202
6.4 The checks (step 5) . . . 221
6.5 The report (step 6) . . . 226
6.6 The next measurement (step 7) . . . 228
6.7 What to retain . . . 231
6.8 Exercises . . . 234

7 Failure-time distributions and survival analysis 236
7.1 Lifetimes from a record with survivors . . . 236
7.2 The posterior of the state (step 3) . . . 250
7.3 The answers (step 4) . . . 255
7.4 The checks (step 5) . . . 264
7.5 The report (step 6) . . . 272
7.6 The next measurement (step 7) . . . 274
7.7 What to retain . . . 277
7.8 Exercises . . . 279

8 Repair rates and component availability 281
8.1 Repair rates and the availability of a repairable component . . . 281
8.2 The posterior of the state (step 3) . . . 289
8.3 The answers (step 4) . . . 293
8.4 The checks (step 5) . . . 305
8.5 The report (step 6) . . . 310
8.6 The next measurement (step 7) . . . 312
8.7 What to retain . . . 316
8.8 Exercises . . . 317

III Predicting Future Reliability 320

9 Mission reliability 321
9.1 Mission reliability under an uncertain failure rate . . . 321
9.2 The posterior of the state (step 3) . . . 327
9.3 The answers (step 4) . . . 331
9.4 The checks (step 5) . . . 344
9.5 The report (step 6) . . . 348
9.6 The next measurement (step 7) . . . 349
9.7 What to retain . . . 354
9.8 Exercises . . . 355

10 Remaining life and conditional survival 358
10.1 Remaining life of units that have already survived . . . 358
10.2 The posterior of the state (step 3) . . . 365
10.3 The answers (step 4) . . . 368
10.4 The checks (step 5) . . . 385
10.5 The report (step 6) . . . 389
10.6 The next measurement (step 7) . . . 391
10.7 What to retain . . . 395
10.8 Exercises . . . 398

11 Predicting future failure counts 400
11.1 Failure counts over a future period . . . 400
11.2 The posterior of the state (step 3) . . . 406
11.3 The answers (step 4) . . . 409
11.4 The checks (step 5) . . . 425
11.5 The report (step 6) . . . 428
11.6 The next measurement (step 7) . . . 430
11.7 What to retain . . . 435
11.8 Exercises . . . 437

12 Aging and time-varying failure rates 439
12.1 A failure rate that changes with time . . . 439
12.2 The posterior of the state (step 3) . . . 446
12.3 The answers (step 4) . . . 450
12.4 The checks (step 5) . . . 465
12.5 The report (step 6) . . . 470
12.6 The next measurement (step 7) . . . 472
12.7 What to retain . . . 476
12.8 Exercises . . . 477

IV System Reliability and Dependencies 480

13 Series, parallel and k-out-of-n systems 481
13.1 Services provided by several units . . . 481
13.2 The posterior of the state (step 3) . . . 488
13.3 The answers (step 4) . . . 490
13.4 The checks (step 5) . . . 512
13.5 The report (step 6) . . . 516
13.6 The next measurement (step 7) . . . 518

13.7 What to retain . . . 522
13.8 Exercises . . . 524

14 Fault trees and minimal cut sets **527**
14.1 The combinations that take a service down . . . 527
14.2 The posterior of the state (step 3) . . . 533
14.3 The answers (step 4) . . . 535
14.4 The checks (step 5) . . . 553
14.5 The report (step 6) . . . 557
14.6 The next measurement (step 7) . . . 559
14.7 What to retain . . . 563
14.8 Exercises . . . 564

15 Dependent failures and common causes **567**
15.1 Dependent failures and the common-cause share . . . 567
15.2 The posterior of the state (step 3) . . . 576
15.3 The answers (step 4) . . . 581
15.4 The checks (step 5) . . . 593
15.5 The report (step 6) . . . 600
15.6 The next measurement (step 7) . . . 602
15.7 What to retain . . . 607
15.8 Exercises . . . 609

16 Repairable and multistate systems **612**
16.1 Repair capacity and the states of a redundant pair . . . 612
16.2 The posterior of the state (step 3) . . . 619
16.3 The answers (step 4) . . . 621
16.4 The checks (step 5) . . . 640
16.5 The report (step 6) . . . 646
16.6 The next measurement (step 7) . . . 648
16.7 What to retain . . . 651
16.8 Exercises . . . 652

A Atlas of probabilistic questions **655**

Notation and vocabulary

Every symbol below is used with the same meaning in every chapter. The unknown always stands on the left of a conditioning bar and the evidence on the right. Vectors and matrices are bold.

Symbol	Unit	Meaning
The plant graph and its state		
$\mathbf{z}$, z_i	varies	State vector of the plant graph: every free quantity at once (temperatures, flows, conductances, log-rates); its i-th entry.
$\mathbf{y}$, y	varies	The evidence: all readings, counts and prior information treated as known; a single reading.
$r_k(\mathbf{z})$	varies	Row k of the model: a physical relation that must hold ($r_k = 0$), a reading compared with the state ($r_k = z_i - y$, in the units of the reading), or a prior ($r_k = z_i - z_i^0$).
w_k, σ_k	varies	Weight (precision) of row k and the standard deviation it corresponds to, $w_k = 1/\sigma_k^2$. A row may equivalently be written normalized, r_k/σ_k, with unit weight; both put the same term $\frac{1}{2}w_k r_k^2$ in the objective.
T, T_a	°C	Temperature of the equipment hall; ambient temperature (pinned).
Q, G	kW, kW/K	Heat load supplied to the hall; effective thermal conductance of the heat-removal path.
g, G_{ref}	–, kW/K	Log-conductance $\ln(G/G_{\text{ref}})$, the entry of the state that carries the conductance; its reference value.
y_T, σ_T	°C, K	A temperature reading of the hall and the standard deviation of its error.
$T_{\max}$, f^*	varies	A limit on a temperature; a requirement level on a state entry or on a derived quantity $f(\mathbf{z})$: the event $f \geq f^*$ is *the requirement holds*, and an upper limit is the complementary event (Chapter 2). From Part V a starred plant quantity (Q_1^*, T_1^*) is a scenario value: the quantity at a load not yet applied.
ω, ξ	–	Log-rate $\ln(\lambda/\lambda_{\text{ref}})$ and logit of a probability, $\ln(p/(1-p))$: the unbounded entries that carry a rate or a probability in the state.
u, v	–	Log-scale $\ln\eta$ and log-shape $\ln\beta$ of a Weibull lifetime, the two entries of the state in a lifetime fit.
Probability		
$\Pr(A)$, $\Pr(A \mid B)$	–	Probability of event A; conditional on event B.

Symbol	Unit	Meaning
$p(\cdot)$	varies	A probability density; $p(\mathbf{z} \mid \mathbf{y})$ is the posterior of the state, $p(\mathbf{y} \mid \mathbf{z})$ the likelihood, $p(\mathbf{z})$ the prior.
$\mathbb{E}$, Var, Cov	varies	Expectation, variance, covariance.
$\mathcal{N}(\mu, \sigma^2)$	–	Normal distribution with mean μ and variance σ^2; Φ its standard cumulative distribution function.
Beta(a, b), Gamma(a, b)	–	Beta distribution with shape parameters a, b; gamma distribution with shape a and rate b.
Binomial(n, p), Poisson(μ)	–	Number of failures in n trials with failure probability p; count with mean μ.
Inference objects		
$\mathcal{L}(\mathbf{z})$, $\mathcal{L}_k(\mathbf{z})$	–	The objective, minus the log posterior up to a constant, and the term row k puts in it: $\mathcal{L} = \sum_k \mathcal{L}_k$, with $\mathcal{L}_k = \frac{1}{2} w_k r_k(\mathbf{z})^2$ for a Gaussian row and $\mathcal{L}_k = -\ln p(y_k \mid \mathbf{z})$ for a count, lifetime or demand row (Section 4.2.7).
$\mathcal{L}'$, $\mathcal{L}''$, $\mathcal{L}'_k$, $\mathcal{L}''_k$	varies	For a state of one entry, say ω: the slope and the curvature of the objective, $\mathrm{d}\mathcal{L}/\mathrm{d}\omega$ and $\mathrm{d}^2\mathcal{L}/\mathrm{d}\omega^2$, and of row k's term; they add over the rows, and the Newton step is $\Delta\omega = -\mathcal{L}'/\mathcal{L}''$ (Chapter 5).
Working form	varies	The way a count row joins the Jacobian beside Gaussian rows (equation (4.43)): the line $\sqrt{\mu}\,\partial \ln \mu / \partial \mathbf{z}$ with the working residual $(\mu - k)/\sqrt{\mu}$ for a Poisson count k of mean μ, and $\sqrt{np(1-p)}\,\partial \xi / \partial \mathbf{z}$ with $(np - k)/\sqrt{np(1-p)}$ for a binomial count carried on the logit ξ. The pair reproduces the row's gradient and the $\mathbf{J}^\mathsf{T}\mathbf{J}$ part of its Hessian; its half square is not the row's term.
$\hat{\mathbf{z}}$	varies	The maximum a posteriori (MAP) state, the minimizer of $\mathcal{L}$.
$\mathbf{J}$	varies	Jacobian of the rows with respect to the state, evaluated at $\hat{\mathbf{z}}$: for a Gaussian row the line $\partial r_k / \partial z_i$, for a count row the line of its working form (Section 4.2.7).
$\mathbf{\Lambda}$, $\mathbf{H}$, $\mathbf{C}$	varies	Information matrix $\mathbf{J}^\mathsf{T}\mathbf{W}\mathbf{J}$ with $\mathbf{W} = \mathrm{diag}(w_k)$ ($\mathbf{J}^\mathsf{T}\mathbf{J}$ for normalized rows); the Hessian of $\mathcal{L}$, $\mathbf{H} = \mathbf{J}^\mathsf{T}\mathbf{J} + \mathbf{C}$, with $\mathbf{C}$ the curvature term the rows add beyond the information, zero for linear Gaussian rows and for count rows on their canonical parameter (equation (4.44)).
Λ	varies	For a state of one entry, the information, the curvature of the objective at the MAP, $\Lambda = \mathcal{L}''(\hat{\omega})$ (Chapter 5).
$\mathbf{\Sigma}$, Σ_{ij}	varies	Posterior covariance of the state under the Laplace approximation, $\mathbf{H}^{-1}$ ($\mathbf{\Lambda}^{-1}$ when $\mathbf{C} = \mathbf{0}$); its entries. The posterior spread of entry i is written $\sqrt{\Sigma_{ii}}$; σ_k always names a row's standard deviation and σ_p a physics row's tolerance.
σ^2_ω, σ_ω	varies	For a state of one entry, say ω or ξ: its Laplace variance $1/\Lambda$ and standard deviation (σ^2_ξ, σ_ξ for the logit). A derived quantity with derivative $g = \mathrm{d}f/\mathrm{d}\omega$ at the MAP has variance $g^2\sigma^2_\omega$.

Symbol	Unit	Meaning
ρ	K/kW	Thermal resistance of a path, $1/G$ (Chapter 4); not a correlation.
f_L, f_U	varies	The ends of the credible interval of a derived quantity $f(\mathbf{z})$ (Chapter 2).
y_{new}, $\tilde{y}_a$	varies	A future observation (the predictive) and the outcome of a candidate measurement a before it is taken (the preposterior).
M, M_j, B_{AB}	–	A drawing (model) of the plant; the members of a family; the Bayes factor $p(\mathbf{y} \mid M_A)/p(\mathbf{y} \mid M_B)$, written B_{10} for M_1 against M_0.
$f_t(\mathbf{z})$, $\mathbf{g}(t)$, $\bar{t}$	varies	A quantity read at time t from a state indexed by time, $f_t = \mathbf{g}(t)^{\mathsf{T}}\mathbf{z}$ with the gradient $\mathbf{g}(t) = (1, t - \bar{t})^{\mathsf{T}}$ about the center $\bar{t}$ of the record, the state holding a level entry α and a trend entry γ (the trend form, Chapter 2).
$U(a)$, u	varies	The expected value of a candidate action a over the outcomes it may produce, and the utility u of the updated posterior it is built from (the design form, Chapter 2).
$\mathbf{J}_a$, $\mathbf{\Lambda}_a$	varies	The Jacobian rows of a candidate measurement and its information $\mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a$ (expected counts at $\hat{\mathbf{z}}$ for a count row).
κ	–	Assurance level of an acceptance rule: the posterior probability a requirement must reach, $\Pr(R \geq R_{\min} \mid \mathbf{y}) \geq \kappa$ (Chapter 5).
π_0	–	The prior row's share of the information, $\mathcal{L}_0''/\mathcal{L}''$ for a state of one entry (Chapter 5); in Chapter 16 the stationary probability that no unit is up.
$f(\mathbf{z})$, $\mathbf{g}$	varies	A derived quantity of the state and its gradient $\partial f/\partial \mathbf{z}$ at $\hat{\mathbf{z}}$.
μ_f, σ_f^2	varies	Mean and variance of a derived quantity under the delta method: $f(\hat{\mathbf{z}})$ and $\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}$. From Part V the spread is written $\text{sd}(f \mid \mathbf{y}) = \sqrt{\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}}$.
$p(\mathbf{y} \mid M)$	–	Evidence (marginal likelihood) of model M.
B_{10}	–	Bayes factor of model M_1 against M_0, $p(\mathbf{y} \mid M_1)/p(\mathbf{y} \mid M_0)$; usually reported as a log in nats.
$\mathbf{\Upsilon}$, υ	varies	The gain of the correction form, $\mathbf{\Sigma}_0\mathbf{A}^{\mathsf{T}}(\mathbf{A}\mathbf{\Sigma}_0\mathbf{A}^{\mathsf{T}} + \mathbf{R})^{-1}$, which moves the prior mean toward the readings; its scalar value for one reading, the weight $\sigma_0^2/(\sigma_0^2 + \sigma_T^2)$ (Chapter 3).
$\mathbf{\Sigma}_a$	varies	Preposterior covariance after a candidate measurement a with information $\mathbf{\Lambda}_a$: $(\mathbf{\Sigma}^{-1} + \mathbf{\Lambda}_a)^{-1}$.
EIG	nat or bit	Expected information gain of a candidate measurement.
Reliability quantities (as in the companion book)		
T, $R(t)$, $F(t)$	h, –, –	Time to failure; reliability $\Pr(T > t)$; cumulative failure probability.
$f(t)$, $\lambda(t)$, $\Lambda(t)$	h^{-1}, h^{-1}, –	Failure density; hazard function; cumulative hazard.
λ, μ	h^{-1}	Constant failure rate; constant repair rate.
p, k, n	–	Failure probability per demand; number of failures observed; number of trials.

Symbol	Unit	Meaning
T_{tot}	h	Exposure (total operating time on test) behind a failure count.
β, η	–, h	Weibull shape and scale (characteristic life).
t_i, δ_i, c_i	h, –, h	Observed time of unit i; its event indicator (1 failed, 0 censored); its censoring time.
t_{50}, t_0, T_{new}	h	Median lifetime; an age already reached; the lifetime of a new unit.
ℓ, ℓ_p	–	Log-likelihood; profile log-likelihood (maximized over the other entries).
r, D	–, h	Number of repairs observed; their total duration.
U, A_{req}	–	Unavailability $1 - A$; a required availability.
α, γ	–, yr^{-1}	The log-rate case of the level and trend entries of the trend form: log-rate at the record's center and its trend per year in $\ln \lambda(t) = \alpha + \gamma\,(t - \bar{t})$; t^* the time the rate crosses a limit.
K_f, τ	–, h or yr	A future failure count; the future period or horizon it is counted over.
q_i, $\mathbf{s}$, $\phi(\mathbf{s})$	–	Unavailability of unit i; a configuration of units (up or down); the structure function, 1 when the configuration delivers the service.
R_{sys}, U_{sys}	–	Service availability and unavailability of several units under a drawing, by the 2^n enumeration.
β_{cc}	–	Common-cause share: the fraction of a unit's failures that fail all units sharing the cause (β is the Weibull shape).
A, MTTF, MTTR	–, h, h	Steady-state availability; mean time to failure; mean time to repair.
$L(a, \mathbf{z})$, $\mathcal{R}(a)$, C, C_e, c_a	R$	The loss of action a when the state is $\mathbf{z}$, its consequence; its expected value over the posterior, the risk $\mathcal{R}(a) = \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]$; the consequence of an event, of event e; the fixed cost of the action (the expected-consequence form, Chapter 2).

Vocabulary

Plant graph The representation of the physical system as nodes (places where a conserved quantity balances), edges (paths along which it flows) and sources, with the relations that must hold among them. Built before any question is asked.

State vector The collection of free quantities the inference estimates at once; outputs that are entries of the state need no further transformation.

Pinned, free A quantity treated as known (boundary condition) against one carried in the state.

Residual, row A relation written so that it equals zero when satisfied; physical rows constrain the state, reading rows compare it with evidence, prior rows carry handbook knowledge.

Prior, likelihood, posterior Knowledge before the evidence; the probability of the evidence given a state; knowledge after the evidence, by Bayes' theorem.

Marginalization Integrating a joint distribution over the unknowns not asked about; the law of

total probability.

Aleatory, epistemic Variability in outcomes the model treats as random; incomplete knowledge of quantities that have one true value.

Maximum a posteriori (MAP) estimate The state at which the posterior density is highest; the minimizer of the objective.

Laplace approximation Replacing the posterior by the Gaussian whose covariance is the inverse curvature of the objective at the MAP state.

Information matrix The matrix $\mathbf{J}^\mathsf{T}\mathbf{W}\mathbf{J}$, the Hessian of the objective without the curvature term $\mathbf{C}$ of $\mathbf{H} = \mathbf{J}^\mathsf{T}\mathbf{J} + \mathbf{C}$ (equation (4.44)); precisions of independent rows add. For a state of one entry it is the number Λ.

Delta method Linearizing a derived quantity around the MAP state so that its variance is $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$.

Change of variables Carrying a density to a new coordinate $u = h(x)$:

$$p_u(u) = p_x(h^{-1}(u))\,|\mathrm{d}x/\mathrm{d}u|,$$

the Jacobian factor keeping the probability of each interval; for several entries the factor is $|\det \partial\mathbf{z}/\partial\mathbf{u}|$. It gives the shape of a derived quantity, where the delta method gives only its spread (Section 4.3.3).

Variance budget The expansion of $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$ into one term per state entry plus the cross terms from their covariance.

Threshold probability The probability that a state entry or derived quantity is below (or above) a limit; a normal tail under the Laplace approximation.

Posterior predictive The distribution of a future observation obtained by averaging its model over the posterior.

Credible interval An interval that contains the unknown with a stated posterior probability.

Evidence, Bayes factor The normalizing constant of one model's posterior; the ratio of two models' evidences, which compares the models rather than their parameters.

Preposterior The posterior that would result from a measurement not yet taken; the object that values a measurement before buying it.

Identifiability Whether the evidence determines the state entries separately or only in combination; read from the shape of the posterior, not from marginal spreads.

The five question categories Estimation, prediction under a threshold, hypothesis comparison, diagnosis and design, answered by a marginal, a tail of the posterior or of the trend, an evidence ratio, a ranking of evidences, and a preposterior or an expected consequence.

The five computations The objects that answer every category (Chapter 2): the *posterior* $p(\mathbf{z} \mid \mathbf{y})$, read for an entry or a derived quantity; the *trend*, the posterior over a state indexed by time or window with a trend entry, read as the probability that the trend is positive, as a crossing time with its band, or as a forecast; the *evidence* $p(\mathbf{y} \mid M)$ and its ratios; the *preposterior*, the expected posterior under a measurement not yet taken; and the *expected consequence* $\mathcal{R}(a) = \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]$ of an action over the posterior, with the choice of the action that minimizes it (risk is probability times consequence).

Factor graph A graph of the state's entries and the factors (the rows) that bear on them; the plant graph with its rows is one, and the posterior is the normalized product of its factors (Chapters 1 and 4).

Variable elimination Integrating the entries a question does not name out of the joint distribution one at a time; for a Gaussian posterior it is the Schur complement, and its cost on a large sparse graph is set by the order of elimination (Chapter 3).

Belief propagation Computing the marginals of a graphical model by passing messages between neighboring factors and entries (the sum-product algorithm); exact on a graph without loops, where on a Gaussian graph it returns the Schur complements. Not used in this book, whose posteriors are factored jointly.

Kalman filter and smoother Variable elimination on a state indexed by time, taken in time order: the filter forward, the smoother back.

Importance sampling (reweighted draws) Drawing from an approximation, here the Laplace normal, and weighting each draw by the ratio of the exact posterior to it, so that the draws check the approximation (Chapter 20, in Volume II).

Markov chain Monte Carlo Drawing from a posterior by a Markov chain whose stationary distribution it is (Gibbs sampling, the Metropolis algorithm); the route when neither the Laplace normal nor reweighting is adequate, as for a hierarchical model (Chapter 33, in Volume III).

Explaining away Two causes independent before an effect is seen become dependent once it is seen: evidence for one lowers the probability of the other, as a biased sensor and a changed conductance compete for one disagreeing reading (Chapters 22 and 23, in Volume II).

Part I

Thinking in Probabilities

Chapter 1

Engineering Systems as Graphs

1.1 One system and several questions

The owner of a data center equipment hall wants to double its electrical load, from 30 to 60 kW, and the hall must then stay below 27.5°C. The record is thin: one temperature sensor, with an error of about half a kelvin, reads 23.2°C at the present load, and an engineering handbook gives the conductance of the hall's cooling path to the 18°C ambient as 7.5 kW/K, right to within a factor of about 1.35. The handbook value gives the first answer. Heat leaves the hall in proportion to its excess over the ambient, so at 7.5 kW/K the hall sits at 22°C under 30 kW, comfortably below its present limit of 24°C, and at 26°C under 60 kW, well below 27.5°C: the full load passes with room to spare.

The sensor tells another story. Taken as the truth, 23.2°C at 30 kW means a conductance of 5.77 kW/K, and at that conductance the full load puts the hall at 28.4°C, over the limit. One source says the hall passes and the other that it fails; each treats its own number as exact, neither is weighed against the other, and neither says how sure the owner can be. The decision needs a probability, that the hall stays below 27.5°C at a load at which it has never run, and that probability can only come from holding the handbook value, the reading and the energy balance that ties them together in one description of the hall, each with its own uncertainty. The question of this chapter is therefore

how a physical engineering system is represented so that the same representation serves every probabilistic question asked of it, and what must be fixed before any inference is performed.

Jargon: plant graph, node, path, source, pinned and free quantities, state vector, residual, row, observation, prior, posterior, model discrepancy, scenario node, Jacobian, normal equations, information matrix, covariance, Laplace approximation, factor graph

A *plant graph* represents a physical system as *nodes*, places where a conserved quantity such as energy or mass must balance; *paths* (edges), along which the quantity flows; and *sources*, which supply it. A quantity is *pinned* when it is treated as known, as a boundary condition, and *free* when it is unknown and must be inferred. The free quantities, collected in one vector, form the *state vector* $\mathbf{z}$. A *residual* is a relation written so that it equals zero when it is satisfied; each residual is a *row* of the model. A physical row states a relation that must hold among entries of the state; an *observation* (a reading) is a measured value related to an entry of the state with an error, and its row compares the two. The *prior* expresses what is known about the state before the readings, for example a handbook value of a conductance; the *posterior* $p(\mathbf{z} \mid \mathbf{y})$ expresses what is known after the evidence $\mathbf{y}$. A *model discrepancy* is the difference between the physical system and the model of it that would remain if every entry of the state were known exactly. A *scenario node* is a node added to the graph for an operating condition that has not occurred, so that a question can be asked about it. The *Jacobian* $\mathbf{J}$ is the matrix of derivatives of every row with respect to every entry of the state, one row per row of the model and one column per entry. The *normal equations* $(\mathbf{J}^\mathsf{T}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^\mathsf{T}\mathbf{r}$ give the step that best reduces all the rows at once when each is replaced by its tangent. The *information matrix* $\mathbf{J}^\mathsf{T}\mathbf{J}$ states how sharply the rows together fix the state, direction by direction, and its inverse is the posterior *covariance* $\boldsymbol{\Sigma}$, whose diagonal holds the variances of the entries and whose off-diagonal entries hold how they vary together. The *Laplace approximation* replaces the posterior by the normal distribution centered on its most probable state with that covariance. A *factor graph* draws the entries of the state and the rows that act on them, one node each.

Key idea

A graph describes the engineering system; a probabilistic question states what is to be learned about it. The plant graph, its pinned and free quantities and its physical rows are built once, and each question is a probability over the state of that one graph, with the unknown on the left of the conditioning bar and the evidence on the right. The posterior of the state is carried by three objects, the most probable state $\hat{\mathbf{z}}$, the covariance $\boldsymbol{\Sigma}$ and the Jacobian $\mathbf{J}$ of the rows, and every answer is an entry of the state read off them or a function of the state carried through them.

1.1.1 Several questions about one equipment hall

An engineer responsible for the hall might ask any of the following:

1. What is the actual temperature inside the hall?
2. What is the effective conductance of the cooling path?
3. Will the temperature exceed its limit if the electrical load doubles?
4. Is the temperature sensor biased?
5. Would installing another sensor reduce the uncertainty?
6. What happens if the cooling system partially fails?

The questions are different. The first two ask for the value of an unknown quantity; the third asks for the probability that a future temperature stays within a limit; the fourth asks which of two descriptions of the sensor is true; the fifth asks what a measurement not yet taken would be worth; the sixth asks about an operating condition that has not occurred. All six, however, concern the same physical system: the same hall, the same cooling path, the same sensor.

Constructing a separate mathematical model for every question would repeat the same physics six times, and the six models could disagree about it. The method of this book begins instead with one common representation of the equipment and its physical relations, the *plant graph*. It is step 0 of the method stated in the preface: the system as it is, before any question is asked, built once and kept for every question that follows. Section 1.1.6 writes the questions that this chapter answers over the graph.

1.1.2 The system, the question and the evidence

Three things must be fixed before any probabilistic inference is performed, and the plant graph is where they are fixed. The first is the division of the quantities: which are known and pinned (the ambient temperature, the heat load), which are unknown and free (the hall temperature, the conductance), and which are observed (the sensor reading). The field often separates the unknowns into *states*, such as temperatures and flows, and *parameters*, such as conductances and rates; in this book both kinds are entries of one state vector, and the inference treats them alike. The second is the set of physical relations that must hold among the quantities, such as conservation of energy at the hall. The third is the evidence: the readings, the records and the prior knowledge, attached to the graph as rows of their own.

The question is a separate object. It is stated over the graph after the graph is built, and it decides what is computed, not what the system is. The separation makes complex reliability assessments systematic: when a new question arrives, the graph is already in hand, and the work is to write the question over its state.

> **Principle 1.1.** *The physical model is reusable; the probabilistic question determines what is calculated from it. A plant graph is built once from the equipment and its physics, and every question is a probability or an expectation over the state of that graph.*

1.1.3 Larger plant graphs and factor graphs

The equipment hall of this chapter has one thermal path and two unknown quantities, the smallest graph with every kind of element. The same construction serves a plant of many components, and the graph on which an inference operates is a second object, drawn from the plant graph and the evidence attached to it.

A cooling network. Real engineering systems may include hundreds or thousands of components, each with its own equations and uncertainties. Consider a cooling network with two chillers, two pumps, a common electrical supply and several temperature sensors. Its plant graph may contain temperatures, pressures, flows and electrical states at its nodes; physical components represented by paths or by component relations; heat, mass and power conservation equations; sensor measurements and historical observations; failure and repair states of equipment; and shared parameters and common-cause dependencies.

Figure 1.1 draws a reduced version. The green boxes are free nodes: the two chillers and the two pumps, each carrying its status and capacity or flow, the supply header with its water temperature T_s and flow, and the hall with its temperature T. The gray box on the left is the pinned ambient, to which the chillers reject heat; the white box in the middle is the common electrical supply, which feeds all four machines; the white box at the upper right is the heat load of the hall. Solid green arrows are physical paths of heat and water; thin gray arrows are electrical feeds; dashed arrows are observations, two temperature readings and the maintenance record of a chiller. The common supply is a single node on which four components depend: if it fails, all four stop together, and a graph that omitted it would treat their failures as independent.

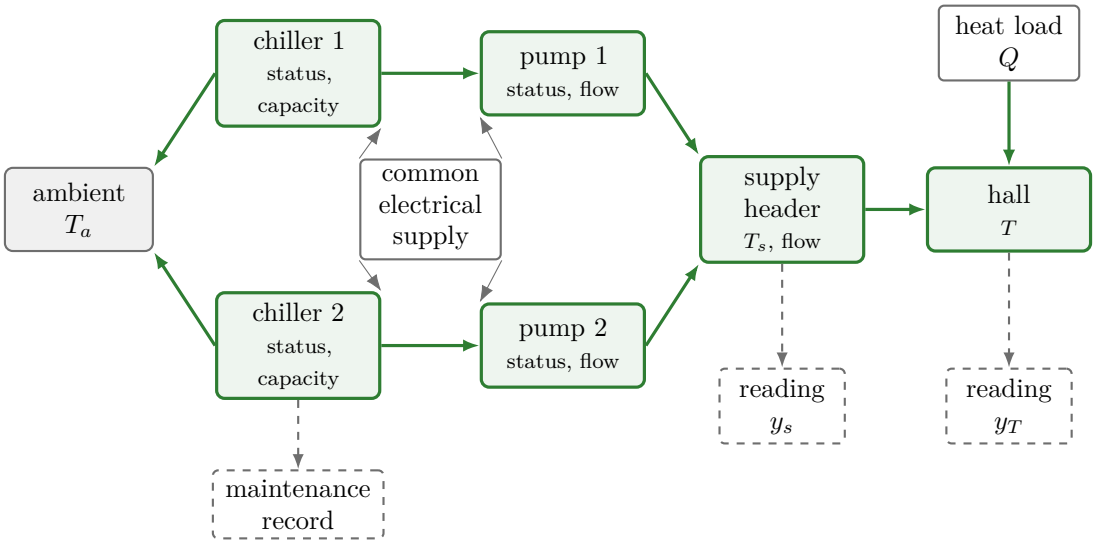


Figure 1.1: Plant graph of a reduced cooling network (illustrative; return piping omitted). Free: the status and capacity of each chiller, the status and flow of each pump, the header water temperature T_s and flow, and the hall temperature T. Pinned: the ambient temperature T_a and the heat load Q; the common electrical supply feeds all four machines and is the shared dependency. Observed: T_s and T through two readings, and chiller 2 through its maintenance record.

The mathematical principle is the one the single hall shows below. The network is a collection of physical relations,

$$r_1(\mathbf{z}) = 0, \quad r_2(\mathbf{z}) = 0, \quad \ldots, \quad r_m(\mathbf{z}) = 0,$$

one per balance and per component relation, over a state vector $\mathbf{z}$ that now has many entries. Observations and probability models for the uncertain quantities are attached as rows of their own. What is ultimately sought is the joint posterior of the state,

$$p(\mathbf{z} \mid \mathbf{y}, M),$$

where M represents the assumed system model, including its structure: which components exist, how they are connected and which dependencies are represented. Writing M explicitly matters because the structure itself may be uncertain, and a question about the structure is answered by comparing two such models, not by the posterior inside either.

Physical graphs and factor graphs. A physical graph and a probabilistic factor graph are related, but they are not the same object. The physical graph represents equipment and connections: halls, paths, chillers, pumps, sensors. A *factor graph* represents the variables of the inference and the mathematical relations among them, including the observations, the priors and the dependencies: each entry of the state is one node, each row of the model is one factor node, and each factor is linked to the entries its row involves. It is constructed from the physical graph and the evidence attached to it, and it shows which row acts on which unknown. The distinction becomes important when competing system structures and diagnosis are studied in later chapters. Two physical drawings of the same plant, one with a single cooling path and one with two, give two factor graphs with different variables and rows, and comparing them is comparing two models M. A sensor bias adds a variable to the factor graph without changing the equipment at all.

1.1.4 The plant graph (step 0)

The simplest plant graph with every kind of element has one free node, one pinned node, one path, one source and one observation, and the equipment hall is that graph. It is built before any question is asked and before any number is fitted.

Nodes, a path, a source and an observation. Consider a simplified equipment hall with one heat-removal path to the ambient air. Figure 1.2 draws it as a plant graph. The green box on the left is the hall, a free node whose temperature T is unknown; the gray box on the right is the ambient air, a pinned node whose temperature $T_a = 18°\text{C}$ is known. The solid green arrow between them is the heat-removal path, characterized by its effective thermal conductance G. The box above the hall is the heat load Q supplied by the equipment, a source. The dashed box below is the temperature sensor, and the dashed arrow from the hall to it is an observation: the sensor reports a value related to T, not T itself. Solid connections are physical relations; the dashed connection is evidence.

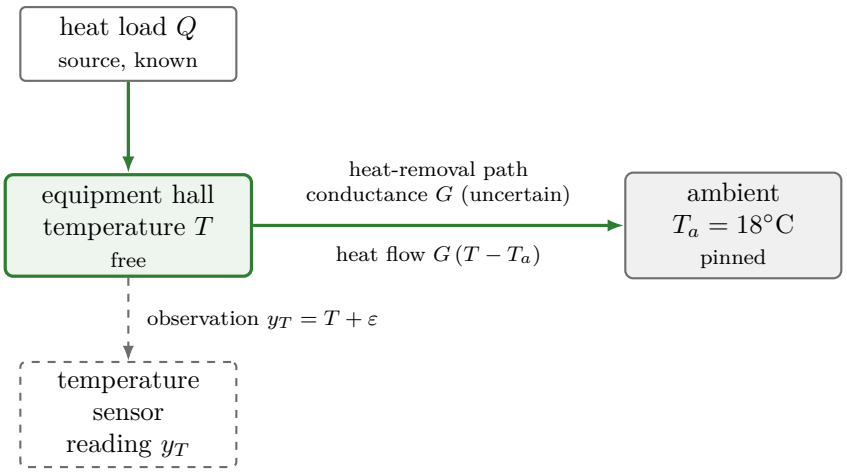


Figure 1.2: Plant graph of the equipment hall in steady state. Free: the hall temperature T and the conductance G of the heat-removal path. Pinned: the ambient temperature $T_a = 18°\text{C}$ and the heat load Q. Observed: T, through the reading y_T with error ε. Solid arrows are physical relations; the dashed arrow is an observation.

Table 1.1 lists the elements of the graph. The first column names each element, the second its role in the graph, the third its symbol and the fourth its unit; the last column states how the element enters this first example.

Table 1.1: The elements of the plant graph of Figure 1.2.

Element	Role	Symbol	Unit	In the example
Equipment hall	free node	T	°C	temperature unknown
Ambient air	pinned node	T_a	°C	known, 18°C
Heat load	source	Q	kW	thermal power supplied by the equipment, known
Thermal path	path (edge)	G	kW/K	effective conductance, uncertain
Temperature sensor	observation	y_T	°C	a reading related to the hall temperature

For now the ambient temperature and the heat load are assumed known, and the effective conductance is uncertain. The conductance lumps every mechanism by which heat leaves the

hall, the cooling equipment and the envelope alike, into one number; how well that lumping represents the hall is a modeling assumption, examined in Section 1.6.5. This is the starting physical model.

Conservation of energy at the hall node. A graph becomes useful for computation when the mathematical relations between its elements are specified. Assume steady state: the hall stores no more and no less heat from one moment to the next, so that its temperature is constant. Conservation of energy at the hall node then requires the heat entering to equal the heat leaving,

$$Q_{\text{in}} = Q_{\text{out}},$$

where Q_{in} and Q_{out} are thermal powers in kilowatts. The heat entering is the load, $Q_{\text{in}} = Q$. The heat leaving is modeled as proportional to the temperature difference across the path,

$$Q_{\text{out}} = G\,(T - T_a),$$

the defining relation of the effective conductance: G is the heat flow per kelvin of difference between the hall and the ambient. Combining the two,

$$Q = G\,(T - T_a), \tag{1.1}$$

where Q is the heat generated by the equipment in kW, G the effective thermal conductance in kW/K, T the hall temperature in °C and T_a the ambient temperature in °C. A temperature difference has the same value in kelvin and in degrees Celsius, so the units balance: kW/K times K is kW.

Equation (1.1) is a *physical relation*, not a probability statement. It states what must hold if the simplified thermal model is correct, whatever is or is not known about its quantities. Since $G > 0$, dividing by G and adding T_a gives

$$T = T_a + \frac{Q}{G}, \tag{1.2}$$

which gives the temperature when the other three quantities are specified. The hall is warmer than the ambient by Q/G: more load or less conductance raises it.

The forward calculation. The hall is operated at the load $Q = 30\,\text{kW}$, and its cooling path was designed for a conductance of $G = 6\,\text{kW/K}$. With $T_a = 18°\text{C}$, equation (1.2) gives

$$T = 18 + \frac{30\,\text{kW}}{6\,\text{kW/K}} = 18 + 5 = 23°\text{C},$$

the steady-state temperature of the hall at its design conductance. The calculation used no probability, estimated no uncertainty and used no observational evidence. It is a deterministic *forward calculation*: from pinned inputs and a specified conductance to the temperature they imply. Every later chapter needs it, because every inference compares temperatures computed this way with what the sensor reports. The temperature falls as the conductance grows, with the slope $\partial T/\partial G = -Q/G^2$, which at the design point is $-30/36 = -0.833\,\text{K}$ per kW/K: an error of 1 kW/K in the conductance moves the computed temperature by almost a kelvin. Doubling the load doubles the excess $T - T_a$ at every conductance: at $G = 6\,\text{kW/K}$ the hall at 60 kW would be at 28°C.

The state vector. The forward calculation required the conductance to be known. In the inference problems of this book it is not, and neither is the hall temperature. The *state vector* is the collection of the unknown quantities that the inference estimates at once. In the equipment hall the unknowns are the hall temperature and the thermal conductance, and one possible representation is

$$\mathbf{z} = \begin{bmatrix} T \\ G \end{bmatrix}.$$

The state holds everything the inference treats as free; the pinned quantities T_a and Q are not in it, and neither is the reading y_T, which is evidence. The concept recurs throughout the book: the graph fixes the state, the evidence constrains it, and each question is a statement about it. Quantities that are entries of the state, such as T here, are answered directly by the inference; quantities that are functions of the state, such as the heat flow through the path, require one further transformation.

The log-conductance. A conductance is positive, and an estimate of it must be positive too. In practice the conductance is therefore often represented by a logarithmic entry,

$$g = \ln\!\left(\frac{G}{G_{\mathrm{ref}}}\right),$$

where $G_{\mathrm{ref}} > 0$ is a reference conductance in kW/K and g is dimensionless. Inverting,

$$G = G_{\mathrm{ref}}\, \mathrm{e}^{g},$$

which is positive for every real g: an inference free to move g anywhere on the real line can never produce a negative or zero conductance. A change in g is a relative change in G, since $\mathrm{d}G = G_{\mathrm{ref}}\, \mathrm{e}^{g}\, \mathrm{d}g = G\, \mathrm{d}g$, so that $\mathrm{d}g = \mathrm{d}G/G$; uncertainty in g reads directly as a fractional uncertainty in G. The reference value only sets the origin of g. In this chapter it is taken as $G_{\mathrm{ref}} = 6\,\mathrm{kW/K}$, the design conductance of the forward calculation, so that $g = 0$ there, $G = 7.5\,\mathrm{kW/K}$ corresponds to $g = \ln 1.25 = 0.2231$, and a conductance 10 % above the reference to $g = \ln 1.1 = 0.0953$. The state vector becomes

$$\mathbf{z} = \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} = \begin{bmatrix} T \\ g \end{bmatrix}, \tag{1.3}$$

with entries $z_1 = T$ in °C and $z_2 = g$, dimensionless. It is the one state of this chapter: every question below is a statement about these two entries.

The physics row. The thermal relation can now be written in terms of the state. Substituting $G = G_{\mathrm{ref}}\, \mathrm{e}^{g}$ into equation (1.1),

$$Q = G_{\mathrm{ref}}\, \mathrm{e}^{g}\, (T - T_a),$$

or, moving everything to one side,

$$r_{\mathrm{phys}}(\mathbf{z}) = G_{\mathrm{ref}}\, \mathrm{e}^{g}\, (T - T_a) - Q = 0. \tag{1.4}$$

The function $r_{\mathrm{phys}}(\mathbf{z})$ is a *residual*, in kW. A residual measures how far a candidate state is from satisfying a relation: it is zero when the relation holds exactly and grows as the candidate departs from it. With $G_{\mathrm{ref}} = 6\,\mathrm{kW/K}$, $T_a = 18°\mathrm{C}$ and $Q = 30\,\mathrm{kW}$, the candidate $\mathbf{z} = (23°\mathrm{C},\ 0)$ gives $r_{\mathrm{phys}} = 6 \times 1 \times 5 - 30 = 0$: it is the forward calculation. The candidate $\mathbf{z} = (23.2°\mathrm{C},\ 0)$ gives $6 \times 5.2 - 30 = 1.2\,\mathrm{kW}$: at that temperature and conductance the path would remove 1.2 kW more than the equipment supplies, so the hall could not be in steady state. The candidate $(23.2°\mathrm{C},\ \ln(30/31.2)) = (23.2°\mathrm{C},\ -0.0392)$ again gives zero: the temperature 23.2°C

is consistent with the physics if the conductance is $30/5.2 = 5.769\,\mathrm{kW/K}$. The zero set of r_{phys} is the curve $T = T_a + Q/(G_{\mathrm{ref}}\,\mathrm{e}^{g})$ in the plane of (g, T); the physics alone allows every point of that curve and no point off it.

The same row has a second form, in the units of the quantity it fixes. Dividing equation (1.4) by $G_{\mathrm{ref}}\,\mathrm{e}^{g} = G$, which is positive, leaves the zero set unchanged and gives the row in kelvin, the departure of a candidate temperature from the temperature the physics implies. Divided further by a *model tolerance* σ_p, the temperature error in kelvin that the physics is allowed, it becomes a dimensionless row,

$$r_1(\mathbf{z}) = \frac{T - T_a - (Q/G_{\mathrm{ref}})\,\mathrm{e}^{-g}}{\sigma_p}, \tag{1.5}$$

which counts the violation of the physics in multiples of its tolerance. This is the form in which the physics row enters every computation of the book. The formulation extends to systems of any size. A larger physical system is represented by a collection of residuals, one per relation: an electrical network has rows for Kirchhoff's current and voltage laws, a fluid network rows for mass conservation at its junctions and for the pressure losses along its pipes, and a cooling plant rows for the energy balance at each of its nodes. Each row has the form of equation (1.4), and the inference treats them alike.

Exact constraints and model discrepancy. Physical rows and reading rows should not automatically be treated as equivalent kinds of noisy measurement. Conservation of energy is a law: if the hall is in steady state and the model has captured every path, equation (1.4) holds exactly, and it can be imposed as an *exact constraint*, $r_{\mathrm{phys}}(\mathbf{z}) = 0$, which restricts the state to the zero set of the row. In a computation the exact constraint is carried as the physics row (1.5) with a model tolerance σ_p far smaller than every other spread of the problem, and exactness is the limit $\sigma_p \to 0$; the limit is checked by showing that the results stop moving as σ_p shrinks, never by using the row to eliminate an entry of the state. A reading row, by contrast, is meant to be violated by an amount of the order of the sensor's error: the reading is not the temperature.

Approximate physics is a third case. The single effective conductance of equation (1.1) is an idealization; the real hall has several paths, a conductance that varies with airflow, and stored heat. When the relation is only approximately true, the honest model is

$$G_{\mathrm{ref}}\,\mathrm{e}^{g}\,(T - T_a) - Q = \delta,$$

with an explicit *model discrepancy* δ in kW, an entry of its own with a stated prior, for example $\delta \sim \mathcal{N}(0, \sigma_\delta^2)$. Algebraically this is a physics row with the weight $1/\sigma_\delta^2$, and in the objective it looks like a reading row. It behaves differently. A measurement error is new at each reading and averages out over many readings; a discrepancy is a property of the model and repeats wherever the same operating condition repeats, so that many readings at one condition do not reduce it. Treating a discrepancy as measurement noise therefore makes the inference overconfident, and treating a law as noisy lets the inference trade physics against data. The choice between an exact constraint and an explicit discrepancy is made when the graph is built, and it is stated with the model. In this chapter conservation of energy is exact.

1.1.5 The record and the handbook value

The record of the hall is one reading. A temperature sensor installed in the hall reports, at the load of 30 kW,

$$y_T = 23.2^{\circ}\mathrm{C}.$$

The sensor is not perfectly accurate. Its measurement model is

$$y_T = T + \varepsilon, \qquad \varepsilon \sim \mathcal{N}(0, \sigma_T^2),$$

where ε is the measurement error in kelvin, assumed normal with zero mean, and σ_T its standard deviation, which the sensor's specification gives as

$$\sigma_T = 0.5\,\mathrm{K}.$$

The reading is the only observation: the ambient temperature and the load are pinned, and nothing measures the conductance.

Knowledge of the conductance that does not come from the reading is also in hand. An engineering handbook, the design documents or earlier tests of similar halls give an estimate of the effective conductance with an uncertainty; for this hall (illustrative) the handbook gives 7.5 kW/K, with the true conductance within a factor of about 1.35 of it with probability of about 0.68. The handbook value is not a measurement of this hall. It is information held before the reading, and it enters the model as a prior.

1.1.6 The questions (step 1)

The engineer's questions can now be written over the graph, each as a probability or an expectation over the state, with the unknown on the left of the bar and the evidence on the right. Table 1.2 lists the five that this chapter answers: the first column is the question in engineering language; the second the general form that every question of its category takes, written in the book's notation for a quantity $f(\mathbf{z})$ derived from the state, a threshold f^* on it and the posterior density $p(\mathbf{z} \mid \mathbf{y})$; the third its mathematical form for the hall; the fourth its category among the five of the preface, which Chapter 2 develops; and the last the subsection where it is answered. The first two are estimation, of an entry of the state and of a quantity derived from one, the conductance $G = G_{\mathrm{ref}}\,\mathrm{e}^{g}$. The third and fourth are predictions under a threshold, tail probabilities of the posterior: the third at the present load, the fourth at a load of 60 kW at which the hall has never run, whose temperature T_2 is a quantity the graph acquires in step 2. The fifth is a design question about a measurement not yet taken; it asks nothing new of the present evidence, and it is answered as the next measurement for the first four.

Table 1.2: The questions asked of the equipment hall, the general form of their category, their mathematical form for the hall, their category and where each is answered. T_2 is the hall temperature at the load $Q_2 = 60\,\mathrm{kW}$ (Section 1.1.7); $\boldsymbol{\Sigma}_a$ the covariance expected after a candidate row with Jacobian row $\mathbf{J}_a$.

Engineering question	General form	Mathematical form	Category	Answered in
What is the hall temperature now?	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$, $\mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y})$	$p(T \mid \mathbf{y})$, $T = z_1$	estimation	1.3.1
What is the effective conductance of the path?	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$, $\mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y})$	$p(G \mid \mathbf{y})$, $G = G_{\mathrm{ref}}\,\mathrm{e}^{z_2}$	estimation of a derived quantity	1.3.1
Will the hall stay below 24°C?	$\Pr(f(\mathbf{z}) < f^* \mid \mathbf{y})$	$\Pr(T < 24^\circ\mathrm{C} \mid \mathbf{y})$	prediction under a threshold	1.3.2
Will the hall stay below 27.5°C at 60 kW?	$\Pr(f(\mathbf{z}) < f^* \mid \mathbf{y})$	$\Pr(T_2 < 27.5^\circ\mathrm{C} \mid \mathbf{y})$, $T_2 = T_a + (Q_2/G_{\mathrm{ref}})\,\mathrm{e}^{-z_2}$	prediction under a threshold	1.3.2
Would another sensor narrow the answers?	preposterior covariance $(\boldsymbol{\Sigma}^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a)^{-1}$	$\boldsymbol{\Sigma}_a$ for each candidate row	design	1.6.1

The questions differ in what they compute and agree in what they are computed over: the same two entries of the state, under the same posterior. Two of the engineer's six questions

are not in the table. Whether the sensor is biased compares two complete models of the same reading, the plant with an unbiased sensor and the plant with a bias entry on the reading row, and its answer is a ratio of their evidences rather than a property of either posterior; when every sensor of a larger plant is a candidate, the same comparison run over the candidates is a diagnosis. Chapter 22, in Volume II, answers it, and the residual check of Section 1.4.3 takes a first, coarse look at it. What happens if the cooling partially fails is a question about a transient that the steady-state graph cannot represent, and Section 1.6.5 states what the graph would need for it.

1.1.7 The evidence as rows (step 2)

Step 2 attaches the reading and the handbook value to the graph as rows on the state $\mathbf{z} = (T, g)^{\mathsf{T}}$, and extends the graph with the operating condition the fourth question names. The three rows, the physics row of Section 1.1.4 among them, are Gaussian, and each is written normalized, with a residual and a Jacobian row.

The modeling choices. One choice makes the physics row, and it buys something and can fail. The energy balance is taken as exact, the limit $\sigma_p \to 0$ of the physics row: this is what confines the posterior to the physics curve, so that a reading of the temperature informs the conductance, and it fails if heat leaves the hall by a mechanism the single conductance does not lump or if the hall is not in steady state (Section 1.6.5); Section 1.4.1 checks that the tolerance used has reached the limit.

The reading row. The reading of Section 1.1.5 becomes the normalized observation residual

$$r_{\text{obs}}(\mathbf{z}) = \frac{T - y_T}{\sigma_T}, \tag{1.6}$$

which is dimensionless: the difference between the candidate temperature and the reading, in units of the sensor's standard deviation. At $T = 23.2°\text{C}$ it is zero; at the design temperature $T = 23°\text{C}$ it is $(23 - 23.2)/0.5 = -0.4$. The row carries the probability of the reading. Given the state, the reading differs from T only by the error, so its density is the normal density of ε evaluated at $y_T - T$:

$$p(y_T \mid \mathbf{z}) = \frac{1}{\sqrt{2\pi}\,\sigma_T} \exp\left[-\frac{(y_T - T)^2}{2\sigma_T^2}\right] = \frac{1}{\sqrt{2\pi}\,\sigma_T} \exp\left[-\tfrac{1}{2}\, r_{\text{obs}}(\mathbf{z})^2\right].$$

Read as a function of the state for the fixed reading, this is the *likelihood* of the state. A candidate state with a small $|r_{\text{obs}}|$ makes the reading probable; the ratio of the likelihoods of $T = 23°\text{C}$ and $T = 23.2°\text{C}$ is $\mathrm{e}^{-0.4^2/2} = \mathrm{e}^{-0.08} = 0.923$, so the reading alone hardly distinguishes the two. The likelihood does not depend on g: the sensor sees the temperature, and it informs the conductance only through the physics row that ties the two together.

The prior row. The handbook value becomes a row of the same form. Written as a normal distribution on the log-conductance, $g \sim \mathcal{N}(g_0, \sigma_g^2)$, it is the row $(g - g_0)/\sigma_g$: a handbook value is treated as a reading of g taken before the test. The handbook conductance of $7.5\,\text{kW/K}$ gives $g_0 = \ln(7.5/6) = 0.2231$, and $\sigma_g = 0.3$ states that the true conductance is within a factor $\mathrm{e}^{0.3} = 1.35$ of it with probability of about 0.68. The prior is written $p(g)$ in general; its form is a modeling choice that must be stated and justified, and Section 1.4.4 tests how much this one weighs.

Table 1.3: The three kinds of row in the equipment hall model, with $G = G_{\text{ref}}\,\mathrm{e}^{g}$.

Row	Equation	Residual	Role
Physics	$G\,(T - T_a) - Q = 0$	$G_{\text{ref}}\,\mathrm{e}^{g}(T - T_a) - Q$	constrains what is physically possible
Reading	$y_T = T + \varepsilon$	$(T - y_T)/\sigma_T$	provides imperfect information about the state
Prior	$g \sim \mathcal{N}(g_0, \sigma_g^2)$	$(g - g_0)/\sigma_g$	carries knowledge held before the readings

Three kinds of row. The model now holds three very different kinds of relation. Table 1.3 sets them side by side. Each row of the table is one kind of relation; the columns give its equation for the hall, its residual and its role.

The distinction between the first two rows is foundational. The physical equation constrains the system: it says which combinations of temperature and conductance can occur at all. The reading provides evidence about the unknown quantities: it says which of the possible combinations are plausible in light of what was measured. The prior row has the form of a reading row and the role of knowledge held in advance. A probabilistic inference method combines the three kinds of relation to obtain a distribution over the plausible states.

Rows and weights. Chapter 3 shows that, for normal errors, the negative logarithm of the posterior is, up to a constant, the objective $\mathcal{L}(\mathbf{z}) = \frac{1}{2}\sum_k w_k\, r_k(\mathbf{z})^2$, a weighted sum of the squared rows with weights $w_k = 1/\sigma_k^2$. A reading row may be written either normalized, as in equation (1.6), with unit weight, or in kelvin, as $T - y_T$ with weight $1/\sigma_T^2$; both put the same term $(T - y_T)^2/\sigma_T^2$ into the objective. Every row of this chapter is written normalized: the physics row in the form (1.5) with the tolerance σ_p, the reading row (1.6) and the prior row, so that every row counts in standard deviations and none dominates because of its unit.

The scenario node. The fourth question doubles the electrical load from 30 kW to 60 kW, an operating condition at which the hall has never run. A second operating condition is added to the graph, sharing the same unknown conductance. Figure 1.3 draws it. The upper row is the test condition, at which the hall has been operated and observed: the free node at $Q_1 = Q = 30\,\text{kW}$ with temperature T, its path to the pinned ambient, and its sensor. The lower row is the full-load condition, a scenario node at $Q_2 = 60\,\text{kW}$ with temperature T_2 and its own path to the same ambient. The vertical dashed line between the two paths marks that they carry one and the same conductance. The scenario node has no sensor: it is in the graph only because the question asks about it.

The two operating conditions satisfy the energy balance of equation (1.1) each at its own load,

$$Q_1 = G\,(T - T_a), \tag{1.7}$$

$$Q_2 = G\,(T_2 - T_a), \tag{1.8}$$

with $Q_1 = Q$. The two temperatures are different, but the conductance is shared, and that connection is what makes the second condition predictable from the first. The scenario node adds no evidence and no free quantity: by equation (1.8) its temperature is fixed by the conductance, $T_2 = T_a + (Q_2/G_{\text{ref}})\,\mathrm{e}^{-g}$, a function of the state $\mathbf{z} = (T, g)^{\mathsf{T}}$, which keeps its two entries and its three rows. Chapter 4 carries T_2 instead as a third entry with a second physics row, the same formulation; with the same rows it gives the same answers.

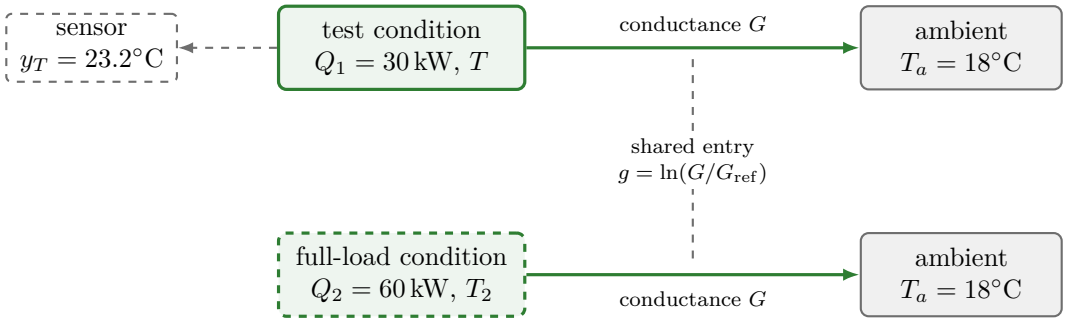


Figure 1.3: Two operating conditions of the equipment hall linked through the same physical conductance. Free: the test-condition temperature T and the shared log-conductance g, the two entries of the state; the full-load temperature T_2 (dashed node) is a function of g. Pinned: the ambient temperature and the two loads. Observed: T only, through the reading y_T; the scenario node is unobserved and enters the graph because the question asks about it.

Which row acts on which entry. The three rows act on the two entries of the state in a pattern fixed by the graph, before any number is computed. The physics row ties the temperature to the conductance and involves both; the sensor sees the temperature and not the conductance; the handbook speaks about the conductance and not the temperature; the scenario node adds no row. The Jacobian of the rows, the matrix of their derivatives with respect to the entries, can therefore be nonzero only where marked $*$:

$$\mathbf{J} = \begin{bmatrix} * & * \\ * & 0 \\ 0 & * \end{bmatrix} \begin{matrix} \text{physics row} \\ \text{reading row,} \\ \text{prior row} \end{matrix} \qquad \text{columns } z_1 = T,\ z_2 = g. \tag{1.9}$$

The zeros are structural: they hold at every state. Only the physics row has entries in both columns, and it is the one row through which a reading of the temperature reaches the conductance.

The factor graph of the hall. The pattern (1.9) is what a factor graph (Section 1.1.3) draws. Figure 1.4 draws it for the hall with its scenario node. The two solid circles are the entries of the state, T and g. The black squares are the rows: the physics row of equation (1.5), linked to both entries; the reading row, linked to T alone; and the prior row, linked to g alone. The dashed circle is the full-load temperature T_2, joined to g by a dashed line: it is a function of g through equation (1.8), not an entry, and no row acts on it. The figure makes visible what the physical graph only implies: T_2 has no reading of its own and is connected to the evidence only through g, and g is connected to the reading only through the physics row. The physical drawing has two ambient boxes and two paths; the factor graph has no node for the ambient at all, since T_a is pinned and appears only as a constant inside the physics row.

The residual vector. The state is the column vector of equation (1.3),

$$\mathbf{z} = \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} = \begin{bmatrix} T \\ g \end{bmatrix},$$

with T in °C and g dimensionless. Three rows act on it, the physics row (1.5), the reading row (1.6) and the prior row, each divided by its standard deviation, so that every row counts in standard deviations and none dominates because of its unit. Stacked into one residual vector,

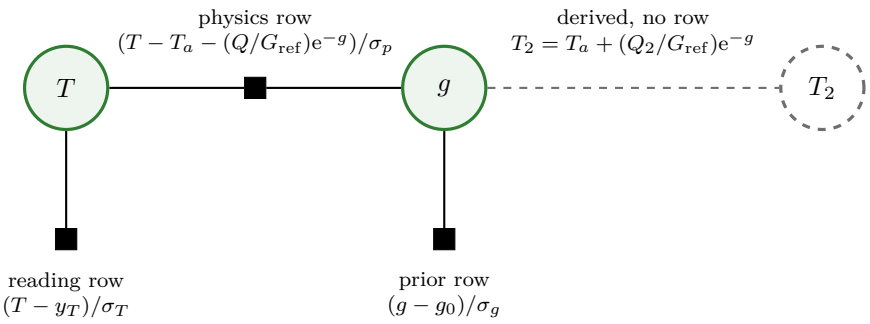


Figure 1.4: Factor graph of the equipment hall with its scenario node. Solid circles: the entries of the state, T and g. Squares: the rows of the model, each linked to the entries it involves, the pattern (1.9) of the Jacobian. Dashed: the full-load temperature T_2, a function of g with no row of its own, which reaches the evidence only through the shared g.

the three rows are

$$\mathbf{r}(\mathbf{z}) = \begin{bmatrix} r_1(\mathbf{z}) \\ r_2(\mathbf{z}) \\ r_3(\mathbf{z}) \end{bmatrix} = \begin{bmatrix} (T - T_a - (Q/G_{\mathrm{ref}})\,\mathrm{e}^{-g})/\sigma_p \\ (T - y_T)/\sigma_T \\ (g - g_0)/\sigma_g \end{bmatrix}, \tag{1.10}$$

with $T_a = 18°\mathrm{C}$, $Q = 30\,\mathrm{kW}$, $G_{\mathrm{ref}} = 6\,\mathrm{kW/K}$, the reading $y_T = 23.2°\mathrm{C}$ with $\sigma_T = 0.5\,\mathrm{K}$, and the handbook prior $g_0 = \ln(7.5/6) = 0.2231$ with $\sigma_g = 0.3$. Each entry of $\mathbf{r}$ is one row: a dimensionless number stating by how many standard deviations a candidate state violates one relation. The objective is half the sum of their squares,

$$\mathcal{L}(\mathbf{z}) = \tfrac{1}{2}\,\mathbf{r}(\mathbf{z})^{\mathsf{T}}\,\mathbf{r}(\mathbf{z}) = \tfrac{1}{2}\sum_{k=1}^{3} r_k(\mathbf{z})^2, \tag{1.11}$$

which Chapter 3 shows to be minus the logarithm of the posterior up to a constant, so that the most probable state $\hat{\mathbf{z}}$, the maximum a posteriori (MAP) state, is the minimizer of $\mathcal{L}$. Conservation of energy is exact in this model, and an exact constraint is the limit $\sigma_p \to 0$ of the physics row (Section 1.1.4). The computation takes $\sigma_p = 0.01\,\mathrm{K}$, fifty times tighter than the sensor, and Section 1.4.1 checks that the limit has been reached.

The Jacobian. The Jacobian $\mathbf{J}$ collects the derivative of every row with respect to every entry of the state: its entry in row k and column i is $\partial r_k/\partial z_i$, so it has one row per row of the model and one column per entry of the state, here three by two. Differentiating equation (1.10) entry by entry,

$$\mathbf{J}(\mathbf{z}) = \frac{\partial \mathbf{r}}{\partial \mathbf{z}} = \begin{bmatrix} 1/\sigma_p & (Q/G_{\mathrm{ref}})\,\mathrm{e}^{-g}/\sigma_p \\ 1/\sigma_T & 0 \\ 0 & 1/\sigma_g \end{bmatrix}, \tag{1.12}$$

the pattern (1.9) with its entries filled in. Each entry is a sensitivity in standard deviations: how far one row moves when one entry of the state moves by one unit. Only the entry in the first row and second column depends on the state; the others are constants. The first row is the physics row's own sensitivities in its physical unit, rescaled. Differentiating the kilowatt form of equation (1.4),

$$\frac{\partial r_{\mathrm{phys}}}{\partial T} = G_{\mathrm{ref}}\,\mathrm{e}^{g} = G, \qquad \frac{\partial r_{\mathrm{phys}}}{\partial g} = G_{\mathrm{ref}}\,\mathrm{e}^{g}\,(T - T_a) = G\,(T - T_a),$$

in kW/K and kW; on the zero set the second equals the load Q by equation (1.1). A one-kelvin error in the temperature unbalances the hall by G kilowatts, and a unit change in g, a factor e

in the conductance, by about Q kilowatts to first order. Divided by G, to reach the temperature form, and by σ_p, the pair becomes $(1/\sigma_p,\ (Q/G)/\sigma_p)$, the first row of equation (1.12) on the zero set.

1.2 The posterior of the state (step 3)

All the questions of Table 1.2 are read from one posterior: the posterior of the state that the physics row, the reading row and the prior row of step 2 give, approximately the Laplace normal of (T, g) and exactly a density along the physics curve. This section derives it and computes it for the hall; Section 1.3 reads every answer from it.

Four objects of the fit recur in the specialized forms, and they are named here by their roles. The *MAP state* $\hat{\mathbf{z}} = (\hat{T}, \hat{g})^{\mathsf{T}}$, the maximum a posteriori state, is the minimum of the objective that the three rows put on the state, the most probable state; the *information matrix* $\mathbf{\Lambda} = \mathbf{J}^{\mathsf{T}}\mathbf{J}$ is the curvature that the rows give that objective at the minimum, with $\mathbf{J}$ the 3×2 Jacobian of the normalized rows; the *Laplace covariance* $\mathbf{\Sigma}$, a 2×2 matrix, is the inverse of that curvature after a term from the physics row's own curvature is added to it, and it makes the posterior of the state approximately normal, $\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$; and a derived quantity f is carried by its *gradient* $\mathbf{g} = \partial f/\partial \mathbf{z}$ at $\hat{\mathbf{z}}$ (bold, to tell it from the entry g), with the variance $\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}$. Section 1.2.1 derives the four from the rows. Beside them stands the exact posterior that this problem admits in the limit $\sigma_p \to 0$: on the physics curve the state has one free direction, and the posterior is the density of one entry, evaluated by quadrature. The answers are reported by the Laplace route, a normal distribution centered on the MAP state: this is what lets every answer be read from $\hat{\mathbf{z}}$ and $\mathbf{\Sigma}$ by operations that serve a plant of any size, and it misleads where the posterior is skewed along the curved ridge, which Section 1.4.2 measures against the exact posterior. The chapter reports the Laplace values, the route that carries over unchanged to plants of many entries, sets the exact value beside each answer with the gap, and collects the comparison in Section 1.4.

1.2.1 The posterior from the rows

The answers use the objects of one posterior: the MAP state $\hat{\mathbf{z}}$, the covariance $\mathbf{\Sigma}$ and the Jacobian $\mathbf{J}$ of the rows. The Jacobian is written out in Section 1.1.7; the other two are derived here for the hall. Chapter 3 derives these objects from the probability rules for Gaussian rows over any state, and Chapter 4 dissects the same objects in full on the hall with its scenario node carried as an entry, including the Laplace approximation and its curvature term. Here each is written out once for a reader meeting it for the first time.

The posterior. Four ingredients are in hand: the unknown state $\mathbf{z} = (T, g)^{\mathsf{T}}$; the observation $y_T = 23.2°\mathrm{C}$; the physical relation $Q = G\,(T - T_a)$; and the prior information $p(g)$. What is ultimately wanted is

$$p(\mathbf{z} \mid \mathbf{y}), \tag{1.13}$$

the *posterior distribution* of the state, conditional on the evidence $\mathbf{y}$, here the reading y_T and the handbook value, and on the assumed physical model. It states which values of the state remain plausible after the physics, the measurement and the prior information have been combined. Chapter 3 derives Bayes' theorem, by which the posterior is proportional to the likelihood times the prior, $p(\mathbf{z} \mid \mathbf{y}) \propto p(y_T \mid \mathbf{z})\, p(\mathbf{z})$; here the reading row plays the part of the likelihood, and the physics row and the prior row on g together play the part of the prior $p(\mathbf{z})$.

The normal equations and the Gauss–Newton step. Near a state $\mathbf{z}$, each row is replaced by its tangent, $\mathbf{r}(\mathbf{z} + \Delta\mathbf{z}) \approx \mathbf{r}(\mathbf{z}) + \mathbf{J}\,\Delta\mathbf{z}$, and the objective by $\frac{1}{2}\,\|\mathbf{r} + \mathbf{J}\,\Delta\mathbf{z}\|^2$, a quadratic in the

step $\Delta\mathbf{z}$. Its gradient with respect to the step is $\mathbf{J}^\mathsf{T}(\mathbf{r} + \mathbf{J}\,\Delta\mathbf{z})$, and setting it to zero gives the *normal equations*

$$(\mathbf{J}^\mathsf{T}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^\mathsf{T}\mathbf{r}, \tag{1.14}$$

a linear system of two equations in the two entries of the step. The Gauss–Newton method solves it, moves to $\mathbf{z} + \Delta\mathbf{z}$, evaluates $\mathbf{r}$ and $\mathbf{J}$ there, and repeats until the step no longer changes the state at the precision wanted. The vector $\mathbf{J}^\mathsf{T}\mathbf{r}$ is the gradient of the objective (1.11), so the iteration stops where that gradient vanishes, at the MAP state. The iteration starts here from the design state $\mathbf{z}^0 = (23°\mathrm{C},\ 0)^\mathsf{T}$ of the forward calculation, at which the physics row is zero.

The information matrix. The matrix $\mathbf{\Lambda} = \mathbf{J}^\mathsf{T}\mathbf{J}$ of the normal equations is the *information matrix* of the rows. It is a sum over the rows: writing $\mathbf{J}_k$ for the k-th row of $\mathbf{J}$,

$$\mathbf{J}^\mathsf{T}\mathbf{J} = \sum_{k=1}^{3} \mathbf{J}_k^\mathsf{T}\mathbf{J}_k,$$

so that each row contributes the outer product of its own sensitivities, and the precisions of independent rows add. The information matrix of normalized rows is $\mathbf{J}^\mathsf{T}\mathbf{J}$; when the rows are written in their physical units with weights $w_k = 1/\sigma_k^2$, as in Section 1.1.7, the same matrix is $\mathbf{J}^\mathsf{T}\mathbf{W}\mathbf{J}$, with $\mathbf{W} = \mathrm{diag}(w_k)$ and $\mathbf{J}$ the Jacobian of the unnormalized rows. For the hall the three outer products are of two kinds. The reading row and the prior row each have one nonzero entry, so each contributes one diagonal entry, $1/\sigma_T^2$ in the T–T position and $1/\sigma_g^2$ in the g–g position. The physics row has two, and its outer product $\mathbf{J}_1^\mathsf{T}\mathbf{J}_1$ fills the whole matrix, off-diagonal entries included: it is the only source of the coupling between the two entries, and it has rank one, with all its information along its own row and none along the tangent of the curve. The eigenvectors of $\mathbf{J}^\mathsf{T}\mathbf{J}$ give the directions in which the rows fix the state, and its eigenvalues how sharply: a large one across the curve, set by $1/\sigma_p^2$, and a small one along it, set by the reading and the prior, as the ridge below makes exact.

The Laplace covariance and the curvature term. The inverse of the information matrix is the covariance of the posterior: exactly so for rows linear in the state (Chapter 3), and as the *Laplace approximation*, the normal distribution fitted at the MAP state, for curved rows (Chapter 4). The Gauss–Newton matrix $\mathbf{J}^\mathsf{T}\mathbf{J}$ drops the curvature of the rows themselves. The full Hessian of the objective (1.11), its matrix of second derivatives, is

$$\frac{\partial^2 \mathcal{L}}{\partial \mathbf{z}\,\partial \mathbf{z}^\mathsf{T}} = \mathbf{J}^\mathsf{T}\mathbf{J} + \sum_{k=1}^{3} r_k \frac{\partial^2 r_k}{\partial \mathbf{z}\,\partial \mathbf{z}^\mathsf{T}}, \tag{1.15}$$

and of the three rows only the physics row is curved, through e^{-g}, with the single nonzero second derivative

$$\frac{\partial^2 r_1}{\partial g^2} = -\frac{(Q/G_\mathrm{ref})\,\mathrm{e}^{-g}}{\sigma_p}.$$

The *curvature term* $r_1\,\partial^2 r_1/\partial g^2$ therefore adds to the g–g entry alone. It does not vanish in the exact-physics limit, because r_1 shrinks in proportion to σ_p while $\partial^2 r_1/\partial g^2$ grows as $1/\sigma_p$: at the MAP state the gradient $\mathbf{J}^\mathsf{T}\mathbf{r}$ is zero, its T entry reads $r_1/\sigma_p + r_2/\sigma_T = 0$, so $r_1 = -(\sigma_p/\sigma_T)\,r_2$, and

$$r_1 \frac{\partial^2 r_1}{\partial g^2} = \frac{r_2}{\sigma_T}\,\frac{Q}{G_\mathrm{ref}}\,\mathrm{e}^{-\hat{g}}$$

at every σ_p. It is the curvature of the physics curve itself, and it acts in the one direction, along the curve, in which the information is small. The *Laplace covariance* is the inverse of the full Hessian at the MAP state,

$$\mathbf{\Sigma} = \Big(\mathbf{J}^\mathsf{T}\mathbf{J} + \sum_k r_k \frac{\partial^2 r_k}{\partial \mathbf{z}\,\partial \mathbf{z}^\mathsf{T}}\Big)^{-1}_{\hat{\mathbf{z}}}, \qquad \mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma}); \tag{1.16}$$

a Newton iteration, which uses the full Hessian in place of $\mathbf{J}^\mathsf{T}\mathbf{J}$ in equation (1.14), reaches the same $\hat{\mathbf{z}}$.

The ridge. The spreads of the two entries are not independent of each other, and the reason is read off the physics row. Its Jacobian row, the derivative of equation (1.5) with respect to (T, g), and the vector it annihilates are

$$\mathbf{J}_1 = \frac{1}{\sigma_p}\begin{bmatrix} 1 & (Q/G_{\mathrm{ref}})\,\mathrm{e}^{-g} \end{bmatrix}, \qquad \mathbf{n} = \begin{bmatrix} (Q/G_{\mathrm{ref}})\,\mathrm{e}^{-g} \\ -1 \end{bmatrix}, \qquad \mathbf{J}_1\,\mathbf{n} = 0. \tag{1.17}$$

A step along $\mathbf{n}$ leaves the physics row unchanged to first order, so $\mathbf{n}$ is the tangent to the zero set of the row, a change of $(Q/G_{\mathrm{ref}})\,\mathrm{e}^{-g}$ kelvin of temperature per unit of g with the opposite sign. With $\mathbf{J}_k$ the k-th row of $\mathbf{J}$, the information is a sum over rows, $\mathbf{J}^\mathsf{T}\mathbf{J} = \sum_k \mathbf{J}_k^\mathsf{T}\mathbf{J}_k$. The information the physics row contributes, $\mathbf{J}_1^\mathsf{T}\mathbf{J}_1$, is a matrix of rank one: it grows as $1/\sigma_p^2$ across the curve and is zero along $\mathbf{n}$. Along the curve only the reading row and the prior row inform the state, and by the pattern (1.9) each acts on one entry, with the contributions $\mathbf{J}_2^\mathsf{T}\mathbf{J}_2 = \mathrm{diag}(1/\sigma_T^2,\ 0)$ and $\mathbf{J}_3^\mathsf{T}\mathbf{J}_3 = \mathrm{diag}(0,\ 1/\sigma_g^2)$; the physics row's own curvature adds the term $\mathbf{C} = \mathrm{diag}(0,\ n_1 r_2/\sigma_T)$, derived above, with r_2 the reading row at the MAP state. With $\mathbf{u} = \mathbf{n}/\|\mathbf{n}\|$ the unit tangent, the information along the curve is

$$\lambda_\parallel = \mathbf{u}^\mathsf{T}(\mathbf{J}_2^\mathsf{T}\mathbf{J}_2 + \mathbf{J}_3^\mathsf{T}\mathbf{J}_3 + \mathbf{C})\mathbf{u} = \frac{n_1^2/\sigma_T^2 + 1/\sigma_g^2 + n_1 r_2/\sigma_T}{n_1^2 + 1}, \qquad n_1 = (Q/G_{\mathrm{ref}})\,\mathrm{e}^{-\hat{g}}, \tag{1.18}$$

which does not depend on σ_p. The tangent $\mathbf{u}$ is orthogonal to $\mathbf{J}_1$, so in the orthonormal basis $(\mathbf{J}_1^\mathsf{T}/\|\mathbf{J}_1\|,\ \mathbf{u})$ the curvature of the objective has the entry $\lambda_\parallel$ along $\mathbf{u}$, an entry across the curve that grows as $1/\sigma_p^2$, and a coupling that stays bounded; its inverse, by the Schur complement of Chapter 3, leaves $1/\lambda_\parallel$ along $\mathbf{u}$ and nothing else as $\sigma_p \to 0$. The covariance therefore tends to the rank-one matrix $\boldsymbol{\Sigma} \to \mathbf{u}\,\mathbf{u}^\mathsf{T}/\lambda_\parallel$: the posterior collapses onto a segment of the tangent, a *ridge* of the objective. Then $\sqrt{\Sigma_{11}} \to |u_1|/\sqrt{\lambda_\parallel}$ and $\sqrt{\Sigma_{22}} \to |u_2|/\sqrt{\lambda_\parallel}$ are one spread along the curve seen through the two components of the tangent, and the correlation $\Sigma_{12}/\sqrt{\Sigma_{11}\Sigma_{22}}$ tends to the sign of $u_1 u_2$, which is -1. The two entries are not known separately: a hall slightly warmer than $\hat{T}$ is a hall with a slightly lower conductance, and the posterior allows the one only together with the other. The reading row has a zero in the column of g, and it narrows the conductance all the same, through the first component of $\mathbf{u}$.

> **Principle 1.2.** *Physical rows constrain the system; reading rows provide evidence about its unknown quantities. The posterior is where the two agree, weighted by the prior, and neither kind of row alone determines it. An exact physical row leaves the posterior only along its own zero set, and its signature in $\boldsymbol{\Sigma}$ is a correlation near -1 or $+1$.*

The exact posterior along the curve. In the limit $\sigma_p \to 0$ the posterior lies on the physics curve, and along the curve it can be computed without approximation. On the curve the temperature is a function of the conductance entry, $T(g) = T_a + (Q/G_{\mathrm{ref}})\,\mathrm{e}^{-g}$. The rows define the posterior relative to a base flat in (T, g). Integrated across the curve, the physics row contributes a factor that does not depend on g, because the row is written in kelvin and its slope in T, $1/\sigma_p$, does not depend on g; the base is therefore flat in g along the curve, and what remains is the density of g along the curve, the reading row and the prior row evaluated on it:

$$p(g \mid \mathbf{y}) \propto \exp\left[-\frac{1}{2}\left(\frac{T(g) - y_T}{\sigma_T}\right)^2 - \frac{1}{2}\left(\frac{g - g_0}{\sigma_g}\right)^2\right]. \tag{1.19}$$

One-dimensional quadrature of this density over g gives every probability and moment of a quantity derived from the state exactly, for T through $T(g)$ and for G through $G_{\mathrm{ref}}\,\mathrm{e}^{g}$. The chapter computes every answer from the Laplace normal and sets the exact value from this density beside it, with the gap between the two.

The answer read off. An answer that is an entry of the state is read from the two objects directly: the entry $\hat{z}_i$ of $\hat{\mathbf{z}}$, with the spread $\sqrt{\Sigma_{ii}}$. The relation between two entries is the correlation $\Sigma_{ij}/\sqrt{\Sigma_{ii}\Sigma_{jj}}$, and the directions of the posterior are the eigenvectors of $\boldsymbol{\Sigma}$. A derived quantity f is carried by its gradient, with the spread $\sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}}$, and a threshold on an entry is a normal tail of its marginal. Section 1.2.2 computes $\hat{\mathbf{z}}$, $\mathbf{J}$ and $\boldsymbol{\Sigma}$ for the hall.

1.2.2 The posterior of the hall

The hall's inputs are the pinned values $T_a = 18°\mathrm{C}$, $Q = Q_1 = 30\,\mathrm{kW}$ and $G_{\mathrm{ref}} = 6\,\mathrm{kW/K}$; the reading $y_T = 23.2°\mathrm{C}$ with $\sigma_T = 0.5\,\mathrm{K}$; the handbook prior $g_0 = 0.2231$ with $\sigma_g = 0.3$; and the tolerance $\sigma_p = 0.01\,\mathrm{K}$ of the physics row.

The first step. At the design state $\mathbf{z}^0 = (23°\mathrm{C},\ 0)^{\mathsf{T}}$ the Jacobian (1.12) is

$$\mathbf{J}(\mathbf{z}^0) = \begin{bmatrix} 100 & 500 \\ 2 & 0 \\ 0 & 3.333 \end{bmatrix} :$$

a one-kelvin change in T moves the physics row by 100 and the reading row by 2, and a unit change in g moves the physics row by 500 and the prior row by 3.33. The first row is the pair of kilowatt sensitivities $(G,\ Q) = (6,\ 30)$ divided by $G = 6\,\mathrm{kW/K}$ and by σ_p. The residual vector there is $\mathbf{r} = (0,\ -0.400,\ -0.744)$: the physics holds, the temperature lies 0.4 standard deviations below the reading, and g lies 0.74 standard deviations below the handbook value. The two sides of the normal equations (1.14) are

$$\mathbf{J}(\mathbf{z}^0)^{\mathsf{T}}\mathbf{J}(\mathbf{z}^0) = \begin{bmatrix} 10004 & 50000 \\ 50000 & 250011 \end{bmatrix}, \qquad -\mathbf{J}(\mathbf{z}^0)^{\mathsf{T}}\mathbf{r}(\mathbf{z}^0) = \begin{bmatrix} 0.800 \\ 2.479 \end{bmatrix},$$

and the first step is their solution,

$$\Delta\mathbf{z} = \begin{bmatrix} 0.0684 \\ -0.0137 \end{bmatrix}.$$

The iterations. Table 1.4 lists the iterations. Each row is one iterate $\mathbf{z}^k$; the columns give its two entries, its three rows and the objective of equation (1.11).

Table 1.4: Gauss–Newton iterations for the posterior of the equipment hall (equations (1.10) to (1.14)), with $\sigma_p = 0.01\,\mathrm{K}$, from the design state $(23°\mathrm{C},\ 0)$.

k	T (°C)	g [–]	r_1 [–]	r_2 [–]	r_3 [–]	$\mathcal{L}$ [–]
0	23.0000	0.0000	0.000	−0.400	−0.744	0.3566
1	23.0684	−0.0137	−0.042	−0.263	−0.789	0.3471
2	23.0701	−0.0139	0.005	−0.260	−0.790	0.3460
3	23.0701	−0.0139	0.005	−0.260	−0.790	0.3460

The first step does almost all the work, because over a step this short the rows are nearly linear. It leaves the physics row slightly violated, $r_1 = -0.042$, through the curvature of e^{-g}

that the tangent ignores; the second step restores it, and the third changes nothing at the printed precision. The MAP state is

$$\hat{\mathbf{z}} = \begin{bmatrix} \hat{T} \\ \hat{g} \end{bmatrix} = \begin{bmatrix} 23.070\,^{\circ}\mathrm{C} \\ -0.0139 \end{bmatrix}, \qquad \hat{G} = 6\,\mathrm{e}^{-0.0139} = 5.917\,\mathrm{kW/K}. \tag{1.20}$$

There the physics holds to $r_1\,\sigma_p = 5 \times 10^{-5}\,\mathrm{K}$, and the reading and the prior share the disagreement: the temperature lies 0.26 standard deviations below the reading and g lies 0.79 standard deviations below the handbook value.

The information matrix. At the MAP state the Jacobian (1.12) is

$$\mathbf{J}(\hat{\mathbf{z}}) = \begin{bmatrix} 100 & 507.0 \\ 2 & 0 \\ 0 & 3.333 \end{bmatrix}, \qquad \frac{(Q/G_{\mathrm{ref}})\,\mathrm{e}^{-\hat{g}}}{\sigma_p} = \frac{5\,\mathrm{e}^{0.0139}}{0.01} = 507.0,$$

and the information matrix is the sum of the three outer products,

$$\mathbf{J}^{\mathsf{T}}\mathbf{J} = \underbrace{\begin{bmatrix} 10000 & 50701 \\ 50701 & 257056 \end{bmatrix}}_{\text{physics}} + \underbrace{\begin{bmatrix} 4 & 0 \\ 0 & 0 \end{bmatrix}}_{\text{reading}} + \underbrace{\begin{bmatrix} 0 & 0 \\ 0 & 11.11 \end{bmatrix}}_{\text{prior}} = \begin{bmatrix} 10004 & 50701 \\ 50701 & 257067 \end{bmatrix}.$$

The physics term has rank one: all its information lies along its own row (100, 507.0), across the physics curve, and none along the curve; only the reading's 4 and the prior's 11.11 inform along it. The eigenvalues of $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ are 267067 across the curve and 4.27 along it, with the eigenvector $(0.9811,\ -0.1935)^{\mathsf{T}}$ of the small one pointing along the curve; the first two terms of equation (1.18) give it directly, $(5.070^2 \times 4 + 11.11)/(5.070^2 + 1) = 4.27$. The Gauss–Newton covariance is the inverse,

$$\mathbf{\Sigma}_{\mathrm{GN}} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1} = \begin{bmatrix} 0.2256 & -0.0445 \\ -0.0445 & 0.0088 \end{bmatrix},$$

with the standard deviations $\sqrt{0.2256} = 0.475\,\mathrm{K}$ and $\sqrt{0.0088} = 0.094$.

The curvature term and the Laplace covariance. The second derivative of the physics row at the MAP state is $\partial^2 r_1/\partial g^2 = -507.0$; multiplied by $r_1 = 0.0052$, it adds -2.63 to the g–g entry, $257067 \to 257064$, the same value as $(r_2/\sigma_T)(Q/G_{\mathrm{ref}})\,\mathrm{e}^{-\hat{g}} = (-0.260/0.5) \times 5.070$. Projected on the along-curve eigenvector $(0.9811,\ -0.1935)^{\mathsf{T}}$, the term is $-2.63 \times 0.1935^2 = -0.099$ against the along-curve information 4.27, a loss of 2.3 % in the one direction in which the information is small; it raises the variances by 2.4 % and the standard deviations by 1.2 %. The Laplace covariance (1.16) is

$$\mathbf{\Sigma} = \begin{bmatrix} 0.23096 & -0.045552 \\ -0.045552 & 0.0089882 \end{bmatrix}, \qquad \sqrt{\Sigma_{11}} = 0.4806\,\mathrm{K}, \quad \sqrt{\Sigma_{22}} = 0.09481. \tag{1.21}$$

These are the Laplace values used in the rest of the chapter, quoted as 0.481 K and 0.0948; Gauss–Newton alone gives 0.475 K and 0.094. The correlation of the two entries, computed from the unrounded $\mathbf{\Sigma}$, is

$$\frac{\Sigma_{12}}{\sqrt{\Sigma_{11}\,\Sigma_{22}}} = \frac{-0.045552}{0.48058 \times 0.094806} = -0.9998.$$

The ridge. The physics row of $\mathbf{J}(\hat{\mathbf{z}})$, $(100,\ 507.0)$, is annihilated by $\mathbf{n} = (5.070,\ -1)^{\mathsf{T}}$, equation (1.17) at $\hat{g}$, and the long axis of $\boldsymbol{\Sigma}$ is $\mathbf{u} = (0.9811,\ -0.1935)^{\mathsf{T}}$, the same direction normalized. By equation (1.18),

$$\lambda_{\parallel} = \frac{5.070^2 \times 4 + 11.11 + 5.070 \times (-0.260)/0.5}{5.070^2 + 1} = 4.168,$$

$$\frac{0.9811}{\sqrt{4.168}} = 0.481\,\mathrm{K}, \qquad \frac{0.1935}{\sqrt{4.168}} = 0.0948,$$

the two diagonal entries of equation (1.21).

Figure 1.5 shows the geometry. The horizontal axis is the log-conductance g and the vertical axis the hall temperature T, the two entries of the state, so that every point of the plane is a candidate state. The green curve is the zero set of the physics row: the states the physics allows. The gray band is the set of temperatures within one standard deviation of the reading, $23.2 \pm 0.5°\mathrm{C}$, and the dashed line inside it the reading itself. Points are labeled by their plane coordinates (g, T), horizontal first. The three black dots are the design state, $(g, T) = (0, 23)$; the state on the curve that matches the reading exactly, $(g, T) = (-0.039, 23.2)$; and, between them, the MAP state of equation (1.20), $(g, T) = (-0.014, 23.07)$. The purple segment is the one-standard-deviation ellipse of the covariance of equation (1.21), centered on the MAP state; it is drawn as a thick line because its half-width across the curve, about 0.002 in g, is below the resolution of the figure. The dotted red line marks the threshold g^* of the full-load question, drawn at the value computed in Section 1.3.2.

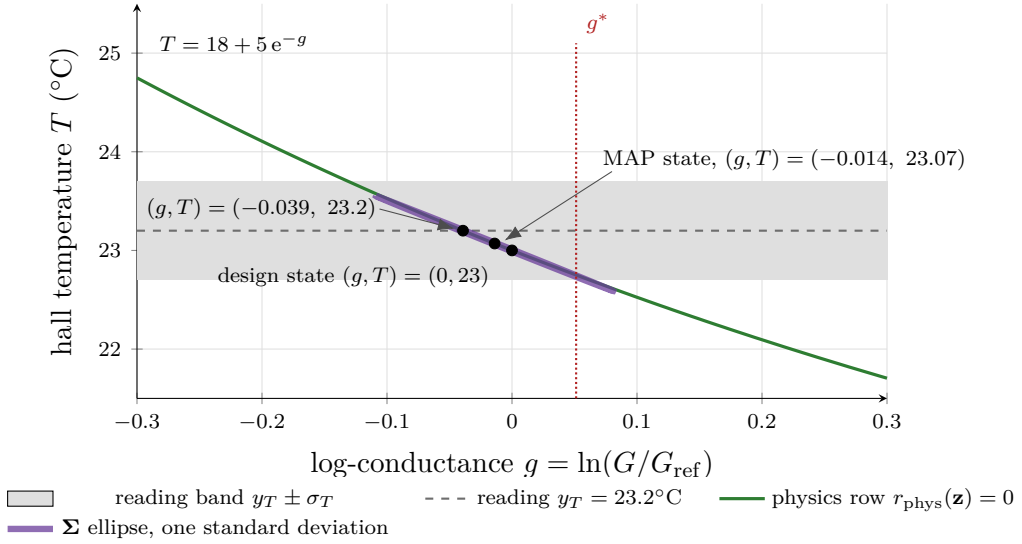


Figure 1.5: The state plane of the equipment hall with $T_a = 18°\mathrm{C}$, $Q = 30\,\mathrm{kW}$ and $G_{\mathrm{ref}} = 6\,\mathrm{kW/K}$. Green: the zero set of the physics row, $T = 18 + 5\,\mathrm{e}^{-g}$ (equation (1.4)). Gray: the reading $y_T = 23.2°\mathrm{C}$ with its band of one standard deviation, $\sigma_T = 0.5\,\mathrm{K}$ (equation (1.6)). The plausible states lie on the arc of the curve inside the band. Purple: the one-standard-deviation ellipse of the Laplace covariance (equation (1.21)) around the MAP state, $(g, T) = (-0.014,\ 23.07)$, collapsed onto the curve by the physics row. Points are labeled (g, T), horizontal coordinate first. Red: the threshold g^* of the full-load question (Section 1.3.2), above which the hall meets the limit at 60 kW.

The plausible states lie on the arc of the curve inside the band, from $g = -0.131$ to $g = 0.062$ ($G = 5.26$ to $6.38\,\mathrm{kW/K}$), and the MAP state sits on it between the reading-match point and the design state, pulled toward the handbook value $g_0 = 0.2231$ by the prior row.

The pair $(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of equations (1.20) and (1.21) is what step 4 uses: every answer of the chapter is computed from it. Because the physics row is exact in the limit, the posterior also

has the exact form (1.19) along the curve, and each answer has an exact value beside it, with the gap between the two reported.

1.3 The answers (step 4)

The first four questions of Table 1.2 take two general forms, written in the book's notation: $p(\mathbf{z} \mid \mathbf{y})$ is the posterior density of the state, $f(\mathbf{z})$ a quantity derived from the state and f^* a threshold on it. The fifth, a design question about a row not yet attached, is derived with the next measurement in Section 1.6.1. Each is answered from the posterior of Section 1.2: its engineering question, the general form of its category, the form it takes for the hall, the quantities of the posterior it uses, and the hall's numbers. The questions name the limit of 24°C at the present load and the limit of 27.5°C at $Q_2 = 60\,\mathrm{kW}$. Every answer is computed from the MAP state $\hat{\mathbf{z}}$ and the Laplace covariance $\boldsymbol{\Sigma}$ of step 3, and that is the value the chapter reports; the exact density along the physics curve, equation (1.19), gives the value set beside it, with the gap.

1.3.1 The hall temperature and the conductance

Engineering questions. What is the hall temperature now? What is the effective conductance of the path? The first asks for the temperature of the hall at the test load, given its one reading and the handbook value; the second for the conductance of its heat-removal path, the quantity a cooling engineer compares with the design value.

General form. An estimation question about a quantity $f(\mathbf{z})$ derived from the state asks for its posterior distribution, the probability that the quantity lies at or below each value f^*,

$$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y}) = \int \mathbf{1}\big[f(\mathbf{z}) \le f^*\big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \tag{1.22}$$

where $\mathbf{1}[\cdot]$ is one when the event holds and zero otherwise, summarized by a location and a spread,

$$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}] = \int f(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \qquad \mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y}) = \mathbb{E}\big[(f(\mathbf{z}) - \mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}])^2 \mid \mathbf{y}\big]. \tag{1.23}$$

An entry of the state is the case $f(\mathbf{z}) = z_i$.

Specialized form. The two questions are equations (1.22) and (1.23) with the state $\mathbf{z} = (T, g)^{\mathsf{T}}$ of the hall, the derived quantity $f(\mathbf{z}) = z_1 = T$ for the temperature and $f(\mathbf{z}) = z_2 = g$, or $f(\mathbf{z}) = G = G_{\mathrm{ref}}\, \mathrm{e}^{z_2}$, for the conductance, and $p(\mathbf{z} \mid \mathbf{y})$ the posterior of step 3, the Laplace normal $\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of equations (1.20) and (1.21). Each answer is read from $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$; the exact density (1.19) gives the value set beside it.

The temperature. Here $f(\mathbf{z}) = z_1 = T$, the first entry of the state. Under the Laplace posterior the marginal of an entry of a normal vector is normal, with the entry's mean and the diagonal entry of the covariance, so

$$T \mid \mathbf{y} \approx \mathcal{N}(\hat{z}_1, \Sigma_{11}),$$

and the answer is read off without any transformation: the first entry of $\hat{\mathbf{z}}$, with the spread $\sqrt{\Sigma_{11}}$.

The conductance. The second entry $z_2 = g$ is read off in the same way, $g \mid \mathbf{y} \approx \mathcal{N}(\hat{z}_2, \Sigma_{22})$. The conductance itself, $G = G_{\mathrm{ref}}\, \mathrm{e}^{z_2}$, is a derived quantity of the state. Linearized at the MAP state, its variance is $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ with the gradient

$$\mathbf{g} = \frac{\partial G}{\partial \mathbf{z}} = \begin{bmatrix} \partial G/\partial T \\ \partial G/\partial g \end{bmatrix}_{\hat{\mathbf{z}}} = \begin{bmatrix} 0 \\ G_{\mathrm{ref}}\, \mathrm{e}^{\hat{g}} \end{bmatrix} = \begin{bmatrix} 0 \\ \hat{G} \end{bmatrix}, \qquad \sigma_G^2 = \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} = \hat{G}^2\, \Sigma_{22}, \tag{1.24}$$

the delta-method spread σ_f of the notation table for $f = G$, written σ_G. The ratio $\sigma_G/\hat{G} = \sqrt{\Sigma_{22}}$: the spread of the log-conductance is the fractional spread of the conductance, as Section 1.1.4 anticipated.

The value beside each comes from the exact posterior: the moments (1.23) under the density (1.19), evaluated by one-dimensional quadrature over g, with T through $T(g) = T_a + (Q/G_{\text{ref}})\,\mathrm{e}^{-g}$ and G through $G_{\text{ref}}\,\mathrm{e}^{g}$. Its location is the posterior mean; the Laplace location is the entry of $\hat{\mathbf{z}}$, the mode, and the two differ by the skew of the posterior along the curve.

Building blocks. The MAP state $\hat{\mathbf{z}} = (23.070^\circ\text{C},\ -0.0139)^{\mathsf{T}}$ and the Laplace covariance $\boldsymbol{\Sigma}$, curvature term included, with $\sqrt{\Sigma_{11}} = 0.4806\,\text{K}$ and $\sqrt{\Sigma_{22}} = 0.09481$ (equations (1.20) and (1.21)); the gradient (1.24); and, for the exact values beside them, the density (1.19) and a one-dimensional quadrature over g.

The hall temperature now. The hall temperature is the first entry of the state: $\hat{T} = 23.070^\circ\text{C}$ with $\sqrt{\Sigma_{11}} = 0.481\,\text{K}$ (Gauss–Newton alone, 0.475 K). The value beside it comes from the exact posterior: the mean of T along the curve is 23.030°C, 0.040 K below $\hat{T}$, with the standard deviation 0.4813 K, 0.0007 K wider.

The effective conductance. The conductance entry is $\hat{g} = -0.0139$ with $\sqrt{\Sigma_{22}} = 0.0948$. By equation (1.24), with $\mathbf{g} = (0,\ 5.917)^{\mathsf{T}}$,

$$\hat{G} = 5.917\,\text{kW/K}, \qquad \sigma_G = \sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{g}} = 5.917\,\sqrt{0.0089882} = 0.561\,\text{kW/K},$$

a band of 9.5 % of its value, against a factor $\mathrm{e}^{0.3} = 1.35$ in the handbook before the reading. The values beside them come from the exact posterior: the conductance entry has the mean −0.0013, 0.013 above $\hat{g}$, and the standard deviation 0.0967, 0.0019 wider; the conductance has the mean 6.020 kW/K, 0.10 kW/K above $\hat{G}$, and the standard deviation 0.591 kW/K, 0.030 kW/K or 5 % wider. The exact posterior extends further toward large conductances than the normal at $\hat{\mathbf{z}}$ does, the skew that the full-load question below meets in its tail.

1.3.2 The limits at the present load and at 60 kilowatts

Engineering questions. Will the hall stay below 24°C? Will the hall stay below 27.5°C at 60 kW? The first asks about the hall as it runs now. The second is the owner's question,

What is the probability that the equipment hall will remain below 27.5°C *when the load increases to* 60 kW*?*

Its unknown is the temperature T_2 of the scenario node at the full load of 60 kW, and over the state of the graph it is

$$\Pr(T_2 < 27.5^\circ\text{C} \mid \mathbf{y}), \tag{1.25}$$

with the unknown, the full-load temperature, on the left of the bar and the evidence, the reading of the test condition and the handbook value, on the right; the graph and its pinned inputs are conditions of the model throughout.

General form. A prediction under a threshold is a tail probability of a derived quantity under the posterior, the posterior mass of the states that satisfy the event,

$$\Pr(f(\mathbf{z}) < f^* \mid \mathbf{y}) = \int \mathbf{1}\big[f(\mathbf{z}) < f^*\big]\, p(\mathbf{z} \mid \mathbf{y})\,\mathrm{d}\mathbf{z}. \tag{1.26}$$

Specialized form. The two questions are equation (1.26) with the state $\mathbf{z} = (T, g)^{\mathsf{T}}$ of the hall, $p(\mathbf{z} \mid \mathbf{y})$ the posterior of step 3, the Laplace normal $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of equations (1.20) and (1.21), and the derived quantity $f(\mathbf{z})$ and the threshold f^* given for each below: $f(\mathbf{z}) = T$ with $f^* = 24°\mathrm{C}$, and $f(\mathbf{z}) = T_2$ with $f^* = T_{\max} = 27.5°\mathrm{C}$. Two choices enter here. The full-load condition is taken to share the conductance of the test condition, with the load and the ambient pinned at both: this is what makes a condition the hall has never run predictable from a reading at 30 kW, and it fails if the conductance changes with the load, for instance through an airflow that rises with it (Section 1.1.4), which no check of this chapter can test because nothing has been observed at 60 kW. The full-load answer is the steady temperature at 60 kW: this is what lets the event be written on the conductance alone, and it says nothing about the period just after the load rises, while the hall's stored heat still matters (Section 1.6.5).

The present load. Here $f(\mathbf{z}) = z_1 = T$ and $f^* = 24°\mathrm{C}$. Under the Laplace posterior the tail is a normal tail of the first entry,

$$\Pr(T < 24°\mathrm{C} \mid \mathbf{y}) \approx \Phi\!\left(\frac{24 - \hat{z}_1}{\sqrt{\Sigma_{11}}}\right), \tag{1.27}$$

where Φ is the standard normal cumulative distribution function. The value beside it comes from the exact posterior: on the curve the event $T < 24°\mathrm{C}$ is the event $T(g) < 24°\mathrm{C}$, that is $g > -\ln\!\big((24 - T_a)\, G_{\mathrm{ref}}/Q\big) = -\ln 1.2$, and its probability is the integral of equation (1.19) above that value.

The full load. Here $f(\mathbf{z}) = T_2$, the temperature of the scenario node. By equation (1.8) it is a function of the second entry alone, $T_2 = T_a + (Q_2/G_{\mathrm{ref}})\, \mathrm{e}^{-z_2}$, and on the zero set of the physics row it is also a function of the first: dividing equation (1.8) by equation (1.7) removes the conductance,

$$\frac{T_2 - T_a}{T - T_a} = \frac{Q_2}{Q_1} \qquad \Longrightarrow \qquad T_2 = T_a + \frac{Q_2}{Q_1}\,(T - T_a). \tag{1.28}$$

At twice the load the excess over the ambient is twice as large, whatever the conductance: a test temperature of 23°C predicts $T_2 = 18 + 2 \times 5 = 28°\mathrm{C}$, and a test temperature of 23.2°C predicts 28.4°C. Measurements taken during the first operating condition therefore inform the second, although the hall has never been operated at full load; this is how an engineering model links evidence collected today to predictions under operating conditions that have not yet occurred. In the vector form the link is the tangent $\mathbf{n}$ of equation (1.17): a step along the curve moves T by $(Q_1/G_{\mathrm{ref}})\, \mathrm{e}^{-g}$ and T_2 by $(Q_2/G_{\mathrm{ref}})\, \mathrm{e}^{-g}$ per unit of $-g$, in the ratio Q_2/Q_1.

The event itself needs no posterior to be rewritten. The full-load temperature $T_2 = T_a + Q_2/G$ decreases in G, so for a limit $T_{\max} > T_a$ the event $T_2 < T_{\max}$ is, in turn, an event on the conductance, on the conductance entry and, by equation (1.28), on the test temperature:

$$T_2 < T_{\max} \iff G > G^* \iff g > g^* = \ln \frac{G^*}{G_{\mathrm{ref}}} \iff T < T^*, \tag{1.29}$$

with the thresholds

$$G^* = \frac{Q_2}{T_{\max} - T_a}, \tag{1.30}$$

$$T^* = T_a + \frac{Q_1}{Q_2}\,(T_{\max} - T_a). \tag{1.31}$$

For $T_{\max} \le T_a$ the probability is zero, since $T_2 > T_a$ for every positive conductance. The inequalities of equation (1.29) describe the same set of states on the physics curve, so their probabilities are equal under any distribution on the curve; the equivalence holds under the assumed steady-state model, with known load and ambient temperature. The question about a condition that has never been operated has become a question about the unknown conductance, an entry of the state that the test condition informs, before any data are processed.

Figure 1.6 draws the equivalence (1.29) on the forward relation (1.2). The horizontal axis is the conductance G in kW/K and the vertical axis the hall temperature in °C. The solid blue curve is the hall at the load of 30 kW, the dashed green curve the same hall at 60 kW; the dot on the blue curve is the forward calculation of Section 1.1.4, $T = 23$°C at $G = 6$ kW/K, and the dot above it the same conductance at 60 kW, 28°C. The red line is the limit of 27.5°C, and the dotted vertical line marks the threshold G^* of equation (1.30), where the dashed curve crosses the limit. The dashed green curve lies below the red line exactly to the right of G^*: because the curve is strictly decreasing, the event "below the limit" on the vertical axis is the event "above G^*" on the horizontal axis.

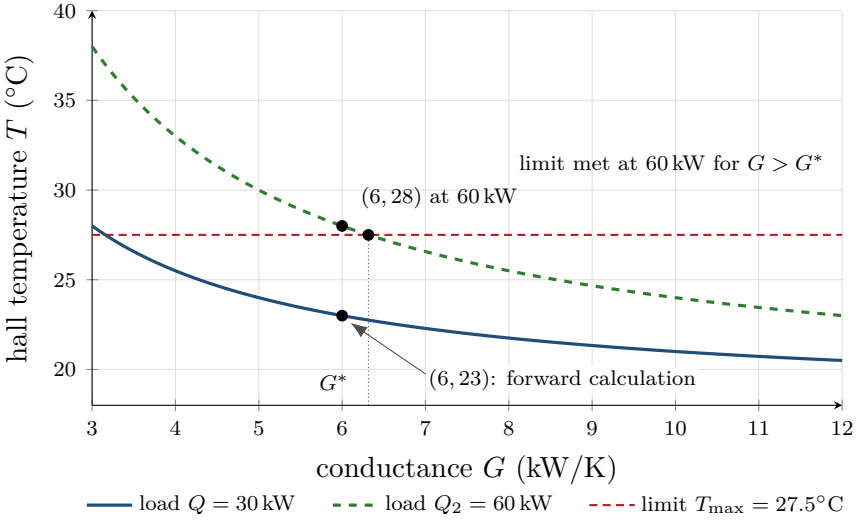


Figure 1.6: The forward relation $T = T_a + Q/G$ (equation (1.2)) with $T_a = 18$°C, for the loads $Q = 30$ kW and $Q_2 = 60$ kW. Dots: the forward calculation $T = 23$°C at $G = 6$ kW/K; the same conductance at 60 kW, 28°C; and the threshold $G^* = Q_2/(T_{\max} - T_a)$ of equation (1.30) at which the 60 kW hall reaches the limit of 27.5°C.

Under the Laplace posterior the probability is a normal tail of the second entry, the entry onto which the physics maps the event,

$$\Pr(T_2 < T_{\max} \mid \mathbf{y}) = \Pr(g > g^* \mid \mathbf{y}) \approx \Phi\left(\frac{\hat{z}_2 - g^*}{\sqrt{\Sigma_{22}}}\right). \tag{1.32}$$

The same event written on the first entry, $T < T^*$, has the same exact probability, but under the Laplace normal, whose ridge is the straight tangent rather than the curve, its tail differs slightly; the chapter reports the tail on g, the entry on which the full-load temperature depends alone. The value beside it comes from the exact posterior: the probability is the integral of equation (1.19) above g^*.

Principle 1.3. *When a quantity in a question is a monotone function of one entry of the state, a threshold on the quantity is a threshold on that entry, and the two events have the same probability under every posterior. The transformation uses the physics only, and it can be made before any data are processed.*

Building blocks. The entries $\hat{z}_1 = 23.070$°C and $\hat{z}_2 = -0.0139$ of the MAP state and the standard deviations $\sqrt{\Sigma_{11}} = 0.481$ K and $\sqrt{\Sigma_{22}} = 0.0948$ of the covariance (equations (1.20) and (1.21)); the threshold of equation (1.30) for the owner's $T_{\max} = 27.5$°C and $Q_2 = 60$ kW, and for other limits and loads, with its equivalent on the test temperature, equation (1.31); the normal cumulative distribution Φ; and, for the exact values beside them, the density (1.19) integrated over g above each threshold.

The hall below 24 degrees. By equation (1.27),

$$\Pr(T < 24^\circ\mathrm{C} \mid \mathbf{y}) \approx \Phi\left(\frac{24 - 23.070}{0.481}\right) = \Phi(1.93) = 0.973.$$

The value beside it comes from the exact posterior: the mass of equation (1.19) above $g = -\ln 1.2 = -0.1823$ is 0.977, 0.004 higher.

The hall at 60 kilowatts. For the owner's limit $T_{\max} = 27.5^\circ\mathrm{C}$ at $Q_2 = 60\,\mathrm{kW}$, equations (1.30) and (1.31) give

$$G^* = \frac{60}{27.5 - 18} = \frac{60}{9.5} = 6.316\,\mathrm{kW/K}, \qquad g^* = \ln\frac{6.316}{6} = 0.0513,$$

$$T^* = 18 + \tfrac{1}{2} \times 9.5 = 22.75^\circ\mathrm{C}:$$

the temperature at 60 kW stays below 27.5°C if and only if the temperature at 30 kW is below 22.75°C; on the curve of Figure 1.5, the states to the right of the red line are those below 22.75°C. By equation (1.32),

$$\Pr(T_2 < 27.5^\circ\mathrm{C} \mid \mathbf{y}) \approx \Phi\left(\frac{-0.0139 - 0.0513}{0.0948}\right) = \Phi(-0.688) = 0.246,$$

with the MAP conductance 5.917 kW/K to the left of G^* in Figure 1.6. The value beside it comes from the exact posterior: the mass of equation (1.19) above $g^* = 0.0513$ is 0.281, 0.036 higher. The event lies in the tail of g toward large conductances, the side on which the exact posterior extends further than the normal at $\hat{\mathbf{z}}$, so the state route understates it. The gap does not change the decision: on both routes the hall meets the limit at 60 kW with probability of about one in four.

Table 1.5 evaluates the same blocks for other limits and loads. Each row is a limit $T_{\max}$ and a future load Q_2; the columns give the threshold conductance G^*, the threshold log-conductance g^*, the threshold on the test temperature T^*, the probability of equation (1.32), which the chapter reports, and beside it the exact value, the mass of equation (1.19) above g^*. One row of the arithmetic, for $T_{\max} = 27^\circ\mathrm{C}$ and $Q_2 = 60\,\mathrm{kW}$: $G^* = 60/9 = 6.667\,\mathrm{kW/K}$, $g^* = \ln(6.667/6) = 0.1054$, $T^* = 18 + 0.5 \times 9 = 22.5^\circ\mathrm{C}$, and $\Phi((-0.0139 - 0.1054)/0.0948) = \Phi(-1.258) = 0.104$.

Table 1.5: Equivalent thresholds and the full-load probability, with $T_a = 18^\circ\mathrm{C}$, $Q_1 = 30\,\mathrm{kW}$ and $G_{\mathrm{ref}} = 6\,\mathrm{kW/K}$ (equations (1.30) and (1.31)); the probability by the Laplace tail (1.32), with $\hat{g} = -0.0139$ and $\sqrt{\Sigma_{22}} = 0.0948$, and exactly, by quadrature of equation (1.19). The row 27.5°C at 60 kW is the owner's question; the others are added to show the trend.

Limit $T_{\max}$ (°C)	Load Q_2 (kW)	G^* (kW/K)	g^* [–]	T^* (°C)	Laplace [–]	Exact [–]
26	60	7.500	0.2231	22.00	0.0062	0.0155
27	60	6.667	0.1054	22.50	0.1042	0.1354
27.5	60	6.316	0.0513	22.75	0.2458	0.2815
29	60	5.455	−0.0953	23.50	0.8047	0.8355
27.5	45	4.737	−0.2364	24.33	0.9905	0.9964
27.5	75	7.895	0.2744	21.80	0.0012	0.0049

A stricter limit, or a larger future load, demands a larger conductance and a cooler test temperature; the first five columns follow from the physics alone. The Laplace tail is below the exact one in every row, and the gap grows in ratio toward the far tail: a factor of 2.5 to 4 in the two smallest rows.

1.4 The checks (step 5)

The fit rests on two approximations and one assumption: a physics row with a finite tolerance in place of an exact law, a normal distribution in place of the posterior along a curved ridge, and a reading and a handbook value that are compatible with each other. Each is set against a full calculation, and the weight of the handbook value is measured by removing it.

1.4.1 The exact-physics limit

Table 1.6 repeats the computation at three model tolerances. Each row is one σ_p; the columns are the two entries of $\hat{\mathbf{z}}$, their Laplace standard deviations and their correlation. From 0.1 to 0.01 K the conductance entry moves by 0.001, its spread by 0.002 and the correlation from -0.979 to -0.9998; from 0.01 to 0.001 K nothing moves at the printed precision but the last digit of the correlation. The limit is reached at $\sigma_p = 0.01\,\mathrm{K}$, and the correlation tends to -1 as equation (1.18) requires.

Table 1.6: The posterior of the equipment hall at three model tolerances σ_p of the physics row: MAP state, Laplace standard deviations (with the curvature term) and correlation. The exact-physics limit $\sigma_p \to 0$ is reached at 0.01 K.

σ_p (K)	$\hat{T}$ (°C)	$\hat{g}$ [–]	$\sqrt{\Sigma_{11}}$ (K)	$\sqrt{\Sigma_{22}}$ [–]	correlation [–]
0.1	23.0705	−0.0130	0.4806	0.0966	−0.9790
0.01	23.0701	−0.0139	0.4806	0.0948	−0.9998
0.001	23.0701	−0.0139	0.4806	0.0948	−1.0000

1.4.2 The Laplace posterior against the exact posterior

In the limit $\sigma_p \to 0$ the exact posterior is the density (1.19) along the physics curve, and one-dimensional quadrature of it gives every quantity of the questions; the event $T < 24°\mathrm{C}$ is the event $g > -\ln 1.2 = -0.1823$ on the curve, and the full-load event is $g > g^* = 0.0513$. Table 1.7 collects the pairs that step 4 computed, the Laplace value it reports beside the exact one. Each row is one quantity of the posterior, with its unit; the third column is its Laplace value, read from $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$ of equations (1.20) and (1.21), and the fourth its exact value from the quadrature. For the location of each quantity the Laplace column gives the value at $\hat{\mathbf{z}}$ and the exact column the posterior mean.

Table 1.7: The Laplace approximation of the posterior of the equipment hall against the exact posterior along the physics curve, in the limit $\sigma_p \to 0$.

Quantity	Unit	Laplace, at $\hat{\mathbf{z}}$	Exact, along the curve
Location of T	°C	23.070	23.030
Standard deviation of T	K	0.4806	0.4813
$\Pr(T < 24°\mathrm{C} \mid \mathbf{y})$	–	0.973	0.977
Location of g	–	−0.0139	−0.0013
Standard deviation of g	–	0.0948	0.0967
Location of G	kW/K	5.917	6.020
Standard deviation of G	kW/K	0.561	0.591
$\Pr(T_2 < 27.5°\mathrm{C} \mid \mathbf{y})$ at 60 kW	–	0.246	0.281

For the temperature the approximation is adequate: the tail probability differs by 0.004 and the standard deviation by 0.0007 K. The posterior is slightly skewed along the curve, so its

mean and its mode differ: the mean of g lies 0.013 above $\hat{g}$, the mean conductance 0.10 kW/K above $\hat{G}$, and the exact standard deviation of G is 0.591 kW/K, 5 % wider than the Laplace 0.561 kW/K. The skew matters most for the full-load question, whose event lies in the tail of g toward large conductances, the side on which the exact posterior extends further: the Laplace tail 0.246 is low by 0.036 against the exact 0.281. Further into the tail the understatement grows: in Table 1.5 the Laplace tail on g is smaller than the exact one by a factor of 2.5 to 4 in the two smallest rows, 0.0062 against 0.0155 and 0.0012 against 0.0049. Written on the first entry instead, as $T < T^* = 22.75°\mathrm{C}$, the full-load event has the Laplace tail 0.253, because the normal follows the straight tangent rather than the curve; the exact value is 0.281 on either entry. The exact values rest on the base of equation (1.19), flat in g along the curve; with a base flat in T along the curve instead, the full-load value is 0.249, a reminder that an exact tail depends on the base measure as well as on the rows. The conductance band and the full-load probability are therefore slightly narrow and slightly low, and the report says so.

A scalar check. Eliminating T through the physics row and differentiating the remaining one-entry objective in g twice gives the same $\sqrt{\Sigma_{22}} = 0.0948$ and, through the slope of the curve, the same $\sqrt{\Sigma_{11}} = 0.481\,\mathrm{K}$. The vector form reaches them without the substitution, and it is the form that carries over unchanged to the cooling network of Section 1.1.3.

1.4.3 The residual check

Three rows act on two entries, and with the physics exact the reading and the handbook both inform the one direction along the curve: one row is redundant, and the rows at the MAP state measure how well the two agree. They are $r_1 = 0.005$, $r_2 = -0.260$ and $r_3 = -0.790$ (Table 1.4), and their squares sum to $2\mathcal{L}(\hat{\mathbf{z}}) = 0.692$. When the model and the standard deviations of the rows are right, that sum follows a chi-square distribution with as many degrees of freedom as there are rows less entries, here one; Chapter 22, in Volume II, derives the result. A sum of 0.692 or more then has probability 0.41, while a sum above 3.84 would have probability below 0.05 and would say that the reading or the handbook value is wrong. The two agree, which is a first, coarse look at the question of a biased sensor.

1.4.4 The weight of the handbook prior

The handbook value is the one input of the fit that the hall did not supply, and removing its row shows how much the answers lean on it. Without the prior row the residual vector holds the rows r_1 and r_2 of equation (1.10): two rows, two entries, a square Jacobian, both rows zero at the MAP state ($23.2°\mathrm{C},\ -0.0392$), and $\Sigma_{11} = \sigma_T^2$ whatever σ_p, and $\sqrt{\Sigma_{22}} = 0.0962$ (the last exercise of the chapter derives both). On the conductance entry, the route of equation (1.32), the full-load probability is then

$$\Pr(T_2 < 27.5°\mathrm{C} \mid y_T) \approx \Phi\!\left(\frac{-0.0392 - 0.0513}{0.0962}\right) = \Phi(-0.941) = 0.173,$$

against 0.246 with the handbook row, a rise of 0.073. Exactly, with the same base flat in g along the curve and the prior term removed from equation (1.19), it is 0.212 against 0.281, a rise of 0.069. Figure 1.7 draws the two Laplace marginals of the conductance entry. The horizontal axis is g and the vertical axis the probability density; the dashed curve is the marginal under the reading alone, $\mathcal{N}(-0.0392,\ 0.0962^2)$, and the solid curve the marginal with the handbook row, $\mathcal{N}(-0.0139,\ 0.0948^2)$. The dotted red line is the threshold $g^* = 0.0513$, and the shaded areas to its right are the two probabilities. The handbook value, which places the conductance above what the reading alone implies, moves the marginal toward the threshold with nearly the same width.

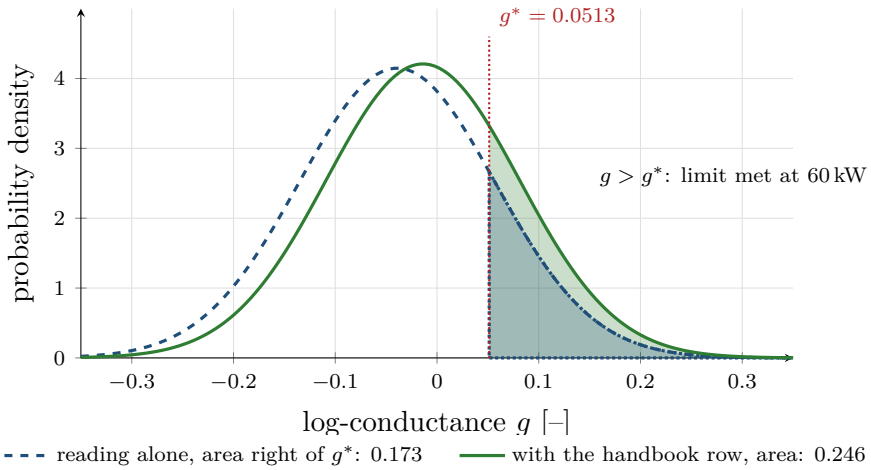


Figure 1.7: The weight of the handbook prior in the full-load answer: Laplace marginals of the conductance entry, $\mathcal{N}(-0.0392,\ 0.0962^2)$ under the reading alone (dashed) and $\mathcal{N}(-0.0139,\ 0.0948^2)$ with the handbook row (solid), densities $\varphi((g-\hat{g})/s)/s$ with φ the standard normal density. The shaded areas right of $g^* = 0.0513$ are the full-load probabilities, 0.173 and 0.246.

The same holds in every row of Table 1.5, on both routes: without the handbook row each probability is lower than with it. The full-load answer therefore depends on the handbook value, and the handbook value is stated with the report.

1.5 The report (step 6)

The record does not support the handbook's answer. Carried forward, the handbook conductance of 7.5 kW/K put the hall at 26°C at 60 kW, well below the limit of 27.5°C; held against the one reading through the energy balance, it gives way to a conductance of 5.92 kW/K, below the 6.316 kW/K the full load needs, and the hall stays below 27.5°C at 60 kW with probability of about one in four, 0.246 by the Laplace route and 0.281 exactly: neither the comfortable pass of the handbook nor the certain failure of the reading taken alone. A heat-flow meter on the cooling path, the instrument that seems to measure the very heat in question, would not change that answer (Section 1.6.2). The answers to the first four questions of Table 1.2 are given below in the units in which they were asked, and the owner's full-load decision follows from the fourth.

1.5.1 The answers and their bands

Table 1.8 gives the answers. Each row is one question; the columns give the answer computed from $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$, which the chapter reports, the exact value beside it from the density along the curve, and the band (one standard deviation for the temperature and the conductance, Laplace and exact; for the full-load probability, its value without the handbook row, the check of Section 1.4.4).

The reading has placed the hall temperature to within about half a kelvin, slightly better than the sensor's own 0.5 K because the handbook value adds a little through the physics row. The conductance is known to within 9.5 % of its value, against a factor $e^{0.3} = 1.35$ in the handbook before the reading; the exact posterior widens that band to 0.59 kW/K. The hall is below 24°C with probability of about 0.97. The fifth question, whether another sensor would narrow these answers, is answered in Section 1.6.

Table 1.8: The report on the equipment hall from the reading $y_T = 23.2°\mathrm{C}$ and the handbook conductance of 7.5 kW/K, with the physics exact.

Question	Answer (state route)	Exact	Band
What is the hall temperature now?	$\hat{T} = 23.07°\mathrm{C}$	mean 23.03°C	±0.48 K (0.481 K exact)
What is the effective conductance of the path?	$\hat{G} = 5.92\,\mathrm{kW/K}$	mean 6.02 kW/K	±0.56 kW/K (0.59 kW/K exact)
Will the hall stay below 24°C?	$\Pr(T < 24°\mathrm{C} \mid \mathbf{y}) =$ 0.973, about 0.97	0.977; gap 0.004	–
Will the hall stay below 27.5°C at 60 kW?	$\Pr(T_2 < 27.5°\mathrm{C} \mid \mathbf{y}) =$ 0.246, about 0.25	0.281; gap 0.036	without the handbook value 0.173 (0.212 exact)

1.5.2 The full-load decision

The owner's question has the answer of the fourth row: at 60 kW the hall stays below 27.5°C with probability of about one in four, 0.246 by the Laplace route and 0.281 exactly. The present cooling path does not support the full load. The reason is read from the equivalent threshold: the limit is met only if the conductance exceeds $G^* = 6.316\,\mathrm{kW/K}$, and the posterior places the conductance at $5.92 \pm 0.56\,\mathrm{kW/K}$; in the conductance entry the threshold $g^* = 0.0513$ lies 0.69 standard deviations above $\hat{g} = -0.0139$. The answer leans on the handbook value: without it the probability would be 0.173 by the Laplace route and 0.212 exactly. A measurement cannot raise the conductance; it can only show where it lies, and with the estimate below the threshold, further evidence is more likely to confirm the shortfall than to remove it. Running the hall at 60 kW under this limit calls for more cooling, a larger conductance, rather than for more measurement.

1.6 The next measurement (step 7)

Whether the bands of Table 1.8 are narrow enough is the engineer's decision. If they are not, the fifth question of Table 1.2 asks which measurement would narrow them most, and a measurement not yet taken is a row not yet attached.

1.6.1 The preposterior covariance

The value a candidate measurement will return is unknown, but its Jacobian row at $\hat{\mathbf{z}}$, its sensitivity to each entry divided by its standard deviation, is known before the measurement is made, and that is enough. The information matrix is a sum of one outer product per row (Section 1.2.1), so a candidate row with the Jacobian row $\mathbf{J}_a$ adds $\mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a$ to the inverse of the present covariance, and the covariance expected after the measurement is

$$\boldsymbol{\Sigma}_a = \left(\boldsymbol{\Sigma}^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a\right)^{-1}, \tag{1.33}$$

with $\boldsymbol{\Sigma}$ of equation (1.21), curvature term included. The value the measurement will return does not appear in it. Chapter 2 derives $\boldsymbol{\Sigma}_a$ as the preposterior covariance, and Chapter 20, in Volume II, ranks candidate rows by it for quantities derived from the state. The band on the temperature after the measurement is $\sqrt{\Sigma_{a,11}}$, and the band on the conductance, the quantity on which the full-load answer rests, is $\sigma_{G,a} = \sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_a\mathbf{g}}$ with the gradient of equation (1.24).

1.6.2 The candidate measurements

Four candidates are considered, each written as a row over $\mathbf{z} = (T, g)^{\mathsf{T}}$:

- a second hall sensor of the same accuracy, with an error independent of the first: the row $(T - y'_T)/\sigma_T$, with the Jacobian row $\mathbf{J}_a = (2,\ 0)$, the same as the present reading row;
- a reference sensor of 0.2 K in the hall: $\mathbf{J}_a = (5,\ 0)$;
- a reading of the ambient temperature, or a meter on the heat load: T_a and Q are pinned, not entries of the state, so the row has no entry in either column, $\mathbf{J}_a = (0,\ 0)$;
- a heat-flow meter on the cooling path, with a standard deviation $\sigma_q = 1\,\mathrm{kW}$, reading the heat removed $G\,(T - T_a)$: $\mathbf{J}_a = \big(\hat{G},\ \hat{G}\,(\hat{T} - T_a)\big)/\sigma_q = (5.917,\ 30.00)$.

Table 1.9 evaluates equation (1.33) for each. Each row is one candidate; the columns give its Jacobian row, the standard deviation of T after it, $\sqrt{\Sigma_{a,11}}$, and the width it removes from the present 0.481 K, then the conductance band after it, $\sigma_{G,a}$ with the gradient $\mathbf{g} = (0,\ 5.917)^{\mathsf{T}}$, and the width it removes from the present 0.561 kW/K. One row of the arithmetic, for the second sensor: $\mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a$ has the single entry 4, in its T–T position, and

$$\boldsymbol{\Sigma}_a = \begin{bmatrix} 0.12005 & -0.023678 \\ -0.023678 & 0.0046738 \end{bmatrix}, \qquad \sqrt{0.12005} = 0.346\,\mathrm{K}, \quad 5.917\,\sqrt{0.0046738} = 0.405\,\mathrm{kW/K}.$$

Table 1.9: Candidate measurements for the equipment hall ranked by the preposterior covariance of equation (1.33), evaluated at $\hat{\mathbf{z}}$. Present bands: 0.481 K on T and 0.561 kW/K on G.

Candidate	$\mathbf{J}_a$ (K^{-1}, [–])	$\sqrt{\Sigma_{a,11}}$ (K)	removed (K)	$\sigma_{G,a}$ (kW/K)	removed (kW/K)
Second hall sensor, 0.5 K	(2, 0)	0.346	0.134	0.405	0.156
Reference hall sensor, 0.2 K	(5, 0)	0.185	0.296	0.216	0.345
Ambient reading or load meter	(0, 0)	0.481	0	0.561	0
Path heat-flow meter, 1 kW	(5.917, 30.00)	0.481	$< 10^{-6}$	0.561	$< 10^{-6}$

Only the two sensors in the hall remove anything, and the reason is visible in their Jacobian rows. The posterior is wide in one direction only, along the physics curve, $(0.9811,\ -0.1935)^{\mathsf{T}}$, where the present information is 4.17: the 4.27 of $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ less the 0.099 of the curvature term. The second sensor's row projects on that direction as $2 \times 0.9811 = 1.962$ and adds $1.962^2 = 3.85$, nearly doubling the information, so that the variances along the curve fall by the factor $4.17/8.02 = 0.52$ and both standard deviations by the factor 0.72. The ambient reading and the load meter read pinned quantities: in this model they are known already, their rows are zero, and they would inform only a graph in which T_a or Q were free. The heat-flow meter is the subtler case. Its row (5.917, 30.00) is proportional to the physics row (100, 507.0), both in the ratio 5.070, so it points across the curve, where the physics already fixes the state to within σ_p; it measures again what the energy balance asserts, and it adds nothing along the curve. A measurement is worth what it adds in the direction in which the posterior is wide.

1.6.3 The recommendation

The recommendation, and the answer to the fifth question, is a second sensor in the hall: it removes 0.134 K of the 0.481 K on the temperature and 0.156 kW/K of the 0.561 kW/K on the conductance, 28 % of each, and a 0.2 K reference sensor, if one is available, removes 62 %. A reading of the ambient, a load meter and a heat-flow meter on the path are not worth installing

for these answers. Two conditions apply. The second sensor's error must be independent of the first: a second reading from the same sensor at the same condition would share any bias of the first, and a shared error is a discrepancy that repeats (Section 1.1.4), which a repeated reading does not reduce. And the gain is computed at $\hat{\mathbf{z}}$ under the Laplace approximation, which the checks found adequate for T and slightly narrow for G. Once the measurement is taken, its row joins the three of equation (1.10), and the method returns to step 3 with four rows. For the full-load decision the narrower conductance band sharpens the answer but does not move the threshold: the decision of Section 1.5.2 turns on the cooling, not on the sensor.

1.6.4 What the analysis showed

The handbook value, carried forward, put the hall at 26°C at 60 kW, a comfortable pass, and the reading, taken as the truth, at 28.4°C, a plain failure; the analysis, which holds the two against each other through the energy balance, finds neither, but a probability of about one in four that the hall stays below 27.5°C, 0.246 from the posterior of the state (0.281 exact). It also overturns the obvious instrument: a heat-flow meter on the path, although it measures the very heat the question is about, removes nothing, because it repeats what the energy balance already asserts. One reading and a handbook value pin the hall temperature to about half a degree but the conductance of its cooling path only to about ten percent, and the two are not known separately: the posterior lives on the physics curve, a ridge along which a warmer hall is a hall with a weaker path, which is why the correlation of the two entries is -0.9998. The owner's question about the full load is a question about that conductance. The hall meets 27.5°C at 60 kW only if its path conducts more than 6.316 kW/K; the evidence places it below that, and the answer leans on the handbook value. What would narrow the answers is a second sensor in the hall, which removes 28 % of both bands. What would change the full-load answer is more cooling, not more measurement.

1.6.5 Limits of the model

The posterior is exact only under the model, the steady-state graph of one hall with one path, and that model can be wrong before any inference is run.

Modeling mistakes. Several errors can be made before any inference algorithm is run, in the construction of the graph and its rows. Table 1.10 lists five. Each row names the mistake, gives its form in the equipment hall, and states the correction.

The first two mistakes confuse the roles of Table 1.3. A sensor reporting 23.2°C does not establish that the true temperature is exactly 23.2°C: the observation has an uncertainty, and the physics row of Section 1.1.4 showed that 23.2°C and 23°C imply conductances of 5.77 and 6 kW/K, a difference of 4 %. The conductance is a physical property of the path; the reading is evidence about the temperature, which is connected to the conductance through the model. The fourth mistake mixes two kinds of uncertainty that require different treatments. Uncertainty about the numerical value of G within the model of Figure 1.2 is represented by the posterior of g. Uncertainty about whether the hall has one or two effective cooling paths is uncertainty about the model M itself, and it is represented by comparing the evidence of the two drawings. The fifth mistake is the converse of the first. A graph provides relations; the posterior also requires evidence, and some entries may remain poorly identified even when the model is perfectly specified. In the two-path drawing, for example, a reading of the hall temperature determines only the total conductance of the two paths and leaves their split to the prior: the Jacobian of the physics and reading rows has a null vector along the split, and only a prior row adds information in that direction.

Table 1.10: Five modeling mistakes made before any inference is run.

Mistake	In the equipment hall	Correction
Treating a measurement as the true state	the sensor reports 23.2°C, and $T = 23.2$°C is assumed	the reading is evidence with error σ_T; T stays a free entry of the state
Confusing parameters and observations	the conductance and the reading are treated alike	G is a physical property, an entry of the state; y_T is evidence about T, connected to G through the physics row
Assuming the physical model is complete	the steady-state relation is used just after a cooling failure	stored heat matters in a transient, $C_{\mathrm{th}}\,\mathrm{d}T/\mathrm{d}t = Q - G\,(T - T_a)$
Confusing structural with parameter uncertainty	doubt whether the hall has one or two effective cooling paths is folded into the spread of G	the value of G is a parameter question within one model; the number of paths is a question between two models
Believing the graph alone determines the posterior	a correct graph is expected to give a narrow answer	useful inference also needs suitable evidence, assumptions and constraints; some entries may remain poorly identified

The steady-state assumption. The third mistake concerns the physics itself. The model of equation (1.1) assumes steady conditions and a single effective conductance. During transient operation the hall stores and releases heat, and the energy balance gains a storage term,

$$C_{\mathrm{th}}\,\frac{\mathrm{d}T}{\mathrm{d}t} = Q - G\,(T - T_a),$$

where C_{th} is the thermal capacitance of the hall, its air, structure and equipment, in kJ/K, and t is time in seconds. For constant Q, G and T_a, write the departure from the steady temperature as $u = T - T_a - Q/G$. Then $\mathrm{d}u/\mathrm{d}t = \mathrm{d}T/\mathrm{d}t$ and the balance becomes

$$C_{\mathrm{th}}\,\frac{\mathrm{d}u}{\mathrm{d}t} = -G\,u,$$

whose solution is $u(t) = u(0)\,\mathrm{e}^{-t/\tau}$ with the time constant $\tau = C_{\mathrm{th}}/G$. Therefore

$$T(t) = T_a + \frac{Q}{G} + \Big(T(0) - T_a - \frac{Q}{G}\Big)\,\mathrm{e}^{-t/\tau}, \qquad \tau = \frac{C_{\mathrm{th}}}{G}. \tag{1.34}$$

The steady-state relation of equation (1.2) is the limit $t \to \infty$. Suppose (illustratively) that at $t = 0$ half of the cooling fails, so that the conductance falls from 6 to 3 kW/K while the hall is at 23°C under 30 kW. The new steady temperature is $18 + 30/3 = 28$°C, and equation (1.34) gives $T(t) = 28 - 5\,\mathrm{e}^{-t/\tau}$: after one time constant the hall is at $28 - 5\,\mathrm{e}^{-1} = 26.16$°C. A steady-state model would place it at 28°C at once. It is therefore inadequate for predicting temperatures immediately after a cooling failure, where the time constant decides whether the operators have minutes or seconds; the graph for the engineer's sixth question, the partial cooling failure, needs the storage term and C_{th} as a further entry or a pinned input.

The chapter has built the representation on which every question will be asked. The next chapter, Chapter 2, turns to the questions themselves: how an engineering question stated in ordinary language is recognized as estimation, prediction, comparison, diagnosis or design, and written as a precise probability or expectation over the state of such a graph.

1.7 What to retain

Definition 1.1 (Plant graph)**.** A plant graph represents a physical system as free and pinned nodes, paths and sources, with the physical relations that must hold among them written as

residual rows; its free quantities, physical states and parameters alike, form the state vector $\mathbf{z}$; readings, records and prior knowledge are attached as further rows; and every engineering question is then a probability or an expectation over $\mathbf{z}$, conditional on the evidence $\mathbf{y}$, computed from the posterior $p(\mathbf{z} \mid \mathbf{y})$ that the rows define.

The forward calculation, the posterior and the question are three different things computed over the same graph. The forward calculation uses no evidence and gives a number; the posterior combines the physics, the readings and the prior; the question picks what is read from the posterior. The posterior is computed in state-vector form: the rows, each divided by its standard deviation, form one residual vector over $\mathbf{z}$; the normal equations $(\mathbf{J}^\mathsf{T}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^\mathsf{T}\mathbf{r}$ give the Gauss–Newton steps to the MAP state $\hat{\mathbf{z}}$; the information matrix $\mathbf{J}^\mathsf{T}\mathbf{J}$ is a sum of one outer product per row; the covariance $\boldsymbol{\Sigma}$ is the inverse of $\mathbf{J}^\mathsf{T}\mathbf{J}$ plus the rows' own curvature, here the curvature of the physics curve; and an answer is an entry of $\hat{\mathbf{z}}$ with its $\sqrt{\Sigma_{ii}}$, or a derived quantity with $\sqrt{\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\,\mathbf{g}}$. An exact physical law is the limit $\sigma_p \to 0$ of its row, and its signature in $\boldsymbol{\Sigma}$ is a correlation near -1: a ridge along the zero set of the row, in the direction of the null vector of the row's Jacobian. A question about a quantity that is a monotone function of one entry of the state is transformed into a threshold on that entry by the physics alone, as the full-load question became $G > 6.316\,\mathrm{kW/K}$, and its probability is a normal tail of that entry's marginal. The answer is computed from $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$, with the exact value along the curve set beside it and the gap stated, checked against the full calculation (the limit $\sigma_p \to 0$, the exact posterior along the curve, the agreement of the redundant rows, the weight of the prior), reported with its band in the units of the question, and followed by the next measurement, ranked by the preposterior covariance $\boldsymbol{\Sigma}_a = (\boldsymbol{\Sigma}^{-1} + \mathbf{J}_a^\mathsf{T}\mathbf{J}_a)^{-1}$ of a row not yet taken; a reading of a pinned quantity, or a row parallel to a physics row, removes nothing. Table 1.11 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 1.11: The terms of the plant graph and where each is defined. Units for the equipment hall.

Term	Expression	Units	Meaning	Defined in
Plant graph	nodes, paths, sources, rows	–	the system before any question	1.1.4
Free, pinned	T, g free; T_a, Q pinned	–	inferred against known	1.1.4
Energy balance	$Q = G\,(T - T_a)$	kW	conservation at the hall node	1.1.4
Forward calculation	$T = T_a + Q/G$	°C	temperature from known inputs	1.1.4
State vector	$\mathbf{z} = (T, g)^{\mathsf{T}}$	°C, –	all free quantities at once	1.1.4
Log-conductance	$g = \ln(G/G_{\mathrm{ref}})$	–	keeps G positive	1.1.4
Physics row	$(T - T_a - (Q/G_{\mathrm{ref}})\mathrm{e}^{-g})/\sigma_p$	–	a relation that must hold, tolerance σ_p	1.1.4
Model discrepancy	$r_{\mathrm{phys}}(\mathbf{z}) = \delta$	kW	model error, not sensor error	1.1.4
Reading row	$(T - y_T)/\sigma_T$	–	evidence about the state	1.1.7
Prior row	$(g - g_0)/\sigma_g$	–	knowledge before the readings	1.1.7
Jacobian pattern	which row acts on which entry	–	a property of the graph	1.1.7
Scenario node	$T_2 = T_a + (Q_2/G_{\mathrm{ref}})\mathrm{e}^{-g}$	°C	an unobserved condition sharing g	1.1.7
Factor graph	entries and rows	–	the graph of the inference	1.1.3
Delta-method band	$\sigma_G^2 = \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$, $\mathbf{g} = (0, \hat{G})^{\mathsf{T}}$	kW/K	spread of a derived quantity	1.3.1
Ridge	$\mathbf{J}_1\mathbf{n} = 0$, $\boldsymbol{\Sigma} \to \mathbf{u}\mathbf{u}^{\mathsf{T}}/\lambda_{\parallel}$, $\lambda_{\parallel} = (n_1^2/\sigma_T^2 + 1/\sigma_g^2 + n_1 r_2/\sigma_T)/(n_1^2 + 1)$	–	posterior along the physics curve	1.3.1
Exact posterior on the curve	$p(g \mid \mathbf{y})$, equation (1.19)	–	the limit $\sigma_p \to 0$ by quadrature	1.3.1
Threshold probability	$\Pr(T_2 < T_{\max} \mid \mathbf{y})$	–	prediction under a threshold	1.3.2
Equivalent threshold	$G^* = Q_2/(T_{\max} - T_a)$	kW/K	the same event on an entry	1.3.2
Posterior	$p(\mathbf{z} \mid \mathbf{y})$	–	plausible states after the evidence	1.2.1
Residual vector	$\mathbf{r}(\mathbf{z})$, $\mathcal{L} = \frac{1}{2}\mathbf{r}^{\mathsf{T}}\mathbf{r}$	–	all rows in standard deviations	1.1.7
Jacobian	$\mathbf{J} = \partial\mathbf{r}/\partial\mathbf{z}$	varies	sensitivity of each row to each entry	1.1.7
Normal equations	$(\mathbf{J}^{\mathsf{T}}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^{\mathsf{T}}\mathbf{r}$	varies	the Gauss–Newton step to $\hat{\mathbf{z}}$	1.2.1
MAP state	$\hat{\mathbf{z}} = \arg\min_{\mathbf{z}} \mathcal{L}(\mathbf{z})$	°C, –	the most probable state	1.2.1
Information matrix	$\boldsymbol{\Lambda} = \mathbf{J}^{\mathsf{T}}\mathbf{J} = \sum_k \mathbf{J}_k^{\mathsf{T}}\mathbf{J}_k$	varies	precisions of the rows, added	1.2.1
Laplace covariance	$\boldsymbol{\Sigma} = \left(\mathbf{J}^{\mathsf{T}}\mathbf{J} + \sum_k r_k\,\partial^2 r_k/\partial\mathbf{z}\,\partial\mathbf{z}^{\mathsf{T}}\right)^{-1}$	varies	spreads and correlation of the entries	1.2.1
Preposterior covariance	$\boldsymbol{\Sigma}_a = (\boldsymbol{\Sigma}^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a)^{-1}$	varies	spreads after a row not yet taken	1.6.1
Time constant	$\tau = C_{\mathrm{th}}/G$	s	transient omitted by steady state	1.6.5

1.8 Exercises

Exercise 1.1. (a) Which statement best describes the role of the temperature sensor: it defines the physical law; it reports the exact true state; or it supplies evidence about an uncertain state? (The third: the sensor provides noisy evidence, while the physical model relates the state to the inputs and the conductance.) (b) A hall loses heat to the same ambient through two

parallel paths, its cooling equipment and its envelope, with conductances G_1 and G_2. Name the free and pinned quantities, write the physics row, and say what a temperature reading alone can determine. (Free: T, g_1, g_2; pinned: T_a, Q; the physics row, normalized as in equation (1.5), is $(T - T_a - Q/(G_{\text{ref}}(\mathrm{e}^{g_1} + \mathrm{e}^{g_2})))/\sigma_p$; on its zero set its Jacobian row over $\mathbf{z} = (T, g_1, g_2)^\mathsf{T}$ is $(1,\ (T - T_a)G_1/(G_1 + G_2),\ (T - T_a)G_2/(G_1 + G_2))/\sigma_p$ and the reading row's is $(1/\sigma_T,\ 0,\ 0)$, so the two-by-three Jacobian of the two rows has the null vector $(0,\ G_2,\ -G_1)^\mathsf{T}$, the step $\Delta G_1 = -\Delta G_2$ that leaves the total unchanged; $\mathbf{J}^\mathsf{T}\mathbf{J}$ is singular along it, the reading determines only $G_1 + G_2$, and the split rests on the priors.) (c) Write "will the hall stay below 27°C if the load rises to 45 kW?" as a probability, and say what the graph needs. ($\Pr(T_3 < 27°\text{C} \mid \mathbf{y})$, with a scenario node $T_3 = T_a + (45/G_{\text{ref}})\,\mathrm{e}^{-g}$ at 45 kW sharing g, unobserved.)

Exercise 1.2. (a) A room has an ambient temperature of 20°C, a heat load of 24 kW and a conductance of 4 kW/K. Compute its steady-state temperature. ($T = 20 + 24/4 = 26$°C.) (b) A room has a load of 40 kW, an ambient temperature of 19°C and a strict limit of 27°C. What conductance keeps it below the limit? ($19 + 40/G < 27$ gives $40/G < 8$, so $G > 5$ kW/K.) (c) For the hall of the chapter, with $G_{\text{ref}} = 6$ kW/K, evaluate r_{phys} and r_{obs} at $\mathbf{z} = (23.2°\text{C},\ 0)$ and at $\mathbf{z} = (23°\text{C},\ 0)$. (1.2 kW and 0; 0 and -0.4: the first state matches the reading and violates the physics, the second satisfies the physics and lies 0.4 standard deviations from the reading.)

Exercise 1.3. Suppose both the ambient temperature T_a and the conductance G are uncertain. Starting from $T = T_a + Q/G$, write the probability that the temperature exceeds a limit $T_{\max}$, and state what joint distribution is needed to evaluate it. ($\Pr(T > T_{\max} \mid \mathbf{y}) = \iint \mathbf{1}\{T_a + Q/G > T_{\max}\}\, p(T_a, G \mid \mathbf{y})\, \mathrm{d}T_a\, \mathrm{d}G$, which needs the joint posterior of T_a and G; if the two are independent given $\mathbf{y}$, it reduces to $\int p(T_a \mid \mathbf{y})\, \Pr(G < Q/(T_{\max} - T_a) \mid \mathbf{y})\, \mathrm{d}T_a$ over $T_a < T_{\max}$, plus $\Pr(T_a \geq T_{\max} \mid \mathbf{y})$. A later chapter evaluates this integral.)

Exercise 1.4. For the two-condition hall of Figure 1.3, with a proposed load of $Q_2 = 50$ kW and a limit of 28°C: (a) derive the threshold conductance and the equivalent threshold on the test temperature. ($T_2 = 18 + 50/G < 28$ gives $G > 5$ kW/K; by equation (1.31), $T < 18 + (30/50) \times 10 = 24$°C.) (b) Show that the test temperature determines the full-load temperature without knowledge of G. (On the physics curve both temperatures are functions of g, $T = T_a + (Q_1/G_{\text{ref}})\,\mathrm{e}^{-g}$ and $T_2 = T_a + (Q_2/G_{\text{ref}})\,\mathrm{e}^{-g}$, so a step along the null vector of the physics row moves T_2 by $Q_2/Q_1 = 5/3$ times the move of T; in closed form, dividing the two balances gives $T_2 - T_a = (Q_2/Q_1)(T - T_a)$, equation (1.28).) (c) If the reading $y_T = 23.2$°C with $\sigma_T = 0.5$ K were the only information about T, what would the probability of meeting the limit be? ($\Phi((24 - 23.2)/0.5) = \Phi(1.6) = 0.945$.)

Exercise 1.5. The hall is at 23°C under 30 kW with $T_a = 18$°C when its conductance falls from 6 to 3 kW/K; its thermal capacitance is $C_{\text{th}} = 3600$ kJ/K (illustrative). (a) Compute the time constant and the new steady temperature. ($\tau = 3600/3 = 1200\,\text{s} = 20\,\text{min}$; 28°C.) (b) Compute the temperature after 10 and 60 minutes, and the time at which it reaches 27.5°C. ($28 - 5\,\mathrm{e}^{-0.5} = 24.97$°C; $28 - 5\,\mathrm{e}^{-3} = 27.75$°C; $5\,\mathrm{e}^{-t/\tau} = 0.5$ gives $t = \tau \ln 10 = 46.1$ min.) (c) State what a steady-state model would predict and when it is adequate. (It predicts 28°C immediately; it is adequate for conditions held for several time constants, not for the first minutes after the failure.)

Exercise 1.6. Remove the prior row from the hall of Section 1.2.1, so that the residual vector holds the rows r_1 and r_2 of equation (1.10) over $\mathbf{z} = (T, g)^\mathsf{T}$, with $\sigma_p = 0.01$ K and $\sigma_T = 0.5$ K. (a) Find the MAP state and write the two-by-two Jacobian there. (Both rows vanish at $T = y_T = 23.2$°C and $g = \ln(5/5.2) = -0.0392$; there $\mathbf{J} = \left[\begin{smallmatrix} 100 & 520 \\ 2 & 0 \end{smallmatrix}\right]$, since $5\,\mathrm{e}^{0.0392}/0.01 = 520$.) (b) Show that $\mathbf{\Sigma} = (\mathbf{J}^\mathsf{T}\mathbf{J})^{-1}$ has $\Sigma_{11} = \sigma_T^2$ for every σ_p, and compute $\sqrt{\Sigma_{22}}$. (The Jacobian is square and invertible, so $\mathbf{\Sigma} = \mathbf{J}^{-1}\mathbf{J}^{-\mathsf{T}}$. Writing $\mathbf{J} = \left[\begin{smallmatrix} a & b \\ c & 0 \end{smallmatrix}\right]$, $\mathbf{J}^{-1} = \left[\begin{smallmatrix} 0 & 1/c \\ 1/b & -a/(bc) \end{smallmatrix}\right]$; its first row gives $\Sigma_{11} = 1/c^2 = \sigma_T^2 = 0.25\,\text{K}^2$, and its second $\Sigma_{22} = (\sigma_p^2 + \sigma_T^2)/(5\,\mathrm{e}^{-\hat{g}})^2$, so

$\sqrt{\Sigma_{22}} = \sqrt{0.2501}/5.2 = 0.0962$.) (c) Explain why no curvature term is needed here, and how the result gives the reading-alone probability 0.173 of Section 1.4.4. (Both rows are zero at $\hat{\mathbf{z}}$, so $\sum_k r_k\, \partial^2 r_k / \partial \mathbf{z}\, \partial \mathbf{z}^{\mathsf{T}}$ vanishes; the entry g then has the Laplace marginal $\mathcal{N}(-0.0392,\ 0.0962^2)$, and its tail above $g^* = 0.0513$ is $\Phi((-0.0392 - 0.0513)/0.0962) = \Phi(-0.941) = 0.173$.)

Chapter 2

The Language of Probabilistic Questions

2.1 Engineering questions and their mathematical form

Chapter 1 drew the equipment hall as a plant graph: one hall node whose temperature balances a heat load against one heat-removal path to a pinned ambient, with the conductance of that path unknown and a temperature sensor attached to the hall. The graph was built before any question was asked of it. The operator of the hall now asks several questions, and they are not the same kind of question even though they concern the same plant: what the conductance is, whether the hall will exceed a temperature limit, whether the sensor is biased, and whether another sensor is worth installing. An inference algorithm can return a mathematically correct result to a question nobody asked, and the step that prevents it is the translation of the engineering question into the mathematical quantity that answers it. The question of this chapter is

how a practical engineering question about a plant is translated into a mathematical expression that probabilistic inference can evaluate, and how that expression determines what must be computed.

The answer is short. Once a probability statement names its event, its evidence and its model, every question about the plant is a probability or an expectation over the state of its plant graph, with the unknown on the left of the conditioning bar and the evidence on the right. Engineering questions about a plant then fall into five categories, estimation, prediction under a threshold, hypothesis comparison, diagnosis and design, each with its general forms, integrals or ratios written without reference to any particular plant. Five computations supply every quantity those forms contain: the posterior of the state; the trend, which reads that posterior along time; the evidence of a model; the preposterior of a measurement not yet taken; and the expected consequence of an action. Two of them, the trend and the expected consequence, have general forms of their own, which the categories call on. What differs from one plant to the next is the specialization: which entry or function of the state the form is applied to, which event or family of models it names, and which objects of the fit carry the answer.

Jargon: event, conditional probability, indicator function, posterior predictive, trend, crossing time, evidence, Bayes factor, expected consequence, risk, value of information

An *event* is a statement about the plant that is either true or false, such as "the hall temperature exceeds 28°C"; a probability assigns it a number between zero and one. A *conditional probability* $\Pr(A \mid B)$ is the probability of event A when the information B is treated as known. The *indicator function* $\mathbf{1}_A(\mathbf{z})$ of an event equals one for the states at which the event is true and zero elsewhere, so that the probability of the event is the expectation of its indicator. A *posterior predictive* probability is the probability of a future event that stays random even when the state is known, averaged over the posterior of the state. A *trend* is the change per unit time of a quantity the state carries, and the *crossing time* is the time at which a quantity that moves reaches a stated level. The *evidence*, or marginal likelihood, of a model is the probability of the observed evidence averaged over the prior of that model's unknowns; the *Bayes factor* is the ratio of the evidences of two models. The *expected consequence* of an action, also called its expected loss, is the probability-weighted average of the consequences that may follow it; summed over the events that may follow, it is each event's probability times its consequence, the *risk* of the action. The *value of information*, here the expected value of sample information (EVSI), is the reduction in expected loss that a measurement is expected to bring because the decision can be taken after its result is known.

Key idea

Every engineering question about a plant is answered by a specific mathematical object computed from the same plant graph and its rows: a marginal, a tail probability, a predictive average, the posterior read along time, a ratio of evidences, a preposterior or an expected consequence minimized over actions. Five computations produce these objects, the posterior, the trend, the evidence, the preposterior and the expected consequence. The question, not the inference tool, decides which object, and so which computation, is required.

The questions are put to one plant, the equipment hall of Chapter 1, so that the differences between them lie in the questions alone.

2.1.1 The plant graph of the equipment hall

Figure 2.1 draws the plant graph of Chapter 1, with the scenario node that the hall example of Section 2.4.1 adds to it. The upper green box is the free node of the hall at the test load, whose temperature T_1 is an entry of the state; the lower green box is the same hall at the full load, a scenario node whose temperature T_2 is also an entry and which has no sensor; the gray box on the right is the pinned node, the ambient at $T_a = 18$°C; the two solid green arrows are the heat-removal path at the two loads, which carry one and the same effective thermal conductance G, held in the state as the log-conductance $g = \ln(G/G_{\text{ref}})$; the white boxes are the sources, the heat loads Q_1 and Q_2; and the dashed box on the left, joined to the hall by a dashed arrow, is the observation, the reading y_T of the test-load temperature. Each hall node carries one physical row, its steady-state energy balance, written for a node of temperature T and load Q as

$$Q = G\,(T - T_a), \tag{2.1}$$

where Q is the heat load in kW, G the conductance in kW/K, T the hall temperature and T_a the ambient temperature, both in °C. Solved for the hall temperature, it gives the relation used throughout the chapter,

$$T = T_a + \frac{Q}{G}. \tag{2.2}$$

The relation is deterministic: once G is known, T is known. What makes the hall a probabilistic problem is that G is not known, and that a reading y_T differs from T by a measurement error.

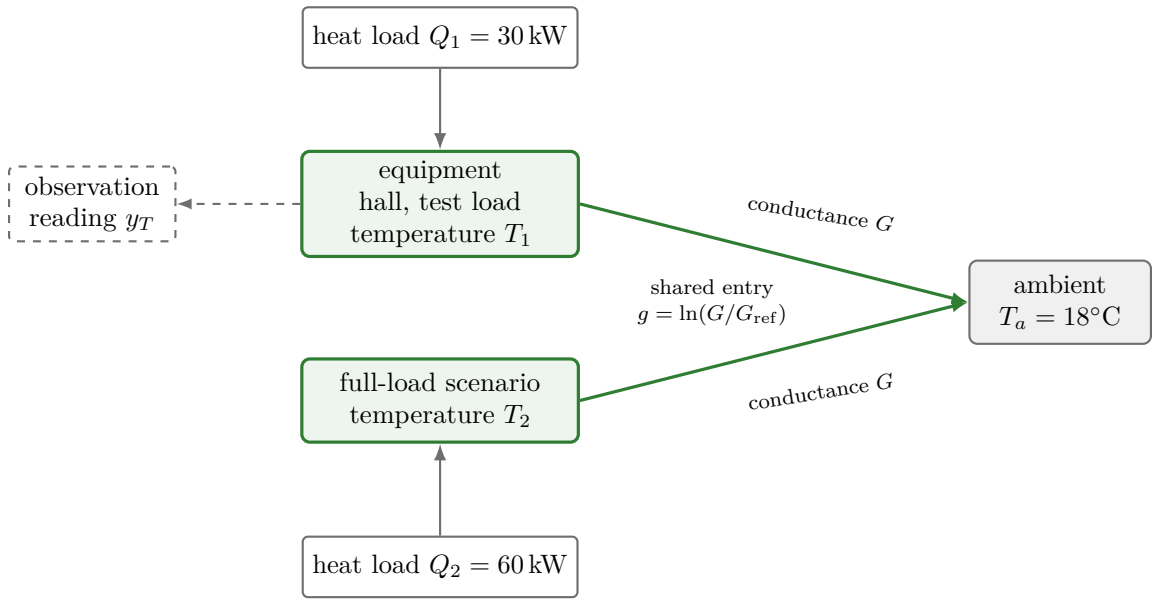


Figure 2.1: The plant graph of the equipment hall with the full-load scenario node. Free: the test-load temperature T_1 and the full-load temperature T_2 (green nodes) and the log-conductance g shared by the two heat-removal paths (solid green edges). Pinned: the ambient temperature $T_a = 18°\mathrm{C}$ (gray node). Sources: the heat loads $Q_1 = 30\,\mathrm{kW}$ and $Q_2 = 60\,\mathrm{kW}$. Observed: the reading y_T of the test-load temperature (dashed box and arrow); the scenario node has no sensor.

2.1.2 Four questions about one plant

Suppose the operator of the hall asks the four questions of Table 2.1. Each row is a question in the operator's words; the second column names the kind of calculation it requires and the third column the quantity that answers it, in the notation that Section 2.2 defines.

Table 2.1: Four questions about the equipment hall and the calculation each requires.

Question	Calculation	Answering quantity
A. What is the effective conductance of the cooling path?	estimation of an unknown physical quantity	$p(G \mid \mathbf{y})$
B. What is the probability that the hall stays below 27.5°C at the full load of 60 kW?	probability of an event	$\Pr(T_2 < 27.5°\mathrm{C} \mid \mathbf{y})$
C. Is the temperature sensor biased?	comparison of competing explanations	$\Pr(M_{\mathrm{bias}} \mid \mathbf{y})$ against $\Pr(M_{\mathrm{unbiased}} \mid \mathbf{y})$
D. Would another temperature sensor be worth installing?	prediction of what a future measurement would change	the preposterior, and the expected loss after the measurement

Question A estimates an unknown physical quantity. Question B asks for the probability of an event. Question C distinguishes between two explanations of the same readings, one in which the sensor reads the hall temperature with zero-mean error and one in which it carries a bias. Question D predicts how a measurement that does not yet exist would change an engineering

decision. All four concern the same graph, and all four are computed from the same physical model; none of them is answered by the same calculation as another.

Principle 2.1. *The same physical graph does not imply the same probabilistic calculation. The formulation of the question determines which calculation is needed.*

2.2 The grammar of probability statements

A question can be classified only when the statement that expresses it is complete: an event or an unknown quantity, the evidence it is conditioned on, and the model under which its probability is computed.

2.2.1 Events

An event is a statement that can be true or false. Three events about the plant are

$$T > 28^{\circ}\mathrm{C}, \qquad G < 5\,\mathrm{kW/K}, \qquad N_f \geq 2,$$

where N_f is the number of failures during an operating period. A probability assigns to an event a number between zero and one. The statement

$$\Pr(T > 28^{\circ}\mathrm{C}) = 0.10$$

means that the probability assigned to the temperature-exceedance event is 10 %. The statement is incomplete in one respect: it does not say how the probability was calculated or what information it incorporates. A probability computed before any reading of the hall was taken and one computed after thirty readings are different numbers for the same event, and the notation must say which is meant. That is the role of conditioning.

2.2.2 Conditional probability and the evidence

The probability of event A given the information B is written

$$\Pr(A \mid B), \tag{2.3}$$

and read "the probability of A, given B". In reliability engineering the information on the right of the bar is the evidence: temperature and pressure measurements, historical failure and repair records, inspection results, operating conditions, manufacturer specifications and prior engineering knowledge. The book collects all of it in the bold symbol $\mathbf{y}$, and keeps a plain y with a subscript, such as y_T, for a single reading. The statement

$$\Pr(T > 28^{\circ}\mathrm{C} \mid \mathbf{y})$$

asks for the probability that the hall temperature exceeds 28°C, conditional on the available evidence. It differs from $\Pr(T > 28^{\circ}\mathrm{C})$, which does not condition on that evidence explicitly. The unknown, the event about the plant, stands on the left of the bar; everything treated as known stands on the right. Keeping that order fixed is what prevents questions from becoming entangled when several are asked of one plant.

2.2.3 The role of the model

A complete probability statement also names the model under which it is computed:

$$\Pr(A \mid \mathbf{y}, M), \tag{2.4}$$

where A is the event of interest, $\mathbf{y}$ the evidence, and M the assumed mathematical model, which comprises the physical relations of the plant graph and the probability assumptions about errors and unknowns. The thermal model of equation (2.2) assumes steady state; a model that adds the thermal storage of the hall would assign a different probability to the same temperature event under the same evidence, because it predicts a different temperature after a change of load. The book sometimes suppresses M to keep the notation readable. The dependence on the model remains in every probability, and a probability is a statement about the plant only to the extent that the model is.

2.2.4 Unknown quantities that are not events

An unknown conductance G is a numerical quantity, not an event; it has a value, not a truth value. What can be asked of it is its probability distribution, described for a continuous quantity by a probability density,

$$p(G \mid \mathbf{y}, M),$$

where the lowercase p denotes a density. Alternatively, a question can ask for the probability that the conductance falls within a range, for example

$$\Pr(5 < G < 7 \mid \mathbf{y}, M),$$

and that is an event probability, because "$5 < G < 7$" is a statement that is true or false. The two are connected by integration of the density over the range that defines the event:

$$\Pr(5 < G < 7 \mid \mathbf{y}, M) = \int_5^7 p(G \mid \mathbf{y}, M)\,\mathrm{d}G, \tag{2.5}$$

with the limits in kW/K. A distribution describes uncertainty across all possible values of a quantity; an event probability describes how much of that probability falls within a specified set of values. A density is not a probability: it has units, $\mathrm{kW}^{-1}\,\mathrm{K}$ here, and it can exceed one. Only its integral over a set is a probability.

2.2.5 A reusable notation

The questions of this chapter, and of the book, are written with the symbols of Table 2.2. The first column is the symbol, the second its meaning. The state $\mathbf{z}$ holds every free quantity of the plant graph at once, the physical quantities such as the log-conductance as well as the temperatures; there is no separate parameter vector. The action a stands for an engineering decision such as inspecting, repairing, replacing, or adding a sensor; the symbol $\tilde{y}_a$ stands for the outcome of a candidate measurement a that has not yet been made, and y_{new} for a future outcome of the plant itself, such as the time of its next failure.

With these symbols, questions as different as the conductance of a cooling path, the survival of a pump and the choice of a sensor are written in one mathematical language, and the differences between them become visible as differences in the expression, not in the words.

Table 2.2: The reusable notation of probabilistic questions.

Symbol	Meaning
$\mathbf{z}$	state vector of the plant graph: every unknown quantity, physical parameters included
$\mathbf{y}$	observed evidence: readings, counts, records, priors
M, M_j	assumed model or system structure; one of several competing models
$p(\mathbf{y} \mid M)$	evidence of model M: the probability of the evidence with the model's state integrated out
B_{AB}	Bayes factor of model M_A against M_B, $p(\mathbf{y} \mid M_A)/p(\mathbf{y} \mid M_B)$; B_{10} is M_1 against M_0
A	an event
$f(\mathbf{z})$, f^*	a quantity derived from the state; a requirement or limit on it
κ	assurance level: the posterior probability a requirement must reach
y_{new}	a future observation that stays random when the state is known
a	an engineering action or decision, or a candidate measurement
$\tilde{y}_a$	the outcome of candidate measurement a, not yet observed
u, $U(a)$	a quantity of the posterior after $\tilde{y}_a$; its expectation over the outcomes, which ranks the candidates
$\mathbf{\Sigma}_a$	covariance of the state after candidate measurement a (the preposterior)
$L(a, \mathbf{z})$	loss or cost of action a when the state is $\mathbf{z}$: its consequence
$\mathcal{R}(a)$	expected loss of action a over the posterior, $\mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]$
c_a, C_e	the fixed cost of action a; the consequence of event e
$\mathbf{g}_a$, $\mathbf{H}_a$	gradient and curvature of $L(a, \mathbf{z})$ in the state at $\hat{\mathbf{z}}$
$f_t(\mathbf{z})$, $\mathbf{g}(t)$	a quantity at time t carried by a state indexed by time; its gradient
α, γ	the level and the trend entries of such a state: the quantity at the record's center $\bar{t}$ and its change per unit time
t^*	crossing time: the time at which f_t reaches a level f^*

2.3 The probability over the state, the general forms and the five computations

Every question of Section 2.1 is an integral over the posterior of the state, a ratio of such integrals, or the minimum of one over a set of actions. The categories of question differ in the integrand, and the integrands draw on five computations: the posterior, the trend, the evidence, the preposterior and the expected consequence.

2.3.1 The probability of an event as an integral over the state

Suppose that inference has produced the posterior distribution of the state,

$$p(\mathbf{z} \mid \mathbf{y}, M),$$

and that a question asks whether the state lies in some region $\mathcal{A}$ of the state space. For the hall, $\mathcal{A}$ might be the set of all combinations of temperature and log-conductance that satisfy an operating limit. The probability of the region is the posterior integrated over it,

$$\Pr(\mathbf{z} \in \mathcal{A} \mid \mathbf{y}, M) = \int_{\mathcal{A}} p(\mathbf{z} \mid \mathbf{y}, M)\, \mathrm{d}\mathbf{z}, \tag{2.6}$$

which is equation (2.5) with a vector in place of one quantity. The same expression is written with the indicator function of the event,

$$\mathbf{1}_A(\mathbf{z}) = \begin{cases} 1 & \text{if event } A \text{ holds at } \mathbf{z}, \\ 0 & \text{otherwise,} \end{cases}$$

which extends the integral to the whole state space and lets the indicator select the region:

$$\Pr(A \mid \mathbf{y}, M) = \int \mathbf{1}_A(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y}, M)\, \mathrm{d}\mathbf{z} = \mathbb{E}\big[\mathbf{1}_A(\mathbf{z}) \mid \mathbf{y}, M\big]. \tag{2.7}$$

The last form reads the probability of an event as the posterior expectation of its indicator. The equation states the central rule of the chapter: *the probability of an event is the posterior integrated over the states at which the event is true.* Many reliability calculations are variations of it, differing only in the event and in how the integral is evaluated.

2.3.2 Future events and the posterior predictive

Some events remain random even when the state is known exactly. A pump whose failure rate is known still fails at a random time; knowing the rate fixes only the probability of surviving a mission. For such a future event A_{future}, the model supplies the conditional probability $\Pr(A_{\text{future}} \mid \mathbf{z}, M)$, and the question asks for its probability given the evidence. The derivation uses three rules. By the law of total probability, the probability of the event is the integral of its joint probability with the state:

$$\Pr(A_{\text{future}} \mid \mathbf{y}, M) = \int \Pr(A_{\text{future}}, \mathbf{z} \mid \mathbf{y}, M)\, \mathrm{d}\mathbf{z}.$$

By the product rule, the joint factors into a conditional and a marginal:

$$\Pr(A_{\text{future}} \mid \mathbf{y}, M) = \int \Pr(A_{\text{future}} \mid \mathbf{z}, \mathbf{y}, M)\, p(\mathbf{z} \mid \mathbf{y}, M)\, \mathrm{d}\mathbf{z}.$$

Finally, the model assumes that the future event depends on the past evidence only through the state: once $\mathbf{z}$ is known, the readings and records of the past add nothing about it. Under that conditional independence, $\Pr(A_{\text{future}} \mid \mathbf{z}, \mathbf{y}, M) = \Pr(A_{\text{future}} \mid \mathbf{z}, M)$, and

$$\Pr(A_{\text{future}} \mid \mathbf{y}, M) = \int \Pr(A_{\text{future}} \mid \mathbf{z}, M)\, p(\mathbf{z} \mid \mathbf{y}, M)\, \mathrm{d}\mathbf{z}. \tag{2.8}$$

This is the posterior predictive probability, the foundation of every prediction of a future failure. Equation (2.7) is the special case in which the conditional probability is an indicator, zero or one at each state. The difference between the two matters as soon as the event is a future failure rather than a property of the present plant, which Section 2.4 shows with a pump.

2.3.3 The general forms of the five categories

Engineering questions about a plant are many, but the mathematical objects that answer them are few, and the book classifies questions by the object rather than by the domain word in the question. Five categories cover every question the book asks, each with its general forms, written here in the notation of Table 2.2: $p(\mathbf{z} \mid \mathbf{y})$ is the posterior density of the state; $f(\mathbf{z})$ is a quantity derived from the state, an entry being the case $f(\mathbf{z}) = z_i$, and f^* a value of it, a requirement in a threshold question; y_{new} is a future observation; M_j is a complete model of the plant, a drawing with its rows; and $\mathbf{\Sigma}_a$ is the covariance of the state after a candidate measurement a. Nothing of any particular plant appears in the forms. Two further forms, of a trend (Section 2.3.4) and of an expected consequence (Section 2.3.5), belong to computations that the categories call on rather than to a category of their own. Every later chapter opens its derivation with one of these forms and then specializes it to its plant.

Estimation. An estimation question asks what a quantity is and how well it is known. Its general form is the posterior distribution of the quantity, the probability that it lies at or below each value f^*,

$$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y}) = \int \mathbf{1}\big[f(\mathbf{z}) \le f^*\big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \tag{2.9}$$

which is equation (2.7) for the event $f(\mathbf{z}) \le f^*$, read as a function of f^*. It is summarized by a named location, the mean, the median or the mode, and a spread; the mean and the variance are

$$\begin{aligned} \mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}] &= \int f(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \\ \mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y}) &= \int \big(f(\mathbf{z}) - \mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]\big)^2\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \end{aligned} \tag{2.10}$$

the median is the value f^* at which equation (2.9) equals $\frac{1}{2}$, and the mode is the value at which the density of $f(\mathbf{z})$ is largest. The uncertainty is also stated as the central 95 % credible interval $[f_L, f_U]$, whose endpoints leave 2.5 % of posterior probability in each tail:

$$\Pr(f(\mathbf{z}) \le f_L \mid \mathbf{y}) = 0.025, \qquad \Pr(f(\mathbf{z}) \le f_U \mid \mathbf{y}) = 0.975. \tag{2.11}$$

The quantiles of a monotone function of a quantity are the images of that quantity's quantiles, in the same order for an increasing function and in reverse order for a decreasing one; the median is a quantile and maps the same way, while the mean and the mode in general do not. Under the Laplace posterior $\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$ an entry is read off directly, with location $\hat{z}_i$ and spread $\sqrt{\Sigma_{ii}}$, and a derived quantity through its gradient $\mathbf{g} = \partial f / \partial \mathbf{z}$ at $\hat{\mathbf{z}}$, with location $f(\hat{\mathbf{z}})$ and spread $\sqrt{\mathbf{g}^\top \mathbf{\Sigma} \mathbf{g}}$ by the delta method. The form is the posterior seen through one quantity: a distribution, of which any single number reported is a named summary.

Prediction under a threshold. A prediction under a threshold asks whether the plant meets a requirement. The requirement is written as the event $f(\mathbf{z}) \ge f^*$, "the requirement holds", with f^* the required value, and its general form is the posterior probability of that event, compared with an assurance level κ, the posterior probability the requirement must reach:

$$\Pr(f(\mathbf{z}) \ge f^* \mid \mathbf{y}) = \int \mathbf{1}\big[f(\mathbf{z}) \ge f^*\big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \qquad \text{accept if } \Pr(f(\mathbf{z}) \ge f^* \mid \mathbf{y}) \ge \kappa. \tag{2.12}$$

A requirement that a quantity stay below an upper limit, such as a temperature $T \le T_{\max}$, is the same form with $f(\mathbf{z}) = -T$ and $f^* = -T_{\max}$; equivalently, the probability of the requirement is one minus the probability of the complementary event, the exceedance $T > T_{\max}$, and the rule is applied to the requirement, never to the exceedance. A joint event is the same integral with the indicator of an intersection. The required value is the engineer's and the assurance level is the customer's, and the form keeps the two apart. When the event lies in the future and stays random even when the state is known, the tail is taken of the predictive distribution of the future observation,

$$p(y_{\text{new}} \mid \mathbf{y}) = \int p(y_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \tag{2.13}$$

of which equation (2.8) is the case of an event, with y_{new} the indicator of A_{future}. The predictive is the randomness the model gives the future outcome, averaged over what is not known about the state. When the quantity moves in time and the requirement concerns a date, the event is a crossing, and its form is the trend form of Section 2.3.4.

Hypothesis comparison. A hypothesis comparison asks which of two descriptions of the plant is true. Its general form is the ratio of the evidences of two complete models fitted to the same evidence, the Bayes factor,

$$B_{AB} = \frac{p(\mathbf{y} \mid M_A)}{p(\mathbf{y} \mid M_B)}, \tag{2.14}$$

where the evidence $p(\mathbf{y} \mid M)$ of a model is the probability of the observed evidence with the model's own state integrated out over its prior. The ratio multiplies the prior odds of the two models into their posterior odds, and because each evidence integrates its own state out, it compares the drawings, not the values of their entries. The Bayes factor B_{10} of the notation table is the case $M_A = M_1$ against $M_B = M_0$.

Diagnosis. A diagnosis asks where and what a fault is. Its general form is hypothesis comparison run over a family $M_1, \ldots, M_J$ of candidate faults generated from the plant graph, one for each component or reading that could be at fault, with the answer returned as the ranking of

$$\Pr(M_j \mid \mathbf{y}) \propto p(\mathbf{y} \mid M_j)\, \Pr(M_j), \qquad j = 1, \ldots, J. \tag{2.15}$$

The answer is the order of the candidates with the posterior probability of each, not a single verdict.

Design. A design question asks what to measure or do next, before the measurement exists. Its general form is the preposterior expectation: for a candidate measurement a whose outcome $\tilde{y}_a$ is not yet known, the expected value, over the outcomes weighted by their predictive distribution, of a quantity u of the posterior that the outcome would produce,

$$U(a) = \mathbb{E}_{\tilde{y}_a \mid \mathbf{y}}\Big[\, u\big(p(\mathbf{z} \mid \mathbf{y}, \tilde{y}_a)\big)\Big] = \int u\big(p(\mathbf{z} \mid \mathbf{y}, \tilde{y}_a)\big)\, p(\tilde{y}_a \mid \mathbf{y})\, \mathrm{d}\tilde{y}_a, \tag{2.16}$$

where $p(\tilde{y}_a \mid \mathbf{y})$ is the predictive (2.13) of the outcome; for an outcome that takes discrete values, such as a number of failures in a test campaign, the integral is a sum over those values weighted by their predictive probabilities. The quantity u is one of four: the spread of a derived quantity after the outcome; the information gained; the indicator that an acceptance rule (2.12) is met after the outcome, which makes $U(a)$ the probability that the campaign ends with the requirement accepted; or the expected loss of the best action once the outcome is known. The candidates are ranked by $U(a)$, the largest first when u is an information gain or an acceptance indicator, the smallest first when u is a spread or a loss, and a candidate's cost enters as a loss of its own. When the posterior is Gaussian and the candidate is a linear Gaussian row,

$r_a(\mathbf{z}) = (\mathbf{h}_a\mathbf{z} - \tilde{y}_a)/\sigma_a$ with $\mathbf{h}_a$ the row vector of its sensitivities and σ_a the standard deviation of its noise, the covariance of the updated posterior does not depend on the outcome, and for a spread or an information gain the form needs only two objects: the preposterior covariance,

$$\boldsymbol{\Sigma}_a = \left(\boldsymbol{\Sigma}^{-1} + \boldsymbol{\Lambda}_a\right)^{-1}, \qquad \boldsymbol{\Lambda}_a = \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a = \frac{\mathbf{h}_a^{\mathsf{T}}\mathbf{h}_a}{\sigma_a^2}, \tag{2.17}$$

with $\mathbf{J}_a = \mathbf{h}_a/\sigma_a$ the Jacobian of the candidate row, and the expected information gain, in nats,

$$\mathrm{EIG}(a) = \tfrac{1}{2}\ln\frac{\det\boldsymbol{\Sigma}}{\det\boldsymbol{\Sigma}_a}. \tag{2.18}$$

The spread of a derived quantity f after the measurement is then $\sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_a\mathbf{g}}$, with the same gradient $\mathbf{g}$ at $\hat{\mathbf{z}}$. For a count, demand or lifetime row the candidate's information is the expected information: $\boldsymbol{\Lambda}_a = \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a$ evaluated at $\hat{\mathbf{z}}$ with the expected counts of the campaign in place of observed ones, a count row entering $\mathbf{J}_a$ through the line $\sqrt{\mu_k}\,\partial \ln \mu_k/\partial\mathbf{z}$ of its working form (Section 4.2.7); equation (2.17) is then the Laplace preposterior, exact only to the order of the Laplace approximation, and it still does not depend on the outcome. A design question values a posterior that does not exist yet, through what the measurement would do to it. When the candidate is an action that brings no new evidence, a repair or a change of design, there is no outcome to average over, the posterior after the action is the present one, and the form reduces to the expected consequence of Section 2.3.5.

2.3.4 The general form of a trend

Some questions concern a quantity that moves in time: whether a failure rate is rising, when it will cross a limit written into a maintenance contract, what it will be next year. By its integrand each is an estimation or a prediction under a threshold; what sets it apart is that the quantity is indexed by time, so that the state must carry how the quantity changes as well as where it stands. Write the quantity, on the scale on which the state carries it (a log-rate, a logit, a temperature), as $f_t(\mathbf{z})$, its value at time t. In the simplest drawing it is a straight line in time about the center $\bar{t}$ of the record, with a level entry α, the quantity at $\bar{t}$, and a *trend entry* γ, its change per unit time:

$$\mathbf{z} = \begin{bmatrix}\alpha\\ \gamma\end{bmatrix}, \qquad f_t(\mathbf{z}) = \alpha + \gamma\,(t-\bar{t}) = \mathbf{g}(t)^{\mathsf{T}}\mathbf{z}, \qquad \mathbf{g}(t) = \frac{\partial f_t}{\partial \mathbf{z}} = \begin{bmatrix}1\\ t-\bar{t}\end{bmatrix}. \tag{2.19}$$

The evidence arrives in windows, one row per window. Window k, centered at the time t_k, gives the row $r_k(\mathbf{z})$, which sees the state only through the quantity at its own time, $f_{t_k}(\mathbf{z})$, so its Jacobian row is the gradient $\mathbf{g}(t_k)^{\mathsf{T}}$ times the row's sensitivity $d_k = \partial r_k/\partial f_{t_k}$: $1/\sigma_k$ for a reading row of standard deviation σ_k, and $\sqrt{\mu_k}$, the square root of the window's expected count, for a count row in the working form of Section 4.2.7. The rows of the record stack to

$$\begin{gathered}\mathbf{J} = \mathbf{D}\,\mathbf{X}, \qquad \mathbf{X} = \begin{bmatrix}\mathbf{g}(t_1)^{\mathsf{T}}\\ \mathbf{g}(t_2)^{\mathsf{T}}\\ \vdots\end{bmatrix}, \qquad \mathbf{D} = \mathrm{diag}(d_1, d_2, \ldots),\\ \boldsymbol{\Lambda} = \mathbf{J}^{\mathsf{T}}\mathbf{J} = \begin{bmatrix}\sum_k d_k^2 & \sum_k d_k^2\,(t_k-\bar{t})\\ \sum_k d_k^2\,(t_k-\bar{t}) & \sum_k d_k^2\,(t_k-\bar{t})^2\end{bmatrix}, \qquad \boldsymbol{\Sigma} = \boldsymbol{\Lambda}^{-1}.\end{gathered} \tag{2.20}$$

The trend entry of the information, $\sum_k d_k^2(t_k-\bar{t})^2$, weighs each window's distance from the center squared, while the windows and their precisions d_k^2 add linearly: widening the span of a record raises the information on the trend faster than adding windows inside it.

The posterior of this state is read along time in three ways, and the three make up the general form of a trend:

$$\begin{aligned}
\Pr(\gamma > 0 \mid \mathbf{y}) &= \int \mathbf{1}[\gamma > 0]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \\
\Pr(t^* \le s \mid \mathbf{y}) &= \int \mathbf{1}\big[f_s(\mathbf{z}) \ge f^*\big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \qquad s \ge \bar{t}, \\
\Pr(f_t(\mathbf{z}) \le f^* \mid \mathbf{y}) &= \int \mathbf{1}\big[f_t(\mathbf{z}) \le f^*\big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \qquad t \text{ beyond the record.}
\end{aligned} \tag{2.21}$$

The first line is the sign of the trend, the probability that the quantity is rising: the threshold form on the trend entry. The second is the distribution function of the crossing time $t^*(\mathbf{z})$, the time at which the quantity reaches the level f^*, for an upward crossing of a level that the quantity stands below at $\bar{t}$: such a quantity, if it rises, has crossed by a time $s \ge \bar{t}$ if and only if it is at or above f^* at s, so each value of the distribution function is a threshold probability at one time, and a state with $\gamma \le 0$ never crosses, a mass $\Pr(\gamma \le 0 \mid \mathbf{y})$ at $t^* = \infty$. The third is the forecast, the estimation form at a time beyond the last window of the record. Nothing of a particular plant appears in the three lines; a plant enters through the scale of f_t, the rows of its windows and the level f^*.

Under the Laplace posterior $\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ the quantity at any time is linear in the state, so each line is an exact normal tail in two numbers, the center $\mathbf{g}(t)^{\mathsf{T}}\hat{\mathbf{z}}$ and the variance $\mathbf{g}(t)^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{g}(t)$:

$$\Pr(\gamma > 0 \mid \mathbf{y}) \approx \Phi\left(\frac{\hat{\gamma}}{\sqrt{\Sigma_{\gamma\gamma}}}\right), \qquad \Pr(t^* \le s \mid \mathbf{y}) \approx \Phi\left(\frac{\mathbf{g}(s)^{\mathsf{T}}\hat{\mathbf{z}} - f^*}{\sqrt{\mathbf{g}(s)^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{g}(s)}}\right), \tag{2.22}$$
$$\mathbf{g}(t)^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{g}(t) = \Sigma_{\alpha\alpha} + 2\,(t - \bar{t})\,\Sigma_{\alpha\gamma} + (t - \bar{t})^2\,\Sigma_{\gamma\gamma}.$$

The sign is the same calculation with the gradient $(0,\ 1)^{\mathsf{T}}$, which reads the trend entry alone, so that its variance is $\Sigma_{\gamma\gamma}$; the forecast is the same with the time asked. The variance of the quantity grows with the square of the distance $t - \bar{t}$, the lever arm: a forecast pays for the uncertainty of the trend with that square. The crossing time itself, $t^* = \bar{t} + (f^* - \alpha)/\gamma$ for $\gamma > 0$, is a ratio with the trend entry in its denominator. When the posterior of γ has a positive density at zero, as a Gaussian posterior always has, t^* has no posterior mean and no variance, so it is reported by its median, the plug-in date $t^*(\hat{\mathbf{z}})$ at which the argument of the normal tail is zero, by the band whose ends solve the second line at 0.025 and 0.975, and by the mass that never crosses; the delta band $t^*(\hat{\mathbf{z}}) \pm 1.96\sqrt{\mathbf{g}_*^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{g}_*}$, with $\mathbf{g}_* = \partial t^*/\partial \mathbf{z}$ at $\hat{\mathbf{z}}$, is a finite number standing in for an infinite one. For a fleet of forty cooling units with six yearly failure counts, Chapter 12 finds the band of the crossing of its contract limit from 5.90 to 28.3 years by inversion, against a delta band from 4.44 to 11.45 years that is wrong at both ends.

2.3.5 The general form of an expected consequence

Other questions ask what to do: whether to maintain a pump, which of two design changes to buy, how many spares to hold. Each compares actions a from a set of candidates by what they are expected to cost. Let $L(a, \mathbf{z})$ be the loss of taking action a when the state is $\mathbf{z}$, its consequence in whatever unit the decision is taken in: reais, hours without cooling, lost production. It is a derived quantity of the state, one for each action. The general form of an expected consequence is its posterior expectation, the expected loss $\mathcal{R}(a)$ of the action, and the action that minimizes it. When the action costs c_a for certain and leaves events e that occur with probability $\Pr(e \mid \mathbf{z}, a)$ and have the consequence C_e each, the loss given the state is the cost of the action plus the consequence of each event weighted by its probability, and by the

predictive (2.8) the posterior expectation turns each conditional probability into its predictive one:

$$\begin{aligned} \mathcal{R}(a) &= \mathbb{E}[L(a,\mathbf{z}) \mid \mathbf{y}] = \int L(a,\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \qquad a^* = \arg\min_a \mathcal{R}(a), \\ L(a,\mathbf{z}) &= c_a + \sum_e C_e \Pr(e \mid \mathbf{z}, a) \quad \Longrightarrow \quad \mathcal{R}(a) = c_a + \sum_e C_e \Pr(e \mid \mathbf{y}, a). \end{aligned} \tag{2.23}$$

The second line is the rule that risk is probability times consequence, summed over the events and added to the certain cost of the action. Nothing of a particular plant appears in it either; a plant enters through the events, the probabilities its posterior gives them and the consequence schedule c_a, C_e that the operator supplies.

Under the Laplace posterior the expected consequence is read from the same objects as every other answer. With $\mathbf{g}_a = \partial L(a,\mathbf{z})/\partial \mathbf{z}$ and $\mathbf{H}_a = \partial^2 L(a,\mathbf{z})/\partial \mathbf{z}\, \partial \mathbf{z}^{\mathsf{T}}$ at $\hat{\mathbf{z}}$, the expansion of the loss to second order about $\hat{\mathbf{z}}$, averaged over $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$, keeps the constant and the curvature term, since $\mathbb{E}[\mathbf{z} - \hat{\mathbf{z}} \mid \mathbf{y}] = \mathbf{0}$ and $\mathbb{E}[(\mathbf{z} - \hat{\mathbf{z}})^{\mathsf{T}} \mathbf{H}_a (\mathbf{z} - \hat{\mathbf{z}}) \mid \mathbf{y}] = \operatorname{tr}(\mathbf{H}_a \boldsymbol{\Sigma})$; the spread follows from the gradient, and the comparison of two actions from the gradient of their difference:

$$\begin{aligned} \mathcal{R}(a) &\approx L(a,\hat{\mathbf{z}}) + \tfrac{1}{2} \operatorname{tr}(\mathbf{H}_a \boldsymbol{\Sigma}), \qquad \operatorname{sd}(L(a,\mathbf{z}) \mid \mathbf{y}) \approx \sqrt{\mathbf{g}_a^{\mathsf{T}} \boldsymbol{\Sigma}\, \mathbf{g}_a}, \\ \Pr(L(a,\mathbf{z}) < L(a',\mathbf{z}) \mid \mathbf{y}) &\approx \Phi\left(\frac{\mathcal{R}(a') - \mathcal{R}(a)}{\sqrt{(\mathbf{g}_a - \mathbf{g}_{a'})^{\mathsf{T}} \boldsymbol{\Sigma}\, (\mathbf{g}_a - \mathbf{g}_{a'})}} \right). \end{aligned} \tag{2.24}$$

The loss at the most probable state, $L(a,\hat{\mathbf{z}})$, misses the curvature term. The curvature is taken in the coordinates of the state, and its sign there, not in a physical quantity, sets the direction of the second-order term. The pump of Section 2.4.4 shows the difference. Its mission survival $\mathrm{e}^{-\lambda\tau}$ is convex in the rate λ, so by Jensen's inequality the predictive 0.7725 lies above the plug-in 0.7659; but the state carries the log-rate ω, in which the survival is concave, with curvature -0.150 at $\hat{\omega}$, and the second-order reading with $\sigma_\omega^2 = 0.25$ moves the plug-in down to 0.747, close to the average 0.748 over the Laplace normal of the log-rate and away from the exact 0.7725. The term is right about the Laplace posterior and wrong about the pump, because the error lies in the Laplace posterior itself; where the second-order term and an exact route move the plug-in in opposite directions, the exact route is the one to report. The second line of equation (2.24) asks how firm a ranking is: the probability that action a costs less than a' at the true state. The variance of the difference is the sum of the two variances minus $2\, \mathbf{g}_a^{\mathsf{T}} \boldsymbol{\Sigma}\, \mathbf{g}_{a'}$, so the band of the comparison is narrower than the two bands taken apart exactly when $\mathbf{g}_a^{\mathsf{T}} \boldsymbol{\Sigma}\, \mathbf{g}_{a'} > 0$, when the two losses move together across the posterior. Chapter 14 ranks two design changes of a cooling plant by draws, the fifth cooling unit removing more lost cooling than a redundant controller with posterior probability 0.865, and Chapter 16 finds a second repair crew and a halved common-cause share nearly equal in expected hours without cooling, 0.73 and 0.75 hours a year, so that the choice between them rests on what else each changes.

2.3.6 The five computations beneath the forms

Five computations sit underneath the forms, each defined by the object it returns: the posterior of the state, the trend, the evidence, the preposterior and the expected consequence.

The posterior. The posterior of the state follows from Bayes' rule, the likelihood of the evidence times the prior divided by their integral over the state. For Gaussian rows divided by their standard deviations it is the exponential of minus the objective of the rows,

$$p(\mathbf{z} \mid \mathbf{y}) = \frac{p(\mathbf{y} \mid \mathbf{z})\, p(\mathbf{z})}{p(\mathbf{y})} \propto \exp[-\mathcal{L}(\mathbf{z})], \qquad \mathcal{L}(\mathbf{z}) = \tfrac{1}{2}\, r(\mathbf{z})^{\mathsf{T}} r(\mathbf{z}), \tag{2.25}$$

and a count, demand or lifetime row puts its negative log-probability term into $\mathcal{L}$ instead. Under the Laplace approximation the posterior is the normal $\mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$ centered at the MAP state $\hat{\mathbf{z}}$, the minimum of $\mathcal{L}$, with the covariance $\mathbf{\Sigma} = (\mathbf{J}^\mathsf{T}\mathbf{J})^{-1}$ built from the Jacobian $\mathbf{J}$ of the rows; Chapter 3 derives the normal equations and the covariance, and Chapter 4 the Laplace approximation. The estimation, threshold and predictive forms are integrals against this posterior, read for an entry $\hat{z}_i$ with $\sqrt{\Sigma_{ii}}$ or for a derived quantity through $\mathbf{g}^\mathsf{T}\mathbf{\Sigma}\mathbf{g}$.

The trend. The trend is the posterior over a state indexed by time or by window, with a trend entry among its entries, read along time: as the probability that the trend entry is positive, as a crossing time with its band, or as a forecast, the three lines of equation (2.21). Its building blocks are those of the posterior, $\hat{\mathbf{z}}$ and $\mathbf{\Sigma}$ from the rows of the windows (2.20), with the gradient $\mathbf{g}(t)$ of the quantity at the time asked. It is a reading of the posterior, and it is named as a computation of its own because its object, its check and its failures are its own: the object is a function of the time asked rather than one number; the check orders the record's residuals in time and inverts the threshold line for the band of the crossing instead of linearizing the crossing time; and it fails where no other reading does, in a forecast that pays for the uncertainty of the trend with the square of its lever arm, in a crossing time without a mean when the posterior of the trend reaches zero, and in the shape of the line, which the record tests only inside its own span.

The evidence. The denominator of equation (2.25), written with its model explicit, is the evidence of the model, the probability of the observed evidence with the model's own unknowns integrated out over their prior:

$$p(\mathbf{y} \mid M_j) = \int p(\mathbf{y} \mid \mathbf{z}, M_j)\, p(\mathbf{z} \mid M_j)\, \mathrm{d}\mathbf{z}. \tag{2.26}$$

It is the normalizing constant of the posterior within model M_j, and it is the only place where the comparison between models enters: the posterior inside either model, however sharp, does not say which model is true. Under the Laplace approximation it is read from the model's own fit, as the volume of its posterior times the height of the integrand at its MAP state, $p(\mathbf{y} \mid M_j) \approx (2\pi)^{d_j/2} \det(\mathbf{\Sigma}_j)^{1/2}\, p(\mathbf{y} \mid \hat{\mathbf{z}}_j, M_j)\, p(\hat{\mathbf{z}}_j \mid M_j)$ with d_j the number of entries of its state, exact when the objective is quadratic (Section 4.2.2). The comparison and diagnosis forms are a ratio and a ranking of it.

The preposterior. The preposterior is the expected posterior under a measurement not yet taken: the posterior that each outcome $\tilde{y}_a$ would produce, averaged over the predictive distribution of the outcomes, the inner object of the design form (2.16). For a linear Gaussian candidate row its covariance does not depend on the outcome, and equation (2.17) follows from the rows. Let the current posterior of the state be approximately Gaussian with mean $\hat{\mathbf{z}}$ and covariance $\mathbf{\Sigma} = (\mathbf{J}^\mathsf{T}\mathbf{J})^{-1}$, and let the proposed measurement depend on the state through its sensitivities, $\tilde{y}_a = \mathbf{h}_a\mathbf{z} + \varepsilon$ to first order, with independent Gaussian noise ε of variance σ_a^2. Divided by its standard deviation, the candidate row is $r_a(\mathbf{z}) = (\mathbf{h}_a\mathbf{z} - \tilde{y}_a)/\sigma_a$, and its Jacobian is the row vector $\mathbf{J}_a = \partial r_a/\partial \mathbf{z} = \mathbf{h}_a/\sigma_a$. The objective after the measurement is the current one plus that row,

$$\mathcal{L}_a(\mathbf{z}) = \tfrac{1}{2}\, r(\mathbf{z})^\mathsf{T} r(\mathbf{z}) + \tfrac{1}{2}\, r_a(\mathbf{z})^2,$$

and its curvature, the information, is the curvature of the stacked Jacobian $[\mathbf{J};\ \mathbf{J}_a]$, the sum $\mathbf{J}^\mathsf{T}\mathbf{J} + \mathbf{J}_a^\mathsf{T}\mathbf{J}_a$: the information of independent rows adds. With $\mathbf{J}^\mathsf{T}\mathbf{J} = \mathbf{\Sigma}^{-1}$, the covariance after the measurement under the linear-Gaussian approximation is the inverse of that sum, which is equation (2.17). When the current rows include exact constraints, $\mathbf{\Sigma}$ is singular in the limit

$\sigma_p \to 0$ and $\boldsymbol{\Sigma}^{-1}$ does not exist there. The limit is then taken through the equivalent update form, which uses $\boldsymbol{\Sigma}$ alone,

$$\boldsymbol{\Sigma}_a = \boldsymbol{\Sigma} - \boldsymbol{\Sigma}\mathbf{J}_a^{\mathsf{T}}\left(\mathbf{I} + \mathbf{J}_a\boldsymbol{\Sigma}\mathbf{J}_a^{\mathsf{T}}\right)^{-1}\mathbf{J}_a\boldsymbol{\Sigma}, \tag{2.27}$$

where $\mathbf{I}$ is the identity of the size of the number of candidate rows. The value $\tilde{y}_a$ of the measurement does not appear in either form, only its Jacobian, its sensitivity divided by its noise. That is why equation (2.17) predicts how the uncertainty would shrink before the measurement is taken, and why it is one foundation of sensor placement and experimental design. The information gain (2.18) is the standard result for a Gaussian prior and a linear Gaussian measurement; where both determinants vanish in the limit of exact rows, the determinant identity $\det\boldsymbol{\Sigma}/\det\boldsymbol{\Sigma}_a = \det(\mathbf{I} + \mathbf{J}_a\boldsymbol{\Sigma}\mathbf{J}_a^{\mathsf{T}})$, which also uses $\boldsymbol{\Sigma}$ alone, evaluates it.

The expected consequence. This computation returns the expected loss of each candidate action over the posterior, $\mathcal{R}(a) = \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]$, and the choice of the action that minimizes it, equation (2.23); summed over events, it is probability times consequence. Its building blocks are $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$ with the gradient $\mathbf{g}_a$ and the curvature $\mathbf{H}_a$ of each action's loss, equation (2.24). It is the expectation of one more derived quantity, and it is named as a computation of its own because its object is a choice among actions rather than a description of the state, because its check is the curvature term that separates it from the loss at the most probable state together with the probability that the ranking holds, and because its failures are its own: a second-order reading whose sign follows the coordinates of the state, a ranking that is clear in expectation and uncertain at the true state, and a consequence schedule that the posterior does not contain, whose own uncertainty no band of the state carries.

Table 2.3 sets the five categories against the five computations: its first column is the category, its second the computation with the kind of question in the category that calls on it, and its third the object the computation returns. A category calls on more than one computation when its questions differ in kind.

In this book the posterior is used in every chapter; the trend in Chapter 12; the evidence in Chapters 4, 6, 7, 10, 11, 12 and 15, and in Chapters 22 to 24 of Volume II; the preposterior at step 7 of Chapters 1 and 4 to 16, and in Chapter 20 of Volume II; and the expected consequence in Chapters 14 and 16, and in Chapter 25 of Volume II, for an expected loss and for the choice of a design.

2.3.7 The finer question types inside the five categories

A finer list of question types is useful in practice, because each type has a characteristic wording. Table 2.4 lists twelve of them. The first column is the type, the second an example in engineering language, the third the mathematical target, and the last two the category of the book it belongs to and the computation beneath it. The list does not limit the applications; it collects the mathematical patterns that recur across them.

Most rows fall into a category directly; five need a sentence. A *joint event* such as "both pumps fail" is an event like any other, a region of the state space, so its probability is equation (2.7) with the indicator of an intersection; the wording "both" changes the region, not the category. A *trend* question asks for threshold events, on the trend entry and on the quantity at a date, so its category is prediction under a threshold; its computation is the trend, because the state is indexed by time and the crossing time is a ratio that the threshold line must be inverted to bound. An *expectation* of a consequence, $\mathbb{E}[L \mid \mathbf{y}] = \int L(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$, is the mean of the marginal of the derived quantity $L(\mathbf{z})$: its category is estimation, applied to a loss, and its computation is the expected consequence, whose curvature term separates it from the loss at the most probable state; its tail $\Pr(L > l \mid \mathbf{y})$ would be a prediction under a threshold. An *importance* measure is a derivative of a derived quantity with respect to an entry

Table 2.3: The five categories of question, the computations that answer them, and the object each computation returns.

Category	Computation	Object returned
estimation	posterior	an entry, $\hat{z}_i$ with $\sqrt{\Sigma_{ii}}$; a derived quantity, $f(\hat{\mathbf{z}})$ with $\sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}}$
	trend, for the quantity at a stated time and for the distribution of a crossing time	$\mathbf{g}(t)^{\mathsf{T}}\hat{\mathbf{z}}$ with $\mathbf{g}(t)^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{g}(t)$; the median crossing time $t^*(\hat{\mathbf{z}})$ and its band by inversion
	expected consequence, for the expected loss of one course of action	$\mathcal{R}(a) = \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]$
prediction under a threshold	posterior	a tail of an entry or of a derived quantity; a predictive
	trend, for the sign of a trend and for a crossing by a date	$\Pr(\gamma > 0 \mid \mathbf{y})$; $\Pr(t^* \le s \mid \mathbf{y})$
hypothesis comparison	evidence	a Bayes factor $B_{AB} = p(\mathbf{y} \mid M_A)/p(\mathbf{y} \mid M_B)$
diagnosis	evidence	the ranking of $\Pr(M_j \mid \mathbf{y})$ over a family generated from the graph
design	preposterior, for the next measurement	$\boldsymbol{\Sigma}_a$, the information gain, the value of information
	expected consequence, for the choice of an action	$a^* = \arg\min_a \mathcal{R}(a)$; the probability that the ranking holds

of the state, the gradient that the delta method already uses, and it belongs with estimation. A *decision* among actions is a design question answered by the expected consequence, with the preposterior as well when a measurement comes before the action.

2.4 Worked examples of estimation, prediction and joint events

Three pairs of questions are often confused: the estimate of a quantity and the probability of a threshold event, the estimate of a rate and the prediction of a mission, a joint failure and the product of two marginal failure probabilities. In each pair the two questions are put to the same model and differ in the form put to it.

2.4.1 Estimation of the conductance of the hall

The operator has measured the hall temperature at the test load of 30 kW and wants to know what the conductance is (question A) and what will happen at the future load of $Q_2 = 60\,\mathrm{kW}$ (question B). Step 2 of the workflow has extended the graph with the scenario node of Figure 2.1, the hall at full load, whose temperature T_2 obeys the balance of equation (2.2),

$$T_2 = T_a + \frac{Q_2}{G}, \qquad T_a = 18^{\circ}\mathrm{C}, \quad Q_2 = 60\,\mathrm{kW}. \tag{2.28}$$

The conductance enters the state as the log-conductance $g = \ln(G/G_{\mathrm{ref}})$ with $G_{\mathrm{ref}} = 1\,\mathrm{kW/K}$. (The symbol η that sometimes names this quantity is reserved in this book for the Weibull scale.) For this example, suppose that an earlier inference from the readings has been summarized as one row on the log-conductance, the value $\hat{g}$ with the standard deviation σ_g,

$$r_g(\mathbf{z}) = \frac{g - \hat{g}}{\sigma_g}, \qquad \hat{g} = \ln 6.5 = 1.871802, \quad \sigma_g = 0.05. \tag{2.29}$$

Table 2.4: Twelve question types, their mathematical targets, and the category and computation each belongs to.

Question type	Example	Mathematical target	Category	Computation
Estimation	What is the failure rate?	$p(\lambda \mid \mathbf{y})$	estimation	posterior
Threshold probability	Will the temperature stay within its limit?	$\Pr(T \le T_{\max} \mid \mathbf{y})$	prediction under a threshold	posterior
Joint and conditional events	Will both pumps fail?	$\Pr(F_A \cap F_B \mid \mathbf{y})$	prediction under a threshold	posterior
Prediction	Will this component survive 100 h?	$\Pr(T > 100\,\mathrm{h} \mid \mathbf{y})$	prediction under a threshold	posterior
Trend and crossing	Is the failure rate rising, and when will it pass its limit?	$\Pr(\gamma > 0 \mid \mathbf{y})$, $\Pr(t^* \le s \mid \mathbf{y})$	prediction under a threshold	trend
Derived quantities	What is the system availability?	$p(R(\mathbf{z}) \mid \mathbf{y})$	estimation	posterior
Model comparison	Is the system series or parallel?	$\Pr(M_j \mid \mathbf{y})$	hypothesis comparison	evidence
Diagnosis	Which component caused the problem?	$\Pr(F_j \mid \mathbf{y})$	diagnosis	evidence
Risk and expectation	What is the expected failure cost?	$\mathbb{E}[L \mid \mathbf{y}]$	estimation	expected consequence
Importance	Which component contributes most to risk?	$\partial R/\partial z_i$ or another importance measure	estimation	posterior
Information and experiment design	Which sensor should be installed?	expected information gain or value of information	design	preposterior
Decision optimization	Should the pump be repaired or replaced?	$\arg\min_a \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]$	design	expected consequence, with the preposterior when a measurement comes first

These values are illustrative; how readings are reduced to a row of this kind is the subject of later chapters. The row is the only one that acts on g alone, and the fit below returns the posterior of the third entry with center $\hat{g}$ and spread $\sqrt{\Sigma_{33}} = \sigma_g$.

Question A is an estimation question, of the general form (2.9). The quantity asked for is the conductance, a derived quantity of the log-conductance, $f(\mathbf{z}) = G = G_{\text{ref}}\,\mathrm{e}^{g}$, with g the third entry of the state; it is increasing in that one entry. The answer is therefore read from three objects of the fit of the hall: the MAP state $\hat{\mathbf{z}}$, the Laplace covariance $\boldsymbol{\Sigma}$ and the gradient of f. The gradient has a single nonzero entry, $\mathbf{g} = (0,\ 0,\ G)^{\mathsf{T}}$, so the band $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} = G^2\,\Sigma_{33}$ uses only the third diagonal entry of $\boldsymbol{\Sigma}$; and because f is increasing in one entry, its quantiles are the images of that entry's quantiles under the exponential. When the third entry is exactly normal, as it is here, G is exactly lognormal, with median $G_{\text{ref}}\,\mathrm{e}^{\hat{z}_3}$, mean $G_{\text{ref}}\,\mathrm{e}^{\hat{z}_3+\Sigma_{33}/2}$ and central 95 % interval $G_{\text{ref}}\,\mathrm{e}^{\hat{z}_3 \pm 1.960\sqrt{\Sigma_{33}}}$, and these exact summaries are reported beside the Laplace band. The fit of the state of the extended graph is the building block.

The posterior of the whole state of the extended graph follows from the row (2.29) and the two balance rows as a fit. The state holds the hall temperature T_1 at the test load $Q_1 = 30\,\text{kW}$, the full-load temperature T_2 and the log-conductance,

$$\mathbf{z} = \begin{bmatrix} T_1 \\ T_2 \\ g \end{bmatrix}. \tag{2.30}$$

The rows are the energy balances of the two hall nodes, written in temperature form as in Chapter 1 and divided by a model tolerance σ_p in kelvin, and the earlier inference, which enters as a row on g of the same form as a prior row:

$$r(\mathbf{z}) = \begin{bmatrix} (T_1 - T_a - (Q_1/G_{\text{ref}})\,\mathrm{e}^{-g})/\sigma_p \\ (T_2 - T_a - (Q_2/G_{\text{ref}})\,\mathrm{e}^{-g})/\sigma_p \\ (g - \hat{g})/\sigma_g \end{bmatrix}, \qquad \mathcal{L}(\mathbf{z}) = \tfrac{1}{2}\,r(\mathbf{z})^{\mathsf{T}} r(\mathbf{z}). \tag{2.31}$$

The exact balance of equation (2.28) is the limit $\sigma_p \to 0$. The Jacobian has one row per row of r and one column per entry of $\mathbf{z}$, so it is 3×3; each entry is the sensitivity of a row to an entry of the state, in standard deviations of the row per unit of the entry:

$$\mathbf{J} = \frac{\partial r}{\partial \mathbf{z}} = \begin{bmatrix} 1/\sigma_p & 0 & (Q_1/G_{\text{ref}})\,\mathrm{e}^{-g}/\sigma_p \\ 0 & 1/\sigma_p & (Q_2/G_{\text{ref}})\,\mathrm{e}^{-g}/\sigma_p \\ 0 & 0 & 1/\sigma_g \end{bmatrix} = \begin{bmatrix} 100 & 0 & 461.5 \\ 0 & 100 & 923.1 \\ 0 & 0 & 20 \end{bmatrix}$$

at the most probable state with $\sigma_p = 0.01\,\text{K}$, in K^{-1} in the first two columns. The third column says that a unit of g moves the full-load balance twice as far as the test-load balance, because the full load is twice the test load. The normal equations $(\mathbf{J}^{\mathsf{T}}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^{\mathsf{T}} r$ give the Gauss–Newton step. The table below lists the iterations from the starting state (23°C, 28°C, $\ln 6$), the forward solution of both balances at $G = 6\,\text{kW/K}$: the columns are the iteration, the three entries and the objective.

iteration	T_1 (°C)	T_2 (°C)	g	$\mathcal{L}$
0	23.0000	28.0000	1.7918	1.281
1	22.5998	27.1996	1.8718	6.083
2	22.6154	27.2308	1.8718	2×10^{-26}
3	22.6154	27.2308	1.8718	2×10^{-26}

At the start both balances hold, so the objective 1.281 is the prior row's alone, $\tfrac{1}{2}\big((\ln 6 - \ln 6.5)/0.05\big)^2$. The row on g is linear, so the first step puts g at $\hat{g}$; the objective rises at that step because the step is taken on the linearization of e^{-g}, and the balance residuals it leaves,

-0.016 and $-0.031\,\mathrm{K}$, are magnified a hundredfold by the division by σ_p. The second step closes both balances. At the most probable state

$$\hat{\mathbf{z}} = \begin{bmatrix} 22.615^\circ\mathrm{C} \\ 27.231^\circ\mathrm{C} \\ 1.871802 \end{bmatrix}$$

every row is zero, $T_1 = 18 + 30/6.5$ and $T_2 = 18 + 60/6.5$.

For rows divided by their standard deviations, as here, the information matrix is $\mathbf{\Lambda} = \mathbf{J}^\mathsf{T}\mathbf{J}$. When the rows are left in their physical units, with Jacobian $\mathbf{J}_u$ and weights $\mathbf{W} = \mathrm{diag}(1/\sigma_k^2)$, it is $\mathbf{J}_u^\mathsf{T}\mathbf{W}\mathbf{J}_u$, the same matrix, since $\mathbf{J} = \mathbf{W}^{1/2}\mathbf{J}_u$. Its inverse is the Laplace covariance of the state. The curvature term $\sum_k r_k\, \partial^2 r_k/\partial \mathbf{z}^2$ that Newton's method would add vanishes here, because every row is zero at $\hat{\mathbf{z}}$, so

$$\mathbf{\Sigma} = (\mathbf{J}^\mathsf{T}\mathbf{J})^{-1} = \begin{bmatrix} 0.05335 & 0.10651 & -0.01154 \\ 0.10651 & 0.21312 & -0.02308 \\ -0.01154 & -0.02308 & 0.00250 \end{bmatrix} \qquad (\sigma_p = 0.01\,\mathrm{K}). \tag{2.32}$$

The ridge of the posterior. The eigenvalues of this $\mathbf{\Sigma}$ are 0.2689, 1.0×10^{-4} and 9.3×10^{-7}: one direction carries almost all of the uncertainty, and the other two shrink with σ_p^2. In the limit $\sigma_p \to 0$ the covariance is a single outer product of the vector $\mathbf{v} = (-Q_1/G,\ -Q_2/G,\ 1)^\mathsf{T}$,

$$\mathbf{\Sigma} \to \sigma_g^2\, \mathbf{v}\mathbf{v}^\mathsf{T} = \begin{bmatrix} 0.05325 & 0.10651 & -0.01154 \\ 0.10651 & 0.21302 & -0.02308 \\ -0.01154 & -0.02308 & 0.00250 \end{bmatrix}, \qquad \mathbf{v} = \begin{bmatrix} -4.615 \\ -9.231 \\ 1 \end{bmatrix}, \tag{2.33}$$

where $\mathbf{v}$ is the direction along which both balance rows stay zero (it is orthogonal to the first two rows of $\mathbf{J}$) and σ_g^2 the variance the third row leaves along it. The three entries move as one along $\mathbf{v}$: a unit rise of g cools the hall by 4.615 K at the test load and 9.231 K at full load. The limit is reached quickly: $\sqrt{\Sigma_{22}}$ is 1.1014 K at $\sigma_p = 1\,\mathrm{K}$, 0.4722 K at 0.1 K, 0.46165 K at 0.01 K, 0.46154 K at 0.001 K and 0.461538 K at 10^{-4} K, and $\sqrt{\Sigma_{33}} = 0.05$ at every σ_p. The answers below are read off the limit.

The fit supplies the objects that question A was reduced to. The log-conductance is the third entry of the state, $\hat{z}_3 = 1.871802$ with $\sqrt{\Sigma_{33}} = 0.05$, and the conductance has the gradient and band

$$\mathbf{g} = \frac{\partial G}{\partial \mathbf{z}} = \begin{bmatrix} 0 \\ 0 \\ G \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 6.5 \end{bmatrix} \mathrm{kW/K}, \qquad \sqrt{\mathbf{g}^\mathsf{T}\mathbf{\Sigma}\mathbf{g}} = 6.5 \times 0.05 = 0.325\,\mathrm{kW/K},$$

so that $G = 6.5 \pm 0.325\,\mathrm{kW/K}$ under the Laplace approximation. The exact lognormal summaries follow from the normal of the third entry: the median is $G_{\mathrm{ref}}\, \mathrm{e}^{\hat{z}_3} = 6.5\,\mathrm{kW/K}$; the mean, by the moment-generating function of the normal, is $G_{\mathrm{ref}}\, \mathrm{e}^{\hat{z}_3 + \Sigma_{33}/2} = 6.5\, \mathrm{e}^{0.00125} = 6.508\,\mathrm{kW/K}$; and the central 95 % credible interval is $6.5\, \mathrm{e}^{\pm 1.960 \times 0.05}$, from 5.893 to 7.169 kW/K. Each of these is a summary of a distribution; the answer to question A is the distribution itself, equation (2.9) with $f(\mathbf{z}) = G$, and any single number reported from it must say which summary it is.

2.4.2 The threshold probability of the hall at full load

Question B asks whether the hall will remain below $T_{\max} = 27.5^\circ\mathrm{C}$ at full load, that is, for

$$\Pr(T_2 < 27.5^\circ\mathrm{C} \mid \mathbf{y}).$$

It is a prediction under a threshold, of the general form (2.12). The requirement is the upper limit $T_2 < T_{\max}$ on the second entry of the state, so the form applies with $f(\mathbf{z}) = -T_2 = -z_2$ and $f^* = -T_{\max}$; equivalently, the probability of the requirement is one minus that of the exceedance, $1 - \Pr(T_2 \geq T_{\max} \mid \mathbf{y})$, and the assurance rule is applied to the requirement. The plant enters through the scenario node: the event is a bound on an entry that no sensor reads, tied to the third entry by the full-load balance row. Two routes read the tail from the fit. Under the Laplace approximation the second entry is normal with center $\hat{z}_2$ and variance Σ_{22}, the second diagonal entry of $\boldsymbol{\Sigma}$, so the probability is a normal tail in those two numbers, read off $\hat{\mathbf{z}}$ and the limit of equation (2.33):

$$\Pr(T_2 < T_{\max} \mid \mathbf{y}) \approx \Phi\Big(\frac{T_{\max} - \hat{z}_2}{\sqrt{\Sigma_{22}}}\Big) = \Phi\Big(\frac{27.5 - 27.231}{0.4615}\Big) = \Phi(0.5833) = 0.720. \tag{2.34}$$

The exact value along the ridge. The second route, which the chapter reports, uses the ridge. Along $\mathbf{v}$ the full-load temperature is a decreasing function of the third entry, $T_2 = T_a + (Q_2/G_{\text{ref}})\,\mathrm{e}^{-g}$, so the bound on T_2 is a bound on g, at the conductance where the second row meets the limit,

$$\Pr(T_2 < 27.5^\circ\text{C} \mid \mathbf{y}) = \Pr\big(G > G^* \mid \mathbf{y}\big), \qquad G^* = \frac{Q_2}{T_{\max} - T_a} = 6.3158\,\text{kW/K}, \tag{2.35}$$

that is $g > \ln(G^*/G_{\text{ref}}) = 1.843053$; and the third entry is exactly normal, so

$$\Pr(T_2 < 27.5^\circ\text{C} \mid \mathbf{y}) = \Phi\Big(\frac{\hat{z}_3 - \ln(G^*/G_{\text{ref}})}{\sqrt{\Sigma_{33}}}\Big) = \Phi\Big(\frac{1.871802 - 1.843053}{0.05}\Big), \tag{2.36}$$

$$\Pr(T_2 < 27.5^\circ\text{C} \mid \mathbf{y}) = \Phi(0.5750) = 0.717. \tag{2.37}$$

The normal tail of the entry and the exact value differ by 0.003, because T_2 is not linear in g along the ridge. The probability that the hall stays below the limit at full load is 71.7 %, and the probability that it exceeds the limit is 28.3 %, conditional on the row of equation (2.29), the operating conditions and the steady-state thermal model. Whether a 28.3 % probability of exceeding the limit is acceptable is not a probability question; it depends on the consequences of exceedance and the operating criteria that apply. The assurance rule of equation (2.12) states such a criterion as a number, applied to the probability that the requirement holds: at an assurance level $\kappa = 0.95$ the requirement $T_2 < 27.5^\circ\text{C}$ is not accepted, since $\Pr(T_2 < 27.5^\circ\text{C} \mid \mathbf{y}) = 0.717 < 0.95$.

Figure 2.2 shows the two questions side by side. In panel (a) the horizontal axis is the conductance G and the vertical axis its posterior density; the shaded area to the right of the dashed line at G^* is 0.717. In panel (b) the horizontal axis is the full-load temperature T_2 and the vertical axis its posterior density; the solid curve is the exact density implied by the lognormal G, the dashed curve the normal of the entry used in equation (2.34), and the shaded area to the left of the limit is the same 0.717. The two shaded areas are one probability, seen once on the conductance and once on the temperature: the event $T_2 < 27.5^\circ\text{C}$ and the event $G > G^*$ are the same set of states.

The check against the exact value and against draws. Equation (2.37) is exact under the normal posterior of g, because the map from g to T_2 is monotone along the ridge; equation (2.34) is the route that every derived quantity has. Had T_2 not been carried in the state, the same band would come from the gradient of $f(\mathbf{z}) = T_a + (Q_2/G_{\text{ref}})\,\mathrm{e}^{-g}$, $\mathbf{g} = (0,\ 0,\ -9.231)^{\mathsf{T}}$ K, and it agrees with Σ_{22} because in the limit of equation (2.33) $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} = \sigma_g^2(\mathbf{g}^{\mathsf{T}}\mathbf{v})^2$ and $\mathbf{g}$ has the same projection on $\mathbf{v}$, -9.231, as the unit vector on T_2. Table 2.5 compares the exact and the approximate probabilities for the limit of the example and three others. The first column is

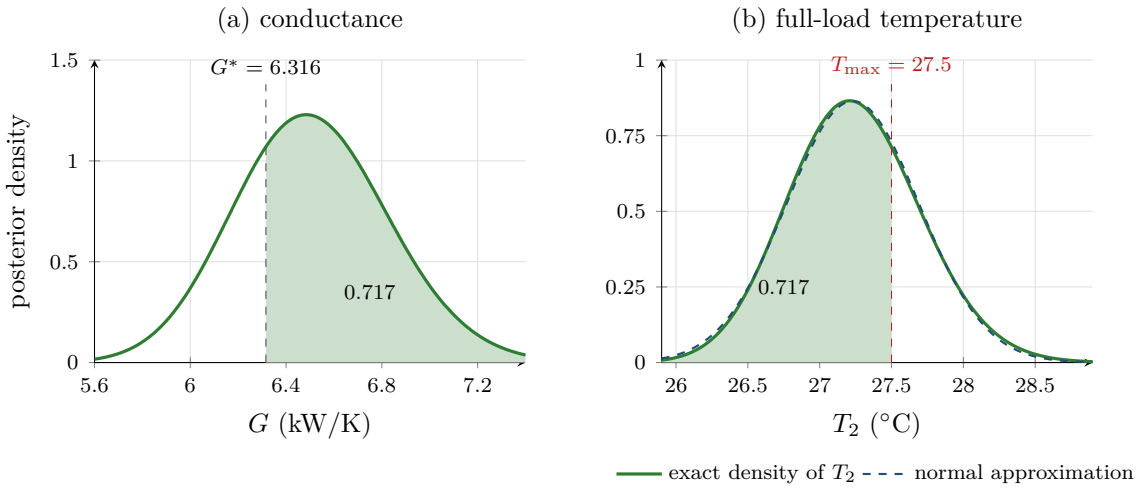


Figure 2.2: One posterior, two questions. (a) Lognormal posterior of the conductance, density $(G\,s\sqrt{2\pi})^{-1}\exp[-(\ln G - \ln 6.5)^2/(2s^2)]$ with $s = \sqrt{\Sigma_{33}} = 0.05$; shaded, $\Pr(G > 6.316\,\text{kW/K} \mid \mathbf{y}) = 0.717$. (b) Posterior of the full-load temperature $T_2 = 18 + 60/G$, lognormal in $T_2 - 18$ with median 9.231 K and log-spread 0.05 (solid), and the normal of the entry, $\mathcal{N}(\hat{z}_2, \Sigma_{22}) = \mathcal{N}(27.231, 0.4615^2)$ (dashed); shaded, $\Pr(T_2 < 27.5°\text{C} \mid \mathbf{y}) = 0.717$.

the limit, the second the equivalent conductance threshold of equation (2.35), the third the standardized distance in g, and the last two the exact probability and its normal approximation with mean $\hat{z}_2$ and variance Σ_{22}. The approximation errs by at most 0.007 over these limits and by 0.003 at 27.5°C, because the curvature of $60/G$ across a posterior of 5 % spread is small; panel (b) of Figure 2.2 shows the slight right skew it misses. A second check draws 200 000 values of g from its posterior, computes T_2 for each, and counts the fraction below the limit: 0.7170, against the exact 0.7174, within the sampling error of 0.001.

Table 2.5: The threshold probability $\Pr(T_2 < T_{\max} \mid \mathbf{y})$ at $Q_2 = 60\,\text{kW}$, $T_a = 18°\text{C}$, $g \mid \mathbf{y} \sim \mathcal{N}(\ln 6.5, 0.05^2)$: exact (equation (2.36)) and by the normal $\mathcal{N}(\hat{z}_2, \Sigma_{22})$ of the entry T_2 (equation (2.34)).

$T_{\max}$ (°C)	$G^* = Q_2/(T_{\max} - T_a)$ (kW/K)	$(\hat{z}_3 - \ln G^*)/\sqrt{\Sigma_{33}}$	exact	normal, Σ_{22}
27.0	6.6667	−0.5064	0.3063	0.3085
27.5	6.3158	0.5750	0.7174	0.7202
28.0	6.0000	1.6009	0.9453	0.9522
28.5	5.7143	2.5767	0.9950	0.9970

What the two questions share. Both questions used the same posterior. Estimation, $p(G \mid \mathbf{y})$, describes the uncertainty about the conductance; the threshold probability, $\Pr(T_2 < T_{\max} \mid \mathbf{y})$, integrates that uncertainty over the states that satisfy the temperature limit. Neither required a new posterior; each asked a different question of the same one. That reuse is one of the principal advantages of a probabilistic model of the plant: the expensive step, the posterior, is paid once, and each new question is an integral over it.

2.4.3 The posterior of the pump's failure rate

A different reliability problem: an industrial pump has failed twice in 1000 h of observed operation, and the reliability engineer asks for the probability that it survives the next 100 h without failing. This is a prediction question, not only a question about the failure rate. Let λ be the unknown constant failure rate, $k = 2$ the observed number of failures, $T_{\text{tot}} = 1000\,\text{h}$ the observed exposure, and $\tau = 100\,\text{h}$ the duration of the future mission. Assume that failures follow a homogeneous Poisson process, so that the count in the exposure is Poisson with mean λT_{tot},

$$N(T_{\text{tot}}) \mid \lambda \sim \text{Poisson}(\lambda T_{\text{tot}}), \qquad \Pr(k \mid \lambda, T_{\text{tot}}) = \frac{(\lambda T_{\text{tot}})^k \, \mathrm{e}^{-\lambda T_{\text{tot}}}}{k!}.$$

As a function of λ, the likelihood is proportional to $\lambda^k \, \mathrm{e}^{-\lambda T_{\text{tot}}}$; the factor $T_{\text{tot}}^k / k!$ does not depend on the rate. The prior is a gamma distribution in the shape–rate parameterization, $\lambda \sim \text{Gamma}(a, b)$, with density proportional to $\lambda^{a-1} \, \mathrm{e}^{-b\lambda}$ and, here,

$$a = 2, \qquad b = 500\,\text{h}.$$

(The book writes the gamma parameters as a and b, so that β keeps its meaning as the Weibull shape.) Bayes' rule gives the posterior up to a constant as likelihood times prior, $p(\lambda \mid k, T_{\text{tot}}) \propto \Pr(k \mid \lambda, T_{\text{tot}}) \, p(\lambda)$.

The rate enters the state, as every rate does in this book, through its log-rate $\omega = \ln(\lambda / \lambda_{\text{ref}})$ with $\lambda_{\text{ref}} = 0.001\,\text{h}^{-1}$, so that the state is the one number

$$\omega = \ln \frac{\lambda}{\lambda_{\text{ref}}}, \qquad \lambda = \lambda_{\text{ref}} \, \mathrm{e}^{\omega}.$$

In the general forms $\mathbf{z}$ is here the one number ω, and every quantity of the fit is a number. The rate is asked for first, an estimation question of the general form (2.9) with $f(\omega) = \lambda = \lambda_{\text{ref}} \, \mathrm{e}^{\omega}$, increasing in ω. It is read from the MAP log-rate $\hat{\omega}$, its variance σ_ω^2 and the derivative $g = \mathrm{d}\lambda/\mathrm{d}\omega = \hat{\lambda}$, whose band is $g \, \sigma_\omega = \hat{\lambda} \, \sigma_\omega$; the exact gamma posterior that the objective identifies is reported beside it. Neither row is Gaussian, and each enters the objective as its negative log-probability term. The count row is the negative log of the Poisson probability; the prior row is the negative log of the gamma density carried over to ω, where the change of variable multiplies the density by $\mathrm{d}\lambda/\mathrm{d}\omega = \lambda$ and so adds $-\ln \lambda$ (the rule, that a density carried to a new coordinate is multiplied by the derivative of the old coordinate with respect to the new, is derived in Section 4.3.3):

$$\mathcal{L}(\omega) = \underbrace{T_{\text{tot}} \lambda - k \ln \lambda}_{\text{count row, } \mathcal{L}_c} + \underbrace{b\lambda - a \ln \lambda}_{\text{prior row, } \mathcal{L}_0} + \text{const.} \tag{2.38}$$

With primes for derivatives in ω and $\mathrm{d}\lambda/\mathrm{d}\omega = \lambda$, each row has a slope and a curvature, and they add:

$$\mathcal{L}_c' = T_{\text{tot}} \lambda - k, \quad \mathcal{L}_c'' = T_{\text{tot}} \lambda, \qquad \mathcal{L}_0' = b\lambda - a, \quad \mathcal{L}_0'' = b\lambda, \qquad \mathcal{L}' = \mathcal{L}_c' + \mathcal{L}_0', \quad \mathcal{L}'' = \mathcal{L}_c'' + \mathcal{L}_0''.$$

The Newton step $\Delta\omega = -\mathcal{L}'/\mathcal{L}''$ moves to the minimum of the parabola with the current slope and curvature. The table below lists the iterations from the maximum likelihood rate $k/T_{\text{tot}} = 0.002\,\text{h}^{-1}$, that is $\omega = \ln 2$, where the count row's slope is zero: the columns are the iteration, the log-rate, the rate, the slopes of the two rows, the slope and curvature of the objective, and the step.

iteration j	ω_j	λ (h^{-1})	$\mathcal{L}_c'$	$\mathcal{L}_0'$	$\mathcal{L}'$	$\mathcal{L}''$	$\Delta\omega$
0	0.69315	0.002000	0.0000	−1.0000	−1.0000	3.0000	0.33333
1	1.02648	0.002791	0.7912	−0.6044	0.1868	4.1868	−0.04462
2	0.98186	0.002669	0.6694	−0.6653	0.0041	4.0041	−0.00103
3	0.98083	0.002667	0.6667	−0.6667	2×10^{-6}	4.0000	-5×10^{-7}

The slope vanishes at $\lambda = (a+k)/(b+T_{\text{tot}})$, so the mode is $\hat{\omega} = \ln(8/3) = 0.98083$ and $\lambda_{\text{ref}}\,\mathrm{e}^{\hat{\omega}} = 4/1500\,\text{h} = 0.002667\,\text{h}^{-1}$. The curvature there is the information, $\Lambda = \mathcal{L}''(\hat{\omega}) = T_{\text{tot}}\hat{\lambda} + b\hat{\lambda} = 2.667 + 1.333 = 4 = a + k$: the count row contributes the expected number of failures in the exposure, the prior row $b\hat{\lambda}$, and the two add as the information of independent rows does. Then $\sigma_\omega^2 = 1/\Lambda = 0.25$, the log-rate has $\sigma_\omega = 0.5$, and the rate, a derived quantity with derivative $g = \mathrm{d}\lambda/\mathrm{d}\omega = \hat{\lambda}$, has $g\,\sigma_\omega = 0.002667 \times 0.5 = 0.001333\,\text{h}^{-1}$.

The objective also identifies the exact posterior. Back in λ, without the $-\ln\lambda$ of the change of variable, the negative log posterior is $(b+T_{\text{tot}})\lambda - (a+k-1)\ln\lambda$, the negative log of the kernel $\lambda^{a+k-1}\,\mathrm{e}^{-(b+T_{\text{tot}})\lambda}$ of a gamma density. Therefore

$$\lambda \mid k, T_{\text{tot}} \sim \text{Gamma}(a', b'), \qquad a' = a + k, \quad b' = b + T_{\text{tot}}. \tag{2.39}$$

For the pump, $a' = 2 + 2 = 4$ and $b' = 500 + 1000 = 1500\,\text{h}$, so $\lambda \mid \mathbf{y} \sim \text{Gamma}(4, 1500\,\text{h})$. The posterior mean and standard deviation are

$$\mathbb{E}[\lambda \mid \mathbf{y}] = \frac{a'}{b'} = \frac{4}{1500\,\text{h}} = 0.002667\,\text{h}^{-1}, \qquad \sqrt{\text{Var}[\lambda \mid \mathbf{y}]} = \frac{\sqrt{a'}}{b'} = 0.001333\,\text{h}^{-1},$$

the two numbers the Newton fit returned: the mode of the log-rate maps to the posterior mean of the rate, and the linearized band is the exact standard deviation, as it is for every gamma posterior. The posterior mean lies between the prior mean $a/b = 0.004\,\text{h}^{-1}$ and the maximum likelihood estimate $k/T_{\text{tot}} = 0.002\,\text{h}^{-1}$, as the weighted form $a'/b' = \frac{b}{b'}\,\frac{a}{b} + \frac{T_{\text{tot}}}{b'}\,\frac{k}{T_{\text{tot}}} = \frac{1}{3}(0.004) + \frac{2}{3}(0.002)$ shows. An estimation question has been answered. The engineer's question, the survival of the next 100 h, has not.

2.4.4 The predictive reliability of the pump's mission

Let T be the time to the next failure, counted from the start of the mission. (In this example T, in hours, is the time to failure; in the hall example, T_1 and T_2 are hall temperatures in °C. The units distinguish them.) Under the Poisson process the time to the next failure does not depend on how long the pump has already run, so the question is

$$\Pr(T > \tau \mid \mathbf{y}), \qquad \tau = 100\,\text{h}. \tag{2.40}$$

For a known constant rate the survival probability is exponential, $\Pr(T > \tau \mid \lambda) = \mathrm{e}^{-\lambda\tau}$. The rate is not known, and the event is a future one that stays random even when the rate is known, so the question is a prediction of the general form (2.13), in its event form (2.8): y_{new} is the event $T > \tau$, and $p(y_{\text{new}} \mid \mathbf{z}) = \mathrm{e}^{-\lambda\tau}$ with $\lambda = \lambda_{\text{ref}}\,\mathrm{e}^{\omega}$. The integral is a function of the one number ω averaged over its posterior: under the Laplace approximation a normal expectation in $\hat{\omega}$ and σ_ω^2, and exactly an integral against the gamma posterior, which has a closed form and is the value reported. The conditional survival probability is averaged over the posterior of the rate,

$$\Pr(T > \tau \mid \mathbf{y}) = \int_0^\infty \mathrm{e}^{-\lambda\tau}\, p(\lambda \mid \mathbf{y})\,\mathrm{d}\lambda.$$

Substitute the density of the gamma posterior of equation (2.39),

$$p(\lambda \mid \mathbf{y}) = \frac{{b'}^{a'}}{\Gamma(a')}\,\lambda^{a'-1}\,\mathrm{e}^{-b'\lambda},$$

and combine the two exponentials:

$$\Pr(T > \tau \mid \mathbf{y}) = \frac{{b'}^{a'}}{\Gamma(a')}\int_0^\infty \lambda^{a'-1}\,\mathrm{e}^{-(b'+\tau)\lambda}\,\mathrm{d}\lambda.$$

The remaining integral is the standard gamma integral, $\int_0^\infty \lambda^{a'-1}\mathrm{e}^{-c\lambda}\,\mathrm{d}\lambda = \Gamma(a')/c^{a'}$, obtained from the definition of the gamma function by the substitution $x = c\lambda$; with $c = b' + \tau$ the gamma functions cancel, and

$$\Pr(T > \tau \mid \mathbf{y}) = \left(\frac{b'}{b'+\tau}\right)^{a'}. \tag{2.41}$$

For the pump,

$$\Pr(T > 100\,\mathrm{h} \mid \mathbf{y}) = \left(\frac{1500}{1600}\right)^{4} = 0.9375^4 = 0.7725.$$

The predicted probability of completing the 100 h mission without a failure is 77.25 %. The calculation needed no Gaussian or Laplace approximation: with a gamma posterior and an exponential survival function, the predictive probability has an exact closed form. That is why the exact gamma is kept here although the Newton fit of the log-rate supplied its mean and spread: averaging $\mathrm{e}^{-\lambda\tau}$ over the normal $\mathcal{N}(\hat\omega, \sigma_\omega^2)$ of the log-rate instead gives 0.748 at 100 h, because the lognormal it implies for λ has the larger mean $\hat\lambda\,\mathrm{e}^{\sigma_\omega^2/2} = 0.00302\,\mathrm{h}^{-1}$ and a heavier upper tail than the gamma.

2.4.5 The error of the plug-in estimate for the pump

A common shortcut substitutes the posterior mean rate into the exponential survival function:

$$R_{\mathrm{plug}}(100\,\mathrm{h}) = \mathrm{e}^{-\mathbb{E}[\lambda|\mathbf{y}]\times 100\,\mathrm{h}} = \mathrm{e}^{-0.2667} = 0.7659.$$

It gives 76.59 % against the predictive 77.25 %. The two differ because averaging and exponentiation do not commute,

$$\mathbb{E}\big[\mathrm{e}^{-\lambda\tau} \mid \mathbf{y}\big] \neq \mathrm{e}^{-\mathbb{E}[\lambda|\mathbf{y}]\,\tau} \quad \text{in general,} \tag{2.42}$$

since survival is a nonlinear function of the rate. The direction of the difference follows from the shape of that function. For $\tau > 0$ the map $\lambda \mapsto \mathrm{e}^{-\lambda\tau}$ has second derivative $\tau^2\mathrm{e}^{-\lambda\tau} > 0$, so it is convex, and Jensen's inequality, the standard result that the expectation of a convex function is at least the function of the expectation, gives $\mathbb{E}[\mathrm{e}^{-\lambda\tau} \mid \mathbf{y}] \geq \mathrm{e}^{-\mathbb{E}[\lambda|\mathbf{y}]\tau}$: the plug-in value understates the predictive reliability. The low rates the posterior admits raise the survival probability by more than the high rates lower it. The plug-in value is wrong not because it is far from the predictive value at 100 h, but because it answers a different question: the survival probability of a pump whose rate is known to be $0.002667\,\mathrm{h}^{-1}$, which is not the pump in question.

Table 2.6 shows how the difference changes with the mission. The first column is the mission duration; the second the predictive reliability of equation (2.41); the third the plug-in reliability at the posterior mean; the fourth their difference; and the last the plug-in reliability at the maximum likelihood estimate $0.002\,\mathrm{h}^{-1}$, which ignores the prior as well as the uncertainty. For a short mission the two Bayesian values agree to three decimals. Their ratio grows with the mission, from 1.0005 at 24 h to 1.87 at 1000 h and 7.0 at 2000 h; the absolute difference peaks near 800 h, at 0.0625, and then falls, to 0.0289 at 2000 h, because both values tend to zero. Figure 2.3 draws the ingredients. In panel (a) the horizontal axis is the failure rate in failures per 1000 h and the vertical axis the probability density per unit of that axis; the blue curve is the prior, the green curve the posterior, and the dotted lines mark the maximum likelihood estimate, the posterior mean and the prior mean. In panel (b) the horizontal axis is the mission duration and the vertical axis the survival probability; the green curve is the predictive reliability and the dashed red curve the plug-in reliability, and the vertical dotted line marks the 100 h mission. The predictive curve lies above the plug-in curve at every duration, and their ratio grows as the mission lengthens, because the longer the mission, the more the survival probability bends across the spread of the posterior.

Table 2.6: Predictive and plug-in mission reliability of the pump. Posterior Gamma$(4, 1500\,\text{h})$; predictive $(1500/(1500+\tau))^4$; plug-in $e^{-4\tau/1500}$; plug-in at the maximum likelihood estimate $e^{-0.002\tau}$.

τ (h)	predictive	plug-in, posterior mean	difference	plug-in, maximum likelihood
24	0.9385	0.9380	0.0005	0.9531
100	0.7725	0.7659	0.0065	0.8187
500	0.3164	0.2636	0.0528	0.3679
1000	0.1296	0.0695	0.0601	0.1353
2000	0.0337	0.0048	0.0289	0.0183

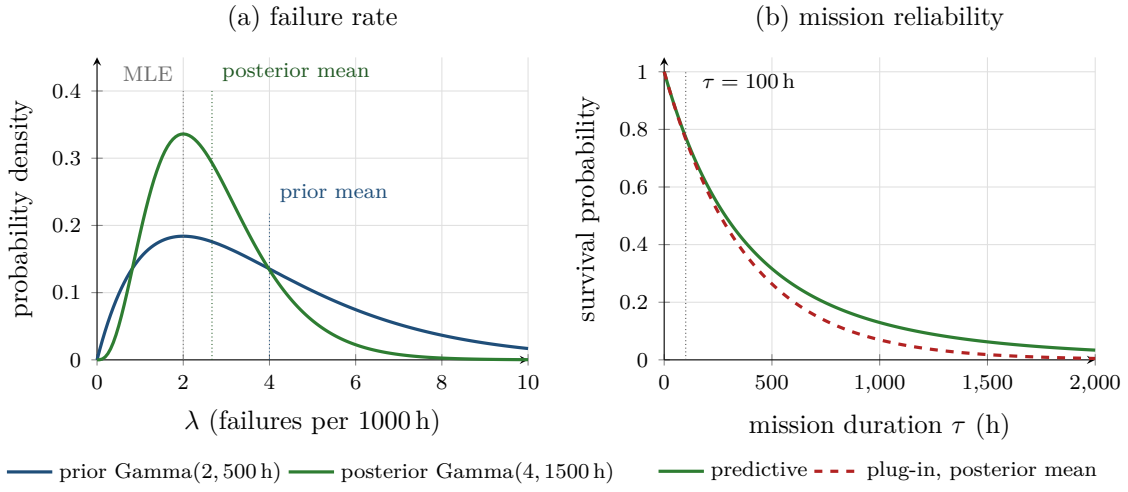


Figure 2.3: The pump failure rate and the mission reliability. (a) Prior $0.25\,x\,e^{-0.5x}$ and posterior $(1.5^4/3!)\,x^3e^{-1.5x}$ in $x = 1000\lambda$; dotted lines at the maximum likelihood estimate 2.0, the posterior mean 2.67 and the prior mean 4.0 failures per 1000 h. (b) Predictive reliability $(1500/(1500+\tau))^4$ of equation (2.41) and plug-in reliability $e^{-4\tau/1500}$; at $\tau = 100$ h they are 0.7725 and 0.7659.

Principle 2.2. *Predicting the reliability of a future mission requires averaging the survival probability over the uncertainty in the rate. Substituting a point estimate of the rate answers the question for a different item, one whose rate is known.*

2.4.6 The reliability of two pumps as a joint event

A cooling system contains two redundant pumps, A and B, either of which can maintain adequate flow by itself. Figure 2.4 draws the arrangement: the white box on the left is the electrical supply that both pumps draw on, the two green boxes are the pumps, each with its own failure mechanisms, and the gray box on the right is the cooling flow they deliver; the solid arrows are the paths along which power and flow pass. Define the events

$$F_A = \{\text{pump A fails during the mission}\}, \qquad F_B = \{\text{pump B fails during the mission}\}.$$

The system fails only if both pumps fail, so the system-failure event is the intersection

$$F_{\text{system}} = F_A \cap F_B, \qquad \Pr(F_{\text{system}} \mid \mathbf{y}) = \Pr(F_A \cap F_B \mid \mathbf{y}).$$

It is the general form (2.12) with the indicator of an intersection in place of the indicator of one bound, here with the marginal failure probabilities given directly rather than computed from a posterior; every probability of this example is conditional on the same evidence, and the bar is left out below to keep the products readable. Suppose $\Pr(F_A) = 0.05$ and $\Pr(F_B) = 0.03$. If the two failures are independent, the probability of the intersection is the product of the marginal probabilities,

$$\Pr(F_A \cap F_B) = \Pr(F_A)\,\Pr(F_B) = 0.05 \times 0.03 = 0.0015,$$

and the system reliability is the probability of the complement,

$$R_{\text{system}} = 1 - 0.0015 = 0.9985, \tag{2.43}$$

or 99.85 %.

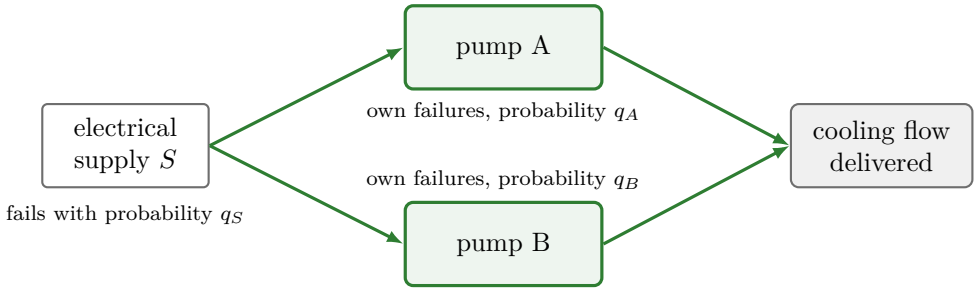


Figure 2.4: Two redundant pumps on a shared electrical supply. Either pump delivers the required flow alone; the system fails when both pump paths are lost. A supply failure (probability q_S) removes both paths at once; the pumps' own failures (probabilities q_A, q_B) are independent given a working supply. With $q_S = 0$ the pump failures are independent.

2.4.7 A shared failure mechanism of the two pumps

Suppose both pumps depend on the same electrical supply. Then independence may not hold, and the probability of the intersection must be written with the product rule, which holds without any independence assumption:

$$\Pr(F_A \cap F_B) = \Pr(F_A)\,\Pr(F_B \mid F_A). \tag{2.44}$$

The conditional probability $\Pr(F_B \mid F_A)$ may be much larger than the marginal $\Pr(F_B)$: a failure of pump A raises the probability that the supply has failed, and a failed supply takes pump B with it. A shared failure mechanism can make the two pump failures strongly dependent.

An illustrative model shows by how much. Let the supply fail during the mission with probability q_S, which takes both pumps out, and let each pump otherwise fail by its own mechanisms, independently, with probabilities q_A and q_B. The law of total probability over the two states of the supply gives the marginal and the joint failure probabilities:

$$\Pr(F_A) = q_S + (1 - q_S)\, q_A, \qquad \Pr(F_B) = q_S + (1 - q_S)\, q_B,$$

$$\Pr(F_A \cap F_B) = q_S \cdot 1 + (1 - q_S)\, q_A q_B. \tag{2.45}$$

Holding the marginals at 0.05 and 0.03, the own-failure probabilities are $q_A = (0.05 - q_S)/(1 - q_S)$ and $q_B = (0.03 - q_S)/(1 - q_S)$. For $q_S = 0.01$, $q_A = 0.04/0.99 = 0.040404$ and $q_B = 0.02/0.99 = 0.020202$, so

$$\Pr(F_A \cap F_B) = 0.01 + 0.99 \times 0.040404 \times 0.020202 = 0.01 + 0.000808 = 0.010808,$$

and equation (2.44) gives $\Pr(F_B \mid F_A) = 0.010808/0.05 = 0.216$, seven times the marginal 0.03. Table 2.7 repeats the calculation for five supply failure probabilities. The first column is q_S; the next two the own-failure probabilities that keep the marginals fixed; then the joint failure probability, the conditional probability of the second failure, the ratio of the joint probability to the independent value 0.0015, and the system reliability. With the marginals unchanged, a supply failure probability of 0.01 multiplies the system failure probability by 7.2, and one of 0.02 by 13.5: the supply, not the redundancy, then governs the system.

Table 2.7: System failure of two redundant pumps on a shared supply, with the marginals held at $\Pr(F_A) = 0.05$ and $\Pr(F_B) = 0.03$ (illustrative model of equation (2.45)).

q_S	q_A	q_B	$\Pr(F_A \cap F_B)$	$\Pr(F_B \mid F_A)$	ratio to 0.0015	R_{system}
0	0.05000	0.03000	0.001500	0.0300	1.00	0.99850
0.001	0.04905	0.02903	0.002422	0.0484	1.61	0.99758
0.005	0.04523	0.02513	0.006131	0.1226	4.09	0.99387
0.01	0.04040	0.02020	0.010808	0.2162	7.21	0.98919
0.02	0.03061	0.01020	0.020306	0.4061	13.54	0.97969

The example shows why the wording "both pumps fail" matters: it names an intersection of events, and an intersection directs the formulation to joint probability and to the modeling of dependence, not to a product of marginals. A system can also fail through several pathways, and in larger systems the failure event is a union of several cut-set events, which may overlap; later chapters on system structure, fault trees and dependence develop those cases.

2.5 Model comparison, diagnosis and decisions

The questions so far were asked of the state of one model. A question about which model holds or which fault is present is put to several models at once, and its answer is a ratio or a ranking of evidences; a question about which action to take is put to several actions, and its answer is a comparison of expected consequences.

2.5.1 Posterior probabilities of competing models

Consider a cooling system whose temperature readings are inconsistent with its expected behavior, and two hypotheses about it,

M_0 : the cooling path operates normally, M_1 : the cooling path is partially obstructed.

The question is which explanation the evidence supports better. This is not a threshold probability on a fixed, known model: the engineer is comparing models. It is a hypothesis comparison, of the general form (2.14) with $M_A = M_1$ and $M_B = M_0$. The plant enters through the two drawings: each is the graph of the cooling path with its own rows, and M_1 carries an entry for the obstruction, a reduction of the conductance of the path, that M_0 pins at zero. Each evidence is the integral (2.26) over its own model's state, read under the Laplace approximation from that model's own fit, its MAP state $\hat{\mathbf{z}}_j$ and covariance $\mathbf{\Sigma}_j$. The extra entry of M_1 is not rewarded merely for fitting the readings better: the evidence multiplies the height of the fit by the posterior volume $\det(\mathbf{\Sigma}_j)^{1/2}$, and the prior density $p(\hat{\mathbf{z}}_j \mid M_j)$ in that height is low when the prior spreads the entry over a wide range. Treat the model index as an unknown taking the values $M_0, M_1, \ldots$, with prior probabilities $\Pr(M_j)$. Bayes' rule for a discrete unknown, with the law of total probability in the denominator, gives

$$\Pr(M_j \mid \mathbf{y}) = \frac{p(\mathbf{y} \mid M_j)\, \Pr(M_j)}{\sum_k p(\mathbf{y} \mid M_k)\, \Pr(M_k)}, \tag{2.46}$$

where $\Pr(M_j)$ is the prior probability of model M_j, $\Pr(M_j \mid \mathbf{y})$ its posterior probability, and $p(\mathbf{y} \mid M_j)$ its evidence, equation (2.26). Dividing equation (2.46) for M_1 by the same equation for M_0 cancels the denominator and gives the odds form,

$$\frac{\Pr(M_1 \mid \mathbf{y})}{\Pr(M_0 \mid \mathbf{y})} = B_{10}\, \frac{\Pr(M_1)}{\Pr(M_0)}, \qquad B_{10} = \frac{p(\mathbf{y} \mid M_1)}{p(\mathbf{y} \mid M_0)} : \tag{2.47}$$

posterior odds equal the Bayes factor (2.14) times the prior odds.

2.5.2 A numerical comparison of two explanations

Suppose the computed evidences are $p(\mathbf{y} \mid M_0) = 0.01$ and $p(\mathbf{y} \mid M_1) = 0.04$, and the prior probabilities are $\Pr(M_0) = 0.8$ and $\Pr(M_1) = 0.2$. Equation (2.46) gives

$$\Pr(M_1 \mid \mathbf{y}) = \frac{0.04 \times 0.2}{0.04 \times 0.2 + 0.01 \times 0.8} = \frac{0.008}{0.016} = 0.50.$$

The Bayes factor is

$$B_{10} = \frac{p(\mathbf{y} \mid M_1)}{p(\mathbf{y} \mid M_0)} = \frac{0.04}{0.01} = 4, \qquad \ln B_{10} = \ln 4 = 1.386 \text{ nats},$$

and the odds form checks the result: posterior odds $4 \times (0.2/0.8) = 1$, a posterior probability of $1/(1+1) = 0.50$. The evidence is four times more probable under the obstruction model, yet its posterior probability is only 50 %, because the prior probability assigned to obstruction was low. The Bayes factor expresses the relative support that the evidence provides; it is not by itself a posterior model probability, and turning it into one requires the prior probabilities of the models.

Figure 2.5 makes the distinction visible. In panel (a) the two pairs of bars are the prior and the posterior probabilities of the two models: the evidence moves probability from M_0 to M_1 until they are equal. In panel (b) the horizontal axis is the prior probability of the obstruction model and the vertical axis its posterior probability; the green curve is $\Pr(M_1 \mid \mathbf{y}) = 4\,\Pr(M_1)/(3\,\Pr(M_1) + 1)$, which equation (2.47) gives for $B_{10} = 4$, and the dashed diagonal is the case $B_{10} = 1$, evidence that does not discriminate. The same Bayes factor of 4 turns a prior of 0.05 into 0.174, a prior of 0.2 into 0.50 (the marked point) and a prior of 0.5 into 0.80: the factor is fixed by the evidence, the conclusion is not.

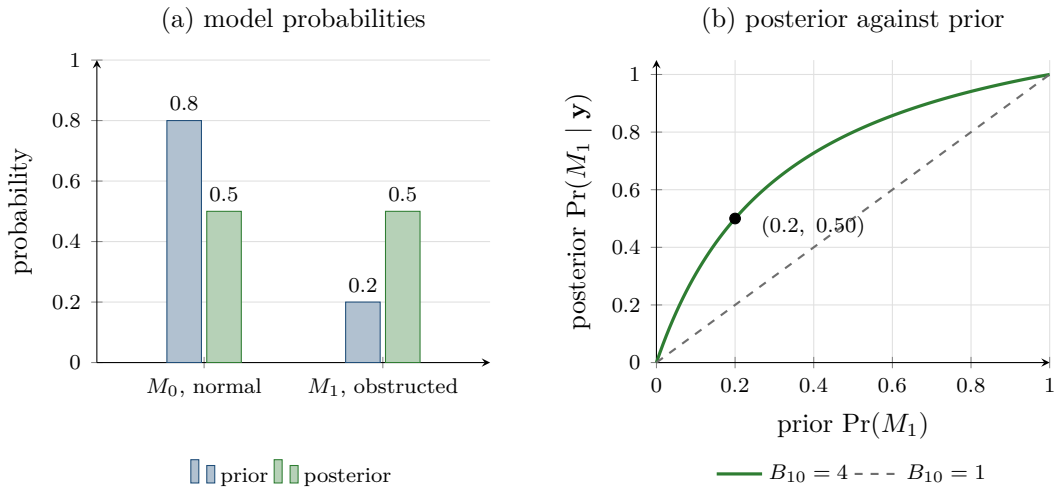


Figure 2.5: Model comparison for the cooling path. (a) Prior and posterior probabilities of the normal (M_0) and obstructed (M_1) models, with evidences 0.01 and 0.04. (b) Posterior probability of M_1 as a function of its prior, $4x/(3x+1)$ for a Bayes factor of 4 (equation (2.47)); the dashed diagonal is a Bayes factor of 1; the point is the example.

2.5.3 Diagnosis over a family of candidate faults

Fault diagnosis usually involves many competing explanations, not two. For a cooling loop whose flow has dropped they might be a pump obstruction, a valve partially closed, a sensor bias, reduced motor performance or an incorrect flow measurement. The mathematical target is the general form (2.15), normalized as in equation (2.46): the posterior probability of each candidate,

$$\Pr(F_j \mid \mathbf{y}) = \frac{p(\mathbf{y} \mid F_j)\,\Pr(F_j)}{\sum_k p(\mathbf{y} \mid F_k)\,\Pr(F_k)},$$

where F_j is a candidate fault condition, which is the drawing M_j of equation (2.15): the plant graph with that fault written into its rows, so that $p(\mathbf{y} \mid F_j) = p(\mathbf{y} \mid M_j)$. The output is the ranking of the candidates by it, not a single verdict. A fault condition can be represented in two ways: as a separate model, whose evidence enters equation (2.46), or as a latent quantity within one common model, such as a bias entry of the state that is zero when the sensor is sound. The second representation means that diagnosis does not always require fitting a separate model for every possible fault. Diagnosis earns its own category for two reasons: the candidates are generated systematically from the plant graph, one for each component or reading that could be at fault, rather than written by hand; and the answer is a ranking. The critical step in either representation is to define the candidate fault events or hypotheses clearly enough that their probabilities can be calculated and compared.

For the hall of Figure 2.1 the family is built on the graph with its reading row kept as a row of its own, $(T_1 - y_T)/\sigma_T$, rather than folded, as in the fit of Section 2.4.1, into the summary row on g; a fault must have a row to act on. One candidate is generated for each element that could be at fault, each the extended graph with one entry added: M_{bias}, a bias entry Δ_T on the reading row, which becomes $(T_1 + \Delta_T - y_T)/\sigma_T$; M_{path}, a shift Δ_g of the log-conductance away from its prior value, an entry added to both balance rows; and M_{load}, an offset Δ_Q of the test load Q_1 from its nominal value, an entry added to the test-load balance row. Each candidate is fitted to the same readings, its evidence is read from its own $\hat{\mathbf{z}}_j$ and $\boldsymbol{\Sigma}_j$, and the ranking orders them. Which candidates the readings can tell apart is read from the Jacobian before any fit. With the one reading of the hall, each added entry can absorb the misfit of that reading alone, so the three candidates fit it equally well and their ranking rests on their prior rows and prior

probabilities. A second reading at another load separates them, because the three added entries then move the two readings in different proportions, which are different columns of $\mathbf{J}$: a bias by equal amounts, a shift of the path in proportion to the loads, and an offset of the test load at the test load only.

2.5.4 Expected loss and the evaluation of an intervention

Questions about a plant also lead to decisions. To decide whether it is worth replacing an unreliable pump, the probability of failure is not enough: the consequences of failure and the cost of the intervention enter as well. Let L be the financial loss caused by a failure, a random quantity because the failure is. Suppose the system-failure probability over the mission is $\Pr(F \mid \mathbf{y}) = 0.0015$, the value of equation (2.43), that a failure would cause a loss $L_F = \text{R\$}\,100\,000$, and that no failure causes no loss, $L_{\neg F} = 0$. The failure is a future event that stays random when the state is known, so the loss given the state is itself an expectation, the derived quantity

$$f(\mathbf{z}) = \mathbb{E}[L \mid \mathbf{z}] = L_F \Pr(F \mid \mathbf{z}) + L_{\neg F} \Pr(\neg F \mid \mathbf{z}) = L_F \Pr(F \mid \mathbf{z}),$$

and the expected loss is its location (2.10), an estimation question about a consequence, answered by the expected consequence (2.23) of leaving the pump as it is. Averaging over the posterior with equation (2.8) turns each conditional probability into its predictive one, and the expected loss is the probability-weighted average of the two outcomes:

$$\mathbb{E}[L \mid \mathbf{y}] = \Pr(F \mid \mathbf{y})\, L_F + \Pr(\neg F \mid \mathbf{y})\, L_{\neg F}, \tag{2.48}$$

the second line of equation (2.23) with one event, consequence $C_F = L_F$, and no cost of action: probability times consequence. Substituting,

$$\mathbb{E}[L \mid \mathbf{y}] = 0.0015 \times \text{R\$}\,100\,000 + 0.9985 \times 0 = \text{R\$}\,150$$

per mission, under these assumptions. It does not mean that the operator will lose exactly R$ 150; the actual loss is either zero or R$ 100 000. It means that R$ 150 is the probability-weighted average of the possible losses, the amount that would be lost per mission on average over many missions like this one.

Now suppose a maintenance action costs R$ 80 and reduces the failure probability from 0.0015 to 0.0005. The expected total cost with maintenance is the certain cost of the action plus the expected loss that remains,

$$\mathbb{E}[L(\text{maintain}) \mid \mathbf{y}] = \text{R\$}\,80 + 0.0005 \times \text{R\$}\,100\,000 = \text{R\$}\,130,$$

against $\mathbb{E}[L(\text{no maintenance}) \mid \mathbf{y}] = \text{R\$}\,150$. The expected net benefit of maintenance is $\text{R\$}\,150 - \text{R\$}\,130 = \text{R\$}\,20$ per mission, and under these assumptions maintenance is the preferred decision. The comparison rests on three numbers, two failure probabilities and one loss, and its conclusion changes if any of them moves: the break-even cost of the action is $(0.0015 - 0.0005) \times \text{R\$}\,100\,000 = \text{R\$}\,100$.

2.5.5 The general decision formulation

Let a be an engineering action from a set of candidates, and $L(a, \mathbf{z})$ the loss of taking it when the state is $\mathbf{z}$. The expected loss of the action is its expected consequence (2.23), $\mathcal{R}(a) = \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]$, the location (2.10) of the derived quantity $L(a, \mathbf{z})$, one for each candidate, and the preferred action is the one that minimizes it:

$$a^* = \arg\min_a \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]. \tag{2.49}$$

The maintenance example of Section 2.5.4 is this rule with one event, the failure, and two actions. The shift is an important one. The questions of the earlier sections asked what will happen; equation (2.49) asks what should be done, given what might happen. The transition from inference to decisions is a principal objective of the engineering of a plant.

2.6 Measurements, the classification of questions and the derivation workflow

A design question asks whether more information is worth obtaining, and it is answered before the information exists.

2.6.1 Uncertainty reduction by a measurement

Suppose the reliability engineer is uncertain about a quantity because the available record is short, and an additional inspection or sensor measurement could reduce that uncertainty. The question "is the information worth its cost?" is a design question, of the general form (2.16). Its first part is how much the measurement would narrow the posterior: the form with u the posterior covariance, which for linear Gaussian rows is the preposterior covariance (2.17), with the information gain (2.18) as its summary. The plant enters through two objects: the Jacobian $\mathbf{J}_a$ of the candidate rows, which says which entries of the state the measurement acts on and how strongly, and the current covariance $\boldsymbol{\Sigma}$, which says along which directions the posterior can still shrink.

For the hall, consider buying n readings of the temperature at the test load $Q_1 = 30\,\mathrm{kW}$, each with an error of standard deviation $\sigma_T = 0.5\,\mathrm{K}$ (an illustrative value). Each reading is a row on the first entry of the state of equation (2.30),

$$r_{a,j}(\mathbf{z}) = \frac{T_1 - \tilde{y}_j}{\sigma_T}, \qquad \mathbf{J}_{a,j} = \frac{\partial r_{a,j}}{\partial \mathbf{z}} = \begin{bmatrix} 1/\sigma_T & 0 & 0 \end{bmatrix} = \begin{bmatrix} 2 & 0 & 0 \end{bmatrix} \mathrm{K}^{-1},$$

so that n readings stack to an $n \times 3$ Jacobian $\mathbf{J}_a$ with $\boldsymbol{\Lambda}_a = \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a = \mathrm{diag}(4n, 0, 0)$ in K^{-2}. The current $\boldsymbol{\Sigma}$ is the ridge of equation (2.33), $\sigma_g^2 \mathbf{v}\mathbf{v}^{\mathsf{T}}$, so in the update form (2.27) every product runs through $\mathbf{v}$: $\boldsymbol{\Sigma}\mathbf{J}_{a,j}^{\mathsf{T}} = \sigma_g^2\,\mathbf{v}\,(2v_1)$ and $\mathbf{J}_{a,j}\boldsymbol{\Sigma}\mathbf{J}_{a,k}^{\mathsf{T}} = 4\sigma_g^2 v_1^2 = 4\Sigma_{11} = 0.21302$ for every pair of readings. The update therefore removes a multiple of $\mathbf{v}\mathbf{v}^{\mathsf{T}}$ itself, and

$$\boldsymbol{\Sigma}_a = \frac{\boldsymbol{\Sigma}}{1 + 0.21302\,n}, \qquad \boldsymbol{\Sigma}_a = \frac{\boldsymbol{\Sigma}}{1.21302} = \begin{bmatrix} 0.04390 & 0.08780 & -0.00951 \\ 0.08780 & 0.17561 & -0.01902 \\ -0.00951 & -0.01902 & 0.00206 \end{bmatrix} \quad (n = 1). \tag{2.50}$$

The reading sees only T_1, yet every entry of $\boldsymbol{\Sigma}$ shrinks by the same factor, because the posterior has only the one direction $\mathbf{v}$ to shrink along: one reading removes $1 - 1.21302^{-1/2} = 9.2\,\%$ of the spread of T_1, T_2 and g alike, so that $\sqrt{\Sigma_{a,33}}$ falls from 0.05 to 0.0454 and $\sqrt{\Sigma_{a,22}}$ from 0.462 to 0.419 K. Table 2.8 gives the result for 0 to 64 readings. The columns are the number of readings, the precision $1/\Sigma_{a,33}$ of g, its posterior standard deviation $\sqrt{\Sigma_{a,33}}$, the standard deviation $\sqrt{\Sigma_{a,22}}$ of T_2, the fraction of that spread removed, and the expected information gain. The precision of g is $400\,(1 + 0.21302\,n) = 400 + 85.21\,n$ by equation (2.50). The expected information gain is equation (2.18); in the limit both determinants are zero, and the determinant identity $\det\boldsymbol{\Sigma}/\det\boldsymbol{\Sigma}_a = \det(\mathbf{I} + \mathbf{J}_a\boldsymbol{\Sigma}\mathbf{J}_a^{\mathsf{T}})$, with every entry of the $n \times n$ matrix $\mathbf{J}_a\boldsymbol{\Sigma}\mathbf{J}_a^{\mathsf{T}}$ equal to 0.21302, gives $\mathrm{EIG} = \frac{1}{2}\ln(1 + 0.21302\,n)$. Equation (2.17) evaluated directly at $\sigma_p = 10^{-4}\,\mathrm{K}$, where $\boldsymbol{\Sigma}^{-1}$ exists, gives the same entries, the precision of 64 readings to within 0.02. The reduction is steep at first and then slows, because the precision grows linearly in n and the spread only as its inverse square root: halving the spread requires a precision of 1600, about fourteen readings, and halving it again a precision of 6400, about seventy.

Table 2.8: Preposterior spread of the hall after n readings at $Q_1 = 30\,\mathrm{kW}$ with $\sigma_T = 0.5\,\mathrm{K}$ (illustrative), from the posterior $\boldsymbol{\Sigma}$ of equation (2.33) and equation (2.50).

n	precision $1/\Sigma_{a,33}$	$\sqrt{\Sigma_{a,33}}$	$\sqrt{\Sigma_{a,22}}$ (K)	spread removed	EIG (nat)
0	400.0	0.0500	0.462	0	0
1	485.2	0.0454	0.419	9.2 %	0.097
4	740.8	0.0367	0.339	26.5 %	0.308
16	1763.3	0.0238	0.220	52.4 %	0.742
64	5853.3	0.0131	0.121	73.9 %	1.342

2.6.2 The value of information

Reducing uncertainty is not necessarily the same as making a better decision. Suppose there are two possible actions, maintain and do not maintain. The best action now minimizes the expected loss, and the expected loss of acting now is

$$\bar{L}_{\text{now}} = \min_a \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}].$$

(The book writes $\bar{L}$ for these minimum expected losses because C names a consequence.) If a candidate measurement is collected first, written m here because a ranges over the actions, with outcome $\tilde{y}_m$, the best action can be chosen again after its value is known, with expected loss $\min_a \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}, \tilde{y}_m]$. Before the measurement is collected its result is unknown, so that expected loss is averaged over the possible results, weighted by their predictive distribution given the present evidence:

$$\bar{L}_{\text{after}} = \mathbb{E}_{\tilde{y}_m|\mathbf{y}}\Big[\min_a \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}, \tilde{y}_m]\Big].$$

This is the general form (2.16) for the candidate m, with u the minimum expected loss under the updated posterior and the candidate ranked smallest first, and here, unlike for the spread, u depends on the outcome itself and not only on its covariance. The expected value of sample information is the difference,

$$\mathrm{EVSI} = \bar{L}_{\text{now}} - \bar{L}_{\text{after}}, \tag{2.51}$$

the expected improvement in the decision, measured as a reduction in expected loss. It is never negative. For each fixed action, averaging the post-measurement expected loss over the results returns the present one, $\mathbb{E}_{\tilde{y}_m|\mathbf{y}}\big[\mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}, \tilde{y}_m]\big] = \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]$, by the law of total expectation; and the minimum over actions inside the average is at most the value of any one action, so $\bar{L}_{\text{after}} \le \min_a \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}] = \bar{L}_{\text{now}}$. Equality holds when no possible result would change the action: a measurement that cannot change the decision has no value for it, however much it narrows the posterior. If the measurement costs more than its EVSI, obtaining it is not economically justified on this criterion.

An illustrative case shows the calculation. Suppose the failure probability 0.0015 of Section 2.5.4 is an average over two conditions of the pump: worn, with probability 0.25 and failure probability 0.0045, and sound, with probability 0.75 and failure probability 0.0005; indeed $0.25 \times 0.0045 + 0.75 \times 0.0005 = 0.0015$. Maintenance costs R$ 80 and leaves either pump at 0.0005. An inspection would reveal the condition exactly. Table 2.9 lists, for each condition and for the present state of knowledge, the expected loss of each action and the better one. Without the inspection the best action is to maintain, at R$ 130. With it, a worn pump is maintained at R$ 130 and a sound one left alone at R$ 50, so

$$\bar{L}_{\text{after}} = 0.25 \times \mathrm{R\$}\,130 + 0.75 \times \mathrm{R\$}\,50 = \mathrm{R\$}\,70, \qquad \mathrm{EVSI} = \mathrm{R\$}\,130 - \mathrm{R\$}\,70 = \mathrm{R\$}\,60.$$

An inspection that reveals the condition is worth up to R$ 60 per mission; because it reveals the condition exactly, this is the expected value of perfect information, an upper bound on what any imperfect inspection could be worth.

Table 2.9: Expected loss per mission of the two actions for a pump of unknown condition (illustrative): maintenance R$ 80, failure loss R$ 100 000.

Condition	probability	Pr(F), no action	no maintenance (R$)	maintain (R$)	better action
worn	0.25	0.0045	450	130	maintain
sound	0.75	0.0005	50	130	no maintenance
unknown (present)	1	0.0015	150	130	maintain

Three quantities are now on the table, and they are related but not interchangeable. The reduction of the posterior covariance, equation (2.17), measures how much uncertainty disappears. The expected information gain, equation (2.18), measures how much is learned. The value of information, equation (2.51), measures how much better the resulting decision is expected to be. A measurement can score well on the first two and be worth nothing on the third, which is why the design of experiments and the optimization of decisions both need the preposterior but do not ask the same question of it. The value of information is where two computations meet: the preposterior supplies the outcomes and their weights, and the expected consequence, minimized over the actions, is what it averages.

2.6.3 Identifying the mathematical form of a question

The examples supply a systematic procedure: before computing anything, identify the mathematical objective of the question. Table 2.10 is the reference for that step. Each row is a form of question in engineering language, the general form it takes, by equation number, its mathematical form, what it derives to, and its category; the table mirrors the atlas of Appendix A, which collects the questions of every chapter with the chapter that derives each.

Many variations of these forms exist. Reliability importance, availability forecasting, maintenance scheduling and spare-parts optimization each add mathematical structure of their own: gradients, integrals of failure probability over an interval, consequences that accrue over time. The rows of the table, over the general forms of the categories, of a trend and of an expected consequence, already cover a substantial range of inference problems, and a new question is usually classified by recognizing which of them it is, before any computation.

2.6.4 The general derivation workflow

The classification is one step of the derivation workflow. Figure 2.6 draws its eight steps: the top row runs from left to right, the bottom row from right to left, and the arrow beneath the bottom row is the return from step 7 to step 3 when a new measurement is added. In the words of the preface:

The plant graph (step 0). The physical system as it is, before any question is asked: its nodes, paths and sources; which quantities are pinned and which are free; the physical relations that must hold among them. The free quantities form the state vector $\mathbf{z}$.

The question (step 1). A probability or an expectation of a function of the state, with the unknown on the left of the conditioning bar and everything treated as known on the right.

Table 2.10: Classification of questions by their mathematical form, with the general form of each category from Section 2.3.3 and the forms of a trend and of an expected consequence from Sections 2.3.4 and 2.3.5.

Engineering question	General form	Mathematical form	Derives to	Category
What is unknown, and how well is it known?	(2.9), (2.10)	$p(z_i \mid \mathbf{y})$, $\mathbb{E}[z_i \mid \mathbf{y}]$	a distribution, a mean, a credible interval: $\hat{z}_i$ and Σ_{ii}	estimation
Does a requirement hold, with the stated assurance?	(2.12)	$\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}) \geq \kappa$, f^* the requirement	the posterior integral of the indicator of the requirement; a normal tail under linearization	prediction under a threshold
What happens in the future?	(2.13)	$\int \Pr(A_{\text{future}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	the posterior predictive, e.g. $(b'/(b'+\tau))^{a'}$	prediction under a threshold
Is it rising, and when does it cross its limit?	(2.21)	$\Pr(\gamma > 0 \mid \mathbf{y})$; $\Pr(t^* \leq s \mid \mathbf{y}) = \Pr(f_s(\mathbf{z}) \geq f^* \mid \mathbf{y})$	the posterior read along time, $\mathbf{g}(t)^\mathsf{T}\hat{\mathbf{z}}$ and $\mathbf{g}(t)^\mathsf{T}\boldsymbol{\Sigma}\,\mathbf{g}(t)$; a crossing band by inversion	prediction under a threshold for the sign and a crossing by a date; estimation for a forecast and the distribution of a crossing time
Which explanation is correct?	(2.14), (2.15)	$\Pr(M_j \mid \mathbf{y}) \propto p(\mathbf{y} \mid M_j)\, \Pr(M_j)$	the evidences and the prior model probabilities; the Bayes factor	hypothesis comparison; diagnosis over a family of faults
What is the expected consequence?	(2.23)	$\mathcal{R}(a) = \mathbb{E}[L(a,\mathbf{z}) \mid \mathbf{y}] = \int L(a,\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	probability times consequence; $L(a,\hat{\mathbf{z}}) + \frac{1}{2}\operatorname{tr}(\mathbf{H}_a\boldsymbol{\Sigma})$	estimation of a derived quantity
What should be done or measured?	(2.23), (2.16)	$\arg\min_a \mathcal{R}(a)$; EVSI	the expected consequence of each action; the preposterior $\boldsymbol{\Sigma}_a$	design

The graph extended for the question (step 2). Scenario nodes, a second drawing when two structures compete, and the evidence attached as rows: readings, counts, handbook values as priors, each with the modeling assumption it rests on, what that assumption buys and how it can fail.

The posterior of the state (step 3). The rows combined into one objective; its minimum, the MAP state; its curvature, the information; and the covariance, derived and then computed for the plant, with the exact posterior beside them where the model admits one. It is built once and serves every question.

The answers (step 4). Each question read from that posterior, from the engineering question through its general and specialized forms down to its number: an entry of the state or the gradient of a derived quantity, the quantity at a time asked included; when drawings compete, a ratio of evidences; when actions compete, an expected consequence.

Evaluation and checks (step 5). The closed forms evaluated, and their approximations tested against the full calculation.

The report (step 6). The probability, the Bayes factor or the expected consequence with its band, or a crossing time with its band by inversion.

The next measurement (step 7). If the band is too wide, the candidate measurements ranked by the width each would remove, and a return to step 3 with the new row.

The examples of this chapter sit in it as follows. Stating precisely what the engineer needs to know or decide, and defining the events, uncertain quantities, future outcomes, hypotheses or actions that express it, is step 1; adding the full-load scenario node of the hall example, and attaching the readings and priors as rows with the assumptions behind them, is step 2. Step 3 builds the posterior the questions are read from: a fit for the most probable state and its Laplace covariance, a numerical integration, the evidence of each competing model, or another operation appropriate to the problem; for the hall it is the Gauss–Newton fit of three rows, and for the pump the Newton fit of the log-rate, whose objective is a gamma kernel and so gives the update in closed form as well. Step 4 writes each question in its general form and reduces it, through the specialized form, to posterior marginals, integrals, gradients, evidences or decision objectives, and evaluates them. Step 5 checks assumptions, identifiability, approximations, numerical convergence and the sensitivity of the answer to modeling choices, as the tolerance, exact-value and sampling checks of the hall example did. Step 6 communicates the answer, its uncertainty, its limitations and its engineering implications, and step 7 is the design question of Section 2.6.1.

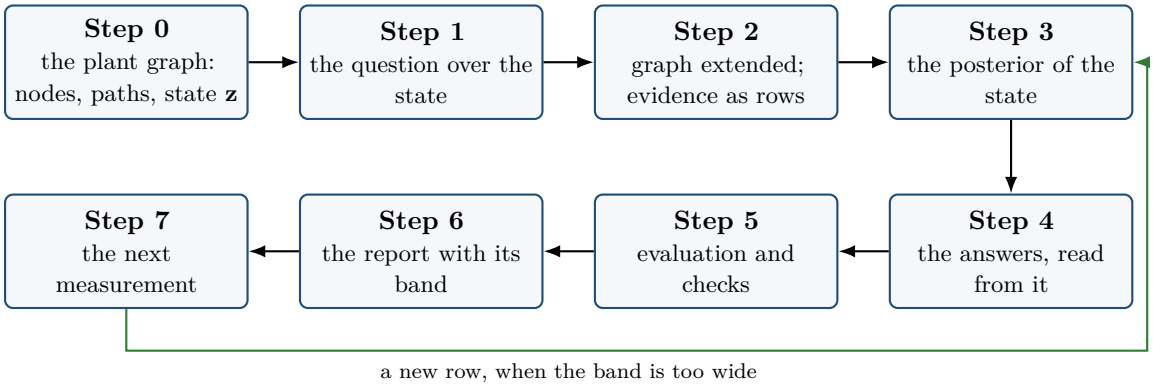


Figure 2.6: The derivation workflow of the book. Step 0 is built once per plant; steps 1 to 7 are one pass for one question; the green return arrow adds the measurement chosen at step 7 as a new row and rebuilds the posterior from step 3.

The workflow is broader than any particular software implementation. The derivation fixes what has to be calculated before a computational tool is chosen, and the tool then calculates it.

Principle 2.3. *The derivation tells the computational tool what must be calculated. A computational tool should not dictate which engineering question is asked.*

2.6.5 Common mistakes in probabilistic formulation

Five mistakes of formulation recur, and each is a confusion between two questions that look alike.

Common mistake

Confusing uncertainty about reliability with reliability itself. For a component with an unknown rate λ, the mission reliability at a given rate is $R(\tau \mid \lambda) = \mathrm{e}^{-\lambda\tau}$. The question $\mathbb{E}[R(\tau \mid \lambda) \mid \mathbf{y}]$ is the posterior predictive survival probability: how likely the mission is to succeed. The question $\Pr(R(\tau \mid \lambda) > 0.99 \mid \mathbf{y})$ is the posterior probability that the component's underlying mission reliability exceeds 99 %: how much support there is for a reliability requirement. They are different numbers and answer different questions.

Common mistake

Confusing prediction with diagnosis. When a motor overheats, the prediction $\Pr(T_{\text{motor}} > T_{\max} \mid \mathbf{y})$ concerns a temperature event, and the diagnosis $\Pr(F_{\text{cooling}} \mid \mathbf{y})$ concerns a fault explanation. They can use the same observations but require different calculations: a tail of a posterior in one model against a comparison of evidences across models.

Common mistake

Ignoring the time horizon. $\Pr(T > 100\,\text{h} \mid \mathbf{y})$ and $\Pr(T > 10\,000\,\text{h} \mid \mathbf{y})$ are not interchangeable statements about the reliability of the equipment; for the pump of Section 2.4.4 the first is 0.77 and the second, $(1500/11\,500)^4$, is 0.0003. A mission probability is incomplete without its time horizon and its operating conditions.

Common mistake

Forgetting dependence. In general $\Pr(A \cap B) \neq \Pr(A)\,\Pr(B)$; the product is justified only under independence. Shared power supplies, environmental conditions, maintenance procedures and common-cause failures create dependence between components, and Table 2.7 shows that a small shared cause can multiply a joint failure probability several times over.

Common mistake

Treating the model assumptions as certain. A calculated probability may be numerically precise while depending strongly on uncertain modeling assumptions. A temperature-exceedance probability computed under the steady-state model of equation (2.2) may be unsuitable for the temperature transient immediately after an outage of the cooling, which the steady-state model does not describe. Every probability is conditional on its model, $\Pr(A \mid \mathbf{y}, M)$, and must be interpreted in that context.

The questions of this chapter were all written with conditional probabilities and the posterior was taken as given; the next chapter, Chapter 3, derives conditional probability and Bayes' rule, separates aleatory variability from epistemic uncertainty, and shows how observations change the uncertainty about physical states and failure mechanisms through marginalization, dependence and conditional independence on engineering graphs.

2.7 What to retain

Definition 2.1 (Probabilistic question)**.** A probabilistic question about a plant is a probability or an expectation of a function of the state of its plant graph, written with the unknown on the left of the conditioning bar and the evidence and the model on the right, $\Pr(A \mid \mathbf{y}, M) = \int \mathbf{1}_A(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y}, M)\, \mathrm{d}\mathbf{z}$ for a present event and $\int \Pr(A_{\text{future}} \mid \mathbf{z}, M)\, p(\mathbf{z} \mid \mathbf{y}, M)\, \mathrm{d}\mathbf{z}$ for a future one. Its form decides its category, estimation, prediction under a threshold, hypothesis comparison, diagnosis or design, each with its general forms, equations (2.9) to (2.16); the question decides which of five computations answers it, the posterior of the state, the trend, the evidence, the preposterior or the expected consequence (Table 2.3), two of which carry general forms of their own, the trend form (2.21) and the expected-consequence form (2.23); and the plant enters only in the specialized form, through the entry or function of the state, the event or family of models, and the objects of the fit that carry the answer.

Table 2.11 lists the chapter's terms with their expressions, units and the subsection that defines each, grouped by their role.

Table 2.11: The terms of probabilistic questions and where each is defined.

Term	Expression	Units	Meaning	Defined in
Event	A, e.g. $T > 28\,^{\circ}\mathrm{C}$	–	a statement true or false	2.2.1
Conditional probability	$\Pr(A \mid \mathbf{y}, M)$	–	probability given evidence and model	2.2.3
Event probability from a density	$\int_5^7 p(G \mid \mathbf{y})\, \mathrm{d}G$	–	probability of the range $5 < G < 7$	2.2.4
Probability over the state	$\int \mathbf{1}_A(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	–	posterior expectation of the indicator	2.3.1
Posterior predictive	$\int \Pr(A_{\text{future}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	–	future event averaged over the state	2.3.2
General forms				
Estimation	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$; $\mathbb{E}[f \mid \mathbf{y}]$, $\mathrm{Var}(f \mid \mathbf{y})$	–	posterior of a derived quantity	2.3.3
Credible interval	$\Pr(f \le f_L \mid \mathbf{y}) = 0.025$, $\Pr(f \le f_U \mid \mathbf{y}) = 0.975$	–	central 95 % band; quantiles map through monotone f	2.3.3
Threshold and assurance	$\Pr(f(\mathbf{z}) \ge f^* \mid \mathbf{y}) \ge \kappa$	–	probability that the requirement $f \ge f^*$ holds, against κ	2.3.3
Predictive	$\int p(y_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	–	future observation averaged over the state	2.3.3
Comparison	$B_{AB} = p(\mathbf{y} \mid M_A)/p(\mathbf{y} \mid M_B)$	–	evidence ratio of two drawings	2.3.3
Diagnosis	$\Pr(M_j \mid \mathbf{y}) \propto p(\mathbf{y} \mid M_j)\, \Pr(M_j)$, ranked	–	comparison over a family of faults	2.3.3
Design	$U(a) = \mathbb{E}_{\tilde{y}_a \mid \mathbf{y}}[u(p(\mathbf{z} \mid \mathbf{y}, \tilde{y}_a))]$	varies	preposterior expectation	2.3.3
Preposterior covariance	$\boldsymbol{\Sigma}_a = (\boldsymbol{\Sigma}^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a)^{-1}$	varies	spread after a future measurement	2.3.3
Expected information gain	$\frac{1}{2}\ln(\det \boldsymbol{\Sigma} / \det \boldsymbol{\Sigma}_a)$	nat	what a measurement teaches	2.3.3
Trend	$\Pr(\gamma > 0 \mid \mathbf{y})$; $\Pr(t^* \le s \mid \mathbf{y}) = \Pr(f_s(\mathbf{z}) \ge f^* \mid \mathbf{y})$; $p(f_t(\mathbf{z}) \mid \mathbf{y})$; read with $\mathbf{g}(t)^{\mathsf{T}}\hat{\mathbf{z}}$, $\mathbf{g}(t)^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{g}(t)$	–	sign, crossing and forecast: the posterior read along time	2.3.4

Term	Expression	Units	Meaning	Defined in
Rows of a record in windows	$\mathbf{J} = \mathbf{D}\mathbf{X}$, $[\mathbf{\Lambda}]_{\gamma\gamma} = \sum_k d_k^2 (t_k - \bar{t})^2$	varies	a window's distance enters squared, windows add linearly	2.3.4
Expected consequence	$\mathcal{R}(a) = \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]$, $a^* = \arg\min_a \mathcal{R}(a)$; spread $\sqrt{\mathbf{g}_a^\mathsf{T} \mathbf{\Sigma}\, \mathbf{g}_a}$	varies	probability times consequence, minimized over actions	2.3.5
Curvature term	$L(a, \hat{\mathbf{z}}) + \frac{1}{2} \operatorname{tr}(\mathbf{H}_a \mathbf{\Sigma})$	varies	what the plug-in loss misses, in the state's coordinates	2.3.5
Computations				
Posterior	$p(\mathbf{z} \mid \mathbf{y}) \propto \mathrm{e}^{-\mathcal{L}(\mathbf{z})}$; $\mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$	varies	what estimation and prediction integrate	2.3.6
Evidence	$p(\mathbf{y} \mid M) = \int p(\mathbf{y} \mid \mathbf{z}, M)\, p(\mathbf{z} \mid M)\, \mathrm{d}\mathbf{z}$	–	normalizing constant of a model	2.3.6
Preposterior	$p(\mathbf{z} \mid \mathbf{y}, \tilde{y}_a)$ averaged over $p(\tilde{y}_a \mid \mathbf{y})$	varies	the expected posterior under a measurement not yet taken	2.3.6
Information adds	$\mathbf{J}^\mathsf{T}\mathbf{J} + \mathbf{J}_a^\mathsf{T}\mathbf{J}_a$	varies	stacked rows of a candidate	2.3.6
Update form	$\mathbf{\Sigma} - \mathbf{\Sigma}\mathbf{J}_a^\mathsf{T}(\mathbf{I} + \mathbf{J}_a\mathbf{\Sigma}\mathbf{J}_a^\mathsf{T})^{-1}\mathbf{J}_a\mathbf{\Sigma}$	varies	$\mathbf{\Sigma}_a$ when $\mathbf{\Sigma}$ is singular	2.3.6
Building blocks of the fit				
State	$\mathbf{z} = (T_1, T_2, g)^\mathsf{T}$	varies	every free quantity of the graph at once	2.4.1
Rows and objective	$r_k(\mathbf{z})/\sigma_k$; $\mathcal{L} = \frac{1}{2} r^\mathsf{T} r$	–	physics, reading and prior rows in standard deviations	2.4.1
Jacobian	$\mathbf{J} = \partial r / \partial \mathbf{z}$	varies	sensitivity of each row to each entry	2.4.1
Normal equations	$(\mathbf{J}^\mathsf{T}\mathbf{J})\, \Delta\mathbf{z} = -\mathbf{J}^\mathsf{T} r$	varies	Gauss–Newton step to $\hat{\mathbf{z}}$	2.4.1
Laplace covariance	$\mathbf{\Sigma} = (\mathbf{J}^\mathsf{T}\mathbf{J})^{-1}$	varies	band of every entry of the state	2.4.1
Ridge of exact rows	$\mathbf{\Sigma} \to \sigma_g^2 \mathbf{v}\mathbf{v}^\mathsf{T}$ as $\sigma_p \to 0$	varies	entries tied by the physics move as one	2.4.1
Band of a derived quantity	$\mathbf{g}^\mathsf{T}\mathbf{\Sigma}\mathbf{g}$, $\mathbf{g} = \partial f / \partial \mathbf{z}$	varies	delta method	2.4.1
Specialized forms and answers				
Requirement probability of the hall	$\Phi\big((\hat{z}_3 - \ln(G^*/G_{\mathrm{ref}}))/\sqrt{\Sigma_{33}}\big) = 0.717$	–	exact, the bound on T_2 mapped along the ridge to the third entry	2.4.2
Gamma update	$\mathrm{Gamma}(a + k,\ b + T_{\mathrm{tot}})$	–, h	posterior of a rate	2.4.3
Predictive reliability	$(b'/(b' + \tau))^{a'}$	–	survival averaged over the rate	2.4.4
Plug-in reliability	$\mathrm{e}^{-\mathbb{E}[\lambda \mid \mathbf{y}]\tau}$	–	survival at a point estimate; too low	2.4.5
Joint event	$\Pr(F_A \mid \mathbf{y})\, \Pr(F_B \mid F_A, \mathbf{y})$	–	intersection; product rule	2.4.7
Model posterior	$\Pr(M_j \mid \mathbf{y}) \propto p(\mathbf{y} \mid M_j)\, \Pr(M_j)$	–	probability of a model	2.5.1
Bayes factor of the cooling path	$B_{AB} = p(\mathbf{y} \mid M_1)/p(\mathbf{y} \mid M_0)$, $M_A = M_1$, $M_B = M_0$	–	relative support from the evidence	2.5.1
Expected loss	$\mathbb{E}[L \mid \mathbf{y}] = \int L(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	R$	probability-weighted loss	2.5.4
Optimal action	$a^* = \arg\min_a \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]$	–	decision by expected loss	2.5.5

Term	Expression	Units	Meaning	Defined in
Readings on the hall	$\boldsymbol{\Sigma}_a = \boldsymbol{\Sigma}/(1 + 0.21302\,n)$	varies	every entry shrinks along the ridge	2.6.1
Value of information	$\mathrm{EVSI} = \bar{L}_{\mathrm{now}} - \bar{L}_{\mathrm{after}}$	R$	what a measurement improves	2.6.2

2.8 Exercises

Exercise 2.1 (Formulation)**.** For each statement, name the category of the question, the expression that answers it and the computation beneath it. (a) What is the probability that a transformer survives the next 12 months? ($\Pr(T > 12\ \text{months} \mid \mathbf{y})$, a posterior predictive probability over the stated horizon; prediction under a threshold; the posterior.) (b) What is the failure rate of a cooling pump, given its maintenance records? ($p(\lambda \mid \mathbf{y})$ or a summary of it; estimation; the posterior.) (c) Are two simultaneous equipment failures more likely to have been caused by a common fault? ($\Pr(M_{\mathrm{common}} \mid \mathbf{y})$ against $\Pr(M_{\mathrm{independent}} \mid \mathbf{y})$, or their Bayes factor; hypothesis comparison; the evidence.) (d) Which inspection interval minimizes the expected operating cost? ($\arg\min_a \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]$ over the candidate intervals; a decision, in the design category; the expected consequence.) (e) Is the failure rate of a fleet rising, and in which year will it pass the limit written into its maintenance contract? ($\Pr(\gamma > 0 \mid \mathbf{y})$ and $\Pr(t^* \le s \mid \mathbf{y})$ over a state that carries the log-rate and its trend; prediction under a threshold; the trend.) (f) Of the questions A to D of Table 2.1, which two are answered from the same posterior without a new calculation of it? (A and B: an estimation and a threshold probability read from one posterior.)

Exercise 2.2 (Formulation)**.** A component has a constant failure rate λ with a known posterior, and a customer requires at least a 99 % probability of surviving a 24 h mission. (a) Formulate the posterior predictive probability of mission survival and the probability that the component's underlying mission reliability meets the requirement. ($\mathbb{E}[\mathrm{e}^{-24\lambda} \mid \mathbf{y}] = \int_0^\infty \mathrm{e}^{-24\lambda} p(\lambda \mid \mathbf{y})\,\mathrm{d}\lambda$; $\Pr(\mathrm{e}^{-24\lambda} \ge 0.99 \mid \mathbf{y}) = \Pr(\lambda \le -\ln 0.99/24 \mid \mathbf{y}) = \Pr(\lambda \le 4.188 \times 10^{-4}\,\mathrm{h}^{-1} \mid \mathbf{y})$.) (b) Evaluate both for the pump posterior $\mathrm{Gamma}(4, 1500\,\mathrm{h})$ of the pump example, using for the second the integer-shape identity $\Pr(\lambda \le x \mid \mathbf{y}) = 1 - \sum_{j=0}^{a'-1} \mathrm{e}^{-b'x}(b'x)^j/j!$. ($(1500/1524)^4 = 0.9385$; $b'x = 0.6281$ and the four Poisson terms 0.5336, 0.3352, 0.1053 and 0.0220 leave 0.0039.) (c) Explain why they differ. (The first averages survival over the uncertainty in the rate and asks how likely this mission is to succeed; the second asks how much posterior probability supports the requirement on the rate itself. Here the mission is likely to succeed, 0.94, while the requirement is almost certainly not met, because the posterior puts almost no mass below $4.2 \times 10^{-4}\,\mathrm{h}^{-1}$.)

Exercise 2.3 (Derivation)**.** The hall obeys $T = T_a + Q/G$ with $T_a = 18°\mathrm{C}$ and $Q = 60\,\mathrm{kW}$. (a) With the state, rows and $\boldsymbol{\Sigma}$ of equations (2.30) to (2.33), find the probability that the full-load temperature remains below 28°C from the normal of the entry T_2, then the exact value from the decreasing map between G and T_2 along the ridge. ($\Pr(T_2 < 28°\mathrm{C} \mid \mathbf{y}) \approx \Phi\big((28 - \hat{z}_2)/\sqrt{\Sigma_{22}}\big) = \Phi\big((28 - 27.231)/0.4615\big) = \Phi(1.667) = 0.952$. Exactly, along $\mathbf{v}$ the bound on T_2 is $G > Q_2/(28 - T_a) = 6\,\mathrm{kW/K}$, that is $g > \ln 6$ on the third entry, so $\Pr(T_2 < 28°\mathrm{C} \mid \mathbf{y}) = \int_6^\infty p(G \mid \mathbf{y})\,\mathrm{d}G = \Phi\big((\hat{z}_3 - \ln 6)/\sqrt{\Sigma_{33}}\big) = \Phi(1.601) = 0.945$; both are rows of Table 2.5.) (b) Now let both G and T_a be uncertain, with a joint posterior $p(G, T_a \mid \mathbf{y})$, and derive an integral for $\Pr(T > T_{\max} \mid \mathbf{y})$ without assuming that G and T_a are independent. (The event is $T_a + Q/G > T_{\max}$, so $\Pr(T > T_{\max} \mid \mathbf{y}) = \iint \mathbf{1}\big[T_a + Q/G > T_{\max}\big]\, p(G, T_a \mid \mathbf{y})\,\mathrm{d}G\,\mathrm{d}T_a$; the ambient becomes an entry of the state, and the joint density is required because the readings can leave it and the conductance correlated, a nonzero off-diagonal entry of $\boldsymbol{\Sigma}$.)

Exercise 2.4 (Derivation)**.** Two redundant cooling units, either of which can provide the required service, have mission failure probabilities $\Pr(F_A) = 0.08$ and $\Pr(F_B) = 0.04$. (a) Assuming independence, derive the probability that both fail and the system reliability. ($\Pr(F_A \cap F_B) = 0.08 \times 0.04 = 0.0032$; $R_{\text{system}} = 1 - 0.0032 = 0.9968$.) (b) Now let both units share a supply that fails with probability $q_S = 0.01$, with the marginals unchanged. Derive the joint failure probability from equation (2.45), the conditional probability $\Pr(F_B \mid F_A)$ and the system reliability. ($q_A = 0.07/0.99 = 0.07071$, $q_B = 0.03/0.99 = 0.03030$; $\Pr(F_A \cap F_B) = 0.01 + 0.99 \times 0.07071 \times 0.03030 = 0.01212$; $\Pr(F_B \mid F_A) = 0.01212/0.08 = 0.152$; $R_{\text{system}} = 0.98788$, the failure probability 3.8 times the independent value.)

Exercise 2.5 (Derivation)**.** The log failure rate of a fleet is carried as a straight line in time, $f_t(\mathbf{z}) = \alpha + \gamma(t - \bar{t})$ with $\bar{t} = 3.5$ years, and a fit to six yearly counts gives $\hat{\mathbf{z}} = (-1.5768,\ 0.1989)^{\mathsf{T}}$ and $\boldsymbol{\Sigma}$ with entries $\Sigma_{\alpha\alpha} = 0.021811$, $\Sigma_{\alpha\gamma} = -0.004328$ and $\Sigma_{\gamma\gamma} = 0.007259$, the fit of Chapter 12. The contract limit is 0.5 failures per unit-year, $f^* = \ln 0.5$. (a) From the trend form (2.21) under the Laplace posterior, find the probability that the rate is rising. ($\Phi(0.1989/\sqrt{0.007259}) = \Phi(2.335) = 0.990$.) (b) Find the crossing time at the MAP state and say why it is the median of the crossing time. ($t^*(\hat{\mathbf{z}}) = 3.5 + (\ln 0.5 + 1.5768)/0.1989 = 7.94$ years; the argument of the normal tail in equation (2.22) is zero at $s = t^*(\hat{\mathbf{z}})$, so the distribution function is one half there.) (c) Find the probability that the rate has crossed the limit by $t = 8$. ($\mathbf{g}(8) = (1,\ 4.5)^{\mathsf{T}}$, $\mathbf{g}(8)^{\mathsf{T}}\hat{\mathbf{z}} = -0.6816$, $\mathbf{g}(8)^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{g}(8) = 0.021811 - 9 \times 0.004328 + 20.25 \times 0.007259 = 0.1299$, so $\Phi\big((-0.6816 + 0.6931)/0.3604\big) = \Phi(0.032) = 0.513$.) (d) Explain why the crossing time has no posterior mean. ($t^* = \bar{t} + (f^* - \alpha)/\gamma$ divides by the trend entry, whose Gaussian posterior has a positive density at $\gamma = 0$, and the integral of $1/\gamma$ against such a density diverges.)

Exercise 2.6 (Application)**.** A plant has a probability of shutdown of 0.004 over the next operating period, and each shutdown costs R$ 250 000. An intervention costs R$ 400 and reduces the shutdown probability to 0.001. (a) Derive the expected total cost with and without the intervention, and state which action minimizes it. (The expected consequence (2.23) of each action, probability times consequence plus the certain cost. Without: $0.004 \times$ R$ 250 000 = R$ 1000; with: R$ 400 $+ 0.001 \times$ R$ 250 000 = R$ 650; the intervention minimizes expected cost, with an expected saving of R$ 350.) (b) Find the cost of the intervention at which the decision would reverse. ($(0.004 - 0.001) \times$ R$ 250 000 = R$ 750.)

Exercise 2.7 (Application)**.** Two models are competing explanations of a system fault, with $\Pr(M_0) = \Pr(M_1) = 0.5$ and evidences $p(\mathbf{y} \mid M_0) = 0.02$ and $p(\mathbf{y} \mid M_1) = 0.08$. (a) Derive the Bayes factor B_{10} and the posterior probability $\Pr(M_1 \mid \mathbf{y})$. ($B_{10} = 0.08/0.02 = 4$, $\ln B_{10} = 1.386$ nats; $\Pr(M_1 \mid \mathbf{y}) = 0.08 \times 0.5/(0.08 \times 0.5 + 0.02 \times 0.5) = 0.80$, assuming that the two models exhaust the possibilities under consideration.) (b) Repeat with $\Pr(M_1) = 0.2$, and say what the comparison of the two answers shows. ($0.08 \times 0.2/(0.08 \times 0.2 + 0.02 \times 0.8) = 0.50$; the Bayes factor is the same 4 in both cases, because it depends only on the evidence, while the posterior model probability also depends on the prior.)

Chapter 3

Uncertainty and Conditional Probability

3.1 Uncertainty and evidence in reliability questions

Chapter 1 represented a physical system as a plant graph whose free quantities form the state vector $\mathbf{z}$, and Chapter 2 wrote engineering questions as probabilities and expectations over that state, with the unknown on the left of the conditioning bar and the evidence on the right. This chapter treats the mechanism that connects the two: the way a probability changes when information arrives. Four situations recur through it. A pump has an estimated failure rate of $0.002\,\mathrm{h}^{-1}$, and the probability that it survives a mission is wanted. A cooling unit carries a high-temperature alarm, the alarm has sounded, and the probability that the unit is degraded is wanted. Two pumps share one electrical supply, and the probability that both fail in the same mission is wanted. The equipment hall of the earlier chapters has a temperature sensor, and the probability that the hall exceeds its limit after one reading is wanted. The question of the chapter is

how new evidence changes what is known about an engineering system, and how dependencies between components and between physical quantities affect a probability calculation.

The chapter answers it in five steps. It separates the variability of outcomes from the incompleteness of knowledge and names the distributions that represent each. It derives conditional probability from counts of missions, the multiplication rule from the definition, and Bayes' theorem from the multiplication rule, and works a fault diagnosis from an alarm. It derives marginalization and the law of total probability, and shows with a common cause how an independence assumption can understate a joint failure probability by a factor of nine. It writes the hall as a factorized probability model, derives the posterior of a state vector under Gaussian rows, by the normal equations and by the conditional of a joint Gaussian, applies it to the hall temperature after one reading and to the two-entry hall of Chapter 1, and answers a threshold question with the result. Finally, it derives the law of total variance, which separates the randomness of future failures from the uncertainty of a rate, and the covariance of a joint posterior, which decides how uncertain a derived quantity is.

Jargon: aleatory and epistemic uncertainty, conditional probability, prior, likelihood, evidence, posterior, marginalization, nuisance quantity, conditional independence, common cause

Aleatory uncertainty is the variability of outcomes that the model treats as random even when its parameters are known, such as the time at which a pump with a known failure rate next fails. *Epistemic* uncertainty is incomplete knowledge of quantities that have one true value, such as the failure rate itself; more evidence can reduce it. A *conditional probability* $\Pr(A \mid B)$ is the probability of event A among the situations in which event B occurs. In Bayes' theorem the *prior* is the probability of a hypothesis before an observation, the *likelihood* is the probability of the observation under the hypothesis, the *evidence* (also *marginal likelihood*) is the probability of the observation averaged over the hypotheses, and the *posterior* is the probability of the hypothesis after the observation. *Marginalization* removes an unknown from a joint distribution by summing or integrating over it; an unknown removed this way because the question does not ask about it is a *nuisance quantity*. Two quantities are *conditionally independent* given a third when, once the third is known, learning one says nothing more about the other. A *common cause* is a shared condition, such as an electrical disturbance, that raises the failure probability of several components at once.

Key idea

Conditioning on evidence restricts attention to the situations consistent with it and renormalizes; Bayes' theorem does this for a hypothesis, marginalization averages over every unknown the question does not name, and the dependence among unknowns, through a common cause or a shared physical relation, is what lets one observation inform another quantity and what an independence assumption silently discards.

This section separates the two kinds of uncertainty that a reliability calculation carries, shows why the separation changes a mission reliability, and names the distributions that represent what is uncertain.

3.1.1 Aleatory variability

Suppose the failure rate of an industrial pump is known exactly and constant, $\lambda = 0.002\,\mathrm{h}^{-1}$. Even then the time of the next failure cannot be predicted under a stochastic failure-time model. Let T_f be the time to failure, measured from the start of a mission, and let it be exponential given the rate, $T_f \mid \lambda \sim \text{Exponential}(\lambda)$, with density $\lambda\,\mathrm{e}^{-\lambda t}$ for $t \geq 0$. The probability that the pump survives beyond t hours is the integral of that density from t to infinity,

$$\Pr(T_f > t \mid \lambda) = \int_t^\infty \lambda\,\mathrm{e}^{-\lambda s}\,\mathrm{d}s = \mathrm{e}^{-\lambda t}, \tag{3.1}$$

where s is the variable of integration. For a 100 h mission,

$$\Pr(T_f > 100\,\mathrm{h} \mid \lambda = 0.002\,\mathrm{h}^{-1}) = \mathrm{e}^{-0.002\times 100} = \mathrm{e}^{-0.2} = 0.8187.$$

With a perfectly known rate there is still a probability $1 - 0.8187 = 0.1813$, or 18.13 %, that the pump fails during the mission. This uncertainty arises from the variability that the failure process itself represents, and it is called *aleatory*. Random equipment failures, variations in operating demand and environmental disturbances that cannot be forecast are uncertainties of this kind.

Figure 3.1(a) shows it. The horizontal axis numbers twelve pumps that share the same known rate, and the vertical axis is the time at which each fails. The bars are the twelve quantiles of the exponential distribution with $\lambda = 0.002\,\mathrm{h}^{-1}$, placed in an arbitrary order, and

the dashed line marks the mean life $1/\lambda = 500\,\mathrm{h}$. Identical pumps with an identical, exactly known rate fail at times that range from 21 to 1589 h: knowing the rate fixes the distribution of the outcome, not the outcome.

3.1.2 Epistemic uncertainty

Now suppose the failure rate is not known. The pump may have operated for only a short period, or its maintenance records may be incomplete. What is known about the rate is then a probability distribution given the evidence, $p(\lambda \mid \mathbf{y})$. The rate might plausibly be 0.001, 0.002 or $0.004\,\mathrm{h}^{-1}$, with different posterior probabilities, or a density, assigned to those values. This is *epistemic* uncertainty: uncertainty caused by incomplete knowledge of the characteristics of the system, which additional observations can reduce.

Figure 3.1(b) draws such a distribution. The horizontal axis is the failure rate in units of $10^{-3}\,\mathrm{h}^{-1}$ and the vertical axis the posterior density. The curve is the gamma posterior of Chapter 2, obtained from the prior $\mathrm{Gamma}(2, 500\,\mathrm{h})$ and a record of $k = 2$ failures in an exposure $T_{\mathrm{tot}} = 1000\,\mathrm{h}$, which is $\mathrm{Gamma}(4, 1500\,\mathrm{h})$; the dashed line is its mode, $(4-1)/1500\,\mathrm{h} = 0.002\,\mathrm{h}^{-1}$, and the three dots are the three plausible values named above. The rate is a single number; the spread of the curve is the spread of knowledge about it.

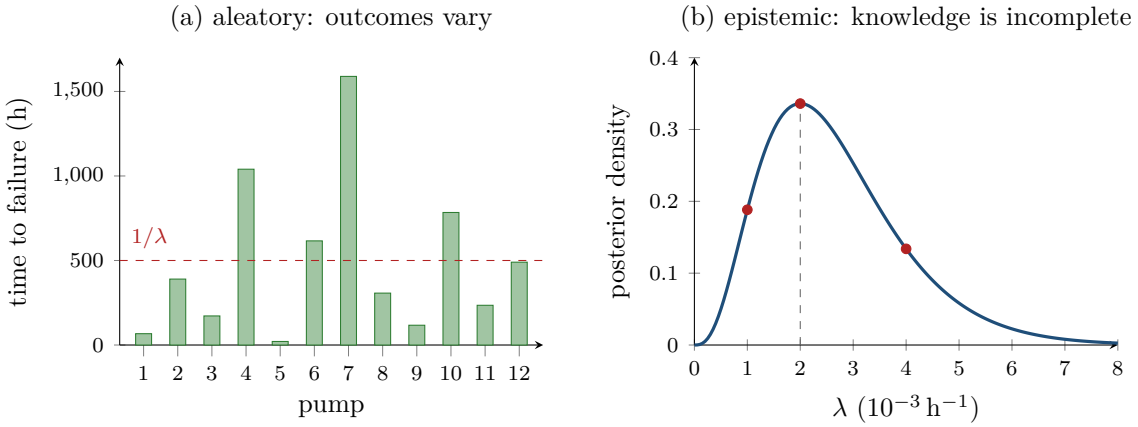


Figure 3.1: The two kinds of uncertainty. (a) Times to failure of twelve pumps with the same known rate $\lambda = 0.002\,\mathrm{h}^{-1}$: the exponential quantiles $-\ln(1-(i-0.5)/12)/\lambda$, $i = 1, \ldots, 12$, in an arbitrary order; dashed: the mean life $1/\lambda = 500\,\mathrm{h}$. (b) The posterior $\mathrm{Gamma}(4, 1500\,\mathrm{h})$ of an unknown rate, density $(1.5^4/3!)\,x^3 \mathrm{e}^{-1.5x}$ with $x = 1000\lambda$; dashed: its mode $0.002\,\mathrm{h}^{-1}$; dots: the plausible values 0.001, 0.002 and $0.004\,\mathrm{h}^{-1}$.

3.1.3 The distinction and why it matters

In a real system the distinction can depend on what information is available and on how the model is defined. A future load may be modeled as random, aleatory, while the load at the present moment is a fixed but poorly measured quantity, epistemic. The distinction is therefore a statement about the model, and it changes the calculation.

If the rate is known, the mission reliability is equation (3.1), $R(t) = \mathrm{e}^{-\lambda t}$. If the rate is uncertain, the probability that the pump survives to t is the conditional reliability averaged over what is known about the rate, which is the law of total probability derived in Section 3.4.2:

$$R_{\mathrm{pred}}(t \mid \mathbf{y}) = \int_0^\infty \Pr(T_f > t \mid \lambda)\, p(\lambda \mid \mathbf{y})\, \mathrm{d}\lambda = \int_0^\infty \mathrm{e}^{-\lambda t}\, p(\lambda \mid \mathbf{y})\, \mathrm{d}\lambda. \tag{3.2}$$

For three plausible values the integral is a sum. With the illustrative probabilities 0.25, 0.50 and 0.25 on 0.001, 0.002 and $0.004\,\mathrm{h}^{-1}$, the 100 h predictive reliability is

$$R_{\mathrm{pred}}(100\,\mathrm{h}) = 0.25\,\mathrm{e}^{-0.1} + 0.50\,\mathrm{e}^{-0.2} + 0.25\,\mathrm{e}^{-0.4} = 0.2262 + 0.4094 + 0.1676 = 0.8032,$$

against $\mathrm{e}^{-0.225} = 0.7985$ at the mean rate $0.00225\,\mathrm{h}^{-1}$. For the gamma posterior $\mathrm{Gamma}(a, b)$, with density $b^a\lambda^{a-1}\mathrm{e}^{-b\lambda}/\Gamma(a)$, the integral has a closed form. Combining the two exponentials and using the standard gamma integral $\int_0^\infty \lambda^{a-1}\mathrm{e}^{-s\lambda}\,\mathrm{d}\lambda = \Gamma(a)/s^a$ with $s = b + t$,

$$R_{\mathrm{pred}}(t \mid \mathbf{y}) = \frac{b^a}{\Gamma(a)} \int_0^\infty \lambda^{a-1}\,\mathrm{e}^{-(b+t)\lambda}\,\mathrm{d}\lambda = \frac{b^a}{\Gamma(a)}\,\frac{\Gamma(a)}{(b+t)^a} = \left(\frac{b}{b+t}\right)^a. \tag{3.3}$$

For $\mathrm{Gamma}(4, 1500\,\mathrm{h})$ and $t = 100\,\mathrm{h}$ it gives $(1500/1600)^4 = 0.7725$, against $\mathrm{e}^{-100\times 4/1500} = \mathrm{e}^{-0.2667} = 0.7659$ at the posterior mean rate. In both cases the average of $\mathrm{e}^{-\lambda t}$ exceeds its value at the average rate, because the function is convex in λ.

The future failure remains random even when the uncertainty about the rate is reduced. As evidence accumulates, $p(\lambda \mid \mathbf{y})$ concentrates near one value λ_0 and equation (3.2) tends to $\mathrm{e}^{-\lambda_0 t}$; the probability of failure in the mission tends to $1 - \mathrm{e}^{-\lambda_0 t}$, not to zero.

> **Principle 3.1.** *Collecting more data can reduce the uncertainty about a failure rate without eliminating the randomness of future failures: evidence removes epistemic uncertainty, and the aleatory part remains as long as the failure process is modeled as random.*

3.1.4 Distributions of what is uncertain

Three mathematical objects represent what is uncertain, and Figure 3.2 draws one of each.

A discrete unknown. Suppose a pump is either healthy or degraded, and define the indicator $F = 1$ if it is degraded and $F = 0$ if it is healthy. The assignment $\Pr(F = 1) = 0.10$, $\Pr(F = 0) = 0.90$ is a discrete probability distribution: each value receives a probability, and the probabilities sum to one. Panel (a) of the figure has the two values on the horizontal axis and their probabilities as bars.

A continuous unknown. Suppose the effective efficiency η of the pump is unknown, and represent it by a density $p(\eta)$. A probability over a range is the integral of the density over that range,

$$\Pr(0.75 < \eta < 0.85) = \int_{0.75}^{0.85} p(\eta)\,\mathrm{d}\eta.$$

For a continuous variable the density itself is not a probability: it is a probability per unit of the variable, and it can exceed one. Panel (b) uses the illustrative density $\mathcal{N}(0.80, 0.04^2)$. The horizontal axis is the efficiency and the vertical axis the density, whose peak is $1/(0.04\sqrt{2\pi}) = 9.97$; the shaded area between 0.75 and 0.85 is $\Phi(1.25) - \Phi(-1.25) = 0.789$, the probability of that range.

A joint distribution. Suppose a cooling system has two unknown quantities, the thermal conductance G of its heat-removal path and the ambient temperature T_a. Their *joint* distribution $p(G, T_a)$ describes the uncertainty about both together, including any dependence between them. The joint distribution is essential when an engineering output depends on several uncertain inputs, as the hall temperature does,

$$T = T_a + \frac{Q}{G},$$

with the heat load Q in kW and G in kW/K. The individual distributions $p(G)$ and $p(T_a)$ are not always enough to determine the distribution of T; how G and T_a vary together is needed as well. Panel (c) draws two joint distributions with the same marginals, G with mean 3 kW/K and standard deviation 0.15 kW/K, T_a with mean 20°C and standard deviation 1 K, one with correlation $\mathrm{corr}(G, T_a) = 0$ and one with $\mathrm{corr}(G, T_a) = 0.8$. The horizontal axis is G, the vertical axis T_a, and each ellipse encloses 95 % of its distribution. Both ellipses span the same ranges of G and of T_a, so the marginals cannot tell them apart; Section 3.6.4 shows that at $Q = 30$ kW the hall temperature has a standard deviation of 1.12 K under the first and 0.67 K under the second.

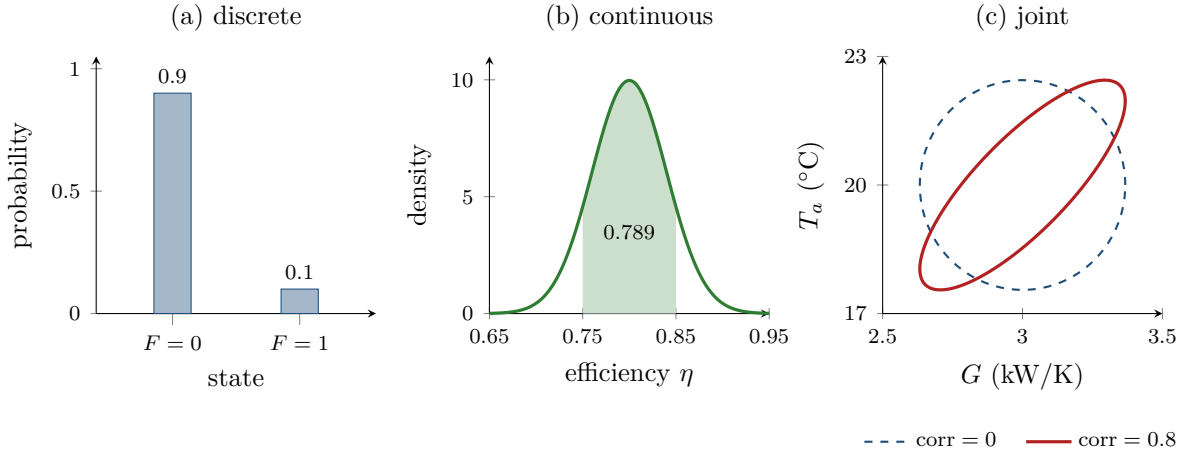


Figure 3.2: Distributions of what is uncertain (illustrative). (a) A discrete unknown: $\Pr(F = 1) = 0.10$ for a degraded pump. (b) A continuous unknown: the density $\mathcal{N}(0.80, 0.04^2)$ of an efficiency; shaded, $\Pr(0.75 < \eta < 0.85) = 0.789$. (c) Two joint distributions of G and T_a with the same normal marginals (means 3 kW/K and 20°C, standard deviations 0.15 kW/K and 1 K) and correlations 0 and 0.8; the ellipses enclose 95 % of each ($\chi^2_{2,0.95} = 5.991$).

3.2 Conditional probability and Bayes' theorem

Conditional probability is the foundation of every update in this book. This section derives it from counts of missions, states the general definition and the multiplication rule that follows from it, warns against reversing the conditioning, and derives Bayes' theorem.

3.2.1 Conditional probability from mission counts

Let F_A be the event that pump A fails during a mission and F_B the event that pump B fails during the same mission. The question *given that pump A failed, what is the probability that pump B also failed* is written $\Pr(F_B \mid F_A)$. The phrase "given that" restricts attention to the situations in which F_A occurred.

Imagine 1000 comparable missions involving the two pumps, recorded in Table 3.1. The rows say whether pump A failed, the columns whether pump B failed, and the margins are the totals: both failed in 20 missions, only pump A in 60, only pump B in 20, and neither in 900.

From these counts,

$$\Pr(F_A) = \frac{80}{1000} = 0.08, \qquad \Pr(F_B) = \frac{40}{1000} = 0.04, \qquad \Pr(F_A \cap F_B) = \frac{20}{1000} = 0.02.$$

Among the 80 missions in which pump A failed, pump B also failed in 20. Therefore

$$\Pr(F_B \mid F_A) = \frac{20}{80} = 0.25.$$

Table 3.1: Outcomes of 1000 comparable missions of two pumps (illustrative).

	pump B fails	pump B survives	total
pump A fails	20	60	80
pump A survives	20	900	920
total	40	960	1000

The conditional probability is 25 %, against an unconditional probability of 4 %: in these observations the two failures are strongly associated.

3.2.2 The definition

Dividing the numerator and the denominator of 20/80 by the total number of missions turns the count ratio into a ratio of probabilities,

$$\Pr(F_B \mid F_A) = \frac{20/1000}{80/1000} = \frac{\Pr(F_A \cap F_B)}{\Pr(F_A)}.$$

For any two events A and B with $\Pr(B) > 0$ the definition is

$$\Pr(A \mid B) = \frac{\Pr(A \cap B)}{\Pr(B)}. \tag{3.4}$$

Conditioning on B restricts the sample space to the outcomes in B and renormalizes, so that the restricted probabilities again sum to one. The definition requires the conditioning event to have positive probability. Two events are *independent* when $\Pr(A \cap B) = \Pr(A)\,\Pr(B)$, which by equation (3.4) is the same as $\Pr(A \mid B) = \Pr(A)$: learning B does not change the probability of A.

Table 3.2 works the definition for four joint probabilities with the marginals of the missions held fixed, $\Pr(F_A) = 0.08$ and $\Pr(F_B) = 0.04$. The first column is the probability that both pumps fail, the next two are the conditional probabilities in the two directions, and the last says what relation between the failures the row describes. The third row is the record of Table 3.1: $\Pr(F_B \mid F_A) = 0.02/0.08 = 0.25$ and $\Pr(F_A \mid F_B) = 0.02/0.04 = 0.50$. The second row is the independent case, in which $\Pr(F_B \mid F_A) = 0.0032/0.08 = 0.04 = \Pr(F_B)$. The joint probability cannot exceed the smaller marginal, 0.04; at that bound every failure of pump B occurs together with a failure of pump A.

Table 3.2: Conditional probabilities from equation (3.4) for four joint failure probabilities, with $\Pr(F_A) = 0.08$ and $\Pr(F_B) = 0.04$ held fixed.

$\Pr(F_A \cap F_B)$	$\Pr(F_B \mid F_A)$	$\Pr(F_A \mid F_B)$	Relation between the failures
0	0	0	mutually exclusive: never together
0.0032	0.04	0.08	independent: $\Pr(F_B \mid F_A) = \Pr(F_B)$
0.02	0.25	0.50	the recorded missions: positively associated
0.04	0.50	1.00	pump B fails only together with pump A

3.2.3 The multiplication rule

Multiplying equation (3.4), written for $\Pr(B \mid A)$, by $\Pr(A)$ gives the multiplication rule,

$$\Pr(A \cap B) = \Pr(B \mid A)\ \Pr(A). \tag{3.5}$$

Conditioning in the other direction gives $\Pr(A \cap B) = \Pr(A \mid B)\ \Pr(B)$, and the two expressions for the same intersection are equal:

$$\Pr(B \mid A)\ \Pr(A) = \Pr(A \mid B)\ \Pr(B). \tag{3.6}$$

For the missions, $0.25 \times 0.08 = 0.02 = 0.50 \times 0.04$. This relation is the starting point of Bayes' theorem.

Reversing the conditioning. The two conditional probabilities of the missions differ: $\Pr(F_B \mid F_A) = 0.25$, but $\Pr(F_A \mid F_B) = 20/40 = 0.50$. The first asks about pump B among the missions in which pump A failed; the second asks about pump A among the missions in which pump B failed. Equation (3.6) shows that they are equal only when $\Pr(A) = \Pr(B)$.

Common mistake

Treating $\Pr(A \mid B)$ and $\Pr(B \mid A)$ as the same number. Reversing the conditioning changes the question. In fault diagnosis the probability of an alarm given a fault is known from the design of the alarm, and the probability of a fault given an alarm is what the operator needs; the second follows from the first only through Bayes' theorem and the prior probability of the fault.

3.2.4 Bayes' theorem

Dividing both sides of equation (3.6) by $\Pr(B) > 0$ gives Bayes' theorem for discrete events,

$$\Pr(A \mid B) = \frac{\Pr(B \mid A)\ \Pr(A)}{\Pr(B)}. \tag{3.7}$$

For the missions, $\Pr(F_A \mid F_B) = 0.25 \times 0.08/0.04 = 0.50$, the value read from the counts. In words,

$$\text{posterior} = \frac{\text{likelihood} \times \text{prior}}{\text{evidence}},$$

and the four terms have distinct roles when A is a hypothesis about the system and B an observation. The *prior* $\Pr(A)$ is what was believed about the hypothesis before the observation. The *likelihood* $\Pr(B \mid A)$ is how probable the observation is when the hypothesis holds. The *evidence* $\Pr(B)$ is the overall probability of the observation, averaged over the hypotheses; it is the normalizing constant of the posterior and, for a whole model M, the quantity the notation writes $p(\mathbf{y} \mid M)$. The *posterior* $\Pr(A \mid B)$ is what is believed about the hypothesis after the observation is accounted for.

3.3 Bayesian diagnosis of a cooling fault

Bayes' theorem turns the known behavior of an alarm into the probability of the fault the alarm is meant to detect. This section works the diagnosis step by step, explains why the answer is far below the detection probability of the alarm, and treats the case in which the alarm stays silent.

3.3.1 The diagnosis problem

An industrial cooling unit carries a high-temperature alarm designed to detect degradation of the cooling system. Define the events

$$F = \{\text{the cooling system is degraded}\}, \qquad H = \{\text{the high-temperature alarm occurs}\}.$$

Engineering records suggest that 5 % of the relevant operating periods have a degraded cooling system, $\Pr(F) = 0.05$. The alarm activates in 90 % of degraded conditions, $\Pr(H \mid F) = 0.90$, its detection probability; it also produces false alarms, in 10 % of healthy conditions, $\Pr(H \mid \neg F) = 0.10$, its false-alarm probability. An alarm is observed, and the question is the probability that the cooling system is in fact degraded, $\Pr(F \mid H)$.

The first step writes Bayes' theorem, equation (3.7), for the hypothesis F and the observation H:

$$\Pr(F \mid H) = \frac{\Pr(H \mid F)\ \Pr(F)}{\Pr(H)}.$$

The numerator is known, $\Pr(H \mid F)\ \Pr(F) = 0.90 \times 0.05 = 0.045$. The denominator, the probability of an alarm, is not yet known.

3.3.2 The total probability of an alarm

An alarm can occur in two mutually exclusive situations: the system is degraded and the alarm activates, or the system is healthy and the alarm activates falsely. In this simplified model the two exhaust the cases, so the alarm event is the union of two disjoint events and its probability is their sum,

$$\Pr(H) = \Pr(H \cap F) + \Pr(H \cap \neg F).$$

The multiplication rule, equation (3.5), writes each intersection as a likelihood times a prior:

$$\begin{aligned} \Pr(H) &= \Pr(H \mid F)\ \Pr(F) + \Pr(H \mid \neg F)\ \Pr(\neg F) \\ &= 0.90 \times 0.05 + 0.10 \times 0.95 = 0.045 + 0.095 = 0.14. \end{aligned} \tag{3.8}$$

Alarms occur in 14 % of operating periods; two thirds of them, 0.095 of the 0.14, come from healthy periods.

3.3.3 The posterior fault probability

The third step divides,

$$\Pr(F \mid H) = \frac{0.045}{0.14} = 0.3214. \tag{3.9}$$

Figure 3.3 compares the probability of degradation before and after the observation. Each row is one state of knowledge, the horizontal scale is a probability from 0 to 1, and the green part of each bar is the probability of degradation. The alarm raises it from 5 % to 32.1 %; it increases suspicion substantially, but it does not establish that a fault is present. The third row is the case without an alarm, derived in Section 3.3.4.

The base rate. The answer is not 90 %. The alarm detects 90 % of degraded conditions, but degraded conditions are uncommon, and many alarms occur while the cooling system is healthy. Table 3.3 makes this concrete for 1000 operating periods. The rows are the actual condition of the system, the columns whether the alarm sounded, and each entry is the expected number of periods: of the 50 degraded periods, 90 %, or 45, raise the alarm; of the 950 healthy periods, 10 %, or 95, raise it falsely. Among the 140 periods with an alarm only 45 have a degraded cooling system, so $\Pr(F \mid H) = 45/140 = 0.3214$. The table and Bayes' theorem give the same result, because the table is equation (3.8) multiplied by 1000.

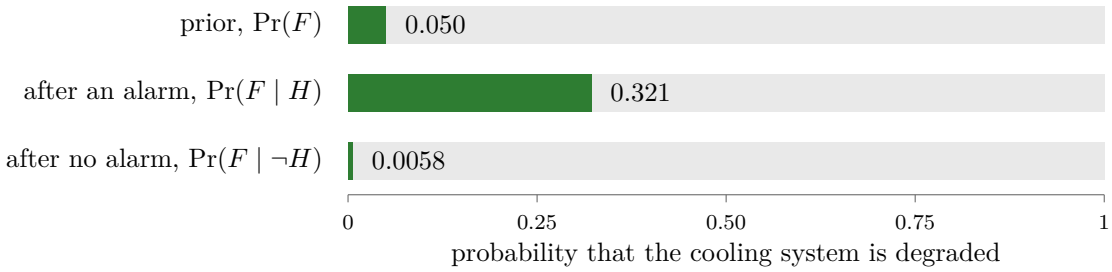


Figure 3.3: What an alarm tells about degradation. Prior $\Pr(F) = 0.05$; alarm with detection probability $\Pr(H \mid F) = 0.90$ and false-alarm probability $\Pr(H \mid \neg F) = 0.10$. After an alarm, $0.045/0.14 = 0.321$ (equation (3.9)); after no alarm, $0.005/0.86 = 0.0058$ (equation (3.11)).

Table 3.3: Expected outcomes of 1000 operating periods with $\Pr(F) = 0.05$, $\Pr(H \mid F) = 0.90$ and $\Pr(H \mid \neg F) = 0.10$.

Actual condition	alarm	no alarm	total
degraded	45	5	50
healthy	95	855	950
total	140	860	1000

The odds form. Writing Bayes' theorem for F and for $\neg F$ with the same observation and dividing the two, the evidence $\Pr(H)$ cancels:

$$\frac{\Pr(F \mid H)}{\Pr(\neg F \mid H)} = \frac{\Pr(H \mid F)}{\Pr(H \mid \neg F)} \cdot \frac{\Pr(F)}{\Pr(\neg F)}. \tag{3.10}$$

Posterior odds equal the likelihood ratio times the prior odds. Here the likelihood ratio is $0.90/0.10 = 9$ and the prior odds are $0.05/0.95 = 0.05263$, so the posterior odds are 0.4737 and the posterior probability is $0.4737/(1 + 0.4737) = 0.3214$, as before. The alarm multiplies the odds of degradation by nine whatever the base rate; the base rate decides where that multiplication starts.

3.3.4 The case without an alarm

The same derivation applies to the observation that the alarm does not activate. Bayes' theorem gives

$$\Pr(F \mid \neg H) = \frac{\Pr(\neg H \mid F)\ \Pr(F)}{\Pr(\neg H)},$$

with $\Pr(\neg H \mid F) = 1 - 0.90 = 0.10$, the probability that a degraded system goes undetected, and, by the same total probability as equation (3.8),

$$\Pr(\neg H) = 0.10 \times 0.05 + 0.90 \times 0.95 = 0.005 + 0.855 = 0.86.$$

Therefore

$$\Pr(F \mid \neg H) = \frac{0.005}{0.86} = 0.00581. \tag{3.11}$$

The absence of an alarm reduces the fault probability from 5 % to about 0.58 %, the third row of Figure 3.3; in odds, the silence multiplies the prior odds by $0.10/0.90 = 1/9$. In this example an alarm substantially raises suspicion, and its absence is considerable evidence that the system is healthy.

Interpretation. Bayesian diagnosis does not only ask whether a sensor crossed a threshold. It evaluates how strongly the observation supports competing explanations, weighted by their prior plausibility. The distinction is central to reliability engineering because many abnormal measurements are not uniquely associated with one physical fault: a high hall temperature can be caused by a reduced cooling capacity, an increased load, a change of ambient temperature or a biased sensor. Each of these is a hypothesis with its own prior and its own likelihood for the same reading, and the inference becomes a comparison among them, the diagnosis category of the preface, which later chapters on fault isolation develop.

3.4 Marginalization and dependence

The probability of an alarm in Section 3.3.2 was computed without knowing whether the system was degraded. This section derives that operation in general, for discrete and continuous unknowns and for nuisance quantities, contrasts it with optimization, and then shows how a common cause makes two failures dependent and what an independence assumption costs.

3.4.1 Discrete marginalization

Suppose a cooling unit can be in one of three operating conditions, healthy, degraded or failed, and let S denote the condition present. Table 3.4 gives, for each state, its probability and the probability that the alarm H occurs in it; the last two columns are computed below.

Table 3.4: Three operating conditions of a cooling unit: probabilities, alarm likelihoods, the joint probabilities of state and alarm, and the probabilities of the states given an alarm.

State s	$\Pr(S = s)$	$\Pr(H \mid S = s)$	$\Pr(H \cap S = s)$	$\Pr(S = s \mid H)$
healthy	0.85	0.05	0.0425	0.272
degraded	0.12	0.70	0.0840	0.538
failed	0.03	0.99	0.0297	0.190
total	1.00		0.1562	1.000

An alarm can arise under any of the three states. Because the states are mutually exclusive and exhaustive, the events $H \cap \{S = s\}$ are disjoint and their union is H, so $\Pr(H) = \sum_s \Pr(H \cap S = s)$. The multiplication rule writes each term as a likelihood times a prior:

$$\Pr(H) = \sum_s \Pr(H \mid S = s)\ \Pr(S = s). \tag{3.12}$$

This is the *law of total probability*. Substituting,

$$\Pr(H) = 0.05 \times 0.85 + 0.70 \times 0.12 + 0.99 \times 0.03 = 0.0425 + 0.0840 + 0.0297 = 0.1562.$$

The probability of an alarm has been computed without knowing which state is present. That is what marginalization does: it accounts for the uncertainty in variables that are not part of the final question. The same sum is the denominator of Bayes' theorem over the three states, and dividing each joint probability by it gives the last column of the table: given an alarm, the unit is degraded with probability $0.0840/0.1562 = 0.538$, healthy with 0.272 and failed with 0.190.

3.4.2 Continuous marginalization

When the unknown is continuous, the sum becomes an integral. Let the failure rate λ of a component be uncertain and let $\mathbf{y}$ be the observed evidence; their joint distribution is

$p(\lambda, \mathbf{y}) = p(\mathbf{y} \mid \lambda)\, p(\lambda)$. The probability density of the observations is obtained by integrating over every possible value of the rate,

$$p(\mathbf{y}) = \int p(\mathbf{y} \mid \lambda)\, p(\lambda)\, \mathrm{d}\lambda, \tag{3.13}$$

the *marginal likelihood* of the observations. Bayes' theorem for a continuous unknown then reads

$$p(\lambda \mid \mathbf{y}) = \frac{p(\mathbf{y} \mid \lambda)\, p(\lambda)}{\int p(\mathbf{y} \mid \lambda')\, p(\lambda')\, \mathrm{d}\lambda'}, \tag{3.14}$$

where λ' is a dummy variable of integration. The numerator evaluates the relative plausibility of one particular rate; the denominator ensures that the posterior integrates to one.

For the pump record of Chapter 2, $k = 2$ failures in $T_{\mathrm{tot}} = 1000\,\mathrm{h}$ under the Poisson model $\Pr(k \mid \lambda) = (\lambda T_{\mathrm{tot}})^k \mathrm{e}^{-\lambda T_{\mathrm{tot}}}/k!$ and the prior $\mathrm{Gamma}(a, b) = \mathrm{Gamma}(2, 500\,\mathrm{h})$, the gamma integral of Section 3.1.3 evaluates equation (3.13) in closed form:

$$\Pr(k \mid T_{\mathrm{tot}}) = \frac{T_{\mathrm{tot}}^k}{k!}\, \frac{b^a}{\Gamma(a)}\, \frac{\Gamma(a+k)}{(b + T_{\mathrm{tot}})^{a+k}} = \frac{1000^2}{2} \cdot \frac{500^2}{1} \cdot \frac{3!}{1500^4} = 0.1481.$$

Before the record, the model gave a probability of 0.148 to exactly two failures in 1000 h; dividing the numerator of equation (3.14) by it gives the posterior $\mathrm{Gamma}(4, 1500\,\mathrm{h})$. The predictive reliability of equation (3.2) is the same operation applied to the event $T_f > t$: the conditional probability averaged over the posterior of the rate.

3.4.3 Nuisance quantities

Suppose the failure rate λ of a motor is being estimated, but the model also contains an uncertain environmental sensitivity ψ, the change of the motor's rate per unit of an environmental stress. The full posterior is $p(\lambda, \psi \mid \mathbf{y})$, and the engineer wants the distribution of λ alone. It is obtained by integrating ψ out,

$$p(\lambda \mid \mathbf{y}) = \int p(\lambda, \psi \mid \mathbf{y})\, \mathrm{d}\psi. \tag{3.15}$$

The sensitivity does not have to be treated as known; its uncertainty enters the answer through the integration. This matters most in physics-constrained inference, where many unknown state entries influence the quantity studied: the marginal of one entry z_i is the posterior of the whole state $\mathbf{z}$ integrated over all its other entries.

Marginalization against optimization. These two operations are different and must not be confused. Marginalization, equation (3.15), accumulates probability across all values of ψ. Optimization over the nuisance quantity, $\max_\psi p(\lambda, \psi \mid \mathbf{y})$, selects the best-fitting value of ψ for each λ. They generally give different results. A discrete illustration makes the difference plain. Let λ be either low or high and ψ take one of three values, with joint posterior probabilities 0.30, 0.05, 0.05 for a low rate and 0.20, 0.20, 0.20 for a high rate (illustrative). Marginalization gives $\Pr(\text{low}) = 0.40$ and $\Pr(\text{high}) = 0.60$; maximization over ψ gives 0.30 against 0.20 and prefers the low rate. The low rate owes its single peak to one value of ψ; the high rate is supported by every value. The distinction returns when later chapters compare a maximum a posteriori estimate with the marginals of the posterior.

3.4.4 A common cause

Many classical reliability formulas assume that component failures are independent. If two redundant components each fail with probability 0.05, independence gives $\Pr(F_A \cap F_B) = 0.05^2 = 0.0025$. The conclusion is invalid if the components share an important failure mechanism.

Consider two pumps connected to a common electrical supply, and let C be the event of a common electrical disturbance, with $\Pr(C) = 0.05$. Figure 3.4 draws the dependency: the box at the top is the shared condition, the two boxes below are the failure events of the pumps, and each arrow says that the probability of the event at its head depends on the condition at its tail. There is no arrow between the pumps.

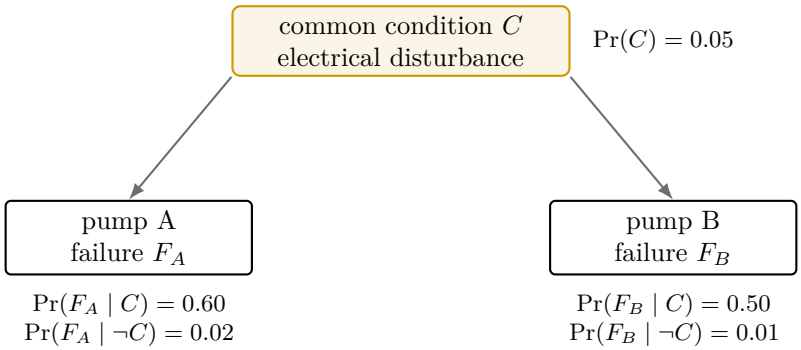


Figure 3.4: A probabilistic dependency graph for two pumps on a common electrical supply. The shared condition C influences both failure events; the absence of an arrow between the pumps states that the failures are conditionally independent given C. The probabilities are those of Table 3.5.

The failure probabilities of the pumps depend on whether the disturbance occurs. Table 3.5 lists them: the rows are the two conditions and the columns the failure probability of each pump in that condition.

Table 3.5: Failure probabilities of two pumps with and without a common electrical disturbance (illustrative); $\Pr(C) = 0.05$.

Condition	$\Pr(F_A \mid \cdot)$	$\Pr(F_B \mid \cdot)$
common disturbance occurs, C	0.60	0.50
no common disturbance, $\neg C$	0.02	0.01

Assume that the two failures are *conditionally independent given* C:

$$\Pr(F_A \cap F_B \mid C) = \Pr(F_A \mid C)\,\Pr(F_B \mid C), \qquad \Pr(F_A \cap F_B \mid \neg C) = \Pr(F_A \mid \neg C)\,\Pr(F_B \mid \neg C).$$

Once the condition of the supply is known, the failure of one pump says nothing more about the other. This does not mean that the failures are unconditionally independent, as the next subsection shows.

3.4.5 The joint failure probability and the independence error

The law of total probability, equation (3.12), applied to the event $F_A \cap F_B$ with the two conditions C and $\neg C$ as the states, gives

$$\Pr(F_A \cap F_B) = \Pr(F_A \cap F_B \mid C)\,\Pr(C) + \Pr(F_A \cap F_B \mid \neg C)\,\Pr(\neg C).$$

Under conditional independence each conditional joint probability is a product,

$$\Pr(F_A \cap F_B) = 0.60 \times 0.50 \times 0.05 + 0.02 \times 0.01 \times 0.95 = 0.015 + 0.00019,$$

so that

$$\Pr(F_A \cap F_B) = 0.01519. \tag{3.16}$$

The joint failure probability is about 1.52 %, and almost all of it, 0.015 of 0.01519, comes from the disturbance.

The marginal failure probabilities follow from the same law,

$$\Pr(F_A) = 0.60 \times 0.05 + 0.02 \times 0.95 = 0.049, \qquad \Pr(F_B) = 0.50 \times 0.05 + 0.01 \times 0.95 = 0.0345.$$

If the failures are incorrectly treated as independent, the joint probability is the product of the marginals,

$$\Pr(F_A \cap F_B) \overset{\text{incorrect}}{=} 0.049 \times 0.0345 = 0.00169,$$

or 0.169 %. The correct value, 1.519 %, is 8.99 times larger: ignoring the common cause underestimates the probability that both pumps fail by a factor of about nine in this illustrative system, and the independent product is only 11 % of the correct value.

Common mistake

Multiplying marginal failure probabilities to obtain the probability that redundant components fail together. The product is the joint probability only when the failures are independent; a shared supply, a shared environment or a shared maintenance crew makes them dependent, and the redundant pair is then much less reliable than the product suggests.

Principle 3.2. *Understanding the dependency structure of a system can matter more than estimating the failure probabilities of its individual components with greater precision.*

3.4.6 Conditional independence

Two random variables X_1 and X_2 are conditionally independent given C when, for every value c of C,

$$p(x_1, x_2 \mid c) = p(x_1 \mid c)\, p(x_2 \mid c), \tag{3.17}$$

written $X_1 \perp X_2 \mid C$. Dividing both sides by $p(x_1 \mid c)$ and using the definition of conditional probability on the left gives the equivalent statement $p(x_2 \mid x_1, c) = p(x_2 \mid c)$: after conditioning on C, learning X_1 gives no additional information about X_2 under the assumed model.

When C is unknown, observing X_1 informs C, and that changes what is believed about X_2. If pump A fails, the shared supply becomes suspect: by Bayes' theorem,

$$\Pr(C \mid F_A) = \frac{\Pr(F_A \mid C)\ \Pr(C)}{\Pr(F_A)} = \frac{0.60 \times 0.05}{0.049} = 0.612,$$

up from 0.05. The probability that pump B also fails is then the law of total probability conditioned on F_A, with conditional independence used to drop F_A from $\Pr(F_B \mid C, F_A)$:

$$\begin{aligned}\Pr(F_B \mid F_A) &= \Pr(F_B \mid C)\ \Pr(C \mid F_A) + \Pr(F_B \mid \neg C)\ \Pr(\neg C \mid F_A) \\ &= 0.50 \times 0.612 + 0.01 \times 0.388 = 0.310,\end{aligned}$$

which equals $\Pr(F_A \cap F_B)/\Pr(F_A) = 0.01519/0.049$, against the unconditional $\Pr(F_B) = 0.0345$. The failure of one pump raises the failure probability of the other ninefold, through the cause they share. Conditional independence is therefore not unconditional independence: the variables are independent only once the common cause is known.

3.5 Conditional independence and a continuous update in a physical model

The same rules apply when the unknowns are physical quantities tied by a conservation relation. This section writes the equipment hall as a factorized probability model, derives the posterior of a state vector under Gaussian rows, works the numbers for the hall temperature after one reading, answers the threshold question with the posterior, and then carries the same derivation to the two-entry hall of Chapter 1, in which the reading reaches the conductance through the physics row.

3.5.1 The hall as a probability model

The hall of the earlier chapters satisfies the steady heat balance

$$T = T_a + \frac{Q}{G}, \tag{3.18}$$

with the hall temperature T and the ambient temperature T_a in °C, the heat load Q in kW and the conductance G of the heat-removal path in kW/K. A sensor reads the hall temperature with an additive error,

$$y_T = T + \epsilon, \qquad \epsilon \sim \mathcal{N}(0, \sigma_T^2),$$

where ϵ is the measurement error and σ_T its standard deviation in K. Figure 3.5 draws the system twice. Panel (a) is the plant graph: the source supplies the load Q to the hall node, which carries the free entry T; the green path to the ambient node carries the conductance, held in the state as the log-conductance $g = \ln(G/G_{\text{ref}})$; the ambient node is pinned at T_a; and the dashed arrow is the reading y_T. Panel (b) is the dependency graph of the probability model of this section, in which Q, G and T_a are all treated as uncertain: the load, the conductance and the ambient temperature determine the modeled hall temperature, and the reading depends on that temperature alone.

If the measurement error is independent of the physical inputs, the reading is conditionally independent of those inputs once the true temperature is known:

$$p(y_T \mid T, G, Q, T_a) = p(y_T \mid T). \tag{3.19}$$

This does not make the reading uninformative about the conductance. The reading informs T, and the physical relation (3.18) connects T to G: evidence propagates from the sensor reading to the temperature and from the temperature to the conductance. An observation of one quantity updates the beliefs about every quantity connected to it through shared physical constraints, and this is one of the main reasons why inference over a graph is useful.

The factorized probability model. Suppose G, Q and T_a are uncertain inputs with a joint prior $p(G, Q, T_a)$, which relaxes the pins of the plant graph and carries Q and T_a as free entries. The chain rule, the multiplication rule applied twice, and equation (3.19) give the joint model

$$p(G, Q, T_a, T, y_T) = p(G, Q, T_a)\, p(T \mid G, Q, T_a)\, p(y_T \mid T). \tag{3.20}$$

The physics factor is written as a Gaussian around the heat balance, $p(T \mid G, Q, T_a) = \mathcal{N}(T;\ T_a + Q/G,\ \sigma_p^2)$, with a model tolerance σ_p in K: a physical discrepancy is a finite σ_p, and the exact deterministic balance of equation (3.18) is the limit $\sigma_p \to 0$, in which the factor becomes the Dirac delta $\delta(T - T_a - Q/G)$. The joint distribution is the complete probability model, and the posterior wanted is $p(G, Q, T_a, T \mid y_T)$, over the whole state with T kept in it, which answers questions about the conductance, the temperature, the operating margins and every other quantity of the model. Each of its three factors is a row over that

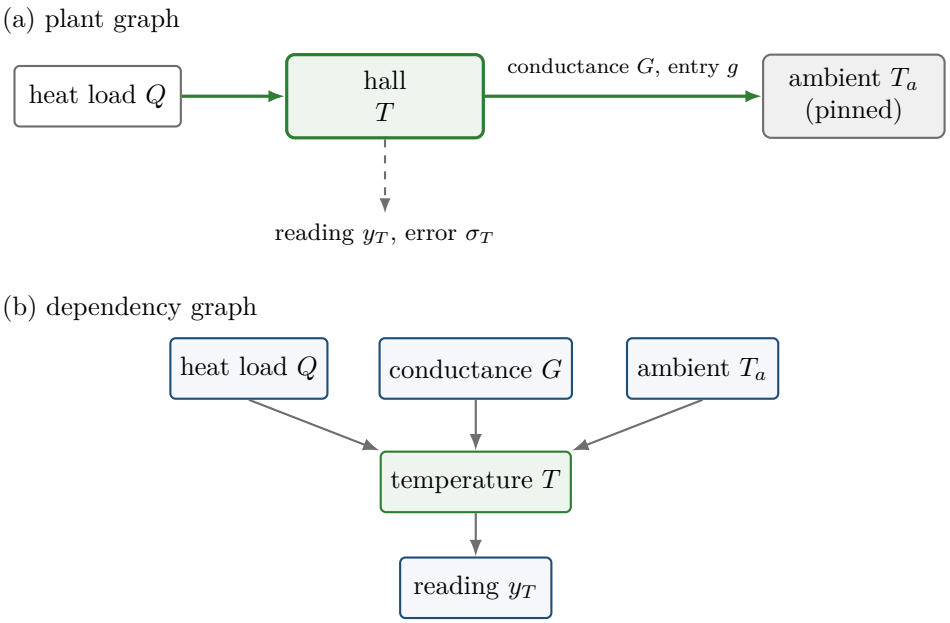


Figure 3.5: The hall in two drawings. (a) The plant graph: free entry T at the hall node and $g = \ln(G/G_{\text{ref}})$ on the heat-removal path; ambient temperature T_a pinned; source Q; reading y_T observed (dashed). (b) The dependency graph of the factorized model of equation (3.20), in which Q, G and T_a carry priors: arrows point from a quantity to what depends on it, and the reading depends on the inputs only through T.

state: the prior, the physics row with tolerance σ_p and the reading row; the derivations of this section keep all of them as rows and never remove an entry by substitution. In the limit $\sigma_p \to 0$, integrating T out against the delta gives the same posterior of the inputs as $p(G, Q, T_a \mid y_T) \propto p(G, Q, T_a)\, p(y_T \mid T = T_a + Q/G)$: every combination of inputs is weighted by how well the temperature it implies explains the reading.

3.5.2 Prior and reading for the hall temperature

The simplest continuous case has one unknown, the true hall temperature T, so that the state is this one number,

$$T \in \mathbb{R}.$$

The general forms of the book write the state as a vector $\mathbf{z}$; here it has the single entry T, so every quantity of the fit below is a number: each row's residual and its slope in T, the objective's slope and curvature, the information and the posterior variance. Previous thermal modeling gives an uncertain prior estimate,

$$T \sim \mathcal{N}(\mu_0, \sigma_0^2) = \mathcal{N}(26, 2^2), \qquad \mu_0 = 26\,^{\circ}\text{C}, \quad \sigma_0 = 2\,\text{K}.$$

A sensor reports $y_T = 27\,^{\circ}\text{C}$, and its error is Gaussian with $\sigma_T = 0.5\,\text{K}$, so that $y_T \mid T \sim \mathcal{N}(T, 0.5^2)$. The question is the posterior distribution of the true temperature after the reading. Bayes' theorem for a continuous unknown, with the evidence dropped because it does not depend on the state, gives

$$p(T \mid y_T) \propto p(y_T \mid T)\, p(T),$$

with the Gaussian prior and the Gaussian likelihood

$$p(T) \propto \exp\left[-\frac{(T-\mu_0)^2}{2\sigma_0^2}\right], \qquad p(y_T \mid T) \propto \exp\left[-\frac{(y_T - T)^2}{2\sigma_T^2}\right].$$

Multiplying,

$$p(T \mid y_T) \propto \exp\left[-\frac{1}{2}\left(\frac{(T-\mu_0)^2}{\sigma_0^2}+\frac{(y_T-T)^2}{\sigma_T^2}\right)\right]. \tag{3.21}$$

Each term of the exponent is the square of a difference divided by its standard deviation. The two differences are the residuals of the problem, one prior row and one reading row, each already divided by its standard deviation,

$$r_1(T)=\frac{T-\mu_0}{\sigma_0}, \qquad r_2(T)=\frac{T-y_T}{\sigma_T}, \tag{3.22}$$

and equation (3.21) reads $p(T \mid y_T) \propto \exp[-\frac{1}{2}(r_1^2+r_2^2)] = \mathrm{e}^{-\mathcal{L}(T)}$, with the objective $\mathcal{L}(T) = \frac{1}{2}(r_1^2+r_2^2)$ of the notation chapter. Both rows are linear in T, each with a constant slope $j_k = \mathrm{d}r_k/\mathrm{d}T$: $j_1 = 1/\sigma_0$ and $j_2 = 1/\sigma_T$.

Gaussian rows over any state. The same step holds for a state of any size and any number of rows. The joint model factorizes by the chain rule, as equation (3.20) does for the hall: a prior factor over the inputs, a physics factor for each physical relation, and a reading factor for each reading, which by conditional independence, equation (3.19), depends on the state only through the quantity read. Row k is one such factor written as a residual: the difference between the quantity computed from the state and what was observed, or what the prior or the physics expects, divided by its standard deviation σ_k. A Gaussian factor is $\exp[-\frac{1}{2}r_k(\mathbf{z})^2]$ times a normalizer, $1/(\sqrt{2\pi}\,\sigma_k)$, that is free of the state. Multiplying the factors and dropping the normalizers and the evidence, Bayes' theorem gives

$$p(\mathbf{z} \mid \mathbf{y}) \propto \prod_k \exp\left[-\tfrac{1}{2}r_k(\mathbf{z})^2\right] = \exp\left[-\tfrac{1}{2}r(\mathbf{z})^{\mathsf{T}}r(\mathbf{z})\right] = \mathrm{e}^{-\mathcal{L}(\mathbf{z})}, \tag{3.23}$$

where $r(\mathbf{z})$ stacks the reading rows, the prior rows and the physics rows with their tolerances σ_p. For the one-entry hall the product has the two factors of equation (3.21) and the sum of squares is the number $r_1^2+r_2^2$; the vector is needed when the state has several entries, as in Section 3.5.6 and in the hall of Chapter 4. A Gaussian prior over several entries, $\mathcal{N}(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$ with correlated entries, enters as the rows $\mathbf{L}_0(\mathbf{z}-\boldsymbol{\mu}_0)$ with $\mathbf{L}_0^{\mathsf{T}}\mathbf{L}_0 = \boldsymbol{\Sigma}_0^{-1}$, since then $(\mathbf{z}-\boldsymbol{\mu}_0)^{\mathsf{T}}\boldsymbol{\Sigma}_0^{-1}(\mathbf{z}-\boldsymbol{\mu}_0)$ is the sum of their squares.

3.5.3 Completion of the square in the state

Suppose the rows are linear in the state, as both rows of equation (3.22) are. Then, for any point $\mathbf{z}_0$,

$$r(\mathbf{z}) = r(\mathbf{z}_0) + \mathbf{J}\,(\mathbf{z}-\mathbf{z}_0)$$

holds exactly, with the Jacobian $\mathbf{J} = \partial r/\partial \mathbf{z}$, a matrix of m rows and n columns for m rows and n entries. Write $\boldsymbol{\Delta} = \mathbf{z}-\mathbf{z}_0$ and $r_0 = r(\mathbf{z}_0)$. The objective expands as

$$\mathcal{L}(\mathbf{z}) = \tfrac{1}{2}(r_0+\mathbf{J}\boldsymbol{\Delta})^{\mathsf{T}}(r_0+\mathbf{J}\boldsymbol{\Delta}) = \tfrac{1}{2}r_0^{\mathsf{T}}r_0 + \boldsymbol{\Delta}^{\mathsf{T}}\mathbf{J}^{\mathsf{T}}r_0 + \tfrac{1}{2}\boldsymbol{\Delta}^{\mathsf{T}}\mathbf{J}^{\mathsf{T}}\mathbf{J}\,\boldsymbol{\Delta}.$$

Its gradient, $\mathbf{J}^{\mathsf{T}}r_0 + \mathbf{J}^{\mathsf{T}}\mathbf{J}\,\boldsymbol{\Delta}$, vanishes at the step $\hat{\boldsymbol{\Delta}}$ that solves the *normal equations*,

$$(\mathbf{J}^{\mathsf{T}}\mathbf{J})\,\hat{\boldsymbol{\Delta}} = -\mathbf{J}^{\mathsf{T}}r(\mathbf{z}_0), \qquad \hat{\mathbf{z}} = \mathbf{z}_0 + \hat{\boldsymbol{\Delta}}. \tag{3.24}$$

Writing $\mathbf{J}^{\mathsf{T}}r_0 = -\mathbf{J}^{\mathsf{T}}\mathbf{J}\,\hat{\boldsymbol{\Delta}}$ in the expansion completes the square,

$$\mathcal{L}(\mathbf{z}) = \tfrac{1}{2}(\boldsymbol{\Delta}-\hat{\boldsymbol{\Delta}})^{\mathsf{T}}\mathbf{J}^{\mathsf{T}}\mathbf{J}\,(\boldsymbol{\Delta}-\hat{\boldsymbol{\Delta}}) + \text{constant} = \tfrac{1}{2}(\mathbf{z}-\hat{\mathbf{z}})^{\mathsf{T}}\mathbf{J}^{\mathsf{T}}\mathbf{J}\,(\mathbf{z}-\hat{\mathbf{z}}) + \text{constant},$$

where the constant collects the terms free of $\mathbf{z}$. By equation (3.23) the posterior is the exponential of minus this quadratic form, the density of a multivariate Gaussian with mean $\hat{\mathbf{z}}$ and covariance

$$\boldsymbol{\Sigma} = (\mathbf{J}^\mathsf{T}\mathbf{J})^{-1}. \tag{3.25}$$

The matrix $\mathbf{J}^\mathsf{T}\mathbf{J}$ is the information matrix $\boldsymbol{\Lambda}$ of the normalized rows. When the rows are written unnormalized, each in its own unit with the weight $w_k = 1/\sigma_k^2$ and the weights collected in $\mathbf{W} = \operatorname{diag}(w_k)$, the same matrix is $\boldsymbol{\Lambda} = \mathbf{J}^\mathsf{T}\mathbf{W}\mathbf{J}$, with $\mathbf{J}$ there the Jacobian of the unnormalized rows. For rows linear in the state the normal equations reach $\hat{\mathbf{z}}$ in one step from any $\mathbf{z}_0$, and equations (3.24) and (3.25) are the exact posterior mean and covariance. For nonlinear rows the same step is the Gauss–Newton step, repeated from the new point until it vanishes, and the result is a local Gaussian approximation, the Laplace approximation, whose derivation and curvature term Chapter 4 gives, with the three-entry hall as its example; Section 3.5.6 reads the two-entry hall of Chapter 1 in the terms of this section.

Precisions add. Split the rows into prior rows $\mathbf{L}_0(\mathbf{z} - \boldsymbol{\mu}_0)$, whose Jacobian $\mathbf{L}_0$ satisfies $\mathbf{L}_0^\mathsf{T}\mathbf{L}_0 = \boldsymbol{\Sigma}_0^{-1}$, and reading rows $\mathbf{R}^{-1/2}(\mathbf{A}\mathbf{z} - \mathbf{y})$, where the matrix $\mathbf{A}$ maps the state to the quantities read and the diagonal matrix $\mathbf{R}$ holds the variances of the readings. The Jacobian of the reading rows is $\mathbf{J}_y = \mathbf{R}^{-1/2}\mathbf{A}$; stacking the two blocks of rows stacks their Jacobians, and $\mathbf{J}^\mathsf{T}\mathbf{J}$ is the sum of the two blocks' products,

$$\boldsymbol{\Sigma}^{-1} = \boldsymbol{\Sigma}_0^{-1} + \mathbf{J}_y^\mathsf{T}\mathbf{J}_y = \boldsymbol{\Sigma}_0^{-1} + \mathbf{A}^\mathsf{T}\mathbf{R}^{-1}\mathbf{A}. \tag{3.26}$$

Starting the normal equations at $\mathbf{z}_0 = \boldsymbol{\mu}_0$, where the prior rows vanish, gives $\mathbf{J}^\mathsf{T} r(\boldsymbol{\mu}_0) = \mathbf{A}^\mathsf{T}\mathbf{R}^{-1}(\mathbf{A}\boldsymbol{\mu}_0 - \mathbf{y})$, and equation (3.24) gives

$$\hat{\mathbf{z}} = \boldsymbol{\mu}_0 + \boldsymbol{\Sigma}\,\mathbf{A}^\mathsf{T}\mathbf{R}^{-1}(\mathbf{y} - \mathbf{A}\boldsymbol{\mu}_0). \tag{3.27}$$

The posterior information is the prior information plus the information of the readings, and the posterior mean is moved from the prior mean toward the readings in proportion to their information.

The conditional of a joint Gaussian. The same posterior follows from the joint distribution of the state and the readings. Before the readings, $\mathbf{z} \sim \mathcal{N}(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$ and $\mathbf{y} = \mathbf{A}\mathbf{z} + \boldsymbol{\epsilon}$ with $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{R})$ independent of the state, so the stacked vector of $\mathbf{z}$ and $\mathbf{y}$ is Gaussian with means $\boldsymbol{\mu}_0$ and $\mathbf{A}\boldsymbol{\mu}_0$ and covariance blocks

$$\boldsymbol{\Sigma}_{zz} = \boldsymbol{\Sigma}_0, \qquad \boldsymbol{\Sigma}_{zy} = \boldsymbol{\Sigma}_0\mathbf{A}^\mathsf{T} = \boldsymbol{\Sigma}_{yz}^\mathsf{T}, \qquad \boldsymbol{\Sigma}_{yy} = \mathbf{A}\boldsymbol{\Sigma}_0\mathbf{A}^\mathsf{T} + \mathbf{R}.$$

Let $\mathbf{P}$ be the inverse of the joint covariance and $\mathbf{S} = \boldsymbol{\Sigma}_{zz} - \boldsymbol{\Sigma}_{zy}\boldsymbol{\Sigma}_{yy}^{-1}\boldsymbol{\Sigma}_{yz}$ the *Schur complement* of $\boldsymbol{\Sigma}_{yy}$. The block-inversion formula, checked by multiplying back, gives $\mathbf{P}_{zz} = \mathbf{S}^{-1}$ and $\mathbf{P}_{zy} = -\mathbf{S}^{-1}\boldsymbol{\Sigma}_{zy}\boldsymbol{\Sigma}_{yy}^{-1}$. With the readings fixed, the exponent of the joint density, times -1, is a quadratic in the state,

$$\tfrac{1}{2}(\mathbf{z}-\boldsymbol{\mu}_0)^\mathsf{T}\mathbf{P}_{zz}(\mathbf{z}-\boldsymbol{\mu}_0) + (\mathbf{z}-\boldsymbol{\mu}_0)^\mathsf{T}\mathbf{P}_{zy}(\mathbf{y}-\mathbf{A}\boldsymbol{\mu}_0) + \text{constant} = \tfrac{1}{2}(\mathbf{z}-\mathbf{m})^\mathsf{T}\mathbf{S}^{-1}(\mathbf{z}-\mathbf{m}) + \text{constant},$$

with $\mathbf{m} = \boldsymbol{\mu}_0 + \boldsymbol{\Sigma}_{zy}\boldsymbol{\Sigma}_{yy}^{-1}(\mathbf{y} - \mathbf{A}\boldsymbol{\mu}_0)$. The conditional density is the joint density divided by $p(\mathbf{y})$, which is free of the state, so

$$\mathbf{z} \mid \mathbf{y} \sim \mathcal{N}(\boldsymbol{\mu}_0 + \boldsymbol{\Upsilon}(\mathbf{y} - \mathbf{A}\boldsymbol{\mu}_0),\ \boldsymbol{\Sigma}_0 - \boldsymbol{\Upsilon}\mathbf{A}\boldsymbol{\Sigma}_0), \qquad \boldsymbol{\Upsilon} = \boldsymbol{\Sigma}_0\mathbf{A}^\mathsf{T}(\mathbf{A}\boldsymbol{\Sigma}_0\mathbf{A}^\mathsf{T} + \mathbf{R})^{-1}, \tag{3.28}$$

since $\boldsymbol{\Sigma}_{zy}\boldsymbol{\Sigma}_{yy}^{-1}\boldsymbol{\Sigma}_{yz} = \boldsymbol{\Upsilon}\mathbf{A}\boldsymbol{\Sigma}_0$. This is the correction form: the prior mean is moved toward the readings by the gain $\boldsymbol{\Upsilon}$ times the innovation $\mathbf{y} - \mathbf{A}\boldsymbol{\mu}_0$, and the covariance is the Schur complement. The Kalman literature writes the gain $\mathbf{K}$; this book keeps $\mathbf{K}$ for the hat matrix and κ for an assurance level. It is the same posterior as equations (3.26) and (3.27), by the matrix inversion identity $(\boldsymbol{\Sigma}_0^{-1} + \mathbf{A}^\mathsf{T}\mathbf{R}^{-1}\mathbf{A})^{-1} = \boldsymbol{\Sigma}_0 - \boldsymbol{\Upsilon}\mathbf{A}\boldsymbol{\Sigma}_0$; the information form inverts a matrix of the size of the state, the correction form one of the size of the readings.

Remark 3.1. For one entry and one reading, $\mathbf{A} = 1$, $\mathbf{\Sigma}_0 = \sigma_0^2$ and $\mathbf{R} = \sigma_T^2$, and the three equations reduce to the scalar rules $1/\sigma_{\text{post}}^2 = 1/\sigma_0^2 + 1/\sigma_T^2$, $\mu_{\text{post}} = \sigma_{\text{post}}^2(\mu_0/\sigma_0^2 + y_T/\sigma_T^2)$ and $\mu_{\text{post}} = \mu_0 + \upsilon(y_T - \mu_0)$ with $\upsilon = \sigma_0^2/(\sigma_0^2 + \sigma_T^2)$: precisions add, and the posterior mean is the precision-weighted average of the prior mean and the reading.

3.5.4 The numbers and what changed

For the hall the state is the one number T, and the two rows of equation (3.22) have the slopes

$$j_1 = \frac{1}{\sigma_0} = 0.5\,\mathrm{K}^{-1}, \qquad j_2 = \frac{1}{\sigma_T} = 2\,\mathrm{K}^{-1}.$$

The row terms are $\mathcal{L}_k = \frac{1}{2}r_k^2$, with the slopes $\mathcal{L}_k' = j_k r_k$ and the curvatures $\mathcal{L}_k'' = j_k^2$ in T, the primes denoting $\mathrm{d}/\mathrm{d}T$; slopes and curvatures add over the rows,

$$\mathcal{L}' = j_1 r_1 + j_2 r_2, \qquad \mathcal{L}'' = j_1^2 + j_2^2.$$

Because the rows are linear, the curvature has no further term and is the same at every T; it is the information

$$\Lambda = j_1^2 + j_2^2 = 0.5^2 + 2^2 = 0.25 + 4 = 4.25\,\mathrm{K}^{-2},$$

the prior precision $1/\sigma_0^2$ plus the reading precision $1/\sigma_T^2$, as equation (3.26) states in the scalar form of the remark above, with the prior variance $\sigma_0^2 = 4\,\mathrm{K}^2$ and the reading variance $\sigma_T^2 = 0.25\,\mathrm{K}^2$. It is the one-entry case of $\mathbf{J}^{\mathsf{T}}\mathbf{J}$, and equation (3.25) gives the posterior variance as its reciprocal,

$$\sigma_{\text{post}}^2 = \frac{1}{\Lambda} = \frac{1}{4.25} = 0.2353\,\mathrm{K}^2, \qquad \sigma_{\text{post}} = 0.4851\,\mathrm{K}.$$

With one entry the normal equations (3.24) are the one equation $\Lambda\,\Delta T = -\mathcal{L}'$, the Newton step $\Delta T = -\mathcal{L}'/\mathcal{L}''$. They start at the prior mean, $T_0 = 26°\mathrm{C}$, where the prior row is 0 and the reading row $(26 - 27)/0.5 = -2$, so that $\mathcal{L}' = 0.5 \times 0 + 2 \times (-2) = -4\,\mathrm{K}^{-1}$, and the step is $\Delta T = 4/4.25 = 0.9412\,\mathrm{K}$. The iterations, with the row slopes $\mathcal{L}_1'$ and $\mathcal{L}_2'$ in K^{-1}, are

iteration	T (°C)	$\mathcal{L}(T)$	$\mathcal{L}_1'$	$\mathcal{L}_2'$	$\mathcal{L}'$	$\mathcal{L}''$ (K^{-2})	ΔT (K)
0	26	2.0000	0	−4	−4	4.25	0.9412
1	26.9412	0.1176	0.2353	−0.2353	0	4.25	0

At the second point the residuals are $r_1 = 0.4706$ and $r_2 = -0.1176$, the slope is $\mathcal{L}' = 0.5 \times 0.4706 + 2 \times (-0.1176) = 0$, and the step vanishes: the rows are linear, and one step reaches the minimum. The correction form (3.28) confirms both moments: the variance of the reading before it arrives is $\sigma_0^2 + \sigma_T^2 = 4 + 0.25 = 4.25\,\mathrm{K}^2$, $\upsilon = 4/4.25 = 0.9412$, the mean $26 + 0.9412 \times (27 - 26) = 26.9412°\mathrm{C}$ and the Schur complement $4 - 0.9412 \times 4 = 0.2353\,\mathrm{K}^2$. The answer is the minimum $\hat{T} = 26.9412°\mathrm{C}$, written μ_{post} below, with the standard deviation σ_{post}:

$$T \mid y_T \sim \mathcal{N}(26.9412, 0.4851^2). \tag{3.29}$$

Figure 3.6(a) draws the update. The horizontal axis is the temperature and the vertical axis the probability density. The gray curve is the prior, the dashed blue curve the likelihood of the reading as a function of T (normalized to integrate to one), and the green curve the posterior. The posterior mean moved from 26 to 26.94°C, and the standard deviation fell from 2 to 0.49 K. The reading was more precise than the prior, so it received the greater weight, $\upsilon = 0.94$; but the posterior mean is not exactly 27°C, because the prior also carries information.

Table 3.6 shows the trend for sensors of four precisions, all reading 27°C against the same prior. The first column is the sensor's standard deviation, the second the weight υ of the

reading, and the remaining columns the posterior mean, the posterior standard deviation and the exceedance probability derived in Section 3.5.5. Each row is the same computation with the slopes $j_1 = 1/\sigma_0$ and $j_2 = 1/\sigma_T$, so that $\sigma^2_{\text{post}} = 1/\Lambda = 1/(0.25 + 1/\sigma_T^2)$. The second row is the case worked above. A sensor as uncertain as the prior, $\sigma_T = 2\,\text{K}$, receives half the weight and leaves the posterior standard deviation at $2/\sqrt{2} = 1.41\,\text{K}$.

Table 3.6: The hall-temperature update for four sensor precisions: prior $\mathcal{N}(26, 2^2)$, reading $y_T = 27°\text{C}$, equations (3.26) and (3.28), and the exceedance probability of $T_{\max} = 28°\text{C}$ derived in Section 3.5.5.

σ_T (K)	weight υ	μ_{post} (°C)	σ_{post} (K)	$\Pr(T > T_{\max} \mid y_T)$
0.25	0.9846	26.9846	0.2481	2×10^{-5}
0.5	0.9412	26.9412	0.4851	0.0145
1.0	0.8000	26.8000	0.8944	0.0899
2.0	0.5000	26.5000	1.4142	0.1444

3.5.5 The reliability question

Suppose the equipment temperature limit is $T_{\max} = 28°\text{C}$, and the question is the probability that the hall exceeds it, $\Pr(T > T_{\max} \mid y_T)$. Since the posterior is Gaussian, standardizing with the mean $\mu_{\text{post}} = \hat{T}$ and the standard deviation $\sigma_{\text{post}} = 1/\sqrt{\Lambda}$ and using the standard normal cumulative distribution function Φ,

$$\begin{aligned} \Pr(T > T_{\max} \mid y_T) &= 1 - \Phi\left(\frac{T_{\max} - \mu_{\text{post}}}{\sigma_{\text{post}}}\right) \\ &= 1 - \Phi\left(\frac{28 - 26.9412}{0.4851}\right) = 1 - \Phi(2.183) = 0.0145, \end{aligned} \tag{3.30}$$

or about 1.45 %. Before the reading the same calculation with the prior gives $\Pr(T > 28°\text{C}) = 1 - \Phi((28 - 26)/2) = 1 - \Phi(1) = 0.1587$. The estimated exceedance probability fell from 15.9 % to 1.45 %, for the same fixed operating condition and the assumed Gaussian measurement model. Figure 3.6(b) shades the two tails: the axes are those of panel (a), the dashed line is the limit, the gray area to its right under the prior is 0.159, and the green area under the posterior is 0.0145.

The physical system did not necessarily change when the reading arrived; the information about it changed, and that is the mechanism of Bayesian updating. The calculation also walks the method of the preface on its smallest example. The plant graph of Figure 3.5(a) is step 0. The question $\Pr(T > T_{\max} \mid y_T)$, unknown on the left and reading on the right, is step 1. Attaching the prior and the reading as the two rows of equation (3.22) is step 2. The posterior, derived through the normal equations and the variance as the reciprocal of the information, equations (3.24) and (3.25) with one entry, and computed by the one Newton step $-\mathcal{L}'/\mathcal{L}''$ with $\Lambda = 4.25\,\text{K}^{-2}$, is step 3. Reading the normal tail from it is step 4, and checking that tail by the correction form and its Schur complement is step 5, and the sentence "the hall exceeds 28°C with probability 1.45 %, posterior $26.94 \pm 0.49\,\text{K}$" is step 6. Step 7 reads Table 3.6: if 1.45 % were too high a residual doubt, a sensor with $\sigma_T = 0.25\,\text{K}$ would, for a reading near the same value, reduce it to the order of 10^{-5}. The posterior is inferred first, and the engineering question is then answered from it:

$$\underbrace{p(T \mid y_T)}_{\text{inference}} \quad \longrightarrow \quad \underbrace{\Pr(T > T_{\max} \mid y_T)}_{\text{reliability question}}.$$

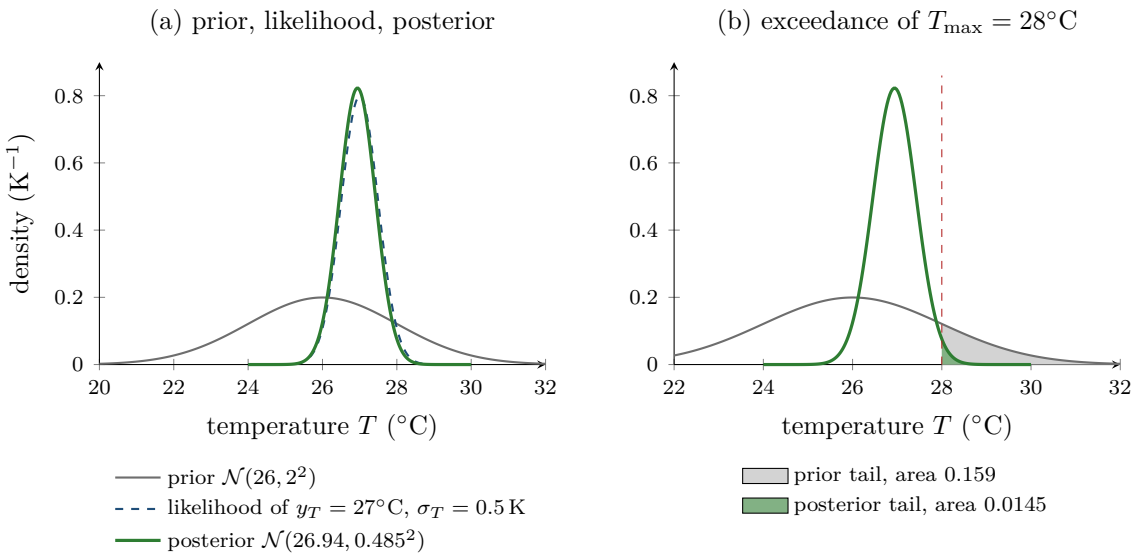


Figure 3.6: The hall-temperature update of equations (3.26)–(3.29). (a) Prior $\mathcal{N}(26, 2^2)$, likelihood of the reading $y_T = 27°\mathrm{C}$ with $\sigma_T = 0.5\,\mathrm{K}$ drawn as the density $\mathcal{N}(27, 0.5^2)$ in T, and posterior $\mathcal{N}(26.9412, 0.4851^2)$. (b) The probability that the hall exceeds $T_{\max} = 28°\mathrm{C}$ (dashed): $1 - \Phi(1) = 0.159$ under the prior and $1 - \Phi(2.183) = 0.0145$ under the posterior (equation (3.30)).

3.5.6 The two-entry hall and its physics row

The reading of the one-entry hall informed only the entry it measured; Section 3.5.1 claimed that a reading also reaches the conductance through the physics row. The hall of Chapter 1 shows it. Its state is $\mathbf{z} = (T,\ g)^{\mathsf{T}}$, and its three rows, equation (1.10), are the physics row with tolerance $\sigma_p = 0.01\,\mathrm{K}$, the reading row with $\sigma_T = 0.5\,\mathrm{K}$ and a handbook prior row on g with standard deviation 0.3; Chapter 1 solves their normal equations by Gauss–Newton (Table 1.4) and reaches $\hat{\mathbf{z}} = (23.070°\mathrm{C},\ -0.0139)^{\mathsf{T}}$. Equation (3.26) holds row by row: with $\mathbf{J}_k$ the k-th row of $\mathbf{J}$, $\mathbf{J}^{\mathsf{T}}\mathbf{J} = \sum_k \mathbf{J}_k^{\mathsf{T}}\mathbf{J}_k$, each row contributing the outer product of its own sensitivities. At $\hat{\mathbf{z}}$ the rows of $\mathbf{J}$ are $(100,\ 507.0)$, $(2,\ 0)$ and $(0,\ 3.333)$, and

$$\mathbf{J}^{\mathsf{T}}\mathbf{J} = \underbrace{\begin{bmatrix} 10\,000 & 50\,701 \\ 50\,701 & 257\,056 \end{bmatrix}}_{\text{physics}} + \underbrace{\begin{bmatrix} 4 & 0 \\ 0 & 0 \end{bmatrix}}_{\text{reading}} + \underbrace{\begin{bmatrix} 0 & 0 \\ 0 & 11.11 \end{bmatrix}}_{\text{prior}} = \begin{bmatrix} 10\,004 & 50\,701 \\ 50\,701 & 257\,067 \end{bmatrix}.$$

The reading's and the prior's terms are diagonal, because each of those rows carries one entry. The only off-diagonal entry, 50 701, comes from the physics row, the one row with sensitivities in both columns: this is equation (3.19) in vector form, the reading informing T and the physics row carrying that information to g. Without the reading row, the physics and prior terms alone invert to $\sqrt{\Sigma_{22}} = 0.3$, the prior spread; with it, the covariance of Chapter 1, equation (1.21), has $\sqrt{\Sigma_{11}} = 0.481\,\mathrm{K}$, $\sqrt{\Sigma_{22}} = 0.0948$ and correlation -0.9998. The reading has cut the conductance spread from 0.3 to 0.0948 although its own row has a zero in the column of g, and the posterior lies along the physics curve, on the short arc of the state plane of Figure 1.5 inside the reading band. Because the physics row is nonlinear, these spreads include the curvature term that Chapter 4 derives; $(\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1}$ alone gives $0.475\,\mathrm{K}$ and 0.094.

3.5.7 Dependence inside and outside the physics graph

The common cause of Section 3.4.4 cost a factor of nine because it was left out of the product. A plant graph built on conservation relations does not leave out a shared cause that has a

physical path: the cause is an entry of the state, and every row that depends on it carries it. The two-entry hall above shows this. The reading and the conductance share no row, yet the physics row couples them, and the posterior correlation of -0.9998 comes out of $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ without being assumed. Had the pumps' electrical supply been a node of the plant graph, with a row for each pump's failure given its state, the joint failure probability (3.16) would have followed from the rows, and the independent product of Section 3.4.5 could not have been formed. Conditional independence given the state is then the statement of Section 3.4.6 with the common cause known: it holds because the cause is in the state.

What the graph does not supply is the independence of the errors. Equation (3.19) assumes that a reading's error is independent of the inputs and, when there are several readings, of the other readings' errors, so that the posterior (3.23) is a product over rows. No conservation relation states that, and none checks it directly. Suppose the hall of Section 3.5.4 carries two sensors from one calibration batch, each with $\sigma_T = 0.5\,\mathrm{K}$, both reading 27°C, with errors correlated through the offset they share, $\mathrm{corr}(\epsilon_1, \epsilon_2) = 0.5$. Treated as independent rows, the pair adds the precision $2/\sigma_T^2 = 8\,\mathrm{K}^{-2}$ to the prior's 0.25. With the correlation, equation (3.26) holds with $\mathbf{R}$ a full matrix, its off-diagonal entry $0.5\,\sigma_T^2$ and its rows whitened by $\mathbf{R}^{-1/2}$, and the pair adds only

$$\mathbf{A}^{\mathsf{T}}\mathbf{R}^{-1}\mathbf{A} = \frac{2}{\sigma_T^2\,(1 + 0.5)} = 5.33\,\mathrm{K}^{-2}, \qquad \mathbf{A} = \begin{bmatrix} 1 \\ 1 \end{bmatrix},$$

the precision of the pair's average, whose error has the variance $\sigma_T^2\,(1 + 0.5)/2$. Table 3.7 compares the two treatments with the single sensor of Table 3.6. Its columns are the precision the readings add, the posterior mean, the posterior standard deviation and the probability of exceeding $T_{\max} = 28$°C.

Table 3.7: Two hall sensors of $\sigma_T = 0.5\,\mathrm{K}$, both reading 27°C, against the prior $\mathcal{N}(26, 2^2)$ and the limit $T_{\max} = 28$°C: the pair treated as independent, the pair with $\mathrm{corr}(\epsilon_1, \epsilon_2) = 0.5$, and one sensor, which is also the limit of a correlation of one.

Readings	precision added (K^{-2})	mean (°C)	σ_{post} (K)	$\Pr(T > T_{\max} \mid \mathbf{y})$
two, treated as independent	8.00	26.9697	0.3482	0.00154
two, correlation 0.5	5.33	26.9552	0.4232	0.00678
one sensor	4.00	26.9412	0.4851	0.0145

The independent treatment reports an exceedance probability of 0.00154 against the 0.00678 the correlated pair supports, 4.4 times too small, and a posterior standard deviation of 0.348 K against 0.423 K. The error has the direction of Section 3.4.5's: a posterior narrower than the evidence supports, and a limit that looks safer than it is. At a correlation of one the second sensor adds nothing, and the posterior is that of one sensor.

The shared offset is a common cause, and the remedy of Section 3.4.6 applies: put it in the state. A bias entry $b \sim \mathcal{N}(0,\ 0.5\,\sigma_T^2)$ carried by both reading rows, with independent remainders of variance $0.5\,\sigma_T^2$, makes the two rows conditionally independent given (T, b), and the state (T, b) gives the posterior of T of the correlated row of Table 3.7 to every printed digit. The posterior of b, however, has the standard deviation 0.348 K against its prior's 0.354 K, with $\mathrm{corr}(T, b \mid \mathbf{y}) = -0.81$. The readings barely inform it: a shared offset moves both readings together, as a change of T does, and only the prior on b bounds it.

The same assumption hides in four other places a physics graph does not reach: a prior shared by several entries, such as two conductances taken from one handbook; a physics tolerance

σ_p treated as independent from row to row where the model's error is systematic; repeated readings averaged as if their errors were independent when they are correlated in time; and a path of the conserved quantity that the graph omits. Each makes the posterior narrower than the evidence supports. What the graph adds is redundancy. When more rows bear on the state than it has entries, a misfit of the rows taken together, a reading far from what the other rows predict, or a comparison of drawings with and without the suspected cause can reveal the dependence: Chapter 15 carries a shared cause among failures as an entry, and Chapters 18, 21, 22 and 24, in Volume II, derive the misfit test, the held-out reading, the bias entry and the comparison of drawings. No test reveals a dependence that moves the readings along a direction the physics also allows, as the shared offset moves them along T.

Principle 3.3. *A physics graph carries dependence through shared physical quantities as entries of the state, so the independent product cannot be formed for them. The independence it still assumes is that of the errors: it must be stated, tested where the rows are redundant, and, where it fails, removed by putting the shared cause in the state.*

3.6 Predictive variance and the covariance of a joint posterior

Two variances complete the chapter. This section derives the law of total variance, which separates the randomness of future failures from the uncertainty of a rate, and then the covariance of a posterior over a joint state, which decides how uncertain a derived quantity is and why the shape of the posterior must be read rather than its marginals.

3.6.1 The law of total variance

Suppose the number of failures of a pump during a future period of length τ is Poisson given the rate, $K_f \mid \lambda \sim \text{Poisson}(\lambda\tau)$, and the rate itself is uncertain, with posterior $p(\lambda \mid \mathbf{y})$. The predictive variance of K_f is wanted. Two facts give it. The law of total expectation, which is the law of total probability applied inside an expectation, states that $\mathbb{E}[K_f \mid \mathbf{y}] = \mathbb{E}\big[\mathbb{E}[K_f \mid \lambda] \mid \mathbf{y}\big]$, the conditional mean averaged over the rate; and the variance of any quantity is the mean of its square minus the square of its mean. Then

$$\text{Var}(K_f \mid \mathbf{y}) = \mathbb{E}\big[\mathbb{E}[K_f^2 \mid \lambda] \mid \mathbf{y}\big] - \big(\mathbb{E}\big[\mathbb{E}[K_f \mid \lambda] \mid \mathbf{y}\big]\big)^2,$$

and writing the conditional second moment as $\mathbb{E}[K_f^2 \mid \lambda] = \text{Var}(K_f \mid \lambda) + \mathbb{E}[K_f \mid \lambda]^2$,

$$\text{Var}(K_f \mid \mathbf{y}) = \mathbb{E}\big[\text{Var}(K_f \mid \lambda) \mid \mathbf{y}\big] + \Big\{\mathbb{E}\big[\mathbb{E}[K_f \mid \lambda]^2 \mid \mathbf{y}\big] - \big(\mathbb{E}\big[\mathbb{E}[K_f \mid \lambda] \mid \mathbf{y}\big]\big)^2\Big\}.$$

The brace is the variance of the conditional mean, which gives the *law of total variance*,

$$\text{Var}(K_f \mid \mathbf{y}) = \mathbb{E}\big[\text{Var}(K_f \mid \lambda) \mid \mathbf{y}\big] + \text{Var}\big(\mathbb{E}[K_f \mid \lambda] \mid \mathbf{y}\big). \tag{3.31}$$

The first term is the expected variability of the failure count when the rate is specified, the aleatory part. The second is the additional predictive variability because the rate is uncertain, the epistemic part. For a Poisson count the mean and the variance are equal, $\mathbb{E}[K_f \mid \lambda] = \text{Var}(K_f \mid \lambda) = \lambda\tau$, so that $\text{Var}(K_f \mid \mathbf{y}) = \mathbb{E}[\lambda\tau \mid \mathbf{y}] + \text{Var}(\lambda\tau \mid \mathbf{y})$, and since τ is fixed,

$$\text{Var}(K_f \mid \mathbf{y}) = \tau\,\mathbb{E}[\lambda \mid \mathbf{y}] + \tau^2\,\text{Var}(\lambda \mid \mathbf{y}). \tag{3.32}$$

The predictive variance has been separated into a process term and a rate term.

3.6.2 Failure counts over a future period

Take the gamma posterior $\text{Gamma}(a, b) = \text{Gamma}(4, 1500\,\text{h})$ of Chapter 2, whose mean and variance are the standard gamma moments,

$$\mathbb{E}[\lambda \mid \mathbf{y}] = \frac{a}{b} = \frac{4}{1500\,\text{h}} = 2.667 \times 10^{-3}\,\text{h}^{-1}, \qquad \text{Var}(\lambda \mid \mathbf{y}) = \frac{a}{b^2} = \frac{4}{1500^2\,\text{h}^2} = 1.778 \times 10^{-6}\,\text{h}^{-2}.$$

Over the next $\tau = 100\,\text{h}$ the expected failure count is $\mathbb{E}[K_f \mid \mathbf{y}] = 100 \times 4/1500 = 0.2667$, and by equation (3.32) the predictive variance is

$$\text{Var}(K_f \mid \mathbf{y}) = 100 \left(\frac{4}{1500}\right) + 100^2 \left(\frac{4}{1500^2}\right) = 0.2667 + 0.0178 = 0.2844.$$

The process term is 93.75 % of the predictive variance and the rate term 6.25 %. In general the share of the rate term is

$$\frac{\tau^2 a/b^2}{\tau a/b + \tau^2 a/b^2} = \frac{\tau}{b + \tau},$$

which for $\tau = 100\,\text{h}$ and $b = 1500\,\text{h}$ is $100/1600 = 0.0625$.

The result can be checked against the full predictive distribution. A Poisson count whose mean is averaged over a gamma rate has, by a standard result, the negative binomial distribution with shape a and success probability $\pi = b/(b + \tau)$, whose mean is $a(1 - \pi)/\pi = a\tau/b$ and whose variance is $a(1 - \pi)/\pi^2 = a\tau/b + a\tau^2/b^2$: exactly the two terms of equation (3.32). For $\pi = 1500/1600 = 0.9375$ they are 0.2667 and 0.2844.

Table 3.8 extends the calculation to longer periods. The first column is the period, the second the expected count, which is also the process term, the third the rate term, the fourth the total, and the last the share of the rate term. Over 100 h most of the predictive variance comes from the random failure process, and more reliability data, which shrinks only the rate term, would barely change the prediction. Over 1500 h, as long as the exposure behind the posterior, the two terms are equal, and over 3000 h the rate term is two thirds of the total. Whether additional data collection will meaningfully improve a reliability prediction depends on this split: more evidence shrinks the second term, and the first remains as long as the future count is modeled as a Poisson process.

Table 3.8: Predictive mean and variance of the failure count over a future period τ, equation (3.32), for the posterior $\text{Gamma}(4, 1500\,\text{h})$.

τ (h)	$\mathbb{E}[K_f \mid \mathbf{y}]$ = process term	rate term $\tau^2 a/b^2$	$\text{Var}(K_f \mid \mathbf{y})$	rate share $\tau/(b+\tau)$
100	0.2667	0.0178	0.2844	0.0625
500	1.3333	0.4444	1.7778	0.2500
1500	4.0000	4.0000	8.0000	0.5000
3000	8.0000	16.0000	24.0000	0.6667

3.6.3 One posterior over a joint state

The examples so far had one unknown at a time. A plant graph carries many at once. Suppose the state vector holds two hall temperatures, the log-conductance of a heat-removal path and the failure rates of two units,

$$\mathbf{z} = \begin{bmatrix} T_1 & T_2 & g & \lambda_A & \lambda_B \end{bmatrix}^{\mathsf{T}},$$

and the evidence $\mathbf{y}$ gathers temperature observations, equipment failure records and inspection results. Everything known after the evidence is the joint posterior $p(\mathbf{z} \mid \mathbf{y}, M)$, where M names the model: the graph, its rows and its priors. Different questions are asked of this same distribution:

State estimation, $p(T_1 \mid \mathbf{y}, M)$: the marginal of one entry.

Threshold probability, $\Pr(T_2 > T_{\max} \mid \mathbf{y}, M)$: a tail of another marginal.

Parameter comparison, $\Pr(\lambda_A > \lambda_B \mid \mathbf{y}, M)$: a tail of the derived quantity $\lambda_A - \lambda_B$, which depends on the joint distribution of the two rates.

Expected derived reliability, $\mathbb{E}[R(\mathbf{z}) \mid \mathbf{y}, M]$: the average of a function of the state.

Prediction, $\Pr(T_f > \tau \mid \mathbf{y}, M)$: a posterior predictive probability, equation (3.2) with the rate taken from the state.

The methods of computation differ, but every one of them depends on a consistent probability model of the physical system and the evidence, and none needs a separate inference.

3.6.4 The posterior covariance and the variance budget

When the posterior can be approximated locally by a multivariate Gaussian, $\mathbf{z} \mid \mathbf{y}, M \approx \mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$, the covariance matrix $\boldsymbol{\Sigma}$ describes the uncertainty of each entry and how the entries vary together. The approximation takes the most probable state as the mean and the inverse curvature of the objective there as the covariance. For rows linear in the state it is exact, $\boldsymbol{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1}$ by Section 3.5.3; for nonlinear rows it is the Laplace approximation, which Chapter 4 derives with its curvature term. For two entries,

$$\boldsymbol{\Sigma} = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{12} & \Sigma_{22} \end{bmatrix}, \qquad \operatorname{corr}(z_1, z_2 \mid \mathbf{y}) = \frac{\Sigma_{12}}{\sqrt{\Sigma_{11}\Sigma_{22}}}.$$

The diagonal entries are the marginal variances, $\Sigma_{11} = \operatorname{Var}(z_1 \mid \mathbf{y})$; the off-diagonal entry is the covariance, and $\operatorname{corr}(z_1, z_2 \mid \mathbf{y})$ the correlation. The off-diagonal entries matter because physical relations and shared observations introduce posterior dependence between unknowns; they are not zero merely because the priors of the entries were independent. The hall of Section 3.5.6 is one case. Another carries the ambient temperature as an entry: the state $\mathbf{z} = (T,\ T_a,\ g)^{\mathsf{T}}$ for the hall of Figure 3.2(c), with $Q = 30\,\mathrm{kW}$, $G_{\mathrm{ref}} = 3\,\mathrm{kW/K}$, independent priors $T_a \sim \mathcal{N}(20, 1^2)$ and $g \sim \mathcal{N}(0, 0.05^2)$ (the $0.15\,\mathrm{kW/K}$ of the figure), a reading with $\sigma_T = 0.5\,\mathrm{K}$ equal to the prior prediction $30^{\circ}\mathrm{C}$, and the physics row $(T - T_a - 10\,\mathrm{K}\,\mathrm{e}^{-g})/\sigma_p$ with $\sigma_p = 0.01\,\mathrm{K}$. At the prior means every residual is zero and the rows of $\mathbf{J}$ are $(1,\ -1,\ 10)/\sigma_p$ for the physics, $(1,\ 0,\ 0)/\sigma_T$ for the reading, and $(0,\ 1,\ 0)$ and $(0,\ 0,\ 1)/0.05$ for the priors; equation (3.25) gives

$$\boldsymbol{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1} = \begin{bmatrix} 0.2083 & 0.1667 & -0.0042 \\ 0.1667 & 0.3334 & 0.0167 \\ -0.0042 & 0.0167 & 0.0021 \end{bmatrix},$$

and the correlation of T_a and g is $0.0167/\sqrt{0.3334 \times 0.0021} = 0.632$ (unrounded entries; 0.6325 at $\sigma_p = 0.001\,\mathrm{K}$). Without the reading row the same computation gives zero. A single reading leaves T_a and G positively correlated, since a higher ambient temperature and a larger conductance explain the same reading, and the correlation enters through the physics row, the only row with entries in the columns of both.

Suppose a derived quantity $f(\mathbf{z})$ has the gradient $\mathbf{g} = \partial f/\partial \mathbf{z}$ at $\hat{\mathbf{z}}$, a column with one entry per entry of the state. With $\boldsymbol{\delta} = \mathbf{z} - \hat{\mathbf{z}}$ the deviation of the state from its mean, the first-order

change is $f - f(\hat{\mathbf{z}}) \approx \mathbf{g}^{\mathsf{T}}\boldsymbol{\delta}$, a scalar equal to its own transpose, and the covariance is $\boldsymbol{\Sigma} = \mathbb{E}[\boldsymbol{\delta}\boldsymbol{\delta}^{\mathsf{T}}]$, so that

$$\mathrm{Var}(f) \approx \mathbb{E}[\mathbf{g}^{\mathsf{T}}\boldsymbol{\delta}\,\boldsymbol{\delta}^{\mathsf{T}}\mathbf{g}] = \mathbf{g}^{\mathsf{T}}\,\mathbb{E}[\boldsymbol{\delta}\boldsymbol{\delta}^{\mathsf{T}}]\,\mathbf{g} = \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}. \tag{3.33}$$

This is the *variance budget* of the output. Written out entry by entry, $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ has one term per entry of the state, the squared sensitivity of the output to the entry times the variance of the entry, plus a cross term $2g_k g_l \Sigma_{kl}$ for every pair of entries that move together. For the sum $f = z_1 + z_2$ the gradient is $\mathbf{g} = (1,\ 1)^{\mathsf{T}}$, and the budget is exact because f is linear,

$$\mathrm{Var}(z_1 + z_2) = \begin{bmatrix} 1 & 1 \end{bmatrix} \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{12} & \Sigma_{22} \end{bmatrix} \begin{bmatrix} 1 \\ 1 \end{bmatrix} = \Sigma_{11} + \Sigma_{22} + 2\Sigma_{12}. \tag{3.34}$$

With $\Sigma_{11} = \Sigma_{22} = 1\,\mathrm{K}^2$ (illustrative), the variance of the sum is $0.4\,\mathrm{K}^2$ for $\Sigma_{12} = -0.8\,\mathrm{K}^2$, $2\,\mathrm{K}^2$ for independent entries and $3.6\,\mathrm{K}^2$ for $\Sigma_{12} = 0.8\,\mathrm{K}^2$. A negative covariance makes the output more precisely known than the individual entries suggest; a positive one makes it less.

Table 3.9 works the budget for the hall temperature of equation (3.18), with the load known, $Q = 30\,\mathrm{kW}$, and the two unknowns of Figure 3.2(c) as the state $\mathbf{z} = (T_a,\ G)^{\mathsf{T}}$: T_a with standard deviation $1\,\mathrm{K}$ and G with mean $3\,\mathrm{kW/K}$ and standard deviation $0.15\,\mathrm{kW/K}$. The gradient of T and the covariance are

$$\mathbf{g} = \begin{bmatrix} \partial T/\partial T_a \\ \partial T/\partial G \end{bmatrix} = \begin{bmatrix} 1 \\ -Q/G^2 \end{bmatrix} = \begin{bmatrix} 1 \\ -3.333 \end{bmatrix}, \qquad \boldsymbol{\Sigma} = \begin{bmatrix} 1 & 0.15\,c \\ 0.15\,c & 0.0225 \end{bmatrix}, \qquad c = \mathrm{corr}(T_a, G),$$

with $\mathbf{g}$ in K per unit of each entry and $\boldsymbol{\Sigma}$ in the products of the entries' units, so that $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} = (1.00 + 0.25 - 1.00\,c)\,\mathrm{K}^2$: the ambient term $1^2 \times 1$, the conductance term $3.333^2 \times 0.0225$ and the cross term $2 \times 1 \times (-3.333) \times 0.15\,c$. The columns of the table are the correlation, the three terms, the variance and the standard deviation of T. The two marginals are the same in every row, and the standard deviation of the hall temperature ranges from 0.67 to 1.43 K with the correlation alone.

Table 3.9: Variance budget of the hall temperature $T = T_a + Q/G$, equation (3.33), with $Q = 30\,\mathrm{kW}$, G of mean $3\,\mathrm{kW/K}$ and standard deviation $0.15\,\mathrm{kW/K}$, T_a of standard deviation $1\,\mathrm{K}$, and correlation $c = \mathrm{corr}(T_a, G)$ (illustrative). Terms in K^2.

c	ambient term	conductance term	cross term	$\mathrm{Var}(T)$ (K^2)	σ (K)
-0.8	1.00	0.25	$+0.80$	2.05	1.43
0	1.00	0.25	0	1.25	1.12
0.8	1.00	0.25	-0.80	0.45	0.67

Two consequences follow from equation (3.33). An entry can be very uncertain and still not matter: if the output is insensitive to it, its whole term is near zero, and narrowing its variance buys nothing. And an entry can be poorly known while the output is well known: when the sensitivities and the covariance have opposite signs, the cross term subtracts from the diagonal terms. A measurement is worth making when it shrinks a variance Σ_{kk} whose sensitivity is large, and worthless when the sensitivity is near zero.

3.6.5 Reading the shape of the posterior

The second consequence deserves a worked case, because it is where marginal spreads mislead. Carry the hall temperature, the load and the conductance as entries, the last two through the logarithms $\ell_Q = \ln(Q/Q_{\mathrm{ref}})$ and $g = \ln(G/G_{\mathrm{ref}})$, each with an independent prior of mean 0 and standard deviation 0.10, that is about 10 %, and with the ambient temperature pinned. The

references are chosen so that $Q_{\text{ref}}/G_{\text{ref}} = 10\,\text{K}$, a temperature rise of 10 K at the prior means, and a sensor with $\sigma_T = 0.2\,\text{K}$ reads that temperature, $y_T = T_a + 10\,\text{K}$. The state, its four rows and their 4×3 Jacobian are

$$\mathbf{z} = \begin{bmatrix} T \\ \ell_Q \\ g \end{bmatrix}, \qquad r(\mathbf{z}) = \begin{bmatrix} (T - y_T)/\sigma_T \\ \ell_Q/0.10 \\ g/0.10 \\ (T - T_a - 10\,\text{K}\,\mathrm{e}^{\ell_Q - g})/\sigma_p \end{bmatrix},$$

$$\mathbf{J} = \begin{bmatrix} 1/\sigma_T & 0 & 0 \\ 0 & 10 & 0 \\ 0 & 0 & 10 \\ 1/\sigma_p & -10\,\mathrm{e}^{\ell_Q - g}/\sigma_p & 10\,\mathrm{e}^{\ell_Q - g}/\sigma_p \end{bmatrix},$$

with the physics row of equation (3.18) carrying the tolerance $\sigma_p = 0.001\,\text{K}$. At the prior means, $\mathbf{z}_0 = (T_a + 10\,\text{K},\ 0,\ 0)^{\mathsf{T}}$, every residual is zero, the normal equations (3.24) give a zero step, and $\hat{\mathbf{z}} = \mathbf{z}_0$; the curvature term that Chapter 4 adds for nonlinear rows is a sum of residuals times second derivatives and vanishes with the residuals. There the Jacobian is

$$\mathbf{J} = \begin{bmatrix} 5 & 0 & 0 \\ 0 & 10 & 0 \\ 0 & 0 & 10 \\ 1000 & -10\,000 & 10\,000 \end{bmatrix},$$

whose last two columns carry the relation between load and conductance: both are zero in the reading row and equal and opposite in the physics row, so the direction $(0,\ 1,\ 1)^{\mathsf{T}}$, load and conductance raised together, gives $\mathbf{J}\,(0,\ 1,\ 1)^{\mathsf{T}} = (0,\ 10,\ 10,\ 0)^{\mathsf{T}}$ and is seen by the prior rows alone. Multiplying by $\mathbf{J}^{\mathsf{T}}$,

$$\mathbf{J}^{\mathsf{T}}\mathbf{J}\,(0,\ 1,\ 1)^{\mathsf{T}} = 100\,(0,\ 1,\ 1)^{\mathsf{T}},$$

so $(0,\ 1,\ 1)^{\mathsf{T}}$ is an eigenvector of $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ with the eigenvalue $100 = 1/0.10^2$, the prior precision alone, and of $\boldsymbol{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1}$ with the eigenvalue 0.0100, the prior variance: the reading has left that direction untouched. Inverting $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ numerically, equation (3.25) gives

$$\boldsymbol{\Sigma} = \begin{bmatrix} 0.039216 & 0.001961 & -0.001961 \\ 0.001961 & 0.005098 & 0.004902 \\ -0.001961 & 0.004902 & 0.005098 \end{bmatrix}.$$

For a Gaussian, the integral of equation (3.15) over T leaves the block of $\boldsymbol{\Sigma}$ in the rows and columns of ℓ_Q and g, the lower right 2×2 block. Its precision is not the corresponding block of $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ but, by the block-inversion formula of Section 3.5.3 applied to $\boldsymbol{\Lambda} = \mathbf{J}^{\mathsf{T}}\mathbf{J}$, the Schur complement of the T block, $\boldsymbol{\Lambda}_{qq} - \boldsymbol{\Lambda}_{qT}\boldsymbol{\Lambda}_{TT}^{-1}\boldsymbol{\Lambda}_{Tq}$ with q the indices of ℓ_Q and g:

$$\boldsymbol{\Lambda}_{qq} - \boldsymbol{\Lambda}_{qT}\boldsymbol{\Lambda}_{TT}^{-1}\boldsymbol{\Lambda}_{Tq} = \begin{bmatrix} 2599.94 & -2499.94 \\ -2499.94 & 2599.94 \end{bmatrix} \longrightarrow \begin{bmatrix} 2600 & -2500 \\ -2500 & 2600 \end{bmatrix} \quad (\sigma_p \to 0),$$

against the block $\boldsymbol{\Lambda}_{qq}$ of $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ itself, whose entries are $10^8 + 100$ on the diagonal and -10^8 off it. Each marginal standard deviation is $\sqrt{0.005098} = 0.0714$, down only from 0.100, and the correlation is $0.004902/0.005098 = 0.962$. The eigenvectors of the block are $(1,\ 1)^{\mathsf{T}}/\sqrt{2}$ and $(1,\ -1)^{\mathsf{T}}/\sqrt{2}$, with the eigenvalues 0.005098 ± 0.004902, that is 0.010000 and 0.000196: the posterior is as wide as the prior along the first and 7.1 times narrower along the second. The budget (3.33) with the gradients $\mathbf{g} = (0,\ 1,\ -1)^{\mathsf{T}}$ and $(0,\ 1,\ 1)^{\mathsf{T}}$ gives the variance of the difference, $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} = 0.000392$, a standard deviation of 0.0198, and that of the sum, 0.0200, a standard deviation of 0.1414, exactly its prior value $\sqrt{2} \times 0.10$. Repeated at $\sigma_p = 0.01\,\text{K}$, the block entries print the same, 0.005098 and 0.004902, and the correlation prints 0.961 against 0.962 at 0.001 K; the limit is reached at 0.001 K.

Remark 3.2. Eliminating T through the physics row turns the reading into the linearized row $(\ell_Q - g)/0.02$, since $\sigma_T/(T - T_a) = 0.02$, and the same block follows from a 3×2 Jacobian; along the difference alone, the prior precision 50 plus the reading's 2500 gives $1/2550 = 0.000392$.

Figure 3.7 draws the two distributions in the plane of the two entries. The horizontal axis is ℓ_Q and the vertical axis g; the dashed gray circle encloses 95 % of the prior, the green ellipse 95 % of the posterior, and the thin line is $\ell_Q - g = 0$, the value of the difference that the reading fixed. The posterior is a narrow ellipse lying along that line, its axes the two eigenvectors of the block: it is long in the direction the reading cannot see, $(1,\ 1)^{\mathsf{T}}$, where it reaches the prior circle because the reading added nothing there, and short in the direction it measured, $(1,\ -1)^{\mathsf{T}}$.

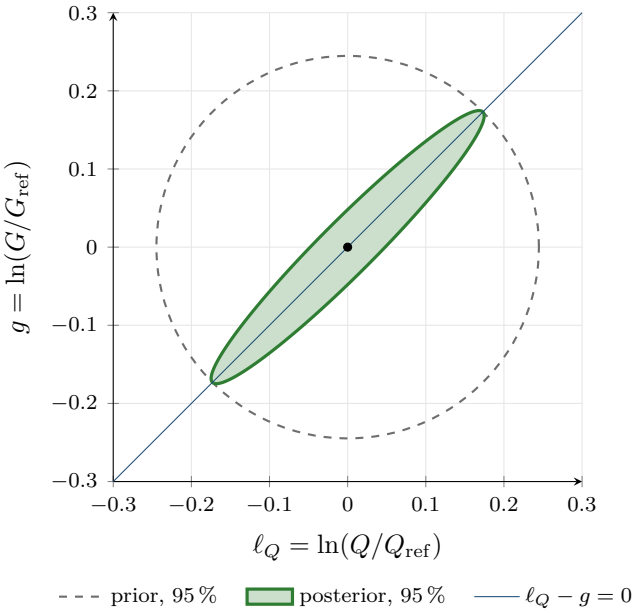


Figure 3.7: Prior and posterior of the log-load ℓ_Q and the log-conductance g after one hall-temperature reading that fixes $\ell_Q - g$ within 0.02 (illustrative): the (ℓ_Q, g) block of the posterior of $\mathbf{z} = (T, \ell_Q, g)^{\mathsf{T}}$. Prior: independent, standard deviations 0.10. Posterior: standard deviations 0.0714, correlation 0.962; standard deviation of $\ell_Q - g$, 0.0198; of $\ell_Q + g$, 0.1414. Contours enclose 95 % ($\chi^2_{2,0.95} = 5.991$).

Read by its marginals, this posterior says that neither the load nor the conductance is well known, about 7 % each. Read by its shape, it says that their ratio, and with it the hall temperature at the present load, is known to 2 %. A prediction at a different, specified load Q', $T' = T_a + Q'/G$, has a gradient with respect to (T, ℓ_Q, g) proportional to $(0,\ 0,\ 1)^{\mathsf{T}}$, so $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ picks out Σ_{gg} and the prediction inherits its 7 % spread; a prediction that scales the present load, such as the hall at twice the unknown Q, has a gradient proportional to $(0,\ 1,\ -1)^{\mathsf{T}}$, the short axis of the ellipse, and its temperature rise is known to the same 2 % relative as the present one; the absolute spread doubles. Which statement holds depends on the gradient of the question, and the covariance, not the list of marginal variances, carries the answer.

Principle 3.4. *Read the posterior's shape, not its marginals. Wide marginal spreads with a strong correlation mean that the evidence determines a combination of the entries; a quantity that depends on that combination is well known, and the entries are not unidentified merely because each alone is uncertain.*

3.6.6 Mixtures over structure

A Gaussian approximation is most useful for continuous uncertainty near a well-identified posterior mode. Conditional probability is more general. If the system may operate under either of two structures, $M \in \{M_1, M_2\}$ (two drawings of the plant), the law of total probability, applied over the models with $\mathbf{z}$ conditioned on the evidence throughout, gives

$$p(\mathbf{z} \mid \mathbf{y}) = \sum_j p(\mathbf{z}, M_j \mid \mathbf{y}) = \sum_j p(\mathbf{z} \mid \mathbf{y}, M_j)\ \Pr(M_j \mid \mathbf{y}), \tag{3.35}$$

where the second step is the multiplication rule. The weights follow from Bayes' theorem over the models, $\Pr(M_j \mid \mathbf{y}) = p(\mathbf{y} \mid M_j)\ \Pr(M_j) / \sum_i p(\mathbf{y} \mid M_i)\ \Pr(M_i)$, so that the evidence of each model, the normalizing constant of its own posterior, sets its share. The expression assumes that the relevant state $\mathbf{z}$ has a common interpretation across the models. Such a posterior need not be Gaussian, even when each model's posterior is.

A two-structure illustration on the hall temperature shows what is lost. Let the drawing M_1, the heat-removal path as designed, give $T \mid \mathbf{y}, M_1 \sim \mathcal{N}(26.9, 0.5^2)$, and a drawing M_2 with a degraded path give $\mathcal{N}(28.5, 0.5^2)$, with posterior model probabilities 0.7 and 0.3 (illustrative). By the law of total expectation the mixture mean is $0.7 \times 26.9 + 0.3 \times 28.5 = 27.38°\mathrm{C}$; by the law of total variance, equation (3.31) with the model in place of the rate, its variance is the mean of the within-model variances plus the variance of the model means,

$$\mathrm{Var}(T \mid \mathbf{y}) = 0.5^2 + 0.7 \times 0.3 \times (28.5 - 26.9)^2 = 0.25 + 0.5376 = 0.7876\,\mathrm{K}^2,$$

a standard deviation of 0.887 K. The probability of exceeding 28°C is the weighted sum of the two tails,

$$\Pr(T > 28°\mathrm{C} \mid \mathbf{y}) = 0.7\,\big(1 - \Phi(2.2)\big) + 0.3\,\big(1 - \Phi(-1)\big) = 0.7 \times 0.0139 + 0.3 \times 0.8413 = 0.262.$$

The most probable drawing alone would report 0.014, and a single Gaussian with the mixture's mean and variance would report $1 - \Phi(0.699) = 0.242$ and hide the two separate modes, 3.2 standard deviations apart. The law of total probability combines the uncertainty over the structure with the uncertainty over the numerical state, and this combination becomes central in fault diagnosis and in the identification of a system's structure. The next chapter assembles these pieces into the complete workflow of probabilistic inference, from the plant graph and its likelihood to the posterior, its marginals, its expectations and the probabilities of functions of the state.

3.7 What to retain

Definition 3.1 (Conditioning, marginalization and dependence)**.** A conditional probability restricts attention to the situations consistent with the evidence and renormalizes, $\Pr(A \mid B) = \Pr(A \cap B) / \Pr(B)$; Bayes' theorem reverses the conditioning through the prior and the evidence; marginalization sums or integrates over every unknown the question does not name; and the posterior covariance of a joint state, not its marginal variances, decides how uncertain a derived quantity is.

Four ideas carry forward. Reliability problems contain two kinds of uncertainty: aleatory uncertainty is the variability of outcomes, and epistemic uncertainty is incomplete knowledge of the system; evidence reduces the second and leaves the first, and the law of total variance separates them in a prediction. Conditional probability is how information is incorporated: Bayes' theorem updates a probability when an observation arrives, and the base rate decides how far an alarm moves it. Engineering quantities are often dependent: a shared electrical

supply, a conservation relation or a common environment creates relations among components and unknown states, and conditional independence simplifies a calculation only where the model justifies it. A physics graph puts the shared physical causes in the state; the independence it still assumes is that of the errors, and a shared error, such as a calibration offset, is removed by carrying it as an entry too. Marginalization answers a focused question without pretending that the other unknowns are known, whether they are nuisance quantities, competing conditions or competing structures of the plant.

Together these establish the mechanism of inference: the prior $p(\mathbf{z})$ and the observation model $p(\mathbf{y} \mid \mathbf{z})$ combine into the posterior $p(\mathbf{z} \mid \mathbf{y})$, the update multiplying the prior by the likelihood and normalizing the product. For Gaussian rows the product is $\exp(-\frac{1}{2} r^{\mathsf{T}} r)$ over the whole state, physics rows included; the normal equations give its most probable state, $\boldsymbol{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1}$ its covariance, exact when the rows are linear, and the conditional of the joint Gaussian of state and readings, through the Schur complement, gives the same answer. Relations between quantities are read in that space: a reading informs an entry it does not measure through a row of $\mathbf{J}$ that carries both, and a combination of entries is well known when it lies along a short eigenvector of $\boldsymbol{\Sigma}$, which $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ measures. Once the posterior is available, the engineering question decides the next calculation: a marginal probability, an expectation, a prediction, a model comparison or a decision. Table 3.10 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 3.10: The terms of uncertainty and conditional probability and where each is defined.

Term	Expression	Units	Meaning	Defined in
Aleatory uncertainty	$\Pr(T_f > t \mid \lambda) = \mathrm{e}^{-\lambda t}$	–	outcomes vary with known parameters	3.1.1
Epistemic uncertainty	$p(\lambda \mid \mathbf{y})$	h	knowledge of a fixed quantity	3.1.2
Predictive reliability	$\int \mathrm{e}^{-\lambda t} p(\lambda \mid \mathbf{y})\,\mathrm{d}\lambda$	–	survival averaged over the rate	3.1.3
Conditional probability	$\Pr(A \cap B)/\Pr(B)$	–	probability of A among cases of B	3.2.2
Multiplication rule	$\Pr(A \cap B) = \Pr(B \mid A)\Pr(A)$	–	joint from conditional	3.2.3
Bayes' theorem	$\Pr(B \mid A)\Pr(A)/\Pr(B)$	–	reverses the conditioning	3.2.4
Posterior odds	likelihood ratio $\times$ prior odds	–	Bayes without the evidence	3.3.3
Law of total probability	$\sum_s \Pr(H \mid s)\Pr(s)$	–	marginal of an observation	3.4.1
Marginal likelihood	$\int p(\mathbf{y} \mid \lambda) p(\lambda)\,\mathrm{d}\lambda$	–	evidence; normalizer	3.4.2
Nuisance marginal	$\int p(\lambda, \psi \mid \mathbf{y})\,\mathrm{d}\psi$	varies	integrate out, do not maximize	3.4.3
Conditional independence	$X_1 \perp X_2 \mid C$	–	joint given c factorizes	3.4.6
Common-cause joint	$\sum_c \Pr(F_A \mid c)\Pr(F_B \mid c)\Pr(c)$	–	joint failure with a shared cause	3.4.5
Correlated reading errors	$\mathbf{A}^{\mathsf{T}}\mathbf{R}^{-1}\mathbf{A}$, $\mathbf{R}$ full; or a shared bias entry	K^{-2}	the precision the readings truly add	3.5.7
Gaussian rows	$p(\mathbf{z} \mid \mathbf{y}) \propto \exp(-\frac{1}{2} r^{\mathsf{T}} r)$	–	posterior of the whole state	3.5.2
Normal equations	$(\mathbf{J}^{\mathsf{T}}\mathbf{J})\,\boldsymbol{\Delta} = -\mathbf{J}^{\mathsf{T}} r$, $\boldsymbol{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1}$	varies	most probable state, covariance	3.5.3
Precisions add	$\boldsymbol{\Sigma}^{-1} = \boldsymbol{\Sigma}_0^{-1} + \mathbf{J}_y^{\mathsf{T}}\mathbf{J}_y$	varies	prior plus reading information	3.5.3
Gaussian conditional	$\boldsymbol{\Sigma}_{zz} - \boldsymbol{\Sigma}_{zy}\boldsymbol{\Sigma}_{yy}^{-1}\boldsymbol{\Sigma}_{yz}$	varies	Schur complement	3.5.3
Exceedance probability	$1 - \Phi\big((T_{\max} - \mu_{\mathrm{post}})/\sigma_{\mathrm{post}}\big)$	–	tail of the posterior	3.5.5
Law of total variance	$\mathbb{E}[\mathrm{Var}(K_f \mid \lambda)] + \mathrm{Var}(\mathbb{E}[K_f \mid \lambda])$	–	process plus rate terms	3.6.1
Posterior covariance	$\boldsymbol{\Sigma}$, $\mathrm{corr}(z_1, z_2) = \Sigma_{12}/\sqrt{\Sigma_{11}\Sigma_{22}}$	varies	joint spread of the state	3.6.4
Variance budget	$\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$	varies	sensitivity^2 $\times$ variance, plus cross terms	3.6.4
Mixture posterior	$\sum_j p(\mathbf{z} \mid \mathbf{y}, M_j)\Pr(M_j \mid \mathbf{y})$	varies	uncertainty over structure	3.6.6

3.8 Exercises

Exercise 3.1 (Formulation)**.** (a) During 2000 comparable missions pump A failed in 100, pump B in 60, and both in 25. Compute $\Pr(F_A)$, $\Pr(F_B)$, $\Pr(F_A \cap F_B)$, $\Pr(F_B \mid F_A)$ and $\Pr(F_A \mid F_B)$, and state whether the failures are independent. (0.05, 0.03, 0.0125, $25/100 = 0.25$, $25/60 = 0.417$; not independent, since $0.0125 \neq 0.05 \times 0.03 = 0.0015$.) (b) Two redundant components have $\Pr(F_A) = 0.05$, $\Pr(F_B) = 0.04$ and $\Pr(F_A \cap F_B) = 0.015$. Compute the joint probability under an incorrect independence assumption, the reliability of a system that fails only when both fail, and $\Pr(F_B \mid F_A)$, and state the engineering significance. ($0.05 \times 0.04 = 0.002$;

$1 - 0.015 = 0.985$, against 0.998 under independence; $0.015/0.05 = 0.30$; the pair fails together 7.5 times as often as independence predicts, so the redundancy protects far less than the product suggests.)

Exercise 3.2 (Formulation)**.** A hall satisfies $T = T_a + Q/G$ with T_a known, $G > 0$, and both Q and G uncertain. Starting from $\Pr(T > T_{\max} \mid \mathbf{y})$, write an equivalent event in Q and G and express the probability as an integral over their joint posterior, without assuming independence. (Since $G > 0$, $T_a + Q/G > T_{\max} \iff Q > G(T_{\max} - T_a)$, so $\Pr(T > T_{\max} \mid \mathbf{y}) = \iint \mathbf{1}[Q > G(T_{\max} - T_a)]\, p(Q, G \mid \mathbf{y})\, \mathrm{d}Q\, \mathrm{d}G$; the joint posterior keeps the correlation between Q and G that a reading of T creates.)

Exercise 3.3 (Derivation)**.** (a) A vibration alarm monitors a bearing with $\Pr(F) = 0.02$, $\Pr(H \mid F) = 0.95$ and $\Pr(H \mid \neg F) = 0.05$. Derive the probability that the bearing is faulty when the alarm activates, and explain why it is not 95 %. ($\Pr(H) = 0.95 \times 0.02 + 0.05 \times 0.98 = 0.068$; $\Pr(F \mid H) = 0.019/0.068 = 0.279$; 0.95 is $\Pr(H \mid F)$, not $\Pr(F \mid H)$, and the low base rate makes most alarms false.) (b) With prior $\Pr(F) = 0.10$, sensor A has $\Pr(H_A \mid F) = 0.9$, $\Pr(H_A \mid \neg F) = 0.1$, and sensor B has $\Pr(H_B \mid F) = 0.8$, $\Pr(H_B \mid \neg F) = 0.2$; the alarms are conditionally independent given the fault state. Derive $\Pr(F \mid H_A, H_B)$ and explain why the assumption is needed. ($\Pr(H_A, H_B \mid F) = 0.9 \times 0.8 = 0.72$ and $\Pr(H_A, H_B \mid \neg F) = 0.1 \times 0.2 = 0.02$, so $\Pr(F \mid H_A, H_B) = 0.072/(0.072 + 0.018) = 0.80$; sequentially, sensor A gives 0.5 and sensor B then 0.8; the likelihoods multiply only under conditional independence given each fault state, and a shared sensor bias or environmental disturbance would invalidate it.)

Exercise 3.4 (Derivation)**.** Future failures satisfy $K_f \mid \lambda \sim \text{Poisson}(\lambda\tau)$ with $\lambda \mid \mathbf{y} \sim \text{Gamma}(3, 1000\,\text{h})$ (shape and rate) and $\tau = 200\,\text{h}$. Derive the expected count, the predictive variance and its two contributions, and state what happens to each when much more evidence supporting a similar rate arrives. ($\mathbb{E}[K_f \mid \mathbf{y}] = 200 \times 0.003 = 0.6$; by equation (3.32), $\text{Var}(K_f \mid \mathbf{y}) = 0.6 + 200^2 \times 3/1000^2 = 0.6 + 0.12 = 0.72$; process 0.6, rate 0.12, a share $200/1200 = 0.167$. With a and b growing in proportion, the rate term $\tau^2(a/b)/b$ tends to zero while the process term stays near 0.6.)

Exercise 3.5 (Application)**.** A system operates under normal conditions with probability 0.90 and failure probability 0.01, and under severe conditions with probability 0.10 and failure probability 0.15. Derive the unconditional failure probability, and the probability that severe conditions were present given that the system failed. ($\Pr(F) = 0.90 \times 0.01 + 0.10 \times 0.15 = 0.024$; $\Pr(\text{severe} \mid F) = 0.015/0.024 = 0.625$.)

Exercise 3.6 (Application)**.** (a) A hall temperature has prior $\mathcal{N}(25, 1.5^2)$, and a sensor with $\sigma_T = 0.5\,\text{K}$ reads $y_T = 26°\text{C}$. For the state, the one number T, write the two residuals and their slopes j_k, the information Λ, take the Newton step from the prior mean, and compute σ^2_{post} and $\Pr(T > 27°\text{C} \mid y_T)$. ($r_1 = (T - 25)/1.5$ and $r_2 = (T - 26)/0.5$, $j_1 = 0.6667$ and $j_2 = 2\,\text{K}^{-1}$, $\Lambda = 0.4444 + 4 = 4.4444\,\text{K}^{-2}$; at $T = 25$, $r_1 = 0$ and $r_2 = -2$, so $\mathcal{L}' = j_1 r_1 + j_2 r_2 = -4$, $\hat{\Delta} = 4/4.4444 = 0.9$ and $\hat{T} = 25.90°\text{C}$; $\sigma^2_{\text{post}} = 1/4.4444 = 0.225\,\text{K}^2$, $\sigma_{\text{post}} = 0.4743\,\text{K}$; $1 - \Phi(1.10/0.4743) = 1 - \Phi(2.319) = 0.0102$.) (b) For the two-entry hall of equation (1.10), whose Jacobian at the forward state (23°C, 0) has the rows (100, 500), (2, 0) and (0, 3.333), compute the three outer products $\mathbf{J}_k^{\mathsf{T}}\mathbf{J}_k$ and their sum, and state which row creates the off-diagonal entry and what the rank of its term is. (Physics $\begin{bmatrix} 10\,000 & 50\,000 \\ 50\,000 & 250\,000 \end{bmatrix}$, reading diag(4, 0), prior diag(0, 11.11); sum $\begin{bmatrix} 10\,004 & 50\,000 \\ 50\,000 & 250\,011 \end{bmatrix}$; only the physics row has entries in both columns, and its term has rank one, determinant $10^4 \times 2.5 \times 10^5 - (5 \times 10^4)^2 = 0$.)

Chapter 4

The Anatomy of Probabilistic Inference

4.1 From a plant graph to an engineering probability

An equipment hall with one heat-removal path to the outside air has been run once at a test load of 30 kW, and its temperature sensor read 23.2°C. The operator proposes to run it at 60 kW, and the hall must stay below 28°C. The engineer who must approve the change has that one reading, a handbook value for the thermal resistance of the cooling path and nothing at all from full load. The handbook gives one quick answer: a resistance of 0.160 K/kW puts the hall at $18 + 60 \times 0.160 = 27.6$°C at full load, under the limit, so the load can be raised. The reading, its rise above ambient doubled with the load, gives a second quick answer: a hall just above the limit.

The two quick answers disagree, and neither says how sure the operator can be: the handbook value ignores the reading, and the scaled reading ignores the handbook and the error of the sensor. What neither settles is what one reading at the test condition, the handbook value and the physics of the hall together say about a condition that has never been run. The question of this chapter is

how a posterior distribution over the state of a plant graph is obtained from physical relations, readings and prior information, and how that posterior is turned into the probability that answers an engineering question.

Chapter 1 built the plant graph, Chapter 2 classified the questions and Chapter 3 derived the rules one at a time; here they are used together for the first time, on a state of three entries of which the reading touches only one.

Jargon: maximum a posteriori (MAP) estimate, Hessian, Laplace approximation, delta method, change of variables, credible interval, prior and posterior predictive distributions, marginal likelihood, Bayes factor, rank-one update

The *maximum a posteriori (MAP) estimate* is the value of the unknowns at which the posterior density is highest; it is found by minimizing the negative logarithm of the posterior, the *objective*. The *Hessian* is the matrix of second derivatives of the objective; at the MAP estimate it measures the curvature of the objective, and for rows linear in the state and divided by their standard deviations it is $\mathbf{J}^\mathsf{T}\mathbf{J}$, the information matrix. The *Laplace approximation* replaces the posterior by the Gaussian whose mean is the MAP estimate and whose covariance is the inverse of the Hessian there. The *delta method* approximates the variance of a function of the unknowns by linearizing the function at the estimate. A *change of variables* carries a density from one coordinate to another, multiplying it by the derivative of the old coordinate with respect to the new so that each interval keeps its probability. A *credible interval* contains the unknown with a stated posterior probability. The *prior predictive* and *posterior predictive* distributions are the distributions of a reading not yet seen, averaged over the prior and over the posterior respectively. The *marginal likelihood*, or *evidence*, of a model is the probability density of the observed readings averaged over the prior; the *Bayes factor* is the ratio of the evidences of two models. A *rank-one update* adds one new reading to an existing Gaussian posterior by a closed-form correction of its mean and covariance, without solving the problem again.

Key idea

A reliability probability is computed in two operations. Inference turns the plant graph, the readings and the priors into a posterior over the state: its mean is the MAP state and its covariance is the inverse curvature of the objective. Query evaluation turns that posterior into the engineering answer: an entry of the state and its variance, or a derived quantity $f(\mathbf{z})$ with variance $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$, and then a normal tail at the limit.

4.1.1 The complete inference workflow

The method fixed in the preface is a progression: the plant graph (step 0), then the questions written over its state (step 1), the graph extended for them with the evidence attached (step 2), the posterior of the state that the rows give, derived and computed once (step 3), the answers read from that posterior (step 4), their checks (step 5), the report with its bands (step 6), and the next measurement when a band is too wide (step 7). Figure 4.1 draws the progression. The wide box at the top is the plant graph, built once; the seven boxes beneath are one pass that carries every question, read left to right along the top row and right to left along the bottom row; the second line of each box names what the step produces. The dashed arrow from step 7 back to step 3 is the loop taken when a new measurement is bought: its row is added and the posterior is rebuilt.

The distinction between steps 3 and 4 carries the method, and it is the distinction between the two operations of every inference. Step 3 is *inference*: the rows of step 2 are combined into one posterior of the state, its MAP state, information and covariance, derived and then computed for the plant; it does not depend on which question is asked. Step 4 is *query evaluation*: each question is reduced, from its engineering statement through its general and specialized forms, to quantities read from that posterior, in this chapter the posterior mean and variance of one state entry, the full-load temperature, and the standard normal distribution function evaluated at the limit. Keeping them apart is what lets one posterior serve every question, and a new reading change step 3 alone.

The order within the chapter follows from it. What the answers rest on comes first: the

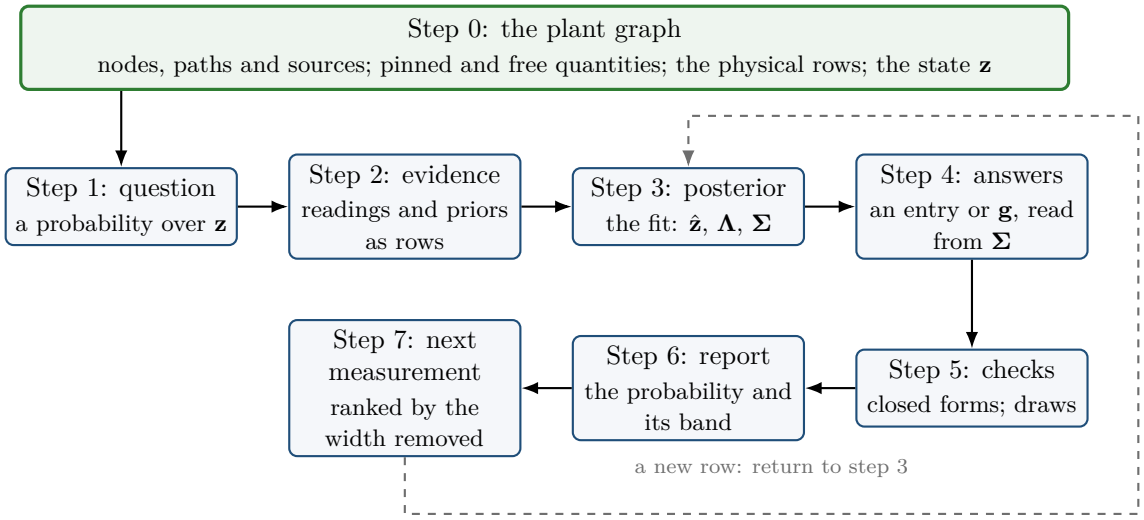


Figure 4.1: The inference workflow. Step 0 is built once; steps 1 to 7 are one pass, carrying every question of the plant. The second line of each box names what the step produces.

modeling assumptions and the rows they produce in step 2, the posterior in step 3, each stated and computed before anything uses it. Each answer is then derived top-down in step 4, from the question the engineer asked to the quantities it needs, and every one of those points back to the posterior, never forward.

4.1.2 The posterior and the questions put to it

The calculation divides into two operations of different kinds:

$$\underbrace{(\text{graph},\ \mathbf{y}) \longrightarrow p(\mathbf{z} \mid \mathbf{y})}_{\text{inference}}, \qquad \underbrace{p(\mathbf{z} \mid \mathbf{y}) \longrightarrow \text{engineering probability or decision}}_{\text{query evaluation}}, \tag{4.1}$$

the second applied to the result of the first. Inference is done once for a given graph and a given body of evidence. Under a Gaussian posterior, exact or approximate, its result is the triple of the MAP state $\hat{\mathbf{z}}$, the posterior covariance $\boldsymbol{\Sigma}$, and the Jacobian of the model at $\hat{\mathbf{z}}$, from which the gradient $\mathbf{g}$ of any derived quantity is formed. Query evaluation is then a formula on the triple. An output that is itself an entry z_i of the state, such as a temperature at a node, has mean $\hat{z}_i$ and variance Σ_{ii}, read straight off the posterior. An output that is a function $f(\mathbf{z})$ of several entries needs its gradient and the quadratic form $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$. A probability that the output stays below a limit is a normal tail on either pair.

Two kinds of further work follow from this division and must be kept apart. A new question asked of the same evidence needs no new inference: a new function of the state is pushed through the existing $\boldsymbol{\Sigma}$. New evidence changes the posterior itself: either the inference is repeated with the new row, or the existing Gaussian is conditioned on the new row by a rank-one update, which Section 4.6.4 derives. The posterior represents what is known about the plant under the model; the queries represent what the engineer wants to determine.

4.1.3 The inference in the terms of graphical models

The operations of this book are those of inference in a probabilistic graphical model, under the names a plant engineer uses. The plant graph with its rows is a factor graph (Section 1.1.3): each row is a factor over the entries it touches, and the posterior $\mathrm{e}^{-\mathcal{L}(\mathbf{z})}$ is the normalized product of the factors. Table 4.1 sets each operation beside its name in the literature on graphical

models. Its first column is the operation as this book writes it, the second the standard name, and the third the chapters that derive it.

Table 4.1: The operations of the book and their names in the literature on probabilistic graphical models.

Operation in this book	Name in graphical models	Derived in
a row: prior, physics or reading	a factor of a factor graph	Chapters 1, 3
the posterior $e^{-\mathcal{L}(\mathbf{z})}$	the normalized product of the factors	Chapter 3
the normal equations, $\mathbf{\Lambda} = \mathbf{J}^\mathsf{T}\mathbf{J}$, $\mathbf{\Sigma} = \mathbf{\Lambda}^{-1}$	exact inference in a linear-Gaussian factor graph, in information form	Chapter 3; this chapter
Gauss–Newton and the Laplace covariance	the Gaussian approximation at the mode of a nonlinear factor graph (factor-graph optimization)	this chapter; Chapter 7
the Schur complement; the marginal of a joint Gaussian	variable elimination of the entries not asked about	Chapter 3
the enumeration of the 2^n configurations; total probability	exact inference in a discrete Bayesian network by summation	Chapters 3, 13
the Beta and gamma updates	exact inference with conjugate factors	Chapters 5, 6
the evidence $p(\mathbf{y} \mid M)$	the normalizing constant (partition function) of the product	Chapter 2
the reweighted draws	importance sampling, with the Laplace normal as the proposal	Chapter 20, in Volume II

The general algorithms of that literature are not needed for the plants of this book, and the table shows why: every posterior here is solved jointly, by factoring the information matrix of the whole state, which is exact for linear-Gaussian rows and costs little at a few dozen entries. Message passing, belief propagation (the sum-product algorithm) and the junction tree compute the same marginals by local operations on the graph, and they earn their place on graphs too large to factor jointly or with discrete entries of many levels; on a Gaussian graph without loops, belief propagation returns exactly the Schur complements of the joint factorization. Factoring a large sparse information matrix is itself variable elimination, and its cost is set by the order in which the entries are eliminated: eliminated in time order, a state indexed by time gives the Kalman filter and smoother. Chapter 35, in Volume III, meets a time-indexed state on which the order matters. Markov chain Monte Carlo (Gibbs sampling and the Metropolis algorithm) draws from a posterior that the Laplace normal cannot approximate and reweighting cannot correct: one with several modes, or the funnel of a hierarchical model, in which the spread of a population of rates is itself unknown. Chapter 33, in Volume III, meets it in a fleet of pumps. Variational inference is not used in this book.

These algorithms differ in what they give up. Variable elimination, the junction tree, belief propagation on a graph without loops and the Kalman smoother on linear-Gaussian rows are exact: they compute the posterior the joint factorization computes, in another order. The Laplace approximation, loopy belief propagation, the linearized Kalman filters and variational inference are approximate, and their error is a bias that more computation does not remove; loopy belief propagation and variational inference in particular tend to report variances too small, which for a tail probability against a limit is the error that matters most. Importance sampling and Markov chain Monte Carlo are exact in the limit of many draws, and their error is random and can be estimated, although a chain that never visits a second mode can look converged. The two kinds of error check each other, which is why the draws serve as the other route against the Laplace answer in step 5. The route also decides which objects come cheaply. The joint factorization gives $\hat{\mathbf{z}}$, $\mathbf{\Sigma}$, the gradient and the preposterior covariance $\mathbf{\Sigma}_a$ of a rank-one

update in one pass, and the evidence in closed form under the Laplace approximation; draws give the posterior as a sample, from which a variance is an estimate, a preposterior needs the candidate's outcomes simulated and the posterior refitted or reweighted, and the evidence needs a further computation, since the draws do not carry the normalizing constant. The method is unchanged by the choice: the plant graph, the questions, the rows and the objective $\mathcal{L}(\mathbf{z})$ are the same for every algorithm, the general forms of Chapter 2 are expectations over the posterior whatever computes it, and only the computations of steps 3 and 4, and the checks of step 5 that test them, depend on the route.

4.1.4 The plant graph under two operating conditions (step 0)

The plant is an equipment hall with a heat-generating load and a single effective heat-removal path to the ambient air at $T_a = 18°\mathrm{C}$. Two operating conditions are considered: a test condition with a heat load of $Q_1 = 30\,\mathrm{kW}$, and a proposed full-load condition with $Q_2 = 60\,\mathrm{kW}$. At steady state the heat supplied to the hall leaves it through the path, and energy conservation at the hall node reads

$$Q = G\,(T - T_a), \tag{4.2}$$

where T is the hall temperature in °C, Q the heat load in kW, and G the effective thermal conductance of the path in kW/K. The balance is taken at steady state through one effective path: this reduces the cooling of the hall to a single unknown, and it fails in the minutes after a load change, while heat stored in the hall and its equipment is still moving, which Section 4.6.6 and Exercise 4.6 take up. It is convenient to work with the thermal resistance of the path,

$$\rho = \frac{1}{G}, \tag{4.3}$$

in K/kW, with which equation (4.2) becomes

$$T = T_a + Q\rho. \tag{4.4}$$

Carrying the resistance in the state, rather than the conductance, is a choice: the temperature is then a linear function of the state, which makes every row of the example linear, the posterior exactly Gaussian and every result below exact rather than approximate, and it fails only where a Gaussian on the resistance puts real probability below zero, which Section 4.1.7 measures for the handbook prior and Section 4.6.6 discusses. The notation chapter carries the conductance in the state as the log-conductance $g = \ln(G/G_{\mathrm{ref}})$, which keeps it positive but makes the temperature nonlinear in the state; a later chapter on physical models uses it.

At the two operating conditions equation (4.4) gives

$$T_1 = 18 + 30\rho, \tag{4.5}$$

$$T_2 = 18 + 60\rho, \tag{4.6}$$

with T_1 the hall temperature during the test and T_2 the temperature at full load, both in °C with ρ in K/kW. Figure 4.2 draws the graph. Each row is one operating condition: the box on the left is the heat source, the green box in the middle the hall node carrying a free temperature, and the gray box on the right the ambient node, pinned at 18°C. The solid arrows are the physical paths, the heat entering the hall and the heat leaving it through the resistance. The green node between the rows is the resistance ρ, free and shared by both paths, as the dashed links show. The dashed arrows leave the graph for the evidence: the reading of the test temperature and the handbook prior on ρ, attached in Section 4.1.7. The full-load node carries no reading; it is the scenario node that the question of Section 4.1.6 requires, drawn here with the rest of the graph so that the figure shows the graph as extended for that question.

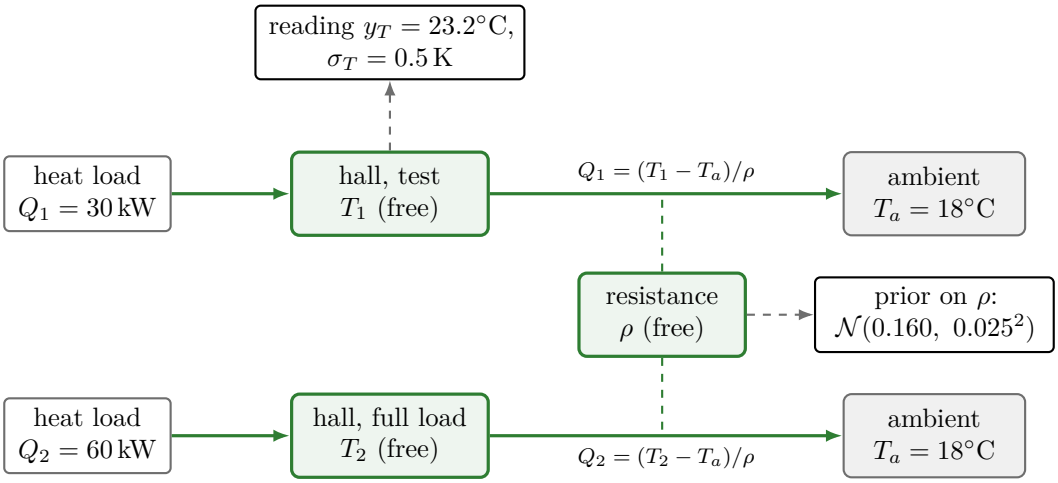


Figure 4.2: The plant graph of the equipment hall under its two operating conditions. Free: the resistance ρ and the temperatures T_1 and T_2 (green). Pinned: the ambient temperature T_a and the loads Q_1, Q_2 (gray and white). Observed: T_1, by one reading (dashed arrow, top); ρ carries the handbook prior (dashed arrow, right). The full-load node T_2 is the scenario node of the question and has no reading; it is linked to the test only through the shared resistance.

The resistance is taken to be the same at 60 kW as at 30 kW, one free entry shared by both rows: this is what lets a reading taken at the test condition carry information about the full-load condition, although the hall has never been operated at full load, and it fails if the path changes with load, for example when fans or chillers change state, which one test reading cannot detect (Section 4.6.6) and a reading at full load would test (Section 4.6.3). The link between the two rows of Figure 4.2 is not a device of the calculation; it is the physical statement that both conditions are the same hall with the same heat-removal path, and the prediction at full load rests on it.

4.1.5 The state vector and the physical rows

The free quantities of the graph form the state vector,

$$\mathbf{z} = \begin{bmatrix} \rho \\ T_1 \\ T_2 \end{bmatrix}, \tag{4.7}$$

with entries in K/kW, °C and °C. The physical relations are written as rows that vanish when they hold:

$$r_1(\mathbf{z}) = T_1 - 18 - 30\rho = 0, \tag{4.8}$$
$$r_2(\mathbf{z}) = T_2 - 18 - 60\rho = 0. \tag{4.9}$$

These are the conservation rows of the graph, one for each hall node. The balance is taken to hold at both conditions up to a small model tolerance σ_p, a standard deviation that represents model discrepancy, and the answers are reported in the limit $\sigma_p \to 0$ of exact balances: the finite tolerance keeps every matrix of the calculation invertible while the reported numbers are those of an exact balance, and the limit understates the spread if the real discrepancy is finite, which Section 4.4.5 measures over three tolerances and Section 4.6.6 discusses. Section 4.2.3 puts the rows, with their tolerance, into one residual vector with the reading and the prior.

The three entries of the state are therefore not three independent unknowns. The two physical rows leave one direction of the state free, a change of ρ with both temperatures following

it; in the limit of exact rows the state has three entries but one degree of freedom, and this fact governs the shape of the posterior covariance derived in Section 4.2.5.

4.1.6 The engineering questions (step 1)

The operator proposes to raise the heat load from 30 kW to 60 kW. The maximum acceptable hall temperature is $T_{\max} = 28°\mathrm{C}$. In ordinary language the question is: given the existing temperature reading and what is known of the cooling path, what is the probability that the hall stays below 28°C at full load? As a probability over the state, with the unknown on the left of the bar and everything treated as known on the right,

$$\Pr(T_2 < T_{\max} \mid \mathbf{y}, M), \tag{4.10}$$

where T_2 is the full-load entry of the state, $\mathbf{y}$ the evidence (the test reading and the handbook prior), and M the assumed steady-state thermal model, equation (4.4) with the rows (4.8)–(4.9). The model is suppressed from the notation below when it is not in question. This is a question of the category *prediction under a threshold*: a tail probability of one marginal of the posterior.

Three estimation questions are asked of the same evidence, one for each entry of the state: the thermal resistance of the cooling path, the hall temperature at the test condition, and the hall temperature at full load. Table 4.2 lists the four questions: the first column is the question in engineering language; the second the general form that every question of its category takes, written for a quantity $f(\mathbf{z})$ derived from the state and a threshold f^* on it; the third its mathematical form for the hall; the fourth its category; and the last the subsections where it is answered. Which further reading would narrow the answers most asks nothing new of the present evidence: it is the next measurement for these four questions, answered in Section 4.6.

Table 4.2: The questions asked of the hall's test reading, the general form of their category, their mathematical form for the hall, their category and where each is answered. $\mathbf{y}$ is the test reading and the handbook prior; $T_{\max} = 28°\mathrm{C}$.

Engineering question	General form	Mathematical form	Category	Answered in
What is the thermal resistance of the cooling path?	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$, moments	$p(\rho \mid \mathbf{y})$: $\hat{z}_1$, $\sqrt{\Sigma_{11}}$	estimation	4.2.9, 4.3.6
What is the hall temperature at the test condition?	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$, moments	$p(T_1 \mid \mathbf{y})$: $\hat{z}_2$, $\sqrt{\Sigma_{22}}$	estimation	4.3.6
What is the hall temperature at full load?	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$, moments, quantiles	$p(T_2 \mid \mathbf{y})$: $\hat{z}_3$, $\sqrt{\Sigma_{33}}$, credible interval	estimation	4.3.1, 4.3.4, 4.3.6
Will the hall stay below 28°C at 60 kW?	$\Pr(f(\mathbf{z}) < f^* \mid \mathbf{y})$	$\Pr(T_2 < T_{\max} \mid \mathbf{y})$	prediction under a threshold	4.3.5, 4.3.7

Because the full-load temperature is carried as an entry of the state, the threshold question needs exactly one marginal of the posterior: the mean $\hat{z}_3$ and the variance Σ_{33} of the third entry: the question is written over the state, and the position of T_2 in the state shows which part of the posterior it needs. Each estimation question decomposes in the same way, into one entry of $\hat{\mathbf{z}}$ and one diagonal entry of $\boldsymbol{\Sigma}$.

4.1.7 The evidence as rows (step 2)

The temperature reading. During the 30 kW test a sensor in the hall reports $y_T = 23.2°\mathrm{C}$, with a measurement error of standard deviation $\sigma_T = 0.5\,\mathrm{K}$. The observation model writes the

random reading Y as the true test temperature, the second entry of the state, plus an error,

$$Y = T_1 + \epsilon = \mathbf{h}^\mathsf{T}\mathbf{z} + \epsilon, \qquad \mathbf{h} = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \qquad \epsilon \sim \mathcal{N}(0, \sigma_T^2), \tag{4.11}$$

where $\mathbf{h}$ selects the entry the sensor measures. The error is taken as Gaussian, unbiased, of known spread and independent of the state: this makes the reading one linear row of fixed weight $1/\sigma_T^2$, and it fails if the sensor is biased, which would shift every inferred quantity; the prior predictive check of Section 4.4.2 would flag a large bias, and the comparison of models of Section 4.2.2 is the test built for it. The reading lands on T_1 only; it is the physical row of the test hall, not a substitution, that will carry it to ρ.

Prior information. An engineering handbook gives a preliminary value of the resistance, represented by the Gaussian prior

$$\rho \sim \mathcal{N}(\mu_0, \sigma_0^2), \qquad \mu_0 = 0.160\,\mathrm{K/kW}, \quad \sigma_0 = 0.025\,\mathrm{K/kW}. \tag{4.12}$$

The prior mean corresponds to a conductance $G_0 = 1/0.160 = 6.25\,\mathrm{kW/K}$. Taking the handbook value as the prior for this hall is a choice: it is what places the full-load temperature at 27.6°C before any reading and what keeps the estimate from following the one reading alone, and it misleads if the handbook describes paths unlike this one, a conflict that the prior predictive check of Section 4.4.2 would show as a reading far from 22.8°C. A Gaussian prior on a resistance assigns some probability to negative, nonphysical values; here it is $\Phi(-0.160/0.025) = \Phi(-6.4) \approx 8 \times 10^{-11}$, negligible, and the Gaussian is kept for its analytical simplicity. A production model enforces positivity through a logarithmic parameterization, such as the log-conductance g of the notation chapter, or a prior with positive support.

The rows and the entries they act on. The physics and the evidence put four rows on the state: the physical rows of the test hall and of the full-load hall, equations (4.8)–(4.9), the reading row on T_1, and the prior row on ρ. Which row acts on which entry is a property of the graph, and it fixes where the Jacobian of the rows can be nonzero, marked $*$:

$$\mathbf{J} = \begin{bmatrix} * & * & 0 \\ * & 0 & * \\ 0 & * & 0 \\ * & 0 & 0 \end{bmatrix} \begin{matrix} \text{physics, test hall} \\ \text{physics, full-load hall} \\ \text{reading} \\ \text{prior} \end{matrix}, \qquad \text{columns } \rho, T_1, T_2. \tag{4.13}$$

A line of the pattern lists the entries one row sees, and a column lists the rows that see one entry. The third column has a single mark: the full-load temperature is seen by its own physical row and by nothing else. The second column shows that the reading and the test hall's balance are the two rows that see T_1, and the first that the two balances and the prior see ρ. No row joins the reading to T_2; the two physical rows, through the column ρ they share, are the only path from the one to the other.

Knowns and unknowns. Table 4.3 lists every quantity of the problem. The first column is the symbol, the second its value, and the third its role in the graph: pinned boundary condition, evidence, or free entry of the state. The graph and the probabilistic model are now specified completely enough for the inference to be carried out.

Remark 4.1. Four spreads appear in what follows and must not be confused. The reading error has standard deviation σ_T (no index), a property of the sensor. The prior on the resistance has standard deviation σ_0. The physical rows have the model tolerance σ_p, a property of the steady-state model. The posterior standard deviation of the i-th entry of the state is $\sqrt{\Sigma_{ii}}$, the square root of a diagonal entry of the posterior covariance, written $\mathrm{sd}(T_2 \mid \mathbf{y})$ when the entry is named.

Table 4.3: Knowns and unknowns of the equipment-hall example.

Quantity	Value	Role
T_a	18°C	pinned: ambient boundary condition
Q_1	30 kW	pinned: test load
Q_2	60 kW	pinned: proposed load
y_T	23.2°C	evidence: the test reading
σ_T	0.5 K	evidence: standard deviation of the reading error
μ_0	0.160 K/kW	evidence: prior mean of the resistance
σ_0	0.025 K/kW	evidence: prior standard deviation of the resistance
σ_p	0.01 K, then $\to 0$	model: tolerance of the two physical rows
ρ	unknown	free: thermal resistance, the first entry of the state
T_1, T_2	unknown	free: test and full-load temperatures, tied to ρ by the physical rows

4.2 The posterior of the state (step 3)

Every question of Table 4.2 is a property of one density, the posterior of the state given the reading and the handbook prior. Each row of Section 4.1.7 contributes one factor to it, and the MAP state, the information matrix and the covariance are properties of that product. This section derives the posterior and computes it for the hall; Section 4.3 reads every answer from it.

4.2.1 The likelihood of the reading and the physical rows

The likelihood states how probable the observed reading would be under a proposed state. By equation (4.11), for a fixed state the reading is the test temperature T_1 plus a Gaussian error, so $Y \mid \mathbf{z} \sim \mathcal{N}(T_1,\ \sigma_T^2)$, and its density at the observed value is

$$p(y_T \mid \mathbf{z}) = \frac{1}{\sqrt{2\pi\sigma_T^2}} \exp\left[-\frac{(T_1 - y_T)^2}{2\sigma_T^2}\right]. \tag{4.14}$$

For the reading $y_T = 23.2$°C and $\sigma_T = 0.5\,\mathrm{K}$, so that $2\sigma_T^2 = 0.5\,\mathrm{K}^2$,

$$p(y_T \mid \mathbf{z}) \propto \exp\left[-\frac{(T_1 - 23.2)^2}{0.5}\right]. \tag{4.15}$$

The likelihood is a function of the state with the reading held fixed; it is not a density in $\mathbf{z}$ and need not integrate to one over it. It depends on the first temperature only and is flat in ρ and T_2: the sensor measures a temperature and says nothing directly about the resistance or about the full-load hall.

What connects the three entries is the pair of physical rows. A steady-state balance that holds up to a model tolerance σ_p (a standard deviation in kelvin, the amount by which the balance of equation (4.4) is allowed to fail) is the density

$$p(T_1 \mid \rho) = \mathcal{N}(T_1;\ 18 + 30\rho,\ \sigma_p^2), \qquad p(T_2 \mid \rho) = \mathcal{N}(T_2;\ 18 + 60\rho,\ \sigma_p^2), \tag{4.16}$$

where $\mathcal{N}(x; m, v)$ denotes the Gaussian density with mean m and variance v evaluated at x. The exact rows (4.8)–(4.9) are the limit $\sigma_p \to 0$, in which each density concentrates on its row. The derivation below keeps σ_p finite, carries all three entries, and takes the limit at the end, where it is checked numerically; no row is used to eliminate an entry.

4.2.2 The prior density and Bayes' theorem

By equation (4.12) the prior density of the resistance is

$$p(\rho) = \frac{1}{\sqrt{2\pi\sigma_0^2}} \exp\left[-\frac{(\rho-\mu_0)^2}{2\sigma_0^2}\right] \propto \exp\left[-\frac{(\rho-0.160)^2}{2(0.025)^2}\right]. \tag{4.17}$$

The prior of the state is the product of this density with the two physical densities (4.16), $p(\mathbf{z}) = p(T_1 \mid \rho)\, p(T_2 \mid \rho)\, p(\rho)$, the factorization of Chapter 3 read along the graph: each hall temperature depends on the resistance, and on nothing else. Bayes' theorem, derived in that chapter, gives the posterior of the state as

$$p(\mathbf{z} \mid y_T) = \frac{p(y_T \mid \mathbf{z})\, p(\mathbf{z})}{p(y_T)}, \qquad p(y_T) = \int p(y_T \mid \mathbf{u})\, p(\mathbf{u})\, \mathrm{d}\mathbf{u},$$

with $\mathbf{u}$ a variable of integration over the state. The denominator does not depend on $\mathbf{z}$, so

$$p(\mathbf{z} \mid y_T) \propto p(y_T \mid \mathbf{z})\, p(\mathbf{z}).$$

The evidence of the model. The denominator has a name of its own. Writing the model M explicitly, it is the evidence, or marginal likelihood, of the model,

$$p(y_T \mid M) = \int p(y_T \mid \mathbf{z}, M)\, p(\mathbf{z} \mid M)\, \mathrm{d}\mathbf{z}\,: \tag{4.18}$$

the density of the observed reading averaged over the prior of the state. For the linear-Gaussian model it is the prior predictive density of the reading evaluated at the observed value (Section 4.4.2). It also follows from the objective and its curvature, which is how it is computed when no closed form exists. The integrand is $\exp[-\mathcal{L}(\mathbf{z})]$, with the objective $\mathcal{L}$ of equation (4.24), times the normalizing constants $(2\pi\sigma_k^2)^{-1/2}$ of the four rows; when the objective is quadratic about its minimum $\hat{\mathbf{z}}$ with matrix $\mathbf{J}^{\mathsf{T}}\mathbf{J}$, whose inverse is the covariance $\boldsymbol{\Sigma}$ of Section 4.2.5, the Gaussian integral over the three entries gives

$$p(y_T \mid M) = \frac{(2\pi)^{3/2}\, \det(\boldsymbol{\Sigma})^{1/2}}{(2\pi)^2\, \sigma_p^2\, \sigma_T\, \sigma_0}\, \mathrm{e}^{-\mathcal{L}(\hat{\mathbf{z}})}. \tag{4.19}$$

The two directions of the state that the physical rows fix have posterior widths proportional to σ_p, and they cancel the two physical rows' normalizing constants: $\det(\boldsymbol{\Sigma})^{1/2}/\sigma_p^2$ tends to the width of the posterior along the free direction. The result is the volume of the posterior, $\det(\boldsymbol{\Sigma})^{1/2}$, times the height of the integrand at the MAP state. The formula is exact when the objective is quadratic, as it is here; in general it is the Laplace approximation to the evidence.

The evidence normalizes the posterior, and its second use is to compare models. Alternatives for the hall include M_1, a single thermal resistance; M_2, two distinct heat-removal paths; and M_3, a model with a biased temperature sensor. Each gives its own evidence for the same reading, and their ratio is a Bayes factor,

$$B_{21} = \frac{p(y_T \mid M_2)}{p(y_T \mid M_1)}, \tag{4.20}$$

which multiplies the prior odds of the two models to give their posterior odds, $\Pr(M_2 \mid y_T)/\Pr(M_1 \mid y_T) = B_{21} \cdot \Pr(M_2)/\Pr(M_1)$. That comparison is the category *hypothesis comparison*, and a later chapter on model selection develops it.

The posterior up to its constant. The data $\mathbf{y}$, the evidence in the sense of what is known about the hall, consists of the reading and the handbook prior; the likelihood involves only the reading, and the prior enters through $p(\mathbf{z})$, so $p(\mathbf{z} \mid \mathbf{y})$ and $p(\mathbf{z} \mid y_T)$ name the same density. Substituting equations (4.14), (4.16) and (4.17) and collecting the four exponents,

$$p(\mathbf{z} \mid \mathbf{y}) \propto \exp\left[-\frac{1}{2}\left(\frac{(T_1 - 18 - 30\rho)^2}{\sigma_p^2} + \frac{(T_2 - 18 - 60\rho)^2}{\sigma_p^2} + \frac{(T_1 - y_T)^2}{\sigma_T^2} + \frac{(\rho - \mu_0)^2}{\sigma_0^2}\right)\right]. \tag{4.21}$$

This is the posterior of the state up to its normalizing constant. It combines the physical relations, the reading, the reading's uncertainty and the prior knowledge in one probability model, one term per row. Every later result of the chapter is a property of this expression.

4.2.3 The residual vector, the Jacobian and the information matrix

The maximum a posteriori estimate is the state at which the posterior density is greatest,

$$\hat{\mathbf{z}} = \arg\max_{\mathbf{z}} \, p(\mathbf{z} \mid \mathbf{y}). \tag{4.22}$$

The logarithm is increasing, so maximizing the posterior is equivalent to minimizing its negative logarithm. Dropping the terms that do not depend on $\mathbf{z}$ (the normalizing constants of the four densities and $\ln p(y_T)$) leaves half the sum of the squared exponents of (4.21). Each exponent is a row divided by its standard deviation, and the four rows stack into one residual vector,

$$\mathbf{r}(\mathbf{z}) = \begin{bmatrix} (T_1 - 18 - 30\rho)/\sigma_p \\ (T_2 - 18 - 60\rho)/\sigma_p \\ (T_1 - y_T)/\sigma_T \\ (\rho - \mu_0)/\sigma_0 \end{bmatrix} \begin{matrix} \text{physics, test hall} \\ \text{physics, full-load hall} \\ \text{reading} \\ \text{prior} \end{matrix} \tag{4.23}$$

whose entries are dimensionless, so that the objective is

$$\mathcal{L}(\mathbf{z}) = \tfrac{1}{2}\,\mathbf{r}(\mathbf{z})^{\mathsf{T}}\mathbf{r}(\mathbf{z}), \qquad p(\mathbf{z} \mid \mathbf{y}) \propto \exp\bigl[-\mathcal{L}(\mathbf{z})\bigr]. \tag{4.24}$$

In the notation of the book it is $\mathcal{L} = \frac{1}{2}\sum_k w_k r_k^2$ over the unnormalized rows r_1, r_2 of equations (4.8)–(4.9), $T_1 - y_T$ and $\rho - \mu_0$, with weights $w_k = 1/\sigma_k^2$; the normalized rows carry their weights inside them. The first two terms penalize a state that unbalances a hall, in units of the model tolerance; the third, a state that disagrees with the reading, in units of the sensor spread; the fourth, a resistance far from its handbook value, in units of the prior spread. The weight of a physical row has a statistical meaning: $1/\sigma_p^2$ is the inverse variance of a model discrepancy, the amount by which the steady-state balance may fail.

The Jacobian. The derivatives of the four rows with respect to the three entries (ρ, T_1, T_2) form the 4×3 Jacobian $\mathbf{J} = \partial\mathbf{r}/\partial\mathbf{z}$, one line per row and one column per entry:

$$\mathbf{J} = \begin{bmatrix} -30/\sigma_p & 1/\sigma_p & 0 \\ -60/\sigma_p & 0 & 1/\sigma_p \\ 0 & 1/\sigma_T & 0 \\ 1/\sigma_0 & 0 & 0 \end{bmatrix} = \begin{bmatrix} -3000 & 100 & 0 \\ -6000 & 0 & 100 \\ 0 & 2 & 0 \\ 40 & 0 & 0 \end{bmatrix} \quad (\sigma_p = 0.01\,\mathrm{K}), \tag{4.25}$$

with the first column in kW/K and the other two in K^{-1}. The entry J_{ki} is the change in row k, counted in units of that row's standard deviation, per unit change of entry i. Its nonzero entries sit where the pattern (4.13) of the graph places them. The entries $2 = 1/0.5$ and $40 = 1/0.025$ are the normalizing factors of the reading and the prior. Every row is linear in the state, so $\mathbf{J}$ is the same at every linearization point. The tolerance $\sigma_p = 0.01\,\mathrm{K}$, one-fiftieth of the sensor's spread, is tight enough that the answers no longer move (the limit is checked in Section 4.2.9) and loose enough, within two decades of σ_T, to keep the normal equations conditioned well enough for double precision.

The information matrix. For rows normalized by their standard deviations the information matrix is $\mathbf{\Lambda} = \mathbf{J}^\mathsf{T}\mathbf{J}$, equation (3.25). Written with the unnormalized rows, whose Jacobian is $\tilde{\mathbf{J}}$, and the diagonal weight matrix $\mathbf{W} = \operatorname{diag}(1/\sigma_p^2,\ 1/\sigma_p^2,\ 1/\sigma_T^2,\ 1/\sigma_0^2)$, it is the same matrix in the form used throughout the book,

$$\mathbf{\Lambda} = \mathbf{J}^\mathsf{T}\mathbf{J} = \tilde{\mathbf{J}}^\mathsf{T}\mathbf{W}\tilde{\mathbf{J}} = \mathbf{J}_\mathrm{d}^\mathsf{T}\mathbf{W}_\mathrm{d}\mathbf{J}_\mathrm{d} + \mathbf{\Lambda}_\mathrm{prior}, \tag{4.26}$$

where the last form keeps the prior rows apart: $\mathbf{J}_\mathrm{d}$ and $\mathbf{W}_\mathrm{d}$ hold the physical and reading rows, and $\mathbf{\Lambda}_\mathrm{prior}$, here $\operatorname{diag}(1/\sigma_0^2, 0, 0)$, is the prior precision. For rows that are nonlinear in the state, equation (4.26) is the Gauss–Newton approximation to the Hessian of $\mathcal{L}$: it drops the curvature term $\sum_k r_k\, \partial^2 r_k/\partial\mathbf{z}\,\partial\mathbf{z}^\mathsf{T}$, the second derivatives of the rows weighted by their residuals, which vanishes for linear rows. Writing $\mathbf{j}_k$ for the k-th line of $\mathbf{J}$, the product is a sum of one outer product per row, so the information of independent rows adds, as equation (3.26) states for priors and readings; with the numbers of equation (4.25),

$$\mathbf{J}^\mathsf{T}\mathbf{J} = \sum_{k=1}^{4} \mathbf{j}_k^\mathsf{T}\mathbf{j}_k = \begin{bmatrix} 45\,001\,600 & -300\,000 & -600\,000 \\ -300\,000 & 10\,004 & 0 \\ -600\,000 & 0 & 10\,000 \end{bmatrix}. \tag{4.27}$$

The entry in the corner is $3000^2 + 6000^2 + 40^2$: the two physical rows and the prior. The reading contributes only the $4 = 1/\sigma_T^2$ of the (T_1, T_1) entry. The reading and prior rows alone would give $\operatorname{diag}(1600, 4, 0)$: no information at all would reach T_2, and none would pass from the reading to ρ. It is the off-diagonal entries of the physical rows that carry the reading at T_1 to ρ and from ρ to T_2. The matrix does not contain the reading's value y_T, only its standard deviation: the information is fixed by the design of the measurement and by the prior, not by the value observed. Its eigenvalues are 1.155, 1.0003×10^4 and 4.5012×10^7, all positive, with a condition number of 3.9×10^7 that grows as $1/\sigma_p^2$.

4.2.4 The normal equations and the MAP state

The stationary condition. By the chain rule the gradient of the objective (4.24) is

$$\nabla\mathcal{L}(\mathbf{z}) = \mathbf{J}^\mathsf{T}\mathbf{r}(\mathbf{z}), \tag{4.28}$$

and the MAP state is the point where it vanishes, $\mathbf{J}^\mathsf{T}\mathbf{r}(\hat{\mathbf{z}}) = \mathbf{0}$: three equations, one per entry.

The Gauss–Newton step. The step that sets the gradient of the linearized rows to zero solves the normal equations of equation (3.24),

$$(\mathbf{J}^\mathsf{T}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^\mathsf{T}\mathbf{r}(\mathbf{z}), \qquad \mathbf{z}^{(k+1)} = \mathbf{z}^{(k)} + \Delta\mathbf{z}, \tag{4.29}$$

the Gauss–Newton iteration. The rows of this example are linear, so the linearization is exact and one step lands on $\hat{\mathbf{z}}$ from any starting state.

The free direction. The two physical lines of $\mathbf{J}$ annihilate one direction of the state,

$$\mathbf{J}_\mathrm{phys}\,\mathbf{v} = \mathbf{0}, \qquad \mathbf{J}_\mathrm{phys} = \begin{bmatrix} -3000 & 100 & 0 \\ -6000 & 0 & 100 \end{bmatrix}, \quad \mathbf{v} = \begin{bmatrix} 1 \\ 30 \\ 60 \end{bmatrix}, \tag{4.30}$$

the free direction: a change of ρ with T_1 and T_2 following by $Q_1 = 30\,\mathrm{kW}$ and $Q_2 = 60\,\mathrm{kW}$ times that change, which unbalances neither hall. A pull of the reading on T_1 is answered by a move along $\mathbf{v}$, because any move off it would cost the physical rows far more than it gains on the reading; from a state that balances both halls, the step of the normal equations therefore lies, to within the tolerance, along $\mathbf{v}$.

4.2.5 The posterior covariance from the information matrix

The point estimate does not answer the reliability question, which depends on how uncertain the state remains.

The covariance. The rows are linear and Gaussian, so by equation (3.25) the posterior is exactly $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ with

$$\boldsymbol{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1},$$

the inverse of the matrix (4.27), with rows and columns ordered (ρ, T_1, T_2), so that its entries carry $(\mathrm{K/kW})^2$, K^2, and K·K/kW for the products. A larger information means a narrower posterior: the covariance is the inverse of the steepness of the objective.

The answers read off. The answers to the estimation questions are entries. The marginals of a Gaussian are read off its mean and covariance,

$$\rho \mid \mathbf{y} \sim \mathcal{N}(\hat{z}_1,\ \Sigma_{11}), \qquad T_1 \mid \mathbf{y} \sim \mathcal{N}(\hat{z}_2,\ \Sigma_{22}), \qquad T_2 \mid \mathbf{y} \sim \mathcal{N}(\hat{z}_3,\ \Sigma_{33}), \tag{4.31}$$

so the resistance is $\hat{\rho} = \hat{z}_1$ with standard deviation $\sqrt{\Sigma_{11}}$, the test temperature $\hat{z}_2$ with $\sqrt{\Sigma_{22}}$, and the full-load temperature, the entry the threshold question asks about, $\hat{T}_2 = \hat{z}_3$ with $\sqrt{\Sigma_{33}}$. How the entries move together is read off the correlations $\Sigma_{ij}/\sqrt{\Sigma_{ii}\Sigma_{jj}}$, which the physical rows drive towards one.

Coordinates along and across the physical rows. The shape of $\boldsymbol{\Sigma}$, and its limit as the tolerance shrinks, are read by measuring the state along the free direction and across the two physical rows. Write

$$\mathbf{z} = \mathbf{c} + \mathbf{B}\mathbf{u}, \qquad \mathbf{u} = \begin{bmatrix} \rho \\ d_1 \\ d_2 \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} 1 & 0 & 0 \\ 30 & 1 & 0 \\ 60 & 0 & 1 \end{bmatrix}, \quad \mathbf{c} = \begin{bmatrix} 0 \\ 18 \\ 18 \end{bmatrix}, \tag{4.32}$$

where $d_1 = T_1 - 18 - 30\rho$ and $d_2 = T_2 - 18 - 60\rho$ are the two hall imbalances in kelvin and the first column of $\mathbf{B}$ is $\mathbf{v}$. The rows are linear in $\mathbf{u}$ with Jacobian $\mathbf{JB}$, whose columns are $\mathbf{Jv}$, $\mathbf{J}\mathbf{e}_2$ and $\mathbf{J}\mathbf{e}_3$,

$$\mathbf{JB} = \begin{bmatrix} 0 & 1/\sigma_p & 0 \\ 0 & 0 & 1/\sigma_p \\ 30/\sigma_T & 1/\sigma_T & 0 \\ 1/\sigma_0 & 0 & 0 \end{bmatrix} = \begin{bmatrix} 0 & 100 & 0 \\ 0 & 0 & 100 \\ 60 & 2 & 0 \\ 40 & 0 & 0 \end{bmatrix}.$$

The first column, $\mathbf{Jv} = (0,\ 0,\ 60,\ 40)^{\mathsf{T}}$, is seen only by the reading (with slope $2 \times 30 = 60$) and the prior (slope 40); a move along $\mathbf{v}$ costs the information $\|\mathbf{Jv}\|^2 = 60^2 + 40^2 = 5200\,(\mathrm{kW/K})^2$. The other two columns carry the entries $1/\sigma_p$, so every direction that unbalances a hall costs about $1/\sigma_p^2 = 10^4\,\mathrm{K}^{-2}$: a tiny tolerance leaves the state free to move cheaply only along $\mathbf{v}$, a narrow ridge in the objective. The third column is orthogonal to the other two, so the information matrix in these coordinates is block diagonal,

$$(\mathbf{JB})^{\mathsf{T}}(\mathbf{JB}) = \begin{bmatrix} 1/\sigma_0^2 + 30^2/\sigma_T^2 & 30/\sigma_T^2 & 0 \\ 30/\sigma_T^2 & 1/\sigma_p^2 + 1/\sigma_T^2 & 0 \\ 0 & 0 & 1/\sigma_p^2 \end{bmatrix} = \begin{bmatrix} 5200 & 120 & 0 \\ 120 & 10\,004 & 0 \\ 0 & 0 & 10\,000 \end{bmatrix},$$

and its inverse, the covariance $\boldsymbol{\Sigma}_u$ of $\mathbf{u}$, is

$$\boldsymbol{\Sigma}_u = \begin{bmatrix} 1.9236 \times 10^{-4} & -2.3074 \times 10^{-6} & 0 \\ -2.3074 \times 10^{-6} & 9.9988 \times 10^{-5} & 0 \\ 0 & 0 & 1.0000 \times 10^{-4} \end{bmatrix} \quad (\sigma_p = 0.01\,\mathrm{K}).$$

Three results follow. First, the full-load imbalance d_2 is seen by its own row only: $(\boldsymbol{\Sigma}_u)_{33} = \sigma_p^2$, the variance that row allows, and it is uncorrelated with ρ and d_1. Second, by the block-inversion formula of Section 3.5.3, the first diagonal entry of $\boldsymbol{\Sigma}_u$ is the inverse of the Schur complement of the (d_1, d_1) entry in the upper 2×2 block,

$$\Sigma_{11} = (\boldsymbol{\Sigma}_u)_{11} = \left[\frac{1}{\sigma_0^2} + \frac{30^2}{\sigma_T^2} - \frac{(30/\sigma_T^2)^2}{1/\sigma_p^2 + 1/\sigma_T^2} \right]^{-1} = \left[\frac{1}{\sigma_0^2} + \frac{30^2}{\sigma_T^2 + \sigma_p^2} \right]^{-1}, \tag{4.33}$$

the first row of $\mathbf{B}$ being $\mathbf{e}_1^\mathsf{T}$; a tolerance acts on the reading as extra sensor variance and widens the spread of ρ. Third, the covariance of the state is $\boldsymbol{\Sigma} = \mathbf{B}\boldsymbol{\Sigma}_u\mathbf{B}^\mathsf{T}$; its third diagonal entry, with $T_2 = 18 + 60\rho + d_2$ and d_2 uncorrelated with ρ, is

$$\mathrm{Var}(T_2 \mid \mathbf{y}) = \Sigma_{33} = \mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g} + \sigma_p^2 = 60^2\,\Sigma_{11} + \sigma_p^2, \qquad \mathbf{g} = (60,\ 0,\ 0)^\mathsf{T}, \tag{4.34}$$

with $\mathbf{g}$ the gradient of the prediction $18 + 60\rho$. A model discrepancy therefore enters the answer twice: it widens Σ_{11}, and it adds its own variance at the node that has no reading.

The limit of exact rows. As $\sigma_p \to 0$, $(\boldsymbol{\Sigma}_u)_{22}$ and $(\boldsymbol{\Sigma}_u)_{33}$, both at most σ_p^2, vanish and $(\boldsymbol{\Sigma}_u)_{11}$ tends to $1/5200$ by equation (4.33), so $\boldsymbol{\Sigma}_u \to \mathrm{diag}(1/5200,\ 0,\ 0)$ and

$$\boldsymbol{\Sigma} = \mathbf{B}\boldsymbol{\Sigma}_u\mathbf{B}^\mathsf{T} \ \longrightarrow\ \frac{\mathbf{B}\mathbf{e}_1\mathbf{e}_1^\mathsf{T}\mathbf{B}^\mathsf{T}}{5200} = \frac{\mathbf{v}\mathbf{v}^\mathsf{T}}{5200} = \frac{1}{5200}\begin{bmatrix} 1 & 30 & 60 \\ 30 & 900 & 1800 \\ 60 & 1800 & 3600 \end{bmatrix}, \tag{4.35}$$

of rank one. The matrix has two zero eigenvalues and no ordinary inverse, and the posterior correlation of T_1 and T_2 is exactly one. This is not a failure of the inference; it is the consequence of exact physical rows, which leave the state one degree of freedom. The state-vector form needs no special treatment of it: the normal equations are solved at a finite tolerance, where $\mathbf{J}^\mathsf{T}\mathbf{J}$ is invertible, and the limit is derived in the coordinates of equation (4.32) and read from the sequence of tolerances. The distinction between exact constraints and uncertain physical relations matters whenever a Hessian and a covariance are constructed over the full state, and the tolerance is the device that keeps both defined.

The objective along the free direction. On the line $\mathbf{z} = \hat{\mathbf{z}} + t\,\mathbf{v}$ through the MAP state, in the limit of exact rows, the physical rows keep their values at $\hat{\mathbf{z}}$ (zero in the limit) and only the reading and prior rows change, as $r_3(\hat{\mathbf{z}}) + 60t$ and $r_4(\hat{\mathbf{z}}) + 40t$, the slopes being the entries of $\mathbf{J}\mathbf{v} = (0,\ 0,\ 60,\ 40)^\mathsf{T}$. The objective on the line is the reading term $\frac{1}{2}(r_3(\hat{\mathbf{z}}) + 60t)^2$ plus the prior term $\frac{1}{2}(r_4(\hat{\mathbf{z}}) + 40t)^2$. At the minimum the slopes of the two terms cancel, $\mathbf{v}^\mathsf{T}\mathbf{J}^\mathsf{T}\mathbf{r}(\hat{\mathbf{z}}) = 60\,r_3 + 40\,r_4 = 0$: the reading pulls with weight 60 and the prior with weight 40. The gradient vanishes at $\hat{\mathbf{z}}$, so the objective on the line is

$$\mathcal{L}(\hat{\mathbf{z}}{+}t\mathbf{v}) = \mathcal{L}(\hat{\mathbf{z}}){+}\tfrac{1}{2}\,t^2\,\|\mathbf{J}\mathbf{v}\|^2, \qquad \frac{\mathrm{d}^2}{\mathrm{d}t^2}\mathcal{L}(\hat{\mathbf{z}}{+}t\mathbf{v}) = \mathbf{v}^\mathsf{T}\mathbf{J}^\mathsf{T}\mathbf{J}\,\mathbf{v} = 60^2{+}40^2 = 5200\,(\mathrm{kW/K})^2, \tag{4.36}$$

a parabola whose curvature is the information along $\mathbf{v}$, equal to $1/\Sigma_{11}$ in the limit; the rise of $\frac{1}{2}$ is reached at $t = \pm 1/\sqrt{5200} = \pm\sqrt{\Sigma_{11}}$.

Remark 4.2. Substituting the exact rows $T_1 = 18 + 30\rho$ and $T_2 = 18 + 60\rho$ into the objective gives the same curve as a function of the resistance alone,

$$\mathcal{L}(\rho) = \frac{1}{2}\left[\frac{(\rho - \mu_0)^2}{\sigma_0^2} + \frac{(y_T - 18 - 30\rho)^2}{\sigma_T^2} \right], \tag{4.37}$$

whose second derivative is the 5200 of equation (4.36) and whose minimum is $\hat{\rho} = 880/5200$.

4.2.6 The Gaussian form of the posterior

In this example the prior is Gaussian, the reading error and the model tolerance are Gaussian, and every row is linear in the state. The negative log posterior is therefore exactly quadratic in $\mathbf{z}$, the curvature term $\sum_k r_k\,\partial^2 r_k/\partial\mathbf{z}\,\partial\mathbf{z}^\mathsf{T}$ is zero, the Gauss–Newton matrix $\mathbf{J}^\mathsf{T}\mathbf{J}$ is the exact Hessian, and its inverse is the exact posterior covariance under the stated model. Nothing in equations (4.29)–(4.31) is an approximation; the only limit taken, $\sigma_p \to 0$, is derived in equation (4.35).

For a nonlinear physical model, for instance one written in the log-conductance g, the negative log posterior is not quadratic. The same construction is then made locally: the objective is expanded to second order about the MAP state,

$$\mathcal{L}(\mathbf{z}) \approx \mathcal{L}(\hat{\mathbf{z}}) + \tfrac{1}{2}(\mathbf{z}-\hat{\mathbf{z}})^\mathsf{T}\mathbf{H}(\mathbf{z}-\hat{\mathbf{z}}), \qquad \mathbf{H} = \left.\frac{\partial^2\mathcal{L}}{\partial\mathbf{z}\,\partial\mathbf{z}^\mathsf{T}}\right|_{\hat{\mathbf{z}}} = \mathbf{J}^\mathsf{T}\mathbf{J} + \sum_k r_k \left.\frac{\partial^2 r_k}{\partial\mathbf{z}\,\partial\mathbf{z}^\mathsf{T}}\right|_{\hat{\mathbf{z}}},$$

with no first-order term because the gradient vanishes at the minimum, and the posterior is replaced by the Gaussian that this expansion defines, the Laplace approximation:

$$\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}\big(\hat{\mathbf{z}},\ \mathbf{H}^{-1}\big). \tag{4.38}$$

It holds when the posterior is close to Gaussian over its own width and when $\mathbf{H}$ is positive definite in the chosen free coordinates. The distinction matters throughout the book: the calculation of this chapter is exact, while a Laplace approximation in general is not, and its adequacy is one of the checks of step 5.

4.2.7 Rows that are not Gaussian

Every row of the hall is Gaussian, and its term in the objective is half its squared residual. Later chapters add rows of another kind: a count of failures in a number of demands or hours, a lifetime, a Beta or gamma prior. The construction above carries over to them unchanged once the objective is defined by the posterior rather than by a sum of squares.

The objective. The objective is the negative logarithm of the posterior up to a constant. Each row, a likelihood factor or a prior factor, contributes one term, its negative log-probability:

$$\mathcal{L}(\mathbf{z}) = \sum_k \mathcal{L}_k(\mathbf{z}), \qquad \mathcal{L}_k(\mathbf{z}) = -\ln p_k(\cdot \mid \mathbf{z}), \qquad p(\mathbf{z} \mid \mathbf{y}) \propto \exp\big[-\mathcal{L}(\mathbf{z})\big]. \tag{4.39}$$

A Gaussian row is the case $\mathcal{L}_k = \frac{1}{2}r_k^2$, and the objective (4.24) of the hall is this sum for its four rows.

What the fit uses. The MAP state, the Newton step and the Laplace covariance use the objective only through its gradient and its Hessian, and both are sums over the rows:

$$\nabla\mathcal{L} = \sum_k \nabla\mathcal{L}_k, \qquad \mathbf{H} = \sum_k \nabla^2\mathcal{L}_k, \qquad \mathbf{H}\,\Delta\mathbf{z} = -\nabla\mathcal{L}, \qquad \mathbf{\Sigma} = \mathbf{H}(\hat{\mathbf{z}})^{-1}. \tag{4.40}$$

The value of $\mathcal{L}$ enters only the evidence, equation (4.19). When the state has one entry ξ, the gradient is the slope $\mathcal{L}'(\xi)$, the Hessian the curvature $\mathcal{L}''(\xi)$, the step $\Delta\xi = -\mathcal{L}'/\mathcal{L}''$ and the variance $1/\mathcal{L}''(\hat{\xi})$; nothing else changes.

The Gaussian row. For $\mathcal{L}_k = \frac{1}{2}r_k^2$ with Jacobian line $\mathbf{j}_k = \partial r_k/\partial\mathbf{z}^\mathsf{T}$, the chain rule gives

$$\nabla\mathcal{L}_k = \mathbf{j}_k^\mathsf{T} r_k, \qquad \nabla^2\mathcal{L}_k = \mathbf{j}_k^\mathsf{T}\mathbf{j}_k + r_k\,\frac{\partial^2 r_k}{\partial\mathbf{z}\,\partial\mathbf{z}^\mathsf{T}}. \tag{4.41}$$

Summed over the rows these are $\mathbf{J}^\mathsf{T}\mathbf{r}$ and $\mathbf{J}^\mathsf{T}\mathbf{J}$ plus the curvature term of Section 4.2.6.

The count row. A record of k events whose expected number is $\mu(\mathbf{z})$, Poisson in hours or binomial in demands, has the term $\mathcal{L}_k = \mu - k \ln \mu$ in the Poisson case. Its gradient and Hessian are

$$\nabla \mathcal{L}_k = (\mu - k)\, \nabla \ln \mu, \qquad \nabla^2 \mathcal{L}_k = \mu\, \nabla \ln \mu\, \nabla \ln \mu^{\mathsf{T}} + (\mu - k)\, \nabla^2 \ln \mu. \tag{4.42}$$

The first term of the Hessian is the product of one line with itself, so it is $\mathbf{j}_k^{\mathsf{T}} \mathbf{j}_k$ for the line and the residual

$$\mathbf{j}_k = \sqrt{\mu}\, \nabla \ln \mu^{\mathsf{T}}, \qquad r_k = \frac{\mu - k}{\sqrt{\mu}}, \tag{4.43}$$

which also give $\mathbf{j}_k^{\mathsf{T}} r_k = \nabla \mathcal{L}_k$. This pair, the *working residual* and the *working Jacobian line*, is how a count row joins the Jacobian beside Gaussian rows: one more line of $\mathbf{J}$, one more entry of $\mathbf{r}$, and the same products (4.40). The remaining term $(\mu - k)\, \nabla^2 \ln \mu$ plays the part of $r_k\, \partial^2 r_k$ in a Gaussian row. It vanishes when $\ln \mu$ is linear in the state, which is the case when the state carries the log-rate, the canonical parameter of the Poisson distribution. A binomial count of k failures in n demands has the term $-k \ln p - (n - k) \ln(1 - p)$; carried on the logit ξ of the failure probability p, its canonical parameter, it has the slope $np - k$ and the curvature $np(1 - p)$, the variance of the count, with no remaining term, and its working pair is $r_k = (np - k)/\sqrt{np(1-p)}$, $\mathbf{j}_k = \sqrt{np(1-p)}\, \partial \xi / \partial \mathbf{z}^{\mathsf{T}}$.

The working pair reproduces the row's gradient and the Gauss–Newton part of its Hessian, and nothing more. Its half square is not the row's term, $\frac{1}{2} r_k^2 \neq \mathcal{L}_k$, and $\mathbf{j}_k$ is not the derivative of r_k except where $r_k = 0$; the fit never uses either, because it uses only the products (4.40). For a Gaussian row that is linear in the state, the working pair is the row's own residual and Jacobian line, which is why the two forms are the same in the hall.

The information and its curvature term. With every row in this form, the Hessian is

$$\mathbf{H} = \mathbf{J}^{\mathsf{T}} \mathbf{J} + \mathbf{C}, \qquad \mathbf{\Lambda} = \mathbf{J}^{\mathsf{T}} \mathbf{J}, \tag{4.44}$$

with $\mathbf{C}$ the sum of the rows' remaining terms. The matrix $\mathbf{\Lambda}$ is the information; $\mathbf{H}$ at the recorded data is the *observed information*, and $\mathbf{\Lambda}$, its average over the records the model could have produced, the *expected* or Fisher information. The two are equal when $\mathbf{C} = \mathbf{0}$: for linear Gaussian rows, and for count rows carried on their canonical parameter. Where $\mathbf{C}$ is not zero, the chapters name it, keep it in the Newton step and report the covariance from $\mathbf{H}$. Once $\hat{\mathbf{z}}$ and $\mathbf{\Sigma}$ are in hand, the kind of row no longer matters: every answer of step 4 is read from them in the same way.

4.2.8 The fit of the hall state

The fit is run at the tolerance $\sigma_p = 0.01\,\mathrm{K}$ of Table 4.3, and the limit of exact rows is read off equation (4.35).

The step from the handbook state. The iteration starts at the forward calculation with the handbook resistance, $\mathbf{z}^{(0)} = (0.160,\ 22.8,\ 27.6)^{\mathsf{T}}$, where both physical rows and the prior row vanish and the reading row is $(22.8 - 23.2)/0.5 = -0.8$, so that $\mathcal{L}(\mathbf{z}^{(0)}) = 0.32$ and $\mathbf{J}^{\mathsf{T}} \mathbf{r}(\mathbf{z}^{(0)}) = (0,\ -1.6,\ 0)^{\mathsf{T}}$: only the reading is out of balance, and it pulls on T_1 alone. The normal equations (4.29) are

$$\begin{bmatrix} 45\,001\,600 & -300\,000 & -600\,000 \\ -300\,000 & 10\,004 & 0 \\ -600\,000 & 0 & 10\,000 \end{bmatrix} \Delta \mathbf{z} = \begin{bmatrix} 0 \\ 1.6 \\ 0 \end{bmatrix}, \qquad \Delta \mathbf{z} = \begin{bmatrix} 0.0092296 \\ 0.27694 \\ 0.55378 \end{bmatrix}.$$

Divided by its first entry, the step is $\Delta \mathbf{z} = 0.0092296 \times (1,\ 30.005,\ 60.000)^{\mathsf{T}}$: it lies, to within the tolerance, along the free direction $\mathbf{v}$ of equation (4.30). Table 4.4 lists the iteration. Each

line is one iterate: its three entries, the objective, and the largest entry of the gradient (4.28). After the first step the gradient is at round-off and the second step moves no entry by more than 10^{-13}.

Table 4.4: Gauss–Newton iterations for the hall state from the handbook state (0.160, 22.8, 27.6), with $\sigma_p = 0.01\,\mathrm{K}$.

Iterate k	ρ [K/kW]	T_1 [°C]	T_2 [°C]	$\mathcal{L}$ [–]	$\max_i \lvert(\mathbf{J}^\mathsf{T}\mathbf{r})_i\rvert$ [kW/K or K^{-1}]
0	0.160000	22.8000	27.6000	0.32000	1.6
1	0.169230	23.0769	28.1538	0.09845	$< 10^{-8}$
2	0.169230	23.0769	28.1538	0.09845	$< 10^{-8}$

The MAP state. The fit returns

$$\hat{\mathbf{z}} = \begin{bmatrix} \hat{\rho} \\ \hat{T}_1 \\ \hat{T}_2 \end{bmatrix} = \begin{bmatrix} 0.16923 \\ 23.0769 \\ 28.1538 \end{bmatrix} : \tag{4.45}$$

a resistance of 0.16923 K/kW, a test temperature of 23.077°C and a full-load temperature of 28.154°C, each read straight off its entry, with the effective conductance $\hat{G} = 1/\hat{\rho} = 5.909\,\mathrm{kW/K}$ evaluated at the estimated resistance.

4.2.9 The covariance and the resistance

The covariance. Inverting the matrix (4.27),

$$\boldsymbol{\Sigma} = \left(\mathbf{J}^\mathsf{T}\mathbf{J}\right)^{-1} = \begin{bmatrix} 1.9236 \times 10^{-4} & 5.7685 \times 10^{-3} & 1.1542 \times 10^{-2} \\ 5.7685 \times 10^{-3} & 0.17309 & 0.34611 \\ 1.1542 \times 10^{-2} & 0.34611 & 0.69260 \end{bmatrix} \quad (\sigma_p = 0.01\,\mathrm{K}). \tag{4.46}$$

Its eigenvalues are 0.866 along (nearly) $\mathbf{v}$, and 1.0×10^{-4} and 2.2×10^{-8} in the two directions the physics fixes, both shrinking as σ_p^2.

The entries read off. The resistance has

$$\hat{\rho} = \hat{z}_1 = 0.16923\,\mathrm{K/kW}, \qquad \sqrt{\Sigma_{11}} = \sqrt{1.9236 \times 10^{-4}} = 0.01387\,\mathrm{K/kW},$$

the test temperature $\sqrt{\Sigma_{22}} = 0.416\,\mathrm{K}$, and the full-load temperature $\hat{T}_2 = \hat{z}_3 = 28.154$°C with $\sqrt{\Sigma_{33}} = \sqrt{0.69260} = 0.8322\,\mathrm{K}$. The correlations $\Sigma_{ij}/\sqrt{\Sigma_{ii}\Sigma_{jj}}$ are 0.99971 between ρ and T_1, 0.99993 between ρ and T_2, and 0.99964 between T_1 and T_2.

The limit of exact rows. The limit covariance (4.35) gives $\Sigma_{11} = 1/5200$, so that

$$\sqrt{\Sigma_{11}} = 0.013868\,\mathrm{K/kW}, \qquad \hat{\rho} = 880/5200 = 0.169231\,\mathrm{K/kW}, \qquad \hat{T}_2 = 18{+}60\hat{\rho} = 28.1538°\mathrm{C},$$

the resistance being the minimum of equation (4.37); the rest of the chapter carries these limit values.

The objective along the free direction. Figure 4.3 draws the objective (4.36) on the line $\mathbf{z} = \hat{\mathbf{z}} + t\,\mathbf{v}$ in the limit of exact rows, with t read on the horizontal axis as the resistance $\rho = \hat{z}_1 + t$ in K/kW and the dimensionless objective on the vertical axis: the solid curve is the objective, the dashed curve the prior term and the dotted curve the reading term; the dot marks the minimum and the double arrow joins the two points where the objective has risen by $\frac{1}{2}$. At the minimum $r_3(\hat{\mathbf{z}}) = -0.2462$ and $r_4(\hat{\mathbf{z}}) = 0.3692$, and the slopes cancel, $60 \times (-0.2462) + 40 \times 0.3692 = 0$.

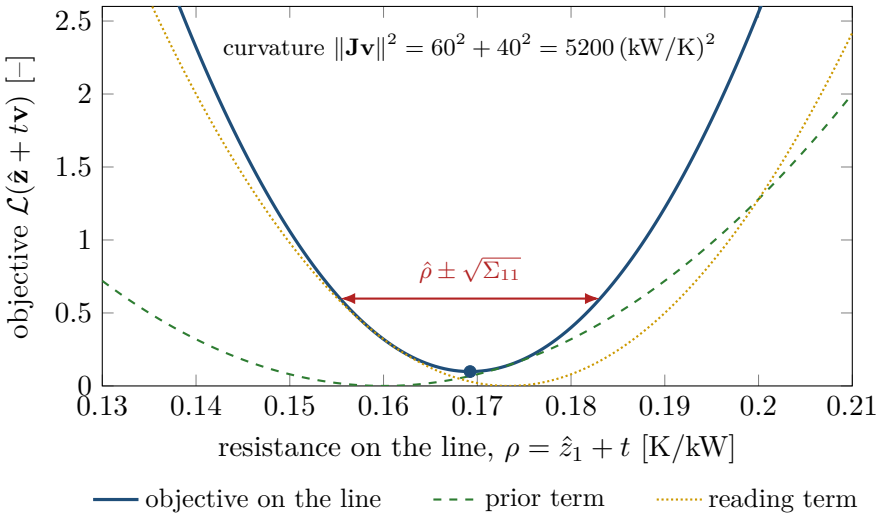


Figure 4.3: The objective on the line $\mathbf{z} = \hat{\mathbf{z}} + t\,\mathbf{v}$ along the free direction $\mathbf{v} = (1, 30, 60)^\mathsf{T}$, in the limit of exact physical rows, with its prior term $\frac{1}{2}(r_4(\hat{\mathbf{z}}) + 40t)^2$ and reading term $\frac{1}{2}(r_3(\hat{\mathbf{z}}) + 60t)^2$; the horizontal axis is the resistance $\hat{z}_1 + t$. Dot: the minimum at $\hat{\rho} = 0.16923\,\mathrm{K/kW}$, $\mathcal{L} = 0.0985$. Double arrow: the points $\hat{\rho} \pm \sqrt{\Sigma_{11}}$, $\sqrt{\Sigma_{11}} = 0.01387\,\mathrm{K/kW}$, where the objective has risen by $\frac{1}{2}$.

The posterior of the resistance. By the first marginal of equation (4.31),

$$\rho \mid \mathbf{y} \sim \mathcal{N}(\hat{z}_1,\ \Sigma_{11}) = \mathcal{N}(0.169231,\ 0.013868^2), \tag{4.47}$$

a standard deviation of $0.013868/0.169231 = 0.082$ of the estimate. Figure 4.4 compares it with the prior. The horizontal axis is the resistance and the vertical axis the probability density in kW/K; the dashed curve is the prior of equation (4.12) and the solid curve the posterior of equation (4.47). The reading has shifted the resistance upward, from 0.160 to 0.169, and narrowed it, from a standard deviation of 0.025 to 0.0139.

4.3 The answers (step 4)

Each question of Table 4.2 is read from the posterior of Section 4.2 with the tools derived here: the propagation of its covariance through the gradient of a derived quantity, the variance budget, the density of a transformed quantity, credible intervals and threshold probabilities. The hall's numbers follow; every quantity in them is a formula of this section or of Section 4.2 evaluated on the result of the fit.

4.3.1 Uncertainty propagation and the delta method

The rows of the hall are linear, so the entries of $\boldsymbol{\Sigma}$ are exact variances. In most engineering models an output is a nonlinear function of the state, $f(\mathbf{z})$, and the posterior is Gaussian at best approximately, $\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$.

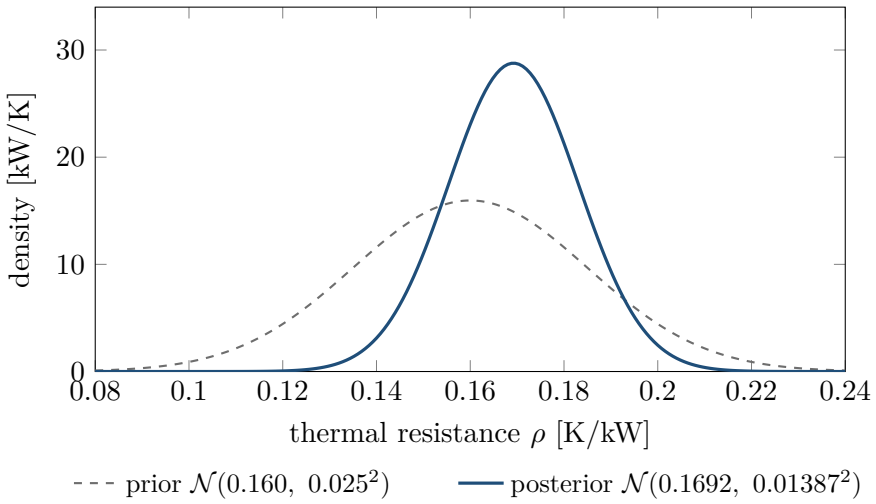


Figure 4.4: Prior and posterior densities of the thermal resistance, the first entry of the state: $\mathcal{N}(0.160, 0.025^2)$ and $\mathcal{N}(0.169231, 0.013868^2)$. The reading moves the resistance up and reduces its spread.

Linearization of the output. The first-order Taylor expansion of f about $\hat{\mathbf{z}}$ is

$$f(\mathbf{z}) \approx f(\hat{\mathbf{z}}) + \nabla f(\hat{\mathbf{z}})^{\mathsf{T}}(\mathbf{z} - \hat{\mathbf{z}}).$$

Writing $\mathbf{g} = \nabla f(\hat{\mathbf{z}})$ for the column gradient,

$$f(\mathbf{z}) \approx f(\hat{\mathbf{z}}) + \mathbf{g}^{\mathsf{T}}(\mathbf{z} - \hat{\mathbf{z}}). \tag{4.48}$$

Taking the posterior expectation, with $\mathbb{E}[\mathbf{z} - \hat{\mathbf{z}} \mid \mathbf{y}] = \mathbf{0}$ because $\hat{\mathbf{z}}$ is the mean of the Gaussian,

$$\mu_f = \mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}] \approx f(\hat{\mathbf{z}}).$$

Taking the variance, the constant $f(\hat{\mathbf{z}})$ drops out, and the variance of the linear combination $\mathbf{g}^{\mathsf{T}}(\mathbf{z} - \hat{\mathbf{z}}) = \sum_i g_i(z_i - \hat{z}_i)$ is the double sum of its covariances,

$$\operatorname{Var}(f(\mathbf{z}) \mid \mathbf{y}) \approx \operatorname{Var}\Big(\sum_i g_i z_i \;\Big|\; \mathbf{y}\Big) = \sum_i \sum_j g_i\, g_j\, \operatorname{Cov}(z_i, z_j \mid \mathbf{y}),$$

which in matrix form is

$$\operatorname{Var}(f(\mathbf{z}) \mid \mathbf{y}) \approx \mathbf{g}^{\mathsf{T}} \boldsymbol{\Sigma}\, \mathbf{g}. \tag{4.49}$$

This is the first-order delta method, one of the most reused formulas in probabilistic engineering. Its only approximation is the linearization (4.48); when f is linear and the posterior Gaussian, it is exact.

Application to the thermal system. The full-load temperature is the third entry of the state, so its gradient is the third unit vector and the quadratic form picks one diagonal entry,

$$\mathbf{g} = \frac{\partial T_2}{\partial \mathbf{z}} = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}, \qquad \mathbf{g}^{\mathsf{T}} \boldsymbol{\Sigma}\, \mathbf{g} = \Sigma_{33}.$$

The same temperature can also be predicted as a derived quantity of the resistance alone, $f(\mathbf{z}) = T_a + Q_2 \rho = 18 + 60\rho$, with

$$\mathbf{g} = \frac{\partial f}{\partial \mathbf{z}} = \begin{bmatrix} 60 \\ 0 \\ 0 \end{bmatrix}, \qquad \mathbf{g}^{\mathsf{T}} \boldsymbol{\Sigma}\, \mathbf{g} = 60^2\, \Sigma_{11}.$$

In the limit the two are equal, since the limit covariance (4.35) has $\Sigma_{33} = 3600/5200 = 60^2\,\Sigma_{11}$; at a finite tolerance they differ by σ_p^2, equation (4.34), because the entry carries the discrepancy of the full-load balance and the prediction from ρ does not. The two calculations, reading the variance off the covariance and propagating the resistance's uncertainty through a gradient, agree in the limit because they are the same transformation of the same uncertainty: an output that is an entry of the state needs no propagation, and an output that is not needs its gradient.

4.3.2 The gradient and the variance budget

Physical meaning of the gradient. The first entry of the gradient $\mathbf{g} = (60,\ 0,\ 0)^\mathsf{T}$ of the prediction $18 + 60\rho$, 60 kW, means that a change of 0.01 K/kW in the resistance changes the full-load temperature by $\Delta T_2 = 60 \times 0.01 = 0.6\,\mathrm{K}$. The gradient is the link between physical uncertainty and engineering consequence. A state entry can have a large posterior variance and still contribute little to an output that is insensitive to it, with $g_i \approx 0$; conversely, a modest uncertainty in a highly influential entry can dominate the uncertainty of a prediction.

Several uncertain entries. Suppose the ambient temperature is uncertain as well, so that the full-load temperature $T = T_a + Q\rho$ depends on the sub-vector $\boldsymbol{\theta} = (T_a, \rho)^\mathsf{T}$ of the state, with posterior covariance $\boldsymbol{\Sigma}_\theta$. The gradient is $\mathbf{g} = (\partial T/\partial T_a,\ \partial T/\partial \rho)^\mathsf{T} = (1,\ Q)^\mathsf{T}$, and equation (4.49) gives

$$\mathrm{Var}(T) = \mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}_\theta\,\mathbf{g} = \begin{bmatrix} 1 & Q \end{bmatrix} \begin{bmatrix} \mathrm{Var}(T_a) & \mathrm{Cov}(T_a,\rho) \\ \mathrm{Cov}(T_a,\rho) & \mathrm{Var}(\rho) \end{bmatrix} \begin{bmatrix} 1 \\ Q \end{bmatrix}.$$

Multiplying out,

$$\mathrm{Var}(T) = \mathrm{Var}(T_a) + Q^2\ \mathrm{Var}(\rho) + 2Q\ \mathrm{Cov}(T_a,\rho). \tag{4.50}$$

This is the variance budget of the output: one term per entry, each its sensitivity squared times its variance, plus a cross term from their covariance. The first term is the spread the ambient alone puts on the temperature; the second, the spread the resistance alone puts on it, $Q^2\ \mathrm{Var}(\rho)$; the third corrects for the two moving together in the posterior. If the readings cannot separate the two entries, the posterior leaves them correlated, and the cross term either adds to the diagonal terms or cancels part of them. With a negative correlation and same-signed sensitivities it is negative, and the output can be well known although each entry is poorly known: the readings then fix the combination $T_a + Q\rho$ that the output depends on, not the entries separately. Adding the two individual variances without the cross term is wrong in either direction, which is why the joint posterior, not its marginals, must be kept; Exercise 4.2 works a case in which ignoring the cross term overstates the spread by more than half.

4.3.3 The density of a transformed quantity

The delta method gives a derived quantity a mean and a variance. Some questions need more: the whole density of the quantity, or a tail probability far from the center, where the shape matters and a Gaussian with the right variance can still be wrong. The tool for these is the *change of variables*, and it is used throughout the book whenever a quantity is computed in one coordinate and reported in another.

Probability is conserved; density is not. Let x have the density $p_x(x)$, and let $u = h(x)$ be a monotone function with inverse $x = h^{-1}(u)$. A short interval $[x, x + \mathrm{d}x]$ is mapped onto a short interval $[u, u + \mathrm{d}u]$, and the two carry the same probability, because they hold the same outcomes:

$$p_u(u)\,|\mathrm{d}u| = p_x(x)\,|\mathrm{d}x|.$$

Dividing by $|\mathrm{d}u|$ gives the density of the transformed quantity,

$$p_u(u) = p_x\big(h^{-1}(u)\big) \left|\frac{\mathrm{d}x}{\mathrm{d}u}\right|. \tag{4.51}$$

The first factor evaluates the old density at the point that maps to u. The second, the *Jacobian factor*, records how much the map stretches or squeezes the axis there: where a small change of u corresponds to a large change of x, the probability of the long x-interval is packed into a short u-interval, and the density of u must be larger to hold it. Without the factor, $p_x(h^{-1}(u))$ is not a density of u at all: it does not in general integrate to one. For a state of several entries and an invertible map $\mathbf{u} = \mathbf{h}(\mathbf{z})$, the factor becomes the absolute determinant of the Jacobian matrix of the inverse map,

$$p_{\mathbf{u}}(\mathbf{u}) = p\big(\mathbf{h}^{-1}(\mathbf{u})\big) \left|\det \frac{\partial \mathbf{z}}{\partial \mathbf{u}}\right|, \tag{4.52}$$

the ratio by which the map changes a small volume. A monotone map also carries quantiles over unchanged: if h is increasing, $\Pr(u \le h(x^*)) = \Pr(x \le x^*)$, so the median of u is h of the median of x, while the mode, which depends on the Jacobian factor, moves.

The conductance. For the conductance $G = 1/\rho$ of the hall, the inverse map is $\rho = 1/G$ and $|\mathrm{d}\rho/\mathrm{d}G| = 1/G^2$, so equation (4.51) gives

$$p_G(G) = \frac{1}{G^2}\, p\big(\rho = 1/G \mid \mathbf{y}\big),$$

the solid curve of Figure 4.6 in Section 4.4.7. It is skewed to the right of $\hat{G}$ because the factor $1/G^2$ is larger at small G, which squeezes the left of the curve and stretches its right.

Two routes to a derived quantity. The delta method and the change of variables answer the same question with different approximations. For a scalar entry with posterior $x \mid \mathbf{y} \approx \mathcal{N}(\hat{x}, \Sigma)$ and a monotone output $u = h(x)$:

- the *delta method* linearizes h at $\hat{x}$, with gradient $g = \mathrm{d}u/\mathrm{d}x$ evaluated there, and reports $u \mid \mathbf{y} \approx \mathcal{N}\big(h(\hat{x}),\ g^2\Sigma\big)$, a Gaussian in u;
- the *change of variables* keeps h exact and reports $p_u(u) = p_x(h^{-1}(u))\,|\mathrm{d}x/\mathrm{d}u|$, with the Gaussian of x as its only approximation.

The two factors are reciprocals: the delta method's gradient $\mathrm{d}u/\mathrm{d}x$ is the inverse of the Jacobian factor $\mathrm{d}x/\mathrm{d}u$, but the delta method evaluates it once, at $\hat{x}$, while the change of variables evaluates it at every u. The delta method is the change of variables with the map replaced by its tangent, which is why the two agree on the spread near the center and differ in the tails. When h is linear the factor is constant and the two coincide. When h bends across the posterior, the delta method puts a symmetric Gaussian on a quantity whose density is skewed, and can put probability where the quantity cannot go: below zero for a conductance or a rate, outside $(0, 1)$ for a probability. Chapter 5 meets the second case in the valve, where the failure probability is computed through its logit and the delta method puts 2 % of the posterior on $p < 0$ (Section 5.3.1).

Principle 4.1. *The gradient carries a spread; the Jacobian factor carries a density. Report $\sqrt{\mathbf{g}^\top \boldsymbol{\Sigma} \mathbf{g}}$ when a standard deviation is the answer, and transform the density when the answer is a shape, a tail probability or a bound near the edge of what the quantity can be.*

4.3.4 Credible intervals

Engineers often need an interval for an uncertain quantity as well as its mean. For the Gaussian marginal $T_2 \mid \mathbf{y} \sim \mathcal{N}(\hat{z}_3, \Sigma_{33})$ of equation (4.31) the central 95 % credible interval is the interval that leaves 2.5 % of the posterior in each tail. Since $\Phi(1.960) = 0.975$,

$$\left[\hat{z}_3 - 1.960\,\sqrt{\Sigma_{33}},\ \hat{z}_3 + 1.960\,\sqrt{\Sigma_{33}}\right]. \tag{4.53}$$

The same form serves any entry, with $\hat{z}_i$ and $\sqrt{\Sigma_{ii}}$, and a derived quantity, with $f(\hat{\mathbf{z}})$ and $\sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}}$ from the delta method (4.49).

The interval and a threshold probability carry different information. The interval describes the range that holds most of the posterior probability; a threshold probability answers directly whether the plant meets a stated requirement. For an operating acceptance decision the threshold probability is usually the more actionable of the two. Neither is meaningful apart from the model assumptions, the quality of the measurement and the sources of uncertainty that the model includes; Section 4.6.6 lists those it leaves out.

4.3.5 The temperature-threshold probability

The question of equation (4.10) needs one marginal of the posterior. The marginal of the full-load temperature is read off the third entry of the state, equation (4.31): $T_2 \mid \mathbf{y} \sim \mathcal{N}(\hat{z}_3,\ \Sigma_{33})$.

The standardized temperature. For a Gaussian variable the standardized variable $(T_2 - \hat{z}_3)/\sqrt{\Sigma_{33}}$ has the standard normal distribution, and subtracting a constant and dividing by a positive constant preserve the inequality, so

$$\Pr(T_2 < T_{\max} \mid \mathbf{y}) = \Pr\!\left(\frac{T_2 - \hat{z}_3}{\sqrt{\Sigma_{33}}} < \frac{T_{\max} - \hat{z}_3}{\sqrt{\Sigma_{33}}} \;\middle|\; \mathbf{y}\right) = \Phi\!\left(\frac{T_{\max} - \hat{z}_3}{\sqrt{\Sigma_{33}}}\right). \tag{4.54}$$

The probability of exceeding the limit is its complement, $1 - \Pr(T_2 < T_{\max} \mid \mathbf{y})$.

The prior marginal. Before the test reading the state had the prior formed by the physical and prior rows alone. Its Jacobian $\mathbf{J}_0$ is $\mathbf{J}$ of equation (4.25) without the reading's line, its mean is the handbook state $\mathbf{z}^{(0)} = (0.160,\ 22.8,\ 27.6)^{\mathsf{T}}$, and its covariance $\boldsymbol{\Sigma}_0 = (\mathbf{J}_0^{\mathsf{T}}\mathbf{J}_0)^{-1}$ tends, by the argument of equation (4.35) with $\mathbf{J}_0\mathbf{v} = (0,\ 0,\ 40)^{\mathsf{T}}$ in place of $\mathbf{J}\mathbf{v}$, to $\mathbf{v}\mathbf{v}^{\mathsf{T}}/40^2 = \sigma_0^2\,\mathbf{v}\mathbf{v}^{\mathsf{T}}$. Its third diagonal entry is $(\boldsymbol{\Sigma}_0)_{33} = 60^2 \times 0.025^2 = 2.25\,\mathrm{K}^2$, so the full-load temperature had the prior marginal $\mathcal{N}(27.6,\ 1.5^2)$, the third entry of $\mathbf{z}^{(0)}$ with $\sqrt{(\boldsymbol{\Sigma}_0)_{33}} = 1.5\,\mathrm{K}$.

The question through the resistance. In the limit of exact rows the posterior lies on the line $\mathbf{z} = \hat{\mathbf{z}} + t\,\mathbf{v}$ through the MAP state along the free direction, $\boldsymbol{\Sigma} = \mathbf{v}\mathbf{v}^{\mathsf{T}}/5200$ by equation (4.35). On that line the third entry is $\hat{z}_3 + 60t$ and the first $\hat{z}_1 + t$, so the event $\{T_2 < T_{\max}\}$ is the event $\{t < t^*\}$ with $t^* = (T_{\max} - \hat{z}_3)/60$, which is the event that the resistance stays below

$$\rho^* = \hat{z}_1 + t^* = \hat{z}_1 + \frac{T_{\max} - \hat{z}_3}{60}. \tag{4.55}$$

With the marginal (4.31) of the first entry,

$$\Pr(\rho < \rho^* \mid \mathbf{y}) = \Phi\!\left(\frac{\rho^* - \hat{z}_1}{\sqrt{\Sigma_{11}}}\right) = \Phi\!\left(\frac{t^*}{\sqrt{\Sigma_{11}}}\right),$$

which is the probability (4.54): both standardized distances are t^* divided by the spread of t, since the third entry moves 60 times as far as the first along $\mathbf{v}$ and $\sqrt{\Sigma_{33}} = 60\sqrt{\Sigma_{11}}$ in the limit. Reading the third entry and propagating the first one along the free direction are the same calculation; the full-load question is a question about the resistance.

The dependence on the value read. The information matrix (4.27), and with it $\mathbf{\Sigma}$, does not contain the value observed, so a different reading leaves every standard deviation where it is and moves only the mean, along one vector. The reading enters the normal equations only through the part $-y_T/\sigma_T$ of the third entry of the residual vector that depends on y_T; that row's line of the Jacobian is $\mathbf{h}^{\mathsf{T}}/\sigma_T$, so differentiating $\hat{\mathbf{z}} = \mathbf{z} - \mathbf{\Sigma}\mathbf{J}^{\mathsf{T}}\mathbf{r}(\mathbf{z})$ with respect to y_T gives the gain vector

$$\frac{\partial \hat{\mathbf{z}}}{\partial y_T} = \frac{\mathbf{\Sigma}\mathbf{h}}{\sigma_T^2}, \tag{4.56}$$

the second column of $\mathbf{\Sigma}$ divided by σ_T^2. The threshold probability (4.54) is computed with a spread fixed before the reading and a mean set by its value, so the reading can move it in either direction.

4.3.6 The posterior state and the full-load temperature

The posterior state. Table 4.5 collects it: each row is one entry of $\hat{\mathbf{z}}$, with its posterior mean and its standard deviation $\sqrt{\Sigma_{ii}}$, in the limit $\sigma_p \to 0$; the full-load spread is twice the test spread, the ratio 60/30 of the entries of $\mathbf{v}$.

Table 4.5: The posterior state of the equipment hall after the test reading, in the limit of exact physical rows.

State entry	Posterior mean	Standard deviation
ρ	$0.16923\,\mathrm{K/kW}$	$0.01387\,\mathrm{K/kW}$
T_1	$23.077^{\circ}\mathrm{C}$	$0.416\,\mathrm{K}$
T_2	$28.154^{\circ}\mathrm{C}$	$0.832\,\mathrm{K}$

The diagonal of the limit covariance. The limit covariance (4.35) gives $\Sigma_{11} = 1/5200 = 1.9231 \times 10^{-4}\,(\mathrm{K/kW})^2$, $\Sigma_{22} = 900/5200 = 0.17308\,\mathrm{K}^2$ ($\sqrt{\Sigma_{22}} = 0.4160\,\mathrm{K}$) and

$$\Sigma_{33} = \mathrm{Var}(T_2 \mid \mathbf{y}) = \frac{3600}{5200} = 0.69231\,\mathrm{K}^2, \qquad \sqrt{\Sigma_{33}} = 0.8321\,\mathrm{K}. \tag{4.57}$$

The credible interval. By equation (4.53), $28.1538 \pm 1.960 \times 0.8321 = 28.1538 \pm 1.6308$, so that

$$T_2 \in [26.52,\ 29.78]^{\circ}\mathrm{C}$$

with 95 % posterior probability under the stated model; the prior interval was $[24.66,\ 30.54]^{\circ}\mathrm{C}$, 1.8 times as wide.

4.3.7 The threshold probability at full load

The probability. With $\hat{z}_3 = 28.1538^{\circ}\mathrm{C}$ and $\sqrt{\Sigma_{33}} = 0.8321\,\mathrm{K}$, the standardized distance in equation (4.54) is $(28 - 28.1538)/0.8321 = -0.1849$, and

$$\Pr(T_2 < 28^{\circ}\mathrm{C} \mid \mathbf{y}) = \Phi(-0.1849) = 0.4267. \tag{4.58}$$

The probability of exceeding the limit is $1 - 0.4267 = 0.5733$.

Against the prior. Under the prior marginal $\mathcal{N}(27.6,\ 1.5^2)$ of Section 4.3.5, before the test reading, the probability that the full-load temperature stays below the limit was

$$\Pr(T_2 < 28) = \Phi\left(\frac{28 - 27.6}{1.5}\right) = \Phi(0.2667) = 0.6051.$$

Figure 4.5 draws the answer. The horizontal axis is the full-load temperature and the vertical axis its density in K^{-1}. The solid curve is the posterior marginal, the dashed curve the prior marginal, and the vertical line the limit $T_{\max}$. The shaded area under the posterior to the left of the line is the probability (4.58); the unshaded area under it to the right is the probability of exceeding the limit. The posterior is narrower than the prior and centered to the right of it, across the limit.

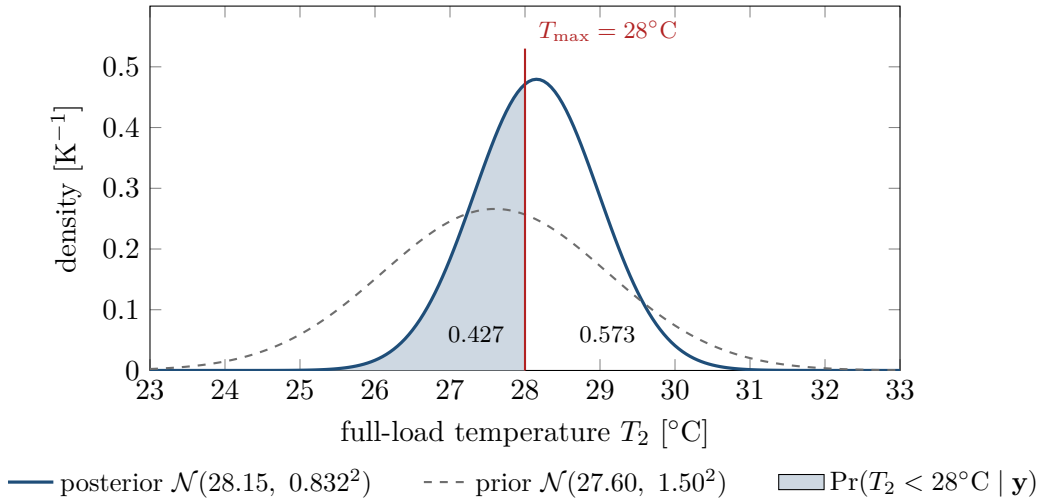


Figure 4.5: The threshold probability of equation (4.54). Posterior marginal of the full-load temperature, $\mathcal{N}(28.1538, 0.8321^2)$, with the area below $T_{\max} = 28°\mathrm{C}$ shaded ($\Phi(-0.1849) = 0.4267$), and the prior marginal $\mathcal{N}(27.6, 1.5^2)$, under which the same probability is $\Phi(0.2667) = 0.6051$.

4.4 The checks (step 5)

Obtaining a posterior is not the same as establishing that it is credible.

4.4.1 Algebraic and physical verification

Algebraic verification. At the fitted state of equation (4.45), with $\sigma_p = 0.01\,\mathrm{K}$, the residual vector (4.23) is $\mathbf{r}(\hat{\mathbf{z}}) = (0.0049225,\ 0,\ -0.24612,\ 0.36919)^{\mathsf{T}}$, and the gradient (4.28), column by column of $\mathbf{J}$,

$$\mathbf{J}^{\mathsf{T}}\mathbf{r}(\hat{\mathbf{z}}) = \begin{bmatrix} -3000 \times 0.0049225 + 40 \times 0.36919 \\ 100 \times 0.0049225 + 2 \times (-0.24612) \\ 100 \times 0 \end{bmatrix} = \begin{bmatrix} -14.767 + 14.767 \\ 0.4922 - 0.4922 \\ 0 \end{bmatrix} = \mathbf{0},$$

to numerical precision: the state satisfies the first-order optimality condition. Each line reads as a balance of rows on one entry: the prior's pull on ρ is held by the first physical row, which in turn is held at T_1 by the reading row; nothing pulls on T_2 but its own physical row. The information matrix (4.27) has positive eigenvalues, so the objective is strictly convex and the stationary point is its unique minimum.

Physical consistency. The physical rows at the fitted state, in kelvin, are $\sigma_p r_1 = 4.9 \times 10^{-5}\,\mathrm{K}$ and $\sigma_p r_2 = 0$: both halls balance to well within the tolerance, and the first imbalance shrinks with σ_p. The difference between the observed and the fitted test temperature is $23.2 - 23.0769 = 0.1231\,\mathrm{K}$, a quarter of the assumed sensor standard deviation of $0.5\,\mathrm{K}$. The model is consistent with this reading. Agreement with one reading does not validate the physical model: with one reading and one free direction, the fit has nearly enough freedom to reproduce any value.

4.4.2 Prior predictive check

Before the reading was used, the model and the prior predicted a distribution for it. By equation (4.11), $Y = \mathbf{h}^\mathsf{T}\mathbf{z} + \epsilon$, with the state under its prior $\mathcal{N}(\mathbf{z}^{(0)}, \mathbf{\Sigma}_0)$ of Section 4.3.5 and ϵ independent of it, so Y is Gaussian with

$$\mathbb{E}[Y] = \mathbf{h}^\mathsf{T}\mathbf{z}^{(0)} = 22.8^\circ\mathrm{C}, \qquad \mathrm{Var}(Y) = \mathbf{h}^\mathsf{T}\mathbf{\Sigma}_0\mathbf{h} + \sigma_T^2 = (\mathbf{\Sigma}_0)_{22} + \sigma_T^2 = 0.5625 + 0.25 = 0.8125\,\mathrm{K}^2,$$

the variances adding because the two terms are independent, and $(\mathbf{\Sigma}_0)_{22} = 30^2 \times 0.025^2$ in the limit of exact rows. Thus

$$Y \sim \mathcal{N}(22.8,\ 0.9014^2). \tag{4.59}$$

The observed reading lies $(23.2 - 22.8)/0.9014 = 0.444$ standard deviations from the predicted mean, an unremarkable distance: the reading is compatible with the prior information and the model. This check is distinct from computing the posterior. It asks whether the evidence was plausible under the assumptions before those assumptions were used to interpret it; a reading several predictive standard deviations away would have signaled a wrong prior, a wrong model or a faulty sensor before any posterior was formed.

4.4.3 The evidence by two routes

The evidence (4.18) of the model is the prior predictive density (4.59) evaluated at the reading,

$$p(y_T \mid M) = \mathcal{N}(y_T;\ \mathbf{h}^\mathsf{T}\mathbf{z}^{(0)},\ \mathbf{h}^\mathsf{T}\mathbf{\Sigma}_0\mathbf{h} + \sigma_T^2),$$

with $\mathcal{N}(y; m, v)$ the Gaussian density of Section 4.2.1. With $m = 22.8^\circ\mathrm{C}$ and $v = 0.8125\,\mathrm{K}^2$,

$$p(y_T = 23.2 \mid M) = \frac{1}{\sqrt{2\pi \times 0.8125}} \exp\left[-\frac{(23.2 - 22.8)^2}{2 \times 0.8125}\right] = 0.44259 \times 0.90623 = 0.4011\,\mathrm{K}^{-1}. \tag{4.60}$$

The Laplace form (4.19) gives the same number from the fit. At $\sigma_p = 0.01\,\mathrm{K}$, $\det(\mathbf{\Sigma})^{1/2} = 1.3867 \times 10^{-6}$ and $\mathcal{L}(\hat{\mathbf{z}}) = 0.09845$ (Table 4.4), so

$$p(y_T \mid M) = 0.39894 \times \frac{\det(\mathbf{\Sigma})^{1/2}}{\sigma_p^2} \times \frac{\mathrm{e}^{-\mathcal{L}(\hat{\mathbf{z}})}}{\sigma_T\,\sigma_0} = 0.39894 \times 0.013867 \times 0.90624 \times 80 = 0.4011\,\mathrm{K}^{-1},$$

with $\det(\mathbf{\Sigma})^{1/2}/\sigma_p^2 = 0.013867$ against its limit $1/\sqrt{5200} = 0.013868$, the width of the posterior along the free direction.

Common mistake

Reading an evidence value as a probability that the model is correct. The density $0.4011\,\mathrm{K}^{-1}$ is not a 40 % confidence in the thermal model: it is a density, in inverse units of the reading, whose value changes by a factor of 1000 if the reading is expressed in millikelvin. A posterior probability of a model requires the evidence of the competing models and their prior probabilities. Likewise, the posterior probability 0.4267 is a confidence in the claim about T_2 under the model, while the evidence ratio is a confidence in the model the claim was computed under; the two answer different questions.

4.4.4 Posterior predictive validation with a new observation

Suppose a further, independent reading is to be taken at the same 30 kW condition. Before its value is seen, its distribution given the evidence already in hand is the posterior predictive. The new reading is $Y_{\text{new}} = \mathbf{h}^{\mathsf{T}}\mathbf{z} + \epsilon_{\text{new}}$, with $\mathbf{z} \mid \mathbf{y} \sim \mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ and $\epsilon_{\text{new}} \sim \mathcal{N}(0, \sigma_T^2)$ independent of it, so

$$Y_{\text{new}} \mid \mathbf{y} \sim \mathcal{N}(\mathbf{h}^{\mathsf{T}}\hat{\mathbf{z}},\ \mathbf{h}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{h} + \sigma_T^2) = \mathcal{N}(23.0769,\ 0.17308 + 0.25) = \mathcal{N}(23.0769,\ 0.6504^2), \quad (4.61)$$

with $\mathbf{h}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{h} = \Sigma_{22}$ from equation (4.35). The predictive variance has two parts: the remaining uncertainty about the test temperature, $0.173\,\text{K}^2$, and the error of the new reading itself, $0.25\,\text{K}^2$. A further reading of the same sensor can therefore never be predicted more precisely than σ_T, however well the resistance becomes known.

An independent new reading compared with this distribution assesses the predictive performance of the model: a reading a fraction of a predictive standard deviation from 23.0769°C passes, and one several standard deviations away does not. Repeated large discrepancies, or a systematic sign in them, would indicate that the model, the prior, the measurement-error assumptions or the representation of the operating condition need revision.

4.4.5 The limit of exact rows

Table 4.6 repeats the fit at three tolerances and in the limit $\sigma_p \to 0$. Each column is one tolerance; the rows give the MAP resistance and full-load temperature, their standard deviations, the correlation of the two hall temperatures, the threshold probability of Section 4.3.5, and the condition number of the information matrix. By $\sigma_p = 0.01\,\text{K}$ the resistance, its spread, the full-load temperature and the threshold probability have stopped moving at the printed precision; the full-load spread still moves in its fourth digit, 0.8322 against 0.8321, and at $\sigma_p = 0.001\,\text{K}$ it has stopped too. The fit at $\sigma_p = 0.001\,\text{K}$ reproduces the limit values of Section 4.2.9 to every printed digit.

Table 4.6: The fit of the hall state as the model tolerance of the physical rows shrinks. The last column is the limit of exact rows.

Quantity	$\sigma_p = 0.1\,\text{K}$	$0.01\,\text{K}$	$0.001\,\text{K}$	$\sigma_p \to 0$
$\hat{\rho}$ [K/kW]	0.169119	0.169230	0.169231	0.169231
$\sqrt{\Sigma_{11}}$ [K/kW]	0.014056	0.013869	0.013868	0.013868
$\hat{T}_2$ [°C]	28.1471	28.1538	28.1538	28.1538
$\sqrt{\Sigma_{33}}$ [K]	0.84926	0.83223	0.83205	0.83205
correlation of T_1, T_2 [–]	0.96522	0.99964	0.999996	1
$\Pr(T_2 < 28°\text{C})$ [–]	0.4312	0.4267	0.4267	0.4267
condition number of $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ [–]	4.0×10^5	3.9×10^7	3.9×10^9	–

The $\sqrt{\Sigma_{11}}$ row widens with the tolerance as equation (4.33) requires, the tolerance acting on the reading as extra sensor variance, and the limit column is the matrix of equation (4.35). The fitted covariances satisfy equation (4.34) at every tolerance of the table ($0.69260 = 0.69250 + 0.0001$ at $\sigma_p = 0.01\,\text{K}$, $0.72125 = 0.71125 + 0.01$ at $0.1\,\text{K}$). A tighter tolerance than $0.001\,\text{K}$ buys nothing and costs conditioning: the condition number grows a hundredfold per decade, and at $\sigma_p = 10^{-6}\,\text{K}$ it reaches 4×10^{15} and round-off already moves the fifth digit of $\sqrt{\Sigma_{33}}$.

4.4.6 Second routes to the full-load answers

The propagated full-load temperature. As the third entry, with $\mathbf{g} = (0,\ 0,\ 1)^{\mathsf{T}}$, the full-load temperature has $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} = \Sigma_{33} = 0.69231\,\text{K}^2$; as the prediction $18 + 60\rho$, with $\mathbf{g} = (60,\ 0,\ 0)^{\mathsf{T}}$,

it has $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} = 60^2\,\Sigma_{11} = 3600/5200 = 0.69231\,\mathrm{K}^2$, the same number, as Section 4.3.1 requires in the limit.

The threshold through the resistance. By equation (4.55), $t^* = (28 - 28.1538)/60 = -0.0025641\,\mathrm{K/kW}$ and

$$\rho^* = \hat{z}_1 + t^* = 0.169231 - 0.002564 = 0.166667\,\mathrm{K/kW},$$
$$\Pr(\rho < \rho^* \mid \mathbf{y}) = \Phi\Big(\frac{-0.002564}{0.013868}\Big) = \Phi(-0.1849) = 0.4267,$$

the number of equation (4.58).

4.4.7 The propagation check against draws

For a linear output the delta method is exact, and the check is the algebraic agreement of the entry and the gradient already shown in Section 4.4.6. For a nonlinear output the linearized result must be compared with the full nonlinear evaluation at draws from the posterior. The procedure is the following.

1. Factor the covariance as $\boldsymbol{\Sigma} = \mathbf{L}\mathbf{L}^{\mathsf{T}}$ with $\mathbf{L}$ lower triangular (the Cholesky factor, which exists at any finite tolerance), and draw N independent states $\mathbf{z}^{(j)} = \hat{\mathbf{z}} + \mathbf{L}\boldsymbol{\varepsilon}^{(j)}$, $j = 1, \ldots, N$, each $\boldsymbol{\varepsilon}^{(j)}$ a vector of three independent standard normal numbers, so that $\mathbf{z}^{(j)} \sim \mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$. At $\sigma_p = 0.01\,\mathrm{K}$,
$$\mathbf{L} = \begin{bmatrix} 0.013869 & 0 & 0 \\ 0.41592 & 0.0099980 & 0 \\ 0.83217 & 0 & 0.010000 \end{bmatrix},$$
whose first column is $\sqrt{\Sigma_{11}}$ times (nearly) $\mathbf{v}$ and whose other two columns have the width σ_p of the physical rows. (In the limit of exact rows the factor has the single column $\mathbf{v}/\sqrt{5200}$, and a draw is $\hat{\mathbf{z}} + \mathbf{v}\,\varepsilon^{(j)}/\sqrt{5200}$, along the free direction.)
2. Evaluate the output at each drawn state through the full model, for example $T_2^{(j)} = z_3^{(j)}$ or $G^{(j)} = 1/z_1^{(j)}$.
3. Estimate a threshold probability as the fraction of draws satisfying the event, $\hat{P} = N^{-1}\sum_j \mathbf{1}[f(\mathbf{z}^{(j)}) < f^*]$, whose standard error is $\sqrt{\hat{P}(1-\hat{P})/N}$; estimate quantiles as the empirical quantiles of the $f(\mathbf{z}^{(j)})$.
4. Compare with the closed forms of the delta method. Agreement within the standard error means the closed form can be trusted.

A few hundred draws resolve a probability near one half to about ± 0.03; $N = 10^5$ resolves it to about ± 0.002.

Table 4.7 applies the procedure to the full-load temperature, which is linear in ρ, and to the conductance $G = 1/\rho$, which is not. The first column names the quantity; the second gives the delta-method value, from a Gaussian with mean $f(\hat{\mathbf{z}})$ and standard deviation $\sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}}$; the third gives the exact value of the push-forward of the posterior, the limit to which the draw estimates converge, computed here in closed form because both outputs are monotone in ρ (the q-quantile of G is $1/(\hat{\rho} - z_q\sqrt{\Sigma_{11}})$ with z_q the standard normal quantile); the fourth gives one run of $N = 10^5$ drawn states at $\sigma_p = 0.01\,\mathrm{K}$, generated with a seeded pseudo-random generator so that the run can be repeated exactly. For the conductance $G(\mathbf{z}) = 1/z_1$ the delta method uses

$$\mathbf{g} = \frac{\partial G}{\partial \mathbf{z}} = \begin{bmatrix} -1/\hat{z}_1^2 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} -34.92 \\ 0 \\ 0 \end{bmatrix} (\mathrm{kW/K})^2, \qquad \mathrm{Var}(G \mid \mathbf{y}) \approx \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} = \frac{\Sigma_{11}}{\hat{z}_1^4},$$

so that $\mathrm{sd}(G \mid \mathbf{y}) \approx 0.484\,\mathrm{kW/K}$.

Table 4.7: The propagation check: delta method against the exact push-forward of the posterior, whose first entry is $\mathcal{N}(0.169231, 0.013868^2)$, and one run of 10^5 states drawn as $\hat{\mathbf{z}} + \mathbf{L}\boldsymbol{\varepsilon}$ (standard error of a probability about 0.0015). The threshold $G = 6.25\,\mathrm{kW/K}$ is the handbook conductance.

Quantity	Delta method	Exact	10^5 draws
T_2, median [°C]	28.154	28.154	28.152
T_2, 5 % and 95 % [°C]	26.785, 29.522	26.785, 29.522	26.786, 29.516
$\Pr(T_2 < 28°\mathrm{C})$ [–]	0.4267	0.4267	0.4269
G, median [kW/K]	5.909	5.909	5.910
G, 5 % and 95 % [kW/K]	5.113, 6.706	5.207, 6.830	5.210, 6.828
$\Pr(G > 6.25\,\mathrm{kW/K})$ [–]	0.2407	0.2528	0.2527

For the temperature the three columns agree to within the sampling error, and the closed form needs no correction. For the conductance the median is right in all three columns, since $1/\rho$ is monotone and maps the median of ρ to the median of G, but the delta-method band is shifted: the exact band $[5.207, 6.830]$ is skewed to the right of $\hat{G}$, while the linearized band $[5.113, 6.706]$ is symmetric about it, and the delta method understates the probability that the conductance exceeds its handbook value by about 0.012. The draws reproduce the exact column to within their standard error. Figure 4.6 shows the cause. The horizontal axis is the conductance and the vertical axis its density; the solid curve is the exact density of $G = 1/\rho$, obtained by the change of variables $p_G(G) = p(\rho = 1/G \mid \mathbf{y})/G^2$ of Section 4.3.3, and the dashed curve the delta-method Gaussian. The exact density has its mode to the left of 5.909 and a longer right tail, and the Gaussian cannot represent the asymmetry.

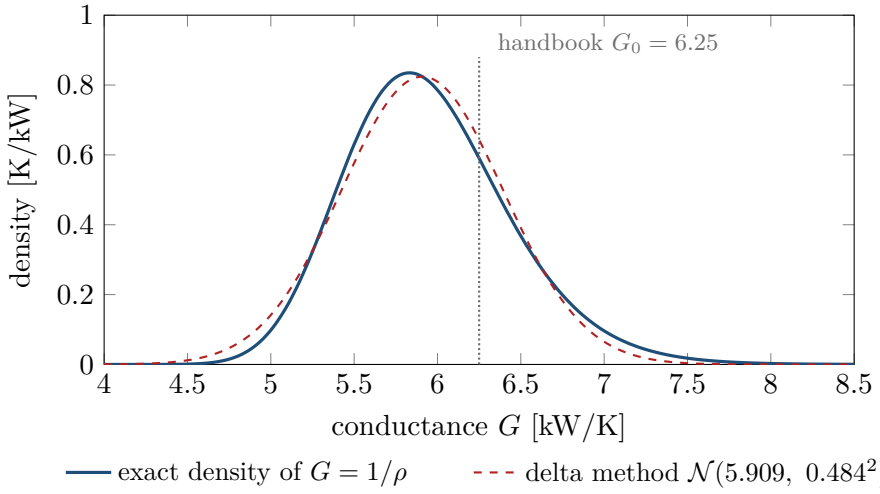


Figure 4.6: Posterior density of the conductance $G = 1/\rho$ under $\rho \mid \mathbf{y} \sim \mathcal{N}(0.169231, 0.013868^2)$: exact, by change of variables, $p_G(G) = \phi((1/G - 0.169231)/0.013868)/(0.013868\,G^2)$ with ϕ the standard normal density; and the delta-method Gaussian $\mathcal{N}(1/\hat{\rho}, \Sigma_{11}/\hat{\rho}^4)$. Dotted line: the handbook conductance.

The conductance also shows why "the posterior mean of G" must be handled with care. Under a Gaussian posterior for ρ, the mean of $1/\rho$ does not exist in the strict sense, because the Gaussian has positive density at $\rho = 0$, where $1/\rho$ is not integrable; the mass involved here is about 10^{-34} and has no practical effect, but it is a reminder that moments of a transformed quantity are properties of the parameterization. Quantiles and threshold probabilities, which the push-forward preserves exactly for a monotone output, are the robust summaries, and they are what the check compares.

4.4.8 Sensitivity to the value read

Table 4.8 repeats the assessment for other values of the single test reading. The first numeric column is the prior; each further column is the posterior after a different value of the reading, the actual one, 23.2°C, among them. The rows give the MAP resistance from equation (4.29), the mean and standard deviation of the full-load temperature, and the probability (4.54). The standard deviation is the same for every reading, and the mean moves along the gain vector (4.56), which with the limit covariance (4.35) is

$$\frac{\partial \hat{\mathbf{z}}}{\partial y_T} = \frac{\boldsymbol{\Sigma}\mathbf{h}}{\sigma_T^2} = 4 \begin{bmatrix} 0.0057692 \\ 0.17308 \\ 0.34615 \end{bmatrix} = \begin{bmatrix} 0.023077 \\ 0.69231 \\ 1.38462 \end{bmatrix},$$

four times the second column of the limit covariance, in K/kW per K and K per K. One column of the table, for $y_T = 22.4$°C, is the actual state moved by -0.8 K along it: $\hat{\rho} = 0.169231 - 0.8 \times 0.023077 = 0.150769$, $\mathbb{E}[T_2 \mid \mathbf{y}] = 28.1538 - 0.8 \times 1.38462 = 27.046$°C, and $\Phi((28 - 27.046)/0.8321) = \Phi(1.1464) = 0.8742$.

Table 4.8: The full-load assessment before the test and after a single test reading of various values ($\sigma_T = 0.5$ K, prior $\mathcal{N}(0.160, 0.025^2)$). The column 23.2°C is the observed reading.

Quantity	Prior	$y_T = 22.4$	22.8	23.2	23.6	24.0
$\hat{\rho}$ [K/kW]	0.16000	0.15077	0.16000	0.16923	0.17846	0.18769
$\mathbb{E}[T_2]$ [°C]	27.60	27.05	27.60	28.15	28.71	29.26
sd(T_2) [K]	1.500	0.832	0.832	0.832	0.832	0.832
$\Pr(T_2 < 28$°C) [–]	0.6051	0.8742	0.6846	0.4267	0.1975	0.0647

A reading of 22.8°C, equal to what the prior predicted, leaves the mean where it was and only narrows the distribution; because the mean then lies below the limit, narrowing alone raises the probability from 0.605 to 0.685. The observed 23.2°C moves the mean across the limit and lowers the probability to 0.427. The posterior probability equals the prior one for a reading of 22.93°C and falls below one half for any reading above 23.09°C.

4.5 The report (step 6)

One reading at the test condition overturns the handbook's answer. The handbook resistance put the hall at 27.6°C at full load, under the limit, with probability 0.6051 of staying below it; the reading of 23.2°C moves the resistance up, from 0.160 to 0.169 K/kW, and leaves the probability at 0.4267, near even odds, although it narrows the full-load spread from 1.5 to 0.832 K. The answers to the questions of Table 4.2 are given in the units in which the questions were asked, conditional on the steady-state model, the test reading and the handbook prior; the operating decision follows from them.

4.5.1 The answers and their bands

Table 4.9 gives the answers. Each row is a question; the second column gives the answer in the limit of exact physical rows, and the third its band: the posterior standard deviation of an estimate, with the central 95 % credible interval of the full-load temperature, and for the probability the range over the model tolerances of Table 4.6.

The observed temperature was slightly warmer than the prior model predicted (22.8°C, the second entry of $\mathbf{z}^{(0)}$), so the inferred resistance rose above its handbook value and the implied conductance fell below 6.25 kW/K. The inferred test temperature is not the observed 23.2°C: the reading is noisy, and the prior, carried to T_1 by the physical row, also contributes

Table 4.9: The report for the equipment hall at 60 kW, after the test reading of 23.2°C. Answers in the limit of exact physical rows; bands are posterior standard deviations, a credible interval, or the range over model tolerances σ_p from 0 to 0.1 K.

Question	Answer	Band
What is the thermal resistance of the cooling path?	0.16923 K/kW (conductance 5.909 kW/K)	standard deviation 0.01387 K/kW
What is the hall temperature at the test condition?	23.077°C	standard deviation 0.416 K
What is the hall temperature at full load?	28.15°C	standard deviation 0.83 K; 26.52 to 29.78°C, central 95 % interval
Will the hall stay below 28°C at 60 kW?	probability 0.427; of exceeding, 0.573	0.4267 to 0.4312 over the model tolerances

information. The conductance in the first row is $1/\hat{\rho}$, the conductance evaluated at the estimated resistance; it is not the posterior mean of the conductance (Section 4.4.7). With correlations above 0.9996 the three entries are known almost only as one combination, the free direction $\mathbf{v}$, and the full-load interval contains temperatures both below and above the limit.

4.5.2 The operating decision

Conditional on the steady-state model, the reading and the prior, the estimated full-load temperature lies slightly above the limit, with substantial uncertainty on both sides of it. The reading reduced the uncertainty and, at the same time, lowered the probability of satisfying the operating limit, from 0.605 before it to 0.427. The evidence does not support a high-confidence claim that the proposed operating condition stays within the limit: the posterior probability of satisfying it is about 0.43. A decision to operate at full load would require an explicit risk-acceptance criterion and a consideration of the consequences of exceeding the limit. Further measurements or improved cooling capacity may be appropriate, but their value should be computed, not assumed. That is what an engineering inference report provides: the physical interpretation of a mathematical result, its uncertainty, and its implications for a decision.

Table 4.8 shows that the direction of the change was set by the value read, not by the act of reading. New evidence does not necessarily make the system look safer: it can make it look less safe, or move the probability in either direction, depending on the value observed. The purpose of inference is to improve the knowledge of the plant, not to produce reassuring probabilities.

> **Principle 4.2.** *A measurement narrows the posterior by an amount fixed before it is taken, but the direction in which it moves a threshold probability is set by the value it returns. Evidence improves knowledge; it does not promise reassurance.*

4.6 The next measurement (step 7)

The report leaves the operating question near even odds, with a band that a single reading at the test condition has set. A further reading is one more row of the residual vector, and the information it adds is known before its value is.

4.6.1 The preposterior covariance

A further reading at the test condition is the row $(T_1 - y)/\sigma_T$, whose line of the Jacobian is $\mathbf{j} = (0,\ 2,\ 0)$, the line of the first reading; a reading at full load, by a sensor of the same

spread, is the row $(T_2 - y)/\sigma_T$, with line $\mathbf{j}_\mathrm{f} = (0,\ 0,\ 2)$. The information that a reading adds, $\mathbf{j}^\mathsf{T}\mathbf{j}$, is known before the reading is taken, because its Jacobian depends on the sensor and on where it is placed, not on the value. For a candidate measurement a with information $\mathbf{\Lambda}_a$ the preposterior covariance is

$$\mathbf{\Sigma}_a = \left(\mathbf{\Sigma}^{-1} + \mathbf{\Lambda}_a\right)^{-1}, \tag{4.62}$$

formed at a finite tolerance, where $\mathbf{\Sigma}^{-1} = \mathbf{J}^\mathsf{T}\mathbf{J}$ exists, the object on which candidate measurements are ranked. In the limit of exact rows the posterior stays on the free direction, and a candidate adds to the information $\|\mathbf{J}\mathbf{v}\|^2 = 5200$ along it the square of its own line times $\mathbf{v}$: $(\mathbf{j}\mathbf{v})^2 = 60^2 = 3600$ for a reading at 30 kW and $(\mathbf{j}_\mathrm{f}\mathbf{v})^2 = 120^2 = 14\,400$ for a reading at 60 kW. The full-load standard deviation after the candidate is 60 divided by the square root of the total. For n further readings at 30 kW, $\mathbf{\Lambda}_a = n\,\mathbf{j}^\mathsf{T}\mathbf{j}$, which adds $4n$ to the (T_1, T_1) entry, and the third diagonal entry of $\mathbf{\Sigma}_a$ gives the standard deviation of the full-load temperature: 0.640 K for $n = 1$, 0.539 K for $n = 2$, 0.429 K for $n = 4$ and 0.325 K for $n = 8$. More readings resolve the question; they do not decide its answer. If the full-load temperature truly lies above 28°C, the probability tends to zero as the readings accumulate, and the decision then rests on the cooling path, not on the measurement.

4.6.2 The candidate measurements

Table 4.10 ranks three candidates. Each row is a measurement; the columns give the information along $\mathbf{v}$ after it, the standard deviations of the resistance and of the full-load temperature that the preposterior covariance (4.62) predicts, and the width it removes from the current full-load standard deviation of 0.8321 K. One reading at full load carries four times the information of one at the test condition, because the full-load temperature moves twice as far as the test temperature along $\mathbf{v}$, and it removes exactly the width of four test readings.

Table 4.10: Candidate measurements for the hall, in the limit of exact physical rows, from the posterior after the test reading. Each reading has $\sigma_T = 0.5$ K; the width removed is taken from the current $\sqrt{\Sigma_{33}} = 0.8321$ K.

Candidate	Information along $\mathbf{v}$ [(kW/K)^2]	sd(ρ) after [K/kW]	sd(T_2) after [K]	Width removed [K]
One further reading at 30 kW	8800	0.01066	0.640	0.192
Four further readings at 30 kW	19 600	0.00714	0.429	0.403
One reading at 60 kW	19 600	0.00714	0.429	0.403

4.6.3 The recommendation

Four further readings at the test condition and one reading at full load narrow the answer equally, and the choice between them rests on what else each tests. Readings at 30 kW need no operation at the proposed condition, but repeated readings of one sensor share its bias, and no number of them can show whether the resistance at 60 kW is the resistance at 30 kW. A reading at full load tests exactly that assumption, on which the prediction rests (Section 4.6.6), at the cost of running the hall at the condition in question. The next measurement is therefore one reading at full load, taken in a supervised run that can be stopped if the hall approaches its limit; where such a run is not acceptable, it is four further readings at the test condition, by more than one sensor where possible.

4.6.4 A second observation as an update

Before any run at full load, the first candidate of Table 4.10 arrives: a second, independent reading at the test condition, $y_{T,2} = 23.0°\mathrm{C}$, with the same standard deviation $\sigma_T = 0.5\,\mathrm{K}$; the first reading is now written $y_{T,1} = 23.2°\mathrm{C}$. Against the posterior predictive (4.61) it lies $(23.0 - 23.0769)/0.6504 = -0.12$ standard deviations from the predicted mean and passes that check. The posterior after the first reading serves as the prior for the second.

The precision form. Information from independent rows adds. The new reading is one more row of the residual vector, $(T_1 - y_{T,2})/\sigma_T$, whose line of the Jacobian is $\mathbf{j} = (0,\ 2,\ 0)$, the same as the first reading's. Its contribution to the information matrix is the outer product $\mathbf{j}^\mathsf{T}\mathbf{j}$, which changes the (T_1, T_1) entry alone,

$$\boldsymbol{\Lambda}_2 = \mathbf{J}^\mathsf{T}\mathbf{J} + \mathbf{j}^\mathsf{T}\mathbf{j} = \mathbf{J}^\mathsf{T}\mathbf{J} + \begin{bmatrix} 0 & 0 & 0 \\ 0 & 4 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \qquad \boldsymbol{\Sigma}_2 = \boldsymbol{\Lambda}_2^{-1}.$$

The old rows are balanced at $\hat{\mathbf{z}}$, $\mathbf{J}^\mathsf{T}\mathbf{r}(\hat{\mathbf{z}}) = \mathbf{0}$, so the gradient of the enlarged objective at $\hat{\mathbf{z}}$ comes from the new row alone, $\mathbf{j}^\mathsf{T}(23.0769 - 23.0)/0.5 = (0,\ 0.30769,\ 0)^\mathsf{T}$, and one Gauss–Newton step (4.29) from $\hat{\mathbf{z}}$ gives $\hat{\mathbf{z}}_2 = \hat{\mathbf{z}} - 0.30769\,(\boldsymbol{\Sigma}_2)_{\cdot 2}$, with $(\boldsymbol{\Sigma}_2)_{\cdot 2}$ the second column of $\boldsymbol{\Sigma}_2$. In the limit of exact rows this column is $(3.4091 \times 10^{-3},\ 0.10227,\ 0.20455)^\mathsf{T}$, and

$$\hat{\mathbf{z}}_2 = \begin{bmatrix} 0.169231 \\ 23.0769 \\ 28.1538 \end{bmatrix} - 0.30769 \begin{bmatrix} 0.0034091 \\ 0.10227 \\ 0.20455 \end{bmatrix} = \begin{bmatrix} 0.168182 \\ 23.0455 \\ 28.0909 \end{bmatrix}.$$

The first entry of $\boldsymbol{\Sigma}_2$ is $1.1364 \times 10^{-4}\,(\mathrm{K/kW})^2$, so that $\sqrt{(\boldsymbol{\Sigma}_2)_{11}} = 0.010660\,\mathrm{K/kW}$.

The rank-one form. The same result follows from the first posterior alone, without forming and inverting $\boldsymbol{\Lambda}_2$. The second reading measures $\mathbf{h}^\mathsf{T}\mathbf{z} = T_1$ with $\mathbf{h} = (0,\ 1,\ 0)^\mathsf{T}$; its predicted value under the first posterior is $\mathbf{h}^\mathsf{T}\hat{\mathbf{z}} = 23.0769°\mathrm{C}$ and its predictive variance, from equation (4.61), is $S = \mathbf{h}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{h} + \sigma_T^2 = \Sigma_{22} + \sigma_T^2 = 0.42308\,\mathrm{K}^2$. Conditioning the Gaussian on the new row gives the update

$$\hat{\mathbf{z}}_2 = \hat{\mathbf{z}} + \mathbf{K}\,(y_{T,2} - \mathbf{h}^\mathsf{T}\hat{\mathbf{z}}), \qquad \boldsymbol{\Sigma}_2 = \boldsymbol{\Sigma} - S\,\mathbf{K}\mathbf{K}^\mathsf{T}, \qquad \mathbf{K} = \frac{\boldsymbol{\Sigma}\mathbf{h}}{S}, \tag{4.63}$$

where $\mathbf{K}$ is the gain vector, one entry per state entry, and $y_{T,2} - \mathbf{h}^\mathsf{T}\hat{\mathbf{z}}$ the innovation, the part of the reading the first posterior did not predict. Because $\boldsymbol{\Sigma}_2\mathbf{h} = \boldsymbol{\Sigma}\mathbf{h} - \mathbf{K}\,\mathbf{h}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{h} = (\sigma_T^2/S)\,\boldsymbol{\Sigma}\mathbf{h}$, the gain is also $\mathbf{K} = \boldsymbol{\Sigma}_2\mathbf{h}/\sigma_T^2$, the same vector as the gain (4.56) of the posterior that includes the new reading. The gain follows from writing the posterior mean as $\hat{\mathbf{z}} + \mathrm{Cov}(\mathbf{z}, Y_{\mathrm{new}})\,(y_{T,2} - \mathbb{E}[Y_{\mathrm{new}}])/S$ for the jointly Gaussian state and new reading, with $\mathrm{Cov}(\mathbf{z}, Y_{\mathrm{new}}) = \boldsymbol{\Sigma}\mathbf{h}$, the second column of $\boldsymbol{\Sigma}$. With the limit covariance (4.35),

$$\mathbf{K} = \frac{1}{0.42308} \begin{bmatrix} 0.0057692 \\ 0.17308 \\ 0.34615 \end{bmatrix} = \begin{bmatrix} 0.013636 \\ 0.40909 \\ 0.81818 \end{bmatrix},$$

in kW^{-1} for the first entry and dimensionless for the other two; the innovation is $23.0 - 23.0769 = -0.0769\,\mathrm{K}$, and $\hat{\mathbf{z}}_2 = \hat{\mathbf{z}} - 0.0769\,\mathbf{K} = (0.168182,\ 23.0455,\ 28.0909)^\mathsf{T}$, the state of the precision form. The covariance update subtracts the outer product $S\mathbf{K}\mathbf{K}^\mathsf{T}$; on its first and last diagonal entries,

$$(\boldsymbol{\Sigma}_2)_{11} = 1.9231 \times 10^{-4} - 0.42308 \times 0.013636^2 = 1.1364 \times 10^{-4}\,(\mathrm{K/kW})^2,$$
$$(\boldsymbol{\Sigma}_2)_{33} = 0.69231 - 0.42308 \times 0.81818^2 = 0.40909\,\mathrm{K}^2,$$

the values of the precision form. Because $S\,\mathbf{K}^{\mathsf{T}} = \mathbf{h}^{\mathsf{T}}\boldsymbol{\Sigma}$, the covariance update is also

$$\boldsymbol{\Sigma}_2 = \left(\mathbf{I} - \mathbf{K}\mathbf{h}^{\mathsf{T}}\right)\boldsymbol{\Sigma},$$

with $\mathbf{I}$ the identity matrix: each column of $\boldsymbol{\Sigma}$ loses the gain vector times that column's second entry. The correction is the outer product of one vector, which is the reason for the name rank-one update. For a linear-Gaussian model it is exact, and one more independent reading costs one product of $\boldsymbol{\Sigma}$ with a vector rather than a new inference. For a strongly nonlinear model, or after a shift of the estimate large enough to make the Jacobian at the old estimate stale, the model is linearized again or the full fit is repeated.

The updated assessment. Table 4.11 follows the assessment through the two readings. Each column is a state of knowledge, before any reading, after the first, after both; the rows are the information about the resistance, the reciprocal of the first diagonal entry of the covariance, the MAP resistance, the mean and standard deviation of the full-load temperature, the threshold probability, and the 95 % credible interval.

Table 4.11: The full-load assessment before the test, after the first reading (23.2°C) and after the second (23.0°C).

Quantity	Prior	One reading	Two readings
$1/\Sigma_{11}$ [(kW/K)^2]	1600	5200	8800
$\hat{\rho}$ [K/kW]	0.160000	0.169231	0.168182
$\mathbb{E}[T_2]$ [°C]	27.600	28.154	28.091
sd(T_2) [K]	1.500	0.832	0.640
Pr($T_2 < 28$°C) [–]	0.6051	0.4267	0.4435
95 % interval [°C]	24.66, 30.54	26.52, 29.78	26.84, 29.34

After the second reading the full-load temperature is the third entry of $\hat{\mathbf{z}}_2$, 28.091°C, with standard deviation $\sqrt{(\boldsymbol{\Sigma}_2)_{33}} = \sqrt{0.40909} = 0.640$ K, and the probability of staying below the limit becomes $\Phi(-0.1421) = 0.4435$. The interval narrowed by a quarter, yet the probability moved only from 0.427 to 0.443, because two effects acted in opposite directions. The narrowing alone, with the mean held at 28.154°C, would have lowered the probability to 0.405, since the mean lies above the limit; the shift of the mean alone, with the spread held at 0.832 K, would have raised it to 0.456. The effect of actual new evidence cannot be described by its reduction of the covariance alone: the value observed matters.

4.6.5 What the analysis showed

The handbook said the hall passes at 27.6°C, and the reading scaled to full load said it fails; the analysis says neither. With the reading, the handbook and the physics combined, the probability that the hall stays below 28°C at 60 kW is 0.4267, less than even, and the full-load temperature is not known well enough to call it either way. One reading at the test condition fixes the thermal resistance of the cooling path to about ±8 % and the full-load temperature to about ±0.8 K, one standard deviation each, although the hall has never run at 60 kW: the physics carries the reading along the free direction $\mathbf{v}$, from T_1 to ρ and from ρ to T_2. The posterior is a ridge along that direction, the three entries known almost only as one combination, so every question asked of the hall is a question about where it sits on the ridge. The 60 kW question in particular is a resistance question: the hall stays below 28°C exactly when the resistance stays below ρ^*, and the reading at 30 kW places the resistance close to ρ^*, which leaves the question near even odds. Further readings at the test condition narrow the ridge and resolve the question without deciding its answer. A reading at full load would act on the third column

of $\mathbf{J}$ directly, and its worth is what its preposterior covariance $\boldsymbol{\Sigma}_a$ says before it is taken: as much as four readings at the test condition, with the shared resistance itself put to the test.

4.6.6 Limits of the model

The steady-state balance (4.2) neglects the heat stored in the hall and its equipment, and it describes the temperature long after a load change, not minutes after it; Exercise 4.6 formulates the transient model that a question about the first minutes of full load would need. The single effective path assumes that the resistance is the same at 60 kW as at 30 kW; a path whose resistance changes with load, for example because fans or chillers change state, would break the link drawn in Figure 4.2, and one test reading cannot detect it. The ambient temperature is pinned at 18°C; if it is uncertain, the budget (4.50) adds its variance to the answer.

The probabilistic assumptions have limits of their own. The Gaussian prior on a positive resistance is harmless here, with a negative mass of 8×10^{-11}, but a wide prior on a small quantity would place real probability below zero; a logarithmic parameterization avoids this at the price of a nonlinear model and an approximate posterior. The sensor error is taken as Gaussian, unbiased and of known spread; a bias would shift every inferred quantity, and Section 4.2.2 names the comparison of models that would detect it. The physical rows are taken in the limit of zero tolerance; a real model discrepancy is a finite σ_p, which widens Σ_{11} and adds σ_p^2 to the full-load variance by equation (4.34). Each of these is a statement about the model, and the probability 0.4267 is conditional on all of them.

Chapter 5 turns from continuous temperatures to discrete failure events and applies the same workflow to a component that succeeds or fails on each demand.

4.7 What to retain

Definition 4.1 (The anatomy of an inference)**.** A reliability probability over a plant graph is obtained in two operations. Inference sums one term per row into the objective $\mathcal{L} = -\ln p(\mathbf{z} \mid \mathbf{y})$ up to a constant, and uses only its gradient and Hessian. A Gaussian row's term is half its squared residual; a count row's is its negative log-probability, which joins the Jacobian through its working residual and line, reproducing the row's gradient and information but not its value. The physical rows, the readings and the priors, each divided by its standard deviation, stack into one residual vector $\mathbf{r}(\mathbf{z})$ with Jacobian $\mathbf{J}$; the normal equations $(\mathbf{J}^\mathsf{T}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^\mathsf{T}\mathbf{r}$ lead to the MAP state $\hat{\mathbf{z}}$, and the inverse of the information matrix, $\boldsymbol{\Sigma} = (\mathbf{J}^\mathsf{T}\mathbf{J})^{-1}$, is the posterior covariance, exactly for a linear-Gaussian model and under the Laplace approximation otherwise. An exact physical relation is the limit of a row's tolerance, never a substitution. Query evaluation then reads the answer off the posterior: an entry of the state with its variance Σ_{ii}, or a derived quantity $f(\mathbf{z})$ with variance $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$, and a threshold probability as the normal tail $\Phi\big((f^* - f(\hat{\mathbf{z}}))/\sqrt{\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}}\big)$, which for an entry is $\Phi\big((f^* - \hat{z}_i)/\sqrt{\Sigma_{ii}}\big)$.

Three conclusions carry beyond the example. Physical relations connect the quantities that are measured to those that are not, which is what allows an unobserved property, here a full-load temperature never reached, to be inferred from indirect evidence. Uncertainty is not an error bar attached to a point estimate but a probability distribution, and it must be propagated through the physical model to answer an engineering question. And the posterior is reusable: once the joint posterior is in hand, many questions are answered without repeating the inference, and new evidence updates it rather than replacing it. A derived quantity is propagated in one of two ways: its spread by the gradient, $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$, and its density by the change of variables, which multiplies the density by the Jacobian factor of the map and so keeps the skew and the bounds that a linearization loses.

The inference of the hall returns the posterior state $\hat{\mathbf{z}}$ of equation (4.45), its covariance $\boldsymbol{\Sigma}$ of equations (4.46) and (4.35), and the Jacobian $\mathbf{J}$ of equation (4.25), from which the gradients

of derived quantities are formed. Table 4.12 lists what the chapter has computed from them without repeating the inference: the four answers of Table 4.2, and three further quantities that later chapters reuse. Each row is a quantity, the calculation on the posterior that gives it, and the place where it is carried out.

Table 4.12: The four answers of Table 4.2 and three further quantities reused later, all computed from one posterior of the equipment hall.

Quantity	Calculation on the posterior	Section
The answers		
Thermal resistance of the cooling path	the entry $\hat{z}_1$ and $\sqrt{\Sigma_{11}}$	4.2.9
Hall temperature at the test condition	the entry $\hat{z}_2$ and $\sqrt{\Sigma_{22}}$	4.3.6
Hall temperature at full load	the entry $\hat{z}_3$, $\sqrt{\Sigma_{33}}$ and the credible interval	4.3.6
Probability of staying below 28°C at 60 kW	the Gaussian tail $\Phi\big((T_{\max} - \hat{z}_3)/\sqrt{\Sigma_{33}}\big)$	4.3.7
Further quantities reused later		
Spread of the conductance	$\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ with $\mathbf{g} = (-1/\hat{z}_1^2, 0, 0)^{\mathsf{T}}$, and the exact push-forward	4.4.7
The entry that drives the temperature	the gradient and the variance budget	4.3.2
Effect of another reading	the information matrix plus the reading's $\mathbf{j}^{\mathsf{T}}\mathbf{j}$	4.6.1

The table is one inference followed by a column of cheap queries on its result, the division of Section 4.1.2. Two kinds of question fall outside it: a question that compares two models, which needs the evidence of each (Section 4.2.2), and the arrival of new evidence, which changes the posterior itself (Section 4.6.4).

Table 4.13 lists the five mathematical objects of the chapter that recur throughout the book, with the form of each and its purpose. The first three are the working objects when the posterior is well represented by a local Gaussian. The posterior distribution itself is the more fundamental object: it may be Gaussian, multimodal, discrete, a mixture or otherwise non-Gaussian, and its structure decides which computational method is appropriate.

Table 4.13: The five mathematical objects of an inference.

Object	Form	Purpose
State estimate	$\hat{\mathbf{z}}$	the most probable (or otherwise selected) central state
Posterior covariance	$\boldsymbol{\Sigma}$	local uncertainty of the state and the dependence among its entries
Output gradient	$\mathbf{g} = \nabla f(\hat{\mathbf{z}})$	connects state uncertainty to a derived quantity
Posterior distribution	$p(\mathbf{z} \mid \mathbf{y}, M)$	the uncertainty about the state given the evidence
Model evidence	$p(\mathbf{y} \mid M)$	normalizes the posterior and compares models

Table 4.14 lists the terms of the chapter with their expressions, units and the subsection

that defines each.

Table 4.14: The terms of the anatomy of an inference and where each is defined. Units for the equipment hall.

Term	Expression	Units	Meaning	Defined in
Thermal resistance	$\rho = 1/G$, $T = T_a + Q\rho$	K/kW	makes every row linear in the state	4.1.4
Physical rows	$r_1 = T_1 - 18 - 30\rho$, r_2	K	conservation at each hall node	4.1.5
Decomposition	$\Pr(T_2 < T_{\max})$ from $\hat{z}_3$, Σ_{33}	–	the question as one marginal	4.1.6
Model tolerance	σ_p; exact rows as $\sigma_p \to 0$	K	model discrepancy of a physical row	4.2.1
Likelihood	$p(y_T \mid \mathbf{z})$, eq. (4.14)	K^{-1}	the reading as a function of the state	4.2.1
Residual vector	$\mathbf{r}(\mathbf{z})$, eq. (4.23)	–	four rows, each over its standard deviation	4.2.3
Objective	$\mathcal{L} = \frac{1}{2}\mathbf{r}^{\mathsf{T}}\mathbf{r}$	–	minus log posterior	4.2.3
Jacobian	$\mathbf{J} = \partial\mathbf{r}/\partial\mathbf{z}$, 4×3	mixed	one line per row, one column per entry	4.2.3
Information matrix	$\mathbf{J}^{\mathsf{T}}\mathbf{J} = \sum_k \mathbf{j}_k^{\mathsf{T}}\mathbf{j}_k = \tilde{\mathbf{J}}^{\mathsf{T}}\mathbf{W}\tilde{\mathbf{J}}$	mixed	the Hessian when $\mathbf{C} = \mathbf{0}$	4.2.3
General objective	$\mathcal{L} = \sum_k \mathcal{L}_k$, $\mathcal{L}_k = -\ln p_k$	–	one term per row, any distribution	4.2.7
Count row, working form	$r_k = (\mu - k)/\sqrt{\mu}$, $\mathbf{j}_k = \sqrt{\mu}\,\nabla \ln \mu^{\mathsf{T}}$	–	gradient and information, not the term	4.2.7
Hessian	$\mathbf{H} = \mathbf{J}^{\mathsf{T}}\mathbf{J} + \mathbf{C}$	mixed	observed information; $\mathbf{C}$ named	4.2.7
Normal equations	$(\mathbf{J}^{\mathsf{T}}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^{\mathsf{T}}\mathbf{r}$	mixed	Gauss–Newton step to $\hat{\mathbf{z}}$	4.2.4
Posterior covariance	$\mathbf{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1}$	mixed	spread is inverse information	4.2.5
Laplace approximation	$\mathcal{N}(\hat{\mathbf{z}}, \mathbf{H}^{-1})$	–	local Gaussian posterior	4.2.6
Free direction	$\mathbf{v}$, $\mathbf{J}_{\text{phys}}\mathbf{v} = \mathbf{0}$	mixed	$\mathbf{\Sigma} \to \mathbf{v}\mathbf{v}^{\mathsf{T}}/5200$, rank one	4.2.4
Delta method	$\mathrm{Var}(f) \approx \mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}$	varies	linearized propagation	4.3.1
Change of variables	$p_u(u) = p_x(h^{-1}(u))\,\lvert \mathrm{d}x/\mathrm{d}u \rvert$, eq. (4.51)	varies	density of a transformed quantity	4.3.3
Variance budget	eq. (4.50)	K^2	one term per entry plus cross terms	4.3.2
Threshold probability	$\Phi\big((T_{\max} - \hat{z}_3)/\sqrt{\Sigma_{33}}\big)$	–	normal tail at the limit	4.3.5
Credible interval	$\hat{z}_3 \pm 1.960\sqrt{\Sigma_{33}}$	°C	central 95 % posterior range	4.3.4
Prior predictive	$\mathcal{N}(22.8, 0.9014^2)$	°C	reading predicted before use	4.4.2
Posterior predictive	$\mathcal{N}(23.077, 0.6504^2)$	°C	next reading predicted	4.4.4
Evidence	$p(y_T \mid M)$, eq. (4.18)	K^{-1}	normalizer; compares models	4.2.2
Bayes factor	$B_{21} = p(y_T \mid M_2)/p(y_T \mid M_1)$	–	evidence ratio of two models	4.2.2
Rank-one update	eq. (4.63)	–	adds one reading to a posterior	4.6.4
Preposterior covariance	$\mathbf{\Sigma}_a = (\mathbf{\Sigma}^{-1} + \mathbf{\Lambda}_a)^{-1}$, eq. (4.62)	mixed	ranks candidate measurements	4.6.1

4.8 Exercises

Exercise 4.1 (Application: a complete inference)**.** A different equipment hall obeys $T = T_a + Q\rho$ with $T_a = 20°\mathrm{C}$ and a test load $Q = 20\,\mathrm{kW}$. The prior on the resistance is $\rho \sim \mathcal{N}(0.20, 0.04^2)$ in K/kW, and a sensor with $\sigma_T = 0.4\,\mathrm{K}$ reads $y_T = 24.5°\mathrm{C}$. The question is the temperature at a proposed load of 40 kW. (a) Write the state $\mathbf{z} = (\rho, T_1, T_2)^\mathsf{T}$, with T_1 the test temperature and T_2 the temperature at 40 kW, the four normalized rows, and the 4×3 Jacobian at $\sigma_p = 0.01\,\mathrm{K}$. ($\mathbf{r} = ((T_1 - 20 - 20\rho)/\sigma_p,\ (T_2 - 20 - 40\rho)/\sigma_p,\ (T_1 - 24.5)/0.4,\ (\rho - 0.20)/0.04)^\mathsf{T}$; $\mathbf{J}$ has the lines $(-2000, 100, 0)$, $(-4000, 0, 100)$, $(0, 2.5, 0)$ and $(25, 0, 0)$.) (b) Take one Gauss–Newton step from the handbook state and compute $\boldsymbol{\Sigma}$ in the limit of exact rows. (At $\mathbf{z}^{(0)} = (0.20,\ 24.0,\ 28.0)^\mathsf{T}$ only the reading row is nonzero, -1.25, and $\mathbf{J}^\mathsf{T}\mathbf{r} = (0,\ -3.125,\ 0)^\mathsf{T}$; the normal equations give $\hat{\mathbf{z}} = (0.2200,\ 24.400,\ 28.800)^\mathsf{T}$, and a second step does not move it. The free direction is $\mathbf{v} = (1,\ 20,\ 40)^\mathsf{T}$, with $\mathbf{J}\mathbf{v} = (0,\ 0,\ 50,\ 25)^\mathsf{T}$ and information $50^2 + 25^2 = 3125$, so $\boldsymbol{\Sigma} \to \mathbf{v}\mathbf{v}^\mathsf{T}/3125$: $\Sigma_{11} = 3.2 \times 10^{-4}$, $\sqrt{\Sigma_{11}} = 0.01789\,\mathrm{K/kW}$; at $\sigma_p = 0.01\,\mathrm{K}$ the fit gives $\Sigma_{11} = 3.2016 \times 10^{-4}$ and a correlation of 0.9996 between ρ and T_1.) (c) Read off the posterior mean and standard deviation of the temperature at 40 kW, and the probability that it stays below 30°C. ($\hat{z}_3 = 28.8°\mathrm{C}$ and $\sqrt{\Sigma_{33}} = \sqrt{1600/3125} = 0.7155\,\mathrm{K}$; $\Pr(T_2 < 30 \mid \mathbf{y}) = \Phi(1.2/0.7155) = \Phi(1.677) = 0.9532$, against $\Phi(2/1.6) = 0.8944$ under the prior alone.)

Exercise 4.2 (Derivation: a variance budget with correlated entries)**.** For $T = T_a + Q\rho$ at the known load $Q = 50\,\mathrm{kW}$, the posterior covariance of (T_a, ρ) is $\boldsymbol{\Sigma} = \begin{bmatrix} 0.25 & -0.003 \\ -0.003 & 0.0001 \end{bmatrix}$, in K^2, K·K/kW and $(\mathrm{K/kW})^2$. (a) Derive the temperature variance from $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$ and compute the standard deviation. ($\mathbf{g} = (1, 50)^\mathsf{T}$; $\mathrm{Var}(T) = 0.25 + 50^2 \times 0.0001 + 2 \times 50 \times (-0.003) = 0.25 + 0.25 - 0.30 = 0.20\,\mathrm{K}^2$, $\mathrm{sd}(T) = 0.4472\,\mathrm{K}$.) (b) State the result if the covariance term were ignored, and explain the difference. ($\sqrt{0.50} = 0.7071\,\mathrm{K}$, 58 % too large; the posterior correlation is $-0.003/\sqrt{0.25 \times 0.0001} = -0.6$, and with sensitivities of the same sign the negative cross term cancels part of the two diagonal terms: the evidence pins the combination $T_a + 50\rho$ better than either entry.)

Exercise 4.3 (Derivation: a nonlinear derived quantity)**.** With $\rho \mid \mathbf{y} \sim \mathcal{N}(0.169231, 0.013868^2)$ from the chapter's example, derive the first-order approximation to the variance of the conductance $G = 1/\rho$, compute its standard deviation, and explain why the result is only an approximation. (With $G(\mathbf{z}) = 1/z_1$, $\mathbf{g} = (-1/\hat{z}_1^2,\ 0,\ 0)^\mathsf{T}$, so $\mathrm{Var}(G \mid \mathbf{y}) \approx \mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g} = \Sigma_{11}/\hat{z}_1^4$; $\hat{G} = 1/0.169231 = 5.909\,\mathrm{kW/K}$ and $\mathrm{sd}(G \mid \mathbf{y}) \approx 0.013868/0.169231^2 = 0.484\,\mathrm{kW/K}$. The map $1/\rho$ is nonlinear, so G is not Gaussian: its exact 5–95 % band is $[5.207, 6.830]$, skewed, against the linearized $[5.113, 6.706]$; strictly, the moments of $1/\rho$ do not exist under a Gaussian ρ, which has positive density at zero, so positivity and the transformed distribution must be handled explicitly when exact conductance statements are needed.)

Exercise 4.4 (Formulation: identifiability and the plant graph)**.** In $T = T_a + Q\rho$ suppose that both T_a and ρ are unknown and that there is a single temperature reading at one known load. (a) Write the observation Jacobian with respect to (T_a, ρ) and state whether the reading can determine both unknowns without further information. ($y = T_a + Q\rho + \epsilon$, so $\mathbf{J} = [1 \quad Q]$, of rank one for two unknowns: the reading determines only the combination $T_a + Q\rho$, and every pair with the same combination explains it equally well.) (b) Name a measurement or operating condition that resolves the ambiguity. (A second reading at a different known load $Q_2 \neq Q_1$ gives $\mathbf{J} = \begin{bmatrix} 1 & Q_1 \\ 1 & Q_2 \end{bmatrix}$, with determinant $Q_2 - Q_1 \neq 0$ and full rank; alternatively, a direct reading of the ambient temperature, a new node observed in the graph, removes the ambiguity.)

Exercise 4.5 (Application: Bayesian evidence)**.** A model predicts $Y \sim \mathcal{N}(24, 1^2)$ in °C, and the reading $y = 25°\mathrm{C}$ arrives. (a) Compute the marginal likelihood $p(y \mid M)$. ($\frac{1}{\sqrt{2\pi}} \exp[-(25 -$

$24)^2/2] = 0.39894 \times 0.60653 = 0.2420\,\mathrm{K}^{-1}$.) (b) Explain why this value is not the probability that M is correct. (It is a probability density in inverse temperature units, 2.420×10^{-4} per millikelvin for the same reading, not $\Pr(M \mid y)$; a posterior model probability requires prior probabilities and the evidence of the alternative models under comparison.)

Exercise 4.6 (Formulation: model validity and thermal storage)**.** The chapter assumed the steady-state balance $Q = G(T - T_a)$. The load now steps from 30 kW to 60 kW, and the engineer asks for the temperature five minutes later. (a) Explain why the steady-state equation may be insufficient. (It neglects the energy stored in the air, structure and equipment of the hall, so it gives only the temperature reached long after the change; five minutes may be short against the thermal time constant.) (b) Starting from an energy balance, formulate a model with thermal storage and name the new parameter and the condition it needs. (With an effective thermal capacitance C_{th} in kJ/K, $C_{\mathrm{th}}\,\mathrm{d}T/\mathrm{d}t = Q - G(T - T_a)$; the new parameter is C_{th} and an initial temperature $T(0)$ is required. For constant Q, G and T_a after the step, $T(t) = T_a + Q/G + [T(0) - T_a - Q/G]\,\mathrm{e}^{-Gt/C_{\mathrm{th}}}$, with time constant C_{th}/G; the probability for the temperature at five minutes propagates the posterior of G, C_{th} and $T(0)$ through this time-dependent model rather than through the steady-state equation.)

Part II

Estimating Reliability from Data

Chapter 5

Estimating Failure Probability from Observed Data

5.1 Failure probability from a record of tests

An emergency cooling valve has been demanded fifty times on the test stand. It opened forty-eight times and failed to open twice. The customer's specification asks for a valve that opens on at least 95 % of demands, and the first reading of the record is simple: forty-eight in fifty is 96 %, above the specification, so the valve passes. The engineer who must sign the acceptance asks something else: how sure can the plant be that the valve itself, not the fifty tests, opens on 95 % of its demands? Fifty tests are not many, and earlier valves of the same type failed on about one demand in ten.

The two readings answer different questions. The first is a statistic of the record. The second is a statement about an unknown, the probability that this valve fails on a demand, and it depends on how many tests there were, on what was known before them and on how sure the decision must be. The question of this chapter is the second,

how the probability that a component fails on demand is estimated from a limited number of tests, how uncertain that estimate is, and whether the component satisfies a stated reliability requirement.

It is answered with the method of the thermal example of Chapter 4: a graph, the evidence attached to it as rows, a posterior over the state, and each question read from that posterior. What changes is the evidence, zeros and ones in place of temperatures with Gaussian errors, and with it the likelihood. The first reading, 96 % against a specification of 95 %, is worth keeping in mind: it looks like a pass.

Jargon: per-demand reliability, Bernoulli trial, binomial distribution, maximum-likelihood estimate, Beta distribution, conjugate prior, pseudo-counts, credible interval, one-sided credible bound, assurance level, posterior predictive distribution, Beta–binomial distribution, plug-in estimate, reliability demonstration test, operating characteristic

Per-demand reliability is the probability that a component performs its function when demanded once, under specified conditions. A *Bernoulli trial* is a single test with two outcomes, failure and success; the *binomial distribution* is the distribution of the number of failures in a fixed number of independent Bernoulli trials with a common failure probability. The *maximum-likelihood estimate* (MLE) is the parameter value that maximizes the probability of the observed data. The *Beta distribution* is a family of densities on the interval $[0, 1]$, the natural prior for a probability; a prior is *conjugate* to a likelihood when the posterior belongs to the same family, and the Beta parameters then act as *pseudo-counts* of prior failures and successes. A *credible interval* contains the unknown with a stated posterior probability; a *one-sided credible bound* is a single limit with that property. The *assurance level* is the posterior probability that a requirement demands before the component is accepted. The *posterior predictive distribution* is the distribution of future outcomes averaged over the posterior; for future failure counts under a Beta posterior it is the *Beta–binomial distribution*. A *plug-in estimate* substitutes a point estimate for the unknown parameter instead of averaging over it. A *reliability demonstration test* is a test campaign designed in advance to show that a requirement is met; its *operating characteristic* is the probability that the campaign passes, as a function of the true failure probability.

Key idea

A record of k failures in n comparable tests and a Beta(a, b) prior give the exact posterior Beta$(a + k,\ b + n - k)$ for the failure probability per demand, and its Laplace approximation in the logit, the state the rest of the book would carry. The posterior is not the engineering answer; it is the reusable object from which the estimate, its uncertainty, the probability that a requirement is met, the prediction of future tests and the size of a demonstration test are each derived, and these are different numbers because they answer different questions.

5.1.1 Failure probability, reliability and failure rate

Three quantities are confused in practice, and the chapter keeps them apart. The *failure probability per demand* p is the probability that a component fails when it is demanded once; it is dimensionless and lies between zero and one. The *per-demand reliability* is its complement, the probability of success on one demand,

$$R = 1 - p. \tag{5.1}$$

The *failure rate* λ, by contrast, is a number of failures per unit of operating time, in h^{-1}, and belongs to a component that operates continuously. The rate and the probability are connected only through a lifetime model and a mission duration: under the exponential lifetime model, the probability of failure over a mission of τ hours is $1 - \mathrm{e}^{-\lambda\tau}$, which depends on τ (Section 5.6.5). A valve that is demanded and a pump that runs for a thousand hours therefore produce different kinds of evidence.

This chapter treats the first kind. The evidence is a set of demands, each with a binary outcome, and the unknown is the single number p that the outcomes share. The failure rate and its estimation from operating records are the subject of the next chapter.

5.1.2 Binary evidence and its likelihood

The outcome of one test is a binary variable,

$$\text{outcome} = \begin{cases} 1, & \text{failure}, \\ 0, & \text{success}, \end{cases}$$

and a test record is a sequence of such outcomes. The Gaussian reading rows of Chapter 4 do not describe this evidence: a reading there was a temperature plus a normal error, whereas here the probability of each outcome is the unknown itself. The probability rules are unchanged, posterior proportional to likelihood times prior, but the likelihood is derived from the binary outcomes, and the prior must live on the interval $[0, 1]$. The posterior is carried, as everywhere in the book, by a state vector, here with one entry; with the Beta prior it also has a closed form in two numbers, from which the answers are read, with the fit set beside it.

5.1.3 The physical system and its reliability graph (step 0)

A manufacturer evaluates an electrically operated cooling valve. The valve must open whenever it receives an emergency cooling command. The manufacturer performs repeated functional tests under standardized conditions, and each test has one of two outcomes: the valve opens correctly, a success, or it fails to open correctly, a failure. Each test is assumed to be performed under comparable operating conditions, and the component population or the test process is assumed to have a stable failure probability,

$$p = \Pr(\text{the valve fails on a demand}).$$

The reliability $R = 1 - p$ of equation (5.1) is then the probability that the valve responds successfully to one demand under the specified test conditions. It is a per-demand reliability, not a reliability over a specified amount of operating time, and the distinction returns in Section 5.6.5.

Figure 5.1(a) draws the test as a graph. The box on the left is the emergency command, the source that issues one standardized demand per test; it is pinned, since the test conditions are fixed by the procedure. The green box in the middle is the valve, the free node of the graph; it carries the one entry of the state, the failure probability p carried as its logit, which no test observes directly. The box on the right is the observed outcome X_i of test i, reached by a dashed observation arrow, with its two values beneath it. The physical edge from command to valve is the demand; the observation is the only place where evidence enters.

5.1.4 The observed record and the state

The manufacturer performs $n = 50$ tests and observes $k = 2$ failures, and therefore $n - k = 48$ successes. Figure 5.1(b) shows the record as a grid of fifty squares, one per test in the order performed; the two red squares are the failures. The record is illustrative. Its immediate question is the probability that the valve fails on a demand, given the test results. Written as a distribution over the unknown,

$$f(p \mid k = 2,\, n = 50), \tag{5.2}$$

the posterior density of the failure probability given the record. The conventional notation $p(p \mid k, n)$ uses the letter p twice, once for a density and once for the unknown probability; to avoid that, a density of the failure probability is written $f(\cdot)$ in this chapter. It is the density $p(\cdot)$ of the notation table applied to the one quantity p, and it is not the failure density $f(t)$ of a lifetime, which this chapter does not use.

The failure probability is the only unknown of this graph. It must lie between zero and one, and an entry with such a bound is better carried on an unbounded scale, as the conductance of

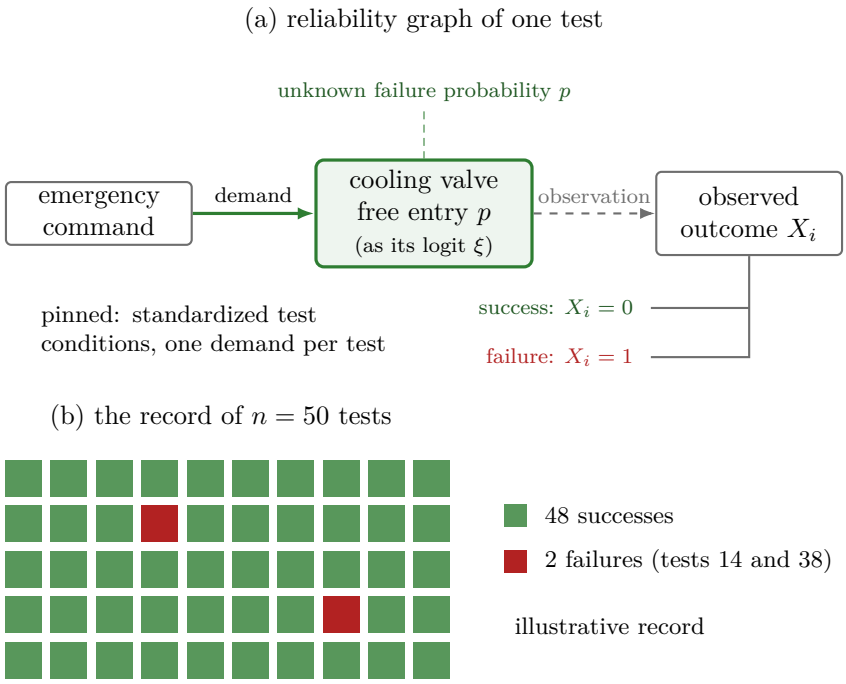


Figure 5.1: The valve test. (a) Reliability graph of one test: the emergency command (pinned source) demands the valve; the valve is the free node carrying the state entry, the failure probability p carried as its logit ξ (Section 5.1.4); the outcome X_i is observed (dashed). (b) The illustrative record: fifty tests read row by row, each square one test, green for a success and red for a failure.

the thermal examples is carried by its logarithm g. The state entry is therefore the logit of the failure probability,

$$\xi = \ln \frac{p}{1-p}, \qquad p = \frac{1}{1+\mathrm{e}^{-\xi}}, \qquad \frac{\mathrm{d}p}{\mathrm{d}\xi} = p(1-p), \tag{5.3}$$

which takes every real value as p runs over $(0, 1)$. The state is this one number,

$$\xi \in \mathbb{R}. \tag{5.4}$$

The general forms of the book write the state as a vector $\mathbf{z}$; here it has the single entry ξ, so every quantity of the fit is a number: the objective's slope and curvature in ξ, the information Λ, the posterior variance $\sigma_\xi^2 = 1/\Lambda$ of the logit, and the derivative g of a derived quantity. The failure probability, the reliability and every quantity the questions ask for are derived quantities, functions of the state. Carrying p itself as the state entry is possible, and Section 5.4.1 states in one sentence what it costs. The evidence $\mathbf{y}$ is the record (k, n) together with the prior.

5.1.5 The probabilistic questions (step 1)

The manufacturer may ask several questions of the same record, and step 1 of the method writes each as a probability or an expectation over the state, with the unknown on the left of the bar and the record on the right. Table 5.1 lists them: the first column is the question in engineering language; the second the general form that every question of its category takes, written in the book's notation for a quantity $f(\mathbf{z})$ derived from the state, a threshold f^* on it, the posterior density $p(\mathbf{z} \mid \mathbf{y})$, a future observation y_{new} or the outcome K_n of a planned campaign of n tests; the third its mathematical form for the valve; the fourth its category; and the last the subsection where it is answered. All seven concern the same unknown failure probability, and they are not

mathematically equivalent. The first three are estimation, answered by the posterior itself and its moments. The fourth is a prediction under a threshold, a tail probability of the posterior, with the assurance rule that turns it into an acceptance decision. The fifth and sixth are predictions of future outcomes, which Chapter 2 also files under prediction under a threshold: each is answered by a predictive distribution, an expectation over the posterior of the probability of a future outcome given p. The seventh is a design question, which chooses the size of a test before it is run. The atlas of Appendix A lists these questions.

Table 5.1: The questions asked of the valve record, the general form of their category, their mathematical form for the valve, their category and where each is answered. X_{next} is the outcome of the next test; K_f the number of failures in M future tests; $R_{\min}$ the minimum reliability, $p_{\max}$ the largest acceptable failure probability and κ the required assurance.

Engineering question	General form	Mathematical form	Category	Answered in
What is the failure probability?	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$	$f(p \mid k, n)$	estimation	5.3.1
What is the expected reliability?	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$	$\mathbb{E}[1 - p \mid k, n]$	estimation	5.3.1
How uncertain is the estimate?	$\mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y})$, quantiles	$\mathrm{Var}(p \mid k, n)$, credible interval	estimation	5.3.1
Is reliability at least 95 %?	$\Pr(f(\mathbf{z}) \ge f^* \mid \mathbf{y}) \ge \kappa$	$\Pr(p \le 0.05 \mid k, n)$; rule $\Pr(R \ge R_{\min} \mid k, n) \ge \kappa$	prediction under a threshold	5.3.2
Will the next demand succeed?	$\int p(y_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	$\Pr(X_{\text{next}} = 0 \mid k, n)$	prediction under a threshold	5.3.3
Will ten future demands all succeed?	$\int p(y_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	$\Pr(K_f = 0 \mid k, n)$, $M = 10$	prediction under a threshold	5.3.3
How many tests establish a requirement?	preposterior $\sum_j \Pr(K_n = j \mid \mathbf{y})\, \mathbf{1}[\text{rule met}]$	smallest n with $\Pr(p \le p_{\max} \mid 0, n) \ge \kappa$	design	5.3.4

5.1.6 The evidence as rows (step 2)

Step 2 attaches the evidence to the graph as rows on the state: a count row for the record of k failures in n demands, and a prior row for the engineering knowledge held before the tests. Neither row is Gaussian; each enters the objective as its negative log-probability term. What a row contributes to the posterior is its slope and its curvature, written with primes, $\mathcal{L}_k' = \mathrm{d}\mathcal{L}_k/\mathrm{d}\xi$ and $\mathcal{L}_k'' = \mathrm{d}^2\mathcal{L}_k/\mathrm{d}\xi^2$; each row below is given by its term, its slope and its curvature.

The modeling choices. Three choices make the two rows, and each buys something and can fail. The fifty tests are treated as comparable, one failure probability shared by all of them: this is what lets the record be summarized by k and n alone, and it fails if the stand conditions drifted during the campaign or if the two failures share a cause. The outcomes are taken as independent given p: this makes the count binomial, and it fails if one failure makes the next more likely. The earlier experience is carried as the prior $\mathrm{Beta}(2, 18)$, two failures in twenty tests: this gives the posterior a closed form and lets the experience weigh as much as twenty tests, and it misleads if that experience came from valves unlike this one, which Section 5.4.3 measures and Section 5.6.5 discusses.

The binomial likelihood. A single functional test has two possible outcomes. The outcome X_i of test i (Figure 5.1) is the random variable

$$X_i = \begin{cases} 1, & \text{if test } i \text{ fails,} \\ 0, & \text{if test } i \text{ succeeds.} \end{cases}$$

If the failure probability is p, then $\Pr(X_i = 1 \mid p) = p$ and $\Pr(X_i = 0 \mid p) = 1 - p$. The two expressions combine into one formula in the observed value $x_i \in \{0, 1\}$,

$$\Pr(X_i = x_i \mid p) = p^{x_i} \, (1 - p)^{1 - x_i}, \tag{5.5}$$

the Bernoulli probability mass function. The formula is verified on both outcomes. For a failure, $x_i = 1$ and $p^1(1-p)^0 = p$; for a success, $x_i = 0$ and $p^0(1-p)^1 = 1 - p$. Equation (5.5) therefore represents both possible outcomes, and a whole test record can be written with one expression instead of separate formulas for successes and failures. That is the foundation of the likelihood.

From Bernoulli trials to the binomial distribution. Suppose n tests are conducted with results $x_1, x_2, \ldots, x_n$, and assume that the outcomes are independent *conditional on the common failure probability* p. The joint probability of the record is then the product of the single-test probabilities:

$$\Pr(x_1, \ldots, x_n \mid p) = \prod_{i=1}^{n} \Pr(x_i \mid p).$$

Substituting equation (5.5),

$$\Pr(x_1, \ldots, x_n \mid p) = \prod_{i=1}^{n} p^{x_i} \, (1 - p)^{1 - x_i}.$$

A product of powers of the same base adds the exponents:

$$\Pr(x_1, \ldots, x_n \mid p) = p^{\sum_i x_i} \, (1 - p)^{\sum_i (1 - x_i)}.$$

Define the number of failures $k = \sum_i x_i$. Since there are n observations, $\sum_i (1 - x_i) = n - k$, and

$$\Pr(x_1, \ldots, x_n \mid p) = p^k \, (1 - p)^{n - k}. \tag{5.6}$$

This is the probability of one particular ordered sequence of successes and failures. It depends on the sequence only through k: any two sequences with the same number of failures are equally probable.

The number of failures alone. An engineering record often holds only the count k, not the sequence. The probability of the count is the sum of equation (5.6) over all the sequences with k failures, and since each has the same probability, it is that probability times the number of such sequences. The number follows from a counting argument. Choose the positions of the k failures among the n tests one after another: the first can be placed in n ways, the second in $n - 1$, down to $n - k + 1$ for the last, which gives $n(n-1)\cdots(n-k+1) = n!/(n-k)!$ ordered choices. Each set of k positions has been counted once for every order in which its positions could have been chosen, that is $k!$ times, so the number of distinct sets is

$$\binom{n}{k} = \frac{n!}{k! \, (n - k)!}.$$

Therefore the number of failures K has the probability

$$\Pr(K = k \mid p, n) = \binom{n}{k} p^k \, (1 - p)^{n - k}, \qquad K \mid p, n \sim \text{Binomial}(n, p), \tag{5.7}$$

the binomial distribution.

The likelihood. Read as a function of p with the record held fixed, equation (5.7) is the likelihood $L(p)$ of the record: it states which values of the failure probability make the observed record more probable. It is not a posterior distribution over p. It does not integrate to one over p, and it carries no prior information; both are added in Section 5.2.1.

The count row. The count row is the term the record puts in the objective. A Gaussian reading would enter the objective as its residual divided by its standard deviation. The count row is not Gaussian, so it enters as its negative log-probability term, the negative logarithm of the likelihood (5.7), a function of the state through $p = p(\xi)$ of equation (5.3):

$$\mathcal{L}_c(\xi) = -\ln L(p) = -\ln \binom{n}{k} - k \ln p - (n-k)\ln(1-p), \qquad p = p(\xi). \tag{5.8}$$

The binomial coefficient is a constant that drops out of every derivative. Written instead with one row per test, each test contributing equation (5.5), the record's term is

$$-\sum_{i=1}^{n} \ln \Pr(X_i = x_i \mid p) = -k \ln p - (n-k)\ln(1-p), \tag{5.9}$$

by equation (5.6): the n test rows sum to the count row up to its constant, and their slopes and curvatures sum to its slope and curvature. A Gaussian observation row is not equivalent to this Bernoulli and binomial term.

The slope and the curvature of the count row. The chain rule with $\mathrm{d}p/\mathrm{d}\xi = p(1-p)$ of equation (5.3) gives the slope, $\mathcal{L}_c' = \left(-k/p + (n-k)/(1-p)\right) p(1-p) = -k(1-p) + (n-k)p$, and the slope's own derivative gives the curvature:

$$\mathcal{L}_c'(\xi) = np - k, \qquad \mathcal{L}_c''(\xi) = np(1-p). \tag{5.10}$$

The slope is the expected count np minus the observed count k: the row pulls ξ down where the model expects more failures than were seen and up where it expects fewer. The curvature is the variance of the count under the binomial model (5.7), the binomial Fisher information. It does not contain the observed k, and that is a property of the logit, the canonical parameter of the binomial distribution, the quantity in which its log-probability is linear: in another coordinate, p itself among them, the curvature carries an extra term proportional to $np - k$.

The maximum-likelihood estimate. Before the prior is added, the maximum-likelihood estimate $\hat{p}_{\mathrm{MLE}} = \arg\max_p L(p)$ is the minimum of the count row alone: the logarithm is decreasing in L, so the minimizer of $\mathcal{L}_c$ is the maximizer of L. Its slope $\mathcal{L}_c' = np - k$ vanishes where the expected count equals the observed one, at

$$\hat{p}_{\mathrm{MLE}} = \frac{k}{n}, \qquad \hat{\xi}_{\mathrm{MLE}} = \ln \frac{k}{n-k}, \tag{5.11}$$

the observed failure fraction, reached without iteration. The curvature there, $n\hat{p}(1-\hat{p})$, is the information of the record and sets the spread of the likelihood around its maximum: a width $1/\sqrt{np(1-p)}$ in ξ and, since a small change of ξ moves p by $p(1-p)$ times as much, a width $p(1-p)/\sqrt{np(1-p)} = \sqrt{p(1-p)/n}$ in p. The information grows in proportion to the number of tests and the width shrinks as $1/\sqrt{n}$: a failure fraction based on fifty tests is different evidence from the same fraction based on fifty thousand. The point estimate alone does not say how uncertain the failure probability is. When $k = 0$ the MLE is $\hat{p}_{\mathrm{MLE}} = 0$ on the boundary, where the logit is $-\infty$ and the count row alone has no minimum, an estimate no engineer would report as the failure probability of a safety valve. A prior, and the posterior it yields, represents the uncertainty that the point estimate hides.

The prior row. The unknown failure probability satisfies $0 \le p \le 1$, so its prior must be a distribution on this interval. A natural choice is the Beta distribution,

$$p \sim \mathrm{Beta}(a, b), \qquad f(p) = \frac{1}{B(a,b)}\, p^{a-1}\,(1-p)^{b-1}, \qquad 0 < p < 1,\ a, b > 0, \tag{5.12}$$

where a and b are the two shape parameters, often written α and β elsewhere; this book writes them a and b, keeping β for the Weibull shape. The normalizing constant is the Beta function,

$$B(a,b) = \int_0^1 p^{a-1}\,(1-p)^{b-1}\,\mathrm{d}p = \frac{\Gamma(a)\,\Gamma(b)}{\Gamma(a+b)}, \tag{5.13}$$

where $\Gamma(c) = \int_0^\infty x^{c-1}\mathrm{e}^{-x}\,\mathrm{d}x$ is the gamma function, with $\Gamma(c+1) = c\,\Gamma(c)$ and $\Gamma(m) = (m-1)!$ for a positive integer m. The integral defines B so that the density integrates to one; the expression in gamma functions is a standard result. For integer shapes it reduces to factorials, $1/B(a,b) = (a+b-1)!/((a-1)!\,(b-1)!)$.

Mean and variance. Every moment of the Beta distribution is a ratio of Beta functions, because multiplying the density by p raises the first exponent by one:

$$\mathbb{E}[p] = \frac{1}{B(a,b)} \int_0^1 p^{a}\,(1-p)^{b-1}\,\mathrm{d}p = \frac{B(a+1,b)}{B(a,b)} = \frac{\Gamma(a+1)\,\Gamma(a+b)}{\Gamma(a)\,\Gamma(a+b+1)} = \frac{a}{a+b},$$

using equation (5.13) and $\Gamma(c+1) = c\,\Gamma(c)$ twice. In the same way $\mathbb{E}[p^2] = B(a+2,b)/B(a,b) = a(a+1)/((a+b)(a+b+1))$, and $\mathrm{Var}(p) = \mathbb{E}[p^2] - \mathbb{E}[p]^2$ simplifies to

$$\mathbb{E}[p] = \frac{a}{a+b}, \qquad \mathrm{Var}(p) = \frac{a\,b}{(a+b)^2\,(a+b+1)}. \tag{5.14}$$

The ratio $a : b$ sets the mean, and the sum $a + b$ sets the concentration: with the ratio held fixed, a larger sum gives a smaller variance, roughly in proportion to $1/(a+b)$.

Interpretation of a and b. In the posterior of Section 5.2.1, the observed failures add to a and the observed successes add to b. Algebraically, therefore, a and b act like prior failure and success counts, and $a + b$ like a prior number of tests. The pseudo-count reading is convenient for judging how strong a prior is, but it is a feature of the algebra: a prior need not come from literal historical tests, and a Beta prior with a noninteger parameter corresponds to no count at all.

The prior as a row on the logit. The prior row is the negative logarithm of the density of the state entry ξ, which is the density of p times the change-of-variable factor $\mathrm{d}p/\mathrm{d}\xi = p(1-p)$ of equation (5.3), by equation (4.51) with the roles of p and ξ exchanged,

$$f_\xi(\xi) = f(p)\,p(1-p) = \frac{p^a\,(1-p)^b}{B(a,b)}, \qquad p = p(\xi),$$

so that, with the constant dropped, the prior row's term is

$$\mathcal{L}_0(\xi) = -a\ln p - b\ln(1-p), \qquad p = p(\xi). \tag{5.15}$$

This has the binomial form of a failures in $a+b$ trials, so its slope and curvature are equation (5.10) with (k, n) replaced by $(a, a+b)$:

$$\mathcal{L}_0'(\xi) = (a+b)\,p - a, \qquad \mathcal{L}_0''(\xi) = (a+b)\,p(1-p).$$

In the logit the pseudo-count reading is therefore exact: the prior row is the row of a pseudo-failures in $a+b$ pseudo-trials. The count row, by contrast, is a probability of the record, not a density over the state, and takes no change-of-variable factor.

The two priors. Previous engineering experience suggests a failure probability around 10 %, with substantial uncertainty, and the prior chosen is the Beta density (5.12)

$$p \sim \text{Beta}(2, 18), \tag{5.16}$$

which, read as pseudo-counts, weighs as much as two failures in twenty tests. The prior is not a statement that the failure probability is exactly 10 %; it represents uncertainty around that value. Prior choices should be documented and justified, especially when they materially influence a safety-related decision; Section 5.4.3 tests how much this one does.

The seventh question concerns a new valve rather than the tested one, and it brings its own data. The new valve is required to have $R \geq 0.99$, equivalently $p \leq 0.01$, with at least 95 % posterior probability that the requirement is satisfied, and the engineer asks how many tests would be needed if no failures are observed. Its prior is uniform over the failure probability,

$$p \sim \text{Beta}(1, 1), \qquad f(p) = 1 \quad (0 < p < 1),$$

which assigns equal density to every value between zero and one. A uniform prior is not appropriate for every engineering assessment, but it gives a clear reference case; Section 5.4.3 returns to the choice.

5.2 The posterior of the state (step 3)

All seven questions are read from one posterior: the posterior of the state that the prior row and the count row of step 2 give, exactly a Beta density of p and approximately the Laplace normal of the logit. This section derives it and computes it for the valve; Section 5.3 reads every answer from it.

5.2.1 The posterior from the two rows

The two rows' terms, equations (5.15) and (5.8) with the constants dropped, add to the objective, the negative log posterior up to a constant,

$$\mathcal{L}(\xi) = \mathcal{L}_0(\xi) + \mathcal{L}_c(\xi) = -a_n \ln p - b_n \ln(1-p), \qquad p = p(\xi), \qquad a_n = a + k, \quad b_n = b + n - k, \tag{5.17}$$

and their slopes and curvatures add with them:

$$\mathcal{L}'(\xi) = \underbrace{(a+b)\,p - a}_{\text{prior}} + \underbrace{np - k}_{\text{count}} = (a_n + b_n)\,p - a_n, \qquad \mathcal{L}''(\xi) = (a_n + b_n)\,p(1-p). \tag{5.18}$$

Both rows are binomial terms in the logit, so $\mathcal{L}''$ is the exact curvature at every ξ. Newton's method replaces the objective at the current ξ_j by the parabola with the same slope and curvature and moves to its minimum,

$$\Delta\xi = -\frac{\mathcal{L}'(\xi_j)}{\mathcal{L}''(\xi_j)}, \qquad \xi_{j+1} = \xi_j + \Delta\xi. \tag{5.19}$$

The iteration starts from the record's own estimate, $\xi_0 = \hat{\xi}_{\text{MLE}}$ of equation (5.11), where the count row's slope is zero and the prior row alone sets the slope.

The MAP state and the Laplace variance. The slope $(a_n + b_n)p - a_n$ vanishes at $p = a_n/(a_n + b_n)$, so the minimum of the objective is

$$\hat{\xi} = \ln \frac{a_n}{b_n}, \qquad p(\hat{\xi}) = \frac{a_n}{a_n + b_n}. \tag{5.20}$$

There neither row's slope is zero: the prior row pulls ξ up toward the prior mean and the count row pulls it down toward the observed fraction, with equal and opposite slopes. The curvature at the minimum is the information Λ, and its reciprocal is the Laplace variance of the logit,

$$\Lambda = \mathcal{L}''(\hat{\xi}) = \underbrace{(a+b)\,\hat{p}(1-\hat{p})}_{\text{prior}} + \underbrace{n\,\hat{p}(1-\hat{p})}_{\text{count}}, \qquad \sigma_\xi^2 = \frac{1}{\Lambda} = \frac{1}{(a_n+b_n)\,\hat{p}(1-\hat{p})}, \tag{5.21}$$

with $\hat{p} = p(\hat{\xi})$: the prior row's information plus the record's Fisher information. Since $\hat{p}(1-\hat{p}) = a_n b_n/(a_n+b_n)^2$ at the minimum (5.20), the variance is also $\sigma_\xi^2 = (a_n+b_n)/(a_n b_n)$. The fit returns the pair $(\hat{\xi}, \sigma_\xi)$.

The weight of the prior. Which row carries the state is read in the two terms of Λ: the prior row's share of the information is

$$\pi_0 = \frac{\mathcal{L}_0''}{\mathcal{L}''} = \frac{(a+b)\,p(1-p)}{(a+b+n)\,p(1-p)} = \frac{a+b}{a+b+n}, \tag{5.22}$$

the same at every ξ, since both curvatures carry the factor $p(1-p)$. As tests accumulate, $\pi_0 \to 0$ and the record carries the state. The same share splits the failure probability at the minimum exactly,

$$\frac{a+k}{a+b+n} = \underbrace{\frac{a+b}{a+b+n}}_{\pi_0}\,\frac{a}{a+b} + \underbrace{\frac{n}{a+b+n}}_{1-\pi_0}\,\frac{k}{n}, \tag{5.23}$$

a weighted average of the prior mean and the MLE.

The exact posterior. This problem also admits a closed form. Bayes' theorem, $f(p \mid k, n) \propto \Pr(K = k \mid p, n)\, f(p)$, with the likelihood (5.7) and the prior (5.12), gives, once the factors $\binom{n}{k}$ and $1/B(a,b)$ that do not depend on p are absorbed, the kernel $p^{a+k-1}(1-p)^{b+n-k-1}$ of a Beta density, whose normalizing constant is fixed by equation (5.13):

$$p \mid k, n \sim \text{Beta}(a_n, b_n) = \text{Beta}(a+k,\ b+n-k), \tag{5.24}$$

$$f(p \mid k, n) = \frac{p^{a_n-1}\,(1-p)^{b_n-1}}{B(a_n, b_n)}. \tag{5.25}$$

This is the *Beta–binomial conjugate update*: the failures add to the first parameter and the successes to the second, and it is the exact posterior under the model assumptions. The fit above is its Laplace approximation: the objective (5.17) is the negative logarithm of this density written as a density of ξ, and the minimum (5.20) lies at its mean, the first of equations (5.14) with the posterior parameters. The same integral gives the denominator of Bayes' theorem, the evidence of the model, which compares two priors or two models fitted to the same record and answers none of the questions here. The record enters only through the sufficient statistic (k, n), so campaigns can be added one at a time, in any order and grouping, by two additions each, which is the statement that rows add.

5.2.2 The posterior of the valve

The valve's inputs are the prior parameters $a = 2$ and $b = 18$ and the record of $k = 2$ failures in $n = 50$ tests.

The prior. The engineering prior Beta(2, 18) of equation (5.16) has the density $f(p) = 342\,p\,(1-p)^{17}$, since $1/B(2,18) = 19!/(1!\,17!) = 19 \times 18 = 342$. By equation (5.14) its mean is $2/20 = 0.10$ and its standard deviation is

$$\sqrt{\mathrm{Var}(p)} = \sqrt{\frac{2 \times 18}{20^2 \times 21}} = \sqrt{0.0042857} = 0.0655,$$

about 6.55 percentage points. In the logit its row is that of two pseudo-failures in twenty pseudo-trials, equation (5.15).

The record. For fifty tests with two failures, equation (5.7) gives $\Pr(K = 2 \mid p, n = 50) = \binom{50}{2} p^2(1-p)^{48}$, and since $\binom{50}{2} = 50 \times 49/2 = 1225$, the likelihood of the record is

$$L(p) = 1225\,p^2\,(1-p)^{48}. \tag{5.26}$$

Figure 5.2(a) draws equation (5.26). The horizontal axis is the failure probability per demand and the vertical axis the probability of the observed record, two failures in fifty, at that value of p. The curve has its maximum of 0.276 at the MLE of equation (5.11), $2/50 = 0.04$. The width of the likelihood around its maximum is $\sqrt{0.04 \times 0.96/50} = 0.0277$, about 0.028. Figure 5.2(b) shows how that width depends on the number of tests: the horizontal axis is again p, the vertical axis the likelihood divided by its maximum, so that both curves peak at one. The solid curve is the valve record; the dashed curve is an illustrative record of twenty failures in five hundred tests, with the same MLE. The larger record gives a likelihood narrower by the factor $\sqrt{10}$, a width of 0.0088; fifty thousand tests would narrow it by $\sqrt{1000}$, to a width of 0.0009.

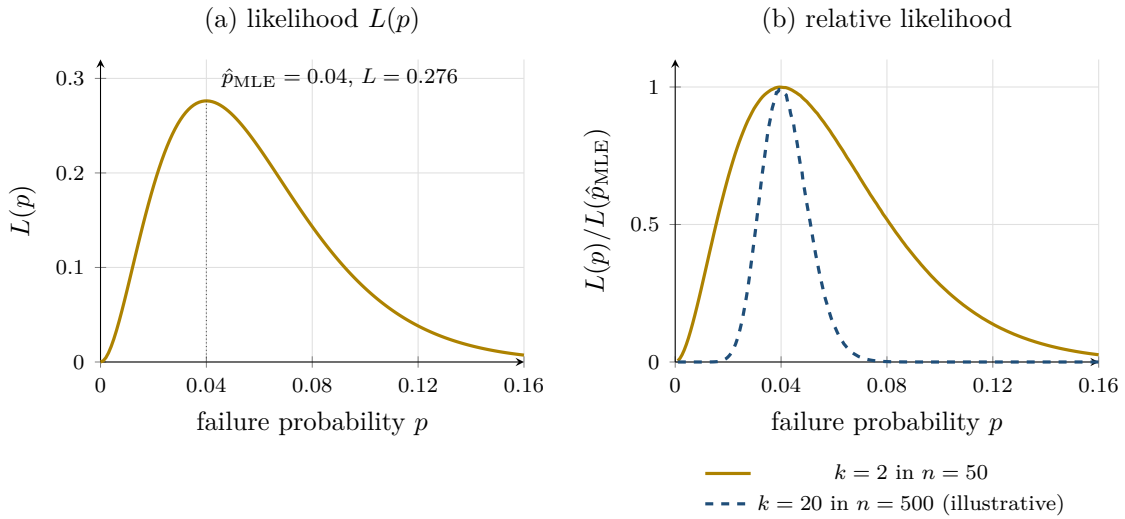


Figure 5.2: The binomial likelihood of the valve record. (a) Equation (5.26), $L(p) = 1225\,p^2(1-p)^{48}$, maximal at $\hat{p}_{\mathrm{MLE}} = 0.04$. (b) Likelihood relative to its maximum, $(p/0.04)^k((1-p)/0.96)^{n-k}$, for the record and for an illustrative record ten times larger with the same failure fraction; the information of the count row (Section 5.1.6) gives widths $\sqrt{p(1-p)/n} = 0.028$ and 0.0088.

The Newton iterations. With $a = 2$, $b = 18$, $k = 2$ and $n = 50$, the slopes of equation (5.18) are $20p - 2$ for the prior row and $50p - 2$ for the count row, and the curvature is $70\,p(1-p)$. At the starting point $\xi_0 = \hat{\xi}_{\mathrm{MLE}} = -3.1781$, where $p = 0.04$, the count row's slope is zero, since ξ_0 is its own minimum, and the prior row's slope -1.2000 alone pulls the state up toward the prior. Table 5.2 runs the step (5.19) from that point. Each row is an iteration; the columns

give the state, the failure probability, the two rows' slopes, their sum, the curvature and the step. The first step overshoots the minimum and the next two return to it, at $\xi = -2.80336$, where the two slopes are -0.8571 and $+0.8571$ and cancel: $\hat{\xi} = \ln(4/66) = -2.8034$ at $p = 4/70$, equation (5.20).

Table 5.2: Newton iterations, equation (5.19), for the valve in the logit ξ: the slopes of the two rows and the curvature of equation (5.18), from $\xi_0 = \ln(0.04/0.96)$.

Iteration j	ξ_j	p_j	Prior slope $\mathcal{L}_0'$	Count slope $\mathcal{L}_c'$	Slope $\mathcal{L}'$	Curvature $\mathcal{L}''$	$\Delta\xi$
0	-3.17805	0.04000	-1.2000	0.0000	-1.2000	2.6880	0.44642
1	-2.73163	0.06113	-0.7773	1.0566	0.2793	4.0177	-0.06952
2	-2.80114	0.05726	-0.8548	0.8631	0.0084	3.7788	-0.00222
3	-2.80336	0.05714	-0.8571	0.8571	0.0000	3.7714	-2.2×10^{-6}

The Laplace variance. At $\hat{\xi}$ the two rows' curvatures are their information, 1.0776 for the prior and 2.6939 for the count, and equation (5.21) gives

$$\Lambda = 1.0776 + 2.6939 = 3.7714, \qquad \sigma_\xi^2 = \frac{1}{3.7714} = 0.26515, \qquad \sigma_\xi = 0.5149.$$

The fit returns $\hat{\xi} = -2.8034$ with $\sigma_\xi = 0.5149$. The prior's share of the information, equation (5.22), is $\pi_0 = 1.0776/3.7714 = 0.2857 = 20/70$, and the weighted average of equation (5.23) is $0.2857 \times 0.10 + 0.7143 \times 0.04 = 0.05714$.

The exact posterior. With the prior Beta$(2, 18)$ and the record $k = 2$, $n = 50$, equation (5.24) gives $a_n = 2 + 2 = 4$ and $b_n = 18 + 48 = 66$:

$$p \mid k = 2,\, n = 50 \sim \text{Beta}(4, 66), \qquad f(p \mid k, n) = 3\,458\,004\; p^3\,(1-p)^{65}, \tag{5.27}$$

since $1/B(4, 66) = 69!/(3!\,65!) = 69 \times 68 \times 67 \times 66/6 = 3\,458\,004$. The two failures and forty-eight successes have moved the prior Beta$(2, 18)$, of mean 10 %, to the posterior Beta$(4, 66)$, of mean 5.71 %.

Figure 5.3 draws the three densities that enter the update. The horizontal axis is the failure probability per demand, the vertical axis probability density, not the probability of an exact value. The blue curve is the prior Beta$(2, 18)$; the dashed gold curve is the likelihood of equation (5.26) rescaled to unit area, which is the Beta$(3, 49)$ density and the posterior that a uniform prior would give; the green curve is the posterior Beta$(4, 66)$. The dotted lines mark the MLE, the posterior mean and the prior mean. The observations shift the distribution toward lower failure probabilities and concentrate it: the posterior lies between the prior and the likelihood and is narrower than either, its peak more than twice the height of the prior's.

The density of the failure probability. The state is the logit, not p, so the Laplace approximation makes the logit normal, and the density of p is obtained from it in three steps by the change of variables of Section 4.3.3. First, the Laplace posterior of the state is $\xi \mid \mathbf{y} \approx \mathcal{N}(\hat{\xi}, \sigma_\xi^2)$, whose density, written with the standard normal density $\varphi(z) = \mathrm{e}^{-z^2/2}/\sqrt{2\pi}$, is

$$p_\xi(\xi \mid \mathbf{y}) \approx \frac{1}{\sigma_\xi}\, \varphi\!\left(\frac{\xi - \hat{\xi}}{\sigma_\xi}\right),$$

the factor $1/\sigma_\xi$ being the normalization of a normal density with standard deviation σ_ξ. Second, a density is carried from ξ to p by equation (4.51): it is evaluated at the logit $\xi(p) = \ln(p/(1-p))$

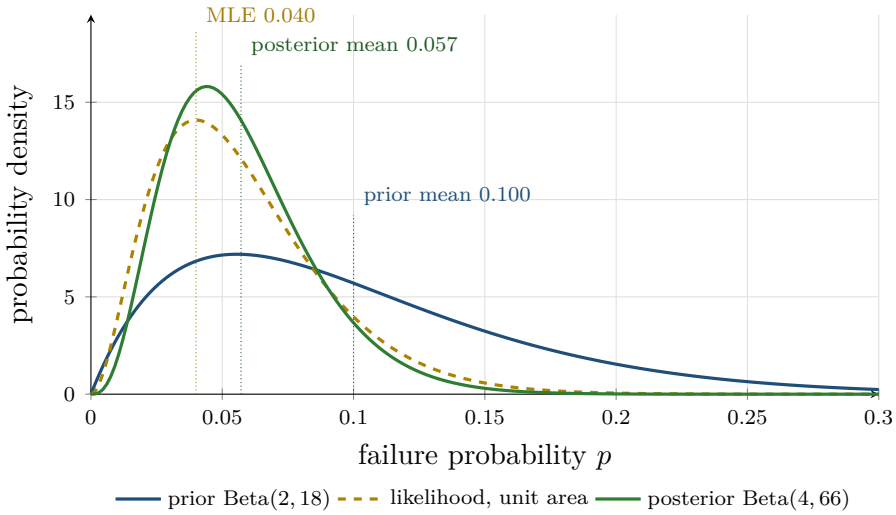


Figure 5.3: Prior, normalized likelihood and posterior of the valve's failure probability per demand, from equation (5.12): prior $342\,p(1-p)^{17}$, likelihood rescaled to unit area $62\,475\,p^2(1-p)^{48}$ (the Beta(3, 49) density), posterior $3\,458\,004\,p^3(1-p)^{65}$. The vertical axis is density, not probability at an exact value.

of each p and multiplied by the Jacobian factor $|\mathrm{d}\xi/\mathrm{d}p|$. Third, that factor is

$$\frac{\mathrm{d}\xi}{\mathrm{d}p} = \frac{\mathrm{d}}{\mathrm{d}p}(\ln p - \ln(1-p)) = \frac{1}{p} + \frac{1}{1-p} = \frac{1}{p(1-p)} > 0,$$

large near $p = 0$ and $p = 1$, where a small change of p is a large change of the logit. Together,

$$f(p \mid \mathbf{y}) \approx \frac{1}{\sigma_\xi\, p(1-p)}\, \varphi\!\left(\frac{\xi(p) - \hat{\xi}}{\sigma_\xi}\right), \qquad \xi(p) = \ln\frac{p}{1-p}, \tag{5.28}$$

a density on $(0, 1)$, the *logit-normal* density. Its mode is not $\hat{p}$, because the factor $1/(p(1-p))$ tilts the curve. The factor is not optional: for the valve, the normal density of the logit evaluated at $\xi(p)$ without it integrates over $(0, 1)$ to 0.0586, not to one.

The delta method of Section 4.3.1 gives a second approximation to the same density, and comparing the two shows what the Jacobian factor buys. Linearizing $p(\xi)$ at $\hat{\xi}$, the derivative is $g = \mathrm{d}p/\mathrm{d}\xi = \hat{p}(1-\hat{p})$, the reciprocal of the factor at $\hat{p}$, and the delta method reports $p \mid \mathbf{y} \approx \mathcal{N}(\hat{p},\ g^2\sigma_\xi^2)$. For the valve, $\hat{p} = 0.0571$, $\sigma_\xi = 0.5149$ and $g = 0.0571 \times 0.9429 = 0.0539$, so $\mathrm{sd}(p \mid \mathbf{y}) \approx g\sigma_\xi = 0.0277$, within 0.7 % of the exact 0.0275: on the spread the two routes agree. On the shape they do not. The delta Gaussian is symmetric about $\hat{p}$, with the central 95 % interval [0.0028, 0.1115], and it puts $\Pr(p < 0 \mid \mathbf{y}) = 0.020$ on failure probabilities that cannot exist. The logit-normal density (5.28) keeps the map exact, stays inside $(0, 1)$ and is skewed to the right like the exact Beta density, with the interval [0.0216, 0.1426] against the exact [0.0160, 0.1218]. The chapter therefore reports the delta method's standard deviation and reads intervals and threshold probabilities from the density, which is the rule of Section 4.3.3.

The pair $(\hat{\xi}, \sigma_\xi) = (-2.8034,\ 0.5149)$ is what step 4 uses: every answer of the chapter is computed from it. Because this problem also admits the exact posterior, each answer has an exact value beside it, a function of the two numbers $(a_n, b_n) = (4, 66)$, and the gap between the two is reported with it.

5.3 The answers (step 4)

The questions of Table 5.1 fall into three categories: estimation (the first three), prediction under a threshold (the fourth, and the fifth and sixth, which predict future outcomes) and design

(the seventh). Each is answered from the posterior of Section 5.2: its engineering question, the general form of its category, the form it takes for the valve, the quantities of the posterior it uses, and the valve's numbers. The questions name the requirement $R_{\min} = 0.95$, the $M = 10$ future demands, and the new valve's $p_{\max} = 0.01$ at assurance $\kappa = 0.95$ under the uniform prior. Every answer is computed from the pair of step 3, and that is the value the chapter reports; the exact Beta posterior gives the value set beside it, with the gap.

5.3.1 The failure probability, the expected reliability and the spread

Engineering question. What is the failure probability? What is the expected reliability? How uncertain is the estimate? The three ask for the failure probability per demand of the valve given its record, equation (5.2), for the mean of its complement, and for the spread of the estimate.

General form. An estimation question about a quantity $f(\mathbf{z})$ derived from the state asks for its posterior distribution, the probability that the quantity lies at or below each value f^*,

$$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y}) = \int \mathbf{1}\big[f(\mathbf{z}) \le f^*\big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \tag{5.29}$$

summarized by a location, the mode or the mean,

$$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}] = \int f(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \tag{5.30}$$

and by a spread, the posterior variance $\mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y}) = \mathbb{E}[(f(\mathbf{z}) - \mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}])^2 \mid \mathbf{y}]$ and the central 95 % credible interval $[f_L, f_U]$, which leaves 2.5 % of posterior probability in each tail of the distribution (5.29):

$$\Pr(f(\mathbf{z}) < f_L \mid \mathbf{y}) = 0.025, \qquad \Pr(f(\mathbf{z}) > f_U \mid \mathbf{y}) = 0.025, \qquad \Pr(f_L \le f(\mathbf{z}) \le f_U \mid \mathbf{y}) = 0.95. \tag{5.31}$$

Specialized form. The three questions are equations (5.29), (5.30) and (5.31) with the state $\mathbf{z}$ the one number ξ, the derived quantity $f(\mathbf{z}) = p(\xi) = 1/(1 + \mathrm{e}^{-\xi})$ of equation (5.3) (and $f(\mathbf{z}) = R = 1 - p(\xi)$ for the expected reliability), and $p(\mathbf{z} \mid \mathbf{y})$ the posterior of step 3, the Laplace normal $\xi \mid \mathbf{y} \approx \mathcal{N}(\hat{\xi}, \sigma_\xi^2)$ of equations (5.20) and (5.21).

For the location, the logistic map is increasing, so equation (5.29) at $f^* = \hat{p}$ is the normal probability $\Pr(\xi \le \hat{\xi} \mid \mathbf{y}) = \frac{1}{2}$: the failure probability at the MAP state, $\hat{p}$, is the median of p under the Laplace posterior, and it is the point estimate the chapter reports. Its density in p is the logit-normal density (5.28). Beside it stand two estimates that do not use the state. The maximum-likelihood estimate k/n uses the record alone. The exact posterior mean is $a_n/(a_n + b_n)$, the same number as $\hat{p}$, because $\hat{\xi} = \ln(a_n/b_n)$. The exact posterior mode, the most probable value under the density of p, is a property of the Beta density itself: the derivative of $\ln f(p \mid \mathbf{y}) = (a_n - 1)\ln p + (b_n - 1)\ln(1 - p) - \ln B(a_n, b_n)$ vanishes at

$$\hat{p}_{\mathrm{mode}} = \frac{a_n - 1}{a_n + b_n - 2}, \qquad a_n, b_n > 1. \tag{5.32}$$

The condition ensures that the stationary point is an interior maximum; when $a_n \le 1$ the density is largest at $p = 0$. The mode and the MAP state differ because the MAP state depends on the entry chosen as the state: the density of ξ carries the change-of-variable factor $p(1 - p)$ that the density of p does not, and the factor moves its maximum from the mode of p to the mean. The three estimates are not supposed to be identical; they summarize the same inference in different ways, and the mean exceeds the mode whenever the posterior is skewed to the right.

For a future binary outcome the posterior mean is the important one, because it equals the predictive probability of failure on the next comparable demand (the next question, below).

For the expected reliability, equation (5.30) with $f(\mathbf{z}) = 1 - p(\xi)$ is an integral over the one entry against the Laplace normal, which a Gauss–Hermite rule of Q nodes t_q and weights w_q evaluates as a weighted sum,

$$\mathbb{E}[h(\xi) \mid \mathbf{y}] \approx \int h(\xi)\, \mathcal{N}(\xi;\, \hat{\xi}, \sigma_\xi^2)\, \mathrm{d}\xi \approx \frac{1}{\sqrt{\pi}} \sum_{q=1}^{Q} w_q\, h(\hat{\xi} + \sqrt{2}\, \sigma_\xi\, t_q), \tag{5.33}$$

here with $h(\xi) = 1 - p(\xi)$: the nodes $\xi_q = \hat{\xi} + \sqrt{2}\, \sigma_\xi t_q$ are values of the state spread over its posterior, and the sum weighs the reliability at each. It does not equal $1 - p(\hat{\xi})$, because $1 - p(\xi)$ is not linear in ξ. Exactly, the expectation is linear, $\mathbb{E}[R \mid \mathbf{y}] = 1 - \mathbb{E}[p \mid \mathbf{y}]$, and with the mean $a/(a+b)$ of a Beta(a, b) density (derived with the prior row, Section 5.1.6) applied to the posterior parameters,

$$\mathbb{E}[p \mid \mathbf{y}] = \frac{a_n}{a_n + b_n}, \qquad \mathbb{E}[R \mid \mathbf{y}] = 1 - \mathbb{E}[p \mid \mathbf{y}] = \frac{b_n}{a_n + b_n}. \tag{5.34}$$

For the spread, $f(\mathbf{z}) = p(\xi)$ again and $[f_L, f_U] = [p_L, p_U]$. The variance of p follows from the derivative of the map at the MAP state by the delta method,

$$g = \frac{\mathrm{d}p}{\mathrm{d}\xi}\Big|_{\hat{\xi}} = \hat{p}(1 - \hat{p}), \qquad \mathrm{Var}(p \mid \mathbf{y}) \approx g^2 \sigma_\xi^2 = \hat{p}^2 (1 - \hat{p})^2\, \sigma_\xi^2, \tag{5.35}$$

and the reliability $R = 1 - p$, with $g = -\hat{p}(1 - \hat{p})$, has the same spread. The interval of equation (5.31) is an interval for the state entry mapped to p: the logit is increasing, so the quantiles of ξ map to the quantiles of p,

$$p_L \approx p(\hat{\xi} - 1.96\sigma_\xi), \qquad p_U \approx p(\hat{\xi} + 1.96\sigma_\xi), \tag{5.36}$$

and only the normal tail is an approximation. The variance describes uncertainty about the failure probability, an epistemic quantity, not the variability of the next binary outcome, which is aleatory: even with p known exactly, the next test would fail with probability p. The two kinds of uncertainty combine in the predictions of the next question. Exactly, with the variance of a Beta density (Section 5.1.6) applied to the posterior parameters,

$$\mathrm{Var}(p \mid \mathbf{y}) = \frac{a_n b_n}{(a_n + b_n)^2 (a_n + b_n + 1)}, \tag{5.37}$$

and the exact endpoints are quantiles of the Beta posterior, obtained from its cumulative distribution function,

$$I_x(a, b) = \frac{1}{B(a, b)} \int_0^x t^{a-1} (1 - t)^{b-1}\, \mathrm{d}t, \tag{5.38}$$

the regularized incomplete Beta function, which every statistical library evaluates and inverts. For integer parameters it also has a closed form, a standard identity worth having for hand checks: $I_x(a, b)$ equals the probability that a Binomial$(a + b - 1, x)$ count is at least a,

$$I_x(a, b) = \sum_{j=a}^{a+b-1} \binom{a + b - 1}{j} x^j (1 - x)^{a+b-1-j}, \tag{5.39}$$

which follows by integrating equation (5.38) by parts $a - 1$ times. The endpoints of the central interval satisfy $I_{p_L}(a_n, b_n) = 0.025$ and $I_{p_U}(a_n, b_n) = 0.975$, so that $p_L = I^{-1}_{0.025}(a_n, b_n)$ and $p_U = I^{-1}_{0.975}(a_n, b_n)$, with I^{-1} the inverse of the cumulative distribution in its first argument. A central credible interval is a probability statement about the unknown conditional on the record and the assumptions; it does not establish that a component has passed a reliability requirement, which needs the explicit acceptance threshold of the next question.

Building blocks. The pair $(\hat{\xi}, \sigma_\xi) = (-2.8034,\ 0.5149)$ of Section 5.2.2; the derivative g of equation (5.35); the Gauss–Hermite nodes and weights of equation (5.33); and, for the exact values beside them, the posterior parameters $a_n = 4$, $b_n = 66$, the inverse of the incomplete Beta function and the maximum-likelihood estimate of the count row (Section 5.1.6).

The point estimates of the valve. The failure probability at the MAP state, the Laplace median, is

$$p(\hat{\xi}) = \frac{1}{1 + \mathrm{e}^{2.8034}} = 0.05714,$$

about 5.71 %. The exact posterior median is 0.05296, a gap of 0.00418. The exact posterior mean, $4/70 = 0.05714$, is a different summary that falls on the reported median, the Laplace mean being 0.0636; the exact mode, equation (5.32), is $3/68 = 0.04412$; and the MLE is 0.04. Table 5.3 collects them: each row is an estimator, the middle column its expression and the right column its value.

Table 5.3: Point estimates of the valve's failure probability per demand. Prior Beta$(2, 18)$; record $k = 2$, $n = 50$; exact posterior Beta$(4, 66)$.

Estimator	Expression	Failure probability
Laplace median (reported)	$\hat{p}$	5.71 %
Exact posterior median	$I^{-1}_{1/2}(a_n, b_n)$	5.30 %
Exact posterior median	$I^{-1}_{1/2}(a_n, b_n)$	5.30 %
Exact posterior mean	$a_n/(a_n + b_n)$	5.71 %
Exact posterior mode	$(a_n - 1)/(a_n + b_n - 2)$	4.41 %
Maximum likelihood	k/n	4.00 %

The expected reliability of the valve. Table 5.4 evaluates equation (5.33) with five nodes. Each row is a node: the first two columns give the node t_q and its weight $w_q/\sqrt{\pi}$, the third the state $\xi_q = \hat{\xi} + \sqrt{2}\,\sigma_\xi t_q$ at which the sum looks, the fourth the failure probability $p_q = p(\xi_q)$ there, and the last two the two functions the questions of this chapter average, the reliability $1 - p_q$ and the clean run $(1 - p_q)^{10}$ of the next question. The middle node is the MAP state itself; the others reach about two standard deviations either side of it.

Table 5.4: The Gauss–Hermite rule (5.33) with five nodes over the Laplace posterior $\xi \mid \mathbf{y} \approx \mathcal{N}(-2.8034,\ 0.5149^2)$ of the valve.

Node t_q	Weight $w_q/\sqrt{\pi}$	State ξ_q	p_q	$1 - p_q$	$(1 - p_q)^{10}$
−2.0202	0.01126	−4.2745	0.01373	0.98627	0.87090
−0.9586	0.22208	−3.5014	0.02927	0.97073	0.74298
0	0.53333	−2.8034	0.05714	0.94286	0.55521
0.9586	0.22208	−2.1053	0.10858	0.89142	0.31682
2.0202	0.01126	−1.3322	0.20879	0.79121	0.09614
Weighted sum				0.9364	0.5424

The weighted sum of the fifth column is

$$\mathbb{E}[R \mid \mathbf{y}] \approx 0.01126 \times 0.98627 + 0.22208 \times 0.97073 + \cdots + 0.01126 \times 0.79121 = 0.9364,$$

a mean reliability of about 93.64 %; forty nodes change no digit shown. The exact value, equation (5.34), is $66/70 = 0.94286$, 0.0065 higher: the normal in the logit gives more weight to

large failure probabilities than the Beta posterior does, and the mean of p under it is 0.0636, not 0.0571.

The posterior spread and the credible interval of the valve. By equation (5.35),

$$g = \hat{p}(1-\hat{p}) = 0.053878, \qquad g^2\sigma_\xi^2 = 0.053878^2 \times 0.26515 = 7.697 \times 10^{-4},$$

a standard deviation of 0.02774, about 2.77 percentage points, for p and, with $g = -0.053878$, for R, down from 6.55 for the prior. The exact value, equation (5.37), is $\sqrt{4 \times 66/(70^2 \times 71)} = \sqrt{0.00075884} = 0.02755$, within 0.7 %. By equation (5.36), the logit interval $-2.8034 \pm 1.96 \times 0.5149 = [-3.8126,\ -1.7941]$ maps to

$$p_L = 0.0216, \qquad p_U = 0.1426, \qquad p \in [0.0216,\ 0.1426] \text{ with } 95\,\% \text{ posterior probability}, \tag{5.40}$$

so the reliability lies in $[0.8574,\ 0.9784]$ with the same probability. Inverting $I_x(4, 66)$ at 0.025 and 0.975 gives the exact interval $[0.01602,\ 0.12185]$, and $[0.8782,\ 0.9840]$ for the reliability: the Laplace interval is shifted toward larger failure probabilities, by 0.006 at the lower end and 0.021 at the upper, the same bias as the mean. Figure 5.4(a) shows both intervals: the horizontal axis is the failure probability, the vertical axis the posterior density, the dotted blue curve the Laplace density (5.28) with its interval shaded, and the green curve the exact density with its interval between the dashed lines; panel (b) is read in Section 5.3.2.

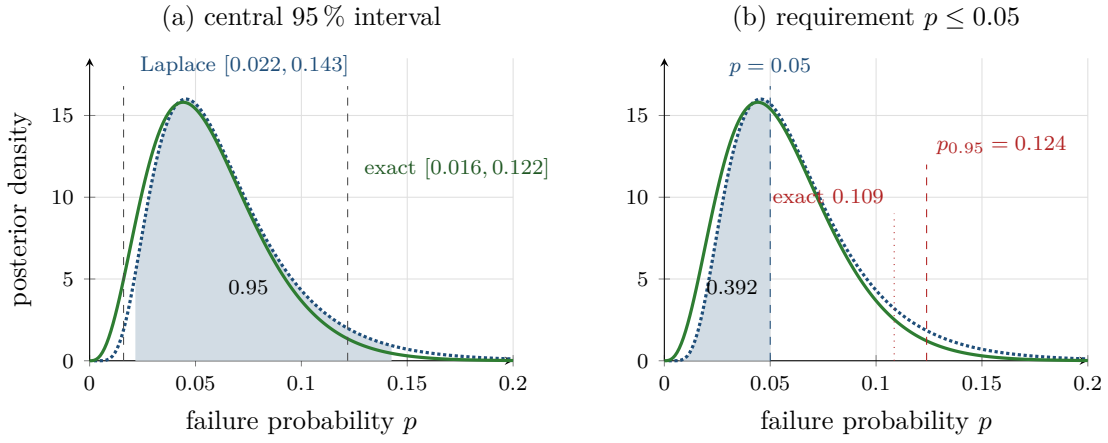


Figure 5.4: Reading the posterior of the valve. Dotted blue: the Laplace density (5.28), which the answers report; green: the exact Beta(4, 66), $3\,458\,004\,p^3(1-p)^{65}$. (a) Central 95 % credible interval, shaded for the Laplace density, 0.0216 to 0.1426 (equation (5.40)), dashed for the exact one, 0.0160 to 0.1218. (b) Posterior probability of the 95 % reliability requirement, shaded area 0.392 under the Laplace density (0.455 exact), and the one-sided 95 % upper bound $p_{0.95} = 0.124$ (0.109 exact; Section 5.3.2).

5.3.2 The requirement probability and the assurance rule

Engineering question. Is reliability at least 95 %? The manufacturer requires that the valve operate successfully on a demand with probability at least $R_{\min} = 0.95$, and a customer may demand that this hold with a stated posterior probability.

General form. A prediction under a threshold is a tail probability of a derived quantity under the posterior, compared with an assurance level κ:

$$\Pr(f(\mathbf{z}) \ge f^* \mid \mathbf{y}) = \int \mathbf{1}[f(\mathbf{z}) \ge f^*]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \qquad \text{accept if } \Pr(f(\mathbf{z}) \ge f^* \mid \mathbf{y}) \ge \kappa. \tag{5.41}$$

The same decision can be stated through a one-sided bound: for the event $f(\mathbf{z}) \geq f^*$ the bound is the $(1-\kappa)$ quantile $f_{1-\kappa}$ of the derived quantity, $\Pr(f(\mathbf{z}) \leq f_{1-\kappa} \mid \mathbf{y}) = 1-\kappa$, and the rule holds exactly when $f^* \leq f_{1-\kappa}$.

Specialized form. One choice enters here: the requirement is read per demand, as the specification states it, so the event is a tail of p alone. A specification over a mission of operating hours would turn it into a question about a failure rate, the subject of Chapter 6.

The question is equation (5.41) with $f(\mathbf{z}) = R = 1 - p(\xi)$, $f^* = R_{\min}$ and $p(\mathbf{z} \mid \mathbf{y})$ the posterior of step 3. Since $R = 1 - p$ and the logit is increasing, the event becomes an event on the one entry of the state:

$$R \geq R_{\min} \iff p \leq p_{\max} = 1 - R_{\min} \iff \xi \leq \xi_{\max} = \ln \frac{p_{\max}}{1 - p_{\max}},$$

so the integral of equation (5.41) is a normal tail in the MAP state and the variance, and the one-sided bound $f_{1-\kappa}$ is the MAP state moved by 1.6449 standard deviations and mapped to p:

$$\Pr(p \leq p_{\max} \mid \mathbf{y}) \approx \Phi\left(\frac{\xi_{\max} - \hat{\xi}}{\sigma_\xi}\right), \qquad p_{0.95} \approx p(\hat{\xi} + 1.6449\sigma_\xi), \tag{5.42}$$

with Φ the standard normal cumulative distribution; only the tail of the normal is an approximation, since the threshold maps exactly. The value beside it comes from the exact posterior: since the posterior of p is $\mathrm{Beta}(a_n, b_n)$, the tail is the incomplete Beta function of equation (5.38) at $p_{\max}$,

$$\Pr(R \geq R_{\min} \mid \mathbf{y}) = \Pr(p \leq p_{\max} \mid \mathbf{y}) = \int_0^{p_{\max}} \frac{p^{a_n - 1}(1-p)^{b_n - 1}}{B(a_n, b_n)}\, \mathrm{d}p = I_{p_{\max}}(a_n, b_n), \tag{5.43}$$

and the assurance rule of equation (5.41) reads, on either route,

$$\Pr(R \geq R_{\min} \mid \mathbf{y}) \geq \kappa, \tag{5.44}$$

a minimum reliability $R_{\min}$ together with a minimum posterior probability κ, the assurance level, supporting it; for example $\Pr(R \geq 0.90 \mid \mathbf{y}) \geq 0.95$, or the analogous rule of 95 % reliability at 95 % assurance. The one-sided bound is the 0.95 quantile of the posterior of p, the second of equations (5.42) under the Laplace posterior and, exactly,

$$\Pr(p \leq p_{0.95} \mid \mathbf{y}) = 0.95, \qquad p_{0.95} = I^{-1}_{0.95}(a_n, b_n), \qquad R_L = 1 - p_{0.95}, \tag{5.45}$$

so that R_L is the 0.05 quantile of R, the bound $f_{1-\kappa}$ above: the reliability is at least R_L with 95 % posterior probability, and the rule (5.44) with $\kappa = 0.95$ is met exactly when $R_{\min} \leq R_L$. It is a one-sided 95 % bound, not the upper end of the central 95 % interval: $p_{0.95}$ leaves 5 % of the posterior above it, whereas p_U leaves 2.5 %. The two are different summaries of the same posterior.

Building blocks. The pair $(\hat{\xi}, \sigma_\xi)$ of Section 5.2.2; the threshold $\xi_{\max} = \ln(0.05/0.95)$ for the valve's $p_{\max} = 0.05$; the normal cumulative distribution Φ; and, for the exact values beside them, the posterior parameters $a_n = 4$, $b_n = 66$ and the incomplete Beta function and its inverse.

The requirement of 95 percent reliability. By the first of equations (5.42) with $\xi_{\max} = \ln(0.05/0.95) = -2.9444$,

$$\Pr(R \geq 0.95 \mid \mathbf{y}) \approx \Phi\left(\frac{-2.9444 + 2.8034}{0.5149}\right) = \Phi(-0.2740) = 0.3921,$$

about 39.2 %: the shaded area of Figure 5.4(b), left of the dashed line at $p = 0.05$. The exact value, equation (5.43), is $I_{0.05}(4, 66) = 0.4550$, 0.063 higher. For 90 % reliability, $\xi_{\max} = \ln(0.10/0.90) = -2.1972$ and $\Pr(R \geq 0.90 \mid \mathbf{y}) \approx \Phi(1.1771) = 0.8804$, against the exact $I_{0.10}(4, 66) = 0.9238$. Table 5.5 gives five requirement levels: the first column is the minimum reliability $R_{\min}$, the second the equivalent bound on p, the third the probability under the prior, and the last two the probability under the posterior, from the Laplace normal and exactly. The record raises the support for every requirement up to 96 % and lowers it for 99 %, on both routes; the Laplace route is lower at every level, because its normal in the logit puts too much weight on large failure probabilities.

Table 5.5: Probability that the valve meets a reliability requirement, $\Pr(p \leq 1 - R_{\min})$: under the prior Beta(2, 18), and under the posterior from the Laplace normal of the logit, equation (5.42), and from the exact Beta(4, 66), equation (5.38).

Requirement $R_{\min}$	Bound on p	Prior	Posterior, Laplace	Posterior, exact
0.85	$p \leq 0.15$	0.8015	0.9810	0.9950
0.90	$p \leq 0.10$	0.5797	0.8804	0.9238
0.95	$p \leq 0.05$	0.2453	0.3921	0.4550
0.96	$p \leq 0.04$	0.1751	0.2334	0.2980
0.99	$p \leq 0.01$	0.0153	0.0003	0.0052

The assurance rule (5.44) at $\kappa = 0.95$ fails for 90 % reliability, $0.8804 < 0.95$ (0.9238 exact), and for 95 %, $0.3921 < 0.95$ (0.4550 exact): the two routes differ in the support, not in the decision. The one-sided bound, the second of equations (5.42), is

$$p_{0.95} \approx p(-2.8034 + 1.6449 \times 0.5149) = p(-1.9564) = 0.12386, \qquad R_L = 1 - p_{0.95} = 0.87614,$$

a reliability of at least 87.6 % with 95 % probability, below the 90 % requirement; exactly, equation (5.45), $p_{0.95} = I^{-1}_{0.95}(4, 66) = 0.10855$ and $R_L = 0.89145$. Both are drawn in red in Figure 5.4(b), against the upper ends 0.1426 and 0.1218 of the central intervals.

5.3.3 The next demand and future failures

Engineering question. Will the next demand succeed? Will ten future demands all succeed? The questions now change from whether the unknown reliability exceeds a fixed requirement to what the valve will do on its next comparable demand, and on the $M = 10$ further functional tests the manufacturer plans.

General form. A prediction of a future observation y_{new} is its probability given the state, averaged over the posterior of the state,

$$p(y_{\text{new}} \mid \mathbf{y}) = \int p(y_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}. \tag{5.46}$$

Specialized form. One choice enters here: the coming demands are taken to share the failure probability of the fifty already made, so the valve neither ages nor is repaired between them. That makes every future demand a draw from the same posterior; if the valve aged, the predictive would need the trend model of Chapter 12.

For the next demand, the question is equation (5.46) with $y_{\text{new}} = X_{\text{next}}$, the outcome of the next comparable demand, $p(y_{\text{new}} \mid \mathbf{z}) = \Pr(X_{\text{next}} = 0 \mid \xi) = 1 - p(\xi)$, and $p(\mathbf{z} \mid \mathbf{y})$ the posterior of step 3. The integral is then the Gauss–Hermite sum (5.33) with $h(\xi) = 1 - p(\xi)$,

the sum that gave the expected reliability: the probability that the next demand succeeds equals the posterior mean reliability, because $1 - p$ is what both average. Exactly, by linearity of expectation and the posterior mean of equation (5.34),

$$\Pr(X_{\text{next}} = 0 \mid \mathbf{y}) = 1 - \frac{a_n}{a_n + b_n} = \frac{b_n}{a_n + b_n}. \tag{5.47}$$

It is not the requirement probability of equation (5.43): the first predicts a future outcome; the second expresses uncertainty about whether the underlying reliability meets a requirement.

For the future demands, the question is equation (5.46) again, with $y_{\text{new}} = K_f$, the number of failures in the M future tests, and, since the tests are independent given the state, $p(y_{\text{new}} \mid \mathbf{z}) = \Pr(K_f = j \mid \xi) = \binom{M}{j} p(\xi)^j (1 - p(\xi))^{M-j}$, the binomial of equation (5.7). Under the Laplace posterior the integral is the Gauss–Hermite sum (5.33) with

$$h_j(\xi) = \binom{M}{j} p(\xi)^j \big(1 - p(\xi)\big)^{M-j}, \qquad h_0(\xi) = \big(1 - p(\xi)\big)^M \tag{5.48}$$

for a clean run, the route that remains for a predictive without a closed form. Exactly, the same integral written in p averages the binomial over the Beta posterior,

$$\Pr(K_f = j \mid \mathbf{y}) = \int_0^1 \Pr(K_f = j \mid p)\, f(p \mid \mathbf{y})\, \mathrm{d}p;$$

substitute the binomial probability and the Beta posterior,

$$\Pr(K_f = j \mid \mathbf{y}) = \binom{M}{j} \int_0^1 p^j (1-p)^{M-j}\, \frac{p^{a_n - 1}(1-p)^{b_n - 1}}{B(a_n, b_n)}\, \mathrm{d}p,$$

and combine the powers:

$$\Pr(K_f = j \mid \mathbf{y}) = \frac{\binom{M}{j}}{B(a_n, b_n)} \int_0^1 p^{a_n + j - 1}(1-p)^{b_n + M - j - 1}\, \mathrm{d}p.$$

The integral is, by the definition of the Beta function (Section 5.1.6), $B(a_n + j,\, b_n + M - j)$. Therefore

$$\Pr(K_f = j \mid \mathbf{y}) = \binom{M}{j} \frac{B(a_n + j,\ b_n + M - j)}{B(a_n, b_n)}, \qquad j = 0, 1, \ldots, M, \tag{5.49}$$

the *Beta–binomial posterior predictive distribution*. It carries both kinds of uncertainty: the random variability of the outcomes of future tests for a given p, and the uncertainty about the common failure probability, as the sum (5.48) does on the Laplace route.

With $j = 0$, the question's clean run, equation (5.49) is a ratio of Beta functions, which the gamma-function form of the Beta function turns into a finite product:

$$\Pr(K_f = 0 \mid \mathbf{y}) = \frac{B(a_n,\ b_n + M)}{B(a_n, b_n)} = \frac{\Gamma(b_n + M)\,\Gamma(a_n + b_n)}{\Gamma(b_n)\,\Gamma(a_n + b_n + M)} = \prod_{i=0}^{M-1} \frac{b_n + i}{a_n + b_n + i}. \tag{5.50}$$

The product has a reading of its own: its first factor is the probability that the next test succeeds, equation (5.47), and each later factor is the same probability after the posterior has been updated with the successes before it. The probability of at least two failures follows by the complement, $\Pr(K_f \geq 2 \mid \mathbf{y}) = 1 - \Pr(K_f = 0 \mid \mathbf{y}) - \Pr(K_f = 1 \mid \mathbf{y})$, with $\Pr(K_f = 1 \mid \mathbf{y}) = M\, B(a_n + 1,\, b_n + M - 1)/B(a_n, b_n)$.

Remark 5.1. *The plug-in shortcut.* A common shortcut substitutes the estimate $\hat{p} = \mathbb{E}[p \mid \mathbf{y}]$ into the binomial and computes $\Pr(K_f = 0) \approx (1 - \hat{p})^M$. The two differ because

$$\mathbb{E}\big[(1-p)^M \mid \mathbf{y}\big] \neq \big(1 - \mathbb{E}[p \mid \mathbf{y}]\big)^M.$$

The function $(1-p)^M$ is convex in p, its second derivative $M(M-1)(1-p)^{M-2}$ being positive, and the expectation of a convex function is at least the function of the expectation (Jensen's inequality), so the plug-in value understates the probability of a clean run. The full predictive calculation accounts for the uncertainty in p: the smaller failure probabilities the posterior admits raise the chance of M successes by more than the larger ones lower it. The same uncertainty fattens the tail of the failure count.

Predictive mean and variance. The predictive moments follow from the conditional ones by the laws of total expectation and total variance, without summing equation (5.49). Given p, $\mathbb{E}[K_f \mid p] = Mp$ and $\mathrm{Var}(K_f \mid p) = Mp(1-p)$. Write $\mu_p = \mathbb{E}[p \mid \mathbf{y}]$. The mean is

$$\mathbb{E}[K_f \mid \mathbf{y}] = \mathbb{E}\big[\mathbb{E}[K_f \mid p] \mid \mathbf{y}\big] = M\mu_p = M\,\frac{a_n}{a_n + b_n}. \tag{5.51}$$

The variance is

$$\mathrm{Var}(K_f \mid \mathbf{y}) = \mathbb{E}\big[\mathrm{Var}(K_f \mid p) \mid \mathbf{y}\big] + \mathrm{Var}\big(\mathbb{E}[K_f \mid p] \mid \mathbf{y}\big) = M\,\mathbb{E}[p - p^2 \mid \mathbf{y}] + M^2\,\mathrm{Var}(p \mid \mathbf{y}).$$

Since $\mathbb{E}[p^2 \mid \mathbf{y}] = \mathrm{Var}(p \mid \mathbf{y}) + \mu_p^2$, the first term is $M\mu_p(1-\mu_p) - M\,\mathrm{Var}(p \mid \mathbf{y})$, and

$$\mathrm{Var}(K_f \mid \mathbf{y}) = M\mu_p(1-\mu_p) + M(M-1)\,\mathrm{Var}(p \mid \mathbf{y}). \tag{5.52}$$

Written in the posterior parameters, with $s = a_n + b_n$, the same variance is the standard Beta–binomial form,

$$\mathrm{Var}(K_f \mid \mathbf{y}) = \frac{M a_n b_n}{s^2}\left(1 + \frac{M-1}{s+1}\right) = \frac{M\,a_n b_n\,(a_n + b_n + M)}{(a_n + b_n)^2\,(a_n + b_n + 1)}. \tag{5.53}$$

If the posterior mean failure probability were treated as exactly known, the variance would be the binomial first term alone. The second term is the additional predictive variance caused by the uncertainty about the failure probability; it vanishes for a single future test, $M = 1$, which is why the next-test probability, equation (5.47), needed only the mean.

Dependence of future trials. The trials were assumed independent conditional on p. Averaged over an uncertain common p, future outcomes are dependent. For two distinct future failure indicators X_i and X_j, conditional independence gives $\mathbb{E}[X_i X_j \mid p] = p^2$, so that $\mathbb{E}[X_i X_j \mid \mathbf{y}] = \mathbb{E}[p^2 \mid \mathbf{y}]$, while $\mathbb{E}[X_i \mid \mathbf{y}] = \mathbb{E}[X_j \mid \mathbf{y}] = \mu_p$. Therefore

$$\mathrm{Cov}(X_i, X_j \mid \mathbf{y}) = \mathbb{E}[p^2 \mid \mathbf{y}] - \mu_p^2 = \mathrm{Var}(p \mid \mathbf{y}), \qquad i \neq j, \tag{5.54}$$

under the shared-p model, with the correlation $\mathrm{Var}(p \mid \mathbf{y})/(\mu_p(1-\mu_p)) = 1/(a_n + b_n + 1)$. A failure in one future test is evidence that p is larger, and so raises the probability of failure in the next. Conditional independence does not imply marginal independence when both outcomes depend on the same uncertain parameter, the distinction of Chapter 3; the $M(M-1)$ cross terms of equation (5.52) are exactly these covariances, summed over the ordered pairs of future tests.

Common mistake

Substituting a point estimate of p into the binomial and reporting the result as the probability of a future test outcome. The plug-in value ignores the uncertainty in p; the probability of a future outcome is the binomial averaged over the posterior, equation (5.49).

Building blocks. The pair $(\hat{\xi}, \sigma_\xi)$ of Section 5.2.2; the five Gauss–Hermite nodes of Table 5.4; the number $M = 10$ of future tests; and, for the exact values beside them, the posterior parameters $a_n = 4$, $b_n = 66$ and the Beta function, with the predictive mean, variance and covariance of equations (5.51), (5.52) and (5.54).

The probability that the next demand succeeds. The sum over the fifth column of Table 5.4 gives $\Pr(X_{\text{next}} = 0 \mid \mathbf{y}) \approx 0.9364$, about 93.6 %, the expected reliability of the valve; exactly, equation (5.47), $66/70 = 0.94286$. Both stand against a requirement probability of 0.3921 (0.4550 exact).

Failures in ten future demands. The sum over the last column of Table 5.4 is the sum (5.48) for a clean run of ten,

$$\begin{aligned}\Pr(K_f = 0 \mid \mathbf{y}) \approx{}& 0.01126 \times 0.87090 + 0.22208 \times 0.74298 + 0.53333 \times 0.55521 \\ &+ 0.22208 \times 0.31682 + 0.01126 \times 0.09614 = 0.5424,\end{aligned}$$

about 54.2 %. Each node is a value of the failure probability the posterior admits, from 1.4 % to 20.9 %, and the sum weighs the clean run at each: the middle node alone, 0.55521, is the plug-in value at $\hat{p}$, and the outer nodes spread the answer over the posterior. With the same nodes and h_1, $\Pr(K_f = 1 \mid \mathbf{y}) \approx 0.3191$, so $\Pr(K_f \geq 2 \mid \mathbf{y}) \approx 1 - 0.5424 - 0.3191 = 0.1385$, about 13.9 %. Exactly, by equation (5.50) with $M = 10$,

$$\Pr(K_f = 0 \mid \mathbf{y}) = \frac{B(4, 76)}{B(4, 66)} = \frac{\Gamma(76)\,\Gamma(70)}{\Gamma(66)\,\Gamma(80)} = \prod_{i=0}^{9} \frac{66 + i}{70 + i} = \frac{66}{70} \times \frac{67}{71} \times \cdots \times \frac{75}{79} = 0.5754,$$

and with $\Pr(K_f = 1 \mid \mathbf{y}) = 10\,B(5, 75)/B(4, 66) = 0.3069$, $\Pr(K_f \geq 2 \mid \mathbf{y}) = 1 - 0.5754 - 0.3069 = 0.1178$: the Laplace route gives the clean run 0.033 less and two or more failures 0.021 more, the same bias toward large failure probabilities as the mean. Table 5.6 lists the predictive for every count, from the Laplace sum (5.48) and exactly from equation (5.49), beside the plug-in binomial at $\hat{p} = 0.0571$; the row above the rule aggregates five or more failures, the row below it two or more. Figure 5.5 draws the same numbers: the horizontal axis is the number of failures in the next ten demands, the vertical axis its probability in percent, blue bars the Laplace predictive, green bars the exact one and gray bars the plug-in, whose tail is thinnest. The plug-in clean run is $0.94286^{10} = 0.5552$, and its probability of at least two failures 0.1083. The Laplace clean run lies below the plug-in at $\hat{p}$, not above it as Jensen's inequality would suggest, because the inequality compares the predictive with the plug-in at the posterior mean, and the Laplace mean of p is 0.0636, not 0.0571: $(1 - 0.0636)^{10} = 0.5182 < 0.5424$.

The predictive mean and variance, equations (5.51) and (5.52), need only the mean and variance of p. Under the Laplace posterior the same Gauss–Hermite sum with $h = p$ and $h = p^2$ gives $\mathbb{E}[p \mid \mathbf{y}] = 0.0636$ and $\mathrm{Var}(p \mid \mathbf{y}) = 0.000999$, so the mean is $10 \times 0.0636 = 0.636$ failures and the variance $10 \times 0.0636 \times 0.9364 + 90 \times 0.000999 = 0.596 + 0.090 = 0.685$, of which 0.090 comes from the uncertainty about p. The variance of p here is that of the logit-normal density, larger than the delta method's 7.70×10^{-4} of the spread question, which linearizes the map. Exactly, the mean is $10 \times 4/70 = 0.5714$ failures and the variance, by equation (5.52) or (5.53), $10 \times 0.05714 \times 0.94286 + 90 \times 0.00075884 = 0.5388 + 0.0683 = 0.6071$. The covariance of two future outcomes, equation (5.54), is the variance of p: 0.000999 under the Laplace posterior, a correlation of 0.017, and 0.00075884 exactly, a correlation of $1/71 = 0.014$.

5.3.4 The zero-failure demonstration size

Engineering question. How many tests establish a requirement? The new valve of Section 5.1.6 must show $R \geq 0.99$ with 95 % assurance, and the engineer asks how many tests are needed if none fails.

Table 5.6: Predicted number of failures in the next ten comparable valve demands. Laplace predictive, the Gauss–Hermite sum (5.48); exact Beta–binomial predictive from Beta$(4, 66)$, equation (5.49); and the plug-in Binomial$(10, 0.05714)$.

Failures j in ten tests	Laplace predictive	Exact predictive	Plug-in binomial
0	0.54236	0.57537	0.55521
1	0.31914	0.30687	0.33649
2	0.10614	0.09330	0.09177
3	0.02611	0.02045	0.01483
4	0.00523	0.00348	0.00157
5 or more	0.00103	0.00053	0.00012
2 or more	0.13850	0.11776	0.10830

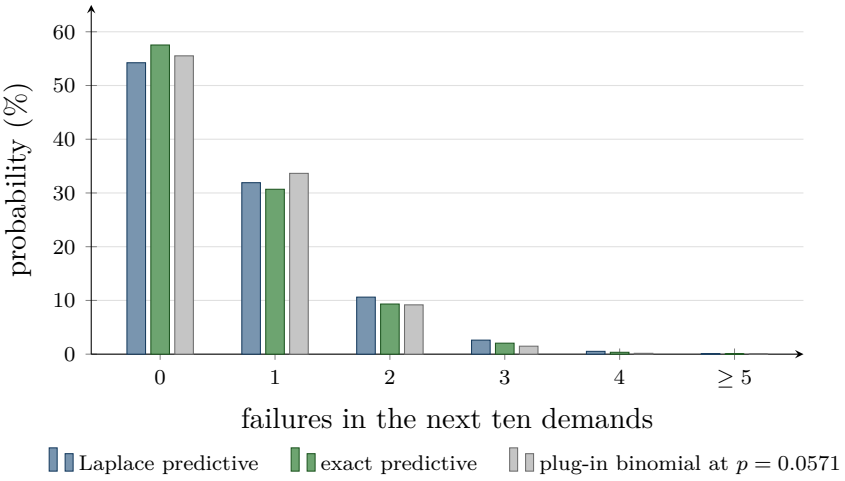


Figure 5.5: Predicted number of failures in the next ten comparable valve demands (Table 5.6). Blue: Laplace predictive, the sum (5.48) over $\mathcal{N}(-2.8034,\ 0.5149^2)$. Green: exact Beta–binomial predictive, equation (5.49) with $a_n = 4$, $b_n = 66$, $M = 10$. Gray: binomial with $\hat{p} = 0.0571$ substituted for p. The last category aggregates five or more.

General form. A design question about a test campaign is answered by the preposterior. Before a campaign of n tests is run, its outcome, the number K_n of failures, is unknown, and the probability that the campaign ends with the rule of equation (5.41) met averages the rule over the predictive distribution of that outcome,

$$\Pr(\text{rule met} \mid \mathbf{y}) = \sum_{j=0}^{n} \Pr(K_n = j \mid \mathbf{y})\, \mathbf{1}\big[\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}, K_n = j) \geq \kappa\big]; \tag{5.55}$$

the design chooses the campaign by this probability. A demonstration test fixes its accepted outcome in advance, $K_n = 0$, and its size is the smallest campaign whose planned outcome meets the rule,

$$n_{\min} = \min\{\, n : \Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}, K_n = 0) \geq \kappa \,\}. \tag{5.56}$$

Specialized form. Two choices enter here. The new valve has no record and is given the uniform prior: this makes the campaign size a property of the requirement alone, and it ignores whatever is known about similar valves (Section 5.4.3 gives the size under the engineering prior). The accepted outcome, zero failures, is fixed before the campaign starts: this is what makes the tail calculation valid, and it is broken by deciding, after a failure, to extend the test until the record looks good.

For the new valve, the question is equation (5.56) with $f(\mathbf{z}) = R = 1 - p(\xi)$, $f^* = 0.99$, $\kappa = 0.95$, and $p(\mathbf{z} \mid \mathbf{y}, K_n = 0)$ the posterior of the state after a planned record $\mathbf{y}_n$ of n successes under the uniform prior $\text{Beta}(1, 1)$, there being no earlier record. By the tail of the previous question, the acceptance criterion is

$$\Pr(p \leq p_{\max} \mid \mathbf{y}_n) \geq \kappa, \tag{5.57}$$

for the new valve $\Pr(p \leq 0.01 \mid \mathbf{y}_n) \geq 0.95$. The planned record has two rows on the logit, the uniform prior's one pseudo-failure in two pseudo-trials and the count row of 0 failures in n, so step 3's construction gives, for each n, the MAP state $\hat{\xi}_n = \ln(1/(n+1))$, at $\hat{p}_n = 1/(n+2)$, and the information of the prior row's two pseudo-trials and the n count rows,

$$\Lambda_n = (2+n)\,\hat{p}_n(1-\hat{p}_n) = \frac{n+1}{n+2}, \qquad \sigma_\xi^2 = \frac{1}{\Lambda_n} = \frac{n+2}{n+1}, \tag{5.58}$$

which stays near one for every n: each further success adds the information $\hat{p}_n(1-\hat{p}_n)$, which shrinks as $\hat{p}_n$ does. The tail of equation (5.42) at $\xi_{\max} = \ln(p_{\max}/(1-p_{\max}))$ is then

$$\Pr(p \leq p_{\max} \mid \mathbf{y}_n) \approx \Phi\left((\xi_{\max} + \ln(n+1)) \sqrt{\frac{n+1}{n+2}} \right), \tag{5.59}$$

and the campaign size is the smallest n at which it reaches κ. The tail improves with n because $\hat{\xi}_n$ falls as $-\ln(n+1)$, not because σ_ξ shrinks.

The value beside it comes from the exact posterior. With $k = 0$ failures the posterior parameters are $a_n = 1$ and $b_n = n + 1$, so the posterior is $\text{Beta}(1, n+1)$. Its normalizing constant is $B(1, n+1) = \int_0^1 (1-p)^n \, \mathrm{d}p = 1/(n+1)$, so the density is

$$f(p \mid k = 0, n) = (n+1)(1-p)^n. \tag{5.60}$$

The probability of meeting a requirement $p \leq p_{\max}$ is the integral of this density up to $p_{\max}$. The antiderivative of $(n+1)(1-p)^n$ is $-(1-p)^{n+1}$, and evaluating it between the limits,

$$\Pr(p \leq p_{\max} \mid k = 0, n) = \Big[-(1-p)^{n+1}\Big]_0^{p_{\max}} = 1 - (1-p_{\max})^{n+1}, \tag{5.61}$$

the acceptance-probability formula for zero failures under a uniform prior. It is also equation (5.39) with $a = 1$ and $b = n + 1$: the probability that a Binomial$(n + 1, p_{\max})$ count is at least one. The criterion (5.57) then requires

$$1 - (1 - p_{\max})^{n+1} \geq \kappa \quad \Longleftrightarrow \quad (1 - p_{\max})^{n+1} \leq 1 - \kappa.$$

Taking logarithms, $(n + 1)\ln(1 - p_{\max}) \leq \ln(1 - \kappa)$. Because $\ln(1 - p_{\max}) < 0$, dividing by it reverses the inequality, and

$$n \geq \frac{\ln(1 - \kappa)}{\ln(1 - p_{\max})} - 1, \qquad n_{\min} = \left\lceil \frac{\ln(1 - \kappa)}{\ln(1 - p_{\max})} - 1 \right\rceil. \tag{5.62}$$

The two routes part here more than anywhere else in the chapter. After zero failures the exact posterior of ξ is skewed, with a long tail toward small failure probabilities and a steep side toward large ones, and its variance stays well above σ_ξ^2; the normal fitted to the curvature at the MAP state puts too much probability above $\xi_{\max}$, understates the tail and asks for more tests than the exact posterior. The size of the campaign is set by the tail of the posterior at $p_{\max}$, not by its width; the next tests of the valve already tested are ranked by the same tail in Section 5.6.

Building blocks. The MAP state $\hat{\xi}_n$ and the variance σ_ξ^2 of equation (5.58) for each campaign size n, the threshold $\xi_{\max} = \ln(0.01/0.99) = -4.5951$, $\kappa = 0.95$ and Φ in the tail (5.59); and, for the exact value beside it, the Beta tail at $p_{\max}$ in the closed form (5.61) and the logarithms of equation (5.62). Two sensitivities of the plan are computed with the same blocks, because the result states how many tests are necessary *if the campaign produces zero failures* and nobody knows before the campaign whether the valve will fail during it. The first is the *operating characteristic*, the probability of acceptance $(1 - p)^n$ of a plan of n tests that passes only with zero failures, as a function of the true p; a good valve can fail the demonstration with appreciable probability, so the size of the campaign is a trade between that risk and the cost of testing, and a later chapter on acceptance testing develops it. The second is the campaign after a failure: for $k > 0$ the posterior is Beta$(1 + k,\ 1 + n - k)$ and the tail $I_{p_{\max}}(1 + k, 1 + n - k)$ must be recalculated. A demonstration program should therefore specify in advance how failures are handled, what counts as a valid test and whether additional testing is allowed, and it should not select new rules after observing the results merely to obtain a passing outcome.

The demonstration campaign for the new valve. With $p_{\max} = 0.01$ and $\kappa = 0.95$, the tail (5.59) first reaches 0.95 at $n = 513$, where

$$\hat{\xi}_{513} = -\ln 514 = -6.2422, \quad \sigma_\xi = 1.0010, \quad \Phi\!\left(\frac{-4.5951 + 6.2422}{1.0010}\right) = 0.95007,$$

whereas 512 tests give 0.94987: the Laplace route asks for 513 zero-failure tests. The exact route, equation (5.62), requires

$$n + 1 \geq \frac{\ln 0.05}{\ln 0.99} = \frac{-2.99573}{-0.01005} = 298.07,$$

that is $n \geq 297.07$, and rounding up to the next whole number,

$$n_{\min} = 298 \text{ tests.} \tag{5.63}$$

Zero failures in 298 tests give $1 - 0.99^{299} = 0.95046$, or 95.05 %, whereas 297 tests give $1 - 0.99^{298} = 0.94996$. The Laplace route asks for 215 more tests, 72 % more than the exact one: at 298 tests its tail is 0.865, against the exact 0.950. This is a gap that changes the plan.

The reported 513 is what the posterior of the state gives, the size a model without a closed form would ask for; the exact 298 beside it measures what the normal in the logit costs after a record with no failures, 215 tests. Figure 5.6 draws both tails against the number of tests: the horizontal axis is n, the vertical axis the posterior probability of the requirement after n successes, the dotted curves the Laplace tail (5.59) and the solid and dashed curves the exact tail (5.61), for $p \le 0.01$ and $p \le 0.02$, and the horizontal dashed line the assurance 0.95. The exact curves cross it at 298 and 148 tests, the Laplace curves at 513 and 254.

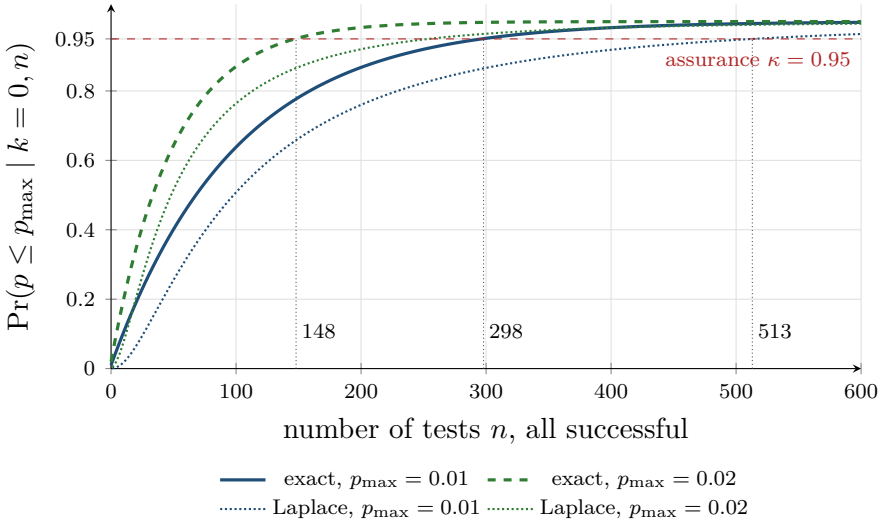


Figure 5.6: Posterior probability that a zero-failure campaign of n tests demonstrates $p \le p_{\max}$ under the uniform prior, for $p_{\max} = 0.01$ and 0.02. Dotted: the Laplace tail (5.59), crossing the assurance 0.95 at 513 and 254 tests. Solid and dashed: the exact tail $1 - (1 - p_{\max})^{n+1}$, equation (5.61), crossing it at 298 and 148 tests.

Table 5.7 applies the tail (5.59) and, beside it, equation (5.62) to four requirements (rows) and three assurance levels (columns). Each cell gives the Laplace size, then the exact size in parentheses. The exact size scales roughly as $-\ln(1-\kappa)/p_{\max}$, ten times larger for a tenfold smaller $p_{\max}$ and twice as large from $\kappa = 0.90$ to 0.99; the Laplace size is larger in every cell, by 56 % to 73 % at $\kappa = 0.90$ and 0.95 and by more than double at $\kappa = 0.99$, where the tail of the normal matters most.

Table 5.7: Smallest zero-failure campaign demonstrating $\Pr(p \le p_{\max} \mid 0, n) \ge \kappa$ under the uniform prior Beta(1, 1): Laplace tail (5.59), with the exact size of equation (5.62) in parentheses.

Requirement	$\kappa = 0.90$	$\kappa = 0.95$	$\kappa = 0.99$
$p_{\max} = 0.05$ $(R \ge 0.95)$	69 (44)	99 (58)	195 (89)
$p_{\max} = 0.02$ $(R \ge 0.98)$	177 (113)	254 (148)	502 (227)
$p_{\max} = 0.01$ $(R \ge 0.99)$	357 (229)	513 (298)	1014 (458)
$p_{\max} = 0.001$ $(R \ge 0.999)$	3599 (2301)	5175 (2994)	10231 (4602)

The operating characteristic $(1-p)^{513}$ of the 513-test plan is 0.599 at $p = 0.001$, 0.214 at 0.003, 0.076 at 0.005 and 0.006 at the limit $p = 0.01$: the larger plan fails a good valve more often. For the exact 298-test plan it is 0.742, 0.408, 0.225 and 0.050. After one failure the planned record has the two rows of one failure in n tests and the uniform prior, the logit posterior is that of Beta$(2, n)$, and the Laplace tail reaches 0.95 at $n = 635$ tests in all; at 298 tests it is 0.718. Exactly, the posterior after one failure in 298 tests is Beta$(2, 298)$, $I_{0.01}(2, 298) = 0.8009$, and the criterion needs $n = 472$ tests in all.

5.4 The checks (step 5)

This section sets the Laplace approximation and the closed-form evaluations beside the full calculation, which for the valve is the exact Beta posterior, and then varies the prior. The valve has no physics row and no exact constraint, so no limit $\sigma_p \to 0$ of a model tolerance arises and no limit table is needed; the bound $0 < p < 1$ is not a row either, and it is handled by the logit state.

5.4.1 The Laplace posterior against the exact posterior

Figure 5.7 compares the Laplace approximation of the fit with the exact posterior. Panel (a) has the state entry ξ on the horizontal axis and the objective of equation (5.17), less its minimum, on the vertical axis; since the state has one entry, the curve is the whole objective, not a slice of it. The solid curve is the exact objective and the dashed curve the parabola $\frac{1}{2}\Lambda\,(\xi - \hat{\xi})^2$ with the information of equation (5.21), which the Laplace approximation puts in its place. The two touch at the minimum with the same curvature and separate away from it: toward small failure probabilities, on the left, the exact objective rises more slowly than the parabola, and toward large ones it rises faster. Panel (b) has the failure probability on the horizontal axis and the probability density on the vertical axis; the green curve is the exact Beta$(4, 66)$ and the dotted blue curve the Laplace normal in the logit, transformed back to p, equation (5.28). The transformed normal stays inside $(0, 1)$ and is skewed to the right like the exact posterior, but it lies to the right of it, as panel (a) predicts: it gives too little weight to small failure probabilities and too much to large ones.

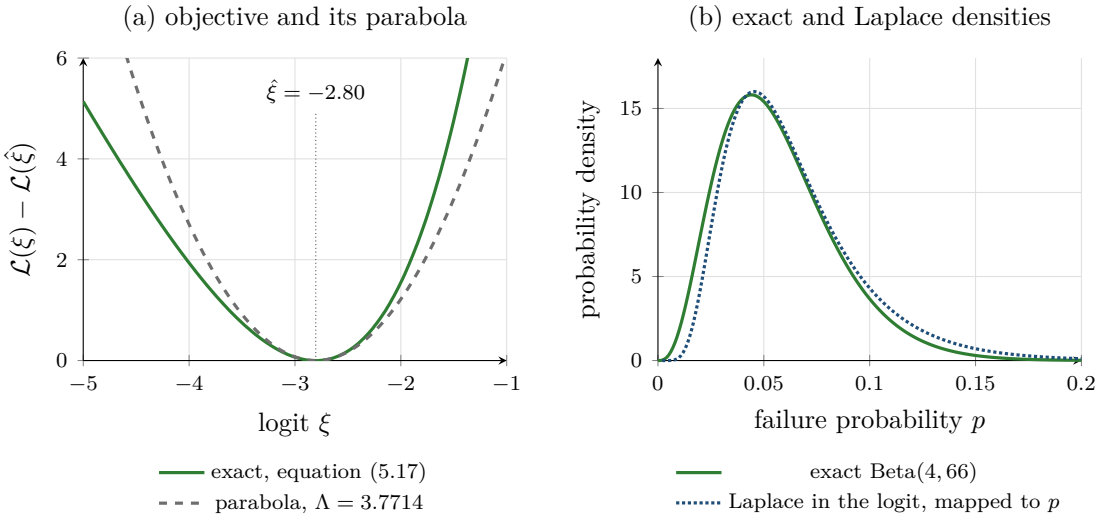


Figure 5.7: The Laplace approximation of the valve posterior. (a) The objective $-4\ln p - 66\ln(1-p)$, with $p = p(\xi)$, less its minimum, against the parabola $\frac{1}{2} \times 3.7714\,(\xi + 2.8034)^2$. (b) Exact density $3\,458\,004\,p^3(1-p)^{65}$, and the normal $\mathcal{N}(-2.8034,\ 0.5149^2)$ on $\xi = \ln(p/(1-p))$ mapped to p, density $\varphi((\xi + 2.8034)/0.5149)/(0.5149\,p(1-p))$ with φ the standard normal density.

Table 5.8 turns the figure into numbers for the questions of the chapter. Each row is a summary; the columns give it under the exact posterior and under the Laplace approximation, whose standard deviation comes from equation (5.35) and whose interval and probabilities come from the normal in ξ, equations (5.36) and (5.42), and whose expectations and predictive probabilities come from the Gauss–Hermite sum (5.33). The approximation respects the bounds and the skew and carries the standard deviation of p to within 1 %; it moves the interval and

the one-sided bound toward larger failure probabilities and understates the support for the 95 % requirement by 0.06. The expectations carry the same bias: the mean reliability and the next-demand probability are 0.936 against 0.943, a clean run of ten demands 0.542 against 0.575, and at least two failures in ten 0.139 against 0.118, because the normal in ξ puts too much weight on large failure probabilities. For the new valve's zero-failure campaign the gap is largest, as Section 5.3.4 shows: at 298 tests the information of equation (5.58) is $\Lambda_n = 299/300$, so $\sigma_\xi^2 = 1.003$, while the exact posterior of ξ has the variance 1.65; the Laplace tail is 0.865 against the exact 0.950, and the Laplace route asks for 513 tests against 298. Every answer of the chapter is reported from the posterior of the state, the route that remains when the binary evidence enters a larger model that has no closed form; the table shows, question by question, what that route costs on a record this small, and that it costs most where the answer is a tail after a record with no failures. Carried in p itself as the state entry, the count row and the prior written as a density of p, without the change-of-variable factor, give a fit that ends at the exact mode $3/68 = 0.04412$ with a standard deviation of 0.02490 in p, whose normal gives 0.593 for the 95 % requirement and puts 3.8 % of its probability on negative failure probabilities, and whose first Newton step from the prior mean lands at $p = -0.0110$, outside the admissible interval, where the first step in the logit from the same start, $\xi_0 = \ln(0.1/0.9) = -2.1972$, lands at $\xi = -2.6734$, $p = 0.0646$; the logit avoids all three. In p the curvature of each row also carries a term proportional to its expected-minus-observed count, so it depends on the record; in the logit that term vanishes at every state (Section 5.1.6).

Table 5.8: The valve posterior Beta$(4, 66)$ against its Laplace approximation in the logit ξ, equation (5.21), mapped to p; expectations by the Gauss–Hermite sum (5.33) with $Q = 40$. The last two rows are the new valve's zero-failure campaign under the uniform prior, equation (5.58).

Summary	Exact	Laplace in ξ
Center	0.0441 (mode) / 0.0571 (mean)	0.0571 (median)
Standard deviation of p	0.0275	0.0277 (delta method)
Central 95 % interval	[0.0160, 0.1218]	[0.0216, 0.1426]
$\Pr(p \le 0.05 \mid \mathbf{y})$	0.455	0.392
One-sided bound $p_{0.95}$	0.109	0.124
$\mathbb{E}[R \mid \mathbf{y}] = \Pr(X_{\text{next}} = 0 \mid \mathbf{y})$	0.943	0.936
$\Pr(K_f = 0 \mid \mathbf{y})$, $M = 10$	0.575	0.542
$\Pr(K_f \ge 2 \mid \mathbf{y})$, $M = 10$	0.118	0.139
New valve, tail at $n = 298$	0.950	0.865
New valve, zero-failure size $n_{\min}$	298	513

5.4.2 The exact identities

Four results of Section 5.3 have independent checks: the binomial tail, the sufficiency of the record, the end of the fit and the demonstration size. The requirement probability $I_{0.05}(4, 66) = 0.4550$ is checked by hand with equation (5.39): it is the probability that a Binomial$(69, 0.05)$ count is at least four, one minus the four terms $0.0290 + 0.1054 + 0.1887 + 0.2218 = 0.5450$, which gives 0.4550.

The sufficiency of (k, n) is checked by splitting the record. Suppose the first campaign gives one failure in twenty tests. Starting from Beta$(2, 18)$, the posterior is Beta$(2 + 1, 18 + 19) =$ Beta$(3, 37)$. A later campaign gives one more failure in thirty tests; with Beta$(3, 37)$ as its prior,

the posterior is $\text{Beta}(3+1,\ 37+29) = \text{Beta}(4, 66)$, exactly the posterior obtained by updating the original prior with all fifty tests at once. In the state-vector form the split puts one count row per campaign in the objective, and the two terms sum to the count row of the whole record, equation (5.9); addition does not depend on the order or the grouping. For a monitoring system that receives test results continually, the posterior is updated by two additions and no stored record needs to be revisited.

Finally the fit itself: the Newton iterations of Table 5.2 end at $\hat{\xi} = \ln(4/66)$, which maps to the exact posterior mean $4/70$, as equation (5.20) states, and the demonstration size of 298 tests is confirmed by the acceptance probability at 298 and 297 tests in Section 5.3.4.

5.4.3 Sensitivity to the prior

Bayesian inference incorporates prior assumptions, and part of validating a result is determining whether the engineering conclusion changes substantially under reasonable alternative priors. Table 5.9 returns to the original record, $n = 50$, $k = 2$, and repeats the analysis under five priors. Each row is a prior; the columns give the posterior from equation (5.24), its mean, the probabilities of the 95 % and 90 % reliability requirements, the one-sided 95 % upper bound and the probability of no failure in ten further tests. The first row is the engineering prior; the second the uniform prior; the third the Jeffreys prior $\text{Beta}(\frac{1}{2}, \frac{1}{2})$, a standard reference prior; the fourth an optimistic prior of mean 2 % worth fifty pseudo-tests; the fifth a strong prior of mean 10 % worth a hundred pseudo-tests. One row of the arithmetic, the uniform prior: $\text{Beta}(1+2,\ 1+48) = \text{Beta}(3, 49)$, mean $3/52 = 0.0577$, and $I_{0.05}(3, 49) = 0.4725$.

Table 5.9: Sensitivity of the valve conclusions to the prior; record $k = 2$ in $n = 50$. Each cell gives the Laplace value from the posterior of the logit, then the exact value from the Beta posterior in parentheses. $\hat{p}$ is the Laplace median, equal to the exact mean; $p_{0.95}$ is the one-sided 95 % upper bound; the last column is the clean run of $M = 10$ tests, equations (5.48) and (5.49).

ootnotesize

Prior	Posterior	$\hat{p}$	$\Pr(R \geq 0.95)$	$\Pr(R \geq 0.90)$	$p_{0.95}$	$\Pr(K_f = 0)$
$\text{Beta}(2, 18)$, engineering	$\text{Beta}(4, 66)$	5.71 %	39.2 (45.5) %	88.0 (92.4) %	0.124 (0.109)	0.542 (0.575)
$\text{Beta}(1, 1)$, uniform	$\text{Beta}(3, 49)$	5.77 %	40.0 (47.3) %	84.2 (89.6) %	0.140 (0.118)	0.536 (0.579)
$\text{Beta}(\frac{1}{2}, \frac{1}{2})$, Jeffreys	Beta(2.5, 48.5)	4.90 %	51.3 (59.0) %	88.2 (93.5) %	0.130 (0.106)	0.582 (0.630)
$\text{Beta}(1, 49)$, optimistic	$\text{Beta}(3, 97)$	3.00 %	81.8 (87.8) %	98.5 (99.8) %	0.075 (0.062)	0.714 (0.747)
$\text{Beta}(10, 90)$, strong	Beta(12, 138)	8.00 %	4.8 (7.1) %	79.2 (82.2) %	0.125 (0.119)	0.433 (0.445)

The engineering and the uniform priors produce similar conclusions on both routes: both leave the 95 % requirement below even odds and put the 90 % requirement below a 95 % assurance. The optimistic prior produces a materially different result, 81.8 % support for the 95 % requirement (87.8 % exact), because it expresses a strong belief that the failure probability is low; the strong prior pulls the other way, to 4.8 % (7.1 %). The Laplace value is below the exact one in every row of the two requirement columns, by 1 to 8 points: the bias of the normal in the logit has the same sign under every prior. The prior's weight is the prior row's share of the information, equation (5.22), the split of the information between the two rows: 0.29 for the engineering prior, 0.04 for the uniform, 0.50 for the optimistic and 0.67 for the strong one. The share measures how much a prior weighs against the record, not the direction it pulls:

the optimistic and strong priors, with the two largest shares, move the conclusions furthest, in opposite directions set by their means. A small share does not make a tail probability insensitive when the record holds only two failures: the Jeffreys prior, whose share is smaller still than the uniform prior's, raises the support for the 95 % requirement from 40.0 % to 51.3 % (47.3 % to 59.0 % exact). The demonstration test of Section 5.3.4 is just as sensitive: with the engineering prior Beta$(2, 18)$ instead of the uniform one, a zero-failure campaign needs 617 tests rather than 513 (454 rather than 298 exact), because the prior's mean of 10 % has to be overcome; with the optimistic prior Beta$(1, 49)$ it needs 465 (250 exact, exactly 48 fewer than 298, the prior's extra pseudo-successes).

> **Principle 5.1.** *When the observed reliability data are limited, prior assumptions can have an important influence on engineering conclusions. A responsible assessment documents the prior, identifies its source and examines whether alternative defensible priors change the acceptance decision; the prior is never chosen to obtain a desired result.*

5.5 The report (step 6)

The record does not support the first reading of it. Forty-eight successes in fifty is 96 %, above the specification, but the posterior probability that the valve meets the 95 % specification is 39.2 % (45.5 % exactly), less than even, and it stays near even at best under the uniform and the Jeffreys priors. The fifty tests show that the valve is good; they do not show that it is good enough. The answers below are given in the units in which the questions were asked.

5.5.1 The answers and their bands

Table 5.10 gives the answers. Each row is a question; the second column gives the answer under the engineering prior, from the posterior of the state; the third the exact value from Beta$(4, 66)$; and the last the band of the reported answer: the central 95 % credible interval for an estimate, and for a probability the range over the engineering, the uniform and the Jeffreys priors of Table 5.9, the three that carry no strong claim of their own.

5.5.2 The acceptance decision

The acceptance decision follows from the fourth row, Table 5.5 and the assurance rule of equation (5.44). The record does not demonstrate 95 % reliability at 95 % assurance, under any of the three priors. It does not demonstrate 90 % reliability at 95 % assurance either: the support is 0.880 against the required 0.95, and 84.2 % to 88.2 % over the three priors; exactly it is 0.924, a narrower miss, and 89.6 % to 93.5 %. The two routes reach the same decision on both requirements. What the record does support is the one-sided statement of the fourth row, a per-demand reliability of at least 87.6 % with 95 % posterior probability (89.1 % exactly).

Three distinct reliability statements. The record has produced three numbers associated with reliability, and Table 5.11 sets them side by side: each row is a statement, the middle column its value and the right column what kind of quantity it is. The observed 96 % success fraction is a statistic of the test record. The posterior mean reliability of 93.6 % (94.3 % exactly) is an estimate of the unknown per-demand reliability. The 39.2 % (45.5 %) is the posterior support for one specific engineering requirement. The distinction is fundamental to demonstration and acceptance testing.

Table 5.10: The report for the valve record, $k = 2$ failures in $n = 50$ tests. Answers from the posterior of the state, with the exact value from Beta$(4, 66)$ beside each; bands are credible intervals or the range over the engineering, uniform and Jeffreys priors (Table 5.9).

Question	Answer	Exact	Band
What is the failure probability?	median 5.71 %	median 5.30 %, mean 5.71 %, mode 4.41 %	2.16 % to 14.26 %, central 95 % interval (1.60 % to 12.18 % exact)
What is the expected reliability?	93.64 %	94.29 %	85.74 % to 97.84 %, central 95 % interval
How uncertain is the estimate?	standard deviation 2.77 percentage points	2.75 points	2.16 % to 14.26 %, central 95 % interval
Is reliability at least 95 %?	probability 39.2 %; reliability at least 87.6 % with 95 % probability	45.5 %; 89.1 %	39.2 % to 51.3 % over the three priors
Will the next demand succeed?	probability 93.64 %	94.29 %	93.36 % to 94.18 % over the three priors
Will ten future demands all succeed?	probability 54.2 %	57.5 %	53.6 % to 58.2 % over the three priors
How many tests establish a requirement?	513 zero-failure tests for $R \geq 0.99$ at 95 % assurance, uniform prior	298 tests	617 under the engineering prior (454 exact); 635 in all if one failure occurs (472)

Table 5.11: Three numbers drawn from the valve record. Posterior of the state, exact values from Beta$(4, 66)$ in parentheses.

Statement	Value	Kind of quantity
Observed success fraction $(n - k)/n$	96.0 %	a statistic of the record
Posterior mean reliability $\mathbb{E}[R \mid \mathbf{y}]$	93.6 (94.3) %	an estimate of the unknown reliability
$\Pr(R \geq 0.95 \mid \mathbf{y})$	39.2 (45.5) %	posterior support for a stated requirement

Common mistake

Reading a test record that shows 96 % successes as 96 % confidence that the component satisfies a 95 % reliability requirement. The success fraction is a statistic of the record; the support for the requirement is a posterior probability, here 39.2 % (45.5 % exactly), and it depends on the number of tests as well as on their outcomes.

5.6 The next measurement (step 7)

The report leaves one decision open, the 90 % requirement at 95 % assurance, missed by 0.070 (0.026 exactly). The candidate measurements are further demands of the same valve under the same procedure, M more tests.

5.6.1 The preposterior variance

Each future test adds a count row; its slope depends on the unseen outcome, but in the logit its curvature does not: each adds $\hat{p}(1-\hat{p})$ to the curvature, so the information the M tests add at the current minimum is known in advance, and the informations add,

$$\Lambda_a = M\,\hat{p}(1-\hat{p}), \qquad \frac{1}{\sigma^2_{\xi,a}} = \frac{1}{\sigma^2_{\xi}} + \Lambda_a = (a_n + b_n + M)\,\hat{p}(1-\hat{p}), \qquad \frac{\sigma^2_{\xi,a}}{\sigma^2_{\xi}} = \frac{a_n + b_n}{a_n + b_n + M}, \tag{5.64}$$

with $\sigma^2_{\xi,a}$ the Laplace preposterior variance of the logit, the one-entry case of equation (2.17). For the $M = 10$ tests of Section 5.3.3, $\Lambda_a = 0.53878$ and $\sigma^2_{\xi,a} = 1/(3.7714 + 0.5388) = 0.23201$, exactly $70/80 = 0.875$ of σ^2_{ξ}: a standard deviation of 0.4817 in the logit and, through the same g, 0.02595 in p. The exact counterpart is the posterior variance of p averaged over the Beta–binomial predictive of equation (5.49), $\sum_j \Pr(K_f = j \mid \mathbf{y})\ \mathrm{Var}(p \mid \mathrm{Beta}(4 + j,\ 76 - j)) = 6.640 \times 10^{-4}$, whose square root is 0.02577. Its ratio to the current exact variance is also $70/80$, and for $M = 50$ both ratios are $70/120$, so the Laplace preposterior shrinks the variance by exactly the right factor. The gap between 0.02595 and 0.02577, and between 0.02119 and 0.02104 for $M = 50$, comes entirely from the current Laplace standard deviation, 0.02774 against the exact 0.02755. The size of a demonstration campaign is set by another quantity, the zero-failure tail at the requirement $p_{\max}$, equation (5.62).

5.6.2 The candidate campaigns

A campaign of M further tests is a design question of the form of equation (5.55), with the record of M outcomes $K_f = j$ appended to the valve's own. Each outcome adds a count row, and step 3's construction gives the posterior of the state after it: the MAP state $\hat{\xi}_j = \ln((4 + j)/(66 + M - j))$ and the variance $\sigma^2_{\xi,j} = (70 + M)/((4 + j)(66 + M - j))$, the record's (a_n, b_n) grown by the outcome. The rule is met after outcome j when the normal tail of equation (5.42) at the requirement reaches 0.95, and the probability of each outcome is the Laplace predictive (5.48), so

$$\Pr(\text{rule met} \mid \mathbf{y}) \approx \sum_{j=0}^{M} \Pr(K_f = j \mid \mathbf{y})\ \mathbf{1}\left[\,\Phi\left(\frac{\xi_{\max} - \hat{\xi}_j}{\sigma_{\xi,j}}\right) \geq 0.95\right],$$

for $p_{\max} = 0.10$ and 0.05. Exactly, the same sum runs over the Beta–binomial predictive (5.49) with the indicator $I_{p_{\max}}(4 + j,\ 66 + M - j) \geq 0.95$.

Table 5.12 ranks six candidate campaigns. Each row is a number M of further tests. The next two columns give the standard deviation of p expected after the campaign, from the Laplace preposterior (5.64) through g and exactly; the fourth the width that the campaign

removes from the current 0.02774; and the last three, on the Laplace route with the exact value in parentheses, the predictive probability that all M tests succeed, equations (5.48) and (5.50), and the probability that the campaign ends with the 90 % or the 95 % requirement met at 95 % assurance. The rows $M = 18$ and $M = 109$ are the smallest zero-failure campaigns that meet each rule on the Laplace route: after 18 successes the posterior of the state is $\hat{\xi} = \ln(4/84) = -3.0445$, $\sigma_\xi = 0.5118$, and the tail at $p_{\max} = 0.10$ is $\Phi(1.6556) = 0.9511$, while after 17 it is $\Phi(1.6318) = 0.9486$; after 109 successes the tail at $p_{\max} = 0.05$ is 0.9505, after 108 it is 0.9493. The exact route needs 7 and 84: $I_{0.10}(4, 73) = 0.9530$ while $I_{0.10}(4, 72) = 0.9496$; $I_{0.05}(4, 150) = 0.9506$ while $I_{0.05}(4, 149) = 0.9488$.

Table 5.12: Candidate test campaigns for the valve. Standard deviations of p after M further tests, Laplace (equation (5.64)) and exact; the width removed from the current 0.02774; and, from the posterior of the state with the exact Beta–binomial value in parentheses, the probability that all M succeed and that the campaign ends with $R \geq 0.90$ or $R \geq 0.95$ demonstrated at 95 % assurance.

Further tests M	Laplace sd of p	Exact sd of p	Width removed	All succeed	90 % rule met	95 % rule met
7	0.02645	0.02627	0.00129	0.645 (0.674)	0 (0.674)	0 (0)
10	0.02595	0.02577	0.00179	0.542 (0.575)	0 (0.575)	0 (0)
18	0.02474	0.02457	0.00300	0.352 (0.388)	0.352 (0.388)	0 (0)
50	0.02119	0.02104	0.00655	0.085 (0.109)	0.451 (0.675)	0 (0)
84	0.01870	0.01857	0.00904	0.026 (0.039)	0.592 (0.646)	0 (0.039)
109	0.01735	0.01723	0.01039	0.012 (0.021)	0.626 (0.680)	0.012 (0.021)

The standard deviation shrinks slowly, by the factor $\sqrt{70/(70 + M)}$ of equation (5.64), on both routes alike: fifty further tests remove less than a quarter of it. No campaign in the table is likely to demonstrate the 95 % requirement: on the Laplace route only a run of 109 successes demonstrates it, and the posterior gives such a run a probability of 0.012 (84 successes and 0.039 exactly). The 90 % requirement behaves differently. Eighteen successes demonstrate it, and the posterior gives eighteen successes a probability of 0.352; seven or ten tests cannot meet it at all on this route, because even a clean run leaves the tail below 0.95. Fifty tests tolerate some failures and end with the rule met with probability 0.451, eighty-four with 0.592: more tests raise the chance of settling the rule, at a cost that grows faster than the chance. The ranking is set by the tail at the requirement, not by the width: the campaign that removes the most width is not the one most likely to settle the decision, the point of equation (5.62) carried over to a valve that already has a record.

The exact route ranks the campaigns differently. Seven successes already demonstrate the 90 % requirement there, with probability 0.674, and fifty tests do no better, 0.675, at seven times the cost. The gap is the one the requirement question showed, 0.880 against 0.924 for the 90 % rule today: the normal in the logit puts too much weight on large failure probabilities, so it needs eleven more successes, 18 against 7, to push the tail past 0.95. Here the gap changes the plan, not only its numbers.

5.6.3 The recommendation

The next measurement is a campaign of eighteen further demands under the same procedure, specified in advance to pass only with zero failures. Under the engineering prior it ends with the 90 % requirement demonstrated at 95 % assurance with probability 0.352; a failure among the eighteen leaves the support at 0.908, and the rule then needs thirteen further zero-failure tests, thirty-one in all. The exact posterior gives the same requirement after only seven zero-failure demands, with probability 0.674; a failure among those seven leaves the exact support at $I_{0.10}(5, 72) = 0.888$, and the rule then needs twenty further tests in all. The two plans differ

by eleven demands; the eighteen-demand plan is the one the posterior of the state supports, and the seven-demand plan the one the exact posterior supports where, as here, it exists. The 95 % requirement is not a testing question for this valve on either route: the record already places more than half of the posterior above $p = 0.05$, and the run of successes that would demonstrate it has a posterior probability of 0.012 (0.039 exactly). Meeting it calls for a change to the valve or to its operating conditions, after which this record is no longer the evidence for the changed design.

5.6.4 What the analysis showed

The first reading of the record, 96 % successes against a 95 % specification, said the valve passes; the analysis overturns it. Two failures in fifty demands, read with the engineering prior, cannot demonstrate 95 % reliability at any practical assurance: the record leaves that requirement at 39.2 % support (45.5 % exactly), and the run of successes that would carry it to 95 % assurance has a posterior probability of only 0.012 (0.039). What the record does support is a reliability of about 90 % with fair confidence, at least 87.6 % with 95 % probability (89.1 %) and 88.0 % support for the 90 % requirement itself (92.4 %). The two requirements therefore call for different actions on site. The 90 % rule is a testing question, settled with probability 0.352 by eighteen further demands fixed in advance to pass only with zero failures; the exact posterior would settle it with seven, and the eleven demands between the two plans are what the normal in the logit costs on a record with two failures. The 95 % rule is not settled by any campaign of further demands; it is a design question, answered by a change to the valve or to its operating conditions. Meanwhile the probability that the valve opens on its next demand, 93.6 % (94.3 %), is read directly from the posterior and needs no further test.

5.6.5 Limits of the model

The derivation is exact under the assumptions of the Beta–binomial model. That does not make the assumptions appropriate, and five of them need to be checked before a result is used.

Comparability of the tests. A failure probability measured under normal operating conditions may not describe performance under severe heat, vibration or electrical stress. If the test conditions differ, a model in which each test has its own failure probability, $p_i = \Pr(X_i = 1 \mid T_i, V_i, \text{age}_i, \ldots)$, with T_i the temperature, V_i the supply voltage and age_i the age of the component at test i, may be needed.

Conditional independence of the tests. Repeated tests on the same component may not be independent if failures cause permanent damage, if aging changes the reliability, or if previous tests alter the component's condition. Shared manufacturing defects can also create dependence across a population of components.

Exposure hidden in binary observations. A valve tested for one demand and a pump operated continuously for 1000 h are different reliability observations. Under the exponential lifetime model with a constant failure rate λ, the failure probability over a mission of duration τ is

$$p_\tau = 1 - \mathrm{e}^{-\lambda\tau}. \tag{5.65}$$

For $\lambda = 0.001\,\mathrm{h}^{-1}$, the failure probability over 100 h is $p_{100} = 1 - \mathrm{e}^{-0.1} = 0.0952$, and a mission of 200 h has a different one, $1 - \mathrm{e}^{-0.2} = 0.1813$. Mission durations therefore cannot be ignored when time-dependent failures are modeled, and missions of different lengths are not Bernoulli trials with one common p.

Imperfect recording of failures. If an inspection fails to detect some genuine valve failures, or reports failures that did not occur, the observed binary outcome differs from the true physical outcome. The observation model must then account for missed detections and false indications; otherwise a Beta–binomial model applied directly to recorded failure counts may produce biased conclusions. Exercise 5.2 derives the likelihood of a reported count.

Change of the failure probability with time. Aging, maintenance, environmental exposure and operational stress can change reliability, and a constant-p model does not represent those dynamics. Time-dependent failure rates, survival distributions and models of aging are the subject of later chapters.

The next chapter moves from the failure probability per demand to failures in operating time, the Poisson process and the gamma posterior of a failure rate.

5.7 What to retain

Definition 5.1 (Beta–binomial estimation of a failure probability)**.** Under comparable tests, independent given a common failure probability p, a record of k failures in n demands has the binomial likelihood $\binom{n}{k}p^k(1-p)^{n-k}$. The state is the one number ξ, the logit of p; the prior $\mathrm{Beta}(a,b)$ and the record are two rows, a prior row of a pseudo-failures in $a+b$ trials and a count row, whose slopes $(a+b)p-a$ and $np-k$ and curvatures $(a+b)p(1-p)$ and $np(1-p)$ add, the curvature exact in the logit. The Newton step $\Delta\xi=-\mathcal{L}'/\mathcal{L}''$ reaches $\hat{\xi}=\ln(a_n/b_n)$, and the Laplace variance $\sigma_\xi^2=1/\Lambda$, with $g^2\sigma_\xi^2$ for p, is the fit's approximation of the posterior. The problem admits the exact posterior $\mathrm{Beta}(a+k,\,b+n-k)$, whose mean is $\hat{p}$; the answers are computed from the fit, and the exact values are set beside them with the gap. The estimate, its credible interval, the probability that a requirement $p\le 1-R_{\min}$ holds, the prediction of future tests and the size of a demonstration test are each a different function of the posterior, evaluated from $(\hat{\xi},\sigma_\xi)$ by a normal tail, a mapped quantile or a Gauss–Hermite sum and, beside it, in closed form in the two posterior parameters, and they must not be substituted for one another. The MAP state depends on which entry is the state: the mode of the density of p is 0.0441 for the valve, while the minimum in the logit maps to the posterior mean 0.0571.

Table 5.13 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 5.13: The terms of failure-probability estimation from test records and where each is defined. All quantities are dimensionless probabilities or counts unless stated.

Term	Expression	Units	Meaning	Defined in
Per-demand reliability	$R = 1 - p$	–	success on one demand	5.1.1
Logit state	$\xi = \ln(p/(1-p))$	–	unbounded entry for p	5.1.4
Questions				
Estimation	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$, $\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$, Var	–	posterior of a derived quantity	5.3.1
Threshold and assurance	$\Pr(f(\mathbf{z}) \ge f^* \mid \mathbf{y}) \ge \kappa$	–	tail of a derived quantity	5.3.2
Prediction	$\int p(y_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	–	future observation	5.3.3
Design	$\sum_j \Pr(K_n = j \mid \mathbf{y})\, \mathbf{1}[\text{rule met}]$	–	preposterior of a campaign	5.3.4
Answers				
Laplace density of p	$\varphi((\xi(p) - \hat{\xi})/\sigma_\xi)/(\sigma_\xi\, p(1-p))$	–	median $\hat{p}$	5.3.1
Posterior mode	$(a_n - 1)/(a_n + b_n - 2)$	–	maximum of the density of p	5.3.1
Gauss–Hermite expectation	$\pi^{-1/2} \sum_q w_q\, h(\hat{\xi} + \sqrt{2}\,\sigma_\xi\, t_q)$	–	Laplace expectation over ξ	5.3.1
Band of p	$g^2\sigma_\xi^2$, $g = \hat{p}(1-\hat{p})$	–	delta-method spread	5.3.1
Credible interval	$I_{p_L} = 0.025$, $I_{p_U} = 0.975$	–	95 % posterior probability	5.3.1
Requirement probability	$I_{1-R_{\min}}(a_n, b_n)$	–	support for $R \ge R_{\min}$	5.3.2
One-sided bound	$p_{0.95} = I^{-1}_{0.95}(a_n, b_n)$	–	$R_L = 1 - p_{0.95}$	5.3.2
Next-test success	$b_n/(a_n + b_n)$	–	posterior mean reliability	5.3.3
Beta–binomial predictive	$\binom{M}{j} B(a_n + j, b_n + M - j)/B(a_n, b_n)$	–	failures in M future tests	5.3.3
Predictive variance	$M\mu_p(1-\mu_p) + M(M-1)\,\mathrm{Var}(p)$	–	outcome plus parameter spread	5.3.3
Zero-failure sample size	$\lceil \ln(1-\kappa)/\ln(1-p_{\max}) - 1 \rceil$	tests	uniform prior, no failure	5.3.4
Operating characteristic	$(1-p)^n$ for a zero-failure plan	–	probability the plan passes	5.3.4
Rows and posterior				
Binomial likelihood	$\binom{n}{k} p^k (1-p)^{n-k}$, from $p^{x_i}(1-p)^{1-x_i}$	–	k failures in n tests	5.1.6
Count row	$\mathcal{L}_c = -k \ln p - (n-k)\ln(1-p)$	–	record as a negative log-probability term	5.1.6
Slope and curvature	$\mathcal{L}_c' = np - k$, $\mathcal{L}_c'' = np(1-p)$	–	curvature exact in the logit	5.1.6
Maximum-likelihood estimate	$\hat{p}_{\text{MLE}} = k/n$	–	observed failure fraction	5.1.6
Beta prior	$\mathrm{Beta}(a, b)$, mean $a/(a+b)$	–	a, b as pseudo-counts	5.1.6
Prior row	$\mathcal{L}_0 = -a \ln p - b \ln(1-p)$	–	density of ξ, a in $a+b$ trials	5.1.6
Newton step	$\Delta\xi = -\mathcal{L}'/\mathcal{L}''$	–	slopes and curvatures of the rows add	5.2.1
Information and variance	$\Lambda = (a_n + b_n)\hat{p}(1-\hat{p})$, $\sigma_\xi^2 = (a_n + b_n)/(a_n b_n)$	–	at $\hat{\xi} = \ln(a_n/b_n)$	5.2.1
Prior share	$\pi_0 = \mathcal{L}_0''/\mathcal{L}'' = (a+b)/(a+b+n)$	–	weight of the prior	5.2.1
Beta–binomial update	$\mathrm{Beta}(a+k,\, b+n-k)$	–	failures to a, successes to b	5.2.1
Laplace preposterior	$\sigma^2_{\xi,a} = (1/\sigma_\xi^2 + \Lambda_a)^{-1}$, $\Lambda_a = M\,\hat{p}(1-\hat{p})$	–	after M further tests	5.6.1
Preposterior acceptance	$\sum_j \Pr(K_f = j \mid \mathbf{y})\, \mathbf{1}[\text{rule met}]$	–	probability a campaign ends accepted	5.6.2
Mission failure probability	$1 - \mathrm{e}^{-\lambda\tau}$	–	rate and duration to p	5.6.5

5.8 Exercises

Exercise 5.1. *Formulation.* For the valve posterior $p \mid \mathbf{y} \sim \mathrm{Beta}(4, 66)$, write mathematical expressions for (a) the probability that the reliability exceeds 95 %, (b) the probability that it exceeds 90 %, (c) the probability that the next demand succeeds, and (d) the probability that all ten future demands succeed; then state why these are different questions. ((a) $\Pr(p < 0.05 \mid \mathbf{y}) = I_{0.05}(4, 66) = 0.4550$; (b) $I_{0.10}(4, 66) = 0.9238$; (c) $\Pr(X_{\text{next}} = 0 \mid \mathbf{y}) = 66/70 = 0.9429$; (d) $\Pr(K_f = 0 \mid \mathbf{y}) = B(4, 76)/B(4, 66) = 0.5754$. The first two concern the unknown underlying reliability, the last two future observable outcomes.)

Exercise 5.2. *Formulation.* (a) A continuously operating component follows an exponential lifetime with $\lambda = 0.001\,\mathrm{h}^{-1}$. Compute its failure probability over a 50 h and a 200 h mission, and state why the two mission outcomes should not automatically be modeled as Bernoulli tests with one common failure probability. (b) A test system detects 90 % of actual failures and falsely reports a failure for 2 % of successful operations. With p the probability of an actual failure and $Y = 1$ a reported failure, derive $\Pr(Y = 1 \mid p)$, and state why the ordinary Beta–binomial update applied to the reported count is generally incorrect. ((a) $1 - \mathrm{e}^{-0.05} = 0.04877$ and $1 - \mathrm{e}^{-0.2} = 0.18127$; the failure probability depends on the mission duration, so the two outcomes are not identically distributed unless their missions and operating conditions are comparable. (b) By total probability, $\Pr(Y = 1 \mid p) = 0.90\,p + 0.02\,(1 - p) = 0.02 + 0.88\,p$; the reported count is binomial with probability $0.02 + 0.88p$, its likelihood $(0.02 + 0.88p)^k (0.98 - 0.88p)^{n-k}$ is not proportional to $p^k (1 - p)^{n-k}$, and the conjugate update for the true p no longer applies.)

Exercise 5.3. *Derivation.* A component is tested ten times and fails twice. (a) Write the likelihood of one particular ordered sequence with two failures, and (b) the probability of exactly two failures regardless of their positions. (c) Derive the maximum-likelihood estimate in the logit ξ of p: write the count row's term, its slope and its curvature, find the minimum from the slope, and give the variance $1/\Lambda$ of ξ and the width in p by $g^2 \sigma_\xi^2$. (d) Explain why the binomial coefficient does not change the estimate. ((a) $p^2 (1 - p)^8$; (b) $\binom{10}{2} p^2 (1 - p)^8 = 45\, p^2 (1 - p)^8$; (c) $\mathcal{L}_c = -2 \ln p - 8 \ln(1 - p)$ with $p = p(\xi)$, slope $\mathcal{L}_c' = 10p - 2$ and curvature $\mathcal{L}_c'' = 10p(1 - p)$; the slope vanishes at $\hat{p} = 0.2$, $\hat{\xi} = \ln(2/8) = -1.3863$, where $\Lambda = 10 \times 0.2 \times 0.8 = 1.6$ and $\sigma_\xi^2 = 1/1.6 = 0.625$; with $g = 0.16$, $g^2 \sigma_\xi^2 = 0.16^2 \times 0.625 = 0.016$, a width of 0.126 in p; (d) the factor 45 does not depend on p, so it adds the constant $-\ln 45$ to $\mathcal{L}_c$ and changes neither its slope nor its minimizer.)

Exercise 5.4. *Derivation.* The prior is $p \sim \mathrm{Beta}(1, 9)$. (a) Twenty tests give two failures. Derive the posterior parameters, the posterior mean failure probability, the mean reliability and the mode. (b) Starting again from $\mathrm{Beta}(1, 9)$, a first campaign gives one failure in twenty tests and a second zero failures in thirty. Derive the posterior after each campaign and show that it equals the update with all fifty tests at once. (c) For the record of (a), in the logit ξ of p, give the slopes and curvatures of the prior row and the count row at the minimum, the Laplace variance σ_ξ^2, and the standard deviation of p by $g^2 \sigma_\xi^2$, against the exact one. ((a) $\mathrm{Beta}(1 + 2,\, 9 + 18) = \mathrm{Beta}(3, 27)$; mean $3/30 = 0.10$; mean reliability 0.90; mode $2/28 = 0.0714$. (b) $\mathrm{Beta}(2, 28)$, then $\mathrm{Beta}(2, 58)$; directly, $\mathrm{Beta}(1 + 1,\, 9 + 49) = \mathrm{Beta}(2, 58)$. (c) The minimum is $\hat{\xi} = \ln(3/27) = -2.1972$, $\hat{p} = 0.10$; both slopes vanish there, $10\hat{p} - 1 = 20\hat{p} - 2 = 0$, since the prior mean equals the observed fraction; the curvatures are 0.9 and 1.8, so $\Lambda = 2.7$ and $\sigma_\xi^2 = 0.3704$, $\sigma_\xi = 0.6086$; with $g = 0.09$ the standard deviation of p is 0.0548, against the exact 0.0539.)

Exercise 5.5. *Application.* A component has the posterior $p \mid \mathbf{y} \sim \mathrm{Beta}(3, 27)$, and five further tests are planned. Derive the predictive distribution of the number of failures K_f, and compute the expected number of failures and the probabilities of zero, exactly one and two or more failures. ($\Pr(K_f = j \mid \mathbf{y}) = \binom{5}{j} B(3 + j,\, 32 - j)/B(3, 27)$; $\mathbb{E}[K_f \mid \mathbf{y}] = 5 \times 3/30 = 0.5$;

$\Pr(K_f = 0) = \prod_{i=0}^{4}(27+i)/(30+i) = 0.6106$; $\Pr(K_f = 1) = 0.2955$; $\Pr(K_f \geq 2) = 1 - 0.6106 - 0.2955 = 0.0939$.)

Exercise 5.6. *Application.* The prior is uniform, $\mathrm{Beta}(1,1)$, the requirement is $R \geq 0.98$, and the customer demands at least 95 % posterior probability that it is satisfied. Assuming zero failures, derive the minimum number of tests. ($\Pr(p \leq 0.02 \mid 0, n) = 1 - 0.98^{n+1} \geq 0.95$ gives $n \geq \ln 0.05 / \ln 0.98 - 1 = 147.28$, so $n_{\min} = 148$; then $1 - 0.98^{149} = 0.9507$, while 147 tests give 0.9497.)

Chapter 6

Estimating Failure Rates from Operating Records

6.1 Failure rates from operating records

A cooling-water circulation pump has run 2000 operating hours and failed four times, each time repaired and returned to service. The facility's limit for the pump is a failure intensity below 0.003 failures per operating hour, and the first reading of the log is simple: four failures in two thousand hours is $0.002\,\mathrm{h}^{-1}$, comfortably below the limit, so the pump complies. The logs of three comparable pumps invite a second quick reading: pump C failed three times in 1200 h and pump B once in 800 h, so pump C is the worse pump. Two more logs sit in the same file: a pump that failed twice in 1000 h split between normal and high load, with no note of the load each failure fell in, and two logs with the pump's totals that keep the failure times, one spread evenly, the other bunched at the end.

The reliability engineer who must confirm the limit asks something else: how sure can the facility be that the pump itself, not its first two thousand hours, fails at less than 0.003 per hour? Four failures are few, and the experience with similar pumps, about $0.002\,\mathrm{h}^{-1}$, is itself uncertain. The first reading is a statistic of the log; the decision needs a statement about an unknown rate, and it depends on how much exposure there was, on what was known before it and on how sure the decision must be. The second reading takes a difference of counts for a difference of pumps, and whether one failure against three says anything about the pumps is itself a question for the logs. The question of this chapter is

how the rate at which a repairable component fails is estimated from a record of failures and operating exposure, how uncertain that estimate is, whether it meets a stated limit, what it predicts for future failures, and whether comparable components share one rate.

It is answered with the method of Chapter 5: a graph, the evidence attached to it as rows, a posterior over the state, and each question read from that posterior. What changes is the evidence, a count of failures in hours of service in place of successes and failures on demand, and with it the likelihood, now a Poisson count, and the state, which carries the logarithm of a rate.

Jargon: failure intensity, rate of occurrence of failures (ROCOF), hazard rate, mission reliability, exposure, homogeneous Poisson process, gamma distribution, shape and rate, Gamma–Poisson update, negative-binomial distribution, Lomax distribution, mean time between failures (MTBF), sufficient statistic, effective exposure, tie row, nonhomogeneous Poisson process, Laplace trend test

The *failure intensity* of a repairable item, also called the *rate of occurrence of failures* (ROCOF), is the rate of change of the expected number of failures with operating time. The *hazard rate* is the instantaneous failure tendency of one item that has survived to a given age. *Mission reliability* is the probability of no failure during a specified mission. *Exposure* is the accumulated operating time behind a failure count, excluding the time spent under repair. A *homogeneous Poisson process* is a model of events in time with a constant intensity and independent counts in disjoint intervals. The *gamma distribution* is a family of densities on the positive numbers, the natural prior for a rate; it is written here with a *shape* a and a *rate* b, and the *Gamma–Poisson update* is its conjugate update by a Poisson count. The *negative-binomial distribution* is the distribution of a Poisson count whose mean is itself gamma distributed; the *Lomax distribution* is the corresponding distribution of the time to the first event. The *mean time between failures* (MTBF) is the average operating interval between successive failures of a repairable item. A *sufficient statistic* is a summary of the record through which alone the likelihood depends on the unknown. *Effective exposure* is operating time weighted by the relative intensity of the conditions it was spent in. A *tie row* is a Gaussian row that holds two entries of the state close to each other within a stated tolerance. A *nonhomogeneous Poisson process* lets the intensity change with time. The *Laplace trend test* checks a record for a trend in its event times.

Key idea

A record of k failures in an exposure of T_{tot} operating hours is one count row on the log-rate ω, and a Gamma(a, b) prior is a count row of a pseudo-failures in b pseudo-hours; the information the two rows give the state is the sum of their expected counts, $\Lambda = a + k$ at the MAP state, and the exact posterior of the rate is Gamma$(a + k,\ b + T_{\text{tot}})$. The count and the exposure are all the record contributes to the rate. The estimate, its band, the probability that a rate requirement is met, the predicted number of failures, the mission reliability and the time to the next failure are each a different function of the same posterior, and they are different numbers because they answer different questions. Where pumps may differ or a load multiplier is uncertain, the same rows on a larger state carry answers that no gamma form gives.

6.1.1 Failure probability, failure intensity and hazard

Three quantities are connected through a constant rate λ under certain assumptions, and they are not interchangeable in general.

Failure probability. A failure probability concerns an event within a specified interval. For example, $\Pr(T \le 100\,\text{h})$ is the probability that a component fails within its first 100 operating hours, where T is its time to failure. A probability is dimensionless and lies between zero and one, and it means nothing until the interval is stated.

Failure intensity. For a repairable item the record is a counting process. Let $N(t)$ be the number of failures recorded by the operating time t, and let

$$m(t) = \mathbb{E}[N(t)]$$

be the expected number of failures by that time. The rate of change of the expected count,

$$\nu(t) = \frac{\mathrm{d}m(t)}{\mathrm{d}t}, \tag{6.1}$$

is the *failure intensity*, or rate of occurrence of failures, wherever the derivative exists. It has units of inverse time, failures per operating hour. For the homogeneous Poisson process of Section 6.1.3, the intensity is a constant, $\nu(t) = \lambda$, and the expected count grows linearly,

$$\mathbb{E}[N(t)] = \lambda t. \tag{6.2}$$

Hazard rate. The hazard belongs to one time to failure T, not to a sequence of failures. It is defined as the limit

$$\lambda(t) = \lim_{\Delta t \to 0^+} \frac{\Pr(t \le T < t + \Delta t \mid T \ge t)}{\Delta t}, \tag{6.3}$$

the probability of failing in a short interval after t, given survival to t, per unit length of the interval. It describes the instantaneous conditional failure tendency of an item that has survived to age t. A hazard is not itself a probability: its units are inverse time, and its numerical value, which depends on the unit of time, may exceed one.

6.1.2 Hazard, survival and the constant rate

The reliability, or survival function, of the time to failure is $R(t) = \Pr(T > t)$, and the failure density is its negative derivative, $f(t) = -\mathrm{d}R(t)/\mathrm{d}t$. Two general relations connect them with the hazard of equation (6.3):

$$\lambda(t) = \frac{f(t)}{R(t)}, \qquad R(t) = \exp\!\left[-\int_0^t \lambda(u)\,\mathrm{d}u\right] = \mathrm{e}^{-\Lambda(t)}, \tag{6.4}$$

with $\Lambda(t) = \int_0^t \lambda(u)\,\mathrm{d}u$ the cumulative hazard. They hold for every continuous lifetime distribution, and Chapter 7 derives both from the definition of the hazard; this chapter needs only the case of a constant hazard.

The constant hazard. If the hazard is constant, $\lambda(t) = \lambda$, then $\Lambda(t) = \lambda t$ and

$$R(t) = \mathrm{e}^{-\lambda t}, \qquad \Pr(T \le t) = 1 - \mathrm{e}^{-\lambda t}, \tag{6.5}$$

the exponential lifetime relations. The mean of a nonnegative time is the integral of its survival function, a standard identity obtained by integrating $\mathbb{E}[T] = \int_0^\infty t\, f(t)\,\mathrm{d}t$ by parts, so that under the constant hazard

$$\mathrm{MTTF} = \int_0^\infty \mathrm{e}^{-\lambda t}\,\mathrm{d}t = \frac{1}{\lambda}. \tag{6.6}$$

Where intensity and hazard coincide. The same constant λ can describe the event intensity of a homogeneous Poisson process and the hazard of the time to the next failure, provided each repair restores the pump to a condition statistically identical to the one before the failure. Section 6.1.3 shows that under that process the probability of no failure in any interval of length t is $\mathrm{e}^{-\lambda t}$, which is equation (6.5): the time from any instant to the next failure is exponential with hazard λ. For more general repairable systems, whose repairs leave them better or worse than before or whose intensity drifts with age, the failure intensity of the process and the hazard of a single lifetime need not coincide.

Table 6.1 sets the five quantities side by side: each row is a quantity, the middle column its mathematical meaning and the right column its units. Two of them are probabilities, two are rates and one is a time. For repairable equipment the mean time between failures (MTBF) requires a clear definition of the intervals measured: whether repair downtime is included changes its interpretation, and only when it is excluded does it equal $1/\lambda$ under a constant intensity.

Table 6.1: Five quantities that the words *rate* and *reliability* are used for. The middle column states each one's meaning; under a constant hazard they are connected through λ by equations (6.2), (6.5) and (6.6).

Quantity	Mathematical meaning	Units
Failure probability	probability of a failure within a specified period, $\Pr(T \le t)$	dimensionless
Hazard rate $\lambda(t)$	conditional instantaneous failure tendency at age t	h^{-1}
Failure intensity $\nu(t)$	rate of change of the expected number of failures, $\mathrm{d}m/\mathrm{d}t$	h^{-1}
Mission reliability $R(\tau)$	probability of no failure during a mission of length τ	dimensionless
MTTF	mean time to failure, $1/\lambda$ under a constant hazard	h

6.1.3 The homogeneous Poisson process

Let $N(t)$ be the number of failures in t hours of operating exposure. A homogeneous Poisson process with intensity λ is defined by four assumptions:

1. $N(0) = 0$;
2. counts in disjoint intervals are independent;
3. the distribution of the count in an interval depends only on the length of the interval;
4. in a short interval of length Δt the probability of one failure is $\lambda\Delta t + o(\Delta t)$ and the probability of two or more is $o(\Delta t)$, where $o(\Delta t)$ denotes a term with $o(\Delta t)/\Delta t \to 0$ as $\Delta t \to 0$.

The distribution of the count follows from them.

Division into short intervals. Divide an interval of length t into J equal pieces of duration $\Delta t = t/J$. By assumption 4, for large J each piece contains one failure with probability close to

$$p_J = \lambda \frac{t}{J},$$

and two or more with a probability that is negligible; by assumptions 2 and 3 the pieces are independent and alike. The count is then approximately the number of pieces that contain a failure, a binomial count of J trials with probability p_J:

$$\Pr(N(t) = k) \approx \binom{J}{k} \left(\frac{\lambda t}{J}\right)^k \left(1 - \frac{\lambda t}{J}\right)^{J-k}.$$

The approximation makes two errors in each piece: it replaces the one-failure probability $\lambda\Delta t + o(\Delta t)$ by $\lambda\Delta t$, and it ignores two or more failures, whose probability is $o(\Delta t)$. Summed over the J pieces, both are at most $J\,o(t/J)$, which vanishes as $J \to \infty$.

The limit. Write the binomial coefficient as a product and separate the factors that depend on J:

$$\binom{J}{k} \left(\frac{\lambda t}{J}\right)^k \left(1 - \frac{\lambda t}{J}\right)^{J-k} = \frac{(\lambda t)^k}{k!} \underbrace{\frac{J(J-1)\cdots(J-k+1)}{J^k}}_{\to 1} \underbrace{\left(1 - \frac{\lambda t}{J}\right)^{J}}_{\to \mathrm{e}^{-\lambda t}} \underbrace{\left(1 - \frac{\lambda t}{J}\right)^{-k}}_{\to 1}.$$

The first underbraced factor is $\prod_{i=0}^{k-1}(1 - i/J)$, a fixed number k of factors each tending to one; the second tends to $\mathrm{e}^{-\lambda t}$ by the limit $(1 + x/J)^J \to \mathrm{e}^x$; the third tends to one because its exponent is fixed. Therefore

$$\Pr(N(t) = k \mid \lambda) = \frac{(\lambda t)^k \, \mathrm{e}^{-\lambda t}}{k!}, \qquad k = 0, 1, 2, \ldots, \tag{6.7}$$

the Poisson distribution, written

$$N(t) \mid \lambda \sim \text{Poisson}(\lambda t). \tag{6.8}$$

The probabilities sum to one, since $\sum_k (\lambda t)^k / k! = \mathrm{e}^{\lambda t}$.

The moments of the count. Write $\mu_N = \lambda t$ for the Poisson mean. The mean of the count follows by cancelling k against $k!$ and shifting the index:

$$\mathbb{E}[N(t) \mid \lambda] = \sum_{k=1}^{\infty} k \, \frac{\mu_N^k \mathrm{e}^{-\mu_N}}{k!} = \mu_N \sum_{k=1}^{\infty} \frac{\mu_N^{k-1} \mathrm{e}^{-\mu_N}}{(k-1)!} = \mu_N.$$

The same step twice gives the factorial moment $\mathbb{E}[N(t)(N(t) - 1) \mid \lambda] = \mu_N^2$, and the variance is $\mathbb{E}[N(N-1)] + \mathbb{E}[N] - \mathbb{E}[N]^2 = \mu_N^2 + \mu_N - \mu_N^2$. Hence

$$\mathbb{E}[N(t) \mid \lambda] = \lambda t, \qquad \mathrm{Var}(N(t) \mid \lambda) = \lambda t. \tag{6.9}$$

Under the model, the conditional mean and variance of the count are equal. That equality is a first check of a record against the homogeneous Poisson model: counts over equal exposures that scatter much more than their mean signal a rate that is not constant (Section 6.4.6).

The probability of no failure. Setting $k = 0$ in equation (6.7), with $(\lambda t)^0 = 0! = 1$,

$$\Pr(N(t) = 0 \mid \lambda) = \mathrm{e}^{-\lambda t}. \tag{6.10}$$

This is the exponential survival function of equation (6.5). By assumptions 2 and 3 it holds for any interval of length t, wherever it starts, so the probability of no failure in an interval and the probability that the time to the next failure exceeds the interval are the same mathematical quantity.

6.1.4 The physical system and its graph (step 0)

A cooling-water circulation pump in an industrial facility draws from a water supply and delivers to a cooling network. It operates continuously when available and is repaired or replaced after each failure. The interest here is the pump's failure process, not yet the hydraulic behavior of the network it serves. The record is the counting process

$$N(t) = \text{number of recorded pump failures by operating exposure } t.$$

The initial model rests on five assumptions: the pump is restored to service after each failure; failures occur at a constant intensity during accumulated operating time; the failure mechanism does not change systematically with age or operating history; the record contains correctly identified failure events; and only actual operating exposure is counted, repair downtime being excluded. These assumptions make the record a homogeneous Poisson process, Section 6.1.3. They are modeling assumptions, not facts implied by the existence of a failure record, and Section 6.4.6 tests them against the record itself.

Figure 6.1(a) draws the pump as a graph. The gray boxes on the left and right are the water supply and the cooling network, pinned boundaries of the hydraulic path, which is drawn

as solid arrows. The green box in the middle is the pump, the free node of the graph; it carries the one entry of the state. The gray box above it is the operating exposure, pinned because the hours of operation are read from the plant log and are not uncertain. The two dashed arrows leave the pump for the evidence: the failure record below, a count of failures in that exposure, and the prior on the rate to the right, which carries the experience with similar pumps. Figure 6.1(b) draws the fleet of three comparable pumps described below: the green box at the top carries the shared log-rate of the fleet, and the gray arrows carry it into each pump's rate; each pump below carries its own deviation from that rate; the dotted boxes beside the pumps are the tie rows, which act on the deviations alone and hold them near zero; and each pump's count record is reached by a dashed arrow, a count row on the pump's log-rate, the shared entry plus its deviation.

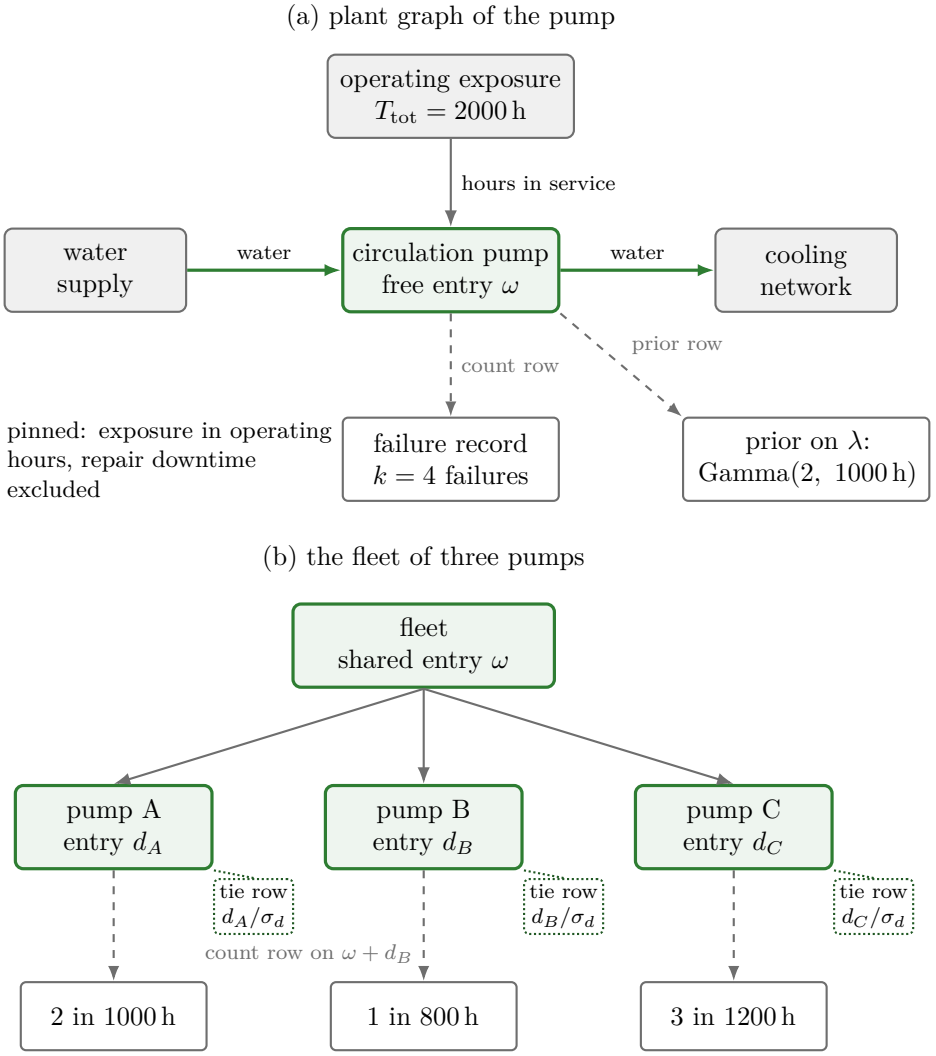


Figure 6.1: The plant graphs. (a) The pump. Free: the pump node, carrying the log-rate $\omega = \ln(\lambda/\lambda_{\text{ref}})$ (green). Pinned: the water supply, the cooling network and the operating exposure T_{tot} (gray). Observed: the number of failures k in that exposure (dashed arrow, below); the rate carries the prior from similar pumps (dashed arrow, right). (b) The fleet of three comparable pumps. Free: the shared log-rate ω and the deviation d_i of each pump (green); pump i fails at the rate $\lambda_{\text{ref}} e^{\omega + d_i}$, the shared entry entering through the gray arrows. Observed: each pump's failures in its own exposure, a count row on $\omega + d_i$ (dashed). Dotted: the tie rows d_i/σ_d, on the deviations alone; the prior row on ω is not drawn.

The state entry. A rate must be positive, and the graph carries it on an unbounded scale, as the notation of this book carries a conductance by the log-conductance g. The entry of the state is the log-rate

$$\omega = \ln \frac{\lambda}{\lambda_{\text{ref}}}, \qquad \lambda = \lambda_{\text{ref}}\, \mathrm{e}^{\omega}, \tag{6.11}$$

with $\lambda_{\text{ref}} = 10^{-3}\,\text{h}^{-1}$ a fixed reference rate, one failure per thousand operating hours, and the whole state of the graph is this one number. The general forms of the book write the state as a vector $\mathbf{z}$; here it has the single entry ω, so every quantity of the pump's fit is a number: the objective's slope and curvature in ω, the information Λ, the posterior variance $\sigma_\omega^2 = 1/\Lambda$ of the log-rate, and the derivative g of a derived quantity. Since ω and λ determine each other one to one, a probability statement about one is a probability statement about the other. The rate itself, under a gamma prior, has an exact posterior in closed form, which this problem admits; it stands beside the posterior of the state (Section 6.2.2). Every quantity the questions ask for, the rate among them, is a function of the state.

Three extended graphs. The fleet of Figure 6.1(b) gives each pump its own log-rate $\omega + d_i$, the shared log-rate ω plus the deviation d_i of pump i, so that pump i fails at the rate $\lambda_i = \lambda_{\text{ref}} \mathrm{e}^{\omega + d_i}$, and the state is the four-entry column vector

$$\mathbf{z} = \begin{pmatrix} \omega & d_A & d_B & d_C \end{pmatrix}^{\mathsf{T}}. \tag{6.12}$$

The pump that ran under two loads has the graph of Figure 6.1(a) with the exposure split by condition. Equipment may operate under different levels of stress, and its failure intensity may be larger during high-temperature or high-load operation; one model scales a reference intensity by a factor for each condition,

$$\lambda_i = \lambda_0\, c_i, \tag{6.13}$$

where λ_0 is the reference failure intensity, c_i the dimensionless relative intensity of condition i, and h_i the operating exposure spent in condition i. With the factors known, the state is the one entry $\omega_0 = \ln(\lambda_0/\lambda_{\text{ref}})$; with the high-load factor c_2 known only roughly, the state is $\mathbf{z} = (\omega_0,\ \ln c_2)^{\mathsf{T}}$. The third extension lets the intensity change with operating age. A common model of aging is the power law

$$\nu(t) = \theta\, \beta\, t^{\beta - 1}, \qquad m(t) = \theta\, t^{\beta}, \qquad \theta > 0,\ \beta > 0, \tag{6.14}$$

whose intensity increases with operating age when $\beta > 1$, decreases when $\beta < 1$, and is the constant θ when $\beta = 1$. Both unknowns are positive, and the state carries them on unbounded scales chosen so that each row informs one entry: the log of the expected number of failures over the record, $\ln m(T_{\text{tot}}) = \ln(\theta T_{\text{tot}}^{\beta})$, and the log-shape $v = \ln \beta$, as for the Weibull shape in the notation chapter, $\mathbf{z} = (\ln m(T_{\text{tot}}),\ v)^{\mathsf{T}}$.

6.1.5 The operating records

The maintenance records of the pump show $k = 4$ failures in $T_{\text{tot}} = 2000$ operating hours, and the period of interest for the questions is the next $\tau = 500$ operating hours. The records are illustrative. Table 6.2 lists the pump's record and two of the others: each row is a period, a pump or a condition, the second column its operating exposure h_i, the third its failures k_i and the last the relative intensity c_i of a condition. The top block divides the pump's record into its four operating periods. The middle block holds the fleet of three comparable pumps of Figure 6.1(b). The bottom block is a pump that ran 800 h under normal load and 200 h under high load, where the intensity is taken to be twice the normal one, and failed twice; the record does not say under which load each failure occurred.

Table 6.2: The operating records. Top: the pump's four operating periods. Middle: three comparable pumps. Bottom: one pump under two operating conditions, with the relative intensity c_i of each condition.

Record	Operating exposure h_i [h]	Failures k_i [–]	Relative intensity c_i [–]
Period 1	500	1	–
Period 2	700	2	–
Period 3	300	0	–
Period 4	500	1	–
Total	2000	4	–
Pump A	1000	2	–
Pump B	800	1	–
Pump C	1200	3	–
Total	3000	6	–
Normal load	800	–	1.0
High load	200	–	2.0
Total	1000	2	–

The pump's log keeps counts by period, not the failure times. Two illustrative records of failure times with the pump's totals, four failures in 2000 operating hours, show what the times carry. Figure 6.2 draws them. In panel (a) each horizontal line is the operating time of one record from 0 to 2000 h and each tick a failure: record A has its failures at 320, 770, 1320 and 1830 h; record B at 1440, 1550, 1650 and 1760 h. Panel (b) plots the cumulative number of failures against operating time for both records, with the straight line through the records' totals, four failures at 2000 h.

The evidence $\mathbf{y}$ of the pump is its record (k, T_{tot}) together with the prior; that of each further record is its own counts or times with the priors of Section 6.1.7. Densities of the rate are written $p(\lambda)$ and $p(\lambda \mid \mathbf{y})$.

6.1.6 The probabilistic questions (step 1)

The reliability engineer asks eleven questions of the records, and each is written as a probability, an expectation or a ratio of evidences, with the unknown on the left of the bar and the record on the right. Table 6.3 lists them: the first column is the question in engineering language; the second the general form that every question of its category takes, written in the book's notation for a quantity $f(\mathbf{z})$ derived from the state, a threshold f^* on it, an assurance level κ, the posterior density $p(\mathbf{z} \mid \mathbf{y})$, a future observation y_{new} and the evidence $p(\mathbf{y} \mid M)$ of a drawing M; the third its mathematical form for the record; the fourth its category, the predictions of future failures being filed, as in Chapter 2, under prediction under a threshold; and the last the subsection where it is answered. The atlas of Appendix A lists the same forms.

Whether the pump's rate is constant in time is not a row of the table: the power law carries an implicit flat prior on its log-shape, so its evidence is undefined and the comparison with the constant rate waits for a proper prior on the shape; the chapter checks the constant rate against the failure times instead (Section 6.4.6), and estimates the shape in the eleventh question.

6.1.7 The evidence as rows (step 2)

Step 2 attaches the evidence to the graphs as rows. On the pump's one entry two rows act: a prior row, which carries the experience with similar pumps as a gamma density on

Table 6.3: The questions asked of the operating records, the general form of their category, their mathematical form, their category and where each is answered. $\lambda_{\max}$ is the largest acceptable rate; K_f the number of failures in the next $\tau = 500$ operating hours; T the time from now to the next failure; λ_B pump B's own rate; λ_0 and λ_{high} the rates under normal and high load; M_1 and M_0 the drawings of the fleet with one shared rate and with a rate per pump; β the power-law shape.

Engineering question	General form	Mathematical form	Category	Answered in
What is the failure intensity?	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$, moments	$p(\lambda \mid k, T_{\text{tot}})$, $\mathbb{E}[\lambda \mid k, T_{\text{tot}}]$	estimation	6.3.1
How certain is the estimate?	$\mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y})$, quantiles	$\mathrm{Var}(\lambda \mid k, T_{\text{tot}})$, credible interval	estimation	6.3.1
Is the intensity below a target?	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y}) \ge \kappa$	$\Pr(\lambda < 0.003\,\text{h}^{-1} \mid k, T_{\text{tot}}) \ge 0.95$	prediction under a threshold	6.3.1
How many failures in the next 500 h?	$\int p(y_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	$\Pr(K_f = j \mid k, T_{\text{tot}})$	prediction under a threshold	6.3.2
What is the expected number?	$\mathbb{E}[y_{\text{new}} \mid \mathbf{y}]$, variance	$\mathbb{E}[K_f \mid k, T_{\text{tot}}]$, $\mathrm{Var}(K_f \mid k, T_{\text{tot}})$	prediction under a threshold	6.3.2
No failure in the next 500 h?	$\int p(y_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	$\Pr(K_f = 0 \mid k, T_{\text{tot}}) = \Pr(T > 500\,\text{h} \mid k, T_{\text{tot}})$	prediction under a threshold	6.3.2
How long until the next failure?	$\int p(y_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$, moments	$p(T \mid k, T_{\text{tot}})$, $\mathbb{E}[T \mid k, T_{\text{tot}}]$	prediction under a threshold	6.3.2
What is the rate of the three pumps, and pump B's own?	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$, $f = \mathbf{g}^{\mathsf{T}}\mathbf{z}$	$p(\omega \mid \mathbf{y})$; $p(\omega + d_B \mid \mathbf{y})$, λ_B	estimation	6.3.3
Are the three pumps alike?	$p(\mathbf{y} \mid M_1)/p(\mathbf{y} \mid M_0)$	$p(\mathbf{y} \mid \text{shared})/p(\mathbf{y} \mid \text{separate})$	hypothesis comparison	6.3.3
What is the rate under normal and under high load?	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$, quantiles	$p(\lambda_0 \mid \mathbf{y})$, $p(\lambda_{\text{high}} \mid \mathbf{y})$	estimation	6.3.3
What is the power-law shape of each record?	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$, quantiles	$p(\beta \mid t_1, \ldots, t_4)$; $\beta = 1$ in the band?	estimation	6.3.4

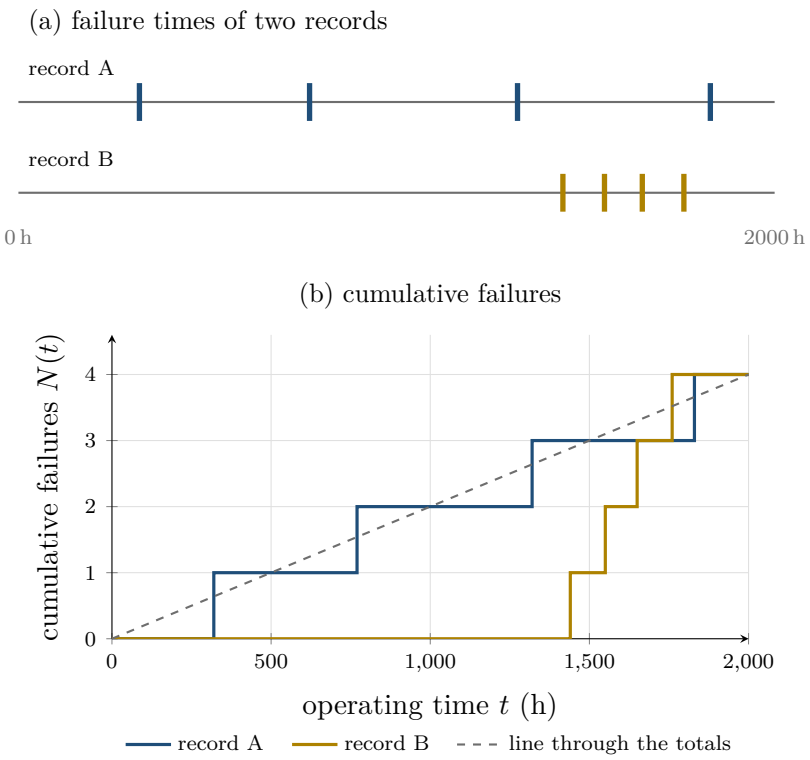


Figure 6.2: Two illustrative records with the same count and exposure, four failures in 2000 operating hours. (a) Failure times: record A (blue), record B (gold). (b) Cumulative failures against operating time, with the line through the records' totals, $4t/2000\,\mathrm{h}$.

the rate, and a count row, which carries the record of k failures in T_{tot} hours through the Poisson probability (6.7). Neither row is Gaussian; each enters the objective as its negative log-probability term, the row of Chapter 4, Section 4.2.7, a function of the one number ω. What a row contributes to the posterior is its slope and its curvature in ω, written with primes, $\mathcal{L}_k' = \mathrm{d}\mathcal{L}_k/\mathrm{d}\omega$ and $\mathcal{L}_k'' = \mathrm{d}^2\mathcal{L}_k/\mathrm{d}\omega^2$; each row of the pump below is given by its term, its slope and its curvature.

The modeling choices. Three choices make the rows, and each buys something and can fail. The four failures are treated as events of one constant intensity, each repair returning the pump to the condition it had before: this is what lets the record be summarized by its count and its exposure alone, and it fails if the intensity grows with age or a repair leaves the pump worse, which the trend and dispersion checks of Section 6.4.6 test. The exposure counts hours in service only, and the count only correctly identified failures: this makes λ a rate per operating hour, and it is broken by repair downtime logged as service or by failures missing from the log (Section 6.6.5). The experience with similar pumps is carried as the prior $\mathrm{Gamma}(2, 1000\,\mathrm{h})$, two failures in a thousand hours: this gives the posterior a closed form and lets the experience weigh half as much as the record, and it misleads if those pumps were unlike this one, which Section 6.4.3 measures.

The gamma prior. The unknown rate satisfies $\lambda > 0$, so its prior must be a distribution on the positive numbers. The natural choice is the gamma distribution in the shape–rate parameterization used throughout this chapter,

$$\lambda \sim \mathrm{Gamma}(a, b), \qquad p(\lambda) = \frac{b^a}{\Gamma(a)}\,\lambda^{a-1}\,\mathrm{e}^{-b\lambda}, \qquad \lambda > 0,\ a, b > 0, \tag{6.15}$$

where the shape a is dimensionless and the rate b has units of time, so that $b\lambda$ is dimensionless. The parameters are often written α_0 and β_0 elsewhere; this book writes them a and b, as for the Beta distribution of the previous chapter, keeping β for a shape of a lifetime or an intensity. Every exact result of the chapter rests on one identity. Substituting $u = b\lambda$, so that $\lambda = u/b$ and $\mathrm{d}\lambda = \mathrm{d}u/b$,

$$\int_0^\infty \lambda^{a-1}\,\mathrm{e}^{-b\lambda}\,\mathrm{d}\lambda = \frac{1}{b^a}\int_0^\infty u^{a-1}\,\mathrm{e}^{-u}\,\mathrm{d}u = \frac{\Gamma(a)}{b^a}, \tag{6.16}$$

where $\Gamma(a) = \int_0^\infty u^{a-1}\mathrm{e}^{-u}\,\mathrm{d}u$ is the gamma function of the previous chapter, with $\Gamma(a+1) = a\,\Gamma(a)$ and $\Gamma(m) = (m-1)!$ for a positive integer m; the identity fixes the constant $b^a/\Gamma(a)$ in equation (6.15). Multiplying the density by λ raises the exponent by one, and equation (6.16) with $a+1$ in place of a gives $\mathbb{E}[\lambda] = (b^a/\Gamma(a))\,\Gamma(a+1)/b^{a+1} = a/b$; in the same way $\mathbb{E}[\lambda^2] = a(a+1)/b^2$, and $\mathrm{Var}(\lambda) = \mathbb{E}[\lambda^2] - \mathbb{E}[\lambda]^2$ simplifies to

$$\mathbb{E}[\lambda] = \frac{a}{b}, \qquad \mathrm{Var}(\lambda) = \frac{a}{b^2}, \qquad \frac{\mathrm{sd}(\lambda)}{\mathbb{E}[\lambda]} = \frac{1}{\sqrt{a}}. \tag{6.17}$$

The ratio a/b sets the mean and the shape alone sets the coefficient of variation: a gamma density is sharp in proportion to $\sqrt{a}$, whatever its mean. In the update of Section 6.2.2 the observed failures add to a and the observed exposure adds to b, so that a acts like a prior number of failures and b like a prior number of operating hours; the reading is a convenience for judging how strong a prior is, not a claim that a literal historical record lies behind it. Experience with similar pumps suggests an average failure intensity of about $0.002\,\mathrm{h}^{-1}$, with substantial uncertainty, and the prior chosen is

$$\lambda \sim \mathrm{Gamma}(2,\ 1000\,\mathrm{h}), \tag{6.18}$$

which as pseudo-data weighs as much as two failures in a thousand hours.

The likelihood of the record. Under the homogeneous Poisson process, by equation (6.8), the count model is $K \mid \lambda, T_{\mathrm{tot}} \sim \mathrm{Poisson}(\lambda T_{\mathrm{tot}})$, and the likelihood of a record of k failures in T_{tot} operating hours is

$$p(k \mid \lambda, T_{\mathrm{tot}}) = \Pr(K = k \mid \lambda, T_{\mathrm{tot}}) = \frac{(\lambda T_{\mathrm{tot}})^k\,\mathrm{e}^{-\lambda T_{\mathrm{tot}}}}{k!}, \tag{6.19}$$

read as a function of λ. The data constrain the rate through both numbers, the count and the exposure: the same failures in twice the hours would point to half the rate. The likelihood is a function of λ with the record held fixed; it does not integrate to one over λ and it carries no prior information.

The maximum-likelihood estimate. The log-likelihood is

$$\ell(\lambda) = k\ln\lambda + k\ln T_{\mathrm{tot}} - \lambda T_{\mathrm{tot}} - \ln(k!).$$

Differentiate with respect to λ; the terms in T_{tot} and $k!$ alone do not depend on it:

$$\frac{\mathrm{d}\ell}{\mathrm{d}\lambda} = \frac{k}{\lambda} - T_{\mathrm{tot}}.$$

Setting the derivative to zero gives $k/\hat{\lambda}_{\mathrm{MLE}} = T_{\mathrm{tot}}$, and therefore

$$\hat{\lambda}_{\mathrm{MLE}} = \frac{k}{T_{\mathrm{tot}}}. \tag{6.20}$$

The second derivative,

$$\frac{\mathrm{d}^2\ell}{\mathrm{d}\lambda^2} = -\frac{k}{\lambda^2}, \qquad \left.\frac{\mathrm{d}^2\ell}{\mathrm{d}\lambda^2}\right|_{\lambda=k/T_{\mathrm{tot}}} = -\frac{T_{\mathrm{tot}}^2}{k},$$

is negative everywhere, so the stationary point is the maximum, and its magnitude corresponds to the width $k^{1/2}/T_{\mathrm{tot}} = \hat{\lambda}_{\mathrm{MLE}}/\sqrt{k}$: k events fix the rate only to within about $1/\sqrt{k}$ of its value. A point estimate is not a perfectly known rate. When $k = 0$ the estimate is $\hat{\lambda}_{\mathrm{MLE}} = 0$, on the boundary of the admissible rates. The posterior represents the uncertainty that the point estimate hides.

The count row. On the state, with $\lambda = \lambda_{\mathrm{ref}}\mathrm{e}^{\omega}$ of equation (6.11), the count row puts in the objective the negative logarithm of the likelihood (6.19), without the terms $k\ln(T_{\mathrm{tot}}\lambda_{\mathrm{ref}})$ and $\ln k!$ that do not depend on the state: $\mathcal{L}_c(\omega) = T_{\mathrm{tot}}\lambda - k\omega$. It is the count row of Chapter 4, equation (4.42), the term $\mu - k\ln\mu$ of a record of k events with the expected count $\mu = T_{\mathrm{tot}}\lambda$. Written for a count row i of k_i failures in an exposure h_i, the pump's record being the case $h_i = T_{\mathrm{tot}}$, $k_i = k$, its expected count is $\mu_i = h_i\lambda_{\mathrm{ref}}\mathrm{e}^{\omega}$, whose logarithm $\ln(h_i\lambda_{\mathrm{ref}}) + \omega$ has the slope one in ω and no curvature, so that the row's slope and curvature are

$$\mathcal{L}_i' = \mu_i - k_i, \qquad \mathcal{L}_i'' = \mu_i, \qquad \mu_i = h_i\lambda_{\mathrm{ref}}\,\mathrm{e}^{\omega}. \tag{6.21}$$

The slope is the expected count minus the observed count: the row pulls ω down where the model expects more failures than were seen and up where it expects fewer. The curvature is the expected count, the variance of the Poisson count and its Fisher information in the log of its mean. It does not contain the observed k_i, and that is a property of the log-rate, the canonical parameter of the Poisson distribution, the quantity in which its log-probability is linear: in another coordinate, λ itself among them, the curvature carries the extra term $(\mu_i - k_i)$ times the second derivative of $\ln\mu_i$ of equation (4.42).

The count row on a state of several entries. The fleet and the two loads of Section 6.1.4 carry several entries, and there a count row joins the Jacobian beside the Gaussian tie and factor rows in the working form of Section 4.2.7, equation (4.43), which for a count row of this chapter reads

$$r_i = \frac{\mu_i - k_i}{\sqrt{\mu_i}}, \qquad \mathbf{J}_i = \sqrt{\mu_i}\,\frac{\partial \ln\mu_i}{\partial \mathbf{z}}^{\mathsf{T}}. \tag{6.22}$$

The pair reproduces the row's gradient and the Gauss–Newton part of its Hessian and nothing more; its half square is not the row's term, and the fit still minimizes the sum of the terms. In the fleet $\ln\mu_i = \ln(h_i\lambda_{\mathrm{ref}}) + \omega + d_i$ is linear in the state, and the line has the entry $\sqrt{\mu_i}$ in the column of ω and in that of d_i; for the two loads with an uncertain factor $\ln\mu$ is not linear in the state, and the remaining term $\mathbf{C}$ of equation (4.44) is named and kept (equation (6.73)). On the pump's one entry neither is needed: the rows' slopes and curvatures add as numbers.

The count row is not a Gaussian reading. A frequent substitute for the count row is a Gaussian reading row on the log-rate, a reading $\ln(k/(T_{\mathrm{tot}}\lambda_{\mathrm{ref}}))$ of ω with variance $1/k$. It reproduces the curvature k of the count row at the MLE when $k > 0$, but it has no value at all when $k = 0$, a record of failure-free hours, where the count row still carries the information that T_{tot} hours passed without failure.

The prior row. The gamma prior (6.15) has the density $p(\lambda) \propto \lambda^{a-1}\mathrm{e}^{-b\lambda}$. Carried on ω, the density gains the change-of-variable factor $\mathrm{d}\lambda/\mathrm{d}\omega = \lambda$ of equation (4.51), $p_\omega(\omega) = p(\lambda)\,\lambda \propto \lambda^a\mathrm{e}^{-b\lambda}$, so that its exponent is a rather than $a-1$, and the prior row's term is $\mathcal{L}_0(\omega) = b\lambda - a\omega$: a count row of a pseudo-failures with the expected count $b\lambda$, whose log-mean $\ln(b\lambda_{\mathrm{ref}}) + \omega$ is

linear in the state, so that by equation (6.21) its slope is $\mathcal{L}_0' = b\lambda - a$ and its curvature $\mathcal{L}_0'' = b\lambda$. The count row, by contrast, is a probability of the record, not a density over the state, and takes no change-of-variable factor.

The rows of each record. Which row acts on which entry is a property of the graphs. On the pump's one entry both rows act on ω, and divided into its four periods the pump's record is four count rows on the same number, each with its own exposure. On the states of several entries the graphs fix where the Jacobian of the rows can be nonzero, marked $*$, with the rows labeled at the right of each pattern:

$$\underbrace{\begin{bmatrix} * & 0 \\ * & * \\ 0 & * \end{bmatrix}\begin{matrix} \text{prior} \\ \text{count} \\ \text{factor prior} \end{matrix}}_{\text{two loads}} \qquad \underbrace{\begin{bmatrix} * & 0 \\ 0 & * \end{bmatrix}\begin{matrix} \text{count} \\ \text{timing} \end{matrix}}_{\text{power law}}$$

$$\underbrace{\begin{bmatrix} * & 0 & 0 & 0 \\ * & * & 0 & 0 \\ * & 0 & * & 0 \\ * & 0 & 0 & * \\ 0 & * & 0 & 0 \\ 0 & 0 & * & 0 \\ 0 & 0 & 0 & * \end{bmatrix}\begin{matrix} \text{prior on } \omega \\ \text{count A} \\ \text{count B} \\ \text{count C} \\ \text{tie A} \\ \text{tie B} \\ \text{tie C} \end{matrix}}_{\text{three pumps}} \tag{6.23}$$

with the columns ω_0, $\ln c_2$ for the two loads; $\ln m(T_{\text{tot}})$, v for the power law; and ω, d_A, d_B, d_C for the three pumps. In the fleet the column of ω is shared by the prior and every count row, each count row has a second entry in its own pump's column, and the tie rows d_i/σ_d, Gaussian rows with the standard deviation σ_d, act on the deviations alone; they state that the pumps' rates agree to within a factor of about e^{σ_d}. The fleet's rate is reported on the pooled drawing, the limit $\sigma_d \to 0$, and pump B's own rate at $\sigma_d = 0.5$, a factor of about 1.6. The comparison of the ninth question sets the pooled drawing against a drawing with a rate per pump, each rate under the prior (6.18). For the two loads the count row's expected count adds the two conditions, so it alone acts on both columns when the high-load factor is uncertain, and the factor's Gaussian prior row $(\ln c_2 - \ln 2)/\sigma_c$, with $\sigma_c = 0.3$, about thirty percent, acts on $\ln c_2$. The power law has a count row and a timing row, the density of the failure times given their number, and no prior row (Section 6.3.4). None of the rows is a physics row; the tie rows come closest, an exact equality of the rates carried as a row whose tolerance can be taken to zero. The count rows, the timing row and the gamma prior rows are not Gaussian: each enters the objective as its negative log-probability term, with the slope and curvature of equation (6.21) on the pump's one entry and in the working form (6.22) where it is stacked with other rows on several entries.

6.2 The posterior of the state (step 3)

Four objects of the fit recur in the specialized forms, and they are named here by their roles, as in Chapter 4, Section 4.2.7. The *MAP state* is the minimum of the objective that the rows put on the state; the *information* is the curvature of that objective at the minimum; its inverse, the *Laplace variance*, makes the posterior of the state approximately normal; and a derived quantity f is carried by its derivative at the MAP state. On the pump's one entry the four objects are numbers: $\hat{\omega}$, the information $\Lambda = \mathcal{L}''(\hat{\omega})$, the variance $\sigma_\omega^2 = 1/\Lambda$ and the derivative $g = \mathrm{d}f/\mathrm{d}\omega$, with the variance $g^2\sigma_\omega^2$. On the four entries of the fleet and the two of the pump under two

loads they are the vector $\hat{\mathbf{z}}$, the information matrix $\mathbf{\Lambda} = \mathbf{J}^{\mathsf{T}}\mathbf{J}$ of the working rows (6.22) with the pattern (6.23), the Hessian $\mathbf{H} = \mathbf{J}^{\mathsf{T}}\mathbf{J} + \mathbf{C}$ of equation (4.44) where a row's log-mean is not linear in the state, the covariance $\mathbf{\Sigma} = \mathbf{H}^{-1}$, whose off-diagonal entries come from the shared columns of the pattern, and the gradient $\mathbf{g} = \partial f/\partial \mathbf{z}$, with the variance $\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}$. Beside them stands the exact posterior that the one-entry problems admit: the rate has the gamma posterior $\text{Gamma}(a_n, b_n)$, whose shape $a_n = a + k$ is the prior shape plus the failures and whose rate parameter $b_n = b + T_{\text{tot}}$ is the prior rate parameter plus the hours (Section 6.2.2), the evidence of a gamma drawing is a closed form, and the power-law shape has an exact band of its own.

All eleven questions are answered from the posterior of the state that the rows of step 2 give: for the pump, the Laplace normal of the log-rate, with the exact gamma density of the rate beside it; for the fleet, the two loads and the power law, the same rows on their extended states. This section derives it and computes it for the pump; Section 6.3 computes every answer from it.

6.2.1 The posterior from the rows

The objective and the Newton step. The two rows act on the one number ω. Their terms, $\mathcal{L}_0 = b\lambda - a\omega$ and $\mathcal{L}_c = T_{\text{tot}}\lambda - k\omega$ of Section 6.1.7, add to the objective,

$$\mathcal{L}(\omega) = \mathcal{L}_0(\omega) + \mathcal{L}_c(\omega) = b_n\lambda - a_n \ln\lambda, \qquad \lambda = \lambda_{\text{ref}}\, \mathrm{e}^{\omega}, \tag{6.24}$$

up to a constant, with $a_n = a + k$ and $b_n = b + T_{\text{tot}}$: the exact negative log posterior of ω. Their slopes and curvatures, equation (6.21), add with them:

$$\mathcal{L}'(\omega) = \underbrace{b\lambda - a}_{\text{prior}} + \underbrace{T_{\text{tot}}\lambda - k}_{\text{count}} = b_n\lambda - a_n, \qquad \mathcal{L}''(\omega) = b\lambda + T_{\text{tot}}\lambda = b_n\lambda.$$

Both rows are count rows in the log of their mean, so $\mathcal{L}''$ is the exact curvature at every ω. Newton's method replaces the objective at the current ω_j by the parabola with the same slope and curvature and moves to its minimum,

$$\Delta\omega = -\frac{\mathcal{L}'(\omega_j)}{\mathcal{L}''(\omega_j)}, \qquad \omega_{j+1} = \omega_j + \Delta\omega. \tag{6.25}$$

The MAP state and the variance. The slope $b_n\lambda - a_n$ vanishes at $\lambda = a_n/b_n$, so the minimum of the objective is

$$\hat{\omega} = \ln\frac{a_n}{b_n\lambda_{\text{ref}}}, \qquad \hat{\lambda} = \lambda_{\text{ref}}\, \mathrm{e}^{\hat{\omega}} = \frac{a_n}{b_n}, \tag{6.26}$$

the logarithm of the exact posterior mean (Section 6.2.2), not of its mode, because a mode is not invariant under a change of variable. The curvature at the minimum is the information Λ, and its reciprocal is the Laplace variance of the log-rate,

$$\Lambda = \mathcal{L}''(\hat{\omega}) = \underbrace{b\,\hat{\lambda}}_{\text{prior}} + \underbrace{T_{\text{tot}}\,\hat{\lambda}}_{\text{count}} = a_n, \qquad \sigma_\omega^2 = \frac{1}{\Lambda} = \frac{1}{a_n}. \tag{6.27}$$

The two curvatures are the rows' precisions, and they add: the prior row contributes $b\hat{\lambda}$, the count row $T_{\text{tot}}\hat{\lambda}$, the expected number of failures in the exposure and the Poisson Fisher information in the log of its mean, as the precisions of independent readings add. Their shares of the information, $b/(b + T_{\text{tot}})$ and $T_{\text{tot}}/(b + T_{\text{tot}})$, are the weights with which equation (6.42) averages the prior mean and the MLE, and the total, a_n, is why every band of the rate narrows as $1/\sqrt{a_n}$. The answer is read off as $\hat{\omega}$ with $\sigma_\omega = 1/\sqrt{a_n}$, and the rate as a derived quantity with the derivative of equation (6.40), whose delta-method spread $\hat{\lambda}/\sqrt{a_n} = \sqrt{a_n}/b_n$ is the exact posterior standard deviation of equation (6.43): the agreement holds for every gamma posterior, since $\hat{\lambda}^2/a_n = a_n/b_n^2$. It belongs to the linearized value only (Section 6.4.1).

When the rows' slopes do not vanish. At the MAP state the slopes of the prior row and the count row are both zero only when each row sits at its own optimum, as when the prior mean a/b equals the MLE k/T_{tot}. Otherwise neither is zero: the prior row pulls ω toward the prior mean and the count row toward the observed rate, with equal and opposite slopes that cancel in $\mathcal{L}'$ (Section 6.4.4 works a record where they do not vanish).

Several count rows on one entry. Suppose a record is divided into I periods, period i with exposure h_i and k_i failures. Under a shared constant intensity λ, the count in period i is $K_i \mid \lambda, h_i \sim \text{Poisson}(\lambda h_i)$, and if counts in disjoint periods are conditionally independent, as assumption 2 of Section 6.1.3 states,

$$p(k_1, \ldots, k_I \mid \lambda) = \prod_{i=1}^{I} \frac{(\lambda h_i)^{k_i}\, \mathrm{e}^{-\lambda h_i}}{k_i!}.$$

Collect the terms that involve λ:

$$p(k_1, \ldots, k_I \mid \lambda) \propto \lambda^{\sum_i k_i}\, \mathrm{e}^{-\lambda \sum_i h_i} = \lambda^{k}\, \mathrm{e}^{-\lambda T_{\text{tot}}}, \qquad k = \sum_i k_i, \quad T_{\text{tot}} = \sum_i h_i. \tag{6.28}$$

The likelihood depends on the record only through the total count and the total exposure, which are sufficient statistics for λ. The same holds when the individual failure times $0 < t_1 < \cdots < t_k \le T_{\text{tot}}$ are recorded. By assumptions 2 and 3 and equation (6.10), the intervals between successive failures are independent exponential times with rate λ, density $\lambda \mathrm{e}^{-\lambda(t_j - t_{j-1})}$ with $t_0 = 0$, and the interval after the last failure contributes the probability $\mathrm{e}^{-\lambda(T_{\text{tot}} - t_k)}$ that no failure followed. The density of the record is their product,

$$p(t_1, \ldots, t_k \mid \lambda) = \prod_{j=1}^{k} \lambda\, \mathrm{e}^{-\lambda(t_j - t_{j-1})} \times \mathrm{e}^{-\lambda(T_{\text{tot}} - t_k)} = \lambda^{k}\, \mathrm{e}^{-\lambda T_{\text{tot}}}, \tag{6.29}$$

because the exponents telescope to λT_{tot}. It is proportional in λ to the count likelihood (6.19). On the state the periods are I count rows beside the prior row on the same number ω, each with the expected count $\mu_i = h_i \lambda_{\text{ref}} \mathrm{e}^{\omega}$, and by equation (6.21) their slopes and curvatures add to

$$\mathcal{L}' = (b\lambda - a) + \sum_i (h_i \lambda - k_i) = (b + T_{\text{tot}})\lambda - (a + k), \qquad \mathcal{L}'' = b\lambda + \sum_i h_i \lambda = (b + T_{\text{tot}})\lambda,$$

which involve the periods only through the totals: sufficiency is read in the sums. As functions of ω they are the slope and the curvature of the single record, so the iteration and its end are the same. The individual slopes need not vanish at the MAP state. They are what a shared rate leaves unexplained, and they carry the evidence if the periods' rates differed, which is the dispersion check of Section 6.4.6.

The fit on any state. A larger model has the same terms: a state that holds the log-rate and, where the question needs them, deviations, load factors or other parameters; one count row per record, period, pump or condition, each with its own exposure; the gamma prior rows; and the Gaussian rows beside them. Its fit is that of Chapter 4, Section 4.2.7: it stacks the working rows (6.22) with the normalized Gaussian rows, keeps the remaining term $\mathbf{C}$ of equation (4.44) wherever $\ln \mu$ is not linear in the state, takes the Newton steps of equation (4.40) to the MAP state, and inverts the Hessian there.

6.2.2 The exact posterior

The Gamma–Poisson posterior. Bayes' theorem with the Poisson likelihood (6.19) and the gamma prior (6.15) gives a gamma posterior, the conjugate update derived in Chapter 2, equation (2.39): the factors $T_{\text{tot}}^k/k!$ and $b^a/\Gamma(a)$ drop into the proportionality, the powers of λ and the exponents collect into $\lambda^{a+k-1}\mathrm{e}^{-(b+T_{\text{tot}})\lambda}$, and the gamma integral (6.16) supplies the constant. In this chapter's notation,

$$a_n = a + k, \qquad b_n = b + T_{\text{tot}}, \tag{6.30}$$

$$\lambda \mid k, T_{\text{tot}} \sim \text{Gamma}(a_n, b_n), \qquad p(\lambda \mid k, T_{\text{tot}}) = \frac{b_n^{a_n}}{\Gamma(a_n)}\,\lambda^{a_n-1}\,\mathrm{e}^{-b_n\lambda}. \tag{6.31}$$

This is the *Gamma–Poisson conjugate update*: the failures add to the shape and the exposure, in operating hours, adds to the rate parameter, which is why b carries units of time. It is an exact posterior under the model assumptions.

Pooled records. By the sufficiency (6.28), a record divided into periods, pumps that share one rate, or the failure times of equation (6.29) give the posterior of the aggregated count,

$$\lambda \mid k_1, \ldots, k_I \sim \text{Gamma}\Big(a + \sum_i k_i,\ b + \sum_i h_i\Big), \tag{6.32}$$

and the update is incremental, in any order and grouping, which in the state-vector form is the statement that rows add.

The fit and the exact posterior. The objective (6.24) is the negative logarithm of the posterior (6.31) written as a density of ω, and the minimum (6.26) lies at the logarithm of its mean; the fit is its Laplace approximation. The entire operating record is absorbed into the two numbers (a_n, b_n), and once they are known every question of the pump follows without revisiting the record, as long as the model assumptions remain appropriate.

6.2.3 The posterior of the pump

The pump's inputs are the prior Gamma$(2, 1000\,\text{h})$ and the record of $k = 4$ failures in $T_{\text{tot}} = 2000$ operating hours, with its four periods.

The prior and the record. The prior (6.18) has the density $p(\lambda) = 10^6\,\lambda\,\mathrm{e}^{-1000\lambda}$, since $b^a/\Gamma(a) = 1000^2/1! = 10^6\,\text{h}^2$. By equation (6.17) its mean is $2/1000 = 0.002\,\text{h}^{-1}$, its standard deviation $\text{sd}(\lambda) = \sqrt{2}/1000 = 0.001414\,\text{h}^{-1}$ and its coefficient of variation $1/\sqrt{2} = 0.71$; its central 95 % interval runs from 0.00024 to $0.0056\,\text{h}^{-1}$, a factor of twenty-three. By equation (6.19) the record has the likelihood

$$p(k = 4 \mid \lambda, T_{\text{tot}} = 2000\,\text{h}) = \frac{(2000\,\lambda)^4\,\mathrm{e}^{-2000\,\lambda}}{4!}, \tag{6.33}$$

with λ in h^{-1}, and by equation (6.20) the maximum-likelihood estimate $4/2000 = 0.002\,\text{h}^{-1}$, two failures per thousand operating hours. The likelihood at its maximum is $4^4\mathrm{e}^{-4}/24 = 0.195$, the probability of exactly four failures when four are expected; the curvature there, $T_{\text{tot}}^2/k = 10^6\,\text{h}^2$, corresponds to the width $0.001\,\text{h}^{-1}$: four events fix the rate only to within about half its value. Treated as an exact constant hazard, the estimate would give MTTF $= 1/0.002 = 500\,\text{h}$ by equation (6.6).

The Newton iteration. With $a = 2$, $b = 1000\,\mathrm{h}$, $k = 4$ and $T_{\mathrm{tot}} = 2000\,\mathrm{h}$, the slopes of the two rows are $1000\lambda - 2$ for the prior row and $2000\lambda - 4$ for the count row, and the curvature is 3000λ. At the starting state $\omega = 0$, $\lambda = \lambda_{\mathrm{ref}}$, the slopes are -1.000 and -2.000 and the curvature $1.000 + 2.000 = 3.000$, so that equation (6.25) gives $\Delta\omega = 3/3 = 1$. Table 6.4 lists the iteration: each row is an iterate, with the log-rate, the rate, the slopes of the prior row and the count row, their sum, the curvature and the step. In four steps the log-rate reaches the MAP state (6.26),

$$\hat{\omega} = \ln \frac{6}{3000 \times 0.001} = \ln 2 = 0.6931, \qquad \hat{\lambda} = 0.002\,\mathrm{h}^{-1}.$$

Table 6.4: Newton iterations, equation (6.25), for the pump record in the log-rate ω, from $\omega = 0$: the slopes of the prior row and the count row, equation (6.21), their sum and the curvature. Rates in $10^{-3}\,\mathrm{h}^{-1}$.

Iteration j	ω_j	λ_j	Prior slope $\mathcal{L}_0'$	Count slope $\mathcal{L}_c'$	Slope $\mathcal{L}'$	Curvature $\mathcal{L}''$	$\Delta\omega$
0	0	1.000	−1.000	−2.000	−3.000	3.000	1.0000
1	1.0000	2.718	0.718	1.437	2.155	8.155	−0.2642
2	0.7358	2.087	0.087	0.174	0.261	6.261	−0.0417
3	0.6940	2.002	0.002	0.004	0.005	6.005	−0.0009
4	0.6931	2.000	0.000	0.000	0.000	6.000	0.0000

The information and the variance. At the MAP state the two rows' curvatures are their information, $1000 \times 0.002 = 2$ for the prior row and the expected number of failures $2000 \times 0.002 = 4$ for the count row, shares of $2/6$ and $4/6$, and equation (6.27) gives

$$\Lambda = 2 + 4 = 6, \qquad \sigma_\omega^2 = \frac{1}{6} = 0.1667, \qquad \sigma_\omega = 0.4082.$$

Both rows' slopes are zero at the MAP state, because the prior mean $2/1000$ equals the MLE $4/2000$.

The exact posterior. With the prior $\mathrm{Gamma}(2, 1000\,\mathrm{h})$ and the record $k = 4$, $T_{\mathrm{tot}} = 2000\,\mathrm{h}$, equation (6.30) gives $a_n = 2 + 4 = 6$ and $b_n = 1000 + 2000 = 3000\,\mathrm{h}$:

$$\lambda \mid \mathbf{y} \sim \mathrm{Gamma}(6,\ 3000\,\mathrm{h}), \qquad p(\lambda \mid \mathbf{y}) = \frac{3000^6}{5!}\,\lambda^5\,\mathrm{e}^{-3000\lambda} = 6.075 \times 10^{18}\,\lambda^5\,\mathrm{e}^{-3000\lambda}. \quad (6.34)$$

The normalizing constant (6.68) is $\Gamma(6)/(\Gamma(2)\,4!) \times (1/3)^2(2/3)^4 = 5 \times 16/729 = 0.1097$: before the record, the prior assigned 11.0 % to exactly four failures in two thousand hours.

The four periods. In the order prior, periods 1 to 4, the rows' slopes $h_i\hat{\lambda} - k_i$ at the MAP state are 0, 0, -0.6, 0.6 and 0, and their curvatures $h_i\hat{\lambda}$ add to $\Lambda = 2 + 1.0 + 1.4 + 0.6 + 1.0 = 6$, with $\sigma_\omega^2 = 1/6$, the values of the single record. Period 2, two failures where 1.4 were expected, has the slope -0.6, and period 3, none where 0.6 were expected, has $+0.6$; they cancel in $\mathcal{L}'$. Applied one after another to the prior, the four periods give the sequence of gamma parameters

$$(a, b) = (2, 1000) \to (3, 1500) \to (5, 2200) \to (5, 2500) \to (6, 3000),$$

which ends at the posterior of the whole record taken at once.

The pair $(\hat{\omega}, \sigma_\omega) = (0.6931,\ 0.4082)$ is what step 4 uses for the pump: every answer about the pump is computed from it, and those about the fleet, the two loads and the power law from $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$ of the same rows on their extended states. Because the one-entry problems also admit the exact posterior, each such answer has an exact value beside it, a function of the two numbers $(a_n, b_n) = (6,\ 3000\,\mathrm{h})$ for the pump, and the gap between the two is reported with it.

6.3 The answers (step 4)

The questions of Table 6.3 fall into three categories: estimation (the first two, the eighth, the tenth and the eleventh), prediction under a threshold (the third to the seventh) and hypothesis comparison (the ninth). Each is answered from the posterior of Section 6.2: its engineering question, the general form of its category, the form it takes for the record, the quantities of the posterior it uses, and the record's numbers. The questions name the fleet's records and the two loads of Table 6.2, with the tie tolerance $\sigma_d = 0.5$ and the factor's tolerance $\sigma_c = 0.3$; the failure times of records A and B; the rate limit $\lambda_{\max} = 0.003\,\mathrm{h}^{-1}$ at the assurance $\kappa = 0.95$; and the future exposure $\tau = 500\,\mathrm{h}$. Every answer is computed from the posterior of the state of step 3, the MAP state $\hat{\omega}$ and the standard deviation σ_ω for the pump and $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$ on the extended states, and that is the value the chapter reports. Where an exact posterior exists, the gamma posterior of a rate or the chi-square band of the power-law shape, it gives the value set beside it, with the gap. Pump B's own rate when the pumps may differ, and the rates under an uncertain load factor, have no exact counterpart. The comparison of the pumps is answered by the exact evidences of the two drawings; the chapter has no state route for an evidence.

6.3.1 The failure intensity, its band and the rate requirement

Engineering questions. What is the failure intensity? How certain is the estimate? Is the intensity below a target? The first asks for the rate at which the pump fails given its record; the second for the spread of that estimate and an interval that contains the rate with a stated probability; the third whether the rate meets the facility's limit $\lambda < \lambda_{\max} = 0.003\,\mathrm{h}^{-1}$ with the posterior probability $\kappa = 0.95$, the assurance level of Chapter 5.

General form. An estimation question about a quantity $f(\mathbf{z})$ derived from the state asks for its posterior distribution, the probability that the quantity lies at or below each value f^*,

$$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y}) = \int \mathbf{1}\big[f(\mathbf{z}) \le f^*\big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \tag{6.35}$$

summarized by a location, the mode or the mean,

$$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}] = \int f(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \tag{6.36}$$

and by a spread, the posterior variance $\mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y})$ and the central 95 % credible interval $[f_L, f_U]$, which leaves 2.5 % of posterior probability in each tail of the distribution (6.35):

$$\Pr(f(\mathbf{z}) < f_L \mid \mathbf{y}) = 0.025, \qquad \Pr(f(\mathbf{z}) > f_U \mid \mathbf{y}) = 0.025. \tag{6.37}$$

An entry of the state is the case $f(\mathbf{z}) = z_i$. When f is a monotone function of a single entry, its quantiles are the images of that entry's quantiles. A prediction under a threshold compares a tail of the same distribution with an assurance level, the form of Chapter 2, equation (2.12), here with f^* a requirement level that the quantity must stay below:

$$\text{accept if } \Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y}) \ge \kappa, \quad \text{equivalently if } f^* \ge f_\kappa, \qquad \Pr(f(\mathbf{z}) \le f_\kappa \mid \mathbf{y}) = \kappa, \tag{6.38}$$

with f_κ the one-sided bound, the κ quantile of the quantity.

Specialized form. The three questions are equations (6.35), (6.36) and (6.37), and for the third the rule (6.38), with the state $\mathbf{z}$ the one number ω, the derived quantity $f(\mathbf{z}) = \lambda = \lambda_{\mathrm{ref}} \mathrm{e}^{\omega}$ of equation (6.11), an increasing function of the one entry, the requirement $f^* = \lambda_{\max}$, and $p(\mathbf{z} \mid \mathbf{y})$ the posterior of step 3, the Laplace normal $\omega \mid \mathbf{y} \approx \mathcal{N}(\hat{\omega}, \sigma_\omega^2)$ of equations (6.26)

and (6.27). Each answer is read from the two numbers $\hat{\omega}$ and σ_ω, and the exact gamma posterior (6.31) gives the value set beside it.

For the location, the exponential is increasing, so equation (6.35) at $f^* = \hat{\lambda} = \lambda_{\text{ref}} \mathrm{e}^{\hat{\omega}}$ is the normal probability $\Pr(\omega \le \hat{\omega} \mid \mathbf{y}) = \frac{1}{2}$: the rate at the MAP state is the median of λ under the Laplace posterior, and it is the point estimate the chapter reports. The central interval of the rate is the interval of ω mapped through the exponential,

$$\lambda \in \left[\lambda_{\text{ref}} \mathrm{e}^{\hat{\omega} - 1.96\sigma_\omega},\ \lambda_{\text{ref}} \mathrm{e}^{\hat{\omega} + 1.96\sigma_\omega}\right], \tag{6.39}$$

while the spread of the rate follows from the derivative of the map at the MAP state by the delta method of Chapter 4 (Section 4.3.1),

$$g = \left.\frac{\mathrm{d}\lambda}{\mathrm{d}\omega}\right|_{\hat{\omega}} = \hat{\lambda}, \qquad \operatorname{sd}(\lambda \mid \mathbf{y}) \approx |g|\, \sigma_\omega = \hat{\lambda}\, \sigma_\omega, \tag{6.40}$$

with $\hat{\lambda} = \lambda_{\text{ref}} \mathrm{e}^{\hat{\omega}}$. The requirement becomes an event on the one entry, $\lambda < \lambda_{\max} \iff \omega < \omega^* = \ln(\lambda_{\max}/\lambda_{\text{ref}})$, so its tail is a normal tail and the one-sided bound the MAP state moved by 1.6449 standard deviations,

$$\Pr(\lambda < \lambda_{\max} \mid \mathbf{y}) \approx \Phi\!\left(\frac{\omega^* - \hat{\omega}}{\sigma_\omega}\right), \qquad \lambda_{0.95} \approx \lambda_{\text{ref}}\, \mathrm{e}^{\hat{\omega} + 1.6449\, \sigma_\omega}, \tag{6.41}$$

with Φ the standard normal cumulative distribution; only the tail is an approximation, since the threshold maps exactly.

The values beside them come from the exact posterior: the rate has the gamma posterior $\text{Gamma}(a_n, b_n)$. Its mean and variance are the moments (6.17) with the posterior parameters. The mean is

$$\mathbb{E}[\lambda \mid \mathbf{y}] = \frac{a_n}{b_n} = \frac{b}{b + T_{\text{tot}}}\, \frac{a}{b} + \frac{T_{\text{tot}}}{b + T_{\text{tot}}}\, \frac{k}{T_{\text{tot}}}, \tag{6.42}$$

the average of the prior mean a/b and the maximum-likelihood estimate k/T_{tot} (Section 6.1.7) weighted by the prior's share $b/(b+T_{\text{tot}})$ of the total exposure and the record's share $T_{\text{tot}}/(b+T_{\text{tot}})$; as exposure accumulates, the prior's share tends to zero and the mean approaches the observed rate. The variance is

$$\operatorname{Var}(\lambda \mid \mathbf{y}) = \frac{a_n}{b_n^2}, \qquad \operatorname{sd}(\lambda \mid \mathbf{y}) = \frac{\sqrt{a_n}}{b_n}, \qquad \frac{\operatorname{sd}(\lambda \mid \mathbf{y})}{\mathbb{E}[\lambda \mid \mathbf{y}]} = \frac{1}{\sqrt{a_n}}. \tag{6.43}$$

It describes uncertainty about the rate, an epistemic quantity; it is not the variability of the future count, which is aleatory and remains even when λ is known (Section 6.3.2). The most probable value under the exact density of λ follows from the log posterior $\ln p(\lambda \mid \mathbf{y}) = (a_n - 1)\ln\lambda - b_n\lambda + C$, where C does not depend on λ; its derivative $(a_n - 1)/\lambda - b_n$ vanishes at

$$\hat{\lambda}_{\text{mode}} = \frac{a_n - 1}{b_n}, \qquad a_n > 1, \tag{6.44}$$

and for $a_n \le 1$ the density is largest at $\lambda = 0$. Three point estimates summarize the posterior: the maximum-likelihood estimate k/T_{tot}, which uses the record alone, the mode and the mean. They need not agree, since they summarize the same inference in different ways, and the mean exceeds the mode because the gamma posterior is skewed to the right. For the prediction of future failures the mean is the important one, because the expected count is linear in λ.

The endpoints of the central interval are quantiles of the posterior, obtained from its cumulative distribution function

$$F_\Gamma(x; a, b) = \int_0^x \frac{b^a}{\Gamma(a)}\, \lambda^{a-1} \mathrm{e}^{-b\lambda}\, \mathrm{d}\lambda = \frac{1}{\Gamma(a)} \int_0^{bx} u^{a-1} \mathrm{e}^{-u}\, \mathrm{d}u, \tag{6.45}$$

the regularized lower incomplete gamma function evaluated at bx, which every statistical library evaluates and inverts. For an integer shape it has a closed form worth having for hand checks:

$$F_\Gamma(x; a, b) = \Pr\big(\text{Poisson}(bx) \geq a\big) = 1 - \sum_{j=0}^{a-1} \frac{(bx)^j \, \mathrm{e}^{-bx}}{j!}. \tag{6.46}$$

The identity follows by integrating equation (6.45) by parts $a - 1$ times, and it has a direct reading: a Gamma(a, b) variable with integer a is distributed as the time of the a-th event of a Poisson process of rate b, and that event has occurred by x exactly when at least a events have. A second standard form is that $2b\lambda$ follows the chi-square distribution with $2a$ degrees of freedom, so that

$$\lambda_q = \frac{\chi^2_{q;\,2a_n}}{2b_n} \tag{6.47}$$

is the q quantile of the posterior, with $\chi^2_{q;\,2a_n}$ the q quantile of the chi-square distribution with $2a_n$ degrees of freedom, and the interval (6.37) is $[\lambda_{0.025},\ \lambda_{0.975}]$. The requirement's tail is the gamma cumulative distribution at the limit, which for an integer shape is the Poisson sum (6.46) with $bx = b_n\lambda_{\max}$,

$$\Pr(\lambda < \lambda_{\max} \mid \mathbf{y}) = F_\Gamma(\lambda_{\max}; a_n, b_n) = 1 - \mathrm{e}^{-b_n\lambda_{\max}} \sum_{j=0}^{a_n-1} \frac{(b_n\lambda_{\max})^j}{j!}, \tag{6.48}$$

and the same formula with the prior parameters (a, b) in place of (a_n, b_n) gives the support under the prior alone. The one-sided bound of the rule (6.38) is the 0.95 quantile,

$$\lambda_{0.95} = \frac{\chi^2_{0.95;\,2a_n}}{2b_n}, \qquad \Pr(\lambda \leq \lambda_{0.95} \mid \mathbf{y}) = 0.95, \tag{6.49}$$

and the rule with $\kappa = 0.95$ is met exactly when $\lambda_{\max} \geq \lambda_{0.95}$. It is a one-sided bound, not the upper end of the central interval, which leaves 2.5 % above it rather than 5 %.

Building blocks. The MAP state $\hat{\omega} = \ln 2 = 0.6931$ and the standard deviation $\sigma_\omega = 1/\sqrt{6} = 0.4082$ of equations (6.26) and (6.27) (Section 6.2.3); the derivative $g = \hat{\lambda}$ of equation (6.40); the threshold $\omega^* = \ln 3$ and the normal cumulative distribution Φ; and, for the exact values beside them, the posterior parameters $a_n = 6$ and $b_n = 3000\,\mathrm{h}$ (Section 6.2.3), the chi-square quantiles, the Poisson sum (6.46) and the maximum-likelihood estimate of the count row (Section 6.1.7).

The failure intensity. The rate at the MAP state, the Laplace median, is

$$\hat{\lambda} = \lambda_{\text{ref}}\, \mathrm{e}^{0.6931} = 10^{-3} \times 2 = 0.002\,\mathrm{h}^{-1},$$

two failures per thousand operating hours. The value beside it comes from the exact posterior, whose median is $0.00189\,\mathrm{h}^{-1}$, a gap of $0.00011\,\mathrm{h}^{-1}$. Its mean, equation (6.42), is a different summary that falls on the reported median,

$$\mathbb{E}[\lambda \mid \mathbf{y}] = \frac{6}{3000} = 0.002\,\mathrm{h}^{-1},$$

because $\hat{\omega} = \ln(a_n/(b_n\lambda_{\text{ref}}))$, while the Laplace mean is $0.00217\,\mathrm{h}^{-1}$. The exact mean equals the MLE because the prior mean and the observed rate agree: the shares of the prior and the record are 1/3 and 2/3, and $\frac{1}{3} \times 0.002 + \frac{2}{3} \times 0.002 = 0.002$. By equation (6.44) the exact mode is $5/3000 = 0.001667\,\mathrm{h}^{-1}$, below the mean. Table 6.5 collects the point estimates and the standard deviations: each row is a summary, the middle column its expression and the right column its value for the pump; the first two rows are the reported ones.

Table 6.5: Point estimates and standard deviations of the pump's failure rate. Prior Gamma$(2, 1000\,\mathrm{h})$; record $k = 4$, $T_{\mathrm{tot}} = 2000\,\mathrm{h}$; Laplace posterior $\omega \mid \mathbf{y} \approx \mathcal{N}(0.6931,\ 0.4082^2)$; exact posterior Gamma$(6, 3000\,\mathrm{h})$.

Summary	Expression	Value [h^{-1}]
Laplace median (reported)	$\lambda_{\mathrm{ref}} \mathrm{e}^{\hat{\omega}}$	0.002000
Laplace standard deviation (reported)	$\hat{\lambda}\,\sigma_\omega$	0.000816
Maximum likelihood	k/T_{tot}	0.002000
Exact posterior mode	$(a_n - 1)/b_n$	0.001667
Exact posterior median	$F_\Gamma^{-1}(1/2; a_n, b_n)$	0.001890
Exact posterior median	$F_\Gamma^{-1}(1/2; a_n, b_n)$	0.001890
Exact posterior mean	a_n/b_n	0.002000
Prior standard deviation	$\sqrt{a}/b$	0.001414
Exact posterior standard deviation	$\sqrt{a_n}/b_n$	0.000816

Figure 6.3 draws the densities that enter the update. The horizontal axis is the failure rate in units of $10^{-3}\,\mathrm{h}^{-1}$, failures per thousand operating hours, and the vertical axis is probability density per that unit. The blue curve is the prior Gamma$(2, 1000)$; the dashed gold curve is the likelihood (6.33) rescaled to unit area, which is the Gamma$(5, 2000)$ density and the posterior that a flat prior on λ would give; the dotted blue curve is the Laplace posterior, the normal $\mathcal{N}(0.6931,\ 0.4082^2)$ on ω carried to the rate; the green curve is the exact posterior Gamma$(6, 3000)$ of equation (6.34). The dotted lines mark the common value of the Laplace median, the prior mean, the MLE and the exact posterior mean at 2.0, and the exact posterior mode at 1.67. Both posteriors lie between the prior and the likelihood and are narrower than either; the record has left the best estimate where it was and has concentrated the distribution around it. The Laplace density peaks where the exact one does, near 1.7, and puts more probability on large rates: 0.044 above $4 \times 10^{-3}\,\mathrm{h}^{-1}$, against 0.020.

The certainty of the estimate. By equation (6.40),

$$g = \hat{\lambda} = 0.002\,\mathrm{h}^{-1}, \qquad g^2\sigma_\omega^2 = 0.002^2 \times \frac{1}{6} = 6.667 \times 10^{-7}\,\mathrm{h}^{-2},$$

$$\mathrm{sd}(\lambda \mid \mathbf{y}) \approx 0.002 \times 0.4082 = 0.000816\,\mathrm{h}^{-1},$$

down from $0.001414\,\mathrm{h}^{-1}$ for the prior, and the coefficient of variation is $\sigma_\omega = 0.41$, down from 0.71: the record has not moved the estimate but it has narrowed the distribution. The value beside it comes from the exact posterior: by equation (6.43), $\mathrm{Var}(\lambda \mid \mathbf{y}) = 6/3000^2 = 6.667 \times 10^{-7}\,\mathrm{h}^{-2}$ and $\mathrm{sd}(\lambda \mid \mathbf{y}) = \sqrt{6}/3000 = 0.000816\,\mathrm{h}^{-1}$, with the coefficient of variation $1/\sqrt{6} = 0.41$: the gap is zero, the identity $\hat{\lambda}^2/a_n = a_n/b_n^2$ of Section 6.2.1. By equation (6.39), the log-rate interval $0.6931 \pm 1.96 \times 0.4082 = [-0.1070,\ 1.4933]$ maps to

$$\lambda_L = 10^{-3}\,\mathrm{e}^{-0.1070} = 0.000899\,\mathrm{h}^{-1}, \qquad \lambda_U = 10^{-3}\,\mathrm{e}^{1.4933} = 0.004452\,\mathrm{h}^{-1}, \tag{6.50}$$

so that $\lambda \in [0.000899,\ 0.004452]\,\mathrm{h}^{-1}$ with 95 % posterior probability, conditional on the homogeneous Poisson model and the gamma prior. The interval runs from 0.45 to 2.23 times the median $0.002\,\mathrm{h}^{-1}$, and four observed failures are not enough to determine the rate precisely; it is symmetric about $\hat{\omega}$ in the log-rate, and so not about $\hat{\lambda}$ in the rate. The exact interval follows from the chi-square quantiles $\chi^2_{0.025;12} = 4.404$ and $\chi^2_{0.975;12} = 23.337$ and equation (6.47),

$$\lambda_L = \frac{4.404}{6000} = 0.000734\,\mathrm{h}^{-1}, \qquad \lambda_U = \frac{23.337}{6000} = 0.003889\,\mathrm{h}^{-1},$$

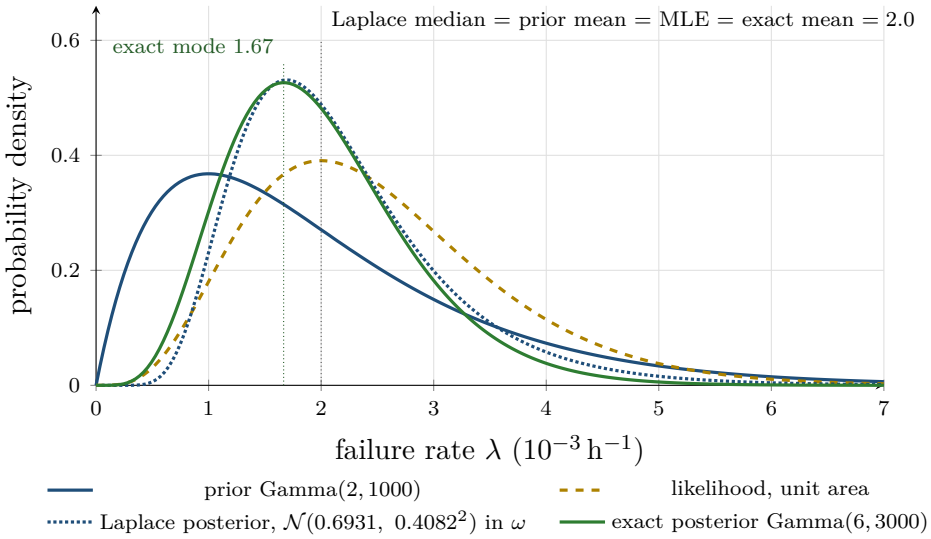


Figure 6.3: Prior, normalized likelihood and posterior of the pump's failure rate, from equation (6.15) with $x = 1000\lambda$: prior $x\,\mathrm{e}^{-x}$; likelihood rescaled to unit area $\frac{32}{24}\,x^4\mathrm{e}^{-2x}$ (the Gamma$(5, 2000)$ density); dotted blue, the Laplace posterior, the normal $\mathcal{N}(\ln 2,\ 1/6)$ on ω carried to the rate, which the answers report; green, the exact posterior $\frac{729}{120}\,x^5\mathrm{e}^{-3x}$. The vertical axis is density per $10^{-3}\,\mathrm{h}^{-1}$, not probability at an exact value.

from 0.37 to 1.94 times the mean; the mean lies 0.00127 above its lower end and 0.00189 below its upper one, because the gamma posterior has the longer tail on the right. The Laplace interval is shifted toward larger rates, by $0.000165\,\mathrm{h}^{-1}$ at the lower end and $0.000563\,\mathrm{h}^{-1}$ at the upper: the normal in the log-rate is symmetric, while the exact posterior of ω has its longer tail toward small rates (Section 6.4.1).

Figure 6.4(a) shows both intervals. The horizontal axis is the rate in $10^{-3}\,\mathrm{h}^{-1}$ and the vertical axis the posterior density; the dotted blue curve is the Laplace density with its interval shaded, area 0.95, and the green curve the exact density with its interval between the dashed lines. Panel (b) belongs to the requirement below.

The requirement of 0.003 per hour. By the first of equations (6.41) with $\omega^* = \ln 3 = 1.0986$,

$$\Pr(\lambda < 0.003 \mid \mathbf{y}) \approx \Phi\left(\frac{1.0986 - 0.6931}{0.4082}\right) = \Phi(0.9932) = 0.8397,$$

about 84.0 % posterior probability that the failure intensity lies below the limit: the shaded area of Figure 6.4(b), left of the dashed line at $3 \times 10^{-3}\,\mathrm{h}^{-1}$. The prior row alone, with $\hat{\omega} = \ln 2$ and $\Lambda = 2$, gives $\Phi(0.4055 \times \sqrt{2}) = \Phi(0.5734) = 0.717$. By the second, the one-sided bound is

$$\lambda_{0.95} \approx 10^{-3}\,\mathrm{e}^{0.6931 + 1.6449 \times 0.4082} = 10^{-3}\,\mathrm{e}^{1.3646} = 0.003914\,\mathrm{h}^{-1},$$

above the limit $0.003\,\mathrm{h}^{-1}$, so the rule at the assurance 0.95 is not met. The values beside them come from the exact posterior. With the integer shape $a_n = 6$ and $b_n\lambda_{\max} = 3000 \times 0.003 = 9$, equation (6.48) evaluates by hand:

$$\begin{aligned}\Pr(\lambda < 0.003 \mid \mathbf{y}) = F_\Gamma(0.003; 6, 3000) &= 1 - \mathrm{e}^{-9}\left(1 + 9 + \frac{9^2}{2} + \frac{9^3}{6} + \frac{9^4}{24} + \frac{9^5}{120}\right) \\ &= 1 - 0.1157 = 0.8843,\end{aligned} \tag{6.51}$$

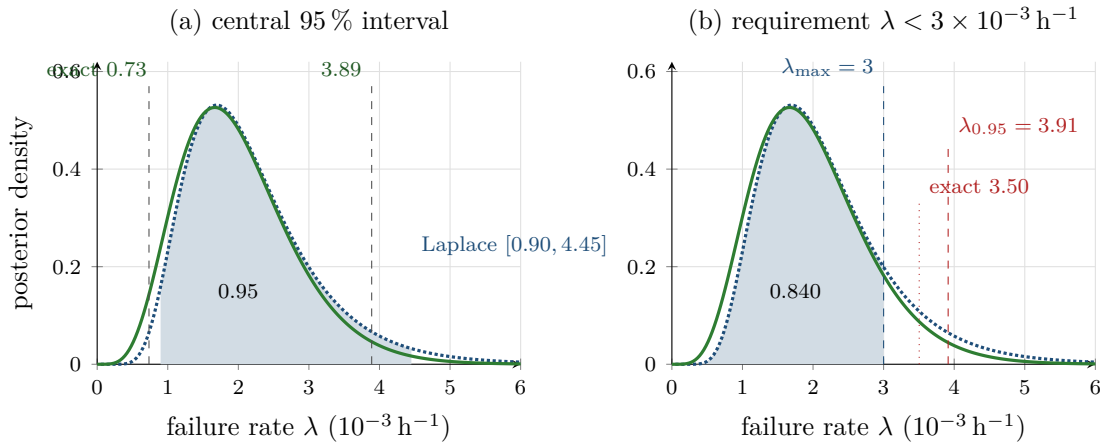


Figure 6.4: Reading the posterior of the pump, with $x = 1000\lambda$. Dotted blue: the Laplace density, $\omega \mid \mathbf{y} \approx \mathcal{N}(\ln 2,\ 1/6)$ carried to the rate, which the answers report; green: the exact Gamma$(6, 3000\,\mathrm{h})$, density $\frac{729}{120}\,x^5 \mathrm{e}^{-3x}$. (a) Central 95 % credible interval, shaded for the Laplace density, 0.899 to $4.452 \times 10^{-3}\,\mathrm{h}^{-1}$ (equation (6.50)), dashed for the exact one, 0.734 to $3.889 \times 10^{-3}\,\mathrm{h}^{-1}$. (b) Posterior probability of the requirement $\lambda < 0.003\,\mathrm{h}^{-1}$, shaded area 0.840 under the Laplace density ($F_\Gamma(0.003; 6, 3000) = 0.884$ exact), and the one-sided 95 % upper bound $\lambda_{0.95} = 0.003914\,\mathrm{h}^{-1}$ ($0.003504\,\mathrm{h}^{-1}$ exact).

since the bracket is 937.45 and $\mathrm{e}^{-9} = 1.234 \times 10^{-4}$: 88.4 %, against $1 - \mathrm{e}^{-3}(1+3) = 0.801$ under the prior alone, and 0.044 higher than the Laplace value. The exact one-sided bound (6.49) is

$$\lambda_{0.95} = \frac{\chi^2_{0.95;12}}{6000} = \frac{21.026}{6000} = 0.003504\,\mathrm{h}^{-1},$$

$0.000410\,\mathrm{h}^{-1}$ below the Laplace bound and still above the limit: the two routes differ in the support, not in the decision, and the Laplace route is the lower for the reason the interval is shifted, the weight its normal puts on large rates. Figure 6.4(b) draws both posteriors again; the blue area to the left of the dashed line at $3 \times 10^{-3}\,\mathrm{h}^{-1}$ is the Laplace probability, and the red lines are the two bounds, both to the right of the limit.

6.3.2 The count in a window, the mission and the next failure

Engineering questions. How many failures in the next 500 h? What is the expected number? No failure in the next 500 h? How long until the next failure? The questions now change from the unknown rate to what the pump will do: the number of failures in the next $\tau = 500$ operating hours, its mean, the probability that the pump completes those hours without a failure, which the maintenance manager asks for, and the time from now to the next failure.

General form. A prediction of a future observation y_{new} is its probability given the state, averaged over the posterior of the state,

$$p(y_{\text{new}} \mid \mathbf{y}) = \int p(y_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \tag{6.52}$$

and its moments are the conditional moments averaged as in equation (6.36), $\mathbb{E}[y_{\text{new}} \mid \mathbf{y}] = \int \mathbb{E}[y_{\text{new}} \mid \mathbf{z}]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$.

Specialized form. One choice enters here: the next 500 operating hours are taken to share the rate of the 2000 already run, the pump keeping its duty and being repaired as before. That

makes every future count a Poisson count at a rate drawn from the same posterior; if the intensity rose with age, the prediction would need the power law of Section 6.3.4 or the trend model of Chapter 12.

The count in the window. The question is equation (6.52) with $y_{\text{new}} = K_f$, the number of failures in the next τ operating hours, the state $\mathbf{z} = \omega$, $p(y_{\text{new}} \mid \mathbf{z}) = \Pr(K_f = j \mid \omega)$, the Poisson probability of equation (6.8) with the mean $\tau\lambda(\omega)$, $\lambda(\omega) = \lambda_{\text{ref}}\mathrm{e}^{\omega}$, and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal $\mathcal{N}(\hat{\omega}, \sigma_\omega^2)$ of step 3. The integral over the one entry is the Gauss–Hermite sum of Chapter 5, equation (5.33), with ω in the place of the logit: a rule of Q nodes t_q and weights w_q gives

$$\Pr(K_f = j \mid \mathbf{y}) \approx \frac{1}{\sqrt{\pi}} \sum_{q=1}^{Q} w_q \, h_j(\omega_q), \qquad h_j(\omega) = \frac{\bigl(\tau\lambda(\omega)\bigr)^j \, \mathrm{e}^{-\tau\lambda(\omega)}}{j!}, \qquad \omega_q = \hat{\omega} + \sqrt{2}\,\sigma_\omega t_q, \tag{6.53}$$

the nodes ω_q being values of the log-rate spread over its posterior and the sum weighing the Poisson probability at each; it is the route for a predictive without a closed form. The value beside it comes from the exact posterior. Written in λ, the same integral averages the Poisson probability over the gamma posterior,

$$\Pr(K_f = j \mid \mathbf{y}) = \int_0^\infty \Pr(K_f = j \mid \lambda)\, p(\lambda \mid \mathbf{y})\, \mathrm{d}\lambda;$$

substitute the Poisson probability and the gamma posterior,

$$\Pr(K_f = j \mid \mathbf{y}) = \int_0^\infty \frac{(\lambda\tau)^j \mathrm{e}^{-\lambda\tau}}{j!} \, \frac{b_n^{a_n}}{\Gamma(a_n)} \, \lambda^{a_n - 1} \mathrm{e}^{-b_n\lambda} \, \mathrm{d}\lambda,$$

and take the factors that do not depend on λ outside the integral:

$$\Pr(K_f = j \mid \mathbf{y}) = \frac{\tau^j \, b_n^{a_n}}{j!\,\Gamma(a_n)} \int_0^\infty \lambda^{a_n + j - 1} \, \mathrm{e}^{-(b_n + \tau)\lambda} \, \mathrm{d}\lambda.$$

The integral is the gamma integral (6.16) with the shape $a_n + j$ and the rate $b_n + \tau$. Therefore

$$\Pr(K_f = j \mid \mathbf{y}) = \frac{\Gamma(a_n + j)}{\Gamma(a_n)\, j!} \, \frac{\tau^j \, b_n^{a_n}}{(b_n + \tau)^{a_n + j}},$$

and, splitting the denominator between the two powers,

$$\Pr(K_f = j \mid \mathbf{y}) = \frac{\Gamma(a_n + j)}{\Gamma(a_n)\, j!} \, \pi^{a_n} (1 - \pi)^j, \qquad \pi = \frac{b_n}{b_n + \tau}, \quad j = 0, 1, 2, \ldots \tag{6.54}$$

This is the negative-binomial distribution with shape a_n and success probability π, the standard result quoted in Chapter 3; the derivation shows where the future exposure τ enters, as the extra rate $b_n + \tau$ of the gamma integral. It arises whenever a Poisson count is integrated over a gamma distribution of its rate, and it carries both kinds of uncertainty: the random occurrence of failures for a given λ and the uncertainty about λ, as the sum (6.53) does on the state route. For an integer shape the leading factor is the binomial coefficient $\binom{a_n + j - 1}{j}$, and successive terms follow from the first, π^{a_n}, by the ratio

$$\frac{\Pr(K_f = j + 1 \mid \mathbf{y})}{\Pr(K_f = j \mid \mathbf{y})} = \frac{a_n + j}{j + 1} \, (1 - \pi). \tag{6.55}$$

The probability of at most one failure is the sum of the first two terms, $\Pr(K_f \le 1 \mid \mathbf{y}) = \pi^{a_n}\bigl(1 + a_n(1 - \pi)\bigr)$, and its complement is the probability of two or more.

The expected number. The question is the moments of equation (6.52), $\mathbb{E}[y_{\text{new}} \mid \mathbf{y}] = \int \mathbb{E}[y_{\text{new}} \mid \mathbf{z}]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$, with $y_{\text{new}} = K_f$, the state ω and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal of step

3. They follow from the conditional ones, $\mathbb{E}[K_f \mid \lambda] = \mathrm{Var}(K_f \mid \lambda) = \lambda\tau$, by the laws of total expectation and total variance (Chapter 3, equations (3.31) and (3.32)):

$$\mathbb{E}[K_f \mid \mathbf{y}] = \tau\,\mathbb{E}[\lambda \mid \mathbf{y}], \tag{6.56}$$

$$\mathrm{Var}(K_f \mid \mathbf{y}) = \tau\,\mathbb{E}[\lambda \mid \mathbf{y}] + \tau^2\,\mathrm{Var}(\lambda \mid \mathbf{y}). \tag{6.57}$$

The first term of the variance is the Poisson variance the count would have if the rate were known; the second is the extra variance caused by the uncertainty about the rate. Both need only the mean and the variance of λ under the posterior. On the state route they are the Gauss–Hermite sums with $h(\omega) = \lambda(\omega)$ and $h(\omega) = \lambda(\omega)^2$, which for a normal in ω also have the closed lognormal form

$$\mathbb{E}[\lambda \mid \mathbf{y}] \approx \hat{\lambda}\,\mathrm{e}^{\sigma_\omega^2/2}, \qquad \mathrm{Var}(\lambda \mid \mathbf{y}) \approx \hat{\lambda}^2\,\mathrm{e}^{\sigma_\omega^2}(\mathrm{e}^{\sigma_\omega^2} - 1), \tag{6.58}$$

the mean of the density the normal in ω induces on the rate, not its median $\hat{\lambda}$, and a variance larger than the delta method's $\hat{\lambda}^2\sigma_\omega^2$, which linearizes the map. Exactly, they are the posterior moments (6.42) and (6.43), with which the rate's share of the total variance is $\tau/(b_n + \tau)$, growing with the future exposure. The mean is an average, not a certainty, and need not be the most probable count.

The mission. The question is equation (6.52) with y_{new} the event $T > \tau$, which by equation (6.10) is the event $K_f = 0$, so that $p(y_{\text{new}} \mid \mathbf{z}) = \mathrm{e}^{-\tau\lambda(\omega)}$ and the state route is the sum (6.53) with $h_0(\omega) = \mathrm{e}^{-\tau\lambda(\omega)}$. Exactly, equation (6.54) at $j = 0$ gives at once

$$R_{\text{pred}}(\tau) = \Pr(T > \tau \mid \mathbf{y}) = \mathbb{E}[\mathrm{e}^{-\lambda\tau} \mid \mathbf{y}] = \left(\frac{b_n}{b_n + \tau}\right)^{a_n}, \tag{6.59}$$

the exact posterior predictive survival function derived as an average of $\mathrm{e}^{-\lambda\tau}$ over the gamma posterior in Chapter 2, equation (2.41). In the exposure reading of this chapter the formula has a direct meaning: the mission adds τ hours to the rate parameter b_n, and its survival probability is the ratio of the record's exposure to the extended exposure, raised to the number of failures and prior pseudo-failures a_n. A common shortcut substitutes the posterior mean rate into the exponential survival function, $R_{\text{plug-in}}(\tau) = \mathrm{e}^{-\mathbb{E}[\lambda|\mathbf{y}]\,\tau}$. The two differ because survival is a nonlinear function of the rate,

$$\mathbb{E}[\mathrm{e}^{-\lambda\tau} \mid \mathbf{y}] \neq \mathrm{e}^{-\mathbb{E}[\lambda|\mathbf{y}]\,\tau}, \tag{6.60}$$

and since $\mathrm{e}^{-\lambda\tau}$ is convex in λ, Jensen's inequality makes the predictive value the larger one, as derived in Chapter 2 (Section 2.4.5). The inequality does not mean that uncertainty improves the pump; it states the difference between averaging the survival probabilities of the rates the posterior admits and evaluating survival at their average rate. The low rates raise the average survival by more than the high rates lower it.

Common mistake

Inserting the estimated rate into $\mathrm{e}^{-\lambda\tau}$ and reporting the result as the probability of a failure-free mission. The plug-in value ignores the uncertainty in λ; the predictive probability is the survival function averaged over the posterior, equation (6.59).

The next failure. The question is equation (6.52) with $y_{\text{new}} = T$, the time from now to the next failure, which given the state is exponential with the rate $\lambda(\omega)$, and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal of step 3. Its predictive survival function and its moments are Gauss–Hermite sums of the form (6.53) with $h(\omega) = \mathrm{e}^{-\lambda(\omega)t}$, $1/\lambda(\omega)$ and $2/\lambda(\omega)^2$, since $\mathbb{E}[T \mid \lambda] = 1/\lambda$ and $\mathbb{E}[T^2 \mid \lambda] = 2/\lambda^2$; for a normal in ω the two moments are again lognormal,

$$R_{\text{pred}}(t) \approx \mathbb{E}[\mathrm{e}^{-\lambda(\omega)t} \mid \mathbf{y}], \qquad \mathbb{E}[T \mid \mathbf{y}] \approx \frac{\mathrm{e}^{\sigma_\omega^2/2}}{\hat{\lambda}}, \qquad \mathbb{E}[T^2 \mid \mathbf{y}] \approx \frac{2\,\mathrm{e}^{2\sigma_\omega^2}}{\hat{\lambda}^2}, \tag{6.61}$$

the predictive median solves $R_{\mathrm{pred}}(t) = 1/2$, and the predictive hazard is the ratio of the sums with $h(\omega) = \lambda(\omega)\mathrm{e}^{-\lambda(\omega)t}$ and $h(\omega) = \mathrm{e}^{-\lambda(\omega)t}$. The values beside them come from the exact posterior. Equation (6.59), read as a function of the time t from now, is the predictive survival function of the time to the next failure, $R_{\mathrm{pred}}(t) = (b_n/(b_n + t))^{a_n}$, from any instant. Its predictive density is the negative derivative,

$$f_{\mathrm{pred}}(t) = -\frac{\mathrm{d}R_{\mathrm{pred}}(t)}{\mathrm{d}t} = \frac{a_n\, b_n^{a_n}}{(b_n + t)^{a_n+1}}, \tag{6.62}$$

the Lomax, or Pareto type II, distribution. Its mean is the integral of the survival function, which exists for $a_n > 1$:

$$\mathbb{E}[T \mid \mathbf{y}] = \int_0^\infty \left(\frac{b_n}{b_n + t}\right)^{a_n} \mathrm{d}t = \left[-\frac{b_n^{a_n}}{(a_n - 1)(b_n + t)^{a_n-1}}\right]_0^\infty = \frac{b_n}{a_n - 1}. \tag{6.63}$$

The same number is the posterior mean of $1/\lambda$, since $\mathbb{E}[T \mid \lambda] = 1/\lambda$ and, by the gamma integral, $\mathbb{E}[1/\lambda \mid \mathbf{y}] = b_n\Gamma(a_n - 1)/\Gamma(a_n) = b_n/(a_n - 1)$. It is not the reciprocal b_n/a_n of the posterior mean rate, because the reciprocal is nonlinear, another instance of equation (6.60). The band of the prediction follows from the second moment: given λ the time is exponential with $\mathbb{E}[T^2 \mid \lambda] = 2/\lambda^2$, the gamma integral gives $\mathbb{E}[1/\lambda^2 \mid \mathbf{y}] = b_n^2/((a_n - 1)(a_n - 2))$ for $a_n > 2$, and

$$\mathrm{Var}(T \mid \mathbf{y}) = \frac{2b_n^2}{(a_n - 1)(a_n - 2)} - \frac{b_n^2}{(a_n - 1)^2}, \tag{6.64}$$

which carries both the exponential scatter of failure times for a known rate and the uncertainty about the rate. The predictive median, the solution of $R_{\mathrm{pred}}(t) = 1/2$, is $b_n(2^{1/a_n} - 1)$, against $\ln 2/\mathbb{E}[\lambda \mid \mathbf{y}]$ for the plug-in exponential. By equation (6.4) the hazard of the predictive distribution is the ratio of equations (6.62) and (6.59),

$$\lambda_{\mathrm{pred}}(t) = \frac{f_{\mathrm{pred}}(t)}{R_{\mathrm{pred}}(t)} = \frac{a_n}{b_n + t}, \tag{6.65}$$

which decreases with elapsed time. The pump does not become healthier as it runs. Every fixed value of λ in the model has a constant hazard; the decrease comes from integrating out the uncertainty about λ. The form of equation (6.65) shows the mechanism exactly: t failure-free hours are one more record, zero failures in an exposure t, and by the gamma update they turn the posterior into Gamma$(a_n, b_n + t)$, whose mean is $a_n/(b_n + t)$. The predictive hazard is the posterior mean rate given survival to t: as the pump survives, the lower rates become relatively more plausible.

Common mistake

Reading the decreasing predictive hazard of equation (6.65) as evidence that the pump becomes more reliable with age. The model assumes a constant hazard for every value of the rate; the decrease is the effect of updating an uncertain rate on the failure-free time, and it vanishes as the posterior concentrates.

Building blocks. The MAP state $\hat{\omega} = 0.6931$ and the standard deviation $\sigma_\omega = 0.4082$ of equations (6.26) and (6.27); the five Gauss–Hermite nodes of Table 6.6, and forty for the reported values; the future exposure $\tau = 500\,\mathrm{h}$; and, for the exact values beside them, the posterior parameters $a_n = 6$ and $b_n = 3000\,\mathrm{h}$, the success probability π, and the posterior moments of equations (6.42) and (6.43).

The rule of five nodes. Table 6.6 evaluates the sum (6.53) with five nodes. Each row is a node: the first two columns give the node t_q and its weight $w_q/\sqrt{\pi}$, the third the log-rate $\omega_q = 0.6931 + \sqrt{2} \times 0.4082\, t_q$ at which the sum looks, the fourth the rate λ_q there and the fifth the expected count $\mu_q = 500\lambda_q$ of the next 500 h, and the last three the functions the questions of this subsection average: the probabilities of no failure and of one failure, $\mathrm{e}^{-\mu_q}$ and $\mu_q \mathrm{e}^{-\mu_q}$, and the mean time to failure $1/\lambda_q$ at that rate. The middle node is the MAP state itself; the others reach about two standard deviations either side of it, rates from 0.62 to $6.4 \times 10^{-3}\,\mathrm{h}^{-1}$.

Table 6.6: The Gauss–Hermite rule (6.53) with five nodes over the Laplace posterior $\omega \mid \mathbf{y} \approx \mathcal{N}(0.6931,\ 0.4082^2)$ of the pump; $\lambda_q = 10^{-3}\,\mathrm{e}^{\omega_q}$ in $10^{-3}\,\mathrm{h}^{-1}$, $\mu_q = 500\lambda_q$.

Node t_q	Weight $w_q/\sqrt{\pi}$	ω_q	λ_q	μ_q	$\mathrm{e}^{-\mu_q}$	$\mu_q \mathrm{e}^{-\mu_q}$	$1/\lambda_q$ [h]
−2.0202	0.01126	−0.4732	0.6230	0.3115	0.73235	0.22813	1605.1
−0.9586	0.22208	0.1397	1.1499	0.5750	0.56272	0.32355	869.6
0	0.53333	0.6931	2.0000	1.0000	0.36788	0.36788	500.0
0.9586	0.22208	1.2466	3.4784	1.7392	0.17566	0.30551	287.5
2.0202	0.01126	1.8595	6.4205	3.2103	0.04035	0.12952	155.8
Weighted sum			2.1738	1.0869	0.36888	0.33993	543.5

Failures in the next 500 hours. The sum over the sixth column of Table 6.6 is the probability of no failure,

$$\Pr(K_f = 0 \mid \mathbf{y}) \approx 0.01126 \times 0.73235 + 0.22208 \times 0.56272 + 0.53333 \times 0.36788 + 0.22208 \times 0.17566 + 0.01126 \times 0.04035 = 0.36888,$$

and over the seventh, $\Pr(K_f = 1 \mid \mathbf{y}) \approx 0.33993$; forty nodes change no digit shown. Each node is a rate the posterior admits, and the sum weighs the Poisson probability at each: the middle node alone, 0.36788, is the plug-in value at $\hat{\lambda}$. The probability of at most one failure is $0.36888 + 0.33993 = 0.7088$, and of two or more 0.2912. The value beside it comes from the exact posterior. With $a_n = 6$, $b_n = 3000\,\mathrm{h}$ and $\tau = 500\,\mathrm{h}$, $\pi = 3000/3500 = 6/7$ and $1 - \pi = 1/7$. The probability of no failure is $\pi^6 = 0.39657$, and the ratio (6.55) gives the rest; for one failure, $6 \times 0.39657/7 = 0.33992$. At most one failure has the probability $0.39657 + 0.33992 = 0.7365$, and two or more 0.2635: the state route gives no failure 0.028 less and two or more 0.028 more, the weight the normal in the log-rate puts on large rates. Table 6.7 lists the predictive for every count, from the sum (6.53) with forty nodes and exactly from equation (6.54), with the plug-in Poisson probabilities at $\hat{\lambda} = 0.002\,\mathrm{h}^{-1}$ beside them; the row of five or more aggregates the remaining counts. Figure 6.5 draws the same numbers: the horizontal axis is the number of failures in the next 500 h, the vertical axis its probability in percent, the blue bars the Laplace predictive, the green bars the exact negative-binomial predictive and the gray bars the plug-in Poisson. The pump is most likely to experience zero or one failure on both routes, although larger counts are possible, and both predictives put more weight on three or more than the plug-in calculation does.

The expected number. By equation (6.58), or the sum over the fourth column of Table 6.6, $\mathbb{E}[\lambda \mid \mathbf{y}] \approx 0.002\,\mathrm{e}^{1/12} = 0.002174\,\mathrm{h}^{-1}$ and $\mathrm{sd}(\lambda \mid \mathbf{y}) \approx 0.002174\sqrt{\mathrm{e}^{1/6} - 1} = 0.000926\,\mathrm{h}^{-1}$, so that by equations (6.56) and (6.57) the mean is $500 \times 0.002174 = 1.087$ failures, not the most probable outcome, which is zero, and the variance is $1.087 + 500^2 \times 0.000926^2 = 1.087 + 0.214 = 1.301$, a predictive standard deviation of 1.14, of which the rate's share of the variance is 0.165. The values beside them come from the exact posterior: the mean is $500 \times 0.002 = 1.0$ failure and

Table 6.7: Predicted number of failures of the pump in the next 500 operating hours. Laplace predictive, the Gauss–Hermite sum (6.53) over $\omega \mid \mathbf{y} \approx \mathcal{N}(0.6931,\ 0.4082^2)$; exact negative-binomial predictive from Gamma$(6, 3000\,\mathrm{h})$, equation (6.54) with $\pi = 6/7$; and the plug-in Poisson$(500 \times 0.002) =$ Poisson(1).

Failures j in 500 h [–]	Laplace predictive [–]	Exact predictive [–]	Plug-in Poisson [–]
0	0.36888	0.39657	0.36788
1	0.33993	0.33992	0.36788
2	0.18049	0.16996	0.18394
3	0.07340	0.06475	0.06131
4	0.02564	0.02081	0.01533
5 or more	0.01166	0.00800	0.00366
at most 1	0.70881	0.73649	0.73576
mean, variance	1.087, 1.301	1.000, 1.167	1.000, 1.000



Figure 6.5: Predicted number of pump failures in the next 500 operating hours (Table 6.7). Blue: Laplace predictive, the sum (6.53) over $\mathcal{N}(0.6931,\ 0.4082^2)$. Green: exact negative-binomial predictive, equation (6.54) with $a_n = 6$, $b_n = 3000\,\mathrm{h}$, $\tau = 500\,\mathrm{h}$. Gray: Poisson with $\hat{\lambda} = 0.002\,\mathrm{h}^{-1}$ substituted for λ. The last category aggregates five or more.

the variance $500 \times 0.002 + 500^2 \times 6.667 \times 10^{-7} = 1 + 0.1667 = 1.1667$, a predictive standard deviation of 1.08, of which the rate's share is $500/3500 = 0.143$. The state route expects 0.087 more failures, because the mean of the rate under the normal in ω lies above its median.

No failure in the next 500 hours. The sum of Section 6.3.2 with h_0 gives

$$\Pr(T > 500\,\mathrm{h} \mid \mathbf{y}) \approx 0.3689,$$

a 36.9 % probability of completing the 500-hour mission and 63.1 % of one or more failures in it, against the plug-in value $\mathrm{e}^{-0.002\times500} = \mathrm{e}^{-1} = 0.3679$ at $\hat{\lambda}$. It lies so close to the plug-in, not above it by the margin Jensen's inequality (6.60) suggests, because the inequality compares the predictive with the plug-in at the posterior mean, and the Laplace mean of the rate is $0.002174\,\mathrm{h}^{-1}$: $\mathrm{e}^{-500\times0.002174} = 0.337 < 0.369$. The value beside it comes from the exact posterior: by equation (6.59),

$$\Pr(T > 500\,\mathrm{h} \mid \mathbf{y}) = \left(\frac{3000}{3500}\right)^6 = 0.3966,$$

39.7 %, 0.028 higher, and 60.3 % of one or more failures.

The time to the next failure. By equation (6.61), or the sum over the last column of Table 6.6, $\mathbb{E}[T \mid \mathbf{y}] \approx \mathrm{e}^{1/12}/0.002 = 543.5\,\mathrm{h}$, against the reciprocal $1/\hat{\lambda} = 500\,\mathrm{h}$ of the median rate, and

$$\mathrm{Var}(T \mid \mathbf{y}) \approx \frac{2\,\mathrm{e}^{1/3}}{0.002^2} - 543.5^2 = 697\,806 - 295\,340 = 402\,466\,\mathrm{h}^2,$$

a predictive standard deviation of 634 h, larger than the mean. The predictive median, the root of $R_{\mathrm{pred}}(t) = 1/2$ by the forty-node sum, is 339 h, against $\ln 2/0.002 = 347\,\mathrm{h}$ for the plug-in exponential. The predictive hazard is $0.00217\,\mathrm{h}^{-1}$ now, $0.00184\,\mathrm{h}^{-1}$ after 500 h without failure and $0.00117\,\mathrm{h}^{-1}$ after 3000 h. The values beside them come from the exact posterior. By equation (6.63), $\mathbb{E}[T \mid \mathbf{y}] = 3000/5 = 600\,\mathrm{h}$, against the reciprocal $1/\mathbb{E}[\lambda \mid \mathbf{y}] = 500\,\mathrm{h}$ of the mean rate. By equation (6.64),

$$\mathrm{Var}(T \mid \mathbf{y}) = \frac{2 \times 3000^2}{20} - 600^2 = 540\,000\,\mathrm{h}^2,$$

a predictive standard deviation of 735 h. The predictive median is $3000(2^{1/6} - 1) = 367\,\mathrm{h}$. By equation (6.65) the predictive hazard is $0.002\,\mathrm{h}^{-1}$ now, $0.00171\,\mathrm{h}^{-1}$ after 500 h without failure and $0.001\,\mathrm{h}^{-1}$ after 3000 h. The state route gives a mean time 57 h shorter, a median 29 h shorter and a standard deviation 100 h smaller: the exact posterior of ω has its longer tail toward small rates, and the long times to failure those rates produce are what the normal in ω leaves out.

Figure 6.6 draws the two predictive functions. Panel (a) has the mission length on the horizontal axis and the probability of no failure on the vertical axis; the dotted blue curve is the Laplace predictive survival, the sum (6.53) at each t, the green curve the exact predictive survival of equation (6.59), and the dashed gray curve the plug-in exponential $\mathrm{e}^{-0.002t}$; the dotted line marks the 500-hour mission. The exact curve lies above the plug-in everywhere; the Laplace curve lies just below it up to about 480 h and above it after, and for both the gap grows with the mission: at 2000 h they are 0.035 and 0.047 against 0.018. Panel (b) has the same horizontal axis and the hazard in $10^{-3}\,\mathrm{h}^{-1}$ on the vertical axis; the blue and green curves are the decreasing predictive hazards and the gray line the constant plug-in hazard. The Laplace hazard starts above the plug-in, at the Laplace mean rate, and falls below it after about 230 h.

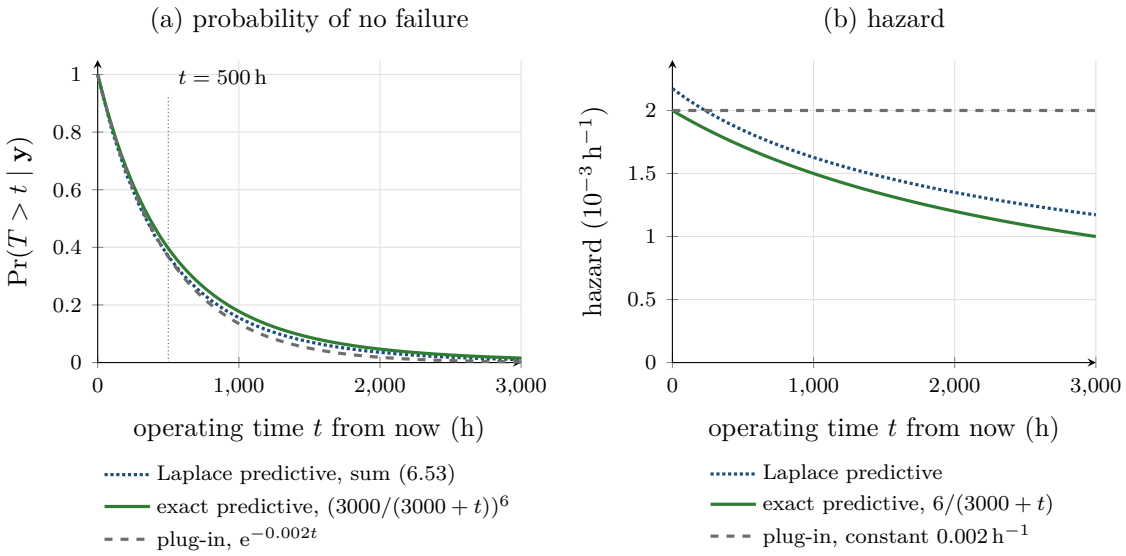


Figure 6.6: The predictive time to the next failure of the pump. Dotted blue: from the Laplace posterior $\omega \mid \mathbf{y} \approx \mathcal{N}(0.6931,\ 0.4082^2)$ by the forty-node sum (6.53), which the answers report; green: from the exact posterior Gamma$(6, 3000\,\mathrm{h})$. (a) Predictive survival against the plug-in exponential with $\lambda = 0.002\,\mathrm{h}^{-1}$; at $t = 500\,\mathrm{h}$ they are 0.3689, 0.3966 and 0.3679. (b) Predictive hazard, exact $a_n/(b_n + t)$ of equation (6.65), against the constant plug-in hazard.

6.3.3 The fleet and the two loads

Engineering questions. What is the rate of the three pumps, and pump B's own? Are the three pumps alike? What is the rate under normal and under high load? The first asks for the rate of the fleet of Figure 6.1(b), from which the expected number of failures of the three pumps together, 3λ per hour, sizes spares and maintenance, and for the rate of the pump with the least exposure if the pumps may differ; the second whether one shared rate or a rate per pump better describes the three records; the third for the reference rate of the pump that ran under two loads and for its rate under high load.

General form. The estimation form (6.35), now for a quantity derived from a state of several entries. When f is a monotone function of one linear combination $\mathbf{g}^\mathsf{T}\mathbf{z}$ of the entries, its quantiles are the images of that combination's quantiles, and under the Laplace posterior the combination is normal,

$$\mathbf{g}^\mathsf{T}\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\mathbf{g}^\mathsf{T}\hat{\mathbf{z}},\ \mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}), \tag{6.66}$$

with no approximation beyond the Gaussian itself. A hypothesis comparison is the ratio of the evidences of two complete drawings fitted to the same records, the Bayes factor $B_{10} = p(\mathbf{y} \mid M_1)/p(\mathbf{y} \mid M_0)$ of Chapter 2, equation (2.14), reported in nats as $\ln B_{10}$; a positive value favors M_1.

Specialized form. Three modeling choices enter here. The three pumps are taken as comparable, their rates held within a factor of about 1.6 of one another by the tie rows at $\sigma_d = 0.5$, or equal in the pooled limit: this is what lends each pump the others' hours, and it fails if the pumps differ in age, duty, maintenance or environment; the dispersion check of Section 6.4.6 and the comparison below test it, though six failures cannot settle it, and σ_d itself is chosen, not learned (Section 6.6.5). Both drawings of the comparison put the prior Gamma$(2, 1000\,\mathrm{h})$ on every rate: this gives each evidence a closed form, and it makes $\ln B_{10}$ move with the prior, since the separate drawing spreads three priors where the shared one

spreads one (Section 6.4.3). The high-load intensity is taken as twice the normal one, known or within thirty percent: this lets two failures of unrecorded load be counted against one effective exposure, and it misleads if the factor is chosen to bring records into agreement rather than taken from a physical or engineering model (Section 6.6.5); the known factor is checked as the limit $\sigma_c \to 0$ in Section 6.4.5.

The fleet. The question is equation (6.66) with the state $\mathbf{z}$ of the fleet, f the exponential of a combination $\mathbf{g}^\mathsf{T}\mathbf{z}$ of its entries, and $p(\mathbf{z} \mid \mathbf{y})$ the posterior of step 3, the Laplace normal $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of the rows of pattern (6.23). If the three pumps genuinely share one failure intensity, their records are disjoint exposures of one process, and the state is the one entry ω with $\mathbf{g} = 1$: the count rows on it add through their totals (Section 6.2.1), so that equations (6.26) and (6.27) with $a_n = a + \sum_i k_i$ and $b_n = b + \sum_i h_i$ give $\hat{\omega}$ and $\sigma_\omega^2 = 1/a_n$, and the band of the shared rate is equation (6.66) mapped through the exponential, the interval (6.39). The value beside it comes from the exact posterior, the gamma posterior of the pooled record, $\text{Gamma}(a + \sum_i k_i,\ b + \sum_i h_i)$ (Section 6.2.2). On the four-entry state (6.12), pump i has the expected count $\mu_i = h_i \lambda_{\text{ref}} \mathrm{e}^{\omega + d_i}$, so that $\ln \mu_i = \ln(h_i \lambda_{\text{ref}}) + \omega + d_i$: the predictor Jacobian of its count row has a one in the ω entry and in the d_i entry, and its curvature term vanishes. The 7×4 Jacobian of pattern (6.23) stacks the prior row $\sqrt{b\lambda}$ in the ω column, the count rows $\sqrt{\mu_i}$ in the ω column and in their own column, and the tie rows $1/\sigma_d$ in their own column. The common rate is the limit $\sigma_d \to 0$: an exact constraint carried as a row whose tolerance vanishes, checked numerically in Section 6.4.5, and never used to eliminate the d_i. The fleet's rate is $\lambda_{\text{ref}} \mathrm{e}^{\omega}$, reported on the pooled drawing, where its exact posterior is the pooled gamma one. Pump B's own log-rate, reported at $\sigma_d = 0.5$, is equation (6.66) again on the four entries, with the derived quantity $\omega + d_B = \mathbf{g}^\mathsf{T}\mathbf{z}$, $\mathbf{g} = (1, 0, 1, 0)^\mathsf{T}$, and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal of the four-entry fit; it has no exact counterpart,

$$\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g} = \Sigma_{\omega\omega} + 2\,\Sigma_{\omega B} + \Sigma_{BB}, \qquad \lambda_B \in \lambda_{\text{ref}}\, \exp\Big(\mathbf{g}^\mathsf{T}\hat{\mathbf{z}} \pm 1.96\sqrt{\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}}\Big), \tag{6.67}$$

the second the Laplace band of equation (6.66) mapped through the exponential. The shared column carries the relation between the entries. The count row of pump B has equal entries in the ω and the d_B columns, so it informs their sum and leaves their difference to the prior and the tie row; the covariance $\Sigma_{\omega B}$ is the trade between the shared level and the deviation that this leaves, and it enters the variance of pump B's log-rate through the cross term.

The comparison of the pumps. The question is the Bayes factor $B_{10} = p(\mathbf{y} \mid M_1)/p(\mathbf{y} \mid M_0)$ of equation (2.14) with M_1 the drawing of one shared rate and M_0 the drawing of a rate per pump, each fitted to the three records. The evidence of one count record under the drawing with a gamma prior on its rate is the probability of the record averaged over the prior, the denominator of Bayes' theorem evaluated in closed form in Chapter 3 (equation (3.13)): with the Poisson probability (6.7) and the prior (6.15), the integrand collects into $\lambda^{a+k-1}\mathrm{e}^{-(b+T_{\text{tot}})\lambda}$ and the gamma integral (6.16) gives

$$\Pr(K = k \mid T_{\text{tot}}) = \frac{T_{\text{tot}}^k}{k!}\,\frac{b^a}{\Gamma(a)}\,\frac{\Gamma(a+k)}{(b+T_{\text{tot}})^{a+k}} = \frac{\Gamma(a+k)}{\Gamma(a)\,k!}\left(\frac{b}{b+T_{\text{tot}}}\right)^a\left(\frac{T_{\text{tot}}}{b+T_{\text{tot}}}\right)^k, \tag{6.68}$$

the negative-binomial distribution (6.54) with the prior in place of the posterior and the recorded exposure in place of the future one. Under the drawing M_0, a rate per pump, each under its own prior $\text{Gamma}(a, b)$, the three records are independent and the evidence is the product of three values of equation (6.68), one per pump. Under the drawing M_1, one shared rate, the three Poisson probabilities share λ and collect, as for the periods of one record (Section 6.2.1), into $\lambda^K \mathrm{e}^{-\lambda H}$ with $K = \sum_i k_i$ and $H = \sum_i h_i$, so that

$$p(\mathbf{y} \mid M_1) = \prod_i \frac{h_i^{k_i}}{k_i!}\,\frac{b^a}{\Gamma(a)}\,\frac{\Gamma(a+K)}{(b+H)^{a+K}}, \qquad p(\mathbf{y} \mid M_0) = \prod_i \Pr(K_i = k_i \mid h_i),$$
$$\ln B_{10} = \ln p(\mathbf{y} \mid M_1) - \ln p(\mathbf{y} \mid M_0). \tag{6.69}$$

Both are exact, and $\ln B_{10}$ is the reported answer. This is the one question of the chapter with no state route: an evidence needs the constants of the rows, which the objective of step 3 drops, and the Laplace evidence that restores them belongs to the later chapter on model selection. The shared drawing spends one rate on three records and the separate drawing three; the ratio weighs how well each predicts the counts against how widely the separate drawing spreads its prior over rates the records did not need.

The two loads. The question is equation (6.66) with the state of the pump that ran under two loads, f the exponential of the reference log-rate or of the high-load log-rate, and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal of its rows. With the relative intensities (6.13) known, the state is the one entry ω_0, $K_i \mid \lambda_0 \sim \text{Poisson}(\lambda_0 c_i h_i)$, and the count rows are rows on one entry with the exposures $c_i h_i$, which add through their totals as the periods of one record do. Define the effective exposure, in equivalent reference hours,

$$T_{\text{eff}} = \sum_i c_i h_i. \tag{6.70}$$

Equations (6.26) and (6.27) with $a_n = a + \sum_i k_i$ and $b_n = b + T_{\text{eff}}$ give $\hat{\omega}_0$ and $\sigma^2_{\omega_0} = 1/a_n$, and the high-load log-rate $\omega_0 + \ln c_2$, with $\mathbf{g} = 1$ and a known shift, has the same variance. The value beside it comes from the exact posterior. The combined likelihood, collected as for the periods of one record, is

$$p(k_1, \ldots, k_I \mid \lambda_0) \propto \lambda_0^{\sum_i k_i} \exp\Big(-\lambda_0 \sum_i c_i h_i\Big) = \lambda_0^{\sum_i k_i} \, \mathrm{e}^{-\lambda_0 T_{\text{eff}}},$$

and the gamma update becomes

$$\lambda_0 \mid \mathbf{y} \sim \text{Gamma}\Big(a + \sum_i k_i, \; b + T_{\text{eff}}\Big). \tag{6.71}$$

The record need not say in which condition each failure occurred. The high-load intensity is $\lambda_{\text{high}} = c_2 \lambda_0$, and a gamma variable multiplied by a constant c remains gamma with its rate parameter divided by c, so that $\lambda_{\text{high}} \mid \mathbf{y}$ is the gamma posterior of λ_0 with the rate parameter $(b + T_{\text{eff}})/c_2$, of c_2 times the mean and the spread.

When the high-load factor is known only roughly, it is an entry of the state, $\mathbf{z} = (\omega_0, \; \ln c_2)^{\mathsf{T}}$, and three rows act on it: the prior row $\mathcal{L}_0 = b\lambda_0 - a\omega_0$ on the reference rate, as for the pump (Section 6.1.7); the one count row of the record, $\mathcal{L}_1 = \mu - k \ln \mu$, with the expected count $\mu = \mu_1 + \mu_2$, $\mu_1 = \lambda_0 h_1$ in normal load and $\mu_2 = \lambda_0 c_2 h_2$ in high load; and the Gaussian prior row on the factor, $r_3 = (\ln c_2 - \ln 2)/\sigma_c$. The known factor of equation (6.71) is the limit $\sigma_c \to 0$, checked in Section 6.4.5. The count row's log-mean now has the gradient and curvature

$$\frac{\partial \ln \mu}{\partial \mathbf{z}} = \begin{pmatrix} 1 \\ \mu_2/\mu \end{pmatrix}, \qquad \frac{\partial^2 \ln \mu}{\partial \mathbf{z}^2} = \begin{pmatrix} 0 & 0 \\ 0 & \mu_1 \mu_2/\mu^2 \end{pmatrix}, \tag{6.72}$$

so the curvature term of a count row whose log-mean is not linear in the state (Section 6.1.7) no longer vanishes in the factor's entry, and the step that keeps it is Newton's. The Laplace covariance is the inverse of the information plus that term,

$$\boldsymbol{\Sigma} = \left(\mathbf{J}^{\mathsf{T}}\mathbf{J} + \mathbf{C}\right)^{-1}, \qquad \mathbf{C} = (\hat{\mu} - k)\, \frac{\partial^2 \ln \mu}{\partial \mathbf{z}^2}, \tag{6.73}$$

evaluated at the MAP state. The count row's entry $\sqrt{\mu}\, \mu_2/\mu$ in the factor's column is the only link between the two entries, the shared row of pattern (6.23), and it produces their correlation: a larger factor explains the same failures with a smaller reference rate. The reference rate is read off as $\lambda_{\text{ref}} \mathrm{e}^{\hat{\omega}_0}$ with the spread $\hat{\lambda}_0 \sqrt{\Sigma_{11}}$, and the high-load intensity $\lambda_{\text{high}} = \lambda_{\text{ref}} \mathrm{e}^{\omega_0 + \ln c_2}$ is a derived quantity with the gradient $\mathbf{g} = \hat{\lambda}_{\text{high}}(1, 1)^{\mathsf{T}}$ and the spread $\sqrt{\mathbf{g}^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g}}$; its band is equation (6.66) with the combination $(1, 1)^{\mathsf{T}}$ of the two entries, mapped through the exponential. This posterior has no exact counterpart; its known-factor limit is the gamma posterior above.

Building blocks. For the fleet, the pooled one-entry fit, $\hat{\omega} = \ln 2$ and $\sigma_\omega = 1/\sqrt{8}$, the four-entry fit with its 7×4 Jacobian, $\hat{\mathbf{z}}$, $\boldsymbol{\Sigma}$ and the gradient $(1, 0, 1, 0)^\mathsf{T}$, and, for the exact value beside the pooled one, the pooled gamma parameters; for the comparison, the four evidences of equation (6.69); for the two loads, the effective exposure and the one-entry fit it gives, the two-entry Newton fit with the curvature term (6.73) and the gradient $\hat{\lambda}_{\text{high}}(1, 1)^\mathsf{T}$ (Section 6.1.7 for the rows), and, for the exact values beside the known-factor ones, the gamma parameters of λ_0 and λ_{high} with the chi-square quantiles. The limits $\sigma_d \to 0$ and $\sigma_c \to 0$ are checked in Section 6.4.5.

The rate of the three pumps. Under the shared rate the pooled record has $a_n = 2 + 6 = 8$ and $b_n = 1000 + 3000 = 4000\,\text{h}$, so the one-entry fit gives $\hat{\omega} = \ln(8/(4000 \times 0.001)) = \ln 2 = 0.6931$ and $\sigma_\omega = 1/\sqrt{8} = 0.3536$: the shared rate has the median $\hat{\lambda} = 0.002\,\text{h}^{-1}$, the delta-method standard deviation $0.002 \times 0.3536 = 0.000707\,\text{h}^{-1}$ and the central 95 % band

$$10^{-3}\,\text{e}^{0.6931 \pm 1.96 \times 0.3536} = [0.00100,\ 0.00400]\,\text{h}^{-1}.$$

The value beside it comes from the exact posterior: by equation (6.32),

$$\lambda \mid \mathbf{y} \sim \text{Gamma}(2 + 6,\ 1000 + 3000) = \text{Gamma}(8,\ 4000\,\text{h}),$$

of mean $8/4000 = 0.002\,\text{h}^{-1}$, with the standard deviation $\sqrt{8}/4000 = 0.000707\,\text{h}^{-1}$ and the central 95 % interval $[0.00086,\ 0.00361]\,\text{h}^{-1}$. The centers and spreads agree; the Laplace band is shifted up by $0.00014\,\text{h}^{-1}$ at the lower end and $0.00039\,\text{h}^{-1}$ at the upper, the pump's bias with eight events in place of six.

Pump B's own rate. At $\sigma_d = 0.5$ the Newton steps of equation (4.40) reach the MAP state from $\mathbf{z} = \mathbf{0}$ in four steps, to four decimals,

$$\hat{\mathbf{z}} = \begin{pmatrix} 0.6820 & 0.0037 & -0.1058 & 0.0965 \end{pmatrix}^\mathsf{T},$$

where the Jacobian is

$$\mathbf{J} = \begin{pmatrix} 1.406 & 0 & 0 & 0 \\ 1.409 & 1.409 & 0 & 0 \\ 1.193 & 0 & 1.193 & 0 \\ 1.617 & 0 & 0 & 1.617 \\ 0 & 2 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 2 \end{pmatrix},$$

in the order prior, counts of pumps A, B and C, and tie rows of the pattern (6.23), and the covariance is

$$\boldsymbol{\Sigma} = (\mathbf{J}^\mathsf{T}\mathbf{J})^{-1} = \begin{pmatrix} 0.1685 & -0.0559 & -0.0442 & -0.0666 \\ -0.0559 & 0.1856 & 0.0147 & 0.0221 \\ -0.0442 & 0.0147 & 0.1960 & 0.0175 \\ -0.0666 & 0.0221 & 0.0175 & 0.1775 \end{pmatrix}.$$

The shared log-rate has $\sqrt{\Sigma_{\omega\omega}} = 0.4105$, against $\sigma_\omega = 1/\sqrt{8} = 0.3536$ for the pooled one-entry fit. By equation (6.67),

$$\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g} = 0.1685 - 0.0884 + 0.1960 = 0.276, \qquad \sqrt{\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}} = 0.525,$$

and pump B's rate $\lambda_{\text{ref}}\text{e}^{0.5761} = 0.00178\,\text{h}^{-1}$ has the band $[0.00064,\ 0.0050]\,\text{h}^{-1}$, with no exact value beside it; it is read against the pooled band of the same route, $[0.00100,\ 0.00400]\,\text{h}^{-1}$,

than which it is wider and centered lower. The negative entries $\Sigma_{\omega i}$ are the trade between the shared level and each deviation: pump B's own record constrains their sum, whose variance 0.276 is smaller than the $0.1685 + 0.1960 = 0.3645$ the two entries would give if they were independent.

The comparison of the pumps. With $a = 2$ and $b = 1000\,\mathrm{h}$, equation (6.68) gives the three pumps the evidences 0.1875, 0.2743 and 0.1341 under their own rates, so that $p(\mathbf{y} \mid M_0) = 0.1875 \times 0.2743 \times 0.1341 = 0.00690$, while equation (6.69) with $K = 6$ and $H = 3000\,\mathrm{h}$ gives $p(\mathbf{y} \mid M_1) = 0.00886$. Then

$$\ln B_{10} = \ln 0.00886 - \ln 0.00690 = -4.726 + 4.976 = +0.25 \text{ nats}, \qquad B_{10} = 1.28,$$

a weak preference for one shared rate: the records do not tell the pumps apart, and they do not show them to be alike either. The value is exact and has no state-route counterpart in this chapter.

The two loads with known factors. By equation (6.70), $T_{\mathrm{eff}} = 800 \times 1.0 + 200 \times 2.0 = 1200$ equivalent reference hours, and with the two failures the one-entry fit has $a_n = 4$ and $b_n = 2200\,\mathrm{h}$, so that $\hat{\omega}_0 = \ln(4/2.2) = 0.5978$ and $\sigma_{\omega_0} = 1/\sqrt{4} = 0.5$. The reference intensity has the median $0.001818\,\mathrm{h}^{-1}$, the delta-method standard deviation $0.001818 \times 0.5 = 0.000909\,\mathrm{h}^{-1}$ and the band $10^{-3}\,\mathrm{e}^{0.5978 \pm 0.98} = [0.000682,\ 0.004844]\,\mathrm{h}^{-1}$: two failures leave the reference intensity uncertain by a factor of seven. The high-load intensity $2\lambda_0$ has the median $0.003636\,\mathrm{h}^{-1}$, the standard deviation $0.001818\,\mathrm{h}^{-1}$ and the band $[0.001365,\ 0.009689]\,\mathrm{h}^{-1}$, twice those of λ_0. The values beside them come from the exact posterior: equation (6.71) gives the reference intensity the posterior Gamma$(4,\ 2200\,\mathrm{h})$, of mean $4/2200 = 0.001818\,\mathrm{h}^{-1}$, standard deviation $\sqrt{4}/2200 = 0.000909\,\mathrm{h}^{-1}$ and central 95 % interval $[0.000495,\ 0.003985]\,\mathrm{h}^{-1}$, from the chi-square quantiles with eight degrees of freedom, a factor of eight. The high-load intensity $2\lambda_0$ has the posterior Gamma$(4,\ 1100\,\mathrm{h})$, of mean $0.003636\,\mathrm{h}^{-1}$, standard deviation $0.001818\,\mathrm{h}^{-1}$ and interval $[0.000991,\ 0.007970]\,\mathrm{h}^{-1}$. The centers and spreads agree; the Laplace bands are shifted up, by $0.000859\,\mathrm{h}^{-1}$ at the upper end of λ_0 and $0.001719\,\mathrm{h}^{-1}$ at that of λ_{high}, a larger shift than the pump's because four events skew the posterior more than six.

The two loads with an uncertain factor. From $\mathbf{z} = (0,\ \ln 2)^{\mathsf{T}}$ the Newton iteration with the curvature term (6.73) takes ω_0 through 0.8192, 0.6214 and 0.5991 to 0.5988 and $\ln c_2$ through 0.6876 to 0.6878. At the MAP state the 3×2 Jacobian of the working rows, in the order prior, count and factor prior, with the count row's log-mean derivatives (6.72), its information, the curvature term and the covariance are

$$\mathbf{J} = \begin{pmatrix} 1.349 & 0 \\ 1.476 & 0.490 \\ 0 & 3.333 \end{pmatrix}, \qquad \mathbf{J}^{\mathsf{T}}\mathbf{J} = \begin{pmatrix} 4.000 & 0.724 \\ 0.724 & 11.352 \end{pmatrix}, \qquad \mathbf{C} = \begin{pmatrix} 0 & 0 \\ 0 & 0.040 \end{pmatrix},$$

$$\mathbf{\Sigma} = \begin{pmatrix} 4.000 & 0.724 \\ 0.724 & 11.392 \end{pmatrix}^{-1} = \begin{pmatrix} 0.2529 & -0.0161 \\ -0.0161 & 0.0888 \end{pmatrix},$$

with the correlation $\Sigma_{12}/\sqrt{\Sigma_{11}\Sigma_{22}} = -0.107$. The reference log-rate is $\hat{\omega}_0 = 0.5988$ with $\sqrt{\Sigma_{11}} = 0.503$, so that $\lambda_0 = 0.001820\,\mathrm{h}^{-1}$ with $\hat{\lambda}_0\sqrt{\Sigma_{11}} = 0.000915\,\mathrm{h}^{-1}$, and the high-load intensity is $\lambda_{\mathrm{high}} = 0.003620\,\mathrm{h}^{-1}$ with $\sqrt{\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}} = 0.002014\,\mathrm{h}^{-1}$. In the log-rate the combination $(1, 1)^{\mathsf{T}}$ has the variance $0.2529 - 2 \times 0.0161 + 0.0888 = 0.3095$, a standard deviation of 0.5563, so that by equation (6.66) the bands are

$$\lambda_0 \in 10^{-3}\,\mathrm{e}^{0.5988 \pm 1.96 \times 0.503} = [0.000679,\ 0.004877]\,\mathrm{h}^{-1},$$

$$\lambda_{\text{high}} \in 10^{-3}\,\mathrm{e}^{1.2866 \pm 1.96 \times 0.5563} = [0.001217,\ 0.010773]\,\mathrm{h}^{-1},$$

with no exact posterior beside them. Against the known-factor bands of the same route, $[0.000682,\ 0.004844]$ and $[0.001365,\ 0.009689]\,\mathrm{h}^{-1}$, the thirty percent on the factor leaves the reference band almost unchanged and widens the high-load band by $0.001232\,\mathrm{h}^{-1}$.

6.3.4 The power-law intensity

Engineering question. What is the power-law shape of each record? Records A and B of Figure 6.2 have the pump's totals, and the question is what their failure times say about the shape β of the power-law intensity (6.14), and so whether the intensity rises with operating age.

General form. The estimation form (6.35) for an entry z_i of the state, summarized by its central 95 % band. For a quantity f that increases with the entry, the band is the image of the entry's quantiles, and under the Laplace posterior the entry is normal with its diagonal variance,

$$z_i \mid \mathbf{y} \approx \mathcal{N}(\hat{z}_i,\ \Sigma_{ii}), \qquad f \in \Big[f\big(\hat{z}_i - 1.96\sqrt{\Sigma_{ii}}\big),\ f\big(\hat{z}_i + 1.96\sqrt{\Sigma_{ii}}\big)\Big], \tag{6.74}$$

the case $\mathbf{g} = \mathbf{e}_i$ of equation (6.66). A hypothesis on the quantity, a value f_0, is answered by whether the band contains it.

Specialized form. Two choices enter here. The intensity is taken to follow the power law (6.14): this gives the trend one parameter, the shape, whose band says whether the rate is constant, and it misreads a record whose intensity jumps, after a disturbance or through a common cause, rather than growing smoothly with age (Section 6.6.5). No prior row is placed on the two entries, an implicit flat prior on $\ln m(T_{\text{tot}})$ and on v: this makes the fit a maximum-likelihood fit with an exact chi-square band, and it leaves the evidence of the power law undefined, so the comparison with the constant rate waits for a proper prior on the shape (Section 6.1.6).

The question is equation (6.74) with the entry $z_i = v$, the log-shape of the state of the power law (6.14), $f = \beta = \mathrm{e}^{v}$, the hypothesis $f_0 = 1$ of a constant rate, and $p(\mathbf{z} \mid \mathbf{y})$ the posterior of the record's two rows, the Laplace normal $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of the fit derived next. The value beside it comes from an exact band, derived after the fit.

When the intensity (6.1) changes with operating time, $\nu = \nu(t)$, the expected count is its integral, the cumulative intensity $m(t) = \int_0^t \nu(u)\,\mathrm{d}u$, and the derivation of Section 6.1.3 carries over with $\lambda\,\Delta t$ replaced by $\nu(u)\,\Delta t$ in each short interval: the count is Poisson with the cumulative intensity as its mean, $N(t) \sim \text{Poisson}(m(t))$. The probability of no failure becomes

$$\Pr(N(t) = 0) = \mathrm{e}^{-m(t)}, \tag{6.75}$$

which generalizes equation (6.10): by equation (6.4) the time to the first failure has the cumulative hazard $\Lambda(t) = m(t)$. The density of a record of event times $0 < t_1 < \cdots < t_k \le T_{\text{tot}}$ follows as for the homogeneous process (Section 6.2.1), each failure contributing its intensity and the whole exposure its probability of no further events (6.75),

$$p(t_1, \ldots, t_k \mid \nu) = \prod_{j=1}^{k} \nu(t_j) \times \mathrm{e}^{-m(T_{\text{tot}})}, \tag{6.76}$$

a standard result for the nonhomogeneous process.

With the power law (6.14) on the state $\mathbf{z} = (\ln m(T_{\text{tot}}),\ v)^{\mathsf{T}}$, $\nu(t_j) = m(T_{\text{tot}})\,\beta\,t_j^{\beta-1}/T_{\text{tot}}^{\beta}$, and equation (6.76) splits into two rows. The *count row* is the Poisson probability of k failures

with mean $m(T_{\text{tot}})$, and the *timing row* the density of the times given the count; their negative logarithms, without the terms independent of the state, are

$$\mathcal{L}_1(\mathbf{z}) = m(T_{\text{tot}}) - k \ln m(T_{\text{tot}}), \qquad \mathcal{L}_2(\mathbf{z}) = \mathrm{e}^v S - kv, \qquad S = \sum_j \ln(T_{\text{tot}}/t_j). \tag{6.77}$$

No prior row acts on either entry: the fit maximizes the likelihood, which is the MAP under an implicit flat prior on $\ln m(T_{\text{tot}})$ and on v, that is, a prior density proportional to $1/\beta$ on the shape; this is why the exact band under that prior, derived below, is the comparison for the Laplace band. The timing row has the form of a count row with k events in the exposure S, and the gradient and curvature of the objective $\mathcal{L} = \mathcal{L}_1 + \mathcal{L}_2$ are

$$\frac{\partial \mathcal{L}}{\partial \mathbf{z}} = \begin{pmatrix} m(T_{\text{tot}}) - k \\ \mathrm{e}^v S - k \end{pmatrix}, \qquad \frac{\partial^2 \mathcal{L}}{\partial \mathbf{z}^2} = \begin{pmatrix} m(T_{\text{tot}}) & 0 \\ 0 & \mathrm{e}^v S \end{pmatrix},$$

the second equal to $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ for the working rows of Section 6.1.7, $\mathbf{J} = \operatorname{diag}(\sqrt{m(T_{\text{tot}})},\ \sqrt{\mathrm{e}^v S})$, 2×2, the diagonal pattern (6.23). At the minimum both gradient entries vanish,

$$\hat{m}(T_{\text{tot}}) = k, \qquad \hat{\beta} = \mathrm{e}^{\hat{v}} = \frac{k}{\sum_j \ln(T_{\text{tot}}/t_j)}, \qquad \mathbf{J} = \operatorname{diag}(\sqrt{k},\ \sqrt{k}), \qquad \boldsymbol{\Sigma} = \operatorname{diag}(1/k,\ 1/k). \tag{6.78}$$

The zero off-diagonal entry is the relation the state makes visible: the count informs only the first entry and the event times only the second. The times enter through $\sum_j \ln t_j$, and they inform β. By equation (6.74), read off as $\mathrm{e}^{\hat{v} \pm 1.96\sqrt{\Sigma_{22}}}$, the Laplace 95 % band on the shape is

$$\beta \in \left[\hat{\beta}\, \mathrm{e}^{-1.96/\sqrt{k}},\ \hat{\beta}\, \mathrm{e}^{1.96/\sqrt{k}}\right]. \tag{6.79}$$

The value beside it comes from an exact band, which exists in closed form. Given k, the failure times of the power-law process over $(0, T_{\text{tot}}]$ are independent with distribution function $(t/T_{\text{tot}})^\beta$, so each $\beta \ln(T_{\text{tot}}/t_j)$ is a unit exponential time and $2\beta \sum_j \ln(T_{\text{tot}}/t_j) = 2k\beta/\hat{\beta}$ follows the chi-square distribution with $2k$ degrees of freedom. The central 95 % band is therefore

$$\beta \in \left[\frac{\chi^2_{0.025;\,2k}}{2k}\,\hat{\beta},\ \frac{\chi^2_{0.975;\,2k}}{2k}\,\hat{\beta}\right]. \tag{6.80}$$

The same band is the central 95 % credible interval of β under a prior density proportional to $1/\beta$, because the density of the times given k, $\prod_j \beta\, t_j^{\beta-1}/T_{\text{tot}}^\beta$, is proportional to $\beta^k \mathrm{e}^{-\beta \sum_j \ln(T_{\text{tot}}/t_j)}$, and with that prior the posterior of β is $\text{Gamma}\big(k,\ \sum_j \ln(T_{\text{tot}}/t_j)\big)$.

Building blocks. The MAP state of the two-entry fit of the rows (6.77), $\hat{v} = \ln \hat{\beta}$ with $\hat{\beta}$ of equation (6.78) from the sums $S = \sum_j \ln(T_{\text{tot}}/t_j)$ of the two records, and its variance $\Sigma_{22} = 1/k = 0.25$; and, for the exact bands beside them, the chi-square quantiles with $2k$ degrees of freedom.

The power-law shape of the two records. For record A the sum $\sum_j \ln(2000/t_j)$ is 3.29 and, by equation (6.78), $\hat{\beta} = 4/3.29 = 1.22$; for record B it is $S = 0.904$ and $\hat{\beta} = 4.43$. With $k = 4$, $\sqrt{\Sigma_{22}} = 0.5$ and the band (6.79) is

$$\beta \in \left[\hat{\beta}\, \mathrm{e}^{-0.98},\ \hat{\beta}\, \mathrm{e}^{0.98}\right] = [0.375\, \hat{\beta},\ 2.664\, \hat{\beta}].$$

For record A it is [0.46, 3.24], which contains the constant rate $\beta = 1$; for record B it is [1.66, 11.8], which excludes it. The values beside them come from the exact band. The 0.025

and 0.975 quantiles of chi-square with eight degrees of freedom are 2.180 and 17.535, and the band (6.80) is

$$\beta \in \left[\frac{2.180}{8}\,\hat{\beta},\ \frac{17.535}{8}\,\hat{\beta}\right] = [0.272\,\hat{\beta},\ 2.192\,\hat{\beta}],$$

[0.33, 2.66] for record A, which contains $\beta = 1$, and [1.21, 9.70] for record B, which excludes it. The two routes reach the same two conclusions; the Laplace bands sit 20 to 40 % higher, because the exact posterior of v is skewed toward small shapes, which its Gaussian in v does not represent. Four failures leave the shape wide in both cases, but only the clustered record points to a rising intensity, which the constant-rate posterior cannot see.

6.4 The checks (step 5)

Every answer of Section 6.3 rests on a route of calculation, a prior and a model of the pump, and each can be tested against an exact result or against the records.

6.4.1 The Laplace posterior against the exact posterior

The pump. Figure 6.7 compares the Laplace approximation in the log-rate with the exact posterior. Panel (a) has the log-rate on the horizontal axis and the objective (6.24), less its minimum, on the vertical axis; the solid curve is the exact objective $3\mathrm{e}^{\omega} - 6\omega$ and the dashed curve the parabola $\frac{1}{2} \times 6\,(\omega - \hat{\omega})^2$ that the Laplace approximation substitutes for it. Panel (b) has the rate on the horizontal axis and probability density on the vertical axis; the green curve is the exact Gamma(6, 3000) and the dotted blue curve the Laplace normal in ω transformed back to λ, a lognormal density. In panel (a) the two curves touch at the minimum with the same curvature; the exact objective rises more steeply toward large rates and more slowly toward small ones. In panel (b) the lognormal stays positive and is skewed to the right like the exact posterior, but it is shifted toward larger rates.

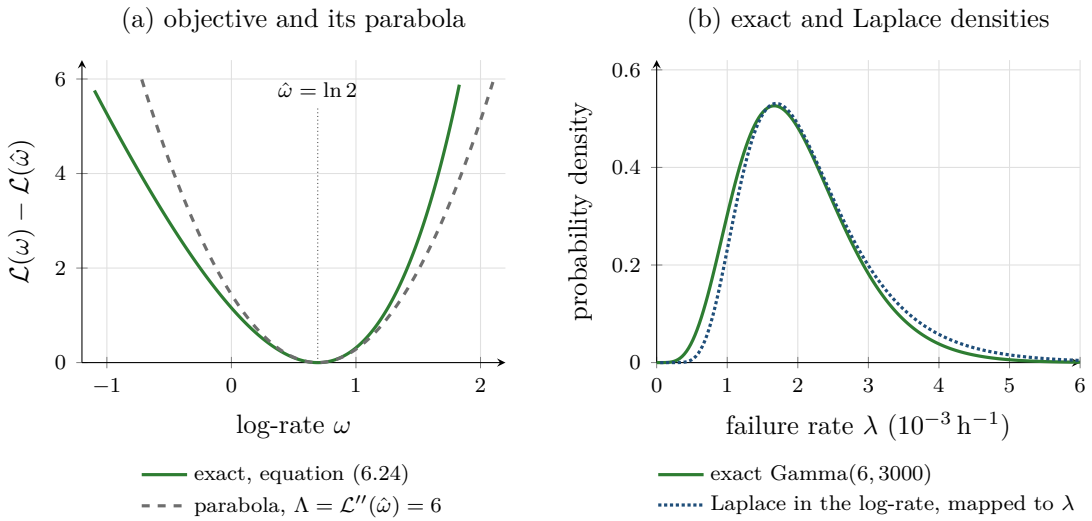


Figure 6.7: The Laplace approximation of the pump posterior, with $\lambda_{\mathrm{ref}} = 10^{-3}\,\mathrm{h}^{-1}$ and $x = 1000\lambda$. (a) The objective $3\mathrm{e}^{\omega} - 6\omega$ less its minimum $6 - 6\ln 2$, against the parabola $3(\omega - \ln 2)^2$. (b) Exact density $\frac{729}{120}x^5\mathrm{e}^{-3x}$, and the normal $\mathcal{N}(\ln 2,\ 1/6)$ on ω mapped to the rate, density $\varphi(\ln(x/2)/0.4082)/(0.4082\,x)$ with φ the standard normal density.

Table 6.8 turns the figure into numbers for the pump's questions, collecting the pairs that Section 6.3 reported question by question. Each row is a summary; the columns give it under the exact posterior and under the Laplace approximation, whose intervals and probabilities come

from the normal in ω mapped to the rate by equations (6.39) and (6.41), and whose predictions are the Gauss–Hermite sums (6.53) and the moments (6.58) and (6.61). The delta-method standard deviation $0.002 \times 0.4082 = 0.000816\,\mathrm{h}^{-1}$ of equation (6.40) matches the exact one, while the lognormal it linearizes has the mean $0.002\,\mathrm{e}^{1/12} = 0.002174\,\mathrm{h}^{-1}$ and the larger standard deviation $0.002174\sqrt{\mathrm{e}^{1/6} - 1} = 0.000926\,\mathrm{h}^{-1}$. The Laplace interval is shifted up at both ends, and the approximation errs by 0.04 on the requirement, 0.03 on the mission, 0.09 failures on the expected count and 57 h on the mean time to the next failure, every error in the direction of a worse pump. A normal placed on the rate itself, centered at the mode $0.001667\,\mathrm{h}^{-1}$ with the inverse curvature $5.556 \times 10^{-7}\,\mathrm{h}^{-2}$ of the exact log posterior there, has the standard deviation $0.000745\,\mathrm{h}^{-1}$, puts 1.3 % of its probability on negative rates and raises the support for the requirement to 0.963, which is why the state carries the log-rate. Every answer about the pump is reported from the Gaussian in the log-rate, the route that remains when a rate enters a larger model that has no closed form; the table shows, question by question, what that route costs on a record of four failures, and that the cost never changes a decision of the chapter.

Table 6.8: The pump posterior Gamma(6, 3000 h) against its Laplace approximation in the log-rate ω, equation (6.27). Rates in $10^{-3}\,\mathrm{h}^{-1}$.

Summary	Exact	Laplace in ω
Center (mode / mean)	1.667 / 2.000	2.000 (median)
Standard deviation of λ	0.816	0.816 (delta), 0.926 (lognormal)
Central 95 % interval	[0.734, 3.889]	[0.899, 4.452]
$\Pr(\lambda < 3 \mid \mathbf{y})$	0.884	0.840
One-sided bound $\lambda_{0.95}$	3.504	3.914
$\mathbb{E}[\mathrm{e}^{-500\lambda} \mid \mathbf{y}]$	0.397	0.369
$\Pr(K_f \geq 2 \mid \mathbf{y})$, $\tau = 500\,\mathrm{h}$	0.264	0.291
$\mathbb{E}[K_f \mid \mathbf{y}]$, $\tau = 500\,\mathrm{h}$	1.000	1.087
$\mathbb{E}[T \mid \mathbf{y}]$ [h]	600	543
Median of T [h]	367	339

The fleet. On the pooled one-entry state the Laplace spread of the shared log-rate is $1/\sqrt{8} = 0.3536$, against the exact standard deviation 0.365 of $\ln\lambda$ under Gamma(8, 4000 h), and its band on the rate is $[0.0010,\ 0.0040]\,\mathrm{h}^{-1}$, against the exact $[0.00086,\ 0.00361]\,\mathrm{h}^{-1}$. Pump B's band $[0.00064,\ 0.0050]\,\mathrm{h}^{-1}$ at $\sigma_d = 0.5$, which has no exact counterpart, is to be read against that pooled band of the same approximation: it is wider and centered lower.

The power-law shape. For record B Newton's iteration from $\mathbf{z} = \mathbf{0}$ takes the first entry through 3.000, 2.199, 1.643, 1.417 and 1.387 to $\ln 4 = 1.386$, and v through 3.427, 2.571, 1.909, 1.565 and 1.491 to 1.488, where the Jacobian is $\mathrm{diag}(2, 2)$ and $\boldsymbol{\Sigma} = \mathrm{diag}(0.25,\ 0.25)$. Read off as $\hat{\beta}\,\mathrm{e}^{\pm 1.96/2}$, the Laplace 95 % bands on the shape are [0.46, 3.24] for record A and [1.66, 11.8] for record B. They reach the same two conclusions as the exact bands [0.33, 2.66] and [1.21, 9.70] and sit 20 to 40 % higher, because the exact posterior of v is skewed toward small shapes, which its Gaussian in v does not represent.

6.4.2 Dependence on the size of the record

The pump's answers need only the count, the exposure and the prior, and the effect of a different record is a different pair (a_n, b_n), and with it a different posterior of the state, $\hat{\omega} = \ln(a_n/(b_n\lambda_{\mathrm{ref}}))$ and $\sigma_{\hat{\omega}}^2 = 1/a_n$. Table 6.9 evaluates both routes for five records under the

same prior Gamma(2, 1000 h). Each row is a record; the columns give the exact posterior from equation (6.30), the rate at the MAP state, which is also the exact mean, and then, as the Laplace value with the exact value in parentheses, the central 95 % interval from equations (6.39) and (6.47), the probability of the requirement $\lambda < 0.003\,\mathrm{h}^{-1}$ from equations (6.41) and (6.48), and the probability of no failure in the next 500 h from the sum (6.53) and equation (6.59). The first row is the pump. One row of the arithmetic, the record with no failures: Gamma(2 + 0, 1000 + 2000) = Gamma(2, 3000), mean $2/3000 = 0.67 \times 10^{-3}\,\mathrm{h}^{-1}$, $\hat{\omega} = \ln(2/3) = -0.4055$ with $\sigma_\omega = 1/\sqrt{2}$, the Laplace tail $\Phi((\ln 3 + 0.4055)\sqrt{2}) = \Phi(2.1272) = 0.983$, and $(3000/3500)^2 = 0.735$ exactly.

Table 6.9: The pump's answers for five records under the prior Gamma(2, 1000 h). Rates in $10^{-3}\,\mathrm{h}^{-1}$. Each cell of the last three columns gives the Laplace value from the posterior of the log-rate, then the exact value from the gamma posterior in parentheses; the last column is the probability of no failure at $\tau = 500$ h.

Record k in T_{tot}	Posterior	$\lambda_{\mathrm{ref}} e^{\hat{\omega}}$	95 % interval	$\Pr(\lambda < 3)$	$\Pr(K_f = 0)$
4 in 2000 h (pump)	Gamma(6, 3000)	2.00	[0.90, 4.45] ([0.73, 3.89])	0.840 (0.884)	0.369 (0.397)
0 in 2000 h	Gamma(2, 3000)	0.67	[0.17, 2.67] ([0.08, 1.86])	0.983 (0.999)	0.682 (0.735)
2 in 1000 h	Gamma(4, 2000)	2.00	[0.75, 5.33] ([0.54, 4.38])	0.791 (0.849)	0.370 (0.410)
8 in 4000 h	Gamma(10, 5000)	2.00	[1.08, 3.72] ([0.96, 3.42])	0.900 (0.930)	0.368 (0.386)
40 in 20 000 h	Gamma(42, 21 000)	2.00	[1.48, 2.71] ([1.44, 2.65])	0.996 (0.998)	0.368 (0.372)

Four of the records share the observed rate $0.002\,\mathrm{h}^{-1}$ and differ only in size. As the record grows, the interval narrows roughly as $1/\sqrt{a_n}$, the support for the requirement rises from 0.79 to 0.996 (0.85 to 0.998 exactly) because less posterior mass remains above the limit, and the exact predictive mission reliability falls toward the plug-in value $e^{-1} = 0.368$, since the Jensen gap of equation (6.60) closes as the posterior concentrates; the Laplace value sits at the plug-in from the start, 0.368 to 0.370, because its mean rate lies above the median by as much as the Jensen gap lifts it. The gap between the two routes closes with the record, from 0.058 to 0.002 in the support: the posterior of ω becomes normal as the events accumulate. The record with no failures shows the prior at work: the MLE would be zero, while the median rate is $0.67 \times 10^{-3}\,\mathrm{h}^{-1}$ and the interval $[0.17,\ 2.67] \times 10^{-3}\,\mathrm{h}^{-1}$ ([0.08, 1.86] exactly) states both how small the rate may be and how large it may still be after two thousand failure-free hours. The width of the band is set by the shape $a_n = a + k$, the prior's pseudo-failures plus the observed failures, which is what Section 6.6 uses to choose the next record.

6.4.3 Sensitivity to the prior

The engineering prior (6.18) weighs as much as two failures in a thousand hours, half the weight of the record, and the decision on the requirement should not hang on it. Table 6.10 repeats the pump's three decision numbers under four other priors: each row is a prior, the columns give the exact posterior from equation (6.30), which fixes the posterior of the state $\hat{\omega} = \ln(a_n/(b_n\lambda_{\mathrm{ref}}))$, $\sigma_\omega^2 = 1/a_n$, and then, as the Laplace value with the exact value in parentheses, the support for $\lambda < 0.003\,\mathrm{h}^{-1}$ from equations (6.41) and (6.48), the one-sided bound from equations (6.41) and (6.49) and the probability of no failure in the next 500 h from the sum (6.53) and equation (6.59). The Jeffreys prior, with density proportional to $\lambda^{-1/2}$, and the flat prior are the limits $b \to 0$ of Gamma$(1/2, b)$ and Gamma$(1, b)$, and add no hours.

The support moves from 0.658 to 0.874 (0.715 to 0.910 exactly) and the mission reliability

Table 6.10: The pump's decision numbers under five priors, record $k = 4$ in $T_{\text{tot}} = 2000\,\text{h}$. The first row is the engineering prior. Each cell of the last three columns gives the Laplace value from the posterior of the log-rate, then the exact value from the gamma posterior in parentheses.

Prior	Posterior	$\Pr(\lambda < 0.003)$ [–]	$\lambda_{0.95}$ [$10^{-3}\,\text{h}^{-1}$]	$\Pr(K_f = 0)$ [–]
Gamma(2, 1000 h)	Gamma(6, 3000 h)	0.840 (0.884)	3.91 (3.50)	0.369 (0.397)
Jeffreys, $\propto \lambda^{-1/2}$	Gamma(4.5, 2000 h)	0.729 (0.787)	4.89 (4.23)	0.331 (0.366)
flat	Gamma(5, 2000 h)	0.658 (0.715)	5.22 (4.58)	0.296 (0.328)
Gamma(1, 500 h)	Gamma(5, 2500 h)	0.818 (0.868)	4.17 (3.66)	0.369 (0.402)
Gamma(4, 2000 h)	Gamma(8, 4000 h)	0.874 (0.910)	3.58 (3.29)	0.368 (0.390)

from 0.296 to 0.369 (0.328 to 0.402), but under every prior the one-sided bound lies above $0.003\,\text{h}^{-1}$ and the support below 0.95, on both routes: the requirement is not met under any of them. The Laplace support is below the exact one in every row, by 0.04 to 0.06: the bias of the normal in the log-rate has the same sign under every prior. The comparison of the pumps moves with the prior as well, since the separate drawing spreads three priors where the shared one spreads one: with every rate under Gamma$(1, 500\,\text{h})$, equation (6.69) gives $\ln B_{10} = +0.57$ nats, and under Gamma$(4, 2000\,\text{h})$, $+0.08$; the preference for one shared rate stays positive and weak.

6.4.4 The rows at the minimum and the hand checks

The predictive moments. The moments of the count in the next 500 h, computed through the laws of total expectation and variance, agree with those of the negative-binomial distribution itself, $a_n(1-\pi)/\pi = 1$ and $a_n(1-\pi)/\pi^2 = 6 \times \frac{1}{7} \times \frac{49}{36} = 1.1667$, the last row of Table 6.7.

The evidence of the shared drawing. One-dimensional quadrature of the product of the three Poisson probabilities against the prior Gamma$(2, 1000\,\text{h})$ returns $p(\mathbf{y} \mid M_1) = 0.008859$, the closed form (6.69) to every printed digit.

A record whose slopes do not vanish. At the pump's MAP state both rows' slopes are zero, because each row sits at its own optimum. For the record with no failures in 2000 h of Table 6.9, the MAP state is $\hat{\omega} = \ln(2/3) = -0.4055$, the rate $a_n/b_n = 0.000667\,\text{h}^{-1}$; there the prior row has the slope $1000 \times 0.000667 - 2 = -1.333$ and the count row the slope $2000 \times 0.000667 - 0 = +1.333$: the prior pulls the rate up toward its mean $0.002\,\text{h}^{-1}$, the failure-free hours pull it down, and the two slopes cancel in $\mathcal{L}' = 0$. The curvatures 0.667 and 1.333 add to $\Lambda = 2 = a_n$, as equation (6.27) states for every record.

The fit with the uncertain factor. Unlike the pump record, the two-load fit, whose count row is stacked with the factor's prior row in a Jacobian of two columns, leaves the working residuals (6.22) $(-0.133,\ 0.122,\ -0.018)$ in its three rows at the MAP state. Their half square is not the objective (Section 4.2.7): minimizing $\frac{1}{2} r^{\mathsf{T}} r$ instead of $\mathcal{L}$ would put ω_0 at 0.6029 rather than at 0.5988. Dropping the curvature term $\mathbf{C}$, the Gauss–Newton covariance, moves the spread of the high-load intensity from 0.002014 to $0.002015\,\text{h}^{-1}$, because the count row's working residual at the MAP state, 0.122, is small.

6.4.5 The tolerances taken to zero

The pooled posterior Gamma$(8, 4000\,\text{h})$ is the claim that the tie rows hold exactly, and the check of the four-entry fit is that it returns that posterior in the limit $\sigma_d \to 0$. Table 6.11 lists

the fit as the tie rows tighten: each row is a tie tolerance, the columns are the MAP values of the shared log-rate and of the two deviations that differ most, the standard deviation of the shared log-rate, and that of pump B's own log-rate $\omega + d_B$. By $\sigma_d = 0.01$ the entries have stopped moving at the pooled one-entry Laplace fit, $\hat{\omega} = \ln 2$ and $\sqrt{\Sigma_{\omega\omega}} = 1/\sqrt{8}$: the $\omega\omega$ entry of $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ is $8.000 = 2.000 + 2.000 + 1.600 + 2.400$, the prior row's $b\hat{\lambda}$ plus the three pumps' expected counts, and the rate, read through $\mathbf{g} = (\hat{\lambda}, 0, 0, 0)^{\mathsf{T}}$, has $\sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}} = 0.002 \times 0.3536 = 0.000707\,\mathrm{h}^{-1}$, the exact standard deviation of Gamma$(8, 4000\,\mathrm{h})$. The check passes: the approach to the limit is smooth, with no entry that jumps or stalls, so the exact constraint is the end of the family of tie rows rather than a separate model, and the pooled band is recovered without eliminating the deviations.

Table 6.11: The three-pump record as a state vector, equation (6.12): the MAP state, the standard deviation of the shared log-rate and that of pump B's own log-rate as the tie rows tighten. The last row is the pooled one-entry Laplace fit; the exact posterior standard deviation of ω under Gamma$(8, 4000\,\mathrm{h})$ is 0.365. The deviation $\hat{d}_A$ stays below 0.014 throughout.

Tie tolerance σ_d [–]	$\hat{\omega}$ [–]	$\hat{d}_B$ [–]	$\hat{d}_C$ [–]	$\sqrt{\Sigma_{\omega\omega}}$ [–]	sd of $\omega + d_B$ [–]
1	0.6723	−0.2367	0.1816	0.5054	0.7058
0.5	0.6820	−0.1058	0.0965	0.4105	0.5254
0.1	0.6926	−0.0059	0.0059	0.3563	0.3644
0.01	0.6931	−0.0001	0.0001	0.3536	0.3537
0.001	0.6931	0.0000	0.0000	0.3536	0.3536
$\to 0$ (pooled fit)	$0.6931 = \ln 2$	0	0	$0.3536 = 1/\sqrt{8}$	0.3536

The known load factor is the limit $\sigma_c \to 0$ of the factor's prior row. At $\sigma_c = 0.01$ the two-entry fit returns $\hat{\omega}_0 = 0.5978 = \ln(4/2.2)$ with $\sqrt{\Sigma_{11}} = 0.5000 = 1/\sqrt{4}$, a reference rate of $0.001818 \pm 0.000909\,\mathrm{h}^{-1}$ and a high-load intensity with standard deviation $0.001818\,\mathrm{h}^{-1}$, the standard deviations of the exact Gamma$(4, 2200\,\mathrm{h})$ and Gamma$(4, 1100\,\mathrm{h})$ posteriors. The thirty percent on the factor raises the spread of the high-load intensity from 0.001818 to $0.002014\,\mathrm{h}^{-1}$, about eleven percent, a posterior for which no gamma form exists.

6.4.6 The dispersion and the trend of the records

The same posterior. By equation (6.29) records A and B both give the pump's posterior Gamma$(6, 3000\,\mathrm{h})$. In Figure 6.2(b), however, record A climbs along the line through the totals, while record B stays flat for more than half the record and then climbs about six times as steeply, four failures in 320 h. Their temporal patterns suggest different physical mechanisms: record B could indicate aging, a temporary environmental disturbance, or clustered or common-cause behavior.

The event times as a check. The event times carry no information about λ within the model, and that is exactly what makes them a check of the model. Dividing the density of the times (6.29) by the probability of the count, equation (6.7) with $t = T_{\mathrm{tot}}$,

$$p(t_1, \ldots, t_k \mid K = k, \lambda) = \frac{\lambda^k \, \mathrm{e}^{-\lambda T_{\mathrm{tot}}}}{(\lambda T_{\mathrm{tot}})^k \, \mathrm{e}^{-\lambda T_{\mathrm{tot}}} / k!} = \frac{k!}{T_{\mathrm{tot}}^k}, \tag{6.81}$$

which does not depend on λ and is the density of k independent times uniform on $(0, T_{\mathrm{tot}}]$, arranged in order. Under a constant rate, the normalized times $u_j = t_j / T_{\mathrm{tot}}$ are therefore independent uniforms, and their sum $S_u = \sum_j u_j$ has mean $k/2$ and variance $k/12$ whatever the

rate. A late cluster makes S_u large; a rate that falls with age makes it small. The statistic is that of the Laplace trend test, and here it serves as a check of the model: because the distribution of S_u given k does not depend on λ, it is exactly what the fitted constant-rate model predicts for the record, and the check asks whether the observed value is typical of that prediction. For four failures the predicted S_u has mean 2 and standard deviation 0.577. Record A gives $S_u = (320 + 770 + 1320 + 1830)/2000 = 2.12$, near the center of the prediction, which puts probability 0.42 on values at least as large. Record B gives $S_u = (1440 + 1550 + 1650 + 1760)/2000 = 3.20$, far in the upper tail: by the symmetry of the sum about 2, the predicted probability of a value at least as large is the probability that four uniforms sum to at most 0.80, which is $0.80^4/4! = 0.0171$ exactly (the volume of a simplex). Record B is not what the constant-rate model produces, and its constant-rate posterior should not be reported; its power-law band of Section 6.3.4 excludes the constant rate for the same reason.

> **Principle 6.1.** *Within the homogeneous Poisson model the count and the exposure are sufficient for the rate, but they are not sufficient for checking the model. The event times, which leave the posterior unchanged, are the evidence about whether the constant rate holds, and they are retained for that purpose.*

The dispersion of the counts. A divided record carries a second check, on the equality of the mean and the variance of a count in equation (6.9). At the pooled fit each count row leaves its Pearson residual $r_i = (\mu_i - k_i)/\sqrt{\mu_i}$, its slope (6.21) divided by the square root of its curvature, the same number as the working residual of equation (6.22), and under a shared constant rate the sum of their squares,

$$X^2 = \sum_i r_i^2 = \sum_i \frac{(\mu_i - k_i)^2}{\mu_i},$$

is approximately chi-square with one degree of freedom fewer than the number of count rows, the one taken by the fitted rate; counts that scatter more than Poisson counts make it large. For the pump's four periods, the slopes of Section 6.2.3 leave the residuals $-0.6/\sqrt{1.4} = -0.507$ in period 2 and $0.6/\sqrt{0.6} = 0.775$ in period 3, zero in the others, and give $X^2 = 0.257 + 0.600 = 0.857$ on three degrees of freedom. For the three pumps of the middle block of Table 6.2, the pooled fit at $\hat{\lambda} = 0.002\,\mathrm{h}^{-1}$ expects 2.0, 1.6 and 2.4 failures, leaves the residuals $(0,\ 0.474,\ -0.387)$ and gives $X^2 = 0.375$ on two. The chi-square reference puts the probabilities 0.836 and 0.829 on values at least as large. With expected counts between 0.6 and 2.4 that reference is rough, and the full calculation is exact: dividing the product of the Poisson probabilities of the counts by the Poisson probability of their sum, as in equation (6.81), leaves the multinomial distribution with probabilities h_i/T_{tot}, which does not depend on λ, and enumerating every split of the four period failures and of the six pump failures gives the tail probabilities 0.921 and 0.886. Both records are typical of a shared constant rate: neither the periods nor the pumps scatter more than their Poisson counts allow, and the pooled rate of the fleet passes this check.

6.5 The report (step 6)

The record does not support the first reading of it. Four failures in 2000 h is $0.002\,\mathrm{h}^{-1}$, well below the limit of $0.003\,\mathrm{h}^{-1}$, but the posterior probability that the pump's rate is below the limit is 0.840 (0.884 exactly), short of the 0.95 assurance, and it stays short under every prior tried, on both routes. Nor do the fleet's logs show pump C to be worse than pump B: they favor one shared rate, weakly, by $\ln B_{10} = +0.25$ nats. Each answer below is stated in the units in which its question was asked.

6.5.1 The answers and their bands

The checks support the constant rate for the pump's record: its four periods scatter as Poisson counts of one rate, and so do the three pumps (Section 6.4.6). A record whose event times failed the trend check, as record B does, would not be reported from the constant-rate posterior. Table 6.12 gives the answers: each row is a question; the second column gives the answer from the posterior of the state, $\omega \mid \mathbf{y} \approx \mathcal{N}(0.6931,\ 0.4082^2)$ for the pump and $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ on the extended states; the third the exact value, where one exists; and the last the band of the reported answer: the central 95 % credible interval for an estimate, and for a probability of the pump the range over the five priors of Table 6.10. The differences between the rows arise from what each question asks, not from inconsistencies in the model.

Table 6.12: The report for the operating records. Answers from the posterior of the state, with the exact value from the gamma posteriors, the closed-form evidences or the chi-square band beside each; bands are central 95 % intervals of the reported route, exact in parentheses, or the range over the five priors of Table 6.10.

Question	Answer	Exact	Band
What is the failure intensity?	median $0.002\,\mathrm{h}^{-1}$	median $0.00189\,\mathrm{h}^{-1}$, mean $0.002\,\mathrm{h}^{-1}$, mode $0.00167\,\mathrm{h}^{-1}$	$[0.00090,\ 0.0045]\,\mathrm{h}^{-1}$ ($[0.00073,\ 0.0039]$ exact)
How certain is the estimate?	standard deviation $0.00082\,\mathrm{h}^{-1}$, coefficient of variation 0.41	the same	$[0.00090,\ 0.0045]\,\mathrm{h}^{-1}$
Is the intensity below $0.003\,\mathrm{h}^{-1}$?	probability 0.840; one-sided bound $0.0039\,\mathrm{h}^{-1}$, so a 0.95 assurance is not met	0.884; $0.0035\,\mathrm{h}^{-1}$	0.658 to 0.874 over the five priors; prior alone 0.717 (0.801)
How many failures in the next 500 h?	$\Pr(K_f \le 1) = 0.709$, $\Pr(K_f \ge 2) = 0.291$	0.736, 0.264	$\Pr(K_f = 0) = 0.369$ (0.397)
What is the expected number?	1.09 failures, predictive standard deviation 1.14	1.0, 1.08	–
No failure in the next 500 h?	0.369; plug-in 0.368	0.397	0.296 to 0.369 over the five priors
How long until the next failure?	mean 543 h, median 339 h, standard deviation 634 h	600 h, 367 h, 735 h	–
What is the rate of the three pumps, and pump B's own?	shared median $0.002\,\mathrm{h}^{-1}$; pump B at $\sigma_d = 0.5$, median $0.00178\,\mathrm{h}^{-1}$	shared median $0.00192\,\mathrm{h}^{-1}$, mean $0.002\,\mathrm{h}^{-1}$; pump B none	shared $[0.00100,\ 0.00400]\,\mathrm{h}^{-1}$ ($[0.00086,\ 0.00361]$ exact); pump B $[0.00064,\ 0.0050]\,\mathrm{h}^{-1}$
Are the three pumps alike?	$\ln B_{10} = +0.25$ nats for one shared rate, from the exact evidences; no state route	the same	weak; +0.57 and +0.08 under the other two gamma priors
What is the rate under normal and under high load?	known factor, medians $0.00182\,\mathrm{h}^{-1}$ and $0.00364\,\mathrm{h}^{-1}$; factor within 30 %, 0.00182 and $0.00362\,\mathrm{h}^{-1}$, the latter with standard deviation $0.00201\,\mathrm{h}^{-1}$	known factor, medians 0.00167 and $0.00334\,\mathrm{h}^{-1}$, means 0.00182 and $0.00364\,\mathrm{h}^{-1}$	known factor $[0.00068,\ 0.0048]$ and $[0.0014,\ 0.0097]\,\mathrm{h}^{-1}$ ($[0.00050,\ 0.00399]$ and $[0.00099,\ 0.00797]$ exact); factor within 30 %, $[0.0012,\ 0.0108]\,\mathrm{h}^{-1}$ for the high load
What is the power-law shape of each record?	record A: $\hat{\beta} = 1.22$; record B: $\hat{\beta} = 4.43$	the same estimates	A $[0.46,\ 3.24]$, contains 1; B $[1.66,\ 11.8]$, excludes 1 ($[0.33,\ 2.66]$ and $[1.21,\ 9.70]$ exact)

The report line. Under the homogeneous Poisson model and the prior from similar pumps, the pump's failure rate has the median $0.002\,\mathrm{h}^{-1}$, with a 95 % band from 0.00090 to $0.0045\,\mathrm{h}^{-1}$

(0.00073 to 0.0039 h^{-1} from the exact posterior); the posterior probability that it is below 0.003 h^{-1} is 0.84 (0.88), so the limit is not yet demonstrated at an assurance of 0.95. Over the next 500 operating hours the pump should expect about one failure, 1.09 (1.0), with a probability of 0.37 (0.40) of none and of 0.29 (0.26) of two or more. The three comparable pumps are consistent with one shared rate of 0.002 h^{-1}, whose band from 0.00100 to 0.00400 h^{-1} (0.00086 to 0.00361 h^{-1}) is narrower than the pump's own; the exact evidences favor the shared rate only weakly, by 0.25 nats.

6.5.2 The requirement decision

The decision follows from the third row of Table 6.12. If the facility demands the rate limit with at least 95 % posterior probability, the record does not meet it: the support is 0.840, and the one-sided bound 0.0039 h^{-1} lies above the limit; exactly, the support is 0.884 and the bound 0.0035 h^{-1}, a narrower miss. The two routes reach the same decision. The point estimate 0.002 h^{-1} is well below the limit, and still a posterior probability of 0.160 (0.116) remains that the true rate is above it. The decision does not depend on the prior: under each prior of Table 6.10 the requirement is not met, on either route.

Common mistake

Reading $\Pr(\lambda < 0.003 \mid \mathbf{y}) = 0.840$ (0.884 exactly) as an 84 % chance that the pump runs the next mission without failure. The first is a statement about the unknown rate; the probability of a failure-free mission is a predictive probability, here 0.369 (0.397) for 500 h, and it depends on the mission length.

6.6 The next measurement (step 7)

The pump's band runs from 0.45 to 2.23 times its median (0.37 to 1.94 times the mean on the exact posterior), and the 95 % assurance on the requirement is not met: the band is too wide for the decision, and what narrows a rate is more operating exposure.

6.6.1 The candidate records

For the pump alone the candidate is the next 500 h of service, with a record of ten times the present exposure carried beside it as a reference. For the fleet the candidates are 1500 pump-hours placed in two ways: the next 500 h on each pump, or all 1500 h on pump B, the pump with the least exposure and the widest band of its own.

6.6.2 The preposterior of the log-rate

A candidate of H further operating hours on a component whose fitted rate is $\hat{\lambda}$ is a count row not yet recorded, with expected count $H\hat{\lambda}$. By equation (6.21) its curvature in the log-rate is that expected count, in which the count itself does not appear, so its information is known before the hours are run. On a state of one entry the preposterior variance of Chapter 2, equation (2.17), is then

$$\Lambda_a = \Lambda + H\hat{\lambda}, \qquad \sigma^2_{\omega,a} = \frac{1}{\Lambda_a},$$

and on a state of several entries, with the candidate's working line (6.22) $\mathbf{J}_a = \sqrt{H\hat{\lambda}}\,\partial \ln \mu_a / \partial \mathbf{z}^{\mathsf{T}}$, the preposterior covariance is

$$\boldsymbol{\Sigma}_a = \left(\boldsymbol{\Sigma}^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a\right)^{-1}.$$

The width the candidate removes is the difference between the band before and the band after.

The pump alone. On the one-entry state, $\Lambda = 6$, and the next 500 h of service add the curvature $500 \times 0.002 = 1$, so that $\Lambda_a = 7$, $\sigma^2_{\omega,a} = 1/7$ and the standard deviation of ω falls from 0.408 to 0.378. The Laplace band on the rate in the log-rate narrows from 3.55 to $3.24 \times 10^{-3}\,\mathrm{h}^{-1}$. Taken outcome by outcome, each count $K_f = j$ of the 500 h gives the posterior of the state $\hat{\omega}_j = \ln((6 + j)/3.5)$, $\sigma^2_{\omega,j} = 1/(6 + j)$, and the width of its band, averaged over the Laplace predictive of Table 6.7, is $3.25 \times 10^{-3}\,\mathrm{h}^{-1}$, the preposterior value to 0.01. The exact check, the width of the gamma interval averaged over the negative-binomial prediction of the count, narrows from 3.16 to $2.92 \times 10^{-3}\,\mathrm{h}^{-1}$: the exact bands are narrower by 0.3 to 0.4, and the 500 h remove 0.30 from the Laplace band and 0.24 from the exact one. The band shrinks as $1/\sqrt{a_n}$: a record that raises the shape from a_n to a'_n narrows it by about $\sqrt{a'_n/a_n}$. Ten times the pump's exposure, 40 failures in 20 000 h in Table 6.9, takes the shape from 6 to 42: the standard deviation of ω falls by exactly $\sqrt{42/6} = 2.65$, and the Laplace band narrows from 3.55 to $1.23 \times 10^{-3}\,\mathrm{h}^{-1}$, a factor of 2.9, because its skew shrinks with it; exactly, the band narrows from 3.16 to $1.21 \times 10^{-3}\,\mathrm{h}^{-1}$, a factor of 2.6, close to $\sqrt{42/6}$ rather than $\sqrt{10} = 3.16$, because the prior's two pseudo-failures do not scale with the record; the factor tends to $\sqrt{10}$ only when the observed failures dominate the prior's, $k \gg a$. One pump's own hours narrow its band slowly.

The fleet. On the pooled one-entry state, $\Lambda = 8$, every count row's log-mean has the slope one in ω, and both candidates add the same curvature, $1500 \times 0.002 = 3$: $\Lambda_a = 11$, $\sigma^2_{\omega,a} = 1/11$, the standard deviation of ω falls from 0.354 to 0.302, and the Laplace band on the shared rate narrows from $[1.00,\ 4.00]$ to $[1.11,\ 3.61] \times 10^{-3}\,\mathrm{h}^{-1}$, a width of $0.50 \times 10^{-3}\,\mathrm{h}^{-1}$ removed; averaged over the outcomes, each weighted by the Laplace predictive of the pooled count, the band after the 1500 pump-hours is $2.51 \times 10^{-3}\,\mathrm{h}^{-1}$ wide, the preposterior 2.50 to 0.01. The exact check agrees on the spread: averaged over the negative-binomial count, the expected posterior variance of Gamma$(8 + K, 5500\,\mathrm{h})$ gives the standard deviation $0.603 \times 10^{-3}\,\mathrm{h}^{-1}$, equal to the delta-method $\hat{\lambda}\,\sigma_{\omega,a} = 0.002 \times 0.302$, and the exact interval narrows in expectation from 2.742 to $2.335 \times 10^{-3}\,\mathrm{h}^{-1}$, a width of $0.407 \times 10^{-3}\,\mathrm{h}^{-1}$ removed, less than the Laplace 0.50 because the exact band is narrower to begin with. On the four-entry state with tie rows at $\sigma_d = 0.5$, the two candidates are different rows. The spread adds three rows $\sqrt{500\hat{\lambda}_i}\,(1, \mathbf{e}_i)$, with $\mathbf{e}_i$ the indicator of pump i's deviation and the fitted rates 1.985, 1.779 and $2.178 \times 10^{-3}\,\mathrm{h}^{-1}$, entries 0.996, 0.943 and 1.044; the concentrated record adds the one row $\sqrt{1500 \times 0.001779}\,(1, 0, 1, 0) = 1.634\,(1, 0, 1, 0)$. Table 6.13 lists the result: each row is a candidate under one model, the columns the standard deviation of the shared log-rate and the width of the 95 % Laplace band on the shared rate, then the same two for pump B's own log-rate $\omega + d_B$, read through $\mathbf{g} = (1, 0, 1, 0)^{\mathsf{T}}$.

Table 6.13: The next record for the three pumps: 1500 pump-hours spread or concentrated on pump B, by the preposterior variance $\sigma^2_{\omega,a} = 1/(\Lambda + H\hat{\lambda})$ of the one-entry pooled state ($\Lambda = 8$, top) and the preposterior covariance $\boldsymbol{\Sigma}_a = (\boldsymbol{\Sigma}^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a)^{-1}$ of the four-entry state with tie rows at $\sigma_d = 0.5$ (tied, bottom). Band widths are those of the 95 % Laplace band in the log-rate, $\hat{\lambda}\,(\mathrm{e}^{1.96s} - \mathrm{e}^{-1.96s})$ with s the standard deviation of the log-rate.

Candidate record	State	sd of ω [–]	shared band $[10^{-3}\,\mathrm{h}^{-1}]$	sd of $\omega + d_B$ [–]	pump B band $[10^{-3}\,\mathrm{h}^{-1}]$
current record	pooled	0.354	3.00	0.354	3.00
500 h on each pump	pooled	0.302	2.50	0.302	2.50
1500 h on pump B	pooled	0.302	2.50	0.302	2.50
current record	tied	0.410	3.54	0.525	4.35
500 h on each pump	tied	0.376	3.19	0.464	3.70
1500 h on pump B	tied	0.380	3.23	0.399	3.07

6.6.3 The recommendation

Under pooling the placement of the hours does not matter, because the pooled state cannot tell the pumps apart. Under the tie rows it does. The spread removes $0.35 \times 10^{-3}\,\mathrm{h}^{-1}$ from the shared band and $0.65 \times 10^{-3}\,\mathrm{h}^{-1}$ from pump B's; the concentrated record removes 0.31 from the shared band and $1.28 \times 10^{-3}\,\mathrm{h}^{-1}$ from pump B's, twice the spread's. These widths are all on the state route; the exact posterior exists only for the pooled rows, where it gives the narrower bands 2.742 and 2.335 in place of 3.00 and 2.50 and the same indifference to placement, so the gap does not change the recommendation. The recommendation follows the question. For the fleet rate, whose triple 3λ sizes spares and maintenance for the three pumps, the next 500 h on each pump is the better record: it removes the most from the shared band, and its three count rows are also the evidence the dispersion check and the question of σ_d need. For a decision about pump B alone, such as whether its own rate meets $0.003\,\mathrm{h}^{-1}$, the 1500 h belong on pump B. For the single pump of the chapter, whose requirement is not yet met, the next record is simply more operating hours with every failure and its time logged. Once the hours are run, their count rows join the fit and the report is redone.

6.6.4 What the analysis showed

The first readings of the logs were that the pump complies, $0.002\,\mathrm{h}^{-1}$ against a limit of $0.003\,\mathrm{h}^{-1}$, and that pump C, with three failures to pump B's one, is the worse pump; the analysis overturns both. The limit has the posterior support 0.840 (0.884 exactly), short of the 0.95 assurance under every prior tried and on both routes, and the three logs favor one shared rate, weakly, by $\ln B_{10} = +0.25$ nats, so the difference in their counts is not evidence that the pumps differ. For this pump the record is small. Its failures put the rate where the experience with similar pumps had placed it, but they fix it only within a wide band, and the shortfall on the limit is a matter of too little exposure rather than of a poor pump. Each prediction differs from its plug-in version in a direction set by the shape of the posterior, and the normal in the log-rate, which every answer is computed from, leans toward a worse pump than the exact gamma posterior: 1.09 expected failures in the next 500 h (1.0), a mean time to the next failure of 543 h (600 h), on a record of four failures where no decision turns on the difference. Divided into periods, the record scatters as a constant rate predicts, and the three pumps of the fleet scatter as one shared rate predicts, so pooling them is defensible and gives the fleet a narrower band than any one pump has; but a record with the same totals whose failures cluster late would be a different pump, one whose intensity rises with age, and its count alone would hide it. Keeping the failure times, not only the counts, is what lets that be seen. Where the next hours go depends on the question: across the fleet for the fleet's rate, on one pump for that pump's own rate.

6.6.5 Limits of the model

The derivation is exact under the assumptions of Section 6.1.4. That does not make the assumptions appropriate, and six of them limit the answers.

Constant rate and repair to as good as new. The failure intensity is constant over operating time and every repair returns the pump to the condition it had before. A rate that changes with age or with maintenance is outside the model; the power law of Section 6.3.4 is one replacement, and Chapter 12 treats a rate that changes with time.

Correctly identified failures. The count is the number of failures of the mechanism the rate describes. Incomplete failure reporting lowers the count and the rate with it, and failures with different underlying mechanisms, counted together, average over mechanisms that a single rate cannot separate.

Operating exposure. The exposure counts time in service only. Repair downtime counted as exposure flatters the pump, since time in which it could not fail is counted as time in which it did not, and standby time belongs to the exposure only if the failure model covers it.

Comparable pumps. The pooled rate assumes that the three pumps share one rate, which the dispersion check and the evidence support but cannot establish with six failures. Different pumps may differ in age, manufacturer, duty cycle, maintenance history or environment, and if those factors matter, pooling conceals the heterogeneity and reports a band narrower than the evidence supports. The tie rows let the rates differ, but their tolerance σ_d is chosen, not learned; a hierarchical model, which gives each pump its own rate drawn from a population distribution $\lambda_i \sim p(\lambda_i \mid \boldsymbol{\psi})$ and learns how much the pumps differ from the records, is the remedy, and the subject of later chapters.

Known load factors. The effective exposure rests on load factors from a physical, empirical or engineering model; they must not be chosen to force records into agreement. An uncertain factor is one more entry of the state, as in the two-load fit.

A shape from four failures. The power-law fit of records A and B rests on four failure times each, and its band on the shape is wide; it shows a rising intensity for record B but does not measure it well. The signs that call the homogeneous model into question are a failure frequency that increases as the equipment ages; failures that cluster during environmental disturbances; maintenance that changes subsequent reliability; an operating load that varies and affects the intensity; failures with different mechanisms; incomplete failure reporting; and repairs that do not restore comparable operating conditions. These situations may require a nonhomogeneous Poisson process, a renewal model, a Weibull lifetime model (Chapter 7), or a more elaborate physical and probabilistic model.

The next chapter moves from counts in an exposure to the failure times themselves, the lifetime distributions whose hazard depends on age, and the censored records they are estimated from.

6.7 What to retain

Definition 6.1 (Gamma–Poisson estimation of a failure rate)**.** Under a homogeneous Poisson process with constant intensity λ, a record of k failures in T_{tot} operating hours has the Poisson likelihood $(\lambda T_{\text{tot}})^k \mathrm{e}^{-\lambda T_{\text{tot}}}/k!$, and a prior Gamma$(a, b)$ gives the exact posterior Gamma$(a + k,\ b + T_{\text{tot}})$; counts from several periods, pumps or conditions combine through their total count and their total, or effective, exposure. On the state, the log-rate ω, the two rows give the MAP state $\hat{\omega} = \ln(a_n/(b_n \lambda_{\text{ref}}))$ and the Laplace variance $\sigma_\omega^2 = 1/a_n$. The estimate, its credible interval, the probability that a rate requirement holds, the prediction of future counts, the mission reliability and the time to the next failure are each a different function of the posterior, computed from $(\hat{\omega}, \sigma_\omega)$ by a mapped quantile, a normal tail or a Gauss–Hermite sum and, beside it, in closed form in the two posterior parameters: the gamma quantiles, the negative binomial, $(b_n/(b_n + \tau))^{a_n}$ and the Lomax distribution. They must not be substituted for one another or for their plug-in versions.

The answers are computed from the fit on the log-rate, the form that reaches the larger models, and the exact gamma posterior is set beside them with the gap, which on four failures leans the state route toward a worse pump without changing a decision. A count is a row on the state with the slope $\mu - k$ and the curvature μ in the log of its mean: on the pump's one entry the rows' slopes and curvatures add as numbers, and on a larger state the count row joins the Gaussian rows in $\mathbf{J}$ in the working form of Chapter 4 while the fit still minimizes

$\sum_i \mathcal{L}_i$; a shared rate across pumps is the limit of tie rows, read off through a shared column, an uncertain multiplier is one more entry, and whether pumps share a rate is a ratio of two closed-form evidences.

The distinction from the previous chapter is the one to carry forward: a failure probability is tied to a defined demand or time horizon, whereas a failure intensity is a frequency of failure events per unit of operating exposure, and the two meet only through a lifetime model and a stated mission length.

Table 6.14 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 6.14: The terms of failure-rate estimation from operating records and where each is defined. Rates in h^{-1}, exposures and times in hours.

Term	Expression	Units	Meaning	Defined in
Failure intensity	$\nu(t) = \mathrm{d}m(t)/\mathrm{d}t$	h^{-1}	rate of expected failures	6.1.1
Hazard rate	$\lambda(t) = f(t)/R(t)$	h^{-1}	conditional failure tendency	6.1.2
Survival function	$R(t) = \mathrm{e}^{-\Lambda(t)}$	–	no failure by t	6.1.2
Poisson count	$(\lambda t)^k \mathrm{e}^{-\lambda t}/k!$	–	k failures in t	6.1.3
General forms				
Estimation	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$, moments, quantiles	–	posterior of a quantity	6.3.1
Threshold and assurance	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y}) \ge \kappa$	–	tail of a quantity	6.3.1
Prediction	$\int p(y_{\mathrm{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	–	future observation	6.3.2
Hypothesis comparison	$B_{10} = p(\mathbf{y} \mid M_1)/p(\mathbf{y} \mid M_0)$, in nats $\ln B_{10}$	–	two complete drawings	6.3.3
Specialized forms and answers				
Laplace band	$\lambda_{\mathrm{ref}} \mathrm{e}^{\hat{\omega} \pm 1.96\sigma_\omega}$	h^{-1}	reported interval	6.3.1
Gauss–Hermite predictive	$\sum_q (w_q/\sqrt{\pi})\, h(\hat{\omega} + \sqrt{2}\sigma_\omega t_q)$	–	prediction from the state	6.3.2
Posterior mode	$(a_n - 1)/b_n$	h^{-1}	most probable rate	6.3.1
Credible interval	$\chi^2_{q;\,2a_n}/(2b_n)$	h^{-1}	quantiles of the posterior	6.3.1
Rate requirement	$F_\Gamma(\lambda_{\max}; a_n, b_n)$	–	support for $\lambda < \lambda_{\max}$	6.3.1
Negative-binomial predictive	$\binom{a_n+j-1}{j} \pi^{a_n} (1-\pi)^j$	–	failures in τ hours	6.3.2
Predictive variance	$\tau \mathbb{E}[\lambda] + \tau^2 \operatorname{Var}(\lambda)$	–	outcome plus rate spread	6.3.2
Mission reliability	$(b_n/(b_n + \tau))^{a_n}$	–	no failure in τ hours	6.3.2
Lomax mean	$b_n/(a_n - 1)$	h	time to the next failure	6.3.2
Predictive hazard	$a_n/(b_n + t)$	h^{-1}	rate given survival to t	6.3.2
Pump B's log-rate	$\mathbf{g}^{\mathsf{T}}\mathbf{z}$, $\mathbf{g} = (1, 0, 1, 0)^{\mathsf{T}}$	–	shared column	6.3.3
Evidence of a count record	$\Gamma(a+k)\, b^a T_{\mathrm{tot}}^k / (\Gamma(a)\, k!\, (b+T_{\mathrm{tot}})^{a+k})$	–	negative binomial at k	6.3.3
Effective exposure	$T_{\mathrm{eff}} = \sum_i c_i h_i$	h	reference-equivalent hours	6.3.3
Cumulative intensity	$m(t) = \int_0^t \nu(u)\, \mathrm{d}u$	–	nonhomogeneous process	6.3.4
Power-law shape	$\hat{\beta} = k/\sum_j \ln(T_{\mathrm{tot}}/t_j)$, band by χ^2_{2k}	–	rising intensity if $\beta > 1$	6.3.4
Building blocks				
Maximum-likelihood estimate	$\hat{\lambda}_{\mathrm{MLE}} = k/T_{\mathrm{tot}}$	h^{-1}	observed rate	6.1.7
Count row on the state	slope $\mu - k$, curvature μ; on several entries $(\mu - k)/\sqrt{\mu}$, $\sqrt{\mu}\, \partial \ln \mu / \partial \mathbf{z}$	–	expected minus observed; working form	6.1.7
Information in the log-rate	$\Lambda = \mathcal{L}''(\hat{\omega}) = b_n \hat{\lambda} = a_n$, $\sigma_\omega^2 = 1/a_n$	–	prior plus Fisher information	6.2.1
Gamma prior	Gamma(a, b), mean a/b	–	a failures in b hours	6.1.7
Gamma–Poisson update	Gamma$(a + k,\, b + T_{\mathrm{tot}})$	–	failures to a, hours to b	6.2.2
Laplace trend test	$S_u = \sum_j t_j/T_{\mathrm{tot}}$, mean $k/2$	–	uniform times under a constant rate	6.4.6
Dispersion check	$X^2 = \sum_i (\mu_i - k_i)^2/\mu_i$	–	counts against one rate	6.4.6

6.8 Exercises

Exercise 6.1. *Formulation.* The posterior of a failure rate is $\lambda \mid \mathbf{y} \sim \text{Gamma}(4,\ 2000\,\text{h})$. The equipment owner wants (a) the posterior probability that $\lambda < 0.003\,\text{h}^{-1}$ and (b) the predictive probability of surviving the next 300 operating hours without failure. Write both expressions, evaluate them, and state why they answer different questions. ((a) $\Pr(\lambda < 0.003 \mid \mathbf{y}) = F_\Gamma(0.003; 4, 2000)$; with $b\lambda_{\max} = 6$ and integer shape, equation (6.46) gives $1-\text{e}^{-6}(1+6+6^2/2+6^3/6) = 1-61\,\text{e}^{-6} = 0.8488$. (b) $\Pr(T > 300 \mid \mathbf{y}) = (2000/2300)^4 = 0.5718$. The first concerns uncertainty about an underlying intensity; the second predicts a future failure-free interval.)

Exercise 6.2. *Formulation.* Two pumps each have four failures in 2000 operating hours. Pump A's failures are roughly evenly dispersed; pump B has all four in the final 400 hours. Explain why the Gamma–Poisson model gives both the same constant-rate posterior under the same prior, and formulate a model in which the failure intensity can increase with operating age. (Under the constant rate the likelihood of either record is $\lambda^4\text{e}^{-2000\lambda}$, equation (6.29), so both give $\text{Gamma}(a+4,\ b+2000)$. A nonhomogeneous Poisson process with the power-law intensity $\nu(t) = \theta\beta t^{\beta-1}$, $\theta, \beta > 0$, has cumulative intensity $m(t) = \theta t^\beta$ and likelihood $\prod_j \nu(t_j)\,\text{e}^{-\theta T_{\text{tot}}^\beta}$; the intensity increases with age when $\beta > 1$, and the event times now inform β through $\sum_j \ln t_j$, which distinguishes the dispersed record from the clustered one.)

Exercise 6.3. *Derivation.* (a) A repairable pump has three failures in 1500 operating hours under a homogeneous Poisson process. Derive the MLE of the rate and, treating it as known, the expected number of failures in the next 200 h, the probability of none and the mean time to the next failure. (b) With the prior $\lambda \sim \text{Gamma}(2,\ 800\,\text{h})$ and a record of three failures in 1200 h, derive the likelihood, the posterior parameters, the posterior mean, variance, standard deviation and mode. (c) For the record of (b), with the log-rate ω as the state, write the slopes and the curvatures of the prior row and the count row at the MAP state, the information Λ, the variance σ_ω^2, and the delta-method standard deviation of the rate. ((a) $\ell'(\lambda) = 3/\lambda - 1500 = 0$ gives $\hat\lambda = 0.002\,\text{h}^{-1}$; $\mathbb{E}[N(200)] = 0.4$; $\text{e}^{-0.4} = 0.6703$; $1/0.002 = 500\,\text{h}$; all plug-in values. (b) $p(k \mid \lambda, T_{\text{tot}}) \propto \lambda^3\text{e}^{-1200\lambda}$; $\text{Gamma}(2+3,\ 800+1200) = \text{Gamma}(5, 2000)$; mean $5/2000 = 0.0025\,\text{h}^{-1}$; variance $5/2000^2 = 1.25\times10^{-6}\,\text{h}^{-2}$, standard deviation $0.001118\,\text{h}^{-1}$; mode $4/2000 = 0.002\,\text{h}^{-1}$. (c) The MAP state is $\hat\omega = \ln(0.0025/0.001) = 0.9163$, where both slopes vanish, $800\times0.0025 - 2 = 0$ and $1200\times0.0025 - 3 = 0$, because the prior mean $2/800$ equals the MLE $3/1200$ (for the record with no failures of Section 6.4.4 they do not); the curvatures are $800\times0.0025 = 2$ and $1200\times0.0025 = 3$, so that $\Lambda = 2+3 = 5 = a_n$, $\sigma_\omega^2 = 0.2$, $\sigma_\omega = 0.4472$, and with $g = \hat\lambda = 0.0025$ the rate has $g\,\sigma_\omega = 0.001118\,\text{h}^{-1}$, the exact standard deviation of (b).)

Exercise 6.4. *Derivation.* The posterior is $\lambda \mid \mathbf{y} \sim \text{Gamma}(3,\ 1000\,\text{h})$, and the number of failures in the next 200 operating hours is to be predicted. Derive the probability of zero failures, of exactly one and of two or more, and the predictive mean and variance. ($\pi = 1000/1200 = 5/6$; equation (6.54) gives $\Pr(K_f = 0) = \pi^3 = 0.5787$, $\Pr(K_f = 1) = 3\pi^3(1-\pi) = 0.2894$ and $\Pr(K_f \geq 2) = 1 - 0.5787 - 0.2894 = 0.1319$; by equations (6.56) and (6.57), $\mathbb{E}[K_f] = 200\times3/1000 = 0.6$ and $\text{Var}(K_f) = 0.6 + 200^2\times3/1000^2 = 0.72$.)

Exercise 6.5. *Application.* The posterior is $\lambda \mid \mathbf{y} \sim \text{Gamma}(5,\ 2000\,\text{h})$. Compute the predictive probability of no failure during the next 100 and the next 500 operating hours, and the expected time to the next failure, and explain why the predictive survival changes with the length of the mission. ($(2000/2100)^5 = (20/21)^5 = 0.7835$; $(2000/2500)^5 = 0.8^5 = 0.3277$; $\mathbb{E}[T \mid \mathbf{y}] = 2000/(5-1) = 500\,\text{h}$, equation (6.63). A longer mission is a longer exposure in which a failure can occur, for every value of λ the posterior admits.)

Exercise 6.6. *Application.* (a) A fleet of comparable motors has the records A: one failure in 600 h, B: two in 800 h, C: one in 1000 h. With the prior Gamma(1, 500 h), derive the posterior under a shared homogeneous intensity and its mean, and give two reasons why the shared intensity might not hold. (b) A pump runs 1000 h under normal conditions and 400 h under high stress with three times the reference intensity, and fails twice. With the prior Gamma(2, 1000 h) for the reference intensity, compute the effective exposure, the posterior, the posterior mean reference intensity and the mean intensity under high stress. (c) For the fleet of (a), with the shared log-rate ω as the state, write the slopes and the curvatures of the prior row and the three count rows at the MAP state, the information Λ, the variance σ_ω^2 and the delta-method standard deviation of the rate. ((a) $k = 1 + 2 + 1 = 4$, $T_{\text{tot}} = 600 + 800 + 1000 = 2400\,\text{h}$; Gamma$(1 + 4,\ 500 + 2400)$ = Gamma$(5, 2900)$, mean $5/2900 = 0.001724\,\text{h}^{-1}$; the motors may differ in age, duty cycle, maintenance quality, manufacturer or operating environment. (b) $T_{\text{eff}} = 1000{\times}1{+}400{\times}3 = 2200\,\text{h}$; Gamma$(2{+}2,\ 1000{+}2200)$ = Gamma$(4, 3200)$; mean $4/3200 = 0.00125\,\text{h}^{-1}$; under high stress $3 \times 0.00125 = 0.00375\,\text{h}^{-1}$. (c) With $\hat\lambda = 5/2900 = 0.001724\,\text{h}^{-1}$ and $\hat\omega = 0.5447$, the rows in the order prior, A, B, C have the slopes $h_i\hat\lambda - k_i = -0.1379$, 0.0345, -0.6207 and 0.7241, which do not vanish but cancel in $\mathcal{L}' = 0$, and the curvatures $h_i\hat\lambda = 0.8621$, 1.0345, 1.3793 and 1.7241, which add to $\Lambda = 2900\,\hat\lambda = 5 = a_n$; $\sigma_\omega^2 = 0.2$, $\sigma_\omega = 0.4472$, and with $g = \hat\lambda$ the rate has $g\,\sigma_\omega = 0.001724 \times 0.4472 = 0.000771\,\text{h}^{-1}$, the exact $\sqrt{5}/2900$.)

Chapter 7

Failure-Time Distributions and Survival Analysis

7.1 Lifetimes from a record with survivors

Eight identical electric motors were put on test under comparable conditions, and the study closed on one calendar date. Five failed, at 120, 230, 310, 450 and 600 h; three were still running when their windows closed at 700, 800 and 900 h. The reliability engineer of the fleet must say how long a motor lasts and whether the motors wear out, since a risk of failure that rises with age would justify replacing them by age. The quickest reading averages the five failure times, 1710 h over five failures, and gives a life of 342 h, less than the 700 to 900 h that three of the motors are known to have exceeded. A better reading fits a Weibull lifetime to all eight motors, survivors included, and takes its characteristic life of 765.6 h and its shape of 1.322, above one, as known: a motor lasts about 765.6 h, and it wears out.

Neither reading says how much eight motors can tell. The first throws the survivors away; the second keeps them but treats a fit to five failures as a measurement, when the same record is consistent with scales and shapes that a single fitted point does not show. The question of this chapter is

how the reliability of a component is estimated from the times at which units fail, including units that have not yet failed, and how uncertain the resulting mission reliability and remaining life are.

What changes from the counts in exposure of Chapter 6 is the evidence, one time per motor that ended either in a failure or with the motor still running, and the state, which now has two entries, the logarithms of the Weibull scale and shape, on which every row of the record bears together.

Jargon: survival function, failure density, hazard function, cumulative hazard, memoryless property, characteristic life, right censoring, noninformative censoring, profile likelihood, observed information, mission reliability, conditional survival, Kaplan–Meier estimator, risk set, nested model, likelihood-ratio statistic, Akaike information criterion

The *survival function* $R(t) = \Pr(T > t)$ is the probability that a unit is still working at age t; reliability engineering calls it the reliability function. The *failure density* $f(t)$ is the probability density of the failure time. The *hazard function* $\lambda(t)$ is the instantaneous failure rate of a unit that has survived to age t, and the *cumulative hazard* $\Lambda(t)$ is its integral from zero to t. A lifetime is *memoryless* when the probability of surviving a further interval does not depend on the age already reached. The *characteristic life* of a Weibull lifetime is its scale, the age by which 63.2 % of units have failed. An observation is *right-censored* when the unit is known to have survived to a time but its failure time is not observed; censoring is *noninformative* when the mechanism that stops the observation carries no information about the lifetime parameters beyond what the model already accounts for. The *profile likelihood* of one parameter is the likelihood maximized over the others at each of its values. The *observed information* is the Hessian of the negative log-likelihood at its minimum. *Mission reliability* is the probability of surviving a stated operating period; *conditional survival* is the same probability for a unit that has already reached a given age, and it describes its remaining life. The *Kaplan–Meier estimator* is a nonparametric estimate of the survival function from censored data; its *risk set* at a time is the set of units still under observation just before it. A model is *nested* in another when it is a special case of it; the *likelihood-ratio statistic* is twice the difference of their maximized log-likelihoods, and the *Akaike information criterion* (AIC) is minus twice the maximized log-likelihood plus twice the number of parameters.

Key idea

A lifetime record is a list of pairs: an observed time and whether it ended in a failure. A failure contributes the failure density at its time to the likelihood, a survivor contributes the survival probability at its time, and every survival quantity follows from one hazard model through $R(t) = \mathrm{e}^{-\Lambda(t)}$. The fitted parameters, carried as $\ln\eta$ and $\ln\beta$ with their Laplace covariance, give the mission reliability and the remaining life with bands through $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$; with few failures those bands are set beside the exact posterior, which for the Weibull model is a one-dimensional integral over the shape. The bands, not the point estimates, decide whether the record supports a claim of wear-out.

7.1.1 Lifetime distributions and failure-count processes

A failure-count process describes how many failures occur in a period of operation. Its unknown is a rate, a number of events per hour, and its evidence is a count k in an exposure T_{tot}. It suits a repairable unit that is restored after each failure and returns to service, so that the same unit produces many events, and it says nothing about the age of the unit at each event when the rate is constant. A lifetime distribution describes something else: the time from the start of service of a unit to its first failure. Its unknown is a whole distribution over that time, and its evidence is one observed time per unit. The distinction matters as soon as the risk of failure depends on age. A motor whose bearings wear has a small chance of failing in its first hundred hours and a larger one after a thousand; a rate fitted to the count of failures in the fleet averages over those ages and cannot say which motors are at risk.

The random variable of this chapter is therefore

$$T = \text{time from the beginning of service until failure,}$$

in hours. It is nonnegative and is taken to have a continuous distribution. Each motor is treated as a nonrepairable item within the study: once it fails, its first lifetime ends, and what happens after a repair is not part of the record. This differs from the pump of the previous chapter, whose repeated failures were the events of a counting process.

Survivors as evidence. A lifetime study rarely runs until every unit has failed. When it closes, some units are still running, and their failure times are unknown. Discarding those units would keep only the lifetimes that were short enough to end inside the study, and so would systematically overrepresent short lifetimes and understate the reliability. Treating the closing time as a failure time would make the same error in a milder form. The surviving units carry definite information: each survived at least as long as its observed exposure. An observation of that kind is called *right-censored*, and the whole of this chapter rests on giving it the correct likelihood. The answer, derived in Section 7.1.6, is that a survivor contributes the probability of surviving to its observed time, a quantity every lifetime model supplies. Before that likelihood can be written, the lifetime distribution must be described by its functions, and the relations among those functions derived.

7.1.2 The functions of a lifetime distribution

Cumulative failure probability. The cumulative distribution function of the lifetime is

$$F(t) = \Pr(T \le t), \tag{7.1}$$

the probability that a unit has failed by age t. A model with $F(500\,\mathrm{h}) = 0.30$ assigns a probability of 30 % to failure within the first 500 h. The function satisfies $0 \le F(t) \le 1$, is nondecreasing in t, and for a nonnegative lifetime starts at $F(0) = 0$.

Survival function. A unit survives beyond age t when $T > t$. The events $\{T \le t\}$ and $\{T > t\}$ are complements for any lifetime distribution, so their probabilities add to one, and the survival function is

$$R(t) = \Pr(T > t) = 1 - F(t). \tag{7.2}$$

Survival analysis writes it $S(t)$; reliability engineering writes it $R(t)$, the reliability function, and this book uses $R(t)$. With $F(500) = 0.30$, $R(500) = 0.70$: the unit has a modeled probability of 70 % of surviving beyond 500 h. Continuity of the distribution adds one more property: $\Pr(T = t) = 0$ for every t, so $\Pr(T \ge t) = \Pr(T > t)$ and whether the boundary is included does not matter.

Failure density. For a continuous lifetime the density is the derivative of the cumulative probability, and because $F = 1 - R$ it is also minus the derivative of the survival function:

$$f(t) = \frac{\mathrm{d}F(t)}{\mathrm{d}t} = -\frac{\mathrm{d}R(t)}{\mathrm{d}t}. \tag{7.3}$$

By the fundamental theorem of calculus, the probability of failure between ages t_1 and t_2 is

$$\Pr(t_1 < T \le t_2) = F(t_2) - F(t_1) = \int_{t_1}^{t_2} f(t)\,\mathrm{d}t. \tag{7.4}$$

The density has units of inverse time, h^{-1}, and its value is not itself a probability; only its integral over an interval is. The three functions are related descriptions of one distribution: $F(t)$ the probability of failure by t, $R(t)$ the probability of survival beyond t, and $f(t)$ the density of the failure times.

The hazard from conditional probability. A unit that has already survived to age t prompts a different question: given that the motor is still operating at age t, how likely is it to fail almost immediately afterward? The answer is the hazard function, defined as a limit of a conditional probability per unit time,

$$\lambda(t) = \lim_{\Delta t \to 0^+} \frac{\Pr(t < T \le t + \Delta t \mid T > t)}{\Delta t}. \tag{7.5}$$

For a small interval, $\Pr(t < T \le t + \Delta t \mid T > t) \approx \lambda(t)\,\Delta t$. The hazard is an instantaneous conditional rate, in h^{-1}, not a probability; survival analysis writes it $h(t)$.

The definition reduces to the three functions above in three steps. By the definition of a conditional probability,

$$\Pr(t < T \le t + \Delta t \mid T > t) = \frac{\Pr(\{t < T \le t + \Delta t\} \cap \{T > t\})}{\Pr(T > t)}.$$

The event $\{t < T \le t + \Delta t\}$ is contained in $\{T > t\}$, so their intersection is the first event, whose probability is $F(t + \Delta t) - F(t)$ by equation (7.4); the denominator is $R(t)$:

$$\Pr(t < T \le t + \Delta t \mid T > t) = \frac{F(t + \Delta t) - F(t)}{R(t)}.$$

Divide by Δt. The denominator $R(t)$ does not depend on Δt, and the difference quotient of F tends to its derivative $f(t)$:

$$\lambda(t) = \frac{1}{R(t)} \lim_{\Delta t \to 0^+} \frac{F(t + \Delta t) - F(t)}{\Delta t} = \frac{f(t)}{R(t)}. \tag{7.6}$$

The density is the rate of failure per unit time among all units at installation; the hazard is the same rate among the units still working at age t, which is why it divides by the probability of being among them.

Cumulative hazard and survival. The identity (7.6) determines the survival function from the hazard. Substitute $f(t) = -R'(t)$ from equation (7.3):

$$\lambda(t) = -\frac{R'(t)}{R(t)}.$$

The right side is minus the derivative of a logarithm, since $\mathrm{d} \ln R(t)/\mathrm{d}t = R'(t)/R(t)$:

$$\frac{\mathrm{d}}{\mathrm{d}t} \ln R(t) = -\lambda(t).$$

Integrate from zero to t, with s the variable of integration and $R(0) = 1$ for a unit that is working when installed, so that $\ln R(0) = 0$:

$$\ln R(t) = -\int_0^t \lambda(s)\,\mathrm{d}s.$$

Define the cumulative hazard

$$\Lambda(t) = \int_0^t \lambda(s)\,\mathrm{d}s, \tag{7.7}$$

dimensionless, written $H(t)$ in survival analysis. Exponentiating,

$$R(t) = \mathrm{e}^{-\Lambda(t)}. \tag{7.8}$$

By equation (7.8) the whole survival distribution follows from a hazard model. Combining it with equations (7.2) and (7.6),

$$R(t) = \mathrm{e}^{-\Lambda(t)}, \qquad F(t) = 1 - \mathrm{e}^{-\Lambda(t)}, \qquad f(t) = \lambda(t)\,\mathrm{e}^{-\Lambda(t)}. \tag{7.9}$$

These are not three unrelated formulas; they are consequences of one hazard model, and a lifetime model may be specified by any one of R, F, f, λ or Λ.

The mean lifetime from the survival function. The mean lifetime $\mathrm{MTTF} = \mathbb{E}[T] = \int_0^\infty t\, f(t)\,\mathrm{d}t$ has a form in the survival function that both models below use. Substitute $f = -R'$ and integrate by parts:

$$\mathrm{MTTF} = \int_0^\infty t\,\big(-R'(t)\big)\,\mathrm{d}t = \Big[-t\,R(t)\Big]_0^\infty + \int_0^\infty R(t)\,\mathrm{d}t = \int_0^\infty R(t)\,\mathrm{d}t, \tag{7.10}$$

where the boundary term vanishes at zero trivially and at infinity whenever the mean is finite, since then $t\,R(t) \to 0$. The mean lifetime is the area under the survival curve.

Conditional survival. For a unit that has survived to age t_0, the probability of surviving a further τ hours is, by the definition of a conditional probability and because the event $\{T > t_0 + \tau\}$ is contained in $\{T > t_0\}$,

$$R(\tau \mid T > t_0) = \Pr(T > t_0 + \tau \mid T > t_0) = \frac{R(t_0 + \tau)}{R(t_0)} = \exp\big[-\big(\Lambda(t_0 + \tau) - \Lambda(t_0)\big)\big], \tag{7.11}$$

the conditional survival: the survival over the next τ hours is set by the cumulative hazard accumulated over that interval alone. When the lifetime distribution depends on the state, the same ratio is written $R(\tau \mid T > t_0, \mathbf{z})$; it describes the remaining life of a unit whose state is known.

7.1.3 The exponential and Weibull lifetimes

The exponential lifetime. The simplest lifetime model assumes a constant hazard $\lambda(t) = \lambda$, in h^{-1}, so that $\Lambda(t) = \lambda t$ and equations (7.9) give

$$R(t) = \mathrm{e}^{-\lambda t}, \qquad F(t) = 1 - \mathrm{e}^{-\lambda t}, \qquad f(t) = \lambda\,\mathrm{e}^{-\lambda t}, \tag{7.12}$$

with the mean $\mathrm{MTTF} = 1/\lambda$ by equation (7.10) and the median $t_{50} = \ln 2/\lambda$, the age at which $R(t_{50}) = 0.5$; Chapter 6 works these quantities for a numerical rate. Its conditional survival (7.11) is

$$\Pr(T > t_0 + \tau \mid T > t_0) = \frac{\mathrm{e}^{-\lambda(t_0+\tau)}}{\mathrm{e}^{-\lambda t_0}} = \mathrm{e}^{-\lambda\tau} = R(\tau). \tag{7.13}$$

Conditional on a fixed λ, the remaining survival probability does not depend on the age already reached: a unit that has run for a thousand hours is, for the model, as good as new. This is the memoryless property. It may be appropriate for failures caused by random external shocks, but it cannot represent a risk of failure that increases with age through wear-out.

The Weibull lifetime. The Weibull model is a more flexible description of equipment lifetimes. It has two parameters: the scale $\eta > 0$, in hours, and the shape $\beta > 0$, dimensionless. Its survival function is

$$R(t) = \exp\left[-\left(\frac{t}{\eta}\right)^{\beta}\right]. \tag{7.14}$$

Comparing it with $R = \mathrm{e}^{-\Lambda}$ identifies the cumulative hazard $\Lambda(t) = (t/\eta)^\beta$. By equation (7.7) the hazard is its derivative; with $\Lambda = \eta^{-\beta}\,t^\beta$, $\mathrm{d}\Lambda/\mathrm{d}t = \beta\,\eta^{-\beta}\,t^{\beta-1}$, which is

$$\lambda(t) = \frac{\beta}{\eta}\left(\frac{t}{\eta}\right)^{\beta-1}, \tag{7.15}$$

and the density follows from $f = \lambda\,R$:

$$f(t) = \frac{\beta}{\eta}\left(\frac{t}{\eta}\right)^{\beta-1}\exp\left[-\left(\frac{t}{\eta}\right)^{\beta}\right]. \tag{7.16}$$

At $\beta = 1$ the hazard is the constant $1/\eta$, and the Weibull lifetime reduces to the exponential with $\lambda = 1/\eta$: the exponential model is nested in the Weibull family.

By equation (7.10) the mean is the integral of equation (7.14). Substitute $x = (t/\eta)^\beta$, so that $t = \eta\, x^{1/\beta}$ and $\mathrm{d}t = (\eta/\beta)\, x^{1/\beta - 1}\, \mathrm{d}x$:

$$\mathrm{MTTF} = \int_0^\infty \mathrm{e}^{-(t/\eta)^\beta}\, \mathrm{d}t = \frac{\eta}{\beta} \int_0^\infty x^{1/\beta - 1}\, \mathrm{e}^{-x}\, \mathrm{d}x = \frac{\eta}{\beta}\, \Gamma\!\left(\frac{1}{\beta}\right) = \eta\, \Gamma\!\left(1 + \frac{1}{\beta}\right), \tag{7.17}$$

with Γ the gamma function and, in the last step, $(1/\beta)\,\Gamma(1/\beta) = \Gamma(1 + 1/\beta)$. The median solves $\exp[-(t_{50}/\eta)^\beta] = 0.5$; taking logarithms once, $(t_{50}/\eta)^\beta = \ln 2$, and

$$t_{50} = \eta\, (\ln 2)^{1/\beta}. \tag{7.18}$$

By equation (7.11) the conditional survival of a Weibull unit is

$$R(\tau \mid T > t_0) = \exp\!\left[-\frac{(t_0 + \tau)^\beta - t_0^\beta}{\eta^\beta}\right], \tag{7.19}$$

which at $\beta = 1$ is the memoryless (7.13).

A numerical example. With $\eta = 800\,\mathrm{h}$ and $\beta = 2$,

$$R(500) = \exp\!\left[-\left(\frac{500}{800}\right)^2\right] = \mathrm{e}^{-0.390625} = 0.6766, \qquad \lambda(500) = \frac{2}{800} \cdot \frac{500}{800} = 0.0015625\,\mathrm{h}^{-1},$$

and at 1000 h, $\lambda(1000) = (2/800)(1000/800) = 0.003125\,\mathrm{h}^{-1}$: the hazard doubles as the age doubles from 500 to 1000 h, which a constant-hazard model cannot represent. The density at 500 h is $f(500) = \lambda(500)\, R(500) = 0.0010572\,\mathrm{h}^{-1}$, smaller than the hazard because a third of the units have already failed by then. The mean lifetime is $800\,\Gamma(1.5) = 800 \times 0.88623 = 709.0\,\mathrm{h}$ and the median $800\sqrt{\ln 2} = 666.0\,\mathrm{h}$. A unit that has survived $t_0 = 400\,\mathrm{h}$ survives a further $\tau = 200\,\mathrm{h}$ with probability

$$R(200 \mid T > 400) = \exp\!\left[-\frac{600^2 - 400^2}{800^2}\right] = \exp\!\left[-\frac{200\,000}{640\,000}\right] = \mathrm{e}^{-0.3125} = 0.7316,$$

while a new unit survives its first 200 h with probability $R(200) = \exp[-(200/800)^2] = \mathrm{e}^{-0.0625} = 0.9394$. Because $\beta = 2$, the hazard increases with age, and the older unit has the lower probability of surviving the next 200 h.

The shape parameter. The exponent $\beta - 1$ in equation (7.15) sets how the hazard changes with age. Figure 7.1 draws the hazard and the survival function for three shapes at the common scale $\eta = 800\,\mathrm{h}$. In panel (a) the horizontal axis is age and the vertical axis the hazard in $10^{-3}\,\mathrm{h}^{-1}$; in panel (b) the vertical axis is the survival probability. With $\beta = 0.6$ the hazard falls with age, steeply at first; with $\beta = 1$ it is the constant $1/800 = 1.25 \times 10^{-3}\,\mathrm{h}^{-1}$; with $\beta = 2$ it rises in proportion to age. All three survival curves pass through the same point at $t = \eta$, the property derived below, and they cross there: the decreasing hazard loses more units early and fewer late.

Table 7.1 states what each regime may mean. Each row is a range of the shape; the columns give the behavior of the hazard and a possible engineering interpretation. The interpretations are possibilities, not diagnoses: a fitted shape parameter does not by itself establish the physical failure mechanism, a point taken up in Section 7.5.2.

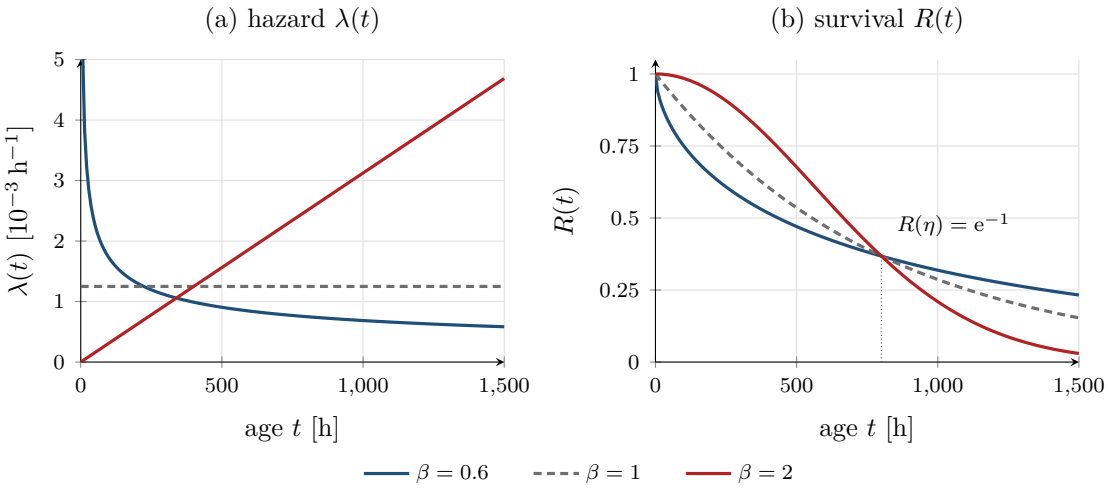


Figure 7.1: Weibull hazard and survival for three shapes at the scale $\eta = 800\,\mathrm{h}$, from equations (7.15) and (7.14): $\lambda(t) = (\beta/800)(t/800)^{\beta-1}$ and $R(t) = \exp[-(t/800)^{\beta}]$ for $\beta = 0.6$, 1 and 2. Illustrative parameters.

Table 7.1: The Weibull shape parameter and the hazard it implies.

Shape [–]	Hazard behavior	Possible engineering interpretation
$\beta < 1$	decreases with age	early-life defects, or the selective removal of weaker units
$\beta = 1$	constant	exponential lifetime, failures from random shocks
$\beta > 1$	increases with age	wear-out, degradation, fatigue, aging

Characteristic life. At $t = \eta$ the ratio t/η is one whatever the shape, so

$$R(\eta) = \mathrm{e}^{-1} = 0.3679, \qquad F(\eta) = 1 - \mathrm{e}^{-1} = 0.6321.$$

The Weibull scale is the age by which 63.2 % of the modeled population has failed, for every shape; this is the common point of the survival curves in Figure 7.1(b). The property returns in Section 7.3.2: at the characteristic life the survival probability does not depend on the shape at all.

7.1.4 The motors and their reliability graph (step 0)

A manufacturer evaluates eight identical electric motors operating under comparable conditions. Each motor is installed and monitored from its installation until it fails or until its observation window ends while it is still operating. The windows differ in length because the motors entered service on different dates and the study closed on one calendar date.

Figure 7.2 draws the system as a graph. The box on the left holds the operating conditions, the load, duty and environment that the study holds comparable across motors; it is pinned. The green box in the middle is the motor, the free node of the graph. It carries the unknown lifetime distribution, described by the two entries of the state: the logarithms $u = \ln \eta$ and $v = \ln \beta$ of the Weibull scale and shape of Section 7.1.3. The box on the right is the observed record of motor i, reached by a dashed observation arrow: an observed time t_i and an indicator δ_i of whether that time ended in a failure. Beneath it the record branches into its two possible outcomes, each with the likelihood factor it contributes. The unknown lifetime parameters, which survival analysis often writes $\boldsymbol{\theta}$, are in this book entries of the state $\mathbf{z}$.

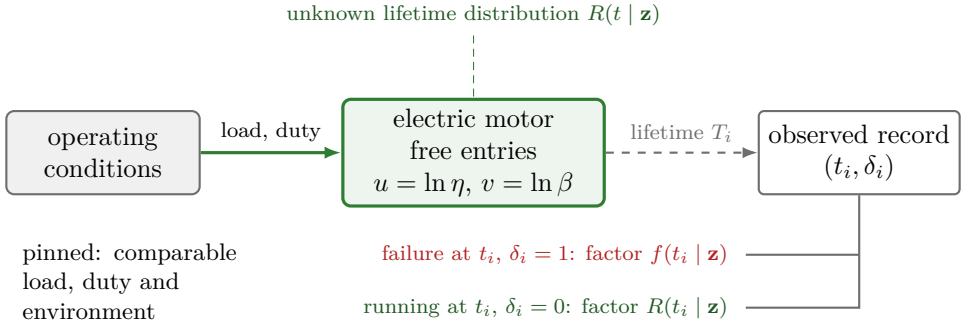


Figure 7.2: The reliability graph of the motor study. The operating conditions are pinned; the motor is the free node carrying the two state entries $u = \ln \eta$ and $v = \ln \beta$ of its lifetime distribution; the record (t_i, δ_i) of motor i is observed (dashed). A failure contributes the density at its time, a survivor the survival probability at its time.

The record. Table 7.2 lists the record. Each row is a motor; the columns give its observed time, the event indicator and the outcome. Five motors, A to E, failed during the study, at times between 120 and 600 h; three, F, G and H, were still running when their windows closed at 700, 800 and 900 h. The record is illustrative.

Figure 7.3 draws the same record against time. The horizontal axis is the age of a motor in hours, and each horizontal line, one per motor, runs from installation to the end of its observation. A filled circle at the end of a line marks an observed failure; a short vertical bar marks the end of observation with the motor still running. The failures are spread over the first 600 h, and the three longest lines all end in bars: the units observed longest are exactly the ones whose lifetimes are not known.

Table 7.2: The observed lifetimes of the eight motors (illustrative). The censored motors were operating at their recorded times.

Motor	Observed time t_i [h]	Indicator δ_i [–]	Outcome
A	120	1	failed
B	230	1	failed
C	310	1	failed
D	450	1	failed
E	600	1	failed
F	700	0	right-censored
G	800	0	right-censored
H	900	0	right-censored

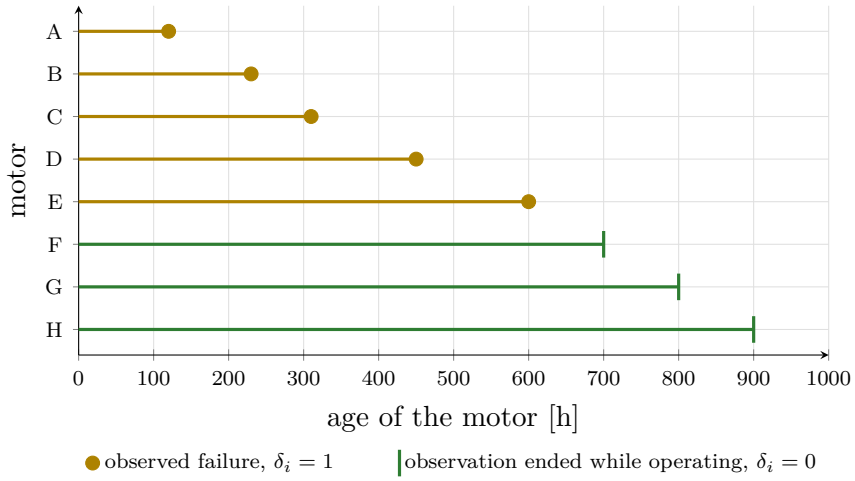


Figure 7.3: The observation record of the eight motors of Table 7.2. Each line is one motor's observed time; circles mark failures and vertical bars right censoring.

The record as pairs. For motor i let T_i be its failure time and c_i its censoring time, the age at which its observation would end if it had not failed. The study observes the smaller of the two,

$$t_i = \min(T_i, c_i),$$

together with the event indicator

$$\delta_i = \begin{cases} 1, & T_i \le c_i \quad \text{(the failure is observed)}, \\ 0, & T_i > c_i \quad \text{(the unit is censored)}. \end{cases}$$

The censoring time is written c_i rather than the more common C_i, which this book reserves for a consequence. The record of motor i is therefore

$$y_i = (t_i, \delta_i), \tag{7.20}$$

and the evidence $\mathbf{y}$ is the set of the eight pairs. If $\delta_i = 1$, the failure time is observed, $T_i = t_i$. If $\delta_i = 0$, only the inequality $T_i > t_i$ is known: the motor will fail at some later age that the study did not see. The record of Table 7.2 has $k = \sum_i \delta_i = 5$ failures, and its total time on test, failures and survivors together, is

$$T_{\text{tot}} = \sum_{i=1}^{8} t_i = 120 + 230 + 310 + 450 + 600 + 700 + 800 + 900 = 4110\,\text{h},$$

of which the five failed motors account for 1710 h and the three survivors for 2400 h. The number of failures is written k and the total time T_{tot}, as for the counts in exposure of the previous chapter.

The state entries. Both Weibull parameters are positive. The graph carries them on unbounded scales, as the conductance is carried by its logarithm in the thermal chapters, through the two entries of the state

$$\mathbf{z} = \begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} \ln \eta \\ \ln \beta \end{bmatrix} \begin{matrix} \text{log-scale} \\ \text{log-shape} \end{matrix}, \qquad \eta = \mathrm{e}^{u}, \quad \beta = \mathrm{e}^{v}, \tag{7.21}$$

with η in hours inside the logarithm. The exponential lifetime is the point $v = 0$ of this state, with $\lambda = \mathrm{e}^{-u}$. The mission reliability, the median lifetime and the remaining life of a motor are derived quantities, functions of the state.

7.1.5 The probabilistic questions (step 1)

The engineer asks eight questions of the same record, and step 1 writes each as a probability or an expectation over the state, with the unknown on the left of the bar and the record on the right. Table 7.3 lists them: the first column is the question in engineering language; the second the general form that every question of its category takes, written in the book's notation for an entry z_i of the state or a quantity $f(\mathbf{z})$ derived from it, a threshold f^*, the posterior density $p(\mathbf{z} \mid \mathbf{y})$, a future observation y_{new} and two models M_0 and M_1; the third its mathematical form for the motors; the fourth its category; and the last the subsection where it is answered. The first estimates the two entries and their correlation; the second to fifth ask for functions of both entries, a predictive, its complement, a tail and the quantiles of the median; the sixth reads the log-shape within the Weibull model and the seventh compares the exponential model with it; the eighth is a predictive conditioned on a survival already observed. The atlas of Appendix A lists the same forms.

Table 7.3: The questions asked of the motor record, the general form of their category, their mathematical form for the motors, their category and where each is answered. T_{new} is the lifetime of a new motor; t_0 the age already reached and τ a further period; $R(500;\mathbf{z})$ and $t_{50}(\mathbf{z})$ are the survival probability and the median lifetime as functions of the state; M_0 and M_1 the exponential and Weibull models; B_{01} the ratio of their evidences.

Engineering question	General form	Mathematical form	Category	Answered in
What are the scale and the shape of the lifetime, and how well are they known together?	$\Pr(z_i \le z^* \mid \mathbf{y})$, $\operatorname{Cov}(\mathbf{z} \mid \mathbf{y})$	$p(u, v \mid \mathbf{y})$; quantiles of η and β, $\operatorname{Cov}(u, v \mid \mathbf{y})$	estimation	7.3.1
Will a motor survive 500 h?	$\int p(y_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	$\Pr(T_{\text{new}} > 500 \mid \mathbf{y})$	prediction under a threshold	7.3.2
What is the probability of failure by 500 h?	$\int p(y_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	$\Pr(T_{\text{new}} \le 500 \mid \mathbf{y})$	prediction under a threshold	7.3.2
Is the 500 h survival probability below one half?	$\Pr(f(\mathbf{z}) < f^* \mid \mathbf{y})$	$\Pr(R(500;\mathbf{z}) < 0.5 \mid \mathbf{y})$	prediction under a threshold	7.3.2
What is the median lifetime?	s_q with $\Pr(f(\mathbf{z}) \le s_q \mid \mathbf{y}) = q$	$\Pr(t_{50}(\mathbf{z}) \le s \mid \mathbf{y}) = 0.5$, and the 0.025 and 0.975 points	estimation of a derived quantity	7.3.2
Does the risk of failure increase with age?	$\Pr(z_i > z^* \mid \mathbf{y})$	$\Pr(\beta > 1 \mid \mathbf{y}) = \Pr(v > 0 \mid \mathbf{y})$	prediction under a threshold	7.3.3
Which lifetime model describes the record?	$p(\mathbf{y} \mid M_0)/p(\mathbf{y} \mid M_1)$	B_{01}, exponential against Weibull	hypothesis comparison	7.3.3
How much longer will a surviving motor last?	$\int p(y_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y}, \text{age})\, \mathrm{d}\mathbf{z}$	$\Pr(T > t_0 + \tau \mid T > t_0, \mathbf{y})$	prediction under a threshold	7.3.4

7.1.6 The evidence as rows (step 2)

Step 2 attaches the record to the graph as rows, one per motor. None of them is Gaussian: each is the negative logarithm of one factor of a censored lifetime likelihood, the term that motor i puts in the objective. A failed motor's term is the negative logarithm of the failure density at its time, a censored motor's the negative logarithm of the survival probability at its time (the paragraphs below derive why):

$$\mathcal{L}_i(\mathbf{z}) = -\ln p(y_i \mid \mathbf{z}) = -\delta_i \ln f(t_i \mid \mathbf{z}) - (1-\delta_i)\ln R(t_i \mid \mathbf{z}), \qquad \mathcal{L}(\mathbf{z}) = \sum_{i=1}^{8} \mathcal{L}_i(\mathbf{z}). \tag{7.22}$$

Each term is also a count row in the working form of Section 4.2.7: the indicator δ_i counts the failures of motor i, whose expected number by its observed age is its cumulative hazard (derived below). Each row therefore has a working residual and a Jacobian row, and $\mathbf{J}$, the 8×2 Jacobian of this chapter, stacks those working rows.

The modeling choices. Four choices make the rows, and each buys something and can fail. The eight motors are treated as comparable, one lifetime distribution shared by all of them: this is what lets one state of two entries describe the whole record, and it fails if the motors ran at different loads or came from different lots (Section 7.6.5). Their lifetimes are taken as independent given the state and their censoring as noninformative: this makes the likelihood a product of one density or one survival probability per motor, and it fails if a shared cause stressed several motors at once or if a motor was withdrawn because it was deteriorating (Sections 7.1.6 and 7.6.5). The lifetime is taken to be Weibull: this gives a hazard that is a power of age, falling, constant or rising with one shape, and a scale that integrates out in closed form, and it fails if the hazard changes direction with age or the record mixes two populations, for which the Kaplan–Meier estimate of Section 7.4.6 is the reference. The prior is flat in u and v: this makes the MAP state the maximum-likelihood estimate and keeps the exact posterior to one integral, and with five failures it is not innocent, as the proper prior of the model comparison shows (Section 7.3.3).

The prior and the pattern of the rows. The prior of the fit is flat in u and v and adds no row, so that the MAP state coincides with the maximum-likelihood estimate. Two questions add rows of their own. The comparison of models gives the log-shape a proper prior, a Gaussian row v/σ_0 that acts on v alone; the exponential model is the Weibull model with that row in the limit $\sigma_0 \to 0$, the exact constraint $v = 0$ reached as a limit and checked numerically, never substituted. The remaining-life question adds a ninth row, the survival to its age of a motor found running, a censored row like those of F, G and H. Which row acts on which entry is a property of the graph, and it fixes where the Jacobian can be nonzero, marked $*$:

$$\mathbf{J} = \begin{bmatrix} * & * \\ \vdots & \vdots \\ * & * \\ * & * \\ 0 & * \end{bmatrix} \begin{matrix} \text{motor A, failure} \\ \vdots \\ \text{motor H, survivor} \\ \text{ninth motor, remaining life} \\ \text{prior row on } v\text{, comparison} \end{matrix}, \qquad \text{columns } u, v. \tag{7.23}$$

Every lifetime row acts on both entries, structurally: the density and the survival probability of a Weibull lifetime at a given age depend on the scale and the shape together, through $(t/\eta)^\beta$. Each lifetime row therefore puts curvature in the (u, v) position as well as on the diagonal, which is why u and v trade; only the prior row on v acts on one entry. No row is a physics row, so no model tolerance σ_p arises; σ_0 is the standard deviation of a prior row. The task of the answers is to reduce each question of Table 7.3 to a function of the posterior of $\mathbf{z}$ under these rows.

An observed failure. Consider first a fixed censoring time c_i, set by the closing date of the study. Motor i produces one of two kinds of observation, and each has its probability under the state. If the motor fails at $t_i \le c_i$, the observation is the event that T_i falls in a short interval $[t_i, t_i + \mathrm{d}t)$, of probability $f(t_i \mid \mathbf{z})\,\mathrm{d}t$. The width $\mathrm{d}t$ is the same for every value of the state and drops out of every comparison between states, so the likelihood contribution, the probability of the record y_i given the state up to that constant, is the density,

$$p(y_i \mid \mathbf{z}) = f(t_i \mid \mathbf{z}). \tag{7.24}$$

It is a density contribution, not a point probability.

A censored observation. If the motor is still running at $t_i = c_i$, the observation is the event $T_i > t_i$, whose probability is the survival function,

$$p(y_i \mid \mathbf{z}) = R(t_i \mid \mathbf{z}). \tag{7.25}$$

This is the key relation for right-censored data: *an observed failure contributes a density; a right-censored observation contributes a survival probability.*

Random censoring times. When the censoring time is itself random, with density $p_{\mathrm{cens}}(c)$ and survival function $S_{\mathrm{cens}}(c)$, two distinct assumptions are needed: the censoring time is independent of the lifetime, and its distribution does not depend on the state. Under the first, a failure at t_i has probability $f(t_i \mid \mathbf{z})\,S_{\mathrm{cens}}(t_i)\,\mathrm{d}t$ and a censoring at t_i the probability $p_{\mathrm{cens}}(t_i)\,R(t_i \mid \mathbf{z})\,\mathrm{d}t$; under the second, the censoring factors do not involve the state and drop out of every maximization and posterior, so that equations (7.24) and (7.25) remain the contributions. Censoring that satisfies both is *noninformative*. A withdrawal triggered by the motor's condition breaks the first assumption; a study whose closing date was chosen from the failures seen so far breaks the second.

The complete censored likelihood. The two contributions combine into one expression in the indicator, $p(y_i \mid \mathbf{z}) = f(t_i \mid \mathbf{z})^{\delta_i}\,R(t_i \mid \mathbf{z})^{1-\delta_i}$, whose negative logarithm is the row term (7.22). If the lifetimes are independent conditional on the state and the censoring is noninformative, the contributions multiply:

$$p(\mathbf{y} \mid \mathbf{z}) = \prod_{i=1}^{n} f(t_i \mid \mathbf{z})^{\delta_i}\,R(t_i \mid \mathbf{z})^{1-\delta_i}, \tag{7.26}$$

with n the number of units, the standard right-censored lifetime likelihood. Substituting $f = \lambda R$ from equation (7.6), the survival factors combine, since $R^{\delta_i}\,R^{1-\delta_i} = R$, and every unit, failed or not, contributes its survival probability:

$$p(\mathbf{y} \mid \mathbf{z}) = \prod_{i} \lambda(t_i \mid \mathbf{z})^{\delta_i}\,R(t_i \mid \mathbf{z}). \tag{7.27}$$

Taking logarithms, and using $\ln R(t) = -\Lambda(t)$ from equation (7.8), the log-likelihood is

$$\ell(\mathbf{z}) = \ln p(\mathbf{y} \mid \mathbf{z}) = \sum_{i} \delta_i \ln \lambda(t_i \mid \mathbf{z}) - \sum_{i} \Lambda(t_i \mid \mathbf{z})\,: \tag{7.28}$$

the failed units contribute the log-hazard at their failure times, and every unit contributes minus the cumulative hazard it accumulated while observed. The censored likelihood assumes that censoring does not selectively conceal failures in a way the modeled variables do not explain. Administrative censoring, the study ending on a predetermined date, is usually compatible with that assumption. If the most degraded motors are removed early for preventive replacement, their withdrawals are not ordinary noninformative censoring: a motor removed because its vibration was rising was more likely than its neighbors to fail soon, and treating its removal as survival overstates the reliability. Competing failure modes, maintenance interventions and condition-dependent withdrawals each require a model of their own.

The rows in the state. Write $x_i = \ln t_i - u$ for the logarithm of the age of unit i relative to the scale and $\Lambda_i = \exp(\beta x_i) = (t_i/\eta)^\beta$ for its cumulative hazard, with $\beta = \mathrm{e}^v$. The log-hazard at t_i is $v - u + (\beta - 1)x_i$ by equation (7.15), and the term that unit i puts in the objective is the negative logarithm of its factor in equation (7.27),

$$\mathcal{L}_i(\mathbf{z}) = -\ln p(y_i \mid \mathbf{z}) = \delta_i\big[u - v - (\beta - 1)x_i\big] + \Lambda_i, \tag{7.29}$$

which is $-\ln f(t_i \mid \mathbf{z})$ for a failure and $-\ln R(t_i \mid \mathbf{z}) = \Lambda_i$ for a survivor. The eight terms of the record are stacked as one vector, and with a prior flat in u and v the objective is their sum:

$$\begin{bmatrix}\mathcal{L}_1(\mathbf{z})\\ \vdots \\ \mathcal{L}_8(\mathbf{z})\end{bmatrix} = \begin{bmatrix}-\ln f(120 \mid \mathbf{z})\\ -\ln f(230 \mid \mathbf{z})\\ -\ln f(310 \mid \mathbf{z})\\ -\ln f(450 \mid \mathbf{z})\\ -\ln f(600 \mid \mathbf{z})\\ \Lambda(700 \mid \mathbf{z})\\ \Lambda(800 \mid \mathbf{z})\\ \Lambda(900 \mid \mathbf{z})\end{bmatrix}, \qquad \mathcal{L}(\mathbf{z}) = \sum_{i=1}^{8} \mathcal{L}_i(\mathbf{z}) = -\ell(\mathrm{e}^u, \mathrm{e}^v). \tag{7.30}$$

The score rows and the gradient. The ingredients of the derivatives are $\partial x_i/\partial u = -1$, $\partial\beta/\partial v = \beta$, $\partial\Lambda_i/\partial u = -\beta\Lambda_i$ and $\partial\Lambda_i/\partial v = \beta x_i \Lambda_i$. Differentiating equation (7.29), the gradient of row i, its score row, is

$$\frac{\partial \mathcal{L}_i}{\partial \mathbf{z}^\mathsf{T}} = \begin{bmatrix}\beta\,(\delta_i - \Lambda_i) & -\delta_i(1 + \beta x_i) + \beta x_i \Lambda_i\end{bmatrix}, \tag{7.31}$$

and the gradient of the objective is the sum of the score rows,

$$\nabla\mathcal{L}(\mathbf{z}) = \sum_i \frac{\partial \mathcal{L}_i}{\partial \mathbf{z}} = \begin{bmatrix}\beta(k - \sum_i \Lambda_i)\\ -k - \beta\sum_i \delta_i x_i + \beta \sum_i x_i \Lambda_i\end{bmatrix}. \tag{7.32}$$

A failure pushes u down with weight $\beta(1 - \Lambda_i)$ and a survivor pushes it up with weight $\beta\Lambda_i$; the second entry carries the lever x_i of each age relative to the scale. A survivor's row at age t is $\Lambda(t \mid \mathbf{z})$, and its score row $(-\beta\Lambda,\ \beta\Lambda\ln(t/\eta))$ is the gradient of the cumulative hazard that gave the mission reliability its gradient (7.59).

The rows as count rows. Since $u + x_i = \ln t_i$, the term $\delta_i[u - (\beta - 1)x_i]$ is $-\delta_i\beta x_i$ plus a constant, and with $\ln\Lambda_i = \beta x_i$

$$\mathcal{L}_i(\mathbf{z}) = (\Lambda_i - \delta_i \ln\Lambda_i) - \delta_i v + \text{const}:$$

the term in parentheses is the negative log-probability of a Poisson count δ_i with mean Λ_i, the count row of Chapter 6, and the rest is linear in v. In the working form of Section 4.2.7, equation (4.43), of which equation (6.22) is the pump's instance, the count row has the residual and the Jacobian row

$$r_i(\mathbf{z}) = \frac{\Lambda_i - \delta_i}{\sqrt{\Lambda_i}}, \qquad \mathbf{J}_i = \sqrt{\Lambda_i}\,\frac{\partial \ln\Lambda_i}{\partial \mathbf{z}^\mathsf{T}} = \sqrt{\Lambda_i}\,\beta\begin{bmatrix}-1 & x_i\end{bmatrix}, \tag{7.33}$$

both entries structurally nonzero, the pattern (7.23); $\mathbf{J}$ in this chapter stacks these rows. The gradient of the objective is $\mathbf{J}^\mathsf{T}\mathbf{r} - (0,\ k)^\mathsf{T}$, the last term from the linear terms $-\delta_i v$, and its Hessian is

$$\nabla^2\mathcal{L}(\mathbf{z}) = \mathbf{J}^\mathsf{T}\mathbf{J} + \sum_i (\Lambda_i - \delta_i)\,\frac{\partial^2 \ln\Lambda_i}{\partial \mathbf{z}\,\partial \mathbf{z}^\mathsf{T}}, \qquad \frac{\partial^2 \ln\Lambda_i}{\partial \mathbf{z}\,\partial \mathbf{z}^\mathsf{T}} = \begin{bmatrix}0 & -\beta\\ -\beta & \beta x_i\end{bmatrix}: \tag{7.34}$$

the Gauss–Newton product of the working rows plus the named curvature term, the matrix $\mathbf{C}$ of equation (4.44). That term does not vanish here, because $\ln\Lambda_i = \mathrm{e}^v(\ln t_i - u)$ is not linear in the state; the log-shape is not the canonical link of these counts. With unnormalized rows and the weights $\mathbf{W}$ the product would be $\mathbf{J}^\mathsf{T}\mathbf{W}\mathbf{J}$.

The best scale for each shape. The first entry of equation (7.32) vanishes when $\sum_i \Lambda_i = k$: for each shape there is one scale at which the cumulative hazards of all eight units, failed and censored, add up to the number of failures. Solving $\sum_i t_i^\beta = k\,\eta^\beta$ for the scale traces this curve of best scale for each shape in the state plane,

$$\hat{\eta}(\beta) = \left(\frac{\sum_i t_i^\beta}{k} \right)^{1/\beta} . \tag{7.35}$$

At $\beta = 1$ the curve passes through $\hat{\eta}(1) = T_{\text{tot}}/k = 1/\hat{\lambda}$, the exponential fit (7.67) at $v = 0$.

The curvature of each row. Differentiating the score row (7.31) once more gives the 2×2 Hessian of row i,

$$\mathbf{\Lambda}_i = \nabla^2 \mathcal{L}_i = \begin{bmatrix} \beta^2 \Lambda_i & \beta(\delta_i - \Lambda_i) - \beta^2 x_i \Lambda_i \\ \beta(\delta_i - \Lambda_i) - \beta^2 x_i \Lambda_i & -\delta_i \beta x_i + \beta x_i \Lambda_i + \beta^2 x_i^2 \Lambda_i \end{bmatrix}, \tag{7.36}$$

which for a survivor, $\delta_i = 0$, is the Hessian of Λ_i alone, the curvature of the ninth row of the remaining-life question. The rows add, and their sum is the Hessian (7.34):

$$\nabla^2 \mathcal{L}(\mathbf{z}) = \sum_{i=1}^{8} \nabla^2 \mathcal{L}_i(\mathbf{z}). \tag{7.37}$$

7.2 The posterior of the state (step 3)

The eight questions of Table 7.3 take five general forms, written in the book's notation: $p(\mathbf{z} \mid \mathbf{y})$ is the posterior density of the state, z_i one of its entries, $f(\mathbf{z})$ a quantity derived from the state, f^* a threshold on it, y_{new} a future observation and $p(\mathbf{y} \mid M)$ the evidence of a model M.

The Weibull lifetime has no conjugate prior: no prior on (u, v) makes the posterior a standard distribution in both entries. One integral nonetheless survives in closed form. Under the prior flat in (u, v) the scale integrates out given the shape, and every question becomes an exact one-dimensional integral over v (Section 7.2.2). The report gives the answers of the Laplace route, the one a model without that integral must take and the one on which the preposterior of step 7 is written, with each exact integral beside it. The route runs through four objects named here by their roles. The *MAP state* $\hat{\mathbf{z}} = (\hat{u}, \hat{v})^\mathsf{T}$ is the minimum of the objective $\mathcal{L} = \sum_i \mathcal{L}_i$; the *information matrix* $\mathbf{\Lambda}$ is its curvature at the minimum, the sum of the row Hessians, which in the working form is the Gauss–Newton product $\mathbf{J}^\mathsf{T}\mathbf{J}$ plus a named curvature term (Section 7.1.6); the *Laplace covariance* $\mathbf{\Sigma} = \mathbf{\Lambda}^{-1}$, a 2×2 matrix with an off-diagonal entry, makes the posterior approximately normal, $\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$; and a derived quantity f is carried by its *gradient* $\mathbf{g} = \partial f / \partial \mathbf{z}$ at $\hat{\mathbf{z}}$, with the variance $\mathbf{g}^\mathsf{T}\mathbf{\Sigma}\mathbf{g}$. Every answer is computed from $\hat{\mathbf{z}}$ and $\mathbf{\Sigma}$; the exact posterior gives the value set beside it.

All eight are read from one posterior: the posterior of the state that the eight lifetime rows of step 2 give, the Laplace normal of (u, v), with the exact posterior through one integral over the shape beside it. This section derives it and computes it for the motors; Section 7.3 reads every answer from it.

7.2.1 The posterior from the rows

The objective and the Laplace posterior. The fit uses a prior flat in u and v, which adds no row, so that the MAP state is the maximum-likelihood estimate; that prior is a choice, the density $1/(\eta\beta)$ on the original parameters, and it keeps the exact posterior of Section 7.2.2 in closed form up to one integral. In the state (7.21) the negative log posterior is the objective

$$\mathcal{L}(\mathbf{z}) = \sum_i \mathcal{L}_i(\mathbf{z}) - \ln p(\mathbf{z}) + \text{const}, \qquad \hat{\mathbf{z}} = \arg\min_{\mathbf{z}} \mathcal{L}(\mathbf{z}), \tag{7.38}$$

and near its minimum the Laplace approximation of Chapter 4 gives

$$\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma}), \qquad \boldsymbol{\Sigma} = \boldsymbol{\Lambda}^{-1}, \qquad \boldsymbol{\Lambda} = \nabla^2_{\mathbf{z}} \mathcal{L}(\mathbf{z})\big|_{\hat{\mathbf{z}}}. \tag{7.39}$$

The Newton step. Linearizing the gradient around the current state $\mathbf{z}_m$ and setting it to zero gives the Newton equations, the normal equations of the working rows with the curvature term kept,

$$\nabla^2 \mathcal{L}(\mathbf{z}_m)\, \Delta\mathbf{z} = -\nabla \mathcal{L}(\mathbf{z}_m), \qquad \mathbf{z}_{m+1} = \mathbf{z}_m + \Delta\mathbf{z}, \tag{7.40}$$

with the gradient (7.32) on the right and the Hessian (7.34), the sum of the row Hessians, on the left. The iteration stops where both entries of the gradient vanish, and there the cumulative hazards add to k: the MAP state lies on the curve of best scale (7.35).

The information in closed form. At $\hat{\mathbf{z}}$ the two gradient entries (7.32) vanish: $\sum_i \Lambda_i = k$ and $\beta \sum_i x_i \Lambda_i - \beta \sum_i \delta_i x_i = k$. The curvature term of equation (7.34) then has the (u, v) entry $-\beta \sum_i (\Lambda_i - \delta_i) = 0$ and the (v, v) entry $\beta \sum_i (\Lambda_i - \delta_i) x_i = k$, and the Gauss–Newton product of the working rows (7.33) is $\beta^2 \sum_i \Lambda_i (1,\ -x_i)^{\mathsf{T}} (1,\ -x_i)$, so that

$$\boldsymbol{\Lambda} = \mathbf{J}^{\mathsf{T}}\mathbf{J} + \begin{bmatrix} 0 & 0 \\ 0 & k \end{bmatrix} = \begin{bmatrix} \hat{\beta}^2 k & -\hat{\beta}^2 \sum_i x_i \Lambda_i \\ -\hat{\beta}^2 \sum_i x_i \Lambda_i & k + \hat{\beta}^2 \sum_i x_i^2 \Lambda_i \end{bmatrix}, \tag{7.41}$$

with x_i and Λ_i at $\hat{\mathbf{z}}$. The (u, u) entry receives $\hat{\beta}^2 \Lambda_i$ from every unit, failed or censored, and the survivors, whose cumulative hazards are the largest, carry much of it; its sum, however, is $\hat{\beta}^2 k$, here and at every point of the curve of best scale, where $\sum_i \Lambda_i = k$. More hours on the survivors move $\hat{\mathbf{z}}$ to a longer scale at which the cumulative hazards again add to k, and the (u, u) entry stays $\hat{\beta}^2$ per failure: the precision of the scale grows with failures, not with hours on test. In the rows, the (v, v) entry of a failure far below the scale is dominated by its term $-\delta_i \beta x_i$, large because its lever is large and negative; a single row need not be convex, only their sum.

A new row on a fitted state. When a row is added to a fitted record, the prior row of the model comparison or the ninth row of the remaining-life question, the old rows have zero gradient at $\hat{\mathbf{z}}$, so the gradient of the enlarged objective there is the new row's alone, and the Newton equations (7.40) start from $\hat{\mathbf{z}}$ with the Hessian $\boldsymbol{\Lambda}$ plus the new row's curvature.

The objective along the best scale. The shape of $\mathcal{L}$ near its minimum is seen by following the curve of best scale. At each shape, put the state at its best scale, $u = \ln \hat{\eta}(\beta)$; there $\sum_i \Lambda_i = k$ and $k\beta u = k \ln(\sum_i t_i^\beta / k)$, and summing the terms (7.29) gives minus the objective as the profile log-likelihood of the shape,

$$\ell_p(\beta) = -\mathcal{L}(\ln \hat{\eta}(\beta),\ \ln \beta) = k \ln \beta - k \ln\left(\frac{1}{k} \sum_i t_i^\beta \right) + (\beta - 1) \sum_i \delta_i \ln t_i - k. \tag{7.42}$$

It is a slice of $\mathcal{L}$ through the minimum at the best scale for each shape, not a separate likelihood. Setting its slope to zero gives one equation in the shape,

$$\frac{1}{\beta} + \frac{1}{k} \sum_i \delta_i \ln t_i - \frac{\sum_i t_i^\beta \ln t_i}{\sum_i t_i^\beta} = 0, \tag{7.43}$$

whose left side decreases monotonically in β and whose root is the MAP shape $\hat{\beta}$: the scalar check of the Newton fit. The Laplace Gaussian read in the shape is a parabola in this slice: the marginal precision of v, $1/\Sigma_{vv}$, is the Schur complement $\Lambda_{vv} - \Lambda_{uv}^2 / \Lambda_{uu}$, the curvature left in v once u is eliminated at its best scale, and since $\mathrm{d}v = \mathrm{d}\beta / \hat{\beta}$ near the maximum the parabola has the curvature $1/(\hat{\beta}^2 \Sigma_{vv})$ in the shape.

7.2.2 The exact posterior through one integral

Under the flat prior in (u, v) the posterior is proportional to the likelihood, and for the Weibull model the scale can be integrated out in closed form. Every question of Table 7.3 then reduces to a one-dimensional integral over the shape.

The scale given the shape. Write $S_\delta = \sum_i \delta_i \ln t_i$ and $A(\beta) = \sum_i t_i^\beta$. In the sum of the rows (7.29), which is $-\ell$, $(\beta-1)\sum_i \delta_i x_i = (\beta-1)S_\delta - (\beta-1)ku$ and $\Lambda_i = t_i^\beta \mathrm{e}^{-\beta u}$, so

$$\ell(u, v) = kv - k\beta u + (\beta-1)S_\delta - \mathrm{e}^{-\beta u} A(\beta).$$

Define $\zeta = \mathrm{e}^{-\beta u} = \eta^{-\beta}$. Then

$$p(\mathbf{y} \mid \mathbf{z}) = \beta^k \, \mathrm{e}^{(\beta-1)S_\delta} \, \zeta^k \, \mathrm{e}^{-A(\beta)\,\zeta}.$$

At fixed v, the change of variable from u to ζ (Section 4.3.3) has $|\mathrm{d}u/\mathrm{d}\zeta| = 1/(\beta\zeta)$, so the posterior of ζ given the shape is proportional to $\zeta^{k-1}\mathrm{e}^{-A(\beta)\zeta}$:

$$\zeta \mid v, \mathbf{y} \sim \mathrm{Gamma}\big(k,\ A(\beta)\big), \tag{7.44}$$

a gamma distribution with shape k and rate $A(\beta)$. For a small number of failures that gamma is skewed, and its logarithm, which is $-\beta u$, is far from normal: the scale given the shape is not Gaussian in u, however well the shape is known. Since $u = -\ln\zeta/\beta$, the mean and the variance of the logarithm of a gamma variable (Chapter 8, with the digamma and trigamma functions ψ and ψ') give

$$\mathbb{E}[u \mid v, \mathbf{y}] = \frac{\ln A(\beta) - \psi(k)}{\beta}, \qquad \mathrm{Var}(u \mid v, \mathbf{y}) = \frac{\psi'(k)}{\beta^2}, \tag{7.45}$$

and averaging over the marginal of v below gives the posterior of the scale: $\mathbb{E}[u \mid \mathbf{y}]$ is the average of the conditional mean, $\mathrm{Var}(u \mid \mathbf{y})$ the average of the conditional variance plus the variance of the conditional mean, $\mathrm{Cov}(u, v \mid \mathbf{y})$ the average of $v\,\mathbb{E}[u \mid v, \mathbf{y}]$ less the product of the means, and $\Pr(u \le x \mid \mathbf{y})$ the average of the gamma upper tail $\Pr(\zeta \ge \mathrm{e}^{-\beta x} \mid v, \mathbf{y})$. Because the conditional mean $\ln A(\beta)/\beta$ falls as the shape rises, the scale and the shape are correlated in the exact posterior even though no single row Hessian enters it.

The marginal of the shape. Integrating the likelihood over u with the same change of variable and $\int_0^\infty \zeta^{k-1}\mathrm{e}^{-A\zeta}\,\mathrm{d}\zeta = \Gamma(k)/A^k$,

$$\Omega(v) = \int p(\mathbf{y} \mid u, v)\,\mathrm{d}u = \Gamma(k)\,\beta^{k-1}\,\mathrm{e}^{(\beta-1)S_\delta}\,A(\beta)^{-k}, \qquad p(v \mid \mathbf{y}) = \frac{\Omega(v)}{\int \Omega(v')\,\mathrm{d}v'}. \tag{7.46}$$

The state-vector form says what integrating the scale out does. At each shape the objective is smallest at the best scale of equation (7.35), where the (u, u) curvature is $\beta^2 k$ (equation (7.41) and the line after it). Integrating $\exp(-\mathcal{L})$ over u around that minimum multiplies the slice $\mathrm{e}^{\ell_p(\beta)}$ of equation (7.42) by $(2\pi/\beta^2 k)^{1/2} \propto 1/\beta$, and for these rows that Laplace step over u is exact up to a constant. The marginal of v is therefore the slice weighted by $1/\beta$, a weight that changes along v and that a single $\boldsymbol{\Sigma}$ at $\hat{\mathbf{z}}$ cannot carry: it moves the mass of v toward small shapes and the marginal mode below the joint MAP $\hat{v}$.

The answers as one integral. Given the shape, the survival probability is $R(t; \mathbf{z}) = \exp(-t^\beta\zeta)$, and the gamma distribution (7.44) has the Laplace transform $\mathbb{E}[\mathrm{e}^{-s\zeta}] = (A/(A+s))^k$. Averaging over ζ and then over v,

$$\Pr(T_{\mathrm{new}} > t \mid \mathbf{y}) = \int p(v \mid \mathbf{y}) \left(\frac{A(\beta)}{A(\beta) + t^\beta}\right)^k \mathrm{d}v. \tag{7.47}$$

The event $R(t;\mathbf{z}) < R^*$ is the event $\zeta > -\ln R^*/t^\beta$, so with $\Upsilon_k(x) = \mathrm{e}^{-x}\sum_{j=0}^{k-1} x^j/j!$ the upper tail of a unit-rate gamma variable of integer shape k,

$$\Pr\bigl(R(t;\mathbf{z}) < R^* \mid \mathbf{y}\bigr) = \int p(v \mid \mathbf{y})\, \Upsilon_k\!\left(\frac{-A(\beta)\ln R^*}{t^\beta}\right) \mathrm{d}v. \tag{7.48}$$

The median lifetime is $t_{50}(\mathbf{z}) = (\ln 2/\zeta)^{1/\beta}$, and in the same way

$$\Pr\bigl(t_{50}(\mathbf{z}) \le s \mid \mathbf{y}\bigr) = \int p(v \mid \mathbf{y})\, \Upsilon_k\!\left(\frac{A(\beta)\ln 2}{s^\beta}\right) \mathrm{d}v, \qquad \Pr(\beta > 1 \mid \mathbf{y}) = \int_0^\infty p(v \mid \mathbf{y})\, \mathrm{d}v.$$

The remaining life (7.72) is the ratio of the predictives (7.47) at $t_0 + \tau$ and at t_0, and the distribution of the conditional survival itself follows as in equation (7.48) on the record with the ninth row $(t_0, 0)$ added. The evidence ratio (7.68) is $\Omega(0)$ against the integral of Ω times the prior density of v. Each is a quadrature over one variable, a few hundred evaluations of $A(\beta)$.

Building blocks. The sums S_δ and $A(\beta)$ of the record, the gamma, digamma and trigamma functions, the tail Υ_k and a quadrature on a grid in v. The exact values set beside the answers are computed in Section 7.3.

7.2.3 The posterior of the eight motors

The motors' inputs are the eight pairs of Table 7.2, with $k = 5$ failures and $T_{\mathrm{tot}} = 4110\,\mathrm{h}$ on test, and a grid of 160 001 points in v for the exact integrals.

The Newton iteration. The start is $\eta = 600\,\mathrm{h}$ and $\beta = 1.5$, $\mathbf{z}_0 = (6.39693,\ 0.40547)^{\mathsf{T}}$. There the gradient and the Hessian of equation (7.40) are

$$\nabla\mathcal{L} = \begin{bmatrix} -2.97681 \\ 1.14214 \end{bmatrix}, \qquad \nabla^2\mathcal{L} = \begin{bmatrix} 15.7152 & -4.2783 \\ -4.2783 & 8.6733 \end{bmatrix}, \qquad \Delta\mathbf{z} = \begin{bmatrix} 0.17739 \\ -0.04418 \end{bmatrix}:$$

the scale is too short for the survivors, and the step lengthens it and lowers the shape. Table 7.4 lists the iterates. Each row is an iteration; the columns give the state, the objective, the two entries of the gradient and the step. The length of the gradient falls from 3.19 to 0.75 and 0.10, then quadratically to 0.0027 and 2×10^{-6}, and four steps reach the MAP state to the printed precision.

Table 7.4: Newton iterations (7.40) for the Weibull fit to the motor record from $\mathbf{z}_0 = (\ln 600,\ \ln 1.5)^{\mathsf{T}}$, with the rows (7.30), the gradient (7.32) and the curvature (7.37).

m [–]	u_m [–]	v_m [–]	$\mathcal{L}(\mathbf{z}_m)$ [–]	$\partial\mathcal{L}/\partial u$ [–]	$\partial\mathcal{L}/\partial v$ [–]	Δu [–]	Δv [–]
0	6.39693	0.40547	38.70002	−2.97681	1.14214	0.17739	−0.04418
1	6.57432	0.36128	38.36882	−0.55603	0.50900	0.05410	−0.07295
2	6.62842	0.28834	38.33233	−0.09325	0.04587	0.01186	−0.00921
3	6.64028	0.27913	38.33156	−0.00245	0.00118	0.00033	−0.00025
4	6.64061	0.27887	38.33156	0.00000	0.00000		

The MAP state. The MAP state and the corresponding parameters are

$$\hat{\mathbf{z}} = \begin{bmatrix} \hat u \\ \hat v \end{bmatrix} = \begin{bmatrix} 6.64061 \\ 0.27887 \end{bmatrix}, \quad \hat\beta = \mathrm{e}^{\hat v} = 1.3216, \quad \hat\eta = \mathrm{e}^{\hat u} = 765.6\,\mathrm{h}, \quad \ell(\hat\eta, \hat\beta) = -\mathcal{L}(\hat{\mathbf{z}}) = -38.3316, \tag{7.49}$$

and the start at the exponential fit, $\mathbf{z}_0 = (\ln 822,\ 0)^{\mathsf{T}}$, reaches the same state in four steps. At $\hat{\mathbf{z}}$ the cumulative hazards add to $\sum_i \Lambda_i = 5.0000$: the MAP state lies on the curve of best scale (7.35), at $\hat\eta(1.3216) = 765.6\,\mathrm{h}$.

The rows at the MAP state. Table 7.5 evaluates every row at $\hat{\mathbf{z}}$. Each row of the table is a motor; the columns give its indicator, its lever x_i, its cumulative hazard Λ_i, its score row (7.31), and the three distinct entries of its Hessian (7.36). The two score columns sum to zero, the gradient at the MAP state; the column of Λ_i sums to 5, the condition of best scale; and the three Hessian columns sum to the entries of $\boldsymbol{\Lambda}$.

Table 7.5: The rows (7.29) of the motor record at the MAP state $\hat{\mathbf{z}} = (\ln 765.56,\ \ln 1.32164)^{\mathsf{T}}$: the lever $x_i = \ln(t_i/\hat{\eta})$, the cumulative hazard $\Lambda_i = (t_i/\hat{\eta})^{\hat{\beta}}$, the score row (7.31) and the row Hessian (7.36). All entries are dimensionless.

Motor	δ_i	x_i	Λ_i	$\partial\mathcal{L}_i/\partial u$	$\partial\mathcal{L}_i/\partial v$	$\partial^2\mathcal{L}_i/\partial u^2$	$\partial^2\mathcal{L}_i/\partial u\,\partial v$	$\partial^2\mathcal{L}_i/\partial v^2$
A	1	−1.8531	0.0864	1.2075	1.2376	0.1509	1.4871	2.7557
B	1	−1.2025	0.2041	1.0519	0.2650	0.3564	1.4806	1.7804
C	1	−0.9040	0.3028	0.9215	−0.1669	0.5288	1.3996	1.2653
D	1	−0.5314	0.4955	0.6668	−0.6457	0.8654	1.1267	0.5987
E	1	−0.2437	0.7247	0.3639	−0.9113	1.2658	0.6723	0.1638
F	0	−0.0895	0.8884	−1.1742	−0.1051	1.5518	−1.0352	−0.0927
G	0	0.0440	1.0599	−1.4008	0.0616	1.8513	−1.4822	0.0652
H	0	0.1618	1.2384	−1.6367	0.2648	2.1632	−1.9867	0.3214
Sum	5		5.0000	0.0000	0.0000	8.7337	1.6621	6.8579

The three survivors carry 5.57 of the 8.73 in the (u,u) column, and the earliest failures, A, B and C, carry 5.80 of the 6.86 in the (v,v) column, where motor F contributes a small negative curvature. With $\hat{\beta}^2 = 1.74673$, $\sum_i x_i\Lambda_i = -0.95155$ and $\sum_i x_i^2\Lambda_i = 1.06363$, the working rows (7.33) give the Gauss–Newton product and equation (7.41) adds the curvature term $k = 5$ to its (v,v) entry:

$$\mathbf{J}^{\mathsf{T}}\mathbf{J} = \begin{bmatrix} 8.7337 & 1.6621 \\ 1.6621 & 1.8579 \end{bmatrix}, \quad \boldsymbol{\Lambda} = \begin{bmatrix} 8.7337 & 1.6621 \\ 1.6621 & 6.8579 \end{bmatrix}, \quad \boldsymbol{\Sigma} = \boldsymbol{\Lambda}^{-1} = \begin{bmatrix} 0.12004 & -0.02909 \\ -0.02909 & 0.15287 \end{bmatrix}, \tag{7.50}$$

the column sums of the table, with $\det\boldsymbol{\Lambda} = 57.132$. The eigenvectors of $\boldsymbol{\Sigma}$ are $(0.864,\ 0.504)^{\mathsf{T}}$, a longer scale together with a larger shape, with standard deviation 0.321, and $(-0.504,\ 0.864)^{\mathsf{T}}$, with standard deviation 0.412.

The two extended records. Two questions of Table 7.3 add a row to the eight, the last two rows of the pattern (7.23), and each extended record has its posterior of the state by the same construction, the Newton steps of a new row on a fitted state started from $\hat{\mathbf{z}}$ (Section 7.2.1). The model comparison adds the prior row $r_\pi = v/0.5$ of a proper prior $\mathcal{N}(0, 0.5^2)$ on the log-shape, of Jacobian row $(0,\ 2)$, whose Gauss–Newton product adds $\mathrm{diag}(0,\ 4)$ to the Hessian. Its gradient at $\hat{\mathbf{z}}$ is $(0,\ 1.11549)^{\mathsf{T}}$, and three Newton steps (7.40), the first $(0.02014,\ -0.10582)$, reach $\hat{\mathbf{z}}_\pi = (6.66428,\ 0.16856)^{\mathsf{T}}$, $\hat{\eta} = 783.9\,\mathrm{h}$ and $\hat{\beta} = 1.1836$: the prior row pulls the shape toward one. The curvature there and its inverse are

$$\boldsymbol{\Lambda}_\pi = \begin{bmatrix} 7.0046 & 1.6787 \\ 1.6787 & 10.1010 \end{bmatrix}, \qquad \boldsymbol{\Sigma}_\pi = \begin{bmatrix} 0.14869 & -0.02471 \\ -0.02471 & 0.10311 \end{bmatrix},$$

so $\sqrt{\Sigma_{\pi,vv}} = 0.3211$. With the wider prior $\mathcal{N}(0, 1^2)$ the same steps reach $(6.64810,\ 0.24114)^{\mathsf{T}}$ with $\sqrt{\Sigma_{vv}} = 0.3716$. The remaining-life question adds the ninth row of a motor found running at $t_0 = 500\,\mathrm{h}$, a censored row with $t_9 = 500\,\mathrm{h}$ and $\delta_9 = 0$. At $\hat{\mathbf{z}}$ it has the score row $(-\beta\Lambda_9,\ \beta x_9\Lambda_9) = (-0.7527,\ -0.3206)$ and the row Hessian (7.36) $\left[\begin{smallmatrix} 0.9947 & -0.3289 \\ -0.3289 & -0.1401 \end{smallmatrix}\right]$, so that the Newton equations (7.40) start from $\hat{\mathbf{z}}$ with the Hessian $\boldsymbol{\Lambda} + \boldsymbol{\Lambda}_9 = \left[\begin{smallmatrix} 9.7284 & 1.3332 \\ 1.3332 & 6.7178 \end{smallmatrix}\right]$; the first step is

$(0.07281,\ 0.03328)$, and three steps converge. The refitted state is $\hat{\mathbf{z}}_9 = (6.71498,\ 0.30388)^{\mathsf{T}}$, that is $\hat{\eta} = 824.7\,\mathrm{h}$ and $\hat{\beta} = 1.3551$, and the curvature there gives

$$\boldsymbol{\Lambda}_9 = \begin{bmatrix} 9.1816 & 2.6018 \\ 2.6018 & 7.2057 \end{bmatrix}, \qquad \boldsymbol{\Sigma}_9 = \boldsymbol{\Lambda}_9^{-1} = \begin{bmatrix} 0.12133 & -0.04381 \\ -0.04381 & 0.15460 \end{bmatrix}, \tag{7.51}$$

so $\sqrt{\Sigma_{9,uu}} = 0.348$, $\sqrt{\Sigma_{9,vv}} = 0.393$ and the correlation is -0.320: the ninth row moves the state toward a longer scale and strengthens the trade between the two entries.

The pair $(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of equations (7.49) and (7.50) is what step 4 uses: every answer of the chapter is computed from it, the model comparison from $(\hat{\mathbf{z}}_\pi, \boldsymbol{\Sigma}_\pi)$ and the remaining life from $(\hat{\mathbf{z}}_9, \boldsymbol{\Sigma}_9)$, the same posterior on the extended records. Because the Weibull record also admits the exact posterior through one integral over the shape (Section 7.2.2), each answer has an exact value beside it, and the gap between the two is reported with it.

7.3 The answers (step 4)

The questions of Table 7.3 fall into three categories: estimation (the first, and the fifth, of a derived quantity), prediction under a threshold (the second to the fourth, the sixth and the eighth) and hypothesis comparison (the seventh). Each is answered from the posterior of Section 7.2: its engineering question, the general form of its category, the form it takes for the motors, the quantities of the posterior it uses, and the motors' numbers. The questions name the mission of 500 h and the threshold 0.5 on its survival probability; the prior $\mathcal{N}(0, 0.5^2)$ on v of the model comparison, with $\mathcal{N}(0, 1^2)$ as the wider alternative; and the age $t_0 = 500\,\mathrm{h}$ and the further period $\tau = 200\,\mathrm{h}$ of the remaining-life question. Every answer is computed from the posterior of the state of step 3, the MAP state $\hat{\mathbf{z}}$ and the covariance $\boldsymbol{\Sigma}$ on the eight rows, and $(\hat{\mathbf{z}}_\pi, \boldsymbol{\Sigma}_\pi)$ and $(\hat{\mathbf{z}}_9, \boldsymbol{\Sigma}_9)$ on the extended records, and that is the value the chapter reports; the exact posterior through one integral over the shape gives the value set beside it, with the gap.

7.3.1 The scale, the shape and their trade

Engineering question. What are the scale and the shape of the lifetime, and how well are they known together? The scale is the characteristic life, the age by which 63.2 % of the motors fail; the shape says whether the hazard falls, stays constant or rises with age. The engineer wants each with its band, and wants to know how an error in one bears on the other.

General form. An estimation question about the entries of the state asks for their posterior distribution, the probability that an entry lies at or below each value z^*,

$$\Pr(z_i \le z^* \mid \mathbf{y}) = \int \mathbf{1}\big[z_i \le z^*\big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \tag{7.52}$$

summarized by a location, a spread and a central 95 % credible interval whose endpoints are the values at which equation (7.52) equals 0.025 and 0.975, and, for two entries, by their covariance $\mathrm{Cov}(\mathbf{z} \mid \mathbf{y})$. When a parameter is a monotone function of one entry, its quantiles are the images of that entry's quantiles.

Specialized form. The question is equation (7.52) with $z_i = u$ and $z_i = v$, the two entries of the state, with $\eta = \mathrm{e}^u$ and $\beta = \mathrm{e}^v$ their monotone images, and $p(\mathbf{z} \mid \mathbf{y})$ the posterior of step 3, the Laplace normal $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of equations (7.39), (7.49) and (7.50). The marginal of u is then $\mathcal{N}(\hat{u}, \Sigma_{uu})$, read off as the first entry of $\hat{\mathbf{z}}$ and the first diagonal entry of $\boldsymbol{\Sigma}$, the marginal of v is $\mathcal{N}(\hat{v}, \Sigma_{vv})$, their correlation is $\Sigma_{uv}/\sqrt{\Sigma_{uu}\Sigma_{vv}}$, and the intervals of the parameters are those of the entries mapped through the exponential:

$$\eta \in \big[\mathrm{e}^{\hat{u}-1.96\sqrt{\Sigma_{uu}}},\ \mathrm{e}^{\hat{u}+1.96\sqrt{\Sigma_{uu}}}\big], \qquad \beta \in \big[\mathrm{e}^{\hat{v}-1.96\sqrt{\Sigma_{vv}}},\ \mathrm{e}^{\hat{v}+1.96\sqrt{\Sigma_{vv}}}\big]. \tag{7.53}$$

The correlation is a result of the rows. By the pattern (7.23) every lifetime row acts on both entries, so each row Hessian $\mathbf{\Lambda}_i$ has a nonzero (u, v) entry, and so has their sum $\mathbf{\Lambda}$. For a 2×2 matrix the inverse has the off-diagonal entry

$$\Sigma_{uv} = -\frac{\Lambda_{uv}}{\Lambda_{uu}\Lambda_{vv} - \Lambda_{uv}^2},$$

of the opposite sign to Λ_{uv}: where the rows' curvature couples the two entries, the posterior trades them. The eigenvectors of $\mathbf{\Sigma} = \sum_{j=1}^{2} \epsilon_j \, \mathbf{e}_j \mathbf{e}_j^{\mathsf{T}}$ say which combinations the record determines: the one with the smaller eigenvalue ϵ_j is the best-determined combination of scale and shape, with standard deviation $\sqrt{\epsilon_j}$, and when it lies near a diagonal of the (u, v) plane the record determines that combination better than either entry alone.

The values beside them come from the exact posterior. The marginal of v is the one integral of Section 7.2.2, from which its mode, mean, spread and the interval of β follow by quadrature. Given the shape, the scale enters through a gamma variable whose logarithm has a closed-form mean and variance; mixing those over the marginal of v gives the mean and the variance of u, its covariance with v, and the quantiles of η (Section 7.2.2). The joint exact posterior is not normal: given a small shape the scale is long and spread, given a large one short and narrow.

Building blocks. The MAP state $\hat{\mathbf{z}} = (6.64061,\ 0.27887)^{\mathsf{T}}$ and the covariance $\mathbf{\Sigma}$ of equations (7.49) and (7.50), with the intervals (7.53); and, for the exact values beside them, the marginal (7.46) of v and the moments of u given v, equation (7.45).

The scale and the shape. The scale is $\hat{\eta} = \mathrm{e}^{6.64061} = 765.6\,\mathrm{h}$ at the MAP state, the Laplace median, and with $\sqrt{\Sigma_{uu}} = 0.3465$, equation (7.53) gives $\eta \in 765.6\,\mathrm{e}^{\pm 1.96 \times 0.3465} = [388,\ 1510]\,\mathrm{h}$ with 95 % posterior probability. The shape is $\hat{\beta} = \mathrm{e}^{0.27887} = 1.3216$, and with $\sqrt{\Sigma_{vv}} = 0.3910$, $\beta \in 1.3216\,\mathrm{e}^{\pm 1.96 \times 0.3910} = [0.61,\ 2.84]$. The correlation of the two entries is $-0.02909/\sqrt{0.12004 \times 0.15287} = -0.215$: a shorter scale goes with a larger shape.

The values beside them come from the exact posterior. On the grid in v, the exact marginal (7.46) has its mode at $v = 0.109$ ($\beta = 1.115$), the mean -0.020 and the standard deviation 0.476, and the shape lies in $[0.34,\ 2.19]$ with posterior probability 95 %. With $\psi(5) = 1.50612$ and $\psi'(5) = 0.22132$, equation (7.45) averaged over that marginal gives $\mathbb{E}[u \mid \mathbf{y}] = 6.911$ and $\mathrm{Var}(u \mid \mathbf{y}) = 0.412 + 0.133 = 0.545$, the average conditional variance and the variance of the conditional mean, a standard deviation of 0.738; the scale has the posterior median 875 h and lies in $[390,\ 5840]\,\mathrm{h}$ with probability 95 %, a band with a long upper tail. The covariance of the two entries is -0.148, a correlation of -0.422. The gaps are large: the scale's median is 109 h lower on the Laplace route and its upper end 4330 h lower, the shape's center 0.207 higher and its band shifted up, by 0.27 at the lower end and 0.65 at the upper, and the correlation half as strong. The normal in (u, v) cannot carry the long tail of the scale at small shapes, for the reason Section 7.2.2 gives: the scale given the shape is a skewed gamma, and integrating it out weights each shape by $1/\beta$.

7.3.2 The mission reliability and the median lifetime

Engineering questions. Will a motor survive 500 h? What is the probability of failure by 500 h? Is the 500 h survival probability below one half? What is the median lifetime? The first two concern one new motor and a mission of 500 h from installation; the third asks about the population the motors come from, whether its probability of surviving 500 h is below one half; the fourth asks for the age by which half the population fails.

General form. A question about a quantity $f(\mathbf{z})$ derived from the state asks for its distribution under the posterior: its expectation, the probability that it lies below a threshold, and its

quantiles,

$$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}] = \int f(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \qquad \Pr(f(\mathbf{z}) < f^* \mid \mathbf{y}) = \int \mathbf{1}\big[f(\mathbf{z}) < f^*\big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \quad (7.54)$$

with the quantile s_q the value at which $\Pr(f(\mathbf{z}) \le s_q \mid \mathbf{y}) = q$. When $f(\mathbf{z}) = \Pr(A \mid \mathbf{z})$ is the probability of a future event given the state, its expectation is the posterior predictive probability of the event, $\int p(y_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$ with y_{new} the indicator of A: the general form of a prediction.

Specialized form. One choice enters here: a new motor is taken to be one of the population the eight came from, run under the same conditions, so that its survival depends on the state alone. It fails if the new motors come from another lot or another duty (Section 7.6.5).

The mission reliability. The question is equation (7.54) with $f(\mathbf{z}) = R(500; \mathbf{z})$, the probability that a motor whose lifetime distribution has the state $\mathbf{z}$ survives 500 h,

$$R(t; \mathbf{z}) = \exp\left[- \left(\frac{t}{\mathrm{e}^{u}} \right)^{\mathrm{e}^{v}} \right],$$

a derived quantity of the state, not an entry of it, and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of step 3. The probability that a new motor survives beyond age t is its expectation, by the law of total probability over the state,

$$\Pr(T_{\text{new}} > t \mid \mathbf{y}) = \int \Pr(T_{\text{new}} > t \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z} = \int \exp\left[- \left(\frac{t}{\mathrm{e}^{u}} \right)^{\mathrm{e}^{v}} \right] p(u, v \mid \mathbf{y})\, \mathrm{d}u\, \mathrm{d}v, \quad (7.55)$$

the posterior predictive survival, and the probability of failure by t is its complement. Because the survival function is nonlinear in the parameters, it is not the survival function at the MAP state: the distinction between estimating a parameter and predicting an outcome of Chapters 5 and 6. Under the Laplace normal the integral has two nested normal directions: the marginal of v, and u given v, normal with mean $\hat{u} + (\Sigma_{uv}/\Sigma_{vv})(v - \hat{v})$ and variance $\Sigma_{uu} - \Sigma_{uv}^2/\Sigma_{vv}$. A Gauss–Hermite rule of Q nodes t_q and weights w_q in each direction evaluates it as a weighted sum,

$$\mathbb{E}[h(\mathbf{z}) \mid \mathbf{y}] \approx \sum_{q=1}^{Q} \frac{w_q}{\sqrt{\pi}} \sum_{r=1}^{Q} \frac{w_r}{\sqrt{\pi}}\, h\big(m_q + \sqrt{2}\, s\, t_r,\ v_q\big), \qquad v_q = \hat{v} + \sqrt{2\Sigma_{vv}}\, t_q, \quad (7.56)$$

with $m_q = \hat{u} + (\Sigma_{uv}/\Sigma_{vv})(v_q - \hat{v})$ and $s^2 = \Sigma_{uu} - \Sigma_{uv}^2/\Sigma_{vv}$, here with $h(\mathbf{z}) = R(t; \mathbf{z})$. The third question is the tail of equation (7.54) of the same f at $f^* = 0.5$, a statement about the population, not about one motor. The event $R(500; \mathbf{z}) < R^*$ is the event $u < \ln 500 - \ln(-\ln R^*)/\beta$, so under the Laplace normal the tail is the normal probability of that event given v, averaged over the normal marginal of v:

$$\Pr\big(R(500; \mathbf{z}) < R^* \mid \mathbf{y}\big) \approx \int \mathcal{N}(v;\ \hat{v},\ \Sigma_{vv})\, \Phi\left(\frac{\ln 500 - \ln(-\ln R^*)\, \mathrm{e}^{-v} - m(v)}{s} \right) \mathrm{d}v, \quad (7.57)$$

with $m(v) = \hat{u} + (\Sigma_{uv}/\Sigma_{vv})(v - \hat{v})$ and Φ the standard normal cumulative distribution, and the band of $R(500; \mathbf{z})$ is the pair of values R^* at which it equals 0.025 and 0.975. The values beside them come from the exact posterior: both are integrals over the shape of closed forms in the scale, equations (7.47) and (7.48).

The linearization is the cheaper reading of the same posterior, and the preposterior of step 7 uses it. It buys a band from one gradient and one covariance, and it fails when the posterior is wide and skewed and the quantity is bounded, as $R(500; \mathbf{z})$ is. The delta method of Chapter 4 linearizes f at $\hat{\mathbf{z}}$,

$$\mu_R \approx R(t; \hat{\mathbf{z}}), \qquad \sigma_R^2 = \mathrm{Var}(R(t; \mathbf{z}) \mid \mathbf{y}) \approx \mathbf{g}^{\mathsf{T}} \boldsymbol{\Sigma}\, \mathbf{g}, \qquad \mathbf{g} = \nabla_{\mathbf{z}} R(t; \mathbf{z})\big|_{\hat{\mathbf{z}}}. \quad (7.58)$$

The gradient is genuinely two-dimensional. Write $\Lambda = (t/\eta)^\beta$, so that $R = \mathrm{e}^{-\Lambda}$ and $\ln\Lambda = \beta(\ln t - u)$. Since $\partial\beta/\partial v = \beta$, the derivatives of the cumulative hazard are

$$\frac{\partial\Lambda}{\partial u} = -\beta\,\Lambda, \qquad \frac{\partial\Lambda}{\partial v} = \beta\,\Lambda \ln\frac{t}{\eta},$$

and since $\partial R/\partial\mathbf{z} = -R\,\partial\Lambda/\partial\mathbf{z}$,

$$\frac{\partial R}{\partial u} = R\,\beta\,\Lambda, \qquad \frac{\partial R}{\partial v} = -R\,\beta\,\Lambda \ln\frac{t}{\eta}, \qquad \mathbf{g} = \begin{bmatrix} g_u \\ g_v \end{bmatrix} = R\,\beta\,\Lambda \begin{bmatrix} 1 \\ -\ln(t/\eta) \end{bmatrix}, \tag{7.59}$$

with $g_u = \partial R/\partial u$ and $g_v = \partial R/\partial v$ the entries of $\mathbf{g}$. Each entry says what one entry of the state carries into the mission reliability. A larger scale postpones failures and raises the survival probability, so g_u is positive at every mission. The sign of g_v depends on the mission: before the characteristic life, $\ln(t/\eta) < 0$, a larger shape lowers the cumulative hazard and raises survival; after it, the reverse. At $t = \eta$ the second entry vanishes, the property of the characteristic life found in Section 7.1.3, and the band of the mission reliability there comes from the scale alone.

Expanding $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$ as in Chapter 4 gives the variance budget, one term per entry and the cross term from their covariance,

$$\sigma_R^2 = g_u^2\,\Sigma_{uu} + g_v^2\,\Sigma_{vv} + 2\,g_u g_v\,\Sigma_{uv}. \tag{7.60}$$

Before the characteristic life both entries of $\mathbf{g}$ are positive, and a negative Σ_{uv} makes the cross term negative: a draw with a shorter scale tends to come with a larger shape, the two errors move the survival probability in opposite directions, and part of the two diagonal terms cancels. The band of the mission reliability is thus set by the angle between $\mathbf{g}$ and the eigenvectors of $\boldsymbol{\Sigma}$, not by the two marginal spreads alone. The first-order band is $R(t;\hat{\mathbf{z}}) \pm 1.96\,\sigma_R$ and the delta tail of the third question $\Phi\big((0.5 - R(500;\hat{\mathbf{z}}))/\sigma_R\big)$; draws of $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ pushed through $R(500;\mathbf{z})$ check them against the tail (7.57). A second linearization works on $\ln\Lambda(t;\mathbf{z}) = \beta(\ln t - u)$, unbounded and closer to linear in the state, with gradient $(-\beta,\ \beta\ln(t/\eta))^\mathsf{T}$; its band maps through $R = \exp(-\mathrm{e}^{\ln\Lambda})$ inside $(0, 1)$.

The median lifetime. The question is the quantile form of equation (7.54) with $f(\mathbf{z}) = t_{50}(\mathbf{z})$, the median (7.18) of the population as a function of the state, at $q = 0.5$, 0.025 and 0.975, and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal of step 3. Its logarithm is simple in the state,

$$\ln t_{50}(\mathbf{z}) = u + \frac{\ln\ln 2}{\beta} = u + \mathrm{e}^{-v}\ln\ln 2, \qquad \mathbf{g} = \frac{\partial \ln t_{50}}{\partial\mathbf{z}} = \begin{bmatrix} 1 \\ -\mathrm{e}^{-v}\ln\ln 2 \end{bmatrix}. \tag{7.61}$$

The scale enters one for one. Since $\ln 2 < 1$, $\ln\ln 2$ is negative, the second entry of $\mathbf{g}$ is positive, and a larger shape lengthens the median: the factor $(\ln 2)^{1/\beta}$ rises toward one as the lifetimes gather about the scale. Given v, $\ln t_{50}$ is normal under the Laplace posterior, so $\Pr(t_{50}(\mathbf{z}) \le s \mid \mathbf{y})$ is the normal probability of $u \le \ln s - \mathrm{e}^{-v}\ln\ln 2$ averaged over the marginal of v, as in equation (7.57), and its quantiles are the answer. The delta-method variance is $\Sigma_{uu} + 2g_v\Sigma_{uv} + g_v^2\Sigma_{vv}$, whose cross term is again negative, and the band $\ln t_{50}(\hat{\mathbf{z}}) \pm 1.96\sqrt{\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}}$, mapped through the exponential, stays positive and reflects the multiplicative way in which the scale enters. The value beside it comes from the exact posterior: the distribution of the median is one more integral over the shape (Section 7.2.2).

Building blocks. The MAP state $\hat{\mathbf{z}} = (6.64061,\ 0.27887)^\mathsf{T}$ and the covariance $\boldsymbol{\Sigma}$ of equations (7.49) and (7.50); the rule (7.56) with the five nodes of Table 7.6; the tail (7.57) with Φ; the gradients (7.59) and (7.61); and, for the exact values beside them, the integrals (7.47) and (7.48) of Section 7.2.2.

The mission reliability and the median. Table 7.6 evaluates the rule (7.56) with five nodes in each direction. Each row is a node of v: the first two columns give the node t_q and its weight $w_q/\sqrt{\pi}$, the next two the log-shape v_q and the shape $\beta_q = \mathrm{e}^{v_q}$ at which the sum looks, the fifth the center m_q of u given that shape, and the last two the inner sum over u, the survival probability averaged at that shape, at 200 and at 500 h. The spread of u given v is $s = \sqrt{0.12004 - 0.02909^2/0.15287} = 0.3384$. The middle row is the shape of the MAP state, and the others reach about two standard deviations either side of it; a small shape comes with a long scale, which is the trade of $\boldsymbol{\Sigma}$.

Table 7.6: The Gauss–Hermite rule (7.56) with five nodes in each direction over the Laplace posterior $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of the motors: one row per node of v, with the inner five-node sum over u given v in the last two columns.

Node t_q	Weight $w_q/\sqrt{\pi}$	v_q	β_q	m_q	$R(200)$ at v_q	$R(500)$ at v_q
−2.0202	0.01126	−0.8382	0.4325	6.8532	0.59865	0.46740
−0.9586	0.22208	−0.2512	0.7779	6.7415	0.71683	0.51119
0	0.53333	0.2789	1.3216	6.6406	0.83213	0.55378
0.9586	0.22208	0.8089	2.2454	6.5397	0.92337	0.58478
2.0202	0.01126	1.3959	4.0386	6.4280	0.97525	0.58364
Weighted sum					0.8258	0.5506

The weighted sum of the last column is

$$\Pr(T_{\text{new}} > 500 \mid \mathbf{y}) \approx 0.01126 \times 0.46740 + 0.22208 \times 0.51119 + \cdots + 0.01126 \times 0.58364 = 0.5506,$$

about 55.1 %, against the plug-in $R(500; \hat{\mathbf{z}}) = 0.5658$ below; forty nodes in each direction change no digit shown. A new motor fails by 500 h with probability $1 - 0.5506 = 0.449$, and at 200 h the predictive is 0.826. The value beside it comes from the exact posterior: by equation (7.47) a new motor survives 500 h with probability 0.574 (0.5743), and fails by 500 h with probability 0.426; at 200 h the predictive is 0.793. The gap is 0.023 at 500 h, with the Laplace route lower, and 0.033 at 200 h, with it higher.

By equation (7.57) at $R^* = 0.5$, with the integral over v taken by the same rule at forty nodes,

$$\Pr(R(500; \mathbf{z}) < 0.5 \mid \mathbf{y}) \approx 0.355,$$

and the population's 500 h survival probability, with the standard deviation 0.153 from the sum over $R(500; \mathbf{z})^2$, lies in [0.222, 0.840] with probability 95 %. By equation (7.48) the exact posterior gives the standard deviation 0.140, the band [0.294, 0.831] and the tail 0.301, so the Laplace tail is 0.054 higher and the lower end of its band 0.072 lower: the normal in (u, v) carries a longer left tail than the exact posterior does.

The median lifetime of the population has, from the same integral over v, the posterior median 568 h and the 95 % band [289, 1109] h; the delta method on the logarithm gives 580 h and [298, 1130] h. The exact posterior gives the median 622 h and the band [222, 2721] h: the Laplace median is 54 h lower and its upper end 1612 h lower, the long tail of the scale at small shapes that the normal does not carry.

The gradient of the mission reliability. The gradient (7.59) and the budget (7.60) are objects of the fit, used again by the preposterior of step 7. At the MAP state

$$R(500; \hat{\mathbf{z}}) = \exp\left[-\left(\frac{500}{765.6}\right)^{1.3216}\right] = \mathrm{e}^{-0.5695} = 0.5658, \tag{7.62}$$

with $\Lambda(500) = 0.56949$, $R = 0.56582$ and $\ln(500/765.56) = -0.42600$. The common factor of equation (7.59) is $R\beta\Lambda = 0.56582 \times 1.32164 \times 0.56949 = 0.42586$, and

$$\mathbf{g} = \begin{bmatrix} 0.42586 \\ 0.42586 \times 0.42600 \end{bmatrix} = \begin{bmatrix} 0.42586 \\ 0.18142 \end{bmatrix},$$

within 7° of the eigenvector $(0.864,\ 0.504)^{\mathsf{T}}$: a change of 0.1 in u, a 10.5 % change in the scale, moves the 500 h survival probability by about 0.043. Table 7.7 evaluates the budget. Each row is a term; the columns give its expression, its value and its share of the total. One row of the arithmetic, the scale term: $0.42586^2 \times 0.12004 = 0.18136 \times 0.12004 = 0.021770$.

Table 7.7: The variance budget of the 500 h survival probability at the MAP state, $R(500; \hat{\mathbf{z}}) = 0.5658$, from the gradient (7.59) and the covariance (7.50).

Term	Expression	Value [–]	Share [–]
Scale, $u = \ln \eta$	$g_u^2\, \Sigma_{uu}$	0.021770	97.6 %
Shape, $v = \ln \beta$	$g_v^2\, \Sigma_{vv}$	0.005031	22.6 %
Cross term	$2\, g_u g_v\, \Sigma_{uv}$	−0.004495	−20.2 %
Total	$\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$	0.022306	100 %
Standard deviation	σ_R	0.1494	

The scale dominates the budget, since the lever of the shape, $\ln(t/\eta) = -0.426$, is small at 500 h; the cross term removes about a sixth of the sum of the diagonal terms.

7.3.3 The hazard trend and the exponential model

Engineering questions. Does the risk of failure increase with age? Which lifetime model describes the record? The first asks whether the motors wear out, which decides whether replacing them by age makes sense; the second whether the record needs the Weibull model's second entry at all, or whether the exponential lifetime of Chapter 6, with its constant hazard, describes it.

General forms. A question whether an entry exceeds a value is a tail of the posterior,

$$\Pr(z_i > z^* \mid \mathbf{y}) = \int \mathbf{1}\big[z_i > z^*\big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \tag{7.63}$$

read inside one model. A question which of two models describes the record is a hypothesis comparison, and its answering object is the ratio of the evidences of the two complete models, each the likelihood integrated over that model's prior,

$$B_{01} = \frac{p(\mathbf{y} \mid M_0)}{p(\mathbf{y} \mid M_1)}, \qquad p(\mathbf{y} \mid M) = \int p(\mathbf{y} \mid \mathbf{z}, M)\, p(\mathbf{z} \mid M)\, \mathrm{d}\mathbf{z}. \tag{7.64}$$

Specialized form. Three choices enter here. The hazard trend is read inside the Weibull family: a shape above one means a rising hazard only if the family describes the motors, and a fitted rise could equally come from a change of stress during the study or a mixture of lots whose weaker members fail first (Section 7.6.5). The exponential model is reached as the limit of a prior row on the log-shape rather than substituted: this keeps one family and one fit for both models, and Section 7.4.1 checks that the limit arrives. The log-shape of the comparison is given the proper prior $\mathcal{N}(0, 0.5^2)$: this is what makes the evidence ratio defined, and the ratio depends on its width, which the wider prior $\mathcal{N}(0, 1^2)$ of Section 7.3.3 measures.

The hazard trend. The Weibull hazard (7.15) rises with age when $\beta > 1$, that is when $v > 0$. The question is equation (7.63) with $z_i = v$, $z^* = 0$ and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal of step 3, whose marginal of v gives, with $f(\mathbf{z}) = v$, $\mathbf{g} = (0,\ 1)^\mathsf{T}$ and $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g} = \Sigma_{vv}$,

$$\Pr(\beta > 1 \mid \mathbf{y}) \approx \Phi\left(\frac{\hat{v}}{\sqrt{\Sigma_{vv}}}\right). \tag{7.65}$$

The value beside it comes from the exact posterior: the mass of the marginal of v to the right of zero (Section 7.2.2).

The exponential model as a row. The comparison is equation (7.64) with M_0 the exponential model, M_1 the Weibull model, $\mathbf{z} = (u, v)^\mathsf{T}$ and each evidence computed from its own posterior of the state. The exponential model M_0 is the Weibull model with the prior row v/σ_0 on the log-shape, of Jacobian row $(0,\ 1/\sigma_0)$, which adds $\mathrm{diag}(0,\ 1/\sigma_0^2)$ to $\boldsymbol{\Lambda}$, in the limit $\sigma_0 \to 0$: the exact constraint $v = 0$ as the limit of a row, approached in Section 7.4.1. In the limit v carries no freedom, $f(t \mid \lambda) = \lambda\,\mathrm{e}^{-\lambda t}$ and $R(t \mid \lambda) = \mathrm{e}^{-\lambda t}$ with $\lambda = \mathrm{e}^{-u}$, and the row terms (7.22) multiply to $\prod_i \lambda^{\delta_i}\,\mathrm{e}^{-\lambda t_i}$, that is, with $k = \sum_i \delta_i$ and $T_{\mathrm{tot}} = \sum_i t_i$,

$$p(\mathbf{y} \mid \lambda) = \lambda^k\,\mathrm{e}^{-\lambda T_{\mathrm{tot}}}. \tag{7.66}$$

The parameter-dependent part has the structure of the Poisson likelihood of Chapter 6, a count k in an exposure T_{tot}, although the observation process is conceptually different: here each unit contributes one lifetime, and the survivors contribute exposure without events. With the prior flat in u the objective is $\mathcal{L} = ku + T_{\mathrm{tot}}\,\mathrm{e}^{-u}$, each failure pulling u down and each survivor pushing it up, with curvature $T_{\mathrm{tot}}\,\mathrm{e}^{-u}$; the gradient $k - T_{\mathrm{tot}}\,\mathrm{e}^{-u}$ vanishes where

$$\hat{\lambda} = \mathrm{e}^{-\hat{u}} = \frac{k}{T_{\mathrm{tot}}}, \qquad \Sigma_{uu} = \frac{1}{k}, \tag{7.67}$$

so that the precision of the log-rate is the number of failures, the survivors contributing exposure but no precision; its maximized log-likelihood is $\ell_0 = k\ln(k/T_{\mathrm{tot}}) - k$.

The Weibull model with a proper prior. The ratio of evidences needs a proper prior on the entry the two models disagree about. Give both models the same flat prior on u, and give the Weibull shape the prior $v \sim \mathcal{N}(0, 0.5^2)$, density $\pi(v)$: a shape between $\mathrm{e}^{-1} = 0.37$ and $\mathrm{e}^{1} = 2.7$ with probability 95 %, centered on the exponential. Write $\Omega(v) = \int p(\mathbf{y} \mid u, v)\,\mathrm{d}u$ for the likelihood with the scale integrated out at fixed shape, a closed form derived in Section 7.2.2. Then $p(\mathbf{y} \mid M_0) = \Omega(0)$ and $p(\mathbf{y} \mid M_1) = \int \Omega(v)\,\pi(v)\,\mathrm{d}v$, the improper constant of the flat prior on u is common to both and cancels, and since the posterior of v under M_1 is $p(v \mid \mathbf{y}, M_1) = \Omega(v)\pi(v)/\int \Omega\,\pi$,

$$B_{01} = \frac{\Omega(0)}{\int \Omega(v)\,\pi(v)\,\mathrm{d}v} = \frac{p(v = 0 \mid \mathbf{y}, M_1)}{\pi(0)}, \tag{7.68}$$

the posterior density of the shape at the nested value divided by its prior density there, the Savage–Dickey ratio.

On the state, the prior of M_1 is the Gaussian row $r_\pi(\mathbf{z}) = v/0.5$ of step 3, whose fit gives the MAP state $\hat{\mathbf{z}}_\pi$ of M_1 and its covariance $\boldsymbol{\Sigma}_\pi$ (Section 7.2.3), the Laplace normal of the state under M_1. The tail of the hazard trend reads, under M_1, $\Phi(\hat{v}_\pi/\sqrt{\Sigma_{\pi,vv}})$. The Savage–Dickey ratio reads the same normal of v at the nested value,

$$B_{01} \approx \frac{\mathcal{N}(0;\ \hat{v}_\pi,\ \Sigma_{\pi,vv})}{\mathcal{N}(0;\ 0,\ 0.5^2)}, \tag{7.69}$$

with $\mathcal{N}(x;\ m,\ s^2)$ the normal density at x; reading a Laplace normal at a point is not the exact identity. This is the ratio the chapter reports; the exact one-integral value of equation (7.68)

stands beside it. A second state route approximates each evidence by the Gaussian integral of $\exp(-\mathcal{L})$ about its own MAP state,

$$p(\mathbf{y} \mid M) \approx \mathrm{e}^{-\mathcal{L}(\hat{\mathbf{z}}_M)}\,(2\pi)^{n/2}\,(\det \boldsymbol{\Sigma}_M)^{1/2}, \tag{7.70}$$

with n the number of state entries, 1 under M_0 and 2 under M_1, and $\mathcal{L}$ carrying the normalizing constant of the prior on v. Two classical summaries use the maximized log-likelihoods ℓ_0 and ℓ_1: the likelihood-ratio statistic $2(\ell_1 - \ell_0)$, approximately chi-squared with one degree of freedom under the exponential model, a standard result, and the Akaike information criterion $-2\ell + 2n$. Neither is a probability of either model, and a more flexible model always achieves the better raw maximum, whether or not its extra parameter describes anything.

Building blocks. The entries $\hat{v} = 0.27887$ and $\Sigma_{vv} = 0.15287$ of equations (7.49) and (7.50) and Φ for the tail; the MAP state $\hat{\mathbf{z}}_\pi$ and the covariance $\boldsymbol{\Sigma}_\pi$ of Section 7.2.3 for the ratio (7.69), and $\boldsymbol{\Sigma}_\pi$ with the objective for the evidences (7.70); the sums k and T_{tot} for the maximized log-likelihood of the exponential; and, for the exact values beside them, the marginal $\Omega(v)$ and its integrals (Section 7.2.2).

The shape and the evidence ratio. By equation (7.65),

$$\Pr(\beta > 1 \mid \mathbf{y}) \approx \Phi\!\left(\frac{0.27887}{0.39098}\right) = \Phi(0.713) = 0.762,$$

and under the proper prior $\mathcal{N}(0, 0.5^2)$ of the comparison, $\Phi(0.16856/0.3211) = 0.700$. The values beside them come from the exact posterior: the mass of the exact marginal of v to the right of zero is $\Pr(\beta > 1 \mid \mathbf{y}) = 0.522$, 0.240 lower; under the proper prior it is 0.542, 0.158 lower. The normal is centered at the joint MAP, to the right of where the exact marginal has its mass, because integrating the scale out weights each shape by $1/\beta$. Both routes stay well short of an assurance that the hazard rises.

The ratio (7.69) is

$$B_{01} \approx \frac{\mathcal{N}(0;\ 0.16856,\ 0.3211^2)}{\mathcal{N}(0;\ 0,\ 0.5^2)} = \frac{1.0825}{0.7979} = 1.36,$$

and with the wider prior $\mathcal{N}(0, 1^2)$, from $(6.64810,\ 0.24114)^{\mathsf{T}}$ and $\sqrt{\Sigma_{vv}} = 0.3716$, $B_{01} \approx 2.18$; the Laplace evidences (7.70) give 1.61 and 2.82. The one integral of equation (7.68) gives the value beside it,

$$B_{01} = 1.47, \qquad \ln B_{01} = 0.39 \text{ nats},$$

and with the wider prior $\mathcal{N}(0, 1^2)$, $B_{01} = 2.39$: the gaps are 0.11 and 0.21, and every route favors the exponential model, weakly. The maximized log-likelihoods are $\ell_1 = -38.3316$ for the Weibull model, equation (7.49), and $\ell_0 = 5\ln(5/4110) - 5 = -38.5587$ for the exponential, by equation (7.67). The likelihood-ratio statistic $2\,(-38.3316 + 38.5587) = 0.454$ is exceeded by a chi-squared variable with one degree of freedom with probability 0.50; the Akaike information criterion is $-2 \times (-38.5587) + 2 = 79.12$ for the exponential and $-2 \times (-38.3316) + 4 = 80.66$ for the Weibull, whose raw maximum is higher by 0.23.

7.3.4 The remaining life of a surviving motor

Engineering question. How much longer will a surviving motor last? An engineer inspects a motor of the same population, not part of the record, that has already operated for t_0 hours, and asks for the probability that it survives a further τ hours. The fleet's decision about which motors to replace first turns on how that probability changes with the age already reached.

General form. A prediction given that an event B has already been observed is the predictive of the future event A under the posterior that includes B,

$$\Pr(A \mid B, \mathbf{y}) = \int \Pr(A \mid B, \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y}, B)\, \mathrm{d}\mathbf{z}, \qquad p(\mathbf{z} \mid \mathbf{y}, B) = \frac{\Pr(B \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})}{\int \Pr(B \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}}. \quad (7.71)$$

Specialized form. One choice enters here: the motor inspected is taken to be one of the same population, outside the record, about which nothing is known but the age it has reached. This lets its survival to t_0 enter as one more censored row, and it fails if the motor was chosen for inspection because of a symptom, a selection that carries information the row does not (Section 7.6.5).

The question is equation (7.71) with $B = \{T > t_0\}$, $A = \{T > t_0 + \tau\}$, $\Pr(A \mid B, \mathbf{z}) = R(\tau \mid T > t_0, \mathbf{z})$, the conditional survival (7.11) of a unit whose state is known, and $p(\mathbf{z} \mid \mathbf{y}, B)$ the posterior of the state on the record with the ninth row, the Laplace normal $\mathcal{N}(\hat{\mathbf{z}}_9, \boldsymbol{\Sigma}_9)$ of equation (7.51). The survival of the motor to age t_0 is itself evidence about the state: a motor known to have reached 500 h makes lifetime distributions with short lives less plausible, and the observation $T > t_0$ is the ninth row of the pattern (7.23), a ninth motor with $t_9 = t_0$ and $\delta_9 = 0$, whose term is $\mathcal{L}_9(\mathbf{z}) = \Lambda(t_0 \mid \mathbf{z})$, whose gradient is that of the cumulative hazard found for the mission reliability, and whose curvature is a survivor's row Hessian (Section 7.1.6). The answer is therefore the rule (7.56) over $\mathcal{N}(\hat{\mathbf{z}}_9, \boldsymbol{\Sigma}_9)$ with $h(\mathbf{z}) = R(\tau \mid T > t_0, \mathbf{z})$. The conditional survival is below c when $\mathrm{e}^{-\beta u}\big((t_0 + \tau)^\beta - t_0^\beta\big) > -\ln c$, an event $u < u^*(v)$ as in equation (7.57), so its quantiles under the same normal follow from one integral over v.

The value beside it comes from the exact posterior. With $\Pr(B \mid \mathbf{z}) = R(t_0 \mid \mathbf{z})$ in equation (7.71),

$$\Pr(T > t_0 + \tau \mid T > t_0, \mathbf{y}) = \int \frac{R(t_0 + \tau \mid \mathbf{z})}{R(t_0 \mid \mathbf{z})}\, \frac{R(t_0 \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})}{\int R(t_0 \mid \mathbf{z}')\, p(\mathbf{z}' \mid \mathbf{y})\, \mathrm{d}\mathbf{z}'}\, \mathrm{d}\mathbf{z}.$$

The factor $R(t_0 \mid \mathbf{z})$ cancels, and

$$\Pr(T > t_0 + \tau \mid T > t_0, \mathbf{y}) = \frac{\int R(t_0 + \tau \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}}{\int R(t_0 \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}}, \quad (7.72)$$

the ratio of two posterior predictive survival probabilities: exactly, the ratio of the exact predictives at $t_0 + \tau$ and at t_0 (Section 7.2.2). It averages the survival over the parameter uncertainty and accounts for the information that the motor reached age t_0; it is not in general the result of substituting estimated parameters into equation (7.19). Averaged over the posterior, the predictive lifetime is a mixture of Weibull lifetimes, and a mixture fails its weak members first: its hazard can fall with age even where every member's rises. For a motor that is already in the record, such as motor F at 700 h, its survival to that age is already a censored row of $\mathbf{y}$; the conditioning on $T > t_0$ must then not be applied a second time, and the answer is the posterior average of equation (7.11) itself.

A plug-in value of the conditional survival must be computed at the state refitted with the ninth row, $\hat{\mathbf{z}}_9$, not at $\hat{\mathbf{z}}$. Its delta band follows from the gradient of $R(\tau \mid T > t_0, \mathbf{z})$, obtained from equation (7.11) as equation (7.59) was,

$$\mathbf{g} = R(\tau \mid T > t_0, \mathbf{z})\, \beta \begin{bmatrix} \Lambda(t_0 + \tau) - \Lambda(t_0) \\ -\big[\Lambda(t_0 + \tau) \ln \frac{t_0 + \tau}{\eta} - \Lambda(t_0) \ln \frac{t_0}{\eta}\big] \end{bmatrix}, \quad (7.73)$$

with $\mathbf{g}^{\mathsf{T}} \boldsymbol{\Sigma}_9 \mathbf{g}$ its variance. The function $\Lambda(t) \ln(t/\eta)$ falls to its minimum at $t = \eta\, \mathrm{e}^{-1/\beta}$ and rises after it, so for every t_0 beyond that age the two entries of $\mathbf{g}$ have opposite signs, and with a negative $\Sigma_{9,uv}$ the cross term adds to the budget. The same linearization applied to $f(\mathbf{z}) = \ln[\Lambda(t_0 + \tau) - \Lambda(t_0)]$, whose first gradient entry is $-\beta$, gives a band mapped through $R = \exp(-\mathrm{e}^f)$ that stays inside $(0, 1)$, and draws of $\mathcal{N}(\hat{\mathbf{z}}_9, \boldsymbol{\Sigma}_9)$ remove the linearization.

Building blocks. The MAP state $\hat{\mathbf{z}}_9 = (6.71498,\ 0.30388)^{\mathsf{T}}$ and the covariance $\boldsymbol{\Sigma}_9$ of equation (7.51); the rule (7.56); the one integral over v for the quantiles; the predictive at 200 h of Table 7.6 for the comparison with a new motor; and, for the exact value beside it, the predictives (7.47) and their ratio (7.72).

The remaining life. The rule (7.56) over $\mathcal{N}(\hat{\mathbf{z}}_9, \boldsymbol{\Sigma}_9)$ with forty nodes in each direction gives

$$\Pr(T > 700 \mid T > 500, \mathbf{y}) \approx 0.712, \tag{7.74}$$

with the standard deviation 0.165; the conditional survival itself has the median 0.760 and the 95 % band [0.234, 0.894]. The value beside it comes from the exact posterior. The exact predictives (7.47) are $\Pr(T_{\text{new}} > 700 \mid \mathbf{y}) = 0.4584$ and $\Pr(T_{\text{new}} > 500 \mid \mathbf{y}) = 0.5743$, and by equation (7.72) their ratio is

$$\Pr(T > 700 \mid T > 500, \mathbf{y}) = 0.798, \tag{7.75}$$

slightly more than the 0.793 with which a new motor survives its first 200 h. On the nine-row record the conditional survival itself has the median 0.817 and the 95 % band [0.538, 0.956]. The gap is 0.086, and here it changes the reading. On the reported route the motor found running is a worse risk than a new one, 0.712 against 0.826 over the next 200 h, and on the exact posterior it is no worse, 0.798 against 0.793: the two routes fall on opposite sides of the new motor. The state route errs for the reason the scale's band does: the normal in (u, v) gives too little weight to the small shapes with long scales, the members of the mixture whose hazard falls. Those members lower the exact survival of a new motor, by their early failures, and their survivors hold it up with age.

7.4 The checks (step 5)

The answers rest on a model, the Weibull lifetime with noninformative censoring under a flat prior, and on a Gaussian for the posterior of the state; the delta bands rest in addition on a linearization of each derived quantity. Each is set against a fuller calculation. The motors have no physics row, and the one exact constraint the chapter uses, $v = 0$ of the exponential model, is reached as the limit $\sigma_0 \to 0$ of a prior row.

7.4.1 The fit and the exponential limit

The fit. The Weibull fit is a minimum and the only one reached from either start (Section 7.2.3). Along the curve of best scale, with $\sum_i \delta_i \ln t_i = \ln 120 + \ln 230 + \ln 310 + \ln 450 + \ln 600 = 28.4683$, the shape equation (7.43) has its root at the same $\hat{\beta} = 1.3216$. At $\beta = 1$, $\hat{\eta}(1) = 822.0$ h and $\ell_p(1) = -38.559$, 0.227 below the maximum $\ell_p(1.3216) = -38.332$; at $\beta = 2$, $\sum_i t_i^2 = 2\,665\,900\,\text{h}^2$, $\hat{\eta}(2) = \sqrt{2\,665\,900/5} = 730.2$ h and $\ell_p(2) = 5 \ln 2 - 5 \ln 533\,180 + 28.4683 - 5 = -38.999$, 0.667 below. The marginal precision of v is $1/\Sigma_{vv} = 6.5416$, the Schur complement $6.8579 - 1.6621^2/8.7337 = 6.5416$, and the Laplace parabola has the curvature $6.5416/\hat{\beta}^2 = 3.745$ in the shape.

Figure 7.4 draws the slice. The horizontal axis is the shape, the vertical axis the profile log-likelihood less its maximum; the solid curve is equation (7.42), the dashed curve the parabola of the Laplace Gaussian with curvature 3.745, and the horizontal line at -1.92 marks where the likelihood has fallen by the amount that defines an approximate 95 % likelihood interval. The maximum is shallow: the curve has fallen only 0.23 at $\beta = 1$, and it stays above the line from $\beta = 0.53$ to 2.57. The profile is also skewed, falling faster toward small shapes than toward large ones, which the parabola does not reproduce.

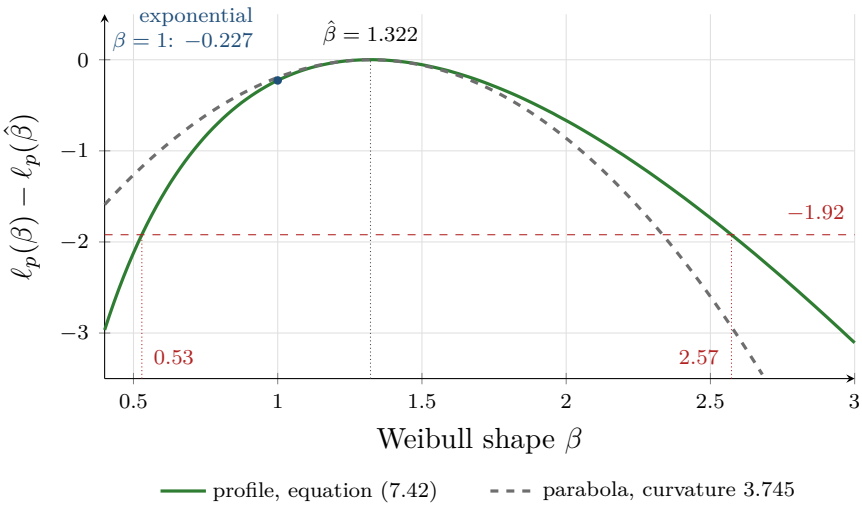


Figure 7.4: The objective of the motor record along its curve of best scale (7.35), as the profile log-likelihood of the Weibull shape, $\ell_p(\beta) = 5\ln\beta - 5\ln(\frac{1}{5}\sum_i t_i^\beta) + 28.4683(\beta - 1) - 5$, less its maximum -38.3316 at $\hat\beta = 1.3216$; the parabola $-\frac{1}{2} \times 3.745\,(\beta - 1.3216)^2$ is the Laplace Gaussian, with $3.745 = 1/(\hat\beta^2\Sigma_{vv})$. The dashed line at -1.92 crosses the profile at 0.53 and 2.57.

The exponential limit. The exponential fit is the Weibull fit with the shape pinned, reached as the limit $\sigma_0 \to 0$ of the prior row v/σ_0 of Section 7.3.3. Table 7.8 follows that limit. Each row is one prior row on the log-shape, from none (the flat prior) to $\sigma_0 = 0.001$, and the last row is the fit of the scale alone, equation (7.67); the columns give the MAP state, the two posterior variances read from $\boldsymbol{\Sigma}$ and the objective at the minimum, the prior term $\frac{1}{2}v^2/\sigma_0^2$ included.

Table 7.8: The Weibull fit to the motor record with the prior row v/σ_0 on the log-shape, from the flat prior to the exact constraint $v = 0$, against the fit of the scale alone, equation (7.67). All entries are dimensionless.

Prior row on v	$\hat u$	$\hat v$	Σ_{uu}	$\sqrt{\Sigma_{vv}}$	$\mathcal{L}(\hat{\mathbf{z}})$
none (flat)	6.64061	0.27887	0.12004	0.3910	38.33156
$\sigma_0 = 1$	6.64810	0.24114	0.12937	0.3716	38.36522
$\sigma_0 = 0.5$	6.66428	0.16856	0.14869	0.3211	38.42643
$\sigma_0 = 0.1$	6.70701	0.01470	0.19517	0.0977	38.54739
$\sigma_0 = 0.01$	6.71169	0.00015	0.19995	0.0100	38.55858
$\sigma_0 = 0.001$	6.71174	0.00000	0.20000	0.0010	38.55870
scale alone, $v = 0$	6.71174	–	0.20000	–	38.55870

As σ_0 falls, the MAP scale moves from $\ln 765.6$ to $\ln 822.0$, the variance of u rises from 0.12004 to $1/k = 0.2$, and the standard deviation of v falls to σ_0 itself: below $\sigma_0 = 0.1$ the prior row, not the record, sets the shape. At $\sigma_0 = 0.001$ every entry agrees with the fit of the scale alone to the printed precision, $\hat\lambda = 5/4110 = 0.0012165\,\mathrm{h}^{-1}$ and $\sqrt{\Sigma_{uu}} = 0.4472 = 1/\sqrt{k}$, and the objective with $-\ell_0 = 38.55870$. The row $\sigma_0 = 0.5$ is the state $\hat{\mathbf{z}}_\pi$ of the model comparison. The variance of u rises because the (u, u) entry of $\boldsymbol{\Lambda}$ is β^2 per failure (equation (7.41)), $1.74673 \times 5 = 8.73$ at $\hat\beta = 1.3216$ and 5 at $\beta = 1$. Under a prior flat in $\ln\lambda$ the exponential model has the exact posterior Gamma$(5, 4110)$ of the rate, and its predictive 500 h survival $(4110/4610)^5 = 0.563$ lies above the plug-in $\mathrm{e}^{-500\times0.0012165} = 0.5443$, with the exact 95 % band $[0.288,\ 0.821]$, because $\mathrm{e}^{-500\lambda}$ is convex in λ.

Table 7.9 compares the two fitted models as plug-in values at their MAP states. Each row is a quantity, each column a model. The two give close survival probabilities at 500 h, 56.6 against 54.4 %, and nearly the same at 600 h, but they separate at longer ages: at 1000 h the Weibull model, whose hazard keeps rising, predicts 24.1 % survival against 29.6 % for the exponential, and at 1500 h 8.8 against 16.1 %: agreement at one mission length does not imply agreement at others when the hazard assumptions differ.

Table 7.9: Plug-in predictions of the two fitted models for the motor record. Exponential: $\hat{\lambda} = 5/4110\,\mathrm{h}^{-1}$; Weibull: $\hat{\eta} = 765.6\,\mathrm{h}$, $\hat{\beta} = 1.3216$.

Quantity	Exponential	Weibull
Maximized log-likelihood ℓ [–]	-38.559	-38.332
$\hat{R}(200)$ [–]	0.784	0.844
$\hat{R}(500)$ [–]	0.544	0.566
$\hat{R}(1000)$ [–]	0.296	0.241
$\hat{R}(1500)$ [–]	0.161	0.088
Median lifetime t_{50} [h]	569.8	580.2
Mean lifetime MTTF [h]	822.0	704.8
Hazard at 500 / 1000 h [$10^{-3}\,\mathrm{h}^{-1}$]	1.22 / 1.22	1.51 / 1.88

7.4.2 The delta method against the draws

The delta method linearizes a bounded, nonlinear function across a posterior that is wide: $\sqrt{\Sigma_{uu}} = 0.35$ changes the scale by a factor of 1.4 in either direction. With $\sigma_R = 0.1494$ from Table 7.7, the first-order band and tail are

$$\begin{aligned} R(500) &\approx 0.5658 \pm 1.96 \times 0.1494 = [0.273,\ 0.859], \\ \Pr(R(500;\mathbf{z}) < 0.5 \mid \mathbf{y}) &\approx \Phi\!\left(\frac{0.5 - 0.5658}{0.1494}\right) = 0.330. \end{aligned} \tag{7.76}$$

The Gauss–Hermite quadrature over the Laplace posterior gives the predictive

$$\Pr(T_{\mathrm{new}} > 500 \mid \mathbf{y}) \approx 0.551, \tag{7.77}$$

(0.5506, with standard deviation 0.1532), below the plug-in 0.566, and 0.826 at 200 h; the one-integral Laplace tail of Section 7.3.2 gives $\Pr(R(500;\mathbf{z}) < 0.5 \mid \mathbf{y}) = 0.355$ and the band [0.222, 0.840]. The delta method on $\ln\Lambda(500) = -0.5630$, with gradient $(-1.3216,\ -0.5630)^{\mathsf{T}}$ and standard deviation 0.4635, gives the band -0.5630 ± 0.9085, which maps to [0.244, 0.795], and the tail 0.336.

Figure 7.5 shows the distribution these summarize, from 400 000 draws of $\mathbf{z}$ from $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ pushed through $R(500;\mathbf{z})$. The horizontal axis is the 500 h survival probability and the vertical axis its density; the bars are the histogram of the draws, the dashed curve the normal of the delta method, $\mathcal{N}(0.5658, 0.1494^2)$, and the vertical lines mark the plug-in value 0.566 and the Laplace mean 0.551. The delta-method normal and the draws agree near the center. The distribution of the draws is skewed to the left, with a long tail toward small survival probabilities produced by draws that combine a small scale with a large shape; the normal is symmetric and misses that tail. The draws give the standard deviation 0.153, the band [0.221, 0.840] and the tail 0.355, the quadrature values within their Monte Carlo errors, about 0.0002 on the mean, 0.0008 on the tail and 0.001 on the quantiles.

Table 7.10 sets the three Laplace-route calculations side by side. Each row is a summary; the columns give it under the delta method on $R(500)$, under the delta method on $\ln\Lambda(500)$

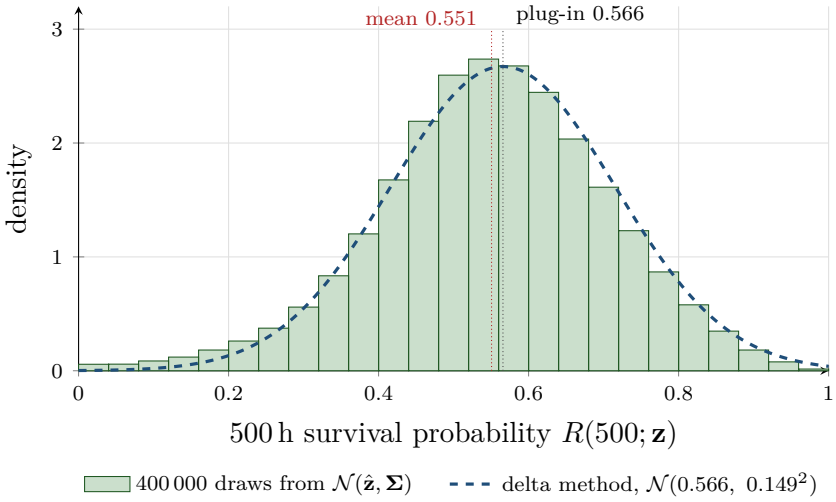


Figure 7.5: The 500 h survival probability of a new motor under the Laplace posterior (7.50): histogram of $R(500;\mathbf{z}) = \exp[-(500/\mathrm{e}^{u})^{\mathrm{e}^{v}}]$ over 400 000 draws of $\mathbf{z}$, against the delta-method normal with mean 0.5658 and standard deviation 0.1494.

Table 7.10: The 500 h survival probability of a new motor by the Laplace route: the delta method on $R(500)$, the delta method on $\ln\Lambda(500)$ mapped back, and the Laplace posterior by Gauss–Hermite and one-integral quadrature. All entries are probabilities [–].

Summary	Delta on R	Delta on $\ln\Lambda$	Laplace posterior
Center	0.566	0.566 (median)	0.551 (mean)
Standard deviation	0.149	–	0.153
95 % band	[0.273, 0.859]	[0.244, 0.795]	[0.222, 0.840]
$\Pr(R(500;\mathbf{z}) < 0.5 \mid \mathbf{y})$	0.330	0.336	0.355

mapped back, and under the Laplace posterior itself, by quadrature. The last row is the third question of Table 7.3.

The delta method gets the spread right, 0.149 against 0.153, but places the band too high at both ends, since it cannot represent the left skew; the transformed delta method moves the lower end toward the Laplace posterior but places the upper end too low. Both linearizations approximate the Laplace posterior, which in turn carries whatever error the Gaussian in (u, v) itself carries.

The median. At $\hat{\mathbf{z}}$, $t_{50} = 765.6\,(\ln 2)^{1/1.3216} = 765.6 \times 0.75781 = 580.2\,\text{h}$, and the gradient (7.61), with $\ln \ln 2 = -0.36651$, is $\mathbf{g} = (1,\ 0.36651/1.32164)^{\mathsf{T}} = (1,\ 0.2773)^{\mathsf{T}}$, so that

$$\sigma^2_{\ln t_{50}} = \mathbf{g}^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g} = \Sigma_{uu} + 2 \times 0.2773\, \Sigma_{uv} + 0.2773^2\, \Sigma_{vv} = 0.12004 - 0.01613 + 0.01176 = 0.11566,$$

a standard deviation of 0.340, and the band $580.2\,\text{e}^{\pm 1.96 \times 0.340} = [298,\ 1130]\,\text{h}$.

7.4.3 The Laplace posterior against the exact posterior

The scale and the shape. Read off $\boldsymbol{\Sigma}$ of equation (7.50), the Laplace posterior has $\sqrt{\Sigma_{uu}} = 0.3465$, $\sqrt{\Sigma_{vv}} = 0.3910$ and the correlation -0.215, and by equation (7.53) $\eta \in 765.6\,\text{e}^{\pm 1.96 \times 0.3465} = [388,\ 1510]\,\text{h}$ and $\beta \in 1.3216\,\text{e}^{\pm 1.96 \times 0.3910} = [0.61,\ 2.84]$. The exact posterior of Section 7.3.1 has a scale more than twice as uncertain, 0.738 against 0.346 in u, whose band reaches 5840 h against 1510 h, and a trade twice as strong, a correlation of -0.422 against -0.215. Both differences come from the shapes the Gaussian underweights: given a small shape the conditional mean $\ln A(\beta)/\beta$ of u is large and its conditional spread $\sqrt{\psi'(5)}/\beta$ wide, so the long upper tail of the scale is the region of small shapes. For the motors the Laplace step over u of Section 7.2.2 differs from $\ln \Omega(v)$ by the constant 0.01664 at every shape, and the weight $1/\beta$ moves the mode of v from the joint MAP $\hat{v} = 0.279$ ($\beta = 1.322$) to $v = 0.109$ ($\beta = 1.115$).

Figure 7.6 compares the exact marginal of v with the Laplace normal $\mathcal{N}(0.279,\ 0.391^2)$. The horizontal axis is $v = \ln \beta$ and the vertical axis the posterior density; the dotted line at $v = 0$ marks the exponential model, $\beta = 1$. The exact marginal is wider, has a long tail toward small shapes, and peaks to the left of the joint MAP; the area to the right of $v = 0$ is 0.522 under it and 0.762, by equation (7.65), $\Phi(0.27887/0.39098) = \Phi(0.713)$, under the normal.

The answers side by side. Table 7.11 sets the two posteriors side by side for every question of the chapter. Each row is a quantity; the middle column gives it under the Laplace approximation in (u, v), by quadrature over $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ unless marked, and the right column under the exact posterior. For the mission reliability the exact distribution is narrower than the Laplace one, by a shorter left tail, and its center lies higher, 0.574 against 0.551.

The center of the shape. With five failures the Laplace approximation in (u, v) is crude for the shape, the scale and the tails, and the errors have two sources. The error in the shape is mainly one of center: the joint MAP is not the marginal mode, because integrating the scale out weights each shape by $1/\beta$ (Section 7.2.2), which shifts the mass of v toward small shapes. A Laplace approximation in v alone, with u integrated exactly so that the weight is kept, centers at 0.109 with standard deviation 0.435 and puts $\Pr(\beta > 1 \mid \mathbf{y}) = 0.599$; the joint Laplace gives 0.762 and the exact marginal 0.522. The weight thus accounts for 0.163 of the gap of 0.240, and the left skew of the exact marginal, -0.61, for the rest. The second source is the gamma distribution (7.44) of the scale given the shape, with shape $k = 5$, skewed and far from normal in its logarithm. That skew integrates to a constant over ζ and leaves the marginal of v untouched, but it shapes every quantity that depends on the scale: the band of η, the tails of $R(500; \mathbf{z})$ and the upper end of the median. More failures shrink both effects, since the gamma distribution approaches a normal as k grows and the factor $1/\beta$ becomes negligible

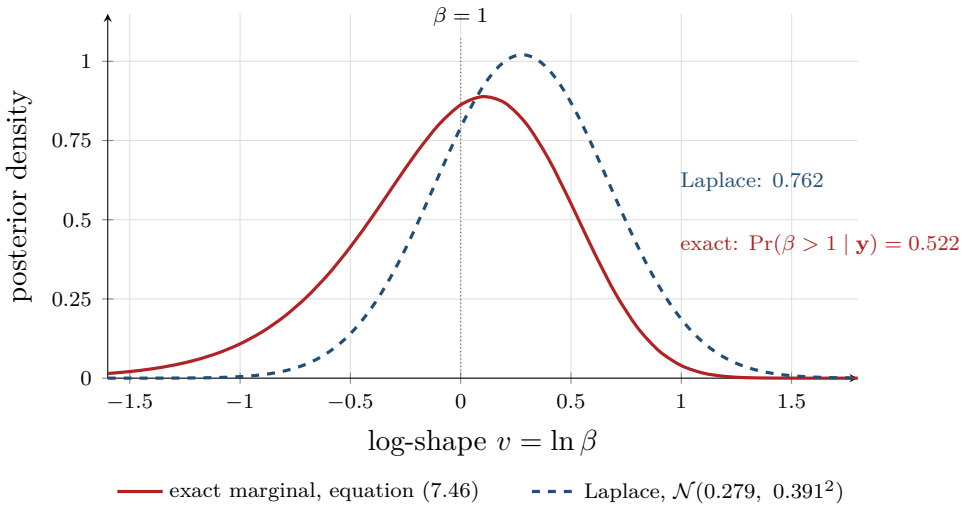


Figure 7.6: The posterior of the log-shape $v = \ln \beta$ under the flat prior in (u, v): the exact marginal $\Omega(v)/\int \Omega$ of equation (7.46), proportional to $\beta^4 e^{28.4683(\beta-1)}(\sum_i t_i^\beta)^{-5}$ with mode 0.109, mean -0.020 and standard deviation 0.476, against the normal $\mathcal{N}(0.2789,\ 0.3910^2)$ of the Laplace approximation (7.50).

Table 7.11: The answers of the chapter under the Laplace approximation in (u, v) and under the exact posterior, both with the flat prior in (u, v); the remaining life is that of a motor outside the record found running at 500 h.

Quantity	Laplace in (u, v)	Exact
Mode of v (shape β) [–]	0.279 (1.322)	0.109 (1.115)
Mean and standard deviation of v [–]	0.279, 0.391	-0.020, 0.476
95 % band on β [–]	[0.61, 2.84]	[0.34, 2.19]
Scale η, center and 95 % band [h]	765.6 (MAP), [388, 1510]	875 (median), [390, 5840]
Standard deviation of u; correlation of u, v [–]	0.346; -0.215	0.738; -0.422
$\Pr(\beta > 1 \mid \mathbf{y})$ [–]	0.762	0.522
$\Pr(T_{\text{new}} > 200 \mid \mathbf{y})$ [–]	0.826	0.793
$\Pr(T_{\text{new}} > 500 \mid \mathbf{y})$ [–]	0.551	0.574
95 % band on $R(500; \mathbf{z})$ [–]	[0.222, 0.840]	[0.294, 0.831]
$\Pr(R(500; \mathbf{z}) < 0.5 \mid \mathbf{y})$ [–]	0.355	0.301
Median lifetime t_{50}, center and band [h]	568, [289, 1109]; delta 580, [298, 1130]	622, [222, 2721]
$\Pr(T > 700 \mid T > 500, \mathbf{y})$ [–]	0.712 (mean over $\mathcal{N}(\hat{\mathbf{z}}_9, \boldsymbol{\Sigma}_9)$), 0.746 (refit plug-in)	0.798

against a sharper slice; with a handful of failures the gaps are those of Table 7.11, and they are reported with the answers.

7.4.4 The evidence ratio against the exact integral

The Laplace ratios. With $\hat{\mathbf{z}}_\pi$ and $\boldsymbol{\Sigma}_\pi$ of Section 7.2.3, $\sqrt{\Sigma_{\pi,vv}} = 0.3211$, and the Savage–Dickey ratio (7.68) read on the Laplace normal of v, the value step 4 reports, is

$$B_{01} \approx \frac{\mathcal{N}(0;\ 0.16856,\ 0.3211^2)}{\mathcal{N}(0;\ 0,\ 0.5^2)} = \frac{1.0825}{0.7979} = 1.36,$$

with $\Pr(\beta > 1 \mid \mathbf{y}, M_1) = \Phi(0.16856/0.3211) = 0.700$. For the wider prior $\mathcal{N}(0, 1^2)$ the same steps give $B_{01} \approx 2.18$, and the Laplace evidences (7.70) give 1.61 and 2.82. The two Laplace routes bracket the exact 1.47 and 2.39 of Section 7.3.3, and both say, as the exact ratio does, that the second parameter is not needed to describe eight motors. Leaving the prior row out and reading the flat-prior normal of equation (7.50) at $v = 0$ would give 0.99 and 1.98: without the row the Gaussian misses the pull of the prior on the posterior.

The shape under the prior. With the prior $\mathcal{N}(0, 0.5^2)$ the exact posterior probability of an increasing hazard is 0.542, against the Laplace 0.700. The spread is not the cause: the exact marginal of v has standard deviation 0.336 against $\sqrt{\Sigma_{\pi,vv}} = 0.321$, and its skewness, -0.23, is small. The center is: the joint MAP $\hat{v} = 0.169$ lies to the right of the exact marginal's mode 0.061, median 0.035 and mean 0.022, by the weight $1/\beta$ again. A Laplace approximation in v alone, with u integrated exactly, keeps the weight: it centers at 0.061 and gives $\Pr(\beta > 1 \mid \mathbf{y}, M_1) = 0.573$ and $B_{01} = 1.472$, against the exact 1.4716. The joint Laplace comes closer for the evidence ratio than for the tail because the two respond differently to a shift of the center. The nested value $v = 0$ sits at the flat top of the exact marginal, whose density there is 1.1742, near its peak. A center shifted by $\Delta = 0.169$, 0.525 standard deviations, scales the normal density at $v = 0$ by $\mathrm{e}^{-\Delta^2/2s^2} = 0.87$ relative to its peak, a change of second order in Δ, so that the Laplace density there, 1.0825, is 8 % low; the tail mass beyond $v = 0$ moves by about $\varphi(0)\Delta/s = 0.21$, a change of first order, and it is 29 % high. Here $s = 0.3211$ is the Laplace standard deviation and φ the standard normal density.

7.4.5 The remaining life by the ninth row

The conditional survival. At the refitted state $\hat{\mathbf{z}}_9$ of Section 7.2.3 $\Lambda(500) = 0.5076$ and $\Lambda(700) = 0.8008$, and the plug-in conditional survival is $\mathrm{e}^{-0.2933} = 0.746$, against 0.727 at the eight-motor fit. The age $t_0 = 500$ h lies beyond $\hat{\eta}\,\mathrm{e}^{-1/\hat{\beta}} = 394$ h, and the gradient (7.73), $(0.2964,\ -0.1240)^\mathsf{T}$, has entries of opposite signs; $\sqrt{\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}_9\mathbf{g}} = 0.1275$, and the delta band $0.746 \pm 0.250 = [0.496,\ 0.996]$ almost reaches one. On the logarithm of the cumulative-hazard increment, $\mathbf{g} = (-1.3551,\ 0.5672)^\mathsf{T}$ and $\sqrt{\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}_9\mathbf{g}} = 0.583$ give $[0.399,\ 0.911]$, and 400 000 draws from $\mathcal{N}(\hat{\mathbf{z}}_9, \boldsymbol{\Sigma}_9)$ give $[0.234,\ 0.894]$ with a mean of 0.712, the lower end moving between 0.230 and 0.236 across seeds, the quadrature values of equation (7.74) within those errors. Equation (7.72) read on the eight-row normal instead, the ratio of the Laplace predictives 0.4011/0.5506, gives 0.729: the identity that makes the two forms of the question one number holds for the exact posterior and not for a Gaussian. All of them lie below the exact 0.798 of equation (7.75) and its band $[0.538,\ 0.956]$: the approximations in (u, v) place the remaining life too low and its band too wide on the low side.

Figure 7.7 draws the conditional survival over the next 200 h at a fixed state as a function of the age already reached. The horizontal axis is the age t_0 and the vertical axis $R(200 \mid T > t_0, \mathbf{z})$. The red curve is the illustrative lifetime of Section 7.1.3, $\eta = 800$ h and $\beta = 2$, falling from 0.939 for a new motor through 0.732 at 400 h, marked by a red dot, to 0.503 at 1000 h; the green

curve is the eight-motor Weibull fit, which falls more slowly because its shape is closer to one, with a green dot at $t_0 = 500\,\text{h}$ marking its value 0.727; the dashed gray line is the exponential fit, flat at $e^{-200\hat{\lambda}} = 0.784$, the memoryless property drawn as a curve. The exact predictive does what no fixed-state curve does: it rises with the age reached, from the new motor's value to the value of equation (7.75), where the green curve falls, the mixture of Section 7.3.4 failing its weak members first.

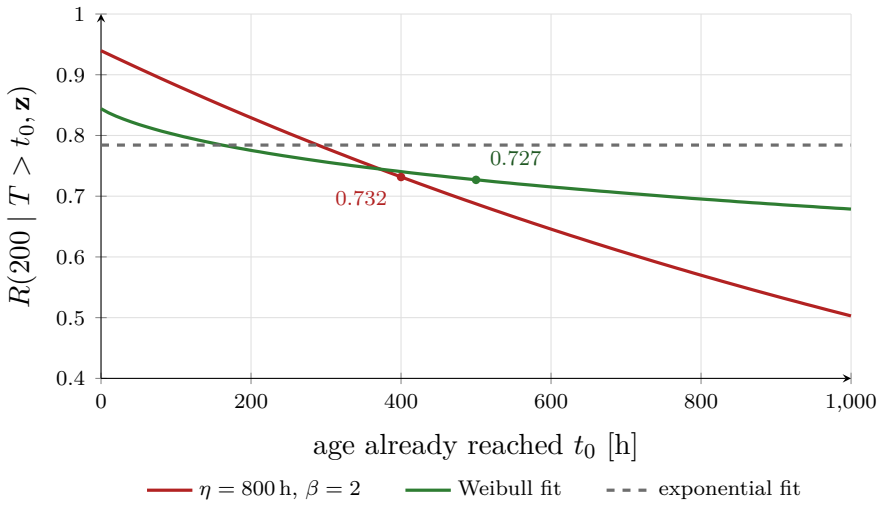


Figure 7.7: The probability of surviving the next 200 h given survival to age t_0 at a fixed state, equation (7.19), for the illustrative lifetime $\eta = 800\,\text{h}$, $\beta = 2$, for the Weibull fit $\hat{\eta} = 765.6\,\text{h}$, $\hat{\beta} = 1.3216$, and for the exponential fit, $\exp(-200 \times 5/4110) = 0.784$ at every age.

7.4.6 The Kaplan–Meier estimate and the discarded survivors

The nonparametric reference. A nonparametric estimate of the survival function is the reference against which the parametric models are judged. The Kaplan–Meier estimator is

$$\hat{R}_{\text{KM}}(t) = \prod_{t_j \le t} \left(1 - \frac{k_j}{n_j}\right), \tag{7.78}$$

where the product runs over the distinct observed failure times t_j up to t, k_j is the number of failures at t_j, and n_j the number of units at risk just before t_j, the units that have neither failed nor been censored earlier. Each factor is the estimated conditional probability of surviving past t_j given survival to just before it, and the product chains those conditional probabilities. It accounts for right censoring through the risk set: a censored unit counts in n_j for every failure time before its censoring and drops out after it.

For the motors, eight are at risk initially. After the first failure, at 120 h, $\hat{R}_{\text{KM}}(120) = 1 - 1/8 = 7/8$; after the second, $\hat{R}_{\text{KM}}(230) = (7/8)(1 - 1/7) = 6/8$; then $(6/8)(1 - 1/6) = 5/8$ at 310 h, $(5/8)(1 - 1/5) = 4/8$ at 450 h and $(4/8)(1 - 1/4) = 3/8 = 0.375$ at 600 h. No unit is censored before the fifth failure, so up to 600 h the estimate equals the fraction of units still running. The three remaining motors are then censored without failing, and the estimate does not drop further. Nothing about the survival beyond 900 h can be inferred from those three observations without a model.

What ignoring censoring does. The five failures occurred at 120, 230, 310, 450 and 600 h, a total of 1710 h. Dividing the five failures by that exposure alone, as if the three survivors had

never been observed, gives

$$\hat{\lambda}_{\text{wrong}} = \frac{5}{1710} = 0.002924\,\text{h}^{-1}, \qquad \hat{R}_{\text{wrong}}(500) = \text{e}^{-0.002924\times 500} = 0.2318.$$

A second error treats the three censoring times as failure times, which keeps the exposure but adds three failures that did not happen: $8/4110 = 0.001946\,\text{h}^{-1}$ and $\hat{R}(500) = 0.3779$. Table 7.12 sets the three calculations side by side: each row is a treatment of the survivors, and the columns give the estimated rate, the mean lifetime it implies and the 500 h survival probability.

Table 7.12: The exponential fit to the motor record under three treatments of the three survivors; only the first uses the censored likelihood (7.66).

Treatment of the survivors	$\hat{\lambda}$ [h^{-1}]	$1/\hat{\lambda}$ [h]	$\hat{R}(500)$ [–]
Censored likelihood, 5/4110	0.0012165	822	54.4 %
Survivors discarded, 5/1710	0.0029240	342	23.2 %
Survivors counted as failures, 8/4110	0.0019465	514	37.8 %

Discarding the survivors gives a substantially more pessimistic estimate, 23.2 % against 54.4 %, and a mean lifetime of 342 h, less than the 700 to 900 h that three of the eight motors are known to have exceeded. The survivors matter because their continued operation is evidence about the lifetime distribution: they add 2400 h of exposure without a failure.

The fits against the estimate. Figure 7.8 compares the Kaplan–Meier estimate with the fitted models. The horizontal axis is age and the vertical axis the survival probability. The black step function is the Kaplan–Meier estimate, with the censoring times marked by vertical ticks; the green curve is the Weibull fit, the gray dashed curve the exponential fit, and the red dotted curve the exponential fitted with the survivors discarded, the second row of Table 7.12. Both censored fits pass through the steps up to 600 h and lie above the last step after it, 0.48 against 0.375 at 600 h, and neither can be rejected by eye on eight units. The curve that ignores censoring lies far below the steps at every age past the first failure: at 700 h, where three of the eight motors are known to be running, it gives 0.129 against the estimate of 0.375.

7.5 The report (step 6)

The characteristic life of 765.6 h that the fit returns is far less certain than the number alone suggests. Its 95 % band runs from about 388 to 1510 h, and the exact posterior beside it reaches much further, from about 390 to 5840 h, and trades it against the shape twice as strongly, a correlation of -0.422 against -0.215; the probability of a rising hazard is 0.762, where the exact value is 0.522, and under the prior of the model comparison it is 0.700, where the exact value is 0.542, little more than even. What the record does pin down is survival inside the characteristic life: a new motor survives a 500 h mission with probability 0.551 (0.574 exact). The answers below are given in the units in which the questions were asked, with the decisions about wear-out and replacement that follow.

7.5.1 The answers and their bands

Table 7.13 is the report. Each row is a question in the engineer's words; the second column gives the answer from the posterior of the state of step 3, under the Weibull model fitted to the eight motors with the prior flat in $\ln\eta$ and $\ln\beta$; the third the exact value beside it; and the last its 95 % credible band, the exact band in parentheses.

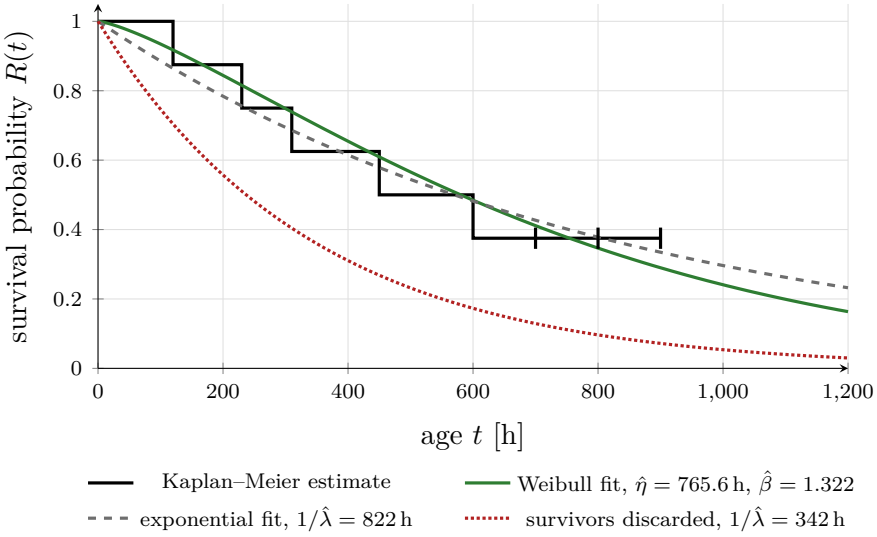


Figure 7.8: The Kaplan–Meier estimate (7.78) of the motors' survival function, with censoring ticks at 700, 800 and 900 h, against the fitted survival functions $\exp[-(t/765.6)^{1.3216}]$, $\exp(-t/822)$ and, ignoring censoring, $\exp(-t/342)$.

Table 7.13: The report of the motor study: the answers from the posterior of the state under the prior flat in $(\ln\eta, \ln\beta)$, the exact value beside each, and the 95 % credible bands of the reported route, exact in parentheses; the evidence ratio uses the prior $\mathcal{N}(0, 0.5^2)$ on $\ln\beta$.

Question	Answer	Exact	Band
What are the scale and the shape of the lifetime, and how well are they known together?	scale: median 765.6 h; shape: MAP 1.32; correlation of $\ln\eta$ and $\ln\beta$ -0.215	scale: median 875 h; shape: mode 1.12; correlation -0.42	$\eta \in [388,\ 1510]$ h ($[390,\ 5840]$ h), $\beta \in [0.61,\ 2.84]$ ($[0.34,\ 2.19]$)
Will a motor survive 500 h?	probability 0.551, averaged over the uncertain state	0.574	–
What is the probability of failure by 500 h?	0.449	0.426	–
Is the 500 h survival probability below one half?	with probability 0.355	0.301	the population's $R(500;\mathbf{z})$ lies in $[0.222,\ 0.840]$ ($[0.294,\ 0.831]$)
What is the median lifetime?	568 h (posterior median)	622 h	$[289,\ 1109]$ h ($[222,\ 2721]$ h); delta 580 h, $[298,\ 1130]$ h
Does the risk of failure increase with age?	not established: $\Pr(\beta > 1 \mid \mathbf{y}) = 0.76$	0.52	–
Which lifetime model describes the record?	the exponential, weakly: $B_{01} = 1.36$	1.47	$B_{01} = 2.18$ (2.39) with the wider prior $\mathcal{N}(0, 1^2)$
How much longer will a surviving motor last?	a motor found running at 500 h survives the next 200 h with probability 0.712	0.798	its conditional survival lies in $[0.234,\ 0.894]$ ($[0.538,\ 0.956]$); plug-in 0.746

The report line. Under a Weibull lifetime fitted to eight motors, five of them failed, a new motor survives a 500 h mission with probability 0.55 (0.57 from the exact posterior), and the population's 500 h survival probability lies between 0.22 and 0.84 (0.29 and 0.83); the median lifetime is about 570 h (620 h), with a band from about 290 to 1110 h (220 to 2700 h), and the characteristic life lies between about 390 and 1500 h (390 and 5800 h). The record does not establish that the hazard rises with age and favors the constant-hazard model weakly. A motor found running at 500 h survives the next 200 h with probability 0.71 against 0.83 for a new one, and from the exact posterior with probability 0.80 against 0.79: on this comparison the two routes fall on opposite sides. The Gaussian in $(\ln\eta, \ln\beta)$ carries the scale's band narrower and the shape higher than the exact posterior, by the gaps of the table.

7.5.2 The wear-out claim and the replacement decision

The question whether the hazard rises is a threshold on the second entry of the state, and its answer depends on how the posterior of that entry is computed: the Laplace Gaussian puts $\Pr(\beta > 1 \mid \mathbf{y})$ at 0.762, the exact marginal at 0.522, and the profile-likelihood interval of Figure 7.4, 0.53 to 2.57, contains both a decreasing and a strongly increasing hazard. The evidence ratio favors the exponential model, weakly, under either prior, on both routes. The data do not provide evidence for a hazard different from constant, and a reliability engineer should not declare a wear-out mechanism on the strength of this record alone. On replacing motors by age the two routes part. A motor found running at 500 h survives the next 200 h with probability 0.712 on the reported route, below the 0.826 of a new motor, which would give replacement some support; on the exact posterior it survives with probability 0.798, above the 0.793 of a new motor, which would give it none. The values fall on opposite sides of the comparison that decides, and both are stated here. Neither makes a case by itself: the conditional survival's band, [0.234, 0.894] ([0.538, 0.956]), covers the new motor's value on both routes.

> **Principle 7.1.** *A fitted lifetime model is a description of a record, not a diagnosis of a mechanism. The shape parameter supports a claim of wear-out only when its uncertainty excludes one, the evidence ratio against the exponential model favors it, and the physical evidence of the failures agrees.*

Common mistake

Discarding the units that have not failed, or recording their censoring times as failure times, biases the estimate toward short lives: the first error turned a 500 h survival probability of 54 % into 23 % (Table 7.12). A related error reports the plug-in survival $R(t;\hat{\mathbf{z}})$ as the probability that a new unit survives, when that probability is the posterior average (7.55), and reports $R(\tau \mid T > t_0, \hat{\mathbf{z}})$ for a surviving unit when the answer is the average of that survival over the posterior on the extended record, equation (7.74), and exactly the ratio (7.72). A third reads $\hat{\beta} > 1$ as proof of wear-out without its band.

7.6 The next measurement (step 7)

The bands of Table 7.13 are wide, and further observation of the fleet could narrow them. The preposterior values each candidate before it is made, by the width it would take off the answers already asked.

7.6.1 The candidate extensions

A candidate a adds rows to the graph, and its expected information $\mathbf{\Lambda}_a$ is the expected Hessian of those rows at $\hat{\mathbf{z}}$, averaged over the outcomes the fitted model predicts; the preposterior

covariance is then that of Chapter 2, equation (2.17), $\boldsymbol{\Sigma}_a = (\boldsymbol{\Sigma}^{-1} + \boldsymbol{\Lambda}_a)^{-1}$, and $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_a\mathbf{g}$ is the variance the answer would have. The (u,u) entry of $\boldsymbol{\Lambda}_a$ has a closed form: by equation (7.41) it is $\hat{\beta}^2$ per failure, and the expected number of failures of a unit observed from age t_0 to age c is $1 - R(c)/R(t_0)$. Four candidates are compared: observing the three survivors F, G and H for a further 500 or 1000 h, or putting five new motors on test for 500 or 1000 h. Five new motors for 1000 h expect $5\,(1 - 0.2409) = 3.80$ failures, which add $1.74673 \times 3.80 = 6.63$ to the (u,u) entry of the information; integrating the row Hessian (7.36) over the outcomes the fit predicts gives the whole matrix, set beside that of the three survivors observed 1000 h longer:

$$\boldsymbol{\Lambda}_{\text{new}} = \begin{bmatrix} 6.630 & 0.101 \\ 0.101 & 4.989 \end{bmatrix}, \qquad \boldsymbol{\Lambda}_{\text{FGH}} = \begin{bmatrix} 4.554 & -5.222 \\ -5.222 & 6.238 \end{bmatrix}.$$

7.6.2 The preposterior and the angle to the gradient

Table 7.14 evaluates the four candidates. Each row is a candidate; the columns give its expected number of failures, the preposterior standard deviations of u, of v and of $R(500;\mathbf{z})$, and the width the candidate removes from the 95 % band $\pm 1.96\sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_a\mathbf{g}}$ of $R(500;\mathbf{z})$, whose present width is $2 \times 1.96 \times 0.1494 = 0.585$.

Table 7.14: Preposterior standard deviations under four candidate extensions of the motor test, $\boldsymbol{\Sigma}_a = (\boldsymbol{\Sigma}^{-1} + \boldsymbol{\Lambda}_a)^{-1}$ with $\boldsymbol{\Lambda}_a$ the expected information at $\hat{\mathbf{z}}$; the first row is the present posterior (7.50). The last column is the width each candidate removes from the 95 % band of $R(500;\mathbf{z})$, now 0.585 wide. All entries but the first column are dimensionless.

Candidate	Expected failures	$\sqrt{\Sigma_{a,uu}}$	$\sqrt{\Sigma_{a,vv}}$	$\sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_a\mathbf{g}}$ for $R(500;\mathbf{z})$	Band width removed
Present record	–	0.346	0.391	0.149	–
F, G, H observed 500 h longer	1.84	0.292	0.314	0.145	0.019
F, G, H observed 1000 h longer	2.61	0.285	0.287	0.144	0.020
Five new motors for 500 h	2.17	0.299	0.327	0.122	0.108
Five new motors for 1000 h	3.80	0.257	0.293	0.115	0.133

The ranking depends on the question. For the 500 h survival probability, new motors buy the most: five of them for 500 h narrow it from 0.149 to 0.122 and remove 0.108 of the band's width, while a further 1000 h on the three survivors barely moves it, removing 0.020. The two information matrices above say why. $\boldsymbol{\Lambda}_{\text{FGH}}$ is nearly of rank one, with eigenvalues 10.69 and 0.11: the old motors inform almost only the direction $(-0.648,\ 0.761)^{\mathsf{T}}$ of the state, a shorter scale with a larger shape or its opposite, and that direction makes an angle of 73° with the gradient $\mathbf{g} = (0.4259,\ 0.1814)^{\mathsf{T}}$ of $R(500;\mathbf{z})$, to which it is nearly blind. The leading direction of $\boldsymbol{\Lambda}_{\text{new}}$ lies within 20° of $\mathbf{g}$. For the shape, the wear-out question, the survivors observed 1000 h longer do as well as five new motors for 1000 h, since they reach the ages where the lever $\ln(t/\eta)$ of the shape is large. Every candidate works through failures, and inspecting the failed motors adds what the times cannot, the failure mode behind each. The preposterior is computed on the posterior of the state and linearized: the widths are those of the delta band, 0.585 now, beside the 0.618 of the band $[0.222,\ 0.840]$ the report gives and the 0.537 of the exact band $[0.294,\ 0.831]$. It ranks the candidates; Chapter 20, in Volume II, writes the same preposterior for a candidate row as a rank-one update of $\boldsymbol{\Sigma}$.

7.6.3 The recommendation

Put five new motors on test for 1000 h: of the four candidates it removes the most from the band of the mission reliability, 0.133 of its 0.585, and it narrows the shape as much as the

survivors observed 1000 h longer, 0.293 against 0.287. Keep F, G and H running as well, since their continued observation needs no new motors; the two extensions together, whose expected informations add, take $\sqrt{\Sigma_{a,vv}}$ to 0.239 and remove 0.141 of the band. Inspect every motor that fails. When the new rows arrive, the fit returns to step 3 with them, and the answers and the exact integrals beside them are recomputed. The replacement question of Section 7.5.2, on which the two routes part, is among those the failures of the new motors bear on.

7.6.4 What the analysis showed

The first readings of the record, a life of 342 h from the five failures alone or a characteristic life of 765.6 h and a rising hazard read off the fit as known, claim more than eight motors can show: the analysis finds the scale between about 388 and 1510 h, and from the exact posterior as far as 5840 h, and a rising hazard at 0.762 by the Gaussian and 0.522 exactly, 0.700 and 0.542 under the prior of the model comparison, which does not establish it on either route. What the record fixes instead is one combination of the two entries, a longer scale together with a larger shape, along (0.864, 0.504) in the state. That combination sets the survival of missions inside the characteristic life, and the gradient of the 500 h survival probability lies within a few degrees of it, which is why the mission reliability is known better than either entry alone. Each entry is poorly known because five failures spread over the first 600 h cannot tell a rising hazard from a constant one, and the small shapes the record cannot exclude go with long scales. The Gaussian of the fit places the shape higher and the scale narrower than the exact integral does, and on the remaining life the two part: a motor still running at 500 h survives the next 200 h with probability 0.712 against 0.826 for a new one, and exactly with 0.798 against 0.793, so the support for replacement by age differs between the routes, and both are stated. The survivors F, G and H, observed longer, would inform almost only the trade between the two entries, to which the mission reliability is blind, while new motors on test inform the combination the question needs; what would settle the wear-out question is failures at the ages the record has not reached, with each failed motor opened to name its failure mode.

7.6.5 Limits of the model

The derivations are exact under their assumptions, and five of them need to be checked before a result is used.

The censoring must be noninformative. Right censoring is acceptable when the censoring process is accounted for. Removing components because their vibration or temperature is deteriorating creates a different observation process from the administrative end of a study, and a likelihood that treats those removals as survivals overstates the reliability (Section 7.1.6).

The units must be comparable. The likelihood multiplies contributions with one common lifetime distribution. Motors operated at different loads or temperatures, or from different lots, need a model in which the state, or the scale, depends on those conditions.

The lifetimes must be independent given the state. A shared cause, a supply transient that stresses every motor at once or a common lubricant defect, makes the failure times dependent, and the product of contributions then overstates the information in the record.

The prior and the Gaussian. The answers hold under the prior flat in $(\ln\eta, \ln\beta)$, and with five failures the prior matters: a proper prior $\mathcal{N}(0, 0.5^2)$ on the log-shape moves $\Pr(\beta > 1 \mid \mathbf{y})$ from 0.762 to 0.700, and exactly from 0.522 to 0.542. The Laplace Gaussian that the answers and the preposterior of step 7 rest on places the shape and the scale away from the exact

posterior (Section 7.4.3), by the gaps stated with each answer; for a lifetime model without the one-integral route, a direct evaluation of the posterior on a grid in the state gives that comparison. A production analysis would also test the goodness of fit and the sensitivity of the answers to the model assumptions.

Age-dependent mechanisms. Even a shape estimated well above one only says that the hazard rises within the Weibull model. The rise may reflect genuine wear-out, a change of operating stress over the study, unmodeled heterogeneity, such as a mixture of motors from two production lots whose weaker members fail first, or simply inadequate data. Physical interpretation requires more than fitting a probability distribution: the failure modes found at teardown, the operating histories and the age-specific exposure have to agree with the mechanism the shape suggests.

The next chapter adds restoration to the lifetime: failure and repair as alternating times, and the availability of a repairable component.

7.7 What to retain

Definition 7.1 (Censored lifetime inference)**.** A lifetime distribution is specified by its hazard $\lambda(t)$, from which $\Lambda(t) = \int_0^t \lambda$, $R(t) = \mathrm{e}^{-\Lambda(t)}$, $F = 1 - R$ and $f = \lambda R$ follow. A record of times t_i with failure indicators δ_i, independent given the state and censored noninformatively, has the log-likelihood $\sum_i \delta_i \ln \lambda(t_i) - \sum_i \Lambda(t_i)$: failures contribute the density and survivors the survival probability. Each unit is one row of the state $\mathbf{z} = (\ln \eta, \ln \beta)^{\mathsf{T}}$, a count row with mean Λ_i in the working form, whose Jacobian row has both entries structurally nonzero; Newton steps on the summed curvature, the Gauss–Newton product plus a named curvature term, give $\hat{\mathbf{z}}$, and its inverse at $\hat{\mathbf{z}}$ gives $\boldsymbol{\Sigma}$, whose off-diagonal entry is the trade between scale and shape. Every answer, the scale's own band, the mission reliability, the median, the remaining life and the evidence ratio against the exponential model, is computed from $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$ and reported, by a normal tail, a Gauss–Hermite sum or $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$, which is also the basis of the preposterior. For the Weibull model under a flat prior in $(\ln \eta, \ln \beta)$ the scale integrates out in closed form given the shape, and each answer has beside it an exact one-dimensional integral over the shape, with the gap between the two.

Four conclusions carry beyond the example. Units that have not failed are evidence, and the likelihood that ignores them is wrong, not merely less precise. The information about a lifetime distribution comes chiefly from failures, so a small record leaves wide bands even when the point estimates look reasonable, and what it determines best may be a combination of the entries rather than any one of them. With a handful of failures the Laplace approximation places the shape high, narrows the scale and distorts the tails, and the exact integral, where it exists, measures by how much: the gap goes with every answer, and where it crosses a decision, as on the remaining life, both values are stated. And the claims that matter to a decision, that the hazard rises, that a surviving unit should be replaced first, rest on the posterior of the shape and on the evidence ratio, not on the estimate. Table 7.15 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 7.15: The terms of lifetime estimation from censored records and where each is defined.

Term	Expression	Units	Meaning	Defined in
Cumulative failure probability	$F(t) = \Pr(T \le t)$	–	failed by age t	7.1.2
Survival function	$R(t) = 1 - F(t)$	–	working at age t	7.1.2
Failure density	$f(t) = -R'(t)$	h^{-1}	density of failure times	7.1.2
Hazard	$\lambda(t) = f(t)/R(t)$	h^{-1}	failure rate of survivors	7.1.2
Cumulative hazard	$\Lambda(t) = \int_0^t \lambda$, $R = \mathrm{e}^{-\Lambda}$	–	accumulated hazard	7.1.2
Mean lifetime	$\mathrm{MTTF} = \int_0^\infty R(t)\,\mathrm{d}t$	h	area under R	7.1.2
Conditional survival	$R(t_0 + \tau)/R(t_0)$	–	remaining life, known state	7.1.2
Exponential lifetime	$R = \mathrm{e}^{-\lambda t}$, $t_{50} = \ln 2/\lambda$	–	constant hazard, memoryless	7.1.3
Weibull lifetime	$R = \exp[-(t/\eta)^\beta]$	–	hazard $\propto t^{\beta-1}$	7.1.3
Weibull mean, median	$\eta\Gamma(1 + 1/\beta)$, $\eta(\ln 2)^{1/\beta}$	h		7.1.3
Characteristic life	$F(\eta) = 0.632$	h	scale, any shape	7.1.3
Censored record	(t_i, δ_i), $t_i = \min(T_i, c_i)$	h, –	time and failure flag	7.1.4
General forms				
Estimation of entries	$\Pr(z_i \le z^* \mid \mathbf{y})$, $\mathrm{Cov}(\mathbf{z} \mid \mathbf{y})$	–	marginals and their trade	7.3.1
Derived quantity, prediction	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$, $\Pr(f(\mathbf{z}) < f^* \mid \mathbf{y})$	–	predictive, tail, quantiles	7.3.2
Tail of an entry, evidence ratio	$\Pr(z_i > z^* \mid \mathbf{y})$, $B_{01} = p(\mathbf{y} \mid M_0)/p(\mathbf{y} \mid M_1)$	–	within and between models	7.3.3
Conditioned prediction	$\int \Pr(A \mid B, \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y}, B)\,\mathrm{d}\mathbf{z}$	–	prediction after an observed event	7.3.4
Specialized forms and answers				
Predictive survival	$\int R(t \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\,\mathrm{d}\mathbf{z}$	–	new unit survives t	7.3.2
Survival gradient	$R\beta\Lambda\,(1, -\ln(t/\eta))^\mathsf{T}$	–	delta method	7.3.2
Median band	$\ln t_{50} = u + \mathrm{e}^{-v} \ln\ln 2$	–	delta method on the log	7.3.2
Exponential estimate	$\hat\lambda = \mathrm{e}^{-\hat u} = k/T_{\mathrm{tot}}$, $\Sigma_{uu} = 1/k$	h^{-1}	limit $\sigma_0 \to 0$	7.3.3
Savage–Dickey ratio	$B_{01} = p(v = 0 \mid \mathbf{y}, M_1)/\pi(0)$	–	exponential against Weibull	7.3.3
Remaining life	$\mathbb{E}[R(t_0 + \tau \mid \mathbf{z})/R(t_0 \mid \mathbf{z})]$ over $\mathcal{N}(\hat{\mathbf{z}}_9, \mathbf{\Sigma}_9)$	–	exactly, the ratio of predictives	7.3.4
Building blocks				
Censored likelihood	$\prod_i f(t_i)^{\delta_i} R(t_i)^{1-\delta_i}$	–	density or survival per unit	7.1.6
Term of unit i	$\mathcal{L}_i = \delta_i[u - v - (\beta - 1)x_i] + \Lambda_i$	–	$-\ln f$ or $-\ln R$ of unit i	7.1.6
Working row	$r_i = (\Lambda_i - \delta_i)/\sqrt{\Lambda_i}$, $\mathbf{J}_i = \sqrt{\Lambda_i}\,\beta(-1, x_i)$	–	count δ_i with mean Λ_i	7.1.6
Best scale for each shape	$\hat\eta(\beta) = (\sum_i t_i^\beta/k)^{1/\beta}$	h	$\sum_i \Lambda_i = k$	7.1.6
Information	$\mathbf{\Lambda} = \mathbf{J}^\mathsf{T}\mathbf{J} + \mathrm{diag}(0, k)$ at $\hat{\mathbf{z}}$	–	$\hat\beta^2$ per failure in (u, u)	7.2.1
Profile likelihood	$\ell_p(\beta) = -\mathcal{L}(\ln\hat\eta(\beta), \ln\beta)$	–	slice at the best scale	7.2.1
Scale given the shape	$\zeta = \eta^{-\beta} \sim \mathrm{Gamma}(k, A(\beta))$	–	conjugate in the scale	7.2.2
Exact marginal of the shape	$\Omega(v) \propto \beta^{k-1}\mathrm{e}^{(\beta-1)S_\delta} A(\beta)^{-k}$	–	scale integrated out	7.2.2
Exact predictive survival	$\int p(v \mid \mathbf{y})\,(A/(A + t^\beta))^k\,\mathrm{d}v$	–	one integral	7.2.2
Kaplan–Meier estimate	$\prod_{t_j \le t}(1 - k_j/n_j)$	–	nonparametric reference	7.4.6
Preposterior	$\mathbf{\Sigma}_a = (\mathbf{\Sigma}^{-1} + \mathbf{\Lambda}_a)^{-1}$	–	value of a test extension	7.6.2

7.8 Exercises

Exercise 7.1. *Formulation.* A motor is observed for 600 h and remains operational. (a) Write the event that was observed and its likelihood contribution. (b) State why recording its failure time as 600 h is incorrect. (c) Write, as a probability over the state, the question whether it will survive another 200 h. ((a) The event is $T > 600$, and its contribution is $R(600 \mid \mathbf{z})$. (b) A failure at 600 h would contribute the density $f(600 \mid \mathbf{z})$, treating the end of the study as a failure time that did not occur. (c) $\Pr(T > 800 \mid T > 600, \mathbf{y})$, which, the motor's survival being part of $\mathbf{y}$, is the posterior average of $R(800 \mid \mathbf{z})/R(600 \mid \mathbf{z})$.)

Exercise 7.2. *Formulation.* A fleet shows most failures after 1000 h, with very few among younger units. Discuss whether an exponential or a Weibull model is more appropriate, and state what further evidence would establish a wear-out mechanism. (The concentration of failures at older ages is consistent with an increasing hazard and favors a Weibull model with $\beta > 1$, to be confirmed by the band of β and the nested comparison with the exponential model. Establishing wear-out needs the age-specific exposure and censoring of every unit, the operating histories, the failure modes found on the failed units, and the exclusion of other explanations: a change of environment, a mixture of lots, selection effects or maintenance practice.)

Exercise 7.3. *Derivation.* (a) A component has the constant hazard $\lambda(t) = 0.0015\,\mathrm{h}^{-1}$. Derive its cumulative hazard, survival function, cumulative failure probability and density, and the probability of failure in the first 200 h. (b) A component has a Weibull lifetime with $\eta = 1000\,\mathrm{h}$ and $\beta = 2$. Compute its survival and failure probabilities at 500 h, its hazard at 500 and 1000 h and its median lifetime, and state why the hazard increases with age. ((a) $\Lambda(t) = 0.0015\,t$, $R(t) = \mathrm{e}^{-0.0015t}$, $F(t) = 1 - \mathrm{e}^{-0.0015t}$, $f(t) = 0.0015\,\mathrm{e}^{-0.0015t}$; $F(200) = 1 - \mathrm{e}^{-0.3} = 0.2592$. (b) $R(500) = \mathrm{e}^{-0.25} = 0.7788$, $F(500) = 0.2212$; $\lambda(500) = (2/1000)(0.5) = 0.001\,\mathrm{h}^{-1}$, $\lambda(1000) = 0.002\,\mathrm{h}^{-1}$; $t_{50} = 1000\sqrt{\ln 2} = 832.6\,\mathrm{h}$; the hazard is proportional to $t^{\beta-1} = t$ because $\beta = 2 > 1$.)

Exercise 7.4. *Derivation.* Five motors are followed: A fails at 100 h, B at 250 h, C at 400 h; D is censored at 500 h and E at 600 h. (a) Derive the exponential likelihood, write the objective of the scale alone, u with $\lambda = \mathrm{e}^{-u}$ and $v = 0$, and compute the total exposure, the number of failures, the MAP rate with Σ_{uu}, the mean lifetime and the 300 h survival probability. (b) Write the five Weibull rows $\mathcal{L}_i(\mathbf{z})$ for $\mathbf{z} = (u, v)^{\mathsf{T}}$, their score rows and the Newton equations (7.40), and give $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$ after Newton steps from $\mathbf{z}_0 = (\ln 500,\ 0)^{\mathsf{T}}$. ((a) $p(\mathbf{y} \mid \lambda) = \lambda^3\,\mathrm{e}^{-1850\lambda}$ with $T_{\mathrm{tot}} = 1850\,\mathrm{h}$ and $k = 3$; $\mathcal{L} = 3u + 1850\,\mathrm{e}^{-u}$, gradient $3 - 1850\,\mathrm{e}^{-u}$, curvature $1850\,\mathrm{e}^{-u}$; $\hat{u} = \ln(1850/3) = 6.4243$, $\hat{\lambda} = 3/1850 = 0.001622\,\mathrm{h}^{-1}$, $\Sigma_{uu} = 1/3$, $1/\hat{\lambda} = 616.7\,\mathrm{h}$, $\hat{R}(300) = \mathrm{e}^{-300\times 3/1850} = 0.6148$. (b) $\mathcal{L}_i = \delta_i[u - v - (\beta - 1)x_i] + \Lambda_i$ with $x_i = \ln t_i - u$ and $\delta = (1, 1, 1, 0, 0)$; the score rows are equation (7.31) and the curvature the sum of equation (7.36); four Newton steps give $\hat{\mathbf{z}} = (6.3320,\ 0.3093)^{\mathsf{T}}$, $\hat{\eta} = 562.3\,\mathrm{h}$, $\hat{\beta} = 1.363$, and $\boldsymbol{\Sigma} = \left[\begin{smallmatrix} 0.1929 & -0.0591 \\ -0.0591 & 0.2622 \end{smallmatrix}\right]$, so $\sqrt{\Sigma_{uu}} = 0.439$, $\sqrt{\Sigma_{vv}} = 0.512$ and the correlation is -0.263; the curve of best scale $\hat{\eta}(\beta) = [(100^\beta + 250^\beta + 400^\beta + 500^\beta + 600^\beta)/3]^{1/\beta}$ passes through $\hat{\mathbf{z}}$.)

Exercise 7.5. *Application.* For the five motors of the previous exercise, the prior is $\lambda \sim \mathrm{Gamma}(2, 500\,\mathrm{h})$. (a) Derive the posterior and its mean. (b) Derive the posterior predictive probability that a new motor survives a 300 h mission, and state why it differs from $\mathrm{e}^{-300\,\mathbb{E}[\lambda|\mathbf{y}]}$. (c) A Weibull component with $\eta = 1000\,\mathrm{h}$ and $\beta = 2$ has survived 600 h; compute the probability that it survives another 200 h, and compare it with a new component's first 200 h. ((a) The likelihood $\lambda^3\mathrm{e}^{-1850\lambda}$ times the prior $\lambda\,\mathrm{e}^{-500\lambda}$ gives $\mathrm{Gamma}(5, 2350)$, mean $5/2350 = 0.002128\,\mathrm{h}^{-1}$. (b) $\int \mathrm{e}^{-300\lambda}\,\mathrm{Gamma}(\lambda; 5, 2350)\,\mathrm{d}\lambda = (2350/2650)^5 = 0.5484$, against the plug-in $\mathrm{e}^{-300\times 0.002128} = 0.5282$; the average of the convex function $\mathrm{e}^{-300\lambda}$ exceeds its value at the average. (c) $\exp[-(800^2 - 600^2)/1000^2] = \mathrm{e}^{-0.28} = 0.7558$, against $\mathrm{e}^{-0.04} = 0.9608$ for a new component: the hazard increases with age.)

Exercise 7.6. *Application.* The log-Weibull parameters have the posterior approximation $\mathbf{z} = (\ln\eta, \ln\beta)^{\mathsf{T}} \sim \mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$. (a) Derive the gradient of the survival probability at age t with respect to $\mathbf{z}$, and write the delta-method variance. (b) For the motors, $\hat{\eta} = 765.6\,\mathrm{h}$, $\hat{\beta} = 1.3216$ and $\boldsymbol{\Sigma}$ of equation (7.50), evaluate the gradient at the characteristic life $t = \hat{\eta}$ and the standard deviation of $R(\hat{\eta})$. (c) State what should be checked before the approximation is trusted. ((a) With $\Lambda = (t/\eta)^{\beta}$ and $R = \mathrm{e}^{-\Lambda}$, $\mathbf{g} = R\beta\Lambda\,(1, -\ln(t/\eta))^{\mathsf{T}}$ and $\mathrm{Var}(R) \approx \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$. (b) At $t = \hat{\eta}$, $\Lambda = 1$ and $\ln(t/\eta) = 0$, so $\mathbf{g} = (\mathrm{e}^{-1} \times 1.3216,\ 0)^{\mathsf{T}} = (0.4862,\ 0)^{\mathsf{T}}$ and $\sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}} = 0.4862\sqrt{0.12004} = 0.168$: only the scale matters there. (c) That the posterior of the state is close to Gaussian, that R is close to linear across its width, how the correlation enters the budget, and that the band stays inside $[0, 1]$; the exact one-integral posterior, or draws of $\mathbf{z}$ pushed through R, are the check.)

Chapter 8

Repair Rates and Component Availability

8.1 Repair rates and the availability of a repairable component

The cooling-water circulation pump of Chapter 6 has failed four times in two thousand operating hours, and each time the maintenance crew returned it to service; the four repairs took 6, 14, 9 and 11 hours. The service agreement for the cooling network asks for a pump that is available at least 97 % of the time, and the first reading of the two logs is simple: four failures in 2000 h is a failure rate of $0.002\,\mathrm{h}^{-1}$, four repairs in 40 h a repair rate of $0.1\,\mathrm{h}^{-1}$, and 0.1/0.102 is an availability of 0.980, above the requirement, so the pump passes.

The engineer who must show the agreement is met asks something else. The agreement wants the 97 % shown with an assurance of 0.90, and the 0.980 is built from four failures and four repairs, neither of them many. The point value is a statistic of the two logs; the long-run availability of the pump is an unknown, a function of two uncertain rates, and how sure the plant can be of it depends on the size of both records as well as on their means. The question of this chapter is

how the repair rate of a component is estimated from a record of repair durations, how the uncertainties of its failure rate and its repair rate combine in the uncertainty of its availability, how probable it is that the availability meets a requirement, and which further record would narrow it most.

It is answered with the method of the previous chapters on the same pump. What changes is the state, which now carries two entries, the log failure rate and the log repair rate, each fed by its own record, and the quantity asked about, the availability, which is a function of both entries rather than an entry itself. The first reading, 0.980 against a requirement of 0.97, looks like a pass.

Jargon: availability, unavailability, repair rate, mean time to repair (MTTR), two-state model, forward equation, time constant, interval availability, outage frequency, inverse-gamma distribution, digamma and trigamma functions, logit, repair drill, alternating renewal process

The *availability* of a repairable item is the probability that it is in its up state; the *steady-state availability* is its long-run value, equal to the long-run fraction of time spent up, and the *instantaneous availability* $A(t)$ is its value at a given time after a known starting state. The *unavailability* is the complement, the probability of being down. The *repair rate* μ is the constant rate at which an item under repair is restored, in h^{-1}, and the *mean time to repair* (MTTR) is the expected duration of one repair. The *two-state model* describes the item as either up or down, failing at rate λ and being repaired at rate μ; its *forward equation* is the differential equation for the probability of the up state, whose solution decays to the steady state with a *time constant*. The *interval availability* is the average of $A(t)$ over a stated interval. The *outage frequency* is the expected number of failures per unit of calendar time. The *inverse-gamma distribution* is the distribution of the reciprocal of a gamma variable. The *digamma function* $\psi(n)$ and the *trigamma function* $\psi'(n)$ are the first and second derivatives of $\ln\Gamma(n)$. The *logit* of a probability p is $\ln(p/(1-p))$, its log-odds. A *repair drill* is a staged repair carried out by the maintenance crew under recorded time to learn the repair duration without waiting for a failure. An *alternating renewal process* is a sequence of up and down intervals, each drawn independently from its own distribution.

Key idea

A record of failures in an operating exposure and a record of repair durations inform two different rates, and for a constant failure rate and a constant repair rate their posteriors are independent gamma distributions. The availability $A = \mu/(\lambda+\mu)$ is a derived quantity: the logit of the unavailability is $\xi = \mathbf{g}^{\mathsf{T}}\mathbf{z}$ with $\mathbf{g} = (1,-1)^{\mathsf{T}}$ on the state of the two log rates, and since the posterior covariance $\boldsymbol{\Sigma}$ of that state is diagonal, its variance $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ is the sum of two terms, one from each record; the exact posterior of the availability is a transformed Beta distribution, the Laplace route reproduces it closely, and among records of comparable size the one that feeds both entries at once narrows the availability most.

8.1.1 Availability as a function of two rates

The companion volume *Reliability Engineering* distinguishes availability from reliability and defines the mean times of a repairable item; those definitions are taken as known here, and only what the uncertainty question needs is restated. A repairable pump with a constant failure rate λ and a constant repair rate μ has a mean time to failure $\mathrm{MTTF} = 1/\lambda$ and a mean time to repair $\mathrm{MTTR} = 1/\mu$. Its steady-state availability, derived from the two-state model in Section 8.3.3, is

$$A = \frac{\mu}{\lambda+\mu} = \frac{\mathrm{MTTF}}{\mathrm{MTTF}+\mathrm{MTTR}}, \tag{8.1}$$

and its unavailability, the long-run fraction of time it is down, is

$$U = 1 - A = \frac{\lambda}{\lambda+\mu}. \tag{8.2}$$

When repairs are much faster than failures, $\mu \gg \lambda$, the unavailability is close to $\lambda/\mu = \mathrm{MTTR}/\mathrm{MTTF}$, the mean repair time as a fraction of the mean time between failures.

Both rates are uncertain. The failure rate is known from four failures, the repair rate from four repairs, and neither record is large. Equation (8.1) turns a pair of uncertain rates into an uncertain availability, and the questions an operator asks about the availability, its value,

its band, whether it meets a contractual level and how many hours per year the pump will be down, are questions about that derived quantity.

8.1.2 Two records and one derived quantity

The availability is not an entry of the state. The state carries the two rates, on the logarithmic scale that keeps them positive, and the availability is a function of both. Chapter 4 gave the general method for such a quantity: linearize it at the MAP state and compute its variance as $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$, the variance budget of its gradient against the posterior covariance. Here the derived quantity has a structure that makes the method unusually transparent. The log-odds of the unavailability is $\mathbf{g}^\mathsf{T}\mathbf{z}$ with the constant gradient $\mathbf{g} = (1, -1)^\mathsf{T}$, and since each row of the evidence acts on one entry only, $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$ has exactly two terms, one per rate (Section 8.3.1).

The same structure gives an exact answer. The two rate posteriors are gamma distributions, and a ratio built from two independent gamma variables has a Beta distribution. Every probability and band of the availability then becomes a Beta tail or quantile, and its mean a one-dimensional integral, which can be compared with the Laplace answer number by number. The chapter carries both and reports the Laplace one, because the Laplace route is the one that survives when the pump is one node of a larger plant whose posterior has no closed form, and it sets the exact value beside each answer, with the gap.

8.1.3 The physical system and its graph (step 0)

The cooling-water circulation pump of Chapter 6 draws from a water supply and delivers to a cooling network. It operates continuously when it is up; when it fails, the maintenance crew repairs it and returns it to service, and the pump is then up again, statistically as good as before. Its history is an alternating sequence of up intervals, ended by failures, and down intervals, ended by completed repairs. The model of this chapter rests on the assumptions of Chapter 6 for the up intervals, a constant failure intensity over operating time, and on three more for the down intervals: repair durations are independent of one another and of the up intervals; they share one constant repair rate, so that each is exponentially distributed; and the operating exposure counts only time in service, never time under repair.

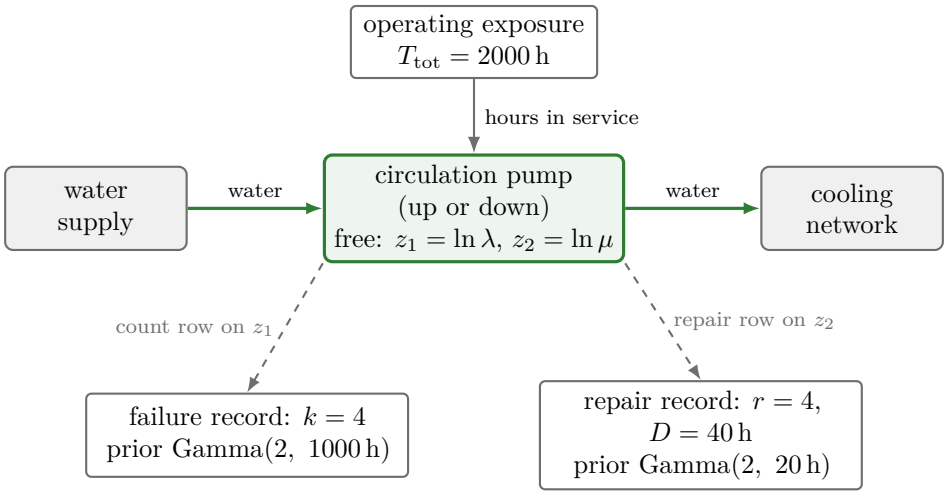
Figure 8.1 draws the pump as a graph. In panel (a) the gray boxes on the left and right are the water supply and the cooling network, pinned boundaries of the hydraulic path drawn as solid arrows. The green box in the middle is the pump, the free node, which is either up or down and carries the two entries of the state, the log failure rate z_1 and the log repair rate z_2. The white box above it is the operating exposure, pinned because the hours in service are read from the plant log. The two dashed arrows leave the pump for the evidence: the failure record on the left, which carries a count of failures in the exposure and the prior on the failure rate, and the repair record on the right, which carries the number and total duration of the repairs and the prior on the repair rate. Panel (b) draws the pump's two states as circles, with the failure transition from up to down at rate λ and the repair transition from down to up at rate μ; it is the model from which Section 8.3.3 derives the availability.

The state entries. Both rates must be positive, and the graph carries each on the logarithmic scale,

$$\mathbf{z} = \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} = \begin{bmatrix} \ln\lambda \\ \ln\mu \end{bmatrix} \quad \begin{matrix} \text{log failure rate} \\ \text{log repair rate'} \end{matrix} \tag{8.3}$$

with both rates in h^{-1}. Chapter 6 carried the failure rate as $\omega = \ln(\lambda/\lambda_{\mathrm{ref}})$ with a reference rate λ_{ref}; the reference shifts the entry by the constant $\ln\lambda_{\mathrm{ref}}$ and changes neither its variance nor any difference of log rates, so this chapter drops it. The availability, the mean time to

(a) plant graph of the repairable pump



(b) the two-state model

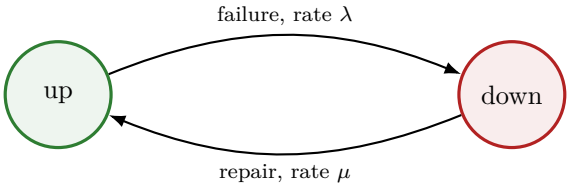


Figure 8.1: The repairable pump. (a) Plant graph: free, the pump node, carrying the log failure rate z_1 and the log repair rate z_2 (green); pinned, the water supply, the cooling network and the operating exposure T_{tot} (gray and white); observed, the failure count in that exposure and the repair durations (dashed arrows), each record with the prior on its rate. (b) The two states of the pump and the transitions between them, failures at rate λ and repairs at rate μ.

repair, the downtime and the outage frequency are derived quantities, functions of the state. The evidence $\mathbf{y}$ is the failure record, the repair record and the two priors.

8.1.4 The two records

The failure record. The pump's record from Chapter 6 is $k = 4$ failures in $T_{\text{tot}} = 2000$ operating hours, under the prior $\lambda \sim \text{Gamma}(a_{\lambda 0}, b_{\lambda 0}) = \text{Gamma}(2, 1000\,\text{h})$, which places the failure rate near $0.002\,\text{h}^{-1}$. The gamma distribution is written, as throughout the book, with a shape and a rate parameter, here $a_{\lambda 0}$ and $b_{\lambda 0}$, the subscript 0 marking the prior's and its absence the posterior's; the rate parameter carries the units of time.

The repair record. Each of the four failures was repaired, and the maintenance log gives the duration of each repair, from the moment the pump stopped to the moment it was returned to service. Table 8.1 lists them: each row is one repair, the middle column its duration and the right column the running total. The four repairs took 6, 14, 9 and 11 hours, $r = 4$ repairs with total duration $D = 40\,\text{h}$, a mean of 10 h per repair; the last rows give the sample standard deviation and the coefficient of variation, which Section 8.4.3 uses to check the repair-time model. The record is illustrative.

Table 8.1: The repair record of the pump: one repair per recorded failure. The coefficient of variation is the sample standard deviation divided by the mean.

Repair j	Duration d_j (h)	Running total (h)
1	6	6
2	14	20
3	9	29
4	11	40
Count and total	$r = 4$	$D = 40$
Mean and standard deviation	10.0 and 3.37	
Coefficient of variation	0.34	

The prior on the repair rate. Experience with similar pumps and the same crew puts the repair rate near $0.1\,\text{h}^{-1}$ with wide uncertainty. The prior is

$$\mu \sim \text{Gamma}(a_{\mu 0}, b_{\mu 0}) = \text{Gamma}(2,\ 20\,\text{h}), \tag{8.4}$$

of the same shape as the prior on the failure rate. Its reciprocal mean rate is not its mean repair duration, a distinction Section 8.3.1 derives.

8.1.5 The probabilistic questions (step 1)

The operator asks five questions of the two records, and step 1 writes each as a probability or an expectation over the state, with the unknown on the left of the bar and the records on the right. Table 8.2 lists them: the first column is the question in engineering language; the second the general form that every question of its category takes, written in the book's notation for a quantity $f(\mathbf{z})$ derived from the state, a threshold f^* on it, an assurance level κ and the posterior density $p(\mathbf{z} \mid \mathbf{y})$; the third its mathematical form for the pump; the fourth its category; and the last the subsection where it is answered. The first two are estimation, of a transformed entry of the state, the log repair rate $\ln \mu$ reported as the repair rate and its reciprocal, and of a quantity derived from both entries, the availability; they share a general form. The third is a

prediction under a threshold, a tail probability of the availability's posterior compared with the assurance level $\kappa = 0.90$ that the service agreement asks for. The fourth and fifth ask for further derived quantities, the annual downtime, the outage frequency and the availability at a time after a repair, through their expectations and quantiles under the posterior. A sixth question, which further record would narrow the availability most, asks nothing new of the present records: it is the next measurement for these five, answered in Section 8.6. The atlas of Appendix A lists the same forms.

Table 8.2: The questions asked of the pump's two records, the general form of their category, their mathematical form for the pump, their category and where each is answered. $\mathbf{y}$ is the failure record, the repair record and the two priors; A_{req} a required availability; $f_t(\mathbf{z})$ the availability $A(t)$ at time t after a repair.

Engineering question	General form	Mathematical form	Category	Answered in
What is the repair rate, and the mean time to repair?	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$, moments	$p(\mu \mid r, D)$, $p(\text{MTTR} \mid r, D)$	estimation	8.3.1
What fraction of the time will the pump be available, and how well is that known?	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$, quantiles	$p(A \mid \mathbf{y})$, $A = \mu/(\lambda + \mu)$	estimation of a derived quantity	8.3.1
Is the availability at least 0.97 with stated assurance?	$\Pr(f(\mathbf{z}) \ge f^* \mid \mathbf{y}) \ge \kappa$	$\Pr(A \ge 0.97 \mid \mathbf{y}) \ge 0.90$	prediction under a threshold	8.3.2
How many hours per year will the pump be down, and how often?	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$, quantiles	$8760\,\text{h} \times U$; $\nu_{\text{out}} = \lambda A$	estimation of derived quantities	8.3.3
How available is the pump in the hours after a repair?	$\mathbb{E}[f_t(\mathbf{z}) \mid \mathbf{y}]$, quantiles	$A(t)$ of the two-state model	estimation of a derived quantity	8.3.3

8.1.6 The evidence as rows (step 2)

Step 2 attaches four rows to the graph. Two act on z_1: the prior row of the failure rate and the count row of the failure record, through the Poisson likelihood of Chapter 6. Two act on z_2: the prior row of the repair rate and the repair row of the repair record, through the likelihood of the repair durations derived below. Which row acts on which entry is a property of the graph, and it fixes where the Jacobian of the rows can be nonzero, marked $*$:

$$\mathbf{J} = \begin{bmatrix} * & 0 \\ * & 0 \\ 0 & * \\ 0 & * \end{bmatrix} \quad \begin{matrix} \text{prior on } \lambda \\ \text{count row} \\ \text{prior on } \mu\text{'} \\ \text{repair row} \end{matrix} \qquad \text{columns } z_1, z_2. \tag{8.5}$$

No row acts on both entries, at any state. No row is a physics row either: the availability enters as a derived quantity of the state, not as an entry tied to the rates by a constraint, so no model tolerance σ_p arises. None of the four rows is Gaussian: each enters the objective as its negative log-probability term, written in the working form of a normalized row with a residual and a Jacobian row, below. Step 3 combines the four rows into the posterior of the state, and step 4 reduces each question of Table 8.2 to a function of that posterior.

The modeling choices. Three modeling choices enter the rows. The repair durations are taken as exponential, with one repair rate shared by all of them: this is what lets the repair

record be summarized by r and D alone and gives the repair rate a gamma posterior, and it fails if the repairs are much more regular than an exponential duration allows, which Section 8.4.3 tests on the four recorded durations. The failure record and the repair record are taken as independent given the two rates, under independent priors: this is what makes the covariance of the two log rates diagonal and the variance of the availability a sum of two terms with no cross term, and it fails when one crew serves several pumps and its repair rate is shared among them (Section 8.6.6). The operating exposure counts time in service only: this keeps the failure rate a rate per hour in which the pump could fail, and counting the hours under repair as well flatters the pump, by an amount Section 8.4.4 measures.

The likelihood of an up-and-down record. The pump's history over the record is a sequence of up intervals of lengths $u_1, \ldots, u_k$, each ended by a failure, a final up interval u_{k+1} still in progress when the record closes, and down intervals of lengths $d_1, \ldots, d_r$, each ended by a completed repair, with $r = k$ when every failure has been repaired. Under the constant failure rate the length of an up interval ended by a failure has the exponential density $\lambda\,\mathrm{e}^{-\lambda u_i}$, and the final interval, which has not ended, contributes the probability of surviving at least u_{k+1}, $\mathrm{e}^{-\lambda u_{k+1}}$. Under the constant repair rate each repair duration has the density $\mu\,\mathrm{e}^{-\mu d_j}$. The intervals are independent given the rates, so the probability density of the whole record is the product

$$p(\text{record} \mid \lambda, \mu) = \prod_{i=1}^{k} \lambda\,\mathrm{e}^{-\lambda u_i} \times \mathrm{e}^{-\lambda u_{k+1}} \times \prod_{j=1}^{r} \mu\,\mathrm{e}^{-\mu d_j}.$$

Collect the powers and the exponents. The up intervals sum to the operating exposure, $\sum_{i=1}^{k+1} u_i = T_{\text{tot}}$, and the repair durations to the total repair time, $\sum_{j=1}^{r} d_j = D$:

$$p(\text{record} \mid \lambda, \mu) = \underbrace{\lambda^k\,\mathrm{e}^{-\lambda T_{\text{tot}}}}_{\text{failure part}} \times \underbrace{\mu^r\,\mathrm{e}^{-\mu D}}_{\text{repair part}}. \tag{8.6}$$

The likelihood factorizes into a part that depends only on λ, through the count k and the exposure T_{tot}, and a part that depends only on μ, through the count r and the total D. The failure part is, up to a factor free of λ, the Poisson likelihood of k failures in T_{tot} hours of Chapter 6; the exposure in it is the sum of the up intervals, which is why repair time must be excluded from it. The two parts share no parameter.

Independence of the two posteriors. With independent priors $p(\lambda)\,p(\mu)$, Bayes' theorem gives

$$p(\lambda, \mu \mid \text{record}) \propto \left[\lambda^k \mathrm{e}^{-\lambda T_{\text{tot}}}\,p(\lambda)\right] \times \left[\mu^r \mathrm{e}^{-\mu D}\,p(\mu)\right],$$

a product of a function of λ and a function of μ. The joint posterior is therefore the product of its marginals: the failure rate and the repair rate are independent a posteriori, and each is updated by its own part of the record. This is the exact statement; Section 8.2.1 shows its Laplace counterpart.

The repair likelihood and its sufficient statistic. The repair part,

$$L(\mu) = \mu^r\,\mathrm{e}^{-\mu D}, \tag{8.7}$$

depends on the durations only through their number and their total, so (r, D) is a sufficient statistic for the repair rate, as (k, T_{tot}) is for the failure rate. The same likelihood follows if only the total is recorded. The sum of r independent exponential durations of rate μ has the gamma density

$$p(D \mid r, \mu) = \frac{\mu^r\,D^{r-1}\,\mathrm{e}^{-\mu D}}{\Gamma(r)}, \qquad D \mid r, \mu \sim \text{Gamma}(r, \mu), \tag{8.8}$$

a standard result obtained by convolving the exponential density with itself $r-1$ times; as a function of μ it is proportional to equation (8.7). The repair row of the graph is this gamma-sum density of the total repair time.

The rows as count rows. The four rows act on the state $\mathbf{z} = (z_1, z_2)^\mathsf{T}$ of equation (8.3). None is Gaussian; each is the negative log-probability term of its record, written in the state, and Chapter 6 derived how such a row takes the form of a normalized Gaussian row. As in that chapter's treatment of the failure rate, the density of a log rate is the density of the rate times the change-of-variable factor $\mathrm{d}\lambda/\mathrm{d}z_1 = \lambda$, so the gamma prior $\mathrm{Gamma}(a_{\lambda 0}, b_{\lambda 0})$ on λ becomes, in z_1, the prior row $b_{\lambda 0}\mathrm{e}^{z_1} - a_{\lambda 0} z_1$, with the shape $a_{\lambda 0}$ rather than $a_{\lambda 0} - 1$: a count row of $a_{\lambda 0}$ pseudo-failures in $b_{\lambda 0}$ pseudo-hours. The prior on μ becomes $b_{\mu 0}\mathrm{e}^{z_2} - a_{\mu 0} z_2$ in the same way. The count row is the negative log of the failure part of equation (8.6), $T_{\mathrm{tot}}\mathrm{e}^{z_1} - k z_1$, and the repair row the negative log of the gamma-sum density (8.8) as a function of z_2, $D\mathrm{e}^{z_2} - r z_2$, with the terms free of the state dropped. The repair row alone has the slope $D\mathrm{e}^{z_2} - r$, which vanishes at $\mathrm{e}^{z_2} = r/D$, the maximum-likelihood repair rate and the reciprocal of the mean repair duration. Indexing the rows by $i = 1, \ldots, 4$, all four have the form of the count rows of Chapter 6,

$$\begin{gathered}\mathcal{L}_i(\mathbf{z}) = \bar{k}_i - k_i \ln \bar{k}_i, \qquad \bar{k}_i = h_i\,\mathrm{e}^{z_1}\ (i = 1, 2), \quad \bar{k}_i = h_i\,\mathrm{e}^{z_2}\ (i = 3, 4), \\ \begin{array}{lll} i = 1: & \text{prior on } \lambda, & (k_1, h_1) = (2,\ 1000\,\mathrm{h}), \\ i = 2: & \text{count row}, & (k_2, h_2) = (4,\ 2000\,\mathrm{h}), \\ i = 3: & \text{prior on } \mu, & (k_3, h_3) = (2,\ 20\,\mathrm{h}), \\ i = 4: & \text{repair row}, & (k_4, h_4) = (4,\ 40\,\mathrm{h}), \end{array}\end{gathered} \tag{8.9}$$

with k_i the observed count of the row (the failures k, the repairs r, or a prior's shape counted as pseudo-events), h_i its exposure (the operating hours, the total repair time, or a prior's rate parameter) and $\bar{k}_i$ the count it expects at the state. The term $\mathcal{L}_i(\mathbf{z})$ is the one row i puts in the objective, $-\ln p$ of its record given the state up to a constant. Summing the four terms gives the objective $\mathcal{L} = \sum_i \mathcal{L}_i$,

$$\mathcal{L}(\mathbf{z}) = \underbrace{1000\,\mathrm{e}^{z_1} - 2z_1}_{\text{prior on } \lambda} + \underbrace{2000\,\mathrm{e}^{z_1} - 4z_1}_{\text{count row}} + \underbrace{20\,\mathrm{e}^{z_2} - 2z_2}_{\text{prior on } \mu} + \underbrace{40\,\mathrm{e}^{z_2} - 4z_2}_{\text{repair row}} = b_\lambda \mathrm{e}^{z_1} - a_\lambda z_1 + b_\mu \mathrm{e}^{z_2} - a_\mu z_2, \tag{8.10}$$

with the exposures and repair times in hours, $a_\lambda = a_{\lambda 0} + k$, $b_\lambda = b_{\lambda 0} + T_{\mathrm{tot}}$, $a_\mu = a_{\mu 0} + r$ and $b_\mu = b_{\mu 0} + D$, up to a constant.

The working residuals and the Jacobian. The log of each expected count is linear in the state, with derivative one with respect to the entry its row acts on and zero with respect to the other, so the constant 4×2 matrix $\partial \ln \bar{\mathbf{k}}/\partial \mathbf{z}$ has the pattern (8.5) with every $*$ equal to one. An entry 1 in the first column says that row i's predicted log count moves one for one with z_1: raising the log failure rate by 0.1 raises the expected number of failures, and the prior's pseudo-failures, by the factor $\mathrm{e}^{0.1}$, and leaves the two rows of the second column unchanged. Each of the four rows is a count row of Section 4.2.7, with expected count $\bar{k}_i$ in place of μ: its gradient and Hessian are those of equation (4.42), with no remaining term because $\ln \bar{k}_i$ is linear in the state, and it joins the Jacobian through the working pair (4.43), of which equation (6.22)

is the pump's instance. Here the pair and its stack are

$$r_i(\mathbf{z}) = \frac{\bar{k}_i - k_i}{\sqrt{\bar{k}_i}}, \qquad \mathbf{J}_i = \sqrt{\bar{k}_i}\,\frac{\partial \ln \bar{k}_i}{\partial \mathbf{z}^\mathsf{T}},$$

$$\mathbf{r}(\mathbf{z}) = \begin{bmatrix} (1000\,\mathrm{e}^{z_1} - 2)/\sqrt{1000\,\mathrm{e}^{z_1}} \\ (2000\,\mathrm{e}^{z_1} - 4)/\sqrt{2000\,\mathrm{e}^{z_1}} \\ (20\,\mathrm{e}^{z_2} - 2)/\sqrt{20\,\mathrm{e}^{z_2}} \\ (40\,\mathrm{e}^{z_2} - 4)/\sqrt{40\,\mathrm{e}^{z_2}} \end{bmatrix}, \qquad \mathbf{J}(\mathbf{z}) = \begin{bmatrix} \sqrt{1000\,\mathrm{e}^{z_1}} & 0 \\ \sqrt{2000\,\mathrm{e}^{z_1}} & 0 \\ 0 & \sqrt{20\,\mathrm{e}^{z_2}} \\ 0 & \sqrt{40\,\mathrm{e}^{z_2}} \end{bmatrix}, \tag{8.11}$$

whose gradient and curvature are those of $\mathcal{L}_i$, so that the gradient and the Hessian of the objective are

$$\frac{\partial \mathcal{L}}{\partial \mathbf{z}} = \mathbf{J}^\mathsf{T}\mathbf{r}, \qquad \frac{\partial^2 \mathcal{L}}{\partial \mathbf{z}\,\partial \mathbf{z}^\mathsf{T}} = \mathbf{J}^\mathsf{T}\mathbf{J}. \tag{8.12}$$

The Hessian is exact. The curvature term $\mathbf{C}$ of equation (4.44), the slope times the second derivative of $\ln \bar{k}_i$, is zero because the log count is linear in the state, the canonical case, so Newton's method and the Gauss–Newton step coincide here. With every row normalized in this way the information matrix is $\mathbf{J}^\mathsf{T}\mathbf{J}$; with the unnormalized log-count Jacobian and the weights $\mathbf{W} = \mathrm{diag}(\bar{k}_1, \ldots, \bar{k}_4)$ it is $(\partial \ln \bar{\mathbf{k}}/\partial \mathbf{z})^\mathsf{T}\,\mathbf{W}\,(\partial \ln \bar{\mathbf{k}}/\partial \mathbf{z})$, the same matrix, since $\mathbf{J} = \mathbf{W}^{1/2}\,\partial \ln \bar{\mathbf{k}}/\partial \mathbf{z}$.

8.2 The posterior of the state (step 3)

In the book's notation, $p(\mathbf{z} \mid \mathbf{y})$ is the posterior density of the state, $f(\mathbf{z})$ a quantity derived from the state and f^* a requirement level on it.

All five are read from one posterior: the posterior of the state that the four rows of step 2 give, exactly two independent gamma densities of the rates and approximately the Laplace normal of the log rates. This section derives it and computes it for the pump; Section 8.3 reads every answer from it.

Four objects of the fit recur in the specialized forms. The *MAP state* $\hat{\mathbf{z}} = (\hat{z}_1, \hat{z}_2)^\mathsf{T}$ is the minimum of the objective that the four rows put on the state; the *information matrix* $\mathbf{\Lambda} = \mathbf{J}^\mathsf{T}\mathbf{J}$ is the curvature of that objective at the minimum, with $\mathbf{J}$ the 4×2 Jacobian of the normalized rows; the *Laplace covariance* $\mathbf{\Sigma} = \mathbf{\Lambda}^{-1}$, a 2×2 matrix, makes the posterior of the state approximately normal, $\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$; and a derived quantity f is carried by its *gradient* $\mathbf{g} = \partial f/\partial \mathbf{z}$ at $\hat{\mathbf{z}}$, with the variance $\mathbf{g}^\mathsf{T}\mathbf{\Sigma}\mathbf{g}$. Beside them stands the exact posterior that this problem admits: the two rates are independent gamma variables, $\lambda \mid \mathbf{y} \sim \mathrm{Gamma}(a_\lambda, b_\lambda)$ and $\mu \mid \mathbf{y} \sim \mathrm{Gamma}(a_\mu, b_\mu)$, whose shapes are the prior shapes plus the counts and whose rate parameters are the prior rate parameters plus the hours, and the availability is a transformed Beta variable. Every answer is computed from $\hat{\mathbf{z}}$ and $\mathbf{\Sigma}$; the exact posterior, or draws of it where no closed form exists, gives the value beside it.

8.2.1 The posterior from the four rows

The state is the two-entry vector of equation (8.3), and the four rows act on it through the residual vector and the 4×2 Jacobian of equation (8.11). Each Newton step solves the normal equations

$$(\mathbf{J}^\mathsf{T}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^\mathsf{T}\mathbf{r}, \qquad \mathbf{J} = \mathbf{W}^{1/2}\,\frac{\partial \ln \bar{\mathbf{k}}}{\partial \mathbf{z}}, \qquad \mathbf{z} \leftarrow \mathbf{z} + \Delta\mathbf{z}, \tag{8.13}$$

with $\mathbf{J}$ and $\mathbf{r}$ evaluated at the current state. Since $\mathbf{J}^\mathsf{T}\mathbf{J}$ is diagonal at every state (Section 8.1.6), each step moves each entry by the Newton step of its own column.

The MAP state. The iteration stops where $\mathbf{J}^\mathsf{T}\mathbf{r} = \mathbf{0}$, that is where $b_\lambda e^{z_1} = a_\lambda$ and $b_\mu e^{z_2} = a_\mu$:

$$\hat{\mathbf{z}} = \begin{bmatrix} \hat{z}_1 \\ \hat{z}_2 \end{bmatrix} = \begin{bmatrix} \ln(a_\lambda / b_\lambda) \\ \ln(a_\mu / b_\mu) \end{bmatrix}, \tag{8.14}$$

the logarithms of the posterior means of the two exact gamma posteriors (Section 8.2.2). In the log rates the mode is at the posterior mean of each rate, not at its mode, because the change-of-variable factor has added one to each shape. When each record agrees with its prior in the center, $k/T_{\text{tot}} = a_{\lambda 0}/b_{\lambda 0}$ and $r/D = a_{\mu 0}/b_{\mu 0}$, the whole residual vector $\mathbf{r}(\hat{\mathbf{z}})$ is zero, not only its projections $\mathbf{J}^\mathsf{T}\mathbf{r}$: each row then expects exactly the count it observed.

The Jacobian, the information and the covariance. At the MAP state the expected counts are $\bar{k}_1 = b_{\lambda 0} a_\lambda / b_\lambda$ and $\bar{k}_2 = T_{\text{tot}}\, a_\lambda / b_\lambda$ for the two rows of z_1, and $\bar{k}_3 = b_{\mu 0} a_\mu / b_\mu$ and $\bar{k}_4 = D\, a_\mu / b_\mu$ for the two rows of z_2, and the nonzero entries of $\mathbf{J}(\hat{\mathbf{z}})$ in the pattern (8.5) are their square roots. The information and covariance matrices are

$$\mathbf{\Lambda} = \mathbf{J}^\mathsf{T}\mathbf{J} = \begin{bmatrix} \bar{k}_1 + \bar{k}_2 & 0 \\ 0 & \bar{k}_3 + \bar{k}_4 \end{bmatrix} = \begin{bmatrix} a_\lambda & 0 \\ 0 & a_\mu \end{bmatrix}, \qquad \mathbf{\Sigma} = \mathbf{\Lambda}^{-1} = \begin{bmatrix} 1/a_\lambda & 0 \\ 0 & 1/a_\mu \end{bmatrix}. \tag{8.15}$$

Each diagonal entry of $\mathbf{\Lambda}$ is the sum of the squared entries of its column of $\mathbf{J}$, the expected counts of the rows that feed it, and since $(b_{\lambda 0} + T_{\text{tot}})\, a_\lambda / b_\lambda = a_\lambda$ it is the posterior shape: $a_{\lambda 0} + k$ for z_1 and $a_{\mu 0} + r$ for z_2. The precisions of independent rows add, and each repair adds as much precision to z_2 as each failure adds to z_1. The Laplace posterior is $\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$; each log rate is read off as its entry $\hat{z}_i$ with the standard deviations $\sqrt{\Sigma_{11}} = 1/\sqrt{a_\lambda}$ and $\sqrt{\Sigma_{22}} = 1/\sqrt{a_\mu}$, and the correlation $\Sigma_{12}/\sqrt{\Sigma_{11}\Sigma_{22}}$ is zero. Minimizing each half of equation (8.10) in its own entry gives the same $\hat{z}_i$ and the same curvatures, as it must when the Hessian is diagonal. The gamma posteriors of Section 8.2.2 remain the exact results; $\mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$ is the Laplace approximation of their joint posterior in the log rates.

8.2.2 The exact posterior of the rates

The gamma posteriors. Write the prior on the repair rate as $\mu \sim \text{Gamma}(a_{\mu 0}, b_{\mu 0})$, with density $b_{\mu 0}^{a_{\mu 0}} \mu^{a_{\mu 0}-1} e^{-b_{\mu 0}\mu} / \Gamma(a_{\mu 0})$. Bayes' theorem with the repair likelihood (8.7) gives $p(\mu \mid r, D) \propto \mu^r e^{-\mu D} \times \mu^{a_{\mu 0}-1} e^{-b_{\mu 0}\mu} = \mu^{a_{\mu 0}+r-1} e^{-(b_{\mu 0}+D)\mu}$, the kernel of a gamma density with shape $a_{\mu 0} + r$ and rate parameter $b_{\mu 0} + D$. Normalizing with the gamma integral $\int_0^\infty x^{n-1} e^{-bx}\, dx = \Gamma(n)/b^n$, valid for any $n, b > 0$,

$$\mu \mid r, D \sim \text{Gamma}(a_\mu, b_\mu), \qquad a_\mu = a_{\mu 0} + r, \qquad b_\mu = b_{\mu 0} + D, \tag{8.16}$$

the update of Chapter 6 with the repairs in the place of the failures and the repair time in the place of the exposure. The failure part of equation (8.6) gives, as that chapter derived,

$$\lambda \mid k, T_{\text{tot}} \sim \text{Gamma}(a_\lambda, b_\lambda), \qquad a_\lambda = a_{\lambda 0} + k, \qquad b_\lambda = b_{\lambda 0} + T_{\text{tot}}, \tag{8.17}$$

and the two posteriors are independent by the factorization of equation (8.6). The objective (8.10) is the negative logarithm of their joint density written as a density of the log rates, and the minimum (8.14) lies at the logarithms of their means.

The share of two independent gamma variables. Let $\lambda \sim \text{Gamma}(a_\lambda, b_\lambda)$ and $\mu \sim \text{Gamma}(a_\mu, b_\mu)$ be independent, and scale each by its rate, $X = b_\lambda \lambda$ and $Y = b_\mu \mu$. A gamma variable times its rate has rate one, so $X \sim \text{Gamma}(a_\lambda, 1)$ and $Y \sim \text{Gamma}(a_\mu, 1)$, with joint density

$$p(x, y) = \frac{x^{a_\lambda - 1}\, y^{a_\mu - 1}\, e^{-(x+y)}}{\Gamma(a_\lambda)\, \Gamma(a_\mu)}, \qquad x, y > 0.$$

Change variables to the sum and the share, $s = x + y$ and $v = x/(x + y)$, so that $x = vs$ and $y = (1 - v)s$. The change-of-variable factor of the inverse map, the absolute value of its determinant, is

$$\left|\det \begin{bmatrix} \partial x/\partial v & \partial x/\partial s \\ \partial y/\partial v & \partial y/\partial s \end{bmatrix}\right| = \left|\det \begin{bmatrix} s & v \\ -s & 1-v \end{bmatrix}\right| = s(1-v) + sv = s,$$

and the joint density of (v, s) is

$$p(v,s) = \frac{(vs)^{a_\lambda - 1}\left((1-v)s\right)^{a_\mu - 1} \mathrm{e}^{-s}\, s}{\Gamma(a_\lambda)\,\Gamma(a_\mu)} = \frac{v^{a_\lambda-1}(1-v)^{a_\mu-1}}{\Gamma(a_\lambda)\,\Gamma(a_\mu)}\; s^{a_\lambda + a_\mu - 1}\,\mathrm{e}^{-s}.$$

It factorizes into a function of v and a function of s. Integrating out s with $\int_0^\infty s^{a_\lambda+a_\mu-1}\mathrm{e}^{-s}\,\mathrm{d}s = \Gamma(a_\lambda + a_\mu)$ and $B(a_\lambda, a_\mu) = \Gamma(a_\lambda)\Gamma(a_\mu)/\Gamma(a_\lambda + a_\mu)$ from Chapter 5,

$$V = \frac{b_\lambda \lambda}{b_\lambda \lambda + b_\mu \mu} \sim \mathrm{Beta}(a_\lambda, a_\mu), \tag{8.18}$$

independent of the sum $X + Y$, which is $\mathrm{Gamma}(a_\lambda + a_\mu, 1)$. The share of the rate-scaled failure rate in the rate-scaled total is a Beta variable whose parameters are the two posterior shapes, and Section 8.3.1 carries it into the availability through equation (8.32).

The moments of the log rates. The Laplace variance $1/a_\lambda$ of the log failure rate is the inverse curvature at the mode. The exact variance of the logarithm of a gamma variable has a closed form as well, and the two differ. Let $\lambda \sim \mathrm{Gamma}(a_\lambda, b_\lambda)$ and $X = b_\lambda \lambda \sim \mathrm{Gamma}(a_\lambda, 1)$, so that $\ln \lambda = \ln X - \ln b_\lambda$. The gamma function $\Gamma(n) = \int_0^\infty x^{n-1}\mathrm{e}^{-x}\,\mathrm{d}x$ is differentiated with respect to its argument, using $\partial x^{n-1}/\partial n = x^{n-1}\ln x$, and evaluated at $n = a_\lambda$:

$$\Gamma'(a_\lambda) = \int_0^\infty \ln x\; x^{a_\lambda - 1}\mathrm{e}^{-x}\,\mathrm{d}x = \Gamma(a_\lambda)\,\mathbb{E}[\ln X], \qquad \Gamma''(a_\lambda) = \Gamma(a_\lambda)\,\mathbb{E}\big[(\ln X)^2\big].$$

The mean is therefore $\mathbb{E}[\ln X] = \Gamma'(a_\lambda)/\Gamma(a_\lambda) = \psi(a_\lambda)$, the digamma function, and the variance is

$$\mathrm{Var}(\ln X) = \frac{\Gamma''(a_\lambda)}{\Gamma(a_\lambda)} - \left(\frac{\Gamma'(a_\lambda)}{\Gamma(a_\lambda)}\right)^2 = \frac{\mathrm{d}}{\mathrm{d}n}\,\frac{\Gamma'(n)}{\Gamma(n)}\bigg|_{n=a_\lambda} = \psi'(a_\lambda),$$

the trigamma function, since the derivative of a quotient Γ'/Γ is $\Gamma''/\Gamma - (\Gamma'/\Gamma)^2$. Subtracting the constant $\ln b_\lambda$ changes the mean and not the variance, and the same holds for the repair rate:

$$\mathbb{E}[\ln\lambda] = \psi(a_\lambda) - \ln b_\lambda, \quad \mathrm{Var}(\ln\lambda) = \psi'(a_\lambda), \qquad \mathbb{E}[\ln\mu] = \psi(a_\mu) - \ln b_\mu, \quad \mathrm{Var}(\ln\mu) = \psi'(a_\mu). \tag{8.19}$$

For large arguments the two functions have the standard expansions

$$\psi(n) = \ln n - \frac{1}{2n} - \frac{1}{12n^2} + \cdots, \qquad \psi'(n) = \frac{1}{n} + \frac{1}{2n^2} + \frac{1}{6n^3} + \cdots, \tag{8.20}$$

whose leading terms are the Laplace values: the center $\ln(a_\lambda/b_\lambda)$ and the variance $1/a_\lambda$. The Laplace approximation therefore understates the variance of each log rate by a term of order $1/(2a_\lambda^2)$, which fades as the records grow, and places its center about $1/(2a_\lambda)$ above the exact mean. In the logit the centers cancel: $\mathbb{E}[\xi] = \psi(a_\lambda) - \psi(a_\mu) + \ln c$, which equals $\ln c$, the Laplace center $\mathbf{g}^\mathsf{T}\hat{\mathbf{z}}$, whenever the two shapes are equal; and the exact variance of ξ is $\psi'(a_\lambda) + \psi'(a_\mu)$, because z_1 and z_2 are independent.

8.2.3 The posterior of the pump

The pump's inputs are the failure record, $k = 4$ failures in $T_{\text{tot}} = 2000$ operating hours under the prior Gamma$(2, 1000\,\text{h})$, and the repair record, $r = 4$ repairs of total duration $D = 40\,\text{h}$ under the prior Gamma$(2, 20\,\text{h})$.

The priors and the records. The prior Gamma$(2, 1000\,\text{h})$ of the failure rate has the mean $2/1000 = 0.002\,\text{h}^{-1}$. The prior (8.4) of the repair rate has the mean $2/20 = 0.1\,\text{h}^{-1}$ and the standard deviation $\sqrt{2}/20 = 0.0707\,\text{h}^{-1}$, with a central 95 % range of 0.012 to $0.279\,\text{h}^{-1}$, mean repair durations of about 3.6 to 83 hours. Both records agree with their priors in the center: the failure record's $4/2000 = 0.002\,\text{h}^{-1}$ equals the prior mean of λ, and the repair record's $r/D = 4/40 = 0.1\,\text{h}^{-1}$, the maximum-likelihood repair rate, equals the prior mean of μ, so the records narrow the two posteriors without moving their centers. With $a_\lambda = 2 + 4 = 6$, $b_\lambda = 1000 + 2000 = 3000\,\text{h}$, $a_\mu = 2 + 4 = 6$ and $b_\mu = 20 + 40 = 60\,\text{h}$, the objective (8.10) is $\mathcal{L}(\mathbf{z}) = 3000\,\mathrm{e}^{z_1} - 6z_1 + 60\,\mathrm{e}^{z_2} - 6z_2$.

The Newton iteration. Start at $\mathbf{z}^{(0)} = (-7,\ -3)^\mathsf{T}$, below both rates ($\lambda = 0.000912\,\text{h}^{-1}$, $\mu = 0.0498\,\text{h}^{-1}$). There every row expects fewer events than it observed, and

$$\bar{\mathbf{k}}(\mathbf{z}^{(0)}) = \begin{bmatrix} 0.9119 \\ 1.8238 \\ 0.9957 \\ 1.9915 \end{bmatrix}, \qquad \mathbf{r}(\mathbf{z}^{(0)}) = \begin{bmatrix} -1.1395 \\ -1.6115 \\ -1.0064 \\ -1.4233 \end{bmatrix}, \qquad \mathbf{J}(\mathbf{z}^{(0)}) = \begin{bmatrix} 0.9549 & 0 \\ 1.3505 & 0 \\ 0 & 0.9979 \\ 0 & 1.4112 \end{bmatrix},$$

so that the first step of equation (8.13) solves

$$\begin{bmatrix} 2.7356 & 0 \\ 0 & 2.9872 \end{bmatrix} \Delta\mathbf{z} = \begin{bmatrix} 3.2644 \\ 3.0128 \end{bmatrix}, \qquad \Delta\mathbf{z} = \begin{bmatrix} 1.1933 \\ 1.0086 \end{bmatrix}.$$

Table 8.3 lists the iteration: each row is one iterate, the columns give the two entries, the two entries of the gradient $\mathbf{J}^\mathsf{T}\mathbf{r}$ and the two diagonal entries of the Hessian. The off-diagonal entry of $\mathbf{J}^\mathsf{T}\mathbf{J}$ is zero at every iterate. The first step overshoots both entries, the second brings them back, and from the third the convergence is quadratic; the gradient is below 10^{-10} at the fifth iterate.

Table 8.3: Newton iteration of equation (8.13) for the pump's log rates from $\mathbf{z}^{(0)} = (-7,\ -3)^\mathsf{T}$. The Hessian $\mathbf{J}^\mathsf{T}\mathbf{J}$ is diagonal at every iterate.

Iteration [–]	z_1 [–]	z_2 [–]	$\partial\mathcal{L}/\partial z_1$ [–]	$\partial\mathcal{L}/\partial z_2$ [–]	$(\mathbf{J}^\mathsf{T}\mathbf{J})_{11}$ [–]	$(\mathbf{J}^\mathsf{T}\mathbf{J})_{22}$ [–]
0	−7.0000	−3.0000	−3.264	−3.013	2.7356	2.9872
1	−5.8067	−1.9914	3.022	2.190	9.0217	8.1899
2	−6.1417	−2.2588	0.454	0.268	6.4540	6.2683
3	−6.2120	−2.3016	1.6×10^{-2}	5.7×10^{-3}	6.0156	6.0057
4	−6.2146	−2.3026	2.0×10^{-5}	2.7×10^{-6}	6.0000	6.0000
5	−6.2146	−2.3026	3.4×10^{-11}	5.9×10^{-13}	6.0000	6.0000

The MAP state, the information and the covariance. By equation (8.14) the iteration ends at

$$\hat{\mathbf{z}} = \begin{bmatrix} \ln(6/3000) \\ \ln(6/60) \end{bmatrix} = \begin{bmatrix} \ln 0.002 \\ \ln 0.1 \end{bmatrix} = \begin{bmatrix} -6.2146 \\ -2.3026 \end{bmatrix}, \tag{8.21}$$

where the whole residual vector $\mathbf{r}(\hat{\mathbf{z}})$ is zero, because both records agree with their priors in the center. The expected counts are $\bar{\mathbf{k}} = (2,\ 4,\ 2,\ 4)^\mathsf{T}$: the prior rows' $1000 \times 0.002 = 2$ and

$20 \times 0.1 = 2$, the expected number of failures $2000 \times 0.002 = 4$ in the exposure, and the repair row's $40 \times 0.1 = 4$, the number of repairs. Equation (8.15) gives

$$\mathbf{J}(\hat{\mathbf{z}}) = \begin{bmatrix} 1.4142 & 0 \\ 2 & 0 \\ 0 & 1.4142 \\ 0 & 2 \end{bmatrix}, \qquad \mathbf{\Lambda} = \mathbf{J}^{\mathsf{T}}\mathbf{J} = \begin{bmatrix} 2+4 & 0 \\ 0 & 2+4 \end{bmatrix} = \begin{bmatrix} 6 & 0 \\ 0 & 6 \end{bmatrix}, \tag{8.22}$$

$$\mathbf{\Sigma} = \mathbf{\Lambda}^{-1} = \begin{bmatrix} 0.16667 & 0 \\ 0 & 0.16667 \end{bmatrix}.$$

The four repairs add exactly as much precision to z_2 as the four failures add to z_1. Each log rate is read off with $\sqrt{\Sigma_{ii}} = 1/\sqrt{6} = 0.408$, and the correlation of the two is 0.

The exact posteriors. Equations (8.17) and (8.16) give

$$\lambda \mid k, T_{\text{tot}} \sim \text{Gamma}(6,\ 3000\,\text{h}), \qquad \mu \mid r, D \sim \text{Gamma}(6,\ 60\,\text{h}),$$

the first of mean $a_\lambda/b_\lambda = 0.002\,\text{h}^{-1}$ and standard deviation $\sqrt{a_\lambda}/b_\lambda = 0.000816\,\text{h}^{-1}$. Table 8.4 sets the two side by side: each row is one rate, and the columns give its prior, its record, its posterior and the posterior mean, standard deviation and 95 % interval. Both posteriors have shape 6. The exact moments of the log rates, equation (8.19), use $\psi(6) = 1.7061$ and $\psi'(6) = 0.1813$.

Table 8.4: The two rate posteriors of the pump. Both are gamma with shape 6; the records are those of Table 8.1 and of Chapter 6.

Rate	Prior	Record	Posterior	Mean (h^{-1})	Std. dev. (h^{-1})	95 % interval (h^{-1})
failure λ	Gamma(2, 1000 h)	$k = 4$, $T_{\text{tot}} = 2000\,\text{h}$	Gamma(6, 3000 h)	0.0020	0.00082	[0.00073, 0.00389]
repair μ	Gamma(2, 20 h)	$r = 4$, $D = 40\,\text{h}$	Gamma(6, 60 h)	0.100	0.041	[0.037, 0.194]

The pair $(\hat{\mathbf{z}}, \mathbf{\Sigma})$ of equations (8.21) and (8.22) is what step 4 uses: every answer of the chapter is computed from it. Because this problem also admits the exact posterior, each answer has an exact value beside it, a function of the two gamma posteriors or of draws from them, and the gap between the two is reported with it.

8.3 The answers (step 4)

The five questions of Table 8.2 take three general forms. Each is answered from the posterior of Section 8.2: its engineering question, the general form of its category, the form it takes for the pump, the quantities of the posterior it uses, and the pump's numbers. The questions name the required availability $A_{\text{req}} = 0.97$ at the assurance $\kappa = 0.90$ and the year of 8760 calendar hours. Every answer is computed from the MAP state $\hat{\mathbf{z}}$ and the covariance $\mathbf{\Sigma}$ of step 3, and that is the value the chapter reports; the exact gamma and Beta posteriors, and two million draws from them for the quantities without a closed form, give the value set beside it, with the gap.

8.3.1 The repair rate and the availability

Engineering questions. What is the repair rate, and the mean time to repair? What fraction of the time will the pump be available, and how well is that known? The first asks for the

rate at which the crew restores this pump and for the expected duration of one repair, which maintenance planning uses; the second for the long-run availability (8.1) with its band.

General form. An estimation question about a quantity $f(\mathbf{z})$ derived from the state asks for its posterior distribution, the probability that the quantity lies at or below each value f^*,

$$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y}) = \int \mathbf{1}\big[f(\mathbf{z}) \le f^*\big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \tag{8.23}$$

summarized by a location, the mean or the median, a spread, and a central 95 % credible interval $[f_L, f_U]$ whose endpoints are the values at which equation (8.23) equals 0.025 and 0.975. An entry of the state is the case $f(\mathbf{z}) = z_i$. When f is a monotone function of a single linear combination of the entries, its quantiles are the images of that combination's quantiles, in the same order for an increasing function and in reverse order for a decreasing one.

Specialized form. *The repair rate.* The question is equation (8.23) with $f(\mathbf{z}) = \mu = \mathrm{e}^{z_2}$, an increasing function of the second entry alone, and $p(\mathbf{z} \mid \mathbf{y})$ the posterior of step 3, the Laplace normal $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of equations (8.14) and (8.15). The marginal of z_2 is $\mathcal{N}(\hat{z}_2, \Sigma_{22})$, read off as the second entry of $\hat{\mathbf{z}}$ and the second diagonal entry of $\boldsymbol{\Sigma}$, and the central interval of the repair rate is the interval of the entry mapped through the exponential; the mean time to repair, $\mathrm{MTTR} = 1/\mu = \mathrm{e}^{-z_2}$, is decreasing in z_2, and its interval is the same one mapped in reverse order:

$$\mu \in \big[\mathrm{e}^{\hat{z}_2 - 1.96\sqrt{\Sigma_{22}}},\ \mathrm{e}^{\hat{z}_2 + 1.96\sqrt{\Sigma_{22}}}\big], \qquad \mathrm{MTTR} \in \big[\mathrm{e}^{-\hat{z}_2 - 1.96\sqrt{\Sigma_{22}}},\ \mathrm{e}^{-\hat{z}_2 + 1.96\sqrt{\Sigma_{22}}}\big]. \tag{8.24}$$

The median of each is the image of the entry's center, $\mathrm{e}^{\hat{z}_2}$ and $\mathrm{e}^{-\hat{z}_2}$, and the mean of a log-normal variable adds half the variance in the exponent,

$$\mathbb{E}[\mu \mid \mathbf{y}] \approx \mathrm{e}^{\hat{z}_2 + \Sigma_{22}/2}, \qquad \mathbb{E}[\mathrm{MTTR} \mid \mathbf{y}] \approx \mathrm{e}^{-\hat{z}_2 + \Sigma_{22}/2}. \tag{8.25}$$

The values beside them come from the exact posterior. The repair rate has the gamma posterior $\mathrm{Gamma}(a_\mu, b_\mu)$, with $a_\mu = a_{\mu 0} + r$ and $b_\mu = b_{\mu 0} + D$ (Section 8.2.2), of mean a_μ / b_μ, standard deviation $\sqrt{a_\mu}/b_\mu$ and mode $(a_\mu - 1)/b_\mu$, with its interval from the gamma quantiles.

Maintenance planning speaks of the mean time to repair rather than of the rate, and its exact posterior is the distribution of the reciprocal of a gamma variable. Let $\mu \sim \mathrm{Gamma}(a_\mu, b_\mu)$ and write $x = 1/\mu$ for a value of the mean time to repair. The map is decreasing, with $|\mathrm{d}\mu/\mathrm{d}x| = 1/x^2$, so the density of x is the gamma density at $1/x$ times that change-of-variable factor:

$$p(x \mid r, D) = \frac{b_\mu^{a_\mu}}{\Gamma(a_\mu)}\, x^{-(a_\mu - 1)}\, \mathrm{e}^{-b_\mu / x}\, \frac{1}{x^2} = \frac{b_\mu^{a_\mu}}{\Gamma(a_\mu)}\, x^{-a_\mu - 1}\, \mathrm{e}^{-b_\mu / x}, \tag{8.26}$$

the inverse-gamma density with shape a_μ and scale b_μ. Its mean follows from the gamma integral after the substitution $\mu = 1/x$:

$$\mathbb{E}[\mathrm{MTTR} \mid r, D] = \mathbb{E}[\mu^{-1}] = \frac{b_\mu^{a_\mu}}{\Gamma(a_\mu)} \int_0^\infty \mu^{a_\mu - 2}\, \mathrm{e}^{-b_\mu \mu}\, \mathrm{d}\mu = \frac{b_\mu^{a_\mu}}{\Gamma(a_\mu)}\, \frac{\Gamma(a_\mu - 1)}{b_\mu^{a_\mu - 1}} = \frac{b_\mu}{a_\mu - 1},$$

using $\Gamma(a_\mu) = (a_\mu - 1)\Gamma(a_\mu - 1)$, and in the same way $\mathbb{E}[\mu^{-2}] = b_\mu^2 / ((a_\mu - 1)(a_\mu - 2))$, so that

$$\mathbb{E}[\mathrm{MTTR} \mid r, D] = \frac{b_\mu}{a_\mu - 1}, \qquad \mathrm{Var}(\mathrm{MTTR} \mid r, D) = \frac{b_\mu^2}{(a_\mu - 1)^2 (a_\mu - 2)}, \tag{8.27}$$

for $a_\mu > 1$ and $a_\mu > 2$ respectively. The mode of equation (8.26) is $b_\mu/(a_\mu + 1)$. Because the reciprocal is monotone, the quantiles of MTTR are the reciprocals of the quantiles of μ in reverse order.

The reciprocal of the posterior mean rate is b_μ/a_μ, while the posterior mean repair time is $b_\mu/(a_\mu - 1)$, which is larger. The two differ because $1/\mu$ is convex, and the expectation of a convex function exceeds the function of the expectation. The same holds for a prior: a prior described as centered on a typical repair time has a prior mean repair time $b_{\mu 0}/(a_{\mu 0} - 1)$, not the reciprocal $b_{\mu 0}/a_{\mu 0}$ of its mean rate, and a prior stated as a typical MTTR must say which of the two it means. Neither number is the availability's input: equation (8.1) needs the rates, and the availability below carries their whole posterior into it.

Common mistake

Reporting the reciprocal of the mean repair rate as the mean time to repair, or a mean time to repair as the reciprocal of a mean rate. Under the posterior Gamma(a_μ, b_μ) the first is b_μ/a_μ and the second $b_\mu/(a_\mu - 1)$; with fewer repairs the gap grows, and with one repair of duration d_1 and the prior Gamma$(2, 20\,\mathrm{h})$ the posterior mean repair time is $b_\mu/(a_\mu - 1) = (20 + d_1)/2$ against $(20 + d_1)/3$.

The availability. The question is equation (8.23) with $f(\mathbf{z}) = A$ of equation (8.1) and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of step 3. The unavailability (8.2) has odds $U/(1 - U) = \lambda/\mu$, so its logit is a linear function of the state, the difference of its two entries:

$$\xi = \ln \frac{U}{1-U} = \ln \lambda - \ln \mu = z_1 - z_2 = \mathbf{g}^\mathsf{T}\mathbf{z}, \qquad \mathbf{g} = \begin{bmatrix} 1 \\ -1 \end{bmatrix}, \qquad U = \frac{1}{1 + \mathrm{e}^{-\xi}}, \qquad A = \frac{1}{1 + \mathrm{e}^{\xi}}. \tag{8.28}$$

The logit ξ ranges over the whole real line while A stays in $(0, 1)$, and A decreases as ξ increases, so a quantile of ξ maps to the opposite quantile of A. The logit is the scale on which the derived quantity is linear in the state, and Chapter 5 already used it for a probability that must stay in $(0, 1)$. Here $f(\mathbf{z}) = A = 1/(1 + \mathrm{e}^{\mathbf{g}^\mathsf{T}\mathbf{z}})$, a decreasing function of one linear combination of the entries, and its posterior is the posterior of $\mathbf{g}^\mathsf{T}\mathbf{z}$.

The covariance enters through the Jacobian. Each of the four rows depends on one entry only, so the Jacobian of the normalized rows has, at every state, the pattern (8.5) of the graph, with each nonzero entry the square root of its row's expected count (Section 8.1.6). The off-diagonal entry of the information matrix is $(\mathbf{J}^\mathsf{T}\mathbf{J})_{12} = \sum_i J_{i1}J_{i2}$, and every product in the sum is zero because no row of $\mathbf{J}$ has two nonzero entries. The information matrix $\boldsymbol{\Lambda} = \mathbf{J}^\mathsf{T}\mathbf{J}$ is therefore diagonal, and so is its inverse $\boldsymbol{\Sigma}$, with $\Sigma_{ii} = 1/\Lambda_{ii}$: the zero covariance of the two log rates is a result of the rows, not an assumption about the rates. At the MAP state the diagonal entries are the posterior shapes, $\boldsymbol{\Lambda} = \mathrm{diag}(a_\lambda, a_\mu)$ (Section 8.2.1).

Because $\xi = \mathbf{g}^\mathsf{T}\mathbf{z}$ is linear, it is exactly normal under the Laplace posterior $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$, and the delta method of Chapter 4 gives its center, written $\hat{\xi}$, and its variance with no approximation beyond the Gaussian itself. The variance is the delta-method variance σ_f^2 of that chapter for $f = \xi$, written σ_ξ^2; it names the Laplace value $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$ only, and the exact spread is written $\mathrm{sd}(\xi \mid \mathbf{y})$. With $\Sigma_{12} = 0$

$$\hat{\xi} = \mathbf{g}^\mathsf{T}\hat{\mathbf{z}} = \hat{z}_1 - \hat{z}_2 = \ln \frac{a_\lambda/b_\lambda}{a_\mu/b_\mu}, \qquad \sigma_\xi^2 = \mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g} = \Sigma_{11} + \Sigma_{22} - 2\Sigma_{12} = \frac{1}{a_\lambda} + \frac{1}{a_\mu}, \tag{8.29}$$

the variance budget of the logit of the unavailability: one term per record and no cross term. Each term is the inverse of one column's information, $1/(a_{\lambda 0} + k)$ for the failure record and $1/(a_{\mu 0} + r)$ for the repair record. The two terms are equal when every failure is followed by a repair, $r = k$, and the prior shapes are equal; under those two conditions they stay equal however long the record runs. A record of failures alone, however long, adds to the column of z_1 only and leaves Σ_{22}, the repair term, where the prior and the repairs put it. The central interval of the availability is the logit's interval $\hat{\xi} \pm 1.96\,\sigma_\xi$ mapped through $A = 1/(1 + \mathrm{e}^{\xi})$ in reverse order, and its median is the image of the center, $1/(1 + \mathrm{e}^{\hat{\xi}})$. Its mean has no such image,

because A is not linear in ξ; the expectation over the one normal variable ξ is a Gauss–Hermite rule of Q nodes t_q and weights w_q, a weighted sum over values of the logit,

$$\mathbb{E}[h(\xi) \mid \mathbf{y}] \approx \int h(\xi)\, \mathcal{N}(\xi;\, \hat{\xi}, \sigma_\xi^2)\, \mathrm{d}\xi \approx \frac{1}{\sqrt{\pi}} \sum_{q=1}^{Q} w_q\, h(\hat{\xi} + \sqrt{2}\, \sigma_\xi\, t_q), \tag{8.30}$$

here with $h(\xi) = 1/(1 + \mathrm{e}^{\xi})$.

Principle 8.1. *The availability of a repairable component is limited by its repair record as much as by its failure record. With constant rates, the variance of the logit of its unavailability is* $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ *with* $\mathbf{g} = (1, -1)^{\mathsf{T}}$ *and a diagonal* $\boldsymbol{\Sigma} = \mathrm{diag}(1/a_\lambda,\ 1/a_\mu)$*, one term per record, and a record that adds failures without the repairs that follow them adds information to one column only and narrows only one of the two terms.*

The same covariance gives a band on U directly. By equation (8.28), $U = 1/(1 + \mathrm{e}^{-\xi})$ has $\mathrm{d}U/\mathrm{d}\xi = U(1 - U)$, so its gradient with respect to the state is the logit's gradient scaled by that factor,

$$\mathbf{g}_U = \frac{\partial U}{\partial \mathbf{z}} = U(1 - U) \begin{bmatrix} 1 \\ -1 \end{bmatrix}, \qquad \sqrt{\mathbf{g}_U^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g}_U} = U(1 - U)\, \sigma_\xi, \tag{8.31}$$

evaluated at $\hat{\mathbf{z}}$, the Laplace spread of U, with $\partial A/\partial \mathbf{z} = -\mathbf{g}_U$ and the same spread for A. The gradient $\mathbf{g}_U$ is not constant: it holds only near $\hat{\mathbf{z}}$, on a scale where U is bounded by zero and one. The band of the availability is therefore computed on the logit, where the gradient is constant, and mapped back; Section 8.4.1 shows what a normal on A built from $\mathbf{g}_U$ gets wrong.

The values beside them come from the exact posterior. The share $V = b_\lambda \lambda/(b_\lambda \lambda + b_\mu \mu)$ of the rate-scaled failure rate in the rate-scaled total has the distribution $\mathrm{Beta}(a_\lambda, a_\mu)$ (Section 8.2.2). The odds of the unavailability are $\lambda/\mu = (b_\mu/b_\lambda)\, V/(1 - V)$, so with the ratio of the two rate parameters

$$c = \frac{b_\mu}{b_\lambda}, \qquad \frac{U}{1 - U} = c\, \frac{V}{1 - V}, \qquad \xi = \ln c + \ln \frac{V}{1 - V}. \tag{8.32}$$

The unavailability is a monotone increasing function of V, and a quantile v_q of the Beta distribution maps to the quantile of the unavailability

$$U_q = \frac{c\, v_q}{1 - v_q + c\, v_q}, \qquad A_{1-q} = 1 - U_q. \tag{8.33}$$

Every band and quantile of the availability is then a Beta quantile, with no approximation beyond the model. The posterior mean is the exception: $\mathbb{E}[A \mid \mathbf{y}]$ is the integral of $1 - U(v)$ against the Beta density, a one-dimensional integral without a closed form in general, evaluated numerically. The density of A follows by the change of variable. With $v(x) = (1 - x)/\delta(x)$ and $1 - v(x) = cx/\delta(x)$, where $\delta(x) = 1 - (1 - c)x$, the derivative is $\mathrm{d}v/\mathrm{d}x = -c/\delta(x)^2$, and the density of A at x is the Beta density at $v(x)$ times $|\mathrm{d}v/\mathrm{d}x|$:

$$p_A(x \mid \mathbf{y}) = \frac{1}{B(a_\lambda, a_\mu)} \left(\frac{1 - x}{\delta(x)}\right)^{a_\lambda - 1} \left(\frac{c\, x}{\delta(x)}\right)^{a_\mu - 1} \frac{c}{\delta(x)^2}, \qquad 0 < x < 1. \tag{8.34}$$

When the two shapes are equal, the Beta distribution is symmetric about one half, so the median of V is 0.5, the median odds of the unavailability are c, and the median availability is $1/(1 + c)$. That is also the plug-in value of equation (8.1) at the posterior means, but only because the shapes are equal: in general the plug-in value has odds $(a_\lambda/b_\lambda)/(a_\mu/b_\mu) = c\, a_\lambda/a_\mu$, while the median odds are c times the median of $V/(1 - V)$, which is one only when $a_\lambda = a_\mu$ (Exercise 8.3 works a case with unequal shapes). Neither is the posterior mean, which lies below both when the posterior has its long tail toward low availabilities. The exact variance of the logit is $\mathrm{Var}(\xi \mid \mathbf{y}) = \psi'(a_\lambda) + \psi'(a_\mu)$ (Section 8.2.2), the counterpart of $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$, because the two log rates are independent; its two terms $\mathrm{Var}(\ln \lambda \mid \mathbf{y})$ and $\mathrm{Var}(\ln \mu \mid \mathbf{y})$ are the exact variance budget.

Building blocks. The MAP state $\hat{\mathbf{z}} = (-6.2146,\ -2.3026)^{\mathsf{T}}$ and the covariance $\boldsymbol{\Sigma} = \operatorname{diag}(1/6,\ 1/6)$ of equations (8.21) and (8.22); the gradient $\mathbf{g} = (1, -1)^{\mathsf{T}}$ of the logit, equation (8.28); the Gauss–Hermite rule (8.30); and, for the exact values beside them, the gamma posterior parameters (a_λ, b_λ) and (a_μ, b_μ), the Beta share V, the trigamma values $\psi'(a)$ (Section 8.2.2) and the gamma and Beta quantile functions.

The repair rate and the mean time to repair. With $\hat{z}_2 = -2.3026$ and $\sqrt{\Sigma_{22}} = 1/\sqrt{6} = 0.4082$, the repair rate at the MAP state, the Laplace median, is $\mathrm{e}^{-2.3026} = 0.100\,\mathrm{h}^{-1}$, and equation (8.24), $-2.3026 \pm 1.96 \times 0.4082 = -2.3026 \pm 0.8002$, gives

$$\mu \in [0.045,\ 0.223]\,\mathrm{h}^{-1}, \qquad \mathrm{MTTR} \in [4.5,\ 22.3]\,\mathrm{h}$$

with 95 % posterior probability, about the median 10.0 h of the mean time to repair. By equation (8.25) the means are $0.100\,\mathrm{e}^{1/12} = 0.1087\,\mathrm{h}^{-1}$ and $10.0\,\mathrm{e}^{1/12} = 10.87\,\mathrm{h}$, with the standard deviations $0.0463\,\mathrm{h}^{-1}$ and 4.63 h.

The values beside them come from the exact posterior. From the building blocks $a_\mu = 6$ and $b_\mu = 60\,\mathrm{h}$,

$$\begin{gathered} \mu \mid r, D \sim \mathrm{Gamma}(6,\ 60\,\mathrm{h}), \\ \mathbb{E}[\mu \mid r, D] = \frac{a_\mu}{b_\mu} = 0.100\,\mathrm{h}^{-1}, \qquad \mathrm{sd}(\mu \mid r, D) = \frac{\sqrt{a_\mu}}{b_\mu} = 0.0408\,\mathrm{h}^{-1}. \end{gathered} \tag{8.35}$$

The posterior mode is $(a_\mu - 1)/b_\mu = 5/60 = 0.0833\,\mathrm{h}^{-1}$. The central 95 % interval, from the gamma quantiles, is

$$\mu \in [0.0367,\ 0.1945]\,\mathrm{h}^{-1} \text{ with } 95\,\% \text{ posterior probability.} \tag{8.36}$$

The standard deviation has fallen by about 40 % from the prior's $0.0707\,\mathrm{h}^{-1}$, the 95 % range has narrowed by a factor of about 1.7, and the relative standard deviation is $1/\sqrt{6} = 0.41$.

The mean time to repair has, by equation (8.27), the posterior mean $\mathbb{E}[\mathrm{MTTR} \mid r, D] = 60/5 = 12.0\,\mathrm{h}$ and the standard deviation $60/(5 \times 2) = 6.0\,\mathrm{h}$. The mode of equation (8.26) is $b_\mu/(a_\mu + 1) = 8.6\,\mathrm{h}$ and the median 10.6 h. The reciprocals of the quantiles (8.36) give

$$\mathrm{MTTR} \in [5.1,\ 27.2]\,\mathrm{h} \text{ with } 95\,\% \text{ posterior probability.} \tag{8.37}$$

The reciprocal of the posterior mean rate is 10 h, against the posterior mean repair time of 12 h; for the prior (8.4) the two are 10 h and $b_{\mu 0}/(a_{\mu 0} - 1) = 20/1 = 20\,\mathrm{h}$. The gaps between the routes are of the size the skew of a log rate gives at a posterior shape of six: the median mean time to repair is 0.6 h lower on the Laplace route and the mean 1.13 h lower, and the band of the repair rate lies higher, by $0.008\,\mathrm{h}^{-1}$ at the lower end and $0.029\,\mathrm{h}^{-1}$ at the upper.

The availability and its band. By equation (8.29) the logit has the center $\hat{\xi} = -6.2146 + 2.3026 = -3.912$ and the variance $\sigma_\xi^2 = 0.16667 + 0.16667 = 0.3333$, so $\sigma_\xi = 0.577$. Its interval $-3.912 \pm 1.96 \times 0.577 = [-5.044,\ -2.780]$ maps in reverse order to the band

$$A \in [0.942,\ 0.994] \text{ with } 95\,\% \text{ posterior probability,} \qquad \text{median } 0.980. \tag{8.38}$$

Table 8.5 evaluates the rule (8.30) with five nodes. Each row is a node: the first two columns give the node t_q and its weight $w_q/\sqrt{\pi}$, the third the logit $\xi_q = \hat{\xi} + \sqrt{2}\,\sigma_\xi t_q$ at which the sum looks, the fourth the unavailability U_q there, and the last two the two functions the questions average, the availability $A_q = 1 - U_q$ and the annual downtime $8760\,\mathrm{h} \times U_q$ of the question below. The middle node is the MAP state.

Table 8.5: The Gauss–Hermite rule (8.30) with five nodes over the Laplace posterior $\xi \mid \mathbf{y} \approx \mathcal{N}(-3.912,\ 0.577^2)$ of the logit of the unavailability.

Node t_q	Weight $w_q/\sqrt{\pi}$	Logit ξ_q	U_q	A_q	Downtime [h/yr]
−2.0202	0.01126	−5.5615	0.00383	0.99617	33.5
−0.9586	0.22208	−4.6947	0.00906	0.99094	79.4
0	0.53333	−3.9120	0.01961	0.98039	171.8
0.9586	0.22208	−3.1294	0.04191	0.95809	367.2
2.0202	0.01126	−2.2626	0.09427	0.90573	825.8
Weighted sum				0.9771	200.4

The weighted sum of the fifth column is

$$\mathbb{E}[A \mid \mathbf{y}] \approx 0.01126 \times 0.99617 + 0.22208 \times 0.99094 + \cdots + 0.01126 \times 0.90573 = 0.9771,$$

below the median; forty nodes change no digit shown.

The values beside them come from the exact posterior. The posteriors Gamma(6, 3000 h) and Gamma(6, 60 h) give $V \sim \text{Beta}(6,6)$ and $c = 60/3000 = 0.02$. The median availability is $1/(1+c) = 3000/3060 = 0.98039$, equal to the plug-in value $0.1/0.102$ because the shapes are equal, and equal to the Laplace median, the gap being zero. By equation (8.33) with the Beta quantiles $v_{0.025} = 0.23379$ and $v_{0.975} = 0.76621$ of Beta(6, 6),

$$U_{0.025} = \frac{0.02 \times 0.23379}{0.76621 + 0.02 \times 0.23379} = 0.006066,$$
$$U_{0.975} = \frac{0.02 \times 0.76621}{0.23379 + 0.02 \times 0.76621} = 0.06151,$$

so that

$$A \in [0.938,\ 0.994] \text{ with } 95\,\% \text{ posterior probability}, \qquad \text{median } 0.980. \tag{8.39}$$

The posterior mean, obtained by integrating $1 - U$ against the Beta(6, 6) density numerically, is $\mathbb{E}[A \mid \mathbf{y}] = 0.9768$, below the median and below the plug-in value 0.980, and 0.0003 below the Laplace mean. The lower end of the Laplace band is 0.004 higher than the exact one and the upper ends agree to the digits shown, because in the logit the center offsets of the two log rates cancel when the shapes are equal. With $1/B(6,6) = 11!/(5!\,5!) = 2772$, the density (8.34) is $2772\,c^6 x^5 (1-x)^5/(1-0.98x)^{12}$. The exact variance budget of the logit is $\text{Var}(\ln\lambda \mid \mathbf{y}) + \text{Var}(\ln\mu \mid \mathbf{y}) = \psi'(6) + \psi'(6) = 0.1813 + 0.1813 = 0.3626$, half from each record, and $\text{sd}(\xi \mid \mathbf{y}) = 0.602$, 0.025 above σ_ξ.

Figure 8.2 draws the posterior of the availability. The horizontal axis is the availability, the vertical axis its posterior density. The dotted blue curve is the Laplace density, the normal on ξ transformed to A; the green curve the exact density of equation (8.34); the dotted gray curve the normal on ξ with the exact variance. The vertical red line is the requirement $A = 0.97$, and the shaded area under the Laplace density to its right is the probability that the requirement is met, computed in Section 8.3.2. The exact density is skewed to the left, bounded above by one, and peaks near 0.986, above the median, with a long tail toward lower availabilities; the two normal curves are read in Section 8.4.1.

8.3.2 The requirement on the availability

Engineering question. Is the availability at least 0.97 with stated assurance? A service agreement requires the pump to be available at least 97 % of the time, and the operator may be asked to show that this holds with the posterior probability $\kappa = 0.90$, the assurance level of Chapter 5.

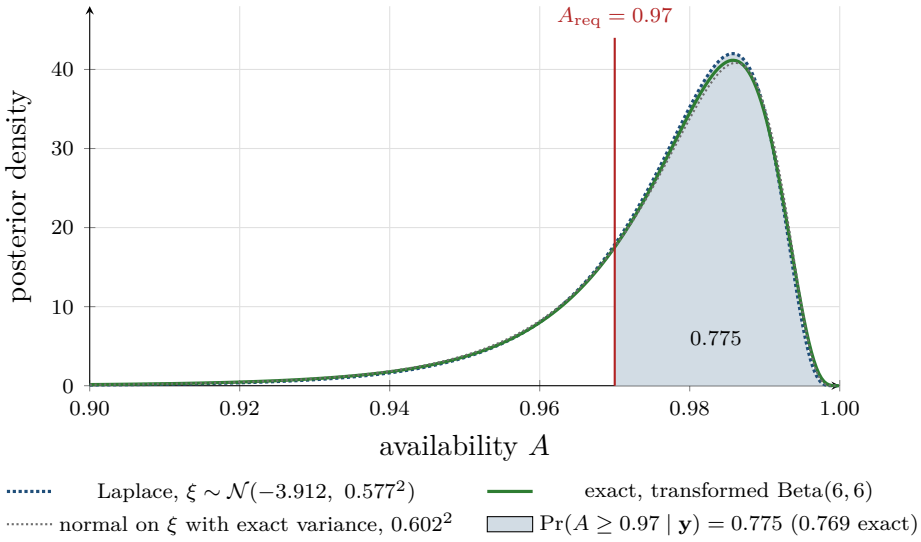


Figure 8.2: The posterior density of the pump's availability. Exact: equation (8.34), $2772\,(0.02x/\delta)^5((1-x)/\delta)^5\,0.02/\delta^2$ with $\delta = 1 - 0.98x$. Laplace: the normal $\mathcal{N}(\ln 0.02,\ 1/3)$ on $\xi = \ln((1-x)/x)$ mapped to x, density $\varphi((\xi + 3.912)/0.5774)/(0.5774\,x(1-x))$ with φ the standard normal density; dotted, the same with standard deviation 0.6022. The shaded area right of 0.97 is 0.775 under the Laplace density; under the exact density it is 0.769, and the area to its left 0.231.

General form. A prediction under a threshold is a tail probability of a derived quantity under the posterior, compared with an assurance level:

$$\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}) = \int \mathbf{1}[f(\mathbf{z}) \geq f^*]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \qquad \text{accept if } \Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}) \geq \kappa. \quad (8.40)$$

Here f^* is the requirement level and the event $f(\mathbf{z}) \geq f^*$ is the requirement holding; a requirement that a quantity stay below an upper limit is the same form with the sign of the quantity reversed, and the probability of the complementary event, the shortfall, is one minus the tail. The rule is applied to the requirement, never to the shortfall.

Specialized form. One choice enters here: the requirement is read on the steady-state availability, the long-run fraction of time up, as the agreement states it. That makes the event a half-plane in the two log rates; an agreement measured over a shift or a day after each repair would need the transient instead, which stays away from the steady state for about two days (Section 8.4.4).

The question is equation (8.40) with $f(\mathbf{z}) = A$, $f^* = A_{\text{req}}$, $\kappa = 0.90$ and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of step 3. The event $A \geq A_{\text{req}}$ is the requirement holding, its complement $A < A_{\text{req}}$ the shortfall. Since A decreases with the logit $\xi = \mathbf{g}^{\mathsf{T}}\mathbf{z}$ of equation (8.28), the requirement becomes an upper limit on that linear combination of the state, still the event of the requirement holding:

$$A \geq A_{\text{req}} \iff \mathbf{g}^{\mathsf{T}}\mathbf{z} = \xi \leq \xi^* = \ln\frac{1 - A_{\text{req}}}{A_{\text{req}}}. \quad (8.41)$$

In the plane of the two entries the event is a half-plane. Its boundary $\mathbf{g}^{\mathsf{T}}\mathbf{z} = \xi^*$ is the line $z_2 = z_1 + \ln(A_{\text{req}}/(1 - A_{\text{req}}))$ of slope one, a line of constant availability, and the gradient $\mathbf{g} = (1, -1)^{\mathsf{T}}$ is its normal: moving both rates by the same factor leaves the availability unchanged, and only the component of the state along $\mathbf{g}$ changes it. Under the Laplace posterior

the probability of the half-plane is the normal tail of the scalar $\mathbf{g}^\top\mathbf{z}$, whose argument is the distance from $\hat{\mathbf{z}}$ to the line, measured along $\mathbf{g}$ in standard deviations,

$$\Pr(A \geq A_{\text{req}} \mid \mathbf{y}) \approx \Phi\left(\frac{\xi^* - \mathbf{g}^\top\hat{\mathbf{z}}}{\sqrt{\mathbf{g}^\top\boldsymbol{\Sigma}\mathbf{g}}}\right), \tag{8.42}$$

with Φ the standard normal cumulative distribution; only the tail is an approximation, since the threshold maps exactly. The same tail with the exact variance $\psi'(a_\lambda) + \psi'(a_\mu)$ in place of $\mathbf{g}^\top\boldsymbol{\Sigma}\mathbf{g}$ separates the error of the variance from the error of the normal shape. The value beside it comes from the exact posterior: the requirement is $U/(1-U) \leq (1-A_{\text{req}})/A_{\text{req}}$, which by equation (8.32) is $cV \leq (1-V)(1-A_{\text{req}})/A_{\text{req}}$, and solving for V,

$$\Pr(A \geq A_{\text{req}} \mid \mathbf{y}) = \Pr(V \leq v^*) = I_{v^*}(a_\lambda, a_\mu), \qquad v^* = \frac{1 - A_{\text{req}}}{(1 - A_{\text{req}}) + c\,A_{\text{req}}}, \tag{8.43}$$

with $I_x(\cdot,\cdot)$ the regularized incomplete Beta function of Chapter 5. The same formula with the prior shapes and the prior ratio $c_0 = b_{\mu 0}/b_{\lambda 0}$ in place of the posterior ones gives the support under the priors alone, against which the contribution of the records is read.

Building blocks. The center $\mathbf{g}^\top\hat{\mathbf{z}} = -3.912$ and the variance $\mathbf{g}^\top\boldsymbol{\Sigma}\mathbf{g} = 0.3333$ of equation (8.29), from $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$ of step 3; the threshold ξ^* of equation (8.41); Φ; and, for the exact value beside it, the ratio c, the shapes a_λ, a_μ (Section 8.2.2) and the incomplete Beta function.

The requirement of 0.97. By equation (8.41), $\xi^* = \ln(0.03/0.97) = -3.4761$, and the tail (8.42) is

$$\Pr(A \geq 0.97 \mid \mathbf{y}) \approx \Phi\left(\frac{-3.4761 + 3.9120}{0.5774}\right) = \Phi(0.755) = 0.775,$$

the shaded area of Figure 8.2, below the assurance 0.90. The value beside it comes from the exact posterior. By equation (8.43) with $A_{\text{req}} = 0.97$ and $c = 0.02$,

$$v^* = \frac{0.03}{0.03 + 0.02 \times 0.97} = \frac{0.03}{0.0494} = 0.60729,$$

and

$$\Pr(A \geq 0.97 \mid \mathbf{y}) = I_{0.60729}(6, 6) = 0.769, \tag{8.44}$$

0.006 lower, and also below the assurance: the two routes differ in the support, not in the decision. Under the priors alone, the same tail with the two prior rows' variance $1/2 + 1/2$ gives $\Phi(0.4359) = 0.669$, and exactly $V \sim \text{Beta}(2, 2)$ with $c_0 = 20/1000 = 0.02$, and the formula (8.43) gives 0.658: the records raise the support without moving the median, by narrowing the posterior.

8.3.3 Downtime, outage frequency and the transient

Engineering questions. How many hours per year will the pump be down, and how often will it fail? How available is the pump in the hours after a repair? The first asks for two long-run quantities, the annual downtime and the annual number of outages; the second for the availability at a time t after the pump returns to service, before the long run is reached.

General form. The expectation of a derived quantity over the posterior,

$$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}] = \int f(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \tag{8.45}$$

with its quantiles from equation (8.23). When f is monotone in a scalar function of the state whose posterior is known, its quantiles are the images of that scalar's quantiles; otherwise the expectation and the quantiles are read from draws $\mathbf{z}^{(s)}$, $s = 1, \ldots, S$, of the posterior, as the mean and the empirical quantiles of $f(\mathbf{z}^{(s)})$.

Specialized form. Two choices enter here. The pump is described by the two-state model with constant rates, every repair returning it as good as new: this is what makes the downtime, the outage frequency and the transient closed functions of the two rates, and it fails if the failure rate changes with age or a repair leaves the pump worse than before (Section 8.6.6). The transient uses the exponential repair time itself, not only its mean: this gives the curve after a repair its single time constant, and it is the answer most changed if the repairs are as regular as the four recorded ones, which Section 8.4.3 examines.

The annual downtime. The question is equation (8.45), with its quantiles from equation (8.23), with $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal of step 3 and $f(\mathbf{z}) = t_{\mathrm{down}}$: over a year of 8760 calendar hours the pump is down, in the long run, for

$$t_{\mathrm{down}} = 8760\,\mathrm{h} \times U, \tag{8.46}$$

a monotone increasing function of U and so of ξ. Its quantiles are 8760 h times the quantiles of U, from the logit band of equation (8.29), and its mean is the rule (8.30) with $h(\xi) = 8760\,\mathrm{h}/(1+\mathrm{e}^{-\xi})$. The values beside them come from the exact posterior: equation (8.33) for the quantiles, and for the mean $8760\,\mathrm{h} \times \mathbb{E}[U \mid \mathbf{y}]$, the one-dimensional integral of $U(v)$ against the Beta density. The band is skewed, and the skew follows from the logit. When U is small, $U \approx \mathrm{e}^{\xi}$, so a band that is nearly symmetric in ξ, of half-width 1.96 standard deviations of the logit, is nearly symmetric on a logarithmic scale of hours: the downtime is known within a factor $\mathrm{e}^{1.96\,\sigma_\xi}$ in either direction, and on the scale of hours that symmetry becomes a long upper tail, with the mean above the median.

The outage frequency. The question is equation (8.45) again, with $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal of step 3 and $f(\mathbf{z}) = \nu_{\mathrm{out}}$. The pump fails only while it is up. In the long run the expected number of failures per calendar hour is therefore the failure rate times the fraction of time up,

$$\nu_{\mathrm{out}} = \lambda A = \frac{\lambda\mu}{\lambda+\mu}. \tag{8.47}$$

The identity $\nu_{\mathrm{out}}/\mu = \lambda/(\lambda+\mu) = U$ connects the two answers: the annual downtime is the number of outages per year times the mean duration of each. The outage frequency is a derived quantity of the state. Its logarithm, $\ln \nu_{\mathrm{out}} = z_1 + z_2 - \ln(\mathrm{e}^{z_1} + \mathrm{e}^{z_2})$, has the gradient

$$\mathbf{g}_\nu = \frac{\partial \ln \nu_{\mathrm{out}}}{\partial \mathbf{z}} = \begin{bmatrix} 1 - \lambda/(\lambda+\mu) \\ 1 - \mu/(\lambda+\mu) \end{bmatrix} = \begin{bmatrix} A \\ U \end{bmatrix}, \tag{8.48}$$

and its band is $\exp(\ln \hat{\nu}_{\mathrm{out}} \pm 1.96\sqrt{\mathbf{g}_\nu^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g}_\nu})$, with $\hat{\nu}_{\mathrm{out}}$ the outage frequency at $\hat{\mathbf{z}}$, the Laplace median. Its mean is a sum over the two entries, since ν_{out} is not a function of ξ alone: the rule (8.30) applied along each entry in turn, with $\Sigma_{12} = 0$,

$$\mathbb{E}[h(\mathbf{z}) \mid \mathbf{y}] \approx \sum_{q=1}^{Q} \sum_{r=1}^{Q} \frac{w_q w_r}{\pi}\, h(\hat{z}_1 + \sqrt{2\Sigma_{11}}\, t_q,\ \hat{z}_2 + \sqrt{2\Sigma_{22}}\, t_r). \tag{8.49}$$

The gradient puts almost all of its weight on z_1 when U is small, and that is where the Laplace normal is least accurate: it centers $\ln\lambda$ at $\ln(a_\lambda/b_\lambda)$, above its exact mean $\psi(a_\lambda)-\ln b_\lambda$, with the variance $1/a_\lambda$ rather than $\psi'(a_\lambda)$ (Section 8.2.2). In the logit $\xi=\ln\lambda-\ln\mu$ the two center offsets are equal when the shapes are equal, and they cancel; in a quantity that depends on one log rate they do not. The value beside it comes from the exact posterior: the outage frequency has no closed-form posterior there, and draws of the two gamma posteriors are pushed through equation (8.47). The gradients also say which rate governs which answer: $\mathbf{g}=(1,-1)^{\mathsf{T}}$ weighs the two entries equally in the downtime, while the second entry of $\mathbf{g}_\nu$ is U, so a crew that repairs faster, a larger z_2, shortens the downtime and changes the log outage frequency by only U per unit of z_2.

The availability after a repair. The question is equation (8.45) with $f(\mathbf{z})=f_t(\mathbf{z})=A(t)$ and $p(\mathbf{z}\mid\mathbf{y})$ the Laplace normal of step 3. The availability at a time t is the probability that the pump is up at t, $A(t)=\Pr(\text{up at } t)$. The two-state model of Figure 8.1(b) gives it from the rates by a balance of probability over a short interval Δt. Write $P_{\text{up}}(t)$ and $P_{\text{down}}(t)$ for the probabilities of the two states at t. The pump is up at $t+\Delta t$ if it was up at t and did not fail in the interval, or was down at t and was repaired in it; it is down at $t+\Delta t$ if it was down and was not repaired, or was up and failed. Under constant rates the probability of a failure in Δt given up is $\lambda\Delta t+o(\Delta t)$, and of a repair given down $\mu\Delta t+o(\Delta t)$, so by the law of total probability

$$\begin{aligned}P_{\text{up}}(t+\Delta t)&=P_{\text{up}}(t)\,(1-\lambda\Delta t)+P_{\text{down}}(t)\,\mu\Delta t+o(\Delta t),\\ P_{\text{down}}(t+\Delta t)&=P_{\text{down}}(t)\,(1-\mu\Delta t)+P_{\text{up}}(t)\,\lambda\Delta t+o(\Delta t).\end{aligned}$$

Subtract the values at t, divide by Δt and let it shrink:

$$\frac{\mathrm{d}P_{\text{up}}}{\mathrm{d}t}=-\lambda P_{\text{up}}+\mu P_{\text{down}},\qquad \frac{\mathrm{d}P_{\text{down}}}{\mathrm{d}t}=\lambda P_{\text{up}}-\mu P_{\text{down}},\tag{8.50}$$

the forward equations of the two-state model. In matrix form, with $\mathbf{P}(t)=(P_{\text{up}},P_{\text{down}})^{\mathsf{T}}$ and the generator

$$\mathbf{Q}=\begin{bmatrix}-\lambda & \lambda\\ \mu & -\mu\end{bmatrix},\qquad \frac{\mathrm{d}\mathbf{P}}{\mathrm{d}t}=\mathbf{Q}^{\mathsf{T}}\mathbf{P},$$

whose rows hold the rates out of each state. $\mathbf{Q}^{\mathsf{T}}$ has the eigenvalues 0 and $-(\lambda+\mu)$, the second equal to its trace: the eigenvector of 0, scaled to sum to one, is $(\mu,\ \lambda)^{\mathsf{T}}/(\lambda+\mu)$, the steady state, and the other mode decays at the rate $\lambda+\mu$. The two right-hand sides sum to zero, so $P_{\text{up}}+P_{\text{down}}$ keeps its initial value of one. With $A(t)=P_{\text{up}}(t)$ and $P_{\text{down}}(t)=1-A(t)$, the first equation becomes one equation in the availability,

$$\frac{\mathrm{d}A(t)}{\mathrm{d}t}=-\lambda\,A(t)+\mu\,(1-A(t))=\mu-(\lambda+\mu)\,A(t).\tag{8.51}$$

It is a linear first-order equation with constant coefficients. Its constant solution is $A=\mu/(\lambda+\mu)$, and the deviation from it, $A(t)-\mu/(\lambda+\mu)$, obeys $\mathrm{d}(\cdot)/\mathrm{d}t=-(\lambda+\mu)(\cdot)$ and therefore decays exponentially. For a pump that is up at $t=0$, just returned from a repair, $A(0)=1$ and the initial deviation is $1-\mu/(\lambda+\mu)=\lambda/(\lambda+\mu)$:

$$A(t)=\frac{\mu}{\lambda+\mu}+\frac{\lambda}{\lambda+\mu}\,\mathrm{e}^{-(\lambda+\mu)t}.\tag{8.52}$$

For a pump that has just failed, $A(0)=0$ and $A(t)=\mu(1-\mathrm{e}^{-(\lambda+\mu)t})/(\lambda+\mu)$. The companion volume treats the general multistate version of this model with its generator matrix; the two-state case is all this chapter needs.

As $t\to\infty$ the exponential in equation (8.52) vanishes, whatever the initial state, and the availability tends to

$$\lim_{t\to\infty}A(t)=\frac{\mu}{\lambda+\mu},$$

the steady-state availability (8.1). Dividing numerator and denominator by $\lambda\mu$ writes it as $(1/\lambda)/(1/\lambda + 1/\mu) = \mathrm{MTTF}/(\mathrm{MTTF} + \mathrm{MTTR})$, the mean up time over the mean cycle. In that second form it is a standard result for any alternating renewal process, whatever the distributions of the up and down durations: the long-run fraction of time up is the mean up duration over the mean cycle. The steady-state availability therefore depends on the repair-time distribution only through its mean, a fact that Section 8.4.3 uses. The deviation from the steady state decays as e^{-t/t_c}, with t_c the reciprocal of the magnitude of the nonzero eigenvalue $-(\lambda + \mu)$ of $\mathbf{Q}^{\mathsf{T}}$,

$$t_c = \frac{1}{\lambda + \mu}. \tag{8.53}$$

When the repair rate is much larger than the failure rate, the time constant is set almost entirely by the repair and is slightly shorter than the mean repair time $1/\mu$. After $t_c \ln 100$ the deviation is one percent of its initial value, from either starting point. For a horizon τ starting just after a repair, the interval availability is the average of equation (8.52),

$$\bar{A}(\tau) = \frac{1}{\tau}\int_0^\tau A(t)\,\mathrm{d}t = A + U\,\frac{t_c}{\tau}(1 - \mathrm{e}^{-\tau/t_c}), \tag{8.54}$$

which tends to A as the horizon grows beyond a few time constants.

The transient at $\hat{\mathbf{z}}$, where the rates are also the two posterior means, is a plug-in curve. The predicted availability at a time t is the expectation of equation (8.45) with $f_t(\mathbf{z}) = A(t)$ of equation (8.52), the probability of the up state at t averaged over the posterior of the rates, $\int \Pr(\text{up at } t \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\,\mathrm{d}\mathbf{z}$, and its band holds the central 95 % of $f_t(\mathbf{z})$ at each time. The average is the sum (8.49) with $h(\mathbf{z}) = A(t)$, and the band follows from the gradient of $A(t)$ at $\hat{\mathbf{z}}$, carried, as for the steady state, on the logit of the unavailability at t, $\xi_t = \ln\big((1 - A(t))/A(t)\big)$:

$$\mathbf{g}_t = \frac{\partial \xi_t}{\partial \mathbf{z}} = -\frac{1}{A(t)(1 - A(t))}\,\frac{\partial A(t)}{\partial \mathbf{z}}, \qquad \xi_t \in \xi_t(\hat{\mathbf{z}}) \pm 1.96\sqrt{\mathbf{g}_t^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_t}, \tag{8.55}$$

mapped back through $A(t) = 1/(1 + \mathrm{e}^{\xi_t})$. As t grows, $\mathbf{g}_t$ tends to the gradient $\mathbf{g} = (1, -1)^{\mathsf{T}}$ of equation (8.28); at $t = 0$ the band has zero width, since a pump just repaired is up whatever its rates. The values beside them come from the exact posterior: neither has a closed form there, and both are read from draws of the two gamma posteriors.

Building blocks. The MAP state $\hat{\mathbf{z}}$ and the covariance $\boldsymbol{\Sigma}$ of equations (8.21) and (8.22), with the gradients $\mathbf{g}$, $\mathbf{g}_\nu$ and $\mathbf{g}_t$ at $\hat{\mathbf{z}}$; the rates at $\hat{\mathbf{z}}$ for the time constant and the plug-in curve; the rule (8.30), its last column in Table 8.5, and the sum (8.49); and, for the exact values beside them, the quantiles of U by equation (8.33), a quadrature for $\mathbb{E}[U \mid \mathbf{y}]$ and draws from the gamma posteriors $\mathrm{Gamma}(a_\lambda, b_\lambda)$ and $\mathrm{Gamma}(a_\mu, b_\mu)$ (Section 8.2.2).

Annual downtime. By equation (8.46) and the logit band of Section 8.3.1, the median downtime is $8760 \times 0.01961 = 172$ h per year and the central 95 % band is

$$t_{\mathrm{down}} \in [8760 \times 0.00641,\ 8760 \times 0.05839] = [56,\ 512]\ \mathrm{h\ per\ year}.$$

The last column of Table 8.5 gives the mean, $0.01126{\times}33.5{+}0.22208{\times}79.4{+}\cdots{+}0.01126{\times}825.8 = 200$ h per year, above the median. With $\sigma_\xi = 0.577$ the factor $\mathrm{e}^{1.96\times0.577} = 3.1$ sets the band: $172/56 = 3.1$ below and $512/172 = 3.0$ above the median.

The values beside them come from the exact posterior. By the quantiles of U of Section 8.3.1, the median downtime is $8760 \times 0.019608 = 172$ h per year and the central 95 % band is

$$t_{\mathrm{down}} \in [8760 \times 0.006066,\ 8760 \times 0.06151] = [53,\ 539]\ \mathrm{h\ per\ year}.$$

The posterior mean is $8760 \times \mathbb{E}[U \mid \mathbf{y}] = 8760 \times 0.02320 = 203\,\mathrm{h}$ per year. With $\mathrm{sd}(\xi \mid \mathbf{y}) = 0.60$ the factor $\mathrm{e}^{1.96\times 0.60} \approx 3.2$ sets the band: $172/53 = 3.2$ below and $539/172 = 3.1$ above the median, 119 h below and 367 h above it in hours. The Laplace median has no gap, the mean is 3 h lower and the band narrower, by 3 h at the lower end and 27 h at the upper.

Outage frequency. At $\hat{\mathbf{z}}$, equation (8.47) gives $0.002 \times 0.98039 \times 8760 = 17.2$ outages per year, the Laplace median, and the identity $\nu_{\mathrm{out}}/\mu = U$ gives back the median downtime, $17.2 \times 10 = 172\,\mathrm{h}$. With $\mathbf{g}_\nu = (0.98039,\ 0.01961)^\mathsf{T}$, $\mathbf{g}_\nu^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}_\nu = 0.1603$, a standard deviation of 0.400 in the log, and the band is $17.2 \times \mathrm{e}^{\pm 1.96\times 0.400} = [7.8,\ 37.6]$ per year. The sum (8.49) with forty nodes in each direction gives the mean 18.5 per year. The values beside them come from the exact posterior: two million draws from the two gamma posteriors give ν_{out} a mean of 17.05 and a median of 16.2 outages per year and the central band [6.4, 32.6]. The Laplace median is 1.0 higher and the mean 1.5 higher, and the band lies above that of the draws, by 1.4 at the lower end and 5.0 at the upper: this quantity depends on one log rate, whose center offset does not cancel.

The availability after a repair. At $\hat{\mathbf{z}}$, where the rates $\lambda = 0.002\,\mathrm{h}^{-1}$ and $\mu = 0.1\,\mathrm{h}^{-1}$ are also the posterior means, equation (8.53) gives

$$t_c = \frac{1}{0.002 + 0.1}\,\mathrm{h} = 9.8\,\mathrm{h}.$$

By equation (8.52) the plug-in curve gives $A(t_c) = 0.98039 + 0.01961/\mathrm{e} = 0.9876$; after 24 h, 0.9821, within 0.0017 of the steady state; and $t_c \ln 100 = 45\,\mathrm{h}$, after which the deviation is one percent of its initial value from either starting point.

The predicted availability, the sum (8.49) with forty nodes in each direction, is 0.9803 at $t = 24\,\mathrm{h}$, against the plug-in 0.9821, and 0.9778 at $t = 48\,\mathrm{h}$. The gradient of equation (8.55) gives $\mathbf{g}_t = (1.003,\ -0.767)^\mathsf{T}$ and $\sqrt{\mathbf{g}_t^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}_t} = 0.515$ at 24 h, so the band [0.952, 0.993]; at 48 h, $\mathbf{g}_t = (1.001,\ -0.963)^\mathsf{T}$, 0.567 and the band [0.943, 0.994]. The repair entry of $\mathbf{g}_t$ grows toward -1 as the transient settles, so the repair record limits the band of the first day less than that of the long run.

Figure 8.3 draws equation (8.52) for the first 48 h after a repair. The horizontal axis is the time since the pump returned to service, the vertical axis the availability. The solid curve is the transient at the MAP state; the dashed line is its steady state; the dotted vertical marks the time constant t_c. The shaded band is the central 95 % band of $A(t)$ at each time when the rates are drawn from their exact posteriors, the values beside the reported ones, and the two triangles mark the Laplace band at 24 and 48 h. The predicted availability, the transient averaged over the draws, is 0.9808 at $t = 24\,\mathrm{h}$, with the central band [0.957, 0.994], against the plug-in 0.9821; at $t = 48\,\mathrm{h}$ it is 0.9779 with the band [0.945, 0.994], so the gaps in the average are 0.0005 and 0.0001 and the Laplace lower ends lie below those of the draws, by 0.005 and 0.002. The band of the draws widens toward the steady-state band [0.938, 0.994] more slowly than the solid curve settles: its upper end settles within a day, while its lower end, 0.945 at 48 h, is still above the steady-state lower end 0.938 after two days. The lower end belongs to draws with slow repairs and long time constants: at the 2.5 % quantile of the repair rate, $0.037\,\mathrm{h}^{-1}$, the time constant is about 26 h, and after 48 h such a pump still keeps a sixth of its initial advantage over its steady state.

The interval availability. By equation (8.54), the interval availability after a repair is 0.9938 for an eight-hour shift, 0.9877 for a day and 0.9815 for a week, against the steady state 0.9804; over a shift the pump is spared 0.0134 of its steady-state unavailability of 0.0196.

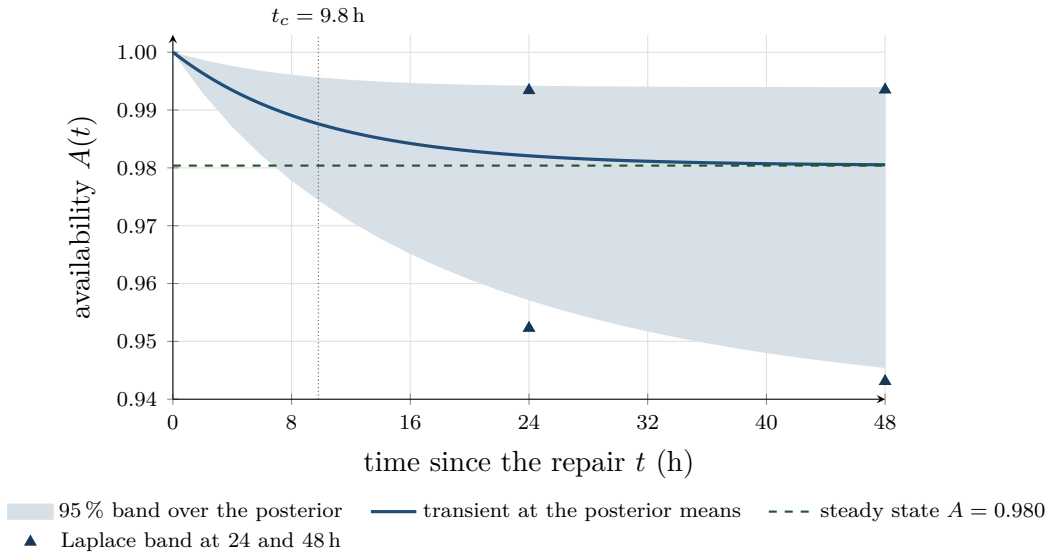


Figure 8.3: The availability of the pump after a repair, equation (8.52): $A(t) = 0.98039 + 0.01961\,\mathrm{e}^{-t/9.804}$, with $\lambda = 0.002\,\mathrm{h}^{-1}$ and $\mu = 0.1\,\mathrm{h}^{-1}$; steady state 0.98039 and time constant $t_c = 1/(\lambda+\mu) = 9.8\,\mathrm{h}$, at which $A(t_c) = 0.9876$. The band holds the central 95 % of $A(t)$ at each time over two million draws of $\lambda \sim \mathrm{Gamma}(6, 3000\,\mathrm{h})$ and $\mu \sim \mathrm{Gamma}(6, 60\,\mathrm{h})$; triangles, the Laplace band (8.55), [0.952, 0.993] at 24 h and [0.943, 0.994] at 48 h.

8.4 The checks (step 5)

The pump has no physics row and no exact constraint, so no limit $\sigma_p \to 0$ of a model tolerance arises and no limit table is needed; the positivity of the rates is not a row either, and the log-rate state handles it.

8.4.1 The Laplace posterior against the exact posterior

The log rates. By the series (8.20),

$$\psi'(6) = 0.16667 + 0.01389 + 0.00077 + \cdots = 0.1813 \quad \text{against} \quad 1/a_\lambda = 0.1667,$$

and $\psi(6) = 1.7061$ against $\ln 6 = 1.7918$. The Laplace approximation understates the variance of each log rate by about 8 % and its standard deviation by 4 %, $1/\sqrt{6} = 0.408$ against $\sqrt{\psi'(6)} = 0.426$, and places its center $\ln 6 - \psi(6) = 0.086$ above the exact mean. For the repair rate, equation (8.24) gives $\hat{z}_2 \pm 1.96\sqrt{\Sigma_{22}} = -2.3026 \pm 0.8002$, so that $\mu \in [0.045,\ 0.223]\,\mathrm{h}^{-1}$ and MTTR $\in [4.5,\ 22.3]\,\mathrm{h}$, against the exact $[0.037,\ 0.194]\,\mathrm{h}^{-1}$ and [5.1, 27.2] h; the same mapping gives $\lambda \in [0.00090,\ 0.00445]\,\mathrm{h}^{-1}$ against the gamma interval [0.00073, 0.00389] of Table 8.4. Both rate intervals are shifted up by the same factors, since both posteriors have shape 6: the lower ends by 22 % and the upper ends by 14 %. The center offset accounts for $\mathrm{e}^{0.086} - 1 = 9\,\%$ of the shift; the rest is the skew of the exact posterior of a log rate, whose left tail is longer than the symmetric normal allows. In the logit the two offsets cancel, and the exact mean of ξ is $\ln c = -3.912$, the Laplace center, which is why the answers below agree more closely than the rates do.

The logit and its budget. From $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$ of equations (8.21) and (8.22), the Laplace center of the logit is $\hat{\xi} = \mathbf{g}^{\mathsf{T}}\hat{\mathbf{z}} = -6.2146 + 2.3026 = \ln 0.02 = -3.912$, and the variance budget (8.29) is

$$\begin{aligned}\sigma_\xi^2 = \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} &= \begin{bmatrix}1 & -1\end{bmatrix}\begin{bmatrix}0.16667 & 0\\ 0 & 0.16667\end{bmatrix}\begin{bmatrix}1\\ -1\end{bmatrix}\\ &= \underbrace{0.16667}_{\text{failure record}} + \underbrace{0.16667}_{\text{repair record}} - \underbrace{2\times 0}_{\text{cross term}} = 0.3333, \qquad \sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}} = 0.577.\end{aligned} \tag{8.56}$$

Table 8.6 sets this budget beside the exact one of Section 8.3.1: each row is a term, and the columns give its Laplace form and value and its exact form and value. Both budgets split equally between the records, with no cross term; the Laplace terms are each 8 % smaller. The logit's Laplace interval $-3.912 \pm 1.960 \times 0.577 = [-5.044,\ -2.780]$ maps through $A = 1/(1+\mathrm{e}^{\xi})$ in reverse order to the band $[0.942,\ 0.994]$, with median 0.980; with $\mathrm{sd}(\xi \mid \mathbf{y}) = 0.602$ in place of $\sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}}$ the same calculation gives $[0.939,\ 0.994]$. At $\hat{\mathbf{z}}$ the unavailability is $U = 0.019608$, and equation (8.31) gives $\mathbf{g}_U = (0.019223,\ -0.019223)^{\mathsf{T}}$ and $\sqrt{\mathbf{g}_U^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_U} = 0.019223 \times 0.5774 = 0.0111$, the same for A.

Table 8.6: The variance budget of the logit of the unavailability. Laplace: the terms $g_i^2\Sigma_{ii}$ of $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ with $\mathbf{g} = (1,-1)^{\mathsf{T}}$ and $\boldsymbol{\Sigma} = \mathrm{diag}(1/a_\lambda,\ 1/a_\mu)$; exact: the variances of the log rates, equation (8.19), which add to $\mathrm{Var}(\xi \mid \mathbf{y})$ because the log rates are independent.

Term	Laplace form	Laplace	Exact form	Exact	Share
Failure record	$g_1^2\Sigma_{11} = 1/a_\lambda$	0.1667	$\mathrm{Var}(\ln\lambda \mid \mathbf{y}) = \psi'(a_\lambda)$	0.1813	50 %
Repair record	$g_2^2\Sigma_{22} = 1/a_\mu$	0.1667	$\mathrm{Var}(\ln\mu \mid \mathbf{y}) = \psi'(a_\mu)$	0.1813	50 %
Cross term	$2g_1g_2\Sigma_{12}$	0	$2\,\mathrm{Cov}(\ln\lambda, \ln\mu \mid \mathbf{y})$	0	0
Variance of ξ	$\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$	0.3333	$\mathrm{Var}(\xi \mid \mathbf{y})$	0.3626	
Spread of ξ	$\sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}}$	0.577	$\mathrm{sd}(\xi \mid \mathbf{y})$	0.602	

The requirement. By equation (8.41) the requirement is $\xi \le \xi^* = \ln(0.03/0.97) = -3.4761$, and the Laplace tail (8.42) is

$$\Pr(A \ge 0.97 \mid \mathbf{y}) \approx \Phi\left(\frac{\xi^* - \hat{\xi}}{\sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}}}\right) = \Phi\left(\frac{-3.4761 + 3.9120}{0.5774}\right) = \Phi(0.755) = 0.775, \tag{8.57}$$

against the exact 0.769; with $\mathrm{sd}(\xi \mid \mathbf{y})$, $\Phi(0.4359/0.6022) = \Phi(0.724) = 0.765$. Figure 8.4 draws the Laplace posterior of the state with the event. The horizontal axis is the log failure rate z_1 and the vertical axis the log repair rate z_2, both with the same scale. The dot is the MAP state (8.21); the two circles are the contours of the Gaussian that hold 39 % and 95 % of its probability, at one and 2.45 standard deviations, 0.408 and 1.00 in each log rate. The straight lines are lines of constant availability, $z_2 = z_1 + \ln(A/(1-A))$, labeled at their ends with the value of A; the shaded half-plane above the line $A = 0.97$ is the event $A \ge 0.97$. The contours are circles with axes parallel to the coordinate axes, the zero cross term and the equal shares of Table 8.6. The line $A = 0.980$ passes through the MAP state, and the line $A = 0.97$ lies below and to the right of it, at the distance $(\xi^* - \mathbf{g}^{\mathsf{T}}\hat{\mathbf{z}})/\sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}} = (-3.476 + 3.912)/0.577 = 0.755$ standard deviations along $\mathbf{g} = (1,-1)^{\mathsf{T}}$, the argument of equation (8.57).

Table 8.7 extends the comparison to five requirement levels. Each row is a required availability; the columns give the threshold v^* of equation (8.43), the probability under the priors alone, the exact posterior probability, the Laplace probability and the probability under the normal on ξ with exact variance. One row of the arithmetic, $A_{\mathrm{req}} = 0.99$: $v^* = 0.01/(0.01 + 0.02 \times 0.99) = 0.3356$, and $I_{0.3356}(6,6) = 0.126$. The exact support falls from 0.95 at the 95 % level to 0.13 at the 99 % level, and passes one half at 0.98, the median.

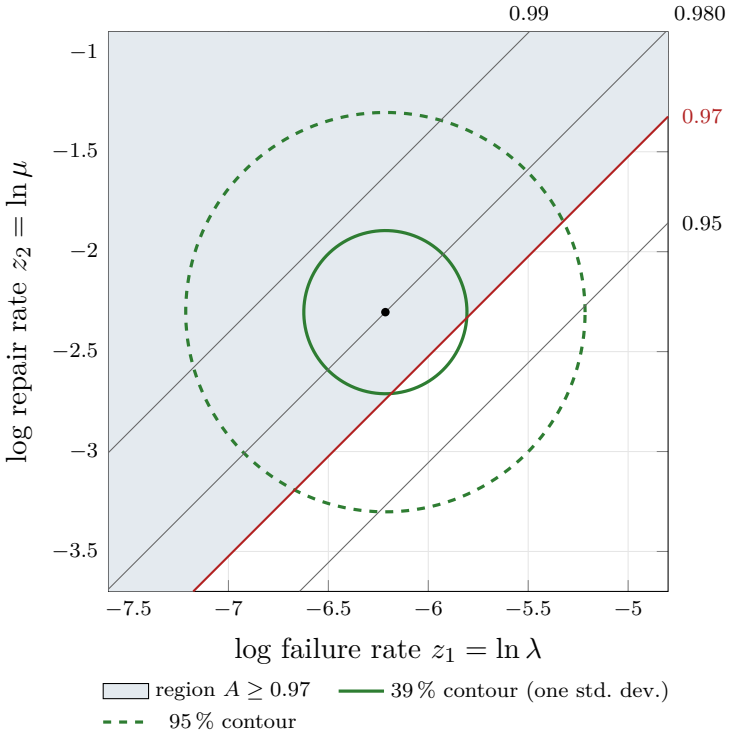


Figure 8.4: The Laplace posterior of the pump's log rates, $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ with $\hat{\mathbf{z}} = (\ln 0.002,\ \ln 0.1) = (-6.215,\ -2.303)$ and $\boldsymbol{\Sigma} = \mathrm{diag}(1/6,\ 1/6)$: the MAP state (dot) and the contours at Mahalanobis distances 1 and 2.448, circles of radius 0.408 and 0.999. Lines of constant availability $z_2 = z_1 + \ln(A/(1-A))$ for $A = 0.95$, 0.97 (red), 0.980 and 0.99; the region $A \geq 0.97$ is shaded.

Table 8.7: Probability that the pump's availability meets a requirement, $\Pr(A \geq A_{\mathrm{req}} \mid \mathbf{y})$. Prior: Beta$(2, 2)$ route with $c = 0.02$; exact: $I_{v^*}(6, 6)$, equation (8.43); Laplace: $\Phi((\xi^* + 3.912)/0.577)$; exact variance: $\Phi((\xi^* + 3.912)/0.602)$.

Requirement A_{req}	v^*	Prior	Exact	Laplace	Exact variance
0.95	0.7246	0.814	0.947	0.953	0.946
0.96	0.6757	0.753	0.891	0.898	0.889
0.97	0.6073	0.658	0.769	0.775	0.765
0.98	0.5051	0.508	0.514	0.514	0.513
0.99	0.3356	0.262	0.126	0.118	0.128

The routes for the availability. Table 8.8 sets the routes side by side for the questions of the chapter. Each row is a summary; the columns give it under the exact posterior, under the Laplace normal on ξ, under the normal on ξ with the exact variance, and under a normal placed directly on A by the delta method, with mean 0.9804 and standard deviation $\sqrt{\mathbf{g}_U^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_U} = 0.0111$ from the gradient (8.31).

Table 8.8: The availability of the pump by four routes. Exact: the transformed Beta$(6, 6)$; Laplace: $\xi \sim \mathcal{N}(-3.912,\ 0.577^2)$; exact variance: $\xi \sim \mathcal{N}(-3.912,\ 0.602^2)$; normal on A: $\mathcal{N}(0.9804,\ 0.0111^2)$. Downtime in hours per year.

Summary	Exact	Laplace	Exact variance	Normal on A
Standard deviation of ξ	0.602	0.577	0.602	–
Median of A	0.9804	0.9804	0.9804	0.9804
Mean of A	0.9768	0.9771	0.9768	0.9804
95 % band on A	[0.938, 0.994]	[0.942, 0.994]	[0.939, 0.994]	[0.959, 1.002]
$\Pr(A \geq 0.97 \mid \mathbf{y})$	0.769	0.775	0.765	0.825
$\Pr(A > 1 \mid \mathbf{y})$	0	0	0	0.039
Downtime, median	172	172	172	172
Downtime, mean	203	200	203	172
Downtime, 95 % band	[53, 539]	[56, 512]	[53, 535]	[−19, 363]

In Figure 8.2 the three densities nearly coincide; the Laplace curve is a little taller and narrower than the exact one. The two normals on the logit respect the bounds of the availability and its skew, and they agree with the exact posterior to within 0.006 on the requirement and 0.004 on the lower end of the band. Their remaining error has two parts that can be told apart: the variance, which is the whole difference between the Laplace column and the exact-variance column, and the shape, the difference between the exact-variance column and the exact one, which is smaller still. Over the five levels of Table 8.7 the two approximations stay within 0.008 of the exact value. The Laplace normal, narrower than the exact posterior, puts more probability near the median: it overstates the support for requirements below the median and understates it for a requirement above it, such as 0.99. The normal placed directly on the availability fails in the way the normal on the failure probability failed in Chapter 5: it is symmetric about 0.980, its band crosses one, it puts 3.9 % of its probability on availabilities above one and a negative lower end on the downtime, and it overstates the support for the requirement by 0.056.

Common mistake

Applying the delta method to the availability itself. The availability is bounded by one and its posterior is skewed toward lower values; a normal on A with the delta-method spread puts 3.9 % of its probability above one and overstates the support for $A \geq 0.97$ by 0.056. The linearization belongs on the logit, where the derived quantity is linear in the state, and the result is mapped back through $A = 1/(1 + \mathrm{e}^{\xi})$.

The outage frequency. At $\hat{\mathbf{z}}$ the gradient (8.48) and its variance are

$$\mathbf{g}_\nu = \begin{bmatrix} 0.98039 \\ 0.01961 \end{bmatrix}, \qquad \mathbf{g}_\nu^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_\nu = 0.98039^2 \times 0.16667 + 0.01961^2 \times 0.16667 = 0.1603,$$

a standard deviation of 0.400 in the log, almost entirely from the failure record, and the Laplace band is $17.2 \times \mathrm{e}^{\pm 1.96 \times 0.400} = [7.8,\ 37.6]$ per year. It lies above the band [6.4, 32.6] of the draws at both ends, the lower end by 23 % and the upper by 15 %, the same pattern as the interval of λ above: the Laplace normal centers $\ln\lambda$ 0.086 above its exact mean, a shift of 9 %, and has no

left skew, while in a quantity that depends on one log rate the offset does not cancel. The band of the draws is narrower than the gamma band of the failure rate times 8760 h and the median availability, [6.3, 33.4] per year, mostly at its upper end, because the draws with the largest failure rates also have the lowest availabilities.

The transient. At $t = 24\,\mathrm{h}$ the gradient (8.55) is $\mathbf{g}_t = (1.003,\ -0.767)^\mathsf{T}$, with $\sqrt{\mathbf{g}_t^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}_t} = 0.515$, and the Laplace band is [0.952, 0.993], against the band [0.957, 0.994] of the draws; at $t = 48\,\mathrm{h}$, $\mathbf{g}_t = (1.001,\ -0.963)^\mathsf{T}$, $\sqrt{\mathbf{g}_t^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}_t} = 0.567$ and the band [0.943, 0.994], against [0.945, 0.994]. The repair entry of $\mathbf{g}_t$ grows toward -1 as the transient settles, so the repair record limits the band of the first day less than that of the long run, and the Laplace lower end lies below that of the draws by up to 0.005.

8.4.2 The hand check and the draws

The Beta–binomial identity of Chapter 5, $I_x(n_1, n_2) = \Pr(\mathrm{Binomial}(n_1 + n_2 - 1, x) \geq n_1)$ for integer parameters, checks the exact probability (8.44): with $n_1 = n_2 = 6$ it is the probability of at least six successes in eleven trials of probability $v^* = 0.60729$, the sum of the six terms $0.2165 + 0.2391 + 0.1849 + 0.0953 + 0.0295 + 0.0041 = 0.7693$.

Two million pairs drawn from the two gamma posteriors and pushed through equation (8.1) give $\Pr(A \geq 0.97) = 0.7694$ and the band [0.9386, 0.9939], the exact values to the third decimal. The draws check the Beta route, which needs no check of its own beyond its algebra; their real use is in a larger model in which the posterior of the rates is only available as a Laplace normal or as draws, where the same comparison tests the delta method.

8.4.3 The exponential repair time

The gamma posterior of the repair rate rests on exponential repair durations. An exponential duration has a standard deviation equal to its mean, a coefficient of variation of one. The four recorded repairs have mean 10 h and sample standard deviation 3.37 h, a coefficient of variation of 0.337 (Table 8.1): they are much less variable than the model says. Four exponential durations give a sample coefficient of variation of 0.337 or less with probability 0.039, found by simulating such samples; the median for four is 0.81.

An evidence ratio. The question whether the exponential model or another describes the repairs is a hypothesis comparison, the subject of later chapters on model selection, and only its result is used here. Under the exponential model with the prior $\mathrm{Gamma}(a_{\mu 0}, b_{\mu 0})$, the evidence of the durations is the likelihood averaged over the prior,

$$p(d_1, \ldots, d_r \mid \text{exponential}) = \int_0^\infty \mu^r \mathrm{e}^{-\mu D}\, \frac{b_{\mu 0}^{a_{\mu 0}} \mu^{a_{\mu 0}-1} \mathrm{e}^{-b_{\mu 0}\mu}}{\Gamma(a_{\mu 0})}\, \mathrm{d}\mu = \frac{b_{\mu 0}^{a_{\mu 0}}\, \Gamma(a_\mu)}{\Gamma(a_{\mu 0})\, b_\mu^{a_\mu}}, \tag{8.58}$$

with $a_\mu = a_{\mu 0} + r$ and $b_\mu = b_{\mu 0} + D$, which for the record is $20^2 \times 120/60^6 = 1.03 \times 10^{-6}\,\mathrm{h}^{-4}$, $\ln p = -13.79$. A lognormal model, in which the log of each duration is normal with mean m_d and standard deviation s_d, with the priors $m_d \sim \mathcal{N}(\ln 10, 1)$, the log of a duration in hours, and s_d uniform on $(0, 2)$, has an evidence found by two-dimensional quadrature of $\ln p = -13.24$. The log Bayes factor of the lognormal against the exponential model is $\ln B_{10} = +0.55$ nats, a weak preference. It depends on the prior of the spread s_d, with m_d's prior kept: 0.16 nats for s_d uniform on $(0, 3)$, 0.82 for a half-normal of scale one, and 0.98 and 1.43 for a lognormal prior with median 0.5 and log standard deviation 1 and 0.5. Every choice favors the lognormal and none by more than about one and a half nats. Four durations are too few to settle the repair-time distribution; they are enough to make the exponential model doubtful.

What the assumption changes. The steady-state availability depends on the repair distribution only through its mean, by the alternating-renewal result of Section 8.3.3, so equation (8.1) with $\mu = 1/\text{MTTR}$ holds whatever the distribution. What the exponential model fixes is the posterior of the mean repair time. Under the lognormal model with the priors above, the posterior standard deviation of $\ln \text{MTTR}$ is 0.43, against the exponential model's 0.41 under the Laplace route and 0.43 exactly: the less variable durations would determine the mean more tightly if their spread were known, but the spread is itself uncertain from four durations, and the two effects nearly cancel. The band on the availability is therefore not materially changed. The transient $A(t)$ of equation (8.52) is a different matter: it uses the exponential repair time directly, and with repair durations as regular as these the pump just failed returns to service more predictably than the exponential curve shows. The repair record should be kept as durations, not as a total, so that the check can be repeated as it grows.

8.4.4 Exposure and horizon

Operating exposure without downtime. The failure part of equation (8.6) has the sum of the up intervals as its exposure. Counting calendar time instead, 2040 h for the record including the 40 h of repairs, would give the failure rate the posterior $\text{Gamma}(6, 3040\,\text{h})$, so that $c = 60/3040 = 0.01974$, the median availability 0.9806 and $\Pr(A \geq 0.97) = 0.776$ instead of 0.769. The error is small here because the pump is down only two percent of the time; it grows with the unavailability, and it always flatters the component, since time in which the pump could not fail is counted as time in which it did not. Standby time in which the pump is not running must likewise be excluded unless the failure model covers it.

The horizon of the steady state. The steady-state availability answers questions about the long run. The transient of Section 8.3.3 reaches its steady state within about two days from either starting point, and any question about a horizon of weeks or a year may use equation (8.1): the interval availability of a week is 0.9815, against the steady state 0.9804. A question about the first shift after a repair may not: over an eight-hour shift the interval availability is 0.9938. A question about the first two days after a repair needs the transient even when its answer is a band rather than a plug-in value, since the lower end of the band, 0.945 after 48 h, is still above the steady-state lower end 0.938.

Common mistake

Reading an availability of 0.98 as a 98 % probability that the pump will not fail over some period. The availability is the fraction of time the pump is up, set by both rates. The probability of no failure in a year of operation is the posterior predictive of Chapter 6, $(b_\lambda/(b_\lambda + 8760\,\text{h}))^{a_\lambda} = (3000/11\,760)^6 = 2.8 \times 10^{-4}$ for this pump, which fails about seventeen times a year and is nonetheless up 98 % of the time; the plug-in value $\mathrm{e}^{-8760 \times 0.002} = 2.5 \times 10^{-8}$ would understate even that by four orders of magnitude. The companion volume draws the distinction; the derivations of the next chapter are about the second quantity.

8.5 The report (step 6)

The records do not support the first reading of them. The availability at the estimated rates is 0.980, above the requirement of 0.97, but the posterior probability that the pump's long-run availability meets 0.97 is 0.775 (0.769 exact), short of the assurance of 0.90 the agreement asks for. The shortfall comes from too little evidence, not from a poor pump: the median availability is 0.980, and it is the width of the band, from 0.942 to 0.994 (0.938 to 0.994 exact), that leaves the requirement undemonstrated, and the failure record carries half of the variance behind that width, so repairs alone cannot close it. The answers rest on the posterior of the state of step 3;

the exact gamma and Beta posteriors and, for the outage frequency and the transient, draws from them give the value beside each.

8.5.1 The answers and their bands

Table 8.9 gives the answers. Each row is a question, the fourth split into its two parts; the second column gives the answer from the posterior of the state, the third the exact value beside it and the last the band of the reported route, exact in parentheses.

Table 8.9: The report for the pump's records: the answers from the normal on the log rates, the exact values from the gamma and Beta posteriors or from draws, and the central 95 % posterior intervals of the reported route, exact in parentheses.

Question	Answer	Exact	Band
What is the repair rate, and the mean time to repair?	μ: median $0.100\,\mathrm{h}^{-1}$, mean $0.1087\,\mathrm{h}^{-1}$; MTTR: median 10.0 h, mean 10.87 h	μ: mean $0.100\,\mathrm{h}^{-1}$; MTTR: median 10.6 h, mean 12.0 h	$\mu \in [0.045,\ 0.223]\,\mathrm{h}^{-1}$ ([0.037, 0.194]), MTTR $\in$ [4.5, 22.3] h ([5.1, 27.2]), from $\hat{z}_2 \pm 1.96\sqrt{\Sigma_{22}}$
What fraction of the time will the pump be available?	median 0.980, mean 0.977	median 0.980, mean 0.977	[0.942, 0.994] ([0.938, 0.994])
Is the availability at least 0.97 with assurance 0.90?	probability 0.775; not demonstrated	0.769; not demonstrated	exact variance 0.765; priors alone 0.669 (0.658)
How many hours per year will the pump be down?	median 172 h, mean 200 h	median 172 h, mean 203 h	[56, 512] h ([53, 539] h)
How often will it fail?	median 17.2, mean 18.5 per year	median 16.2, mean 17.0 per year	[7.8, 37.6] per year ([6.4, 32.6] from draws)
How available is the pump after a repair?	0.980 at 24 h, 0.978 at 48 h; time constant $t_c = 9.8$ h	0.981 at 24 h, 0.978 at 48 h	$A(24\,\mathrm{h})$ in [0.952, 0.993] ([0.957, 0.994]), $A(48\,\mathrm{h})$ in [0.943, 0.994] ([0.945, 0.994])

The report line. Under constant failure and repair rates, the pump's long-run availability has median 0.980 and mean 0.977, with a 95 % band from 0.942 to 0.994 (0.938 to 0.994 from the exact posterior); the posterior probability that it is at least 0.97 is 0.78 (0.77), so a requirement of 0.97 at an assurance of 0.90 is not yet demonstrated on either route. The pump should expect about 17 outages a year, with a band from 8 to 38 (6 to 33), and a downtime of median 172 h and mean 200 h (203 h) per year, with a band from 56 to 512 h (53 to 539 h). The band rests on exponential repair durations, which the four recorded repairs make doubtful; the steady-state availability does not depend on that assumption, and the alternative examined changes the band only slightly, but the check should be repeated as repairs accumulate. The repair record limits the band as much as the failure record.

8.5.2 The requirement decision

The decision follows from the third row of Table 8.9 and from Table 8.7. If the agreement requires an availability of 0.97 at an assurance of 0.90, the records do not meet it: the support is 0.775 (0.769 exact), and the gap does not change the decision. The point value 0.980 is above the requirement, the posterior median is above it, and still about one chance in four remains that the true long-run availability is below 0.97.

Common mistake

Declaring the requirement $A \geq 0.97$ met because the availability computed from the estimated rates, 0.980, exceeds it. The estimate is the center of a posterior whose 95 % band reaches down to 0.942 (0.938 exact); the support for the requirement is the posterior probability 0.775 (0.769), and it depends on the size of both records, not only on their means.

8.6 The next measurement (step 7)

The report leaves the requirement of 0.97 at an assurance of 0.90 undemonstrated: the band $[0.942,\ 0.994]$ ($[0.938,\ 0.994]$ exact) is too wide to support it. Records could be bought to narrow it, and the preposterior values each before its outcome is known.

8.6.1 The candidate records

Two records of comparable effort could be bought to narrow the band:

- 2000 *more operating hours.* The pump keeps running and the log keeps recording. The expected number of failures in that exposure is $2000\,\mathrm{h} \times \mathbb{E}[\lambda \mid \mathbf{y}] = 4$, and every failure is followed by a repair, so the record adds an expected four failures to the count row and four repair durations to the repair row.
- *Ten repair drills.* The crew stages ten repairs under recorded time, which adds ten durations to the repair row and nothing to the count row. The drills assume that a staged repair takes as long as a real one, an assumption to be checked before they are pooled with the record.

A longer operating record of 4000 h is carried beside them as a reference.

8.6.2 The preposterior of the logit

The added information. The preposterior covariance of Chapter 2, equation (2.17), is $\boldsymbol{\Sigma}_a = (\boldsymbol{\Sigma}^{-1} + \boldsymbol{\Lambda}_a)^{-1}$, with $\boldsymbol{\Lambda}_a$ the expected information of the candidate record at the MAP state. A candidate record brings new count rows of the form (8.9), and its information is computed exactly as that of the present rows, $\boldsymbol{\Lambda}_a = \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a$, with $\mathbf{J}_a$ the Jacobian (8.11) of its rows at $\hat{\mathbf{z}}$ with each expected count in place of the realized one. The count row $T'_{\mathrm{tot}}\mathrm{e}^{z_1} - k'z_1$ of a new exposure T'_{tot} has the log-count Jacobian row $(1,\ 0)$ and the expected count $T'_{\mathrm{tot}}\hat{\lambda}$; the repair row $D'\mathrm{e}^{z_2} - r'z_2$ of r' new durations has the log-count Jacobian row $(0,\ 1)$ and the curvature $D'\hat{\mu}$, whose expectation is r' because $\mathbb{E}[D' \mid r', \mu] = r'/\mu$. Each expected failure therefore adds one unit of precision to z_1 and each repair one unit to z_2. For 2000 more operating hours, with $T'_{\mathrm{tot}}\hat{\lambda} = 4$ expected failures and as many repairs, and for ten drills, which feed the repair row only, the Jacobians and the added information are

$$\mathbf{J}_{2000\,\mathrm{h}} = \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix}, \quad \boldsymbol{\Lambda}_{2000\,\mathrm{h}} = \begin{bmatrix} 4 & 0 \\ 0 & 4 \end{bmatrix}; \qquad \mathbf{J}_{\mathrm{drills}} = \begin{bmatrix} 0 & 3.1623 \end{bmatrix}, \quad \boldsymbol{\Lambda}_{\mathrm{drills}} = \begin{bmatrix} 0 & 0 \\ 0 & 10 \end{bmatrix},$$

with $2 = \sqrt{4}$ and $3.1623 = \sqrt{10}$, and for the longer operating record of 4000 h, $\boldsymbol{\Lambda}_{4000\,\mathrm{h}} = \mathrm{diag}(8,\ 8)$. With $\boldsymbol{\Sigma}^{-1} = \mathrm{diag}(6, 6)$ of equation (8.22), the three preposterior covariances are

$$\boldsymbol{\Sigma}_{2000\,\mathrm{h}} = \begin{bmatrix} 0.1000 & 0 \\ 0 & 0.1000 \end{bmatrix}, \quad \boldsymbol{\Sigma}_{\mathrm{drills}} = \begin{bmatrix} 0.1667 & 0 \\ 0 & 0.0625 \end{bmatrix}, \quad \boldsymbol{\Sigma}_{4000\,\mathrm{h}} = \begin{bmatrix} 0.0714 & 0 \\ 0 & 0.0714 \end{bmatrix},$$

each diagonal because no candidate row acts on both entries. The preposterior variance of the logit is the budget of the new covariance, with the gradient $\mathbf{g} = (1, -1)^{\mathsf{T}}$ of equation (8.28),

$$\sigma^2_{\xi,a} = \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_a\mathbf{g} = \frac{1}{a_\lambda + \Delta_\lambda} + \frac{1}{a_\mu + \Delta_\mu}, \tag{8.59}$$

with Δ_λ and Δ_μ the diagonal entries of $\mathbf{\Lambda}_a$. For the three candidates, $0.1000 + 0.1000 = 0.2000$, $0.1667 + 0.0625 = 0.2292$ and $0.0714 + 0.0714 = 0.1429$.

Table 8.10 ranks the candidates. Each row is a candidate record; the columns give the rows it feeds, its added information, the preposterior standard deviation of ξ under the Laplace route and with the exact variances ψ' at the expected shapes, the fraction of the present Laplace width it removes, and the lower end of the 95 % band on A it would leave if the center did not move. Figure 8.5 draws the two standard deviation columns as bars: the horizontal axis lists the present record and the three candidates, the vertical axis the standard deviation of the logit; the blue bars are the Laplace values and the gray bars the exact-variance values. The upper dashed line marks the present Laplace value and the lower dotted line the floor that repair records alone cannot pass.

Table 8.10: The preposterior of the logit ξ for three candidate records, equation (8.59). Laplace: inverse shapes; exact variance: trigamma values ψ' at the expected shapes. Width removed and band end are for the Laplace values, with the center held at $\ln 0.02$.

Record	Rows fed	diag($\mathbf{\Lambda}_a$)	$\sigma_{\xi,a}$ Laplace	sd of ξ, exact variance	Width removed	Lower end of band
present	–	–	0.577	0.602	–	0.942
2000 more hours	both	(4, 4)	0.447	0.459	22.5 %	0.954
ten repair drills	repair	(0, 10)	0.479	0.496	17.1 %	0.951
4000 more hours	both	(8, 8)	0.378	0.385	34.5 %	0.960

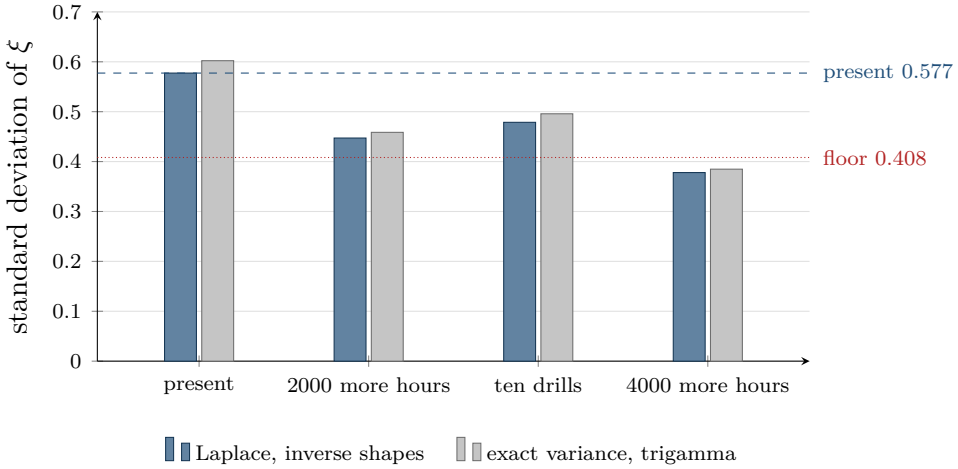


Figure 8.5: Preposterior standard deviation of the logit of the unavailability for the present record and three candidates, equation (8.59) (Table 8.10). Blue: Laplace, $\sqrt{1/(6+\Delta_\lambda) + 1/(6+\Delta_\mu)}$; gray: $\sqrt{\psi'(6+\Delta_\lambda) + \psi'(6+\Delta_\mu)}$. The dotted line is $1/\sqrt{6} = 0.408$, the limit of repair records alone.

8.6.3 The ranking and its reason

Two thousand more operating hours narrow the band more than ten repair drills, although the drills add more than twice as many repair durations, 10 against an expected 4. The reason is in the Jacobians of the candidates. Operation has a Jacobian row in each column, because every failure it produces is followed by a repair, so its $\mathbf{\Lambda}_a$ adds to both diagonal entries and shrinks both entries of $\mathbf{\Sigma}_a$, 0.1667 to 0.1000 each. The drills' Jacobian is (0, 3.1623), with an

empty first column, so their $\mathbf{\Lambda}_a$ adds only to the second diagonal entry: Σ_{22} falls from 0.1667 to 0.0625 while Σ_{11} stays at 0.1667. However many drills are run, the first column of $\mathbf{\Lambda}_a$ stays zero, $\Sigma_{a,11} = 1/a_\lambda = 0.1667$, and $\mathbf{g}^\mathsf{T}\mathbf{\Sigma}_a\mathbf{g} \geq \Sigma_{a,11}$, a standard deviation of 0.408, the floor of Figure 8.5; ten drills already reach 0.479, and the next ten would take it only to 0.453. The ranking is the same under the exact variances. Among records of comparable size, a record that feeds the term with the larger share of the budget, or both terms, is worth more than one that feeds a single term; a larger record of the single term can still reach a given width, as more drills do down to the floor, and whether it does so more cheaply is a question of price, not of width.

The check on the expected count. The preposterior above inserts the expected number of failures, four, into the shapes. The number that will occur is uncertain, negative binomial under the posterior, as Chapter 6 derived, and because $1/(a_\lambda + K)$ is convex in K the average of the preposterior variance over the predictive count is larger than its value at the average count. With K failures the preposterior covariance is $\mathbf{\Sigma}_a(K) = \mathrm{diag}\big(1/(6+K),\ 1/(6+K)\big)$ and $\mathbf{g}^\mathsf{T}\mathbf{\Sigma}_a(K)\mathbf{g} = 2/(6+K)$; averaging it over the negative-binomial distribution of K in 2000 h gives a standard deviation of 0.461 instead of 0.447, and for 4000 h, 0.396 instead of 0.378. Two thousand operating hours still beat ten drills, by a smaller margin. The drills are not affected, since their number is fixed.

Time and cost. The ranking by width is not yet a decision. Two thousand operating hours take about twelve weeks of calendar time, and the drills could be run in a week; each has its cost. Exercise 8.6 puts prices on both and shows that, at those prices, a plan of drills reaches $\sigma_\xi = 0.45$ more cheaply than a plan of operation, although per record the operation removes more.

8.6.4 The recommendation

Of the two candidate records of comparable effort, 2000 more operating hours narrow the logit's standard deviation from 0.577 to 0.447 and ten repair drills to 0.479; operation ranks first because it feeds both terms of the budget, and no number of drills can take the standard deviation below 0.408. On the answer the report gives, with the center held at $\ln 0.02$, the Laplace support $\Phi(0.4359/\sigma_{\xi,a})$ for $A \geq 0.97$ rises from 0.775 to 0.835 after 2000 more hours, to 0.819 after ten drills and to 0.876 after 4000 more hours. None reaches the assurance of 0.90, which needs $\sigma_{\xi,a} \leq 0.4359/1.2816 = 0.340$: by equation (8.59) about 11.3 expected failures, some 5640 more operating hours, or 3810 hours if ten drills are added to them; drills alone stop at the floor of 0.408, a support of at most 0.857.

Keep the operating record running, with every repair logged as a duration, because it is the only record that can demonstrate the requirement of 0.97 at an assurance of 0.90; add the ten drills once a staged repair has been shown to take as long as a real one, since they narrow the repair term in a week and cut the operating hours needed by about a third. How many drills to add beyond the ten is a question of their price (Exercise 8.6). Chapter 20, in Volume II, ranks candidate rows by the same preposterior on graphs whose rows share entries.

8.6.5 What the analysis showed

The first reading of the two logs, an availability of 0.980 at the estimated rates against a requirement of 0.97, said the pump passes; the analysis does not support it. The posterior probability that the long-run availability meets 0.97 is about 0.78 (0.77 exact), short of the assurance of 0.90 the agreement asks for, and the shortfall is a matter of too little evidence rather than of a poor pump: the median availability is 0.980, and it is the width of the band that leaves about one chance in four below the requirement. The two records inform the two

rates separately: the failure log acts only on the log failure rate and the repair log only on the log repair rate, so the two are independent after the records, and the uncertainty of the availability is the sum of two independent terms, one per record, which here are of equal size because four failures were followed by four repairs. The evidence that closes the gap is more operating time with every repair logged, because each failure it brings narrows the failure term and its repair narrows the repair term; repair drills narrow only the repair term, and however many are run the failure term, half of the present variance, sets a floor below which the band cannot fall. A faster crew would shorten the downtime and raise the availability in proportion, but it would leave the number of outages a year almost unchanged, since that number is set by the failure rate: the outage frequency, not the availability, is what better repairs do not improve.

8.6.6 Limits of the model

The derivation is exact under the assumptions of the two-state model of Section 8.1.3. That does not make the assumptions appropriate, and five of them limit the answers.

Constant rates and repair to as good as new. The failure rate is constant over operating time and every repair returns the pump to the condition it had before. A rate that changes with age or with maintenance, or a repair that leaves the pump worse than new, is outside the model; Chapter 12 treats a rate that changes with time.

Exponential repair durations. The posterior of the repair rate, the band of the availability and the transient rest on one exponential repair time, which the four recorded repairs make doubtful (Section 8.4.3). The steady-state availability does not depend on the assumption and the band changes little under the alternative examined; the transient does depend on it.

Operating exposure. The exposure of the failure record counts time in service only, and standby time in which the pump is not running belongs to it only if the failure model covers it (Section 8.4.4).

Independence of the two records. The variance budget has no cross term because $\Sigma_{12} = 0$, and Σ_{12} is zero because no row of $\mathbf{J}$ has entries in both columns. That holds when the repair durations tell nothing about the failure rate and the failure times nothing about the repair rate. Two situations break it. If one crew serves several pumps, its repair rate is a single entry of the state, the repair rows of all the pumps put their Jacobian entries in its one column, and the gradients of the pumps' logits share that entry, so the availabilities of the pumps are correlated through it; when two pumps are down together, the second waits, and the two-state model of each pump no longer holds alone. Exercise 8.5 works the shared entry, and Chapter 16 the chain of states that the queue requires. And if failures come in two modes, a frequent one that is quickly repaired and a rare one that takes long, a single pair of rates averages over the modes; the steady-state availability is still right through the mean repair time, but the posterior band, which assumes one exponential repair time, is not. A model with one pair of rates per mode is the remedy, and its graph has more entries.

Staged repairs. The drills of the recommendation are evidence about real repairs only if a staged repair takes as long as a real one; until that has been shown, their durations belong in a row of their own.

The next chapter turns from the long-run fraction of time a component is up to the probability that it survives a mission of stated length, with the rate uncertainty of this part of the book carried into that probability.

8.7 What to retain

Definition 8.1 (Posterior availability of a repairable component)**.** Under a constant failure rate λ and a constant repair rate μ, an up-and-down record has the likelihood $\lambda^k \mathrm{e}^{-\lambda T_{\mathrm{tot}}}\, \mu^r \mathrm{e}^{-\mu D}$, with T_{tot} the operating exposure excluding downtime and D the total repair time. Independent gamma priors give independent posteriors $\mathrm{Gamma}(a_\lambda, b_\lambda)$ and $\mathrm{Gamma}(a_\mu, b_\mu)$. The steady-state availability $A = \mu/(\lambda+\mu)$ has the logit of its unavailability $\xi = \ln\lambda - \ln\mu$; exactly, $b_\lambda\lambda/(b_\lambda\lambda + b_\mu\mu) \sim \mathrm{Beta}(a_\lambda, a_\mu)$ and every probability and band of A is a Beta tail or quantile, the mean a one-dimensional integral, the exact value set beside each answer. The answers themselves are computed from the posterior of the state. Under the Laplace approximation, the four rows of the state $\mathbf{z} = (\ln\lambda, \ln\mu)^{\mathsf{T}}$, each a count term $\mathcal{L}_i = \bar{k}_i - k_i \ln \bar{k}_i$ with working residual $r_i = (\bar{k}_i - k_i)/\sqrt{\bar{k}_i}$ and Jacobian row $\mathbf{J}_i = \sqrt{\bar{k}_i}\, \partial \ln \bar{k}_i / \partial \mathbf{z}^{\mathsf{T}}$ (equation (4.43)), are fitted by the Newton normal equations $(\mathbf{J}^{\mathsf{T}}\mathbf{J})\Delta\mathbf{z} = -\mathbf{J}^{\mathsf{T}}\mathbf{r}$; $\boldsymbol{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1} = \mathrm{diag}(1/a_\lambda,\ 1/a_\mu)$, diagonal because no row of $\mathbf{J}$ has two nonzero entries; and $\xi = \mathbf{g}^{\mathsf{T}}\mathbf{z}$ with $\mathbf{g} = (1, -1)^{\mathsf{T}}$ is normal with mean $\mathbf{g}^{\mathsf{T}}\hat{\mathbf{z}} = \ln((a_\lambda/b_\lambda)/(a_\mu/b_\mu))$ and variance $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} = 1/a_\lambda + 1/a_\mu$, one term per record.

Three conclusions carry beyond the pump. A derived quantity is asked about through the scale on which it is linear in the state, here the logit, and mapped back; a normal placed on the availability itself crosses its bound. The variance budget names the record that limits the answer, and for a repairable component the repair record limits it as much as the failure record. And among records of comparable size, one that feeds every term of the budget is worth more than one that feeds a single term, which a preposterior shows before either record is bought; the prices of the records then decide which plan reaches a required width more cheaply.

Table 8.11 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 8.11: The terms of repair-rate estimation and component availability and where each is defined.

Term	Expression	Units	Meaning	Defined in
Steady-state availability	$A = \mu/(\lambda+\mu)$	–	long-run fraction up	8.1.1
Unavailability	$U = \lambda/(\lambda+\mu)$	–	long-run fraction down	8.1.1
General forms				
Estimation	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$, moments, quantiles	–	posterior of a derived quantity	8.3.1
Threshold and assurance	$\Pr(f(\mathbf{z}) \ge f^* \mid \mathbf{y}) \ge \kappa$	–	tail of a derived quantity	8.3.2
Expectation	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$, quantiles or draws	–	further derived quantities	8.3.3
Specialized forms and answers				
Mean time to repair	$\mathbb{E}[1/\mu] = b_\mu/(a_\mu - 1)$	h	inverse-gamma mean	8.3.1
Logit of unavailability	$\xi = \ln\lambda - \ln\mu = \mathbf{g}^\mathsf{T}\mathbf{z}$	–	linear in the state	8.3.1
Variance budget	$\sigma_\xi^2 = \mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g} = 1/a_\lambda + 1/a_\mu$	–	one term per record	8.3.1
Exact availability	$U_q = c\,v_q/(1 - v_q + c\,v_q)$, $c = b_\mu/b_\lambda$	–	Beta quantiles mapped	8.3.1
Requirement probability	$I_{v^*}(a_\lambda, a_\mu)$, eq. (8.43)	–	support for $A \ge A_{\text{req}}$	8.3.2
Annual downtime	$8760\,\text{h} \times U$	h	quantiles of U	8.3.3
Outage frequency	$\nu_{\text{out}} = \lambda A$	h^{-1}	failures per calendar hour	8.3.3
Forward equations	$\mathrm{d}P_{\text{up}}/\mathrm{d}t = -\lambda P_{\text{up}} + \mu P_{\text{down}}$; $\mathrm{d}A/\mathrm{d}t = \mu - (\lambda+\mu)A$	h^{-1}	two-state balance	8.3.3
Transient after repair	$A + U\mathrm{e}^{-(\lambda+\mu)t}$	–	from up at $t = 0$	8.3.3
Time constant	$t_c = 1/(\lambda+\mu)$	h	approach to steady state	8.3.3
Interval availability	$A + U(t_c/\tau)(1 - \mathrm{e}^{-\tau/t_c})$	–	average over a horizon	8.3.3
Building blocks				
Record likelihood	$\lambda^k \mathrm{e}^{-\lambda T_{\text{tot}}}\,\mu^r \mathrm{e}^{-\mu D}$	–	factorizes into two rows	8.1.6
Repair likelihood	$\mu^r \mathrm{e}^{-\mu D}$; $D \sim \text{Gamma}(r, \mu)$	–	sufficient statistic (r, D)	8.1.6
Row term	$\mathcal{L}_i = \bar{k}_i - k_i \ln\bar{k}_i$, $r_i = (\bar{k}_i - k_i)/\sqrt{\bar{k}_i}$	–	count, repair or prior row	8.1.6
Newton normal equations	$(\mathbf{J}^\mathsf{T}\mathbf{J})\Delta\mathbf{z} = -\mathbf{J}^\mathsf{T}\mathbf{r}$	–	fit of the log rates	8.2.1
Laplace covariance	$\boldsymbol{\Sigma} = (\mathbf{J}^\mathsf{T}\mathbf{J})^{-1} = \text{diag}(1/a_\lambda,\ 1/a_\mu)$	–	diagonal, no shared row	8.2.1
Repair-rate posterior	$\text{Gamma}(a_{\mu 0} + r,\ b_{\mu 0} + D)$	h^{-1}	repairs to shape, time to rate	8.2.2
Beta route	$b_\lambda\lambda/(b_\lambda\lambda + b_\mu\mu) \sim \text{Beta}(a_\lambda, a_\mu)$	–	exact posterior of A	8.2.2
Log-rate moments	$\psi(a_\lambda) - \ln b_\lambda$, $\psi'(a_\lambda)$	–	exact; Laplace $1/a_\lambda$	8.2.2
Exponential-repair evidence	$b_{\mu 0}^{a_{\mu 0}}\Gamma(a_\mu)/(\Gamma(a_{\mu 0})\,b_\mu^{a_\mu})$	h^{-r}	compares repair models	8.4.3
Preposterior of the logit	$\mathbf{g}^\mathsf{T}(\boldsymbol{\Sigma}^{-1} + \mathbf{J}_a^\mathsf{T}\mathbf{J}_a)^{-1}\mathbf{g}$	–	variance after a record	8.6.2

8.8 Exercises

Exercise 8.1. *Formulation.* A pump has a failure rate of $\lambda = 1/2000\,\text{h}^{-1}$ and a repair rate of $\mu = 1/8\,\text{h}^{-1}$, both taken as known. (a) Compute its steady-state availability and its time constant. (b) The pump has just failed; find the time after which its availability first reaches 99 % of the steady state. (c) Write the question whether its availability is at least 0.99, for uncertain

rates, as an event on the logit ξ. ((a) $A = 0.125/0.1255 = 0.99602$; $t_c = 1/0.1255 = 7.97\,\text{h}$. (b) From $A(t) = A(1 - \mathrm{e}^{-t/t_c})$, $A(t) \geq 0.99A$ requires $\mathrm{e}^{-t/t_c} \leq 0.01$, $t = t_c \ln 100 = 36.7\,\text{h}$. (c) $\Pr(\xi \leq \ln(0.01/0.99) \mid \mathbf{y}) = \Pr(\ln\lambda - \ln\mu \leq -4.595 \mid \mathbf{y})$.)

Exercise 8.2. *Formulation.* The prior on a repair rate is Gamma$(2, 20\,\text{h})$, and three repairs take 4, 5 and 9 hours. Write the likelihood, derive the posterior of μ, and compute the posterior mean repair rate, its reciprocal and the posterior mean time to repair. State which of the last two is the expected repair duration. (Likelihood $\mu^3\mathrm{e}^{-18\mu}$, so the posterior is Gamma$(2+3,\ 20+18)$ = Gamma$(5, 38\,\text{h})$; mean rate $5/38 = 0.132\,\text{h}^{-1}$, reciprocal 7.6 h; posterior mean time to repair $b_\mu/(a_\mu - 1) = 38/4 = 9.5\,\text{h}$, which is the expected duration of a repair; the reciprocal of the mean rate is smaller because $1/\mu$ is convex.)

Exercise 8.3. *Derivation.* (a) A second pump has the failure-rate posterior Gamma$(4, 4000\,\text{h})$, from two failures in 3000 operating hours, and the repair-rate posterior Gamma$(4, 40\,\text{h})$, from two repairs totaling 20 h. Derive $\Pr(A \geq 0.98 \mid \mathbf{y})$ by the Beta route and by the Laplace route. (b) A third pump has the unequal shapes $\lambda \sim$ Gamma$(3, 1500\,\text{h})$ and $\mu \sim$ Gamma$(10, 100\,\text{h})$. Compute its plug-in availability at the posterior means, its posterior median and its posterior mean, and explain why the three differ. ((a) $c = 40/4000 = 0.01$; $v^* = 0.02/(0.02 + 0.01 \times 0.98) = 0.6711$ and $\Pr = I_{0.6711}(4,4) = 0.834$. Laplace: the Newton fit gives $\hat{\mathbf{z}} = (\ln 0.001,\ \ln 0.1)^\mathsf{T}$ and $\mathbf{\Lambda} = \mathbf{J}^\mathsf{T}\mathbf{J} = \text{diag}(4,4)$, so $\mathbf{\Sigma} = \text{diag}(0.25,\ 0.25)$; with $\mathbf{g} = (1,-1)^\mathsf{T}$, $\hat{\xi} = \mathbf{g}^\mathsf{T}\hat{\mathbf{z}} = -4.605$ and $\sigma_\xi = \sqrt{\mathbf{g}^\mathsf{T}\mathbf{\Sigma}\mathbf{g}} = \sqrt{0.5} = 0.707$, $\xi^* = \ln(0.02/0.98) = -3.892$, $\Pr \approx \Phi(0.713/0.707) = \Phi(1.009) = 0.843$; with the exact variance $2\psi'(4) = 0.568$, $\Phi(0.947) = 0.828$. (b) Plug-in $0.1/(0.002 + 0.1) = 0.9804$, the same as the chapter's pump; $c = 100/1500 = 0.0667$ and $V \sim$ Beta$(3, 10)$ with median 0.2167, so the median odds of U are $0.0667 \times 0.2167/0.7833 = 0.0184$ and the median availability 0.9819; the mean, integrating $1 - U(v)$ against the Beta density, is 0.9785. With unequal shapes the Beta distribution is not symmetric, its median is not one half, and the median no longer equals the plug-in value; the mean lies below both because of the long lower tail.)

Exercise 8.4. *Derivation.* For the second pump of part (a) of the previous exercise, derive the median annual downtime and its central 95 % band from the band on U, by the Beta route and by the Laplace route, and state why the band is skewed. (Exact: Beta$(4,4)$ quantiles 0.1841 and 0.8159 give $U \in [0.00225,\ 0.04245]$ by $U_q = cv_q/(1 - v_q + cv_q)$, so the downtime is $[19.7,\ 372]$ h per year about the median $8760 \times 0.01/1.01 = 86.7\,\text{h}$. Laplace: $\xi \in -4.605 \pm 1.96 \times 0.707$ gives $[21.9,\ 337]$ h. The band is nearly symmetric in ξ, and $U \approx \mathrm{e}^\xi$ turns that symmetry into equal ratios, about four, below and above the median, hence a long upper tail in hours.)

Exercise 8.5. *Application.* Two identical pumps are served by one repair crew, whose repair rate μ is the same for both. Each pump has four failures in 2000 operating hours under the prior Gamma$(2, 1000\,\text{h})$, and the crew's eight repairs total 80 h under the prior Gamma$(2, 20\,\text{h})$. (a) Write the state, the log-count Jacobian of the six rows, their expected counts at the MAP state, and the Laplace covariance. (b) Derive the variance of each pump's logit and the correlation between the two logits. (c) State what the shared entry does to the independence used in the chapter, and where the two-state model of each pump breaks. ((a) $\mathbf{z} = (\ln\lambda_1, \ln\lambda_2, \ln\mu)^\mathsf{T}$; the rows are the prior and count rows of each pump and the prior and pooled repair rows of the crew, with the 6×3 log-count Jacobian whose rows are $(1,0,0)$ twice, $(0,1,0)$ twice and $(0,0,1)$ twice; at the MAP state the expected counts are $\bar{\mathbf{k}} = (2,4,2,4,2,8)^\mathsf{T}$, so the working Jacobian carries 1.414, 2, 1.414, 2, 1.414 and 2.828 in those positions, and $\mathbf{J}^\mathsf{T}\mathbf{J} = \text{diag}(6,6,10)$, the exact posteriors being $\lambda_i \sim$ Gamma$(6, 3000\,\text{h})$ and $\mu \sim$ Gamma$(10, 100\,\text{h})$, and $\mathbf{\Sigma} = \text{diag}(0.16667,\ 0.16667,\ 0.10000)$. (b) $\xi_i = z_i - z_3$ with $\mathbf{g}_i = (1,0,-1)^\mathsf{T}$ and $(0,1,-1)^\mathsf{T}$: $\text{Var}(\xi_i) = 1/6 + 1/10 = 0.267$, $\sigma = 0.516$, against 0.577 for a crew per pump; $\text{Cov}(\xi_1, \xi_2) = \mathbf{g}_1^\mathsf{T}\mathbf{\Sigma}\mathbf{g}_2 = 1/10$, correlation $0.1/0.267 = 0.375$. (c) The pooled repairs narrow each pump's band, but the two availabilities are now correlated through the shared entry, so a statement about

both pumps together must use the joint covariance, not the product of the marginals; and when both pumps are down at once the second waits for the crew, so the repair rate of each is no longer μ and a joint model with the crew's queue, a multistate model of the kind treated in later chapters, is needed.)

Exercise 8.6. *Application.* For the chapter's pump, a further operating record costs R$ 50 per operating hour, in monitoring and in the delay of the decision, and a repair drill costs R$ 4 000. (a) Compute the reduction in the Laplace variance of ξ per R$ 100 000 for 2000 more hours and for ten drills. (b) Find the cheapest plan of drills alone and the cheapest plan of operation alone that bring σ_ξ down to 0.45, and state which is cheaper. (c) State the smallest σ_ξ that drills alone can reach, and what a target below it requires. ((a) 2000 hours cost R$ 100 000 and reduce σ_ξ^2 from 0.3333 to 0.2000, a reduction of 0.133; ten drills cost R$ 40 000 and reduce it to 0.2292, 0.104, which is 0.260 per R$ 100 000, so the drills buy twice the variance reduction per real although they remove less in total. (b) Drills: n drills give $\mathbf{\Lambda}_a = \operatorname{diag}(0, n)$, $\mathbf{\Sigma}_a = \operatorname{diag}(1/6,\ 1/(6+n))$, and $\mathbf{g}^\mathsf{T}\mathbf{\Sigma}_a\mathbf{g} \le 0.45^2 = 0.2025$ needs $n = 22$ drills, R$ 88 000; operation with x expected failures gives $\mathbf{\Lambda}_a = \operatorname{diag}(x, x)$, $\mathbf{g}^\mathsf{T}\mathbf{\Sigma}_a\mathbf{g} = 2/(6+x) \le 0.2025$, which needs $x = 3.88$, about 1938 h, R$ 96 900. The drills are cheaper, by R$ 8 900, and much faster. (c) However many drills are run, the first column of $\mathbf{\Lambda}_a$ stays zero, so $\Sigma_{a,11} = 1/6$ and $\sigma_\xi > 1/\sqrt{6} = 0.408$; a target below 0.408 requires operating hours, alone or together with drills.)

Part III

Predicting Future Reliability

Chapter 9

Mission Reliability

9.1 Mission reliability under an uncertain failure rate

The emergency diesel generator of a data center must carry the load through a grid outage, and the design basis plans for an outage of 72 hours: the generator must start when the grid is lost and then run until it returns. The reliability engineer reviewing the design basis has two records. In forty start tests the generator failed to start once. In 480 hours of load runs it failed to run twice. The design basis asks that the generator's reliability over the 72-hour run be at least 0.90, shown with 95 % assurance, and the first reading of the run record is simple: two failures in 480 hours is one failure in 240 hours, and $\mathrm{e}^{-72/240} = 0.741$, about three chances in four that the generator runs through the outage. That is short of 0.90, and the shortfall looks like a matter of a better estimate of one number.

The engineer who must sign the review asks something else. The 0.741 uses a rate estimated from two failures as if it were known, it leaves out the start, and it does not say whether it is the probability that the next outage is carried or the generator's reliability, a property of the unit with one true value that the records know only roughly. The question of this chapter is

how the probability that a standby unit starts and survives a mission is derived when its failure probability and its failure rate are known only through a posterior, and whether the unit meets a mission requirement with stated assurance.

It is answered with the method of the previous chapters on one generator. What changes is the quantity asked about: a survival function of the rate, curved in it, whose average over the posterior is not its value at an estimate; and two records, a start record on the logit and a run record on the log rate, which feed one mission.

Jargon: mission, mission time, failure to start, failure to run, predictive reliability, plug-in reliability, Jensen's inequality, convex function, quantile map, decreasing failure rate (DFR), frailty, complementary log-log, Gauss–Hermite quadrature, reliability demonstration test, assurance

A *mission* is a task with a stated start and a stated duration, the *mission time*; here the mission is to carry the load through a grid outage. A standby unit can fail in two ways during a mission: by a *failure to start*, an event on demand with a probability per demand, and by a *failure to run*, an event in time with a rate per running hour. The *predictive reliability* is the probability of surviving the mission given the evidence, the survival function averaged over the posterior of its parameters; the *plug-in reliability* is the survival function evaluated at a point estimate of the parameters. A function is *convex* when every chord lies above its graph, and *Jensen's inequality* states that the expectation of a convex function is at least the function of the expectation. The *quantile map* is the rule that a monotone function carries the quantiles of a variable to the quantiles of the function, in the same order if it is increasing and in reverse order if it is decreasing. A distribution has a *decreasing failure rate* (DFR) when its hazard falls with age; in survival analysis the selection effect that produces a DFR in a population of units with different rates is called *frailty*. The *complementary log-log* of a probability R is $\ln(-\ln R)$. *Gauss–Hermite quadrature* evaluates an expectation over a normal distribution as a weighted sum of the integrand at a few fixed nodes. A *reliability demonstration test* is a test planned so that, if it ends with a stated outcome, usually zero failures, a reliability requirement is supported with a stated posterior probability, the *assurance*.

Key idea

A mission reliability computed from a posterior is a posterior expectation of a survival function, $\Pr(T > t \mid \mathbf{y}) = \mathbb{E}[\mathrm{e}^{-\lambda t} \mid \mathbf{y}]$, and it is not the survival function at the posterior mean rate: the two differ by an amount that grows with $\frac{1}{2}t^2 \operatorname{Var}(\lambda \mid \mathbf{y})$. The reliability itself, $R(t;\lambda) = \mathrm{e}^{-\lambda t}$, is a derived quantity with a posterior of its own, whose mean is the predictive and whose tail answers a requirement; the two objects meet at the mean and nowhere else.

9.1.1 A reliability averaged over what is not known

For a component with a constant failure rate λ, the probability of surviving a mission of length t is $R(t;\lambda) = \mathrm{e}^{-\lambda t}$, the exponential survival function derived in Chapter 7, equation (7.12). The semicolon in $R(t;\lambda)$ separates the argument, the mission length, from the parameter on which the function depends. If λ were known, that would be the end of the question. It is known only through a record, and the posterior $p(\lambda \mid \mathbf{y})$ of Chapter 6 describes what the record says about it. The law of total probability then gives the probability of surviving the mission, all that is known taken into account, as the survival function averaged over the posterior,

$$\Pr(T > t \mid \mathbf{y}) = \int_0^\infty \Pr(T > t \mid \lambda)\, p(\lambda \mid \mathbf{y})\, \mathrm{d}\lambda = \int_0^\infty \mathrm{e}^{-\lambda t}\, p(\lambda \mid \mathbf{y})\, \mathrm{d}\lambda = \mathbb{E}\big[R(t;\lambda) \mid \mathbf{y}\big]. \quad (9.1)$$

This is the *predictive reliability*, written $R_{\text{pred}}(t)$. It is one number for each mission length, and it contains both kinds of uncertainty: the random occurrence of a failure for a given rate, which is aleatory, and the incomplete knowledge of the rate, which is epistemic. Chapters 6 and 7 each computed it once, in passing; here it is the subject.

The shortcut that replaces equation (9.1) in much practice substitutes an estimate of the rate, usually its posterior mean, into the survival function,

$$R_{\text{plug-in}}(t) = R(t; \mathbb{E}[\lambda \mid \mathbf{y}]) = \mathrm{e}^{-\mathbb{E}[\lambda \mid \mathbf{y}]\, t}. \quad (9.2)$$

It discards the epistemic part. Section 9.3.1 derives how the two differ and by how much; the first reading of the record, the survival at the record's own rate, is the same shortcut at another estimate.

9.1.2 The predictive and the posterior of the reliability

Equation (9.1) averages over the rate and returns a single probability. The reliability $R(t;\lambda)$ is also a quantity in its own right: a function of the rate, which has one true value for this generator, and therefore a quantity with one true value, known only through the posterior. Its distribution $p(R(t;\lambda) \mid \mathbf{y})$ is the posterior of a derived quantity of the state, in the sense of Chapter 4. The two objects answer different questions. The predictive answers "what is the probability that the generator gets through the next outage?" and is the number to multiply by the consequence of a lost outage in an expected loss. The posterior of the reliability answers "what is the generator's reliability, and how well is it known?" and its tail answers "is the reliability at least the required value?". A requirement in a contract or a design basis is a statement about a property of the unit, the reliability, and it is answered from the second object, not from the first. The two meet only at the mean: by equation (9.1) the predictive is the mean of the posterior of $R(t;\lambda)$, and its median, band and tails are not functions of it.

9.1.3 The physical system and its graph (step 0)

The data center is supplied by the utility grid through a transfer switch. When the grid is lost, the emergency diesel generator receives a start demand; if it starts, the switch transfers the load to it, and the generator must then run until the grid returns. The generator's function in an outage therefore has two parts, a start and a run, and the record of each part is kept separately: periodic start tests, each a demand with a binary outcome, and load runs, whose failures to run are counted against the accumulated running hours.

Figure 9.1 draws the generator as a graph. The gray boxes on the top row are the utility grid, the transfer switch and the data-center load, pinned surroundings, and the solid arrows are the power paths. The green box below the switch is the generator, the free node, which carries the two entries of the state: the logit of the start failure probability and the log of the run failure rate. The two dashed arrows leave the generator for the evidence: the start record on the left, with its prior, and the run record on the right, with its prior. The purple box on the right is the question's extension of the graph: the mission, a start demand followed by 72 hours of running, which the plant graph does not contain until the question is asked.

The state entries. The start failure probability must lie in $(0,1)$ and the run failure rate must be positive, and the graph carries each on an unbounded scale, as Chapters 5 and 8 did:

$$\mathbf{z} = \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} = \begin{bmatrix} \xi \\ \ln\lambda \end{bmatrix} \quad \begin{matrix} \text{logit of the start failure probability} \\ \text{log run failure rate} \end{matrix}, \qquad \xi = \ln\frac{p}{1-p}, \tag{9.3}$$

with λ in h^{-1}. Chapter 6 carried the rate with a reference value, $\omega = \ln(\lambda/\lambda_{\mathrm{ref}})$; as in Chapter 8 the reference is dropped, since it shifts the entry by a constant and changes no variance. The probability p, the rate λ, the mission reliability and every quantity the questions ask for are derived quantities, functions of the state; the exact posteriors of p and λ that this problem admits are those of two functions of its entries. The evidence $\mathbf{y}$ is the two records and the two priors.

9.1.4 The two records

The start record. The generator has been demanded $n = 40$ times in start tests and failed to start $k_{\mathrm{s}} = 1$ time. The prior on the start failure probability, from experience with similar

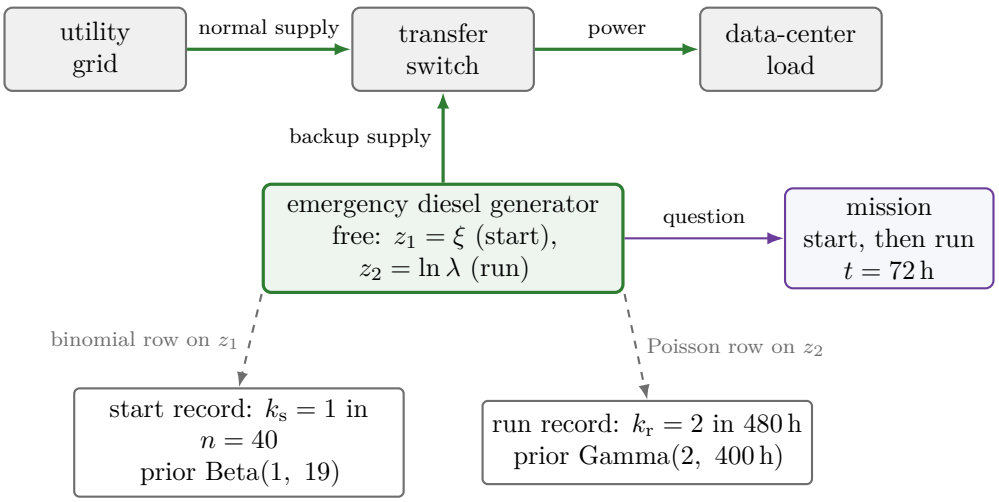


Figure 9.1: The plant graph of the emergency generator. Free: the generator node (green), carrying the logit $z_1 = \xi$ of the start failure probability and the log run failure rate $z_2 = \ln \lambda$. Pinned: the utility grid, the transfer switch and the data-center load (gray). Observed: the start record, one failure in forty demands, and the run record, two failures in 480 running hours (dashed arrows), each with the prior on its parameter. The mission (purple) is the extension of the graph that the question adds.

units, is $p \sim \mathrm{Beta}(a_{p0}, b_{p0}) = \mathrm{Beta}(1, 19)$, of mean $1/20 = 0.05$, which weighs as much as one failure in twenty demands.

The run record. In load runs the generator has accumulated $T_\mathrm{tot} = 480$ running hours and failed to run $k_\mathrm{r} = 2$ times. The prior on the run failure rate is $\lambda \sim \mathrm{Gamma}(a_{\lambda 0}, b_{\lambda 0}) = \mathrm{Gamma}(2, 400\,\mathrm{h})$, of mean $0.005\,\mathrm{h}^{-1}$, which weighs as much as two failures in four hundred hours. As in the earlier chapters the gamma distribution has a shape and a rate parameter, the rate parameter in hours, and the subscript 0 marks the prior's parameters and its absence the posterior's. Both records are illustrative.

9.1.5 The probabilistic questions (step 1)

The engineer asks six questions of the two records, and step 1 writes each as a probability or an expectation over the state, with the unknown on the left of the bar and the records on the right. Table 9.1 lists them: the first column is the question in engineering language; the second the general form that every question of its category takes, in the book's notation for a quantity $f(\mathbf{z})$ derived from the state, a threshold f^* on it, an assurance level κ, the posterior density $p(\mathbf{z} \mid \mathbf{y})$ and a future observation y_new; the third its mathematical form for the generator; the fourth its category; and the last the subsection where it is answered. The first two are the chapter's central distinction, a predictive and the posterior of a derived quantity, and the third is answered from the second object, not from the first. A seventh question, which test would narrow the mission reliability more, twenty start tests or two hundred run hours, asks nothing new of the present records: it is the next measurement for these six, answered in Section 9.6. The atlas of Appendix A lists the same forms.

9.1.6 The evidence as rows (step 2)

Step 2 attaches four rows to the graph. Two act on z_1: the prior row of the start failure probability and the binomial row of the start record, as in Chapter 5. Two act on z_2: the prior row of the run failure rate and the Poisson count row of the run record, as in Chapter 6. Which

Table 9.1: The questions asked of the generator's two records, the general form of their category, their mathematical form for the generator, their category and where each is answered. $\mathbf{y}$ is the start record, the run record and the two priors; t the mission length, 72 h for the outage; R_{req} a required reliability; R_{sr} the probability of starting and then running t hours; T' the length of a further run.

Engineering question	General form	Mathematical form	Category	Answered in
Does the generator, once started, run through a 72 h outage?	$\int p(y_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	$\Pr(T > 72\,\text{h} \mid \mathbf{y}) = \mathbb{E}[\mathrm{e}^{-72\lambda} \mid \mathbf{y}]$	prediction under a threshold	9.3.1
How well is its 72 h reliability known?	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$, quantiles	$p(R(72;\lambda) \mid \mathbf{y})$ and its band	estimation of a derived quantity	9.3.2
Can $R(72) \ge 0.90$ be asserted with 95 % assurance?	$\Pr(f(\mathbf{z}) \ge f^* \mid \mathbf{y}) \ge \kappa$	$\Pr(R(72;\lambda) \ge 0.90 \mid \mathbf{y}) \ge 0.95$	prediction under a threshold	9.3.3
What outage length can be covered at a reliability of 0.95, and of 0.90?	the predictive and the tail, inverted in t	t^* with $\Pr(T > t^* \mid \mathbf{y}) = R_{\text{req}}$; t with $\Pr(R(t;\lambda) \ge R_{\text{req}} \mid \mathbf{y}) = 0.95$	prediction under a threshold	9.3.4
Does the generator start and run 72 h?	$\int p(y_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$; $\mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y})$, quantiles	$\mathbb{E}[(1-p)\,\mathrm{e}^{-72\lambda} \mid \mathbf{y}]$; band and budget of $R_{\text{sr}}(72)$	prediction; estimation of a derived quantity	9.3.5
How many failure-free run hours would demonstrate $R(72) \ge 0.90$ at 95 %?	preposterior $\sum_j \Pr(K = j \mid \mathbf{y})\, \mathbf{1}[\text{rule met}]$	smallest T' with $\Pr(R(72;\lambda) \ge 0.90 \mid \mathbf{y}, 0 \text{ in } T') \ge 0.95$	design	9.3.6

row acts on which entry is a property of the graph, and it fixes where the Jacobian of the rows can be nonzero, marked $*$:

$$\mathbf{J} = \begin{bmatrix} * & 0 \\ * & 0 \\ 0 & * \\ 0 & * \end{bmatrix} \begin{matrix} \text{prior on } p \\ \text{start record} \\ \text{prior on } \lambda \\ \text{run record} \end{matrix}, \qquad \text{columns } z_1, z_2. \tag{9.4}$$

No row acts on both entries, at any state. No row is a physics row either: the mission reliability enters as a derived quantity of the state, not as an entry tied to the parameters by a constraint, so no model tolerance σ_p arises. The mission adds no row at all: it is a question asked of the state, not evidence about it, and it enters the answers only through the function $(1-p)\,\mathrm{e}^{-\lambda t}$ whose expectation and distribution the questions ask for. None of the four rows is Gaussian: each is a count row of Section 4.2.7, entering the objective as its negative log-probability term and the Jacobian through its working pair, below.

The modeling choices. Three modeling choices enter the model. The experience with similar units is carried as the prior Gamma$(2, 400\,\mathrm{h})$ on the run rate, and Beta$(1, 19)$ on the start probability: this gives both posteriors a closed form and lets the experience weigh as much as 400 running hours, nearly as much as the run record itself, and it misleads if that experience came from units unlike this one, which Section 9.4.4 measures. The run failure rate is taken as constant over the run, the same in the first hour as in the seventy-second: this is what makes the survival $\mathrm{e}^{-\lambda t}$ and gives the predictive a closed form in the two gamma parameters, and it fails for an engine whose failures to run crowd into the first hour of loading and warming (Section 9.6.6). The load runs are taken to represent an outage: this lets the 480 hours of the record speak for the 72 hours of the mission, and it fails if the load runs are made at reduced load or on a warm engine (Section 9.6.6).

The rows as negative log-probability terms. The four rows act on the state $\mathbf{z} = (z_1, z_2)^{\mathsf{T}}$ of equation (9.3). None is Gaussian; each enters the objective as its negative log-probability term, $\mathcal{L}_i(\mathbf{z}) = -\ln p(\mathbf{y}_i \mid \mathbf{z})$ up to a constant, and $\mathcal{L} = \sum_i \mathcal{L}_i$. The start rows are those of Chapter 5 in the logit, equation (5.17): the prior Beta(a_{p0}, b_{p0}), written as a density of the logit with the change-of-variable factor $\mathrm{d}p/\mathrm{d}\xi = p(1-p)$ included, is the term $-a_{p0}\ln p - b_{p0}\ln(1-p)$, and the start record is the term $-k_\mathrm{s}\ln p - (n-k_\mathrm{s})\ln(1-p)$. The run rows are the count rows of Chapter 6, equation (6.21): the gamma prior in the log rate, $b_{\lambda 0}\mathrm{e}^{z_2} - a_{\lambda 0} z_2$, with the shape $a_{\lambda 0}$ rather than $a_{\lambda 0} - 1$ because of the change-of-variable factor $\mathrm{d}\lambda/\mathrm{d}z_2 = \lambda$, and the count row $T_\mathrm{tot}\mathrm{e}^{z_2} - k_\mathrm{r} z_2$. Indexing the rows by $i = 1, \ldots, 4$, each is an observed count k_i with an expected count $\bar{k}_i$ at the state,

$$\mathcal{L}_i(\mathbf{z}) = \begin{cases} -k_i \ln p - (N_i - k_i)\ln(1-p), & \bar{k}_i = N_i\, p, \quad i = 1, 2, \\ \bar{k}_i - k_i \ln \bar{k}_i, & \bar{k}_i = T_{\mathrm{tot},i}\,\mathrm{e}^{z_2}, \quad i = 3, 4, \end{cases} \tag{9.5}$$

$$\begin{array}{lll} i = 1: & \text{prior on } p, & (k_1, N_1) = (1,\ 20), \\ i = 2: & \text{start record}, & (k_2, N_2) = (1,\ 40), \\ i = 3: & \text{prior on } \lambda, & (k_3, T_{\mathrm{tot},3}) = (2,\ 400\,\mathrm{h}), \\ i = 4: & \text{run record}, & (k_4, T_{\mathrm{tot},4}) = (2,\ 480\,\mathrm{h}), \end{array}$$

with $p = 1/(1 + \mathrm{e}^{-z_1})$, k_i the observed count of the row (failures to start or to run, or a prior's shape counted as pseudo-failures), N_i the number of demands of a start row (the prior's $a_{p0} + b_{p0} = 20$ pseudo-demands, the forty tests) and $T_{\mathrm{tot},i}$ the exposure T_tot of a run row (the

prior's rate parameter, the running hours). Summing the four terms gives the objective

$$\begin{aligned}\mathcal{L}(\mathbf{z}) &= \underbrace{20\ln(1+\mathrm{e}^{z_1})-z_1}_{\text{prior on } p}+\underbrace{40\ln(1+\mathrm{e}^{z_1})-z_1}_{\text{start record}}+\underbrace{400\,\mathrm{e}^{z_2}-2z_2}_{\text{prior on } \lambda}+\underbrace{480\,\mathrm{e}^{z_2}-2z_2}_{\text{run record}} \\ &= (a_p+b_p)\ln(1+\mathrm{e}^{z_1})-a_pz_1+b_\lambda\mathrm{e}^{z_2}-a_\lambda z_2,\end{aligned} \tag{9.6}$$

up to a constant, with the exposures in hours and the posterior parameters $a_p = a_{p0} + k_\mathrm{s}$, $b_p = b_{p0} + n - k_\mathrm{s}$, $a_\lambda = a_{\lambda0} + k_\mathrm{r}$ and $b_\lambda = b_{\lambda0} + T_\mathrm{tot}$ of Section 9.2.2.

The working residuals and the Jacobian. The count of a start row has the mean $\bar{k}_i = N_i p$ and the variance $v_i = N_i p(1-p)$; that of a run row has the mean and the variance $v_i = \bar{k}_i$. Each row joins the Jacobian through the working pair (4.43) of Section 4.2.7, in its binomial-in-the-logit form for a start row and its Poisson form for a run row; written with the count's variance v_i it is

$$r_i = \frac{\bar{k}_i - k_i}{\sqrt{v_i}}, \qquad \mathbf{J}_i = \frac{1}{\sqrt{v_i}}\,\frac{\partial \bar{k}_i}{\partial \mathbf{z}^\mathsf{T}},$$

$$\mathbf{r}(\mathbf{z}) = \begin{bmatrix} (20\,p-1)/\sqrt{20\,p(1-p)} \\ (40\,p-1)/\sqrt{40\,p(1-p)} \\ (400\,\mathrm{e}^{z_2}-2)/\sqrt{400\,\mathrm{e}^{z_2}} \\ (480\,\mathrm{e}^{z_2}-2)/\sqrt{480\,\mathrm{e}^{z_2}} \end{bmatrix}, \quad \mathbf{J}(\mathbf{z}) = \begin{bmatrix} \sqrt{20\,p(1-p)} & 0 \\ \sqrt{40\,p(1-p)} & 0 \\ 0 & \sqrt{400\,\mathrm{e}^{z_2}} \\ 0 & \sqrt{480\,\mathrm{e}^{z_2}} \end{bmatrix}, \tag{9.7}$$

since $\partial\bar{k}_i/\partial z_1 = N_i p(1-p) = v_i$ for a start row and $\partial\bar{k}_i/\partial z_2 = \bar{k}_i = v_i$ for a run row, so that every nonzero entry of $\mathbf{J}$ is $\sqrt{v_i}$, in the pattern (9.4): the first column holds the two start rows and nothing else, the second the two run rows and nothing else. Then $\mathbf{J}_i^\mathsf{T} r_i = \bar{k}_i - k_i$ in its entry is the gradient of $\mathcal{L}_i$, and the gradient and the Hessian of the objective are

$$\frac{\partial \mathcal{L}}{\partial \mathbf{z}} = \mathbf{J}^\mathsf{T}\mathbf{r}, \qquad \frac{\partial^2 \mathcal{L}}{\partial \mathbf{z}\,\partial \mathbf{z}^\mathsf{T}} = \mathbf{J}^\mathsf{T}\mathbf{J}. \tag{9.8}$$

The Hessian is exact. The curvature term $\mathbf{C}$ of equation (4.44) is zero, because the row's natural parameter, the logit for a start row and the log rate for a run row, is linear in the state, the canonical case, so Newton's method and the Gauss–Newton step coincide. As Section 4.2.7 notes, $\frac{1}{2}\|\mathbf{r}\|^2$ is not the objective (9.6): the two share their gradient and their Hessian, not their value. With every row normalized in this way the information matrix is $\mathbf{J}^\mathsf{T}\mathbf{J}$.

9.2 The posterior of the state (step 3)

In the book's notation, $p(\mathbf{z} \mid \mathbf{y})$ is the posterior density of the state, $f(\mathbf{z})$ a quantity derived from the state, f^* a requirement level on it and y_new a future observation.

All six are read from one posterior: the posterior of the state that the four rows of step 2 give, exactly a Beta density of the start failure probability and an independent gamma density of the run failure rate, and approximately the Laplace normal of the logit and the log rate. This section derives it and computes it for the generator; Section 9.3 reads every answer from it.

Four objects of the fit recur in the specialized forms. The *MAP state* $\hat{\mathbf{z}} = (\hat{z}_1, \hat{z}_2)^\mathsf{T}$ is the minimum of the objective that the four rows put on the state; the *information matrix* $\mathbf{\Lambda} = \mathbf{J}^\mathsf{T}\mathbf{J}$ is the curvature of that objective at the minimum, with $\mathbf{J}$ the 4×2 Jacobian of the normalized rows, of the pattern (9.4); the *Laplace covariance* $\mathbf{\Sigma} = \mathbf{\Lambda}^{-1}$, a 2×2 matrix, makes the posterior of the state approximately normal, $\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$; and a derived quantity f is carried by its *gradient* $\mathbf{g} = \partial f/\partial \mathbf{z}$ at $\hat{\mathbf{z}}$, with the variance $\mathbf{g}^\mathsf{T}\mathbf{\Sigma}\mathbf{g}$. Every answer of step 4 is computed from $\hat{\mathbf{z}}$ and $\mathbf{\Sigma}$. Beside them stands the exact posterior that this problem admits: the start failure probability and the run failure rate are independent a posteriori, $p \mid \mathbf{y} \sim \mathrm{Beta}(a_p, b_p)$

and $\lambda \mid \mathbf{y} \sim \text{Gamma}(a_\lambda, b_\lambda)$, with the shapes the prior shapes plus the counts. It gives, in closed form or by a one-dimensional quadrature, with draws as a check, the value set beside each answer.

9.2.1 The posterior from the four rows

Each Newton step solves the normal equations of Chapter 3, equation (3.24),

$$(\mathbf{J}^\mathsf{T}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^\mathsf{T}\mathbf{r}, \qquad \mathbf{z} \leftarrow \mathbf{z} + \Delta\mathbf{z}, \tag{9.9}$$

with $\mathbf{J}$ and $\mathbf{r}$ evaluated at the current state; since the columns of $\mathbf{J}$ are disjoint, $\mathbf{J}^\mathsf{T}\mathbf{J}$ is diagonal at every state.

The MAP state. The iteration stops where $\mathbf{J}^\mathsf{T}\mathbf{r} = \mathbf{0}$, that is where $(a_p + b_p)\,p = a_p$ and $b_\lambda \mathrm{e}^{z_2} = a_\lambda$:

$$\hat{\mathbf{z}} = \begin{bmatrix} \hat{z}_1 \\ \hat{z}_2 \end{bmatrix} = \begin{bmatrix} \ln(a_p/b_p) \\ \ln(a_\lambda/b_\lambda) \end{bmatrix}, \qquad \hat{p} = \frac{a_p}{a_p + b_p}, \quad \hat{\lambda} = \frac{a_\lambda}{b_\lambda}, \tag{9.10}$$

the posterior means of the exact posteriors of p and λ (Section 9.2.2). In the logit and in the log rate the mode sits at the posterior mean of each parameter, not at its mode, because the change-of-variable factor has added one to each exponent. In each column the prior row and the record row pull with equal and opposite gradients $\bar{k}_i - k_i$; no residual need be zero by itself.

The information and the covariance. At the MAP state the variances of the counts are the binomial variances $N_i\hat{p}(1-\hat{p})$ of the two start rows and the expected counts $T_{\text{tot},i}\hat{\lambda}$ of the two run rows, and the information and covariance matrices of Chapter 3, equation (3.25), are

$$\mathbf{\Lambda} = \mathbf{J}^\mathsf{T}\mathbf{J} = \begin{bmatrix} (a_p + b_p)\,\hat{p}(1-\hat{p}) & 0 \\ 0 & a_\lambda \end{bmatrix} = \begin{bmatrix} \dfrac{a_p b_p}{a_p + b_p} & 0 \\ 0 & a_\lambda \end{bmatrix}, \qquad \mathbf{\Sigma} = \mathbf{\Lambda}^{-1} = \begin{bmatrix} \dfrac{a_p + b_p}{a_p b_p} & 0 \\ 0 & \dfrac{1}{a_\lambda} \end{bmatrix}, \tag{9.11}$$

since $(a_p + b_p)\hat{p}(1-\hat{p}) = a_p b_p/(a_p + b_p)$ and $(b_{\lambda 0} + T_{\text{tot}})\hat{\lambda} = a_\lambda$. Each diagonal entry of $\mathbf{\Lambda}$ is the sum of the squared entries of its column of $\mathbf{J}$: the precisions of independent rows add. The Laplace posterior is

$$\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma}), \qquad z_2 \mid \mathbf{y} \approx \mathcal{N}\!\left(\ln\frac{a_\lambda}{b_\lambda},\ \frac{1}{a_\lambda}\right), \tag{9.12}$$

each entry read off as $\hat{z}_i$ with the spread $\sqrt{\Sigma_{ii}}$, and the correlation $\Sigma_{12}/\sqrt{\Sigma_{11}\Sigma_{22}}$ zero, computed from the rows.

9.2.2 The exact posteriors of the start and the run

The conjugate updates. By the Beta–binomial update of Chapter 5, equation (5.24), and the Gamma–Poisson update of Chapter 6, equation (6.30),

$$p \mid k_\text{s}, n \sim \text{Beta}(a_p, b_p), \qquad a_p = a_{p0} + k_\text{s}, \quad b_p = b_{p0} + n - k_\text{s}, \tag{9.13}$$

$$\lambda \mid k_\text{r}, T_\text{tot} \sim \text{Gamma}(a_\lambda, b_\lambda), \qquad a_\lambda = a_{\lambda 0} + k_\text{r}, \quad b_\lambda = b_{\lambda 0} + T_\text{tot}. \tag{9.14}$$

The two posteriors are independent. The likelihood of the start record depends only on p and that of the run record only on λ, the priors are independent, and so, by the argument of Chapter 8 for the failure and repair records, the joint posterior is the product $p(p \mid \mathbf{y})\,p(\lambda \mid \mathbf{y})$. The objective (9.6) is the negative logarithm of that joint density written as a density of the state, and the minimum (9.10) lies at the logarithms of the two means.

The moments of the log rate. The log of a gamma variable has the exact mean and variance of Chapter 8, equation (8.19): $\mathbb{E}[\ln\lambda \mid \mathbf{y}] = \psi(a_\lambda) - \ln b_\lambda$ and $\mathrm{Var}(\ln\lambda \mid \mathbf{y}) = \psi'(a_\lambda)$, with ψ the digamma function and ψ' the trigamma function, the first and second derivatives of $\ln\Gamma$. For an integer shape both have closed forms, $\psi(a) = \sum_{j=1}^{a-1} 1/j - \gamma_{\mathrm{E}}$ with $\gamma_{\mathrm{E}} = 0.5772$ the Euler–Mascheroni constant, and $\psi'(a) = \pi^2/6 - \sum_{j=1}^{a-1} 1/j^2$. By the expansions (8.20) the Laplace center $\ln(a_\lambda/b_\lambda)$ and variance $1/a_\lambda$ of equation (9.12) are their leading terms: the Laplace normal centers the log rate about $1/(2a_\lambda)$ above its exact mean and understates its variance by about $1/(2a_\lambda^2)$, errors that fade only as the counts grow.

The log moment of a Beta variable. The start term of equation (9.51) follows from the same moments. By the construction of Chapter 8, $1-p = X/(X+Y)$ with independent $X \sim \mathrm{Gamma}(b_p, 1)$ and $Y \sim \mathrm{Gamma}(a_p, 1)$, the share is independent of the sum $X+Y \sim \mathrm{Gamma}(a_p+b_p, 1)$, and $\ln X = \ln(1-p) + \ln(X+Y)$ gives $\psi'(b_p) = \mathrm{Var}(\ln(1-p)) + \psi'(a_p+b_p)$:

$$\mathrm{Var}\big(\ln(1-p) \mid \mathbf{y}\big) = \psi'(b_p) - \psi'(a_p + b_p). \tag{9.15}$$

9.2.3 The posterior of the generator

The generator's inputs are the start record, $k_{\mathrm{s}} = 1$ failure in $n = 40$ demands under the prior $\mathrm{Beta}(1, 19)$, and the run record, $k_{\mathrm{r}} = 2$ failures in $T_{\mathrm{tot}} = 480$ running hours under the prior $\mathrm{Gamma}(2, 400\,\mathrm{h})$.

The Newton iteration. The iteration of equation (9.9) starts at the prior means, $\mathbf{z}^{(0)} = (\ln(1/19),\ \ln 0.005)^{\mathsf{T}}$, that is $(-2.9444,\ -5.2983)^{\mathsf{T}}$. There the expected counts are $\bar{\mathbf{k}} = (1,\ 2,\ 2,\ 2.4)^{\mathsf{T}}$: each prior row expects exactly its own pseudo-failures, the start record expects two failures in forty demands and saw one, and the run record expects 2.4 failures in 480 h and saw two, so that

$$\mathbf{r}(\mathbf{z}^{(0)}) = \begin{bmatrix} 0 \\ 0.7255 \\ 0 \\ 0.2582 \end{bmatrix}, \qquad \mathbf{J}(\mathbf{z}^{(0)}) = \begin{bmatrix} 0.9747 & 0 \\ 1.3784 & 0 \\ 0 & 1.4142 \\ 0 & 1.5492 \end{bmatrix},$$

and the first step solves

$$\begin{bmatrix} 2.85 & 0 \\ 0 & 4.40 \end{bmatrix} \Delta\mathbf{z} = -\begin{bmatrix} 1.0 \\ 0.4 \end{bmatrix}, \qquad \Delta\mathbf{z} = \begin{bmatrix} -0.3509 \\ -0.0909 \end{bmatrix}.$$

Table 9.2 lists the iteration: each row is one iterate; the columns give the two entries, the two entries of the gradient $\mathbf{J}^{\mathsf{T}}\mathbf{r}$ and the two diagonal entries of the Hessian. Both entries move toward the minimum from the first step without overshooting, and the convergence is quadratic: the gradient falls below 10^{-11} at the fourth iterate.

Table 9.2: Newton iteration of equation (9.9) for the generator's state from the prior means, $\mathbf{z}^{(0)} = (\ln(1/19),\ \ln 0.005)^{\mathsf{T}}$. The Hessian $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ is diagonal at every iterate.

Iteration [–]	z_1 [–]	z_2 [–]	$\partial\mathcal{L}/\partial z_1$ [–]	$\partial\mathcal{L}/\partial z_2$ [–]	$(\mathbf{J}^{\mathsf{T}}\mathbf{J})_{11}$ [–]	$(\mathbf{J}^{\mathsf{T}}\mathbf{J})_{22}$ [–]
0	−2.9444	−5.2983	1.000	0.400	2.8500	4.4000
1	−3.2953	−5.3892	0.144	1.8×10^{-2}	2.0673	4.0176
2	−3.3649	−5.3936	4.6×10^{-3}	3.9×10^{-5}	1.9376	4.0000
3	−3.3673	−5.3936	5.0×10^{-6}	1.9×10^{-10}	1.9333	4.0000
4	−3.3673	−5.3936	6.1×10^{-12}	$< 10^{-14}$	1.9333	4.0000

The MAP state. With $a_p = 2$, $b_p = 58$, $a_\lambda = 4$ and $b_\lambda = 880\,\mathrm{h}$, equation (9.10) gives

$$\hat{\mathbf{z}} = \begin{bmatrix} \ln(2/58) \\ \ln(4/880) \end{bmatrix} = \begin{bmatrix} -3.3673 \\ -5.3936 \end{bmatrix}, \qquad \hat{p} = \frac{1}{30}, \quad \hat{\lambda} = 0.004545\,\mathrm{h}^{-1}. \tag{9.16}$$

At $\hat{\mathbf{z}}$ the residual vector is $\mathbf{r}(\hat{\mathbf{z}}) = (-0.4152,\ 0.2936,\ -0.1348,\ 0.1231)^{\mathsf{T}}$: its projections $\mathbf{J}^{\mathsf{T}}\mathbf{r}$ are zero, but no residual is zero by itself. The prior on p expects $20/30 = 0.667$ pseudo-failures against its one and the start record $40/30 = 1.333$ failures against its one; the prior on λ expects 1.818 against two and the run record 2.182 against two.

The Jacobian, the information and the covariance. At the MAP state the variances of the counts are $\mathbf{v} = (0.6444,\ 1.2889,\ 1.8182,\ 2.1818)^{\mathsf{T}}$, and equation (9.11) gives

$$\mathbf{J}(\hat{\mathbf{z}}) = \begin{bmatrix} 0.8028 & 0 \\ 1.1353 & 0 \\ 0 & 1.3484 \\ 0 & 1.4771 \end{bmatrix}, \qquad \mathbf{\Lambda} = \mathbf{J}^{\mathsf{T}}\mathbf{J} = \begin{bmatrix} 0.6444 + 1.2889 & 0 \\ 0 & 1.8182 + 2.1818 \end{bmatrix},$$
$$\mathbf{\Lambda} = \begin{bmatrix} 1.9333 & 0 \\ 0 & 4 \end{bmatrix}, \qquad \mathbf{\Sigma} = \mathbf{\Lambda}^{-1} = \begin{bmatrix} 0.51724 & 0 \\ 0 & 0.25 \end{bmatrix}. \tag{9.17}$$

The start record adds twice the prior's precision to z_1, and the run record 1.2 times the prior's to z_2. The entries are read off with $\sqrt{\Sigma_{11}} = 0.719$ in the logit and $\sqrt{\Sigma_{22}} = 1/\sqrt{a_\lambda} = 0.5$ in the log rate, and the correlation of the two is 0.

The exact posteriors. Equations (9.13) and (9.14) give

$$p \mid k_\mathrm{s}, n \sim \mathrm{Beta}(1 + 1,\ 19 + 39) = \mathrm{Beta}(2,\ 58),$$
$$\lambda \mid k_\mathrm{r}, T_\mathrm{tot} \sim \mathrm{Gamma}(2 + 2,\ 400 + 480\,\mathrm{h}) = \mathrm{Gamma}(4,\ 880\,\mathrm{h}).$$

Table 9.3 sets the two side by side: each row is one part of the generator's function; the columns give the prior, the record, the posterior, its mean, standard deviation and central 95 % interval. The probability of a successful start has mean $\mathbb{E}[1 - p \mid \mathbf{y}] = 58/60 = 0.9667$, and the run rate a coefficient of variation of $1/\sqrt{4} = 0.5$.

Table 9.3: The two records of the emergency generator and their posteriors. The start record updates a Beta prior on the probability per demand, the run record a gamma prior on the rate per running hour. The last three columns carry the unit of their row: [−−] for the start row, [h^{-1}] for the run row.

Part	Prior	Record	Posterior	Mean [row unit]	Std. dev. [row unit]	95 % interval [row unit]
start, p [–]	Beta(1, 19)	$k_\mathrm{s} = 1$ in $n = 40$	Beta(2, 58)	0.0333	0.0230	[0.0041, 0.0909]
run, λ [h^{-1}]	Gamma(2, 400 h)	$k_\mathrm{r} = 2$ in $T_\mathrm{tot} = 480\,\mathrm{h}$	Gamma(4, 880 h)	0.00455	0.00227	[0.00124, 0.00996]

The pair $(\hat{\mathbf{z}}, \mathbf{\Sigma})$ of equations (9.16) and (9.17) is what step 4 uses: every answer about the generator is computed from it, one entry at a time when the question concerns the run alone. Because the two parts also admit exact posteriors, each answer has an exact value beside it, a function of the four numbers $(a_p, b_p) = (2,\ 58)$ and $(a_\lambda, b_\lambda) = (4,\ 880\,\mathrm{h})$, and the gap between the two is reported with it.

9.3 The answers (step 4)

The six questions of Table 9.1 take four general forms of Chapter 2: the predictive, the estimation form with its moments and its credible interval, the threshold with its assurance level, and the design form. The fourth question reads the predictive and the threshold as functions of the mission length and inverts them. Each question is answered from the posterior of Section 9.2: its engineering question, the general form of its category, the form it takes for the generator, the quantities of the posterior it uses, and the generator's numbers. The questions name the mission of $t = 72\,\mathrm{h}$, the requirement $R_{\mathrm{req}} = 0.90$ at the assurance $\kappa = 0.95$, and the reliabilities 0.95 and 0.90 of the fourth question; the check by draws uses 400 000 pairs from the two posteriors. Every answer is computed from the posterior of the state of step 3, the MAP state $\hat{\mathbf{z}}$ and the covariance $\boldsymbol{\Sigma}$, and that is the value the chapter reports. The exact Beta and gamma posteriors, in closed form or by a one-dimensional quadrature, give the value set beside it, with the gap.

9.3.1 The run through the outage

Engineering question. Does the generator, once started, run through a 72 h outage? The question asks for the probability of no failure to run in the next t hours, with $t = 72\,\mathrm{h}$ the outage of the design basis, given the two records: the number an operator multiplies by the consequence of a lost outage.

General form. The predictive of Chapter 2, equation (2.13), for an event A, with y_{new} the indicator of the event:

$$\Pr(A \mid \mathbf{y}) = \int \Pr(A \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}. \tag{9.18}$$

Specialized form. The question is equation (9.18) with A the event of no failure to run in t hours, $\Pr(A \mid \mathbf{z}) = R(t;\lambda) = \exp(-t\,\mathrm{e}^{z_2})$ and $p(\mathbf{z} \mid \mathbf{y})$ the posterior of step 3, the Laplace normal (9.12). The event depends on the second entry only, so the integral over the state is a one-dimensional expectation over $z_2 \mid \mathbf{y} \approx \mathcal{N}(\hat{z}_2, \Sigma_{22})$. With φ the standard normal density, $\varphi(\cdot) = (2\pi)^{-1/2}\exp(-(\cdot)^2/2)$, it is evaluated by Gauss–Hermite quadrature,

$$\Pr(T > t \mid \mathbf{y}) \approx R^{\mathrm{L}}_{\mathrm{pred}}(t) = \int \exp\!\big(-t\,\mathrm{e}^{z_2}\big)\, \frac{1}{\sqrt{\Sigma_{22}}}\varphi\!\left(\frac{z_2 - \hat{z}_2}{\sqrt{\Sigma_{22}}}\right) \mathrm{d}z_2 \approx \sum_j \varpi_j \exp\!\big(-t\,\mathrm{e}^{\hat{z}_2 + \sqrt{\Sigma_{22}}\,x_j}\big), \tag{9.19}$$

with nodes x_j and quadrature weights ϖ_j of the rule for the standard normal, and the superscript marking the value on the state route, the predictive the chapter reports.

The exact value beside it. The value beside it comes from the exact posterior: the integral (9.18) over the state is the integral (9.1) over the posterior of the rate, and Chapter 6 evaluated it for a gamma posterior: with $\mathrm{Gamma}(a_\lambda, b_\lambda)$ in place of $\mathrm{Gamma}(a_n, b_n)$, equation (6.59) gives

$$R_{\mathrm{pred}}(t) = \left(\frac{b_\lambda}{b_\lambda + t}\right)^{a_\lambda}, \tag{9.20}$$

the survival function of the Lomax distribution of the time to the next failure, in which the mission adds t hours to the b_λ hours of evidence and prior. Its predictive hazard, equation (6.65), is $a_\lambda/(b_\lambda + t)$. It falls with t although every rate the posterior admits has a constant hazard, by the mechanism of Section 6.3.2: t failure-free hours are one more record, which turns the posterior into $\mathrm{Gamma}(a_\lambda, b_\lambda + t)$. The fall is the decreasing failure rate of a mixture of exponential lifetimes, called frailty when the mixture is over a population of units; here it is over one unit's unknown rate, a statement about knowledge that vanishes as the posterior concentrates.

The plug-in and the gap. The plug-in of equation (9.2) evaluates the survival at the posterior mean rate $m = a_\lambda / b_\lambda$, and the sign of the gap derived here belongs to that estimate only; a plug-in at another estimate, such as the first reading at the maximum-likelihood rate $k_\mathrm{r}/T_\mathrm{tot}$, can lie on either side of the predictive.

The sign of the gap. The survival function $\mathrm{e}^{-\lambda t}$ is convex in λ for every $t > 0$, and Jensen's inequality for a convex function turns the inequality (6.60) of Chapter 6 into a direction: $R_\mathrm{pred}(t) \geq R_\text{plug-in}(t)$ for any posterior, with equality only when the posterior is a single point. The small rates the posterior admits raise the survival by more than the large rates lower it.

The size of the gap for any posterior. The size follows from a second-order expansion, written for a parameter vector $\boldsymbol{\theta}$ with posterior mean $\mathbf{m}$ and posterior covariance $\boldsymbol{\Sigma}_\theta$, so that it serves the run here and the start and run together in Section 9.3.5. With $\boldsymbol{\delta} = \boldsymbol{\theta} - \mathbf{m}$, the gradient $\mathbf{g}_\theta = \partial R / \partial \boldsymbol{\theta}$ and the Hessian $\mathbf{H} = \partial^2 R / \partial \boldsymbol{\theta}\, \partial \boldsymbol{\theta}^\mathsf{T}$ at $\mathbf{m}$,

$$R(\boldsymbol{\theta}) \approx R(\mathbf{m}) + \mathbf{g}_\theta^\mathsf{T} \boldsymbol{\delta} + \tfrac{1}{2}\, \boldsymbol{\delta}^\mathsf{T} \mathbf{H}\, \boldsymbol{\delta}.$$

The posterior expectation of $\boldsymbol{\delta}$ is zero, which removes the linear term, and the quadratic term is a trace, $\mathbb{E}[\boldsymbol{\delta}^\mathsf{T} \mathbf{H} \boldsymbol{\delta}] = \mathbb{E}[\mathrm{tr}(\mathbf{H} \boldsymbol{\delta} \boldsymbol{\delta}^\mathsf{T})] = \mathrm{tr}(\mathbf{H}\, \mathbb{E}[\boldsymbol{\delta} \boldsymbol{\delta}^\mathsf{T}]) = \mathrm{tr}(\mathbf{H} \boldsymbol{\Sigma}_\theta)$, so that

$$\mathbb{E}\big[R(\boldsymbol{\theta}) \mid \mathbf{y}\big] - R(\mathbf{m}) \approx \tfrac{1}{2}\, \mathrm{tr}\big(\mathbf{H} \boldsymbol{\Sigma}_\theta\big). \tag{9.21}$$

The next term, cubic in $\boldsymbol{\delta}$, carries the third central moments of the posterior, which vanish for a normal one; the error of equation (9.21) under a normal posterior is therefore of fourth order. For the run the vector has one entry, $\boldsymbol{\theta} = (\lambda)$, with $\mathbf{m} = (m)$, $\mathbf{H} = (t^2 \mathrm{e}^{-mt})$ and $\boldsymbol{\Sigma}_\theta = (\mathrm{Var}(\lambda \mid \mathbf{y}))$, and the trace is the leading term of the gap,

$$R_\mathrm{pred}(t) - R_\text{plug-in}(t) \approx \tfrac{1}{2}\, t^2\, \mathrm{Var}(\lambda \mid \mathbf{y})\, \mathrm{e}^{-mt}, \tag{9.22}$$

proportional to the posterior variance of the rate and, through $t^2 \mathrm{e}^{-mt}$, to the square of the mission length only while mt is small.

The coordinate decides the sign. The sign of $\mathbf{H}$ sets the sign of the gap, and it depends on the coordinate in which the expansion is taken. In the rate, $\partial^2 \mathrm{e}^{-\lambda t} / \partial \lambda^2 = t^2 \mathrm{e}^{-\lambda t} > 0$ at every rate: the reliability is convex in λ, which is Jensen's direction. In the log rate $z_2 = \ln \lambda$ of the state, the same function has

$$\frac{\partial^2 R}{\partial z_2^2} = R\, \lambda t\, (\lambda t - 1), \tag{9.23}$$

negative while $\lambda t < 1$: over a mission short enough that the expected number of failures is below one, the reliability is concave in the log rate, and an expansion about the log rate under a normal on $\ln \lambda$ moves the mean the other way. That is the state route: at $\lambda t < 1$ the sum (9.19) lies below the value at $\hat{z}_2$, the plug-in at the posterior mean, and Section 9.3.5 meets the same case for the start and the run together.

A bound under the gamma posterior. For the gamma posterior the gap has a closed form and a bound. The log of the ratio of the two reliabilities is, with $x = t / b_\lambda$,

$$\ln \frac{R_\mathrm{pred}(t)}{R_\text{plug-in}(t)} = -a_\lambda \ln\Big(1 + \frac{t}{b_\lambda}\Big) + \frac{a_\lambda t}{b_\lambda} = a_\lambda\, h(x), \qquad h(x) = x - \ln(1 + x).$$

The function h has $h(0) = 0$ and derivative $h'(x) = 1 - 1/(1 + x) = x/(1 + x)$, which lies between 0 and x for $x \geq 0$. Integrating from 0 to x, $0 \leq h(x) \leq x^2/2$, and since $a_\lambda x^2 / 2 = \frac{1}{2} t^2 a_\lambda / b_\lambda^2 = \frac{1}{2} t^2\, \mathrm{Var}(\lambda \mid \mathbf{y})$,

$$R_\text{plug-in}(t) \leq R_\mathrm{pred}(t) \leq R_\text{plug-in}(t)\, \exp\Big[\tfrac{1}{2}\, t^2\, \mathrm{Var}(\lambda \mid \mathbf{y})\Big]. \tag{9.24}$$

The gap is therefore at most $R_\text{plug-in}(t)\big(\mathrm{e}^{t^2\, \mathrm{Var}(\lambda|\mathbf{y})/2} - 1\big)$, whose first-order expansion is the leading term (9.22); because $h' \geq 0$, the ratio $R_\mathrm{pred} / R_\text{plug-in}$ increases with the mission length without limit.

Building blocks. The entry $\hat{z}_2 = -5.3936$ and the variance $\Sigma_{22} = 0.25$ of equations (9.16) and (9.17) (Section 9.2.3); the five Gauss–Hermite nodes of Table 9.4, and forty for the reported values; the mission $t = 72\,\mathrm{h}$; and, for the exact values beside them, the posterior parameters $a_\lambda = 4$ and $b_\lambda = 880\,\mathrm{h}$ (Section 9.2.3), with the posterior mean $m = a_\lambda/b_\lambda$ and variance a_λ/b_λ^2 of the rate.

The rule of five nodes. Table 9.4 evaluates the sum (9.19) at $t = 72\,\mathrm{h}$ with five nodes. Each row is a node: the first two columns give the node x_j and its weight ϖ_j, the third the log rate $z_{2,j} = -5.3936 + 0.5\,x_j$ at which the sum looks, the fourth the rate λ_j there, the fifth the expected number of failures $\lambda_j t$ in the mission and the last the reliability $\mathrm{e}^{-\lambda_j t}$ the question averages. The middle node is the MAP state itself; the others reach about three standard deviations either side of it, rates from 1.09 to $19.0 \times 10^{-3}\,\mathrm{h}^{-1}$.

Table 9.4: The Gauss–Hermite rule (9.19) with five nodes over the Laplace posterior $z_2 \mid \mathbf{y} \approx \mathcal{N}(-5.3936,\ 0.5^2)$ of the run; $\lambda_j = \mathrm{e}^{z_{2,j}}$ in $10^{-3}\,\mathrm{h}^{-1}$, $t = 72\,\mathrm{h}$.

Node x_j	Weight ϖ_j	$z_{2,j}$	λ_j	$\lambda_j t$	$\mathrm{e}^{-\lambda_j t}$
−2.8570	0.01126	−6.8221	1.0894	0.0784	0.92456
−1.3556	0.22208	−6.0714	2.3078	0.1662	0.84691
0	0.53333	−5.3936	4.5455	0.3273	0.72089
1.3556	0.22208	−4.7158	8.9526	0.6446	0.52488
2.8570	0.01126	−3.9651	18.9653	1.3655	0.25525
Weighted sum			5.1507	0.3708	0.70240

The generator's run through the outage. The sum over the last column of Table 9.4 is

$$R^{\mathrm{L}}_{\mathrm{pred}}(72\,\mathrm{h}) \approx 0.01126 \times 0.92456 + 0.22208 \times 0.84691 + 0.53333 \times 0.72089 + 0.22208 \times 0.52488 + 0.01126 \times 0.25525 = 0.7024,$$

a probability of 0.702 of running through the outage, given that the generator has started; ten or forty nodes change it by less than 10^{-6}. Each node is a rate the posterior admits, and the sum weighs the survival at each. The middle node alone is the plug-in (9.2) at the posterior mean rate, $\mathrm{e}^{-0.004545\times 72} = \mathrm{e}^{-0.32727} = 0.7209$, and the sum lies 0.0185 below it, because at $\hat{\lambda}t = 0.327 < 1$ the reliability is concave in the log rate, equation (9.23).

The value beside it comes from the exact posterior. By equation (9.20) with $a_\lambda = 4$ and $b_\lambda = 880\,\mathrm{h}$, $R_{\mathrm{pred}}(72\,\mathrm{h}) = (880/952)^4 = 0.92437^4 = 0.7301$, which is 0.0277 above the state route: the normal in the log rate puts more weight on large rates, its mean rate $0.00515\,\mathrm{h}^{-1}$ against the exact $0.004545\,\mathrm{h}^{-1}$. The exact value lies above the plug-in, by 0.0092, in the direction of Jensen's inequality in the rate. The plug-in at the maximum-likelihood rate 2/480, the first reading of the record, is $\mathrm{e}^{-72\times 2/480} = 0.7408$, and at the exact posterior median rate 0.7405: both lie above both predictives. The predictive hazard is $4/880 = 0.004545\,\mathrm{h}^{-1}$ at the start of the run and $4/952 = 0.004202\,\mathrm{h}^{-1}$ after 72 h. At 72 h, $\frac{1}{2}t^2\,\mathrm{Var}(\lambda \mid \mathbf{y}) = \frac{1}{2} \times 5184 \times 4/880^2 = 0.01339$, and the bound (9.24) gives a gap of at most $0.7209 \times (\mathrm{e}^{0.01339} - 1) = 0.0097$, against the exact 0.0092.

Table 9.5 evaluates the predictives and the plug-in for six mission lengths, from an eight-hour shift to a month. Each row is a mission length; the columns give the Laplace predictive (9.19) with forty nodes, the exact predictive (9.20), the plug-in (9.2), the exact predictive's difference from the plug-in and its ratio to it, the upper bound on that difference from equation (9.24) and the exact predictive hazard. One row of the arithmetic, $t = 168\,\mathrm{h}$: $(880/1048)^4 = 0.83969^4 = 0.4972$; $\mathrm{e}^{-0.004545\times 168} = \mathrm{e}^{-0.76364} = 0.4660$; the gap is 0.0312 and the bound $0.4660 \times (\mathrm{e}^{0.07289} - 1) = 0.0352$.

Table 9.5: Predictive reliability of the generator's run and its plug-in at the posterior mean rate: the Laplace predictive, equation (9.19) with forty nodes; the exact $R_{\text{pred}}(t) = (880/(880+t))^4$; and $R_{\text{plug-in}}(t) = \mathrm{e}^{-4t/880}$; gap and ratio of the exact predictive to the plug-in; bound from equation (9.24); hazard $4/(880+t)$ in $10^{-3}\,\mathrm{h}^{-1}$.

Mission t [h]	Laplace [–]	Exact [–]	Plug-in [–]	Gap [–]	Ratio [–]	Bound on gap [–]	Hazard $[10^{-3}\,\mathrm{h}^{-1}]$
8	0.95986	0.96445	0.96429	0.00016	1.0002	0.00016	4.505
24	0.88557	0.89796	0.89665	0.00131	1.0015	0.00133	4.425
72	0.70240	0.73010	0.72089	0.00921	1.0128	0.00972	4.202
168	0.45841	0.49715	0.46597	0.03118	1.0669	0.03523	3.817
336	0.23685	0.27428	0.21713	0.05715	1.2632	0.07350	3.289
720	0.06789	0.09151	0.03790	0.05360	2.4142	0.10668	2.500

Figure 9.2 draws the same comparison as curves. Panel (a) has the mission length on the horizontal axis, from zero to 720 h, and the probability of no failure on the vertical axis; the dotted blue curve is the Laplace predictive, the green curve the exact predictive and the dashed gray curve the plug-in; the dotted vertical line marks the 72 h outage. Panel (b) has the same horizontal axis and the gap on the vertical axis; the green curve is the exact gap, the dashed blue curve the bound of equation (9.24), and the dots mark the mission lengths 8, 72, 168 and 720 h.

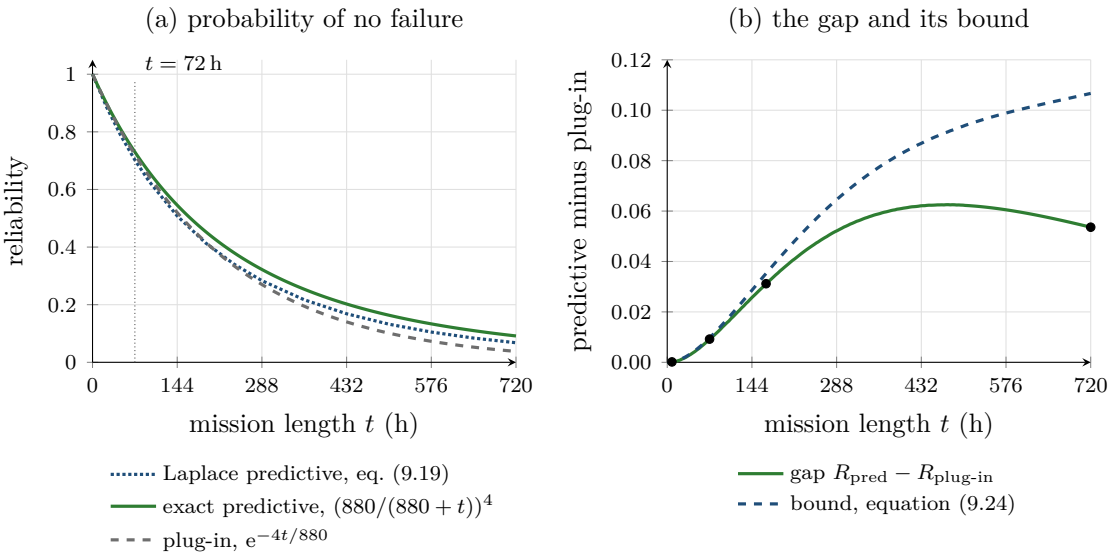


Figure 9.2: The predictive survival of the generator's run against the plug-in. (a) The Laplace predictive (9.19), $R_{\text{pred}}(t) = (880/(880+t))^4$, equation (9.20), and $R_{\text{plug-in}}(t) = \mathrm{e}^{-4t/880}$; at 72 h they are 0.7024, 0.7301 and 0.7209. (b) Their difference and the bound $\mathrm{e}^{-4t/880}(\mathrm{e}^{t^2/387\,200} - 1)$, with $387\,200 = 2b_\lambda^2/a_\lambda$; the dots are the gaps 0.0002, 0.009, 0.031 and 0.054 at 8, 72, 168 and 720 h. The gap is largest, 0.062, at 474 h.

The exact predictive lies above the plug-in at the posterior mean at every mission length; the Laplace predictive lies below it up to about nine days, near $\hat{\lambda}t = 1$, and above it beyond. For a shift to a few days the exact gap is small, 0.0002 at 8 h and 0.009 at 72 h: tripling the mission from 24 to 72 h multiplies it by seven, against the factor $9\,\mathrm{e}^{-48m} = 7.2$ of the leading term (9.22). It reaches 0.031 at a week and its largest value, 0.062, at 474 h, where the plug-in has fallen to 0.116; beyond, both reliabilities approach zero and the gap shrinks while the ratio keeps growing, to 2.4 at 720 h. The bound is tight for short missions and loose for long ones.

9.3.2 The posterior of the run reliability

Engineering question. How well is the generator's 72 h reliability known? The question asks not for the probability of the next outage but for the reliability $R(72;\lambda)$ itself, a property of this generator with one true value, and for how far the records leave it uncertain.

General form. The estimation form of Chapter 2, equation (2.9),

$$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y}) = \int \mathbf{1}\big[f(\mathbf{z}) \le f^*\big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \tag{9.25}$$

summarized by its moments (2.10), its median and its central 95 % credible interval (2.11).

Specialized form. The question is equation (9.25) with $f(\mathbf{z}) = R(t;\lambda) = \exp(-t\,\mathrm{e}^{z_2})$, a strictly decreasing function of the second entry alone, and $p(\mathbf{z} \mid \mathbf{y})$ the posterior of step 3, the Laplace normal (9.12); its moments (2.10) and its interval (2.11) are those of the same posterior.

The state route. The reliability is not linear in the state, but its complementary log-log is:

$$c(\mathbf{z}) = \ln\big(-\ln R(t;\lambda)\big) = \ln(\lambda t) = \mathbf{g}_c^{\mathsf{T}}\mathbf{z} + \ln t, \qquad \mathbf{g}_c = \frac{\partial c}{\partial \mathbf{z}} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}. \tag{9.26}$$

This is the scale that plays the part of the logit in Chapter 8: on it the derived quantity is linear in the state, with a constant gradient, so the delta method is exact for the Gaussian and the only approximation is the Gaussian itself. Under the Laplace posterior, $c \mid \mathbf{y} \approx \mathcal{N}(\hat{c},\ \mathbf{g}_c^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_c)$ with the center $\hat{c} = \hat{z}_2 + \ln t$ and the variance $\mathbf{g}_c^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_c = \Sigma_{22}$: the zero first entry of $\mathbf{g}_c$ removes the start entry from the run reliability. The band maps $\hat{c} \pm 1.96\sqrt{\Sigma_{22}}$ through $R = \exp(-\mathrm{e}^{c})$ in reverse order, and its median is the reliability at $\hat{z}_2$, which is the plug-in at the posterior mean because $\hat{z}_2$ is the log of the posterior mean rate (Section 9.2.1). Its density at a level R_0 is the normal density of z_2 at $\ln(-\ln R_0/t)$ times the factor of the change of variable,

$$p_R(R_0 \mid \mathbf{y}) \approx \frac{1}{\sqrt{\Sigma_{22}}\, R_0\, |\ln R_0|}\, \varphi\!\left(\frac{\ln(-\ln R_0/t) - \hat{z}_2}{\sqrt{\Sigma_{22}}}\right). \tag{9.27}$$

Its mean is the sum (9.19), and since $R(t;\lambda)^2 = \mathrm{e}^{-2\lambda t}$ its second moment is the same sum at twice the mission,

$$\mathbb{E}\big[R(t;\lambda) \mid \mathbf{y}\big] \approx R^{\mathrm{L}}_{\mathrm{pred}}(t), \qquad \mathrm{Var}\big(R(t;\lambda) \mid \mathbf{y}\big) \approx R^{\mathrm{L}}_{\mathrm{pred}}(2t) - R^{\mathrm{L}}_{\mathrm{pred}}(t)^2. \tag{9.28}$$

The exact values beside them. Exactly, the posterior of the reliability follows from the rate's without approximation.

The quantiles. Because the map is decreasing, the q quantile of the reliability is the image of the $1-q$ quantile of the rate,

$$R_q(t) = \exp(-t\,\lambda_{1-q}), \qquad \lambda_{1-q} = F_\Gamma^{-1}(1-q; a_\lambda, b_\lambda), \tag{9.29}$$

with $F_\Gamma^{-1}(q;a,b)$ the q quantile of Gamma(a,b), the notation used for every gamma quantile of this chapter, so that the central 95 % band of $R(t;\lambda)$ is $[\mathrm{e}^{-t\lambda_{0.975}},\ \mathrm{e}^{-t\lambda_{0.025}}]$ and its median is $\mathrm{e}^{-t\lambda_{0.5}}$.

The density. With $\lambda(R_0) = -\ln R_0/t$ and $|\mathrm{d}\lambda/\mathrm{d}R_0| = 1/(tR_0)$, the density of the reliability at a level R_0 is the gamma density at $\lambda(R_0)$ times that change-of-variable factor,

$$p_R(R_0 \mid \mathbf{y}) = \frac{b_\lambda^{a_\lambda}}{\Gamma(a_\lambda)} \left(\frac{-\ln R_0}{t}\right)^{a_\lambda - 1} R_0^{b_\lambda/t}\, \frac{1}{t\,R_0}, \qquad 0 < R_0 < 1, \tag{9.30}$$

since $\mathrm{e}^{-b_\lambda \lambda(R_0)} = R_0^{b_\lambda/t}$.

The moments. The mean of $R(t;\lambda)$ is the predictive (9.20). Its second moment is the predictive at twice the mission, since $R(t;\lambda)^2 = \mathrm{e}^{-2\lambda t}$, and

$$\mathbb{E}\big[R(t;\lambda) \mid \mathbf{y}\big] = R_{\mathrm{pred}}(t), \qquad \mathrm{Var}\big(R(t;\lambda) \mid \mathbf{y}\big) = R_{\mathrm{pred}}(2t) - R_{\mathrm{pred}}(t)^2. \tag{9.31}$$

Building blocks. The center $\hat{c} = \hat{z}_2 + \ln t$ and the variance $\Sigma_{22} = 0.25$ of step 3 (Section 9.2.3); the sum (9.19) at t and $2t$; and, for the exact values beside them, the gamma quantile function at the posterior parameters a_λ and b_λ (Section 9.2.2) and the predictive (9.20) at t and $2t$.

The band of the run reliability. The band of the rate is $\lambda \in 0.004545\,\mathrm{e}^{\pm 0.98}$, that is $[0.001706,\ 0.012111]\,\mathrm{h}^{-1}$, from $\hat{c} \pm 1.96 \times 0.5$, and through $R = \exp(-\mathrm{e}^c)$ in reverse order,

$$R(72;\lambda) \in [\mathrm{e}^{-72\times 0.012111},\ \mathrm{e}^{-72\times 0.001706}] = [0.418,\ 0.884] \text{ with } 95\,\% \text{ posterior probability},$$

with median 0.721, the plug-in value, mean $R^{\mathrm{L}}_{\mathrm{pred}}(72) = 0.702$ and, by equation (9.28), standard deviation $\sqrt{0.50799 - 0.70240^2} = 0.121$. The density (9.27) has its mode near 0.76.

The values beside them come from the exact posterior. The quantiles of Gamma$(4, 880\,\mathrm{h})$ are those of Gamma$(4, 1)$ divided by 880: 1.0899, 3.6721 and 8.7673 at 0.025, 0.5 and 0.975, so that $\lambda_{0.025} = 0.001238$, $\lambda_{0.5} = 0.004173$ and $\lambda_{0.975} = 0.009963\,\mathrm{h}^{-1}$. By equation (9.29) with $t = 72\,\mathrm{h}$,

$$R(72;\lambda) \in [\mathrm{e}^{-72\times 0.009963},\ \mathrm{e}^{-72\times 0.001238}] = [0.488,\ 0.915] \text{ with } 95\,\% \text{ posterior probability}, \tag{9.32}$$

with median $\mathrm{e}^{-72\times 0.004173} = 0.740$, mean $R_{\mathrm{pred}}(72) = 0.730$ and, by equation (9.31), standard deviation $\sqrt{0.54542 - 0.73010^2} = 0.111$. The density (9.30) has its mode near 0.77. The state route's band lies below the exact one by 0.070 at its lower end and 0.031 at its upper end, its median by 0.019 and its mean by 0.028: the same weight on large rates.

9.3.3 The requirement on the run reliability

Engineering question. Can $R(72) \geq 0.90$ be asserted with 95 % assurance? The design basis asks that the generator's reliability over the 72 h run be at least $R_{\mathrm{req}} = 0.90$, shown with the posterior probability $\kappa = 0.95$, the assurance level of Chapter 5.

General form. The threshold of Chapter 2, equation (2.12),

$$\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}) = \int \mathbf{1}\big[f(\mathbf{z}) \geq f^*\big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \qquad \text{accept if } \Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}) \geq \kappa. \tag{9.33}$$

Specialized form. The question is equation (9.33) with $f(\mathbf{z}) = R(t;\lambda)$, $f^* = R_{\mathrm{req}}$, $\kappa = 0.95$ and $p(\mathbf{z} \mid \mathbf{y})$ the posterior of step 3, the Laplace normal (9.12). One choice enters here: the requirement is read on the generator's reliability over the 72 h run, a property of the unit, as the design basis states it. That makes the answer a tail of the posterior of $R(72;\lambda)$, not a statement about the predictive, which a broad posterior can raise above 0.90 while the tail stays small (Section 9.5.2).

Since the reliability decreases with the rate, the event becomes a bound on the rate and on the second entry of the state,

$$R(t;\lambda) \geq R_{\mathrm{req}} \iff \lambda \leq \lambda^* = \frac{-\ln R_{\mathrm{req}}}{t} \iff z_2 \leq \ln \lambda^* \iff c(\mathbf{z}) \leq \ln(\lambda^* t),$$

a half-line in z_2 whatever the value of z_1: a requirement on the reliability over a mission is a requirement on the rate. On the state the tail is the normal tail of the complementary log-log of equation (9.26),

$$\Pr\big(R(t;\lambda) \geq R_{\mathrm{req}} \mid \mathbf{y}\big) \approx \Phi\left(\frac{\ln(\lambda^* t) - \hat{c}}{\sqrt{\mathbf{g}_c^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g}_c}}\right), \tag{9.34}$$

with Φ the standard normal cumulative distribution; only the tail is an approximation, since the threshold maps exactly. The same formula with the Laplace normal of the prior rows,

$\mathcal{N}(\ln(a_{\lambda0}/b_{\lambda0}),\ 1/a_{\lambda0})$, gives the support under the prior alone, against which the contribution of the run record is read.

The value beside it comes from the exact posterior, where the support is a gamma cumulative,

$$\Pr\bigl(R(t;\lambda) \geq R_{\text{req}} \mid \mathbf{y}\bigr) = \Pr\bigl(\lambda \leq \lambda^* \mid \mathbf{y}\bigr) = F_\Gamma(\lambda^*; a_\lambda, b_\lambda), \qquad \lambda^* = \frac{-\ln R_{\text{req}}}{t}, \tag{9.35}$$

with F_Γ the gamma cumulative distribution of Chapter 6, equation (6.45), which for an integer shape takes the Poisson-tail form of equation (6.46), a finite sum of a_λ terms in $b_\lambda\lambda^*$. The same formula with the prior parameters $a_{\lambda0}$ and $b_{\lambda0}$ gives its support under the prior alone. A record that points to a rate well above λ^* lowers the support below the prior's, although it adds information: it narrows the posterior around its own rate and so removes the probability near λ^* that a broad prior had allowed.

Building blocks. The threshold rate λ^*; the center $\hat{c}$ and the variance Σ_{22} of step 3, with Φ; and, for the exact values beside them, the posterior and prior parameters of the rate (Section 9.2.2) and the gamma cumulative distribution in its Poisson-tail form.

The requirement of 0.90. The requirement is $\lambda \leq \lambda^* = -\ln 0.90/72 = 0.105361/72 = 0.0014633\,\text{h}^{-1}$, a mean time to failure in running of at least 683 h. From $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$ of equations (9.16) and (9.17), the center of the complementary log-log (9.26) at $t = 72$ h is $\hat{c} = -5.3936 + 4.2767 = -1.1170$, with $\ln t = 4.2767$, and its variance $\mathbf{g}_c^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_c = \Sigma_{22} = 0.25$. With $\ln\lambda^* = \ln 0.0014633 = -6.5270$ and $\ln(\lambda^* t) = -6.5270 + 4.2767 = -2.2504$, the tail (9.34) is

$$\Pr\bigl(R(72;\lambda) \geq 0.90 \mid \mathbf{y}\bigr) \approx \Phi\biggl(\frac{-2.2504 + 1.1170}{0.5}\biggr) = \Phi(-2.267) = 0.012,$$

far below the required assurance of 0.95: the requirement cannot be asserted. Under the prior alone, with $\ln(2/400) = -5.2983$ and the variance $\frac{1}{2}$, the same probability was $\Phi(-1.738) = 0.041$; the record, whose observed rate $2/480 = 0.0042\,\text{h}^{-1}$ is nearly three times λ^*, has lowered it. At other levels the support is 0.222 for $R(72) \geq 0.80$ and 0.0001 for $R(72) \geq 0.95$.

The value beside it comes from the exact posterior. By equation (9.35), with $b_\lambda\lambda^* = 880 \times 0.0014633 = 1.2877$ and the integer shape 4, the Poisson-tail form of the gamma cumulative gives

$$\begin{aligned}\Pr\bigl(R(72;\lambda) \geq 0.90 \mid \mathbf{y}\bigr) &= 1 - \mathrm{e}^{-1.2877}\left(1 + 1.2877 + \frac{1.2877^2}{2} + \frac{1.2877^3}{6}\right) \\ &= 1 - 0.27590 \times 3.4728 = 0.042,\end{aligned} \tag{9.36}$$

which is 0.030 above the state route, a ratio of three and a half; both are far below 0.95 and the decision is the same on both routes. The state route errs low because its normal in the log rate is centered above the exact mean of $\ln\lambda$ and puts too little weight on small rates, where the requirement lives. Under the prior alone the exact probability was $F_\Gamma(0.0014633; 2, 400) = 0.117$, and at other levels it is 0.292 for $R(72) \geq 0.80$ and 0.0039 for $R(72) \geq 0.95$.

9.3.4 The outage that can be covered

Engineering question. What outage length can be covered at a reliability of 0.95, and of 0.90? The question has two readings, one for each of the two objects: the outage over which the probability of getting through equals the required value, and the outage over which the generator's reliability reaches the required value with 95 % assurance.

General form. The predictive of Chapter 2, equation (2.13), and the threshold, equation (2.12), in their forms here (9.18) and (9.33), each read as a function of the mission length t and inverted:

$$t^* : \Pr(T > t^* \mid \mathbf{y}) = R_{\text{req}}, \qquad t_\kappa = \max\{\, t : \Pr(R(t;\lambda) \geq R_{\text{req}} \mid \mathbf{y}) \geq \kappa \,\}. \tag{9.37}$$

Specialized form. The question is equation (9.37) with $\Pr(T > t \mid \mathbf{y})$ the sum (9.19), the tail the normal tail (9.34), and $p(\mathbf{z} \mid \mathbf{y})$ the posterior of step 3, the Laplace normal (9.12). The sum has no closed inverse, and t^* is found by inverting it numerically. The tail reaches κ at a mission where the threshold rate $-\ln R_{\text{req}}/t$ equals the rate $\mathrm{e}^{\hat{z}_2+\Phi^{-1}(\kappa)\sqrt{\Sigma_{22}}}$, so

$$t_\kappa = \frac{-\ln R_{\text{req}}}{\mathrm{e}^{\hat{z}_2+\Phi^{-1}(\kappa)\sqrt{\Sigma_{22}}}}, \tag{9.38}$$

the one-sided bound on the rate in place of its center. The two inverse questions differ as the first and the third do: one averages over the uncertainty, the other guards against it.

The values beside them come from the exact posterior. With the predictive (9.20),

$$\left(\frac{b_\lambda}{b_\lambda + t^*}\right)^{a_\lambda} = R_{\text{req}} \quad \Longleftrightarrow \quad \frac{b_\lambda + t^*}{b_\lambda} = R_{\text{req}}^{-1/a_\lambda},$$

and therefore

$$t^* = b_\lambda (R_{\text{req}}^{-1/a_\lambda} - 1). \tag{9.39}$$

The inversion of the plug-in at the posterior mean solves $\mathrm{e}^{-mt} = R_{\text{req}}$, with $m = a_\lambda/b_\lambda$: $t_{\text{plug-in}} = -\ln R_{\text{req}}/m = -b_\lambda \ln R_{\text{req}}/a_\lambda$. Writing $\epsilon = -\ln R_{\text{req}}/a_\lambda$, the ratio of the two is

$$\frac{t^*}{t_{\text{plug-in}}} = \frac{\mathrm{e}^\epsilon - 1}{\epsilon} = 1 + \frac{\epsilon}{2} + \frac{\epsilon^2}{6} + \cdots, \tag{9.40}$$

so the predictive allows a mission longer than the plug-in by a relative amount of about $-\ln R_{\text{req}}/(2a_\lambda)$. For a high required reliability this is small whatever the record: the inversion works where both reliabilities are close to one, and there the gap of Section 9.3.1 is still small.

The exact assured length inverts the tail (9.35). At a mission t the requirement holds with probability at least κ when the threshold rate $-\ln R_{\text{req}}/t$ is at least the κ quantile of the rate, so

$$t_\kappa = \frac{-\ln R_{\text{req}}}{\lambda_\kappa}, \qquad \lambda_\kappa = F_\Gamma^{-1}(\kappa; a_\lambda, b_\lambda), \tag{9.41}$$

with the gamma's one-sided bound on the rate.

Building blocks. The sum (9.19) over the entry $\hat{z}_2$ and the variance Σ_{22}, with Φ^{-1}; and, for the exact values beside them, the posterior parameters a_λ and b_λ and the gamma quantile function (Section 9.2.2).

The outage lengths that can be covered. Inverted numerically with forty nodes, the sum (9.19) reaches 0.95 at $t^* = 10.0\,\text{h}$ and 0.90 at 20.8 h. The generator's predictive reliability falls below 0.90 within a day of running, and the 72 h outage of the design basis is more than three times longer. With the one-sided bound $\mathrm{e}^{-5.3936+1.6449\times0.5} = \mathrm{e}^{-4.5711} = 0.010346\,\text{h}^{-1}$, equation (9.38) gives the assured lengths $0.051293/0.010346 = 5.0\,\text{h}$ for a reliability of 0.95 and $0.105361/0.010346 = 10.2\,\text{h}$ for 0.90, about half the predictive values.

The values beside them come from the exact posterior. By equation (9.39) with $a_\lambda = 4$ and $b_\lambda = 880\,\text{h}$, a predictive reliability of 0.95 covers an outage of

$$t^* = 880\,(0.95^{-1/4} - 1) = 880 \times 0.012905 = 11.4\,\text{h},$$

and a predictive reliability of 0.90 covers $880\,(0.90^{-1/4} - 1) = 880 \times 0.026690 = 23.5\,\text{h}$. The plug-in inversion gives $-\ln 0.95/0.004545 = 11.3\,\text{h}$ and $-\ln 0.90/0.004545 = 23.2\,\text{h}$, shorter by 0.6 % and 1.3 %, the relative amounts $-\ln R_{\text{req}}/(2a_\lambda) = 0.0064$ and 0.0132 of equation (9.40); the state route lies below the plug-in, as its sum lies below the value at $\hat{z}_2$, by 1.3 and 2.4 h, and below the exact t^* by 1.3 and 2.7 h, about 12 %. With the one-sided bound $\lambda_{0.95} = F_\Gamma^{-1}(0.95; 4, 880) = 7.7537/880 = 0.008811\,\text{h}^{-1}$, equation (9.41) gives the assured lengths 5.8 h for a reliability of 0.95 and 12.0 h for 0.90, 0.9 and 1.8 h above the state route.

9.3.5 The start and the run together

Engineering question. Does the generator start and run 72 h? The question asks for the probability of carrying the load through the outage with the start included, for how well the start-and-run reliability of the generator is known, and for which of the two records limits it.

General form. The predictive (9.18) for the probability of carrying the outage, and the estimation form (9.25) with its band for the start-and-run reliability as a derived quantity of both entries, whose spread is the posterior variance of equation (2.10),

$$\text{Var}(f(\mathbf{z}) \mid \mathbf{y}) = \int \big(f(\mathbf{z}) - \mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]\big)^2\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}. \tag{9.42}$$

Specialized form. The question is the predictive (9.18) with A the event of starting and then running t hours, and the estimation form (9.25) with its variance (9.42) with $f(\mathbf{z})$ the probability of that event given the state; in both $p(\mathbf{z} \mid \mathbf{y})$ is the posterior of step 3, the Laplace normal $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of equation (9.12). Two choices enter here. The start and the run are taken to fail independently given p and λ: this makes the mission reliability a product and its logarithm a sum with one term per record, and it fails for a unit whose degraded fuel or battery system makes it both more likely to fail to start and more likely to fail early in the run (Section 9.6.6). The start tests are taken to represent the start after an outage: this lets the forty tests speak for the start demand of the mission, and it fails if the tests start a warm engine and the outage a cold one after a long standby (Section 9.6.6).

The derived quantity. The generator completes the outage if it starts and then runs t hours without a failure to run. Given p and λ the two events are independent, and the probability of both is the product

$$R_{\text{sr}}(t; p, \lambda) = (1 - p)\, \mathrm{e}^{-\lambda t}, \tag{9.43}$$

a derived quantity of both entries of the state.

The variance budget on the state. The logarithm of equation (9.43) separates the two parts,

$$\ln R_{\text{sr}} = \ln(1 - p) - \lambda t, \tag{9.44}$$

and in the state it is $f(\mathbf{z}) = -\ln(1 + \mathrm{e}^{z_1}) - t\,\mathrm{e}^{z_2}$, since $1 - p = 1/(1 + \mathrm{e}^{z_1})$ and $\lambda = \mathrm{e}^{z_2}$. Its gradient at the MAP state is

$$\mathbf{g} = \frac{\partial f}{\partial \mathbf{z}}\bigg|_{\hat{\mathbf{z}}} = \begin{bmatrix} -\hat{p} \\ -\hat{\lambda} t \end{bmatrix}, \tag{9.45}$$

because $\partial \ln(1 + \mathrm{e}^{z_1})/\partial z_1 = p$ and $\partial(t\,\mathrm{e}^{z_2})/\partial z_2 = \lambda t$, with $\hat{p}$ and $\hat{\lambda}$ the probability and the rate at $\hat{\mathbf{z}}$. The delta method of Chapter 4 gives the variance budget

$$\mathbf{g}^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g} = \underbrace{g_1^2 \Sigma_{11}}_{\text{start rows}} + \underbrace{g_2^2 \Sigma_{22}}_{\text{run rows}} + \underbrace{2 g_1 g_2 \Sigma_{12}}_{\text{cross term}}, \tag{9.46}$$

and the cross term is zero. The off-diagonal entry of $\boldsymbol{\Lambda} = \mathbf{J}^{\mathsf{T}}\mathbf{J}$ is $\sum_i J_{i1} J_{i2}$, and in the pattern (9.4) every product in the sum is zero because no row of $\mathbf{J}$ has two nonzero entries; so

$\boldsymbol{\Lambda}$ is diagonal, $\boldsymbol{\Sigma}$ with it, and $\Sigma_{12} = 0$ is a result of the rows, not an assumption about the start and the run. No row informs both entries, and the start and the run add their variances. The size of each term is set by an entry of $\mathbf{g}$ as much as by an entry of $\boldsymbol{\Sigma}$: the start entry is $\hat{p}$, the failure probability itself, small for any start that usually succeeds however uncertain its logit, while the run entry is $\hat{\lambda}t$, the expected number of failures in the mission, which grows with the mission length.

The state route for the predictive and the band. Under the normal $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ the two entries are independent, since $\Sigma_{12} = 0$, and the predictive is the product of two one-dimensional sums over the nodes of equation (9.19),

$$\mathbb{E}\big[R_{\mathrm{sr}}(t; p, \lambda) \mid \mathbf{y}\big] \approx \Big(\sum_j \frac{\varpi_j}{1 + \mathrm{e}^{\hat{z}_1 + \sqrt{\Sigma_{11}}\, x_j}}\Big)\, R^{\mathrm{L}}_{\mathrm{pred}}(t). \tag{9.47}$$

Its band is the same normal pushed through equation (9.43): conditioning on z_1, with $1 - p_j = 1/(1 + \mathrm{e}^{\hat{z}_1 + \sqrt{\Sigma_{11}}\, x_j})$, the tail in z_2 is normal,

$$\Pr\big(R_{\mathrm{sr}} \le R_0 \mid \mathbf{y}\big) \approx \sum_j \varpi_j \left[1 - \Phi\left(\frac{\ln\big(-\ln(R_0/(1 - p_j))/t\big) - \hat{z}_2}{\sqrt{\Sigma_{22}}}\right)\right] \tag{9.48}$$

for a reliability level R_0, with the bracket equal to one when $R_0 \ge 1 - p_j$, and it is inverted for the quantiles. The moments follow from the sums at t and $2t$ with the squares of $1 - p_j$. Two linearizations stand beside these. A normal on $\ln R_{\mathrm{sr}}$, centered at $f(\hat{\mathbf{z}})$ with the variance $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ and mapped back, gives the band the budget alone would give; its shape is symmetric, while the log of equation (9.43) has a long lower tail. The delta method on R_{sr} itself uses the gradient $\mathbf{g}_R = \partial R_{\mathrm{sr}}/\partial \mathbf{z} = R_{\mathrm{sr}}(\hat{\mathbf{z}})\, \mathbf{g}$ and the center $R_{\mathrm{sr}}(\hat{\mathbf{z}})$, which is the plug-in at the posterior means, because $\hat{\mathbf{z}}$ maps to the posterior means of p and λ; it returns the plug-in, not the predictive.

The sign of the shift from the plug-in. The expectation of R_{sr} under $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ follows from the expansion (9.21) with the state in place of $\boldsymbol{\theta}$, $\mathbb{E}_{\mathrm{Laplace}}[R_{\mathrm{sr}}] \approx R_{\mathrm{sr}}(\hat{\mathbf{z}}) + \frac{1}{2}\operatorname{tr}(\mathbf{H}\boldsymbol{\Sigma})$. Since $R_{\mathrm{sr}} = \mathrm{e}^{f}$, differentiating $\partial R_{\mathrm{sr}}/\partial \mathbf{z} = R_{\mathrm{sr}}\, \partial f/\partial \mathbf{z}$ once more gives the Hessian

$$\begin{gathered}\mathbf{H} = \frac{\partial^2 R_{\mathrm{sr}}}{\partial \mathbf{z}\, \partial \mathbf{z}^{\mathsf{T}}}\bigg|_{\hat{\mathbf{z}}} = R_{\mathrm{sr}}(\hat{\mathbf{z}})\Big(\mathbf{g}\mathbf{g}^{\mathsf{T}} + \frac{\partial^2 f}{\partial \mathbf{z}\, \partial \mathbf{z}^{\mathsf{T}}}\Big), \qquad \frac{\partial^2 f}{\partial \mathbf{z}\, \partial \mathbf{z}^{\mathsf{T}}}\bigg|_{\hat{\mathbf{z}}} = \begin{bmatrix} -\hat{p}(1 - \hat{p}) & 0 \\ 0 & -\hat{\lambda}t \end{bmatrix}, \\ \mathbf{H} = R_{\mathrm{sr}}(\hat{\mathbf{z}}) \begin{bmatrix} \hat{p}(2\hat{p} - 1) & \hat{p}\hat{\lambda}t \\ \hat{p}\hat{\lambda}t & \hat{\lambda}t(\hat{\lambda}t - 1) \end{bmatrix}.\end{gathered} \tag{9.49}$$

Its first diagonal entry is negative whenever $\hat{p} < \frac{1}{2}$, and its second while $\hat{\lambda}t < 1$, equation (9.23); the off-diagonal entry meets the zero Σ_{12} and drops out of the trace. Where both diagonal entries are negative the Laplace mean lies below the plug-in, on the other side from the exact mean, which by Jensen's inequality in the rate lies above it.

The exact values beside them. The values beside them come from the exact posterior. Its predictive is the posterior expectation of the product (9.43); the two exact posteriors are independent (Section 9.2.2), so the expectation of the product is the product of the expectations:

$$\mathbb{E}\big[R_{\mathrm{sr}}(t; p, \lambda) \mid \mathbf{y}\big] = \mathbb{E}[1 - p \mid \mathbf{y}]\, \mathbb{E}\big[\mathrm{e}^{-\lambda t} \mid \mathbf{y}\big] = \frac{b_p}{a_p + b_p}\left(\frac{b_\lambda}{b_\lambda + t}\right)^{a_\lambda}. \tag{9.50}$$

The plug-in at the posterior means is $(1 - \mathbb{E}[p \mid \mathbf{y}])\, \mathrm{e}^{-mt}$, and since R_{sr} is linear in p the gap between the two is the run's gap scaled by $\mathbb{E}[1 - p \mid \mathbf{y}]$, with nothing from the start.

The exact budget. The exact posterior factorizes for the same reason, and since $\ln R_{\mathrm{sr}}$ is linear in λ the exact variance of the log is

$$\operatorname{Var}\big(\ln R_{\mathrm{sr}} \mid \mathbf{y}\big) = \underbrace{\operatorname{Var}\big(\ln(1 - p) \mid \mathbf{y}\big)}_{\text{start term}} + \underbrace{t^2\, \operatorname{Var}(\lambda \mid \mathbf{y})}_{\text{run term}}, \tag{9.51}$$

with the start term $\psi'(b_p)-\psi'(a_p+b_p)$, the log-moment formula of a Beta variable (Section 9.2.2), and the run term $t^2 a_\lambda/b_\lambda^2$. The Laplace run term $g_2^2\Sigma_{22}$ of equation (9.46) equals the exact one, and the reason is the location of the mode: $\hat{z}_2$ is the log of the posterior mean rate, so the gradient entry $g_2 = -t\,a_\lambda/b_\lambda$ and the variance entry $\Sigma_{22} = 1/a_\lambda$ (Section 9.2.1) multiply to $t^2 a_\lambda/b_\lambda^2 = t^2\,\mathrm{Var}(\lambda \mid \mathbf{y})$. The delta method maps the Laplace variance of the log rate back to the exact variance of the rate, as Chapter 6 found; at any other center it would not.

The exact band and moments. The distribution of R_{sr} needs no normal. Conditioning on p and using the gamma survival for the run,

$$\Pr(R_{\mathrm{sr}} \le R_0 \mid \mathbf{y}) = \int_0^1 \Pr\!\left(\lambda \ge \frac{-\ln(R_0/(1-p))}{t} \;\Big|\; \mathbf{y}\right) p(p \mid \mathbf{y})\,\mathrm{d}p \tag{9.52}$$

for a reliability level R_0, with the inner probability equal to one when $R_0 \ge 1-p$: a one-dimensional integral against the Beta density, evaluated by quadrature and inverted for the quantiles. Its first two moments have closed forms, the mean (9.50) and, since $R_{\mathrm{sr}}^2 = (1-p)^2\mathrm{e}^{-2\lambda t}$,

$$\mathbb{E}[R_{\mathrm{sr}}^2 \mid \mathbf{y}] = \mathbb{E}[(1-p)^2 \mid \mathbf{y}]\; R_{\mathrm{pred}}(2t), \qquad \mathbb{E}[(1-p)^2 \mid \mathbf{y}] = \frac{b_p(b_p+1)}{(a_p+b_p)(a_p+b_p+1)}, \tag{9.53}$$

from the moments of a Beta variable of Chapter 5, which give the exact standard deviation of R_{sr} through the variance (9.42).

Building blocks. The MAP state $\hat{\mathbf{z}}$ and the covariance $\boldsymbol{\Sigma}$ of equations (9.16) and (9.17) (Section 9.2.3), the gradient (9.45) and the nodes of the sum (9.19); and, for the exact values beside them, the posterior parameters (a_p, b_p) and (a_λ, b_λ) and the trigamma function (Section 9.2.2), with a one-dimensional quadrature for the band.

The start and the run together. By equation (9.47) with forty nodes the start factor is 0.9583 and the probability that the generator starts and runs through the 72 h outage is $0.9583 \times 0.7024 = 0.673$. The value at the MAP state, the plug-in, is $(29/30)\,\mathrm{e}^{-0.32727} = 0.697$, and the mean lies 0.0237 below it. With the Hessian (9.49) at $\hat{\mathbf{z}}$,

$$\mathbf{H} = \begin{bmatrix} -0.0217 & 0.0076 \\ 0.0076 & -0.1534 \end{bmatrix},$$

with both diagonal entries negative, since $\hat{p} = 1/30 < \frac{1}{2}$ and $\hat{\lambda}t = 0.327 < 1$, $\frac{1}{2}\,\mathrm{tr}(\mathbf{H}\boldsymbol{\Sigma}) = \frac{1}{2}(H_{11}\Sigma_{11} + H_{22}\Sigma_{22}) = -0.0056 - 0.0192 = -0.0248$, which predicts the mean at $0.697 - 0.025 = 0.672$, and the fourth-order remainder is 0.0011. The value beside it comes from the exact posterior: by equation (9.50) the probability is $0.9667 \times 0.7301 = 0.706$, which is 0.033 above the state route. About 0.006 of that comes from the start, where the normal on the logit gives $\mathbb{E}[1-p] = 0.9583$ against the exact 0.9667, and about 0.027 from the run, 0.7024 against 0.7301. The plug-in $(1 - 0.0333) \times 0.7209 = 0.697$ is lower than the exact mean by 0.009, the run's gap 0.0092 scaled by 0.9667; the exact mean lies 0.0089 above $R_{\mathrm{sr}}(\hat{\mathbf{z}})$, an identity of the closed forms.

With $\mathbf{g} = (-0.0333,\ -0.3273)^{\mathsf{T}}$ and $\boldsymbol{\Sigma}$ of equation (9.17), the budget (9.46) is

$$\begin{aligned} \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} &= \begin{bmatrix} -0.0333 & -0.3273 \end{bmatrix} \begin{bmatrix} 0.51724 & 0 \\ 0 & 0.25 \end{bmatrix} \begin{bmatrix} -0.0333 \\ -0.3273 \end{bmatrix} \\ &= \underbrace{0.00057}_{\text{start rows}} + \underbrace{0.02678}_{\text{run rows}} + \underbrace{2 \times 0}_{\text{cross term}} = 0.02735, \end{aligned}$$

a standard deviation of 0.165, of which the run term holds 98 %: the run alone would leave $\sqrt{0.02678} = 0.164$. The value beside it comes from the exact posterior: the exact budget (9.51)

is

$$\operatorname{Var}(\ln R_{\mathrm{sr}} \mid \mathbf{y}) = \underbrace{\psi'(58) - \psi'(60)}_{0.00058} + \underbrace{72^2 \times 4/880^2}_{0.02678} = 0.02736,$$

a gap of 0.00001.

The quantiles of equation (9.48), with forty nodes, give

$$R_{\mathrm{sr}}(72) \in [0.399,\ 0.855] \text{ with } 95\,\% \text{ posterior probability}, \qquad \text{median } 0.690,$$

$\Pr(R_{\mathrm{sr}}(72) \geq 0.90 \mid \mathbf{y}) = 0.002$ and the standard deviation 0.118. The normal on $\ln R_{\mathrm{sr}}$, at its MAP value $\ln(29/30) - 0.3273 = -0.3612$ with the standard deviation 0.165, gives the band $\exp(-0.3612 \pm 1.96 \times 0.165) = [0.504,\ 0.964]$ and 0.061 for the same probability: the budget linearizes the rate, and the band it alone gives is too high at both ends. The delta method on R_{sr} itself has

$$\mathbf{g}_R = R_{\mathrm{sr}}(\hat{\mathbf{z}})\,\mathbf{g} = 0.6969 \begin{bmatrix} -0.0333 \\ -0.3273 \end{bmatrix} = \begin{bmatrix} -0.0232 \\ -0.2281 \end{bmatrix}, \qquad \mathbf{g}_R^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g}_R = 0.01328,$$

a standard deviation of 0.115, about the plug-in center. The values beside them come from the exact posterior: the quadrature (9.52), inverted for its quantiles, gives

$$R_{\mathrm{sr}}(72) \in [0.470,\ 0.889] \text{ with } 95\,\% \text{ posterior probability}, \qquad \text{median } 0.715, \tag{9.54}$$

and $\Pr(R_{\mathrm{sr}}(72) \geq 0.90 \mid \mathbf{y}) = 0.016$. By equation (9.53), $\mathbb{E}[(1-p)^2 \mid \mathbf{y}] = (58 \times 59)/(60 \times 61) = 0.9350$, and the exact posterior standard deviation of $R_{\mathrm{sr}}(72)$ is $\sqrt{0.9350 \times 0.54542 - 0.70576^2} = 0.109$. The state route's band lies below the exact one by 0.071 at its lower end and 0.034 at its upper end, its median by 0.025, and its probability of reaching 0.90 by 0.014.

9.3.6 The zero-failure demonstration hours

Engineering question. How many failure-free run hours would demonstrate $R(72) \geq 0.90$ at 95 % assurance? The design basis may require the generator's reliability to be shown rather than estimated, and the engineer asks how long a load test, planned in advance to pass only without a failure, would have to run.

General form. The design form of Chapter 2, equation (2.16), with the indicator that the acceptance rule (9.33) is met as the quantity u and the number K of failures in a test of T' hours as the outcome; a demonstration fixes its accepted outcome in advance, $K = 0$, and its size is the smallest test whose planned outcome meets the rule, as in Chapter 5, equation (5.56):

$$T_{\mathrm{demo}} = \min\{\, T' : \Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}, K = 0 \text{ in } T') \geq \kappa \,\}. \tag{9.55}$$

Specialized form. The question is equation (9.55) with $f(\mathbf{z}) = R(t;\lambda)$, $f^* = R_{\mathrm{req}}$, T' the hours of the test, and $p(\mathbf{z} \mid \mathbf{y})$ the posterior of step 3 extended by the zero-failure row of the test, the Laplace normal (9.12) of the extended rows. Two choices enter here. The accepted outcome, zero failures, is fixed before the test starts: this is what makes the tail calculation valid, and it is broken by extending the test after a failure until the record looks good. The test hours are run under the conditions of the record, at the generator's own rate: this lets them add to the hours behind the posterior, and it fails if the test is run at reduced load, or if hours from other units are pooled without the units sharing one rate (Section 9.6.6).

A failure-free run of T' hours adds the count row of exposure T' and count zero to z_2. It moves the MAP entry to $\hat{z}_2 = \ln(a_\lambda/(b_\lambda + T'))$ and leaves $\Sigma_{22} = 1/a_\lambda$ unchanged, since it adds

exposure and no count to the shape; the tail (9.34) reaches κ when $(\ln\lambda^* - \hat{z}_2)\sqrt{a_\lambda} = \Phi^{-1}(\kappa)$, so that

$$T_{\text{demo}}^{\text{Laplace}} = \frac{a_\lambda\, \mathrm{e}^{\Phi^{-1}(\kappa)/\sqrt{a_\lambda}}}{\lambda^*} - b_\lambda, \tag{9.56}$$

with $a_\lambda \mathrm{e}^{\Phi^{-1}(\kappa)/\sqrt{a_\lambda}}$ the κ quantile of a unit-rate gamma variable under the Laplace normal on its log. Before the test, its outcome is uncertain: the probability that a run of $T_{\text{demo}}^{\text{Laplace}}$ hours ends without a failure is the sum (9.19) at that length. A failure during the test leaves the MAP entry $\ln((a_\lambda + j)/(b_\lambda + T'))$ and the variance $1/(a_\lambda + j)$ after j failures, and the hours required in all become the same expression with $a_\lambda + j$ in place of a_λ. The same formula with the prior parameters gives the hours the demonstration would have needed before the present run record.

The values beside them come from the exact posterior. A run of T' hours that ends with zero failures turns the posterior, by equation (6.30), into Gamma$(a_\lambda,\ b_\lambda + T')$, the same shape and a longer exposure, and by the tail (9.35) the requirement is then supported at assurance κ when

$$F_\Gamma(\lambda^*; a_\lambda, b_\lambda + T') \geq \kappa \quad \Longleftrightarrow \quad (b_\lambda + T')\,\lambda^* \geq F_\Gamma^{-1}(\kappa; a_\lambda, 1),$$

with $F_\Gamma^{-1}(\kappa; a, 1)$ the κ quantile of the Gamma$(a, 1)$ distribution, since a gamma variable times its rate has rate one. Hence the demonstration hours are

$$T_{\text{demo}} = \frac{F_\Gamma^{-1}(\kappa; a_\lambda, 1)}{\lambda^*} - b_\lambda = \frac{F_\Gamma^{-1}(\kappa; a_\lambda, 1)\, t}{-\ln R_{\text{req}}} - b_\lambda, \tag{9.57}$$

the analog for a rate of the zero-failure test count of Chapter 5, equation (5.62), with the gamma quantile $F_\Gamma^{-1}(\kappa; a_\lambda, 1)$ in place of the quantile of equation (9.56). The predictive (9.20) at $t = T_{\text{demo}}$ gives the probability of a clean run. A failure during the test leaves the posterior Gamma$(a_\lambda + j,\ b_\lambda + T')$ after j failures, and the hours required in all become $F_\Gamma^{-1}(\kappa; a_\lambda + j, 1)/\lambda^* - b_\lambda$, and the same formula with the prior parameters gives the hours before the present run record.

Building blocks. The threshold rate λ^* of equation (9.35); the shape a_λ and the exposure b_λ behind $\hat{z}_2$ and Σ_{22}, with Φ^{-1}; the sum (9.19); and, for the exact values beside them, the posterior and prior parameters of the rate (Section 9.2.2), the quantile function of the unit-rate gamma distribution and the predictive (9.20).

The demonstration hours. With $\Phi^{-1}(0.95) = 1.6449$ and $\lambda^* = 0.0014633\,\text{h}^{-1}$, equation (9.56) gives

$$T_{\text{demo}}^{\text{Laplace}} = \frac{4\,\mathrm{e}^{1.6449/2}}{0.0014633} - 880 = \frac{9.1041}{0.0014633} - 880 = 5\,341\,\text{h},$$

more than seven months of continuous running without a failure. By the sum (9.19) with forty nodes the probability that 5 341 hours pass without a failure is 3.2×10^{-5}. After one failure during the test the hours required in all rise to 6 250 h, and after two to 7 145 h. Under the prior alone, before the present record, the same demonstration would have needed 3 973 h.

The value beside it comes from the exact posterior: $F_\Gamma^{-1}(0.95; 4, 1) = 7.7537$, the value at which $\Pr(\text{Poisson}(x) \geq 4) = 0.95$, so equation (9.57) gives

$$T_{\text{demo}} = \frac{7.7537}{0.0014633} - 880 = 5\,299 - 880 = 4\,419\,\text{h},$$

half a year of continuous running without a failure. By equation (9.20) the predictive probability that 4 419 hours pass without a failure is $(880/5299)^4 = 0.0008$. After one failure during the

test the hours required in all rise to 5 375 h, and after two to 6 304 h. Under the prior alone, before the present record, the same demonstration would have needed 2 842 h. The state route asks for 922 h more, 21 %; its unit-rate quantile 9.1041 exceeds the gamma's 7.7537 because the normal on the log leans toward large values. Both lengths are out of reach of a test on one unit, and the gap changes no decision.

9.4 The checks (step 5)

The generator has no physics row and no exact constraint, so no limit $\sigma_p \to 0$ of a model tolerance arises and no limit table is needed; the bounds $0 < p < 1$ and $\lambda > 0$ are not rows either, and the logit and the log rate handle them.

Every answer of Section 9.3 is computed from the posterior of the state, with its exact value beside it; this section tests that agreement across the questions, takes the gap apart, and tests the closed forms and the priors.

9.4.1 The Laplace route against the exact route for the run

The run on the two routes. Sections 9.3.2 and 9.3.3 give the Laplace tail 0.012, less than a third of the exact 0.042, and the Laplace band [0.418, 0.884], with median 0.721, which lies below the exact [0.488, 0.915] at both ends.

Figure 9.3 draws the two posteriors of the 72 h reliability. The horizontal axis is the reliability, the vertical axis its posterior density. The green curve is the exact density, the transformed gamma of equation (9.30); the dotted blue curve is the density (9.27) implied by the Laplace normal (9.12) of the log rate. The dotted lines mark the exact mean and median; the red line is the requirement 0.90, and the shaded area under the exact density to its right is the exact probability of the requirement; the blue area under the Laplace density is the reported one. The exact posterior is broad and skewed to the left, its mean below its median because its long tail toward low reliabilities is the image of the gamma's long tail toward high rates. The Laplace density has nearly the same peak but sits slightly to the left and is thinner on the right, where the requirement lives. The gap is large here, and it can be taken apart.

The exact moments of the log rate. For the shape 4, $\psi(4) = 1 + \frac{1}{2} + \frac{1}{3} - 0.5772 = 1.2561$ and $\psi'(4) = \pi^2/6 - 1 - \frac{1}{4} - \frac{1}{9} = 0.2838$ (Section 9.2.2), hence

$$\mathbb{E}[\ln \lambda \mid \mathbf{y}] = 1.2561 - 6.7799 = -5.5238, \qquad \mathrm{Var}(\ln \lambda \mid \mathbf{y}) = 0.2838 = 0.5328^2.$$

The Laplace normal places its center 0.130 above the exact mean, at the log of the mean rate, and its variance 0.25 is 12 % below the exact 0.284. Both errors push probability away from small rates, which is where the requirement lives.

Three layers of error. Table 9.6 sets three routes side by side. Each row is a summary of the 72 h reliability; the columns give it under the exact gamma posterior, under the Laplace normal (9.12), and under a normal on $\ln \lambda$ with the exact mean and variance. The difference between the Laplace column and the exact-moments column is the cost of the mode and the curvature; the difference between the exact-moments column and the exact column is the cost of the normal shape, the left skew of the exact posterior of $\ln \lambda$. On the standardized scale the three routes put the threshold at -2.267 (Laplace), -1.883 (exact moments) and $\Phi^{-1}(0.042) = -1.729$ (exact).

The Laplace route understates every upper-tail probability, by a factor of three and a half at 0.90 and of nearly forty at 0.95; the exact moments recover most of the error at 0.80 and 0.90 and little of it at 0.95, where only the shape of the gamma closes the gap. The predictive

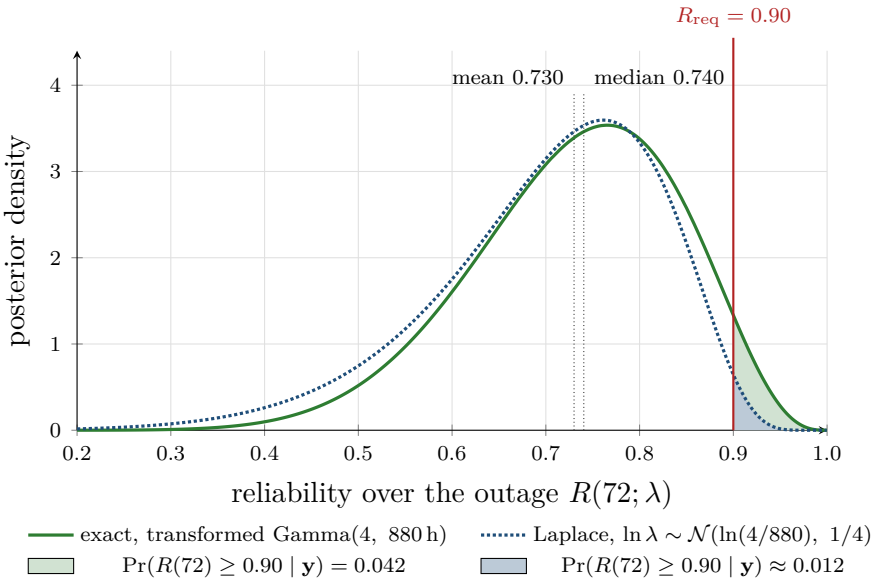


Figure 9.3: The posterior density of the generator's 72 h run reliability. Exact: equation (9.30) with $a_\lambda = 4$, $b_\lambda = 880\,\text{h}$, $t = 72\,\text{h}$, density $(\nu/R_0)\,(-\nu \ln R_0)^3\, R_0^\nu/6$ at the reliability level R_0, with $\nu = b_\lambda/t = 880/72 = 12.22$. Laplace: the normal $\mathcal{N}(\ln(4/880),\ 0.5^2)$ on $z_2 = \ln\lambda$ mapped to $R_0 = \text{e}^{-72\exp(z_2)}$, density $\varphi((\ln(-\ln R_0/72) + 5.3936)/0.5)/(0.5\,R_0\,|\ln R_0|)$ with φ the standard normal density. The shaded areas right of 0.90 are 0.042 under the exact density and 0.012 under the Laplace density.

Table 9.6: The generator's 72 h run reliability by three routes. Exact: the gamma posterior Gamma(4, 880 h) through equations (9.20), (9.29) and (9.35). Laplace: $\ln\lambda \sim \mathcal{N}(-5.3936,\ 0.5^2)$. Exact moments: $\ln\lambda \sim \mathcal{N}(-5.5238,\ 0.5328^2)$. Means by Gauss–Hermite quadrature for the two normals.

Summary	Exact [–]	Laplace, reported [–]	Exact moments [–]
Center and std. dev. of $\ln\lambda$	-5.524, 0.533	-5.394, 0.500	-5.524, 0.533
Mean of $R(72)$, the predictive	0.730	0.702	0.730
Median of $R(72)$	0.740	0.721	0.750
95 % band on $R(72)$	[0.488, 0.915]	[0.418, 0.884]	[0.442, 0.904]
$\Pr(R(72) \geq 0.80 \mid \mathbf{y})$	0.292	0.222	0.318
$\Pr(R(72) \geq 0.90 \mid \mathbf{y})$	0.042	0.012	0.030
$\Pr(R(72) \geq 0.95 \mid \mathbf{y})$	0.0039	0.0001	0.0006

itself, a mean rather than a tail, is the least sensitive: the Laplace route errs by 0.028 on it, the exact moments by less than 0.001. With a posterior shape of four, two recorded failures and two prior pseudo-failures, the posterior of the log rate is skewed enough that a tail is where the Laplace normal at the mode departs furthest; unlike the logit of Chapter 8, a difference of two log rates of equal shape whose offsets cancel, one log rate stands alone here and nothing cancels.

The covered outage and the demonstration. Five nodes give the Laplace predictive at 72 h to 0.7024, so the whole difference from the exact 0.7301 is the Laplace normal, not the quadrature. The covered outages of Section 9.3.4, 10.0 and 20.8 h at 0.95 and 0.90, lie 12 % below the exact 11.4 and 23.5 h. Compared like with like, exact against Laplace, the inversion is far less sensitive than the tail: the outage length moves by 12 % between the two routes, while the probability of the requirement changes by a factor of three and a half. The demonstration, seen from the side of the test, carries the same understated tail: by equation (9.56) the Laplace route asks for $4\,\mathrm{e}^{1.6449/2}/0.0014633 - 880 = 5\,341$ h of failure-free running, against the exact 4 419 h, 21 % more. No decision of the chapter changes on either route.

More counts shrink the errors, whose leading terms are of order $1/(2a_\lambda)$ in the center and $1/(2a_\lambda^2)$ in the variance, equation (8.20).

> **Principle 9.1.** *For a single log rate carried with only a few failures, the Laplace normal at the mode understates the probability that the rate is small, and with it the support for a high mission reliability. Where an exact gamma route exists its value is computed beside the state route, and the gap is reported with the answer.*

9.4.2 The start and the run on the Laplace route

The budget. Table 9.7 lays the Laplace budget of Section 9.3.5 beside the exact budget: each row is a term of equation (9.46); the columns give the gradient entry, the variance of the state entry, the Laplace term, the exact term and the share of the total. The run terms are equal, as Section 9.3.5 derived; the Laplace start term is within 2 % of the exact 0.00058; both totals give the standard deviation 0.165 of $\ln R_{\mathrm{sr}}$. The start term is small because of its gradient entry, not its variance: $\sqrt{\Sigma_{11}} = 0.72$ exceeds $\sqrt{\Sigma_{22}} = 0.5$, so the logit of the start is the less well known entry, but $g_1 = -\hat{p} = -0.0333$ is a tenth of $g_2 = -\hat{\lambda}t = -0.3273$, and the factor of 96 between their squares outweighs the factor of two between the variances.

Table 9.7: The variance budget of the log start-and-run reliability at $t = 72$ h, equation (9.46): $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} = g_1^2\Sigma_{11} + g_2^2\Sigma_{22}$ with $\mathbf{g} = (-\hat{p},\ -\hat{\lambda}t)^{\mathsf{T}}$. Exact: $\psi'(58) - \psi'(60)$ for the start, $t^2 a_\lambda/b_\lambda^2$ for the run, equation (9.51).

Term	Gradient g_i [–]	Σ_{ii}, Laplace [–]	Laplace [–]	Exact [–]	Share [%]
Start, $z_1 = \xi$	−0.0333	$1/1.933 = 0.5172$	0.00057	0.00058	2.1 %
Run, $z_2 = \ln\lambda$	−0.3273	$1/4 = 0.25$	0.02678	0.02678	97.9 %
Cross term		0	0	0	0
Total			0.02735	0.02736	
Std. dev. of $\ln R_{\mathrm{sr}}$			0.165	0.165	

The band and the mean. The band from the normal of the state, [0.399, 0.855], lies below the exact [0.470, 0.889] at both ends, and its probability 0.002 of reaching 0.90 against the exact 0.016 is the run's understated tail again. The linearized band [0.504, 0.964] with its 0.061, against the same exact values, errs the other way: the variance is right, the symmetric shape is

not. The agreement of the budgets 0.02735 and 0.02736 is that of the linearization with the exact moments; under the normal of the state itself the variance of $\ln R_{\text{sr}}$ is 0.0402, a standard deviation of 0.200, because the rate is then lognormal. The delta method's standard deviation 0.115 and the sum's 0.118 compare with the exact 0.109, and the delta method's center is the plug-in 0.697, not the predictive 0.706.

The mean 0.673 is low by 0.033. The exact mean lies above $R_{\text{sr}}(\hat{\mathbf{z}})$, the Jensen gap in the rate, where $\mathrm{e}^{-\lambda t}$ is convex; the Laplace mean lies 0.0237 below it, because in the log rate the reliability is concave over the posterior's center, and only that part is approximated, by the trace. The Laplace route gives the Jensen gap the opposite sign.

9.4.3 The quadrature and the draws

The quadrature against the closed forms. The quadrature (9.52) returns the whole distribution of $R_{\text{sr}}(72)$, and its first two moments have the closed forms (9.50) and (9.53). For a quantity on $(0, 1)$ the mean is $\int_0^1 \Pr(R_{\text{sr}} > R_0 \mid \mathbf{y})\,\mathrm{d}R_0$ and the second moment $\int_0^1 2R_0 \Pr(R_{\text{sr}} > R_0 \mid \mathbf{y})\,\mathrm{d}R_0$. Integrating one minus the cumulative distribution in this way gives 0.70576 and 0.50995; the closed forms give $0.96667 \times 0.73010 = 0.70576$ and $0.93497 \times 0.54542 = 0.50995$, equal to round-off.

The check against draws. The closed forms and the quadrature are checked together with 400 000 independent draws of p from Beta$(2, 58)$ and of λ from Gamma$(4, 880\,\text{h})$ (pseudo-random generator seeded with 9), each pair pushed through equation (9.43). The draws give the mean 0.7057, with a Monte Carlo standard error of 0.0002, against 0.7058 from equation (9.50); the band [0.4705, 0.8893], median 0.7150, against equation (9.54); the probability 0.0164 of reaching 0.90 against 0.0161; and the variance of $\ln R_{\text{sr}}$, 0.02730, split into 0.00058 for the start and 0.02673 for the run, against the budget of Table 9.7. The same draws give the run alone the band [0.488, 0.915] and $\Pr(R(72) \geq 0.90) = 0.0416$, the exact values of Sections 9.3.2 and 9.3.3. On the state route, 400 000 draws of $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ (same seed) give the mean 0.6733, with a Monte Carlo standard error of 0.0002, against the 0.6731 of equation (9.47); the band [0.400, 0.855], median 0.690; the probability 0.0021 of reaching 0.90 against 0.0020; and the variance 0.0401 of $\ln R_{\text{sr}}$ against 0.0402.

9.4.4 Sensitivity to the priors

The run prior Gamma$(2, 400\,\text{h})$ weighs as much as two failures in four hundred hours, nearly as much as the record itself: at the MAP state its row carries $400\hat{\lambda} = 1.818$ of the 4 units of information in z_2, a share of $400/880 = 0.45$, the counterpart for a rate of the prior share of Chapter 5. Table 9.8 repeats the run answers under five priors on the rate, with the start prior held at Beta$(1, 19)$. Each row is a prior; the columns give the posterior, the prior row's share of the information in z_2, and then, each with the Laplace value on the state route first and the exact value in parentheses, the run predictive at 72 h, the band of $R(72)$, the support for $R(72) \geq 0.90$, the demonstration hours of equations (9.56) and (9.57) and the start-and-run predictive. The first row is the engineering prior; the second a weaker prior of the same mean, worth half as much; the third a prior flat in the rate, Gamma$(1, 0)$; the fourth the Jeffreys prior of a Poisson rate, Gamma$(\frac{1}{2}, 0)$; the fifth a strong prior of the same mean as the engineering one, worth four times as much. The flat and the Jeffreys priors are improper and add nothing to the exposure, so the record alone sets $b_\lambda = 480\,\text{h}$ under them.

No prior in the table brings the requirement near reach: its support stays between 0.0001 and 0.028 (0.001 and 0.079 exactly), far below 0.95, and the demonstration hours between 4 355 and 9 416 (3 300 and 8 650). The predictive moves more, from 0.615 to 0.704 (0.658 to 0.739), and the band most, its lower end from 0.248 to 0.526 (0.338 to 0.554): with two recorded failures the width of what is known about the generator is largely the prior's. The two improper priors

Table 9.8: Sensitivity of the generator's run answers to the prior on the run rate; record $k_\mathrm{r} = 2$ in 480 h, start prior Beta$(1, 19)$, mission 72 h. Laplace value, then exact in parentheses; the Laplace normal of each posterior Gamma(a_λ, b_λ) is $\mathcal{N}(\ln(a_\lambda/b_\lambda),\ 1/a_\lambda)$. Prior share: $b_{\lambda 0}/b_\lambda$, the prior row's share of the information in z_2 at the MAP state. T_demo from equations (9.56) and (9.57).

Prior on λ	Posterior	Prior share [–]	Predictive 72 h [–]	Band of $R(72)$ [–]	Pr($R(72)$ ≥ 0.90) [–]	T_demo [h]	Start and run [–]
Gamma(2, 400 h), engineering	Gamma(4, 880 h)	0.45	0.702 (0.730)	[0.418, 0.884] ([0.488, 0.915])	0.012 (0.042)	5 341 (4 419)	0.673 (0.706)
Gamma(1, 200 h), weak	Gamma(3, 680 h)	0.29	0.704 (0.739)	[0.373, 0.903] ([0.465, 0.937])	0.028 (0.079)	4 619 (3 622)	0.675 (0.715)
Gamma$(1, 0)$, flat	Gamma(3, 480 h)	0	0.615 (0.658)	[0.248, 0.865] ([0.338, 0.911])	0.006 (0.034)	4 819 (3 822)	0.590 (0.636)
Gamma$(\frac{1}{2}, 0)$, Jeffreys	Gamma(2.5, 480 h)	0	0.659 (0.705)	[0.274, 0.897] ([0.382, 0.940])	0.022 (0.076)	4 355 (3 303)	0.632 (0.682)
Gamma(8, 1600 h), strong	Gamma(10, 2080 h)	0.77	0.700 (0.712)	[0.526, 0.830] ([0.554, 0.847])	0.0001 (0.001)	9 416 (8 652)	0.670 (0.688)

widen the band toward low reliabilities; the strong prior, surer than the record that the rate is near $0.005\,\mathrm{h}^{-1}$, narrows it and drives the support for 0.90 to 0.0001 (0.001). The start prior matters little: under the uniform Beta$(1, 1)$ the start-and-run predictive is 0.661 (0.695) and under the Jeffreys Beta$(\frac{1}{2}, \frac{1}{2})$ 0.668 (0.703), against 0.673 (0.706), and the start term holds at most 4.6 % (4.7 %) of the budget under any prior of the check.

9.5 The report (step 6)

No estimate closes the gap between the first reading, 0.741, and the 0.90 of the design basis. The posterior probability that the generator's 72 h reliability reaches 0.90 is 0.012 (0.042 exactly), and below 0.03 (0.08) under every prior examined: the two failures to run point to a run rate about three times the one the requirement allows. The answers are computed from the posterior of the state, the Laplace normal of the logit and the log rate, with the exact value beside each from the Beta and gamma posteriors and, for the start and the run together, a one-dimensional quadrature checked by draws.

9.5.1 The answers and their bands

Table 9.9 gives the answers. Each row is a question of Table 9.1, in its order; the second column gives the answer from the posterior of the state under the engineering priors, the third the exact value, and the fourth the band of the reported answer, the other routes and the range over the priors of Table 9.8.

9.5.2 The predictive and the posterior in use

The predictive 0.702 (0.730 exactly) and the posterior of $R(72)$, with its band [0.418, 0.884] ([0.488, 0.915]) and its support 0.012 (0.042) for 0.90, are both correct, and they are not interchangeable. A requirement on the unit is a tail of the posterior (Section 9.3.3); two further uses separate them.

An expected loss over one outage. An operator who must decide whether to rent a second generator for a forecast outage needs the probability that this outage is carried, all uncertainty

Table 9.9: The report for the generator's records. Answers from the posterior of the state, $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$; exact values from the Beta and gamma posteriors and the quadrature (9.52); bands are central 95 % posterior intervals of the reported route, exact in parentheses; prior ranges over Table 9.8, Laplace with exact in parentheses.

Question	Answer (state route)	Exact	Band
Does the generator, once started, run through a 72 h outage?	predictive 0.702	0.730	plug-in 0.721 at the posterior mean, 0.741 at the record's rate; 0.615 to 0.704 (0.658 to 0.739) over the priors
How well is its 72 h reliability known?	median 0.721, mean 0.702, std. dev. 0.121	median 0.740, mean 0.730, std. dev. 0.111	[0.418, 0.884] ([0.488, 0.915] exact); lower end 0.248 to 0.526 (0.338 to 0.554) over the priors
Can $R(72) \geq 0.90$ be asserted with 95 % assurance?	probability 0.012: not asserted	0.042: not asserted	exact moments 0.030; prior alone 0.041 (0.117); 0.0001 to 0.028 (0.001 to 0.079) over the priors
What outage length can be covered at 0.95 and 0.90?	10.0 h and 20.8 h; with 95 % assurance 5.0 and 10.2 h	11.4 h and 23.5 h; with 95 % assurance 5.8 and 12.0 h	plug-in 11.3 and 23.2 h
Does the generator start and run 72 h?	probability (predictive) 0.673; $R_{\mathrm{sr}}(72)$ itself: median 0.690	0.706; median 0.715	[0.399, 0.855] ([0.470, 0.889] exact); linearized normal [0.504, 0.964]; budget 98 % run
How many failure-free run hours would demonstrate $R(72) \geq 0.90$ at 95 %?	5 341 h	4 419 h	clean with probability 3.2×10^{-5} (0.0008); 6 250 h (5 375 h) in all after one failure; 4 355 to 9 416 h (3 300 to 8 650 h) over the priors

included. That is the predictive: the expected loss of the outage is the consequence times $1 - R_{\mathrm{pred}}$, whatever the spread of the reliability.

Many outages over a life. Over many outages, the generator fails to carry them at a long-run frequency set by its true rate, not by the predictive. The outages share the rate exactly as two generators of one design would (Section 9.6.6), and the number of outages carried over a life is a mixture over the posterior of the reliability; its spread is the spread of $R(72)$, not that of a binomial with the predictive probability.

Common mistake

Asserting a reliability requirement from the predictive. The predictive 0.702 (0.730 exactly) is the mean of the posterior of $R(72)$; the support for $R(72) \geq 0.90$ is its upper tail, 0.012 (0.042). Raising the predictive above the requirement does not demonstrate the requirement either: a predictive of 0.92 from a broad posterior can leave the support well below 95 %.

9.6 The next measurement (step 7)

The report leaves the start-and-run reliability known only between 0.40 and 0.86 (0.47 and 0.89 exactly), and the requirement on the run out of reach of the present records. Tests could be bought to narrow the reliability, and the preposterior values each before its outcome is known.

9.6.1 The candidate tests

Two tests of comparable effort could be bought:

- *Twenty start tests.* Twenty more demands add a binomial row of the form (9.5) with $N = 20$, acting on z_1.
- *Two hundred run hours.* A load run of 200 h adds a Poisson count row with the exposure $T'_{\text{tot}} = 200\,\text{h}$, acting on z_2.

A longer run of 480 h, which repeats the present run record, and one of 960 h are added as references, with the 5 341 h of the demonstration of Section 9.3.6 and the 4 419 h of its exact value.

9.6.2 The preposterior of the start-and-run reliability

The added information. The preposterior covariance of Chapter 2, equation (2.17), is $\boldsymbol{\Sigma}_a = (\boldsymbol{\Sigma}^{-1} + \boldsymbol{\Lambda}_a)^{-1}$, with $\boldsymbol{\Lambda}_a$ the expected information of the candidate record at the MAP state. A candidate brings new rows of the form (9.5), and its information is computed exactly as that of the present rows, $\boldsymbol{\Lambda}_a = \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a$, with $\mathbf{J}_a$ the Jacobian rows (9.7) of the new rows at $\hat{\mathbf{z}}$. The binomial row of n' new demands has the Jacobian row $\sqrt{n'\hat{p}(1-\hat{p})}\,(1,\ 0)$, whose square $n'\hat{p}(1-\hat{p})$ is the binomial Fisher information of Chapter 5, which does not depend on the outcome; the count row of a new exposure T'_{tot} has the Jacobian row $\sqrt{T'_{\text{tot}}\hat{\lambda}}\,(0,\ 1)$, whose square is the expected number of failures in it, Chapter 6. For twenty start tests and for 200 run hours,

$$\mathbf{J}_{20\,\text{starts}} = \sqrt{20 \times \tfrac{1}{30} \times \tfrac{29}{30}}\ \begin{bmatrix} 1 & 0 \end{bmatrix} = \begin{bmatrix} 0.8028 & 0 \end{bmatrix},$$

$$\mathbf{J}_{200\,\text{h}} = \sqrt{200 \times 0.004545}\ \begin{bmatrix} 0 & 1 \end{bmatrix} = \begin{bmatrix} 0 & 0.9535 \end{bmatrix},$$

so that

$$\boldsymbol{\Lambda}_{20\,\text{starts}} = \begin{bmatrix} 0.644 & 0 \\ 0 & 0 \end{bmatrix}, \qquad \boldsymbol{\Lambda}_{200\,\text{h}} = \begin{bmatrix} 0 & 0 \\ 0 & 0.909 \end{bmatrix},$$

and the run records of 480 and 960 h have $\boldsymbol{\Lambda}_a = \text{diag}(0,\ 2.182)$ and $\text{diag}(0,\ 4.364)$. With $\boldsymbol{\Sigma}^{-1} = \text{diag}(1.9333,\ 4)$ of equation (9.17), the two preposterior covariances are

$$\boldsymbol{\Sigma}_{20\,\text{starts}} = \begin{bmatrix} 0.38793 & 0 \\ 0 & 0.25 \end{bmatrix}, \qquad \boldsymbol{\Sigma}_{200\,\text{h}} = \begin{bmatrix} 0.51724 & 0 \\ 0 & 0.20370 \end{bmatrix},$$

each diagonal because no candidate row acts on both entries, and each leaving untouched the entry its row does not feed.

The preposterior variance. The preposterior variance of $\ln R_{\text{sr}}$ is the budget of the new covariance, with the gradient (9.45),

$$\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_a\mathbf{g} = \frac{g_1^2}{1.933 + (\boldsymbol{\Lambda}_a)_{11}} + \frac{g_2^2}{4 + (\boldsymbol{\Lambda}_a)_{22}} = \frac{0.00111}{1.933 + (\boldsymbol{\Lambda}_a)_{11}} + \frac{0.10711}{4 + (\boldsymbol{\Lambda}_a)_{22}}, \tag{9.58}$$

with $(\boldsymbol{\Lambda}_a)_{11}$ and $(\boldsymbol{\Lambda}_a)_{22}$ the diagonal entries of $\boldsymbol{\Lambda}_a$, its standard deviation being $(\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_a\mathbf{g})^{1/2}$. For twenty start tests, $0.00111 \times 0.38793 + 0.10711 \times 0.25 = 0.00043 + 0.02678$, $(\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_a\mathbf{g})^{1/2} = 0.1649$; for 200 run hours, $0.00057 + 0.10711 \times 0.20370 = 0.00057 + 0.02182$, $(\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_a\mathbf{g})^{1/2} = 0.1496$.

Averaging over the outcome of the test. The number of failures K the test will produce is uncertain, negative binomial under the posterior for a run (Chapter 6) and Beta–binomial for the starts (Chapter 5); equation (9.58) inserts its expected value. For a run of T'_{tot} hours with K failures the exact posterior is $\text{Gamma}(a_\lambda + K,\ b_\lambda + T'_{\text{tot}})$, whose run term $t^2(a_\lambda + K)/(b_\lambda + T'_{\text{tot}})^2$ is linear in K, so its average over the predictive of K is its value at $\mathbb{E}[K] = T'_{\text{tot}}a_\lambda/b_\lambda$: equation (9.58) exactly, 0.1497 for 200 run hours and 0.1338 for 480. The gradient $-\lambda t$ moves with the posterior, and a run that produces more failures raises the weight of the run term with

the mean rate; holding the gradient at its present value and averaging $1/(a_\lambda + K)$ would give 0.1525 and 0.1389, too large.

Table 9.10 ranks the candidates. Each row is a candidate; the columns give the row it feeds, the diagonal of its information, the preposterior standard deviation of $\ln R_{\mathrm{sr}}$ from equation (9.58), the exact value averaged over the outcome, and the fraction of the present width it removes. Figure 9.4 draws the budget and the ranking. In panel (a) the horizontal axis lists the two terms of Table 9.7 and the vertical axis their variance. In panel (b) the horizontal axis lists the present records and four candidates and the vertical axis the preposterior standard deviation; the dashed line marks the present value.

Table 9.10: The preposterior of the log start-and-run reliability at $t = 72$ h for the candidate tests, equation (9.58). Exact: the posterior variance under the updated Beta or gamma posterior, averaged over the predictive distribution of the number of failures, then its square root. Width removed is for the Laplace column. The last two rows are the zero-failure demonstration run of Section 9.3.6 on the state route and at its exact length, valued like the others by their expected outcome; if they end with zero failures the standard deviations are 0.033 and 0.036.

Test	Row fed	diag($\mathbf{\Lambda}_a$) [–]	$(\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}_a\mathbf{g})^{1/2}$, Laplace [–]	Exact, averaged [–]	Width removed [%]
present records	–	–	0.1654	0.1654	–
twenty start tests	start	(0.644, 0)	0.1649	0.1650	0.3 %
200 run hours	run	(0, 0.909)	0.1496	0.1497	9.5 %
480 run hours	run	(0, 2.182)	0.1338	0.1338	19.1 %
960 run hours	run	(0, 4.364)	0.1157	0.1157	30.0 %
5 341 run hours, the demonstration of Section 9.3.6	run	(0, 24.28)	0.0660	0.0661	60.1 %
4 419 run hours, its exact length	run	(0, 20.08)	0.0709	0.0709	57.1 %

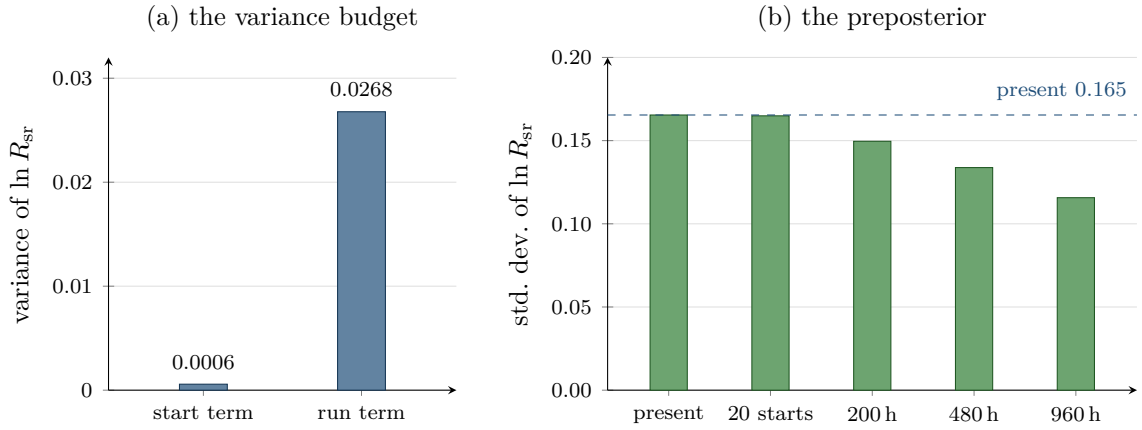


Figure 9.4: The variance budget of the log start-and-run reliability at 72 h and the preposterior of four candidate tests. (a) The two terms of equation (9.46), $g_1^2\Sigma_{11} = 0.00057$ and $g_2^2\Sigma_{22} = 0.02678$ (Table 9.7). (b) The preposterior standard deviation $(\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}_a\mathbf{g})^{1/2}$, equation (9.58), for the present records and the candidates of Table 9.10.

9.6.3 The ranking and its reason

Two hundred run hours narrow the mission reliability; twenty start tests do not. Twenty tests shrink the start term from 0.00057 to 0.00043, a change invisible in the total, and however many are run the standard deviation cannot fall below the run term's 0.164; the run hours act on the term that holds 98 % of the variance, and the width they remove grows with the hours, to 30 % at 960 h. The ranking is the same under the exact averaged values, and the run hours are also the only test that bears on the requirement, a statement about the rate alone.

9.6.4 The recommendation

The next test is run hours, not start tests. Of the two candidates of comparable effort, 200 run hours narrow the standard deviation of the log start-and-run reliability from 0.165 to 0.150, removing 9.5 % of it, and twenty start tests leave it at 0.165. The run hours are bought to narrow the reliability that is reported, and every load run should be logged with its hours and its failures to run, so that the run record keeps growing.

The requirement $R(72) \geq 0.90$ at 95 % assurance is not a testing question for this generator. The record is evidence against it: the two failures to run point to a rate about three times λ^*, and the present posterior gives the 5 341 h demonstration a probability of 3.2×10^{-5} of ending without a failure (0.0008 for the exact 4 419 h); even a generator whose rate is exactly λ^* would pass it with probability only $\mathrm{e}^{-\lambda^* T_{\mathrm{demo}}} = 0.0004$ (0.0016). The gap between the two lengths does not change this. Pooling the hours of several generators that share the rate divides the calendar time of such a test among them and leaves those probabilities as they are. Meeting the requirement calls for a change to the run itself, such as the early-run failures of loading and warming (Section 9.6.6), or to the design basis. The pooled records of a fleet of the same design, under a shared rate or the fleet models of later chapters, settle the question in whichever direction the fleet's rate lies, and for records like this one that direction is against the requirement (Exercise 9.6).

9.6.5 What the analysis showed

The first reading of the record, the survival at the record's own run rate, left the generator short of the design basis by a margin that a better estimate of one number seemed able to close; the analysis overturns that. No estimate closes it: the two failures to run are evidence for a run rate near three times the one the requirement allows, and the support for the 0.90 of the design basis is about one percent (0.012; about four percent, 0.042, exactly), and below 0.03 (0.08) under every prior examined. The probability that the generator starts and runs through the next outage is 0.673 (0.706), close to the first reading, but the generator's reliability over that run is known only to lie between about 0.40 and 0.86 (between about one half and nine tenths exactly). Almost all of that uncertainty is in the run, because the run enters the mission reliability through the expected number of failures in the mission and the start only through its small failure probability; more start tests would not move the answer, and run hours narrow it. On site this sets two different actions: the load runs are extended and logged to narrow the reliability that is reported, and the requirement, which no failure-free test of this unit can realistically show, calls for a change to the run or to the design basis, while the pooled records of generators of the same design settle it in whichever direction their rate lies. And the Laplace route from which every answer is computed, the route a larger model of the data center has, puts the support for the requirement at a third of its exact value and the expected mission reliability below the plug-in, where the exact one lies above it, on a record of two failures where no decision of the chapter turns on the difference.

9.6.6 Limits of the model

The modeling choices were named where they enter the derivation; the derivation is exact under them, and five of them limit the answers.

A constant rate during the mission. Where failures to run crowd into the first hour (Section 9.3.1), the run is split into an early period and the remainder, each with its own rate, as reliability analyses of diesel generators often do: the run reliability is then the product of two exponential survivals, the state gains an entry, and the derivations of this chapter apply to each factor. A rate that rises over a long run, through fuel, cooling or lubrication problems, needs the lifetime models of Chapter 7 or the trend models of Chapter 12.

Independence of start and run. Where a degraded fuel or battery system couples the start and the run (Section 9.3.5), the two entries are correlated in the posterior; a row that acts on both columns represents the coupling, and the variance budget gains a cross term.

Tests that represent demands. Where tests and demands differ (Sections 9.3.1 and 9.3.5), the record estimates the test's p and λ, not the outage's; the effective-exposure device of Chapter 6 weights run hours by their relative stress, and starts after a long standby are kept as a row of their own.

Failures in standby. The model counts failures only on demand and in running. A generator can also degrade while it waits, between tests, so that its start failure probability grows with the time since the last test; the start failure probability of this chapter is then an average over the test interval, and the choice of that interval is the inspection-interval question of a later chapter.

One rate for two units. The model is of one generator. Suppose the data center has two identical generators, either of which can carry the load, and the run record describes their common design: both have the same unknown run rate λ, with the posterior Gamma$(4, 880\,\mathrm{h})$. Given λ, the two runs fail independently, each surviving with probability $s = \mathrm{e}^{-\lambda t}$, and at least one survives with probability $1 - (1 - s)^2 = 2s - s^2$. Averaging over the posterior, and using $s^2 = \mathrm{e}^{-2\lambda t}$,

$$\Pr(\text{at least one runs } t \mid \mathbf{y}) = \mathbb{E}\big[2\mathrm{e}^{-\lambda t} - \mathrm{e}^{-2\lambda t} \mid \mathbf{y}\big] = 2R_{\mathrm{pred}}(t) - R_{\mathrm{pred}}(2t), \tag{9.59}$$

which for the generators is $2 \times 0.7301 - (880/1024)^4 = 1.4602 - 0.5454 = 0.915$. Treating the two generators as independent, each with the predictive reliability 0.7301, gives $1 - (1 - 0.7301)^2 = 0.927$; the plug-in at the posterior mean, $1 - (1 - 0.7209)^2 = 0.922$. Both overstate the redundancy. The independent-predictive shortcut averages over λ separately for each generator, as if each had its own rate drawn from the posterior, when there is one rate. The difference is the covariance of the two survivals through the shared rate,

$$\Pr(\text{both run}) = \mathbb{E}[s^2 \mid \mathbf{y}] = R_{\mathrm{pred}}(2t) = 0.545 \quad \text{against} \quad R_{\mathrm{pred}}(t)^2 = 0.533,$$

a correlation of 0.063 between the two survival events, which the epistemic parameter creates although, for a known rate, they are independent. Redundancy protects against the aleatory failure of one unit, not against the shared uncertainty of their common rate. The coupling lives in the graph: one free entry, the shared log rate, feeds two paths, and the two survivals are correlated because they depend on the same entry of the state. The next part of the book begins there, with system structures, dependent failures and common causes written as graphs whose entries are shared between components.

Common mistake

Multiplying the predictive failure probabilities of two redundant units whose rates are the same unknown quantity. For the two generators, $(1-0.730)^2 = 0.073$ understates the probability that neither runs through the outage, $1-0.915 = 0.085$, by 15 %; the shared rate must be integrated once, over both units together, equation (9.59).

The next chapter turns from the reliability over a mission that starts now to the remaining life of a unit that has already reached a given age, with the uncertainty of its lifetime parameters carried into that prediction.

9.7 What to retain

Definition 9.1 (Mission reliability under an uncertain rate)**.** For a component with a constant failure rate λ and the posterior $\lambda \mid \mathbf{y} \sim \text{Gamma}(a_\lambda, b_\lambda)$, the predictive mission reliability is the posterior expectation $\Pr(T > t \mid \mathbf{y}) = \mathbb{E}[\mathrm{e}^{-\lambda t} \mid \mathbf{y}] = (b_\lambda/(b_\lambda + t))^{a_\lambda}$, with predictive hazard $a_\lambda/(b_\lambda + t)$; it lies between the plug-in at the posterior mean, $\mathrm{e}^{-\mathbb{E}[\lambda]t}$, and that value times $\exp(\frac{1}{2}t^2 \operatorname{Var}(\lambda \mid \mathbf{y}))$. The reliability $R(t;\lambda) = \mathrm{e}^{-\lambda t}$ is a derived quantity whose posterior follows by the quantile map: $\Pr(R(t;\lambda) \geq R_{\text{req}} \mid \mathbf{y}) = F_\Gamma(-\ln R_{\text{req}}/t; a_\lambda, b_\lambda)$, its mean is the predictive and its variance $R_{\text{pred}}(2t) - R_{\text{pred}}(t)^2$. For a standby unit that must start and then run, $R_{\text{sr}} = (1-p)\,\mathrm{e}^{-\lambda t}$, whose log has the two-term variance $\operatorname{Var}(\ln(1-p)) + t^2 \operatorname{Var}(\lambda)$. The answers are computed on the Laplace route, with these closed forms beside them: the state $\mathbf{z} = (\xi,\ \ln\lambda)^\mathsf{T}$ is fitted by Newton's method on four rows with a 4×2 Jacobian of disjoint columns, $\boldsymbol{\Sigma} = (\mathbf{J}^\mathsf{T}\mathbf{J})^{-1} = \operatorname{diag}((a_p + b_p)/(a_p b_p),\ 1/a_\lambda)$, the predictive is the sum (9.19), the requirement the normal tail (9.34), and the same budget is $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$ with $\mathbf{g} = (-\hat{p},\ -\hat{\lambda}t)^\mathsf{T}$.

Three conclusions carry beyond the generator. A reliability computed from uncertain parameters is an expectation, and under the exact posterior of the rate it is never below the reliability at the posterior mean of the parameter; the gap is bounded through $\frac{1}{2}t^2 \operatorname{Var}(\lambda \mid \mathbf{y})$, and the ratio of the two grows without limit with the mission. A requirement on a reliability is a tail of its posterior, a different object from the predictive; every answer is computed from the posterior of the state, and for a single rate with few failures the Laplace normal at the mode understates that tail, so the exact value and the gap are reported with it. And the variance budget of a mission reliability names the record that limits it: when the uncertainty of a rate is multiplied by a long mission, run hours are worth more than any number of demands; but a record that points to a rate above the requirement is evidence against it, and no failure-free test of one unit is a realistic way to overturn it.

Table 9.11 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 9.11: The terms of mission reliability under an uncertain failure rate and where each is defined.

Term	Expression	Units	Meaning	Defined in
Predictive reliability	$\mathbb{E}[\mathrm{e}^{-\lambda t} \mid \mathbf{y}] = (b_\lambda/(b_\lambda + t))^{a_\lambda}$	–	survival averaged over the rate	9.3.1
Plug-in at the posterior mean	$\mathrm{e}^{-\mathbb{E}[\lambda\mid\mathbf{y}]t}$	–	survival at the mean rate	9.1.1
Predictive hazard	$a_\lambda/(b_\lambda + t)$	h^{-1}	mean rate given survival to t	9.3.1
Second-order mean shift	$\mathbb{E}[R] - R(\mathbf{m}) \approx \frac{1}{2}\,\mathrm{tr}(\mathbf{H}\boldsymbol{\Sigma}_\theta)$	–	any coordinate; sign of $\mathbf{H}$	9.3.1
Jensen gap, leading term	$\frac{1}{2}t^2\,\mathrm{Var}(\lambda)\,\mathrm{e}^{-\mathbb{E}[\lambda]t}$	–	one-entry case	9.3.1
Gap bound	$R_{\text{plug-in}} \le R_{\text{pred}} \le R_{\text{plug-in}}\,\mathrm{e}^{t^2\,\mathrm{Var}(\lambda)/2}$	–	gamma posterior	9.3.1
Requirement probability, on the state	$\Phi\big((\ln(\lambda^* t) - \hat{c})/\sqrt{\Sigma_{22}}\big)$	–	the tail reported	9.3.3
Requirement probability, exact	$F_\Gamma(-\ln R_{\text{req}}/t; a_\lambda, b_\lambda)$	–	support for $R(t) \ge R_{\text{req}}$	9.3.3
Quantiles of the reliability	$R_q = \mathrm{e}^{-t\lambda_{1-q}}$	–	quantile map, reversed	9.3.2
Variance of the reliability	$R_{\text{pred}}(2t) - R_{\text{pred}}(t)^2$	–	spread of a derived quantity	9.3.2
Covered mission length	$t^* = b_\lambda(R_{\text{req}}^{-1/a_\lambda} - 1)$	h	inverse of the predictive	9.3.4
Laplace fit of the state	$(\mathbf{J}^\mathsf{T}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^\mathsf{T}\mathbf{r}$, $\boldsymbol{\Sigma} = (\mathbf{J}^\mathsf{T}\mathbf{J})^{-1}$	–	four rows, 4×2 Jacobian	9.2.1
Complementary log-log	$\ln(-\ln R) = \mathbf{g}_c^\mathsf{T}\mathbf{z} + \ln t$	–	linear in the state	9.3.2
Exact log-rate moments	$\psi(a_\lambda) - \ln b_\lambda$, $\psi'(a_\lambda)$	–	Laplace $\ln(a_\lambda/b_\lambda)$, $1/a_\lambda$	9.2.2
Predictive by quadrature	$\sum_j \varpi_j \exp(-t\,\mathrm{e}^{\hat{z}_2 + \sqrt{\Sigma_{22}}\,x_j})$	–	under the Laplace normal	9.3.1
Start-and-run reliability	$R_{\text{sr}} = (1-p)\,\mathrm{e}^{-\lambda t}$	–	start, then run t	9.3.5
Start-and-run budget	$\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$; exact $\mathrm{Var}(\ln(1-p)) + t^2\,\mathrm{Var}(\lambda)$	–	two terms, no cross term	9.3.5
Preposterior variance	$\mathbf{g}^\mathsf{T}(\boldsymbol{\Sigma}^{-1} + \boldsymbol{\Lambda}_a)^{-1}\mathbf{g}$, eq. (9.58)	–	after a candidate test	9.6.2
Demonstration hours	$F_\Gamma^{-1}(\kappa; a_\lambda, 1)/\lambda^* - b_\lambda$	h	zero-failure run to demonstrate	9.3.6
Shared-rate redundancy	$2R_{\text{pred}}(t) - R_{\text{pred}}(2t)$	–	at least one of two runs	9.6.6

9.8 Exercises

Exercise 9.1. *Formulation.* The run failure rate of a fire-pump engine has the posterior $\lambda \mid \mathbf{y} \sim \mathrm{Gamma}(3,\ 1500\,\mathrm{h})$. (a) Compute the predictive reliability and the plug-in reliability at the posterior mean rate for missions of 24 and 168 h. (b) Compute the predictive hazard at the end of each mission. (c) Find the mission length over which the predictive reliability is 0.99. ((a) $(1500/1524)^3 = 0.9535$ against $\mathrm{e}^{-0.048} = 0.9531$; $(1500/1668)^3 = 0.7273$ against $\mathrm{e}^{-0.336} = 0.7146$. (b) $3/1524 = 0.00197\,\mathrm{h}^{-1}$ and $3/1668 = 0.00180\,\mathrm{h}^{-1}$, against the constant

$0.002\,\mathrm{h}^{-1}$. (c) $t^* = 1500(0.99^{-1/3} - 1) = 5.03\,\mathrm{h}$, equation (9.39); the plug-in gives 5.03 h as well, since for a requirement this close to one the ratio (9.40) is 1.0017.)

Exercise 9.2. *Formulation.* The same engine must start on demand; its start failure probability has the posterior Beta$(1, 49)$, independent of the run. (a) Write the probability that it starts and runs 24 h as an expectation over the state, and evaluate it. (b) Write the question "is the probability of starting and running 24 h at least 0.95?" as an event on the state, and say which category of question it is and which of the two objects of Section 9.1.2 answers it. (c) Treating each posterior as one row on its entry of $\mathbf{z} = (\xi,\ \ln\lambda)^\mathsf{T}$, the Beta posterior as a binomial row and the gamma posterior as a count row, write the 2×2 Jacobian, the information matrix and $\boldsymbol{\Sigma}$ at the MAP state, the gradient $\mathbf{g}$ of $\ln R_{\mathrm{sr}}(24)$ there, and $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$. ((a) $\mathbb{E}[(1-p)\,\mathrm{e}^{-24\lambda} \mid \mathbf{y}] = \mathbb{E}[1-p]\mathbb{E}[\mathrm{e}^{-24\lambda}] = (49/50)(1500/1524)^3 = 0.98 \times 0.9535 = 0.934$, the product because the posteriors are independent. (b) $\Pr((1-p)\,\mathrm{e}^{-24\lambda} \geq 0.95 \mid \mathbf{y})$, equivalently $\Pr(\ln(1-p) - 24\lambda \geq \ln 0.95 \mid \mathbf{y})$: a prediction under a threshold on a derived quantity of both entries, answered by the posterior of R_{sr}, not by the predictive 0.934. (c) The rows $50\ln(1 + \mathrm{e}^{z_1}) - z_1$ and $1500\,\mathrm{e}^{z_2} - 3z_2$ give $\hat{\mathbf{z}} = (\ln(1/49),\ \ln 0.002)^\mathsf{T} = (-3.8918,\ -6.2146)^\mathsf{T}$; the working Jacobian (9.7) at $\hat{\mathbf{z}}$ is $\mathbf{J} = \mathrm{diag}(\sqrt{50 \times 0.02 \times 0.98},\ \sqrt{1500 \times 0.002}) = \mathrm{diag}(0.9899,\ 1.7321)$, so $\boldsymbol{\Lambda} = \mathbf{J}^\mathsf{T}\mathbf{J} = \mathrm{diag}(0.98,\ 3)$ and $\boldsymbol{\Sigma} = \mathrm{diag}(1.0204,\ 0.3333)$; $\mathbf{g} = (-\hat{p},\ -24\hat{\lambda})^\mathsf{T} = (-0.02,\ -0.048)^\mathsf{T}$ and $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g} = 0.00041 + 0.00077 = 0.00118$, a standard deviation of 0.034, two thirds of it from the run.)

Exercise 9.3. *Derivation.* For the posterior Gamma$(3, 1500\,\mathrm{h})$ and a mission of 24 h: (a) derive the central 95 % band of $R(24;\lambda)$ by the quantile map, given the Gamma$(3, 1)$ quantiles 0.6187, 2.6741 and 7.2247; (b) derive $\Pr(R(24) \geq 0.97 \mid \mathbf{y})$ exactly and by the Laplace route on the one-number state $z = \ln\lambda$, giving its information Λ, its variance σ_z^2 and the slope g_c of the complementary log-log; (c) state which to trust and why. ((a) $\lambda_{0.025} = 0.000412$, $\lambda_{0.5} = 0.001783$, $\lambda_{0.975} = 0.004816\,\mathrm{h}^{-1}$, so $R(24) \in [\mathrm{e}^{-0.1156},\ \mathrm{e}^{-0.0099}] = [0.891,\ 0.990]$ with median 0.958. (b) $\lambda^* = -\ln 0.97/24 = 0.0012691$, $b\lambda^* = 1.9037$, exact $1 - \mathrm{e}^{-1.9037}(1 + 1.9037 + 1.8120) = 0.297$; Laplace: the posterior as the one row $\mathcal{L}(z) = 1500\,\mathrm{e}^z - 3z$ has the slope $\mathcal{L}' = 1500\,\mathrm{e}^z - 3$, zero at $\hat{z} = \ln 0.002$, and the curvature $\Lambda = \mathcal{L}''(\hat{z}) = 1500\,\mathrm{e}^{\hat{z}} = 3$, so $\sigma_z^2 = 1/3$; with $g_c = 1$ the requirement $z \leq \ln\lambda^*$ has the probability $\Phi((\ln 0.0012691 - \ln 0.002)/\sqrt{g_c^2\sigma_z^2}) = \Phi(-0.788) = 0.215$. (c) The exact value: with a posterior shape of three, failures and prior pseudo-failures together, the Laplace normal at the mode sits $\ln 3 - \psi(3) = 0.176$ above the exact mean of $\ln\lambda$, with variance 1/3 against $\psi'(3) = 0.395$, and understates the probability of small rates.)

Exercise 9.4. *Derivation.* A second generator has the same posterior mean run rate, $0.004545\,\mathrm{h}^{-1}$, but twice the posterior standard deviation. (a) Find its gamma posterior. (b) Derive its predictive and plug-in reliability at 72 h, the gap, and the bound (9.24), and compare with the chapter's generator. (c) Explain from the leading term (9.22) why the gap grows by less than a factor of four. ((a) Mean a/b fixed and standard deviation $\sqrt{a}/b$ doubled require $\sqrt{a}$ halved: $a = 1$, $b = 220\,\mathrm{h}$. (b) $220/292 = 0.7534$ against $\mathrm{e}^{-0.32727} = 0.7209$; gap 0.0325 against 0.0092; bound $0.7209(\mathrm{e}^{0.0536} - 1) = 0.0397$. (c) The leading term is proportional to $\mathrm{Var}(\lambda)$, which is four times larger, giving 0.0386; the higher-order terms of the expansion, negative here, grow with higher powers of the spread, and the exact gap is 3.5 times the chapter's.)

Exercise 9.5. *Application.* Three identical generators share the run rate $\lambda \sim$ Gamma$(4, 880\,\mathrm{h})$, and each has started. (a) Derive the probabilities that at least one and that at least two run through a t-hour outage in terms of $R_{\mathrm{pred}}(t)$, $R_{\mathrm{pred}}(2t)$ and $R_{\mathrm{pred}}(3t)$, and evaluate them at 72 h. (b) Compare with the values obtained by treating the generators as independent with the predictive reliability. ((a) Given λ, with $s = \mathrm{e}^{-\lambda t}$, at least one: $1 - (1-s)^3 = 3s - 3s^2 + s^3$; at least two: $3s^2(1-s) + s^3 = 3s^2 - 2s^3$; averaging, $3R_{\mathrm{pred}}(t) - 3R_{\mathrm{pred}}(2t) + R_{\mathrm{pred}}(3t) = 2.1903 - 1.6363 + 0.4156 = 0.970$ and $3R_{\mathrm{pred}}(2t) - 2R_{\mathrm{pred}}(3t) = 1.6363 - 0.8312 = 0.805$, with $R_{\mathrm{pred}}(216) = (880/1096)^4 = 0.4156$. (b) Independent: $1 - 0.2699^3 = 0.980$ and $3(0.7301)^2 -$

$2(0.7301)^3 = 0.821$; the shared rate lowers both, and the probability that none runs rises from 0.020 to 0.030.)

Exercise 9.6. *Application.* (a) Derive the zero-failure run hours that would demonstrate $R(72) \geq 0.95$ at 95 % assurance for the chapter's generator, Gamma(4, 880 h). (b) Ten generators of the same design are assumed to share one rate; state how the demonstration hours of (a) and of Section 9.3.6 are carried by the fleet. (c) If the fleet's past records are pooled instead, each generator with two failures in 480 h, state what the pooled posterior does to the demonstration. ((a) $\lambda^* = -\ln 0.95/72 = 0.000712\,\mathrm{h}^{-1}$, and equation (9.57) gives $7.7537/0.000712 - 880 = 10\,004$ h. (b) Only the test hours are pooled: under a shared rate the zero-failure running hours of the ten units add in the test exposure, so the 10 004 h for 0.95 and the 4 419 h for 0.90 become about 1 000 and 442 h per generator. (c) Pooling ten such records gives Gamma(22, 5200 h), mean $0.00423\,\mathrm{h}^{-1}$, much narrower around a rate three times λ^* for 0.90: the support for $R(72) \geq 0.90$ falls from 0.042 to 2×10^{-5} and the zero-failure fleet hours rise to 15 465 h; pooling settles the question faster in whichever direction the fleet's rate lies, here against the requirement.)

Chapter 10

Remaining Life and Conditional Survival

10.1 Remaining life of units that have already survived

A data center carries its critical load through a power outage on the batteries of its uninterruptible power supplies (UPS), and the facility engineer answers for twelve battery strings of one design in one UPS room. Each string is given a capacity test once a year. Four have failed theirs, at ages of 2.9, 4.4, 5.8 and 6.9 years, and were replaced; eight are in service, at ages from 1.2 to 6.2 years, and the engineer must decide which of them to replace before the next test and how many spare strings the year needs. The quickest reading fits a Weibull lifetime to the twelve strings and takes the fit as known: a characteristic life of 6.50 years and a shape of 4.73, a steep wear-out, under which the oldest string, at 6.2 years, lasts the coming year with probability 0.44, the eight get through it together without a failure with probability 0.09, and the strings are to be replaced in order of age, oldest first.

None of those numbers says how much four failures can tell. The fit is one point of a posterior that four failures leave wide, and the survival of each string to its age is itself part of the record that produced the fit. The eight strings share the uncertain lifetime distribution, so their chances of lasting the year are not independent. And the ranking by age holds only if the hazard does rise with age, which the record supports to a degree that a fitted shape does not state. The question of this chapter is

what the survival of a unit to its present age says about the unit and about the lifetime parameters it shares with its fleet, how both enter the probability that it survives the next period and its remaining life, and how the survivals of several units that share those parameters combine.

What changes from Chapter 7 is not the state, which is again the log-scale and the log-shape of a Weibull lifetime, but the questions: they are asked of individual rows of the record, the survivors, and of several of them together.

Jargon: battery string, capacity test, conditional survival, remaining life, mean residual life, restricted mean residual life, remaining-life quantile, conditional predictive survival, plug-in, wear-out, shared parameter

A *battery string* is a chain of battery blocks connected in series to give the voltage a UPS needs; a UPS may draw on several strings in parallel. A *capacity test* discharges a string at a stated load and measures how long it sustains it; a string that falls short of the required time has failed the test. The *conditional survival* of a unit of age t_0 is the probability that it survives a further period τ, given that it has survived to t_0; its *remaining life* is the random time $T - t_0$ from its present age to its failure. The *mean residual life* is the expectation of the remaining life, the *restricted mean residual life* its expectation with the remaining life capped at a stated horizon, and a *remaining-life quantile* the time by which the unit has failed with a stated probability, the median for one half. The *conditional predictive survival* is the conditional survival averaged over the posterior of the state. A *plug-in* value is a quantity of the state evaluated at a point estimate, with the uncertainty of the state ignored. *Wear-out* is a hazard that rises with age. A *shared parameter* is an entry of the state on which the lifetimes of several units depend at once.

Key idea

The conditional survival of a unit is the probability of surviving the cumulative hazard of the coming period alone, and it depends on the age already reached only through the slope of the hazard: under a constant hazard the age carries no information about the next year, and under wear-out it carries most of it. When the state is uncertain, the unit's own survival is one of the rows of the record, and the probability that it survives the next period is the posterior average of its conditional survival; for the Weibull model under a flat prior in $(\ln \eta, \ln \beta)$ it is one integral over the shape. Units that share the state must be averaged together, once, never one at a time.

10.1.1 Survival to an age as evidence

A string that has passed its sixth annual test says two things. About itself, it says that its lifetime exceeds six years: whatever happens next, it did not fail early. About the fleet, it says that the lifetime distribution the strings share puts enough probability beyond six years for one of twelve units to reach that age, which makes distributions with short lives less plausible. The first statement alone answers the question of the next year when the lifetime distribution is known; the second matters because it is not. The two enter the answer in different places: the first through the conditioning event $T > t_0$ in the conditional survival, the second through the censored row that the same event contributes to the record.

10.1.2 Conditional survival for a known state

For a known state, the probability that a unit of age t_0 survives a further period τ is, by the definition of a conditional probability and because the event $\{T > t_0 + \tau\}$ is contained in $\{T > t_0\}$, the ratio of two survival probabilities, equation (7.11). Chapter 7 wrote it $R(\tau \mid T > t_0, \mathbf{z})$; this chapter shortens the notation to $R(\tau \mid t_0; \mathbf{z})$, the age reached to the left of the semicolon and the state to the right, and writes every function of the state with a semicolon, $R(t; \mathbf{z})$, $\Lambda(t; \mathbf{z})$:

$$R(\tau \mid t_0; \mathbf{z}) = \Pr(T > t_0 + \tau \mid T > t_0, \mathbf{z}) = \frac{R(t_0 + \tau; \mathbf{z})}{R(t_0; \mathbf{z})}. \tag{10.1}$$

With $R = \mathrm{e}^{-\Lambda}$ and $\Lambda(t) = \int_0^t \lambda(s)\,\mathrm{d}s$, the ratio is the exponential of minus the cumulative hazard accumulated over the coming period alone:

$$R(\tau \mid t_0; \mathbf{z}) = \exp\big[-(\Lambda(t_0+\tau) - \Lambda(t_0))\big] = \exp\Big[-\int_{t_0}^{t_0+\tau} \lambda(s)\,\mathrm{d}s\Big]. \tag{10.2}$$

What happened before t_0 enters only through the fact of survival; the hazard the unit is still to face is that of the interval $(t_0, t_0+\tau]$.

The dependence on age. Differentiate the logarithm of equation (10.2) with respect to t_0. By the fundamental theorem of calculus applied to both limits of the integral,

$$\frac{\partial}{\partial t_0} \ln R(\tau \mid t_0; \mathbf{z}) = -\big[\lambda(t_0+\tau) - \lambda(t_0)\big]. \tag{10.3}$$

The conditional survival falls with age exactly where the hazard rises across the horizon, is flat where the hazard is flat, and rises where the hazard falls. Under the exponential lifetime the hazard is the constant λ, the derivative is zero, and equation (10.2) is $\mathrm{e}^{-\lambda\tau}$ for every age: the memoryless property of equation (7.13). For a known state, a constant hazard is the only one for which the age of a unit carries no information about its next period.

The Weibull lifetime. With $\Lambda(t) = (t/\eta)^\beta = \zeta\, t^\beta$, where $\zeta = \eta^{-\beta}$ is the quantity of Section 7.2.2, equation (10.2) becomes

$$R(\tau \mid t_0; \mathbf{z}) = \exp\big[-\zeta\,\nu(t_0)\big], \qquad \nu(t_0) = (t_0+\tau)^\beta - t_0^\beta, \tag{10.4}$$

with ν in yr^β and ζ in $\mathrm{yr}^{-\beta}$. The increment ν depends on the shape and the ages; the scale enters only through the factor ζ. By equation (10.3), or directly from $\mathrm{d}\nu/\mathrm{d}t_0 = \beta\big[(t_0+\tau)^{\beta-1} - t_0^{\beta-1}\big]$, the conditional survival decreases with age when $\beta > 1$, is constant at $\beta = 1$ and increases when $\beta < 1$. For a horizon short against the age, $\nu \approx \beta\, t_0^{\beta-1}\tau$ and $R(\tau \mid t_0; \mathbf{z}) \approx \exp[-\lambda(t_0)\,\tau]$: the hazard at the present age, times the horizon.

10.1.3 Lifetimes, not degradation signals

The companion volume *Reliability Engineering* treats the remaining useful life of a unit whose degradation is measured: a capacity that fades from test to test, a resistance that grows, a vibration level that rises toward a threshold. This chapter uses lifetimes only. The record of each string is its age and whether it has failed, and the survival of a string to its age is the only information about its condition; the capacity measured at each test is not used. The remaining life computed here is therefore that of a unit known only to belong to its fleet and to have reached its age, and a degradation reading, when one exists, is a further row that narrows it.

10.1.4 The strings and their plant graph (step 0)

The twelve strings are of one design, from one manufacturer, installed in one UPS room at different dates over the life of the facility, and held at the same float charge and room temperature. They are treated as identical units of one population whose lifetime distribution is unknown. Each string is observed from its installation until it fails a capacity test or until the latest test, which it passed.

Figure 10.1 draws the system as a graph. The gray box at the upper left is the UPS room, the pinned operating conditions, with a solid arrow from it to the string node. The green box is the free node: one battery-string design, shared by the whole fleet, which carries the two entries of the state, the log-scale $u = \ln\eta$ and the log-shape $v = \ln\beta$ of a Weibull lifetime. The twelve boxes below are the evidence rows, one per string, each reached by a dashed observation

arrow from the one free node. The four red boxes, strings A to D, are failure rows; the eight green boxes, strings E to L, are survivor rows. The gray bar at the bottom is the set of survivor ages, pinned by the installation and test dates, which enters each survivor row as its censoring time. There is one free node and twelve rows; nothing in the graph distinguishes one string's parameters from another's.

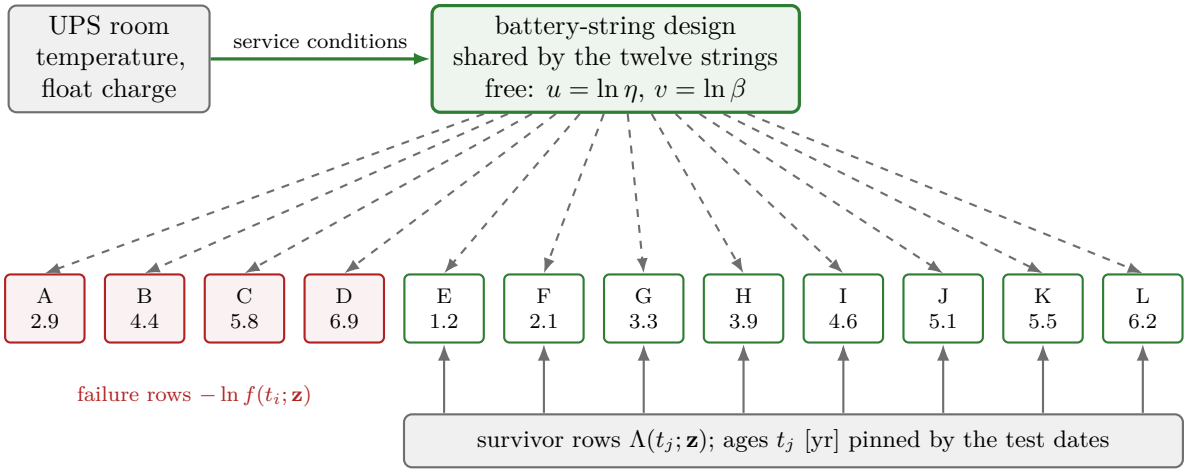


Figure 10.1: The plant graph of the battery strings. Free: the string design, one node shared by the fleet, carrying $u = \ln \eta$ and $v = \ln \beta$ (green). Pinned: the UPS room conditions and the ages of the survivors (gray). Observed: one evidence row per string (dashed), four failure rows at the failure ages (red) and eight survivor rows at the survivor ages (green), each labeled with its string and its age in years.

The record. Table 10.1 lists the record. Each row is a string; the columns give its age at the end of its observation, its event indicator and its outcome. Four strings, A to D, failed their capacity tests at ages 2.9, 4.4, 5.8 and 6.9 years; eight, E to L, are in service at ages from 1.2 to 6.2 years, lettered in order of age, so that L is the oldest. The record is illustrative.

Table 10.1: The record of the twelve battery strings (illustrative). Ages are at the end of observation: the failed test for A to D, the latest passed test for E to L.

String	Age t_i [yr]	Indicator δ_i [–]	Outcome
A	2.9	1	failed its capacity test
B	4.4	1	failed its capacity test
C	5.8	1	failed its capacity test
D	6.9	1	failed its capacity test
E	1.2	0	in service
F	2.1	0	in service
G	3.3	0	in service
H	3.9	0	in service
I	4.6	0	in service
J	5.1	0	in service
K	5.5	0	in service
L	6.2	0	in service
Totals	$T_{\text{tot}} = 51.9$	$k = 4$	

Figure 10.2 draws the same record against age, in the convention of Figure 7.3. The

horizontal axis is the age of a string in years and each horizontal line, one per string, runs from installation to the end of its observation; a filled circle marks a failed test and a vertical bar a string still in service. The failures are spread from the third to the seventh year, and the oldest survivor, at 6.2 years, is younger than the oldest failure. No string failed before 2.9 years, although eleven of the twelve were observed beyond two years.

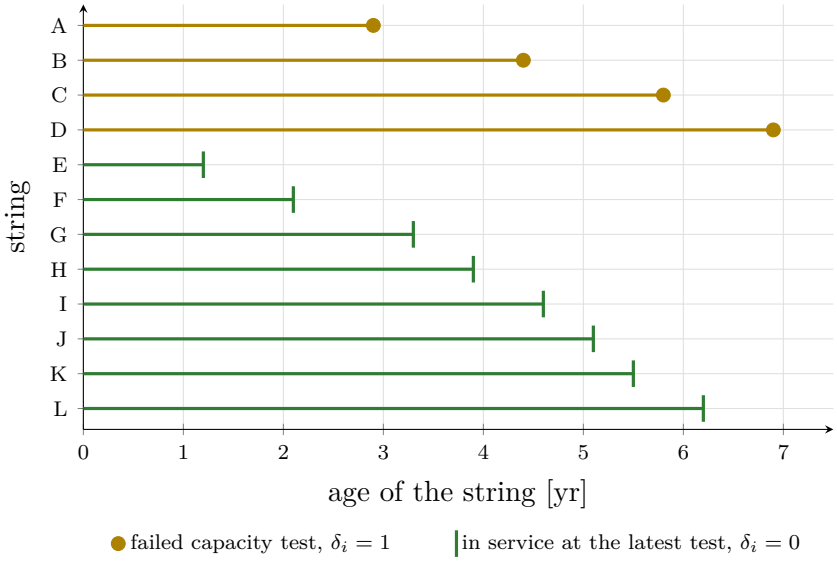


Figure 10.2: The record of the twelve strings of Table 10.1. Each line is one string's observed age; circles mark failed capacity tests and vertical bars strings in service (right censoring).

The summaries of the record. The record has $k = \sum_i \delta_i = 4$ failures. Its total time on test is

$$T_{\text{tot}} = \sum_{i=1}^{12} t_i = (2.9 + 4.4 + 5.8 + 6.9) + (1.2 + 2.1 + 3.3 + 3.9 + 4.6 + 5.1 + 5.5 + 6.2)$$
$$= 20.0 + 31.9 = 51.9\,\text{yr},$$

of which the survivors account for 31.9 yr, more than half. The sum of the log failure ages, which the Weibull likelihood needs, is

$$S_\delta = \sum_i \delta_i \ln t_i = \ln 2.9 + \ln 4.4 + \ln 5.8 + \ln 6.9 = 1.0647 + 1.4816 + 1.7579 + 1.9315 = 6.2357,$$

with the ages in years. Strings are indexed by i when all twelve are meant and by j when only the eight survivors are, so that t_j is a survivor's age and $\delta_j = 0$.

The state. The unknowns that survival analysis writes as a parameter vector $\boldsymbol{\theta}$ are, as in Chapter 7, the two entries of the state of equation (7.21),

$$\mathbf{z} = \begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} \ln \eta \\ \ln \beta \end{bmatrix} \quad \begin{matrix} \text{log-scale} \\ \text{log-shape'} \end{matrix}$$

with η in years inside the logarithm. The graph is the motor graph of Figure 7.2 with the twelve rows drawn out, because the questions here are about individual rows: which survivor fails next, and what the survivors do together. The conditional survival of a string, its remaining life and the number of survivors that fail in the year are derived quantities, functions of the state.

10.1.5 The probabilistic questions (step 1)

The engineer asks six questions of the record, and step 1 writes each as a probability or an expectation over the state, with the unknown on the left of the bar and the record on the right. Table 10.2 lists them: the first column is the question in engineering language; the second the general form that every question of its category takes (Chapter 2), written for a quantity $f(\mathbf{z})$ derived from the state, a level f^* of it, the posterior density $p(\mathbf{z} \mid \mathbf{y})$, a future observation y_{new} and two drawings M_0 and M_1; the third its mathematical form for the strings; the fourth its category; and the last the subsection where it is answered. The horizon is one year, $\tau = 1\,\text{yr}$, the interval between tests; $R(\tau \mid t_0; \mathbf{z})$ is the conditional survival of a string of age t_0 for a known state, equation (10.1); K_f is the number of the eight survivors that fail within the horizon; and M_0 and M_1 are the exponential and the Weibull drawings. The first and fifth questions are predictions of future events, for one string and for the eight together; the second and third ask about the conditional survival as a derived quantity of the state, its band and its tail; the fourth asks for expectations and quantiles of the remaining life; and the sixth rests the replacement order on a comparison of the two drawings. The next annual test is the next measurement for these six. The atlas of Appendix A lists the same forms.

Table 10.2: The questions asked of the strings' record, the general form of their category, their mathematical form for the strings, their category and where each is answered. $\mathbf{y}$ is the record of Table 10.1 with the flat prior in (u, v); $\tau = 1\,\text{yr}$; the oldest string, L, is at $t_0 = 6.2\,\text{yr}$; $\text{MRL}(t_0; \mathbf{z})$ and $q_{1/2}(t_0; \mathbf{z})$ are the mean and the median remaining life for a known state; H is a horizon.

Engineering question	General form	Mathematical form	Category	Answered in
A string has reached its age. What is the probability it lasts until the next test?	$\int p(y_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	$\Pr(T_j > t_j + \tau \mid \mathbf{y})$ for string j of the record; $\Pr(T > t_0 + \tau \mid T > t_0, \mathbf{y})$ for a string outside it	prediction under a threshold	10.3.1
How well is that probability known for the oldest string?	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$, quantiles	$p\big(R(\tau \mid t_0; \mathbf{z}) \mid \mathbf{y}\big)$ and its central 95 % band	estimation of a derived quantity	10.3.2
Is the oldest string's chance of lasting the year at least 0.80?	$\Pr(f(\mathbf{z}) \ge f^* \mid \mathbf{y})$	$\Pr\big(R(\tau \mid t_0; \mathbf{z}) \ge 0.80 \mid \mathbf{y}\big)$	prediction under a threshold	10.3.2
How much life does a string have left?	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$, quantiles	$\mathbb{E}[\min(T_j - t_j, H) \mid \mathbf{y}]$; the predictive median; the bands of $\text{MRL}(t_0; \mathbf{z})$ and $q_{1/2}(t_0; \mathbf{z})$	estimation	10.3.3
Will any of the eight survivors fail before the next test, and how many?	$\int p(y_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	$\Pr(K_f = 0 \mid \mathbf{y})$; $\mathbb{E}[K_f \mid \mathbf{y}]$, $\text{Var}(K_f \mid \mathbf{y})$	prediction under a threshold (joint event)	10.3.4
Which string should be replaced first, and does the answer depend on the hazard shape?	$p(\mathbf{y} \mid M_1)/p(\mathbf{y} \mid M_0)$	the order of the first question's answers; B_{10}, Weibull against exponential	hypothesis comparison	10.3.5

10.1.6 The evidence as rows (step 2)

Step 2 attaches the record to the graph as twelve rows on the one free node, none of them Gaussian. A failed string contributes the negative log failure density at its age, $-\ln \lambda(t_i; \mathbf{z}) + \Lambda(t_i; \mathbf{z})$, and a survivor the negative log survival probability at its age, $\Lambda(t_j; \mathbf{z})$; together they are minus the censored log-likelihood (7.28), and each is also a count row in the working form of Section 4.2.7, with a working residual and a Jacobian row. The prior is flat in u and v, as for the motors, so

that it adds no row and the MAP state is the maximum-likelihood estimate. Two questions add rows of their own. A string of the same design found in service outside the record adds a thirteenth survivor row at its age. The comparison of the two drawings gives the log-shape a proper prior, a Gaussian row v/σ_0 on v alone, whose limit $\sigma_0 \to 0$ is the exponential drawing, the exact constraint $v = 0$ reached as a limit, as Section 7.3.3 did for the motors. The survival of each string of the record is already its row; the questions about the eight survivors add none. Which row acts on which entry is a property of the graph, and it fixes where the Jacobian can be nonzero, marked $*$:

$$\mathbf{J} = \begin{bmatrix} * & * \\ \vdots & \vdots \\ * & * \\ \vdots & \vdots \\ * & * \\ * & * \\ 0 & * \end{bmatrix} \begin{matrix} \text{string A, failure} \\ \vdots \\ \text{string E, survivor} \\ \vdots \\ \text{string L, survivor} \\ \text{a string outside the record} \\ \text{prior row on } v\text{, comparison} \end{matrix}, \qquad \text{columns } u, v. \tag{10.5}$$

Every lifetime row acts on both entries, structurally, because the density and the survival probability at an age depend on the scale and the shape together through $(t/\eta)^\beta$; only the prior row of the comparison acts on one. Every string, failed or in service, therefore bears on the same two entries, which is what makes the survival of one string evidence about every other, and what couples the survivors' coming years. No row is a physics row, so no model tolerance σ_p arises; σ_0 is the standard deviation of a prior row.

The modeling choices. Four choices make the twelve rows, and each buys something and can fail. The twelve strings are treated as comparable, one lifetime distribution shared by all of them: this is what puts every row on one state of two entries, so that the old strings' record informs the young strings' answers, and it fails if strings installed at different dates saw different temperatures or discharge histories (Section 10.6.5). Their lifetimes are taken as independent given the state and their censoring as administrative, set by the calendar of the tests: this makes the likelihood a product of one density or one survival probability per string and each survivor's age a row of the record, and it fails if a string was withdrawn because its test margin was falling (Section 10.6.5). A string found failed at a test is taken to have failed at the age of the test: this gives each failure a density row, and it misplaces the failure, which happened at some time in the year before, by an amount Section 10.4.5 measures. The prior is flat in u and v: this makes the MAP state the maximum-likelihood estimate and keeps the exact posterior to one integral, and with four failures it is not innocent, as Section 10.4.4 shows.

The rows. Each string contributes one lifetime row (7.29), the term $\mathcal{L}_i(\mathbf{z}) = \delta_i[u - v - (\beta - 1)x_i] + \Lambda_i$ with $x_i = \ln t_i - u$ and $\Lambda_i = \mathrm{e}^{\beta x_i}$: the negative log density $-\ln f(t_i;\mathbf{z})$ for A to D and the cumulative hazard $\Lambda(t_i;\mathbf{z})$ for E to L. The objective is the sum of the twelve terms, $\mathcal{L} = \sum_i \mathcal{L}_i$, as in equation (7.30). Each term is a count row with mean Λ_i in the working form of equation (7.33), and the twelve working residuals and Jacobian rows stack, strings A to L, as

$$\mathbf{r}(\mathbf{z}) = \begin{bmatrix} (\Lambda_A - 1)/\sqrt{\Lambda_A} \\ \vdots \\ (\Lambda_D - 1)/\sqrt{\Lambda_D} \\ \sqrt{\Lambda_E} \\ \vdots \\ \sqrt{\Lambda_L} \end{bmatrix}, \qquad \mathbf{J}(\mathbf{z}) = \beta \begin{bmatrix} -\sqrt{\Lambda_A} & x_A\sqrt{\Lambda_A} \\ \vdots & \vdots \\ -\sqrt{\Lambda_D} & x_D\sqrt{\Lambda_D} \\ -\sqrt{\Lambda_E} & x_E\sqrt{\Lambda_E} \\ \vdots & \vdots \\ -\sqrt{\Lambda_L} & x_L\sqrt{\Lambda_L} \end{bmatrix}, \tag{10.6}$$

a 12×2 Jacobian with both entries of every row structurally nonzero, the pattern (10.5); for a survivor the working residual is $\Lambda_j/\sqrt{\Lambda_j} = \sqrt{\Lambda_j}$. The score row of a survivor at age t, its gradient (7.31) with $\delta = 0$, is the gradient $\nabla\Lambda(t;\mathbf{z}) = (-\beta\Lambda(t),\ \beta x(t)\Lambda(t))^{\mathsf{T}}$ of its cumulative hazard, the object from which the gradients of the conditional survival, the remaining life and the test rows of step 7 are built.

10.2 The posterior of the state (step 3)

In the book's notation, $p(\mathbf{z} \mid \mathbf{y})$ is the posterior density of the state, $f(\mathbf{z})$ a quantity derived from the state, f^* a level of it, y_{new} a future observation and $p(\mathbf{y} \mid M)$ the evidence of a drawing M.

All six are read from one posterior: the posterior of the state that the twelve rows of step 2 give. This section derives it and computes it for the strings; Section 10.3 reads every answer from it.

10.2.1 The posterior from the twelve rows

The posterior of the state. The gradient of the objective is $\mathbf{J}^{\mathsf{T}}\mathbf{r} - (0,\ k)^{\mathsf{T}}$, and its Hessian, the sum of the twelve row Hessians (7.36), is $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ plus the named curvature term $\sum_i(\Lambda_i - \delta_i)\,\partial^2 \ln\Lambda_i/\partial\mathbf{z}\,\partial\mathbf{z}^{\mathsf{T}}$ of equation (7.34), which does not vanish because the log-shape is not the canonical link of these counts. The step is Newton's, equation (7.40), from a start near the failure ages, and it stops where both gradient entries vanish. There the cumulative hazards add to the number of failures, $\sum_i \Lambda_i = k$, the curve of best scale (7.35), and the information takes the closed form of equation (7.41),

$$\boldsymbol{\Lambda} = \mathbf{J}^{\mathsf{T}}\mathbf{J} + \begin{bmatrix} 0 & 0 \\ 0 & k \end{bmatrix}, \qquad \boldsymbol{\Sigma} = \boldsymbol{\Lambda}^{-1},$$

whose (u,u) entry is $\hat{\beta}^2 k$: every string takes a share of it in proportion to its cumulative hazard, and the shares add to a precision of the scale set by the number of failures.

The exponential drawing. Under a constant hazard the state is the one number $\omega = \ln\lambda$: the Weibull state with v pinned at zero, where $\lambda = \mathrm{e}^{-u}$, reached as the limit $\sigma_0 \to 0$ of the prior row of the comparison, which Section 7.4.1 followed for the motors. The record is then one count row, k failures against the expected count λT_{tot}, the count row of Section 4.2.7 in the log-rate, with the term, slope and curvature

$$\mathcal{L}(\omega) = \lambda T_{\text{tot}} - k\,\omega, \qquad \mathcal{L}' = \lambda T_{\text{tot}} - k, \qquad \mathcal{L}'' = \lambda T_{\text{tot}},$$

primes being derivatives in ω. At the MAP state $\hat{\lambda} = k/T_{\text{tot}}$ the slope is zero, the information is $\Lambda = \mathcal{L}'' = k$ and the variance is $\sigma_\omega^2 = 1/k$. Under the prior flat in $\ln\lambda$ the exact posterior of the rate is $\text{Gamma}(k, T_{\text{tot}})$.

10.2.2 The exact route through one integral

Under the flat prior in (u,v), Section 7.2.2 showed that the scale integrates out in closed form: given the shape, the quantity $\zeta = \eta^{-\beta}$ has the gamma posterior of equation (7.44),

$$\zeta \mid v, \mathbf{y} \sim \text{Gamma}\big(k,\ A(\beta)\big), \qquad A(\beta) = \sum_{i=1}^{12} t_i^{\beta},$$

and the shape has the marginal $p(v \mid \mathbf{y}) = \Omega(v)/\int\Omega$, with $\Omega(v) \propto \beta^{k-1}\mathrm{e}^{(\beta-1)S_\delta}A(\beta)^{-k}$, equation (7.46). The sum A runs over all twelve strings, failed and surviving, so it already contains the survivor rows.

The average over the scale. By equation (10.4) the conditional survival is $\mathrm{e}^{-\nu\zeta}$, and its average over the gamma distribution of ζ is the Laplace transform of that distribution:

$$\mathbb{E}[\mathrm{e}^{-\nu\zeta} \mid v, \mathbf{y}] = \int_0^\infty \mathrm{e}^{-\nu\zeta}\, \frac{A^k \zeta^{k-1} \mathrm{e}^{-A\zeta}}{\Gamma(k)}\, \mathrm{d}\zeta = \frac{A^k}{\Gamma(k)}\, \frac{\Gamma(k)}{(A+\nu)^k},$$

the integral being that of a gamma density of rate $A+\nu$ up to its constant. Therefore

$$\mathbb{E}[R(\tau \mid t_0; \mathbf{z}) \mid v, \mathbf{y}] = \left(\frac{A(\beta)}{A(\beta)+\nu(t_0)}\right)^k, \tag{10.7}$$

the integrand of equations (10.15) and (10.16). With the sum of the eight increments in place of ν it is the integrand of the fleet, equation (10.31), and with $\nu_i + \nu_j$ that of a pair expectation of equation (10.33).

The distribution given the shape. Given the shape, the conditional survival is a decreasing function of the one variable ζ, and the event $R(\tau \mid t_0; \mathbf{z}) < R^*$ is the upper tail $\zeta > -\ln R^*/\nu$ of its gamma distribution,

$$\Pr\big(R(\tau \mid t_0; \mathbf{z}) < R^* \mid v, \mathbf{y}\big) = \Upsilon_k\!\left(\frac{-A(\beta)\ln R^*}{\nu(t_0)}\right), \qquad \Upsilon_k(\chi) = \mathrm{e}^{-\chi} \sum_{m=0}^{k-1} \frac{\chi^m}{m!},$$

whose average over $p(v \mid \mathbf{y})$ is equation (10.21).

The moments of the state. The conditional mean and variance of u given v, equation (7.45), averaged over the marginal of v, give the exact mean of the state, the variance of u and the covariance of u and v, so that the exact covariance of the state can be set beside $\mathbf{\Sigma}$ and its long axis beside the eigenvector of $\mathbf{\Sigma}$; the gamma upper tail $\Pr(\zeta \geq \mathrm{e}^{-\beta x} \mid v, \mathbf{y})$ averaged over the marginal of v gives $\Pr(u \leq x \mid \mathbf{y})$ and with it the quantiles of the scale.

The draws. The exact posterior is also drawn: v by inverting its cumulative distribution on a grid, and ζ from $\mathrm{Gamma}(k, A(\beta))$ given v. Pushed through $\mathrm{e}^{-\nu\zeta}$, through equations (10.24) and (10.25), and through the product over the eight survivors, the draws carry the quantiles of the remaining life and the distribution of the count K_f, and they check the one-integral values.

Building blocks. The sums S_δ and $A(\beta)$ of the record, the gamma, digamma and trigamma functions, the tail Υ_k, a grid in v and a seeded generator for the draws.

10.2.3 The posterior of the twelve strings

The strings' inputs are the twelve pairs of Table 10.1, with $k = 4$ failures, $T_{\mathrm{tot}} = 51.9\,\mathrm{yr}$ on test and $S_\delta = 6.2357$; a grid of 16 001 points over $v \in [-4, 4]$ for the exact integrals; and 2×10^5 draws of the exact posterior (seed 2026).

The Newton iteration. The start is a scale of 6 yr, near the failure ages, and a shape of 3, $\mathbf{z}_0 = (1.79176,\ 1.09861)^{\mathsf{T}}$, where the gradient and the Hessian of equation (7.40) are

$$\nabla\mathcal{L} = \begin{bmatrix} -7.08513 \\ -2.85150 \end{bmatrix}, \qquad \nabla^2\mathcal{L} = \begin{bmatrix} 57.2554 & -2.1485 \\ -2.1485 & 4.4049 \end{bmatrix}, \qquad \Delta\mathbf{z} = \begin{bmatrix} 0.15080 \\ 0.72090 \end{bmatrix}.$$

Table 10.3 lists the iterates. Each row is an iteration; the columns give the state, the objective, the two entries of the gradient and the step. The first step overshoots the shape, to $\mathrm{e}^{1.81951} = 6.17$, and the second brings it back; from the third the length of the gradient falls quadratically, from 0.78 to 0.10, 0.0021 and 5×10^{-7}, and five steps reach the MAP state to the printed precision.

Table 10.3: Newton iterations (7.40) for the Weibull fit to the twelve strings from $\mathbf{z}_0 = (\ln 6,\ \ln 3)^{\mathsf{T}}$, with the twelve rows (7.29), their gradient and the summed row Hessians (7.36). All entries are dimensionless.

m	u_m	v_m	$\mathcal{L}(\mathbf{z}_m)$	$\partial\mathcal{L}/\partial u$	$\partial\mathcal{L}/\partial v$	Δu	Δv
0	1.79176	1.09861	10.99698	−7.08513	−2.85150	0.15080	0.72090
1	1.94256	1.81951	10.71381	10.54959	3.53253	−0.08048	−0.15733
2	1.86208	1.66218	9.97116	0.03830	0.77437	0.00753	−0.10199
3	1.86961	1.56019	9.93025	−0.09657	0.03789	0.00175	−0.00709
4	1.87135	1.55310	9.93003	−0.00212	0.00026	0.00003	−0.00007
5	1.87138	1.55303	9.93003	0.00000	0.00000		

The MAP state. The iteration ends at

$$\hat{\mathbf{z}} = \begin{bmatrix} \hat{u} \\ \hat{v} \end{bmatrix} = \begin{bmatrix} 1.87138 \\ 1.55303 \end{bmatrix}, \qquad \hat{\eta} = \mathrm{e}^{\hat{u}} = 6.4973\,\mathrm{yr}, \qquad \hat{\beta} = \mathrm{e}^{\hat{v}} = 4.7257,$$

where the cumulative hazards of the twelve strings add to 4.000, the number of failures: the MAP state lies on the curve of best scale (7.35), $\hat{\eta} = (27\,722/4)^{1/4.7257} = 6.4973\,\mathrm{yr}$ with $A(\hat{\beta}) = 27\,722$.

The rows at the MAP state. Table 10.4 evaluates every row at $\hat{\mathbf{z}}$. Each row of the table is a string; the columns give its indicator, its lever $x_i = \ln t_i - \hat{u}$, its cumulative hazard Λ_i, its score row (7.31), and the three distinct entries of its Hessian (7.36). The two score columns sum to zero, the gradient at the MAP state; the column of Λ_i sums to 4, the condition of the best scale; and the three Hessian columns sum to the entries of $\boldsymbol{\Lambda}$.

Table 10.4: The rows (7.29) of the strings' record at the MAP state $\hat{\mathbf{z}} = (1.87138,\ 1.55303)^{\mathsf{T}}$: the lever $x_i = \ln(t_i/\hat{\eta})$, the cumulative hazard $\Lambda_i = (t_i/\hat{\eta})^{\hat{\beta}}$, the score row (7.31) and the row Hessian (7.36). All entries are dimensionless.

String	δ_i	x_i	Λ_i	$\partial\mathcal{L}_i/\partial u$	$\partial\mathcal{L}_i/\partial v$	$\partial^2\mathcal{L}_i/\partial u^2$	$\partial^2\mathcal{L}_i/\partial u\,\partial v$	$\partial^2\mathcal{L}_i/\partial v^2$
A	1	−0.8067	0.0221	4.6213	2.7279	0.4936	5.0195	4.0491
B	1	−0.3898	0.1585	3.9767	0.5500	3.5397	5.3564	2.0878
C	1	−0.1135	0.5848	1.9622	−0.7772	13.0600	3.4448	0.3911
D	1	0.0601	1.3287	−1.5533	−0.9066	29.6732	−3.3378	0.2007
E	0	−1.6891	0.0003	−0.0016	−0.0027	0.0076	0.0113	0.0190
F	0	−1.1294	0.0048	−0.0227	−0.0257	0.1074	0.0986	0.1113
G	0	−0.6775	0.0407	−0.1923	−0.1303	0.9090	0.4234	0.2869
H	0	−0.5104	0.0896	−0.4236	−0.2162	2.0017	0.5981	0.3053
I	0	−0.3453	0.1956	−0.9241	−0.3191	4.3672	0.5840	0.2017
J	0	−0.2421	0.3184	−1.5049	−0.3644	7.1117	0.2172	0.0526
K	0	−0.1666	0.4550	−2.1502	−0.3583	10.1611	−0.4570	−0.0761
L	0	−0.0468	0.8015	−3.7875	−0.1774	17.8986	−2.9492	−0.1381
Sum	4		4.0000	0.0000	0.0000	89.3307	9.0093	7.4912

The information and the covariance. With $k = 4$ added to the (v, v) entry of the Gauss–Newton product of the working rows (10.6), the information is the column sums of Table 10.4:

$$\mathbf{J}^{\mathsf{T}}\mathbf{J} = \begin{bmatrix} 89.331 & 9.009 \\ 9.009 & 3.491 \end{bmatrix}, \quad \boldsymbol{\Lambda} = \begin{bmatrix} 89.331 & 9.009 \\ 9.009 & 7.491 \end{bmatrix}, \quad \boldsymbol{\Sigma} = \boldsymbol{\Lambda}^{-1} = \begin{bmatrix} 0.01274 & -0.01532 \\ -0.01532 & 0.15192 \end{bmatrix}, \tag{10.8}$$

with $\det \boldsymbol{\Lambda} = 588.02$ and $\boldsymbol{\Lambda}_{uu} = \hat{\beta}^2 k = 89.331$. The entries are read off as $\sqrt{\Sigma_{uu}} = 0.113$, $\sqrt{\Sigma_{vv}} = 0.390$ and the correlation $\Sigma_{uv}/\sqrt{\Sigma_{uu}\Sigma_{vv}} = -0.348$. In $\boldsymbol{\Sigma} = \sum_{j=1}^{2} \epsilon_j \, \mathbf{e}_j \mathbf{e}_j^{\mathsf{T}}$ the long direction is $\mathbf{e}_1 = (0.108,\ -0.994)^{\mathsf{T}}$, almost pure log-shape, with $\sqrt{\epsilon_1} = 0.392$, and the short one $\mathbf{e}_2 = (0.994,\ 0.108)^{\mathsf{T}}$, almost pure log-scale, with $\sqrt{\epsilon_2} = 0.105$.

The exponential fit. With $k = 4$ failures in $T_{\text{tot}} = 51.9\,\text{yr}$, $\hat{\lambda} = 4/51.9 = 0.0771\,\text{yr}^{-1}$; the one count row has the curvature $\Lambda = 4$ at the MAP state, so $\sigma_\omega^2 = 0.25$ and the log-rate has $\sigma_\omega = 0.5$. The exact posterior of the rate is $\text{Gamma}(4, 51.9\,\text{yr})$.

The exact posterior. On the grid, the exact marginal of the shape has its mode at $v = 1.383$ ($\beta = 3.99$), its mean at 1.241 and the standard deviation 0.489; the shape has the median 3.64 and the 95 % band [1.15, 7.77], and $\Pr(\beta > 1 \mid \mathbf{y}) = 0.984$, $\Pr(\beta > 2 \mid \mathbf{y}) = 0.872$. By the moments and the gamma tail of Section 10.2.2, the scale has the median 6.93 yr and the band [5.49, 15.5] yr, $\text{sd}(u \mid \mathbf{y}) = 0.323$ and $\text{corr}(u, v \mid \mathbf{y}) = -0.582$; the exact mean of the state is $\mathbb{E}[\mathbf{z} \mid \mathbf{y}] = (2.008,\ 1.241)^{\mathsf{T}}$.

Four objects of the fit carry the answers. The *MAP state* $\hat{\mathbf{z}} = (\hat{u}, \hat{v})^{\mathsf{T}}$ is the minimum of the objective $\mathcal{L} = \sum_i \mathcal{L}_i$ of the twelve rows; the *information matrix* $\boldsymbol{\Lambda}$ is its curvature at the minimum, the sum of the row Hessians, which in the working form is the Gauss–Newton product $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ plus a named curvature term (Section 10.2.1); the *Laplace covariance* $\boldsymbol{\Sigma} = \boldsymbol{\Lambda}^{-1}$, a 2×2 matrix with an off-diagonal entry, makes the posterior approximately normal, $\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$; and a derived quantity f is carried by its *gradient* $\mathbf{g} = \partial f / \partial \mathbf{z}$ at $\hat{\mathbf{z}}$, with the variance $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$. The pair $(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ is what step 4 uses, with $(\hat{\omega}, \sigma_\omega)$ for the exponential drawing: every answer is computed from it. Because the scale integrates out given the shape, each answer also has an exact value beside it, one integral over v or draws of the exact posterior, and the gap between the two is reported with it.

10.3 The answers (step 4)

The questions of Table 10.2 fall into three categories of Chapter 2: prediction under a threshold (the first, third and fifth), estimation (the second and fourth) and hypothesis comparison (the sixth). Each is answered from the posterior of Section 10.2: its engineering question, the general form of its category, the form it takes for the strings, the quantities of the posterior it uses, and the strings' numbers. The questions name the horizon $\tau = 1\,\text{yr}$; the ages of the eight survivors and the ages 0, 2, 4 and 6 yr of a string outside the record; the level 0.80 of the third question; the horizons $H = 15$, 30 and 60 yr of the restricted mean residual life; and the priors $\mathcal{N}(0, 0.5^2)$ and $\mathcal{N}(0, 1^2)$ on v of the comparison.

Every answer is computed from the posterior of the state of step 3, the MAP state $\hat{\mathbf{z}}$ and the covariance $\boldsymbol{\Sigma}$, by a Gauss–Hermite rule over $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ or by the delta method, and that is the value the chapter reports. The Weibull lifetime also admits the one-integral route of Chapter 7: under the prior flat in (u, v) the scale integrates out in closed form given the shape (Section 10.2.2), and every question becomes an exact one-dimensional integral over v, which gives the value set beside it, with the gap. On this record of four failures the gap is large; where it changes a statement of the report, the two values are given and said to fall on opposite sides.

10.3.1 The survival of the coming year

Engineering question. A string has reached its age. What is the probability it lasts until the next test? The engineer asks it for each of the eight strings in service, at its own age, and for a string of the same design found in service at an age t_0 but not part of the record; the next test is one year away.

General form. A prediction of a future event is the probability of the event given the state, averaged over the posterior, equation (2.13) of Chapter 2,

$$p(y_{\text{new}} \mid \mathbf{y}) = \int p(y_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \tag{10.9}$$

here with y_{new} the indicator that the string lasts the year.

Specialized form. The question is equation (10.9) with y_{new} the indicator of the event $T_j > t_j + \tau$ for string j of the record, and $p(\mathbf{z} \mid \mathbf{y})$ the posterior of step 3, the Laplace normal $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ with $\boldsymbol{\Sigma}$ of equation (10.8). When the state is uncertain, the survival of string j to its age t_j is part of the record $\mathbf{y}$: it is the survivor row $R(t_j; \mathbf{z})$ in the likelihood. Write $\mathbf{y}_{-j}$ for the record without that row. Because the rows are independent given the state, the likelihood of $\mathbf{y}$ is that of $\mathbf{y}_{-j}$ times $R(t_j; \mathbf{z})$, and Bayes' theorem gives the posterior with the row as the posterior without it, multiplied by the row and renormalized:

$$p(\mathbf{z} \mid \mathbf{y}) = \frac{R(t_j; \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y}_{-j})}{\int R(t_j; \mathbf{z}')\, p(\mathbf{z}' \mid \mathbf{y}_{-j})\, \mathrm{d}\mathbf{z}'}. \tag{10.10}$$

By the law of total probability over the state,

$$\Pr(T_j > t_j + \tau \mid \mathbf{y}) = \int \Pr(T_j > t_j + \tau \mid \mathbf{z}, \mathbf{y})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}.$$

Given the state, the lifetime of string j is independent of the other strings' rows, and the only part of $\mathbf{y}$ that concerns it is the event $T_j > t_j$; the probability inside the integral is therefore the conditional survival (10.1), and

$$\Pr(T_j > t_j + \tau \mid \mathbf{y}) = \int R(\tau \mid t_j; \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z} = \mathbb{E}\big[R(\tau \mid t_j; \mathbf{z}) \mid \mathbf{y}\big]. \tag{10.11}$$

The conditional predictive survival of a string in the record is the posterior mean of its conditional survival, under the posterior that already contains its row.

The same quantity as a ratio of predictives. Substitute the posterior (10.10) into equation (10.11) and write the conditional survival as the ratio (10.1). The factor $R(t_j; \mathbf{z})$ cancels between the ratio and the posterior, exactly as in the derivation of equation (7.72), and

$$\Pr(T_j > t_j + \tau \mid \mathbf{y}) = \frac{\int R(t_j + \tau; \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y}_{-j})\, \mathrm{d}\mathbf{z}}{\int R(t_j; \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y}_{-j})\, \mathrm{d}\mathbf{z}}, \tag{10.12}$$

the ratio of two posterior predictive survival probabilities, each averaged over the posterior built without string j: equation (7.72), with the record that does not yet contain the survival in the role of the prior knowledge. The two forms are one quantity: equation (10.11) conditions on the survival through the posterior, equation (10.12) through the ratio, and the survival is counted once in each.

A string outside the record. A thirteenth string of the same design, in the same room, found in service at age t_0 but not part of $\mathbf{y}$, has a survival that is new evidence, the thirteenth row of the pattern (10.5), and equation (10.12) applies with $\mathbf{y}$ in place of $\mathbf{y}_{-j}$:

$$\Pr(T_{13} > t_0 + \tau \mid T_{13} > t_0, \mathbf{y}) = \frac{\Pr(T_{13} > t_0 + \tau \mid \mathbf{y})}{\Pr(T_{13} > t_0 \mid \mathbf{y})}, \tag{10.13}$$

where each probability on the right is the posterior predictive survival of a new unit of the population, $\Pr(T_{\text{new}} > t \mid \mathbf{y})$ of equation (7.47), because before its age is known string 13 is such a unit. Applying equation (10.13) to a string already in the record would condition on its survival twice: once in its row and once in the ratio.

The Weibull strings. For the Weibull lifetime the derived quantity is $f(\mathbf{z}) = R(\tau \mid t_j; \mathbf{z}) = \exp[-\zeta\,\nu(t_j)]$ of equation (10.4), bounded in $(0, 1)$ and strongly curved in the state. Under the Laplace posterior the conditional predictive (10.11) is the average of f over $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$, a two-dimensional Gauss–Hermite quadrature on the axes $\mathbf{e}_1$, $\mathbf{e}_2$ of $\boldsymbol{\Sigma}$,

$$\Pr(T_j > t_j + \tau \mid \mathbf{y}) \approx \sum_m \sum_n \varpi_m \varpi_n \, R\big(\tau \mid t_j;\ \hat{\mathbf{z}} + \sqrt{\epsilon_1}\, x_m \mathbf{e}_1 + \sqrt{\epsilon_2}\, x_n \mathbf{e}_2\big), \tag{10.14}$$

with nodes x_m and weights ϖ_m of the rule for the standard normal. For the string outside the record each predictive of the ratio (10.13) is the same sum, with $R(t_0 + \tau; \mathbf{z})$ and $R(t_0; \mathbf{z})$ in place of the conditional survival. The plug-in value $R(\tau \mid t_j; \hat{\mathbf{z}})$ is f at the MAP state, the first reading of the record, and it is not the predictive.

The values beside them come from the exact posterior. The average of $\mathrm{e}^{-\zeta\nu}$ over the gamma posterior of ζ given the shape has the closed form $(A/(A+\nu))^k$ with $A(\beta) = \sum_i t_i^\beta$ (Section 10.2.2), and the conditional predictive of a string of the record is one integral over the shape,

$$\Pr(T_j > t_j + \tau \mid \mathbf{y}) = \int p(v \mid \mathbf{y}) \left(\frac{A(\beta)}{A(\beta) + (t_j + \tau)^\beta - t_j^\beta} \right)^k \mathrm{d}v, \tag{10.15}$$

with $p(v \mid \mathbf{y})$ the exact marginal of the shape, equation (7.46). For the string outside the record the two predictives of equation (10.13) are each of the form of equation (7.47), with $t_0 + \tau$ and t_0 in place of t:

$$\Pr(T_{13} > t_0 + \tau \mid T_{13} > t_0, \mathbf{y}) = \frac{\int p(v \mid \mathbf{y}) \left(A/(A + (t_0 + \tau)^\beta)\right)^k \mathrm{d}v}{\int p(v \mid \mathbf{y}) \left(A/(A + t_0^\beta)\right)^k \mathrm{d}v}. \tag{10.16}$$

The difference between the two lies in A: the record's survivor already has t_j^β inside A, while the thirteenth string's survival enters through the denominator.

The plug-in against the predictive. The state route averages over a normal centered at the joint MAP shape, the plug-in's own, and is narrow in the shape, so it stays close to the plug-in. Two effects separate the exact predictive from the plug-in. The plug-in reads the conditional survival at the joint MAP shape. The marginal of the shape is centered lower, because integrating the scale out weights each shape by $1/\beta$ (Section 7.2.2), and a shallower shape spreads the failures over a wider range of ages, raising the hazard of the coming year for a young string and lowering it for an old one; this shift of the shape's center is the larger effect. The posterior average also averages a curved function of the state across its spread, the effect Chapter 9 derived for an uncertain rate, and that averaging can move the answer either way: over the scale it raises an old string's survival, over the shape it can lower it. Section 10.3.1 splits the gap between the two readings into these parts.

The exponential drawing. Under a constant hazard the conditional survival is $\mathrm{e}^{-\lambda\tau}$ at every age. On the state, $\omega \mid \mathbf{y} \approx \mathcal{N}(\hat{\omega}, \sigma_\omega^2)$ of Section 10.2.1, the predictive of a string of the record is the one-dimensional sum of $\mathrm{e}^{-\tau \exp(\omega)}$ over the nodes, the same for every string; a string outside the record adds its exposure t_0 to T_{tot} in $\hat{\omega}$. The values beside them come from the exact posterior. With the exact posterior $\mathrm{Gamma}(k, T_{\mathrm{tot}})$ of the rate under the prior flat in $\ln\lambda$ (Section 10.2.1), the predictive of a string of the record is $\mathbb{E}[\mathrm{e}^{-\lambda\tau} \mid \mathbf{y}] = (T_{\mathrm{tot}}/(T_{\mathrm{tot}} + \tau))^k$; a string outside the record has the predictive $((T_{\mathrm{tot}} + t_0)/(T_{\mathrm{tot}} + t_0 + \tau))^k$.

Building blocks. The MAP state $\hat{\mathbf{z}}$ and the covariance $\boldsymbol{\Sigma}$ of equation (10.8), with its axes and spreads (Section 10.2.3); the five nodes of Table 10.5, and sixty per axis for the reported values; $\hat{\omega}$ and σ_ω for the exponential drawing; and, for the exact values beside them, the exact

marginal $p(v \mid \mathbf{y})$, the sums $A(\beta)$ and the average over the scale given the shape (Section 10.2.2), with k and T_{tot} for the exponential drawing.

The plug-in. At $\hat{\mathbf{z}}$, $\hat{\eta}^{\hat{\beta}} = 6930.5\,\text{yr}^{\hat{\beta}}$. For a string of six years, equation (10.4) gives $\nu = 7^{\hat{\beta}} - 6^{\hat{\beta}} = 9856.4 - 4757.1 = 5099.3$, a cumulative hazard of the coming year of $5099.3/6930.5 = 0.7358$ and

$$R(1 \mid 6; \hat{\mathbf{z}}) = \mathrm{e}^{-0.7358} = 0.479;$$

at two years, $\nu = 3^{\hat{\beta}} - 2^{\hat{\beta}} = 179.8 - 26.5 = 153.3$, a cumulative hazard of 0.0221 and $R(1 \mid 2; \hat{\mathbf{z}}) = 0.978$. The fitted hazard is $0.009\,\text{yr}^{-1}$ at two years and $0.541\,\text{yr}^{-1}$ at six, a factor of $3^{3.726} = 60$ for a factor of three in age.

The rule of five nodes. Table 10.5 evaluates the sum (10.14) for string L with five nodes on each axis. The rows are the nodes x_m along the long axis $\mathbf{e}_1 = (0.108,\ -0.994)^{\mathsf{T}}$, almost pure log-shape, $\sqrt{\epsilon_1} = 0.392$; the columns the nodes x_n along the short axis $\mathbf{e}_2 = (0.994,\ 0.108)^{\mathsf{T}}$, almost pure log-scale, $\sqrt{\epsilon_2} = 0.105$; each cell is the one-year survival of string L at that state. Down the middle column the shape falls from 14.38 to 1.55, and the survival rises from zero to 0.82.

Table 10.5: The Gauss–Hermite rule (10.14) with five nodes on each axis of $\boldsymbol{\Sigma}$, for string L: $R(1 \mid 6.2; \mathbf{z})$ at $\mathbf{z} = \hat{\mathbf{z}} + 0.392\,x_m\mathbf{e}_1 + 0.105\,x_n\mathbf{e}_2$. The weights are $\varpi = 0.01126$, 0.22208, 0.53333, 0.22208, 0.01126 on each axis. The weighted sum of the twenty-five cells is 0.4248.

x_m	β	$x_n = -2.857$	-1.356	0	1.356	2.857
-2.857	14.38	0.0000	0.0000	0.0000	0.0529	0.7467
-1.356	8.01	0.0000	0.0005	0.0804	0.4417	0.7984
0	4.73	0.0446	0.2093	0.4390	0.6541	0.8197
1.356	2.79	0.4318	0.5699	0.6794	0.7691	0.8449
2.857	1.55	0.7358	0.7813	0.8176	0.8495	0.8797

The survival by age. The sum over the table is 0.4248, and sixty nodes per axis give 0.4247: string L lasts the year with probability 0.425. The middle cell is the plug-in 0.439; the sum lies 0.014 below it, because the survival falls off faster on the steep-shape side than it rises on the shallow side. Table 10.6 gives the one-year survival at ages 0, 2, 4 and 6 years, ages at which the record has no survivor. Each row is an age; the columns give the plug-in value; the posterior average $\mathbb{E}[R(1 \mid t_0; \mathbf{z}) \mid \mathbf{y}]$ under $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$, by sixty nodes per axis, and under the exact posterior, the integral (10.15) at t_0; the conditional predictive of a string of age t_0 that is not in the record, the ratio (10.13) on the state with the exact value (10.16) in parentheses; the exponential predictive of a string of the record, likewise; and the plug-in value, the state-route average and the exact average over two years. The posterior average is the predictive answer to the first question only at the eight survivors' own ages, where the string's survival is a row; for a generic string of age t_0 the answer is the column for a string outside the record.

One row of the arithmetic. At $t_0 = 6\,\text{yr}$ the plug-in is 0.479 and the state-route average 0.457; the exact average, the value beside it, is 0.618, a gap of 0.161. Averaging over the scale alone at the fitted shape, equation (10.7) with $A(\hat{\beta}) = 27\,722$ and $\nu = 5099$, gives $(27\,722/32\,821)^4 = (1 + 0.18394)^{-4} = 0.509$. The same average over the scale alone at the median of the exact marginal of the shape, $\beta = 3.64$, gives 0.627. Averaging over the shape as well, equation (10.15), gives 0.618. Of the difference of 0.139 between the plug-in and the exact posterior average, 0.030 comes from averaging over the scale at the fitted shape and 0.118 from

Table 10.6: One-year survival by age: plug-in at $\hat{\eta} = 6.497\,\text{yr}$, $\hat{\beta} = 4.726$; posterior average $\mathbb{E}[R(1 \mid t_0; \mathbf{z}) \mid \mathbf{y}]$ under $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ and exact; conditional predictive of a string of age t_0 not in the record, on the state with equation (10.16) in parentheses; exponential predictive of a string of the record, on the state with $(51.9/52.9)^4$ in parentheses. Last column: two years, plug-in, state and exact posterior average. All entries are probabilities [–].

Age t_0 [yr]	Plug-in	Average, state	Average, exact	Predictive, string not in the record	Exponential, predictive	Two years, plug-in / state / exact
0	1.000	0.998	0.993	0.998 (0.993)	0.917 (0.927)	0.996 / 0.989 / 0.975
2	0.978	0.969	0.959	0.970 (0.960)	0.917 (0.927)	0.907 / 0.893 / 0.883
4	0.828	0.815	0.850	0.820 (0.854)	0.917 (0.927)	0.557 / 0.538 / 0.641
6	0.479	0.457	0.618	0.521 (0.662)	0.917 (0.927)	0.137 / 0.197 / 0.371

moving the shape from the joint MAP to the center of its marginal; the spread of the shape about that center then takes off 0.009. Over two years at age six the plug-in gives 0.137, the state route 0.197 and the exact average 0.371.

Figure 10.3 draws the one-year conditional survival against the age reached. The horizontal axis is the age t_0 in years and the vertical axis $R(1 \mid t_0)$. The three thinner curves are plug-in values at the fitted scale $\hat{\eta} = 6.497\,\text{yr}$ for three shapes: $\beta = 0.7$, a falling hazard; $\beta = 1$, constant; and the fitted $\hat{\beta} = 4.726$. The heavy red curve is the exact posterior average, equation (10.15), with the eight survivors marked on it at their ages, where it is their conditional predictive; the dotted blue curve is the same average on the state route; the dotted gray line is the exponential plug-in. The curve for $\beta = 0.7$ rises with age, the curve for $\beta = 1$ is flat, and the fitted curve falls from one to 0.29 at seven years, the three cases of equation (10.3). The exact curve falls too, but less steeply: it lies below the plug-in up to 3.36 years, just past string G, and above it after. The state-route curve stays just below the plug-in at every age of the record and crosses above it only past six years.

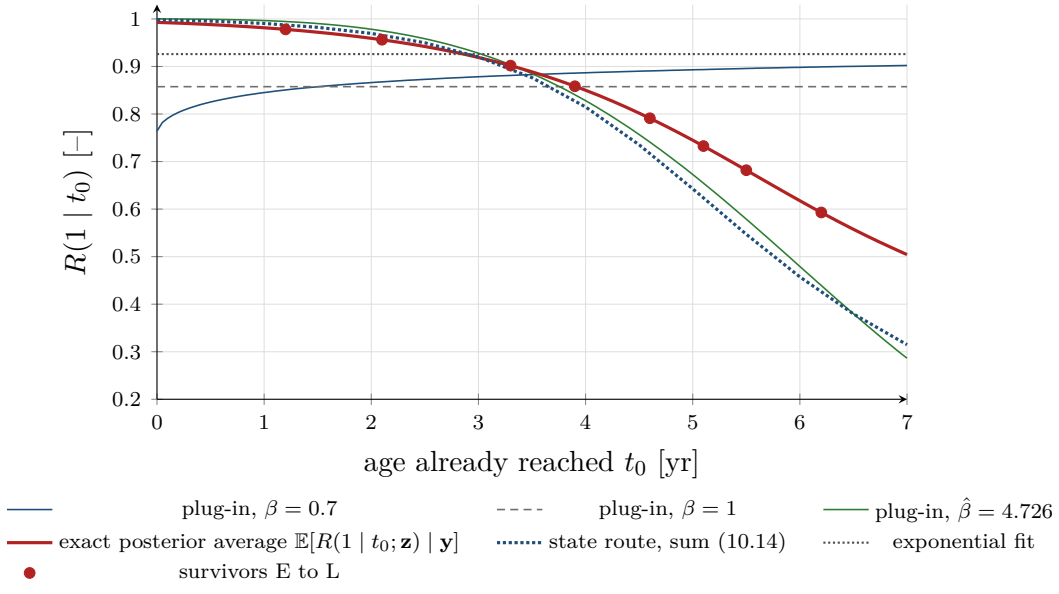


Figure 10.3: One-year conditional survival against the age reached. Thinner curves: $\exp\{-[((t_0+1)/6.497)^\beta - (t_0/6.497)^\beta]\}$ for $\beta = 0.7$, 1 and 4.726, equation (10.4); heavy red: the exact posterior average (10.15) under the flat prior, with the eight survivors marked at their ages, where it is their conditional predictive; dotted blue: the same average on the state, equation (10.14); dotted: the exponential plug-in, $e^{-4/51.9} = 0.926$.

A string outside the record. The predictive for a string of age t_0 not in the record is 0.521 at six years on the state route against the posterior average 0.457; the exact values are 0.662 and 0.618, and the state route is higher than its posterior average at every age past two years, as the exact one is. For string L, whose survival is already a row, the answer is 0.425 against 0.501 for a string of the same age outside the record (exactly 0.593 against 0.650).

The exponential drawing. Every string of the record has the plug-in $\mathrm{e}^{-0.0771} = 0.926$ and, on the state route, the predictive 0.917, whatever its age, by the one-dimensional sum over $\mathcal{N}(\ln 0.0771,\ 0.5^2)$; a string outside the record at six years has 0.925. Exactly, they are $(51.9/52.9)^4 = 0.927$ and $\big((51.9+6)/(51.9+7)\big)^4 = 0.934$. The exponential fit and the exact Weibull average agree near three years and disagree by a factor of five in the failure probability at six, 0.073 by the predictive against 0.382; on the state route the factor is about six, 0.083 against 0.543.

The eight survivors. Table 10.7 lists the eight survivors with their one-year survival, the answer to the first question for the strings of the record. Each row is a survivor; the columns give its age, the plug-in conditional survival, the predictive on the state, equation (10.14), the corresponding failure probability F_j, the exact conditional predictive (10.15) and its failure probability, and the exponential predictive. The failure probabilities rise with age from 0.012 for E to 0.575 for L (0.022 to 0.407 exactly). String L is the case where the gap changes a statement: on the state route it is more likely than not to fail the year, 0.575, and on the exact posterior more likely than not to last it, 0.407; the two fall on opposite sides of one half. The state route errs because a normal in (u, v), centered at the joint MAP shape and narrow in it, gives too little weight to the shallow shapes under which an old string has more life.

Table 10.7: The eight survivors over the coming year: plug-in conditional survival at $\hat{\eta} = 6.497\,\mathrm{yr}$, $\hat{\beta} = 4.726$; predictive on the state, equation (10.14), and its complement; exact conditional predictive, equation (10.15), and its complement; exponential predictive on the state and $(51.9/52.9)^4$. The last row sums the failure probabilities. Probabilities are dimensionless.

String	Age t_j [yr]	Plug-in [–]	State predictive [–]	State F_j [–]	Exact predictive [–]	Exact F_j [–]	Exponential [–]
E	1.2	0.994	0.988	0.012	0.978	0.022	0.917 (0.927)
F	2.1	0.975	0.966	0.034	0.956	0.044	0.917 (0.927)
G	3.3	0.903	0.896	0.104	0.902	0.098	0.917 (0.927)
H	3.9	0.840	0.829	0.171	0.858	0.142	0.917 (0.927)
I	4.6	0.741	0.718	0.282	0.791	0.209	0.917 (0.927)
J	5.1	0.655	0.623	0.377	0.733	0.267	0.917 (0.927)
K	5.5	0.579	0.547	0.453	0.682	0.318	0.917 (0.927)
L	6.2	0.439	0.425	0.575	0.593	0.407	0.917 (0.927)
Sum				2.009		1.507	

10.3.2 The band and the threshold of the oldest string

Engineering questions. How well is that probability known for the oldest string? Is the oldest string's chance of lasting the year at least 0.80? The first asks for the posterior of the one-year survival of string L, at $t_0 = 6.2\,\mathrm{yr}$, as a property of the population at that age, with its band; the second whether that survival is at least the 0.80 the engineer would accept for a string kept in service another year.

General forms. An estimation question about a derived quantity asks for its posterior distribution, with the central band of equation (2.11), and a threshold question for the tail at the required level, equations (2.9) and (2.12):

$$\begin{aligned} \Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y}) &= \int \mathbf{1}\big[f(\mathbf{z}) \le f^*\big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \\ \Pr(f(\mathbf{z}) \ge f^* \mid \mathbf{y}) &= \int \mathbf{1}\big[f(\mathbf{z}) \ge f^*\big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}. \end{aligned} \tag{10.17}$$

Specialized form. The two questions are equation (10.17) with $f(\mathbf{z}) = R(\tau \mid t_0; \mathbf{z})$ at $t_0 = 6.2\,\mathrm{yr}$, $f^* = 0.80$ for the threshold, and $p(\mathbf{z} \mid \mathbf{y})$ the posterior of step 3, $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of equation (10.8). Two choices enter here. The requirement of 0.80 is read as a statement about the population's one-year survival at the string's age, the derived quantity $R(\tau \mid t_0; \mathbf{z})$, not about the predictive: this is what gives the question a band and a tail, and it is not the number the year's replacement decision uses, which is the first question's (Section 10.5.2). The Laplace route linearizes on the double logarithm of the conditional survival rather than on the survival itself: this keeps its band inside $(0, 1)$ from one gradient and one covariance, and it cannot repair a Gaussian in (u, v) that is itself poor, which Section 10.4.2 measures.

The double logarithm. The Laplace route reads the variance of a derived quantity from $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$. By equation (10.2), $-\ln R(\tau \mid t_0; \mathbf{z})$ is the cumulative hazard of the coming period, the difference of two survivor rows, one at the end of the horizon and one at the present age, and its logarithm, the quantity the delta method is applied to, is unbounded:

$$\vartheta(\mathbf{z}) = \ln\big[-\ln R(\tau \mid t_0; \mathbf{z})\big] = \ln\big[\Lambda(t_0 + \tau; \mathbf{z}) - \Lambda(t_0; \mathbf{z})\big], \tag{10.18}$$

which for the Weibull model is $\ln \nu(t_0; \beta) - \beta u$ with $\beta = \mathrm{e}^{v}$, by equation (10.4).

The gradient and the lever. The gradient of a survivor row in the state is its score row (7.31) with $\delta = 0$, $\nabla\Lambda(t; \mathbf{z}) = (-\beta\Lambda(t),\ \beta x(t)\Lambda(t))^{\mathsf{T}}$ with $x(t) = \ln t - u$. Differentiating equation (10.18) gives the gradient of the double logarithm as the difference of the two survivor rows, divided by the cumulative hazard of the period:

$$\mathbf{g} = \frac{\nabla\Lambda(t_0 + \tau) - \nabla\Lambda(t_0)}{\Lambda(t_0 + \tau) - \Lambda(t_0)} = \begin{bmatrix} -\beta \\ \beta\, \bar{x}(t_0) \end{bmatrix}, \qquad \bar{x}(t_0) = \frac{x(t_0 + \tau)\Lambda(t_0 + \tau) - x(t_0)\Lambda(t_0)}{\Lambda(t_0 + \tau) - \Lambda(t_0)}, \tag{10.19}$$

evaluated at $\hat{\mathbf{z}}$. The first entry is $-\beta$ at every age, because the first entry of each survivor row is $-\beta$ times its cumulative hazard. The second carries the lever $\bar{x}$. Since $\mathrm{d}(x\Lambda)/\mathrm{d}t = \lambda(t)\,(x(t) + 1/\beta)$, the numerator of $\bar{x}$ is the integral of $\lambda\,(x + 1/\beta)$ over the coming period and the denominator the integral of λ, so

$$\bar{x}(t_0) = \frac{\int_{t_0}^{t_0+\tau} \lambda(t)\, \ln(t/\eta)\, \mathrm{d}t}{\int_{t_0}^{t_0+\tau} \lambda(t)\, \mathrm{d}t} + \frac{1}{\beta}\,:$$

the log-age relative to the scale averaged over the coming period with the hazard as weight, plus $1/\beta$. The added term is the prefactor β of the hazard, $\ln\lambda = \ln\beta - \beta u + (\beta - 1)\ln t$, whose derivative in v is $1 + \beta x$, not βx. The second entry therefore vanishes when the coming year is centered near $\eta\,\mathrm{e}^{-1/\beta}$, where $\ln(t/\eta) = -1/\beta$, and not where the horizon passes the characteristic life η, as the gradient of the mission reliability does (Section 7.3.2). For a younger string a steeper shape lowers the coming year's hazard, for an older one it raises it.

The variance as a function of the lever. With $\mathbf{g}$ of equation (10.19), the variance of the double logarithm is

$$\sigma_\vartheta^2 = \mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g} = \beta^2 \begin{bmatrix} -1 & \bar{x} \end{bmatrix} \boldsymbol{\Sigma} \begin{bmatrix} -1 \\ \bar{x} \end{bmatrix},$$

a quadratic in the lever with its minimum where the v entry of $\boldsymbol{\Sigma}\mathbf{g}$ vanishes, at $\bar{x}^* = \Sigma_{uv}/\Sigma_{vv}$. There the variance is β^2 times the variance of u with v held fixed, $\sigma_\vartheta^2 = \beta^2/\boldsymbol{\Lambda}_{uu}$; at the MAP state $\boldsymbol{\Lambda}_{uu} = \hat{\beta}^2 k$, by equation (7.41), so the smallest standard deviation the double logarithm can have is $1/\sqrt{k}$, set by the number of failures alone. A string whose coming year sits on either side of $\bar{x}^*$, young or old, has a wider band, and the band widens with the distance of its lever from $\bar{x}^*$.

The band and the threshold. The delta method of Chapter 4 puts ϑ approximately normal with mean $\vartheta(\hat{\mathbf{z}})$ and variance $\sigma_\vartheta^2 = \mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$. The conditional survival $\exp(-\mathrm{e}^\vartheta)$ is decreasing in ϑ, so the 95 % band and the threshold probability map back as

$$\begin{aligned} &\Big[\exp\big(-\mathrm{e}^{\vartheta(\hat{\mathbf{z}})+1.96\sigma_\vartheta}\big),\ \exp\big(-\mathrm{e}^{\vartheta(\hat{\mathbf{z}})-1.96\sigma_\vartheta}\big)\Big], \\ &\Pr\big(R(\tau \mid t_0; \mathbf{z}) \geq R^* \mid \mathbf{y}\big) \approx \Phi\Big(\frac{\ln(-\ln R^*) - \vartheta(\hat{\mathbf{z}})}{\sigma_\vartheta}\Big). \end{aligned} \tag{10.20}$$

The center of the band is the plug-in value $R(\tau \mid t_0; \hat{\mathbf{z}})$, a median of the approximation, not the conditional predictive.

The values beside them come from the exact posterior, where the conditional survival has a closed form given the shape. For a level $R^* \in (0, 1)$, the event $R(\tau \mid t_0; \mathbf{z}) < R^*$ is the event $\nu\zeta > -\ln R^*$, that is $\zeta > -\ln R^*/\nu$, a tail of the gamma posterior of ζ given the shape (Section 10.2.2); with Υ_k the upper tail of a unit-rate gamma variable of integer shape k, as in equation (7.48),

$$\Pr\big(R(\tau \mid t_0; \mathbf{z}) < R^* \mid \mathbf{y}\big) = \int p(v \mid \mathbf{y})\, \Upsilon_k\Big(\frac{-A(\beta)\ln R^*}{\nu(t_0)}\Big)\, \mathrm{d}v. \tag{10.21}$$

Its quantiles give the band, and one minus its value at $R^* = 0.80$ the threshold probability.

Building blocks. The survivor-row gradients $\nabla\Lambda(t; \hat{\mathbf{z}})$ and $\boldsymbol{\Sigma}$ (Section 10.2.3), with Φ; and, for the exact values beside them, the exact marginal $p(v \mid \mathbf{y})$, $A(\beta)$, $\nu(t_0)$ and the tail Υ_k (Section 10.2.2).

The band by age. The gradient (10.19) of the double logarithm, which step 7 also uses, has the lever $\bar{x} = -0.703$ at two years, -0.143 at four, 0.218 at six and 0.248 for string L, where $\mathbf{g}_L = (-4.726,\ 1.173)^\mathsf{T}$. For the six-year string the hazard-weighted mean of $\ln(t/\hat{\eta})$ over the coming year is 0.007 and $1/\hat{\beta} = 0.212$. The second entry changes sign at $t_0 = 4.71$ years, when the coming year is centered near $\hat{\eta}\,\mathrm{e}^{-1/\hat{\beta}} = 5.26$ years, below the characteristic life $\hat{\eta} = 6.50$ years.

The state route. Table 10.8 gives the delta method on the double logarithm, equation (10.20), beside the exact posterior of Table 10.9. Each row is an age; the columns give the standard deviation σ_ϑ, the Laplace band and the Laplace probability of a survival below 0.80, then the exact band and probability.

One row of the Laplace route, at $t_0 = 6\,\mathrm{yr}$: the two survivor rows of equation (10.18) are at 7 and 6 years, with cumulative hazards 1.4222 and 0.6864 at $\hat{\mathbf{z}}$, so $\vartheta(\hat{\mathbf{z}}) = \ln 0.7358 = -0.307$. Their gradients and the gradient (10.19) are

$$\nabla\Lambda(7) = \begin{bmatrix} -6.7209 \\ 0.5009 \end{bmatrix}, \quad \nabla\Lambda(6) = \begin{bmatrix} -3.2438 \\ -0.2583 \end{bmatrix}, \quad \mathbf{g} = \frac{1}{0.7358}\begin{bmatrix} -3.4771 \\ 0.7592 \end{bmatrix} = \begin{bmatrix} -4.726 \\ 1.032 \end{bmatrix},$$

Table 10.8: The one-year conditional survival by the delta method on $\vartheta = \ln[-\ln R]$ with $\mathbf{\Sigma}$ of equation (10.8), centered at the plug-in value, against the exact posterior of Table 10.9. String L is at 6.2 yr. All entries are dimensionless.

	Laplace on ϑ			Exact posterior	
Age t_0 [yr]	σ_ϑ	95 % band	$\Pr(R < 0.8)$	95 % band	$\Pr(R < 0.8)$
2	1.216	[0.787, 0.998]	0.029	[0.877, 0.998]	0.001
4	0.506	[0.601, 0.932]	0.371	[0.683, 0.961]	0.227
6	0.772	[0.035, 0.850]	0.939	[0.143, 0.943]	0.761
L, 6.2	0.815	[0.017, 0.846]	0.945	[0.099, 0.942]	0.778

with the lever $\bar{x} = 1.032/4.726 = 0.218$. With $\mathbf{\Sigma}$ of equation (10.8),

$$\mathbf{\Sigma g} = \begin{bmatrix} -0.0760 \\ 0.2292 \end{bmatrix}, \qquad \sigma_\vartheta^2 = \mathbf{g}^\mathsf{T}\mathbf{\Sigma g} = 0.5957, \qquad \sigma_\vartheta = 0.772,$$

and the band -0.307 ± 1.513 maps through $\exp(-\mathrm{e}^{\vartheta})$ to [0.035, 0.850]. Split by entries, the product has the scale term $g_u^2\Sigma_{uu} = 0.2845$, the shape term $g_v^2\Sigma_{vv} = 0.1617$ and the cross term $2g_ug_v\Sigma_{uv} = 0.1494$, positive because the two gradient entries and the correlation both have opposite signs.

The column of σ_ϑ is the quadratic in the lever of Section 10.3.2. With $\mathbf{\Sigma}$ of equation (10.8) its minimum is at $\bar{x}^* = \Sigma_{uv}/\Sigma_{vv} = -0.101$, the lever of a string of 4.20 years, where $\sigma_\vartheta = 1/\sqrt{4} = 0.500$. A string of four years has $\bar{x} = -0.143$, close to it, and $\sigma_\vartheta = 0.506$. The older strings move away on one side, to $\bar{x} = 0.218$ at six years and 0.248 for string L, whose σ_ϑ are 0.772 and 0.815; a string of two years sits far on the other side, at $\bar{x} = -0.703$, and its σ_ϑ is 1.216, although its survival, close to one, still has a narrow band in R. Past 4.2 years every added year of age moves the lever further from $\bar{x}^*$ and gives the shape entry of $\mathbf{g}$ more weight against the larger Σ_{vv}, and the band of an old string's coming year widens with its age.

For string L the band is [0.017, 0.846], with its center the plug-in 0.439: the record leaves the survival almost undetermined. The probability that it is at least 0.80 is $1 - 0.945 = 0.055$: the requirement of 0.80 is not shown.

The exact values beside them. Table 10.9 gives the exact posterior of the one-year conditional survival, equation (10.21), at three ages and for string L. Each row is an age; the columns give the median, the mean (the posterior average of Table 10.6, which for string L is its conditional predictive), the central 95 % band and the probability that the survival is below 0.80.

Table 10.9: The exact posterior of the one-year conditional survival $R(1 \mid t_0; \mathbf{z})$: the distribution function (10.21) and its quantiles. String L is at 6.2 yr. All entries are probabilities [–].

Age t_0 [yr]	Median	Mean	95 % band	$\Pr(R < 0.80 \mid \mathbf{y})$
2	0.967	0.959	[0.877, 0.998]	0.001
4	0.860	0.850	[0.683, 0.961]	0.227
6	0.650	0.618	[0.143, 0.943]	0.761
L, 6.2	0.626	0.593	[0.099, 0.942]	0.778

String L, exactly. The exact posterior of the oldest string's one-year survival has the median 0.626 and the band [0.099, 0.942]: the record leaves it almost undetermined between one chance in ten and near certainty. The probability that it is at least 0.80 is $1 - 0.778 = 0.222$, odds of about two to seven against. That is 0.167 above the state route, whose band lies below the

exact one by 0.082 at its lower end and 0.096 at its upper end; both probabilities leave the requirement of 0.80 far from shown, and the decision is the same on both routes.

10.3.3 The remaining life

Engineering question. How much life does a string have left? The engineer wants, for each string, the expected remaining life, which plans a replacement program, and the time by which the string is as likely as not to have failed, each with its band.

General forms. An expectation of a derived quantity and a quantile of a distribution under the posterior, equations (2.10) and (2.9):

$$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}] = \int f(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \qquad \Pr(f(\mathbf{z}) \le s_q \mid \mathbf{y}) = q. \tag{10.22}$$

Specialized form. The question is equation (10.22) with $f(\mathbf{z})$ the mean residual life $\mathrm{MRL}(t_0; \mathbf{z})$ or the remaining-life quantile $q_p(t_0; \mathbf{z})$ derived below, and $p(\mathbf{z} \mid \mathbf{y})$ the posterior of step 3, $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$. One choice enters here: the expected remaining life is reported to a horizon H, the remaining life counted up to H and no further. This gives a finite number in the units of the question, and the number depends on the horizon, which Section 10.3.3 gives at three horizons.

The tail-integral identity. For a nonnegative random variable X with a finite mean, the integration by parts that gave equation (7.10) gives $\mathbb{E}[X] = \int_0^\infty \Pr(X > s)\, \mathrm{d}s$. Apply it to the remaining life $X = T - t_0$ under the distribution conditional on $T > t_0$, for which $\Pr(T - t_0 > s \mid T > t_0, \mathbf{z}) = R(s \mid t_0; \mathbf{z})$:

$$\mathrm{MRL}(t_0; \mathbf{z}) = \mathbb{E}[T - t_0 \mid T > t_0, \mathbf{z}] = \int_0^\infty R(s \mid t_0; \mathbf{z})\, \mathrm{d}s = \frac{1}{R(t_0; \mathbf{z})} \int_{t_0}^\infty R(t; \mathbf{z})\, \mathrm{d}t, \tag{10.23}$$

the area under the survival curve beyond the present age, divided by the survival to that age. At $t_0 = 0$ it is the mean lifetime. For the exponential lifetime it is $1/\lambda$ at every age.

The Weibull mean residual life. Substitute $\chi = (t/\eta)^\beta$ in the last integral of equation (10.23), as for equation (7.17), so that $\mathrm{d}t = (\eta/\beta)\, \chi^{1/\beta - 1}\, \mathrm{d}\chi$ and the lower limit becomes the cumulative hazard $\Lambda(t_0) = (t_0/\eta)^\beta$; with $1/R(t_0; \mathbf{z}) = \mathrm{e}^{\Lambda(t_0)}$,

$$\mathrm{MRL}(t_0; \mathbf{z}) = \frac{\eta}{\beta}\, \mathrm{e}^{\Lambda(t_0)}\, \Gamma\!\left(\frac{1}{\beta},\ \Lambda(t_0)\right), \tag{10.24}$$

where the upper incomplete gamma function is $\Gamma(a, c) = \int_c^\infty \chi^{a-1} \mathrm{e}^{-\chi}\, \mathrm{d}\chi$. At $t_0 = 0$ the mean residual life is $(\eta/\beta)\, \Gamma(1/\beta) = \eta\, \Gamma(1 + 1/\beta)$, the mean lifetime of equation (7.17).

The remaining-life quantile. The p-quantile of the remaining life for a known state, $q_p(t_0; \mathbf{z})$, is the time at which the conditional survival falls to $1 - p$; by equation (10.2) it solves

$$\Lambda(t_0 + q_p; \mathbf{z}) - \Lambda(t_0; \mathbf{z}) = -\ln(1 - p), \qquad q_p(t_0; \mathbf{z}) = \eta\big[(t_0/\eta)^\beta - \ln(1 - p)\big]^{1/\beta} - t_0, \tag{10.25}$$

the second form for the Weibull model; $p = 1/2$ gives the median remaining life.

The two quantities as derived quantities. Both are functions of the state, finite for every state. On the state route their gradients are built from survivor rows. Differentiating the first form of equation (10.25) in $\mathbf{z}$, with $\partial\Lambda/\partial t = \lambda$, gives $\nabla\Lambda(t_0 + q_p) - \nabla\Lambda(t_0) + \lambda(t_0 + q_p)\, \nabla q_p = \mathbf{0}$; differentiating the first integral of equation (10.23) under the integral sign, with $\nabla R(s \mid t_0; \mathbf{z}) = -R(s \mid t_0; \mathbf{z})\big[\nabla\Lambda(t_0 + s) - \nabla\Lambda(t_0)\big]$ from equation (10.2), gives the gradient of the mean residual life:

$$\nabla q_p = -\frac{\nabla\Lambda(t_0 + q_p) - \nabla\Lambda(t_0)}{\lambda(t_0 + q_p)}, \qquad \nabla\mathrm{MRL}(t_0) = -\int_0^\infty R(s \mid t_0; \mathbf{z})\big[\nabla\Lambda(t_0 + s) - \nabla\Lambda(t_0)\big]\mathrm{d}s. \tag{10.26}$$

The bands are placed on the logarithms, which keeps them positive: $\mathbf{g} = \nabla q_p / q_p$ or $\nabla\mathrm{MRL}/\mathrm{MRL}$ at $\hat{\mathbf{z}}$, and $\sigma^2 = \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$, as for the double logarithm of the conditional survival.

Exactly, their posteriors follow from draws of the exact posterior (Section 10.2.2) pushed through equations (10.24) and (10.25).

The predictive mean residual life. For a string of the record, the identity applied to the predictive distribution, with equation (10.11) for each s and the order of the integrals over s and over the state exchanged, which Tonelli's theorem allows because the integrand is nonnegative, gives

$$\mathbb{E}[T_j - t_j \mid \mathbf{y}] = \int_0^\infty \Pr(T_j > t_j + s \mid \mathbf{y})\,\mathrm{d}s = \mathbb{E}\big[\mathrm{MRL}(t_j; \mathbf{z}) \mid \mathbf{y}\big]. \tag{10.27}$$

On the state route the integrand is the sum (10.14) with $\tau = s$. Under the exact posterior it is equation (10.15) with $\tau = s$. Given the shape, it behaves for large s as $(A/s^\beta)^k$, a power-law tail of exponent $k\beta$, so its integral over s is finite only when $k\beta > 1$: averaging the Weibull over the gamma distribution of ζ produces a lifetime with a heavy tail when the shape is small. The flat prior in v leaves positive posterior probability on $\beta < 1/k$, and the predictive mean residual life (10.27) is then infinite, whatever the record. The quantity that can be reported is the restricted mean residual life to the horizon H,

$$\mathrm{MRL}_H(t_j) = \mathbb{E}\big[\min(T_j - t_j, H) \mid \mathbf{y}\big] = \int_0^H \Pr(T_j > t_j + s \mid \mathbf{y})\,\mathrm{d}s, \tag{10.28}$$

from the same identity applied to the capped variable, whose survival function is that of $T_j - t_j$ up to H and zero after; for a string outside the record the integrand is equation (10.13) with $\tau = s$. The horizon and a prior are different restrictions. The horizon caps the remaining life counted in the answer and leaves the posterior as it is; a proper prior that excluded very shallow shapes would change the posterior and make the unrestricted mean finite (Section 10.4.4). A quantile of the predictive remaining life needs no horizon: the predictive median of string j is the s at which $\Pr(T_j > t_j + s \mid \mathbf{y}) = 1/2$, and that probability falls to zero as s grows whatever the tail.

Building blocks. The sum (10.14) for each s and a quadrature in s to H; the survivor-row gradients, the hazard and $\boldsymbol{\Sigma}$ (Section 10.2.3); the incomplete gamma function; and, for the exact values beside them, equation (10.15) for each s and draws of the exact posterior (Section 10.2.2).

The mean residual life by age. Table 10.10 gives the mean residual life for ages 0, 2, 4 and 6 years and for string L. Each row is an age; the columns give the plug-in value of equation (10.24) at the fitted state; the state-route average of the restricted mean residual life, which is the same to the digits shown at the three horizons; the exact posterior average for three horizons, which is the predictive (10.28) for a string of the record at its own age, here string L; and, for $H = 30$ years, the predictive for a string of age t_0 not in the record, with equation (10.13) as the integrand, on the state with the exact value in parentheses. The integrals over s are evaluated by adaptive quadrature. One row of the plug-in arithmetic: at $t_0 = 0$, equation (10.24) is $\hat{\eta}\,\Gamma(1 + 1/\hat{\beta}) = 6.497 \times \Gamma(1.2116) = 6.497 \times 0.9152 = 5.95\,\mathrm{yr}$.

Figure 10.4 draws the plug-in and the posterior averages for $H = 30\,\mathrm{yr}$ against age. The horizontal axis is the age reached and the vertical axis the mean residual life in years; the green curve is the plug-in, the dotted blue curve the state-route average and the red curve the exact posterior average of the restricted mean residual life. All fall with age, as a rising hazard requires. The state route stays near the plug-in, 0.16 year above it at six years; the exact gap widens from 1.0 year for a new string to 1.5 years at six years, where the exact predictive is more than twice the plug-in. The gap between the two routes at six years is 1.35 years.

Table 10.10: Mean residual life at age t_0, in years. Plug-in: equation (10.24) at $\hat{\eta} = 6.497\,\text{yr}$, $\hat{\beta} = 4.726$. State: the restricted mean residual life under $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$, the same for the three horizons. Exact: the restricted mean residual life averaged over the exact posterior, for three horizons H; for string L it is its predictive (10.28). Last column: the predictive for a string of age t_0 not in the record, $H = 30\,\text{yr}$, on the state with the exact value in parentheses.

Age t_0 [yr]	Plug-in [yr]	State [yr]	Exact, $H = 15\,\text{yr}$	Exact, $H = 30\,\text{yr}$	Exact, $H = 60\,\text{yr}$	Not in the record, $H = 30\,\text{yr}$
0	5.95	5.98	6.67	6.91	7.06	5.98 (6.91)
2	3.96	4.04	4.90	5.12	5.28	4.04 (5.08)
4	2.23	2.36	3.42	3.63	3.79	2.36 (3.60)
6	1.09	1.25	2.40	2.60	2.77	1.40 (2.99)
L, 6.2	1.01	1.17	2.33	2.53	2.69	–

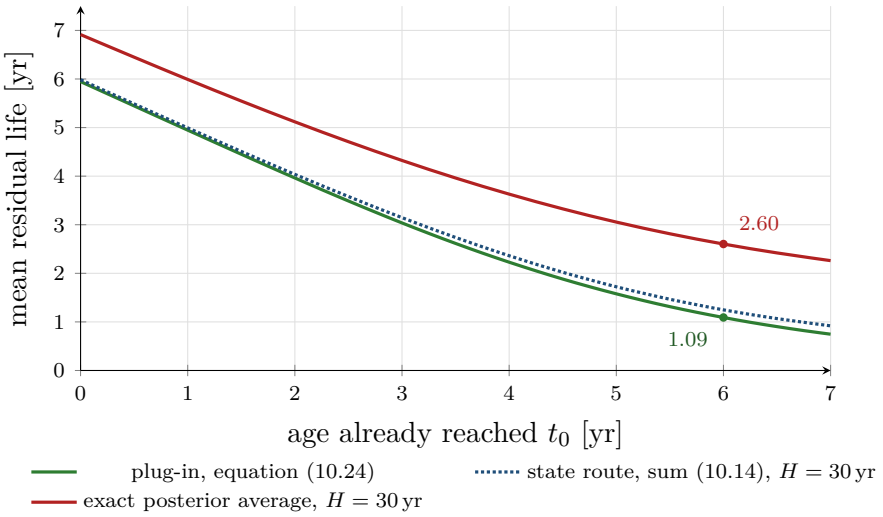


Figure 10.4: Mean residual life against the age reached. Plug-in: $(\hat{\eta}/\hat{\beta})\,\mathrm{e}^{\Lambda}\,\Gamma(1/\hat{\beta}, \Lambda)$ with $\Lambda = (t_0/6.497)^{4.726}$, equation (10.24); state route: the sum (10.14) integrated over s; exact posterior average: $\int_0^{30} \mathbb{E}[R(s \mid t_0; \mathbf{z}) \mid \mathbf{y}]\,\mathrm{d}s$ under the exact posterior, the predictive (10.28) of a string of the record at its own age.

The horizon. On the state route the horizon does not move the answer: $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ puts 2.4×10^{-14} of its probability on shapes below $1/k = 1/4$, and string L has 1.17 yr to every horizon of the table. The exact posterior puts 2.1×10^{-4} of posterior probability on those shapes, which makes the unrestricted predictive mean infinite. At six years the posterior average moves from 2.40 to 2.60 between $H = 15$ and 30 years, and to 2.77 at 60. The shapes below 1/4 contribute only 0.003 and 0.005 of those two increments; shapes between 1/4 and one, a posterior probability of 0.016, contribute 0.10 and 0.12; and shapes above one the remaining 0.10 and 0.04. A string outside the record has more expected life left at six years, 2.99 against 2.60. Under the exponential fit the mean residual life is 13.0 years at every age, plug-in, 14.7 years on the state route, $(1/\hat{\lambda})\, e^{\sigma_\omega^2/2}$, and 17.3 years exactly, the mean of $1/\lambda$ under Gamma(4, 51.9 yr).

The remaining life. Table 10.11 gives the delta method on the logarithms of the mean residual life and of the median remaining life, beside the exact bands of Table 10.12. Each row is a quantity at an age; the columns give the plug-in value at $\hat{\mathbf{z}}$, the gradient $\mathbf{g}$ of its logarithm from equation (10.26), the standard deviation $\sqrt{\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}}$ of the logarithm, the Laplace band and the exact band.

Table 10.11: The mean residual life and the median remaining life, in years, by the delta method on their logarithms at $\hat{\mathbf{z}}$ with $\boldsymbol{\Sigma}$ of equation (10.8), centered at the plug-in, against the exact bands of Table 10.12.

Quantity, age t_0 [yr]	Plug-in [yr]	$\mathbf{g}^\mathsf{T}$ of the log [–]	sd of the log [–]	Laplace band [yr]	Exact band [yr]
MRL, 0	5.95	(1.000, 0.058)	0.107	[4.82, 7.34]	[4.95, 14.6]
MRL, 6	1.09	(3.258, −0.865)	0.579	[0.35, 3.39]	[0.54, 12.7]
MRL, L at 6.2	1.01	(3.343, −0.957)	0.616	[0.30, 3.38]	[0.48, 12.7]
$q_{1/2}$, 0	6.01	(1.000, 0.078)	0.106	[4.88, 7.40]	[4.85, 12.0]
$q_{1/2}$, 6	0.955	(3.659, −0.784)	0.593	[0.30, 3.06]	[0.46, 9.65]
$q_{1/2}$, L at 6.2	0.874	(3.754, −0.892)	0.635	[0.25, 3.03]	[0.40, 9.59]

One row of the Laplace route, the median remaining life of string L: at $\hat{\mathbf{z}}$, $q_{1/2} = 0.874$ yr, the survivor rows at 7.074 and 6.2 years have the gradients (−7.063, 0.601) and (−3.787, −0.177), the hazard at 7.074 years is $0.9985\,\text{yr}^{-1}$, and equation (10.26) gives

$$\nabla q_{1/2} = -\frac{1}{0.9985}\begin{bmatrix}-3.276\\ 0.778\end{bmatrix} = \begin{bmatrix}3.281\\ -0.779\end{bmatrix}, \qquad \mathbf{g} = \frac{\nabla q_{1/2}}{0.874} = \begin{bmatrix}3.754\\ -0.892\end{bmatrix}, \qquad \mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g} = 0.403,$$

a standard deviation of 0.635 and the band $0.874\, e^{\pm 1.244} = [0.25,\ 3.03]$ yr.

The gradients carry the dependence on age. For a new string $\mathbf{g}$ is nearly $(1,\ 0)^\mathsf{T}$: the remaining life is the lifetime, which scales with η, and its standard deviation is close to $\sqrt{\Sigma_{uu}} = 0.113$. For string L the u entry is above three, since a longer scale both stretches the time axis and makes the string younger relative to it, and the v entry is near −0.9, so that the shape, with its larger variance Σ_{vv}, enters the product; the standard deviation of the logarithm rises nearly six-fold. The Laplace bands are centered at the plug-in, at the steep joint MAP shape, and miss the long upper tails of the exact posteriors.

The remaining life as a derived quantity. Table 10.12 gives the exact posteriors of the mean residual life $\text{MRL}(t_0; \mathbf{z})$ and of the median remaining life $q_{1/2}(t_0; \mathbf{z})$ for a new string, a string of six years and string L. Each row is a quantity at an age; the columns give the plug-in value at $\hat{\mathbf{z}}$ and the median and 95 % band over the draws of the exact posterior, pushed through equations (10.24) and (10.25).

The predictive median remaining life of string L, the s at which its conditional predictive survival is one half, is 0.81 yr on the state route, 1.32 yr exactly, against the plug-in 0.87 yr.

Table 10.12: The exact posteriors of the mean residual life and of the median remaining life, in years, as derived quantities of the state: plug-in at $\hat{\mathbf{z}}$, and quantiles over 2×10^5 draws of the exact posterior (seed 2026).

Quantity, age t_0 [yr]	Plug-in [yr]	Exact median [yr]	Exact band [yr]
MRL, 0	5.95	6.31	[4.95, 14.6]
MRL, 6	1.09	1.67	[0.54, 12.7]
MRL, L at 6.2	1.01	1.58	[0.48, 12.7]
$q_{1/2}$, 0	6.01	6.29	[4.85, 12.0]
$q_{1/2}$, 6	0.955	1.46	[0.46, 9.65]
$q_{1/2}$, L at 6.2	0.874	1.37	[0.40, 9.59]

10.3.4 The fleet over the coming year

Engineering question. Will any of the eight survivors fail before the next test, and how many? The answer sets how many spare strings to hold for the year.

General form. A joint event is the predictive (2.13) with y_{new} the indicator of an intersection, and a count of future events is read by its moments, equation (2.10):

$$\Pr\Big(\bigcap_j A_j \;\Big|\; \mathbf{y}\Big) = \int \Pr\Big(\bigcap_j A_j \;\Big|\; \mathbf{z}\Big)\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \qquad \mathbb{E}[K_f \mid \mathbf{y}] = \int \mathbb{E}[K_f \mid \mathbf{z}]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}. \quad (10.29)$$

Specialized form. The question is equation (10.29) with $A_j = \{T_j > t_j + \tau\}$ for the eight survivors, K_f the number of them that fail the year, and $p(\mathbf{z} \mid \mathbf{y})$ the posterior of step 3, $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$. One choice enters here: given the state the eight strings fail independently in the coming year, with no common cause such as a charger fault or an overheated room acting on several at once. This is what makes the joint survival a product given the state, so that every dependence between the strings comes from the state they share, and a common cause would add a shared entry of its own, the subject of Chapter 15.

Given the state, by the argument of Section 10.3.1 applied to each, the probability that all eight survive the coming year is the product of their conditional survivals. Averaging over the posterior once, since all eight depend on the same state,

$$\Pr(K_f = 0 \mid \mathbf{y}) = \int \prod_j R(\tau \mid t_j; \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z} = \mathbb{E}\Big[\prod_j R(\tau \mid t_j; \mathbf{z}) \;\Big|\; \mathbf{y}\Big], \quad (10.30)$$

with the product over the eight survivors. The expectation of a product is not the product of the expectations, $\prod_j \mathbb{E}[R(\tau \mid t_j; \mathbf{z}) \mid \mathbf{y}]$, unless the factors are independent under the posterior, and they are not: all of them are functions of the same two entries of the state. The coupling is a property of the graph, not of the strings: the eight survivors are eight rows hanging from one free node (Figure 10.1), and had each string its own free node, the joint survival would be the product of the marginals, and a record of twelve units would say nothing about any one of them. Chapter 9 found the same coupling for two generators sharing one uncertain rate; here it acts on eight units and a two-entry state, and the derivation, one expectation of a product, is the same.

The state route. Under $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ the expectation (10.30) is the sum (10.14) with the product of the eight conditional survivals as the function of the state. The same nodes, with the count's distribution given the state, a product of eight independent Bernoulli factors, carry the distribution of K_f.

The exact form. The value beside it comes from the exact posterior. By equation (10.4) the product is $\exp(-\zeta \sum_j \nu_j)$, with $\nu_j = (t_j + \tau)^\beta - t_j^\beta$, and the average over the scale given the shape (Section 10.2.2) gives

$$\Pr(K_f = 0 \mid \mathbf{y}) = \int p(v \mid \mathbf{y}) \left(\frac{A(\beta)}{A(\beta) + \sum_j \nu_j(\beta)} \right)^k \mathrm{d}v. \tag{10.31}$$

The eight increments add inside one transform: the fleet faces one unknown ζ and one unknown shape, and the year's cumulative hazards of the eight strings add, whereas the product of the marginals would average a separate copy of the state for each string.

The number of failures. Write $K_f = \sum_j (1 - \mathbf{1}_j)$, with $\mathbf{1}_j$ the indicator that string j survives the year. The expectation is linear whatever the dependence:

$$\mathbb{E}[K_f \mid \mathbf{y}] = \sum_j [1 - \Pr(T_j > t_j + \tau \mid \mathbf{y})], \tag{10.32}$$

the sum of the eight marginal failure probabilities. The variance is not:

$$\mathrm{Var}(K_f \mid \mathbf{y}) = \sum_j F_j(1 - F_j) + \sum_{i \neq j} \Big(\mathbb{E}[R(\tau \mid t_i; \mathbf{z}) R(\tau \mid t_j; \mathbf{z}) \mid \mathbf{y}] - (1 - F_i)(1 - F_j) \Big), \tag{10.33}$$

with $F_j = 1 - \Pr(T_j > t_j + \tau \mid \mathbf{y})$ the marginal failure probability of string j and each pair expectation of the form of equation (10.31) with $\nu_i + \nu_j$ in place of the sum of eight.

The sign of the pair terms. Given the state the indicators are independent, so the covariance of two of them is the posterior covariance of their conditional survivals, $\mathrm{Cov}(\mathbf{1}_i, \mathbf{1}_j \mid \mathbf{y}) = \mathrm{Cov}\big(R(\tau \mid t_i; \mathbf{z}),\ R(\tau \mid t_j; \mathbf{z}) \mid \mathbf{y}\big)$. Linearized at $\hat{\mathbf{z}}$, each conditional survival has the gradient $\mathbf{g}_{R,i} = -R_i \Lambda_i^\tau \, \mathbf{g}_i$, with Λ_i^τ the cumulative hazard of string i's coming year and $\mathbf{g}_i$ the gradient of its double logarithm, equation (10.19), and the covariance is the bilinear form of the two gradients under $\boldsymbol{\Sigma}$:

$$\mathrm{Cov}(\mathbf{1}_i, \mathbf{1}_j \mid \mathbf{y}) \approx \mathbf{g}_{R,i}^\mathsf{T} \boldsymbol{\Sigma} \, \mathbf{g}_{R,j} = R_i \Lambda_i^\tau \, R_j \Lambda_j^\tau \, \beta^2 \begin{bmatrix} -1 & \bar{x}_i \end{bmatrix} \boldsymbol{\Sigma} \begin{bmatrix} -1 \\ \bar{x}_j \end{bmatrix}. \tag{10.34}$$

The factor in front is positive, so the sign is that of the quadratic form, $\Sigma_{uu} - (\bar{x}_i + \bar{x}_j)\Sigma_{uv} + \bar{x}_i \bar{x}_j \Sigma_{vv}$. The first entries of the two gradients, which carry the scale, always agree, and their term Σ_{uu} is positive; the second entries, which carry the shape, have the signs of the two levers. For a young string, $\bar{x}_i < 0$, and an old one, $\bar{x}_j > 0$, the shape term $\bar{x}_i \bar{x}_j \Sigma_{vv}$ is negative, and so is the cross term $-(\bar{x}_i + \bar{x}_j)\Sigma_{uv}$ when the young string's lever is the larger in size, since $\Sigma_{uv} < 0$; together they can outweigh Σ_{uu}. Pairs of old strings are positively correlated; a pair of a young and an old string can be slightly negatively correlated, because a steeper shape raises the young string's survival and lowers the old one's.

The exact split. Under the exact posterior the same covariance splits by the law of total covariance over the shape,

$$\mathrm{Cov}(\mathbf{1}_i, \mathbf{1}_j \mid \mathbf{y}) = \mathbb{E}\big[\mathrm{Cov}(\mathrm{e}^{-\nu_i \zeta}, \mathrm{e}^{-\nu_j \zeta} \mid v, \mathbf{y}) \mid \mathbf{y}\big] + \mathrm{Cov}\big(\mathbb{E}[\mathrm{e}^{-\nu_i \zeta} \mid v, \mathbf{y}],\ \mathbb{E}[\mathrm{e}^{-\nu_j \zeta} \mid v, \mathbf{y}] \mid \mathbf{y}\big), \tag{10.35}$$

each term an integral of the one-integral route. The first, the within-shape term, is nonnegative, since at a fixed shape both survivals decrease in the one variable ζ; the second, the between-shape term, carries the sign of the shape entries in equation (10.34).

Building blocks. The sum (10.14) with the product over the eight survivors, the eight marginals of Section 10.3.1 and the levers $\bar{x}_j$ with $\boldsymbol{\Sigma}$; and, for the exact values beside them, the average over the scale with the increments added, $p(v \mid \mathbf{y})$ and draws of (v, ζ) for the distribution of the count (Section 10.2.2).

The year of the eight survivors. Table 10.13 gives the probability that none of the eight survivors fails before the next test, and the expected number of failures. Each row is a route; the columns give the probability of no failure and the expected number of failures. The state route is the sum (10.14); the exact route is equation (10.31); each product of the marginals multiplies the eight predictives of Table 10.7; the plug-in product multiplies the eight plug-in values.

Table 10.13: The probability that none of the eight survivors fails in the coming year, and the expected number of failures. State: the sum (10.14) over $\mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$. Exact: equation (10.31). Products of marginals: $\prod_j \Pr(T_j > t_j + 1 \mid \mathbf{y})$; plug-in product: $\prod_j R(1 \mid t_j; \hat{\mathbf{z}})$; exponential: on the state, with $(51.9/59.9)^4$ in parentheses, and $\mathrm{e}^{-8\hat{\lambda}}$.

Route	$\Pr(K_f = 0 \mid \mathbf{y})$ [–]	$\mathbb{E}[K_f \mid \mathbf{y}]$ [–]
State route, shared state	0.125	2.01
Exact, shared state	0.235	1.51
Product of the state marginals	0.074	2.01
Product of the exact marginals	0.170	1.51
Plug-in product	0.091	1.87
Exponential, shared rate, state (plug-in)	0.527 (0.540)	0.66
Exponential, shared rate, exact	0.564	0.59

The fleet gets through the year without a failure with probability 0.125, about one year in eight, and expects 2.01 failures. The product of the marginals has the same expected number and puts the probability of no failure at 0.074, 41 % low. By the sum over the nodes the variance of the number of failures is 1.75, against 1.21 for eight independent strings with the same marginal probabilities, the pair terms adding 0.542; the probability of two or more failures is 0.613 and of three or more 0.337. For the two oldest strings, K and L, the probability that both last the year is 0.281, against the product $0.547 \times 0.425 = 0.232$, a correlation of 0.20. The covariance of E and L is -0.0019, and the same seven of the twenty-eight pairs named below are slightly negatively correlated. Under the exponential fit the fleet survives the year with probability 0.527 and expects 0.66 failures.

The values beside them come from the exact posterior. The fleet gets through the year without a failure with probability 0.235, about one year in four, and expects 1.51 failures. The product of the marginals has the same expected number and puts the probability of no failure at 0.170, 28 % low. The year's plug-in cumulative hazards add to 2.400, and $\mathrm{e}^{-2.400} = 0.091$. By equation (10.33) the variance of the number of failures is 1.49, against 1.09 for eight independent strings with the same marginal probabilities. By quadrature over the exact posterior the probability of two or more failures is 0.456 and of three or more 0.206. The pair terms add 0.394 in all, the within-shape term of equation (10.35) giving 0.296 and the between-shape term 0.097. For the two oldest strings, K and L, the probability that both last the year is 0.445, against the product $0.682 \times 0.593 = 0.404$, a correlation of 0.18. For E and L the within-shape term, $+0.0010$, is outweighed by the between-shape term, -0.0036, and their covariance is -0.0025. Seven of the twenty-eight pairs, E with each of I to L and F with each of J to L, are slightly negatively correlated. Under the exponential fit the fleet survives the year with probability 0.564 and expects 0.59 failures. The gap changes a statement here: the plug-in's 1.87 expected failures lies below the state route's 2.01 and above the exact 1.51, so the two routes fall on opposite sides of it. On the state route the plug-in holds too few spares and on the exact route too many.

10.3.5 The replacement order and the two drawings

Engineering question. Which string should be replaced first, and does the answer depend on the hazard shape? If the hazard rises with age the oldest string is first; if it is constant, age says nothing about the coming year.

General form. Which of two drawings describes the record is answered by the ratio of their evidences, equation (2.14):

$$B_{10} = \frac{p(\mathbf{y} \mid M_1)}{p(\mathbf{y} \mid M_0)}. \tag{10.36}$$

Specialized form. The question is equation (10.36) with M_1 the Weibull drawing and M_0 the exponential nested in it at $v = 0$, each evidence the integral of the rows over the posterior of the state of step 3 of that drawing. Two choices enter here. The strings are ranked by the first question's answer, the probability F_j of failing the coming year: this gives one number per string, and it leaves out the consequence of a failure and the cost of replacing a string that still has life, which the replacement decision of Section 10.5.2 weighs. The constant hazard is compared as the drawing nested in the Weibull at $v = 0$, with the Weibull shape under a proper prior $\pi(v)$: this is what makes the ratio defined, and the ratio depends on the width of π, so two widths, $\mathcal{N}(0, 0.5^2)$ and $\mathcal{N}(0, 1^2)$, are carried.

Under the Weibull drawing with $\beta > 1$ the conditional survival falls with age for every state, by equation (10.3), so the order of the strings is the order of their ages; the predictive keeps that order when the posterior of the shape lies mostly above one, and the drawing then decides the gaps between the strings rather than their order. Under the exponential drawing every string of the record has the same predictive, $(T_{\text{tot}}/(T_{\text{tot}} + \tau))^k$ (Section 10.3.1), and the record gives no order. Which drawing the record favors is a hypothesis comparison. The posterior probability $\Pr(\beta > 1 \mid \mathbf{y})$ does not answer it: it is computed inside the Weibull drawing and says where the shape lies if that drawing holds. On the Laplace route the prior row v/σ_0 of the pattern (10.5), with Jacobian row $(0,\ 1/\sigma_0)$, is added to the twelve rows and refitted from $\hat{\mathbf{z}}$, and each evidence is the Gaussian integral of $\exp(-\mathcal{L})$ about its own MAP state, equation (7.70), with $\hat{\mathbf{z}}_M$ and $\det \boldsymbol{\Sigma}_M$ of each drawing: the one-entry fit of Section 10.2.1 for the exponential, the refit with the prior row for the Weibull. The value beside it comes from the exact posterior. With the exponential nested at $v = 0$ and the scale under the same flat prior in both drawings, the improper constant cancels, and the Savage–Dickey ratio (7.68), inverted, gives

$$B_{10} = \frac{\int \Omega(v)\, \pi(v)\, \mathrm{d}v}{\Omega(0)}, \tag{10.37}$$

with $\Omega(v)$ the likelihood with the scale integrated out at fixed shape, equation (7.46).

Building blocks. The failure probabilities F_j of Section 10.3.1; the refit with the prior row and its $\boldsymbol{\Sigma}$, and the two priors; and, for the exact values beside them, $\Omega(v)$ on the grid in v (Section 10.2.2).

The replacement order and the evidence ratio. By the failure probabilities of Table 10.7, string L is first, with 0.58, then K with 0.45 and J with 0.38: the order of age. The exact predictives give the same order, 0.41, 0.32 and 0.27. The plug-in gives the same order with different gaps, 0.56 for L against 0.0056 for E, a ratio of a hundred; the state route 0.575 against 0.012, a ratio of forty-seven; and the exact predictive 0.41 against 0.022, a ratio of eighteen. Under the exponential drawing all eight have the predictive failure probability 0.083 on the state route, 0.073 exactly, and there is no first string. The refitted shapes are 2.19 for $\sigma_0 = 0.5$ and 3.74 for $\sigma_0 = 1$, and the Laplace evidences (7.70) of the two drawings give $B_{10} = 1.96$

and 2.97. The value beside it comes from the exact posterior: on the grid of v, the evidence ratio (10.37) is

$$B_{10} = 2.22 \quad (\ln B_{10} = 0.80 \text{ nats})$$

for $\pi = \mathcal{N}(0, 0.5^2)$, and $B_{10} = 3.70$ (1.31 nats) for $\mathcal{N}(0, 1^2)$: odds of two to four to one for the Weibull drawing, against the state route's two to three, a gap that changes no statement. Under these priors the Weibull posterior puts 0.91 and 0.97 on $\beta > 1$, against 0.984 under the flat prior.

10.4 The checks (step 5)

Every answer of Section 10.3 rests on a model, the Weibull lifetime with administrative censoring under a flat prior, and is computed on a Gaussian for the posterior of the state with a linearization of some derived quantities; each has its exact value beside it. This section sets each route against a fuller calculation, takes the gap apart, and varies the two choices the record cannot defend by itself, the prior on the shape and the reading of a failed test as a failure at the test age. The strings have no physics row, so no limit $\sigma_p \to 0$ of a model tolerance arises.

10.4.1 The fit and the exact route against the draws

Table 10.14 sets each number of the fit and of the one-integral route beside a calculation that does not use the formula that produced it. Each row is one check; the columns give the quantity, its value by the route it belongs to, its value by the independent calculation, and the verdict. The first three rows check the fit of Section 10.2.3; the last four check the one-integral route against the 2×10^5 draws of the exact posterior.

Table 10.14: The checks of the fit and of the one-integral route. Route: the fit or the one integral; independent: a calculation that does not use it. Draws: 2×10^5 of the exact posterior under the flat prior, seed 2026.

Quantity	Route	Independent calculation	Verdict
MAP shape $\hat{\beta}$	Newton fit, Table 10.3: 4.7257	shape equation (7.43) on the best-scale curve: 4.7257	agree to 10^{-14}
Information $\mathbf{\Lambda}$	summed row Hessians: 89.331, 9.009, 7.491	central difference of $\mathcal{L}$ at $\hat{\mathbf{z}}$: the same	agree to 7×10^{-6}
Best-scale condition	$\sum_i \Lambda_i(\hat{\mathbf{z}}) = 4.000$	the number of failures, $k = 4$	agree
String L's conditional predictive	one integral (10.15): 0.593	mean over the draws: 0.594	agree within 0.001
String L's 95 % band	distribution function (10.21): [0.099, 0.942]	quantiles of the draws: [0.099, 0.942]	agree
$\Pr(R < 0.80)$ for string L	equation (10.21): 0.778	fraction of the draws: 0.777	agree within 0.001
Fleet: $\Pr(K_f = 0 \mid \mathbf{y})$; $\mathbb{E}[K_f \mid \mathbf{y}]$	equations (10.31) and (10.32): 0.235; 1.507	means over the draws: 0.236; 1.505	agree within 0.002

For the four rows of Table 10.9 the draws give the bands [0.877, 0.998], [0.681, 0.961], [0.143, 0.943] and [0.099, 0.942] and the probabilities 0.001, 0.227, 0.760 and 0.777, within 0.002 of the one-integral values. The one integral and the draws use the same exact posterior and differ only in how the integral over the state is evaluated, so their agreement checks the closed forms (10.7), (10.15), (10.21) and (10.31), the grid in v and the quadrature; it does not check

the posterior, which both inherit from the flat prior and the record. Sections 10.4.4 and 10.4.5 vary those.

10.4.2 The Laplace posterior against the exact posterior

The state. Table 10.15 sets the Laplace approximation of the fit beside the exact posterior. Each row is a quantity; the middle column gives it under the Gaussian $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of equation (10.8), the right column under the exact posterior of Section 10.2.3.

Table 10.15: The Weibull fit to the strings: the Laplace approximation in (u, v) with covariance (10.8), and the exact posterior under the same flat prior, equations (7.44) and (7.46). Bands are central 95 % intervals.

Quantity	Laplace in (u, v)	Exact
Scale η [yr]	6.50, [5.21, 8.11]	median 6.93, [5.49, 15.5]
Standard deviation of $u = \ln \eta$ [–]	0.113	0.323
Correlation of u and v [–]	-0.348	-0.582
Mode of v (shape β) [–]	1.553 (4.726)	1.383 (3.99)
Mean and standard deviation of v [–]	1.553, 0.390	1.241, 0.489
Shape β, median and band [–]	4.73, [2.20, 10.14]	3.64, [1.15, 7.77]
$\Pr(\beta > 1 \mid \mathbf{y})$ [–]	1.000	0.984
$\Pr(\beta > 2 \mid \mathbf{y})$ [–]	0.986	0.872

The two errors that Section 7.4.3 found for the motors recur: the Gaussian centers the shape at the joint MAP 4.73 when the marginal mode is 3.99, and it ignores the skew of the gamma distribution of the scale given the shape, which with $k = 4$ is strong. Here it is also too narrow, most of all in the scale, whose exact spread in u is nearly three times the Laplace one. The reason is the direction of the long axis: the exact covariance has it along $(0.453,\ -0.892)^{\mathsf{T}}$, with standard deviation 0.534, where $\boldsymbol{\Sigma}$ has $\mathbf{e}_1 = (0.108,\ -0.994)^{\mathsf{T}}$ and $\sqrt{\epsilon_1} = 0.392$. The exact long axis is longer and tilted toward the scale, so that along it a shallower shape travels with a much longer scale, and its projection on u is what widens the scale; the upper end of the exact scale band, 15.5 years, is that tail. The displacement of the exact mean from $\hat{\mathbf{z}}$, $(0.137,\ -0.312)^{\mathsf{T}}$, lies nearly along the same axis.

The survival of the coming year. The Laplace posterior average of the one-year survival, by a 60×60 Gauss–Hermite rule under $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$, is 0.998, 0.969, 0.815 and 0.457 at ages 0, 2, 4 and 6 years, against the exact 0.993, 0.959, 0.850 and 0.618 of Table 10.6. It is closer to the plug-in than to the exact value, because the Gaussian it averages over is centered at the same steep joint MAP shape and is too narrow in the shape. For the eight survivors, E to L, the same rule gives 0.988, 0.966, 0.896, 0.829, 0.718, 0.623, 0.547 and 0.425, against the exact predictives 0.978 to 0.593 of Table 10.7. For string L the Laplace value, 0.425, lies below even the plug-in 0.439: averaging over a Gaussian centered at the steep shape lowers the old string's survival further, where the exact posterior raises it.

The band and the threshold. The band and threshold of Section 10.3.2 are the delta method on the double logarithm. The normal $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ itself, pushed through the conditional survival by 2×10^5 draws (seed 2026), gives for string L the mean 0.425, as the sum (10.14) does, the median 0.462, the band [0.0002, 0.787] and 0.983 below 0.80: the linearization on ϑ lies between it and the exact posterior, and the fleet by those draws is 0.126, against the 0.125 of the sum.

Figure 10.5 draws the posterior of string L's one-year survival. The horizontal axis is $R(1 \mid 6.2; \mathbf{z})$ and the vertical axis its posterior density. The bars are the histogram of the two hundred thousand draws of the exact posterior, pushed through $e^{-\nu\zeta}$; the dotted blue curve is the density implied by the Laplace approximation, the normal on ϑ mapped through $\exp(-e^{\vartheta})$; the shaded band on the right is the region $R \geq 0.80$ of the third question; and the dotted lines mark the plug-in value 0.439 and the exact conditional predictive 0.593. The posterior of the draws is spread over almost the whole interval, rising from the left to a broad peak near 0.75 and falling steeply toward one. The Laplace density places its mass to the left, near the plug-in, and has a second rise toward zero that the draws do not show.

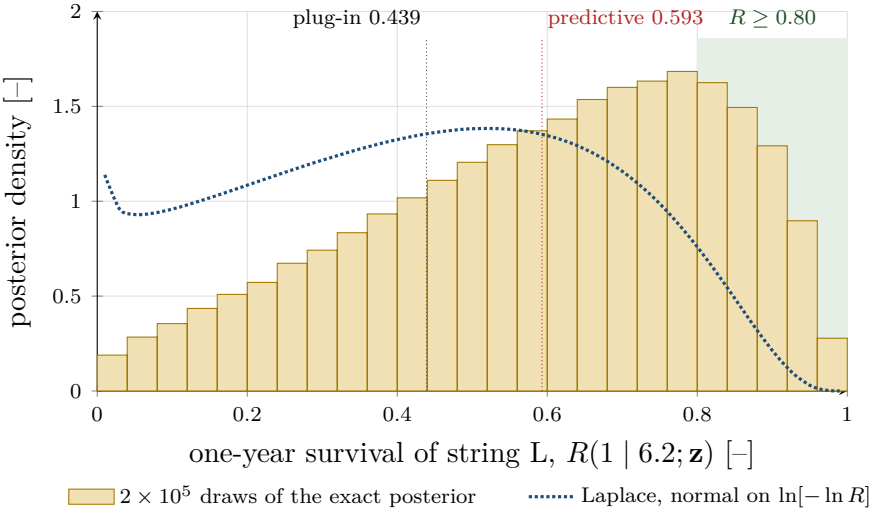


Figure 10.5: The posterior of string L's one-year survival at $t_0 = 6.2\,\text{yr}$: histogram of $\exp[-\zeta(7.2^{\beta} - 6.2^{\beta})]$ over 2×10^5 draws of the exact posterior under the flat prior (seed 2026), against the Laplace density, $\vartheta = \ln[-\ln R] \sim \mathcal{N}(-0.195,\ 0.815^2)$ mapped to R. Shaded: the region $R \geq 0.80$, of posterior probability 0.222 exact and 0.055 under the Laplace approximation.

The delta method on the double logarithm keeps the band inside $(0, 1)$, as it was built to, but it inherits the Gaussian in (u, v): its band for string L is $[0.017,\ 0.846]$ against the exact $[0.099,\ 0.942]$, and it puts 0.945 rather than 0.778 on a survival below 0.80, so that the threshold probability of the third question is 0.055 against 0.222. As for the availability of Chapter 8 and the mission reliability of Chapter 9, a linearization on the right scale cannot repair a posterior Gaussian built on four failures. That cost is reported with each answer, and the requirement of 0.80 is not shown on either route.

10.4.3 The remaining life, the fleet and the evidence ratio by the Laplace route

The remaining life. Section 10.3.3 gives the state-route bands beside the exact ones: the Laplace bands are centered at the plug-in, at the steep joint MAP shape, and miss the long upper tails of the exact posteriors. For string L, 1.17 yr to any horizon stands against the exact 2.53 yr to thirty years.

The fleet. Under $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ the Gauss–Hermite average of the product of the eight conditional survivals gives $\Pr(K_f = 0 \mid \mathbf{y}) = 0.125$, against the exact 0.235. The bilinear form (10.34) with the levers $\bar{x}_E = -1.046$ and $\bar{x}_L = 0.248$ gives -0.0018 for the covariance of E and L, against the exact -0.0025: in its quadratic form the cross term, -0.0122, nearly cancels $\Sigma_{uu} = 0.0127$

by itself, and the shape term, -0.0395, decides the sign. The quadratic form is negative for the same seven pairs as the exact route: the sign pattern is a property of the two gradients and $\boldsymbol{\Sigma}$, and the Laplace route gets it right while it gets the size of the fleet's probability wrong.

The evidence ratio. With the prior row v/σ_0 added to the twelve rows and refitted from $\hat{\mathbf{z}}$, the refitted shapes are 2.19 for $\sigma_0 = 0.5$ and 3.74 for $\sigma_0 = 1$, and the Laplace evidences (7.70) of the two drawings give $B_{10} = 1.96$ and 2.97, against the exact 2.22 and 3.70 of Section 10.3.5: the same modest preference for the Weibull drawing, slightly understated. Reading the Laplace normal of v at the nested value $v = 0$ instead, the Savage–Dickey ratio gives 5.4 and 66: that reading evaluates the Gaussian far in its tail, where it is least like the exact marginal, and it overstates the preference many times over.

10.4.4 Sensitivity to the prior on the shape

The flat prior leaves the shape to four failure ages. Under a proper prior $\pi(v)$ on the log-shape, with the scale flat as before, the marginal of the shape becomes $\Omega(v)\pi(v)$ normalized, and every one-integral formula keeps its form. Table 10.16 repeats the answers under the two priors of the comparison. Each row is a quantity: the posterior median of the shape, $\Pr(\beta > 1 \mid \mathbf{y})$, string L's conditional predictive, its 95 % band and the probability that it is at least 0.80, string E's conditional predictive, the fleet's probability of no failure and its expected count, and string L's restricted mean residual life to 30 years; the columns give it under the flat prior and the two proper priors. The prior $\mathcal{N}(0, 1^2)$ gives the shape a range of 0.14 to 7.1 with probability 95 %; $\mathcal{N}(0, 0.5^2)$, centered on a constant hazard, a range of 0.37 to 2.7.

Table 10.16: The answers for the strings under three priors on $v = \ln \beta$, the scale flat in every row: the flat prior of the chapter and the two proper priors of the comparison. Laplace value of the fit with the prior row, the exact value in parentheses.

Quantity	Flat	$\mathcal{N}(0, 1^2)$	$\mathcal{N}(0, 0.5^2)$
Median of the shape β [–]	4.73 (3.64)	3.74 (2.85)	2.19 (1.80)
$\Pr(\beta > 1 \mid \mathbf{y})$ [–]	1.000 (0.984)	0.999 (0.968)	0.974 (0.910)
String L, conditional predictive [–]	0.425 (0.593)	0.525 (0.683)	0.722 (0.820)
String L, 95 % band of $R(1 \mid 6.2; \mathbf{z})$ [–]	[0.017, 0.846] ([0.099, 0.942])	[0.054, 0.897] ([0.194, 0.957])	[0.321, 0.951] ([0.484, 0.971])
String L, $\Pr(R(1 \mid 6.2; \mathbf{z}) \geq 0.80 \mid \mathbf{y})$ [–]	0.055 (0.222)	0.136 (0.353)	0.465 (0.664)
String E, conditional predictive [–]	0.988 (0.978)	0.978 (0.968)	0.947 (0.947)
Fleet, $\Pr(K_f = 0 \mid \mathbf{y})$ [–]	0.125 (0.235)	0.172 (0.292)	0.299 (0.407)
Fleet, $\mathbb{E}[K_f \mid \mathbf{y}]$ [–]	2.01 (1.51)	1.84 (1.31)	1.39 (0.96)
String L, restricted mean residual life, $H = 30\,\text{yr}$ [yr]	1.17 (2.53)	1.65 (3.55)	3.80 (6.32)

The prior on the shape moves the old strings' answers more than any other check. Under $\mathcal{N}(0, 1^2)$ string L lasts the year with probability 0.525 instead of 0.425 (0.683 instead of 0.593 exactly) and the fleet gets through the year without a failure with probability 0.172 instead of 0.125 (0.292 instead of 0.235); under $\mathcal{N}(0, 0.5^2)$ the two become 0.722 and 0.299 (0.820 and 0.407), the probability that L's survival is at least 0.80 rises from 0.055 to 0.465 (0.222 to 0.664), and L's restricted mean residual life is three times longer (2.5 times exactly). Both priors pull the shape toward one, which flattens the conditional survival across ages: the old strings gain and the young lose, string E from 0.988 to 0.947 (0.978 to 0.947). Under either proper prior

string L is more likely than not to last the year on both routes; only the flat prior's state route has it on the other side of one half. The order of the strings by age survives the three priors, since the failure probabilities still rise with age in each row, from E to L; their gaps, the fleet's count and the band of string L do not. With four failures the shape is not yet the record's alone, and what is assumed about it before the record is part of every answer about the old strings.

10.4.5 The failure ages as intervals

A string is found to have failed at an annual test; its capacity fell below the requirement at some time in the year before. The fit takes the age at the failed test as the lifetime. The exact row of a failed string is the probability of failing within the year before the test, $R(t_i-1;\mathbf{z})-R(t_i;\mathbf{z})$, and with failures known only to within a year at lifetimes of three to seven years the difference is not negligible: the maximum-likelihood fit with the four interval rows is $\hat{\beta}=3.93$ and $\hat{\eta}=6.41\,\text{yr}$, against 4.73 and 6.50 yr, a shallower hazard, and the plug-in one-year survival of string L rises from 0.439 to 0.496. The correction moves the answers in the direction of the proper priors of Section 10.4.4, toward a flatter conditional survival; it alters the rows, not the questions, and the derivations of Section 10.3 are unchanged by it.

10.5 The report (step 6)

The fitted Weibull, read as known, gives string L the survival 0.439 and the eight a clean year with probability 0.091. Averaged over what four failures leave uncertain, on the posterior of the state, string L lasts the coming year with probability 0.425 (0.593 exactly) and the eight get through it without a failure with probability 0.125 (0.235). The two routes fall on opposite sides of one half for string L, and on opposite sides of the fit: the exact value says the fit overstates the old string's risk, and the state route does not. The old strings' numbers also depend on the prior on the shape: under a proper prior centered on a constant hazard, string L lasts the year with probability 0.722 (0.820). The answers below are given in the units in which the questions were asked, with the replacement decision that follows.

10.5.1 The answers and their bands

Table 10.17 is the report. Each row is a question in the engineer's words; the second column gives the answer from the posterior of the state, $\mathcal{N}(\hat{\mathbf{z}},\boldsymbol{\Sigma})$, under the Weibull model fitted to the twelve strings with the prior flat in $\ln\eta$ and $\ln\beta$, the third the exact value, and the fourth the band of the reported answer and the values under the proper priors of Section 10.4.4.

The report line. Under a Weibull lifetime fitted to twelve strings with four failures, and averaging over the uncertainty of its scale and shape, the oldest string, at 6.2 years, lasts the coming year with probability 0.42 (0.59 from the exact posterior), the two on opposite sides of one half; its one-year survival as a property of the population is known only to lie between 0.02 and 0.85 (0.10 and 0.94), and it is at least 0.80 with probability 0.06 (0.22), so that claim is not supported. The eight strings in service expect 2.0 (1.5) failures before the next test and get through it without one with probability 0.13 (0.24); they share one lifetime distribution, and multiplying their individual probabilities would understate that chance by more than a quarter on either route. The record favors the Weibull drawing over a constant hazard by an evidence ratio of 2.0 to 3.0 (2.2 to 3.7), depending on the prior on the shape, modest support for ranking the strings by age, which makes the oldest string the first to replace; under a constant hazard no string of the record would be. With four failures the old strings' numbers move with the prior on the shape and between the two routes, and both are given.

Table 10.17: The report of the strings' record: the answers from the posterior of the state, the exact posterior under the prior flat in $(\ln\eta, \ln\beta)$ beside them, the central 95 % bands of the reported route with the exact band in parentheses, and the values under the proper priors $\mathcal{N}(0, 1^2)$ and $\mathcal{N}(0, 0.5^2)$ on $\ln\beta$, Laplace with exact in parentheses.

Engineering question	Answer (state route)	Exact	Band
A string has reached its age. What is the probability it lasts until the next test?	strings E to L: 0.988, 0.966, 0.896, 0.829, 0.718, 0.623, 0.547, 0.425; a string not in the record at 0, 2, 4, 6 yr: 0.998, 0.970, 0.820, 0.521	0.978, 0.956, 0.902, 0.858, 0.791, 0.733, 0.682, 0.593; 0.993, 0.960, 0.854, 0.662	plug-in for E to L 0.994 to 0.439; exponential 0.917 (0.927) for every string of the record; L 0.525 and 0.722 (0.683 and 0.820) under the proper priors
How well is that probability known for the oldest string?	center 0.439	median 0.626	[0.017, 0.846] ([0.099, 0.942]); [0.054, 0.897] and [0.321, 0.951] ([0.194, 0.957] and [0.484, 0.971]) under the proper priors
Is the oldest string's chance of lasting the year at least 0.80?	not shown: probability 0.055	not shown: probability 0.222	0.136 and 0.465 (0.353 and 0.664) under the proper priors
How much life does a string have left?	string L: 1.17 yr to a horizon of 30 yr, predictive median 0.81 yr; a string not in the record at 0, 2, 4, 6 yr: 5.98, 4.04, 2.36, 1.40 yr ($H = 30$ yr)	2.53 yr, 1.32 yr; 6.91, 5.08, 3.60, 2.99 yr	exact 2.33 to 2.69 yr for $H = 15$ to 60 yr, unrestricted mean infinite under the flat prior; MRL$(6.2; \mathbf{z})$ in [0.30, 3.38] yr ([0.48, 12.7] yr); plug-in 1.01 yr
Will any of the eight survivors fail before the next test, and how many?	none fails with probability 0.125; 2.01 failures expected, variance 1.75	0.235; 1.51, variance 1.49	product of marginals 0.074 (0.170); plug-in 0.091 and 1.87; two or more failures 0.613 (0.456); 0.172 and 0.299 (0.292 and 0.407) under the proper priors
Which string should be replaced first, and does the answer depend on the hazard shape?	L, then K, J and so on to E, by age; $B_{10} = 2.0$ to 3.0	the same order; $B_{10} = 2.2$ to 3.7	under a constant hazard no order; modest support; the order holds under the three priors

10.5.2 The replacement decision

The order of replacement is the one answer that every reading of the record shares: oldest first, on the state route, under the exact posterior, the plug-in and each of the three priors on the shape of Section 10.4.4. What the readings do not share is the size of the risk. Whether string L should be replaced now, or held to its next test, weighs the probability 0.58 (0.41 exactly) of a failure in the coming year, with the consequence of a string that does not carry its share of the load when an outage comes, against the cost of a string replaced with remaining life; the probability 0.055 (0.222) that its one-year survival is at least 0.80 says that the record cannot show it to be a safe string to keep, and the predictive 0.425 (0.593) is the number the year's decision uses. The two predictives fall on opposite sides of one half, and a rule that turned on one half would give opposite answers on them; the chapter does not choose between them. That rule, replacement at an age or at a conditional survival, is the subject of the decision chapters of a later part of the book, and it will use the conditional predictive derived here. The spares follow from the fleet: 2.01 failures expected and a probability 0.613 of two or more, or exactly 1.51 and 0.456, against the 1.87 of the plug-in. The two routes fall on opposite sides of the plug-in, the state route calling for more spares than the plug-in and the exact one for fewer.

Principle 10.1. *The age a unit has reached is information about its next period in proportion to the slope of the hazard. Within a lifetime model the record fixes that slope only through the posterior of the shape, and between models only through their evidence ratio; the conditional survival, its band and the replacement ranking must be averaged over that posterior, not read at the fitted shape.*

Common mistake

Three errors recur with remaining life. Reporting the plug-in conditional survival as the probability that a string lasts the year: for string L it gives 0.44 where the state route gives 0.42 and the exact posterior 0.59, and for string E 0.994 where they give 0.988 and 0.978; the plug-in uses one steep shape and ignores what four failures leave uncertain. Multiplying the strings' predictive survivals to answer a joint question: 0.074 against 0.125 on the state route, and 0.170 against 0.235 exactly, for no failure in the fleet, because the strings share one uncertain state. And applying the ratio of predictives (10.13) to a string already in the record, which conditions on its survival twice: for string L it gives 0.501 against 0.425 on the state route, and 0.650 against 0.593 exactly.

10.6 The next measurement (step 7)

The bands of Table 10.17 are wide, and the next annual test, which the program takes anyway, adds rows to the record. The preposterior values the test and its variants before they are taken, by the width each would take off the band of string L's one-year survival.

10.6.1 The candidate tests

The test of survivor j has two outcomes: the string passes, or it fails within the year, with the probability

$$F_j(\mathbf{z}) = 1 - \exp\big[-(\Lambda(t_j + 1; \mathbf{z}) - \Lambda(t_j; \mathbf{z}))\big],$$

the complement of its conditional survival (10.2); the failure probability F_j of Section 10.3.4 is its posterior mean. The test records which, not when within the year a failure occurred, so each of the eight added rows is a count of one trial, $n_j = 1$ for a failure and 0 for a pass, in the

binomial working form of equation (4.43) with one trial, here with the failure probability F_j a function of the state rather than carried on its logit:

$$r_j = \frac{F_j - n_j}{\sqrt{F_j(1-F_j)}}, \qquad \mathbf{J}_{a,j} = \frac{\nabla F_j^{\mathsf{T}}}{\sqrt{F_j(1-F_j)}}, \qquad \nabla F_j = (1-F_j)\big[\nabla\Lambda(t_j+1) - \nabla\Lambda(t_j)\big], \tag{10.38}$$

with $\nabla\Lambda$ the survivor-row gradient of Section 10.2.1. As for a count row, $\mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a$ is the expected information of the test, the curvature of the eight rows averaged over the outcomes the fit predicts; the preposterior of Chapter 2, equation (2.17), adds it to the present information, $\boldsymbol{\Sigma}_a = (\boldsymbol{\Sigma}^{-1} + \boldsymbol{\Lambda}_a)^{-1}$ with $\boldsymbol{\Lambda}_a = \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a$.

The test rows at the present fit. At $\hat{\mathbf{z}}$ the failure probabilities $F_j(\hat{\mathbf{z}})$ are the complements of the plug-in column of Table 10.7, from 0.006 for E to 0.561 for L, and the 8×2 Jacobian of the test rows is, strings E to L from left to right,

$$\mathbf{J}_a^{\mathsf{T}} = \begin{bmatrix} -0.355 & -0.750 & -1.467 & -1.886 & -2.396 & -2.757 & -3.029 & -3.442 \\ -0.371 & -0.500 & -0.450 & -0.310 & -0.049 & 0.197 & 0.422 & 0.854 \end{bmatrix}.$$

Its product and the preposterior precision are

$$\boldsymbol{\Lambda}_a = \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a = \begin{bmatrix} 40.76 & -2.89 \\ -2.89 & 1.64 \end{bmatrix}, \qquad \boldsymbol{\Sigma}^{-1} + \boldsymbol{\Lambda}_a = \begin{bmatrix} 130.09 & 6.11 \\ 6.11 & 9.13 \end{bmatrix},$$

and $\boldsymbol{\Sigma}_a$ has $\sqrt{\Sigma_{a,uu}} = 0.089$ and $\sqrt{\Sigma_{a,vv}} = 0.336$, against 0.113 and 0.390 now. The fit itself treats a failure found at a test as a lifetime ending at the test; with that convention, the row Hessians of equation (7.36) averaged over the predicted outcomes give $\boldsymbol{\Lambda}_a = [53.61,\ -7.45;\ -7.45,\ 1.45]$ and 0.084 and 0.335: the time of a failure would add to the scale, hardly to the shape.

The effect of the year on the shape. The year would narrow the scale by a fifth and the log-shape by a seventh. The columns of $\mathbf{J}_a$ show why. The u column is negative for every string and largest for the old ones, whose failure probabilities are large; its sum of squares, 40.76, is the (u,u) entry. The v column carries the lever of each string's coming year: it is negative for the young strings, near zero for I, whose coming year sits near the lever's change of sign, and positive for J, K and L, and it is small for every string, at most 0.85 for L; the (v,v) entry, 1.64, is about a fifth of the present 7.49. The eigenvectors of $\boldsymbol{\Lambda}_a$ say the same in one line: the test adds 40.98 along $(0.997,\ -0.073)^{\mathsf{T}}$, nearly pure log-scale, and 1.42 along $(0.073,\ 0.997)^{\mathsf{T}}$, nearly pure log-shape.

10.6.2 The preposterior and the outcome

The answer whose band the next measurement should narrow is string L's one-year survival. Its double logarithm (10.18) has the gradient $\mathbf{g}_L = (-4.726,\ 1.173)^{\mathsf{T}}$ at $\hat{\mathbf{z}}$ (Section 10.3.2) and the standard deviation $\sqrt{\mathbf{g}_L^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_L} = 0.815$ now; each candidate replaces $\boldsymbol{\Sigma}$ by its $\boldsymbol{\Sigma}_a$. Table 10.18 lists the candidates. Each row is a candidate; the columns give the number of added rows, the two standard deviations of the state under $\boldsymbol{\Sigma}_a$, the standard deviation $\sqrt{\mathbf{g}_L^{\mathsf{T}}\boldsymbol{\Sigma}_a\mathbf{g}_L}$, the Laplace band on string L's one-year survival that it gives through equation (10.20), and the width that band loses against the present [0.017, 0.846], of width 0.829. A test of one string adds one column of $\mathbf{J}_a^{\mathsf{T}}$, a rank-one $\boldsymbol{\Lambda}_a$; the tests of several add their columns. Two half-year tests give each string two rows of the form (10.38), one for each half-year, the second weighted by the predicted probability of passing the first, since a string that fails in the first half is not tested again; two annual tests do the same over two years.

Table 10.18: The candidate tests ranked by the preposterior $\boldsymbol{\Sigma}_a = (\boldsymbol{\Sigma}^{-1} + \boldsymbol{\Lambda}_a)^{-1}$ at $\hat{\mathbf{z}}$, with $\boldsymbol{\Lambda}_a$ the expected information of the pass/fail rows (10.38). Last three columns: string L's one-year survival through the delta method on ϑ_L, and the width of its Laplace band removed.

Candidate	Rows [–]	$\sqrt{\Sigma_{a,uu}}$ [–]	$\sqrt{\Sigma_{a,vv}}$ [–]	$\sqrt{\mathbf{g}_L^{\mathsf{T}}\boldsymbol{\Sigma}_a\mathbf{g}_L}$ [–]	Laplace band of $R(1 \mid 6.2)$ [–]	Width removed [–]
Present record	–	0.113	0.390	0.815	[0.017, 0.846]	–
Test of E alone	1	0.113	0.386	0.812	[0.018, 0.846]	0.001
Test of K alone	1	0.105	0.376	0.751	[0.028, 0.828]	0.029
Test of L alone	1	0.102	0.357	0.701	[0.039, 0.812]	0.056
Tests of J, K and L	3	0.093	0.347	0.638	[0.056, 0.790]	0.096
Annual test of the eight	8	0.089	0.336	0.626	[0.060, 0.786]	0.104
Two half-year tests of the eight	16	0.089	0.335	0.624	[0.061, 0.785]	0.105
Two annual tests of the eight	16	0.077	0.291	0.497	[0.113, 0.733]	0.210

The information of a test row grows with the string's probability of failing the test. The young strings' rows add almost nothing, 0.001 for E; string L's row alone removes 0.056, more than half of what the whole test removes, and the three oldest strings together remove 0.096 of its 0.104. Testing twice a year removes 0.001 more than testing once: the half-year in which a string fails says little about the shape that the year does not. A second annual test doubles the width removed, to 0.210, and takes the log-shape to 0.291. The preposterior is computed on the Laplace Gaussian, as the answers are: the widths are those of the reported band, not of the exact band [0.099, 0.942] beside it, and the chapter has no exact preposterior to set beside them. It ranks the candidates; Chapter 20, in Volume II, writes the same preposterior for a candidate row as a rank-one update of $\boldsymbol{\Sigma}$.

The outcome against the expectation. Once the test is taken, the rows of its outcome have their own curvature. If all eight pass, each row is the cumulative hazard of the year, and its curvature at $\hat{\mathbf{z}}$ is the difference of two survivor-row Hessians, equation (7.36) with $\delta = 0$; summed over the eight strings,

$$\nabla^2\mathcal{L}_{\text{pass}} = \begin{bmatrix} 53.61 & -16.31 \\ -16.31 & 0.92 \end{bmatrix}, \qquad \nabla^2\mathcal{L}_{\text{pass}} - \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a = \begin{bmatrix} 12.85 & -13.41 \\ -13.41 & -0.71 \end{bmatrix}.$$

The first matrix is the observed curvature of one outcome, not the information of the test; the second is its curvature term, observed minus expected, which averages to zero over the outcomes the fit predicts. The observed curvature has the eigenvalues 58.25 along $(0.962,\ -0.274)^{\mathsf{T}}$ and -3.71 along $(0.274,\ 0.962)^{\mathsf{T}}$, negative in the shape-heavy direction. The sign comes from the survivor rows: by equation (7.36) with $\delta = 0$, the Hessian of one survivor row has the determinant $-\beta^2\Lambda^2(1 + \beta x)$, which is negative once $\Lambda > \mathrm{e}^{-1}$, past $\hat{\eta}\mathrm{e}^{-1/\hat{\beta}} = 5.26$ years, so the rows of the strings that pass beyond that age are indefinite, and a pass of an old string flattens the objective along the shape. Used in place of the information, this one outcome would give 0.086 and 0.353. The outcome also moves $\hat{\mathbf{z}}$: a year with no failure would refit to $\hat{\beta} = 3.84$ and $\hat{\eta} = 7.41\,\text{yr}$, with $\sqrt{\Sigma_{vv}} = 0.415$, and the exact posterior of v would then have the standard deviation 0.528, wider than the present 0.489, since the record would have lost some of its evidence for wear-out: one outcome can widen the exact posterior of the shape while the expected information narrows it. The next test ranks well as a measurement of the scale; as a

measurement of the shape, which is what the ranking of the strings rests on, one year buys little, and it is the failures, if they come, that buy it.

10.6.3 The recommendation

While the program runs the annual test, take it on all eight strings as scheduled: that adds nothing to what the program already pays, and it removes 0.104 of the band's width of 0.829, taking $\sqrt{\mathbf{g}_L^{\mathsf{T}}\boldsymbol{\Sigma}_a\mathbf{g}_L}$ from 0.815 to 0.626. This values the test as a measurement and does not ask whether the test is worth its price, which this chapter does not have. Chapter 30, in Volume III, prices it: beside an age rule the test costs a net R$ 1 211 per string-year (R$ 958 on the exact route), and that decision is taken there. Do not add a half-year test for the sake of the lifetime record, since it removes 0.001 more. Replacing string L before its test gives up the largest single share, 0.056, which is one more item for the replacement decision of Section 10.5.2. The shape, on which the gaps between the strings rest, narrows with years and failures; to narrow it faster needs a row from outside the lifetime record, a proper prior on the shape from the manufacturer's life tests, which Section 10.4.4 shows would move the old strings' answers most, or the capacity readings that the record discards. When the eight rows arrive, the fit returns to step 3 with them, and the answers and the exact integrals beside them are recomputed.

10.6.4 What the analysis showed

The first reading of the record, the fitted Weibull taken as known, said that the oldest string would more likely than not fail within the year, that a year without a failure among the eight was rare, and that the strings should go in order of age; the analysis, on the posterior of the state, keeps string L more likely than not to fail (0.575), a clean year about one time in eight, and the order of age supported, but only modestly. The exact posterior beside it finds string L more likely than not to last the year (0.593) and a clean year about one time in four; the two fall on opposite sides of one half. Four failures between the third and the seventh year fix the scale of the strings' lifetime given its shape, through the number of failures, and leave the shape itself poorly known: a shallower shape with a longer scale fits the record almost as well, and it is those shallower drawings, which the fitted point omits and a normal at the fitted point underweights, that give the old strings more life and the young ones less on the exact posterior. The motors of Chapter 7 sat on the other side of that line: there the shape could not be told from one, so a motor found running was no worse a risk than a new one, while here the shape lies above one and age is the first thing to look at. The strings are rows of one node, so what is learned or not learned about the shape moves all eight together, which is why the fleet's year is safer than the product of the strings' chances and its failure count more dispersed. For the site, the order of replacement is the one thing every reading agrees on, oldest first; the gaps between the strings, the number of spares and the band of the oldest string are not settled, and they move with what is assumed about the shape before the record. The annual test narrows the scale more than the shape. What would settle the shape is failures at the ages the old strings are reaching, a prior from the manufacturer's life tests of the same design, and the capacity readings that the lifetime record discards.

10.6.5 Limits of the model

The derivations are exact under their assumptions, and five of them need to be checked before a result is used.

Comparable strings. The likelihood multiplies twelve rows under one lifetime distribution. Strings installed at different dates have spent their lives under different room temperatures, loads and discharge histories, and battery life is sensitive to temperature; a string that has seen

more discharges is older than its age. A model in which the scale depends on such covariates replaces the one free node with a node per condition and rows that tie them.

Noninformative censoring. The survivors are censored by the calendar, at the latest test, which is administrative censoring. A string withdrawn because its test margin was falling, before it failed, would be a different observation, and treating its withdrawal as a survival would overstate the reliability (Section 7.1.6).

Independence given the state. The joint survival of the fleet multiplies the strings' conditional survivals for a known state. A common cause in the year, a charger fault or an overheated room, fails several strings at once, and the fleet's chance of a clean year is then lower than the shared state alone makes it; Chapter 15 writes such a cause as an entry of the state that the strings share.

The failure ages and the prior. Two choices of the model are measured rather than assumed away. Reading a failed test as a failure at the test age makes the hazard steeper than the interval rows do (Section 10.4.5), and the flat prior on the shape leaves the old strings' answers to four failure ages, where a proper prior moves them substantially (Section 10.4.4); it also makes the unrestricted predictive mean residual life infinite and leaves the restricted one dependent on its horizon. A prior on the shape from the manufacturer's life tests or from other fleets of the same design would close that tail.

Lifetimes only. The capacity measured at each test is a degradation reading that the lifetime record discards. A string whose capacity has fallen close to the requirement is closer to failure than its age says, and one whose capacity is high is further from it; a degradation model would give each string its own remaining life rather than the fleet's at its age.

The next chapter turns from whether each unit survives a period to how many failures a fleet produces in one: the posterior predictive distribution of a future failure count.

10.7 What to retain

Definition 10.1 (Conditional survival under an uncertain lifetime model)**.** For a known state, a unit of age t_0 survives a further τ with probability $R(\tau \mid t_0; \mathbf{z}) = R(t_0+\tau; \mathbf{z})/R(t_0; \mathbf{z}) = \exp[-\int_{t_0}^{t_0+\tau} \lambda]$, which changes with age by minus the change of the hazard across the horizon. When the unit's survival to t_0 is a row of the record $\mathbf{y}$, the probability that it survives the further period is the posterior mean $\mathbb{E}[R(\tau \mid t_0; \mathbf{z}) \mid \mathbf{y}]$, equal to the ratio of the predictive survivals at $t_0+\tau$ and at t_0 under the posterior without its row; for a unit outside the record the ratio is taken under the posterior with all rows. For the Weibull model under a flat prior in $(\ln\eta, \ln\beta)$ it is $\int p(v \mid \mathbf{y})\,(A/(A+\nu))^k \,\mathrm{d}v$ with $A = \sum_i t_i^\beta$ and $\nu = (t_0+\tau)^\beta - t_0^\beta$; the joint survival of several units that share the state is the same integral with their increments added, and the mean residual life is the integral of the conditional survival over the horizon.

Four conclusions carry beyond the strings. The age of a unit is information about its future only in proportion to the slope of its hazard, so the posterior of the shape, and between lifetime models their evidence ratio, decide how much the age is worth; with few failures the prior on the shape is part of that posterior. The survival of a unit to its age is a row of the record, used once: in the posterior for a member of the record, in the ratio for a unit outside it. On the exact posterior, averaging over an uncertain steep shape flattens the conditional survival, raising it for old units and lowering it for young ones, and widens its band; the plug-in fit ranks the units in the right order with the wrong gaps. And units that share the state must be averaged together: the expectation of a product is not the product of expectations, and the fleet's chance of a year

without failure, and the spread of its failure count, depend on it. In the state $\mathbf{z} = (u, v)^{\mathsf{T}}$ every one of these quantities has a gradient built from the gradients of survivor rows, so its band is $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ with the 2×2 $\boldsymbol{\Sigma}$ of the Newton fit, and every answer is computed from that posterior; with four failures the exact one-integral route, set beside it with the gap, differs enough to change which side of one half the oldest string falls on, and both are reported. A test not yet taken is ranked by its expected information, $\mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a$ of its pass/fail rows; the curvature of one outcome, such as every string passing, is observed, not expected, and it can be indefinite, since a survivor row curves the objective down along the shape once its cumulative hazard exceeds e^{-1}. Table 10.19 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 10.19: The terms of remaining life and conditional survival and where each is defined.

Term	Expression	Units	Meaning	Defined in
Conditional survival	$R(\tau \mid t_0; \mathbf{z}) = R(t_0+\tau)/R(t_0)$	–	survive τ more, known state	10.1.2
Hazard over the horizon	$\exp[-\int_{t_0}^{t_0+\tau} \lambda(s)\,\mathrm{d}s]$	–	only the coming hazard counts	10.1.2
Dependence on age	$\partial_{t_0} \ln R(\tau \mid t_0) = -[\lambda(t_0+\tau) - \lambda(t_0)]$	yr^{-1}	falls under wear-out	10.1.2
Weibull increment	$\nu = (t_0+\tau)^\beta - t_0^\beta$, $R = \mathrm{e}^{-\nu\zeta}$	yr^β	$\zeta = \eta^{-\beta}$	10.1.2
General forms				
Prediction	$\int p(y_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\,\mathrm{d}\mathbf{z}$	–	one string, or the fleet's intersection	10.3.1
Estimation and threshold	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$, $\Pr(f(\mathbf{z}) \ge f^* \mid \mathbf{y})$	–	band and tail of a derived quantity	10.3.2
Expectation and quantile	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$, $\Pr(f \le s_q \mid \mathbf{y}) = q$	–	remaining life	10.3.3
Comparison	$B_{10} = p(\mathbf{y} \mid M_1)/p(\mathbf{y} \mid M_0)$	–	Weibull against exponential	10.3.5
Specialized forms and answers				
Survival row	$p(\mathbf{z} \mid \mathbf{y}) \propto R(t_j)\, p(\mathbf{z} \mid \mathbf{y}_{-j})$	–	survival is evidence	10.3.1
Conditional predictive	$\mathbb{E}[R(\tau \mid t_j; \mathbf{z}) \mid \mathbf{y}]$	–	member of the record	10.3.1
Ratio of predictives	$\int R(t_0+\tau)p / \int R(t_0)p$	–	same, or a new unit	10.3.1
Exact distribution	$\int p(v \mid \mathbf{y})\, \Upsilon_k(-A \ln R^*/\nu)\,\mathrm{d}v$	–	band and threshold	10.3.2
Double logarithm	$\vartheta = \ln[\Lambda(t_0+\tau) - \Lambda(t_0)]$, $\mathbf{g} = (-\beta,\ \beta\bar{x})$	–	delta method, $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$	10.3.2
Mean residual life	$\mathrm{MRL}(t_0) = \int_0^\infty R(s \mid t_0)\,\mathrm{d}s$	yr	expected remaining life	10.3.3
Weibull mean residual life	$(\eta/\beta)\mathrm{e}^{\Lambda(t_0)}\Gamma(1/\beta, \Lambda(t_0))$	yr	incomplete gamma	10.3.3
Restricted mean residual life	$\int_0^H \Pr(T_j > t_j + s \mid \mathbf{y})\,\mathrm{d}s$	yr	finite under a flat prior	10.3.3
Remaining-life quantile	$\Lambda(t_0+q_p) - \Lambda(t_0) = -\ln(1-p)$	yr	median at $p = 1/2$	10.3.3
Fleet survival	$\int p(v \mid \mathbf{y})\,(A/(A+\sum_j \nu_j))^k\,\mathrm{d}v$	–	shared state, averaged once	10.3.4
Expected failures	$\sum_j F_j$	–	linear, no independence needed	10.3.4
Variance of the count	$\sum_j F_j(1-F_j) + \sum_{i\neq j} \mathrm{Cov}(\mathbf{1}_i, \mathbf{1}_j)$	–	pair terms of either sign	10.3.4
Evidence ratio	$B_{10} = \int \Omega(v)\pi(v)\,\mathrm{d}v / \Omega(0)$	–	Savage–Dickey, inverted	10.3.5
Building blocks				
Working rows	$r_i = (\Lambda_i - \delta_i)/\sqrt{\Lambda_i}$, $\mathbf{J}_i = \beta\sqrt{\Lambda_i}(-1, x_i)$	–	12×2 Jacobian	10.2.1
Survivor-row gradient	$\nabla\Lambda(t) = (-\beta\Lambda,\ \beta x\Lambda)$	–	gradients of every answer	10.2.1
Information	$\boldsymbol{\Lambda} = \mathbf{J}^\mathsf{T}\mathbf{J} + \mathrm{diag}(0, k)$ at $\hat{\mathbf{z}}$	–	$\hat{\beta}^2 k$ in (u, u)	10.2.1
Exact given the shape	$(A/(A+\nu))^k$	–	gamma Laplace transform	10.2.2
Test row	$\mathbf{J}_{a,j} = \nabla F_j^\mathsf{T} / \sqrt{F_j(1-F_j)}$	–	expected information $\mathbf{J}_a^\mathsf{T}\mathbf{J}_a$	10.6.1
Preposterior	$\boldsymbol{\Sigma}_a = (\boldsymbol{\Sigma}^{-1} + \boldsymbol{\Lambda}_a)^{-1}$	–	value of the next test	10.6.1
Curvature term	$\nabla^2\mathcal{L}_{\text{pass}} - \mathbf{J}_a^\mathsf{T}\mathbf{J}_a$	–	observed minus expected, one outcome	10.6.2

10.8 Exercises

Exercise 10.1. *Formulation.* A battery string has a Weibull lifetime with $\eta = 8\,\text{yr}$ and $\beta = 3$, both known. (a) Compute its probability of surviving the next year at ages 1 and 5 years. (b) Show that under an exponential lifetime the same probability does not depend on the age. (c) For string J of the chapter's record, at 5.1 yr, with the state uncertain, write the question as a probability, say which form of the conditional predictive applies, and give its value. ((a) $\exp[-(2^3-1)/512] = \mathrm{e}^{-0.0137} = 0.986$ and $\exp[-(6^3-5^3)/512] = \mathrm{e}^{-0.1777} = 0.837$. (b) $\mathrm{e}^{-\lambda(t_0+\tau)}/\mathrm{e}^{-\lambda t_0} = \mathrm{e}^{-\lambda\tau}$, free of t_0. (c) $\Pr(T_J > 6.1 \mid \mathbf{y})$; its survival to 5.1 years is a row of $\mathbf{y}$, so it is the posterior mean $\mathbb{E}[R(1 \mid 5.1; \mathbf{z}) \mid \mathbf{y}] = 0.733$, equation (10.15), and not the ratio of predictives under $p(\mathbf{z} \mid \mathbf{y})$, which would give 0.745.)

Exercise 10.2. *Formulation.* For the fitted strings, $\hat{\eta} = 6.497\,\text{yr}$ and $\hat{\beta} = 4.726$, find the age at which the plug-in one-year survival falls below 0.90, and state which survivors are past it. At $t_0 = 4\,\text{yr}$, write the gradient of the double logarithm (10.18) as the difference of the gradients of two survivor rows and compute σ_ϑ with $\boldsymbol{\Sigma}$ of equation (10.8). (The condition $\exp[-((t_0+1)^\beta - t_0^\beta)/\eta^\beta] = 0.9$ is $(t_0+1)^{4.726} - t_0^{4.726} = 0.10536 \times 6930.5 = 730.2$, whose root is $t_0 = 3.34\,\text{yr}$; the exact posterior average crosses 0.90 at 3.33 yr. Strings H to L, at 3.9 years and older, are past it; G, at 3.3 years, is just short of it, with a plug-in value of 0.903 and a posterior average of 0.902. At four years $\nabla\Lambda(5) = (-1.3704,\ -0.3590)^\mathsf{T}$, $\nabla\Lambda(4) = (-0.4774,\ -0.2316)^\mathsf{T}$ and the cumulative hazard of the year is 0.1890, so $\mathbf{g} = (-0.8930,\ -0.1274)^\mathsf{T}/0.1890 = (-4.726,\ -0.674)^\mathsf{T}$, $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g} = 0.256$ and $\sigma_\vartheta = 0.506$, close to the floor $1/\sqrt{k} = 0.5$ because the lever $\bar{x} = -0.143$ is close to $\bar{x}^* = -0.101$.)

Exercise 10.3. *Derivation.* A unit outside a record $\mathbf{y}$ is found working at age t_0. Derive, from Bayes' theorem with its survival as the evidence, the probability that it survives a further τ, and show that applying the same formula to a unit whose survival is already in $\mathbf{y}$ counts that survival twice. (The new row gives $p(\mathbf{z} \mid \mathbf{y}, T > t_0) = R(t_0; \mathbf{z})p(\mathbf{z} \mid \mathbf{y})/\int R(t_0)p$; averaging $R(t_0+\tau)/R(t_0)$ over it cancels $R(t_0)$ and leaves $\int R(t_0+\tau)p(\mathbf{z} \mid \mathbf{y})\,\mathrm{d}\mathbf{z}/\int R(t_0)p(\mathbf{z} \mid \mathbf{y})\,\mathrm{d}\mathbf{z}$. If the survival is already a row of $\mathbf{y}$, $p(\mathbf{z} \mid \mathbf{y})$ contains $R(t_0)$ once, and the formula multiplies by it again, the posterior of $\mathbf{y}$ with the row doubled.)

Exercise 10.4. *Derivation.* (a) Derive the mean residual life of the exponential lifetime and show that it does not depend on age. (b) Derive the mean residual life of a Weibull lifetime with $\beta = 2$ at age zero, and evaluate it for $\eta = 6.5\,\text{yr}$. (c) Under the posterior $\lambda \sim \text{Gamma}(a, b)$ derive the predictive mean residual life of an exponential unit of the record, and evaluate it for the strings. ((a) $\int_0^\infty \mathrm{e}^{-\lambda s}\,\mathrm{d}s = 1/\lambda$, since $R(s \mid t_0) = \mathrm{e}^{-\lambda s}$ for every t_0. (b) $\int_0^\infty \mathrm{e}^{-(t/\eta)^2}\,\mathrm{d}t = \eta\,\Gamma(3/2) = \eta\sqrt{\pi}/2$, which is 5.76 yr. (c) $\mathbb{E}[1/\lambda] = \int \lambda^{-1} b^a \lambda^{a-1}\mathrm{e}^{-b\lambda}/\Gamma(a)\,\mathrm{d}\lambda = b\,\Gamma(a-1)/\Gamma(a) = b/(a-1)$ for $a > 1$; with $\text{Gamma}(4, 51.9\,\text{yr})$, $51.9/3 = 17.3\,\text{yr}$, against the plug-in $1/\hat{\lambda} = 13.0\,\text{yr}$.)

Exercise 10.5. *Application.* Three strings of the same design in the same room, not in the record, are in service at ages 1, 3 and 5 years. Using the exact posterior of the chapter, compute the probability that all three last the coming year, compare it with the product of their marginal conditional predictives, and explain the size of the difference. (All three: $\mathbb{E}[R(2)R(4)R(6) \mid \mathbf{y}]/\mathbb{E}[R(1)R(3)R(5) \mid \mathbf{y}] = \int p(v \mid \mathbf{y})(A/(A + 2^\beta + 4^\beta + 6^\beta))^4\,\mathrm{d}v/\int p(v \mid \mathbf{y})(A/(A + 1 + 3^\beta + 5^\beta))^4\,\mathrm{d}v = 0.690$. Marginals by equation (10.16): 0.982, 0.921 and 0.755, product 0.683; plug-in product 0.622. The difference is small because the five-year string carries most of the failure probability and the other two are nearly certain to survive, so there is little joint variation for the shared state to correlate.)

Exercise 10.6. *Application.* The eight survivors are tested again in a year and all pass. (a) Write the pass/fail row of string K, at 5.5 yr, as a working residual and give its Jacobian row and

its contribution to the expected information at the present fit. (b) Sum the contributions over the eight strings, compute the preposterior standard deviation of v, and compare it with the curvature of the outcome "all pass" and with the posterior after refitting the record with every survivor one year older. ((a) The survivor-row gradients at 6.5 and 5.5 years are $(-4.735,\ 0.002)^{\mathsf{T}}$ and $(-2.150,\ -0.358)^{\mathsf{T}}$; their difference is $(-2.585,\ 0.360)^{\mathsf{T}}$ and the year's cumulative hazard is 0.5470, so $F_K = 0.4213$, $\nabla F_K = 0.5787\,(-2.585,\ 0.360)^{\mathsf{T}} = (-1.496,\ 0.209)^{\mathsf{T}}$ and, with $\sqrt{F_K(1-F_K)} = 0.4938$, the Jacobian row of equation (10.38) is $(-3.029,\ 0.422)$; its contribution $\mathbf{J}_{a,K}^{\mathsf{T}}\mathbf{J}_{a,K}$ has the entries 9.18, -1.28 and 0.18. (b) $\mathbf{\Lambda}_a = \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a = [40.76,\ -2.89;\ -2.89,\ 1.64]$, and $(\mathbf{\Sigma}^{-1} + \mathbf{\Lambda}_a)^{-1}$ has $\sqrt{\Sigma_{a,vv}} = 0.336$, against 0.390 now. The curvature of the outcome "all pass", $[53.61,\ -16.31;\ -16.31,\ 0.92]$, is indefinite and would give 0.353; it is the curvature of one outcome, not the information. Refitted, $\hat{\beta} = 3.84$, $\hat{\eta} = 7.41$ yr, the Laplace $\sqrt{\Sigma_{vv}} = 0.415$ and the exact posterior standard deviation of v is 0.528: the year without failures moves the shape toward shallow values, where the posterior is wider, which the preposterior at the present fit cannot see.)

Chapter 11

Predicting Future Failure Counts

11.1 Failure counts over a future period

A data center's cooling is carried by forty computer-room air-conditioning units, and every failure of a unit's compressor takes a spare compressor from the store and a crew shift from the maintenance plan. The planner who stocks the store for next year has three years of record: seven, nine and eight compressor failures in 36, 40 and 40 unit-years of service. Next year forty-four units will run. The first reading of the record is simple: the record, with the handbook figure for compressors of this class folded in, puts the rate at its posterior mean of about 0.206 failures per unit-year (the record alone gives $24/116 = 0.207$), forty-four unit-years at that rate is 9.08 failures, and a Poisson count of mean 9.08 has a variance of 9.08, a standard deviation of 3.01, and stays at or below fourteen with probability 0.956. The plan is fourteen spare compressors for 95 % cover.

The reliability engineer who reviews the plan asks something else. The rate behind the 9.08 is known from twenty-four failures, not given, and a Poisson count whose variance equals its mean is the count of a fleet whose rate is known exactly. The number of failures next year is uncertain for two reasons at once: failures occur at random even when the rate is known, and the rate is not known. The first reading keeps the first reason and drops the second. The question of this chapter is

how widely the count of next year's failures spreads when the fleet's rate is known only from its record, and how many spares that spread requires.

It is answered with the method of Chapter 6, a log-rate fitted to count rows. What changes is the quantity asked about: not the rate but a count that does not exist yet, whose distribution, tails, quantiles and joint events with other counts are all functions of one posterior. The first reading, fourteen spares at a variance equal to the mean, is worth keeping in mind.

Jargon: computer-room air-conditioning (CRAC) unit, unit-year, posterior predictive distribution, plug-in prediction, negative-binomial distribution, overdispersion, horizon, spares level, service level, shortfall, shared rate, law of total covariance, posterior predictive check, leave-one-out (cross-validated) predictive check, predictive p-value, held-out year, Gauss–Hermite quadrature, digamma and trigamma functions, Poisson floor

A *computer-room air-conditioning (CRAC) unit* is a packaged cooling unit in a data hall: a fan, a coil and, in the units of this chapter, a compressor of its own. A *unit-year* is one unit in service for one year, the unit of exposure of a fleet. The *posterior predictive distribution* of a future count is its distribution given the record, with the unknown rate averaged out; the *plug-in prediction* replaces the rate by an estimate and treats it as known. The *negative-binomial distribution* is the distribution of a Poisson count whose mean is itself gamma distributed; a count whose variance exceeds its mean is *overdispersed*. The *horizon* is the exposure over which a count is predicted. A *spares level* is a number of spares held at the start of the horizon; its *service level* is the probability that it covers every failure of the horizon, and its *shortfall* is the number of failures that find no spare. A *shared rate* is one unknown rate that governs several units, so that what is learned about it from one unit applies to all. The *law of total covariance* splits a covariance into the expected conditional covariance and the covariance of the conditional means. A *posterior predictive check* compares an observed count with its predictive distribution. In a *leave-one-out (cross-validated) predictive check* the predictive is built from the other years only, so that the year checked, the *held-out year*, is not used twice; the check is called a held-out check in this chapter. Its *predictive p-value* is the probability, under that predictive, of a count at least as extreme as the one observed, a tail probability. *Gauss–Hermite quadrature* evaluates an expectation over a normal distribution as a weighted sum at a few fixed nodes. The *digamma function* $\psi(n)$ and the *trigamma function* $\psi'(n)$ are the first and second derivatives of $\ln\Gamma(n)$. The *Poisson floor* is the variance a future count keeps when its rate is known exactly.

Key idea

A future failure count is a Poisson count averaged over the posterior of its rate, and under a gamma posterior it is negative binomial. Its variance is the Poisson variance, which no record removes, plus the square of the horizon times the variance of the rate, a part that grows faster than the first as the horizon lengthens. Counts that share one uncertain rate are positively correlated, so the probability of a joint event is not the product of the marginals. The predictive tail of an observed count checks the model; an evidence ratio compares it with a stated alternative; the two answer different questions.

11.1.1 The count a plan needs and its two parts of spread

Chapter 6 estimated a failure rate, and its questions were about the rate. A maintenance plan asks about a count: a store of spare compressors is sized by how many compressors will be needed next year, a crew plan by how many repairs, a budget by how many repairs at what price each, all questions about the number of failures in a stated future exposure. The spread of that number has the aleatory and the epistemic parts of the law of total variance of Chapter 3. The Poisson part belongs to the future: a fleet with a rate of exactly 0.2 failures per unit-year and 44 unit-years of exposure still has a count with standard deviation $\sqrt{8.8} = 2.97$, no record removes it, and it sets a floor on how well any count can be predicted. The epistemic part enters through the horizon, an error of δ in the rate moving the mean count by δ times the exposure; it shrinks as the record grows, grows as the horizon lengthens, couples the counts of

units that share the rate, and is the part a record can narrow.

11.1.2 The physical system and its graph (step 0)

The data center has forty CRAC units in service, twelve of them in hall B and twenty-eight in the other halls, and four more are being commissioned in those other halls, so that next year forty-four units will be in service, twelve in hall B and thirty-two elsewhere. Each unit removes heat from the air of its hall into the chilled-water loop. The failures counted in this chapter are compressor failures: each takes the unit out of service until its compressor is replaced from the store of spares, after which the unit returns to service as before. The exposure is counted in unit-years in service.

The fleet's graph is drawn in Figure 11.1. The gray boxes on the left and right are the data halls and the chilled-water loop, pinned boundaries of the heat path drawn as solid arrows. The green box in the middle is the fleet, the free node, which carries the one entry of the state, the log failure rate ω. The gray box above it is the exposure of the record, pinned because the unit-years in service are read from the commissioning and outage log. The four dashed arrows below leave the fleet for the evidence: the prior on the rate and the three yearly counts. The box at the upper right is not evidence. It is the question's extension of the graph: the coming year, with its pinned exposure of 44 unit-years and its count K_f, which is unknown and is the quantity the questions ask about.

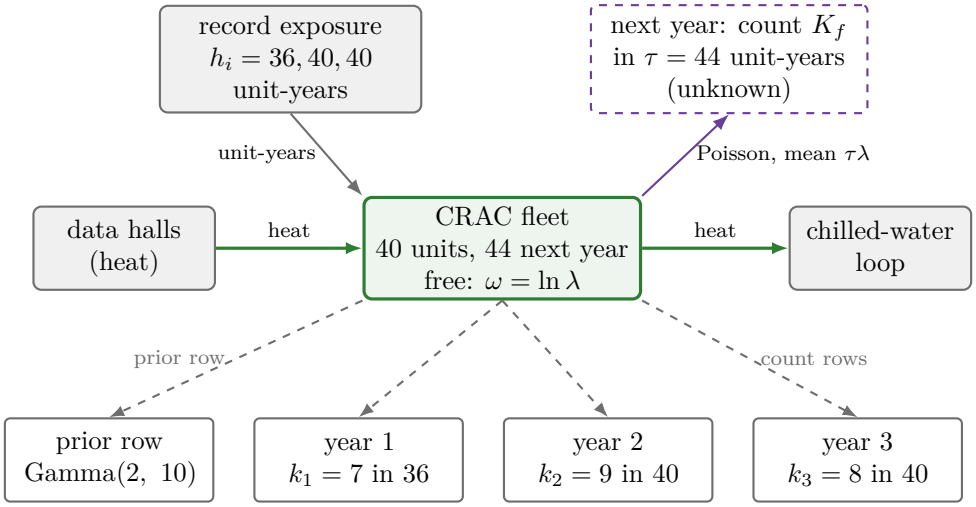


Figure 11.1: The plant graph of the CRAC fleet. Free: the fleet node, carrying the log failure rate $\omega = \ln \lambda$ with λ per unit-year (green). Pinned: the data halls and the chilled-water loop, the boundaries of the heat path, and the exposure of each year of the record (gray). Observed: the three yearly counts in their exposures and the prior on the rate (dashed arrows). The dashed purple box is the question's extension, the coming year's count in 44 unit-years, which carries no reading.

The state entry. The rate is carried on the logarithmic scale, as in Chapter 6:

$$\omega = \ln \frac{\lambda}{\lambda_{\text{ref}}}, \qquad \lambda_{\text{ref}} = 1 \text{ per unit-year}, \qquad \omega \in \mathbb{R}, \tag{11.1}$$

so that $\omega = \ln \lambda$ with the rate in failures per unit-year. The state is this one number. The general forms of the book write the state as a vector $\mathbf{z}$; here it has the single entry ω, so every quantity of the fit is a number: the objective's slope and curvature in ω, the information Λ, the posterior variance $\sigma_\omega^2 = 1/\Lambda$ of the log rate, and the derivative g of a derived quantity. The

rate, the count of a future period, the counts of the two halls and every quantity the questions ask for are derived quantities, functions of the state. The evidence $\mathbf{y}$ is the three yearly counts with their exposures and the prior.

11.1.3 The record of three years

Table 11.1 lists the record: each row is a year, the columns give its exposure h_i in unit-years, its failures k_i and the observed rate k_i/h_i. The first year has 36 unit-years because some units entered service during it. The three observed rates, 0.194, 0.225 and 0.200 per unit-year, are close to one another; Section 11.4.3 checks whether they scatter as counts of one rate would. The record is illustrative.

Table 11.1: The failure record of the CRAC fleet: compressor failures in unit-years of service.

Year i [–]	Exposure h_i [unit-years]	Failures k_i [–]	Observed rate [per unit-year]
1	36	7	0.194
2	40	9	0.225
3	40	8	0.200
Total	$T_{\text{tot}} = 116$	$k = 24$	0.207

The prior. Handbook data for compressors of this class put the rate near 0.2 failures per unit-year, with wide uncertainty. The prior is

$$\lambda \sim \text{Gamma}(a_0, b_0) = \text{Gamma}(2,\ 10 \text{ unit-years}), \tag{11.2}$$

a weak prior, worth two failures in ten unit-years. The gamma distribution is written with a shape and a rate parameter, the rate parameter carrying the units of exposure.

11.1.4 The probabilistic questions (step 1)

The planner asks seven questions of the record, and step 1 writes each as a probability or an expectation over the state, with the unknown on the left of the bar and the record on the right. Table 11.2 lists them: the first column is the question in engineering language; the second the general form that every question of its category takes, written in the book's notation for the posterior density $p(\mathbf{z} \mid \mathbf{y})$, a future observation y_{new}, a level f^* and an assurance κ; the third its mathematical form for the fleet; the fourth its category; and the last the subsection where it is answered. The first two and the fifth are predictions of future outcomes, which Chapter 2 files under prediction under a threshold: each is answered by a predictive distribution, an expectation over the posterior of the probability of a future count given the rate, the fifth for two counts that share a rate and their joint event. The third is a tail of the same predictive. The fourth is a design question, a spares level chosen before the year begins, answered by a quantile of the predictive. The sixth asks whether a year of the record was surprising under the model fitted to the others, a tail of a held-out predictive, and the seventh compares the model with a stated alternative by a ratio of evidences. Which further record would narrow the prediction most asks nothing new of the present record: it is the next measurement for these seven questions, answered in Section 11.6. The atlas of Appendix A lists the same forms.

11.1.5 The evidence as rows (step 2)

Step 2 attaches four rows to the single entry ω: the prior row, which carries the handbook gamma density (11.2), and three count rows, one per year, each the Poisson likelihood of k_i failures

Table 11.2: The questions asked of the fleet's record, the general form of their category, their mathematical form for the fleet, their category and where each is answered. K_f is the fleet's count in the coming year of $\tau = 44$ unit-years; K_B and K_R the counts of hall B (12 unit-years) and of the rest (32); K_3 the count of year 3 and $\mathbf{y}_{1:2}$ the prior with years 1 and 2; s a number of spares and q its service level.

Engineering question	General form	Mathematical form	Category	Answered in
How many failures will the fleet have next year, and within what range?	$\int p(y_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	$\Pr(K_f = j \mid \mathbf{y})$, $j = 0, 1, \ldots$	prediction under a threshold	11.3.1
How does the range grow with the planning horizon?	$\mathbb{E}[y_{\text{new}} \mid \mathbf{y}]$, $\mathrm{Var}(y_{\text{new}} \mid \mathbf{y})$	$\mathrm{Var}(K_f \mid \mathbf{y})$ against τ	prediction under a threshold	11.3.1
What is the probability of more than twelve? of more than fifteen?	$\Pr(y_{\text{new}} > f^* \mid \mathbf{y})$	$\Pr(K_f > 12 \mid \mathbf{y})$, $\Pr(K_f > 15 \mid \mathbf{y})$	prediction under a threshold	11.3.2
How many spare compressors cover next year with 95 % probability? with 99 %?	smallest f^* with $\Pr(y_{\text{new}} \le f^* \mid \mathbf{y}) \ge \kappa$	smallest s with $\Pr(K_f \le s \mid \mathbf{y}) \ge q$	design	11.3.2
How many failures in hall B, and are the halls' counts independent?	$\int p(\mathbf{y}_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$, its covariance	$p(K_B \mid \mathbf{y})$; $\mathrm{Cov}(K_B, K_R \mid \mathbf{y})$; $\Pr(K_B = 0, K_R = 0 \mid \mathbf{y})$	prediction under a threshold (joint event)	11.3.3
Was last year's count consistent with the two years before it?	$\Pr(y_{\text{new}} \ge f^* \mid \mathbf{y})$	$\Pr(K_3 \ge k_3 \mid \mathbf{y}_{1:2})$	prediction under a threshold (held-out check)	11.3.4
Did the rate double in year 3?	$p(\mathbf{y} \mid M_1)/p(\mathbf{y} \mid M_0)$	$\ln B_{10}$ of a doubled rate in year 3	hypothesis comparison	11.3.4

in h_i unit-years. None is Gaussian: each enters the objective as its negative log-probability term, Section 4.2.7, a function of the one number ω. What a row contributes to the posterior is its slope and its curvature in ω, written with primes, $\mathcal{L}'_k = \mathrm{d}\mathcal{L}_k/\mathrm{d}\omega$ and $\mathcal{L}''_k = \mathrm{d}^2\mathcal{L}_k/\mathrm{d}\omega^2$; each row below is given by its term, its slope and its curvature. The questions add quantities of a different kind, without readings: the coming year's count, Poisson with mean $\tau\lambda$, for the first four questions; the same year split into hall B with $\tau_B = 12$ and the rest with $\tau_R = 32$ unit-years, for the fifth; and, for the last two, year 3 removed from the evidence and treated as a prediction from the other two. Which row acts on which entry is a property of the graph, and here the graph has one free entry, so every row and every predicted count is a function of ω:

$$\begin{gathered}\mathcal{L}(\omega) = \mathcal{L}_0(\omega) + \mathcal{L}_1(\omega) + \mathcal{L}_2(\omega) + \mathcal{L}_3(\omega), \\ \mathbb{E}[K_f \mid \omega] = \tau\,\mathrm{e}^{\omega}, \qquad \mathbb{E}[K_B \mid \omega] = \tau_B\,\mathrm{e}^{\omega}, \quad \mathbb{E}[K_R \mid \omega] = \tau_R\,\mathrm{e}^{\omega},\end{gathered} \tag{11.3}$$

with $\mathcal{L}_0$ the prior row and $\mathcal{L}_i$ the count row of year i. Holding year 3 out deletes the term $\mathcal{L}_3$; the doubled-rate alternative changes the exposure of that row and leaves its entry. Every question is then an expectation over the one posterior of a probability that is known once λ is known.

The modeling choices. Three choices make the rows and the predicted counts, and each buys something and can fail. One constant rate is taken to hold over the three years of the record and the year to come: this is what makes the coming year a fourth year of the same process, its count a Poisson count at a rate drawn from the one posterior, and it fails if the rate drifts with the fleet's age, a drift that three yearly counts can hardly show (Sections 11.4.3 and 11.6.6). The four units being commissioned are given the rate of the forty that ran: this lets them borrow the whole record, and it fails if new compressors fail more often in their first months (Section 11.6.6). Failures of different units are taken as independent given the rate: this makes the count given the rate Poisson, its variance equal to its mean, and a cause that strikes several units at once breaks it (Section 11.6.6).

The rows as negative log-probability terms. The state is the one number ω of equation (11.1). Each of the four rows of Section 11.1.5 contributes the term $\mathcal{L}_k(\omega) = -\ln p(y_k \mid \omega)$ of its evidence y_k, equation (4.39), with the terms free of ω dropped. A year's count row is the negative log of the Poisson probability of k_i failures in h_i unit-years, $h_i\mathrm{e}^{\omega} - k_i\omega$. The gamma prior (11.2) written in the log rate, with the change-of-variable factor $\mathrm{d}\lambda/\mathrm{d}\omega = \lambda$ included as in Chapter 6, equation (6.24), is the prior row $b_0\mathrm{e}^{\omega} - a_0\omega$, with the shape a_0 rather than $a_0 - 1$. All four rows have one form, a count n_k against an expected count μ_k that is linear in the state on the log scale:

$$\begin{gathered}\mathcal{L}_k(\omega) = \mu_k(\omega) - n_k \ln \mu_k(\omega), \qquad \mu_k(\omega) = t_k\,\mathrm{e}^{\omega}, \\ \begin{array}{lll} k=0: & \text{prior row,} & (n_0, t_0) = (2,\ 10), \\ k=1: & \text{count row, year 1,} & (n_1, t_1) = (7,\ 36), \\ k=2: & \text{count row, year 2,} & (n_2, t_2) = (9,\ 40), \\ k=3: & \text{count row, year 3,} & (n_3, t_3) = (8,\ 40), \end{array}\end{gathered} \tag{11.4}$$

with n_k the count of the row (the prior's shape acting as a count of pseudo-failures) and t_k its exposure in unit-years (the prior's rate parameter, the year's unit-years in service). Summing the four terms gives the objective $\mathcal{L} = \sum_k \mathcal{L}_k$,

$$\mathcal{L}(\omega) = \underbrace{10\,\mathrm{e}^{\omega} - 2\omega}_{\text{prior row}} + \underbrace{36\,\mathrm{e}^{\omega} - 7\omega}_{\text{year 1}} + \underbrace{40\,\mathrm{e}^{\omega} - 9\omega}_{\text{year 2}} + \underbrace{40\,\mathrm{e}^{\omega} - 8\omega}_{\text{year 3}} = b\,\mathrm{e}^{\omega} - a\,\omega, \tag{11.5}$$

with $b = \sum_k t_k$ and $a = \sum_k n_k$, up to a constant.

The slopes and the curvatures. Each row is the count row of Section 4.2.7, equation (4.42), with $\ln \mu_k = \ln t_k + \omega$, whose slope in ω is one for every row. Its slope and its curvature are therefore

$$\mathcal{L}'_k(\omega) = \mu_k - n_k, \qquad \mathcal{L}''_k(\omega) = \mu_k, \tag{11.6}$$

the expected count minus the observed one, and the expected count. The row pulls ω down where it expects more failures than it holds and up where it expects fewer. The remaining term of equation (4.42), $(\mu_k - n_k)\,\mathrm{d}^2 \ln \mu_k / \mathrm{d}\omega^2$, is zero because the log of each expected count is linear in ω, the log rate being the canonical parameter of the Poisson count: the curvature (11.6) is exact, the observed and the expected information are equal, equation (4.44), and Newton's method and the Gauss–Newton step coincide. The slopes and the curvatures of the four rows add,

$$\begin{aligned} \mathcal{L}'(\omega) &= (10\,\mathrm{e}^{\omega} - 2) + (36\,\mathrm{e}^{\omega} - 7) + (40\,\mathrm{e}^{\omega} - 9) + (40\,\mathrm{e}^{\omega} - 8) = b\,\mathrm{e}^{\omega} - a, \\ \mathcal{L}''(\omega) &= 10\,\mathrm{e}^{\omega} + 36\,\mathrm{e}^{\omega} + 40\,\mathrm{e}^{\omega} + 40\,\mathrm{e}^{\omega} = b\,\mathrm{e}^{\omega}, \end{aligned} \tag{11.7}$$

the derivatives of the objective (11.5): the rows enter the fit only through the sums of their counts and of their exposures, which is the pooling result of Chapter 6 read off the summed slopes.

11.2 The posterior of the state (step 3)

In the book's notation $p(\mathbf{z} \mid \mathbf{y})$ is the posterior density of the state and y_{new} a future observation, here a count.

All seven are read from one posterior: the posterior of the state that the four rows of step 2 give. This section derives it and computes it for the fleet; Section 11.3 reads every answer from it.

11.2.1 The posterior from the four rows

The four rows act on the one number ω through their summed slope and curvature, equation (11.7). Newton's method, equation (4.40) with one entry, replaces the objective at the current ω_m by the parabola with the same slope and curvature and moves to its minimum,

$$\Delta\omega = -\frac{\mathcal{L}'(\omega_m)}{\mathcal{L}''(\omega_m)}, \qquad \omega_{m+1} = \omega_m + \Delta\omega; \tag{11.8}$$

the iteration starts at the prior's center, $\omega_0 = \ln(a_0/b_0)$.

The MAP value and the Laplace variance. The iteration stops where $\mathcal{L}' = 0$, that is where $b\,\mathrm{e}^{\omega} = a$:

$$\hat{\omega} = \ln \frac{a}{b}, \tag{11.9}$$

the logarithm of the posterior mean a/b of the exact gamma posterior, not of its mode, because the change-of-variable factor has added one to the shape. There the expected counts are $\mu_k(\hat{\omega}) = t_k\, a/b$, and the curvature at the minimum is the information Λ, whose reciprocal is the Laplace variance of the log rate:

$$\Lambda = \mathcal{L}''(\hat{\omega}) = \sum_k \mu_k(\hat{\omega}) = b\,\mathrm{e}^{\hat{\omega}} = a, \qquad \sigma_\omega^2 = \frac{1}{\Lambda} = \frac{1}{a}, \tag{11.10}$$

so that the information is the sum of the four rows' curvatures, their expected counts, and it equals the posterior shape. The slope $\mathcal{L}'$ is zero at $\hat{\omega}$, but the rows' slopes are not each zero: a row whose own rate n_k/t_k differs from a/b keeps a slope $\mu_k - n_k$, and the slopes cancel only in their sum. Since the remaining term is zero, the Laplace posterior of Chapter 4, equation (4.38),

is $\omega \mid \mathbf{y} \approx \mathcal{N}(\hat{\omega}, \sigma_\omega^2)$, with the standard deviation $\sigma_\omega = 1/\sqrt{a}$, the relative standard deviation of the gamma posterior. Holding a year out deletes its term: the same iteration on the remaining rows ends at the logarithm of their pooled mean, with the information the sum of their expected counts.

11.2.2 The exact posterior and the Laplace predictive

The gamma posterior. Under a shared constant rate the yearly counts are independent Poisson counts, $K_i \mid \lambda \sim \text{Poisson}(\lambda h_i)$, and by the pooling result of Chapter 6, equation (6.32), the posterior depends on the record only through the total count $k = \sum_i k_i$ and the total exposure $T_{\text{tot}} = \sum_i h_i$:

$$\lambda \mid \mathbf{y} \sim \text{Gamma}(a, b), \qquad a = a_0 + k, \qquad b = b_0 + T_{\text{tot}}, \tag{11.11}$$

with b in unit-years. The objective (11.5) is the negative logarithm of this density written as a density of ω, and the minimum (11.9) lies at the logarithm of its mean; the fit is its Laplace approximation.

Why the delta form is exact here. The delta form (11.19) uses only $\mathrm{e}^{\hat{\omega}}$ and $\mathrm{e}^{2\hat{\omega}}\sigma_\omega^2$, and for a gamma posterior two identities of the fit make these the exact moments of the rate, for every a and b. The slope $\mathcal{L}' = b\,\mathrm{e}^{\omega} - a$ vanishes at $\mathrm{e}^{\hat{\omega}} = a/b = \mathbb{E}[\lambda \mid \mathbf{y}]$: the change-of-variable factor, which added one to the shape, moved the mode in ω from the rate's mode to its mean. And the information is $\Lambda = b\,\mathrm{e}^{\hat{\omega}} = a$, so $\mathrm{e}^{2\hat{\omega}}\sigma_\omega^2 = (a/b)^2/a = a/b^2 = \text{Var}(\lambda \mid \mathbf{y})$. The moments $\mu(\hat{\omega}) = \tau\mathbb{E}[\lambda \mid \mathbf{y}]$ and $g^2\sigma_\omega^2 = \tau^2\,\text{Var}(\lambda \mid \mathbf{y})$ are then those of equation (11.21).

The Laplace predictive as a one-dimensional expectation. The tails and the spares levels need the whole distribution, not two moments. Given the state the count is Poisson with mean $\mu(\omega) = \tau\mathrm{e}^{\omega}$. Averaging over the Laplace normal, with φ the standard normal density,

$$\Pr(K_f = j \mid \mathbf{y}) \approx \int_{-\infty}^{\infty} \frac{(\tau\mathrm{e}^{\omega})^j\,\mathrm{e}^{-\tau\mathrm{e}^{\omega}}}{j!}\,\frac{1}{\sigma_\omega}\,\varphi\!\left(\frac{\omega - \hat{\omega}}{\sigma_\omega}\right)\mathrm{d}\omega.$$

Substitute $\omega = \hat{\omega} + \sigma_\omega x$, so that $\mathrm{d}\omega = \sigma_\omega\,\mathrm{d}x$ and x has the standard normal density:

$$\Pr(K_f = j \mid \mathbf{y}) \approx \int_{-\infty}^{\infty} P_j(x)\,\varphi(x)\,\mathrm{d}x, \qquad P_j(x) = \frac{\mu(x)^j\,\mathrm{e}^{-\mu(x)}}{j!}, \quad \mu(x) = \tau\,\mathrm{e}^{\hat{\omega} + \sigma_\omega x}, \tag{11.12}$$

an expectation of the Poisson probability over a standard normal variable. Gauss–Hermite quadrature for the standard normal replaces such an expectation, the integral (11.12), by a weighted sum over n nodes $x_1, \ldots, x_n$ with quadrature weights $\nu_1, \ldots, \nu_n$ summing to one,

$$\Pr(K_f = j \mid \mathbf{y}) \approx \sum_{m=1}^{n} \nu_m\,P_j(x_m), \tag{11.13}$$

with nodes and weights chosen so that the sum is exact whenever P_j is a polynomial of degree at most $2n - 1$; the nodes are the roots of the n-th Hermite polynomial of the standard normal, a standard result. The Poisson probability is smooth in x, and a handful of nodes suffice. A tail probability is the same sum with P_j replaced by the Poisson tail at each node.

The moments of the Laplace predictive. The mean and variance of the whole Laplace predictive differ from the delta form, which linearizes the conditional mean $\tau\mathrm{e}^{\omega}$ around $\hat{\omega}$. Kept unlinearized, for a normal ω the standard lognormal moments give $\mathbb{E}[\mathrm{e}^{\omega}] = \mathrm{e}^{\hat{\omega} + \sigma_\omega^2/2}$ and $\text{Var}(\mathrm{e}^{\omega}) = \mathrm{e}^{2\hat{\omega} + \sigma_\omega^2}(\mathrm{e}^{\sigma_\omega^2} - 1)$; the law of total variance then gives

$$\mathbb{E}[K_f \mid \mathbf{y}] \approx \tau\,\mathrm{e}^{\hat{\omega} + \sigma_\omega^2/2}, \qquad \text{Var}(K_f \mid \mathbf{y}) \approx \tau\,\mathrm{e}^{\hat{\omega} + \sigma_\omega^2/2} + \tau^2\,\mathrm{e}^{2\hat{\omega} + \sigma_\omega^2}(\mathrm{e}^{\sigma_\omega^2} - 1). \tag{11.14}$$

With $\hat{\omega} = \ln(a/b)$ and $\sigma_\omega^2 = 1/a$ the Laplace mean is the exact mean $a\tau/b$ times $e^{1/(2a)}$, larger by about $1/(2a)$ relatively. The source of the excess is the center: the Laplace normal is centered at $\ln(a/b)$, while the exact mean of $\ln\lambda$ under the gamma posterior is $\psi(a) - \ln b$, smaller by about $1/(2a)$, and its exact variance is the trigamma value $\psi'(a)$ rather than $1/a$, by the expansions of Chapter 8, equations (8.19) and (8.20). The Laplace predictive is therefore shifted slightly upward, and its upper tail slightly heavy. To first order in $1/a$ the factor $e^{\sigma_\omega^2/2}$ and the offset $\ln a - \psi(a)$ of the center cancel, but the exactness of the delta form rests on the two identities above, not on that cancellation.

11.2.3 The posterior of the fleet record

The fleet's inputs are the prior Gamma(2, 10 unit-years) and the record of 7, 9 and 8 failures in 36, 40 and 40 unit-years.

The prior and the posterior. The prior (11.2) has the mean $a_0/b_0 = 0.2$ per unit-year, the standard deviation $\sqrt{2}/10 = 0.141$ and the central 95 % range 0.024 to 0.557. By equation (11.11),

$$a = 2 + 24 = 26, \qquad b = 10 + 116 = 126 \text{ unit-years},$$

so that $\lambda \mid \mathbf{y} \sim \text{Gamma}(26,\ 126)$, with the mean $a/b = 0.2063$ per unit-year, the standard deviation $\sqrt{a}/b = 0.0405$, the relative standard deviation $1/\sqrt{26} = 0.196$ and the central 95 % interval [0.135, 0.293] per unit-year. The maximum-likelihood estimate is $24/116 = 0.207$; the prior has moved the center by less than one percent and has added two failures' worth of information to the record's twenty-four. Table 11.3 lists the three distributions of the rate the chapter uses: each row is one, the columns give its parameters, mean, standard deviation and central 95 % interval. The posterior after years 1 and 2 alone is the one against which year 3 is checked in Section 11.3.4.

Table 11.3: The distributions of the fleet's failure rate used in the chapter. The gamma rate parameter is in unit-years.

Distribution	Parameters (a, b) [–, unit-years]	Mean [per unit-year]	Std. dev. [per unit-year]	95 % interval [per unit-year]
prior	(2, 10)	0.200	0.141	[0.024, 0.557]
posterior after years 1 and 2	(18, 86)	0.209	0.049	[0.124, 0.316]
posterior after years 1 to 3	(26, 126)	0.206	0.040	[0.135, 0.293]

The Newton iteration. Start at the prior's center, $\omega_0 = \ln 0.2 = -1.6094$. There the expected counts of equation (11.4) are (2, 7.2, 8, 8): the prior row and year 3 expect exactly their counts, year 1 expects 0.2 failures more than it had and year 2 one fewer, and equations (11.6) and (11.7) give

$$\mathcal{L}_k' = (0,\ 0.2,\ -1,\ 0), \quad \mathcal{L}' = -0.8, \quad \mathcal{L}'' = 25.2, \quad \Delta\omega = \frac{0.8}{25.2} = 0.0317.$$

Table 11.4 lists the iteration (11.8): each row is one iterate, the columns give the entry ω, the rate e^ω, the slopes of the four rows, the slope $\mathcal{L}'$ of the objective, its curvature $\mathcal{L}''$ and the step. The convergence is quadratic from the first step, and the slope is below 10^{-12} at the third iterate.

Table 11.4: Newton iteration of equation (11.8) for the fleet's log rate from the prior's center $\omega_0 = \ln 0.2$. The row slopes $\mathcal{L}'_k = \mu_k - n_k$ are those of the prior row ($k = 0$) and of years 1 to 3.

Iterate m [–]	ω [–]	e^{ω} [per unit-year]	$\mathcal{L}'_0$ [–]	$\mathcal{L}'_1$ [–]	$\mathcal{L}'_2$ [–]	$\mathcal{L}'_3$ [–]	$\mathcal{L}'$ [–]	$\mathcal{L}''$ [–]	$\Delta\omega$ [–]
0	−1.60944	0.20000	0	0.2000	−1.0000	0	−0.800	25.2000	3.17×10^{-2}
1	−1.57769	0.20645	0.0645	0.4322	−0.7420	0.2580	1.28×10^{-2}	26.0128	-4.93×10^{-4}
2	−1.57819	0.20635	0.0635	0.4286	−0.7460	0.2540	3.17×10^{-6}	26.0000	-1.22×10^{-7}
3	−1.57819	0.20635	0.0635	0.4286	−0.7460	0.2540	1.9×10^{-13}	26.0000	-7.4×10^{-15}

The MAP value and the Laplace variance. By equation (11.9) the iteration ends at

$$\hat{\omega} = \ln(26/126) = \ln 0.2063 = -1.5782,$$

where the expected counts are (2.0635, 7.4286, 8.2540, 8.2540): the rate fitted to the whole record expects 0.43 failures more than year 1 had and 0.75 fewer than year 2 had. The rows' slopes, which cancel in their sum, and, by equation (11.10), the information and the variance are

$$\mathcal{L}'_k(\hat{\omega}) = (0.0635,\ 0.4286,\ -0.7460,\ 0.2540), \qquad \Lambda = \mathcal{L}''(\hat{\omega}) = 26, \qquad \sigma_\omega^2 = \frac{1}{\Lambda} = 0.038462, \tag{11.15}$$

the information the sum $2.0635 + 7.4286 + 8.2540 + 8.2540 = 26$ of the rows' curvatures, their expected counts. The log rate is read off as $\hat{\omega} = -1.5782$ with $\sigma_\omega = 1/\sqrt{26} = 0.1961$.

The pair $(\hat{\omega}, \sigma_\omega) = (-1.5782,\ 0.1961)$ is what step 4 uses: every answer is computed from it, through the delta form or the Gauss–Hermite sum (11.13), and the held-out answers from the same pair of the rows that remain. The exact posterior Gamma(26, 126) gives the value set beside each, with the gap.

11.3 The answers (step 4)

The seven questions of Table 11.2 take three general forms of Chapter 2: the predictive of a future observation, equation (2.13), with its moments; the threshold form, equation (2.12), here applied to a predictive; and the ratio of evidences, equation (2.14). Each is answered from the posterior of Section 11.2: its engineering question, the general form of its category, the form it takes for the fleet, the quantities of the posterior it uses, and the fleet's numbers. The questions name the coming year of $\tau = 44$ unit-years, split into $\tau_B = 12$ for hall B and $\tau_R = 32$ for the rest; the thresholds of 12 and 15 failures; the service levels 0.95 and 0.99, with 0.80, 0.90 and 0.999 beside them; and, for the held-out check, year 3 with $h_3 = 40$ unit-years and the doubling factor 2.

Four objects of the fit recur in the specialized forms. The *MAP value* $\hat{\omega}$ is the minimum of the objective that the four rows put on the one entry; the *information* $\Lambda = \mathcal{L}''(\hat{\omega})$ is the curvature of that objective at the minimum, the sum of the rows' curvatures; the *Laplace variance* $\sigma_\omega^2 = 1/\Lambda$ makes the posterior of the state approximately normal, $\omega \mid \mathbf{y} \approx \mathcal{N}(\hat{\omega}, \sigma_\omega^2)$; and a predicted count enters through the *derivative* g of its conditional mean at $\hat{\omega}$, with $g^2\sigma_\omega^2$. Beside them stands the exact posterior that this problem admits, the gamma posterior Gamma(a, b) of the rate, with $a = a_0 + \sum_i k_i$ the prior shape plus the failures and $b = b_0 + \sum_i h_i$ the prior rate parameter plus the unit-years, under which every predicted count is negative binomial. Every answer is computed from the posterior of the state, and that is the value the chapter reports: the moments by the delta form, which here coincide with the exact ones, and the whole distribution, its tails and its quantiles, by the Gauss–Hermite sum (11.13) over the Laplace normal. The exact value, from the negative-binomial predictive and the gamma posteriors, is set beside each answer with the gap.

11.3.1 The count next year, its variance and the horizon

Engineering questions. How many failures will the fleet have next year, and within what range? How does the range grow with the planning horizon? The first asks for the distribution of the count K_f of the coming year of $\tau = 44$ unit-years, with a central range for it; the second for the same distribution over shorter and longer exposures, from a quarter of a fleet-year to four fleet-years.

General form. The predictive of a future observation, equation (2.13),

$$p(y_{\text{new}} \mid \mathbf{y}) = \int p(y_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z},$$

summarized by its moments, equation (2.10), and by its central 95 % interval, equation (2.11).

Specialized form. The first question is equation (2.13) with $y_{\text{new}} = K_f$, the count in a future exposure of τ unit-years, the state $\mathbf{z} = \omega$, $p(y_{\text{new}} \mid \mathbf{z})$ the Poisson probability, $K_f \mid \omega \sim \text{Poisson}(\mu(\omega))$, with the mean $\mu(\omega) = \tau\, \mathrm{e}^{\omega}$, a derived quantity of the state, and $p(\mathbf{z} \mid \mathbf{y})$ the posterior of step 3, the Laplace normal $\mathcal{N}(\hat{\omega}, \sigma_\omega^2)$. That integral is the Laplace predictive (11.12), and the sum (11.13) evaluates it; its cumulative sums give the central interval, equation (2.11). The value beside it comes from the exact posterior. Written over the rate, the same integral averages the Poisson probability over the gamma posterior:

$$\Pr(K_f = j \mid \mathbf{y}) = \int_0^\infty \frac{(\lambda\tau)^j\, \mathrm{e}^{-\lambda\tau}}{j!}\, \frac{b^a \lambda^{a-1} \mathrm{e}^{-b\lambda}}{\Gamma(a)}\, \mathrm{d}\lambda.$$

Chapter 6 evaluated this integral with the gamma integral, equation (6.54). The result is the negative-binomial distribution with shape a and success probability π:

$$\Pr(K_f = j \mid \mathbf{y}) = \frac{\Gamma(a+j)}{\Gamma(a)\, j!}\, \pi^a (1-\pi)^j, \qquad \pi = \frac{b}{b+\tau}, \qquad j = 0, 1, 2, \ldots \tag{11.16}$$

The future exposure enters only through π, the share of the record's exposure b in the extended exposure $b + \tau$, where b includes the prior's ten pseudo-unit-years. Successive probabilities follow from the ratio

$$\frac{\Pr(K_f = j+1 \mid \mathbf{y})}{\Pr(K_f = j \mid \mathbf{y})} = \frac{a+j}{j+1}\,(1-\pi), \qquad \Pr(K_f = 0 \mid \mathbf{y}) = \pi^a, \tag{11.17}$$

and the cumulative sums give the quantiles. The central 95 % interval of a discrete count runs from its 2.5 % to its 97.5 % quantile and holds at least 0.95 of its probability. The shortcut that replaces the rate by its posterior mean predicts

$$K_f \sim \text{Poisson}(\tau\bar{\lambda}), \qquad \bar{\lambda} = \mathbb{E}[\lambda \mid \mathbf{y}] = \frac{a}{b}, \tag{11.18}$$

the plug-in Poisson of the first reading, a distribution with the same mean as the predictive and a smaller variance. At $j = 0$ the two differ in the direction Chapter 9 derived for mission reliability: $\Pr(K_f = 0 \mid \mathbf{y}) = \mathbb{E}[\mathrm{e}^{-\lambda\tau} \mid \mathbf{y}] = \pi^a$ is never below the plug-in $\mathrm{e}^{-\tau\bar{\lambda}}$, because $\mathrm{e}^{-\lambda\tau}$ is convex in λ and Jensen's inequality applies.

The moments from the fitted state. The second question is the moments (2.10) of equation (2.13), with the same y_{new}, state and posterior. The conditional variance of the count is its conditional mean $\mu(\omega)$, and the law of total variance of Chapter 3 averages over the posterior of the state,

$$\text{Var}(K_f \mid \mathbf{y}) = \mathbb{E}\big[\text{Var}(K_f \mid \omega) \mid \mathbf{y}\big] + \text{Var}\big(\mathbb{E}[K_f \mid \omega] \mid \mathbf{y}\big) = \mathbb{E}\big[\mu(\omega) \mid \mathbf{y}\big] + \text{Var}\big(\mu(\omega) \mid \mathbf{y}\big).$$

The delta method of Chapter 4 linearizes the derived quantity around the MAP value, $\mu(\omega) \approx \mu(\hat{\omega}) + g\,(\omega - \hat{\omega})$, with the derivative

$$g = \left.\frac{\mathrm{d}\mu}{\mathrm{d}\omega}\right|_{\hat{\omega}} = \tau\, \mathrm{e}^{\hat{\omega}},$$

which is $\mu(\hat{\omega})$ itself, and under the Laplace posterior $\mathcal{N}(\hat{\omega}, \sigma_\omega^2)$ the two moments of the predictive become

$$\mathbb{E}[K_f \mid \mathbf{y}] \approx \mu(\hat{\omega}), \qquad \mathrm{Var}(K_f \mid \mathbf{y}) \approx \underbrace{\mu(\hat{\omega})}_{\text{Poisson part}} + \underbrace{g^2\sigma_\omega^2}_{\text{epistemic part}}, \tag{11.19}$$

the first term the variance of the count at the fitted rate and the second the posterior variance of the log rate carried into the count through g. The plug-in Poisson (11.18) is this form with σ_ω^2 set to zero. With the MAP value $\hat{\omega} = \ln(a/b)$ and the variance $\sigma_\omega^2 = 1/a$ of Section 11.2.1, $\mu(\hat{\omega}) = a\tau/b$ and $g^2\sigma_\omega^2 = (a\tau/b)^2/a = a\tau^2/b^2$.

The exact moments and the two parts. The values beside them come from the exact posterior. The exact moments of the negative binomial, equations (6.56) and (6.57) of Chapter 6 with the gamma moments $\mathbb{E}[\lambda \mid \mathbf{y}] = a/b$ and $\mathrm{Var}(\lambda \mid \mathbf{y}) = a/b^2$, are

$$\mathbb{E}[K_f \mid \mathbf{y}] = \frac{a\tau}{b}, \qquad \mathrm{Var}(K_f \mid \mathbf{y}) = \frac{a\tau}{b}\left(1 + \frac{\tau}{b}\right), \tag{11.20}$$

the same two terms that the fitted state gives in equation (11.19): for a gamma posterior the delta form on the log rate returns the exact first two moments, for the reason Section 11.2.2 gives. Read term by term, the variance is

$$\mathrm{Var}(K_f \mid \mathbf{y}) = \underbrace{\tau\,\mathbb{E}[\lambda \mid \mathbf{y}]}_{\mu(\hat{\omega}),\ \text{Poisson part}} + \underbrace{\tau^2\ \mathrm{Var}(\lambda \mid \mathbf{y})}_{g^2\sigma_\omega^2,\ \text{epistemic part}} = \frac{a\tau}{b} + \frac{a\tau}{b}\cdot\frac{\tau}{b}. \tag{11.21}$$

The first term is the variance the count would have if the rate were known to be its posterior mean, the plug-in Poisson variance; the second is the variance added because it is not, and it is the mean times τ/b. The predictive is therefore overdispersed by the factor $1 + \tau/b$, and its standard deviation exceeds the plug-in Poisson standard deviation by the factor $\sqrt{1 + \tau/b}$. The share of the epistemic part in the whole is

$$\frac{g^2\sigma_\omega^2}{\mu(\hat{\omega}) + g^2\sigma_\omega^2} = \frac{\tau^2\ \mathrm{Var}(\lambda \mid \mathbf{y})}{\mathrm{Var}(K_f \mid \mathbf{y})} = \frac{\tau}{b + \tau} = 1 - \pi, \tag{11.22}$$

the share of the future exposure in the extended exposure, the same number that sets π.

Principle 11.1. *The predictive variance of a future count is the Poisson variance of the mean, which no record removes, plus the square of the horizon times the posterior variance of the rate. Under a gamma posterior the second part is the first times τ/b, and its share of the whole is $\tau/(b+\tau)$, the share of the horizon in the record's exposure plus the horizon.*

The horizon. Equation (11.21) has a Poisson part linear in the horizon and an epistemic part quadratic in it. In the form (11.19) the reason is visible in the derivative: $g = \tau\mathrm{e}^{\hat{\omega}}$ grows in proportion to τ, so $g^2\sigma_\omega^2$ grows as τ^2 while $\mu(\hat{\omega})$ grows as τ. The two are equal where $\mu(\hat{\omega}) = g^2\sigma_\omega^2$, that is where $\tau\mathrm{e}^{\hat{\omega}}\sigma_\omega^2 = 1$, at $\tau = 1/(\mathrm{e}^{\hat{\omega}}\sigma_\omega^2) = b$: a horizon as long as the record's exposure, prior included. Below it the Poisson scatter dominates, above it the uncertainty of the rate. As the horizon grows without bound the share (11.22) tends to one, and since $g = \mu(\hat{\omega})$, the coefficient of variation of the count tends to the standard deviation of the log rate,

$$\frac{\sqrt{\mathrm{Var}(K_f \mid \mathbf{y})}}{\mathbb{E}[K_f \mid \mathbf{y}]} = \frac{\sqrt{\mu(\hat{\omega}) + g^2\sigma_\omega^2}}{\mu(\hat{\omega})} = \sqrt{\frac{1}{\mu(\hat{\omega})} + \sigma_\omega^2} \ \longrightarrow\ \sigma_\omega = \frac{1}{\sqrt{a}} \quad (\tau \to \infty). \tag{11.23}$$

Over a long horizon the count is known only as well as the log rate, relatively; over a short one it is known only as well as Poisson scatter allows.

The delta form is the reported moment. The whole Laplace predictive keeps the curvature of e^{ω} that the delta form drops, and has the moments (11.14) of Section 11.2.2, which the values below carry beside it.

Building blocks. The MAP value $\hat{\omega} = -1.5782$ and the standard deviation $\sigma_\omega = 0.1961$ (Section 11.2.3), with the derivative $g = \tau \mathrm{e}^{\hat{\omega}}$; the Gauss–Hermite nodes and weights of the sum (11.13), five for the table below and forty for the reported values; the Poisson probabilities of the plug-in; and, for the exact values beside them, the gamma parameters a and b, the success probability π (Section 11.2.2) and the recursion (11.17) with its cumulative sums.

The moments. With $\tau = 44$ unit-years the derivative is $g = 44 \times 0.2063 = 9.0794$, and equation (11.19) gives

$$\begin{aligned} \mathbb{E}[K_f \mid \mathbf{y}] &= \mu(\hat{\omega}) = \frac{26 \times 44}{126} = 9.079, \\ \mathrm{Var}(K_f \mid \mathbf{y}) &= \underbrace{9.079}_{\text{Poisson part}} + \underbrace{9.0794^2 \times 0.038462}_{\text{epistemic part, } g^2\sigma_\omega^2} = 9.079 + 3.171 = 12.250, \end{aligned}$$

a predictive standard deviation of 3.500 failures against $\sqrt{9.079} = 3.013$ for the plug-in Poisson. The value beside them comes from the exact posterior, equation (11.20), and gives the same numbers: the gap is zero, for the reason of Section 11.2.2. The epistemic share (11.22) is $44/170 = 3.171/12.250 = 25.9\,\%$. The whole Laplace predictive, which does not linearize, has by equation (11.14) the mean $9.079\,\mathrm{e}^{1/52} = 9.256$ and the standard deviation 3.552.

The rule of five nodes. Table 11.5 carries out the rule (11.13) under the Laplace normal $\omega \sim \mathcal{N}(-1.5782,\ 0.1961^2)$. Each row is a node; the columns give the node x_m, its weight ν_m, the rate $\mathrm{e}^{\hat{\omega}+\sigma_\omega x_m}$ at the node, the Poisson mean $\mu(x_m)$ in 44 unit-years, and then the functions that this subsection and the next average: the Poisson probability $\mathrm{e}^{-\mu}$ of no failure, the Poisson tail above twelve with its weighted term, and the Poisson probability of at most fifteen. One row of the arithmetic, the node $x = 1.3556$: the rate is $0.2063\,\mathrm{e}^{0.1961 \times 1.3556} = 0.2692$ per unit-year, the mean 11.84 failures, the Poisson tail 0.4063, and the term $0.22208 \times 0.4063 = 0.0902$. The sum of the terms is 0.1744; ten, twenty or forty nodes give the same value to four decimals, so the quadrature contributes no error at the precision shown.

Table 11.5: The Gauss–Hermite rule (11.13) with five nodes under the Laplace posterior $\omega \sim \mathcal{N}(\ln(26/126),\ 1/26)$; weights for the standard normal, mean $\mu(x_m) = 44\,\mathrm{e}^{\hat{\omega}+\sigma_\omega x_m}$. The last row is the weighted sum of each column.

Node x_m [–]	Weight ν_m [–]	Rate [per unit-yr]	$\mu(x_m)$ [failures]	$\mathrm{e}^{-\mu}$ [–]	$\Pr(\text{Poisson}(\mu) > 12)$ [–]	Term [–]	$\Pr(\text{Poisson}(\mu) \le 15)$ [–]
−2.8570	0.01126	0.1178	5.18	5.60×10^{-3}	0.0027	0.0000	0.9999
−1.3556	0.22208	0.1582	6.96	9.49×10^{-4}	0.0260	0.0058	0.9977
0	0.53333	0.2063	9.08	1.14×10^{-4}	0.1301	0.0694	0.9764
1.3556	0.22208	0.2692	11.84	7.18×10^{-6}	0.4063	0.0902	0.8555
2.8570	0.01126	0.3614	15.90	1.24×10^{-7}	0.8002	0.0090	0.4767
Sum	1			3.36×10^{-4}		0.1744	0.9489

The predictive. Over forty nodes the sum (11.13) gives $\Pr(K_f = 0 \mid \mathbf{y}) \approx 3.36 \times 10^{-4}$, the value of the five nodes, $\Pr(K_f = 1 \mid \mathbf{y}) \approx 2.37 \times 10^{-3}$, and so on. The value beside

it comes from the exact posterior. With $\pi = 126/170 = 0.74118$ and $1 - \pi = 0.25882$, equation (11.17) starts at $\Pr(K_f = 0 \mid \mathbf{y}) = 0.74118^{26} = 4.15 \times 10^{-4}$ and gives $\Pr(K_f = 1 \mid \mathbf{y}) = 26 \times 0.25882 \times 4.15 \times 10^{-4} = 2.79 \times 10^{-3}$, and so on: the state route gives a year without failure 0.79×10^{-4} less, because the normal in ω is centered 0.019 above the mean of the exact log rate. Table 11.6 lists the predictive for counts from zero to twenty: each row is a count j; the columns give the Laplace probability, the exact negative-binomial probability, the plug-in Poisson probability at $\tau\bar{\lambda} = 9.079$, and the three cumulative probabilities $\Pr(K_f \le j)$. The last rows give the probability of more than twenty and the moments.

Table 11.6: The predictive distribution of the fleet's count in the coming year of 44 unit-years. Laplace: the sum (11.13) with forty nodes. Negative binomial: equation (11.16) with $a = 26$, $\pi = 126/170$. Plug-in: Poisson(9.079).

Failures j [–]	Laplace [–]	Negative binomial [–]	Plug-in Poisson [–]	Laplace, $\Pr(K_f \le j)$ [–]	Negative binomial, $\Pr(K_f \le j)$ [–]	Plug-in, $\Pr(K_f \le j)$ [–]
0	0.00034	0.00041	0.00011	0.00034	0.00041	0.00011
1	0.00237	0.00279	0.00103	0.00270	0.00321	0.00115
2	0.00857	0.00976	0.00470	0.01127	0.01296	0.00585
3	0.02134	0.02357	0.01422	0.03261	0.03653	0.02007
4	0.04104	0.04422	0.03228	0.07364	0.08076	0.05234
5	0.06503	0.06868	0.05861	0.13867	0.14943	0.11096
6	0.08843	0.09184	0.08869	0.22710	0.24127	0.19965
7	0.10611	0.10866	0.11504	0.33321	0.34994	0.31468
8	0.11468	0.11601	0.13056	0.44789	0.46595	0.44524
9	0.11338	0.11344	0.13171	0.56126	0.57939	0.57695
10	0.10379	0.10276	0.11958	0.66505	0.68215	0.69654
11	0.08885	0.08704	0.09870	0.75391	0.76919	0.79524
12	0.07171	0.06946	0.07468	0.82562	0.83865	0.86992
13	0.05493	0.05255	0.05216	0.88056	0.89121	0.92208
14	0.04017	0.03789	0.03383	0.92073	0.92910	0.95591
15	0.02818	0.02615	0.02047	0.94890	0.95525	0.97638
16	0.01904	0.01735	0.01162	0.96794	0.97260	0.98800
17	0.01244	0.01109	0.00621	0.98038	0.98369	0.99420
18	0.00789	0.00686	0.00313	0.98827	0.99055	0.99733
19	0.00486	0.00411	0.00150	0.99313	0.99466	0.99883
20	0.00292	0.00239	0.00068	0.99606	0.99705	0.99951
more than 20	0.00394	0.00295	0.00049			
mean, variance	9.256, 12.615	9.079, 12.250	9.079, 9.079			

The three distributions are set side by side in Figure 11.2. The horizontal axis is the number of failures in the coming year, the vertical axis its probability in percent; the blue bars are the Laplace predictive, the green bars the negative-binomial predictive and the gray bars the plug-in Poisson. The two brackets above the bars are the central 95 % intervals, green for the two predictives, which share theirs, and gray for the plug-in, and the dashed red lines separate the counts above the thresholds 12 and 15, read in Section 11.3.2. Both predictives are lower and wider than the plug-in: they put less probability on the counts from seven to twelve and more on both ends, and their upper tails are visibly heavier from fifteen failures on. The blue bars lie a little to the right of the green ones throughout, the upward shift of the Laplace center.

The range. The central 95 % interval of the Laplace predictive is $K_f \in [3,\ 17]$, with predictive probability 0.969. The value beside it comes from the exact posterior:

$$K_f \in [3,\ 17] \text{ with predictive probability } 0.971, \qquad \text{plug-in } [4,\ 15] \text{ with } 0.956,$$

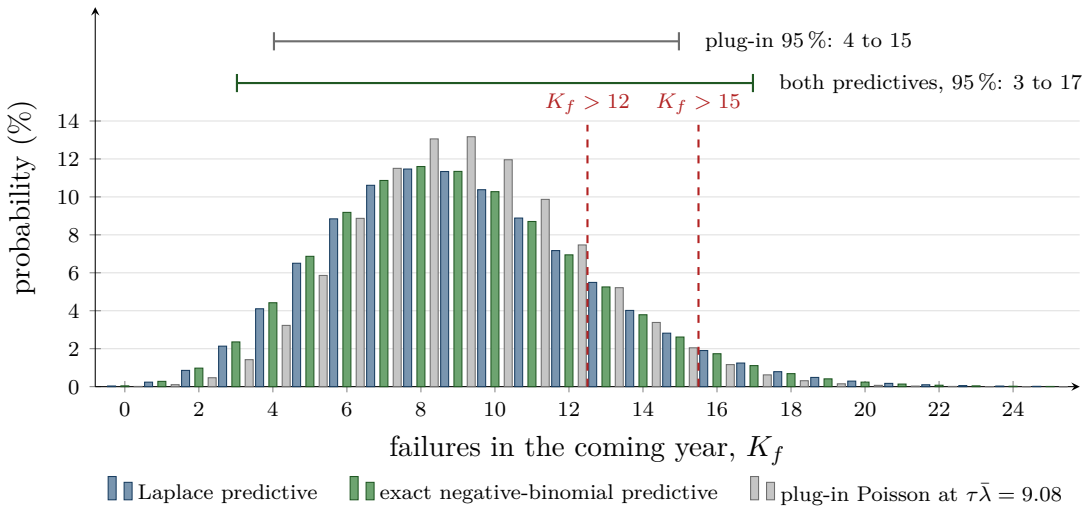


Figure 11.2: Predicted number of compressor failures of the fleet in the coming year of 44 unit-years (Table 11.6). Blue: Laplace predictive, the sum (11.13) over $\mathcal{N}(-1.5782,\ 0.1961^2)$. Green: negative-binomial predictive, equation (11.16) with $a = 26$, $b = 126$, $\tau = 44$. Gray: Poisson with the posterior mean rate, mean 9.079. Brackets: central 95 % intervals. Dashed lines: the thresholds of 12 and 15 failures.

the same interval, with 0.002 more probability inside it. The plug-in interval holds only 0.916 of the Laplace probability (0.919 of the exact). The most probable count is eight on both routes and nine under the plug-in, and the probability of a year without failure is 3.4×10^{-4} (4.1×10^{-4} exactly) against 1.1×10^{-4}.

The range against the horizon. Table 11.7 evaluates the split (11.21) for five horizons: each row is a horizon, from a quarter of a fleet-year (11 unit-years) to four fleet-years (176); the columns give the predictive mean, the Poisson part, the epistemic part, the total variance, the predictive and plug-in standard deviations, the epistemic share and the two central 95 % intervals. One row of the arithmetic, $\tau = 88$: mean $\mu(\hat{\omega}) = 26 \times 88/126 = 18.16$, derivative $g = 18.16$, epistemic part $g^2\sigma_\omega^2 = 18.16^2/26 = 12.68$, total 30.84, standard deviation 5.55 against $\sqrt{18.16} = 4.26$, share $88/214 = 0.411$. The share rises from 8 % for a quarter to 58 % for four years, and the Laplace interval, the same [0, 6] as the Poisson interval for a quarter, is half again as wide as the Poisson interval for four years. From two years on the Laplace interval lies above the exact one, by a failure at each end at two years and by one and two at four.

Table 11.7: The predictive variance of the fleet's count against the horizon, equation (11.21), under the posterior Gamma(26, 126). Poisson part $a\tau/b$, epistemic part $a\tau^2/b^2$, share $\tau/(b+\tau)$. Intervals are central 95 % intervals of the Laplace predictive, with the exact negative binomial in parentheses where it differs, and of the plug-in Poisson.

Horizon τ [unit-years]	Mean [–]	Poisson part [–]	Epistemic part [–]	Total [–]	Std. dev. predictive [–]	Std. dev. plug-in [–]	Share [–]	Predictive interval [–]	Plug-in interval [–]
11	2.27	2.27	0.20	2.47	1.57	1.51	8.0 %	[0, 6]	[0, 6]
22	4.54	4.54	0.79	5.33	2.31	2.13	14.9 %	[1, 10]	[1, 9]
44	9.08	9.08	3.17	12.25	3.50	3.01	25.9 %	[3, 17]	[4, 15]
88	18.16	18.16	12.68	30.84	5.55	4.26	41.1 %	[9, 31] ([8, 30])	[10, 27]
176	36.32	36.32	50.73	87.05	9.33	6.03	58.3 %	[21, 58] ([20, 56])	[25, 49]

How the two parts grow with the horizon is plotted in Figure 11.3. The horizontal axis is the horizon in unit-years, the vertical axis the variance in failures squared. The solid blue line is the Poisson part, the dashed red curve the epistemic part, and the dots mark the five horizons of Table 11.7. The dotted vertical lines mark the coming year, $\tau = 44$, and the crossing at $\tau = b = 126$, where both parts equal $a = 26$. At the coming year the red curve is a third of the blue line; over a two-year horizon, $\tau = 88$, it is two thirds of it, 12.7 against 18.2; and over a four-year horizon, $\tau = 176$, the curve is above the line. By equation (11.23) the coefficient of variation of the count tends to $1/\sqrt{26} = 0.196$ as the horizon grows.

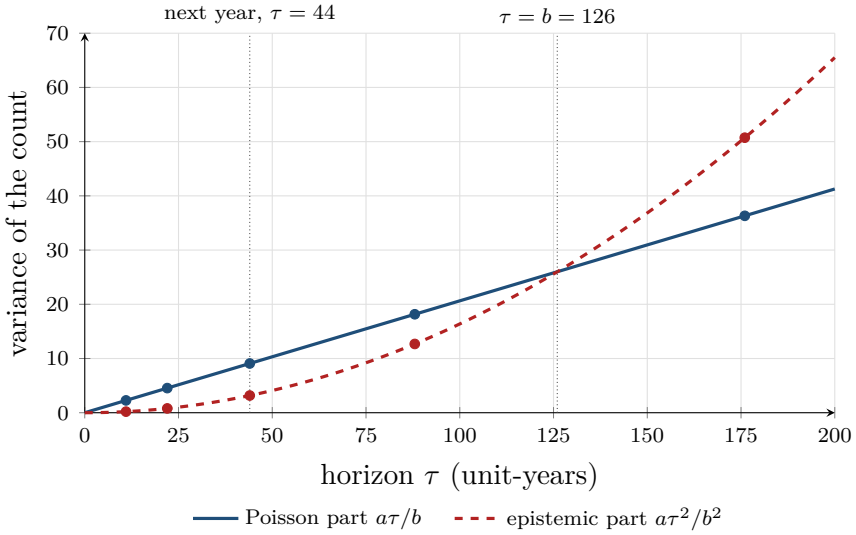


Figure 11.3: The two parts of the predictive variance of the fleet's count against the horizon, equation (11.21) with $a = 26$ and $b = 126$ unit-years: Poisson part $26\tau/126$ and epistemic part $26\tau^2/126^2$. Dots: the horizons of Table 11.7. The parts are equal, at 26, when $\tau = b$.

11.3.2 The thresholds and the spares level

Engineering questions. What is the probability of more than twelve failures next year? Of more than fifteen? How many spare compressors cover next year with 95 % probability? With 99 %? The first two ask for upper tails of the coming year's count, the levels at which the maintenance budget and the crew plan would need revision; the last two for the store that the planner stocks before the year begins.

General form. A prediction under a threshold whose event lies in the future is the threshold form, equation (2.12), taken on the predictive of the future observation,

$$\Pr(y_{\text{new}} \le f^* \mid \mathbf{y}) = \int \Pr(y_{\text{new}} \le f^* \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \qquad \text{accept if } \Pr(y_{\text{new}} \le f^* \mid \mathbf{y}) \ge \kappa,$$

with the exceedance $y_{\text{new}} > f^*$ its complement. A design that chooses the level itself before the observation exists takes the smallest f^* that the rule accepts. No measurement precedes that choice, so it is read from the predictive the record already gives, not from a preposterior of the design form, equation (2.16).

Specialized form. One choice enters here: each failure is taken to consume one compressor from a store that is not replenished during the year. That makes the store's cover the event $K_f \le s$ over the whole year; a store reordered after each use is covered by the count over the reorder lead time instead (Section 11.6.6).

The thresholds. The questions are equation (2.12) with $y_{\text{new}} = K_f$ and $f^* = c$, for $c = 12$ and $c = 15$, $\Pr(y_{\text{new}} \le f^* \mid \mathbf{z})$ the Poisson probability at the mean $\mu(\omega)$, and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal of step 3. The sum (11.13) with the Poisson tail in place of P_j gives

$$\Pr(K_f > c \mid \mathbf{y}) \approx \sum_{m=1}^{n} \nu_m \Pr\bigl(\text{Poisson}(\mu(x_m)) > c\bigr). \tag{11.24}$$

The value beside it comes from the exact posterior: the tails are complements of the cumulative sums of equation (11.16), $\Pr(K_f > c \mid \mathbf{y}) = 1 - \sum_{j=0}^{c} \Pr(K_f = j \mid \mathbf{y})$; the plug-in tails are those of equation (11.18). The predictive's extra variance, the epistemic part of equation (11.21), is the contribution of the high rates the posterior still admits, and those rates carry the far tail, so the plug-in understates a tail by a factor that grows as the threshold rises.

The spares level. The question is equation (2.12) with the level $f^* = s$ chosen and the assurance $\kappa = q$. A store holding s spare compressors at the start of the year covers every failure of the year when $K_f \le s$. The probability of that event is its service level, the planner's counterpart of the assurance κ, and the spares level for a service level q is the smallest s with at least that probability,

$$s_q = \min\{s : \Pr(K_f \le s \mid \mathbf{y}) \ge q\}, \tag{11.25}$$

a quantile of the predictive, read from the cumulative sums of the sum (11.13), and beside it from those of the negative binomial (11.16). The failures that find no spare number $(K_f - s)^+$, with $(x)^+ = \max(x, 0)$, and their expected number is the shortfall

$$\mathbb{E}[(K_f - s)^+ \mid \mathbf{y}] = \sum_{j>s} (j - s) \Pr(K_f = j \mid \mathbf{y}). \tag{11.26}$$

The plug-in sets its level by the Poisson (11.18) and claims the Poisson service level for it; the service level it delivers is the predictive's cumulative probability at that level. The gap between the two levels follows from the two spreads. Both distributions have the mean $\mu = \mu(\hat{\omega}) = a\tau/b$, and to a normal approximation a quantile at level q lies z_q standard deviations above the mean, with z_q the standard normal quantile. By equation (11.19) the two standard deviations differ only by the epistemic part $g^2\sigma_\omega^2$, which the plug-in sets to zero:

$$s_q \approx \mu + z_q\sqrt{\mu + g^2\sigma_\omega^2} \quad \text{(predictive)}, \qquad s_q \approx \mu + z_q\sqrt{\mu} \quad \text{(plug-in)},$$

and the difference is

$$\Delta s_q \approx z_q\Bigl(\sqrt{\mu + g^2\sigma_\omega^2} - \sqrt{\mu}\Bigr). \tag{11.27}$$

The gap grows with the service level, through z_q, and with the horizon, through the epistemic part; the right skew of the negative binomial adds to it beyond the normal approximation.

Building blocks. The MAP value $\hat{\omega}$ and the standard deviation σ_ω through the sums (11.13) and (11.24) and their cumulative sums, the five nodes of Table 11.5 and forty for the reported values; the cumulative sums of the plug-in Poisson; the shortfall sums (11.26); the standard normal quantiles z_q with $\mu(\hat{\omega})$ and $g^2\sigma_\omega^2$ from equation (11.19); and, for the exact values beside them, the cumulative sums of the predictive (11.16).

The thresholds of twelve and fifteen failures. From the sum (11.24), the last rows of Table 11.5 and the Laplace cumulative column of Table 11.6,

$$\Pr(K_f > 12 \mid \mathbf{y}) \approx 1 - 0.82562 = 0.174, \qquad \Pr(K_f > 15 \mid \mathbf{y}) \approx 1 - 0.94890 = 0.051.$$

The values beside them come from the exact posterior, the negative-binomial cumulative column:

$$\Pr(K_f > 12 \mid \mathbf{y}) = 1 - 0.83865 = 0.161, \qquad \Pr(K_f > 15 \mid \mathbf{y}) = 1 - 0.95525 = 0.045,$$

smaller by 0.013 and 0.006, against 0.130 and 0.024 under the plug-in. The ratio of the Laplace tail to the plug-in tail is 1.34 above twelve, 2.16 above fifteen, 3.4 above seventeen and 8.0 above twenty (1.24, 1.89, 2.8 and 6.0 for the exact tail): the plug-in roughly halves the probability of more than fifteen failures.

The node table shows where the tail comes from. The central node, the plug-in rate, contributes only 0.069; the node one and a third standard deviations above it contributes more, 0.090, although its weight is less than half. The upper tail of the count is carried by the upper part of the posterior of the rate, which the plug-in, keeping only the central node, leaves out.

The spares levels. By equation (11.25), the cumulative probability of the Laplace predictive first reaches 0.95 at $s = 16$, with 0.9679; at 15 it is 0.9489, the sum of Table 11.5. The value beside it comes from the exact posterior: its cumulative probability first reaches 0.95 at $s = 15$, with 0.9553; at 14 it is 0.9291. Here the gap changes the level. The two routes fall on opposite sides of 0.95 at fifteen spares, 0.9489 against 0.9553, a difference of 0.006: the normal in ω is centered 0.019 above the mean of the exact log rate and has the heavier upper tail, and at this boundary that is enough. The plug-in reaches 0.95 at 14, with 0.9559. For $q = 0.99$ the Laplace level is 19, with 0.9931, its value at 18 being 0.9883; the exact level is 18, with 0.9906, again on the other side of the line, and the plug-in level is 17, with 0.9942.

Table 11.8 extends the comparison to five service levels. Each row is a service level; the columns give the spares level and its service level, Laplace with exact in parentheses, the plug-in level and the service level the plug-in claims for it, the service level that the plug-in level actually delivers under the predictive, and the expected shortfall (11.26) at the plug-in level under each distribution. Against the Laplace predictive the plug-in is one spare short at 0.90, two at 0.95 and 0.99 and three at 0.999; against the exact one it is one spare short at every level from 0.90 to 0.99 and two short at 0.999. At 0.95 it delivers 0.921 (0.929), and its expected shortfall is 0.210 failures (0.180), more than twice the 0.090 it reports.

Table 11.8: Spares levels for the coming year, equation (11.25), under the Laplace predictive, the exact negative binomial in parentheses, and under the plug-in Poisson. Service levels are $\Pr(K_f \le s)$; the shortfall, equation (11.26), is the expected number of failures beyond the plug-in level s.

Service level q [–]	Level s_q, Laplace (exact) [spares]	Its service level, Laplace (exact) [–]	Plug-in level [spares]	Claimed service level [–]	Delivered service level of the plug-in level, Laplace (exact) [–]	Shortfall, Laplace (exact) [failures]	Shortfall, plug-in [failures]
0.80	12 (12)	0.826 (0.839)	12	0.870	0.826 (0.839)	0.503 (0.451)	0.298
0.90	14 (14)	0.921 (0.929)	13	0.922	0.881 (0.891)	0.329 (0.289)	0.168
0.95	16 (15)	0.968 (0.955)	14	0.956	0.921 (0.929)	0.210 (0.180)	0.090
0.99	19 (18)	0.993 (0.991)	17	0.994	0.980 (0.984)	0.047 (0.037)	0.010
0.999	23 (22)	0.9993 (0.9992)	20	0.9995	0.9961 (0.9971)	0.0088 (0.0063)	0.0008

The size of the gap. Equation (11.27) gives

$$\Delta s_q \approx z_q \times (3.500 - 3.013) = 0.487\, z_q,$$

and with $z_q = 1.645$, 2.326 and 3.090 the gap is 0.80 spares at the 95 % level, 1.13 at 99 % and 1.50 at 99.9 %; rounded to whole spares through the discrete quantiles, it is the two, two and

three spares of the Laplace predictive and the one, one and two of the exact one in Table 11.8. For a four-year horizon, $\tau = 176$, where $g = 36.32$ and $g^2\sigma_\omega^2 = 50.73$, the difference of the standard deviations is $9.330 - 6.026 = 3.30$ instead of 0.487: about five spares at the 95 % level by equation (11.27), eight by the Laplace quantiles, 54 against 46, and seven by the exact quantiles, 53 against 46.

11.3.3 Two halls under a shared rate

Engineering question. How many failures will hall B have, and are the halls' counts independent? The planner keeps a cabinet of spares in hall B and the main store elsewhere, and asks how many failures the hall will see and whether the two counts can be planned separately, the probability of a year without any failure in the fleet taken as the product of the two halls' probabilities.

General form. The predictive of a future observation, equation (2.13), for a vector of two counts,

$$p(\mathbf{y}_{\text{new}} \mid \mathbf{y}) = \int p(\mathbf{y}_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z},$$

with its covariance matrix, the moments (2.10) taken for both entries, and the probability of a joint event as the same integral of the joint conditional probability.

Specialized form. One choice enters here: the units of hall B are given the fleet's one rate. This makes the halls exchangeable and lets hall B borrow the whole fleet's record, and it fails if hall B ran hotter or its units were of another series; its rate would then be an entry of its own, which Section 11.6.6 writes out.

Split the coming year into hall B, with exposure τ_B, and the rest of the fleet, with exposure τ_R, and let K_B and K_R be their counts. The question is equation (2.13) with $\mathbf{y}_{\text{new}} = (K_B,\ K_R)^{\mathsf{T}}$, the state ω, $p(\mathbf{y}_{\text{new}} \mid \mathbf{z})$ the product of two Poisson probabilities, and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal of step 3, with the covariance the moments (2.10) of the pair. Given the rate they are independent Poisson counts with means $\lambda\tau_B$ and $\lambda\tau_R$; each marginal predictive is the sum (11.13) with its own exposure, and the value beside it is equation (11.16) with that exposure. They are not independent given the record alone, because both depend on the same uncertain λ.

The law of total covariance. For any two quantities X and Y and a third quantity λ on which both depend, write the covariance as the mean of the product minus the product of the means, and condition each mean on λ:

$$\operatorname{Cov}(X, Y) = \mathbb{E}\big[\mathbb{E}[XY \mid \lambda]\big] - \mathbb{E}\big[\mathbb{E}[X \mid \lambda]\big]\, \mathbb{E}\big[\mathbb{E}[Y \mid \lambda]\big].$$

Write the conditional mean of the product as $\mathbb{E}[XY \mid \lambda] = \operatorname{Cov}(X, Y \mid \lambda) + \mathbb{E}[X \mid \lambda]\, \mathbb{E}[Y \mid \lambda]$; the expectation of the second term minus the product of the means is the covariance of the conditional means, and

$$\operatorname{Cov}(X, Y) = \mathbb{E}\big[\operatorname{Cov}(X, Y \mid \lambda)\big] + \operatorname{Cov}\big(\mathbb{E}[X \mid \lambda],\ \mathbb{E}[Y \mid \lambda]\big), \tag{11.28}$$

the law of total covariance; with $Y = X$ it is the law of total variance.

The covariance matrix of the two counts. Write the two counts as one vector $\mathbf{K} = (K_B,\ K_R)^{\mathsf{T}}$. Given the state they are independent Poisson counts, so their conditional covariance matrix is diagonal, with the two conditional means on the diagonal. The means are derived quantities of the one entry ω, with their derivatives at $\hat{\omega}$,

$$\mu_B(\omega) = \tau_B\, \mathrm{e}^{\omega}, \qquad \mu_R(\omega) = \tau_R\, \mathrm{e}^{\omega}, \qquad g_B = \tau_B\, \mathrm{e}^{\hat{\omega}}, \qquad g_R = \tau_R\, \mathrm{e}^{\hat{\omega}},$$

the pattern (11.3): both counts depend on the one entry that every row of the record acts on as well. Equation (11.28) written for the pair, with the delta method on its second term, gives the predictive covariance matrix of the two counts,

$$\mathbf{C} = \operatorname{Cov}(\mathbf{K} \mid \mathbf{y}) \approx \underbrace{\begin{bmatrix} \mu_B(\hat{\omega}) & 0 \\ 0 & \mu_R(\hat{\omega}) \end{bmatrix}}_{\text{Poisson part}} + \underbrace{\sigma_\omega^2 \begin{bmatrix} g_B^2 & g_B g_R \\ g_B g_R & g_R^2 \end{bmatrix}}_{\text{epistemic part}}, \qquad C_{BR} = \tau_B \tau_R \, \mathrm{e}^{2\hat{\omega}} \, \sigma_\omega^2 = \frac{\tau_B \tau_R \, a}{b^2}, \tag{11.29}$$

whose diagonal entries are the two marginal variances of equation (11.19) and whose off-diagonal entry, the covariance of the two counts, is positive whenever σ_ω^2 is not zero. Equation (11.28) applied to the exact gamma posterior gives the same entry, $\tau_B \tau_R \operatorname{Var}(\lambda \mid \mathbf{y})$.

The shared entry. The epistemic part has rank one because both derivatives are taken with respect to the one entry: the two conditional means are the same entry of the state scaled by two exposures, so their correlation is $+1$ and its off-diagonal entry $\sigma_\omega^2 g_B g_R$ is the geometric mean of its diagonal entries $\sigma_\omega^2 g_B^2$ and $\sigma_\omega^2 g_R^2$. The Poisson diagonal dilutes it. Each ratio $\sigma_\omega^2 g_i^2 / C_{ii}$ is the epistemic share of that count's variance, $\tau_i/(b+\tau_i)$ by equation (11.22), and the correlation of the counts is

$$\operatorname{corr}(K_B, K_R \mid \mathbf{y}) = \frac{C_{BR}}{\sqrt{C_{BB}\, C_{RR}}} = \sqrt{\frac{\sigma_\omega^2 g_B^2}{C_{BB}} \cdot \frac{\sigma_\omega^2 g_R^2}{C_{RR}}} = \sqrt{\frac{\tau_B}{b+\tau_B} \cdot \frac{\tau_R}{b+\tau_R}}, \tag{11.30}$$

the geometric mean of the two counts' epistemic shares. It grows with both horizons and falls as the record grows, since a longer record shrinks σ_ω^2 and leaves the Poisson diagonal. As a check, the fleet's count is $K_B + K_R$, and $g_B + g_R = (\tau_B + \tau_R)\mathrm{e}^{\hat{\omega}}$ is the derivative of the fleet's conditional mean, so $C_{BB} + 2C_{BR} + C_{RR}$ is the variance $\mu(\hat{\omega}) + g^2\sigma_\omega^2$ of the fleet's count in equation (11.19) with $\tau = \tau_B + \tau_R$, as it must be.

The joint probability of no failure. The probability that neither part of the fleet fails is the same integral at $(0, 0)$, one expectation over the posterior of the state, by the sum (11.13) with the joint conditional probability in place of P_j,

$$\Pr(K_B = 0, K_R = 0 \mid \mathbf{y}) \approx \sum_{m=1}^{n} \nu_m \, \mathrm{e}^{-(\tau_B + \tau_R)\, \mathrm{e}^{\hat{\omega} + \sigma_\omega \, x_m}},$$

and the product of the two marginals is the product of two such sums. The value beside it comes from the exact posterior: conditioning on λ and averaging,

$$\Pr(K_B = 0, K_R = 0 \mid \mathbf{y}) = \mathbb{E}\big[\mathrm{e}^{-\lambda\tau_B} \, \mathrm{e}^{-\lambda\tau_R} \mid \mathbf{y}\big] = \left(\frac{b}{b + \tau_B + \tau_R}\right)^a, \tag{11.31}$$

the predictive probability of no failure in the whole fleet. The product of the two marginals is $(b/(b+\tau_B))^a (b/(b+\tau_R))^a$, and the ratio of the joint probability to the product is

$$\frac{\Pr(K_B = 0, K_R = 0 \mid \mathbf{y})}{\Pr(K_B = 0 \mid \mathbf{y}) \, \Pr(K_R = 0 \mid \mathbf{y})} = \left(\frac{(b+\tau_B)(b+\tau_R)}{b\,(b+\tau_B+\tau_R)}\right)^a = \left(1 + \frac{\tau_B \tau_R}{b\,(b+\tau_B+\tau_R)}\right)^a \geq 1, \tag{11.32}$$

since $(b+\tau_B)(b+\tau_R) = b(b+\tau_B+\tau_R) + \tau_B\tau_R$. The product of the marginals understates the joint event. The reason is the same one that makes the covariance positive: a quiet hall B is evidence of a low rate, and a low rate makes a quiet remainder more likely. The conditional probability makes it explicit,

$$\Pr(K_R = 0 \mid K_B = 0, \mathbf{y}) = \frac{\Pr(K_B = 0, K_R = 0 \mid \mathbf{y})}{\Pr(K_B = 0 \mid \mathbf{y})} = \left(\frac{b+\tau_B}{b+\tau_B+\tau_R}\right)^a, \tag{11.33}$$

the predictive probability of no failure in τ_R under the posterior $\text{Gamma}(a, b+\tau_B)$ that a failure-free τ_B would leave. Under the plug-in Poisson, by contrast, the rate is fixed, the counts

are independent and the joint probability is exactly the product, $\mathrm{e}^{-\tau_B\bar{\lambda}}\mathrm{e}^{-\tau_R\bar{\lambda}}$. The ratio of the predictive joint probability to the plug-in one therefore has two factors that multiply: the Jensen gaps of the two marginals, each predictive probability of no failure exceeding its plug-in value, and the ratio (11.32) that the plug-in removes.

> **Principle 11.2.** *Counts that share one uncertain rate are positively correlated even when they are independent given the rate: their conditional means depend on the one entry of the shared rate, so the epistemic part of their covariance matrix has rank one and the off-diagonal entry $\tau_B\tau_R\,\mathrm{Var}(\lambda \mid \mathbf{y})$. A joint event of several counts is an expectation of the joint conditional probability over the one posterior, not a product of marginal predictives.*

Building blocks. The derivatives g_B and g_R at $\hat{\omega}$ and the variance σ^2_ω (Section 11.2.1); the sums (11.13) with the exposures τ_B, τ_R and their total; and, for the exact values beside them, the marginal predictives (11.16) with those exposures and the gamma parameters a and b of equations (11.31) to (11.33).

Hall B alone. Hall B has $\tau_B = 12$ unit-years next year. Its predictive is equation (11.16) with $\pi_B = 126/138$. By equation (11.19) with $g_B = 12 \times 0.2063 = 2.476$, its mean is $\mu(\hat{\omega}) = 2.48$ failures and its variance $2.476 + 2.476^2/26 = 2.476 + 0.236 = 2.71$, standard deviation 1.65, with an epistemic share of $0.236/2.71 = 12/138 = 8.7\,\%$. By the sum (11.13) with τ_B the probability of no failure in hall B is 0.090 and the probability of at least three 0.456. The values beside them come from the exact posterior: the probability of no failure in hall B is $(126/138)^{26} = 0.094$, and the probability of at least three is $1 - 0.0939 - 0.2124 - 0.2493 = 0.444$, gaps of 0.004 and 0.012. The plug-in Poisson gives 0.084 and 0.450.

The halls together. The rest of the fleet has $\tau_R = 32$ unit-years, a predictive mean of 6.60 and, by the sum, a probability of no failure of 0.00243, with beside it the exact $(126/158)^{26} = 0.00278$. The derivatives of the two conditional means at $\hat{\omega}$ and the two parts of the covariance matrix (11.29) are

$$g_B = 12 \times 0.2063 = 2.4762, \qquad g_R = 32 \times 0.2063 = 6.6032, \qquad \sigma^2_\omega = 0.038462,$$

$$\sigma^2_\omega \begin{bmatrix} g_B^2 & g_B g_R \\ g_B g_R & g_R^2 \end{bmatrix} = \begin{bmatrix} 0.2358 & 0.6289 \\ 0.6289 & 1.6770 \end{bmatrix},$$

$$\mathbf{C} = \begin{bmatrix} 2.4762 & 0 \\ 0 & 6.6032 \end{bmatrix} + \begin{bmatrix} 0.2358 & 0.6289 \\ 0.6289 & 1.6770 \end{bmatrix} = \begin{bmatrix} 2.7120 & 0.6289 \\ 0.6289 & 8.2802 \end{bmatrix}.$$

The epistemic part has rank one, its correlation $0.6289/\sqrt{0.2358 \times 1.6770} = 1$, and the covariance of the two counts is 0.629; the diagonal of $\mathbf{C}$ gives the standard deviations 1.65 and 2.88; the values are those of the exact posterior, whose covariance is the same entry. By equation (11.30) the correlation of the counts is $0.6289/\sqrt{2.7120 \times 8.2802} = \sqrt{0.0870 \times 0.2025} = 0.133$. By the sum over the whole year, the probability that no unit of the fleet fails next year is 3.36×10^{-4}, and the product of the two marginals is $0.0898 \times 0.00243 = 2.18 \times 10^{-4}$, smaller by the factor 1.54; a failure-free hall B raises the probability that the rest of the fleet has no failure from 0.00243 to 0.00375, the ratio of the two. The values beside them come from the exact posterior: by equation (11.31)

$$\Pr(K_B = 0, K_R = 0 \mid \mathbf{y}) = \left(\frac{126}{170}\right)^{26} = 4.15 \times 10^{-4},$$

while the product of the two marginals is $0.0939 \times 0.00278 = 2.61 \times 10^{-4}$, smaller by the factor $(1 + 384/21\,420)^{26} = 1.587$ of equation (11.32). By equation (11.33) a failure-free hall B raises the probability that the rest of the fleet has no failure from 0.00278 to $(138/170)^{26} = 0.00442$.

Table 11.9 collects the numbers: each row is a quantity, the columns give it under the predictive, Laplace with exact in parentheses, and under the plug-in Poisson. Under the plug-in the correlation is zero and the joint probability equals the product, both 1.14×10^{-4}, a factor of 2.95 below the Laplace joint probability, with the two factors of Section 11.3.3,

$$\frac{3.36 \times 10^{-4}}{1.14 \times 10^{-4}} = \underbrace{\frac{0.0898}{0.0841}}_{\text{hall B, } 1.068} \times \underbrace{\frac{0.00243}{0.00136}}_{\text{rest, } 1.792} \times \underbrace{1.54}_{\text{correlation}} = 1.914 \times 1.54 = 2.95,$$

and a factor of 3.64 below the exact one:

$$\frac{4.15 \times 10^{-4}}{1.14 \times 10^{-4}} = \underbrace{\frac{0.0939}{0.0841}}_{\text{hall B, } 1.117} \times \underbrace{\frac{0.00278}{0.00136}}_{\text{rest, } 2.052} \times \underbrace{1.587}_{\text{correlation}} = 2.293 \times 1.587 = 3.64.$$

Table 11.9: Hall B (12 unit-years) and the rest of the fleet (32 unit-years) in the coming year, under the Laplace predictive with the shared posterior $\mathcal{N}(-1.5782,\ 0.1961^2)$, the negative-binomial predictive with the shared posterior Gamma(26, 126) in parentheses where it differs, and under the plug-in Poisson.

Quantity	Predictive, Laplace (exact)	Plug-in Poisson
Hall B: mean and std. dev. of K_B [failures]	2.48, 1.65	2.48, 1.57
Hall B: $\Pr(K_B = 0)$ [–]	0.090 (0.094)	0.084
Hall B: $\Pr(K_B \geq 3)$ [–]	0.456 (0.444)	0.450
Rest: mean and std. dev. of K_R [failures]	6.60, 2.88	6.60, 2.57
Rest: $\Pr(K_R = 0)$ [–]	0.00243 (0.00278)	0.00136
Covariance [failures^2] and correlation [–] of K_B and K_R	0.629, 0.133	0, 0
Joint $\Pr(K_B = 0, K_R = 0)$ [–]	3.36×10^{-4} (4.15×10^{-4})	1.14×10^{-4}
Product $\Pr(K_B = 0)\ \Pr(K_R = 0)$ [–]	2.18×10^{-4} (2.61×10^{-4})	1.14×10^{-4}
$\Pr(K_R = 0 \mid K_B = 0)$ [–]	0.00375 (0.00442)	0.00136

11.3.4 The held-out year and the doubled rate

Engineering questions. Was last year's count consistent with the two years before it? Did the rate double in year 3? The predictive of the coming year assumes one rate over the record; the first question asks whether the record is surprising under that model, the second whether a stated alternative explains it better.

General forms. The first is the threshold form, equation (2.12), taken on the predictive of a year built from the other years, with the observed count as the level: the upper predictive p-value $\Pr(y_{\text{new}} \geq y_{\text{obs}} \mid \mathbf{y}_{-i})$, with $\mathbf{y}_{-i}$ the evidence without year i and y_{obs} its observed count. The second is the ratio of the evidences of two complete models fitted to the same record, equation (2.14),

$$B_{10} = \frac{p(\mathbf{y} \mid M_1)}{p(\mathbf{y} \mid M_0)}.$$

Specialized forms. Two choices enter here. The year checked is held out of its own predictive: this keeps the year from judging itself, and it costs a third of the record, so each held-out predictive is wider than the full one; with only two years left beside it, the check has little power against a slow drift of the rate (Section 11.6.6). The alternative is fixed in advance, a rate doubled in year 3 with the same prior on λ: this gives its evidence a closed form, and it misleads if the alternative is chosen after the count is seen, while a change of unknown size needs a prior on the size, which the comparisons of Chapter 24, in Volume II, carry.

The held-out check. The first question is equation (2.12) with $y_{\text{new}} = K_3$, the level $f^* = k_3$ and $p(\mathbf{z} \mid \mathbf{y})$ the posterior of the rows that remain when year 3 is held out, which Section 11.2.1 gives: the Newton iteration on the three remaining rows returns the held-out value $\hat{\omega}_{1:2} = \ln(a_{1:2}/b_{1:2})$ with the information $\Lambda_{1:2} = a_{1:2}$, the sum of the expected counts of those rows, and the variance $\sigma^2_{\omega,1:2} = 1/a_{1:2}$. The tail is the sum (11.13) over $\mathcal{N}(\hat{\omega}_{1:2}, \sigma^2_{\omega,1:2})$ with the exposure h_3, and equation (11.19) with $\tau = h_3$ gives the predictive mean and variance of K_3. The value beside it comes from the exact posterior. The other two years give the posterior $\text{Gamma}(a_{1:2}, b_{1:2})$ with $a_{1:2} = a_0 + k_1 + k_2$ and $b_{1:2} = b_0 + h_1 + h_2$, and the predictive of year 3's count is equation (11.16) with these parameters and the exposure h_3. Its upper predictive p-value is

$$\Pr(K_3 \geq k_3 \mid \mathbf{y}_{1:2}) = 1 - \sum_{j=0}^{k_3-1} \frac{\Gamma(a_{1:2}+j)}{\Gamma(a_{1:2})\, j!}\, \pi_3^{a_{1:2}} (1-\pi_3)^j, \qquad \pi_3 = \frac{b_{1:2}}{b_{1:2}+h_3}, \tag{11.34}$$

with $\mathbf{y}_{1:2}$ the prior and the record of years 1 and 2, and the lower tail $\Pr(K_3 \leq k_3 \mid \mathbf{y}_{1:2})$ is the companion for a count that is suspiciously low. A small tail says that the count is unusual under the model. It does not say what else would explain it. In the fit, holding out a year deletes its term from the objective and its slope and curvature from the sums (11.7) (Section 11.1.5). The check must be made with the predictive and not with the plug-in Poisson at the held-out mean: a check with the rate fixed at its estimate ignores the uncertainty of the estimate, gives too thin a tail and flags too many years as anomalies.

The comparison as a ratio of evidences. The second question is equation (2.14). Let M_0 be the model of the chapter, one rate for all three years, and M_1 the model in which the rate doubled in year 3: years 1 and 2 at λ, year 3 at 2λ, with the same prior on λ. The evidence of each model is the probability of the whole record with λ integrated out, and the chain rule splits it into the first two years and the third given them:

$$p(\mathbf{y} \mid M) = p(k_1, k_2 \mid M)\; p(k_3 \mid k_1, k_2, M).$$

The two models agree about years 1 and 2, so the first factor is common and cancels in the ratio. The second factor is the evidence of year 3's record under the posterior from years 1 and 2 as its prior, equation (6.68) of Chapter 6, which is the predictive probability of k_3. Under M_0 it is equation (11.16) with exposure h_3; under M_1 the Poisson mean of year 3 is $2\lambda h_3$, which is the same predictive with exposure $2h_3$. In the state the two models share the one entry ω and differ in one row: under M_1 year 3's row has the exposure $2h_3$, the slope $2h_3 \mathrm{e}^{\omega} - k_3$ and the curvature $2h_3 \mathrm{e}^{\omega}$. The Bayes factor is therefore the ratio of two predictives of the same count over one posterior of the state, $\mathcal{N}(\hat{\omega}_{1:2}, \sigma^2_{\omega,1:2})$, and on the state it is the ratio of two sums (11.13), with the exposures $2h_3$ and h_3:

$$B_{10} \approx \frac{\sum_m \nu_m\, P_{k_3}(x_m;\; 2h_3)}{\sum_m \nu_m\, P_{k_3}(x_m;\; h_3)}, \qquad P_j(x;\; h) = \frac{\mu^j\, \mathrm{e}^{-\mu}}{j!}, \quad \mu = h\, \mathrm{e}^{\hat{\omega}_{1:2} + \sigma_{\omega,1:2}\, x}, \tag{11.35}$$

in which the factorials cancel, so that the count k_3 may be read as a real number in the powers and the break-even count is the root of $B_{10} = 1$. The value beside it comes from the exact posterior. The Bayes factor is then the ratio of two negative-binomial probabilities, in which

the gamma-function factors cancel:

$$B_{10} = \frac{\pi_{3,1}^{a_{1:2}}(1-\pi_{3,1})^{k_3}}{\pi_3^{a_{1:2}}(1-\pi_3)^{k_3}}, \qquad \pi_{3,1} = \frac{b_{1:2}}{b_{1:2}+2h_3}.$$

With $1-\pi_3 = h_3/(b_{1:2}+h_3)$ and $1-\pi_{3,1} = 2h_3/(b_{1:2}+2h_3)$, the powers of $b_{1:2}$ and h_3 cancel too, and

$$B_{10} = 2^{k_3}\left(\frac{b_{1:2}+h_3}{b_{1:2}+2h_3}\right)^{a_{1:2}+k_3}, \qquad \ln B_{10} = k_3 \ln 2 - (a_{1:2}+k_3)\ln\frac{b_{1:2}+2h_3}{b_{1:2}+h_3}. \tag{11.36}$$

The first term grows with the count, each failure adding $\ln 2$ in favor of the doubled rate; the second penalizes the doubled rate for the larger count it predicts and did not see. The ratio favors M_1 when the count exceeds the break-even value

$$k^* = \frac{a_{1:2}\,\ln\big((b_{1:2}+2h_3)/(b_{1:2}+h_3)\big)}{\ln 2 - \ln\big((b_{1:2}+2h_3)/(b_{1:2}+h_3)\big)}. \tag{11.37}$$

Two questions. The tail (11.34) and the ratio (11.36) answer different questions, as Chapter 2 classifies them. The first is a tail of one model's predictive, the threshold form applied to the observed count; it needs no alternative and asks whether the model is surprised. The second is a hypothesis comparison, a ratio of evidences between two complete models; it needs a stated alternative and asks which of the two predicted the count better. A count can be unsurprising under M_0 and still favor an alternative a little, or be surprising under M_0 and favor a particular alternative only weakly, because the alternative may be surprised as well.

Building blocks. The fit of the three remaining rows with its $\hat{\omega}_{1:2}$ and $\sigma^2_{\omega,1:2}$, the derivative $g = h_3 e^{\hat{\omega}_{1:2}}$, and the sums (11.13) and (11.35) over them; and, for the exact values beside them, the held-out gamma parameters $a_{1:2}$ and $b_{1:2}$, the cumulative sums of the held-out predictive (11.34) and the two logarithms of equations (11.36) and (11.37). The check is computed for every year in turn.

The held-out fit. Years 1 and 2 give the posterior Gamma(18, 86) of Table 11.3. In the fit, the sums (11.7) lose the slope and the curvature of year 3; Newton's method from $\ln 0.2$ converges in three steps to $\hat{\omega}_{1:2} = \ln(18/86) = -1.5640$, where the expected counts are (2.093, 7.535, 8.372), the rows' slopes (0.093, 0.535, −0.628), the information $\Lambda_{1:2} = 18$ and $\sigma^2_{\omega,1:2} = 0.05556$. Year 3 had $h_3 = 40$ unit-years, so $g = 40\,e^{-1.5640} = 8.372$, and equation (11.19) gives the mean 8.37 failures and the variance $8.372 + 8.372^2 \times 0.05556 = 12.27$, standard deviation 3.50, the moments of the exact predictive with $\pi_3 = 86/126$. The central 95 % interval of the Laplace predictive is [3, 17], and of the exact one [3, 16].

The tails. By the sum (11.13) over $\mathcal{N}(-1.5640,\ 0.05556)$ with the exposure 40,

$$\Pr(K_3 \geq 8 \mid \mathbf{y}_{1:2}) \approx 0.587, \qquad \Pr(K_3 \leq 8 \mid \mathbf{y}_{1:2}) \approx 0.528,$$

neither tail small: year 3 is what the two years before it predicted. The values beside them come from the exact posterior, by equation (11.34),

$$\Pr(K_3 \geq 8 \mid \mathbf{y}_{1:2}) = 0.564, \qquad \Pr(K_3 \leq 8 \mid \mathbf{y}_{1:2}) = 0.551,$$

gaps of 0.023 in each tail and no change in the reading. Had year 3 brought sixteen failures, the upper tail would have been 0.041 on the state route and $\Pr(K_3 \geq 16 \mid \mathbf{y}_{1:2}) = 0.033$ exactly; under the plug-in Poisson at the held-out mean 8.37 the same tail is 0.012, more than three times smaller than the first and nearly three times smaller than the second.

Every year held out. Table 11.10 runs the check on each year in turn: each row holds one year out, the columns give the posterior from the other two years, the predictive mean and 95 % interval of the held-out count, the observed count, its two predictive tails and the log Bayes factor of a doubled rate in that year, which Section 11.3.4 reads, each Laplace with exact in parentheses. The last row is the hypothetical sixteen in year 3. No observed year has a tail below 0.41 on the state route, or below 0.39 on the exact posterior.

Table 11.10: Predictive checks of each year of the record against the other two, and the log Bayes factor of a doubled rate in the held-out year: Laplace, the sums (11.13) and (11.35) over $\omega \sim \mathcal{N}(\ln(a/b),\ 1/a)$; in parentheses exact, equations (11.34) and (11.36). The mean is the delta form, equal to the exact one.

Held-out year	Posterior from the others (a, b)	Predictive mean [failures], 95 % interval	Observed [failures]	$\Pr(K \geq k)$ [–]	$\Pr(K \leq k)$ [–]	$\ln B_{10}$, doubled [nats]
year 1 (36 unit-years)	(19, 90)	7.60, [2, 15]	7	0.621 (0.600)	0.503 (0.525)	−1.81 (−1.68)
year 2 (40)	(17, 86)	7.91, [2, 16] ([2, 15])	9	0.419 (0.395)	0.684 (0.706)	−1.04 (−0.93)
year 3 (40)	(18, 86)	8.37, [3, 17] ([3, 16])	8	0.587 (0.564)	0.528 (0.551)	−1.76 (−1.62)
year 3, hypothetical	(18, 86)	8.37, [3, 17] ([3, 16])	16	0.041 (0.033)	0.974 (0.980)	+1.57 (+1.72)

The doubled rate in year 3. The sums (11.35) over $\mathcal{N}(-1.5640,\ 0.05556)$ give $\ln B_{10} \approx -1.76$ nats at $k_3 = 8$ and $+1.57$ at $k_3 = 16$: the observed eight favors the constant rate by a factor of 5.8, a sixteen would have favored the doubled rate by a factor of 4.8, and the break-even count, the root of $B_{10} = 1$, is 12.1, so that thirteen or more failures would have favored doubling. The values beside them come from the exact posterior. Equation (11.36) with $a_{1:2} = 18$, $b_{1:2} = 86$ and $h_3 = 40$ gives $\ln((86 + 80)/(86 + 40)) = \ln(166/126) = 0.2757$ and

$$\begin{aligned} k_3 = 8: &\quad \ln B_{10} = 8 \ln 2 - 26 \times 0.2757 = 5.545 - 7.168 = -1.62, \\ k_3 = 16: &\quad \ln B_{10} = 16 \ln 2 - 34 \times 0.2757 = 11.090 - 9.374 = +1.72. \end{aligned}$$

The observed eight favors the constant rate by 1.62 nats, a factor of five; a sixteen would have favored the doubled rate by 1.72 nats, a factor of 5.6. The break-even count (11.37) is $18 \times 0.2757/(0.6931 - 0.2757) = 11.9$: twelve or more failures in year 3 would have favored doubling, eleven or fewer the constant rate. The state route asks for one failure more, and leans 0.14 nats toward the constant rate at both counts, because the normal in ω puts its weight on higher rates, and a higher rate helps the constant rate and hurts the doubled one. The two routes differ only at a count of twelve, which was not observed. By the last column of Table 11.10 every observed year favors the constant rate against a doubling, by 1.0 to 1.8 nats (0.9 to 1.7 exactly).

The check and the comparison appear together in Figure 11.4. The horizontal axis is year 3's count, the vertical axis its probability in percent. The bars are the predictive of year 3 from years 1 and 2 under the constant rate; the blue bar is the observed eight, and the red bars the counts of sixteen or more, whose total is the tail 0.033. The purple circles are the predictive under the doubled rate, the same negative binomial with exposure 80 unit-years, and the dotted blue line with small dots is the Laplace predictive of year 3 under the constant rate. At eight the bar is five times the circle; at sixteen the circle is 5.6 times the bar: the tail is an area under one distribution, the Bayes factor a ratio of two distributions at one count.

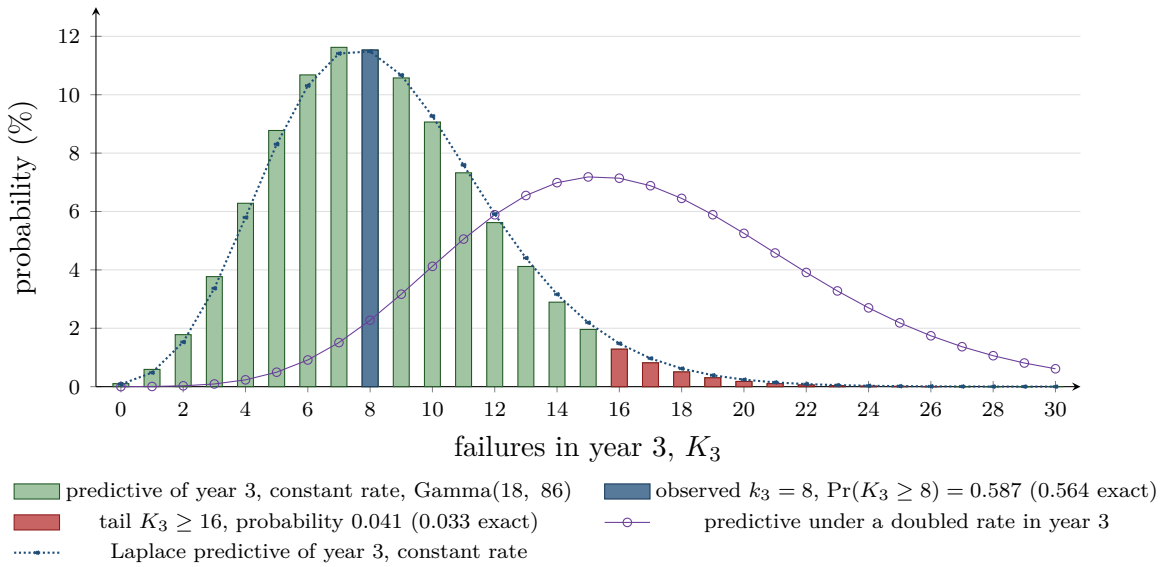


Figure 11.4: The check of year 3 against years 1 and 2. Bars: the negative-binomial predictive of year 3's count, equation (11.16) with $a_{1:2} = 18$, $b_{1:2} = 86$ and $h_3 = 40$ unit-years; blue, the observed count; red, the tail of sixteen or more. Circles: the predictive under a doubled rate in year 3, the same with exposure 80 unit-years. Dotted blue: the Laplace predictive, the sum (11.13) over $\mathcal{N}(-1.5640,\ 0.05556)$. Log Bayes factors of the doubled rate: -1.76 at 8 and $+1.57$ at 16 (-1.62 and $+1.72$ exact).

11.4 The checks (step 5)

The fleet has no physics row and no exact constraint, so no limit $\sigma_p \to 0$ of a model tolerance arises and no limit table is needed; the positivity of the rate is not a row either, and the log-rate state handles it.

11.4.1 The delta form against the summed predictive

The moments by summation. The delta form (11.19) gives the moments from $\hat{\omega}$ and σ_ω^2 alone; the full calculation sums the exact negative-binomial predictive (11.16) term by term, from $j = 0$ to $j = 2000$, where its probabilities add to one to twelve decimals. For the two halls the full calculation is the joint distribution of the pair, the two conditional Poisson probabilities averaged over the posterior by the same gamma integral,

$$\Pr(K_B = i,\ K_R = j \mid \mathbf{y}) = \frac{\Gamma(a+i+j)}{\Gamma(a)\, i!\, j!}\, \pi^a \left(\frac{\tau_B}{b+\tau}\right)^i \left(\frac{\tau_R}{b+\tau}\right)^j, \qquad \tau = \tau_B + \tau_R, \quad \pi = \frac{b}{b+\tau},$$

summed over $i, j < 200$. Table 11.11 sets the two routes side by side: each row is a quantity, the columns give it from the delta form, by summation and the absolute difference. The rows are the mean and the variance of the fleet's count for three horizons, the covariance and the correlation of the halls' counts with their joint probability of no failure, and the moments of year 3's count from years 1 and 2.

The two routes agree to round-off in every row, including the covariance, which the delta form takes from the derivatives g_B and g_R with respect to the one entry, as the two identities of Section 11.2.2 require; the check confirms the arithmetic, and the approximation is tested on the whole distribution in Section 11.4.2.

The hand checks. Setting the derivative $be^\omega - a$ of the one-term objective (11.5) to zero by hand gives the same $\hat{\omega} = \ln(26/126)$, and its second derivative there the same curvature 26.

Table 11.11: The delta form (11.19) and the covariance matrix (11.29), from $\hat{\omega}$ and σ_ω^2, against the moments by summation of the exact predictive. Posterior Gamma(26, 126), except year 3 from Gamma(18, 86); horizons in unit-years.

Quantity	Delta form	By summation	Difference
Mean, $\tau = 11$ [failures]	2.2698	2.2698	$< 10^{-12}$
Variance, $\tau = 11$ [failures^2]	2.4680	2.4680	$< 10^{-12}$
Mean, $\tau = 44$ [failures]	9.0794	9.0794	$< 10^{-12}$
Variance, $\tau = 44$ [failures^2]	12.2499	12.2499	$< 10^{-12}$
Mean, $\tau = 176$ [failures]	36.3175	36.3175	$< 10^{-12}$
Variance, $\tau = 176$ [failures^2]	87.0466	87.0466	$< 10^{-12}$
Covariance C_{BR} of the halls [failures^2]	0.6289	0.6289	$< 10^{-12}$
Correlation of the halls [–]	0.1327	0.1327	$< 10^{-12}$
$\Pr(K_B = 0, K_R = 0)$ [–]	4.1492×10^{-4}	4.1492×10^{-4}	$< 10^{-12}$
Held-out year 3, mean [failures]	8.3721	8.3721	$< 10^{-12}$
Held-out year 3, variance [failures^2]	12.2661	12.2661	$< 10^{-12}$

Updating the prior one year at a time gives Gamma(9, 46) after year 1, Gamma(18, 86) after year 2 and Gamma(26, 126) after year 3, with the posterior means 0.196, 0.209 and 0.206 per unit-year: the posterior of the pooled record, as rows that add must give, in any order and grouping.

11.4.2 The Laplace predictive against the exact predictive

The comparison. Table 11.12 sets four routes side by side for the coming year: the answers of Section 11.3, the one further normal that tests where their error lies, and the plug-in. Each row is a summary of the predictive; the columns give it under the exact negative binomial, under the Laplace normal on ω by quadrature, under a normal on ω with the exact mean $\psi(a) - \ln b$ and variance $\psi'(a)$ of the log rate, and under the plug-in Poisson. The Laplace mean follows from equation (11.14): $9.079\,\mathrm{e}^{1/52} = 9.256$, two percent above the exact mean. The delta form (11.19), which gives only the two moments, returns 9.079 and 3.500 from the same $\hat{\omega}$ and σ_ω^2, the exact values; the table compares whole distributions.

Table 11.12: The predictive of the coming year's count by four routes. Exact: the negative binomial (11.16). Laplace: $\omega \sim \mathcal{N}(-1.5782,\ 0.1961^2)$, quadrature (11.13). Exact moments: $\omega \sim \mathcal{N}(\psi(26) - \ln 126,\ \psi'(26)) = \mathcal{N}(-1.5975,\ 0.1980^2)$. Plug-in: Poisson(9.079).

Summary	Exact	Laplace	Exact moments	Plug-in
Mean of K_f [failures]	9.079	9.256	9.082	9.079
Std. dev. of K_f [failures]	3.500	3.552	3.518	3.013
$\Pr(K_f = 0)$ [–]	4.15×10^{-4}	3.36×10^{-4}	3.91×10^{-4}	1.14×10^{-4}
$\Pr(K_f > 12)$ [–]	0.161	0.174	0.162	0.130
$\Pr(K_f > 15)$ [–]	0.045	0.051	0.046	0.024
95 % interval [failures]	[3, 17]	[3, 17]	[3, 17]	[4, 15]
95 % spares level [spares]	15	16	15	14
99 % spares level [spares]	18	19	19	17

Reading the table. The Laplace route, which the chapter reports, is close to the exact predictive here: its tails exceed the exact ones by 0.013 above twelve and 0.006 above fifteen, and its 95 % interval is the same. The closeness comes from the size of the record, $a = 26$ failures and prior pseudo-failures, which makes the posterior of ω nearly normal. Chapter 8 worked with gamma posteriors of shape six, where the Laplace normal understated the variance

of each log rate by about 8 % and shifted the band of a quantity that depends on one log rate by about 9 %; Exercise 11.4 repeats the present comparison at $a = 4$. The remaining error is almost all the center: the normal with the exact moments of $\ln \lambda$, whose center is 0.019 lower, reproduces the exact mean to 0.003 and the tails to about 0.001; to first order the factor $e^{\sigma_\omega^2/2}$ and the center offset $\ln a - \psi(a)$ cancel, 0.019231 against 0.019354 at $a = 26$. The gap is small in probability and still visible in the decision: the Laplace level at 95 % is 16 against the exact 15, because the exact cumulative probability at 15, 0.9553, is close to the line and the Laplace one, 0.9489, falls below it. The 99 % level is a boundary case for both normals: the exact cumulative probability at 18 is 0.9906, the two normals give 0.9883 and 0.9897, and both call for 19. A spares level at such a boundary moves by one with a change of a few thousandths in the cumulative probability, which is why Section 11.5.2 gives both.

11.4.3 The scatter of the record

The three count rows are redundant for the one entry: three rows, one fitted entry, two degrees of freedom. The scatter statistic is the Pearson sum: each year's squared slope over its curvature at $\hat{\omega}$, $(\mu_k - n_k)^2/\mu_k$ with the slopes of equation (11.15), which is the square of the row's working residual $(\mu_k - n_k)/\sqrt{\mu_k}$ of equation (4.43):

$$\frac{0.4286^2}{7.4286} + \frac{0.7460^2}{8.2540} + \frac{0.2540^2}{8.2540} = 0.1572^2 + (-0.2597)^2 + 0.0884^2 = 0.100,$$

and against the chi-square distribution with two degrees of freedom its upper tail is 0.951 and its lower tail 0.049. With counts of seven to nine the chi-square reference is coarse, and the full calculation conditions on the total: given $k = 24$ failures, the three yearly counts are multinomial with the probabilities $h_i/T_{\text{tot}} = 36/116$, $40/116$ and $40/116$ whatever the rate, and the sum over all 325 splits of the 24 failures gives the exact tails of the same statistic evaluated at the maximum-likelihood rate, 0.0995: the probability of a statistic at least as large is 0.968, and of one at most as large 0.097. The record shows no sign of the extra scatter that a drifting rate or a cause shared within a year would produce, and the closeness of its three years is not unusual either: about one record in ten agrees at least as closely by chance, where the chi-square reference would say one in twenty.

11.4.4 Sensitivity to the prior

The handbook prior (11.2) weighs as much as two failures in ten unit-years, a twelfth of the record, and the spares decision should not hang on it. Table 11.13 repeats the decision numbers under four other priors: each row is a prior, the columns give the posterior from equation (11.11), the predictive mean and standard deviation of next year's count from equation (11.19), the tail above fifteen, the spares levels for 95 % and 99 % by equation (11.25), the plug-in 95 % level with the service level it delivers under that prior's predictive, and the log Bayes factor of a doubled rate in year 3 by equation (11.35), each Laplace with exact in parentheses; the row of a prior of shape a_0 and rate parameter b_0 has the state $\hat{\omega} = \ln(a/b)$, $\sigma_\omega^2 = 1/a$. The Jeffreys prior, with density proportional to $\lambda^{-1/2}$, and the flat prior are the limits $b_0 \to 0$ of Gamma$(1/2, b_0)$ and Gamma$(1, b_0)$, and add no unit-years; the last two rows halve and double the handbook prior's weight at the same mean.

On the state route the three priors of mean 0.2 give the spares levels 16, 16 and 15 at 95 % and 19 at 99 %, the fifteen of the heaviest prior being the one case where the Laplace tail above fifteen, 0.048, falls inside the line; the tail moves between 0.048 and 0.053. On the exact posterior the levels stay at 15 and 18, and the tail moves between 0.043 and 0.046. The two priors that add no unit-years put the posterior mean 2.4 and 4.4 percent higher, since each keeps a pseudo-failure, or half of one, with no unit-years beside it; on the exact posterior they raise both levels by one spare, to 16 and 19, and on the state route the flat prior takes the 99 %

Table 11.13: The fleet's decision numbers under five priors, record $k = 24$ in $T_{\text{tot}} = 116$ unit-years, coming year $\tau = 44$ unit-years. The first row is the handbook prior. Laplace, with exact in parentheses; the mean and standard deviation are the delta form, equal to the exact ones.

Prior	Posterior (a, b)	Mean, sd [failures]	$\Pr(K_f > 15)$ [–]	$s_{0.95}$, $s_{0.99}$ [spares]	Plug-in $s_{0.95}$, delivered [–]	$\ln B_{10}$ [nats]
Gamma(2, 10), handbook	(26, 126)	9.08, 3.50	0.051 (0.045)	16, 19 (15, 18)	14, 0.921 (0.929)	−1.76 (−1.62)
Jeffreys, $\propto \lambda^{-1/2}$	(24.5, 116)	9.29, 3.58	0.060 (0.053)	16, 19 (16, 19)	15, 0.940 (0.947)	−1.87 (−1.71)
flat	(25, 116)	9.48, 3.62	0.067 (0.059)	16, 20 (16, 19)	15, 0.933 (0.941)	−2.03 (−1.86)
Gamma(1, 5)	(25, 121)	9.09, 3.52	0.053 (0.046)	16, 19 (15, 18)	14, 0.919 (0.928)	−1.74 (−1.60)
Gamma(4, 20)	(28, 136)	9.06, 3.46	0.048 (0.043)	15, 19 (15, 18)	14, 0.924 (0.932)	−1.81 (−1.67)

level to 20. Under every prior the plug-in 95 % level is below the predictive one on both routes and delivers between 0.919 and 0.940 (0.928 and 0.947) instead of 0.95, and every prior favors the constant rate against a doubling in year 3, by 1.74 to 2.03 nats (1.60 to 1.86). The shortfall of the plug-in is a property of the record and the horizon, not of the prior.

11.5 The report (step 6)

The record does not support the first reading of it. Nine failures at a Poisson variance of nine put the store at fourteen compressors for 95 % cover, but the count of next year has the variance 12.25, not 9.08: a quarter of it, 3.17, is the uncertainty of a rate known from twenty-four failures, and with it the store needs more than fourteen: sixteen compressors on the posterior of the state, fifteen on the exact posterior, the two routes falling on opposite sides of the 95 % line at fifteen. The fourteen of the first reading cover the year with probability 0.921 (0.929 exactly). The three years of record are consistent with one rate, and the halls' counts are correlated through it. The answers below are given in the units in which the questions were asked.

11.5.1 The answers and their bands

Table 11.14 gives the answers. Each row is a question of Table 11.2, in its order; the second column gives the answer from the posterior of the state, $\omega \mid \mathbf{y} \approx \mathcal{N}(-1.5782,\ 0.1961^2)$ and its held-out counterparts; the third the exact value, from the negative-binomial predictive and the gamma posteriors; and the last its band, with the plug-in and prior-sensitivity values of Section 11.4.4 beside it.

The report line. Under a constant failure rate shared by all units, the fleet of forty-four CRAC units should expect about nine compressor failures next year, with a 95 % predictive range from 3 to 17; the probability of more than fifteen is 5.1 % (4.5 % from the exact posterior). Sixteen spare compressors cover the year with probability 0.968 and nineteen with 0.993; from the exact posterior fifteen cover it with 0.955 and eighteen with 0.991. A store sized from the estimated rate alone would hold fourteen and seventeen and deliver 0.921 and 0.980 (0.929 and 0.984). The halls' counts are positively correlated through the shared rate, and joint events must be computed together. The three years of record are consistent with one rate.

Table 11.14: The answers to the questions asked of the fleet's record, from the posterior of the state, with the exact value beside each. Counts are failures in the coming year of 44 unit-years unless stated; intervals are central 95 % predictive intervals.

Question	Answer (state route)	Exact	Band and comparison
How many failures next year, and within what range?	mean 9.08, std. dev. 3.50, most probable 8; [3, 17]	the same	plug-in [4, 15], std. dev. 3.01
How does the range grow with the horizon?	epistemic share 26 % next year, 58 % over four years; four years [21, 58]	the same shares; four years [20, 56]	plug-in over four years [25, 49]
More than 12; more than 15	0.174; 0.051	0.161; 0.045	plug-in 0.130 above 12, 0.024 above 15; 0.048 to 0.067 (0.043 to 0.059) above 15 over the priors
Spares for 95 %; for 99 %	16; 19	15; 18	opposite sides of the line at 15 and at 18; plug-in 14 and 17, delivering 0.921 and 0.980 (0.929 and 0.984); 15 to 16 and 19 to 20 (15 to 16 and 18 to 19) over the priors
Hall B; halls independent?	mean 2.48, $\Pr(K_B \geq 3) = 0.456$, $\Pr(K_B = 0) = 0.090$; correlation 0.133	0.444, 0.094; 0.133	no failure in the fleet 3.36×10^{-4} (4.15×10^{-4}) against the product 2.18×10^{-4} (2.61×10^{-4})
Year 3 consistent with years 1 and 2?	predictive p-value $\Pr(K_3 \geq 8 \mid \mathbf{y}_{1:2}) = 0.587$	0.564	lower tail 0.528 (0.551); no year below 0.41 (0.39); a sixteen would have given 0.041 (0.033)
Did the rate double in year 3?	$\ln B_{10} = -1.76$, the constant rate favored by a factor of 5.8	-1.62, a factor of five	break-even count 12.1 (11.9); -1.74 to -2.03 (-1.60 to -1.86) over the priors

11.5.2 The spares decision

The decision follows from the fourth row of Table 11.14 and from Table 11.8, and the two routes do not reach the same level. A store for 95 % cover holds sixteen compressors on the state route and fifteen on the exact posterior, and one for 99 % nineteen and eighteen. They fall on opposite sides of the line at the lower level of each pair: fifteen compressors are worth 0.9489 on the state route and 0.9553 on the exact posterior, and eighteen are worth 0.9883 and 0.9906. The normal in the log rate errs upward because its center lies 0.019 above the mean of the exact log rate and its upper tail is the heavier one; the record of twenty-four failures keeps that error to a few thousandths, and at these two boundaries a few thousandths are the whole margin. A plan built on the plug-in holds fourteen and seventeen compressors, below the level of either route, delivers 0.921 (0.929) where it claims 0.956 and 0.980 (0.984) where it claims 0.994, and expects more than twice to more than four times as many failures to wait for a compressor as it reports (two to four times on the exact posterior). The decision does not hang on the prior: on the exact posterior the levels stay at fifteen and eighteen under the priors of mean 0.2 and rise by one under the two that add no exposure, and on the state route they lie between fifteen and sixteen and between nineteen and twenty. Whether the missing compressor is worth buying depends on the price of a compressor in store against the price of a failure that waits for one, a question of expected cost that the decision chapters of the book take up; Exercise 11.5 works one case.

Common mistake

Sizing a store of spares by the Poisson distribution at the estimated rate. The estimated rate is the center of a posterior whose 95 % interval runs from 0.135 to 0.293 per unit-year; the count must be predicted with that uncertainty averaged in. For the fleet the plug-in 95 % level of 14 spares covers the year with probability 0.921 (0.929 exactly), not 0.95, and the shortfall grows with the horizon, since the epistemic share of the variance does.

Common mistake

Multiplying the predictive probabilities of two halls, or of two units, as if their counts were independent. They share the fleet's rate, and given only the record their counts are positively correlated; the probability that the whole fleet has no failure is 3.36×10^{-4} (4.15×10^{-4} exactly), not the product 2.18×10^{-4} (2.61×10^{-4}). A joint event is computed as one expectation over the shared posterior.

11.6 The next measurement (step 7)

The predictive standard deviation of 3.50 failures is 0.49 above the Poisson floor of 3.01, the spread the count would keep at a known rate. Records could be bought to narrow it, and the preposterior values each before its count is known.

11.6.1 The candidate records

Two records could narrow the prediction, and only one of them is a choice:

- *Two more fleet-years.* The fleet keeps running and the log keeps recording, whatever is decided: $\tau_c = 88$ unit-years and an expected $88 \times 0.2063 = 18.2$ failures.
- *Twenty more units for one year.* A record of twenty units of the same model for one year, pooled from a sister site whose units can be taken to share the rate, which has to be requested and vetted: $\tau_c = 20$ unit-years and an expected 4.1 failures.

Neither record exists before the coming year begins, so in both cases the target is the year of 44 unit-years that follows the record, not the coming year. The subscript c marks the candidate,

to keep it apart from the gamma shape a. Four more fleet-years and an unlimited record are carried beside them as references.

11.6.2 The preposterior with the appended row

The appended row. A candidate exposure τ_c adds a further count row to the rows (11.4), with the expected count $\mu_c = \tau_c e^{\omega}$. Its count, and with it its slope, is unknown, but its curvature (11.6) at the present MAP value, $\bar\lambda\tau_c$ with $\bar\lambda = a/b = e^{\hat\omega}$, does not depend on the count. Added to the curvature of equation (11.15), it gives the preposterior information and variance, by equation (2.17) of Chapter 2 with the appended row:

$$\Lambda_c = \bar\lambda\,\tau_c, \qquad \frac{1}{\sigma^2_{\omega,c}} = \frac{1}{\sigma^2_{\omega}} + \Lambda_c = a + \bar\lambda\tau_c, \qquad \sigma^2_{\omega,c} = \frac{1}{a + \bar\lambda\tau_c} = \frac{1}{a\,(1 + \tau_c/b)}. \tag{11.38}$$

The variance $\sigma^2_{\omega,c}$ is the inverse of the expected information, the information the record brings in expectation over its count; it is not the average $\mathbb{E}[1/(a + K_c)]$ of the variances the possible counts would leave. Chapter 20, in Volume II, writes the same preposterior for a candidate row on a state of many entries as a rank-one update of the covariance matrix. For two more fleet-years, $\tau_c = 88$ and $\bar\lambda\tau_c = 18.159$,

$$\frac{1}{\sigma^2_{\omega,c}} = 26 + 18.159 = 44.159, \qquad \sigma^2_{\omega,c} = 0.02265.$$

With $\sigma^2_{\omega,c}$ in place of σ^2_{ω} in equation (11.19) and the same derivative $g = \tau\bar\lambda$ of the target year, the epistemic part of the predictive variance is

$$g^2\sigma^2_{\omega,c} = \frac{\tau^2\bar\lambda^2}{a + \bar\lambda\tau_c} = \frac{a\tau^2}{b\,(b + \tau_c)}, \tag{11.39}$$

for two fleet-years $9.0794^2 \times 0.02265 = 1.867$, against 3.171 for the present record; the Poisson part $\mu(\hat\omega) = \bar\lambda\tau$ is unchanged.

The exact check. After an added exposure τ_c with K_c failures the posterior is Gamma$(a + K_c, b + \tau_c)$, and by equation (11.20) the predictive variance of the target year, $V(K_c) = \big((a + K_c)\tau/(b + \tau_c)\big)\big(1 + \tau/(b + \tau_c)\big)$, is linear in K_c; its expectation over the predictive of K_c, of mean $a\tau_c/b$, is exactly

$$\mathbb{E}[V(K_c) \mid \mathbf{y}] = \underbrace{\bar\lambda\tau}_{\text{Poisson floor}} + \underbrace{\frac{a\tau^2}{b\,(b + \tau_c)}}_{\text{epistemic part}}, \tag{11.40}$$

whose epistemic part equals $g^2\sigma^2_{\omega,c}$ of equation (11.39) exactly, 1.86679 both ways for two fleet-years. In expectation the Poisson part is untouched by the record: averaged over the counts the record may bring, the posterior mean rate stays $\bar\lambda$, and the expected Poisson variance stays $\bar\lambda\tau$, so that the state route and the exact posterior give one value here, with no gap. The record that is actually obtained moves it, down after a quiet record and up after a busy one. Only the epistemic part shrinks, by the factor $b/(b + \tau_c)$, and as $\tau_c \to \infty$ the expected predictive variance falls to the floor $\bar\lambda\tau = 9.079$, a standard deviation of 3.013. The floor too is an expectation: with an unlimited record the predictive variance becomes $\lambda\tau$ at the true rate, below 9.079 whenever $\lambda < \bar\lambda$, which has the probability $\frac{1}{2}$ under the Laplace normal, $\bar\lambda$ being its median, and 0.526 under the exact posterior. Unlike the preposterior of Chapter 8, where the variance was convex in the future count and inserting its mean was an approximation, here the expected variance is exact; the reported standard deviation is the square root of the expected variance.

11.6.3 The ranking and the floor

Table 11.15 ranks the candidates beside the two references. Each row is a record; the columns give its exposure, the added information Λ_c, which is also its expected number of failures, the preposterior standard deviation of ω, the expected epistemic part of the predictive variance and the square root of the expected predictive variance. The last four columns are not expectations but probabilities over the record the candidate may bring: the preposterior distribution of the 95 % spares level, computed on the state by evaluating equation (11.25) under the posterior of the state after each possible count K_c, $\mathcal{N}(\ln((a+K_c)/(b+\tau_c)),\ 1/(a+K_c))$, and weighting by the Laplace predictive of that count; and the probability that the realized predictive standard deviation falls below the floor of 3.01. The exact values in parentheses take Gamma$(a+K_c, b+\tau_c)$ and the negative-binomial probability of K_c. Each route's level is compared with its own present level, 16 on the state route and 15 on the exact posterior. One row of the arithmetic, two fleet-years: $\Lambda_c = 18.16$, $\sigma^2_{\omega,c} = 1/44.16 = 0.02265$, $\sigma_{\omega,c} = 0.150$, epistemic part $g^2\sigma^2_{\omega,c} = 9.0794^2 \times 0.02265 = 1.867$, total $9.079 + 1.867 = 10.946$, standard deviation 3.308, 5.5 % below the present 3.500.

Table 11.15: The preposterior of the target year's count for two candidate records and two references. Expected epistemic part and expected standard deviation: equation (11.40), one value on both routes. Spares level $s_{0.95}$ and realized standard deviation: probabilities over the count the record may bring, K_c with exposure τ_c, Laplace with exact in parentheses; for the unlimited record, over the true rate. The present level is 16 on the state route and 15 on the exact posterior.

Record	Exposure τ_c [unit-years]	Λ_c (expected failures) [–]	$\sigma_{\omega,c}$ [–]	Expected epistemic part [failures^2]	Expected std. dev. [failures]	Probability of $s_{0.95}$, against the present level [–]			Probability of std. dev. [–]
						below	at	above	below 3.01
present	0	0	0.196	3.171	3.500	0	1	0	0
twenty units, one year	20	4.13	0.182	2.736	3.437	0.592 (0.242)	0.259 (0.513)	0.148 (0.245)	0
two fleet-years	88	18.16	0.150	1.867	3.308	0.600 (0.408)	0.218 (0.277)	0.181 (0.315)	0.061 (0.070)
four fleet-years	176	36.32	0.127	1.323	3.225	0.633 (0.492)	0.188 (0.200)	0.179 (0.308)	0.186 (0.206)
unlimited	∞	∞	0	0	3.013	0.695 (0.563)	0.121 (0.158)	0.184 (0.280)	0.500 (0.526)

The expected predictive standard deviation of the target year's count, as a function of the added exposure, is the curve of Figure 11.5. The horizontal axis is the added exposure in unit-years, the vertical axis the standard deviation in failures, starting at 2.9 rather than at zero. The curve is the square root of equation (11.40); the dots mark the present record, the two candidates and the four-year reference; the dashed line is the Poisson floor. Both are expectations over the record the candidate may bring. The curve falls quickly at first and then flattens toward the floor, which it approaches only as the added exposure grows without bound.

The reason for the ranking. Two more fleet-years narrow the prediction more than twenty more units for one year, in expectation 3.31 against 3.44, and the reason is in the information of equation (11.38): under a constant shared rate every candidate row acts on the one entry of the state with its curvature, $\bar{\lambda}$ times its exposure, so the information grows by $\bar{\lambda}\tau_c$ however the exposure is arranged among units and years, and two fleet-years are 88 unit-years against 20. Twenty unit-years add the same 4.13 to the information whether they come from twenty units

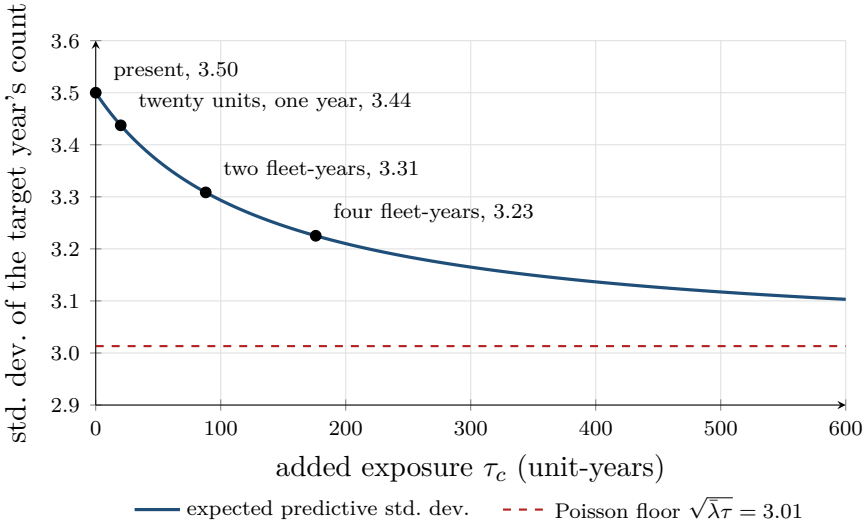


Figure 11.5: The expected predictive standard deviation of the count of a year of 44 unit-years after an added exposure τ_c, the square root of equation (11.40): $\sqrt{9.079\,(1 + 44/(126 + \tau_c))}$. Dots: the present record and the candidates of Table 11.15. Dashed: the floor $\sqrt{9.079} = 3.013$. The vertical axis starts at 2.9.

in one year, as one row or as twenty, or from one unit in twenty. The comparison of price and time, two years of waiting against one year of a sister site's log, is a separate question.

The floor. Both candidates remove little in expectation. The most any record can remove is 0.49 failures of standard deviation, 13.9 %, and two more fleet-years remove 0.19 of it. The floor exists because the count of a future year is random even when the rate is known: the compressors of the coming year fail at random times, and the Poisson variance of their number is a property of the year, not of the knowledge about the rate. The epistemic share of the coming year is only 26 % of the variance, and that share is what a record can buy in expectation.

What the record will actually do. The expected standard deviation hides what matters to the store. The record moves the posterior mean rate as well as narrowing it, and the spares level follows the mean. After two more fleet-years the 95 % level on the state route stays at 16 with probability only 0.22; it falls below with probability 0.60, after a quiet record, and rises above with probability 0.18, after a busy one. The present level of 16 rests on 0.9489 falling just below the line at fifteen, so that any narrowing can bring it down, which is why a fall is the most likely move. On the exact posterior it stays at 15 with probability only 0.28; it falls to 14 or fewer with probability 0.41, and rises to 16 or more with probability 0.32. A quiet record can also bring the realized standard deviation below 3.01, with probability 0.06 (0.07 exactly) after two fleet-years and 0.50 (0.53) for an unlimited record, because the floor $\bar{\lambda}\tau$ is itself the expectation of $\lambda\tau$ over the present posterior. The value of a record for the store is therefore mostly that it settles the rate, and with it the mean count, rather than that it narrows the spread around that mean. For a horizon of one year the record narrows the spread little: two more fleet-years remove in expectation 39 % of the removable standard deviation of the target year, 0.19 of 0.49. For a planning horizon of four years, $\tau = 176$, where the epistemic share is 58 %, the same record takes the expected standard deviation from 9.33 to 8.14 against a floor of 6.03: 1.19 failures, 12.8 % of the present value against 5.5 % for one year.

11.6.4 The recommendation

Do not pursue the sister site's record. Its twenty unit-years remove in expectation 0.06 of the 0.49 failures of standard deviation that any record can remove from the count of the year that follows it, 13 % of the removable part, against 0.19, 39 %, for the fleet's own log over the next two years, which accrues as the fleet runs at no cost beyond keeping it. Keep the log and refit each year: the store for the coming year is sized from the report, 16 spares for 95 % on the state route and 15 on the exact posterior, and each year's count, when it arrives, is a further count row of the form (11.4) with which the fit is run again and the store resized. Resizing matters more than narrowing: after two more fleet-years the 95 % level for the following year moves, by one or more spares up or down, with probability $0.600 + 0.181 = 0.781$, and on the exact posterior $0.408 + 0.315 = 0.723$. The gap does not change the recommendation: the removable standard deviation is one value on both routes.

11.6.5 What the analysis showed

The first reading of the record, a Poisson count of mean 9.08 with its variance taken equal to its mean, put the store at fourteen compressors for 95 % cover; the analysis does not support it. The rate behind that mean is known from twenty-four failures, and its uncertainty makes up a quarter of the variance of next year's count, so the store needs sixteen on the posterior of the state and fifteen on the exact posterior, which fall on opposite sides of the line at fifteen, and fourteen cover the year with probability 0.921 (0.929 exactly), not 0.95. The missing spare is not a rounding accident: the plug-in is short at every service level from 0.90 on and under every prior tried, and its shortfall grows with the horizon, so a store or a budget planned over several years from the estimated rate falls short by more. The two halls cannot be planned apart either. A quiet hall B is evidence of a low rate for the whole fleet, so a year without any failure is more probable than the product of the halls' probabilities, and joint events are computed as one expectation over the shared rate. The three years scatter as counts of one rate, and the last of them is what the first two predicted. What improves the prediction is the fleet's own log, because it adds the most exposure to the one rate; but no record takes the spread of a one-year count below the Poisson floor, and the value of the next two years is that they settle the rate and with it the spares level, which may move either way. The store is therefore sized from the predictive and resized each year as the count arrives.

11.6.6 Limits of the model

The derivation is exact under the modeling choices named in Section 11.1.5 and in the specialized forms of Section 11.3. That does not make the choices appropriate, and four of them limit the answers; the plug-in has a limit of its own.

A constant rate. The predictive treats the coming year as a fourth year of the same process. The held-out checks of Table 11.10 find nothing against that over the three years, but a slow rise of the rate with the fleet's age would not show in three yearly counts and would make the coming year's count larger than predicted, by an amount that grows with the horizon. Chapter 12 frees the trend and predicts the count under it.

Units alike. One rate for all units makes the halls exchangeable. If hall B ran hotter, or its units were of another series, its rate would be a separate entry of the state, with its own count rows and a prior linking it to the fleet's. The record would then have to be split by hall, which the record of Table 11.1 is not. With $\mathbf{z} = (\omega,\ \delta_B)^{\mathsf{T}}$, hall B's log rate $\omega + \delta_B$, a Gaussian tie row δ_B/σ_δ, and record exposures h_B and h_R of hall B and of the other halls, the 4×2 Jacobian of

the record's rows and the gradient matrix of next year's two conditional means take the forms

$$\mathbf{J} = \begin{bmatrix} \sqrt{b_0 \mathrm{e}^{\omega}} & 0 \\ \sqrt{h_R \mathrm{e}^{\omega}} & 0 \\ \sqrt{h_B \mathrm{e}^{\omega+\delta_B}} & \sqrt{h_B \mathrm{e}^{\omega+\delta_B}} \\ 0 & 1/\sigma_\delta \end{bmatrix} \begin{matrix} \text{prior row} \\ \text{count row, other halls} \\ \text{count row, hall B} \\ \text{tie row} \end{matrix}, \qquad \mathbf{G} = \begin{bmatrix} \tau_B \mathrm{e}^{\omega+\delta_B} & \tau_B \mathrm{e}^{\omega+\delta_B} \\ \tau_R \mathrm{e}^{\omega} & 0 \end{bmatrix},$$

each evaluated at the MAP state. The first column is shared by both halls and the second belongs to hall B alone; $\mathbf{G}\boldsymbol{\Sigma}\mathbf{G}^{\mathsf{T}}$ then has rank two, the conditional means are no longer perfectly correlated, and the correlation between the halls comes from the shared column, through the link that the tie row sets, rather than from identity, and is weaker. The four units being commissioned are a sharper case: new compressors may fail more often in their first months than units that have run for years, and the four new unit-years of the coming year belong to a row of their own if the record of commissioning failures says so.

Independence given the rate. Failures of different units are taken as independent once the rate is known. A cause that acts on several units at once, a power disturbance or a contaminated batch of refrigerant, breaks it and makes counts cluster within a year beyond what a shared rate produces; the variance of the yearly counts then exceeds the negative binomial's. Chapter 15 treats such causes.

One spare per failure. A store reordered after each use is sized by the same equations with τ the exposure over the reorder lead time, a shorter horizon with a smaller epistemic share.

The plug-in at a short horizon. The size of the plug-in's error depends on the horizon and the record through the single number τ/b. For the coming year it is 0.35, and the plug-in is two spares short on the state route and one on the exact posterior; for a horizon of four fleet-years it is 1.40 and the plug-in interval is a third narrower than the predictive one. A plug-in prediction may be acceptable for a short horizon against a long record, and it is never acceptable without the comparison that shows it.

The next chapter lets the fleet's rate drift in calendar time, asks whether it is rising, and predicts when it will reach a limit and how many failures a fleet with a trend will have.

11.7 What to retain

Definition 11.1 (Posterior predictive count)**.** Let the posterior of a constant rate λ be Gamma(a, b). The number of failures K_f in a future exposure τ is negative binomial with shape a and success probability $\pi = b/(b+\tau)$. Its mean is $a\tau/b$ and its variance $a\tau/b + a\tau^2/b^2$: the Poisson part, which no record removes, and the epistemic part $\tau^2 \operatorname{Var}(\lambda \mid \mathbf{y})$, whose share $\tau/(b+\tau)$ grows with the horizon. In the fit the log rate is the one entry of the state, fitted by Newton's method on count rows whose slopes $\mu_k - n_k$ and curvatures μ_k add, with $\hat{\omega} = \ln(a/b)$ and $\sigma_\omega^2 = 1/a$; the two moments are $\mu(\hat{\omega})$ and $\mu(\hat{\omega}) + g^2\sigma_\omega^2$ with $g = \tau \mathrm{e}^{\hat{\omega}}$. Two counts with exposures τ_B, τ_R that share the rate have the covariance matrix with the diagonal Poisson part μ_B, μ_R and the epistemic part of entries $\sigma_\omega^2 g_i g_j$, which has rank one and the off-diagonal entry $\tau_B \tau_R a/b^2$, and the probability of a joint event is one expectation over the posterior. A spares level for service level q is the q-quantile of the predictive. A year's count is checked by its predictive p-value under the predictive built from the other years, a leave-one-out check, and compared with an alternative by the ratio of the two predictives at the observed count. Every answer is computed from $(\hat{\omega}, \sigma_\omega)$, by the delta form or a Gauss–Hermite sum over the Laplace normal, and the negative binomial gives the value beside it, with the gap.

Three conclusions carry beyond the fleet. A future count is predicted by averaging the Poisson distribution over the posterior of the rate, and the plug-in that skips the average has the right mean and too little spread, an error that falls on the upper tail where spares and budgets are decided. Units that share one uncertain rate have correlated counts, visible as the one shared entry on which their conditional means depend, so a statement about several of them together is one expectation over the shared posterior, not a product. And the spread of a future count has a floor set by chance, not by knowledge: the preposterior shows how much of the spread a record can remove in expectation, and for a short horizon against a substantial record the answer is little, while the record a candidate actually brings moves the mean and the spares level in either direction.

Table 11.16 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 11.16: The terms of failure-count prediction and where each is defined.

Term	Expression	Units	Meaning	Defined in
Rate posterior	$\text{Gamma}(a_0 + \sum k_i,\ b_0 + \sum h_i)$	–	pooled record	11.2.2
Count row	$\mathcal{L}_k = \mu_k - n_k \ln \mu_k$; slope $\mu_k - n_k$, curvature μ_k	–	slopes and curvatures add	11.1.5
Log-rate fit	$\hat{\omega} = \ln(a/b)$, $\sigma_\omega^2 = 1/\Lambda = 1/a$	–	Newton on the summed slope	11.2.1
Exact predictive count	$\frac{\Gamma(a+j)}{\Gamma(a)\, j!}\pi^a(1-\pi)^j$, $\pi = b/(b+\tau)$	–	negative binomial; the value beside it	11.3.1
Plug-in count	$\text{Poisson}(\tau a/b)$	–	rate fixed at its mean	11.3.1
Predictive mean, variance	$\mu(\hat{\omega})$; $\mu(\hat{\omega}) + g^2\sigma_\omega^2$, $g = \tau \mathrm{e}^{\hat{\omega}}$; exactly $a\tau/b$; $(a\tau/b)(1+\tau/b)$	failures	overdispersed	11.3.1
Variance split	$\tau \mathbb{E}[\lambda] + \tau^2 \operatorname{Var}(\lambda)$	failures^2	Poisson and epistemic	11.3.1
Epistemic share	$\tau/(b+\tau)$	–	grows with horizon	11.3.1
Spares level	$\min\{s : \Pr(K_f \le s) \ge q\}$	spares	predictive quantile	11.3.2
Expected shortfall	$\sum_{j>s}(j-s)\ \Pr(K_f = j)$	failures	failures without a spare	11.3.2
Law of total covariance	$\mathbb{E}[\operatorname{Cov}(X, Y \mid \lambda)] + \operatorname{Cov}(\mathbb{E}[X \mid \lambda], \mathbb{E}[Y \mid \lambda])$	–	shared-rate coupling	11.3.3
Count covariance	$\text{diag}(\mu_B, \mu_R) + \sigma_\omega^2 g_i g_j$; off-diagonal $\tau_B \tau_R a/b^2$	failures^2	rank-one epistemic part	11.3.3
Joint no failure	$\big(b/(b+\tau_B+\tau_R)\big)^a$	–	exceeds the product	11.3.3
Leave-one-out predictive check	$\Pr(K_3 \ge k_3 \mid \mathbf{y}_{1:2})$	–	predictive p-value, held-out year	11.3.4
Doubled-rate Bayes factor	ratio of sums (11.35); exactly $2^{k_3}\big((b_{1:2}+h_3)/(b_{1:2}+2h_3)\big)^{a_{1:2}+k_3}$	–	ratio of predictives	11.3.4
Laplace predictive	$\sum_m \nu_m \text{Poisson}(j;\ \tau \mathrm{e}^{\hat{\omega}+\sigma_\omega x_m})$	–	Gauss–Hermite; the reported route	11.2.2
Preposterior variance	$\sigma_{\omega,c}^2 = (1/\sigma_\omega^2 + \bar{\lambda}\tau_c)^{-1}$	–	added curvature	11.6.2
Expected predictive variance	$\bar{\lambda}\tau(1 + \tau/(b+\tau_c))$	failures^2	preposterior, in expectation	11.6.2
Poisson floor	$\bar{\lambda}\tau$	failures^2	expected limit of any record	11.6.2

11.8 Exercises

Exercise 11.1. *Formulation.* A fleet of pumps has the failure-rate posterior Gamma$(10, 50)$, with the rate parameter in unit-years, and the coming period has 30 unit-years of exposure. (a) Write the predictive of the count and compute its mean, its variance with the split into the Poisson and epistemic parts, and its standard deviation against the plug-in one. (b) Find the 90 % spares level under the predictive and under the plug-in, and the service level the plug-in level delivers. ((a) K_f is negative binomial with shape 10 and $\pi = 50/80 = 0.625$; mean $10 \times 30/50 = 6.0$, variance $6.0 + 6.0 \times 30/50 = 6.0 + 3.6 = 9.6$, standard deviation 3.10 against $\sqrt{6} = 2.45$, epistemic share $30/80 = 0.375$; from the fitted state, $\hat{\omega} = \ln 0.2$, $\sigma_\omega^2 = 1/10$, $g = 30 \times 0.2 = 6.0$ and $g^2\sigma_\omega^2 = 36/10 = 3.6$, the same split. (b) The predictive cumulative probabilities are 0.869 at 9 and 0.915 at 10, so $s_{0.90} = 10$; the plug-in Poisson(6) reaches 0.916 at 9, so its level is 9, which delivers 0.869 under the predictive.)

Exercise 11.2. *Derivation.* Two halls with exposures τ_1 and τ_2 share a rate with posterior Gamma(a, b), and their counts are independent Poisson given the rate. (a) With the one-number state ω, $\hat{\omega} = \ln(a/b)$ and $\sigma_\omega^2 = 1/a$, write the derivatives g_1 and g_2 of the two conditional means, derive the predictive covariance matrix of the two counts and their correlation, and say what the rank of its epistemic part shows. (b) Evaluate the matrix and the correlation for two halls of 22 unit-years each under Gamma(26, 126), and state what the correlation tends to as b grows. ((a) $g_i = \tau_i a/b$; by the law of total covariance with the delta method, $\mathbf{C}$ has the entries $C_{ij} = \delta_{ij}\, a\tau_i/b + \sigma_\omega^2 g_i g_j$, off-diagonal $\tau_1\tau_2 a/b^2$, diagonal $(a\tau_i/b)(1 + \tau_i/b)$; both derivatives are taken with respect to the one entry, so the epistemic part $\sigma_\omega^2 g_i g_j$ has rank one, and $\operatorname{corr}(K_1, K_2 \mid \mathbf{y}) = C_{12}/\sqrt{C_{11}C_{22}}$ is the geometric mean of the shares, $\sqrt{\tau_1\tau_2/((b + \tau_1)(b + \tau_2))}$. (b) $g_1 = g_2 = 4.540$, $\mathbf{C} = [5.332,\ 0.793;\ 0.793,\ 5.332]$, $\operatorname{corr}(K_1, K_2 \mid \mathbf{y}) = 0.793/5.332 = 22/148 = 0.149$, the common epistemic share; it tends to zero as the record grows, since σ_ω^2 shrinks and the Poisson diagonal remains.)

Exercise 11.3. *Derivation.* Show that the epistemic share of the predictive variance tends to one as the horizon grows, find the horizon at which it is one half, and show from g and σ_ω^2 that the coefficient of variation of the count tends to the standard deviation of the log rate. Evaluate both for the fleet of the chapter. (Share $= g^2\sigma_\omega^2/(\mu(\hat{\omega}) + g^2\sigma_\omega^2) = \tau/(b + \tau) \to 1$ as $\tau \to \infty$, and it equals $\frac{1}{2}$ where $\mu(\hat{\omega}) = g^2\sigma_\omega^2$, at $\tau = b = 126$ unit-years, almost three fleet-years; since $g = \mu(\hat{\omega})$, $\mathrm{CV}^2 = 1/\mu(\hat{\omega}) + \sigma_\omega^2 \to \sigma_\omega^2 = 1/a$, so $\mathrm{CV} \to 1/\sqrt{26} = 0.196$, against 0.385 for the coming year and 0.277 at $\tau = 126$.)

Exercise 11.4. *Derivation.* A small fleet has the posterior Gamma$(4, 20)$, with the rate parameter in unit-years, and the coming period has 10 unit-years. Compute the exact predictive mean, $\Pr(K_f = 0)$ and $\Pr(K_f \geq 6)$; fit the log rate with one count row carrying the whole posterior and give its curvature, σ_ω^2 and the delta-form moments; derive the Laplace mean in closed form; and compare the Laplace tail by quadrature with the exact one. State what the comparison shows against the chapter's fleet. (Exact: $\pi = 2/3$, mean $4 \times 10/20 = 2.0$, $\Pr(K_f = 0) = (2/3)^4 = 0.198$, $\Pr(K_f \geq 6) = 0.042$. State form: the row $(n, t) = (4, 20)$ gives $\hat{\omega} = \ln 0.2$, the curvature $\Lambda = 20 \times 0.2 = 4$, $\sigma_\omega^2 = 1/4$; $g = 2.0$, mean 2.0 and variance $2.0 + 4/4 = 3.0$, the exact moments. Laplace: the mean under the full normal is $2.0\,\mathrm{e}^{1/8} = 2.27$; by quadrature $\Pr(K_f = 0) = 0.164$ and $\Pr(K_f \geq 6) = 0.064$, half again the exact tail, and the 95 % spares level 6 against 5. With four failures the posterior of ω is far from normal and its center $\ln(a/b)$ lies $\ln 4 - \psi(4) = 0.13$ above the exact mean; with 26 it is 0.019, about a seventh of that.)

Exercise 11.5. *Application.* For the fleet of the chapter, a spare compressor held for the year costs R$ 3 000, used or not, in capital and storage, and a failure that finds no spare costs R$ 45 000 in an expedited compressor and rented cooling. (a) Write the expected cost of a spares

level s and derive the condition for the level that minimizes it. (b) Find that level under the predictive and under the plug-in, and the expected cost of each level under the predictive. ((a) $C(s) = 3\,000\,s + 45\,000\,\mathbb{E}[(K_f - s)^+]$; raising s by one adds 3 000 and removes 45 000 $\Pr(K_f > s)$, so the best level is the smallest s with $\Pr(K_f > s) \le 3/45$, that is $\Pr(K_f \le s) \ge 0.933$. (b) Predictive: $\Pr(K_f \le 14) = 0.929$, $\Pr(K_f \le 15) = 0.955$, so $s = 15$; plug-in: 0.922 at 13 and 0.956 at 14, so $s = 14$. Under the predictive $C(15) = 45\,000 + 45\,000 \times 0.1095 = \text{R\$}\,49\,928$ and $C(14) = 42\,000 + 45\,000 \times 0.1804 = \text{R\$}\,50\,118$: the missing spare costs R$ 190 a year in expectation, while the plug-in, which reports $C(14) = \text{R\$}\,46\,057$, understates the expected cost of its own plan by R$ 4 061.)

Exercise 11.6. *Application.* In year 4 the fleet of the chapter, with 44 unit-years and the posterior Gamma(26, 126) from years 1 to 3, records seventeen failures. (a) Compute the predictive tail of seventeen. (b) Compute the log Bayes factor of the alternative that the rate rose by half in year 4 against the constant rate, and the posterior probability of the rise if it was given a prior probability of 0.1. (c) State what each number answers. ((a) $\Pr(K_4 \ge 17 \mid \mathbf{y}) = 1 - 0.9726 = 0.027$. (b) By equation (11.36) with the factor 1.5 in place of 2, $\ln B_{10} = 17 \ln 1.5 - 43 \ln(192/170) = 6.893 - 5.233 = +1.66$, $B_{10} = 5.3$; posterior odds $5.3 \times 0.1/0.9 = 0.58$, a probability of 0.37. (c) The tail, a predictive p-value, says that seventeen is unusual under a constant rate, a predictive probability of 0.027 of so many, not a frequency of such years; the Bayes factor says that a rise by half predicted it about five times better; with a prior of 0.1 the rise is still less likely than not, and a fifth year, or the trend model of the next chapter, is needed before the spares plan is changed.)

Chapter 12

Aging and Time-Varying Failure Rates

12.1 A failure rate that changes with time

The forty computer-room air-conditioning units of an older data hall have been in service for six years, and their failures have climbed every year: five in the first year, then six, seven, nine, eleven and fourteen. The maintenance contract provides spares and crew for up to 0.5 failures per unit-year, twenty failures a year for the fleet, and the operator must decide when to renegotiate it or begin replacing units. The first reading of the record is a straight line. Plotted on a logarithmic chart the six yearly rates rise by about 21 % a year, and the line drawn through them reaches the limit about two years after the middle of the last year of record: the contract holds for two more years, and the renegotiation can wait until next year.

The line answers a question the record cannot settle on its face. Six counts of five to fourteen scatter by their square roots, so a part of the rise could be chance, and the slope of the line, which sets the date, is the least certain thing the record holds: a shallower slope puts the crossing years later, a steeper one puts it next year, and a slope near zero puts it never. The question of this chapter is

how a trend in the failure rate of a fleet is estimated from the counts of successive years, whether the record supports a trend against a constant rate, what the rate will be in the coming years and with what band, when it will cross a stated limit, and what the trend means for next year's count.

What changes from the constant rate of Chapter 6 and the fleet of Chapter 11 is the state, which now carries the log-rate and its trend per year, two entries on which every count row bears together, and the quantity the operator cares about most, the crossing time, which divides by the uncertain trend.

Jargon: calendar time, age, fleet failure rate, log-linear trend, rate ratio, doubling time, centering, score equations, lever arm, crossing time, pole, Fieller's construction, Pearson residual, covariate

Calendar time is the time on the plant's clock, here in years counted so that the record's first year is centered at $t = 1$; the *age* of a unit is the time since it entered service. The *fleet failure rate* $\lambda(t)$ is the expected number of failures per unit and per year at calendar time t, the rate of occurrence of failures of a population of repairable units. A *log-linear trend* is a rate whose logarithm is a straight line in time, $\ln \lambda(t) = \alpha + \gamma(t - \bar{t})$; its *rate ratio* per year is e^{γ}, the factor by which the rate grows in one year, and its *doubling time* is $\ln 2/\gamma$. *Centering* measures the time from a reference inside the record, so that the two entries of the line are nearly uncorrelated. The *score equations* are the equations obtained by setting the gradient of the log-likelihood to zero. The *lever arm* of a forecast is its distance in time from the record's center of information, the time at which the rate is best known. The *crossing time* is the time at which the rate reaches a stated limit. A *pole* of a function is a point at which it becomes infinite, here the zero of a denominator. *Fieller's construction* is the standard method that finds the band of a ratio of two normal quantities by inverting a linear event instead of linearizing the ratio. A *Pearson residual* is an observed count minus its fitted mean, divided by the square root of that mean. A *covariate* is an observed condition, such as the load, that enters the rate through a coefficient of its own.

Key idea

A trend in a fleet's failure rate is two entries of the state, the log-rate at the center of the record and the trend per year, fitted to one Poisson count row per window. The rate at any time is linear in those entries on the logarithmic scale, and its band follows exactly from the covariance. The time at which the rate crosses a limit is a ratio with the uncertain trend in its denominator: its distribution is skewed, has no mean, and puts mass on never crossing, so its band comes from the inversion of the threshold probability of the rate.

12.1.1 Calendar time, age and the rate of a fleet

The companion volume *Reliability Engineering* classifies hazard functions by their shape, falling in early life, flat in useful life and rising in wear-out, and describes the physical mechanisms of degradation behind the rising part; those descriptions are taken as known here. Chapter 7 estimated such a shape for one kind of unit from the ages at which units failed or were last seen running. That is a statement about a unit as it ages. The record of a fleet of repairable units is different. Its units are repaired and return to service, its failures are counted per year, and the quantity the record informs is the rate at which the fleet as a whole produces failures in calendar time: the expected number of failures per unit and per year, written $\lambda(t)$ with t in calendar years. For a fleet whose units entered service together it moves with their common age, and it moves also when repairs restore less than a new unit, when operating conditions change, or when the maintenance practice changes.

The nonhomogeneous Poisson process of Chapter 6 (Section 6.3.4) is the model of such a record: failures occur independently in time, with an intensity that changes. That chapter introduced the power-law intensity in operating time, whose logarithm is linear in the logarithm of time. This chapter uses the other standard choice, an intensity whose logarithm is linear in time itself. Both describe a rate that rises or falls monotonically; they differ mostly in how they extrapolate, a point Section 12.6.6 returns to. The record here is grouped into yearly windows, so the event times within a window are not available and the evidence is one count per window.

12.1.2 The physical system and its graph (step 0)

The fleet is an older sister of the fleet of Chapter 11: forty computer-room air-conditioning units in the halls of a second data center, of one design, installed within a year of each other and maintained by one contractor. When a unit fails it is repaired and returned to service. The limit of the opening, the contract's $\lambda_{\max} = 0.5$ failures per unit-year, is the level the fleet's rate is compared with. Failures occur as a nonhomogeneous Poisson process at the fleet's rate $\lambda(t)$, and the exposure of each window, the unit-years in service, is read from the plant log.

Figure 12.1 draws the fleet as a graph. The green box at the top is the free node, the fleet, carrying the two entries of the state. The six gray boxes in the row below are the six yearly windows of the record, each with its exposure pinned; the dashed arrows from the fleet to the windows are the six count rows, each comparing a window's recorded failures with the rate the fleet's state implies for it. The white box on the right of the row, with a dashed border, is the next window, whose count is not yet known; the red box above it is the question's extension of the graph: the limit of 0.5 per unit-year, the rate at a stated time, and the crossing time, all derived from the fleet's state.

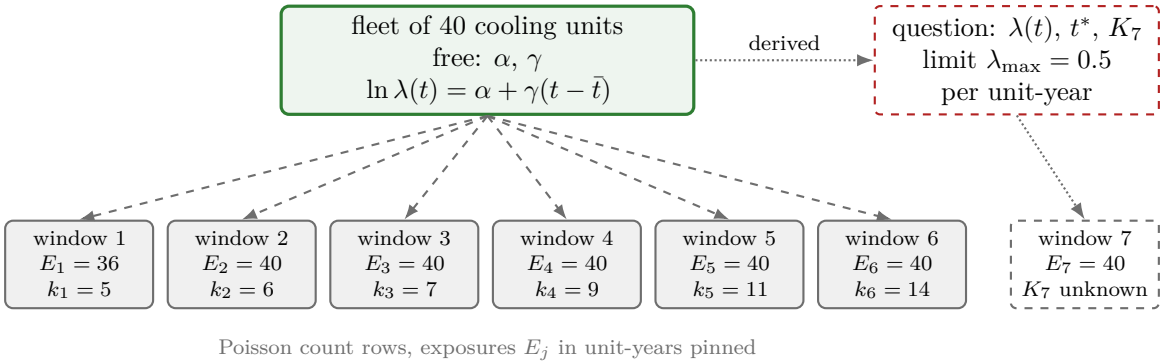


Figure 12.1: Plant graph of the older cooling fleet. Free: the fleet node (green), carrying the log-rate α at the record's mid-time $\bar{t} = 3.5$ and the trend γ per year. Pinned: the exposures E_j of the six yearly windows, in unit-years (gray). Observed: the failure counts k_j of the windows, one Poisson count row each (dashed arrows). The question's extension (red, dashed border): the limit $\lambda_{\max} = 0.5$ per unit-year, the rate $\lambda(t)$ at a stated time, the crossing time t^* and the count K_7 of the next window, all derived from the state.

The state entries. The rate must be positive and is expected to change by a roughly constant factor per year, so its logarithm is written as a straight line in calendar time,

$$\ln \lambda(t) = \alpha + \gamma\,(t - \bar{t}), \qquad \mathbf{z} = \begin{bmatrix} \alpha \\ \gamma \end{bmatrix}, \qquad \bar{t} = 3.5, \tag{12.1}$$

with t in years, counted so that window j occupies the year $t \in [j - \frac{1}{2}, j + \frac{1}{2}]$ and is centered at $t_j = j$; the record begins at $t = \frac{1}{2}$ and its last window is centered at $t = 6$, which the chapter calls now. The entry α is the log-rate at the record's mid-time $\bar{t}$, the mean of the six window centers, with λ in failures per unit-year; the entry γ is the trend, in yr^{-1}, the change of the log-rate per year. The book's notation carries a single log-rate as $\omega = \ln(\lambda/\lambda_{\mathrm{ref}})$; with a trend there are two entries, and this chapter names them α and γ and drops the reference rate, which would shift α by a constant and change nothing else. Section 12.3.1 derives why the time is measured from $\bar{t}$ rather than from the origin of the scale.

12.1.3 The record

Table 12.1 lists the record. Each row is one yearly window; the columns give its exposure in unit-years, its failures, the observed rate, which is their ratio, and the window's centered time $x_j = t_j - \bar{t}$. The first window has 36 unit-years and every later window the full forty. Chapter 11 wrote a year's exposure h_i and a future exposure τ; this chapter writes E_j for both, the windows of record and the next window E_7, because under a trend every window, past or future, enters through its own time t_j and the two belong to one sequence. The counts climb from 5 to 14, and the observed rate from 0.139 to 0.350 per unit-year. In total the record holds 52 failures in 236 unit-years, a pooled rate of 0.220 per unit-year. The record is illustrative. The evidence $\mathbf{y}$ is the six window counts with their exposures.

Table 12.1: The record of the older cooling fleet: six yearly windows with their exposure, failures, observed rate and centered time. The pooled rate is $52/236 = 0.2203$ per unit-year.

Window j	Exposure E_j (unit-yr)	Failures k_j	Observed rate k_j/E_j (per unit-yr)	$x_j = t_j - \bar{t}$ (yr)
1	36	5	0.139	−2.5
2	40	6	0.150	−1.5
3	40	7	0.175	−0.5
4	40	9	0.225	0.5
5	40	11	0.275	1.5
6	40	14	0.350	2.5
Total	236	52	0.220	

12.1.4 The probabilistic questions (step 1)

The owner of the fleet asks whether the rate is rising and when it crosses the limit, and the operator breaks that into six questions. Step 1 writes each as a probability, an expectation or a ratio of evidences over the state, with the unknown on the left of the bar and the record on the right. Table 12.2 lists them, with the owner's question in its own words on the first line: the first column is the question in engineering language; the second the general form that every question of its category takes, written in the book's notation for a quantity $f(\mathbf{z})$ derived from the state, a threshold f^* on it, the posterior density $p(\mathbf{z} \mid \mathbf{y})$, a future observation y_{new} and the evidence $p(\mathbf{y} \mid M)$ of a drawing M; the third its mathematical form for the fleet; the fourth its category; and the last the subsection where it is answered. The first and the fourth are estimation, of the rate at a stated time and of the crossing time, both derived quantities of the state. The second and the third are predictions under a threshold, tails of a quantity linear in the state. The fifth is a hypothesis comparison, a trend against a constant rate. The sixth is a predictive count under the two drawings, which Chapter 2 files under prediction under a threshold. A seventh question, whether one more year of record would settle the crossing time, asks nothing new of the present record: it is the next measurement for these six, answered in Section 12.6. The atlas of Appendix A lists the same forms.

12.1.5 The evidence as rows (step 2)

Step 2 attaches six rows to the graph, one Poisson count row per window, each depending on both entries of the state through equation (12.1): a window's expected count moves with the level α and, through its own time, with the trend γ. None of them is Gaussian: each is a count row of Section 4.2.7, entering the objective as its negative log-probability term and the Jacobian through its working pair, below.

Table 12.2: The questions asked of the fleet's record, the general form of their category, their mathematical form for the fleet, their category and where each is answered. $\mathbf{y}$ is the six window counts with their exposures; $\lambda_{\max} = 0.5$ per unit-year; K_7 is the count of window 7 with $E_7 = 40$ unit-years; M_1 is the trend model and M_0 the constant rate; B_{10} is the Bayes factor of M_1 against M_0.

Engineering question	General form	Mathematical form	Category	Answered in
The owner's question: is the rate rising, and when does it cross the limit?				
What is the rate now and next year, with its band?	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$, quantiles	$p(\lambda(t) \mid \mathbf{y})$ at $t = 6$ and 7	estimation of a derived quantity	12.3.1
Is the failure rate rising?	$\Pr(f(\mathbf{z}) \ge f^* \mid \mathbf{y})$	$\Pr(\gamma > 0 \mid \mathbf{y})$	prediction under a threshold	12.3.2
Will the rate be above the limit two years from now?	$\Pr(f(\mathbf{z}) \ge f^* \mid \mathbf{y})$	$\Pr(\lambda(8) > 0.5 \mid \mathbf{y})$	prediction under a threshold	12.3.2
When does the rate cross 0.5 per unit-year?	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$, quantiles	$\Pr(t^* \le s \mid \mathbf{y})$, $t^* = \bar{t} + (\ln \lambda_{\max} - \alpha)/\gamma$	estimation of a derived quantity	12.3.3
Does the record need a trend, against a constant rate?	$p(\mathbf{y} \mid M_1)/p(\mathbf{y} \mid M_0)$	B_{10}, trend against constant rate	hypothesis comparison	12.3.4
How many failures next year, under the trend and under a constant rate?	$\int p(y_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	$\Pr(K_7 = j \mid \mathbf{y}, M_1)$, $\Pr(K_7 = j \mid \mathbf{y}, M_0)$	prediction under a threshold	12.3.5

The modeling choices. Six modeling choices make the model of the fleet. All forty units are taken to share one failure rate $\lambda(t)$ at each calendar time: this is what lets the record be summarized by six counts acting on two entries, and it fails if some units are worse than others, which makes the counts scatter more than Poisson and the bands too narrow, as the residuals of Section 12.4.4 would show. Failures are taken as independent across units and windows: this makes each window's count Poisson and the six rows' terms add, and it fails if one cause, a bad batch of parts or a cooling event, strikes several units at once (Section 12.6.6). A repaired unit is taken to return to the shared rate: this makes a window's count depend on its exposure alone, not on the history of its units, and it fails if repairs restore less than they should and repaired units fail again soon (Section 12.6.6). The log-rate is taken as a straight line in calendar time, equation (12.1): this makes the rate at any time a function of one linear combination of the two entries, whose band the Laplace posterior gives exactly, and it fails if the rise is not by a constant factor per year, a rate that rises ever more slowly in proportion, as the power law does, or a step at a change of contractor, which the ordered residuals of Section 12.4.4 test and Section 12.6.6 discusses. The line is taken to continue unchanged beyond the record: this is what makes next year's rate a forecast rather than a guess, and it is broken by anything that changes the fleet's course, units replaced or retired, a new contractor, a heavier duty (Section 12.6.6). The posterior of the state is carried by its Laplace Gaussian $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$: this makes the band of every quantity linear in the state an exact normal band, and it misleads when the counts are few and the exact posterior is skewed, which the grid posterior of Section 12.4.1 measures.

The priors on α and γ are flat, so no prior row enters the fit: the record of 52 failures is large enough to determine both entries without one, and the answers are then those of the record alone.

Two questions bring rows or drawings of their own. The comparison of the fifth question needs a proper prior on the trend, which enters the trend model M_1 as one Gaussian prior row,

γ/s_γ, for that answer alone. Its spread s_γ states, before the record is read, how large a trend is plausible for this equipment. A fleet of cooling units may age by a few tens of percent per year, and hardly by a factor of two in one year; $s_\gamma = 0.3\,\mathrm{yr}^{-1}$ puts 95 % of the prior on rate ratios between $\mathrm{e}^{-0.59} = 0.56$ and $\mathrm{e}^{0.59} = 1.80$ per year, which says that, and it is the spread the chapter reports; Section 12.4.3 varies it. The constant rate M_0 has the one-number state α with the same six count rows, each acting on α alone, and the same flat prior on α, for the comparison and for its prediction of next year's count alike. Which row acts on which entry is a property of the graph, and it fixes where the Jacobian of the trend model can be nonzero, marked $*$:

$$\mathbf{J} = \begin{bmatrix} * & * \\ * & * \\ * & * \\ * & * \\ * & * \\ * & * \\ 0 & * \end{bmatrix} \begin{matrix} \text{window 1} \\ \text{window 2} \\ \text{window 3} \\ \text{window 4} \\ \text{window 5} \\ \text{window 6} \\ \text{prior row on } \gamma\text{, comparison only} \end{matrix} , \qquad \text{columns } \alpha, \gamma. \tag{12.2}$$

Every count row acts on both entries, so the two columns of $\mathbf{J}$ overlap in every row, and the information they give has an off-diagonal entry: the level and the trend are fitted together, and their covariance is not zero unless the time origin is chosen to make it so (Section 12.3.1). Only the prior row of the comparison acts on one entry. No row is a physics row, so no model tolerance σ_p arises: the limit, the rate at a stated time and the crossing time are functions of the state, not entries tied to it by a constraint, and the next window's count is a prediction from it.

The count of a window. Under the nonhomogeneous Poisson process of Chapter 6 the number of failures of one unit in an interval is Poisson with mean equal to the integral of the rate over the interval, and counts of different units and intervals are independent. Window j holds E_j unit-years spread evenly over its year, so its count K_j is Poisson with mean

$$\mathbb{E}[K_j \mid \mathbf{z}] = E_j \int_{j-1/2}^{j+1/2} \lambda(t)\,\mathrm{d}t = E_j\,\mathrm{e}^{\alpha} \int_{j-1/2}^{j+1/2} \mathrm{e}^{\gamma(t-\bar{t})}\,\mathrm{d}t,$$

the second form by equation (12.1), since each unit-year of exposure is one unit for one year. The integral of the exponential is elementary:

$$\int_{j-1/2}^{j+1/2} \mathrm{e}^{\gamma(t-\bar{t})}\,\mathrm{d}t = \frac{\mathrm{e}^{\gamma(j+1/2-\bar{t})} - \mathrm{e}^{\gamma(j-1/2-\bar{t})}}{\gamma} = \mathrm{e}^{\gamma(t_j-\bar{t})}\,\frac{\sinh(\gamma/2)}{\gamma/2}.$$

The factor $\sinh(\gamma/2)/(\gamma/2)$ is the ratio of the window's average rate to its rate at the center. It depends on γ but not on j, so it multiplies the means of all six windows equally and is absorbed into α, shifting it by $\ln[\sinh(\gamma/2)/(\gamma/2)]$ and leaving γ unchanged; Section 12.4.4 gives its size at the fitted trend. The chapter therefore writes the count row with the rate at the window's center,

$$K_j \mid \mathbf{z} \sim \mathrm{Poisson}(\lambda_j E_j), \qquad \lambda_j = \lambda(t_j) = \mathrm{e}^{\alpha+\gamma x_j}, \qquad x_j = t_j - \bar{t}. \tag{12.3}$$

The rate ratio and the doubling time. Over one year the log-rate changes by γ, so the rate changes by the factor

$$\frac{\lambda(t+1)}{\lambda(t)} = \mathrm{e}^{\gamma}, \qquad t_2 = \frac{\ln 2}{\gamma}, \tag{12.4}$$

the rate ratio per year and the doubling time t_2, the time over which the rate doubles when $\gamma > 0$. A trend of $\gamma = 0.1\,\mathrm{yr}^{-1}$ is a rise of 10.5 % per year and a doubling in 6.9 years.

The rows as negative log-probability terms. The six count rows are not Gaussian; each is the count row of Section 4.2.7, which puts in the objective the negative logarithm of the probability of its count given the state and enters the normal equations through its working residual. A count row with k_j observed failures and expected count μ_j has the term $\mu_j - k_j \ln \mu_j$ up to terms free of the state, with the gradient and the Hessian of equation (4.42), so that row j puts in the objective the term

$$\begin{aligned} \mathcal{L}_j(\mathbf{z}) &= \mu_j - k_j \ln \mu_j, \qquad \mu_j = \lambda_j E_j = E_j \, \mathrm{e}^{\alpha + \gamma x_j}, \\ \mathcal{L}(\mathbf{z}) &= \sum_{j=1}^{6} \mathcal{L}_j(\mathbf{z}) = \sum_{j=1}^{6} \Big[E_j \, \mathrm{e}^{\alpha + \gamma x_j} - k_j(\alpha + \gamma x_j) \Big], \end{aligned} \tag{12.5}$$

with μ_j the count window j expects at the state and the terms $k_j \ln E_j$, free of the state, dropped from the last form.

The working residuals and the Jacobian. The log of each expected count, $\ln \mu_j = \ln E_j + \alpha + \gamma x_j$, is linear in the state, and its derivatives form the constant 6×2 matrix

$$\mathbf{X} = \frac{\partial \ln \boldsymbol{\mu}}{\partial \mathbf{z}} = \begin{bmatrix} 1 & -2.5 \\ 1 & -1.5 \\ 1 & -0.5 \\ 1 & 0.5 \\ 1 & 1.5 \\ 1 & 2.5 \end{bmatrix} \begin{matrix} \text{window 1} \\ \text{window 2} \\ \text{window 3} \\ \text{window 4'} \\ \text{window 5} \\ \text{window 6} \end{matrix}$$

the pattern (12.2) filled. Its first column, a one for every window, says that every window's expected count moves one for one with the level α; its second, the centered times x_j of Table 12.1, says that the trend raises the late windows and lowers the early ones. The working pair (4.43) of each count row is here

$$\begin{aligned} r_j(\mathbf{z}) &= \frac{\mu_j - k_j}{\sqrt{\mu_j}}, \qquad \mathbf{J}_j = \sqrt{\mu_j} \, \frac{\partial \ln \mu_j}{\partial \mathbf{z}^\mathsf{T}} = \sqrt{\mu_j} \, \begin{bmatrix} 1 & x_j \end{bmatrix}, \\ \mathbf{J} &= \mathbf{W}^{1/2} \mathbf{X}, \qquad \mathbf{W} = \operatorname{diag}(\mu_1, \ldots, \mu_6), \end{aligned} \tag{12.6}$$

so that the Jacobian $\mathbf{J}$ is 6×2, row j being window j's derivatives scaled by the square root of its expected count. The weights μ_j are the precisions of the log counts, $\operatorname{Var}(\ln k_j) \approx 1/\mu_j$, the counterpart of the notation chapter's $w_k = 1/\sigma_k^2$, so that $\mathbf{X}$ is the Jacobian of unnormalized rows in the form $\mathbf{J}^\mathsf{T}\mathbf{W}\mathbf{J}$ of that chapter, and $\mathbf{J} = \mathbf{W}^{1/2}\mathbf{X}$ is its normalized form. As Section 4.2.7 sets out, $\mathbf{J}_j^\mathsf{T} r_j = (\mu_j - k_j) \, \partial \ln \mu_j / \partial \mathbf{z}$ is the gradient of $\mathcal{L}_j$ and $\mathbf{J}_j^\mathsf{T}\mathbf{J}_j$ its curvature, while $\mathbf{J}_j$ is not the derivative of r_j; the gradient and the Hessian of the objective are

$$\nabla \mathcal{L}(\mathbf{z}) = \mathbf{J}^\mathsf{T} \mathbf{r} = \begin{bmatrix} \sum_j (\mu_j - k_j) \\ \sum_j x_j (\mu_j - k_j) \end{bmatrix}, \tag{12.7}$$

$$\boldsymbol{\Lambda}(\mathbf{z}) = \frac{\partial^2 \mathcal{L}}{\partial \mathbf{z} \, \partial \mathbf{z}^\mathsf{T}} = \mathbf{J}^\mathsf{T}\mathbf{J} = \mathbf{X}^\mathsf{T}\mathbf{W}\mathbf{X} = \begin{bmatrix} \sum_j \mu_j & \sum_j x_j \mu_j \\ \sum_j x_j \mu_j & \sum_j x_j^2 \mu_j \end{bmatrix}, \tag{12.8}$$

with $\mathbf{r} = (r_1, \ldots, r_6)^\mathsf{T}$. The Hessian is exact, not an approximation. The curvature term $\mathbf{C}$ of equation (4.44) is, for a count row, the observed-minus-expected difference $(\mu_j - k_j) \, \partial^2 \ln \mu_j / \partial \mathbf{z}^2$ of equation (4.42); it is zero because the log-rate is the canonical link of the Poisson count, so the log count is linear in the state and the count k_j does not appear in the curvature. The observed information, the Hessian at the recorded counts, equals the expected information, its average over the counts, and Newton's method and the Gauss–Newton step are the same

iteration here. Each row's contribution to the information is its expected count, the Poisson information in the logarithm of its mean, which Chapter 6 found as the curvature $T_{\text{tot}}\lambda$ of a single count row (equation (6.27)). With every row normalized in this way the information matrix is $\mathbf{J}^{\mathsf{T}}\mathbf{J}$; with the unnormalized log-count derivatives $\mathbf{X}$ and the weights $\mathbf{W}$ it is $\mathbf{X}^{\mathsf{T}}\mathbf{W}\mathbf{X}$, the same matrix.

The score equations. The MAP state sets the gradient (12.7) to zero, $\mathbf{J}^{\mathsf{T}}\mathbf{r} = \mathbf{0}$, one equation per column. Written as $\mathbf{X}^{\mathsf{T}}\hat{\boldsymbol{\mu}} = \mathbf{X}^{\mathsf{T}}\mathbf{k}$, with $\hat{\boldsymbol{\mu}}$ the fitted expected counts and $\mathbf{k}$ the recorded ones, they are

$$\underbrace{\sum_{j=1}^{6}\hat{\mu}_j = \sum_{j=1}^{6}k_j = 52,}_{\text{column 1}} \qquad \underbrace{\sum_{j=1}^{6}x_j\,\hat{\mu}_j = \sum_{j=1}^{6}x_j k_j = 31,}_{\text{column 2}} \tag{12.9}$$

with $\hat{\mu}_j = \hat{\lambda}_j E_j$, $\hat{\lambda}_j$ the fitted rates, and $\mathbf{X}^{\mathsf{T}}\mathbf{k} = (52,\ 31)^{\mathsf{T}}$ computed from the record, its second entry being $5(-2.5) + 6(-1.5) + 7(-0.5) + 9(0.5) + 11(1.5) + 14(2.5) = 31$. The level column makes the fit reproduce the total count, and the time column makes it reproduce the count-weighted sum of the centered times. The ratio of the two entries of $\mathbf{X}^{\mathsf{T}}\mathbf{k}$ is $t_w - \bar{t}$ of equation (12.18): the fit puts the mean time of its expected failures at the mean time of the recorded ones, later than the mid-time $\bar{t}$ when the later windows hold more failures. The objective (12.5) depends on the counts only through $\mathbf{X}^{\mathsf{T}}\mathbf{k}$, which is therefore the sufficient statistic for the two entries once the exposures are fixed: the record informs the level of the rate through the level column and the trend through the time column, that is through the timing of its failures.

12.2 The posterior of the state (step 3)

All six questions are read from the posterior of the state that the six count rows of step 2 give, approximately the Laplace normal and exactly on a grid of the two entries; the comparison and next year's count add the prior row and the constant-rate drawing that step 2 names. This section derives the posterior and computes it for the fleet; Section 12.3 reads every answer from it.

Four objects of the fit recur in the specialized forms. The *MAP state* $\hat{\mathbf{z}} = (\hat{\alpha}, \hat{\gamma})^{\mathsf{T}}$ is the minimum of the objective that the six count rows put on the state; the *information matrix* $\boldsymbol{\Lambda} = \mathbf{J}^{\mathsf{T}}\mathbf{J}$ is the curvature of that objective at the minimum, with $\mathbf{J}$ the 6×2 Jacobian of the normalized rows, of the pattern (12.2); the *Laplace covariance* $\boldsymbol{\Sigma} = \boldsymbol{\Lambda}^{-1}$, a 2×2 matrix with an off-diagonal entry, makes the posterior approximately normal, $\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$; and a derived quantity f is carried by its *gradient* $\mathbf{g} = \partial f/\partial \mathbf{z}$ at $\hat{\mathbf{z}}$, with the variance $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$. The two-entry problem admits no conjugate closed form, but it admits two exact routes, the posterior evaluated on a grid of the two entries and the exact marginal of the trend (Section 12.2.2).

12.2.1 The posterior from the six rows

The state is the two-entry vector of equation (12.1), and the six rows act on it through the residual vector and the 6×2 Jacobian of equation (12.6). Each Newton step solves the normal equations of Chapter 3, equation (3.24), with the working rows,

$$(\mathbf{J}^{\mathsf{T}}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^{\mathsf{T}}\mathbf{r}, \qquad \mathbf{z}^{(m+1)} = \mathbf{z}^{(m)} + \Delta\mathbf{z}, \tag{12.10}$$

with $\mathbf{J}$ and $\mathbf{r}$ evaluated at $\mathbf{z}^{(m)}$. The iteration starts from the constant-rate fit, $\mathbf{z}^{(0)} = (\ln(\sum_j k_j / \sum_j E_j),\ 0)^{\mathsf{T}}$, where every window has the pooled rate; there the first entry of the gradient is zero, because the constant rate reproduces the total, the first score equation,

and the second is not, because a constant rate places the expected failures evenly while the recorded ones may lean early or late.

The MAP state. The iteration stops where $\mathbf{J}^\mathsf{T}\mathbf{r} = \mathbf{0}$, the score equations (12.9). Unlike the gradient, the residual vector itself does not vanish at $\hat{\mathbf{z}}$: the six working residuals are the Pearson residuals $(k_j - \hat{\mu}_j)/\sqrt{\hat{\mu}_j}$ with their signs reversed, orthogonal to both columns of $\mathbf{J}$, and they are what the trend leaves unexplained (Section 12.4.4).

The information and the covariance. At the MAP state the information of equation (12.8) has two entries known from the record before any fitting, $[\mathbf{\Lambda}]_{11} = \mathbf{1}^\mathsf{T}\mathbf{W}\mathbf{1} = \sum_j k_j$ and $[\mathbf{\Lambda}]_{12} = \mathbf{1}^\mathsf{T}\mathbf{W}\mathbf{x} = \sum_j x_j k_j$, the left-hand sides of the score equations; the third, $[\mathbf{\Lambda}]_{22} = \mathbf{x}^\mathsf{T}\mathbf{W}\mathbf{x} = \sum_j x_j^2 \hat{\mu}_j$, needs the fitted counts against the squared times. The Laplace posterior of Chapter 4 is $\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$ with $\mathbf{\Sigma} = \mathbf{\Lambda}^{-1}$, equation (3.25), and each entry is read off with its standard deviation, the two with their correlation,

$$\sqrt{\Sigma_{\alpha\alpha}}, \qquad \sqrt{\Sigma_{\gamma\gamma}}, \qquad \operatorname{corr}(\alpha, \gamma) = \frac{\Sigma_{\alpha\gamma}}{\sqrt{\Sigma_{\alpha\alpha}\Sigma_{\gamma\gamma}}}.$$

The sign of the correlation is read from the columns of $\mathbf{J}$. Their inner product $\mathbf{1}^\mathsf{T}\mathbf{W}\mathbf{x} = \sum_j x_j k_j$ is positive when the expected failures lean toward the late windows, where $x_j > 0$, and a positive off-diagonal entry of a 2×2 information matrix is a negative off-diagonal entry of its inverse. The sign has a meaning on the line itself. A line pivots about its center of information t_w of equation (12.18), and when the later windows hold more failures the mid-time $\bar{t}$ lies to its left: a steeper line through the same pivot is lower at $\bar{t}$, so larger γ goes with smaller α, and a shallower line, as a prior row pulling the trend toward zero makes it, is higher there.

12.2.2 The exact posterior on the two entries

The Laplace Gaussian is the remaining approximation of every reported answer, and the two-entry state admits two routes that remove it.

The grid posterior. With two entries the posterior can be computed without the Gaussian, by evaluating the likelihood on a grid of the state and normalizing, the exact route that Chapter 15 uses for a posterior with no closed form. The procedure is: lay a grid of 1201 values of α over $\hat{\alpha} \pm 9\sqrt{\Sigma_{\alpha\alpha}}$ and 1601 values of γ over $\hat{\gamma} \pm 10\sqrt{\Sigma_{\gamma\gamma}}$; at each point evaluate the log-likelihood $\sum_j [k_j(\alpha + \gamma x_j) - E_j \mathrm{e}^{\alpha + \gamma x_j}]$, which with flat priors is the log posterior up to a constant; subtract its maximum, exponentiate and divide by the sum, so that the grid weights sum to one; and read every answer as a weighted sum over the grid, a probability as the total weight of the points in the event and a quantile as the point at which the cumulative weight, in the order of the quantity, passes the level.

The exact marginal of the trend. The evidence ratio has a second exact route. The quadratic expansion of the count rows is the only approximation in equation (12.31), and it can be removed, because α integrates out of the likelihood in closed form. Write $m(\gamma)$ for the likelihood of γ with α integrated out under its flat prior; the letter is used, as in Chapter 7, for such a likelihood, and it is not the cumulative intensity $m(t)$ of Chapter 6. With $u = \mathrm{e}^{\alpha}$, $\mathrm{d}\alpha = \mathrm{d}u/u$, $K = \sum_j k_j$ and $S(\gamma) = \sum_j E_j \mathrm{e}^{\gamma x_j}$,

$$\int \exp\Big(\sum_j [k_j(\alpha + \gamma x_j) - E_j \mathrm{e}^{\alpha + \gamma x_j}]\Big) \mathrm{d}\alpha = \mathrm{e}^{\gamma \sum_j k_j x_j} \int_0^\infty u^{K-1} \mathrm{e}^{-uS(\gamma)} \, \mathrm{d}u = \frac{\Gamma(K)\, \mathrm{e}^{31\gamma}}{S(\gamma)^{52}},$$

by the gamma integral (6.16) and the sums of the score equations (12.9), so that exactly $m(\gamma) \propto \mathrm{e}^{31\gamma}/S(\gamma)^{52}$. The evidence of M_1 is the average $\int m(\gamma)\pi(\gamma)\,\mathrm{d}\gamma$ and that of M_0, the same likelihood at no trend, is $m(0)$, with the same constant in both; their ratio is the exact Bayes factor, and the normalized $m(\gamma)$ under a flat prior is the exact marginal posterior of the trend.

12.2.3 The posterior of the fleet

The fleet's inputs are the record of Table 12.1, six windows with 52 failures in 236 unit-years, and the grid of 1201×1601 points of Section 12.2.2.

The Newton iteration. The iteration (12.10) starts from the constant-rate fit, $\mathbf{z}^{(0)} = (\ln(52/236),\ 0)^{\mathsf{T}} = (-1.5126,\ 0)^{\mathsf{T}}$. There every window has the rate 0.2203 per unit-year, the expected counts are $\boldsymbol{\mu} = (7.932,\ 8.814,\ 8.814,\ 8.814,\ 8.814,\ 8.814)^{\mathsf{T}}$, and

$$\mathbf{r}(\mathbf{z}^{(0)}) = \begin{bmatrix} 1.0411 \\ 0.9477 \\ 0.6109 \\ -0.0628 \\ -0.7365 \\ -1.7470 \end{bmatrix}, \qquad \mathbf{J}(\mathbf{z}^{(0)}) = \begin{bmatrix} 2.8164 & -7.0410 \\ 2.9688 & -4.4531 \\ 2.9688 & -1.4844 \\ 2.9688 & 1.4844 \\ 2.9688 & 4.4531 \\ 2.9688 & 7.4219 \end{bmatrix} :$$

the constant rate expects too many failures early and too few late. The gradient is $\mathbf{J}^{\mathsf{T}}\mathbf{r} = (0,\ -28.80)^{\mathsf{T}}$, its second entry the lean of the recorded failures toward the late windows, and the first step solves

$$\begin{bmatrix} 52.000 & 2.203 \\ 2.203 & 148.729 \end{bmatrix} \Delta\mathbf{z} = \begin{bmatrix} 0 \\ 28.797 \end{bmatrix}, \qquad \Delta\mathbf{z} = \begin{bmatrix} -0.0082 \\ 0.1937 \end{bmatrix}.$$

Table 12.3 lists the iteration: each row is one iterate, the columns give the two entries of the state, the two entries of the gradient $\mathbf{J}^{\mathsf{T}}\mathbf{r}$ and the three distinct entries of the Hessian $\mathbf{J}^{\mathsf{T}}\mathbf{J}$. The first step takes nearly the whole trend at once; from the second the convergence is quadratic, the gradient norms falling as 28.8, 2.98, 0.074 and 6×10^{-5}, and the gradient is below 10^{-10} at the fourth iterate. The off-diagonal entry of the Hessian moves from 2.2 at the start to 31 at the end, as the second score equation takes hold.

Table 12.3: Newton iteration of equation (12.10) for the fleet's trend from the constant-rate state $\mathbf{z}^{(0)} = (-1.5126,\ 0)^{\mathsf{T}}$: the state, the gradient $\mathbf{J}^{\mathsf{T}}\mathbf{r}$ and the Hessian $\boldsymbol{\Lambda} = \mathbf{J}^{\mathsf{T}}\mathbf{J}$ at each iterate.

Iterate m [–]	α [–]	γ [yr^{-1}]	$\partial\mathcal{L}/\partial\alpha$ [–]	$\partial\mathcal{L}/\partial\gamma$ [yr]	$[\boldsymbol{\Lambda}]_{11}$ [–]	$[\boldsymbol{\Lambda}]_{12}$ [yr]	$[\boldsymbol{\Lambda}]_{22}$ [yr^2]
0	−1.5126	0.0000	0	−28.797	52.000	2.203	148.729
1	−1.5208	0.1937	2.827	0.930	54.827	31.931	164.373
2	−1.5752	0.1987	7.3×10^{-2}	5.6×10^{-3}	52.073	31.006	156.437
3	−1.5768	0.1989	5.7×10^{-5}	-2.3×10^{-5}	52.000	31.000	156.235
4	−1.5768	0.1989	4.3×10^{-11}	-4.6×10^{-11}	52.000	31.000	156.234

The MAP state. The iteration stops at

$$\hat{\mathbf{z}} = \begin{bmatrix} \hat{\alpha} \\ \hat{\gamma} \end{bmatrix} = \begin{bmatrix} -1.5768 \\ 0.1989\,\mathrm{yr}^{-1} \end{bmatrix}, \qquad \mathrm{e}^{\hat{\alpha}} = 0.2066 \text{ per unit-year}, \qquad \mathrm{e}^{\hat{\gamma}} = 1.220. \tag{12.11}$$

The fleet's rate at the record's mid-time is 0.207 per unit-year and grows by 22.0 % per year, doubling every $\ln 2/0.1989 = 3.48$ years by equation (12.4). The residual vector at the MAP state is $\mathbf{r}(\hat{\mathbf{z}}) = (-0.224,\ 0.054,\ 0.177,\ 0.043,\ 0.042,\ -0.111)^{\mathsf{T}}$, six working residuals orthogonal to both columns of $\mathbf{J}$.

The Jacobian, the information and the covariance. At the MAP state the expected counts are $\hat{\boldsymbol{\mu}} = (4.524,\ 6.133,\ 7.483,\ 9.130,\ 11.139,\ 13.591)^{\mathsf{T}}$, and the Jacobian, 6×2, and the information and covariance matrices are

$$\mathbf{J}(\hat{\mathbf{z}}) = \begin{bmatrix} 2.127 & -5.317 \\ 2.476 & -3.715 \\ 2.735 & -1.368 \\ 3.022 & 1.511 \\ 3.338 & 5.006 \\ 3.687 & 9.217 \end{bmatrix}, \qquad \boldsymbol{\Lambda} = \mathbf{J}^{\mathsf{T}}\mathbf{J} = \begin{bmatrix} 52 & 31 \\ 31 & 156.23 \end{bmatrix}, \tag{12.12}$$

$$\boldsymbol{\Sigma} = \boldsymbol{\Lambda}^{-1} = \begin{bmatrix} 0.021811 & -0.004328 \\ -0.004328 & 0.007259 \end{bmatrix}.$$

The entries $[\boldsymbol{\Lambda}]_{11} = 52$ and $[\boldsymbol{\Lambda}]_{12} = 31$ are the sums of the score equations (12.9); $[\boldsymbol{\Lambda}]_{22} = 156.23$ is the fitted counts against the squared times $x_j^2 = (6.25,\ 2.25,\ 0.25,\ 0.25,\ 2.25,\ 6.25)$. Read off,

$$\sqrt{\Sigma_{\alpha\alpha}} = 0.148, \qquad \sqrt{\Sigma_{\gamma\gamma}} = 0.0852\,\mathrm{yr}^{-1}, \qquad \mathrm{corr}(\alpha,\gamma) = -0.344 : \tag{12.13}$$

the trend is 0.199 ± 0.085 per year, 2.33 standard deviations from zero, and the correlation is negative, the inner product 31 of the columns being positive. The mean failure time of equation (12.18) is $t_w = 3.5 + 31/52 = 4.096$.

The fitted rates. Table 12.4 gives the fitted model window by window. Each row is a window; the columns give its observed and fitted rates, the expected count under the trend and the Pearson residual $(k_j - \hat{\lambda}_j E_j)/\sqrt{\hat{\lambda}_j E_j}$, and then the expected count and the residual under the constant rate $52/236 = 0.2203$ per unit-year. The fitted rate rises from 0.126 in the first window to 0.340 in the sixth. Section 12.4.4 reads the two residual columns as the check of each model.

Table 12.4: The fitted rates of the trend model and of the constant rate, with the expected counts and the Pearson residuals $(k_j - \hat{\lambda}_j E_j)/\sqrt{\hat{\lambda}_j E_j}$. Rates in failures per unit-year. Trend: $\hat{\lambda}_j = \exp(-1.5768 + 0.1989\, x_j)$; constant: $\hat{\lambda} = 0.2203$.

Window	k_j	Observed rate	Fitted rate, trend	Expected, trend	Residual, trend	Expected, constant	Residual, constant
1	5	0.139	0.126	4.52	0.22	7.93	−1.04
2	6	0.150	0.153	6.13	−0.05	8.81	−0.95
3	7	0.175	0.187	7.48	−0.18	8.81	−0.61
4	9	0.225	0.228	9.13	−0.04	8.81	0.06
5	11	0.275	0.278	11.14	−0.04	8.81	0.74
6	14	0.350	0.340	13.59	0.11	8.81	1.75
Total	52			52.00	$\sum r^2 =$ 0.10	52.00	$\sum r^2 =$ 5.95

Figure 12.2 draws the record and both fits. The horizontal axis is the window, the vertical axis the rate in failures per unit-year on a logarithmic scale, on which the trend model is a

straight line. The dots are the observed rates k_j/E_j, the solid blue line the fitted trend and the dashed gold line the constant rate. The dots follow the blue line closely and cross the gold line once, from below to above, between the third and fourth windows. Under a constant rate the dots would scatter about the gold line with no order; here the three early windows lie below it and the three late ones above it, with the gap growing toward both ends.

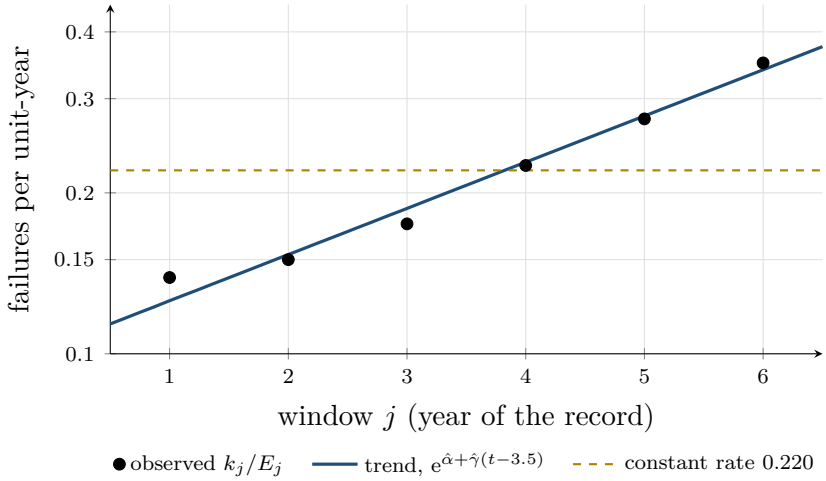


Figure 12.2: The record of the older cooling fleet and two fits, on a logarithmic rate axis. Dots: observed rates k_j/E_j of Table 12.1. Solid: the trend model at the MAP state, $\lambda(t) = \exp(-1.5768 + 0.1989\,(t - 3.5))$ per unit-year. Dashed: the constant rate $52/236 = 0.2203$ per unit-year.

The pair $(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of equations (12.11) and (12.12) is what step 4 uses: every answer is computed from it, the comparison with the prior row added to the same fit and the constant rate with the trend column removed, the one-number fit $\hat{\alpha}_0 = \ln(52/236)$ with $\sigma_\alpha^2 = 1/52$. The exact posterior on the grid, with the exact marginal of the trend, gives the value set beside each, with the gap.

12.3 The answers (step 4)

The six questions of Table 12.2 take four of the general forms of Chapter 2: estimation, prediction under a threshold, hypothesis comparison and the predictive, written in the book's notation, with $p(\mathbf{z} \mid \mathbf{y})$ the posterior density of the state, $f(\mathbf{z})$ a quantity derived from it, f^* a level of that quantity and $p(\mathbf{y} \mid M)$ the evidence of a drawing M. Each is answered from the posterior of Section 12.2: its engineering question, the general form of its category, the form it takes for the fleet, the quantities of the posterior it uses, and the fleet's numbers. The questions name the limit $\lambda_{\max} = 0.5$ failures per unit-year; the times $t = 6$, 7 and 8; the prior spread $s_\gamma = 0.3\,\mathrm{yr}^{-1}$ of the comparison; and the next window's exposure $E_7 = 40$ unit-years. Every answer is computed from the posterior of the state of step 3, $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$: a normal band or tail of a quantity linear in the state, the inversion of that tail for the crossing time, the prior row solved on the same fit for the comparison, and a Gauss–Hermite sum for the count; that is the value the chapter reports. The exact routes of Section 12.2.2, the grid and the exact marginal of the trend, give the value set beside each answer, with the gap. The draws of the Laplace posterior and the delta band of the crossing time are checks of Section 12.4.

12.3.1 The rate at a stated time

Engineering question. What is the rate now and next year, with its band? The contract is written on the fleet's rate, and the operator wants it at the last window of record, $t = 6$, and at the next, $t = 7$, each with an interval that contains it with a stated probability.

General form. An estimation question about a quantity derived from the state asks for its posterior distribution, the estimation form (2.9) of Chapter 2,

$$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y}) = \int \mathbf{1}\big[f(\mathbf{z}) \le f^*\big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z},$$

summarized by its median and its central 95 % credible interval $[f_L, f_U]$ of equation (2.11).

Specialized form. The question is equation (2.9) with $f(\mathbf{z}) = \ln \lambda(t)$, a level f^* on the log-rate, and $p(\mathbf{z} \mid \mathbf{y})$ the posterior of step 3, $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$, summarized by equation (2.11). On the logarithmic scale the rate is linear in the state,

$$\ln \lambda(t) = \mathbf{g}(t)^{\mathsf{T}} \mathbf{z}, \qquad \mathbf{g}(t) = \begin{bmatrix} 1 \\ t - \bar{t} \end{bmatrix}, \tag{12.14}$$

with a gradient that does not depend on the state; at a window's center it is that window's row of the log-count derivatives $\mathbf{X}$ of Section 12.1.5, its Jacobian row divided by $\sqrt{\mu_j}$. The delta method of Chapter 4 is therefore exact for it under the Gaussian posterior: $\ln \lambda(t) \mid \mathbf{y} \sim \mathcal{N}\big(\ln \hat{\lambda}(t),\ V(t)\big)$ with center $\ln \hat{\lambda}(t) = \mathbf{g}(t)^{\mathsf{T}} \hat{\mathbf{z}} = \hat{\alpha} + \hat{\gamma}(t - \bar{t})$ and variance

$$V(t) = \mathbf{g}(t)^{\mathsf{T}} \boldsymbol{\Sigma}\, \mathbf{g}(t) = \Sigma_{\alpha\alpha} + 2(t - \bar{t})\Sigma_{\alpha\gamma} + (t - \bar{t})^2 \Sigma_{\gamma\gamma}, \tag{12.15}$$

a quadratic in t, the variance budget of the rate at t with one term per entry and the cross term in the middle; the entries of $\boldsymbol{\Sigma}$ are written with the names of the two state entries. Since the exponential is increasing, the central 95 % band of the rate is the image of the band $\ln \hat{\lambda}(t) \pm 1.96\sqrt{V(t)}$ of the log-rate,

$$\lambda(t) \in \Big[\hat{\lambda}(t)\, \mathrm{e}^{-1.96\sqrt{V(t)}},\ \hat{\lambda}(t)\, \mathrm{e}^{1.96\sqrt{V(t)}}\Big], \qquad \hat{\lambda}(t) = \mathrm{e}^{\mathbf{g}(t)^{\mathsf{T}} \hat{\mathbf{z}}}, \tag{12.16}$$

with the median $\hat{\lambda}(t)$; the posterior mean of the rate is larger by the factor $\mathrm{e}^{V(t)/2}$. The value beside it comes from the exact posterior: the same quantiles, read as weighted quantiles of $\mathbf{g}(t)^{\mathsf{T}} \mathbf{z}_i$ over the grid of Section 12.2.2.

The origin of time as a change of the Jacobian. Where the band is narrowest depends on the cross term of equation (12.15), and the cross term depends on where time is measured from. Measuring time from an origin c instead of $\bar{t}$ changes only the time column, from x_j to $t_j - c = x_j - (c - \bar{t})$, and the state from $(\alpha, \gamma)^{\mathsf{T}}$ to $\mathbf{z}_c = (\alpha_c, \gamma)^{\mathsf{T}}$, with $\alpha_c = \alpha + \gamma(c - \bar{t})$ the log-rate at c: the same line, written about another point. As matrices,

$$\begin{gathered} \mathbf{J}_c = \mathbf{J}\, \mathbf{U}_c, \qquad \mathbf{U}_c = \begin{bmatrix} 1 & -(c - \bar{t}) \\ 0 & 1 \end{bmatrix}, \qquad \mathbf{z}_c = \mathbf{U}_c^{-1} \mathbf{z}, \\ \boldsymbol{\Lambda}_c = \mathbf{U}_c^{\mathsf{T}} \boldsymbol{\Lambda}\, \mathbf{U}_c, \qquad \boldsymbol{\Sigma}_c = \mathbf{U}_c^{-1} \boldsymbol{\Sigma}\, \mathbf{U}_c^{-\mathsf{T}}, \end{gathered} \tag{12.17}$$

so that the log counts $\mathbf{X}\mathbf{U}_c\mathbf{z}_c = \mathbf{X}\mathbf{z}$ are unchanged and the information and the covariance transform with $\mathbf{U}_c$. With $\mathbf{1}$ the level column and $\mathbf{x}$ the time column of $\mathbf{X}$, and $\mathbf{W}$ the diagonal

of expected counts of Section 12.1.5, the off-diagonal entry of $\mathbf{\Lambda}_c$ is the inner product of the level column with the new time column,

$$[\mathbf{\Lambda}_c]_{12} = \mathbf{1}^\mathsf{T}\mathbf{W}\big(\mathbf{x} - (c - \bar{t})\,\mathbf{1}\big) = [\mathbf{\Lambda}]_{12} - (c - \bar{t})\,[\mathbf{\Lambda}]_{11},$$

which vanishes at

$$c = t_w = \bar{t} + \frac{[\mathbf{\Lambda}]_{12}}{[\mathbf{\Lambda}]_{11}} = \bar{t} + \frac{\sum_j x_j k_j}{\sum_j k_j}, \tag{12.18}$$

the second form because at the MAP state the score equations of Section 12.1.5 make $[\mathbf{\Lambda}]_{11} = \sum_j \hat{\mu}_j$ equal to the total count $\sum_j k_j$ and $[\mathbf{\Lambda}]_{12} = \sum_j x_j \hat{\mu}_j$ equal to $\sum_j x_j k_j$. The time t_w is the mean time of the recorded failures. Centering at t_w makes the two columns of the Jacobian orthogonal, so the information is diagonal and the two entries are uncorrelated.

The minimum of the quadratic. The change of the Jacobian finds the minimum of equation (12.15) without completing a square. In the state written about t_w, the gradient of $\ln\lambda(t)$ is $\mathbf{U}_{t_w}^\mathsf{T}\mathbf{g}(t) = (1,\ t - t_w)^\mathsf{T}$, and the covariance is diagonal, with first entry $1/[\mathbf{\Lambda}]_{11} = 1/\sum_j k_j$ and second entry $\Sigma_{\gamma\gamma}$, which $\mathbf{U}_c$ leaves alone, so that

$$V(t) = \begin{bmatrix}1 & t - t_w\end{bmatrix}\begin{bmatrix}1/\sum_j k_j & 0\\ 0 & \Sigma_{\gamma\gamma}\end{bmatrix}\begin{bmatrix}1\\ t - t_w\end{bmatrix} = \frac{1}{\sum_j k_j} + \Sigma_{\gamma\gamma}\,(t - t_w)^2. \tag{12.19}$$

Completing the square in equation (12.15) with the entries of $\mathbf{\Sigma}$ gives the same quadratic. The band of the rate is narrowest at the mean failure time of the record, where its variance is $1/\sum_j k_j$, exactly the variance of the log-rate of a constant-rate fit to all the failures pooled; away from it the variance grows with the square of the lever arm $t - t_w$, at the rate set by the uncertainty of the trend. The minimum is at $\bar{t}$ only when the two entries are uncorrelated. Everything in the record beyond the total count goes into the trend, and the trend's uncertainty is what a forecast pays for.

What the choice of origin changes. It changes nothing in any answer. The origins describe the same line, and a derived quantity with gradient $\mathbf{g}$ in the state $\mathbf{z}$ has the gradient $\mathbf{U}_c^\mathsf{T}\mathbf{g}$ in $\mathbf{z}_c$, so that $(\mathbf{U}_c^\mathsf{T}\mathbf{g})^\mathsf{T}\mathbf{\Sigma}_c(\mathbf{U}_c^\mathsf{T}\mathbf{g}) = \mathbf{g}^\mathsf{T}\mathbf{\Sigma}\mathbf{g}$ whichever origin is used. It changes how the answer reads and how well the computation behaves. An intercept far outside the record is an extrapolation, poorly known and nearly a linear function of the trend; its variance and the cross term are large and of opposite sign, and the variance of any rate inside the record is then a small difference of large numbers. Near the center of the record each entry has a meaning of its own, the level of the rate and its trend, and the covariance is nearly diagonal.

Building blocks. The MAP state $\hat{\mathbf{z}}$ and the covariance $\mathbf{\Sigma}$ (Section 12.2.1); the gradient $\mathbf{g}(t)$ at $t = 6$ and 7; the sums $\sum_j k_j$ and $\sum_j x_j k_j$ of the score equations (Section 12.1.5); the normal quantile 1.96; and, for the exact values beside them, the grid points and weights (Section 12.2.2).

The origin of time. Table 12.5 writes the fit about three origins by equation (12.17). Each column is an origin c; the rows give the time column $t_j - c$, the intercept $\hat{\alpha}_c$ at the MAP state, the three distinct entries of the information $\mathbf{\Lambda}_c$ and of the covariance $\mathbf{\Sigma}_c$, the standard deviation of the intercept and the correlation of the two entries. At the origin of the scale, $t = 0$, just before the record, the time column $(1, \ldots, 6)$ points nearly along the level column, $[\mathbf{\Lambda}_c]_{12} = 213.0$, and the correlation is -0.93, with the intercept's standard deviation 0.376 against 0.148 for α. At the record's mid-time $\bar{t} = 3.5$ the correlation is -0.34. At t_w the information and the covariance are diagonal: the intercept $\hat{\alpha}_{t_w} = -1.4582$, a rate of 0.2327 per unit-year, has the variance $1/52$ of the total count, and the correlation is zero. In every column $\Sigma_{\gamma\gamma} = 0.007259$.

Table 12.5: The fit written about three origins of time c by equation (12.17): the time column $t_j - c$ of the log-count derivatives, the intercept at the MAP state, the information and the covariance at the MAP state, and the standard deviation and correlation read off $\mathbf{\Sigma}_c$.

Quantity [unit]	$c = 0$	$c = \bar{t} = 3.5$	$c = t_w = 4.096$
Time column $t_j - c$ [yr]	$1, 2, \ldots, 6$	$-2.5, -1.5, \ldots, 2.5$	$-3.096, -2.096, \ldots, 1.904$
Intercept $\hat{\alpha}_c$ [–]	-2.273	-1.577	-1.458
$[\mathbf{\Lambda}_c]_{11}$ [–]	52	52	52
$[\mathbf{\Lambda}_c]_{12}$ [yr]	213.0	31	0
$[\mathbf{\Lambda}_c]_{22}$ [yr^2]	1010.23	156.23	137.75
$[\mathbf{\Sigma}_c]_{11}$ [–]	0.141031	0.021811	0.019231
$[\mathbf{\Sigma}_c]_{12}$ [yr^{-1}]	-0.029735	-0.004328	0
$[\mathbf{\Sigma}_c]_{22}$ [yr^{-2}]	0.007259	0.007259	0.007259
$\sqrt{[\mathbf{\Sigma}_c]_{11}}$ [–]	0.376	0.148	0.139
corr(α_c, γ) [–]	-0.93	-0.34	0

The variance of the log-rate. With $\mathbf{\Sigma}_{t_w} = \text{diag}(0.019231,\ 0.007259)$, equation (12.19) is

$$V(t) = \frac{1}{52} + 0.007259\,(t - 4.096)^2,$$

smallest, 0.139^2, at t_w, which lies 0.60 years after $\bar{t}$ because $\text{corr}(\alpha, \gamma) = -0.34$. Figure 12.3 draws the standard deviation $\sqrt{V(t)}$. The horizontal axis is calendar time from 0 to 12 years, the vertical axis the posterior standard deviation of $\ln \lambda(t)$; the shaded band marks the six years of the record. The solid curve is equation (12.19), with its minimum 0.139 at $t_w = 4.10$ marked; the dashed curve is the same quadratic with the cross term omitted, $\sqrt{\Sigma_{\alpha\alpha} + \Sigma_{\gamma\gamma}(t - \bar{t})^2}$. At the window centers of the record the standard deviation lies between 0.14 and 0.30; beyond the record the curve becomes nearly straight with slope $\sqrt{\Sigma_{\gamma\gamma}} = 0.085$ per year, reaching 0.28 at the next window, 0.36 two years out and 0.52 four years out. The dashed curve is too wide by up to 0.05 to the right of $\bar{t}$ and too narrow to the left: omitting the correlation moves the narrow point of the band to the wrong place.

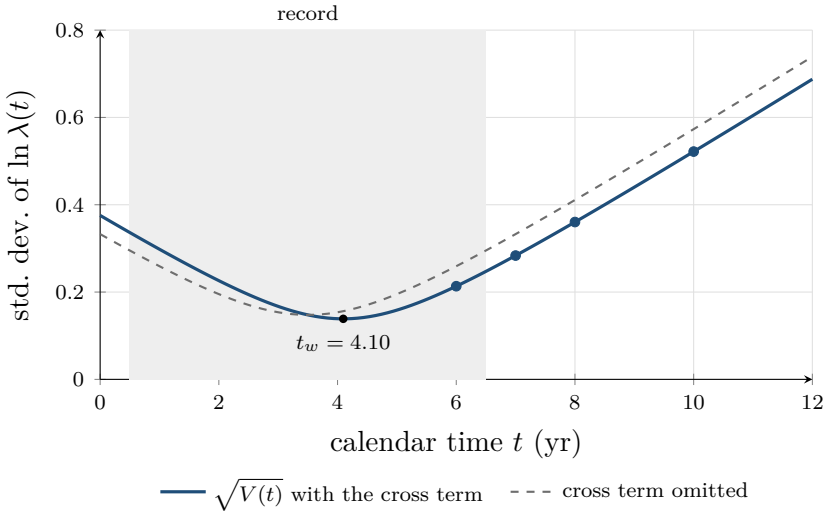


Figure 12.3: Posterior standard deviation of the log-rate at calendar time t, equation (12.19): $V(t) = 1/52 + 0.007259\,(t - 4.096)^2$, minimum 0.1387 at the mean failure time $t_w = 4.096$. Dots: windows 6, 7, 8 and 10. Dashed: $0.021811 + 0.007259\,(t - 3.5)^2$, the cross term $2(t - \bar{t})\Sigma_{\alpha\gamma}$ omitted. Shaded: the years of the record.

The rate now and next year. For the last window of the record, $t = 6$, the gradient is $\mathbf{g}(6) = (1,\ 2.5)^{\mathsf{T}}$, so that $\ln \hat{\lambda}(6) = \mathbf{g}(6)^{\mathsf{T}}\hat{\mathbf{z}} = -1.0795$ and $V(6) = \mathbf{g}(6)^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{g}(6) = 0.04554$, and by equation (12.16)

$$\lambda(6) \in \left[0.340\,\mathrm{e}^{-1.96\times 0.2134},\ 0.340\,\mathrm{e}^{1.96\times 0.2134}\right] = [0.224,\ 0.516] \text{ per unit-year.}$$

For next year, $t = 7$, the median is 0.415 and the band [0.238, 0.723]. The band's width, as a ratio of its ends, grows from 1.7 at $t = 4$ to 2.3 now and 7.7 four years out. The values beside them come from the exact posterior on the grid: now the median 0.337 and the band [0.216, 0.501], next year 0.412 and [0.230, 0.704]. The state route is higher by 0.003 at both medians, by 0.008 at both lower ends and by 0.015 and 0.019 at the upper ends, because the exact posterior puts the rate a little lower.

Figure 12.4 draws the forecast. In panel (a) the horizontal axis is calendar time from the first window to $t = 10$, the vertical axis the rate on a logarithmic scale; the shaded region is the central 95 % band of $\lambda(t)$, the solid line its median, the dashed green lines the ends of the grid's band, the dots the observed rates and the red line the limit of 0.5 per unit-year. The dotted verticals mark where the median and the upper edge of the band reach the limit. On the logarithmic axis the band is symmetric about the median, narrowest near $t = 4$ and widening steadily beyond the record. Panel (b) draws the same median and band edges as lines out to $t = 32$, with dots where the edges meet the limit. The upper edge of the rate band reaches the limit at $t = 5.90$, the median at 7.94; the lower edge rises so slowly that it reaches the limit only at $t = 28.3$. These three times return in Section 12.3.3 as the median and the ends of the band of the crossing time.

12.3.2 The sign of the trend and the limit two years ahead

Engineering questions. Is the failure rate rising? Will the rate be above the limit two years from now? The first asks whether the rise of the record is a rise of the fleet; the second whether, at $t = 8$, the middle of window 8, the rate is above the contract's 0.5 failures per unit-year.

General form. A prediction under a threshold is the posterior probability of the event $f(\mathbf{z}) \geq f^*$, the threshold form (2.12) of Chapter 2,

$$\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}) = \int \mathbf{1}\big[f(\mathbf{z}) \geq f^*\big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z};$$

the two questions ask for the probability itself, with no assurance level attached.

Specialized form. Three choices enter here. "Rising" is read inside the trend model, as the event $\gamma > 0$: this makes the answer the tail of one entry, and it says nothing about whether the trend belongs in the model at all, which is the comparison of Section 12.3.4. The limit is read on the rate at a point in time, as the contract states it: this makes the event a half-line of a quantity linear in the state; a contract written on the number of failures in a year would need the predictive count of Section 12.3.5 instead. A tail probability is where the skew of the exact posterior acts, so the grid of Section 12.2.2 gives the value set beside each.

The questions are equation (2.12) with $p(\mathbf{z} \mid \mathbf{y})$ the posterior of step 3, $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$, and two quantities $f(\mathbf{z})$. Within the trend model the rate is rising if and only if $\gamma > 0$, the event $f(\mathbf{z}) = \mathbf{g}^{\mathsf{T}}\mathbf{z} \geq f^* = 0$ with $\mathbf{g} = (0,\ 1)^{\mathsf{T}}$, and the rate at a time t is above the limit if and only if $\ln \lambda(t) > \ln \lambda_{\max}$, a half-line of the quantity $f(\mathbf{z}) = \mathbf{g}(t)^{\mathsf{T}}\mathbf{z}$ of equation (12.14), with $f^* = \ln \lambda_{\max}$. Under the Laplace posterior the marginal of γ is $\mathcal{N}(\hat{\gamma}, \Sigma_{\gamma\gamma})$, the second entry of $\hat{\mathbf{z}}$ and the second diagonal entry of $\boldsymbol{\Sigma}$, and the probability of a positive trend is the normal tail of Chapter 4 at the normalized entry,

$$\Pr(\gamma > 0 \mid \mathbf{y}) = \Phi\left(\frac{\hat{\gamma}}{\sqrt{\Sigma_{\gamma\gamma}}}\right), \tag{12.20}$$

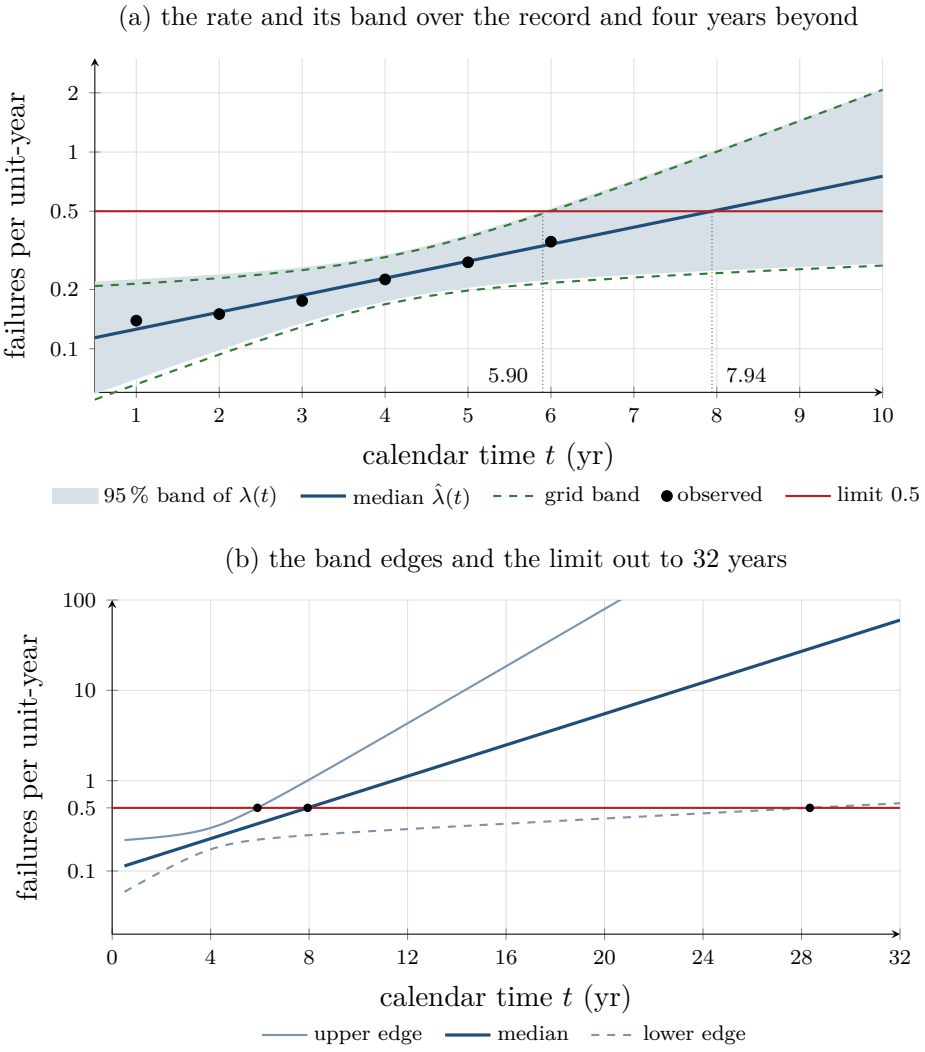


Figure 12.4: The forecast of the fleet's rate. (a) Median $\hat{\lambda}(t) = \exp(-1.5768 + 0.1989\,(t - 3.5))$ and central 95 % band $\hat{\lambda}(t)\exp(\pm 1.96\sqrt{V(t)})$ with $V(t) = 1/52 + 0.007259\,(t - 4.096)^2$; dots: observed rates; red: the limit $\lambda_{\max} = 0.5$ per unit-year; dashed green: the ends of the band on the grid, the exact posterior. The upper edge of the band reaches the limit at $t = 5.90$ and the median at $t = 7.94$. (b) The same curves to $t = 32$; the dots on the limit mark where the upper edge, the median and the lower edge reach it, at $t = 5.90$, 7.94 and 28.3.

with Φ the standard normal cumulative distribution, and the probability that the rate exceeds the limit is again a normal tail, with the center and the variance of equation (12.15):

$$\Pr(\lambda(t) > \lambda_{\max} \mid \mathbf{y}) = \Phi\left(\frac{\ln \hat{\lambda}(t) - \ln \lambda_{\max}}{\sqrt{V(t)}}\right). \tag{12.21}$$

The two answers are one calculation with two gradients: $(0,\ 1)^{\mathsf{T}}$, which reads the trend entry alone, and $(1,\ t - \bar{t})^{\mathsf{T}}$, which combines the level and the trend at the time asked. Both are exact under the Gaussian posterior, since the threshold maps exactly through the logarithm and the quantity is linear in the state; the Laplace Gaussian itself is the only approximation. The value beside it comes from the exact posterior: on the grid, with $\mathbf{z}_i$ its points and w_i their normalized weights, each tail is the total weight of the points in the event,

$$\Pr(\gamma > 0 \mid \mathbf{y}) = \sum_{i:\ \gamma_i > 0} w_i, \qquad \Pr(\lambda(t) > \lambda_{\max} \mid \mathbf{y}) = \sum_{i:\ \mathbf{g}(t)^{\mathsf{T}}\mathbf{z}_i > \ln \lambda_{\max}} w_i, \tag{12.22}$$

the integral of the threshold form evaluated as a weighted sum.

Building blocks. The entry $\hat{\gamma}$ and the variance $\Sigma_{\gamma\gamma}$ (Section 12.2.1); the gradients $(0,\ 1)^{\mathsf{T}}$ and $\mathbf{g}(t)$ at $t = 6$, 7 and 8, with $V(t)$; $\ln 0.5$; Φ; and, for the exact values beside them, the grid points and weights (Section 12.2.2).

The two tails for the fleet. By equation (12.20),

$$\Pr(\gamma > 0 \mid \mathbf{y}) = \Phi\left(\frac{0.1989}{0.0852}\right) = \Phi(2.335) = 0.990 :$$

within a model that allows a trend, the record makes a falling or flat rate improbable. The value beside it comes from the exact posterior: by the grid sum (12.22), the weight of the grid points with $\gamma > 0$ is

$$\Pr(\gamma > 0 \mid \mathbf{y}) = 0.992,$$

0.002 higher. At $t = 8$, with $\mathbf{g}(8) = (1,\ 4.5)^{\mathsf{T}}$, $\ln \hat{\lambda}(8) = \mathbf{g}(8)^{\mathsf{T}}\hat{\mathbf{z}} = -0.6816$ against $\ln 0.5 = -0.6931$, and $\sqrt{\mathbf{g}(8)^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{g}(8)} = 0.3604$, so that by equation (12.21)

$$\Pr(\lambda(8) > 0.5 \mid \mathbf{y}) = \Phi\left(\frac{-0.6816 + 0.6931}{0.3604}\right) = \Phi(0.032) = 0.513.$$

On the grid, the weight of the points with $\mathbf{g}(8)^{\mathsf{T}}\mathbf{z}_i > \ln 0.5 = -0.6931$ is

$$\Pr(\lambda(8) > 0.5 \mid \mathbf{y}) = 0.509,$$

0.004 lower. The median rate two years from now, 0.506 under the Laplace posterior, sits almost exactly on the limit, and whether it is above or below is a coin toss on both routes. The rate is already above the limit now with probability 0.035 (0.026 on the grid), and next year with probability 0.254 (0.246).

Table 12.6 collects the band and the exceedance probability from the middle of the record to four years beyond it. Each row is a time; the columns give the standard deviation of the log-rate, the median rate and the 95 % band of equation (12.16), and the probability that the rate exceeds the limit, the normal tail (12.21) with the grid sum (12.22) in parentheses.

12.3.3 The crossing time

Engineering question. When does the rate cross 0.5 per unit-year? The operator asks for the date at which the contract stops covering the fleet, and for how early and how late it could be.

Table 12.6: The fleet's rate at calendar time t, in failures per unit-year: median $\hat{\lambda}(t)$, central 95 % band $\hat{\lambda}(t)\,\mathrm{e}^{\pm 1.96\sqrt{V(t)}}$ under the Laplace posterior, and the probability of exceeding the limit, equation (12.21), with the grid sum (12.22) in parentheses.

Time t (yr)	$\sqrt{V(t)}$	Median	95 % band	$\Pr(\lambda(t) > 0.5 \mid \mathbf{y})$
4 (record)	0.139	0.228	[0.174, 0.300]	0.0000 (0.0000)
6 (now)	0.213	0.340	[0.224, 0.516]	0.035 (0.026)
7 (next year)	0.284	0.415	[0.238, 0.723]	0.254 (0.246)
8	0.360	0.506	[0.250, 1.025]	0.513 (0.509)
9	0.440	0.617	[0.260, 1.463]	0.684 (0.681)
10	0.522	0.753	[0.271, 2.094]	0.784 (0.782)

General form. The crossing time t^* is a quantity derived from the state, and the question takes the estimation form (2.9),

$$\Pr(t^* \le s \mid \mathbf{y}) = \int \mathbf{1}\big[t^*(\mathbf{z}) \le s\big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z},$$

read as a function of the time s and summarized by its median and the central 95 % band of equation (2.11).

Specialized form. Two choices enter here. A state with $\gamma \le 0$ is taken never to cross, its crossing time infinite: this keeps every state of the posterior in the distribution, with an atom at infinity, the probability of never crossing, and it would mislead if a flat or falling rate could turn upward later, which the log-linear line excludes (Section 12.6.6). The rate at the record's mid-time is taken to be below the limit, so that a line with $\gamma > 0$ crosses once, in the future; a state whose rate is already above the limit at $\bar{t}$ is outside this reading, and Section 12.3.3 measures how improbable it is.

The crossing time. The question is equation (2.9) with $f(\mathbf{z}) = t^*(\mathbf{z})$, $f^* = s$, and $p(\mathbf{z} \mid \mathbf{y})$ the posterior of step 3, $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$. The rate reaches the limit at the time t^* at which $\ln \lambda(t^*) = \ln \lambda_{\max}$. By equation (12.1), $\alpha + \gamma(t^* - \bar{t}) = \ln \lambda_{\max}$, and solving for t^*,

$$t^* = \bar{t} + \frac{\ln \lambda_{\max} - \alpha}{\gamma} \quad (\gamma > 0), \qquad t^* = \infty \quad (\gamma \le 0). \tag{12.23}$$

At the MAP state it gives the plug-in date $\hat{t}^* = t^*(\hat{\mathbf{z}})$, the date at which the line of the MAP state reaches the limit.

The crossing event is a threshold event. For $\gamma > 0$ the rate increases, so it has crossed the limit by time s if and only if it is above the limit at s:

$$t^* \le s \iff \ln \lambda(s) \ge \ln \lambda_{\max} \qquad (\gamma > 0).$$

For $\gamma \le 0$ the crossing never happens, and the rate at any $s > \bar{t}$ is below its value at $\bar{t}$, which is below the limit except with negligible probability; the event $\ln \lambda(s) \ge \ln \lambda_{\max}$ then has negligible probability as well. Up to that negligible term the distribution function of the crossing time is the exceedance probability of the rate, equation (12.21), read as a function of the time:

$$\Pr(t^* \le s \mid \mathbf{y}) = \Pr(\lambda(s) \ge \lambda_{\max} \mid \mathbf{y}) = \Phi\left(\frac{\ln \hat{\lambda}(s) - \ln \lambda_{\max}}{\sqrt{V(s)}}\right). \tag{12.24}$$

The crossing time is a nonlinear function of the state, but each event about it is an event about a quantity that is linear in the state, and that event has an exact normal tail under the Gaussian

posterior. The fourth question and the third are the same calculation: the probability that the rate is above the limit at $t = 8$ is the probability that the crossing has happened by then. This inversion is the route the chapter reports, and its density is the derivative of equation (12.24) in s, plotted beside the draws in Section 12.4.2.

The median and the mass at infinity. The distribution function equals one half where its argument is zero, $\ln \hat{\lambda}(s) = \ln \lambda_{\max}$, which is $s = \hat{t}^*$: the median of the crossing time is the plug-in date exactly. The probability of never crossing is the posterior mass at $\gamma \le 0$, $\Phi(-\hat{\gamma}/\sqrt{\Sigma_{\gamma\gamma}})$, the complement of the sign probability (12.20).

The band in closed form. The ends of the central 95 % band solve $\Pr(t^* \le s \mid \mathbf{y}) = 0.025$ and 0.975, that is $\mathbf{g}(s)^\mathsf{T}\hat{\mathbf{z}} - \ln \lambda_{\max} = \mp c\sqrt{\mathbf{g}(s)^\mathsf{T}\boldsymbol{\Sigma}\,\mathbf{g}(s)}$ with $c = \Phi^{-1}(0.975) = 1.960$. Write $\mathbf{d} = \hat{\mathbf{z}} - (\ln \lambda_{\max},\ 0)^\mathsf{T}$, the MAP state measured from the limit, so that the left side is $\mathbf{g}(s)^\mathsf{T}\mathbf{d}$. Both ends satisfy the squared equation, which is a quadratic form in the gradient,

$$\begin{aligned} \mathbf{g}(s)^\mathsf{T}\mathbf{M}\,\mathbf{g}(s) = 0, \qquad \mathbf{M} = \mathbf{d}\mathbf{d}^\mathsf{T} - c^2\boldsymbol{\Sigma}, \qquad \mathbf{g}(s) = \begin{bmatrix} 1 \\ s - \bar{t} \end{bmatrix}, \\ A(s-\bar{t})^2 + B(s-\bar{t}) + C = 0, \qquad A = M_{22}, \quad B = 2M_{12}, \quad C = M_{11}, \end{aligned} \tag{12.25}$$

the second line being the first written out in $s - \bar{t}$. This is Fieller's construction for the ratio of two normal quantities, a standard result, here as a 2×2 matrix: the band is the set of times at which the form $\mathbf{g}(s)^\mathsf{T}\mathbf{M}\,\mathbf{g}(s)$ is not positive. Its geometry is that of the rate band (12.16): the lower end of the crossing band is the time at which the upper edge of the rate band reaches the limit, and the upper end the time at which the lower edge does.

When the band has no upper end. The upper root exists only if $A > 0$. The entry $A = M_{22} = \hat{\gamma}^2 - c^2\Sigma_{\gamma\gamma}$ is positive only if the normalized trend entry $\hat{\gamma}/\sqrt{\Sigma_{\gamma\gamma}}$ exceeds c: the trend must be more than c standard deviations from zero. The crossing band at a two-sided level q is therefore bounded only for $q < 2\Phi(\hat{\gamma}/\sqrt{\Sigma_{\gamma\gamma}}) - 1$, that is only when the trend is established at level q, the two-sided form of the sign probability (12.20). At a higher level the band has a lower end and no upper end: the record cannot exclude, at that level, that the rate never reaches the limit.

The delta route. The delta method of Chapter 4 linearizes the crossing time at the MAP state, with the gradient

$$\mathbf{g}_* = \frac{\partial t^*}{\partial \mathbf{z}} = \begin{bmatrix} -1/\gamma \\ -(\ln \lambda_{\max} - \alpha)/\gamma^2 \end{bmatrix} \tag{12.26}$$

at $\hat{\mathbf{z}}$, in years per unit of each entry: a higher level of the rate brings the crossing earlier, and so does a steeper trend. Its variance $\mathbf{g}_*^\mathsf{T}\boldsymbol{\Sigma}\,\mathbf{g}_*$ is the variance budget of equation (4.50), one term per entry and the cross term, and its band is

$$t^* \in \hat{t}^* \pm 1.96\sqrt{\mathbf{g}_*^\mathsf{T}\boldsymbol{\Sigma}\,\mathbf{g}_*}, \tag{12.27}$$

symmetric about the plug-in date, with no mass on never crossing.

Why the delta band fails. The crossing time (12.23) divides by γ, which has a pole at $\gamma = 0$, and the Gaussian posterior of γ has a positive density there. Near the pole t^* behaves as $(\ln \lambda_{\max} - \alpha)/\gamma$, and the integral of $1/\gamma$ against a density that is positive at zero diverges, as $\int_0 \mathrm{d}\gamma/\gamma$ does. The crossing time therefore has no posterior mean and no posterior variance, and the sample standard deviation of draws of it grows with their number because there is no value for it to settle on. The delta method replaces $(\ln \lambda_{\max} - \alpha)/\gamma$ by its tangent at the MAP state and returns the standard deviation of that tangent, a finite number standing in for an infinite one. The Hessian of the crossing time in the state,

$$\frac{\partial^2 t^*}{\partial \mathbf{z}\,\partial \mathbf{z}^\mathsf{T}} = \begin{bmatrix} 0 & 1/\gamma^2 \\ 1/\gamma^2 & 2(\ln \lambda_{\max} - \alpha)/\gamma^3 \end{bmatrix},$$

is indefinite, its determinant $-1/\gamma^4$ being negative, so the tangent does not lie below the ratio in every direction of the state. It does along the trend direction at fixed α: there the curvature $2(\ln\lambda_{\max}-\alpha)/\gamma^3$ is positive for every $\gamma>0$ while the line starts below the limit, and the tangent lies below the ratio over the whole range $\gamma>0$. For steep trends the tangent places crossings earlier than the ratio ever does, so the delta band reaches back to dates that the inversion (12.24) excludes. For shallow trends the tangent stays finite while the ratio rises without bound, so the delta band stops where the posterior still holds much of its mass, part of it on never crossing. The band is wrong at both ends. A derived quantity with a pole within the reach of the posterior is therefore reported by its median, its quantile band and the probability of the event that never happens, never by a mean with a standard deviation.

Draws. The check of Chapter 4 draws the state from $\mathcal{N}(\hat{\mathbf{z}},\boldsymbol{\Sigma})$ and pushes each draw through equation (12.23), a draw with $\gamma\le 0$ counting as a crossing that never happens. The draws and the inversion rest on the same Gaussian, so their agreement checks the algebra of the inversion; the grid of Section 12.2.2 removes the Gaussian itself and gives the value beside each answer, as the total weight of the grid points whose t^* lies in the event.

Building blocks. The MAP state $\hat{\mathbf{z}}$ and the covariance $\boldsymbol{\Sigma}$ (Section 12.2.1); the vector $\mathbf{d}$ and the matrix $\mathbf{M}$ of equation (12.25) with $c=1.960$, and the roots of its quadratic; the sign probability (12.20) for the mass at infinity; and, for the exact values beside them, the grid points and weights (Section 12.2.2).

The plug-in date and the start below the limit. At the MAP state equation (12.23) gives

$$\hat{t}^* = 3.5 + \frac{-0.6931+1.5768}{0.1989} = 3.5+4.442 = 7.942,$$

a crossing in the middle of the eighth year, about two years from now, consistent with the coin toss at $t=8$, and the median of the crossing time. The line does start below the limit: $\ln\lambda_{\max}-\hat{\alpha}=0.884$ is six standard deviations of α, so the case of a rate already above the limit at $\bar{t}$ has negligible probability, below 10^{-8}.

The band by inversion. For the fleet, $\mathbf{d}=(-0.88365,\ 0.19893)^{\mathsf{T}}$ and

$$\mathbf{M} = \underbrace{\begin{bmatrix} 0.780835 & -0.175782 \\ -0.175782 & 0.039572 \end{bmatrix}}_{\mathbf{d}\mathbf{d}^{\mathsf{T}}} - \underbrace{\begin{bmatrix} 0.083785 & -0.016625 \\ -0.016625 & 0.027886 \end{bmatrix}}_{c^2\boldsymbol{\Sigma}} = \begin{bmatrix} 0.697050 & -0.159157 \\ -0.159157 & 0.011686 \end{bmatrix},$$

so that $A=0.01169$, $B=-0.31831$ and $C=0.69705$ in equation (12.25), whose roots are $s-\bar{t}=(0.31831\mp\sqrt{0.06874})/0.02337=2.402$ and 24.84, and

$$t^*\in[5.90,\ 28.34]\ \text{with 95 \% posterior probability,}\qquad \text{median } 7.94. \tag{12.28}$$

The ends are the times 5.90 and 28.3 at which the upper and the lower edge of the rate band meet the limit in Figure 12.4(b). The value beside it comes from the exact posterior on the grid: the median 7.96 and the band [5.99, 27.04], the state route earlier by 0.02 at the median and 0.09 at the lower end and later by 1.30 at the upper end.

The mass at infinity and the early years. The probability of never crossing is $\Phi(-2.335)=0.0098$ (0.0085 on the grid). By equation (12.24) the probability of a crossing before $t=9$, the middle of window 9, is 0.684 (0.681), the Laplace exceedance probability at $t=9$, and before window 9 begins, at $t=8.5$, it is 0.61, on both routes.

Other levels. The normalized trend entry 2.335 exceeds $c = 1.960$, so the 95 % band is bounded. It stays bounded up to the level $2\Phi(2.335) - 1 = 98.0\,\%$, but only just: at 98 %, with $c = 2.326$, $A = 0.0003$ and the upper end is near $t = 1070$. A band of 99 %, with $c = 2.576$, has $A = -0.0086 < 0$, a lower end at 5.54 and no upper end: the record cannot exclude, at that level, that the rate never reaches the limit.

12.3.4 The evidence of a trend

Engineering question. Does the record need a trend, against a constant rate? The sign probability is read inside the trend model; whether the trend entry belongs in the model at all is a comparison of two complete drawings, M_0, the constant rate $\ln \lambda(t) = \alpha$, and M_1, the trend of equation (12.1).

General form. A hypothesis comparison is the ratio of the evidences of two complete models fitted to the same record, the comparison form (2.14) of Chapter 2,

$$B_{10} = \frac{p(\mathbf{y} \mid M_1)}{p(\mathbf{y} \mid M_0)}.$$

Specialized form. Two choices enter here. Both drawings are given the same flat prior on α: its improper constant is common to the two evidences and cancels from the ratio, which is what lets a flat prior stand in a comparison; a prior on α that differed between the drawings would enter the ratio. The trend of M_1 is given the proper prior

$$\gamma \mid M_1 \sim \mathcal{N}(0, s_\gamma^2), \qquad \text{density } \pi(\gamma),$$

centered on no trend with the spread $s_\gamma = 0.3\,\mathrm{yr}^{-1}$ of Section 12.1.5: this makes the evidence of M_1 defined, and the ratio depends on s_γ, falling without bound as it widens, which Section 12.4.3 measures.

The question is equation (2.14) with M_1 the trend of equation (12.1) under the prior row and M_0 the constant rate, and the posterior of step 3 with the prior row added. The constant rate is nested in the trend model at $\gamma = 0$, so the Savage–Dickey ratio of Chapter 7, equation (7.68), applies:

$$B_{10} = \frac{\pi(0)}{p(\gamma = 0 \mid \mathbf{y}, M_1)}, \tag{12.29}$$

the prior density of the trend at zero divided by its posterior density there, the posterior being the one under the prior π.

The prior as a row of the trend model. Under M_1 the prior on the trend is the Gaussian row of the pattern (12.2), normalized by its spread,

$$r_\pi(\mathbf{z}) = \frac{\gamma - 0}{s_\gamma}, \qquad \mathbf{J}_\pi = \frac{\partial r_\pi}{\partial \mathbf{z}} = \begin{bmatrix} 0 & 1/s_\gamma \end{bmatrix}, \qquad \mathbf{\Lambda}_\pi = \mathbf{J}_\pi^{\mathsf{T}} \mathbf{J}_\pi = \begin{bmatrix} 0 & 0 \\ 0 & 1/s_\gamma^2 \end{bmatrix},$$

a 1×2 Jacobian with an empty first column, since the prior says nothing of the level. The objective of M_1 is $\mathcal{L}(\mathbf{z}) + \frac{1}{2} r_\pi(\mathbf{z})^2$. Under the Laplace approximation the count rows enter through their quadratic expansion at the flat-prior MAP state, $\frac{1}{2}(\mathbf{z} - \hat{\mathbf{z}})^{\mathsf{T}} \mathbf{\Lambda} (\mathbf{z} - \hat{\mathbf{z}})$ up to a constant, and on that quadratic objective one Newton step from $\hat{\mathbf{z}}$ is exact. At $\hat{\mathbf{z}}$ the gradient of the count rows is zero and that of the prior row is $\mathbf{J}_\pi^{\mathsf{T}} r_\pi(\hat{\mathbf{z}}) = \mathbf{\Lambda}_\pi \hat{\mathbf{z}} = (0,\ \hat{\gamma}/s_\gamma^2)^{\mathsf{T}}$, so the normal equations, with the precisions added as in equation (3.26), and the posterior of the trend model are

$$(\mathbf{\Lambda} + \mathbf{\Lambda}_\pi)\, \Delta \mathbf{z} = -\mathbf{\Lambda}_\pi \hat{\mathbf{z}}, \qquad \tilde{\mathbf{z}} = \hat{\mathbf{z}} + \Delta \mathbf{z}, \qquad \tilde{\mathbf{\Sigma}} = (\mathbf{\Lambda} + \mathbf{\Lambda}_\pi)^{-1}, \tag{12.30}$$

with $\tilde{\mathbf{z}} = (\tilde{\alpha}, \tilde{\gamma})^{\mathsf{T}}$ and $\tilde{\mathbf{\Sigma}}$ the posterior mean and covariance under M_1. The marginal of the trend is its entry, $\gamma \mid \mathbf{y}, M_1 \sim \mathcal{N}(\tilde{\gamma},\ \tilde{\Sigma}_{\gamma\gamma})$, and its density at zero is the denominator of equation (12.29).

The ratio in closed form. Because $\mathbf{\Lambda}_\pi$ touches only the trend entry, the trend entries of the solve (12.30) are $1/\tilde{\Sigma}_{\gamma\gamma} = 1/\Sigma_{\gamma\gamma} + 1/s_\gamma^2$, the prior's precision added to the Schur complement of the level, and $\tilde{\gamma} = (\tilde{\Sigma}_{\gamma\gamma}/\Sigma_{\gamma\gamma})\,\hat{\gamma}$. With $\varphi(\cdot;\ \mu, v)$ the normal density of mean μ and variance v, the Savage–Dickey ratio $\ln \pi(0) - \ln \varphi(0;\ \tilde{\gamma},\ \tilde{\Sigma}_{\gamma\gamma})$ is then

$$\ln B_{10} = \frac{1}{2} \ln \frac{\Sigma_{\gamma\gamma}}{\Sigma_{\gamma\gamma} + s_\gamma^2} + \frac{\hat{\gamma}^2}{2} \left(\frac{1}{\Sigma_{\gamma\gamma}} - \frac{1}{\Sigma_{\gamma\gamma} + s_\gamma^2} \right). \tag{12.31}$$

The first term is negative and is the price of the prior's spread: the trend model spread its prediction over trends of size s_γ, most of which the record rules out. The second is positive and is the reward for the trend the record shows, at most half the square of the normalized entry, $\frac{1}{2}(\hat{\gamma}/\sqrt{\Sigma_{\gamma\gamma}})^2$, which it approaches as the prior widens.

Why the evidence ratio depends on the prior. The evidence of a model is the probability it assigned to the record before seeing it, averaged over its prior. A trend model with a very wide prior predicted large trends as readily as the modest one observed; it spread its prediction thin and earns less for the part that came true than the constant rate earns for being close. As s_γ grows, the first term of equation (12.31) falls like $-\ln(s_\gamma/\sqrt{\Sigma_{\gamma\gamma}})$ while the second approaches its limit, so $\ln B_{10}$ falls without bound; as s_γ shrinks to zero, the trend model becomes the constant rate and $\ln B_{10} \to 0$. Between, it has a maximum where the derivative of equation (12.31) in s_γ^2 vanishes,

$$-\frac{1}{2(\Sigma_{\gamma\gamma} + s_\gamma^2)} + \frac{\hat{\gamma}^2}{2(\Sigma_{\gamma\gamma} + s_\gamma^2)^2} = 0 \quad \Longrightarrow \quad s_\gamma^2 = \hat{\gamma}^2 - \Sigma_{\gamma\gamma},$$

with the value $\ln B_{10} = \frac{1}{2}(z^2 - 1) - \ln z$, where $z = \hat{\gamma}/\sqrt{\Sigma_{\gamma\gamma}}$, when $z > 1$. Under the Laplace form no prior centered on zero gives the trend more than that bound.

Why the sign probability hardly depends on it. The probability $\Pr(\gamma > 0 \mid \mathbf{y})$ is computed inside the trend model, after its posterior is normalized, as $\Phi(\tilde{\gamma}/\sqrt{\tilde{\Sigma}_{\gamma\gamma}})$; the prior's spread affects it only through the shift of the posterior (12.30) toward zero, which is small once s_γ exceeds $\sqrt{\Sigma_{\gamma\gamma}}$ several times, and as $s_\gamma \to \infty$ it is the flat-prior value (12.20).

Which question each answers. The sign probability answers *given that the rate follows a log-linear trend, how sure is it that the trend is upward.* The Bayes factor answers *does the record need a trend at all, against a constant rate*, and with even prior odds it gives the trend model the posterior probability $B_{10}/(1 + B_{10})$. The Bayes factor is not itself the weight of either model: it turns prior odds into posterior odds.

Two classical summaries. The likelihood-ratio statistic compares the two models through the largest value each gives the log-likelihood: twice the difference between the objective of the constant-rate fit and that of the trend fit, which under the constant rate is approximately chi-squared with one degree of freedom, a standard result. The square of the normalized entry, $(\hat{\gamma}/\sqrt{\Sigma_{\gamma\gamma}})^2$, is the same statistic in its Wald form, read off the fit without refitting the constant rate. As in Chapter 7, neither is a probability of either model.

Building blocks. The trend entry $\hat{\gamma}$, its variance $\Sigma_{\gamma\gamma}$ and the information $\mathbf{\Lambda}$ (Section 12.2.1); the prior row of spread s_γ; the objectives of the two fits for the likelihood-ratio statistic; and, for the exact value beside them, the exact marginal of the trend (Section 12.2.2) under the same prior.

The solve at the reported spread. The spread $s_\gamma = 0.3\,\mathrm{yr}^{-1}$ of Section 12.1.5 adds $1/0.09 = 11.11$ to the trend entry of the information, and the solve (12.30) gives

$$\boldsymbol{\Lambda} + \boldsymbol{\Lambda}_\pi = \begin{bmatrix} 52 & 31 \\ 31 & 167.35 \end{bmatrix}, \qquad \tilde{\boldsymbol{\Sigma}} = \begin{bmatrix} 0.021618 & -0.004005 \\ -0.004005 & 0.006718 \end{bmatrix}, \qquad \tilde{\mathbf{z}} = \begin{bmatrix} -1.5679 \\ 0.1841 \end{bmatrix}.$$

The trend's posterior is $\mathcal{N}(0.184,\ 0.082^2)$, pulled a little toward zero, and the level rises by 0.009 with it, through the negative $\Sigma_{\alpha\gamma}$. The trend entries obey the closed form: $137.75+11.11 = 148.86$ is the inverse of the entry 0.006718 the solve returned. The prior density at zero is $\pi(0) = 1/(0.3\sqrt{2\pi}) = 1.330$, the posterior density of the trend entry at zero is 0.391, and the Savage–Dickey ratio (12.29) is $B_{10} = 3.40$, $\ln B_{10} = 1.22$. Equation (12.31) gives $-1.298 + 2.522 = 1.225$, the value of the solve; the reward term can never exceed $\frac{1}{2} \times 2.335^2 = 2.73$. With even prior odds the trend model has the posterior probability $3.40/4.40 = 0.77$, and the sign probability under this prior is 0.988. The value beside it comes from the exact marginal of the trend: one-dimensional quadrature of $m(\gamma)$ against $\pi(\gamma)$, divided by $m(0)$, gives $\ln B_{10} = 1.33$, a Bayes factor of $\mathrm{e}^{1.33} = 3.8$ and at even prior odds 0.79: the state route favors the trend 0.11 nats less, the error of the quadratic expansion of the count rows.

The bound. The largest value any prior centered on zero can give is reached at $s_\gamma = (0.1989^2 - 0.007259)^{1/2} = 0.180$, where $\ln B_{10} = \frac{1}{2}(2.335^2 - 1) - \ln 2.335 = 1.378$, a Bayes factor of 4.0.

The classical summaries. The constant-rate model fits $\hat{\lambda} = 52/236$ and leaves the objective at 130.655; the trend model reaches 127.827. The likelihood-ratio statistic is twice the difference, 5.66, which a chi-squared variable with one degree of freedom exceeds with probability 0.017. Its Wald form is $2.335^2 = 5.45$.

12.3.5 Next year's count under the two models

Engineering question. How many failures next year, under the trend and under a constant rate? The spares, the crew roster and the budget of next year are set from the count of window 7, forty unit-years of exposure.

General form. A prediction of a future observation is its probability given the state averaged over the posterior, the predictive (2.13) of Chapter 2,

$$p(y_{\mathrm{new}} \mid \mathbf{y}) = \int p(y_{\mathrm{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z},$$

with its moments the conditional moments averaged by the laws of total expectation and total variance of Chapter 3.

Specialized form. Two choices enter here. Window 7 is taken to hold forty unit-years at the trend's rate at its center, $t = 7$, with no unit retired or replaced: this makes the count Poisson with mean $40\lambda(7)$ given the state, and it is broken by a change of the fleet during the year (Section 12.6.6). The constant-rate drawing predicts with the same flat prior on α that it carries in the comparison: this keeps one drawing M_0 for both answers, so that the two predictions can be averaged with the weights the Bayes factor gives, and its posterior then has an exact conjugate form.

Under the trend. The question is equation (2.13) with $y_{\mathrm{new}} = K_7$ and $p(\mathbf{z} \mid \mathbf{y}) = \mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$. The count of window 7 is Poisson with mean $40\lambda(7)$ given the state, and the predictive averages

it over the posterior of $\ln\lambda(7)$, which is $\mathcal{N}(\ln\hat{\lambda}(7),\ V(7))$ by equations (12.14) and (12.15):

$$\Pr(K_7 = j \mid \mathbf{y}, M_1) = \int \frac{(40\,\mathrm{e}^{\omega_7})^j\,\mathrm{e}^{-40\mathrm{e}^{\omega_7}}}{j!}\ \varphi(\omega_7;\ \ln\hat{\lambda}(7),\ V(7))\,\mathrm{d}\omega_7, \tag{12.32}$$

with $\omega_7 = \ln\lambda(7)$, the log-rate of the book's notation at $t = 7$ with a reference rate of one per unit-year, a one-dimensional integral evaluated by the Gauss–Hermite quadrature of Chapter 9. Its mean follows without quadrature, from the mean of a Poisson count and the mean of a log-normal rate,

$$\mathbb{E}[K_7 \mid \mathbf{y}, M_1] = 40\,\mathbb{E}[\lambda(7) \mid \mathbf{y}] = 40\,\mathrm{e}^{\ln\hat{\lambda}(7)+V(7)/2}, \tag{12.33}$$

and its variance by the law of total variance: the Poisson part, equal to the mean, plus the epistemic part, the variance of $40\lambda(7)$, which for a log-normal rate is $(40\,\mathbb{E}[\lambda(7) \mid \mathbf{y}])^2(\mathrm{e}^{V(7)} - 1)$. The value beside it comes from the exact posterior: the same average over the grid points.

Under a constant rate. The question is equation (2.13) again, under M_0. In the form of this chapter M_0 has the one-number state α, its curvature $\mathcal{L}'' = \sum_j \mu_j$ is the level entry of $\mathbf{\Lambda}$, equal to $\Lambda = \sum_j k_j$ at the MAP state, and its Laplace variance under the flat prior is $\sigma_\alpha^2 = 1/\sum_j k_j$. The predictive is equation (12.32) with the center $\hat{\alpha}_0 = \ln(\sum_j k_j/\sum_j E_j)$ and the variance σ_α^2 in place of $\ln\hat{\lambda}(7)$ and $V(7)$, and its mean is equation (12.33) with them. The value beside it comes from the exact posterior. The flat prior on $\alpha = \ln\lambda$ is the density $1/\lambda$ on the rate, the gamma prior with shape and rate parameter zero, so the exact posterior is the gamma posterior of the pooled counts (6.32) with $a = b = 0$, Gamma$(\sum_j k_j,\ \sum_j E_j)$. Next year's count is then the negative binomial of Chapter 11, derived in Chapter 6 (equation (6.54)), with the shape $\sum_j k_j$ and the success probability $\sum_j E_j/(\sum_j E_j + E_7)$.

What the difference is worth. The two drawings' predictions can be averaged with their posterior probabilities from the comparison of Section 12.3.4. At even prior odds $\Pr(M_1 \mid \mathbf{y}) = B_{10}/(1 + B_{10})$, and the probability of a count of at least j next year is

$$\Pr(K_7 \geq j \mid \mathbf{y}) = \Pr(M_1 \mid \mathbf{y})\ \Pr(K_7 \geq j \mid \mathbf{y}, M_1) + \Pr(M_0 \mid \mathbf{y})\ \Pr(K_7 \geq j \mid \mathbf{y}, M_0).$$

Building blocks. The center $\ln\hat{\lambda}(7)$ and the variance $V(7)$ (Section 12.3.1); the one-number fit of the constant rate, $\hat{\alpha}_0$ and σ_α^2; a Gauss–Hermite rule of sixty nodes; the posterior probabilities of the two drawings from the Bayes factor; and, for the exact values beside them, the grid and the gamma posterior of the constant rate with its negative binomial.

Under the trend. The posterior of $\ln\lambda(7)$ is $\mathcal{N}(-0.8805,\ 0.2836^2)$, and equation (12.32) with sixty Gauss–Hermite nodes gives the predictive. By equation (12.33) its mean is

$$\mathbb{E}[K_7 \mid \mathbf{y}, M_1] = 40 \times 0.4146 \times 1.041 = 17.26;$$

the Poisson part of the variance is 17.26 and the epistemic part 24.96, so the predictive standard deviation is $\sqrt{42.2} = 6.50$ and the central 95 % range [7, 32] failures, against [9, 25] for the plug-in Poisson at the median rate, $40 \times 0.4146 = 16.58$: the uncertainty of the rate contributes more of the spread than the Poisson scatter does. The log-rate at $t = 7$ has variance $V(7) = 0.080$, four times the $1/52 = 0.019$ at the record's center of information, because it lies 2.9 years beyond t_w. The value beside it comes from the exact posterior on the grid: the mean 17.06, the standard deviation 6.39 and the range [7, 31], a mean 0.20 lower.

Under a constant rate. The one-number fit gives $\hat{\alpha}_0 = \ln(52/236) = -1.5126$ and $\sigma_\alpha^2 = 1/52 = 0.019231$, a log-rate of standard deviation 0.1387, and the sum over sixty nodes gives next year's count the mean $8.81\,\mathrm{e}^{1/104} = 8.90$, the standard deviation 3.23 and the central 95 % range [3, 16]. The value beside it comes from the exact posterior. The rate's posterior is

Gamma(52, 236 unit-yr), of mean 0.2203, and next year's count has mean $52 \times 40/236 = 8.81$, standard deviation 3.21 and central 95 % range [3, 16], a mean 0.08 lower and the same range. Its epistemic variance is only 1.5 against its Poisson part of 8.8: forty unit-years of future exposure against a record six times as long leave little to the rate's uncertainty unless the rate is extrapolated.

Figure 12.5 draws the two predictive distributions. The horizontal axis is the number of failures in window 7, the vertical axis its predictive probability; the blue bars are the trend model's predictive (12.32) and the gold bars the constant rate's, both on the state route, and the green lines the exact predictives, on the grid and the negative binomial. The gold distribution is compact, peaking at eight; the blue one is centered near seventeen and spread twice as wide, with a long upper tail; each green line lies close to its bars. Table 12.7 gives the numbers. Each row is a predictive distribution; the columns give its mean, standard deviation, central 95 % range and the probability of seventeen or more failures, the level at which the trend model's prediction is about even.

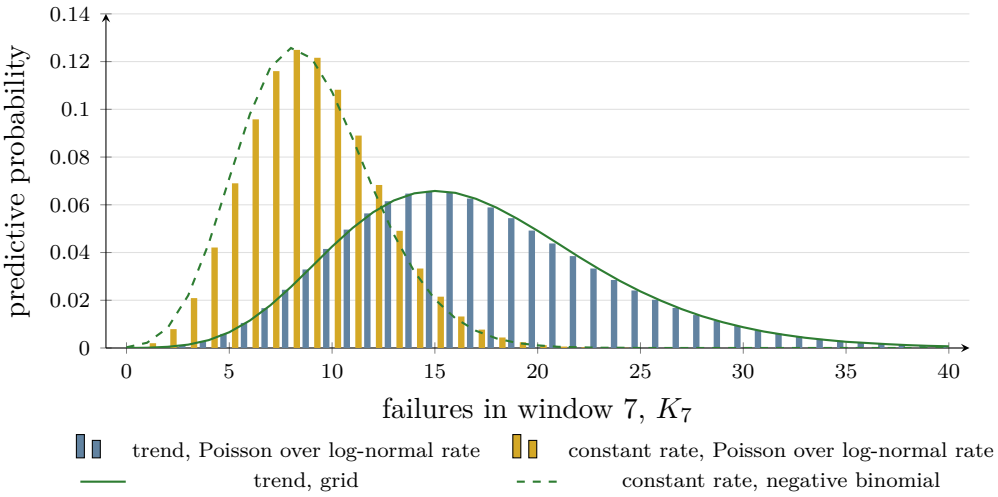


Figure 12.5: Predictive distributions of next year's count, window 7 with $E_7 = 40$ unit-years. Blue: equation (12.32), the Poisson with mean $40e^{\omega_7}$ averaged over $\omega_7 \sim \mathcal{N}(-0.8805,\ 0.2836^2)$; mean 17.26, range [7, 32]. Gold: the same with the log-rate $\mathcal{N}(-1.5126,\ 1/52)$; mean 8.90, range [3, 16]. Green, solid: the trend model's predictive on the grid; mean 17.06, range [7, 31]. Green, dashed: the negative binomial with shape 52 and success probability 236/276 from the posterior Gamma(52, 236 unit-yr) of the constant rate under the flat prior on α; mean 8.81, range [3, 16].

Table 12.7: Next year's failure count under the trend and under a constant rate. Ranges are central 95 % predictive ranges.

Predictive	Mean	Std. dev.	95 % range	$\Pr(K_7 \geq 17 \mid \mathbf{y})$
Trend, equation (12.32)	17.26	6.50	[7, 32]	0.501
Trend, grid, the exact value beside it	17.06	6.39	[7, 31]	0.492
Trend, plug-in Poisson at $\hat{\lambda}(7)$	16.58	4.07	[9, 25]	0.492
Constant rate, the same sum on $(\hat{\alpha}_0, \sigma_\alpha^2)$	8.90	3.23	[3, 16]	0.017
Constant rate, negative binomial, the exact value beside it	8.81	3.21	[3, 16]	0.015

What the difference is worth. The two drawings disagree by a factor of two about next year. Their 95 % ranges overlap only between seven and sixteen failures, the probability mass the two distributions share is 0.38 (0.37 with the exact negative binomial), and seventeen failures or more, about even under the trend, has probability 0.017 (0.015) under the constant rate. At even prior odds the Bayes factor of 3.4 gives the two drawings the posterior probabilities 0.77 and 0.23, and averaged with those weights the probability of seventeen or more failures next year is $0.77 \times 0.501 + 0.23 \times 0.017 = 0.39$. The exact routes, the weights 0.79 and 0.21 of the exact Bayes factor with the grid and the negative binomial, give $0.79 \times 0.492 + 0.21 \times 0.015 = 0.39$: no gap at this precision.

12.4 The checks (step 5)

The fleet's model has no physics row and no exact constraint, so no limit $\sigma_p \to 0$ of a model tolerance arises and no limit table is needed; the positivity of the rate is not a row either, and the log-rate state handles it.

12.4.1 The exact posterior against the Laplace posterior

The tails on the two routes. The reported tails are the normal tails of equations (12.20) and (12.21), and the grid sums set beside them differ as Table 12.9 collects: for the trend, 0.990 against the grid's 0.992; at $t = 8$, 0.513 against 0.509. Now and next year the Laplace tails are 0.035 and 0.254, against the grid's 0.026 and 0.246, and at $t = 9$ and 10 they are 0.684 and 0.784, against 0.681 and 0.782. The Laplace tail is the larger at every time, because the grid puts the rate a little lower: its median rate is 0.337 now and 0.412 next year, against the Laplace 0.340 and 0.415.

The rest of the grid posterior. The grid gives the posterior means $(-1.598,\ 0.2015)$ against the MAP state $(-1.577,\ 0.1989)$, standard deviations 0.149 and 0.0858 against 0.148 and 0.0852, and the correlation -0.35 against -0.34: the exact posterior is slightly skewed toward larger trends. The rate now has the grid band $[0.216,\ 0.501]$ and next year $[0.230,\ 0.704]$, against the Laplace $[0.224,\ 0.516]$ and $[0.238,\ 0.723]$; two years from now the grid gives the median 0.504 and the band $[0.242,\ 1.003]$, against 0.506 and $[0.250,\ 1.025]$. The crossing time has on the grid the median 7.96, the band $[5.99,\ 27.04]$, the probability 0.0085 of never crossing and the probability 0.681 of a crossing before $t = 9$; and next year's count under the trend has the mean 17.06, the standard deviation 6.39, the range $[7,\ 31]$ and the probability 0.492 of seventeen or more, against 17.26, 6.50, $[7,\ 32]$ and 0.501.

The exact marginal of the trend. One-dimensional quadrature of $m(\gamma)$ against $\pi(\gamma)$, divided by $m(0)$, gives $\ln B_{10} = 1.23$, 1.33 and 0.35 for $s_\gamma = 0.1$, 0.3 and 1.0, the values the grid gives, listed with the other prior spreads in Section 12.4.3; the same marginal under a flat prior on γ gives $\Pr(\gamma > 0 \mid \mathbf{y}) = 0.992$, the grid's value. For the reported spread $s_\gamma = 0.3$ the exact Bayes factor is $\mathrm{e}^{1.33} = 3.8$, and at even prior odds the trend model's posterior probability is 0.79, against the Laplace 0.77.

The Laplace route differs from the grid by at most 0.009 in a probability, the probability that the rate is above the limit now; by 0.1 nats in the log Bayes factor; by 0.09 years at the lower end of the crossing band and 1.3 years at the upper end; and by 0.2 failures in next year's expected count. None of these differences changes a decision. The table at the end of Section 12.4.2 collects these values with the draws.

12.4.2 The delta band against the inversion and the draws

The delta band. At the MAP state the gradient (12.26) is $\mathbf{g}_* = (-5.027,\ -22.33)^\mathsf{T}$ years per unit of each entry, and the variance budget of equation (4.50) is

$$\begin{aligned}\mathbf{g}_*^\mathsf{T}\boldsymbol{\Sigma}\,\mathbf{g}_* &= \underbrace{5.027^2 \times 0.021811}_{\alpha:\ 0.551} + \underbrace{22.33^2 \times 0.007259}_{\gamma:\ 3.620} + \underbrace{2 \times 5.027 \times 22.33 \times (-0.004328)}_{\text{cross:}\ -0.972} \\ &= 3.199\ \text{yr}^2,\end{aligned}$$

a standard deviation of 1.79 years and, by equation (12.27), the delta band

$$t^* \in 7.94 \pm 1.96 \times 1.79 = [4.44,\ 11.45].$$

The trend's term is 113 % of the total, the level's 17 %, and the negative correlation removes 30 %: the crossing time is uncertain because the trend is. The band is symmetric about 7.94 and says nothing of a rate that never crosses.

The draws. Two hundred thousand states were drawn from $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ with the covariance (12.12), the pseudo-random generator started from the fixed seed 12 so that the numbers can be reproduced, and each was mapped through equation (12.23). The draws give:

- 0.98 % of the draws have $\gamma \le 0$ and never cross;
- among the draws that cross, the median crossing time is 7.92 and the central 95 % band is [5.90, 21.96];
- over all draws, with a draw that never crosses counted as an infinitely late crossing, the median is 7.94 and the central 95 % band [5.90, 28.35];
- the probability of a crossing before $t = 9$, the middle of window 9, $\Pr(t^* < 9 \mid \mathbf{y})$, is 0.683, or 0.690 among the draws that cross; before window 9 begins, at $t = 8.5$, it is 0.61;
- the sample standard deviation of the crossing times of the draws that cross is 239 years.

Two million draws from the same seed give 0.97 %, a band of [5.90, 21.74] among the crossing draws and [5.90, 28.36] over all, and a sample standard deviation of 688 years. Every quantile is stable to the second decimal; the standard deviation is not, and grows with the number of draws, as a quantity with no variance must.

Figure 12.6 draws the distribution. The horizontal axis is the crossing time in years, the vertical axis the probability density per year. The shaded histogram holds the two hundred thousand draws in half-year bins, normalized by the total number of draws so that the draws that never cross, 1.0 %, and those beyond 26 years, 1.8 %, are missing from its area. The thin black curve is the density of the inversion, the derivative of equation (12.24) in s; the dashed curve the delta-method normal $\mathcal{N}(7.94,\ 1.79^2)$. The two bars below the axis are the two 95 % bands. The histogram rises steeply from nothing at five years to its peak near seven and then decays slowly over decades; the curve of the inversion passes through its bins. The normal is centered at the same place but symmetric: it puts 14 % of its mass before six years, where the posterior has 3.5 %, and its upper tail is gone by twelve years, where the posterior still has one crossing in eight ahead of it.

The failure in numbers. The pole of Section 12.3.3 lies where the posterior density of γ is 0.307, and the Hessian of the crossing time at $\hat{\mathbf{z}}$ is

$$\frac{\partial^2 t^*}{\partial \mathbf{z}\,\partial \mathbf{z}^\mathsf{T}} = \begin{bmatrix} 0 & 25.27 \\ 25.27 & 224.5 \end{bmatrix},$$

indefinite, with the large positive curvature along the trend direction, the direction that carries 113 % of the delta variance.

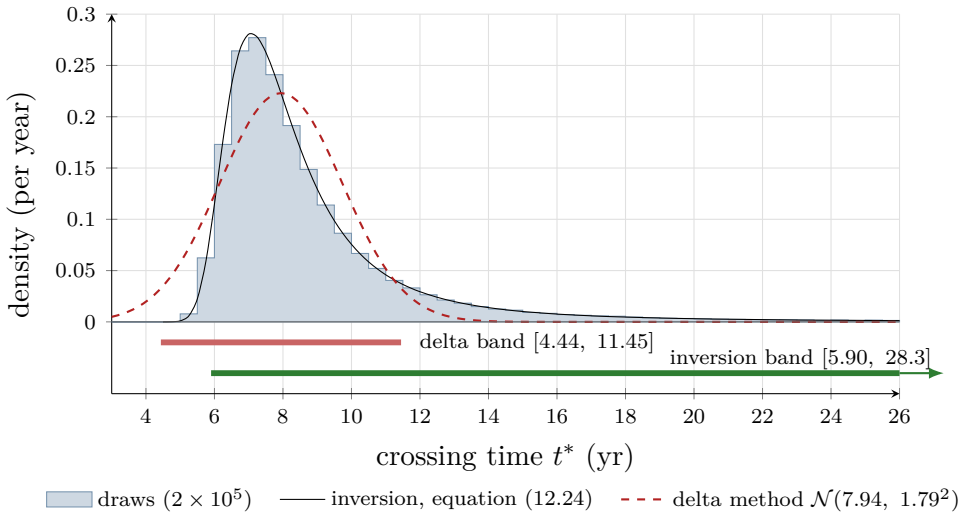


Figure 12.6: The posterior distribution of the crossing time of the limit 0.5 per unit-year. Histogram: 2×10^5 draws of the state from $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ (seed 12) mapped through equation (12.23), half-year bins, density over all draws (0.98 % never cross; 1.8 % cross after 26 years). Thin curve: the derivative of $\Phi\big((\ln \hat{\lambda}(s) - \ln 0.5)/\sqrt{V(s)}\big)$ in s, the density of the inversion. Dashed: the delta-method normal. Bars: the two 95 % bands; the band of the inversion continues to 28.3.

Table 12.8 collects the summaries of the crossing time by every route of the chapter. Each row is a route; the columns give its central value, its central 95 % band, its probability of never crossing and its probability of a crossing before $t = 9$, the middle of window 9. The draws and the inversion agree with each other to the second decimal, which is the check of the inversion's algebra; the grid posterior, the exact route, moves the lower end later by a tenth of a year and the upper end earlier by a year. The delta method agrees with none of them except at the center.

Table 12.8: The crossing time of the limit 0.5 per unit-year by five routes. A draw or grid point with $\gamma \leq 0$ never crosses; in the rows marked "all" it counts as an infinitely late crossing, in the row marked "crossing" it is removed before the quantiles are taken. The time $t = 9$ is the middle of window 9.

Route	Center (yr)	95 % band (yr)	Pr(never)	Pr($t^* < 9$)
Point estimate, equation (12.23)	7.94	–	–	–
Delta method, equation (12.27)	mean 7.94	[4.44, 11.45]	0	0.723
Draws, crossing only	median 7.92	[5.90, 21.96]	0.0098	0.690
Draws, all	median 7.94	[5.90, 28.35]	0.0098	0.683
Inversion, equations (12.24) and (12.25)	median 7.94	[5.90, 28.34]	0.0098	0.684
Grid posterior, all	median 7.96	[5.99, 27.04]	0.0085	0.681

Common mistake

Reporting the crossing time as 7.9 ± 1.8 years. The crossing time has no standard deviation; the figure is that of its tangent. The record puts the crossing between 5.9 and 28.3 years with 95 % probability, with a median of 7.9, a one-in-seven chance of no crossing before $t = 11.45$, and a one-in-a-hundred chance that the rate never reaches the limit at all. A spares plan built on the symmetric band plans for a crossing as early as $t = 4.44$, which the record excludes, and stops planning at $t = 11.45$.

Every answer by every route. Table 12.9 sets each answer of the chapter beside its full calculation. Each row is an answer; the columns give its value by the closed form under the Laplace Gaussian, the reported value, by the 2×10^5 draws where they were taken, and by the exact route, the grid or the exact marginal. The last row is the delta band of the crossing time, which no other route reproduces: it fails under both posteriors.

Table 12.9: The checks of step 5: each answer by its closed form under the Laplace posterior, by draws of that posterior (2×10^5, seed 12) and by an exact route. Rates in failures per unit-year, times in years; $s_\gamma = 0.3\,\mathrm{yr}^{-1}$ for the Bayes factor.

Answer	Closed form, Laplace	Draws	Exact route
$\Pr(\gamma > 0 \mid \mathbf{y})$ [–]	0.990	–	grid and marginal 0.992
$\ln B_{10}$ [nat]	1.22, eq. (12.31)	–	marginal 1.33
Median and band of $\lambda(6)$ [per unit-yr]	0.340, [0.224, 0.516]	–	grid 0.337, [0.216, 0.501]
Median and band of $\lambda(7)$ [per unit-yr]	0.415, [0.238, 0.723]	–	grid 0.412, [0.230, 0.704]
$\Pr(\lambda(6) > 0.5 \mid \mathbf{y})$ [–]	0.035, eq. (12.21)	–	grid 0.026
$\Pr(\lambda(7) > 0.5 \mid \mathbf{y})$ [–]	0.254, eq. (12.21)	–	grid 0.246
$\Pr(\lambda(8) > 0.5 \mid \mathbf{y})$ [–]	0.513, eq. (12.21)	–	grid 0.509
Median of t^* [yr]	7.94, eq. (12.24)	7.94	grid 7.96
95 % band of t^* [yr]	[5.90, 28.34], eq. (12.25)	[5.90, 28.35]	grid [5.99, 27.04]
Pr(never) [–]	0.0098	0.0098	grid 0.0085
$\Pr(t^* < 9 \mid \mathbf{y})$ [–]	0.684	0.683	grid 0.681
$\mathbb{E}[K_7 \mid \mathbf{y}, M_1]$ [failures]	17.26	–	grid 17.06
$\mathbb{E}[K_7 \mid \mathbf{y}, M_0]$ [failures]	8.90	–	gamma 8.81
Delta band of t^* [yr]	[4.44, 11.45], eq. (12.27)	against [5.90, 28.35]	against grid [5.99, 27.04]

12.4.3 Sensitivity to the prior spread

Table 12.10 varies the spread s_γ of the prior on the trend. Each row is a prior spread; the columns give the posterior probability of a positive trend under that prior and the log Bayes factor, each by the Laplace solve (12.30) with its closed form (12.31) and by the exact routes of Section 12.2.2, the grid posterior and the exact marginal of the trend. The first row is the flat prior of the fit, under which the probability is that of equation (12.20) and the Bayes factor is not defined.

Figure 12.7 draws the log Bayes factor against the prior spread. The horizontal axis is s_γ on a logarithmic scale from 0.01 to $10\,\mathrm{yr}^{-1}$, the vertical axis $\ln B_{10}$ in nats. The solid curve is equation (12.31); the squares are the grid values; the dashed line is the value obtained when the flat-prior posterior density at zero, 0.307, is put in the denominator of equation (12.29) in place of the posterior under the prior, the error of the box below. The solid curve starts at zero for a prior so narrow that the two models coincide, rises to its maximum of 1.38 nats at $s_\gamma = 0.18$, and falls without bound as the prior widens, through zero at $s_\gamma = 1.28$. The grid values lie about 0.1 nats above the curve, the error of the quadratic expansion of the count rows at a few tens of failures, and they put the bound near 1.47 nats.

The sign probability and the Bayes factor respond to the spread as Section 12.3.4 derives. The sign probability is 0.962 for the narrow prior 0.1, 0.988 for 0.3 and 0.990, the flat-prior value, from 1.0 on. The log Bayes factor over the same range rises from 1.15 to 1.38 at $s_\gamma = 0.18$, then falls to 1.22, 0.24 and -0.84 at $s_\gamma = 3$, and keeps falling. For spreads from 0.1 to $0.5\,\mathrm{yr}^{-1}$,

Table 12.10: The probability of a rising rate and the evidence ratio of a trend against a constant rate, by the prior spread s_γ of the trend. Laplace: equations (12.30) and (12.31) with $\hat{\gamma} = 0.1989$, $\sqrt{\Sigma_{\gamma\gamma}} = 0.0852$; grid: the exact posterior on the two entries. The spread 0.18 maximizes the Laplace log Bayes factor.

Prior spread s_γ (yr^{-1})	$\Pr(\gamma > 0 \mid \mathbf{y})$, Laplace	same, grid	$\ln B_{10}$ (nats), Laplace	same, grid
flat	0.990	0.992	–	–
0.1	0.962	0.966	1.15	1.23
0.18	0.983	–	1.38	1.47
0.3	0.988	0.989	1.22	1.33
1.0	0.990	0.991	0.24	0.35
3.0	0.990	–	−0.84	−0.73

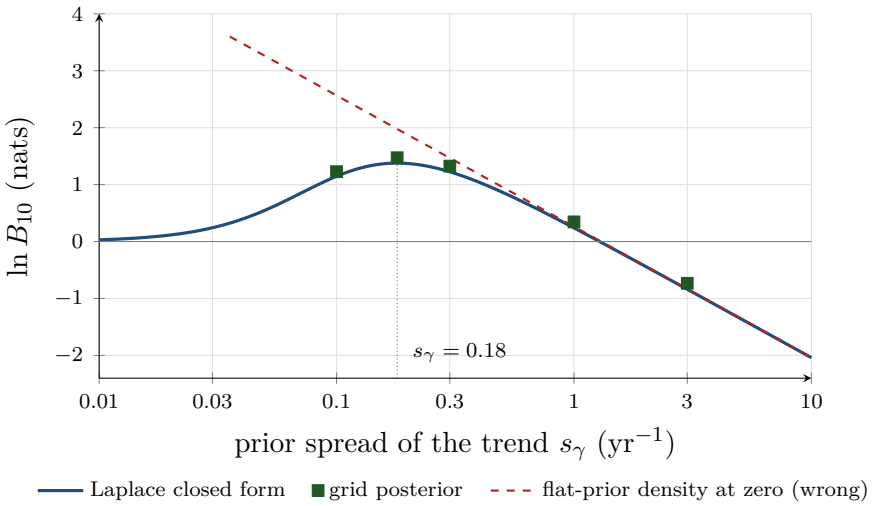


Figure 12.7: The log Bayes factor of the trend model against the constant rate as a function of the prior spread s_γ of the trend. Solid: equation (12.31) with $\hat{\gamma} = 0.1989$ and $\Sigma_{\gamma\gamma} = 0.007259$; maximum 1.378 nats at $s_\gamma = (\hat{\gamma}^2 - \Sigma_{\gamma\gamma})^{1/2} = 0.180$. Squares: the grid posterior. Dashed: $\ln(\sqrt{\Sigma_{\gamma\gamma}}/s_\gamma) + \frac{1}{2}(\hat{\gamma}/\sqrt{\Sigma_{\gamma\gamma}})^2$, the result of using the flat-prior posterior density at zero.

the plausible range for aging equipment, the record favors the trend by a factor of between two and four, a preference that does not depend on the exact choice and is never decisive.

Common mistake

Putting the posterior density of the fit with flat priors, 0.307 at $\gamma = 0$, into the Savage–Dickey ratio in place of the posterior density under the prior π: reading the trend entry of $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$ where the ratio needs that of $\tilde{\mathbf{z}}$ and $\tilde{\boldsymbol{\Sigma}}$, the solve with the prior row. The ratio needs the posterior of the model whose evidence it computes. The shortcut ignores the pull of the prior toward zero and gives $\ln B_{10} = 2.57$, 1.47 and 0.26 for $s_\gamma = 0.1$, 0.3 and 1.0 against the correct 1.15, 1.22 and 0.24, and it grows without bound as the prior narrows, while the correct value falls to zero there, as Chapter 7 found for the Weibull shape.

12.4.4 Residuals, the trend test and the window centers

The residuals of the two fits. A fitted Poisson model predicts each window's count with mean $\hat{\lambda}_j E_j$ and standard deviation $\sqrt{\hat{\lambda}_j E_j}$, so the Pearson residuals of Table 12.4 should look like six independent draws of mean zero and standard deviation about one, in no particular order. Under the trend the six residuals lie between -0.18 and 0.22, and the sum of their squares is 0.10 on four degrees of freedom. Under the constant rate they run in order from -1.04 in the first window to $+1.75$ in the sixth, and the sum of their squares is 5.95 on five degrees of freedom. That sum alone does not reject the constant rate: a chi-squared variable with five degrees of freedom exceeds 5.95 with probability 0.31. The order of the residuals does: six residuals increasing monotonically is the signature of a trend that an omnibus sum of squares cannot see, because it ignores the order. The targeted check is the trend entry itself, read off the fit as its normalized entry, $\hat{\gamma}/\sqrt{\Sigma_{\gamma\gamma}} = 0.1989/0.0852 = 2.335$: the test of zero slope, whose two-sided tail probability $2\Phi(-2.335) = 0.020$ stands beside the likelihood-ratio statistic 5.66 of Section 12.3.4 with its 0.017. The working residuals $\mathbf{r}(\hat{\mathbf{z}})$ of Section 12.2.3 are these Pearson residuals with the sign reversed, so the same check reads the residual vector of the fit directly. The trend fit is, in turn, closer than chance would make it: a chi-squared variable with four degrees of freedom falls below 0.10 with probability 0.001. That is a property of the illustrative record, not of the model; a real record scatters more, and the check would be repeated on it.

The trend test on event times. The yearly counts test the trend only between windows. When the failure times are recorded, the statistic of Chapter 6 (Section 6.4.6) tests it within them as well, the sum of the failure times scaled by the exposure, which a late cluster makes large; this record holds counts only, so the check waits for the failure dates of window 7.

The rate at the window centers. The count row of Section 12.1.5 takes the rate at each window's center in place of its average over the year. At the fitted trend the factor $\sinh(\gamma/2)/(\gamma/2)$ between the two is 1.0017, a shift of 0.0017 in α, about one percent of its posterior standard deviation, and no change in γ: the choice raises every fitted rate by the same 0.17 %, and no answer moves.

12.5 The report (step 6)

The straight line through the yearly rates read the record as a date, the limit two years away at $t = 8.1$; the record does not support a date. The median crossing time is 7.9, close to the line's, but the 95 % band of the crossing runs from 5.9 to 28.3 years, it has no upper end at all at 99 %, and the familiar band of the delta method, 7.9 ± 1.8 years, is wrong at both ends. The early end is what the contract needs, and it has already arrived: the rate is above the limit now with

probability 0.035 and next year with probability 0.254 (0.026 and 0.246 on the grid), so the contract question is one for this year, not for two years from now. Every answer is computed from the posterior of the state, with the exact value beside it, and the grid moves them by too little to change a decision.

12.5.1 The answers and their bands

Table 12.11 gives the answers in the order of Table 12.2. Each row is a question; the second column gives the answer from the posterior of the state, the third the exact value, on the grid, the exact marginal or the gamma posterior, and the last its band, with the routes it was compared with.

Table 12.11: The report for the fleet's record, from the posterior of the state with the exact value beside each answer. Bands are central 95 % posterior or predictive intervals, exact in parentheses; rates in failures per unit-year; times t in years on the scale of equation (12.1), the last window of record being centered at $t = 6$; $s_\gamma = 0.3\,\mathrm{yr}^{-1}$ for the Bayes factor.

Question	Answer (state route)	Exact	Band and comparison
What is the rate now and next year, with its band?	$\lambda(6)$: median 0.340; $\lambda(7)$: median 0.415	0.337; 0.412	$\lambda(6) \in [0.224,\ 0.516]$ ([0.216, 0.501]); $\lambda(7) \in [0.238,\ 0.723]$ ([0.230, 0.704])
Is the failure rate rising?	$\Pr(\gamma > 0 \mid \mathbf{y}) = 0.990$; trend 22 % per year	0.992	0.962 to 0.990 over the prior spreads
Will the rate be above the limit two years from now?	$\Pr(\lambda(8) > 0.5 \mid \mathbf{y}) = 0.513$	0.509	median $\lambda(8) = 0.506$, band [0.250, 1.025] ([0.242, 1.003]); now 0.035 (0.026), next year 0.254 (0.246)
When does the rate cross 0.5 per unit-year?	median 7.94	7.96	[5.90, 28.3] ([5.99, 27.04]); $\Pr(\text{never}) = 0.010$ (0.0085); $\Pr(t^* < 9) = 0.68$, and 0.61 before $t = 8.5$; delta method [4.44, 11.45], not to be used
Does the record need a trend, against a constant rate?	$\ln B_{10} = 1.22$, $B_{10} = 3.4$; trend model 0.77 at even odds	1.33, 3.8; 0.79	1.15 to 1.38 for s_γ from 0.1 to 0.3, falling to 0.24 at 1.0; likelihood-ratio statistic 5.66
How many failures next year, under the trend and under a constant rate?	trend: mean 17.3; constant: mean 8.9; $\Pr(K_7 \geq 17) = 0.39$ over both	17.1; 8.8; 0.39	trend [7, 32] ([7, 31]); constant [3, 16]; shared probability 0.38 (0.37)

The report line. The fleet's failure rate rose from about 0.13 to 0.34 failures per unit-year over six years, a trend of 22 % per year that is positive with probability 0.99 and that the record favors over a constant rate by a Bayes factor of two to four, a preference rather than a proof. Next year the fleet should expect about 17 failures, between 7 and 32, if the trend is real, and about 9, between 3 and 16, if it is not.

12.5.2 The trend and the contract

The first two answers of the trend are different answers. The sign probability says that, given a log-linear trend, the trend is upward with probability 0.99. The Bayes factor says that the record needs a trend at all by a factor of 3.4, which at even prior odds puts the posterior probability of the trend model at 0.77, not 0.99. The two are consistent: six counts make a falling trend improbable, and they are not yet many enough to make a constant rate improbable. The contract decision does not wait for the second: under the trend the rate is above the

limit next year with probability 0.254 (0.246 on the grid), and averaged over the two drawings seventeen or more failures next year, a count the contract would not cover, have probability 0.39. A spares stock, a crew roster or a budget set from the constant rate would be exceeded about half the time if the trend is real.

Common mistake

Pooling a rising record into a constant rate and predicting next year from it. The pooled rate of 0.220 per unit-year is the average of a rate that was 0.126 in the first window and 0.340 in the last; it describes no year of the record and predicts 8.9 failures next year (8.8 exactly) where the trend predicts 17.3. The constant-rate model's own check does not catch this if only the sum of squared residuals is read: the ordered residuals and the trend entry do.

12.6 The next measurement (step 7)

The report leaves the crossing time wide open at its late end, and the operator asks whether one more year of record would settle it. The record grows by forty unit-years a year whatever is decided, so the candidates are lengths of further record, valued before their counts are known by the preposterior of Chapter 2.

12.6.1 The candidate records

Four candidate lengths of record are compared, each read with a refit at its end:

- *the first half of window 7*, twenty unit-years, read at $t = 7.0$;
- *the whole of window 7*, forty unit-years, read at $t = 7.5$;
- *window 7 and the first half of window 8*, read at $t = 8.0$;
- *windows 7 and 8*, read at $t = 8.5$.

A whole window is one count row of forty unit-years with the rate at its center, $t = 7$ or 8; a half-year is one count row of twenty unit-years with the rate at its center, $t = 6.75$ or 7.75. A fifth candidate records the date of every failure in window 7 instead of its count, the same exposure binned more finely.

12.6.2 The preposterior of the trend

The first candidate to work is window 7, forty unit-years at the expected rate. Its count row has the working form of equation (12.6) with the log-count derivatives $\mathbf{g}(7)^{\mathsf{T}} = (1,\ 3.5)$ and the expected count $\hat{\mu}_7 = \hat{\lambda}(7)E_7 = 16.58$ at the MAP state, so its Jacobian row and information are

$$\mathbf{J}_a = \sqrt{16.58}\,\begin{bmatrix}1 & 3.5\end{bmatrix} = \begin{bmatrix}4.072 & 14.252\end{bmatrix}, \qquad \mathbf{\Lambda}_a = \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a = \begin{bmatrix}16.58 & 58.04\\ 58.04 & 203.13\end{bmatrix},$$

and the preposterior covariance of Chapter 2, equation (2.17), with the precisions added as in equation (3.26), is

$$\mathbf{\Sigma}_a = (\mathbf{\Sigma}^{-1} + \mathbf{\Lambda}_a)^{-1} = \begin{bmatrix}68.58 & 89.04\\ 89.04 & 359.37\end{bmatrix}^{-1} = \begin{bmatrix}0.021495 & -0.005326\\ -0.005326 & 0.004102\end{bmatrix}.$$

Read off, $\sqrt{[\mathbf{\Sigma}_a]_{\gamma\gamma}} = 0.0640$ against 0.0852 now, and $\sqrt{[\mathbf{\Sigma}_a]_{\alpha\alpha}} = 0.147$ against 0.148, almost unchanged, with the correlation -0.57. The reason is in the Jacobian row: its second entry carries the lever arm of 3.5 years from $\bar{t}$, which $\mathbf{\Lambda}_a$ squares, so the new window adds 203.1 to the trend entry of the information, now 156.2, and only 16.6 to the level entry, now 52. The trend would then be $\hat{\gamma}/\sqrt{[\mathbf{\Sigma}_a]_{\gamma\gamma}} = 3.11$ standard deviations from zero, and the probability of

never crossing would fall from 0.010 to 0.0009. The delta standard deviation of the crossing time, $\sqrt{\mathbf{g}_*^{\mathsf{T}}\boldsymbol{\Sigma}_a\mathbf{g}_*}$, falls from 1.79 to 1.18.

Table 12.12 carries the calculation to the band of the inversion. Each row is a record: the present one, one more window and two more; the columns give the standard deviation of the trend, the normalized trend entry, the probability of never crossing, the delta band and the band by the inversion (12.25) with $\boldsymbol{\Sigma}_a$ in place of $\boldsymbol{\Sigma}$, and the skew of that band, the distance from the median to its upper end over the distance to its lower end. Two more windows stack two Jacobian rows, $\sqrt{16.58}\,(1,\ 3.5)$ and $\sqrt{20.23}\,(1,\ 4.5)$, the second with the expected count of window 8, and give $[\boldsymbol{\Sigma}_a]_{\gamma\gamma} = 0.002476$. In every row the center is held at the present MAP state; the counts actually recorded will move it.

Table 12.12: The preposterior of the crossing time for one and two more yearly windows of forty unit-years at the expected rate, with the center held at $\hat{t}^* = 7.94$. Band by inversion: equation (12.25) with $\boldsymbol{\Sigma}_a$. Skew: $(\text{upper} - 7.94)/(7.94 - \text{lower})$ of that band.

Record	$\sqrt{\Sigma_{\gamma\gamma}}$ (yr^{-1})	$\hat{\gamma}/\sqrt{\Sigma_{\gamma\gamma}}$	Pr(never)	Delta band (yr)	Band by inversion (yr)	Skew
present, six windows	0.0852	2.33	0.0098	[4.44, 11.45]	[5.90, 28.34]	10.0
one more window	0.0640	3.11	0.0009	[5.63, 10.26]	[6.39, 13.66]	3.7
two more windows	0.0498	4.00	0.00003	[6.36, 9.52]	[6.74, 10.67]	2.3

The delta band shrinks by the same third at both ends, because it is symmetric and scales with one standard deviation. The band by inversion does not. Its lower end moves from 5.90 to 6.39, a quarter of its distance from the median removed; its upper end moves from 28.3 to 13.7, almost three quarters removed. The added window pushes the posterior of γ away from the pole, and the pole is what made the upper end long: the gain of one more year is mostly at the late end of the crossing band, the end that a decision to renegotiate or to replace does not need as urgently as the early one. One more year narrows the band from 22 years wide to 7; it does not settle the crossing time, which two more years would bring to [6.7, 10.7], still skewed. The expected count 16.6 used here is itself uncertain; refitting the seven windows at each possible count of window 7 and averaging the trend entry of the refit covariance over the predictive distribution of that count, as Chapter 8 averaged its preposterior, gives $\sqrt{0.004118} = 0.0642$ instead of 0.0640, no change in the conclusion.

12.6.3 The length of record and the reading time

The rows of each candidate of Section 12.6.1 are stacked, their informations added, and $\boldsymbol{\Sigma}_a$ is the inverse of $\boldsymbol{\Sigma}^{-1}$ plus that sum. Table 12.13 lists the four. Each row is a candidate; the columns give the trend entry $[\boldsymbol{\Lambda}_a]_{\gamma\gamma}$ it adds to the information, the preposterior standard deviation of the trend, the band of the crossing time by the inversion (12.25) with $\boldsymbol{\Sigma}_a$, the width it removes from the present band, 22.44 years wide, and the probability, under the present posterior of the state, that the rate has crossed the limit before the reading is taken, the tail (12.21) at t_{read}, with the grid sum (12.22) in parentheses.

The trend information grows faster than the length of record, because each later half-year sits further from $\bar{t}$ and its lever arm enters squared: a half-year adds 83, a year 203, two years 613. The width removed grows more slowly, 11.3, 15.2, 17.3 and 18.5 years, since the first year already takes most of the long upper end away from the pole. The lower end, the planning date, moves later by only 0.26 to 0.84 years, and by the time the reading is taken the crossing has already happened with probability 0.25 to 0.61, on both routes. Two further checks bound

Table 12.13: Candidate lengths of further record at forty unit-years a year, with the center held at the present MAP state. Band by inversion: equation (12.25) with $\mathbf{\Sigma}_a$; width removed: from the present band [5.90, 28.34]; last column: $\Pr(t^* \le t_{\text{read}} \mid \mathbf{y}) = \Pr(\lambda(t_{\text{read}}) > 0.5 \mid \mathbf{y})$ by equation (12.21), with the grid sum (12.22) in parentheses.

Candidate, read at t_{read}	$[\mathbf{\Lambda}_a]_{\gamma\gamma}$ [yr^2]	$\sqrt{[\mathbf{\Sigma}_a]_{\gamma\gamma}}$ [yr^{-1}]	Band by inversion [yr]	Width removed [yr]	Crossed before reading [–]
half a year, $t_{\text{read}} = 7.0$	83.3	0.0733	[6.16, 17.35]	11.25	0.254 (0.246)
one year, $t_{\text{read}} = 7.5$	203.1	0.0640	[6.39, 13.66]	15.16	0.392 (0.387)
eighteen months, $t_{\text{read}} = 8.0$	377.0	0.0561	[6.58, 11.76]	17.26	0.513 (0.509)
two years, $t_{\text{read}} = 8.5$	612.8	0.0498	[6.74, 10.67]	18.50	0.609 (0.606)

the choice. Recording the date of every failure in window 7 instead of its count, the expected information of the event times $\int 40\,\hat{\lambda}(t)\,\mathbf{g}(t)\mathbf{g}(t)^{\mathsf{T}}\mathrm{d}t$ over the window, adds 206.8 to the trend entry against 203.1 for the count, and the band becomes [6.39, 13.57]: binning the year more finely buys nothing, and the length of record is what counts. And the count itself decides more than its expected value: refitting the seven windows at $K_7 = 25$ gives the band [5.82, 8.52], at $K_7 = 17$ the band [6.35, 13.13], and at nine failures or fewer, which the trend predictive of equation (12.32) gives probability 0.095, the trend falls below 1.96 standard deviations and the band again has no upper end.

12.6.4 The recommendation

Record window 7 in full and refit at its end, $t = 7.5$, with an interim refit on its first half at $t = 7.0$. The full year takes the standard deviation of the trend from 0.085 to 0.064 per year, removes two thirds of the band's width, 15.2 of 22.4 years, and cuts the probability of never crossing tenfold; a second year removes only 3.3 years more, and arrives after the median crossing 7.94 has passed. The interim reading costs nothing, since it is the first half of the same record, and already removes half the width. Neither reading moves the early end of the band by more than half a year, so the contract question is not deferred to them: the renegotiation or the replacement plan starts this year on the present band, and the refit at $t = 7.5$ dates the late end. Logging the failure mode at each repair adds the evidence that the counts lack, the mechanism behind the trend. When the count of window 7 arrives, the fit returns to step 4 with its row; Chapter 20, in Volume II, writes the same preposterior for a candidate row as a rank-one update of $\mathbf{\Sigma}$.

12.6.5 What the analysis showed

The line through the six yearly rates gave the operator a date, the limit about two years away, and with it the sense that the contract could be renegotiated next year; the analysis replaces the date by a distribution that starts now. The fleet is aging: the rise of about a fifth a year is real within any model that allows a trend, and the record already favors a trend over a constant rate, though not yet decisively. What the six years cannot fix is how fast the fleet ages, and the crossing date divides by exactly that: a slope the record cannot rule out puts the crossing decades away or never, a slope it also cannot rule out puts it within the present year. That is why the early end of the band, the end the contract depends on, is already here, and why the familiar plus-or-minus band, which plans for an impossible early crossing and stops planning far too soon, must not be used. For the next year's spares the trend roughly doubles the count a pooled rate would predict. One more year of counts mostly trims the late end of the band, not

the early one, so it dates how long the problem lasts rather than when it starts; the decision to act belongs to this year, and the failure modes logged at repair are the evidence that will say why the fleet is aging.

12.6.6 Limits of the model

The derivation is exact on the logarithmic scale under the assumptions of Section 12.1.2 and the log-linear line. That does not make the assumptions appropriate, and three groups of them bound the answers.

Operating conditions. The rate of a cooling unit depends on how hard it works as well as on how old it is. A covariate enters the log-rate as a third term,

$$\ln \lambda_j = \alpha + \gamma x_j + \zeta\,(\chi_j - \chi_{\mathrm{ref}}),$$

with χ_j the load fraction of window j, the average load of the units as a fraction of their rated capacity, χ_{ref} a reference fraction, and ζ a new entry of the state, the change of the log-rate per unit of load fraction. The derivation of Section 12.2.1 carries over with the row $(1,\ x_j,\ \chi_j - \chi_{\mathrm{ref}})$ of log-count derivatives: the Jacobian $\mathbf{J}$ of equation (12.6) gains a third column, the load deviations scaled by $\sqrt{\mu_j}$, and the information matrix $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ becomes 3×3. That column adds information on ζ only to the extent that the load varies, and it separates aging from loading only to the extent that the load does not move in step with time. If the load rose by a fixed amount each year, the third column would be a multiple of the second plus a multiple of the first, the information matrix singular, and the record would fix only the combination of γ and ζ that the joint movement determines: a rate rising because the units age and a rate rising because the halls run hotter would be the same record. Units in the same window at different loads, or a load that rose and fell, separate them. The record of this fleet carries no load, so the chapter gives no numbers for ζ; Exercise 12.6 works the same structure with an age group in place of the load.

The physics of aging. A log-linear trend describes the rise; it does not explain it. The companion volume's chapters on degradation describe the mechanisms behind a rising rate in equipment of this kind, wear of compressors and fans, fouling of coils, fatigue of joints and the drift of controls, and a physical model of one of them would replace the straight line by a curve with physical entries of its own. The fleet's rise could also come from repairs that restore less each time, from a change of contractor, or from heavier duty; the counts alone cannot tell these apart, and the failure modes recorded at repair are the evidence that can.

The power law as the alternative. The power-law intensity of Chapter 6 has $\ln \nu(t) = \ln(\theta\beta) + (\beta - 1)\ln t$: it is log-linear in $\ln t$, where the model of this chapter is log-linear in t. Over the six years of the record the two can describe the same rise; beyond it the power law grows ever more slowly in proportion and the log-linear trend by a fixed factor per year, so the crossing time of a limit well beyond the record depends on the choice. Which of the two the record supports is a comparison of evidences, the subject of a later chapter on model selection; while it is undecided, the forecast beyond a few years should be reported under both.

Other limits. Six more assumptions bound the answers. The units are taken as identical; if some are worse than others, the counts scatter more than Poisson and the bands derived here are too narrow. Failures are taken as independent; a common cause striking several units at once makes the counts cluster and the bands too narrow. A repaired unit is taken to return to the shared rate; repairs that restore less make a unit's next failure come sooner and the trend

partly a history of repairs. The trend is taken as log-linear forever; the upper end of the crossing band, 28.3 years, belongs to states with a trend near zero and says that the record cannot exclude a crossing decades away, not that the fleet will exist then: at that time the median line of Figure 12.4(b) stands at 29 failures per unit-year, which no fleet would be kept running to see. The rate within a window is taken at the window's center, a choice that Section 12.4.4 showed raises every fitted rate by the same 0.17 % and leaves γ unchanged. And the exposures are pinned; units retired or replaced during the record change them and must be entered as such.

The next chapter begins Part IV of the book: it combines the reliabilities of components, estimated in Parts II and III, into the reliability of systems built from them.

12.7 What to retain

Definition 12.1 (Log-linear trend in a fleet failure rate)**.** A fleet's failure rate that changes in calendar time is written $\ln \lambda(t) = \alpha + \gamma(t - \bar{t})$, with α the log-rate at the record's mid-time and γ the trend per year, and each window of record is a Poisson count row with mean $\mu_j = \lambda_j E_j$, entering the normal equations through its working residual $(\mu_j - k_j)/\sqrt{\mu_j}$ and its Jacobian row $\sqrt{\mu_j}\,(1,\ t_j - \bar{t})$. The information matrix of the six rows is $\mathbf{J}^{\mathsf{T}}\mathbf{J}$, exact because the log-rate is the canonical link, and $\mathbf{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1}$. Moving the time origin changes the time column of $\mathbf{J}$; at the mean failure time t_w the two columns are orthogonal and the entries uncorrelated. The log-rate at a time t is linear in the state, with variance $\mathbf{g}(t)^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}(t) = 1/\sum_j k_j + \Sigma_{\gamma\gamma}(t - t_w)^2$. The crossing time of a limit, $\bar{t} + (\ln \lambda_{\max} - \alpha)/\gamma$, has the distribution function $\Phi((\ln \hat{\lambda}(s) - \ln \lambda_{\max})/\sqrt{V(s)})$, whose quantiles are the roots of Fieller's quadratic form $\mathbf{g}(s)^{\mathsf{T}}(\mathbf{d}\mathbf{d}^{\mathsf{T}} - c^2\mathbf{\Sigma})\mathbf{g}(s) = 0$, and no mean. Every answer is computed from $(\hat{\mathbf{z}}, \mathbf{\Sigma})$; a grid on the two entries and the exact marginal of the trend give the value beside it, with the gap.

Three conclusions carry beyond the fleet. Whether a rate is rising has two answers, a sign probability within the trend model and an evidence ratio between models, and only the second answers whether the trend belongs in the model; it is read from the trend entry of the fit with the prior appended as one more row, depends on the prior's spread, and never exceeds a bound set by the normalized trend entry. A forecast of a rate is exact on the logarithmic scale and widens with the square of its lever arm from the record's center of information. And a derived quantity with an uncertain denominator, the crossing time here, must be computed through an event that is linear in the state, draws, or a grid, and one more year of record shortens mostly the late end of its band.

Table 12.14 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 12.14: The terms of trend estimation and crossing-time prediction and where each is defined.

Term	Expression	Units	Meaning	Defined in
Log-linear trend	$\ln\lambda(t)=\alpha+\gamma(t-\bar t)$	–	state (α,γ)	12.1.2
Count row	$K_j\sim\text{Poisson}(\mu_j)$, $\mu_j=\lambda_j E_j$; $\mathcal{L}_j=\mu_j-k_j\ln\mu_j$	–	one per window	12.1.5
Working row	$r_j=(\mu_j-k_j)/\sqrt{\mu_j}$, $\mathbf{J}_j=\sqrt{\mu_j}\,(1,\ x_j)$	–	6×2 Jacobian	12.1.5
Rate ratio, doubling time	e^{γ}; $\ln 2/\gamma$	–; yr	growth per year	12.1.5
Score equations	$\mathbf{J}^{\mathsf T}\mathbf{r}=\mathbf{0}$: $\sum\hat\mu_j=\sum k_j$, $\sum x_j\hat\mu_j=\sum x_j k_j$	–	total and timing	12.1.5
Information matrix	$\mathbf{J}^{\mathsf T}\mathbf{J}=\mathbf{X}^{\mathsf T}\mathbf{W}\mathbf{X}$, $\mathbf{\Sigma}=(\mathbf{J}^{\mathsf T}\mathbf{J})^{-1}$	–	exact Hessian	12.2.1
Mean failure time	$t_w=\bar t+\sum x_j k_j/\sum k_j$	yr	center of information	12.3.1
Change of origin	$\mathbf{J}_c=\mathbf{J}\mathbf{U}_c$; $[\mathbf{\Lambda}_c]_{12}=0$ at $c=t_w$	–	uncorrelated entries	12.3.1
Variance of the log-rate	$V(t)=1/\sum k_j+\Sigma_{\gamma\gamma}(t-t_w)^2$	–	lever arm	12.3.1
Laplace tails	$\Phi(\hat\gamma/\sqrt{\Sigma_{\gamma\gamma}})$; $\Phi((\ln\hat\lambda(t)-\ln\lambda_{\max})/\sqrt{V(t)})$	–	$\gamma>0$; rate above limit	12.3.2
Exact tails	grid sums, eq. (12.22)	–	the value beside them	12.3.2
Crossing time	$\bar t+(\ln\lambda_{\max}-\alpha)/\gamma$, or ∞	yr	no mean	12.3.3
Fieller's quadratic	$\mathbf{g}(s)^{\mathsf T}(\mathbf{d}\mathbf{d}^{\mathsf T}-c^2\mathbf{\Sigma})\mathbf{g}(s)=0$, eq. (12.25)	yr	band ends	12.3.3
Evidence ratio	prior row γ/s_γ; eqs. (12.30), (12.31)	nat	trend against constant	12.3.4
Trend predictive count	Poisson over log-normal rate, eq. (12.32)	–	next window	12.3.5
Exact marginal of the trend	$m(\gamma)\propto \mathrm{e}^{\gamma\sum k_j x_j}/S(\gamma)^{\sum k_j}$	–	exact evidence	12.2.2
Preposterior of a window	$\mathbf{J}_a=\sqrt{\hat\mu}\,\mathbf{g}^{\mathsf T}$, $\mathbf{\Lambda}_a=\mathbf{J}_a^{\mathsf T}\mathbf{J}_a$	–	one more year	12.6.2

12.8 Exercises

Exercise 12.1. *Formulation.* (a) For the chapter's fleet, $\hat\gamma=0.1989\,\text{yr}^{-1}$; compute the rate ratio per year, the percentage rise per year and the doubling time. (b) A second fleet has the MAP state $\hat\alpha=-1.8$ and $\hat\gamma=0.15\,\text{yr}^{-1}$ with $\bar t=3$, and its contract limit is 0.4 failures per unit-year. Write the crossing time as a function of the state and compute its point estimate. (c) State what the point estimate does not say. ((a) $\mathrm{e}^{0.1989}=1.220$, a rise of 22.0 % per year; doubling time $\ln 2/0.1989=3.48$ years. (b) $t^*=3+(\ln 0.4-\alpha)/\gamma$ for $\gamma>0$, ∞ otherwise; $\hat t^*=3+(-0.9163+1.8)/0.15=8.89$ years. (c) It is the median of the crossing time, not a mean; it says nothing of the band or of the probability that the rate never reaches the limit.)

Exercise 12.2. *Formulation.* For the second fleet of the previous exercise the Laplace covariance is $\Sigma_{\alpha\alpha}=0.03$, $\Sigma_{\alpha\gamma}=-0.005$ and $\Sigma_{\gamma\gamma}=0.01$. (a) Compute the median rate and its central 95 % band at $t=7$. (b) Compute the probability that the rate exceeds 0.4 per unit-year at $t=7$. (c) State whether the crossing band of 95 % is bounded. ((a) $\mathbf{g}(7)=(1,\ 4)^{\mathsf T}$; $\ln\hat\lambda(7)=\mathbf{g}(7)^{\mathsf T}\hat{\mathbf{z}}=-1.2$, median 0.301; $V(7)=\mathbf{g}(7)^{\mathsf T}\mathbf{\Sigma}\,\mathbf{g}(7)=0.150$, standard deviation 0.387; band $0.301\,\mathrm{e}^{\pm 0.759}=[0.141,\ 0.644]$. (b) $\Phi((-1.2+0.9163)/0.387)=\Phi(-0.733)=0.232$. (c) $\hat\gamma/\sqrt{\Sigma_{\gamma\gamma}}=0.15/0.1=1.5<1.96$, so $A<0$ in equation (12.25) and the band has no upper end:

the record cannot exclude, at 95 %, that the rate never reaches 0.4; the delta standard deviation of 3.76 years would hide this.)

Exercise 12.3. *Derivation.* (a) For the Poisson count rows $\mathcal{L}_j = \mu_j - k_j \ln \mu_j$ with $\ln \mu_j = \ln E_j + \alpha + \gamma x_j$, write the working residuals and the 6×2 Jacobian, show that $\mathbf{J}^\mathsf{T}\mathbf{r}$ is the gradient of the objective and $\mathbf{J}^\mathsf{T}\mathbf{J}$ its Hessian, and explain why the Hessian is exact rather than a Gauss–Newton approximation. (b) Move the time origin to c by $\mathbf{J}_c = \mathbf{J}\mathbf{U}_c$, show that the off-diagonal entry of $\mathbf{J}_c^\mathsf{T}\mathbf{J}_c$ vanishes at $c = t_w = \bar{t} + [\mathbf{\Lambda}]_{12}/[\mathbf{\Lambda}]_{11}$, and read the variance of $\ln \lambda(t)$ off the covariance in that origin; state when $t_w = \bar{t}$. ((a) $r_j = (\mu_j - k_j)/\sqrt{\mu_j}$ and $\mathbf{J}_j = \sqrt{\mu_j}\,(1,\ x_j)$; then $\mathbf{J}^\mathsf{T}\mathbf{r} = \sum_j (\mu_j - k_j)(1,\ x_j)^\mathsf{T}$, the gradient, and $\mathbf{J}^\mathsf{T}\mathbf{J} = \sum_j \mu_j (1,\ x_j)^\mathsf{T}(1,\ x_j)$, the Hessian, since $\partial \mu_j / \partial \mathbf{z} = \mu_j (1,\ x_j)^\mathsf{T}$; the term $(\mu_j - k_j)\, \partial^2 \ln \mu_j / \partial \mathbf{z}^2$ that the working form drops is zero because the log mean is linear in the state, the log being the canonical link of the Poisson count, so k_j does not enter the curvature and the observed information equals the expected information. (b) $[\mathbf{J}_c^\mathsf{T}\mathbf{J}_c]_{12} = \mathbf{1}^\mathsf{T}\mathbf{W}(\mathbf{x} - (c - \bar{t})\mathbf{1}) = [\mathbf{\Lambda}]_{12} - (c - \bar{t})[\mathbf{\Lambda}]_{11}$, zero at $c = t_w$; there $\mathbf{\Lambda}_{t_w}$ is diagonal, so $\mathbf{\Sigma}_{t_w} = \operatorname{diag}(1/[\mathbf{\Lambda}]_{11},\ \Sigma_{\gamma\gamma})$, the gradient of $\ln \lambda(t)$ is $(1,\ t - t_w)^\mathsf{T}$, and $V(t) = 1/[\mathbf{\Lambda}]_{11} + \Sigma_{\gamma\gamma}(t - t_w)^2$, smallest at t_w; $t_w = \bar{t}$ if and only if $[\mathbf{\Lambda}]_{12} = 0$, that is if and only if $\Sigma_{\alpha\gamma} = 0$.)

Exercise 12.4. *Derivation.* For the chapter's fleet, with $\hat{\gamma} = 0.1989$ and $\Sigma_{\gamma\gamma} = 0.007259$, take the prior $\gamma \sim \mathcal{N}(0, 0.2^2)$ under the trend model. (a) Write the prior row and its Jacobian row, form $\mathbf{\Lambda} + \mathbf{\Lambda}_\pi$ from equation (12.12), and take the Newton step (12.30) from $\hat{\mathbf{z}}$ to the posterior mean and covariance under that prior. (b) Compute the Savage–Dickey ratio B_{10} from the prior and posterior densities at zero, and check it against equation (12.31). (c) Repeat (b) for $s_\gamma = 0.5$ and state why the answer fell. ((a) $r_\pi = \gamma/0.2$, $\mathbf{J}_\pi = (0,\ 5)$, $\mathbf{\Lambda}_\pi = \operatorname{diag}(0,\ 25)$; $\mathbf{\Lambda} + \mathbf{\Lambda}_\pi = \left[\begin{smallmatrix} 52 & 31 \\ 31 & 181.23 \end{smallmatrix}\right]$; the prior row's gradient at $\hat{\mathbf{z}}$ is $(0,\ 4.973)^\mathsf{T}$, and the step gives $\tilde{\mathbf{z}} = (-1.5586,\ 0.1684)^\mathsf{T}$ and $\tilde{\mathbf{\Sigma}} = \left[\begin{smallmatrix} 0.021414 & -0.003663 \\ -0.003663 & 0.006144 \end{smallmatrix}\right]$, so $\gamma \mid \mathbf{y} \sim \mathcal{N}(0.168,\ 0.0784^2)$. (b) $\pi(0) = 1/(0.2\sqrt{2\pi}) = 1.995$; $p(0 \mid \mathbf{y}) = \varphi(0;\ 0.168,\ 0.0784^2) = 0.507$; $B_{10} = 3.94$, $\ln B_{10} = 1.37$, as equation (12.31) gives. (c) $\tilde{\Sigma}_{\gamma\gamma} = 0.007054$, $\tilde{\gamma} = 0.193$, $\ln B_{10} = 0.86$, $B_{10} = 2.4$: the wider prior spread its prediction over trends the record excludes, and the price term $\frac{1}{2}\ln(\Sigma_{\gamma\gamma}/(\Sigma_{\gamma\gamma} + s_\gamma^2))$ grew faster than the reward.)

Exercise 12.5. *Application.* A new fleet of thirty units has five yearly windows of 30 unit-years each with 12, 9, 7, 6 and 4 failures, a falling rate of early life. (a) Write the log-count derivatives of the five count rows with $\bar{t} = 3$, solve the normal equations (12.10) from the constant-rate state, and give the MAP state, the information matrix, the covariance, the standard deviations and correlation read off it, and the probability that the rate is falling. (b) Give the rate in the last window and next year with their bands. (c) The burn-in is declared over when the rate falls to 0.10 failures per unit-year; give the point estimate of that time, its delta band with $\mathbf{g}_*$ displayed, and its exact band by inversion, and comment. ((a) Rows $(1,\ x_j)$ with $x_j = -2, -1, 0, 1, 2$; from $(-1.373,\ 0)^\mathsf{T}$, where the gradient is $(0,\ 19)^\mathsf{T}$, four Newton steps give $\hat{\mathbf{z}} = (-1.438,\ -0.257)^\mathsf{T}$ (a ratio of 0.773 per year), $\mathbf{\Lambda} = \left[\begin{smallmatrix} 38 & -19 \\ -19 & 79.375 \end{smallmatrix}\right]$ and $\mathbf{\Sigma} = \left[\begin{smallmatrix} 0.029894 & 0.007156 \\ 0.007156 & 0.014311 \end{smallmatrix}\right]$; $\sqrt{\Sigma_{\alpha\alpha}} = 0.173$, $\sqrt{\Sigma_{\gamma\gamma}} = 0.120$, correlation $+0.35$, positive because $[\mathbf{\Lambda}]_{12} = \sum_j x_j \hat{\mu}_j = -19$: the expected failures lean early; $\Pr(\gamma < 0 \mid \mathbf{y}) = \Phi(2.149) = 0.984$. (b) $\lambda(5) = 0.142$, band $[0.073,\ 0.277]$; $\lambda(6) = 0.110$, band $[0.046,\ 0.265]$. (c) $\hat{t}^* = 3 + (\ln 0.1 + 1.438)/(-0.257) = 6.36$; $\mathbf{g}_* = (3.890,\ 13.078)^\mathsf{T}$ and delta $6.36 \pm 1.96 \times 1.91 = [2.63,\ 10.10]$, whose lower end lies inside the record where the rate was well above 0.10; inversion, with $\mathbf{M} = \left[\begin{smallmatrix} 0.63229 & -0.24970 \\ -0.24970 & 0.01111 \end{smallmatrix}\right]$ in equation (12.25), unchanged for a falling line, gives $[4.30,\ 46.6]$, and 1.6 % of the posterior never falls to 0.10; 2×10^5 draws give an upper end of 48.9, the far quantile of a tail this heavy being poorly fixed by draws. The probability that the rate is already below 0.10 next year is 0.42; the record cannot yet date the end of burn-in.)

Exercise 12.6. *Application.* A fleet has two age groups of twenty units each, recorded for three yearly windows ($\bar{t} = 2$): the older group with 6, 8 and 9 failures and the newer with 3, 3 and 5,

each window holding 20 unit-years per group. The model is $\ln \lambda = \alpha + \gamma(t-2) + \zeta\, \mathbf{1}_{\text{old}}$, with the indicator $\mathbf{1}_{\text{old}}$ equal to 1 for the older group and 0 for the newer. (a) Write the log-count derivatives and the Jacobian of the six count rows and state what the third column adds. (b) The fit gives $\hat{\zeta} = 0.738$ with $\sqrt{\Sigma_{\zeta\zeta}} = 0.367$, and $\hat{\gamma} = 0.222$ with $\sqrt{\Sigma_{\gamma\gamma}} = 0.213$; compute the rate ratio of the older group to the newer with its band, and the probability that the older group fails more often. (c) State what the record says about the trend, and why γ and ζ are uncorrelated here. ((a) Rows $(1,\ t_j - 2,\ \mathbf{1}_{\text{old}})$: $(1,-1,1)$, $(1,0,1)$, $(1,1,1)$, $(1,-1,0)$, $(1,0,0)$, $(1,1,0)$, each multiplied by $\sqrt{\hat{\mu}_j}$ in the Jacobian, with $\hat{\boldsymbol{\mu}} = (6.04,\ 7.54,\ 9.42,\ 2.89,\ 3.61,\ 4.51)$; the third column carries the group, and the information it adds on ζ, the squared norm of that column of $\mathbf{J}$, is the older group's expected count, 23. (b) $\mathrm{e}^{0.738} = 2.09$, band $\mathrm{e}^{0.738 \pm 1.96 \times 0.367} = [1.02,\ 4.29]$; $\Pr(\zeta > 0 \mid \mathbf{y}) = \Phi(2.01) = 0.978$. (c) $\Pr(\gamma > 0 \mid \mathbf{y}) = 0.85$: three windows do not establish a trend; both groups were observed over the same windows with the same exposures, and the model scales one group's expected counts by the constant factor e^{ζ}, so the expected-count-weighted mean time is the same in each group as in the whole; the group column of $\mathbf{J}$ is then orthogonal to the time column once the level column is projected out of both, and the covariance between γ and ζ is zero, although the entry $[\mathbf{J}^{\mathsf{T}}\mathbf{J}]_{\gamma\zeta} = 3.38$ is not, while α and ζ are strongly correlated, -0.82, since α is the newer group's level.)

Part IV

System Reliability and Dependencies

Chapter 13

Series, Parallel and k-out-of-n Systems

13.1 Services provided by several units

A data hall is cooled by four computer-room air conditioners, of which the design heat load needs three, and it is fed by one power path. The operator's design basis asks for a cooling unavailability below 10^{-3}, under nine hours a year, and the facilities team holds the fleet's inspection log: at three of a hundred checks a unit was found down. The first reading of the log is the usual one. A unit is down 3 % of the time; with three units needed and four installed the cooling is lost only when two or more are down together; and summing those configurations at 3 % gives a cooling unavailability of 5.2×10^{-3}, about 45 hours a year, five times the design basis. N+1 misses the target, and the team asks for a fifth unit.

The engineer who must approve the purchase asks something else. The 3 % rests on three findings, and the four units are one model, installed together and serviced by one crew, so whatever is uncertain about the unavailability of one unit is uncertain about all four at once. Nor is the drawing of the plant settled: at the present partial load two units may carry the hall, and the design file gives the three-unit drawing a weight of 0.7 and the two-unit drawing 0.3. The number 5.2×10^{-3} is the unavailability at one value of the unknown and under one drawing; it says nothing about how probable it is that the target is missed, or whether a purchase is the right next step. The question of this chapter is

how the availability of a service is derived from the availabilities of the units that provide it, how the uncertainty of the units' records carries into it, whether a redundant arrangement meets an unavailability target with assurance, and what is reported when the drawing itself is uncertain.

What changes from the previous chapters is the quantity asked about. The state is the units' logits, one entry per unit as in Chapter 5, and the service is a function of several entries at once, fixed by the drawing through a sum over the 2^n configurations of up and down units; the drawing itself may be uncertain.

Jargon: computer-room air conditioner (CRAC), N+1 redundancy, configuration, structure function, series system, parallel system, k-out-of-n system, service availability, rare-event approximation, multilinear function, pivotal decomposition, Birnbaum importance, plug-in value, Jensen gap, pooled and separate records, pseudo-checks, drawing, mixture over drawings

A *computer-room air conditioner* (CRAC) is a cooling unit that draws warm air from a data hall and returns it cooled; a hall is served by several. *N+1 redundancy* installs one unit more than the N the load needs, so that the service survives the loss of any one. A *configuration* is a list of which units are up and which are down at one instant. The *structure function* of a service is the indicator, one or zero, of whether a configuration delivers it. A *series system* needs every unit, a *parallel system* needs at least one, and a *k-out-of-n system* needs at least k of its n units. The *service availability* is the probability that the service is delivered at a random instant; for failure probabilities over a mission instead of unavailabilities, the same sum is the system reliability of the companion volume. The *rare-event approximation* replaces the unavailability of a series of subsystems by the sum of their unavailabilities. A function of several variables is *multilinear* when it is of first degree in each variable separately. The *pivotal decomposition* writes the service availability as the average over one unit's two states. The *Birnbaum importance* of a unit is the difference between the service availability with the unit up and with it down. The *plug-in value* of a function evaluates it at the posterior means; the *Jensen gap* is the difference between its posterior expectation and the plug-in value. *Pooled records* treat several units as sharing one unavailability; *separate records* give each unit its own. *Pseudo-checks* are the checks a prior is worth: written in the logit, a Beta(a_0, b_0) prior acts as a_0 units found down in $a_0 + b_0$ checks. A *drawing* is a statement of which configurations deliver the service, and a *mixture over drawings* averages the answers of several candidate drawings with weights.

Key idea

The availability of a service is a sum over the 2^n configurations of its units, and that sum is multilinear in the unit unavailabilities. When the units have independent posteriors, its posterior expectation is exactly its value at the posterior means; when one unavailability is shared by several units, it becomes a polynomial of higher degree in that one entry, and the expected unavailability exceeds the plug-in value for a redundant structure (and falls below it for a series).

13.1.1 From one unit to a service

Chapter 8 derived the unavailability of one repairable pump from its failure and repair rates. A unit's unavailability can also be estimated directly: an operator who inspects the unit at times unrelated to its failures and finds it down at some of the inspections holds a record of Bernoulli trials whose probability is the unit's unavailability q_i, the long-run fraction of time it is down. Chapter 5 gave the posterior of such a probability, a Beta distribution, and its logit, the entry that carries it in the state. The records of this chapter are of that kind, one per unit or one for a fleet, and the question is what they say about a service that several units provide together.

The service is not an entry of the state. It is delivered or not according to which units are up, and the drawing of the plant says which combinations suffice. The companion volume *Reliability Engineering* draws such arrangements as reliability block diagrams and states the closed forms for series, parallel and k-out-of-n blocks. Here each of them is a special case of one sum, because the sum, not the closed form, is what carries over to the arrangements that have no closed form, to the posterior expectation, to the gradient, and to the cut sets and common causes of the next two chapters.

13.1.2 The enumeration and the posterior

Two questions are asked of the sum, and they must be kept apart. The first is the service availability for given unit unavailabilities, $R_{\text{system}}(\mathbf{q})$, a polynomial in $\mathbf{q} = (q_1, \ldots, q_n)$. The second is what the records say about it, and it has two answers, as Chapter 9 kept apart the predictive and the posterior of a reliability. The expectation $\mathbb{E}[R_{\text{system}}(\mathbf{q}) \mid \mathbf{y}]$ is the predictive probability that the service is delivered at a random instant, all uncertainty averaged; the posterior distribution of $R_{\text{system}}(\mathbf{q})$, a derived quantity of the state, carries the band and the probability that the long-run unavailability meets a target. The first is algebra. The second is the derived-quantity problem of Chapter 4, with one feature that the derived quantities of earlier chapters did not have: the function is of first degree in each unit's unavailability, and that property decides when the plug-in value is exact and when it is not. Chapter 9 met the shared-parameter coupling of two redundant generators whose run rate was one unknown; the enumeration gives the general form of that effect for any drawing.

13.1.3 The physical system and its graph (step 0)

A data hall is cooled by four CRAC units in an N+1 arrangement: the design heat load needs three, and the fourth is installed so that any one may be down for a fault or for maintenance. The hall is cooled when at least three of the four run, a 3-out-of-4 drawing. The hall also needs its power feed, the switchgear and distribution path that supplies the units and the equipment; when the feed is down, the hall is down whatever the cooling does. The feed is treated as one more unit in series with the cooling.

Figure 13.1 draws the hall as a graph. The four green boxes in the middle row are the CRAC units, free nodes, each carrying the logit ξ_i of its unavailability q_i. The solid arrows from them to the blue box below are the cooling paths into the hall, the service node whose delivery the drawing decides; the note on it states the drawing. The green box below the hall is the power feed, a free node with the logit ξ_{p} of its unavailability q_{p}, supplied from the utility, the gray pinned box on the right. The white boxes at the top and on the left are the check records, reached from each unit by a dashed observation arrow; they show the records of case B, one per unit. The note on the right states how case A differs: the four unit entries are one shared entry.

The state entries. Each unavailability is a probability and is carried in the state by its logit, as in Chapter 5,

$$\xi_i = \ln \frac{q_i}{1 - q_i}, \qquad q_i = \frac{1}{1 + \mathrm{e}^{-\xi_i}}, \qquad \frac{\mathrm{d}q_i}{\mathrm{d}\xi_i} = q_i(1 - q_i), \tag{13.1}$$

for $i = 1, \ldots, 4$ and for the power feed, $i = \text{p}$. The state is the column vector

$$\mathbf{z}_{\text{B}} = \begin{bmatrix} \xi_1 \\ \xi_2 \\ \xi_3 \\ \xi_4 \\ \xi_{\text{p}} \end{bmatrix} \quad \text{(each unit its own entry)}, \qquad \mathbf{z}_{\text{A}} = \begin{bmatrix} \xi \\ \xi_{\text{p}} \end{bmatrix} \quad \text{(the four units share one entry)}, \tag{13.2}$$

and $\mathbf{q}(\mathbf{z})$ is the vector of unavailabilities it maps to. The second is the first restricted to the direction in which the four unit logits are equal, $\xi_1 = \cdots = \xi_4 = \xi$; the parameterization is the same. The service availability is a polynomial in $\mathbf{q}$, and the enumeration is written there; the fit, the covariance $\boldsymbol{\Sigma}$ and the gradients of the service live in $\mathbf{z}$. The service availability, the band of its unavailability, the importance of each unit and the downtime in hours are derived quantities, functions of the state. Two letters clash with the notation of earlier chapters and are kept apart by context: in k-out-of-n, k is the number of units required and n the number of

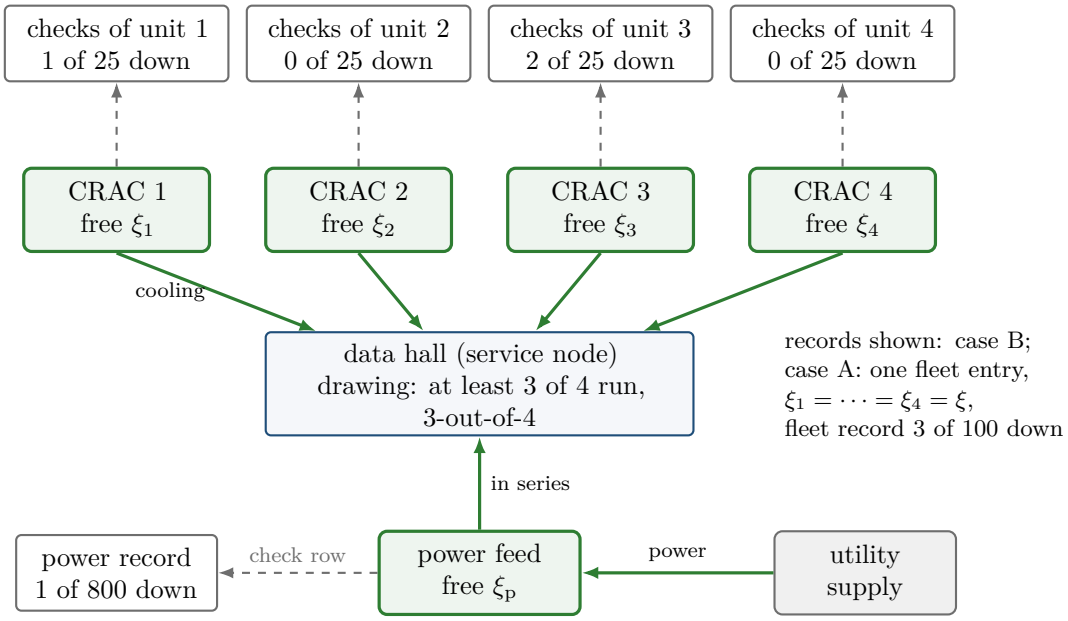


Figure 13.1: Plant graph of the cooled hall. Free: the four CRAC units, each carrying the logit $\xi_i = \ln(q_i/(1-q_i))$ of its unavailability, and the power feed with ξ_p (green). Pinned: the utility supply (gray). The hall (blue) is the service node; its drawing, at least three of the four units in series with the power feed, decides which configurations deliver the service. Observed: the check records of each unit and of the feed (dashed arrows), case B of Section 13.1.4; in case A the four unit entries are one shared entry with the fleet record.

units, while a check record is written with subscripts, k_i units found down in n_i checks for unit i, or k_ℓ in n_ℓ for row ℓ of the fit.

13.1.4 The check records

The evidence is a set of availability checks. At each check an inspector records whether the unit is running or down; a unit found down is a failure in a Bernoulli row of Chapter 5, and the probability of that outcome is q_i. Two evidence cases are carried through the chapter. The records are illustrative.

Case A, one fleet record. The four units are of one model, installed together and maintained by one crew, and the operator treats them as interchangeable: one unavailability q is shared by all four, and the fleet's checks are pooled. Three of a hundred checks found a unit down. The record is taken alone, under the limit of the Beta prior $\mathrm{Beta}(a_0, b_0)$ as its two parameters tend to zero, which is flat in the logit and puts no row in the fit (Section 13.1.6).

Case B, four unit records. Each unit keeps its own record of 25 checks, with $k_i = 1$, 0, 2 and 0 found down, three in a hundred as in case A, and each unit has its own prior $\mathrm{Beta}(1, 19)$, of mean 0.05 and worth twenty checks.

The power feed. The feed's checks found it down once in 800, and its prior, from the switchgear's service history, is $\mathrm{Beta}(1, 199)$, of mean 0.005 and worth two hundred checks. The feed is independent of the cooling units in both cases.

The design file. The design basis asks for a cooling unavailability below $u^* = 10^{-3}$, with the assurance $\kappa = 0.90$ that the design review uses. Whether the hall needs three of its units or

two depends on its heat load against the units' capacity; the design file puts the probability of the 3-out-of-4 drawing at $\pi_3 = 0.7$ and of 2-out-of-4 at $\pi_2 = 0.3$. Three other arrangements are on the table for comparison: the four units with no spare, 4-out-of-4; a fifth unit of the same model with the design load still needing three, 3-out-of-5; and the 2-out-of-4 drawing read as a capacity statement, N+2.

The two cases hold the same three findings in a hundred checks; they differ in the sharing of one entry and in case B's prior, which Section 13.3.3 separates.

13.1.5 The probabilistic questions (step 1)

The operator asks seven questions of the hall, and step 1 writes each as a probability or an expectation over the state, with the unknown on the left of the bar and the records on the right. Table 13.1 lists them: the first column is the question in engineering language; the second the general form that every question of its category takes, in the notation of Chapter 2 for a quantity $f(\mathbf{z})$ derived from the state, a threshold f^* on it, an assurance κ and a candidate action a; the third its mathematical form for the hall; the fourth its category; and the last the subsection where it is answered. The first, third, fourth, fifth and seventh are estimations of derived quantities: the service availability with its band, the same expectation under two posterior structures, the terms of a series, the gradient of the service, and the expectation over an open drawing. The second is a prediction under a threshold, the posterior probability that the long-run unavailability meets the target, compared with the assurance. The sixth is a design question, the choice of an arrangement by the probability that it meets the target. Which further record would narrow the answers most asks nothing new of the present records; it is the next measurement for these seven, answered in Section 13.6. The atlas of Appendix A lists the same forms.

13.1.6 The evidence as rows (step 2)

Step 2 attaches the records to the graph as rows on the entries. In case B, each unit's logit ξ_i receives a prior row from Beta$(1, 19)$ and a count row from its own checks, the binomial likelihood of Chapter 5, and the power feed's logit its own prior and count rows: ten rows on the five entries of $\mathbf{z}_\mathrm{B}$. In case A, the one shared logit ξ receives the count row of the pooled hundred checks, the vague prior putting no row there, and the feed's logit its two rows: three rows on the two entries of $\mathbf{z}_\mathrm{A}$. Which row acts on which entry is a property of the graph, and it fixes where the Jacobian of the rows can be nonzero, marked $*$:

$$\mathbf{J}_\mathrm{B} = \begin{bmatrix} * & 0 & 0 & 0 & 0 \\ 0 & * & 0 & 0 & 0 \\ 0 & 0 & * & 0 & 0 \\ 0 & 0 & 0 & * & 0 \\ * & 0 & 0 & 0 & 0 \\ 0 & * & 0 & 0 & 0 \\ 0 & 0 & * & 0 & 0 \\ 0 & 0 & 0 & * & 0 \\ 0 & 0 & 0 & 0 & * \\ 0 & 0 & 0 & 0 & * \end{bmatrix} \begin{matrix} \text{prior, unit 1} \\ \text{prior, unit 2} \\ \text{prior, unit 3} \\ \text{prior, unit 4} \\ \text{record, unit 1} \\ \text{record, unit 2'} \\ \text{record, unit 3} \\ \text{record, unit 4} \\ \text{prior, feed} \\ \text{record, feed} \end{matrix} \qquad \mathbf{J}_\mathrm{A} = \begin{bmatrix} * & 0 \\ 0 & * \\ 0 & * \end{bmatrix} \begin{matrix} \text{fleet record} \\ \text{feed prior} \\ \text{feed record} \end{matrix}, \tag{13.3}$$

with the columns $\xi_1, \ldots, \xi_4, \xi_\mathrm{p}$ for case B and ξ, ξ_p for case A. In case B no row acts on two entries, at any state. In case A one column carries the record of all four units. No row is a physics row: the service enters as a derived quantity of the state, not as a constraint on it, so no model tolerance σ_p arises. The structure is not a row either: it is the function, fixed by the drawing, through which the state is read when a question about the service is asked. None of

Table 13.1: The questions asked of the hall, the general form of their category, their mathematical form for the hall, their category and where each is answered. $\mathbf{y}$ is the check records and the priors; $R_{\text{system}}(\mathbf{q})$ the service availability for unit unavailabilities $\mathbf{q}$ and $U_{\text{system}} = 1 - R_{\text{system}}$; U_a the unavailability under arrangement a; d indexes the candidate drawings and π_d their weights.

Engineering question	General form	Mathematical form	Category	Answered in
Service availability from components: what is the hall's cooling availability under 3-out-of-4, and its band?	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$, quantiles	the 2^n-state enumeration under the drawing, $\mathbb{E}[R_{\text{system}}(\mathbf{q}) \mid \mathbf{y}]$, band of U_{system}	estimation of a derived quantity	13.3.1
Does N+1 give the hall a cooling unavailability below 10^{-3}, with assurance?	$\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}) \geq \kappa$	$\Pr(U_{\text{system}} < 10^{-3} \mid \mathbf{y}) \geq 0.90$	prediction under a threshold	13.3.2
Is the answer the same when the units' records are pooled and when they are kept apart?	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$	$\mathbb{E}[R_{\text{system}} \mid \mathbf{y}]$ under a shared posterior against independent posteriors	estimation (two posterior structures)	13.3.3
How much does the power feed add, in series?	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$, Var	$\mathbb{E}[U_{\text{series}} \mid \mathbf{y}]$, and $\mathbb{E}$ and Var of each term	estimation of a derived quantity	13.3.4
Which unit matters most?	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$, Var, quantiles	$\mathbb{E}[I_{\text{B},i} \mid \mathbf{y}]$, its band, and each unit's share of $\text{Var}(U_{\text{system}} \mid \mathbf{y})$	estimation of a derived quantity	13.3.4
Which arrangement would meet the target?	$U(a)$, u the acceptance indicator	$\Pr(U_a < 10^{-3} \mid \mathbf{y})$ for each arrangement a	design	13.3.5
If the drawing itself is uncertain, 3-out-of-4 or 2-out-of-4, what availability is reported?	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$ over $(d, \mathbf{z})$	mixture over the open structure, $\sum_d \pi_d \, \mathbb{E}[R_{\text{system}} \mid d, \mathbf{y}]$	estimation under an open structure	13.3.6

the rows is Gaussian: each enters the objective as its negative log-probability term and joins the Jacobian in the working form of Section 4.2.7, written below for these rows.

The modeling choices. Three modeling choices make the model of the hall. The units are taken as up or down independently of one another given their unavailabilities: this is what makes the probability of a configuration a product of one factor per unit, and it fails when a common cause trips several units together or one crew makes the second repair wait for the first (Section 13.6.5). Each unit is taken as either up or down, its unavailability a long-run fraction of time: this makes the service a count of units up and every answer a steady-state fraction, and it fails for a unit running at reduced capacity, for a load that changes through the day, or for checks scheduled after repairs (Section 13.6.5). The four units are carried either as one shared entry or as four entries of their own, cases A and B: this choice decides whether the plug-in value is exact, and the records test it only weakly (Sections 13.3.3 and 13.4.4).

The rows of the check records. Every question reads the posterior of the units' logits, and the check rows give it. Every row of the hall is a binomial row on the logit of one unit, indexed by ℓ. A check record is one such row, k_ℓ units found down in n_ℓ checks, the count row of equation (5.8). It is a probability of the record, not a density over the state, and takes no change-of-variable factor. A prior $\mathrm{Beta}(a_0, b_0)$ is a density over q, and written in the logit it takes the factor $\mathrm{d}q/\mathrm{d}\xi = q(1-q)$, as the prior row of equation (5.15) does:

$$p(\xi) = p(q)\,\frac{\mathrm{d}q}{\mathrm{d}\xi} \propto q^{a_0-1}(1-q)^{b_0-1}\cdot q(1-q) = q^{a_0}(1-q)^{b_0}, \tag{13.4}$$

the binomial term of a_0 found down in $a_0 + b_0$ checks. The prior is therefore a row of *pseudo-checks*: $\mathrm{Beta}(1, 19)$ is the row "1 in 20" and $\mathrm{Beta}(1, 199)$ the row "1 in 200". The limit $\mathrm{Beta}(0, 0)$ is flat in ξ, a row of no checks, which is why the vague prior of case A puts no row in the fit. The shift of one in each exponent is also why the MAP logit is $\ln(a/b)$, which maps to the posterior mean $a/(a+b)$ rather than to the mode of the Beta density. Row ℓ acts on one entry $c(\ell)$ of the state and puts the term

$$\mathcal{L}_\ell(\mathbf{z}) = -k_\ell \ln q_{c(\ell)} - (n_\ell - k_\ell)\ln(1 - q_{c(\ell)})$$

in the objective $\mathcal{L} = \sum_\ell \mathcal{L}_\ell$. Its working residual and Jacobian row are the binomial-in-the-logit pair of equation (4.43), with the selector vector of the entry the row acts on: the deviation of the expected count $n_\ell q_{c(\ell)}$ from the observed one in units of its standard deviation, and the derivative of the expected count in the same units,

$$r_\ell(\mathbf{z}) = \frac{n_\ell q_{c(\ell)} - k_\ell}{\sqrt{n_\ell q_{c(\ell)}(1 - q_{c(\ell)})}}, \qquad \mathbf{J}_\ell = \sqrt{n_\ell q_{c(\ell)}(1 - q_{c(\ell)})}\;\mathbf{e}_{c(\ell)}^{\mathsf{T}}, \tag{13.5}$$

with $\mathbf{e}_{c(\ell)}$ the unit vector of the entry the row acts on. In case B the ten rows are the four unit priors (1 in 20), the four unit records (k_i of 25), the feed's prior (1 in 200) and its record (1 of 800), in the order of the pattern (13.3):

$$\mathbf{r}(\mathbf{z}_\mathrm{B}) = \begin{bmatrix} (20q_1 - 1)/\sqrt{20q_1(1-q_1)} \\ (20q_2 - 1)/\sqrt{20q_2(1-q_2)} \\ (20q_3 - 1)/\sqrt{20q_3(1-q_3)} \\ (20q_4 - 1)/\sqrt{20q_4(1-q_4)} \\ (25q_1 - 1)/\sqrt{25q_1(1-q_1)} \\ (25q_2 - 0)/\sqrt{25q_2(1-q_2)} \\ (25q_3 - 2)/\sqrt{25q_3(1-q_3)} \\ (25q_4 - 0)/\sqrt{25q_4(1-q_4)} \\ (200q_\mathrm{p} - 1)/\sqrt{200q_\mathrm{p}(1-q_\mathrm{p})} \\ (800q_\mathrm{p} - 1)/\sqrt{800q_\mathrm{p}(1-q_\mathrm{p})} \end{bmatrix} \begin{matrix} \text{prior, unit 1} \\ \text{prior, unit 2} \\ \text{prior, unit 3} \\ \text{prior, unit 4} \\ \text{record, unit 1} \\ \text{record, unit 2'} \\ \text{record, unit 3} \\ \text{record, unit 4} \\ \text{prior, feed} \\ \text{record, feed} \end{matrix} \tag{13.6}$$

with $q_i = 1/(1+\mathrm{e}^{-\xi_i})$. In case A the rows are the fleet record (3 of 100) on ξ and the feed's two rows on ξ_{p}.

13.2 The posterior of the state (step 3)

All seven questions are computed from the posterior of the units' logits that the check rows of step 2 give, the Laplace normal; a Beta density per entry, which these rows also admit, gives the exact value beside each answer. This section derives the posterior and computes it for the hall; Section 13.3 reads every answer from it.

Four objects of the fit recur in the specialized forms. The *MAP state* $\hat{\mathbf{z}}$ is the minimum of the objective that the check rows put on the state; the *information matrix* $\mathbf{\Lambda} = \mathbf{J}^{\mathsf{T}}\mathbf{J}$ is the curvature of that objective at the minimum, with $\mathbf{J}$ the Jacobian of the normalized rows, of the pattern (13.3); the *Laplace covariance* $\mathbf{\Sigma} = \mathbf{\Lambda}^{-1}$ makes the posterior of the state approximately normal, $\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$; and a derived quantity f is carried by its *gradient* $\mathbf{g} = \partial f/\partial \mathbf{z}$ at $\hat{\mathbf{z}}$, with the variance $\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}$. Beside them stands the exact posterior that the check rows admit: each unit unavailability has a Beta posterior, independent of the others in case B, and the four units have one shared Beta posterior in case A.

13.2.1 The posterior from the rows

The logit is the canonical parameter of the binomial distribution, so the curvature term that a Gauss–Newton product drops is zero at every state (Chapter 5): $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ is the exact Hessian of $\mathcal{L}$, and Newton's method and the Gauss–Newton step coincide. With every row normalized in this way the information matrix is $\mathbf{\Lambda} = \mathbf{J}^{\mathsf{T}}\mathbf{J}$. Each step solves the normal equations

$$(\mathbf{J}^{\mathsf{T}}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^{\mathsf{T}}\mathbf{r}, \qquad \mathbf{z} \leftarrow \mathbf{z} + \Delta\mathbf{z}, \tag{13.7}$$

with $\mathbf{J}$ and $\mathbf{r}$ at the current state.

The MAP state and the covariance. The iteration stops where $\mathbf{J}^{\mathsf{T}}\mathbf{r} = \mathbf{0}$, that is where each column's expected findings equal its observed ones, the prior's pseudo-checks included: $(a_{i0} + b_{i0} + n_i)\,q_i = a_{i0} + k_i$ for entry i, so that

$$\hat{z}_i = \ln\frac{a_i}{b_i}, \qquad q_i(\hat{z}_i) = \frac{a_i}{a_i + b_i}, \qquad a_i = a_{i0} + k_i, \quad b_i = b_{i0} + n_i - k_i, \tag{13.8}$$

the logits of the posterior means, as equation (5.20) gave for one valve. The residual vector itself is not zero there: a unit's prior row and its record pull in opposite directions when they disagree, and the minimum is where their pulls, weighted by the Jacobian entries, balance. An off-diagonal entry of $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ is $\sum_\ell J_{\ell i}J_{\ell j}$, and every product is zero because no row of $\mathbf{J}$ has two nonzero entries, by the pattern (13.3): the records occupy disjoint columns. The information and the covariance are therefore diagonal,

$$\mathbf{\Lambda} = \mathbf{J}^{\mathsf{T}}\mathbf{J} = \operatorname{diag}\Big(\frac{a_i b_i}{a_i + b_i}\Big), \qquad \mathbf{\Sigma} = \mathbf{\Lambda}^{-1} = \operatorname{diag}\Big(\frac{a_i + b_i}{a_i b_i}\Big), \tag{13.9}$$

each diagonal entry of $\mathbf{\Lambda}$ the sum of the squared entries of its column, the prior's and the record's, $(a_{i0} + b_{i0} + n_i)\,\hat{q}_i(1-\hat{q}_i) = a_i b_i/(a_i + b_i)$, as in equation (5.21). The diagonal $\mathbf{\Sigma}$, and with it the independence of the posteriors of the units, is a result of the rows, not an assumption about the units. In case A one column carries the record of all four units, so any function of the four unavailabilities depends on that one entry through all four at once.

The exact posteriors. The rows of each entry are those of Chapter 5, and so is the exact posterior: by equation (5.24), $q_i \mid \mathbf{y} \sim \mathrm{Beta}(a_i, b_i)$ with the parameters of equation (13.8), independent across the entries because the likelihood and the prior factor by entry. The objective is the negative logarithm of their joint density written in the logits, and $\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ is its Laplace approximation, centered at the logits of the posterior means. The two differ most for short records, and they differ where it matters, in the band of the service (Section 13.4.1).

13.2.2 The posterior of the hall

The hall's inputs are the records of Section 13.1.4: in case A, three of a hundred fleet checks under the vague limit of the prior; in case B, 1, 0, 2 and 0 of 25 checks per unit under $\mathrm{Beta}(1, 19)$ each; and for the feed, 1 of 800 under $\mathrm{Beta}(1, 199)$.

The Newton iteration. Start every unit at the logit of its prior mean, $\xi_i^{(0)} = \ln(0.05/0.95) = -2.9444$, and the feed at $\xi_{\mathrm{p}}^{(0)} = \ln(0.005/0.995) = -5.2933$. Table 13.2 lists the iteration of equation (13.7) for case B: each row is one iterate, the columns give the five entries and the largest entry of the gradient $\mathbf{J}^{\mathsf{T}}\mathbf{r}$ in absolute value. The entries whose records depart most from their priors move furthest: units 2 and 4, with no finding in 25 checks, and the feed, whose one finding in 800 checks lies below its prior's one in 200. The convergence is quadratic from the second step.

Table 13.2: Newton iteration of equation (13.7) for the five logits of case B from the logits of the prior means. The last column is $\max_i |(\mathbf{J}^{\mathsf{T}}\mathbf{r})_i|$.

Iteration [–]	ξ_1 [–]	ξ_2 [–]	ξ_3 [–]	ξ_4 [–]	ξ_{p} [–]	$\max\lvert\mathbf{J}^{\mathsf{T}}\mathbf{r}\rvert$ [–]
0	−2.94444	−2.94444	−2.94444	−2.94444	−5.29330	3.0
1	−3.06140	−3.52923	−2.59356	−3.52923	−5.89632	0.74
2	−3.06803	−3.75574	−2.63818	−3.75574	−6.16767	9.2×10^{-2}
3	−3.06805	−3.78381	−2.63906	−3.78381	−6.21162	2.0×10^{-3}
4	−3.06805	−3.78419	−2.63906	−3.78419	−6.21261	9.8×10^{-7}
5	−3.06805	−3.78419	−2.63906	−3.78419	−6.21261	2.4×10^{-13}

The MAP state, the information and the covariance. By equation (13.8) the iteration ends at

$$\hat{\mathbf{z}}_{\mathrm{B}} = \begin{bmatrix} \ln(2/43) \\ \ln(1/44) \\ \ln(3/42) \\ \ln(1/44) \\ \ln(2/998) \end{bmatrix} = \begin{bmatrix} -3.0681 \\ -3.7842 \\ -2.6391 \\ -3.7842 \\ -6.2126 \end{bmatrix}, \qquad \mathbf{q}(\hat{\mathbf{z}}_{\mathrm{B}}) = \begin{bmatrix} 0.0444 \\ 0.0222 \\ 0.0667 \\ 0.0222 \\ 0.0020 \end{bmatrix}, \tag{13.10}$$

the logits $\ln(a_i/b_i)$ of the Beta posteriors, which map to their posterior means. At $\hat{\mathbf{z}}_{\mathrm{B}}$ the 10×5 Jacobian, with the pattern (13.3) and the rows in the order of equation (13.6), and the

information and covariance of equation (13.9) are

$$\mathbf{J}(\hat{\mathbf{z}}_{\mathrm{B}}) = \begin{bmatrix} 0.9216 & 0 & 0 & 0 & 0 \\ 0 & 0.6592 & 0 & 0 & 0 \\ 0 & 0 & 1.1155 & 0 & 0 \\ 0 & 0 & 0 & 0.6592 & 0 \\ 1.0304 & 0 & 0 & 0 & 0 \\ 0 & 0.7370 & 0 & 0 & 0 \\ 0 & 0 & 1.2472 & 0 & 0 \\ 0 & 0 & 0 & 0.7370 & 0 \\ 0 & 0 & 0 & 0 & 0.6318 \\ 0 & 0 & 0 & 0 & 1.2636 \end{bmatrix}, \tag{13.11}$$

$$\mathbf{\Lambda} = \mathbf{J}^{\mathsf{T}}\mathbf{J} = \mathrm{diag}(1.9111,\ 0.9778,\ 2.8000,\ 0.9778,\ 1.9960),$$

$$\mathbf{\Sigma}_{\mathrm{B}} = \mathbf{\Lambda}^{-1} = \mathrm{diag}(0.52326,\ 1.02273,\ 0.35714,\ 1.02273,\ 0.50100).$$

For unit 1 the two squared entries of its column are $0.8494 + 1.0617 = 1.9111 = 2 \times 43/45$. Each logit is read off as its entry, $\hat{\xi}_i \pm \sqrt{\Sigma_{ii}}$ with $\sqrt{\Sigma_{ii}} = (0.723,\ 1.011,\ 0.598,\ 1.011,\ 0.708)$: the two units with no finding have the widest logits.

Case A. With the four units sharing one entry, the fleet's hundred checks are one row on the first column of a 3×2 Jacobian, and the feed's two rows are on the second. From $(\ln(0.05/0.95),\ \ln(0.005/0.995))^{\mathsf{T}}$ the same iteration reaches, in five steps,

$$\hat{\mathbf{z}}_{\mathrm{A}} = \begin{bmatrix} \ln(3/97) \\ \ln(2/998) \end{bmatrix} = \begin{bmatrix} -3.4761 \\ -6.2126 \end{bmatrix}, \qquad \mathbf{J}(\hat{\mathbf{z}}_{\mathrm{A}}) = \begin{bmatrix} 1.7059 & 0 \\ 0 & 0.6318 \\ 0 & 1.2636 \end{bmatrix},$$
$$\mathbf{\Sigma}_{\mathrm{A}} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1} = \begin{bmatrix} 0.34364 & 0 \\ 0 & 0.50100 \end{bmatrix}, \tag{13.12}$$

with $\mathbf{J}^{\mathsf{T}}\mathbf{J} = \mathrm{diag}(2.9100,\ 1.9960)$ and $\Sigma_{11} = 100/291$.

The exact posteriors. Table 13.3 lists the Beta posteriors that the fit approximates: each row is an entry, and the columns give the record, the prior, the posterior and its mean, standard deviation and 95 % interval. The feed's posterior $\mathrm{Beta}(2, 998)$ carries the weight of a thousand checks in all. Case B's prior, worth twenty checks per unit, pulls each short record toward 0.05, so that the average of its four means is 0.0389 rather than case A's 0.0300.

The pairs $(\hat{\mathbf{z}}_{\mathrm{A}}, \mathbf{\Sigma}_{\mathrm{A}})$ and $(\hat{\mathbf{z}}_{\mathrm{B}}, \mathbf{\Sigma}_{\mathrm{B}})$ of equations (13.12) and (13.11) are what step 4 uses: every answer of the chapter is computed from them. Because this problem also admits the exact posteriors, each answer has an exact value beside it, a function of the Beta parameters of Table 13.3, and the gap between the two is reported with it.

13.3 The answers (step 4)

The seven questions of Table 13.1 take three general forms of Chapter 2: the estimation of a derived quantity, equations (2.9) to (2.11); the threshold, equation (2.12); and the design form, equation (2.16). In each, $f(\mathbf{z})$ is a quantity derived from the units' logits through the enumeration of their configurations, and $p(\mathbf{z} \mid \mathbf{y})$ the posterior density of the state. Each is answered from the posterior of Section 13.2: its engineering question, the general form of its category, the form it takes for the hall, the quantities of the posterior it uses, and the hall's

Table 13.3: The evidence cases of the hall. Case A pools the fleet's checks into one shared unavailability; case B keeps one record and one prior per unit. Intervals are central 95 % posterior intervals of the Beta posteriors.

Entry	Record	Prior	Posterior	Mean	Std. dev.	95 % interval
A: fleet q	3 of 100 down	vague limit	Beta(3, 97)	0.0300	0.0170	[0.0063, 0.0711]
B: unit 1, q_1	1 of 25 down	Beta(1, 19)	Beta(2, 43)	0.0444	0.0304	[0.0055, 0.1202]
B: unit 2, q_2	0 of 25 down	Beta(1, 19)	Beta(1, 44)	0.0222	0.0217	[0.0006, 0.0804]
B: unit 3, q_3	2 of 25 down	Beta(1, 19)	Beta(3, 42)	0.0667	0.0368	[0.0143, 0.1547]
B: unit 4, q_4	0 of 25 down	Beta(1, 19)	Beta(1, 44)	0.0222	0.0217	[0.0006, 0.0804]
Power feed, q_{p}	1 of 800 down	Beta(1, 199)	Beta(2, 998)	0.0020	0.0014	[0.0002, 0.0056]

numbers. The questions name the target $u^* = 10^{-3}$ at the assurance 0.90, the design file's weights 0.7 and 0.3 and the year of 8760 calendar hours. Every answer is computed from the MAP state $\hat{\mathbf{z}}$ and the covariance $\boldsymbol{\Sigma}$ of step 3, and that is the value the chapter reports. The exact posteriors give the value set beside it, with the gap: through the Beta moments and the quantile map where the units share one entry, and through the value at the posterior means and 4×10^5 draws of the posteriors, seed 13, where they are independent.

13.3.1 The cooling availability and its band

Engineering question. Service availability from components: what is the hall's cooling availability under 3-out-of-4, and its band? The question asks for the long-run fraction of time the hall is cooled, a function of the four units' unavailabilities, for its expectation over what the records say, and for the band of its long-run value.

General form. An estimation question about a quantity $f(\mathbf{z})$ derived from the state asks for its posterior distribution, equation (2.9), summarized by its mean and its central 95 % interval, equations (2.10) and (2.11):

$$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}] = \int f(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \qquad \Pr(f(\mathbf{z}) \le f_L \mid \mathbf{y}) = 0.025, \quad \Pr(f(\mathbf{z}) \le f_U \mid \mathbf{y}) = 0.975.$$

Specialized form. The question is equations (2.10) and (2.11) with $f(\mathbf{z}) = U_{\mathrm{system}}(\mathbf{q}(\mathbf{z}))$, the unavailability of the cooling, a function of the four unit logits through the maps (13.1), and $p(\mathbf{z} \mid \mathbf{y})$ the posterior of step 3, the Laplace normal $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of equations (13.12) and (13.11). The service availability is $R_{\mathrm{system}} = 1 - U_{\mathrm{system}}$.

The enumeration. Let $S_i \in \{0, 1\}$ be the status of unit i at a random instant, one when it is up and zero when it is down, so that $\Pr(S_i = 0 \mid q_i) = q_i$. Both cases are covered by one expression,

$$\Pr(S_i = s_i \mid q_i) = q_i^{1-s_i}\,(1 - q_i)^{s_i}, \qquad s_i \in \{0, 1\}. \tag{13.13}$$

A configuration is the list $\mathbf{s} = (s_1, \ldots, s_n)$ of the statuses of all n units, one of 2^n such lists. Given the unavailabilities the units are up or down independently, so the probability of a configuration is the product of the probabilities of its statuses,

$$\Pr(\mathbf{S} = \mathbf{s} \mid \mathbf{q}) = \prod_{i=1}^{n} q_i^{1-s_i}\,(1 - q_i)^{s_i}. \tag{13.14}$$

These probabilities sum to one over the 2^n configurations: the sum over s_1 of the factor of unit 1 is $q_1 + (1 - q_1) = 1$, and summing each unit in turn,

$$\sum_{\mathbf{s}\in\{0,1\}^n} \prod_{i=1}^{n} q_i^{1-s_i}(1-q_i)^{s_i} = \prod_{i=1}^{n} (q_i + (1-q_i)) = 1.$$

The drawing of the plant is a rule that says, for every configuration, whether the service is delivered. Write it as the structure function $\phi(\mathbf{s}) \in \{0, 1\}$, one when configuration $\mathbf{s}$ delivers the service. The configurations partition the outcomes, so by the law of total probability the service availability is

$$R_{\text{system}}(\mathbf{q}) = \Pr(\phi(\mathbf{S}) = 1 \mid \mathbf{q}) = \sum_{\mathbf{s}\in\{0,1\}^n} \phi(\mathbf{s}) \Pr(\mathbf{S} = \mathbf{s} \mid \mathbf{q}), \tag{13.15}$$

and the service unavailability is the sum over the configurations that do not deliver it,

$$U_{\text{system}}(\mathbf{q}) = 1 - R_{\text{system}}(\mathbf{q}) = \sum_{\mathbf{s}} (1 - \phi(\mathbf{s})) \Pr(\mathbf{S} = \mathbf{s} \mid \mathbf{q}). \tag{13.16}$$

Equation (13.15) is the whole of the system model. A drawing enters only through ϕ, the unit records only through $\mathbf{q}$, and any arrangement whatever, with or without a closed form, is evaluated by the same sum. For the cooling of the hall, $n = 4$ and $\phi(\mathbf{s}) = \mathbf{1}[s_1 + s_2 + s_3 + s_4 \geq 3]$.

Series, parallel and k-out-of-n. The three standard arrangements are three structure functions, and their closed forms follow from equation (13.15) in a line each. A series system needs every unit, $\phi(\mathbf{s}) = \prod_i s_i$, which is one only for the configuration with all units up, and the sum has one term:

$$R_{\text{system}} = \prod_{i=1}^{n} (1 - q_i), \qquad U_{\text{system}} = 1 - \prod_{i=1}^{n} (1 - q_i). \tag{13.17}$$

A parallel system needs at least one unit, $\phi(\mathbf{s}) = 1 - \prod_i (1 - s_i)$, which is zero only for the configuration with all units down, and its unavailability has one term:

$$U_{\text{system}} = \prod_{i=1}^{n} q_i, \qquad R_{\text{system}} = 1 - \prod_{i=1}^{n} q_i. \tag{13.18}$$

A k-out-of-n system needs at least k units up, $\phi(\mathbf{s}) = \mathbf{1}[\sum_i s_i \geq k]$. For general unavailabilities the sum has no shorter form than equation (13.15). For equal unavailabilities q, every configuration with j units up has probability $(1 - q)^j q^{n-j}$, there are $\binom{n}{j}$ of them, and grouping the sum by j gives the binomial tail

$$R_{\text{system}}(q) = \sum_{j=k}^{n} \binom{n}{j} (1-q)^j q^{n-j}, \qquad U_{\text{system}}(q) = \sum_{j=0}^{k-1} \binom{n}{j} (1-q)^j q^{n-j}. \tag{13.19}$$

The series system is n-out-of-n and the parallel system 1-out-of-n. For the hall's 3-out-of-4, the terms $j = 0, 1, 2$ are q^4, $4q^3(1 - q)$ and $6q^2(1 - q)^2$; expanding and collecting powers,

$$U_{\text{system}}(q) = q^4 + 4q^3 - 4q^4 + 6q^2 - 12q^3 + 6q^4 = 6q^2 - 8q^3 + 3q^4. \tag{13.20}$$

For small q the lowest power dominates. The service is lost when $n - k + 1$ units are down, so the lowest term of equation (13.19) is $j = k - 1$, and

$$U_{\text{system}}(q) \approx \binom{n}{k-1} q^{n-k+1}, \qquad q \ll 1, \tag{13.21}$$

Table 13.4: Structure functions and their closed forms, each obtained from the enumeration (13.15). The order is the exponent of the leading term of U_{system} for small equal q.

Structure	$\phi(\mathbf{s})$	U_{system}, general q_i	U_{system}, equal q	Order
Series (n-out-of-n)	$\prod_i s_i$	$1-\prod_i(1-q_i)$	$1-(1-q)^n$; for $n=4$: $4q-6q^2+4q^3-q^4$	1
Parallel (1-out-of-n)	$1-\prod_i(1-s_i)$	$\prod_i q_i$	q^n; for $n=4$: q^4	n
k-out-of-n	$\mathbf{1}[\sum_i s_i \geq k]$	equation (13.16)	$\sum_{j<k}\binom{n}{j}(1-q)^j q^{n-j}$	$n-k+1$
2-out-of-4	$\mathbf{1}[\sum_i s_i \geq 2]$	equation (13.16)	$4q^3-3q^4$	3
3-out-of-4	$\mathbf{1}[\sum_i s_i \geq 3]$	equation (13.16)	$6q^2-8q^3+3q^4$	2
3-out-of-5	$\mathbf{1}[\sum_i s_i \geq 3]$	equation (13.16)	$10q^3-15q^4+6q^5$	3

$6q^2$ for 3-out-of-4. The exponent $n-k+1$, the *order* of the structure, is the smallest number of units whose loss takes the service down; the next chapter calls the sets of such units the minimal cut sets and reads them off this enumeration. Table 13.4 collects the forms: each row is a structure, and the columns give its structure function, its unavailability for general and for equal unavailabilities, and its order.

Multilinearity and the gradient. In equation (13.14) each unit contributes one factor, q_i or $1-q_i$, so every term of the sum (13.15) is of first degree in each q_i separately, and so is the sum. The service availability is a multilinear polynomial: it contains no power q_i^2 or higher of any one unit, only products of distinct units' unavailabilities. In particular

$$\frac{\partial^2 R_{\text{system}}}{\partial q_i^2} = 0 \qquad \text{for every unit } i, \tag{13.22}$$

while the mixed derivatives $\partial^2 R_{\text{system}}/\partial q_i \partial q_j$, $i \neq j$, need not vanish. Conditioning on the status of unit i, by the law of total probability,

$$R_{\text{system}}(\mathbf{q}) = q_i\, R_{\text{system}}(0_i, \mathbf{q}) + (1-q_i)\, R_{\text{system}}(1_i, \mathbf{q}), \tag{13.23}$$

the *pivotal decomposition*, where $R_{\text{system}}(1_i, \mathbf{q})$ is the service availability with unit i held up and $R_{\text{system}}(0_i, \mathbf{q})$ with it held down; neither depends on q_i. Equation (13.23) is the multilinearity written for one unit, and it gives the gradient at once:

$$\begin{aligned} \frac{\partial R_{\text{system}}}{\partial q_i} &= R_{\text{system}}(0_i, \mathbf{q}) - R_{\text{system}}(1_i, \mathbf{q}) = -I_{\text{B},i}, \\ I_{\text{B},i} &= R_{\text{system}}(1_i, \mathbf{q}) - R_{\text{system}}(0_i, \mathbf{q}) = \frac{\partial U_{\text{system}}}{\partial q_i}, \end{aligned} \tag{13.24}$$

the same for every value of q_i, because by equation (13.22) the service availability is a straight line in q_i. The quantity $I_{\text{B},i}$ is the Birnbaum importance of unit i, read as the answer to its own question in Section 13.3.4. The fit and the band work in the logits of equation (13.2), and each q_i depends on its own logit only, so the derivative of $\mathbf{q}$ with respect to $\mathbf{z}$ is a diagonal matrix and the gradient of the service unavailability in the state is

$$\mathbf{g} = \frac{\partial U_{\text{system}}}{\partial \mathbf{z}} = \mathbf{D}\,\mathbf{I}_{\text{B}}, \qquad \mathbf{D} = \frac{\partial \mathbf{q}}{\partial \mathbf{z}^{\mathsf{T}}} = \operatorname{diag}(q_i(1-q_i)), \qquad \mathbf{I}_{\text{B}} = \frac{\partial U_{\text{system}}}{\partial \mathbf{q}}, \tag{13.25}$$

an n-vector with one entry per unit, the Birnbaum vector scaled entry by entry by the slope of each logistic map. In case A the four unavailabilities are one entry, and the shared column collects the four units: the gradient has the one entry $q(1-q)\,\mathbf{1}^{\mathsf{T}}\mathbf{I}_{\text{B}} = q(1-q)\,U'_{\text{system}}(q)$, with $\mathbf{1} = (1,1,1,1)^{\mathsf{T}}$. This vector is the object that the cut sets of the next chapter and the importance measures of a later chapter reuse.

The expectation over the state. The integral of equation (2.10) runs over the Laplace normal. For a function h of one entry ξ with MAP value $\hat{\xi}$ and variance σ_ξ^2 the Gauss–Hermite rule of equation (5.33) evaluates it,

$$\mathbb{E}[h(\xi) \mid \mathbf{y}] \approx \frac{1}{\sqrt{\pi}} \sum_{m=1}^{M} w_m \, h(\hat{\xi} + \sqrt{2}\,\sigma_\xi \, t_m), \tag{13.26}$$

with nodes t_m and weights w_m. In case A the cooling is a function of the one entry ξ, so $h = U_{\text{system}}(q(\xi))$. In case B it is a function of four entries, and the next paragraph reduces the integral to one such sum per entry.

The expectation under independent posteriors. Suppose the unit unavailabilities $q_1, \ldots, q_n$ are independent a posteriori, as in case B. Write the factor of unit i in equation (13.14) as $f_{s_i}(q_i)$, with $f_0(q) = q$ and $f_1(q) = 1 - q$, both linear. The expectation of the service availability is

$$\begin{aligned}
\mathbb{E}[R_{\text{system}}(\mathbf{q}) \mid \mathbf{y}] &= \sum_{\mathbf{s}} \phi(\mathbf{s}) \, \mathbb{E}\Big[\prod_{i=1}^{n} f_{s_i}(q_i) \;\Big|\; \mathbf{y}\Big] && \text{linearity of the expectation} \\
&= \sum_{\mathbf{s}} \phi(\mathbf{s}) \prod_{i=1}^{n} \mathbb{E}[f_{s_i}(q_i) \mid \mathbf{y}] && \text{independence of the } q_i \\
&= \sum_{\mathbf{s}} \phi(\mathbf{s}) \prod_{i=1}^{n} f_{s_i}(\mathbb{E}[q_i \mid \mathbf{y}]) && f_{s_i} \text{ linear} \\
&= R_{\text{system}}\big(\mathbb{E}[\mathbf{q} \mid \mathbf{y}]\big). && \text{equation (13.15)}
\end{aligned} \tag{13.27}$$

The posterior expectation of the service availability, the predictive probability that the service is delivered at a random instant, is exactly its value at the posterior means. The second line is where the argument needs independence: each term is a product over distinct units, and the expectation of a product of independent factors is the product of their expectations. The third line is where it needs multilinearity: each factor is linear, so its expectation is its value at the mean. In a Taylor expansion about the posterior means $\bar{\mathbf{q}}$ the same result reads through the second-order term of the expectation,

$$\tfrac{1}{2}\,\operatorname{tr}(\mathbf{H}_q \boldsymbol{\Sigma}_q) = \tfrac{1}{2} \sum_{i,j} (\mathbf{H}_q)_{ij} \, \operatorname{Cov}(q_i, q_j \mid \mathbf{y}), \qquad \mathbf{H}_q = \frac{\partial^2 U_{\text{system}}}{\partial \mathbf{q} \, \partial \mathbf{q}^{\mathsf{T}}}, \tag{13.28}$$

with $\boldsymbol{\Sigma}_q$ the posterior covariance of $\mathbf{q}$; under the Laplace approximation $\boldsymbol{\Sigma}_q \approx \mathbf{D}\boldsymbol{\Sigma}\mathbf{D}$, the covariance of the state carried through the logistic maps (equation (13.25)). Multilinearity puts zeros on the diagonal of $\mathbf{H}_q$, independence puts zeros off the diagonal of $\boldsymbol{\Sigma}_q$, and every product in the sum vanishes. The trace is taken in $\mathbf{q}$ rather than in $\mathbf{z}$ because in the logits the curvature of each logistic map adds a diagonal term to the Hessian that has nothing to do with the structure. The result concerns the expectation only: the band, the probability of meeting a target and every other summary of the distribution of U_{system} depend on the whole posterior.

The argument uses only the independence of the entries and the linearity of the factors, so it holds for the Laplace normal of case B as well, whose covariance (13.11) is diagonal: the state-route expectation is the enumeration at the means $\bar{q}_i^{\,\mathrm{L}} = \mathbb{E}[q(\xi_i) \mid \mathbf{y}]$ of that normal, each from the sum (13.26) with $h = q(\xi_i)$. Those means are not the values $q(\hat{z}_i) = a_i/(a_i + b_i)$ at the MAP state: a normal in the logit, mapped to q, has a long upper tail, and its mean lies above its median $q(\hat{z}_i)$. On the state route, therefore, the expectation exceeds the plug-in value even with independent entries; on the exact route the plug-in value is the expectation.

The expectation under a shared posterior. In case A one entry q stands for all four units. Every configuration with m units down then has probability $q^m(1-q)^{n-m}$, and the service availability becomes the polynomial of degree n in one variable of equation (13.19). Its terms

contain powers of the same unknown, q^2, q^3 and q^4 for the hall, and the expectation of a power is not the power of the expectation. On the state route the polynomial is the function h of the sum (13.26). The value beside it comes from the exact posterior: for $q \sim \mathrm{Beta}(a, b)$ the r-th moment is the Beta integral of Chapter 5 with the exponent of q raised by r:

$$\mathbb{E}[q^r \mid \mathbf{y}] = \frac{B(a+r, b)}{B(a, b)} = \frac{\Gamma(a+r)\,\Gamma(a+b)}{\Gamma(a)\,\Gamma(a+b+r)} = \prod_{c=0}^{r-1} \frac{a+c}{a+b+c}, \tag{13.29}$$

using $\Gamma(x+1) = x\Gamma(x)$ repeatedly. Taking the expectation of equation (13.20) term by term,

$$\mathbb{E}[U_{\text{system}} \mid \mathbf{y}] = 6\mathbb{E}[q^2 \mid \mathbf{y}] - 8\mathbb{E}[q^3 \mid \mathbf{y}] + 3\mathbb{E}[q^4 \mid \mathbf{y}], \tag{13.30}$$

exact under the shared posterior. The second-order term (13.28) gives the leading order of its difference from the plug-in value, the *Jensen gap*, with the matrices of the shared entry. With $\mathbf{q} = \mathbf{1}q$ the covariance of the four unavailabilities is $\mathbf{\Sigma}_q = \mathrm{Var}(q \mid \mathbf{y})\,\mathbf{1}\mathbf{1}^{\mathsf{T}}$, and every entry of $\mathbf{H}_q$ enters, the mixed derivatives included. Under 3-out-of-4 at equal unavailabilities, with units i and j both down the service is lost; with one of them down it is lost when one of the other two is down, with probability $1-(1-q)^2$; with neither down it is lost when both others are, with probability q^2. The mixed derivative is the second difference of these three, $1-2(1-(1-q)^2)+q^2 = (1-q)(1-3q)$, and the diagonal is zero:

$$\begin{gathered}
\mathbf{H}_q = (1-q)(1-3q)\,(\mathbf{1}\mathbf{1}^{\mathsf{T}} - \mathbf{I}), \qquad \mathbf{\Sigma}_q = \mathrm{Var}(q \mid \mathbf{y})\,\mathbf{1}\mathbf{1}^{\mathsf{T}}, \\
\mathbb{E}[U_{\text{system}} \mid \mathbf{y}] - U_{\text{system}}(\bar{q}) \approx \tfrac{1}{2}\,\mathrm{tr}(\mathbf{H}_q\mathbf{\Sigma}_q) = \tfrac{1}{2}\,(\mathbf{1}^{\mathsf{T}}\mathbf{H}_q\mathbf{1})\,\mathrm{Var}(q \mid \mathbf{y}), \\
\mathbf{1}^{\mathsf{T}}\mathbf{H}_q\mathbf{1} = U''_{\text{system}}(q) = 12 - 48q + 36q^2,
\end{gathered} \tag{13.31}$$

with $\mathbf{I}$ the 4×4 identity and the matrices evaluated at $\bar{q}$. The same matrix $\mathbf{H}_q$ gives no gap when it is contracted with the diagonal $\mathbf{\Sigma}_q$ of independent entries and a gap when it is contracted with the rank-one $\mathbf{\Sigma}_q$ of a shared entry: the shared column of the Jacobian turns the mixed derivatives of the enumeration into a curvature in the one entry. The gap is positive because the unavailability of a redundant structure is convex in q for small q, and draws of q above the mean raise U more than draws below it lower it; for the series 4-out-of-4 the curvature is negative and the gap changes sign. That is the gap of the exact expectation (13.30); the state-route sum carries, on top of it, the upper tail of the normal in the logit.

The band. Under the Laplace posterior the band follows from the gradient by the delta method of Chapter 4: with the gradient (13.25) and the covariance of the fit,

$$\sigma_U^2 = \mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}, \qquad \mathbf{g} = \mathbf{D}\,\mathbf{I}_{\mathrm{B}}, \qquad \mathbf{g}_{\mathrm{A}} = \begin{bmatrix} \hat{q}(1-\hat{q})\,U'_{\text{system}}(\hat{q}) \\ 0 \end{bmatrix}, \tag{13.32}$$

the delta-method variance σ_f^2 of that chapter with $f = U_{\text{system}}$, written σ_U^2; it names the Laplace value only, and the exact spread is written $\mathrm{sd}(U_{\text{system}} \mid \mathbf{y})$. Because $\mathbf{\Sigma}$ is diagonal, $\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g} = \sum_i g_i^2\Sigma_{ii}$ is a variance budget with one term per entry. The last entry of $\mathbf{g}$ is zero because the feed is not part of the cooling. A normal with the plug-in value as mean and σ_U as spread cannot describe a quantity that is small, positive and uncertain by a factor: when σ_U is comparable to the value, its band crosses zero. The unavailability is nearly a power of the unit unavailabilities: by equation (13.21), $\ln U_{\text{system}} \approx \ln 6 + 2\ln q$, and for small q the logarithm $\ln q$ differs from the logit ξ by $-\ln(1-q) \approx q$. The logarithm of the service unavailability is therefore nearly linear in the state entries, and the delta method belongs on it. Its gradient is the gradient of U divided by U:

$$\sigma_{\ln U}^2 = \mathbf{g}_{\ln U}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}_{\ln U}, \qquad \mathbf{g}_{\ln U} = \frac{\partial \ln U_{\text{system}}}{\partial \mathbf{z}} = \frac{\mathbf{g}}{U_{\text{system}}}, \tag{13.33}$$

and the band is $U_{\text{system}}(\hat{\mathbf{z}})\,\mathrm{e}^{\pm 1.96\,\sigma_{\ln U}}$, positive and skewed, with the plug-in value $U_{\text{system}}(\hat{\mathbf{z}})$ as its median. That band is the one the chapter reports. The entries of $\mathbf{g}_{\ln U}$ sum to nearly the order of the structure: raising every logit by one raises the unavailability by nearly the factor e^{n-k+1}.

The value beside it comes from the exact posterior. With one shared entry the exact distribution is available. The derivative $U'_{\text{system}}(q) = 12q(1-q)^2$ is positive on $(0,1)$, so the cooling unavailability is an increasing function of q, and its quantiles are the images of the quantiles of q, equation (2.11) read through a monotone map:

$$U_{\text{system},\alpha} = U_{\text{system}}(q_\alpha), \qquad \Pr(U_{\text{system}} < u^* \mid \mathbf{y}) = I_{q^*}(a,b), \quad U_{\text{system}}(q^*) = u^*, \tag{13.34}$$

with q_α the α-quantile of the Beta posterior and $I_x(a,b)$ the regularized incomplete Beta function of Chapter 5. Any structure whose unavailability increases with a single shared entry has this exact route. With independent entries, as in case B, the distribution of U_{system} is that of a polynomial in four independent Beta variables, and its quantiles are read from draws of the four posteriors pushed through the enumeration.

Building blocks. The MAP states and covariances of equations (13.12) and (13.11), $\hat{\xi} = -3.4761$ with $\sigma_\xi = 0.5862$ in case A and $\hat{\mathbf{z}}_{\text{B}}$, $\mathbf{\Sigma}_{\text{B}}$ in case B; the enumeration at $\hat{\mathbf{z}}$, equations (13.14) to (13.16); the sum (13.26); $\mathbf{D}$ and $\mathbf{I}_{\text{B}}$ of equation (13.25), for the delta band; and, for the exact values beside them, the posterior parameters (a_i, b_i), the Beta moments (13.29), the Beta quantiles for the quantile map, and draws of the four posteriors of case B.

The enumeration in case B. Table 13.5 enumerates the sixteen configurations at the point to which $\hat{\mathbf{z}}_{\text{B}}$ maps, $\bar{\mathbf{q}} = (0.0444,\ 0.0222,\ 0.0667,\ 0.0222)$, the plug-in value of case B: each row is one configuration, the columns give the four statuses, the number of units up, the structure function and the probability (13.14). One row of the arithmetic, units 1, 2 and 4 up and unit 3 down: $(1-0.04444)(1-0.02222)(0.06667)(1-0.02222) = 0.95556\times 0.97778\times 0.06667\times 0.97778 = 0.06090$. The five configurations with three or four units up sum to $R_{\text{system}}(\bar{\mathbf{q}}) = 0.991974$, and the eleven others to $U_{\text{system}} = 8.026\times 10^{-3}$, of which the six configurations with two units down carry 7.844×10^{-3}, the four with three down 1.81×10^{-4} and the one with all four down 1.5×10^{-6}.

The state-route expectation is the same enumeration at the means of the Laplace normal. The sum (13.26) with forty nodes gives $\bar{\mathbf{q}}^{\text{L}} = (0.0551,\ 0.0347,\ 0.0764,\ 0.0347)$, and the enumeration at that point gives

$$\mathbb{E}[U_{\text{system}} \mid \mathbf{y}] \approx 1.36\times 10^{-2} \qquad \text{(case B)},$$

the predictive probability that the cooling is lost at a random instant; with the power feed, whose Laplace mean is 0.00256, $1 - 0.99744\times(1-0.01365) = 1.62\times 10^{-2}$. The value beside it comes from the exact posterior: by equation (13.27) the value at the posterior means is the exact expectation,

$$\mathbb{E}[U_{\text{system}} \mid \mathbf{y}] = U_{\text{system}}(\bar{\mathbf{q}}) = 8.03\times 10^{-3} \qquad \text{(case B)}, \tag{13.35}$$

and the same holds for the hall with its power feed, since all five entries are independent: $1 - 0.998\times 0.991974 = 1.001\times 10^{-2}$. The gap on the cooling is 5.6×10^{-3}, a factor of 1.70. It comes from few findings: units 2 and 4, with no finding, have the Laplace mean 0.0347 against the exact 0.0222.

The enumeration in case A. With one shared unavailability, every configuration with m units down has the same probability $q^m(1-q)^{4-m}$, and the sixteen configurations fall into five groups of 1, 4, 6, 4 and 1. Figure 13.2 draws them at the plug-in value $q = 0.03$. Each row of four squares is one configuration, units 1 to 4 from left to right, green for up and red for down;

Table 13.5: The sixteen configurations of the four CRAC units, case B at the posterior means $\bar{q}_i$: status 1 up, 0 down; $\phi = 1$ when at least three units are up. Probabilities by equation (13.14).

No.	s_1	s_2	s_3	s_4	Units up	$\phi(\mathbf{s})$	$\Pr(\mathbf{s} \mid \bar{\mathbf{q}})$
1	1	1	1	1	4	1	0.8527
2	0	1	1	1	3	1	0.03966
3	1	0	1	1	3	1	0.01938
4	1	1	0	1	3	1	0.06090
5	1	1	1	0	3	1	0.01938
6	0	0	1	1	2	0	9.013×10^{-4}
7	0	1	0	1	2	0	2.833×10^{-3}
8	0	1	1	0	2	0	9.013×10^{-4}
9	1	0	0	1	2	0	1.384×10^{-3}
10	1	0	1	0	2	0	4.404×10^{-4}
11	1	1	0	0	2	0	1.384×10^{-3}
12	0	0	0	1	1	0	6.438×10^{-5}
13	0	0	1	0	1	0	2.048×10^{-5}
14	0	1	0	0	1	0	6.438×10^{-5}
15	1	0	0	0	1	0	3.146×10^{-5}
16	0	0	0	0	0	0	1.46×10^{-6}
Sum over $\phi = 1$ (rows 1 to 5): $R_{\text{system}}(\bar{\mathbf{q}})$							0.991974
Sum over $\phi = 0$ (rows 6 to 16): $U_{\text{system}}(\bar{\mathbf{q}})$							8.026×10^{-3}

the five columns hold the configurations with zero to four units down. The green panel on the left contains the configurations that deliver the service and the red panel the eleven that do not. Beneath each column are the probability of each of its configurations and the column's total. By equation (13.20) the red panel is $0.0054 - 0.000216 + 0.0000024 = 5.186 \times 10^{-3}$. The column of one unit down carries 0.1095, twenty times the whole red panel.

Figure 13.3 draws the unavailabilities of the four-unit structures against q, both axes logarithmic. The horizontal axis is the unit unavailability and the vertical axis the service unavailability; the dashed gray line is one unit alone, $U = q$; the four colored curves are 4-out-of-4, 3-out-of-4, 2-out-of-4 and 1-out-of-4; the dotted curve is the hall, 3-out-of-4 in series with the power feed at its posterior mean $q_{\text{p}} = 0.002$. The shaded strip is the central 95 % interval of case A's posterior on q and the vertical line its mean. Each curve has a slope near its order, 1, 2, 3 and 4. Across the shaded strip, a factor of 11 in q, the 3-out-of-4 unavailability moves by a factor of $11^2 \approx 120$. The hall's curve leaves the 3-out-of-4 curve near $q = 0.019$, where the cooling unavailability falls below the feed's 0.002, and is flat at 0.002 below it.

The expectation in case A. Table 13.6 evaluates equation (13.26) with five nodes. Each row is a node: the first two columns give the node t_m and its weight $w_m/\sqrt{\pi}$, the third the state $\xi_m = \hat{\xi} + \sqrt{2}\,\sigma_\xi t_m$ at which the sum looks, the fourth the unavailability there, and the last three the unavailabilities of the structures this chapter asks about, from Table 13.4. The middle node is the MAP state itself.

The weighted sum of the 3-out-of-4 column is

$$\mathbb{E}[U_{\text{system}} \mid \mathbf{y}] \approx 0.01126 \times 1.976 \times 10^{-4} + \cdots + 0.01126 \times 9.891 \times 10^{-2} = 9.15 \times 10^{-3},$$

and forty nodes give 9.145×10^{-3} also (3.95×10^{-4} and 9.28×10^{-4} for the other two columns). The two upper nodes carry 67 % of the sum: the expectation is set by values of q above the MAP state's, which is why the shape of the posterior there matters.

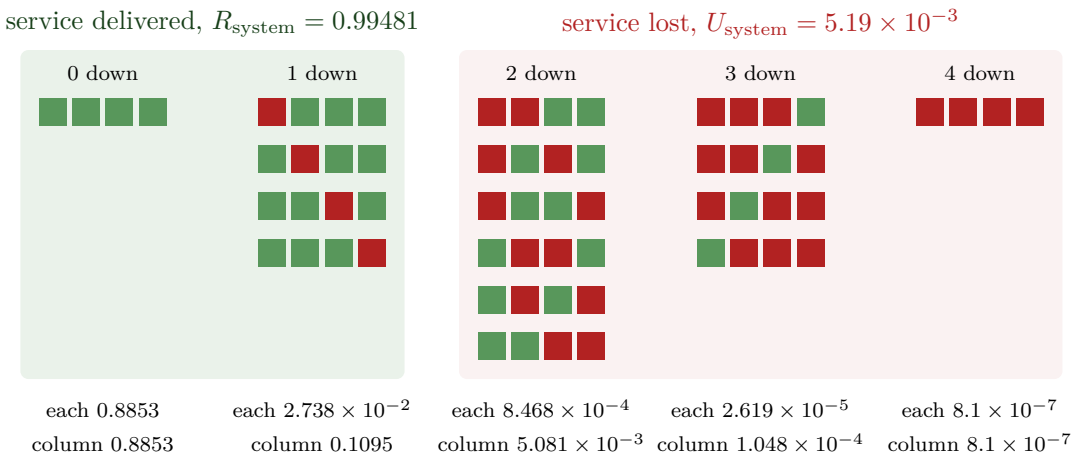


Figure 13.2: The sixteen configurations of the four CRAC units under a shared unavailability at its plug-in value $q = 0.03$. Each row of four squares is one configuration, units 1 to 4 from left to right, green up and red down; columns group the configurations by the number of units down, $\binom{4}{m}$ of them each with probability $q^m(1-q)^{4-m}$. Green panel: $\phi = 1$ under 3-out-of-4; red panel: $\phi = 0$.

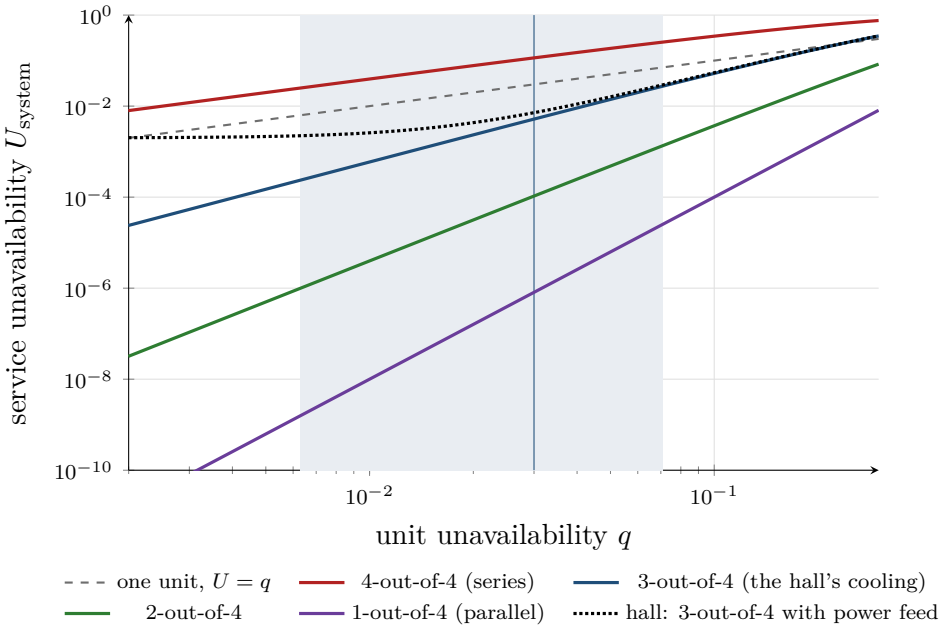


Figure 13.3: Service unavailability of four-unit structures against the unit unavailability, equation (13.19) with $n = 4$: $1-(1-q)^4$, $6q^2-8q^3+3q^4$, $4q^3-3q^4$ and q^4; dashed, $U = q$; dotted, the hall $1-(1-q_{\text{p}})(1-6q^2+8q^3-3q^4)$ with $q_{\text{p}} = 0.002$. Shaded: the 95 % posterior interval $[0.0063,\ 0.0711]$ of q in case A; vertical line: its mean 0.03.

Table 13.6: The Gauss–Hermite rule (13.26) with five nodes over the Laplace posterior $\xi \mid \mathbf{y} \approx \mathcal{N}(-3.4761,\ 0.5862^2)$ of the fleet entry of case A.

Node t_m	Weight $w_m/\sqrt{\pi}$	State ξ_m	$q(\xi_m)$	3-out-of-4	2-out-of-4	3-out-of-5
-2.0202	0.01126	-5.1509	0.00576	1.976×10^{-4}	7.61×10^{-7}	1.90×10^{-6}
-0.9586	0.22208	-4.2708	0.01378	1.118×10^{-3}	1.04×10^{-5}	2.56×10^{-5}
0	0.53333	-3.4761	0.03000	5.186×10^{-3}	1.06×10^{-4}	2.58×10^{-4}
0.9586	0.22208	-2.6814	0.06408	2.258×10^{-2}	1.00×10^{-3}	2.38×10^{-3}
2.0202	0.01126	-1.8013	0.14169	9.891×10^{-2}	1.02×10^{-2}	2.27×10^{-2}
Weighted sum				9.145×10^{-3}	3.96×10^{-4}	9.29×10^{-4}

The value beside it comes from the exact posterior. By equation (13.29), for Beta$(3, 97)$,

$$\mathbb{E}[q^2] = \frac{3 \times 4}{100 \times 101} = 1.1881 \times 10^{-3},$$
$$\mathbb{E}[q^3] = \mathbb{E}[q^2]\,\frac{5}{102} = 5.8241 \times 10^{-5},$$
$$\mathbb{E}[q^4] = \mathbb{E}[q^3]\,\frac{6}{103} = 3.3927 \times 10^{-6},$$

and equation (13.30) gives

$$\mathbb{E}[U_{\text{system}} \mid \mathbf{y}] = 7.1287 \times 10^{-3} - 4.659 \times 10^{-4} + 1.02 \times 10^{-5} = 6.67 \times 10^{-3},$$

against the plug-in value $U_{\text{system}}(0.03) = 5.19 \times 10^{-3}$. The state route is 2.47×10^{-3} above the exact value, a factor of 1.37, because a normal in the logit gives more weight to large q than Beta$(3, 97)$ does: its mean of q is 0.0350, not 0.0300. The plug-in understates the expected cooling unavailability by 43 % against the state route and by 22 % against the exact value; the exact gap is $6.673 \times 10^{-3} - 5.186 \times 10^{-3} = 1.487 \times 10^{-3}$, and the state route's 3.96×10^{-3}.

The band in case A. Equation (13.33) gives the gradients of the logarithm in the two cases,

$$\mathbf{g}_{\ln U,\mathrm{A}} = \begin{bmatrix} 1.9005 \\ 0 \end{bmatrix}, \qquad \mathbf{g}_{\ln U,\mathrm{B}} = \begin{bmatrix} 0.552 \\ 0.332 \\ 0.651 \\ 0.332 \\ 0 \end{bmatrix},$$

whose entries sum to 1.90 in case A and 1.87 in case B, close to the exponent 2 of the structure's order. In case A, $\sigma_{\ln U} = 1.9005 \times \sqrt{0.34364} = 1.11$, and the band is

$$U_{\text{system}} \in 5.19 \times 10^{-3} \times \mathrm{e}^{\pm 1.96 \times 1.114} = [5.8 \times 10^{-4},\ 4.6 \times 10^{-2}],$$

positive, skewed and a factor of 79 wide, with the median 5.19×10^{-3}. The value beside it comes from the exact posterior. By the quantile map (13.34), with the Beta quantiles $q_{0.025} = 0.00629$, $q_{0.5} = 0.02692$ and $q_{0.975} = 0.07108$ of Beta$(3, 97)$, the band of the cooling unavailability is $[2.36 \times 10^{-4},\ 2.75 \times 10^{-2}]$ with median 4.19×10^{-3}; one row of that arithmetic, $U_{\text{system}}(0.07108) = 6 \times 0.005052 - 8 \times 0.000359 + 3 \times 0.0000255 = 0.02751$. The exact standard deviation, from the moments of q up to order eight, is $\mathrm{sd}(U_{\text{system}} \mid \mathbf{y}) = 7.57 \times 10^{-3}$. The state-route band lies above the exact one, by a factor of 2.5 at the lower end and 1.7 at the upper.

Figure 13.4 draws the posterior of the case-A cooling unavailability. The horizontal axis is the unavailability on a logarithmic scale and the vertical axis the posterior density per unit of

$\ln U$, so that equal areas are equal probabilities. The dotted blue curve is the Laplace density, the normal of the delta method on $\ln U$; the green curve is the exact density, the Beta density of q carried through the quantile map; the gray steps are a histogram of 4×10^5 draws of the exact posterior. The shaded area left of the red line is the state route's probability that the unavailability is below 10^{-3}. The vertical lines mark the exact median, the plug-in value, which is the Laplace median, and the two means.

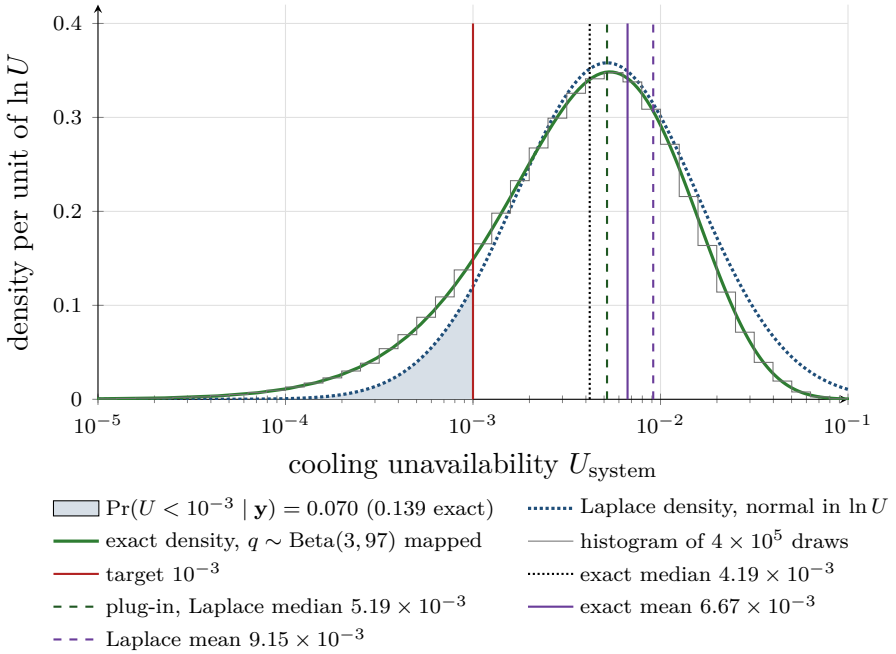


Figure 13.4: Posterior of the cooling unavailability in case A, per unit of $\ln U$. Dotted blue: the Laplace density, $\ln U \approx \mathcal{N}(\ln(5.186 \times 10^{-3}),\ 1.114^2)$, which the answers report. Green, the exact density: with $U(q) = 6q^2 - 8q^3 + 3q^4$, it is $p(q)\,U(q)/U'(q) = 39\,212\,q^3(1-q)^{94}(6 - 8q + 3q^2)$, drawn parametrically in q, with $p(q) = 470\,547\,q^2(1-q)^{96}$ the Beta$(3, 97)$ density. Histogram: 4×10^5 draws of q (seed 13), forty bins equal in $\ln U$. Shaded: $U < 10^{-3}$ under the Laplace density; exactly, $U < 10^{-3}$ is $q < 0.01302$.

On the logarithmic axis the posterior is close to symmetric, with a longer lower tail; the median, the plug-in value and the mean stand a factor of 1.6 apart from first to last, and the band spans more than two decades. The Laplace curve is the same shape moved to the right, its median at the plug-in value and its mean 1.76 times that.

The band in case B. With the gradient of the logarithm above and the diagonal $\mathbf{\Sigma}_\mathrm{B}$, the budget of four terms gives $\sigma_{\ln U} = 0.73$ about the plug-in value 8.03×10^{-3}, and the band $[1.9 \times 10^{-3},\ 3.4 \times 10^{-2}]$. The value beside it comes from the exact posterior. The 4×10^5 draws of the four posteriors pushed through the enumeration give the median 6.47×10^{-3} and the band $[1.15 \times 10^{-3},\ 2.38 \times 10^{-2}]$ about the exact mean 8.03×10^{-3}; the exact standard deviation, 6.05×10^{-3}, follows from the second moments of the four posteriors summed over pairs of configurations.

In hours. At $8760\,\mathrm{h} \times U_\mathrm{system}$, the cooling of the hall is lost for a median of 45 h per year and a mean of 80 h, with a 95 % band from 5 h to 403 h (exactly, a median of 37 h, a mean of 58 h and a band from 2 h to 241 h). In case B the same quantities are a median of 70 h, a mean of 120 h and a band from 17 h to 295 h (exactly, 57 h, 70 h, and 10 h to 209 h).

13.3.2 The N+1 target

Engineering question. Does N+1 give the hall a cooling unavailability below 10^{-3}, with assurance? The design basis asks for less than nine hours a year without cooling, shown with the assurance $\kappa = 0.90$.

General form. A prediction under a threshold is a tail probability of a derived quantity under the posterior, compared with an assurance level, equation (2.12):

$$\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}) = \int \mathbf{1}\big[f(\mathbf{z}) \geq f^*\big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \qquad \text{accept if } \Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}) \geq \kappa.$$

Specialized form. One choice enters here: the target is read on the long-run fraction of time without cooling, as the design basis states it. That makes the event a set of values of the unit unavailabilities; a target on how long a single loss of cooling lasts, or how often one occurs, would need the rates of Chapter 8 (Section 13.6.5).

The question is equation (2.12) with $f(\mathbf{z}) = -U_{\text{system}}(\mathbf{q}(\mathbf{z}))$, $f^* = -u^*$, $u^* = 10^{-3}$, since the requirement is an upper limit, and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal of step 3. The event is $U_{\text{system}} < u^*$, the requirement holding, and on the logarithm, which is normal under the delta method (13.33), its probability is a normal tail,

$$\Pr(U_{\text{system}} < u^* \mid \mathbf{y}) \approx \Phi\left(\frac{\ln u^* - \ln U_{\text{system}}(\hat{\mathbf{z}})}{\sigma_{\ln U}}\right), \tag{13.36}$$

with Φ the standard normal cumulative distribution. The value beside it comes from the exact posterior. With one shared entry the unavailability increases with q, so the event is $q < q^*$ with q^* the root of $U_{\text{system}}(q^*) = u^*$, and by the quantile map (13.34)

$$\Pr(U_{\text{system}} < u^* \mid \mathbf{y}) = \Pr(q < q^* \mid \mathbf{y}) = I_{q^*}(a, b), \qquad 6q^{*2} - 8q^{*3} + 3q^{*4} = u^*, \tag{13.37}$$

the incomplete Beta function at the root of a quartic. The target asks every unit's unavailability to be below q^*, a unit target derived from the system target through the enumeration. With independent entries the event is a region of four unavailabilities, not an interval of one, and its probability is the fraction of draws of the four posteriors whose enumerated unavailability lies below u^*.

Building blocks. The plug-in values at $\hat{\mathbf{z}}$ and the spreads $\sigma_{\ln U} = 1.114$ (case A) and 0.73 (case B) of Section 13.3.1; the normal cumulative distribution Φ; and, for the exact values beside them, the root q^* of the quartic, the Beta parameters (a, b) of case A with the incomplete Beta function, and draws of the four posteriors of case B through the enumeration.

The N+1 target for the hall. By equation (13.36) with $\ln(10^{-3}) = -6.9078$ and $\ln(5.186 \times 10^{-3}) = -5.2617$,

$$\Pr(U_{\text{system}} < 10^{-3} \mid \mathbf{y}) \approx \Phi\left(\frac{-6.9078 + 5.2617}{1.1141}\right) = \Phi(-1.4775) = 0.070 \qquad \text{(case A)},$$

the shaded area of Figure 13.4. For a target of 10^{-4} it is 0.0002. In case B, $\Phi(-2.8437) = 0.002$. All three are far below the assurance of 0.90, and so are the exact values. The root of $6q^{*2} - 8q^{*3} + 3q^{*4} = 10^{-3}$ is $q^* = 0.01302$, and by equation (13.37)

$$\Pr(U_{\text{system}} < 10^{-3} \mid \mathbf{y}) = I_{0.01302}(3, 97) = 0.139 \qquad \text{(case A)},$$

0.069 higher than the state route. For a target of 10^{-4}, $q^* = 0.004094$ and the probability is 0.008. In case B the fraction of the draws below 10^{-3} is 0.018. In the logit the normal gives too little weight to the small q where the event holds, so the state route is lower in all three.

13.3.3 The pooled and the separate records

Engineering question. Is the answer the same when the units' records are pooled and when they are kept apart? The operator holds a hundred checks of four units and may treat them as one fleet record or as four unit records.

General form. The question compares the expectation of one derived quantity, equation (2.10), under two posteriors of the state built from the same checks:

$$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}] = \int f(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}.$$

Specialized form. Two choices enter here. Pooled, the hundred checks are treated as checks of one unit: this is what lets them act as one record on one entry, and it fails if the units differ, which the four records test only weakly (Section 13.4.4). Kept apart, each unit carries its own copy of the prior Beta$(1, 19)$: this lets a short record borrow strength from the prior, and it counts the prior four times against the hundred checks, which Section 13.4.3 weighs against other priors.

The question is equation (2.10) with $f(\mathbf{z}) = U_{\text{system}}$ under 3-out-of-4 and $p(\mathbf{z} \mid \mathbf{y})$, in turn, the Laplace normal of each of four posterior structures for the four units, each expectation by the sum (13.26), with the enumeration where the entries are independent. The first pair gives every unit the same marginal posterior, case A's, and differs only in whether the four unavailabilities are one entry or four independent entries: by equation (13.27) the expectation under independent entries is the value at the means, and on the exact route that is the plug-in value; by equation (13.30) the expectation under the shared entry exceeds it by the Jensen gap. The pair isolates the sharing. The second pair holds case B's records, kept apart as in case B, each unit with its own copy of the unit prior, or pooled into one entry with one copy of that prior, the update (5.24) applied to the hundred checks at once. The two treatments then differ in three ways. The shared entry raises the pooled answer above its plug-in value by the Jensen gap. Four copies of the prior, each worth twenty checks, pull each unit's mean toward 0.05 more than one copy pulls the fleet's. And the four means are unequal, which lowers the separate answer: the leading term of the unavailability is the sum of the pairwise products,

$$\sum_{i<j} q_i q_j = \tfrac{1}{2}\left[(\mathbf{1}^{\mathsf{T}}\mathbf{q})^2 - \|\mathbf{q}\|^2\right], \tag{13.38}$$

and for a fixed $\mathbf{1}^{\mathsf{T}}\mathbf{q}$ it is largest when $\|\mathbf{q}\|$ is smallest, at equal q_i.

The fit objects show where the difference lives. The shared entry is the separate state restricted to the direction $\mathbf{1}$ of the four unit columns: its Jacobian column carries the four records at once, its gradient is $\mathbf{1}^{\mathsf{T}}\mathbf{g}$ and its information $\mathbf{1}^{\mathsf{T}}\boldsymbol{\Sigma}^{-1}\mathbf{1}$, and with equal entries $(\mathbf{1}^{\mathsf{T}}\mathbf{g})^2/(\mathbf{1}^{\mathsf{T}}\boldsymbol{\Sigma}^{-1}\mathbf{1}) = \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$. The delta method therefore gives the same spread under both treatments and cannot see the difference between them; the difference is in the curvature of the enumeration in the shared entry, the Jensen gap of equation (13.31), as Exercise 13.5 works out for a pump set.

The shared entry correlates the units. Given q two units are down together with probability q^2; averaged over the posterior, with probability $\mathbb{E}[q^2 \mid \mathbf{y}] = \bar{q}^2 + \mathrm{Var}(q \mid \mathbf{y})$, more than $\bar{q}^2$. If the fleet unavailability is high, two or more units are often down together; if it is low, rarely; and a redundant structure is lost only through coincidences. The four statuses, independent for a known q, are positively correlated once q is uncertain, with the correlation $\mathrm{Var}(q \mid \mathbf{y})/(\bar{q}(1-\bar{q}))$ between the down-indicators of any two units, as equation (5.54) gave for two future tests of one valve. Redundancy protects against the independent outage of one unit, not against the shared uncertainty of their common unavailability. Whether the four units share one unavailability, have four unrelated ones, or have unavailabilities that differ around a fleet value

is a hypothesis comparison between models, which a later chapter on model selection treats through the evidence of each; here the two extremes are computed side by side.

Building blocks. The MAP states and variances of the structures: $(-3.4761,\ 0.34364)$ for Beta$(3, 97)$ and, by equation (13.8), $(\ln(4/116),\ 120/464) = (-3.3673,\ 0.25862)$ for Beta$(4, 116)$, with $\hat{\mathbf{z}}_{\mathrm{B}}$ and $\boldsymbol{\Sigma}_{\mathrm{B}}$ for the records kept apart; the sum (13.26); the plug-in values of equation (13.20); and, for the exact values beside them, the moments (13.29) of the two Beta posteriors and the enumeration at the four means of case B, exact by equation (13.27).

The four posterior structures of the hall. Table 13.7 sets the four posterior structures of Section 13.3.3 side by side for the cooling of the hall. Each row is a posterior structure for the four units; the columns give the posterior, the plug-in value at the MAP state, the state-route expectation with the exact one beside it, their ratios to the plug-in value, and the route by which the exact value is obtained. The first two rows share the marginal Beta$(3, 97)$; the last two hold case B's records, kept apart with four copies of the unit prior or pooled with one, Beta$(4, 116)$. One copy is worth twenty checks, four copies eighty.

Table 13.7: Expected cooling unavailability of the hall under 3-out-of-4 for shared and independent posteriors. Laplace: the sum (13.26) over the normal in the logits, with the enumeration where the entries are independent. Exact: shared by the moments, equation (13.29); independent by multilinearity, equation (13.27).

Posterior structure	Posterior	Plug-in	Laplace $\mathbb{E}[U]$	Exact $\mathbb{E}[U]$	Ratio, Laplace (exact)	Exact route
One shared entry (case A)	$q \sim$ Beta$(3, 97)$	5.19×10^{-3}	9.15×10^{-3}	6.67×10^{-3}	1.76 (1.29)	moments of q
Four independent entries, same marginals	$q_i \sim$ Beta$(3, 97)$ each	5.19×10^{-3}	6.99×10^{-3}	5.19×10^{-3}	1.35 (1.00)	plug-in is exact
Records kept apart (case B)	Table 13.3	8.03×10^{-3}	1.36×10^{-2}	8.03×10^{-3}	1.70 (1.00)	plug-in is exact
Same records pooled, one copy of the unit prior	$q \sim$ Beta$(4, 116)$	6.37×10^{-3}	9.76×10^{-3}	7.73×10^{-3}	1.53 (1.21)	moments of q

The three effects. Same marginals: the shared entry raises the state-route expectation from 6.99×10^{-3} to 9.15×10^{-3}, a factor of 1.31, and the exact one by 1.29, so the two routes agree on what the sharing does. The plug-in value is 43 % low under the shared entry on the state route; exactly it is 22 % low, a ratio of 1.29. Records held: pooled 9.76×10^{-3}, apart 1.36×10^{-2}, 40 % more; exactly, pooled 7.73×10^{-3}, apart 8.03×10^{-3}, 4 % more. Both routes put the records kept apart higher, by amounts ten times different: the Laplace means of units with no finding stand far above their exact values, and a pooled entry with four findings has less of that. The four means average 0.0502 against the pooled 0.0374 on the state route, and at an equal 0.0502 the cooling would be 1.41×10^{-2}; exactly, the four means average 0.0389 against the pooled 0.0333, and at an equal 0.0389 the plug-in value would be 8.61×10^{-3}.

The coupling of the units. Under the Laplace normal of the shared entry, the sum (13.26) with $h = q^2$ gives $\mathbb{E}[q^2] = 1.660 \times 10^{-3}$, the probability that two units are down together, 84 % more than $\hat{q}^2 = 9.0 \times 10^{-4}$; with the mean 0.0350 and the variance 4.39×10^{-4} the correlation between the down-indicators of any two units is 0.0130. The value beside it comes from the exact

posterior. Under Beta(3, 97), two units are down together with probability $\mathbb{E}[q^2] = 1.188 \times 10^{-3}$, which is 32 % more than $\bar{q}^2 = 9.0 \times 10^{-4}$. With $\mathrm{Var}(q \mid \mathbf{y}) = 3 \times 97/(100^2 \times 101) = 2.881 \times 10^{-4}$ the correlation between the down-indicators of any two units is $\mathrm{Var}(q)/(\bar{q}(1-\bar{q})) = 0.0099$.

13.3.4 The power feed and the importance of each unit

Engineering questions. How much does the power feed add, in series? Which unit matters most? The first asks how the hall's unavailability splits between its feed and its cooling, in the mean and in the uncertainty; the second which unit's unavailability moves the service most.

General form. Both are estimations of derived quantities, the terms of a sum and the entries of a gradient, through their posterior moments, equation (2.10), with the quantiles of equation (2.11) for a band:

$$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}] = \int f(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \qquad \mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y}) = \int \big(f(\mathbf{z}) - \mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]\big)^2\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}.$$

Specialized form. Two choices enter here. The feed and the cooling units share no row and no cause: this is what makes the hall's availability the product of the two and its expectation exact in the two means, and it fails if a disturbance on the feed trips the units as well, a common cause treated in a later chapter of this part. The importance is read as the gradient of the service at the MAP state, with its band from the posterior: this answers which unit's improvement moves the service most, and it does not say which unit's record limits the band, which the variance budget answers; the bands and the ranking from draws (Section 13.3.4) test how far the reading at one point holds.

The two questions are equations (2.10) and (2.11) with $f(\mathbf{z})$ the terms of the hall's unavailability and then the importance of each unit, and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal of step 3.

The power feed in series. The hall is delivered when its power feed is up and its cooling delivers, so its structure function is the product $s_{\mathrm{p}}\, \phi_{\mathrm{c}}(\mathbf{s})$ of the feed's status and the cooling's structure function. Summing over the feed's status first and using the independence of the feed from the units,

$$R^{\mathrm{hall}}_{\mathrm{system}} = (1 - q_{\mathrm{p}})\, R^{\mathrm{c}}_{\mathrm{system}}, \qquad U_{\mathrm{series}} = 1 - (1 - q_{\mathrm{p}})(1 - U_{\mathrm{c}}) = q_{\mathrm{p}} + U_{\mathrm{c}} - q_{\mathrm{p}} U_{\mathrm{c}}, \tag{13.39}$$

with U_{c} the cooling unavailability. A series of subsystems is a series of units whose unavailabilities are the subsystems' unavailabilities, equation (13.17). For L subsystems in series with unavailabilities $U_1, \ldots, U_L$, expand the product of equation (13.17) by powers:

$$1 - \prod_{l=1}^{L} (1 - U_l) = \sum_{l} U_l - \sum_{l<l'} U_l U_{l'} + \sum_{l<l'<l''} U_l U_{l'} U_{l''} - \cdots .$$

The unavailability of the series is the probability of the union of the subsystem-down events, and the expansion is inclusion and exclusion over that union. Truncated after its first term it is the *rare-event approximation*, and the standard Bonferroni inequalities bound its error between the first two truncations,

$$\sum_{l} U_l - \sum_{l<l'} U_l U_{l'} \;\le\; U_{\mathrm{series}} \;\le\; \sum_{l} U_l, \qquad 0 \le \sum_{l} U_l - U_{\mathrm{series}} \le \sum_{l<l'} U_l U_{l'}. \tag{13.40}$$

The approximation always overstates the unavailability, by at most the sum of the pairwise products; for two subsystems the error is exactly $q_{\mathrm{p}} U_{\mathrm{c}}$, by equation (13.39). Equation (13.39) is multilinear in q_{p} and U_{c}, and the two are independent a posteriori on either route, so by the

argument of equation (13.27) its expectation is its value at the two means, on the state route the sums of equation (13.26) over the two entries. The split in the uncertainty is read off the gradient of the hall in the state $\mathbf{z}_{\mathrm{A}}$. By equation (13.39) the derivative of U_{series} with respect to the cooling is $1-q_{\mathrm{p}}$ and with respect to the feed $1-U_{\mathrm{c}}$, so

$$\mathbf{g}_{\mathrm{hall}} = \begin{bmatrix} (1-q_{\mathrm{p}})\, g_{\mathrm{A},1} \\ (1-U_{\mathrm{c}})\, q_{\mathrm{p}}(1-q_{\mathrm{p}}) \end{bmatrix}, \tag{13.41}$$

with $g_{\mathrm{A},1}$ the cooling's entry of equation (13.32), and on the diagonal $\boldsymbol{\Sigma}_{\mathrm{A}}$ the budget $\mathbf{g}_{\mathrm{hall}}^{\mathsf{T}}\boldsymbol{\Sigma}_{\mathrm{A}}\mathbf{g}_{\mathrm{hall}}$ has one term per entry; divided by U_{series} it gives the band and the tail of the hall by equations (13.33) and (13.36).

The values beside them come from the exact posterior: the means $\mathbb{E}[q_{\mathrm{p}} \mid \mathbf{y}]$ and $\mathbb{E}[U_{\mathrm{c}} \mid \mathbf{y}]$ of equation (13.30), the variances $\mathrm{Var}(q_{\mathrm{p}} \mid \mathbf{y})$, from the Beta variance of Chapter 5, and $\mathrm{Var}(U_{\mathrm{c}} \mid \mathbf{y})$, from the moments of order up to eight, and the band of the series from the cooling's quantile map integrated over the feed's posterior: given $q_{\mathrm{p}} = x$, the event $U_{\mathrm{series}} < u$ is $U_{\mathrm{c}} < t(x) = 1-(1-u)/(1-x)$, so

$$\Pr(U_{\mathrm{series}} < u \mid \mathbf{y}) = \int_0^u I_{q^*(t(x))}(a,b)\; p(x \mid \mathbf{y})\, \mathrm{d}x, \tag{13.42}$$

with $q^*(t)$ the root of $U_{\mathrm{c}}(q) = t$ and $p(x \mid \mathbf{y})$ the Beta posterior density of the feed; the quantiles of the series are the roots of this one-dimensional integral.

The importance of each unit. The gradient (13.24) of the service is the answer to the second question. The Birnbaum importance $I_{\mathrm{B},i}$ is the probability that the rest of the system is in a configuration in which unit i decides whether the service is delivered, a unit then called *critical*. Under 3-out-of-4 unit i is critical when exactly two of the other three units are up, that is when exactly one of them is down:

$$I_{\mathrm{B},i} = \sum_{l \neq i} q_l \prod_{l' \neq i,l} (1-q_{l'}). \tag{13.43}$$

For equal unavailabilities each unit has $I_{\mathrm{B}} = 3q(1-q)^2$, and the four together give $12q(1-q)^2$, which is the derivative of equation (13.20), $12q-24q^2+12q^3$, as the chain rule requires. Its family of related measures, and their derivation for decisions, belong to a later chapter on importance measures. The importance (13.43) is itself multilinear in the other units' unavailabilities, so under independent posteriors its expectation is its value at the means, by the argument of equation (13.27): at the Laplace means $\bar{\mathbf{q}}^{\mathrm{L}}$ on the state route. On the exact route those means are the point to which the MAP state maps (Section 13.2), and the importance at the MAP state is the exact expected importance.

The importance depends on the other units, not on the unit itself. Its derivative with respect to the state has column i equal to $\partial I_{\mathrm{B},i}/\partial \mathbf{z}$, which is $\mathbf{D}$ times column i of the matrix $\mathbf{H}_q$ of mixed derivatives of equation (13.28), over the four unit entries,

$$\frac{\partial \mathbf{I}_{\mathrm{B}}^{\mathsf{T}}}{\partial \mathbf{z}} = \mathbf{D}\,\mathbf{H}_q, \qquad (\mathbf{H}_q)_{ii} = 0, \tag{13.44}$$

with the feed's row zero. The diagonal is zero by multilinearity: the importance of a unit does not depend on its own unavailability at all. It is the probability that the others leave it critical, and it moves with the other units' logits. The importances are derived quantities of the state and are as uncertain as the unavailabilities of the other units. Under the Laplace posterior the band of $\ln I_{\mathrm{B},i}$ has the gradient column i of equation (13.44) divided by $I_{\mathrm{B},i}$, and its spread is that gradient against $\boldsymbol{\Sigma}_{\mathrm{B}}$. The ranking of the units is a probability, the fraction of draws of the Laplace normal in which a unit holds a given place. Exactly, the distribution of the importances is that of equation (13.43) under draws of the four posteriors, and the ranking is the same fraction over those draws.

A second ranking answers a second question. The term of unit i in the variance budget $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ of equation (13.32) is $g_i^2\Sigma_{ii}$, its importance scaled by D_{ii} and squared, times the variance of its logit: it says which unit's record limits the band, not which unit's unavailability moves the service. The Laplace variance $D_{ii}^2\Sigma_{ii} = a_ib_i/(a_i+b_i)^3$ of q_i exceeds the exact Beta variance $a_ib_i/((a_i+b_i)^2(a_i+b_i+1))$ by the factor $(a_i+b_i+1)/(a_i+b_i)$, so the shares do not change when that factor is common to the units.

Building blocks. The MAP states and covariances $(\hat{\mathbf{z}}_{\mathrm{A}}, \boldsymbol{\Sigma}_{\mathrm{A}})$ and $(\hat{\mathbf{z}}_{\mathrm{B}}, \boldsymbol{\Sigma}_{\mathrm{B}})$ of step 3, with the feed's $\hat{\xi}_{\mathrm{p}} = -6.2126$ and $\Sigma_{\mathrm{pp}} = 0.50100$; the sum (13.26); the gradient (13.41); the enumeration with one unit held up and held down, equation (13.24); $\mathbf{D}$ and $\mathbf{H}_q$ at $\hat{\mathbf{z}}_{\mathrm{B}}$; draws of the Laplace normal for the ranking; and, for the exact values beside them, the Beta means and variance of the feed, the moments of q up to order eight, the root $q^*(t)$ with a one-dimensional quadrature for the band (13.42), and draws of the four posteriors for the bands and the ranking.

The power feed in series. By equation (13.39) at the two Laplace means, 0.00256 for the feed from the sum (13.26) and 9.145×10^{-3} for the cooling,

$$\mathbb{E}[U_{\mathrm{series}} \mid \mathbf{y}] \approx 1 - (1 - 0.00256)(1 - 9.145 \times 10^{-3}) = 1.17 \times 10^{-2} \qquad \text{(case A)},$$

102 h a year. The value beside it, at the two exact means, is

$$\mathbb{E}[U_{\mathrm{series}} \mid \mathbf{y}] = 1 - (1 - 0.002)(1 - 6.673 \times 10^{-3}) = 8.66 \times 10^{-3},$$

76 h a year, 3.0×10^{-3} lower. Table 13.8 splits both: each row is a term of equation (13.39), and the columns give its posterior mean, its share of the mean, its posterior variance and its share of the variance, each on the state route with the exact value in parentheses. The rare-event sum overstates the state route's value by the product term 2.3×10^{-5}, 0.20 %, and exactly $0.002 + 6.673 \times 10^{-3} = 8.673 \times 10^{-3}$ overstates the exact value by the product term 1.3×10^{-5}, 0.15 %, both inside the bound (13.40).

With $\hat{q}_{\mathrm{p}} = 0.002$ and $U_{\mathrm{c}} = 5.19 \times 10^{-3}$ at $\hat{\mathbf{z}}_{\mathrm{A}}$, the gradient (13.41) and its budget terms are

$$\mathbf{g}_{\mathrm{hall}} = \begin{bmatrix} 9.837 \times 10^{-3} \\ 1.986 \times 10^{-3} \end{bmatrix}, \qquad \left(g_{\mathrm{hall},i}^2 \Sigma_{\mathrm{A},ii}\right)_i = \begin{bmatrix} 3.33 \times 10^{-5} \\ 1.98 \times 10^{-6} \end{bmatrix},$$

so the feed carries 6 % of the variance. Divided by the plug-in $U_{\mathrm{series}} = 7.18 \times 10^{-3}$, the budget gives $\sigma_{\ln U} = 0.83$, the band $[1.42 \times 10^{-3},\ 3.63 \times 10^{-2}]$ with that median, and by equation (13.36) the probability $\Phi(-2.3827) = 0.009$ of an unavailability below 10^{-3}. The value beside them comes from the exact posterior: the band of the series by equation (13.42) is $[1.28 \times 10^{-3},\ 2.96 \times 10^{-2}]$ with median 6.36×10^{-3}, and the hall's probability of an unavailability below 10^{-3} is 0.013. In hours per year the hall is down for a median of 63 h (exactly, a mean of 76 h and a median of 56 h).

The feed's 95 % band is a factor of 16 wide on the state route, the cooling's a factor of 79; exactly, 23 and 117. On both routes the band of the hall is the band of its cooling.

The importance at the MAP state. The importance of a case-B unit at $\hat{\mathbf{z}}_{\mathrm{B}}$ is equation (13.43) at $\bar{\mathbf{q}}$, which on the exact route is its posterior expectation. At the Laplace means $\bar{\mathbf{q}}^{\mathrm{L}}$ the same equation gives the state-route expectation, (0.133, 0.149, 0.115, 0.149), higher throughout and in the same order. For unit 3, the others are down with probabilities 0.0444, 0.0222 and 0.0222, and exactly one of them is down with probability $0.0444 \times 0.9778^2 + 2 \times 0.0222 \times 0.9556 \times 0.9778 = 0.0840$. The four importances are

$$(I_{\mathrm{B},1},\ I_{\mathrm{B},2},\ I_{\mathrm{B},3},\ I_{\mathrm{B},4}) = -\frac{\partial R_{\mathrm{system}}}{\partial \mathbf{q}} = (0.104,\ 0.123,\ 0.084,\ 0.123), \tag{13.45}$$

Table 13.8: The hall's unavailability as a series of its power feed and its cooling, case A, equation (13.39): state route, with the exact value in parentheses. Laplace means by the sum (13.26), variances by the budget of $\mathbf{g}_{\text{hall}}$; exact means, and variances from the Beta posterior of the feed and the exact moments of the cooling.

Term	Posterior mean	Share of mean	Posterior variance	Share of variance
Power feed q_{p}	2.56×10^{-3} (2.00×10^{-3})	22 (23) %	1.98×10^{-6} (1.99×10^{-6})	6 (3) %
Cooling U_{c}, 3-out-of-4	9.15×10^{-3} (6.67×10^{-3})	78 (77) %	3.33×10^{-5} (5.74×10^{-5})	94 (97) %
Product term $-q_{\text{p}}U_{\text{c}}$	-2.3×10^{-5} (-1.3×10^{-5})	-0.20 (-0.15) %		
Hall, U_{series}	1.17×10^{-2} (8.66×10^{-3})	100 %		

against $3q(1-q)^2 = 0.0847$ for each unit under the shared plug-in value 0.03, where the four sum to $12q(1-q)^2 = 0.3387$. In the state, equation (13.25) scales them by $\mathbf{D} = \text{diag}(0.04247,\ 0.02173,\ 0.06222,\ 0.02173,\ 0.001996)$, the slopes of the logistic maps at $\hat{\mathbf{z}}_{\text{B}}$, and gives

$$\mathbf{g}_{\text{B}} = \begin{bmatrix} 4.429 \\ 2.665 \\ 5.228 \\ 2.665 \\ 0 \end{bmatrix} \times 10^{-3}, \qquad \mathbf{g}_{\text{A}} = \begin{bmatrix} 9.857 \\ 0 \end{bmatrix} \times 10^{-3}, \tag{13.46}$$

with $\mathbf{g}_{\text{A}}$ the shared entry's 0.0291×0.3387. The two vectors rank the units differently: in $\mathbf{I}_{\text{B}}$ unit 3 is last, in $\mathbf{g}_{\text{B}}$ first, because a unit with a larger unavailability moves further in q for a unit step in its logit.

Figure 13.5 draws them. Panel (a) has the four units on the horizontal axis and the importance on the vertical axis, with the shared-entry value as a dashed line. Panel (b) has the same axis and the share of each unit's term $g_i^2\Sigma_{ii}$ in the variance budget $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ of equation (13.32), with the equal share of 25 % as a dashed line.

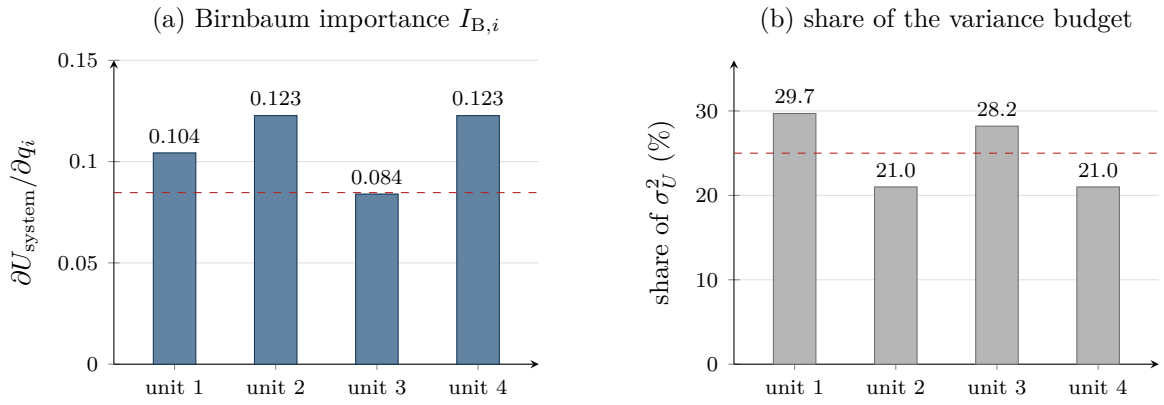


Figure 13.5: The four CRAC units of case B. (a) Birnbaum importance, equation (13.43) at the posterior means; dashed: $3q(1-q)^2 = 0.0847$ at the shared value $q = 0.03$. (b) Share of each unit in $\sigma_U^2 = \sum_i g_i^2\Sigma_{ii} = 3.46 \times 10^{-5}$, with $\mathbf{g}_{\text{B}}$ and $\boldsymbol{\Sigma}_{\text{B}}$ of equations (13.32) and (13.11); dashed: an equal share of 25 %.

The derivative of the importances. Equation (13.44) at $\hat{\mathbf{z}}_\mathrm{B}$ is

$$\frac{\partial \mathbf{I}_\mathrm{B}^\mathsf{T}}{\partial \mathbf{z}} = \mathbf{D}\,\mathbf{H}_q = \begin{bmatrix} 0 & 0.03511 & 0.03876 & 0.03511 \\ 0.01796 & 0 & 0.01890 & 0.01709 \\ 0.05678 & 0.05411 & 0 & 0.05411 \\ 0.01796 & 0.01709 & 0.01890 & 0 \end{bmatrix} \begin{matrix} \xi_1 \\ \xi_2 \\ \xi_3 \\ \xi_4 \end{matrix}. \tag{13.47}$$

Its largest row is the third: the importances move most with ξ_3.

The band on the importances and the ranking. On the logarithmic scale the gradient of $\ln I_{\mathrm{B},i}$ is column i of equation (13.47) divided by $I_{\mathrm{B},i}$, and the products with $\mathbf{\Sigma}_\mathrm{B}$ give spreads of 0.408, 0.364 and 0.463 for unit 1, units 2 and 4, and unit 3, and the 95 % bands [0.047, 0.232], [0.060, 0.250] and [0.034, 0.208]. In 4×10^5 draws of the Laplace normal, seed 13, unit 3 is the least important in 53 % and unit 1 in 25 %; unit 2 or unit 4 is the most important in 86 %; unit 2 is more important than unit 3 in 84 %. The values beside them come from the draws of the exact posteriors, which give the 95 % bands [0.036, 0.199] for unit 1, [0.049, 0.219] for units 2 and 4 and [0.024, 0.173] for unit 3. Unit 3 is the least important in 60 % of the draws and unit 1 in 26 %; unit 2 or unit 4 is the most important in 87 %; unit 2 is more important than unit 3 in 88 %. The two routes rank the units alike.

The budget shares. Units 1 to 4 carry 30, 21, 28 and 21 % of the variance budget (panel (b)), the same under the exact variances since $(a_i + b_i + 1)/(a_i + b_i) = 46/45$ for every unit.

13.3.5 The alternative arrangements

Engineering question. Which arrangement would meet the target? The candidates are the four units with no spare, the present N+1, a fifth unit, and a load that two units can carry.

General form. A design question ranks candidate actions a by the preposterior expectation of a quantity u of the posterior they would produce, equation (2.16):

$$U(a) = \int u\big(p(\mathbf{z} \mid \mathbf{y}, \tilde{y}_a)\big)\, p(\tilde{y}_a \mid \mathbf{y})\, \mathrm{d}\tilde{y}_a.$$

Specialized form. One choice enters here: a fifth unit is taken to be of the same model as the four, sharing their unavailability. This is what lets the present posterior of the one fleet entry carry the fifth unit, and it fails if the new unit is of another type or age, which would need a prior of its own, the pooling question of Section 13.6.5.

The question is equation (2.16) with the action a an arrangement and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal of case A. An arrangement produces no new record, so its outcome $\tilde{y}_a$ is fixed and the integral has one term; with u the indicator that the target is met, $U(a)$ is the posterior probability that arrangement a meets it, by equation (13.36),

$$U(a) = \Pr\big(U_a(q) < u^* \mid \mathbf{y}\big) \approx \Phi\left(\frac{\ln u^* - \ln U_a(\hat{q})}{\sigma_{\ln U,a}}\right), \tag{13.48}$$

with $U_a(q)$ the unavailability polynomial of the arrangement from Table 13.4, each increasing in the shared entry. The same $\hat{\mathbf{z}}_\mathrm{A}$ and the same $\mathbf{\Sigma}_\mathrm{A}$ serve every arrangement; only the derived quantity changes. The arrangements are gradients on one covariance: collecting the gradients of $\ln U_a$ in the state, one column per arrangement, into a matrix $\mathbf{G}$, the log-scale spreads $\sigma_{\ln U,a}$ are the square roots of the diagonal of $\mathbf{G}^\mathsf{T}\mathbf{\Sigma}_\mathrm{A}\mathbf{G}$, for the cooling alone and, with the gradient (13.41) divided by U_series, for the hall. The expectation of each is the sum (13.26) with $h = U_a(q(\xi))$,

and the hall with its feed follows from equation (13.39). The value beside it comes from the exact posterior,

$$U(a) = I_{q_a^*}(a_{\mathrm{A}}, b_{\mathrm{A}}), \qquad U_a(q_a^*) = u^*, \tag{13.49}$$

with $(a_{\mathrm{A}}, b_{\mathrm{A}})$ the parameters of case A's posterior; its expectation is a combination of the moments of equation (13.29) and its band the quantile map (13.34). The higher the order of the arrangement, the higher the power of the shared q in its leading term and the more its expectation exceeds the plug-in value.

Building blocks. The polynomials of Table 13.4; $\hat{\xi} = -3.4761$ and $\sigma_\xi = 0.5862$ with the feed's entry; the matrix $\mathbf{G}$; the sum (13.26); equation (13.39); and, for the exact values beside them, the roots q_a^* and the moments of Beta$(3, 97)$ up to order five with its quantiles.

The arrangements for the hall. Collecting the gradients of $\ln U_{\mathrm{system}}$ in the state, one column per arrangement in the order of Table 13.9, for the cooling alone and for the hall with its feed,

$$\mathbf{G}_{\mathrm{c}} = \begin{bmatrix} 0.926 & 1.901 & 2.865 & 2.888 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \qquad \mathbf{G}_{\mathrm{hall}} = \begin{bmatrix} 0.910 & 1.371 & 0.327 & 0.145 \\ 0.015 & 0.277 & 0.884 & 0.948 \end{bmatrix},$$

and the log-scale spreads of the cooling are the square roots of the diagonal of $\mathbf{G}_{\mathrm{c}}^{\mathsf{T}} \boldsymbol{\Sigma}_{\mathrm{A}} \mathbf{G}_{\mathrm{c}}$: 0.54, 1.11, 1.68 and 1.69. Those of the hall, from the diagonal of $\mathbf{G}_{\mathrm{hall}}^{\mathsf{T}} \boldsymbol{\Sigma}_{\mathrm{A}} \mathbf{G}_{\mathrm{hall}}$, are 0.53, 0.83, 0.65 and 0.68, smaller for the redundant drawings because in the hall a step in either logit moves only its own term's fraction of U, and no entry of $\mathbf{G}_{\mathrm{hall}}$ exceeds 1.4. The first row of $\mathbf{G}_{\mathrm{c}}$ stands near the orders 1, 2, 3 and 3: one uncertainty in the shared logit is raised to the power of the order, so the band on $\ln U$ widens with the redundancy. In $\mathbf{G}_{\mathrm{hall}}$ the weight moves from the first entry to the second as the cooling improves, and for both arrangements of order three the feed's entry is the larger: the state-space form of the statement, read in Table 13.9, that the feed is then the term to improve.

Table 13.9 evaluates equations (13.48) and (13.49) for the four arrangements under the posterior of case A: each row is an arrangement, and the columns give the plug-in value, the expectation, its ratio to the plug-in value, the 95 % band, the probability of meeting 10^{-3}, which is $U(a)$, and the expected hall unavailability with the power feed, each on the state route with the exact value in parentheses. One row of the arithmetic, the fifth unit, 3-out-of-5: with $\sigma_{\ln U} = 1.680$, $\Phi(\ln(10^{-3}/2.58 \times 10^{-4})/1.680) = \Phi(0.806) = 0.790$; and exactly, $\mathbb{E}[U] = 10\mathbb{E}[q^3] - 15\mathbb{E}[q^4] + 6\mathbb{E}[q^5]$, with $\mathbb{E}[q^5] = \mathbb{E}[q^4] \times 7/104 = 2.284 \times 10^{-7}$, gives $5.824 \times 10^{-4} - 5.089 \times 10^{-5} + 1.37 \times 10^{-6} = 5.33 \times 10^{-4}$.

Each spare unit lowers the plug-in value by a factor of twenty to fifty. The ratio of the expectation to the plug-in value is 1.14, 1.76 and about 3.7 for orders one, two and three on the state route, and exactly 0.99, 1.29 and about 2.1: for the series the exact gap is negative while the state route's is positive, because the upper tail of the normal in the logit outweighs the negative curvature. The two arrangements of order three meet the cooling target with 0.790 (a fifth unit) and 0.908 (2-out-of-4), and exactly with 0.855 and 0.956: on both routes 2-out-of-4 clears the assurance of 0.90 and the fifth unit does not, the first by 0.008 on the state route. With the feed in series, neither takes the hall below 2.9×10^{-3}, and exactly neither takes it below 2.2×10^{-3}.

13.3.6 The mixture over drawings

Engineering question. If the drawing itself is uncertain, 3-out-of-4 or 2-out-of-4, what availability is reported? Whether the hall needs three of its units or two depends on its heat load against the units' capacity, and the design file weighs the two.

Table 13.9: Alternative cooling arrangements under the posterior of case A: state route, the Laplace normal of ξ, with the exact value under Beta$(3, 97)$ in parentheses. Expectations by the sum (13.26) (exact: the moments, equation (13.29)); bands and probabilities by the delta method on $\ln U$, equations (13.33) and (13.48) (exact: the quantile map (13.34)); "Meets 10^{-3}" is $\Pr(U_{\text{system}} < 10^{-3} \mid \mathbf{y})$; the last column adds the power feed by equation (13.39). The polynomials are those of Table 13.4.

Drawing	Plug-in	$\mathbb{E}[U]$	Ratio	95 % band	Meets 10^{-3}	With feed
4-out-of-4, no spare	0.115	0.130 (0.113)	1.14 (0.99)	[0.040, 0.33] ([0.025, 0.255])	0 (0)	0.132 (0.115)
3-out-of-4, N+1	5.19×10^{-3}	9.15×10^{-3} (6.67×10^{-3})	1.76 (1.29)	$[5.8 \times 10^{-4}, 4.6 \times 10^{-2}]$ ($[2.36 \times 10^{-4}, 2.75 \times 10^{-2}]$)	0.070 (0.139)	1.17×10^{-2} (8.66×10^{-3})
3-out-of-5, fifth unit	2.58×10^{-4}	9.28×10^{-4} (5.33×10^{-4})	3.60 (2.07)	$[9.6 \times 10^{-6}, 6.9 \times 10^{-3}]$ ($[2.5 \times 10^{-6}, 3.2 \times 10^{-3}]$)	0.790 (0.855)	3.49×10^{-3} (2.53×10^{-3})
2-out-of-4, N+2	1.06×10^{-4}	3.95×10^{-4} (2.23×10^{-4})	3.74 (2.11)	$[3.8 \times 10^{-6}, 2.9 \times 10^{-3}]$ ($[9.9 \times 10^{-7}, 1.4 \times 10^{-3}]$)	0.908 (0.956)	2.96×10^{-3} (2.22×10^{-3})

General form. The expectation of a derived quantity, equation (2.10), with the state extended by a discrete entry, over which the integral is a sum:

$$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}] = \int f(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}.$$

Specialized form. One choice enters here: the weights of the design file are taken as given and independent of the checks, which say nothing about the load. This is what lets them stand unchanged after the records; evidence on the structure would replace them by posterior probabilities through the ratio of the evidences of the two drawings, which a later chapter on model selection computes.

The question is equation (2.10) with the state extended by the discrete entry d, the drawing, with weights π_d, $f(\mathbf{z}) = U_{\text{system}}$ a function of d and of $\mathbf{q}$, and $p(\mathbf{z} \mid \mathbf{y})$ the weights, which the checks do not move, times the Laplace normal of case A. By the law of total expectation over d,

$$\mathbb{E}[U_{\text{system}} \mid \mathbf{y}] = \sum_{d} \pi_d\, \mathbb{E}[U_{\text{system}} \mid d, \mathbf{y}], \tag{13.50}$$

and by the law of total variance, with $\bar{U}_d = \mathbb{E}[U_{\text{system}} \mid d, \mathbf{y}]$ and $\mathrm{sd}_d = \mathrm{sd}(U_{\text{system}} \mid d, \mathbf{y})$ the mean and standard deviation under drawing d and $\bar{U}$ the mixture mean, each $\bar{U}_d$ and sd_d from the sum (13.26) with $h = U_d(q(\xi))$ and $h = U_d(q(\xi))^2$,

$$\mathrm{Var}(U_{\text{system}} \mid \mathbf{y}) = \sum_{d} \pi_d\, \mathrm{sd}_d^2 + \sum_{d} \pi_d (\bar{U}_d - \bar{U})^2 = \pi_3 \mathrm{sd}_3^2 + \pi_2 \mathrm{sd}_2^2 + \pi_3 \pi_2 (\bar{U}_3 - \bar{U}_2)^2, \tag{13.51}$$

the second form for two drawings. The variance has a part within each drawing, from the uncertain unavailabilities, and a part between the drawings, from the uncertain structure. The between part is a product of the same form as a band. Write the drawing as an indicator vector $\mathbf{e}_d$, with mean $\boldsymbol{\pi} = (\pi_3, \pi_2)^{\mathsf{T}}$ and covariance $\mathbf{C}_\pi = \mathrm{diag}(\boldsymbol{\pi}) - \boldsymbol{\pi}\boldsymbol{\pi}^{\mathsf{T}}$, and the drawings' means as the vector $\bar{\mathbf{U}} = (\bar{U}_3, \bar{U}_2)^{\mathsf{T}}$, the derivative of the conditional mean $\bar{\mathbf{U}}^{\mathsf{T}}\mathbf{e}_d$ with respect to the

indicator:

$$\sum_d \pi_d(\bar{U}_d - \bar{U})^2 = \bar{\mathbf{U}}^{\mathsf{T}}\mathbf{C}_\pi\bar{\mathbf{U}}, \qquad \mathbf{C}_\pi = \pi_3\pi_2\begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}. \tag{13.52}$$

The null vector of $\mathbf{C}_\pi$ is $(1,1)^{\mathsf{T}}$: a shift common to both drawings adds nothing between them, and only the difference $\bar{U}_3 - \bar{U}_2$ enters. Unlike a delta-method band, the product is exact, not a linearization: the conditional mean $\bar{\mathbf{U}}^{\mathsf{T}}\mathbf{e}_d$ is linear in the indicator, so its variance over d is exactly $\bar{\mathbf{U}}^{\mathsf{T}}\mathbf{C}_\pi\bar{\mathbf{U}}$. The probability of meeting the target and the distribution function of the mixture follow by the law of total probability over d, with the normal tail (13.36) under each drawing,

$$\Pr(U_{\text{system}} < u \mid \mathbf{y}) = \sum_d \pi_d \Pr(U_{\text{system}} < u \mid d, \mathbf{y}) \approx \sum_d \pi_d \, \Phi\left(\frac{\ln u - \ln U_d(\hat{q})}{\sigma_{\ln U,d}}\right), \tag{13.53}$$

whose roots at 0.025, 0.5 and 0.975 are the mixture's band and median. The values beside them come from the exact posterior: the means and standard deviations by the moments (13.29), and the tail

$$\Pr(U_{\text{system}} < u \mid \mathbf{y}) = \sum_d \pi_d \Pr(U_{\text{system}} < u \mid d, \mathbf{y}) = \pi_3 \, I_{q_3(u)}(a,b) + \pi_2 \, I_{q_2(u)}(a,b), \tag{13.54}$$

with $q_d(u)$ the root of $U_{\text{system}}(q) = u$ under drawing d; at $u = u^*$ it is the weighted sum of the two drawings' tails (13.37). The mixture reports the uncertainty of the drawing; it does not resolve it.

Building blocks. The MAP entry $\hat{\xi} = -3.4761$ and $\sigma_\xi = 0.5862$; the columns of Table 13.6 through the sum (13.26); the spreads $\sigma_{\ln U,3} = 1.114$ and $\sigma_{\ln U,2} = 1.693$ of the matrix $\mathbf{G}_{\text{c}}$ with Φ; the weights $\pi_3 = 0.7$ and $\pi_2 = 0.3$; and, for the exact values beside them, the means and standard deviations of the two drawings from the moments (13.29), up to order eight for the variances, the roots $q_d(u)$ and the incomplete Beta function.

The mixture over the two drawings. Table 13.10 evaluates the mixture of Section 13.3.6 on the posterior of case A: each row is a drawing, and the columns give its weight, its mean and standard deviation, its weighted contribution to the mean and its probability of meeting 10^{-3}, each on the state route with the exact value in parentheses; the last rows give the mixture.

Table 13.10: The mixture over the two candidate drawings of the hall's cooling, case A, equations (13.50) and (13.51), and (13.53) on the state route with (13.54) exactly, in parentheses. The median and band of the mixture are roots of its distribution function.

Drawing	Weight	Mean $\bar{U}_d$	Std. dev. sd_d	$\pi_d\bar{U}_d$	$\Pr(U < 10^{-3})$
3-out-of-4	0.7	9.15×10^{-3}	1.24×10^{-2}	6.40×10^{-3}	0.070
		(6.67×10^{-3})	(7.57×10^{-3})	(4.67×10^{-3})	(0.139)
2-out-of-4	0.3	3.95×10^{-4}	1.17×10^{-3}	1.19×10^{-4}	0.908
		(2.23×10^{-4})	(4.41×10^{-4})	(6.7×10^{-5})	(0.956)
Mixture	1	6.52×10^{-3}	1.11×10^{-2}		0.321
		(4.74×10^{-3})	(7.00×10^{-3})		(0.384)
Mixture median 2.86×10^{-3} (2.10×10^{-3})					
Mixture 95 % band $[1.0 \times 10^{-5},\ 3.87 \times 10^{-2}]$ $([4.3 \times 10^{-6},\ 2.43 \times 10^{-2}])$					

The mixture's expectation is $0.7 \times 9.145 \times 10^{-3} + 0.3 \times 3.953 \times 10^{-4} = 6.52 \times 10^{-3}$. With $\pi_3\pi_2 = 0.21$, the matrix of equation (13.52) has the entries ± 0.21, and the variance

is 1.08×10^{-4} within the 3-out-of-4 drawing, 4.1×10^{-7} within the other, and $\bar{\mathbf{U}}^{\mathsf{T}}\mathbf{C}_{\pi}\bar{\mathbf{U}} = 0.21 \times (8.75 \times 10^{-3})^2 = 1.61 \times 10^{-5}$ between them: the open structure carries 13 % of the variance. The probability of meeting the target is the weighted sum of the two drawings' probabilities, $0.7 \times 0.070 + 0.3 \times 0.908 = 0.321$. The mixture's band, the roots of equation (13.53), is 1.0×10^{-5} to 3.87×10^{-2}, more than three and a half decades.

The values beside them come from the exact posterior. The mixture's expectation is $0.7 \times 6.673 \times 10^{-3} + 0.3 \times 2.228 \times 10^{-4} = 4.74 \times 10^{-3}$, 1.78×10^{-3} below the state route. The variance is 4.02×10^{-5} within the 3-out-of-4 drawing, 5.8×10^{-8} within the other, and $\bar{\mathbf{U}}^{\mathsf{T}}\mathbf{C}_{\pi}\bar{\mathbf{U}} = 0.21 \times (6.45 \times 10^{-3})^2 = 8.74 \times 10^{-6}$ between them: the open structure carries 18 % of the variance. The probability of meeting the target is $0.7 \times 0.139 + 0.3 \times 0.956 = 0.384$, 0.063 above the state route, and neither is near the assurance of 0.90. The mixture's band, 4.3×10^{-6} to 2.43×10^{-2}, spans nearly four decades.

13.4 The checks (step 5)

The hall has no physics row and no exact constraint, so no limit $\sigma_p \to 0$ of a model tolerance arises and no limit table is needed; the bounds $0 < q_i < 1$ are not rows either, and the logit state handles them. Each answer of step 4 is tested here against a fuller calculation: the exact Beta posteriors with their moments and quantile map, or the 4×10^5 draws with seed 13.

13.4.1 The delta routes against the exact calculation

The band of the cooling unavailability. Table 13.11 lists the band of the cooling unavailability by each route, for both evidence cases. Each row is a route; the columns give the spread it assigns, its center, the central 95 % band and the probability below the target 10^{-3}. The rows on $\ln U$ are the reported answers; the exact rows stand beside them in Sections 13.3.1 and 13.3.2.

Table 13.11: The band of the hall's cooling unavailability under 3-out-of-4 by three routes. Delta on U: normal with $\sigma_U = (\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g})^{1/2}$, equation (13.32); delta on $\ln U$: lognormal with $\sigma_{\ln U} = (\mathbf{g}_{\ln U}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_{\ln U})^{1/2}$, equation (13.33), about the plug-in value; both on the Laplace covariance $\boldsymbol{\Sigma}$ of the logits; exact: the quantile map (13.34) for case A, 4×10^5 draws (seed 13) for case B. The last column is $\Pr(U_{\text{system}} < 10^{-3} \mid \mathbf{y})$.

Route	Spread	Center	95 % band	Below 10^{-3}
Case A, one shared entry $q \sim \text{Beta}(3, 97)$				
Delta on U	$\sigma_U = 5.78 \times 10^{-3}$	5.19×10^{-3}	$[-6.1 \times 10^{-3}, 1.65 \times 10^{-2}]$	0.234
Delta on $\ln U$, reported	$\sigma_{\ln U} = 1.11$	5.19×10^{-3}	$[5.8 \times 10^{-4}, 4.6 \times 10^{-2}]$	0.070
Exact	$\text{sd}(U \mid \mathbf{y}) = 7.57 \times 10^{-3}$	median 4.19×10^{-3}, mean 6.67×10^{-3}	$[2.36 \times 10^{-4}, 2.75 \times 10^{-2}]$	0.139
Case B, four independent entries, Table 13.3				
Delta on U	$\sigma_U = 5.88 \times 10^{-3}$	8.03×10^{-3}	$[-3.5 \times 10^{-3}, 1.95 \times 10^{-2}]$	0.116
Delta on $\ln U$, reported	$\sigma_{\ln U} = 0.73$	8.03×10^{-3}	$[1.9 \times 10^{-3}, 3.4 \times 10^{-2}]$	0.002
Exact (draws)	$\text{sd}(U \mid \mathbf{y}) = 6.05 \times 10^{-3}$	median 6.47×10^{-3}, mean 8.03×10^{-3}	$[1.15 \times 10^{-3}, 2.38 \times 10^{-2}]$	0.018

With the gradients (13.46) and the covariances of the fit, the delta method on U, equation (13.32), gives $\sigma_U = 9.857 \times 10^{-3} \times \sqrt{0.34364} = 5.78 \times 10^{-3}$ for case A, and 5.88×10^{-3} for

case B. In case A σ_U is larger than the plug-in value 5.19×10^{-3} itself. A normal with that mean and spread puts $\Phi(-5.19/5.78) = 18\,\%$ of its probability on negative unavailabilities, and its central band is $[-6.1 \times 10^{-3},\ 1.65 \times 10^{-2}]$.

The delta method on U fails as Section 13.3.1 said it would. The delta method on $\ln U$, the reported route, has the right shape, a factor of 79 wide against the exact 117, and sits a factor of 1.7 to 2.5 too high at the two ends in case A, and halves the probability below the target, 0.070 against 0.139; in case B its ends are 1.4 to 1.65 times high and its probability below the target an eighth of the draws', 0.002 against 0.018. Three sources of the case-A error can be told apart. The linearization of $\ln U$ in ξ costs little: mapping the exact quantiles of ξ through the straight line of slope 1.90 instead of through U itself puts the two ends only 8 % and 5 % high. The rest is the normal on ξ. By a standard result, the logit of a Beta(a, b) variable, the difference of the logarithms of two independent gamma variables, has mean $\psi(a) - \psi(b)$ and variance $\psi'(a) + \psi'(b)$, with ψ and ψ' the digamma and trigamma functions of Chapter 8. For Beta$(3, 97)$ the mean is -3.647 against the Laplace center $\hat{\xi} = -3.476$, and the standard deviation 0.637 against $\sqrt{\Sigma_{11}} = 0.586$. The offset of 0.17 in ξ is multiplied by the entry 1.90 of $\mathbf{g}_{\ln U,\mathrm{A}}$, a factor of 1.4 at both ends. A normal on ξ with the exact mean and standard deviation still gives $[3.3 \times 10^{-4},\ 3.7 \times 10^{-2}]$, because the logit posterior is not symmetric: with three findings it is skewed to the left, with skewness -0.60, so its lower tail is longer and its upper tail shorter than any normal's. The tail of case A, computed with the threshold mapped exactly to the logit, $\Phi((\ln(0.01302/0.98698) + 3.4761)/0.5862)$, is 0.073 against the reported 0.070: there too the linearization costs little and the normal the rest. In case B a fuller state route is available, 4×10^5 draws of the normal $\mathcal{N}(\hat{\mathbf{z}}_\mathrm{B}, \boldsymbol{\Sigma}_\mathrm{B})$ pushed through the enumeration, seed 13, and it gives the median 1.04×10^{-2}, the band $[2.3 \times 10^{-3},\ 4.4 \times 10^{-2}]$ and the probability 0.001 below the target. Those draws have the mean 1.367×10^{-2}, within about one Monte Carlo standard error of the enumerated 1.365×10^{-2}, as multilinearity under independent entries requires.

Common mistake

Placing the band of a small unavailability by the delta method on the unavailability itself. For the hall it gives a standard deviation larger than the value and a band whose lower end is -6.1×10^{-3}. A small unavailability is uncertain by a factor, and its band belongs on the logarithm, with the exact distribution, here the quantile map of one shared entry, beside it.

The Jensen gap to leading order. At $q = 0.03$ every off-diagonal entry of $\mathbf{H}_q$ in equation (13.31) is $(1 - q)(1 - 3q) = 0.8827$,

$$\mathbf{H}_q = \begin{bmatrix} 0 & 0.8827 & 0.8827 & 0.8827 \\ 0.8827 & 0 & 0.8827 & 0.8827 \\ 0.8827 & 0.8827 & 0 & 0.8827 \\ 0.8827 & 0.8827 & 0.8827 & 0 \end{bmatrix}, \qquad \mathbf{1}^\mathsf{T}\mathbf{H}_q\mathbf{1} = 12 \times 0.8827 = 10.592,$$

and with $\mathrm{Var}(q) = 2.881 \times 10^{-4}$ (the Laplace $D_{11}\Sigma_{\mathrm{A},11}D_{11} = 0.0291^2 \times 0.34364 = 2.910 \times 10^{-4}$) the leading term is 1.526×10^{-3}, 2.7 % above the exact gap 1.487×10^{-3} of Section 13.3.1. The reported state-route gap, 3.96×10^{-3}, is 2.7 times the exact one. The leading order uses only the variance of q; the sum over the normal in the logit also moves the mean of q from 0.0300 to 0.0350 and raises every higher moment.

The importances. The reported bands of Section 13.3.4, [0.047, 0.232], [0.060, 0.250] and [0.034, 0.208], stand against the draws' [0.036, 0.199], [0.049, 0.219] and [0.024, 0.173]. The delta bands are slightly narrower and sit higher, by a factor of 1.1 to 1.4 at the ends, for the reason found for the band of the service: the Laplace normal centers each logit at the log-odds of the posterior mean, above the mean of the logit.

The budget of the hall. The feed's budget term of Section 13.3.4 matches its exact variance in Table 13.8; the cooling's falls short of the exact 5.74×10^{-5} for the reason of the band, so the delta method gives the feed 6 % of the variance and the exact moments 3 %. Both say that the band of the hall is the band of its cooling.

The arrangements on one covariance. The spreads 0.54, 1.11, 1.68 and 1.69 of Section 13.3.5 set the bands of Table 13.9, whose width in decades they give exactly, and the state route's tails stand below the exact ones at every arrangement but the series, where both are zero.

13.4.2 The enumeration, the moments and the draws

The fit against the Beta posteriors. The fifth Newton iterate of Table 13.2 equals the logits $\ln(a_i/b_i)$ of the exact posteriors to 10^{-12}, and the diagonal of $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ equals $a_i b_i/(a_i + b_i)$ to the printed digits; its off-diagonal entries are exactly zero at every iterate.

The enumeration and the gradient. The sixteen rows of Table 13.5 sum to one, and the column totals of Figure 13.2 are the binomial terms of equation (13.19): $0.8853 + 0.1095 = 0.99481 = 1 - U_{\text{system}}(0.03)$ from equation (13.20). Central differences of the enumeration with a step of 10^{-6} reproduce the four importances of equation (13.45) to five decimals, and for equal q the sum of the four is $U'_{\text{system}}(q)$.

Multilinearity and the moments. Draws of four independent Beta$(3, 97)$ unavailabilities give a mean cooling unavailability of 5.188×10^{-3}, and draws of case B 8.024×10^{-3}, against the plug-in values 5.186 and 8.026×10^{-3} that equation (13.27) declares exact. Draws of the shared entry give a mean of 6.650×10^{-3} against the exact 6.673×10^{-3} of equation (13.30), a difference of twice the Monte Carlo standard error, $7.57 \times 10^{-3}/\sqrt{4 \times 10^5} = 1.2 \times 10^{-5}$.

The quantile map and the series. The draws of the shared entry give the band $[2.34 \times 10^{-4},\ 2.74 \times 10^{-2}]$ and $\Pr(U < 10^{-3}) = 0.141$, against the exact $[2.36 \times 10^{-4},\ 2.75 \times 10^{-2}]$ and 0.139, and their histogram in Figure 13.4 lies on the exact density. Draws of the feed and of the shared entry give the band $[1.28 \times 10^{-3},\ 2.96 \times 10^{-2}]$ of the series, the band of the integral (13.42).

The mixture. Draws of the drawing and of q give $\Pr(U < 10^{-3}) = 0.386$, the median 2.09×10^{-3}, the band $[4.2 \times 10^{-6},\ 2.42 \times 10^{-2}]$ and the mean 4.72×10^{-3}, against the 0.384, 2.10×10^{-3} and $[4.3 \times 10^{-6},\ 2.43 \times 10^{-2}]$ of the distribution function (13.54) and the mean 4.74×10^{-3} of the moments. Each difference is within about two Monte Carlo standard errors: 1.1×10^{-5} on the mean, 8×10^{-4} on the probability and 1.3 % on the lower end of the band, where the mixture's density is thinnest.

13.4.3 Sensitivity to the prior

Case A takes the fleet record alone, under the vague limit of the Beta prior, and that limit is a choice like any other. Table 13.12 repeats the case-A answers under three further priors. Each row is a prior: the vague limit of the chapter; the Jeffreys prior Beta$(\frac{1}{2}, \frac{1}{2})$, a standard reference prior; the uniform prior Beta$(1, 1)$; and the unit prior Beta$(1, 19)$ of case B, used once for the fleet, the pooled row of Table 13.7. The columns give the posterior, the mean of q, the expected cooling unavailability under 3-out-of-4 and its ratio to the plug-in value, and the probability of meeting 10^{-3} under 3-out-of-4 and under the mixture of the two drawings, each on the state route with the exact value in parentheses. One row of the arithmetic, the uniform

prior: Beta$(1+3,\ 1+97)$ = Beta$(4, 98)$, mean $4/102 = 0.0392$, so that $\hat{\xi} = \ln(4/98) = -3.1987$ and $\sigma_\xi^2 = 102/392 = 0.26020$, and the tail (13.36) is 0.011; exactly, $I_{0.01302}(4, 98) = 0.044$.

Table 13.12: Sensitivity of the case-A answers to the prior on the fleet unavailability; record 3 of 100 checks. State route: $\mathbb{E}[U]$ by the sum (13.26), probabilities by equations (13.36) and (13.53); exact, in parentheses: $\mathbb{E}[U]$ by the moments (13.29), probabilities by the quantile map (13.37) and the mixture (13.54), with the weights 0.7 and 0.3.

Prior	Posterior	Mean q	$\mathbb{E}[U]$, 3-out-of-4	Ratio	$\Pr(U < 10^{-3})$, 3-out-of-4	$\Pr(U < 10^{-3})$, mixture
vague limit (case A)	Beta$(3, 97)$	0.0300	9.15×10^{-3} (6.67×10^{-3})	1.76 (1.29)	0.070 (0.139)	0.321 (0.384)
Jeffreys, Beta$(\frac{1}{2}, \frac{1}{2})$	Beta$(3.5, 97.5)$	0.0347	1.11×10^{-2} (8.54×10^{-3})	1.62 (1.24)	0.030 (0.080)	0.284 (0.335)
uniform, Beta$(1, 1)$	Beta$(4, 98)$	0.0392	1.32×10^{-2} (1.06×10^{-2})	1.51 (1.21)	0.011 (0.044)	0.260 (0.299)
unit prior, Beta$(1, 19)$	Beta$(4, 116)$	0.0333	9.76×10^{-3} (7.73×10^{-3})	1.53 (1.21)	0.027 (0.071)	0.291 (0.335)

Each of the three proper priors is a row of pseudo-checks, equation (13.4), with findings at a rate above the fleet's 3 %: half a finding in one pseudo-check for the Jeffreys prior, one in two for the uniform, one in twenty for the unit prior. Each therefore raises the mean of q and lowers the probability of the target. The vague limit is the most favorable of the four to the target. The verdict does not move on either route: under the three-unit drawing the target holds with probability at most 0.070 (0.139 exactly), and under the mixture with probability 0.26 to 0.32 (0.30 to 0.38), never near the assurance of 0.90. The ratio to the plug-in value stays between 1.51 and 1.76 on the state route and between 1.21 and 1.29 exactly, so the plug-in value understates the expectation under every prior. What the prior does move is the size of the shortfall: the probability under 3-out-of-4 varies by a factor of six across the four rows on the state route and of three exactly.

13.4.4 The rows against each other and the pooling

The residuals of case B. Each entry of case B carries two rows, its prior and its record, so the rows are redundant and their residuals at the MAP state test each record against its prior. The ten residuals of equation (13.6) give $\|\mathbf{r}(\hat{\mathbf{z}}_\mathrm{B})\|^2 = 3.87$ on $10 - 5 = 5$ degrees of freedom, a χ^2 tail probability of 0.57; the largest pair, 1.28, is that of units 2 and 4, no finding in 25 checks against a prior of one in twenty. The exact check agrees: under each prior the Beta–binomial probability of a record at least as far out as the one held, equation (5.49), is 0.32 or more for every unit and 0.36 for the feed, so no record conflicts with its prior.

The homogeneity of the pooled records. Case A pools the four unit records under one unavailability, and the records can be asked whether they are consistent with it. At the pooled estimate 0.03 each unit expects 0.75 findings in 25 checks, and the Pearson statistic

$$\sum_{i=1}^{4} \frac{(k_i - 0.75)^2}{25 \times 0.03 \times 0.97} = \frac{2.75}{0.7275} = 3.78$$

on $4 - 1 = 3$ degrees of freedom has a χ^2 tail probability of 0.29. With three findings in all the χ^2 reference is rough. Conditional on the total, the three findings fall on the four units as a draw of three from the four sets of 25 checks, and the probability of a split at least as uneven as the observed 2, 1, 0, 0 is 0.61. Neither test rejects the shared unavailability, and neither

could reject much: under a common unavailability a split this uneven arises more often than not. The pooling is a modeling choice that the records neither support nor contradict, which is why both cases are carried and the difference between them is reported.

13.5 The report (step 6)

The plug-in reading of the log, a cooling unavailability of 5.2×10^{-3} and a clear miss that calls for a fifth unit, is sharply qualified by the analysis. If, as the operator treats them, the four units share one unavailability, the expected cooling unavailability is 9.15×10^{-3}, 80 h a year, 76 % above the plug-in value (6.67×10^{-3}, 58 h and 29 % exactly), because that one unavailability is uncertain; kept per unit, the same checks give 1.36×10^{-2} (8.03×10^{-3}), higher for another reason, the pull of the unit priors, and on the state route for few findings on two units. The band runs over two decades, from 5.8×10^{-4} to 4.6×10^{-2} (2.36×10^{-4} to 2.75×10^{-2}); and with the drawing open, the probability that the cooling meets 10^{-3} is 0.321 (0.384), neither the plug-in's clear miss nor a pass. Under the three-unit drawing the target is met with only 0.070 (0.139): N+1 with units of this unavailability is a design for a few times 10^{-3}, not for 10^{-3}. Under the two-unit drawing it is met with 0.908 (0.956). A fifth unit would meet it with 0.790 (0.855), not with certainty, and the first question for the plant is how many units its load needs. The answers are computed from the posterior of the state, each with the exact value beside it; the gaps are large here, a factor of 1.37 on the mean of case A and 1.70 on case B, because few findings make the normal in the logit a poor copy of a Beta density, and no verdict of the report changes between the two routes.

13.5.1 The answers and their bands

Table 13.13 gives the answers. Each row is a question of Table 13.1, in its order; the second column gives the answer from the posterior of the state, the third the exact value, and the last the band or comparison, on the state route with the exact value in parentheses.

13.5.2 The target decision

The report line. With the fleet's hundred checks pooled into one unavailability, the hall's N+1 cooling is expected to be down 9.1×10^{-3} of the time, 80 h a year (6.7×10^{-3}, 58 h from the exact posterior), with a median of 45 h (37 h) and a 95 % band from 5 to 403 h (2 to 241 h); with its power feed, 102 h a year (76 h), of which the feed is a quarter. The probability that the cooling meets 10^{-3} is 0.07 (0.14) if three units are needed, 0.32 (0.38) with the design file's weights on the two drawings. A fifth unit, or a load that two units can carry, would leave the power feed as the larger term of the hall. Kept per unit, the records give 1.4×10^{-2} (8.0×10^{-3}, exactly the value at the per-unit means); which treatment of the fleet is right is a question for the evidence, not for the arithmetic.

Common mistake

Reading N+1 redundancy as a guarantee of a low unavailability. The redundancy turns single-unit outages into delivered service, but the unavailability that remains is of order $6q^2$, and with units down 3 % of the time it is about 5×10^{-3}, forty-five hours a year at the plug-in value and eighty in expectation (fifty-eight from the exact posterior); the posterior probability that it is below 10^{-3} is 0.07 (0.14). Redundancy raises the order of the structure; the unit unavailability still sets its scale.

Table 13.13: The answers to the questions asked of the hall. Answers from the posterior of the state, with the exact value from the moments and the quantile map, or from draws where the entries are independent, beside each. Bands are central 95 % posterior intervals of the unavailability; prior ranges from Table 13.12.

Question	Answer	Exact	Band and comparison
Cooling availability under 3-out-of-4	case A: $\mathbb{E}[U] = 9.15 \times 10^{-3}$, median 5.19×10^{-3}; case B: $\mathbb{E}[U] = 1.36 \times 10^{-2}$, median 8.03×10^{-3}	A: 6.67×10^{-3}, 4.19×10^{-3}; B: 8.03×10^{-3}, 6.47×10^{-3}	A: $[5.8 \times 10^{-4},\ 4.6 \times 10^{-2}]$ ($[2.36 \times 10^{-4},\ 2.75 \times 10^{-2}]$), 5 to 403 h per year (2 to 241 h); B: $[1.9 \times 10^{-3},\ 3.4 \times 10^{-2}]$ ($[1.15 \times 10^{-3},\ 2.38 \times 10^{-2}]$); plug-in 5.19×10^{-3} (A)
Unavailability below 10^{-3}	case A: 0.070; case B: 0.002; not demonstrated at 0.90	A: 0.139; B: 0.018; not demonstrated	0.011 to 0.070 over the four priors (0.044 to 0.139) (A); below 10^{-4}: 0.0002 (0.008) (A)
Pooled against apart	same marginals: 9.15×10^{-3} shared, 6.99×10^{-3} independent	6.67 shared, 5.19×10^{-3} independent	case B records: 9.76×10^{-3} pooled with one copy of the unit prior, 1.36×10^{-2} apart (7.73×10^{-3}, 8.03×10^{-3}); plug-in exact only when apart, on the exact route
Power feed in series	$\mathbb{E}[U_{\text{series}}] = 1.17 \times 10^{-2}$, 102 h per year	8.66×10^{-3}, 76 h per year	$[1.42 \times 10^{-3},\ 3.63 \times 10^{-2}]$ ($[1.28 \times 10^{-3},\ 2.96 \times 10^{-2}]$); feed 22 % (23 %) of the mean, 6 % (3 %) of the variance; hall below 10^{-3}: 0.009 (0.013)
Most important unit	$\mathbb{E}[I_{\text{B}}] = (0.133, 0.149, 0.115, 0.149)$: units 2 and 4, with probability 0.86	(0.104, 0.123, 0.084, 0.123), also I_{B} at the MAP state; 0.87	bands [0.047, 0.232], [0.060, 0.250], [0.034, 0.208], [0.060, 0.250] ([0.036, 0.199], [0.049, 0.219], [0.024, 0.173], [0.049, 0.219]); unit 3 least important with probability 0.53 (0.60); budget shares (30, 21, 28, 21) %: the next checks on unit 1
Arrangement meeting the target	2-out-of-4: 0.908; 3-out-of-5, a fifth unit: 0.790; 3-out-of-4: 0.070; 4-out-of-4: 0	0.956; 0.855; 0.139; 0	cooling $\mathbb{E}[U] = 3.95 \times 10^{-4}$ and 9.28×10^{-4} for the two of order three (2.23 and 5.33×10^{-4}); with the feed 2.96 and 3.49×10^{-3} (2.22 and 2.53×10^{-3})
Open drawing	$\mathbb{E}[U] = 6.52 \times 10^{-3}$, std. dev. 1.11×10^{-2}	4.74×10^{-3}, 7.00×10^{-3}	$\Pr(U < 10^{-3}) = 0.321$ (0.384), 0.260 to 0.321 over the priors (0.299 to 0.384); 13 % (18 %) of the variance from the drawing; band $[1.0 \times 10^{-5},\ 3.87 \times 10^{-2}]$ ($[4.3 \times 10^{-6},\ 2.43 \times 10^{-2}]$)

13.6 The next measurement (step 7)

The report leaves the target undecided where it can still be decided: under the open drawing the cooling meets 10^{-3} with probability 0.321 (0.384 exactly). Records could narrow the band of the unavailability or move the weights of the drawings, and the preposterior values each before its outcome is known.

13.6.1 The candidate records

Three records are open to the operator:

- *A hundred more checks of the fleet.* The record of case A doubled, one more count row on the shared logit ξ.
- *Twenty-five more checks of one unit.* A record of case B, one more count row on the logit of that unit.
- *A reading of the heat load.* The hall's heat load read against the units' capacity, which bears on the drawing rather than on the units: an observation of the discrete entry d.

Each is judged by what it removes from the answers of Table 13.13: the width of the band and the assurance of meeting 10^{-3}.

13.6.2 The preposterior of the candidate records

Each candidate adds rows to the graph, and its preposterior, the covariance $\mathbf{\Sigma}_a = (\mathbf{\Sigma}^{-1}+\mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a)^{-1}$ of equation (2.17) or the posterior of the state after each outcome the record could return, averaged over the Laplace predictive of those outcomes, is computed before the record is taken. The exact distribution of the answer, averaged over the exact predictive, gives the value beside it.

A hundred more checks of the fleet. The band of case A spans two decades, and the operator asks whether a hundred more checks of the fleet would narrow it enough to settle the target. The expected number of findings in a hundred more checks is $100 \times 0.03 = 3$, and at that count the posterior would be $\mathrm{Beta}(6, 194)$, whose Laplace normal has $\hat{\xi} = \ln(6/194)$ and $\sigma_\xi^2 = 200/(6 \times 194)$. The count is uncertain. Under the Laplace normal of the state its predictive probability, the sum (13.26) with h the binomial probability of the count, is 0.088 for no finding and 0.161 for exactly three; under the exact posterior it is Beta–binomial, as equation (5.49) gives, with probability 0.121 of no finding and 0.159 of exactly three. The preposterior averages over it. Table 13.14 lists the outcomes: each row is a record, present or candidate, and the columns give the posterior, the 95 % band on the cooling unavailability, its width in decades, the probability of meeting 10^{-3}, and the predictive probability of that record, each on the state route with the exact value in parentheses.

Narrowing of the band. At the expected count the band narrows from 1.90 to 1.34 decades, by a factor of 1.9 at each end, and averaged over the possible counts the width is 1.37 decades. The value beside it comes from the exact posterior: from 2.07 to 1.40 decades, from both ends, the lower end rising by a factor of 3.1 and the upper falling by a factor of 1.5, and 1.50 decades on average. The record is worth taking if the question is the size of the cooling's unavailability.

The same narrowing from the covariance. The Laplace preposterior gives the width before any count is considered. The hundred checks add one row to the first column of $\mathbf{J}$, with

Table 13.14: The preposterior of further fleet checks, case A, cooling under 3-out-of-4. State route: bands and probabilities by equations (13.33) and (13.36) at the Laplace normal after the record, predictive probabilities by the sum (13.26). Exact, in parentheses: bands by the quantile map (13.34), predictive probabilities from the Beta–binomial distribution under Beta(3, 97).

Record	Posterior	95 % band	Decades	$\Pr(U < 10^{-3})$	Predictive probability
Present, 3 of 100	Beta(3, 97)	$[5.8 \times 10^{-4},\ 4.6 \times 10^{-2}]$ ($[2.36 \times 10^{-4},\ 2.75 \times 10^{-2}]$)	1.90 (2.07)	0.070 (0.139)	–
100 more, 3 found (expected)	Beta(6, 194)	$[1.11 \times 10^{-3},\ 2.43 \times 10^{-2}]$ ($[7.3 \times 10^{-4},\ 1.84 \times 10^{-2}]$)	1.34 (1.40)	0.018 (0.047)	0.161 (0.159)
100 more, none found	Beta(3, 197)	$[1.4 \times 10^{-4},\ 1.22 \times 10^{-2}]$ ($[5.8 \times 10^{-5},\ 7.3 \times 10^{-3}]$)	1.93 (2.10)	0.403 (0.480)	0.088 (0.121)
100 more, averaged over the count	–	–	1.37 (1.50)	0.082 (0.139)	1
383 more, none found (exact: 309)	Beta(3, 480) (Beta(3, 406))	$[2.4 \times 10^{-5},\ 2.2 \times 10^{-3}]$ ($[1.4 \times 10^{-5},\ 1.8 \times 10^{-3}]$)	1.95 (2.12)	0.900 (0.901)	0.002 (0.014)

entry $\sqrt{100\,\hat{q}(1-\hat{q})} = 1.7059$ at $\hat{\mathbf{z}}_{\mathrm{A}}$, equal to the present row's, so that

$$\boldsymbol{\Lambda}_a = \begin{bmatrix} 2.9100 & 0 \\ 0 & 0 \end{bmatrix}, \qquad \boldsymbol{\Sigma}_a = \left(\boldsymbol{\Sigma}_{\mathrm{A}}^{-1} + \boldsymbol{\Lambda}_a\right)^{-1} = \begin{bmatrix} 0.17182 & 0 \\ 0 & 0.50100 \end{bmatrix},$$

$$\left(\mathbf{g}_{\ln U,\mathrm{A}}^{\mathsf{T}} \boldsymbol{\Sigma}_a \mathbf{g}_{\ln U,\mathrm{A}}\right)^{1/2} = 0.788,$$

half the first entry of $\boldsymbol{\Sigma}_{\mathrm{A}}$ and $1/\sqrt{2}$ of the present $\sigma_{\ln U} = 1.114$. The log-scale band narrows from 1.90 to 1.34 decades, by the factor 1.41, the state route's width at the expected count, against the exact 2.07 to 1.50, a factor of 1.38: the covariance predicts the narrowing well although it places the band itself less well.

The assurance on average. The probability of meeting the target behaves differently. Over the exact predictive count its expectation is exactly its present value, 0.139. This is the law of total expectation: the present posterior probability of an event is the average, over the records that could be observed, of the posterior probability after observing them, so no planned record can raise an assurance on average; it can only spread it. The state route is an approximation of those posteriors and keeps the law only roughly: its average over the Laplace predictive is 0.082 against its present 0.070. At the expected count the assurance falls to 0.018 (0.047), and with no finding at all it rises only to 0.403 (0.480). On neither route does any count of a hundred more checks bring it to 0.9.

Checks needed to demonstrate the target. By equation (13.37) the target needs $q < 0.0130$, and the fleet's posterior mean is 0.030. To demonstrate it at an assurance of 0.9 the fleet would need 383 further checks without a single finding on the state route, an outcome of predictive probability 0.002, and of probability 8.6×10^{-6} if the unavailability is in fact 0.03; exactly, 309 further checks, predictive probability 0.014, and of probability 8×10^{-5} if the unavailability is in fact 0.03. Demonstrating 10^{-3} by checks of the hall itself, without the enumeration, would take 2301 consecutive checks with the cooling never lost, by the zero-failure sample size of Chapter 5 under a uniform prior. The enumeration turns a system target of 10^{-3} into a unit target of 1.3 %, which a few hundred checks can test, and the test it would run here would most likely confirm that the target is not met.

Twenty-five more checks of one unit. In case B a further 25 checks on unit i, at its present mean, add one row to $\mathbf{J}$, $\mathbf{J}_a = \sqrt{25\,\hat{q}_i(1-\hat{q}_i)}\;\mathbf{e}_i^{\mathsf{T}}$ in column i, evaluated at $\hat{\mathbf{z}}_{\mathrm{B}}$ whatever the outcome, since in the logit a binomial row's information depends on its count only through the state at which it is evaluated. Their information and the preposterior covariance are

$$\mathbf{\Lambda}_a = \mathrm{diag}\big(0,\dots,25\,\hat{q}_i(1-\hat{q}_i),\dots,0\big), \qquad \mathbf{\Sigma}_a = \big(\mathbf{\Sigma}_{\mathrm{B}}^{-1} + \mathbf{\Lambda}_a\big)^{-1}, \qquad \sigma_{U,a}^2 = \mathbf{g}_{\mathrm{B}}^{\mathsf{T}}\mathbf{\Sigma}_a\mathbf{g}_{\mathrm{B}}, \quad (13.55)$$

and only the i-th diagonal entry of $\mathbf{\Sigma}$ changes, by the factor 45/70. For unit 1, $\mathbf{\Lambda}_a$ has the one entry 1.0617 and $\Sigma_{a,11} = 1/(1.9111 + 1.0617) = 0.3364$ against 0.5233; for unit 3 the entry is 1.5556, for units 2 and 4 it is 0.5432. The delta-method σ_U falls by 5.5 % for unit 1, 5.2 % for unit 3 and 3.8 % for unit 2 or 4, and the log-scale $\sigma_{\ln U}$ by the same fractions, since $\mathbf{g}_{\ln U} = \mathbf{g}/U_{\mathrm{system}}$: among the units the next checks belong to unit 1, whose term in the budget is the largest. The assurance of case B, 0.002 on the state route, averages 0.003 after the checks of unit 1 or unit 3 and stays at 0.002 after those of unit 2 or 4, the drift of successive Laplace fits and not new information; exactly, 0.018 is unchanged on average, by the law of total expectation as for the fleet.

A reading of the heat load. The mixture of Section 13.3.6 carries 13 % of the variance of the cooling unavailability in the drawing (18 % exactly), and no check of the units moves it. A reading of the hall's heat load against the units' capacity that decides between 3-out-of-4 and 2-out-of-4 is an observation of the discrete entry d. Taken as decisive, it returns 3-out-of-4 with predictive probability $\pi_3 = 0.7$ and 2-out-of-4 with $\pi_2 = 0.3$, on either route, since the weights are not moved by the state. The expected variance after the reading is the within part of equation (13.51), 1.08×10^{-4}: the reading removes the between part, $\bar{\mathbf{U}}^{\mathsf{T}}\mathbf{C}_\pi\bar{\mathbf{U}} = 1.61 \times 10^{-5}$, and the root-mean-square standard deviation falls from 1.11×10^{-2} to 1.04×10^{-2}. The band falls from the mixture's 3.58 decades to 1.90 decades if the reading shows 3-out-of-4 and 2.88 if it shows 2-out-of-4, 2.19 on average. The assurance becomes 0.070 or 0.908, and its average over the two outcomes is the present 0.321, the weighted sum of the two. The value beside them comes from the exact posterior: within part 4.02×10^{-5}, between part 8.74×10^{-6}, the root-mean-square standard deviation from 7.00×10^{-3} to 6.34×10^{-3}; the band from 3.75 decades to 2.07 or 3.14, 2.39 on average; the assurance 0.139 or 0.956, averaging 0.384. On both routes, with probability 0.3 the reading lifts the assurance above 0.9, which no count of a hundred fleet checks does.

Table 13.15 collects the candidates: each row is a record, and the columns give the rows it adds, its effect on the band of the cooling unavailability and the assurance of meeting 10^{-3} after it, on the state route with the exact value in parentheses.

13.6.3 The recommendation

The heat-load reading is the record to take first. It is the only candidate that can change the verdict on the target: with probability 0.3 it shows that two units carry the load, and the cooling then meets 10^{-3} with assurance 0.908 (0.956 exactly), while a planned count of checks raises the assurance on average only by the drift of the Laplace fits, and exactly not at all, and a hundred fleet checks bring it to 0.9 at no count, on either route. If the reading shows 3-out-of-4, further checks would narrow the band but most likely confirm that the target is not met, and the decision is the arrangement: a fifth unit meets the cooling target with probability 0.790 (0.855) (Table 13.9), after which the power feed is the larger term of the hall. Checks of the units are worth taking for the size of the unavailability rather than for the verdict, and then on the fleet in case A or on unit 1 in case B. Chapter 20, in Volume II, ranks candidate rows by the same preposterior on graphs whose rows share entries.

Table 13.15: The candidate records and what each removes from the reported answers. Bands are central 95 % intervals of the cooling unavailability, under 3-out-of-4 for the checks and under the mixture for the load reading; a width in decades is $\log_{10}$ of the ratio of its ends. Assurance is $\Pr(U_{\text{system}} < 10^{-3} \mid \mathbf{y})$ after the record, on average over its outcomes and at its extremes.

Candidate record	Rows added	Effect on the band	Assurance after
100 more fleet checks (case A)	one row on ξ, $\mathbf{J}_a = 1.7059$, information 2.9100	1.90 to 1.37 decades on average, 1.34 by $\mathbf{\Sigma}_a$ (2.07 to 1.50)	0.082 on average (0.139); 0.018 at the expected count, 0.403 with no finding (0.047, 0.480)
25 more checks of unit 1 (case B)	one row on ξ_1, information 1.0617	σ_U and $\sigma_{\ln U}$ down 5.5 %	0.003 on average, from 0.002 (0.018 unchanged)
25 more checks of unit 3, or of unit 2 or 4 (case B)	one row on ξ_3, information 1.5556; or on ξ_2 or ξ_4, 0.5432	down 5.2 %; down 3.8 %	0.003; 0.002 on average, from 0.002 (0.018 unchanged)
Heat-load reading	an observation of the drawing d	3.58 to 2.19 decades on average (3.75 to 2.39); 13 % (18 %) of the variance removed	0.321 on average (0.384); 0.070 (probability 0.7) or 0.908 (probability 0.3) (0.139 or 0.956)

13.6.4 What the analysis showed

The first reading of the log, a unit down 3 % of the time and N+1 lost only to coincidences, gave 5.2×10^{-3} and a verdict: the target is missed, buy a fifth unit. The analysis keeps the miss under the design drawing and takes away the verdict. If, as the operator treats them, the four units share one unavailability, the plug-in value is low: a redundant structure is lost only through coincidences, and a shared uncertainty makes coincidences more likely than the product of the means says, by more the more redundant the structure. Kept per unit, the records give a higher value for another reason, the pull of the unit priors. The band of the cooling is wide, because three findings fix the unit unavailability only to a factor of about eleven and the order-two structure squares that factor. And the quantity that decides the target is not about the units at all: whether three units or two carry the load moves the probability of meeting 10^{-3} from about one in fourteen on the state route (one in seven exactly) to 0.908 (0.956), which no count of further checks can do. On site, the heat load should be read against the units' capacity before a fifth unit is bought; if three units are needed, the fifth unit is the arrangement that comes near the target, and the power feed becomes the term to improve after it. Among the four units, the one whose improvement would help the service most, unit 2 or 4, is neither the one with the worst record, unit 3, nor the one whose next checks would narrow the band most, unit 1.

13.6.5 Limits of the model

The enumeration is exact under the assumptions of Section 13.3.1. That does not make the assumptions appropriate, and four of them limit the answers.

Independence given the unavailabilities. The enumeration multiplies the statuses of the units, equation (13.14). A common cause, a power disturbance that trips several units or a maintenance error repeated across them, makes coincident outages more likely than the product. The shared posterior of case A creates a correlation of 0.01 between two units' down-states through epistemic uncertainty alone; a common cause adds an aleatory correlation on top, and a later chapter of this part frees a shared latent term for it. A single repair crew serving several units makes the second outage wait for the first repair, and the steady-state probabilities of

configurations are then those of a Markov model of the repairable system, the subject of a later chapter.

Checks and steady state. A check samples the status of a unit at one instant, and a record of checks estimates the long-run unavailability only if the checks are made at times unrelated to the failures and far enough apart that consecutive checks do not see the same outage. Checks scheduled after a repair, or repeated during one long outage, bias the estimate. The service unavailability of this chapter is a long-run fraction; it does not answer how long a single loss of cooling lasts or how often it occurs, which need the rates of Chapter 8.

Binary units and a fixed load. Each unit is up or down, and the drawing says how many suffice. A unit running at reduced capacity, or a load that varies through the day, makes the number needed a function of the load, and the structure becomes a capacity margin rather than a count; the mixture of Section 13.3.6 is the simplest form of that dependence.

The pooling choice. Case A and case B are the two extremes of a fleet: one unavailability for all, or unrelated ones. Real fleets lie between, with unavailabilities that differ around a fleet value, and a model with a fleet entry and unit deviations from it interpolates between the two rows of Table 13.7. The difference between the treatments, 9.76×10^{-3} against 1.36×10^{-2} on the state route and 7.73 against 8.03×10^{-3} exactly for the same records, is a statement about modeling that the report should carry, and the records cannot settle it (Section 13.4.4).

Common mistake

Computing the service availability at the posterior means of the units and reporting it as the expected availability when the units share one uncertain unavailability. For the hall the plug-in value is 5.19×10^{-3} and the expectation 9.15×10^{-3} on the state route, 6.67×10^{-3} exactly; for a fifth unit the plug-in value is less than half the expectation on either route. On the exact route the plug-in value is the expectation only when the posteriors of the units are independent.

Common mistake

Reading the Birnbaum importance at the MAP state as a ranking of the units' records, or as a certain ranking at all. Unit 3 has the worst record and, at the MAP state, the smallest importance; the importance says how much the service gains per unit of improvement, while the variance budget says which record limits the band, and the two rank the units differently. The importances are themselves uncertain, with overlapping bands, and unit 3 is the least important with posterior probability 0.53 (0.60 exactly), not with certainty.

The next chapter reads the minimal cut sets off this enumeration: the smallest combinations of units whose loss takes the service down, each with its probability, ranked.

13.7 What to retain

Definition 13.1 (Service availability by enumeration)**.** For n units with unavailabilities $\mathbf{q}$, independent given $\mathbf{q}$, and a drawing with structure function ϕ, the service availability is $R_{\text{system}}(\mathbf{q}) = \sum_{\mathbf{s}} \phi(\mathbf{s}) \prod_i q_i^{1-s_i}(1-q_i)^{s_i}$, a multilinear polynomial whose gradient with respect to $\mathbf{q}$ is minus the vector of Birnbaum importances $R_{\text{system}}(1_i) - R_{\text{system}}(0_i)$. Under independent posteriors $\mathbb{E}[R_{\text{system}}(\mathbf{q}) \mid \mathbf{y}] = R_{\text{system}}(\mathbb{E}[\mathbf{q} \mid \mathbf{y}])$ exactly, and this expectation is the predictive probability that the service is delivered at a random instant; under a shared entry q, R_{system} is a polynomial in q whose expectation is taken through the moments $\mathbb{E}[q^r \mid \mathbf{y}]$, and the expected

unavailability exceeds the plug-in value for a redundant structure (and falls below it for a series). In the logits of the units, the check rows give the Laplace covariance $\mathbf{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1}$, diagonal because each record occupies its own column of $\mathbf{J}$, and the delta-method variance of the service is $\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}$ with $\mathbf{g} = \mathbf{D}\,\mathbf{I}_{\mathrm{B}}$, $\mathbf{D} = \mathrm{diag}(q_i(1-q_i))$. Every answer is computed from the MAP state and that covariance: the expectation by a Gauss–Hermite sum over the state, the band of $U_{\mathrm{system}} = 1 - R_{\mathrm{system}}$ on $\ln U$, and the tail from the normal on $\ln U$; the exact value stands beside it, from the Beta moments and the quantile map when U is monotone in one shared entry, or from draws. An open drawing is a discrete entry of the state, and the mixture over it adds the between-drawing variance $\bar{\mathbf{U}}^{\mathsf{T}}\mathbf{C}_{\pi}\bar{\mathbf{U}}$ to the variance within the drawings.

Three conclusions carry beyond the hall. The system model is the enumeration, and every closed form is a special case of it, so arrangements without a closed form need no new theory. Whether the plug-in value is exact is decided by the posterior's structure, not by the drawing: independent entries make it the exact expectation, a shared entry restores a Jensen gap that grows with the redundancy. And a target well below the plug-in value cannot be demonstrated by records of the units that predict it is not met.

Table 13.16 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 13.16: The terms of system structures and their posterior and where each is defined.

Term	Expression	Units	Meaning	Defined in
Unit unavailability	$q_i = \Pr(S_i = 0)$; logit ξ_i	–	long-run fraction down	13.1.3
Configuration probability	$\prod_i q_i^{1-s_i}(1-q_i)^{s_i}$	–	independent units	13.3.1
Structure function	$\phi(\mathbf{s}) \in \{0,1\}$	–	the drawing	13.3.1
Service availability	$R_{\text{system}} = \sum_{\mathbf{s}} \phi(\mathbf{s}) \Pr(\mathbf{s} \mid \mathbf{q})$	–	the enumeration	13.3.1
Series, parallel	$\prod_i (1-q_i)$; $1 - \prod_i q_i$	–	all needed; one needed	13.3.1
k-out-of-n, equal q	$\sum_{j \geq k} \binom{n}{j} (1-q)^j q^{n-j}$	–	binomial tail	13.3.1
Order of a structure	$U \approx \binom{n}{k-1} q^{n-k+1}$	–	smallest loss that takes it down	13.3.1
Pivotal decomposition	$R_{\text{system}} = q_i R_{\text{system}}(0_i) + (1-q_i) R_{\text{system}}(1_i)$	–	multilinearity	13.3.1
Gradient in the state	$\mathbf{g} = \mathbf{D}\,\mathbf{I}_{\text{B}}$, $\mathbf{D} = \operatorname{diag}(q_i(1-q_i))$	–	one entry per unit	13.3.1
No Jensen gap	$\mathbb{E}[R_{\text{system}}(\mathbf{q})] = R_{\text{system}}(\mathbb{E}[\mathbf{q}])$	–	independent posteriors	13.3.1
Expectation over the state	$\frac{1}{\sqrt{\pi}} \sum_m w_m h(\hat{\xi} + \sqrt{2}\,\sigma_\xi t_m)$	–	Gauss–Hermite sum	13.3.1
Beta moments	$\mathbb{E}[q^r] = \prod_{c<r} (a+c)/(a+b+c)$	–	shared entry, exact	13.3.1
Jensen gap, leading order	$\frac{1}{2} \operatorname{tr}(\mathbf{H}_q \boldsymbol{\Sigma}_q) = \frac{1}{2}\, \mathbf{1}^{\mathsf{T}} \mathbf{H}_q \mathbf{1}\, \operatorname{Var}(q)$	–	shared entry	13.3.1
Quantile map	$U_\alpha = U(q_\alpha)$; $\Pr(U < u^*) = I_{q^*}(a,b)$	–	exact, one shared entry	13.3.1
Delta method on $\ln U$	$\mathbf{g}_{\ln U}^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g}_{\ln U}$, $\mathbf{g}_{\ln U} = \mathbf{g}/U$	–	band on log scale	13.3.1
Rare-event approximation	$U_{\text{series}} \leq \sum_l U_l$, error $\leq \sum_{l<l'} U_l U_{l'}$	–	series of subsystems	13.3.4
Birnbaum importance	$I_{\text{B},i} = R_{\text{system}}(1_i) - R_{\text{system}}(0_i) = \partial U / \partial q_i$	–	unit critical	13.3.4
Design by arrangement	$U(a) = \Pr(U_a < u^* \mid \mathbf{y})$	–	probability an arrangement meets the target	13.3.5
Tail on the logarithm	$\Phi((\ln u^* - \ln U(\hat{\mathbf{z}}))/\sigma_{\ln U})$	–	target, state route	13.3.2
Mixture over drawings	$\sum_d \pi_d \bar{U}_d$; $\sum_d \pi_d \text{sd}_d^2 + \bar{\mathbf{U}}^{\mathsf{T}} \mathbf{C}_\pi \bar{\mathbf{U}}$	–	open structure	13.3.6
Check row	$r_\ell = (n_\ell q - k_\ell)/\sqrt{n_\ell q(1-q)}$, $\mathbf{J}_\ell = \sqrt{n_\ell q(1-q)}\, \mathbf{e}_{c(\ell)}^{\mathsf{T}}$; prior $\text{Beta}(a_0, b_0)$ as a_0 in $a_0 + b_0$ pseudo-checks	–	working residual and Jacobian row	13.1.6
Laplace covariance	$\boldsymbol{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1} = \operatorname{diag}((a_i + b_i)/(a_i b_i))$	–	disjoint columns	13.2
Preposterior	$\boldsymbol{\Sigma}_a = (\boldsymbol{\Sigma}^{-1} + \boldsymbol{\Lambda}_a)^{-1}$, $\mathbf{g}^{\mathsf{T}} \boldsymbol{\Sigma}_a \mathbf{g}$	–	next checks	13.6.2
Annual downtime	$8760\,\text{h} \times U$	h	hours per year	13.3.1

13.8 Exercises

Exercise 13.1. *Formulation.* Three units have unavailabilities $q_1 = 0.01$, $q_2 = 0.02$ and $q_3 = 0.05$, taken as known. (a) Compute the unavailability of the three in series and in parallel, and convert each to hours per year. (b) Write the question whether the series arrangement has an availability of at least 0.95 with an assurance of 0.9, for uncertain unavailabilities, as a

probability over the state. ((a) Series $1 - 0.99 \times 0.98 \times 0.95 = 0.07831$, 686 h per year; parallel $0.01 \times 0.02 \times 0.05 = 1.0 \times 10^{-5}$, 0.088 h per year. (b) $\Pr(1 - \prod_i (1 - q_i) \le 0.05 \mid \mathbf{y}) \ge 0.9$, a threshold probability of a derived quantity of the three logits.)

Exercise 13.2. *Formulation.* For a 2-out-of-3 arrangement, list the eight configurations with their structure function, derive the service availability from the enumeration for general unavailabilities, and reduce it for equal q. Evaluate it for the units of the previous exercise and for $q = 0.02$, in hours per year. (The four configurations with at least two units up give $R_{\text{system}} = R_1R_2R_3 + (1 - R_1)R_2R_3 + R_1(1 - R_2)R_3 + R_1R_2(1 - R_3) = R_1R_2 + R_1R_3 + R_2R_3 - 2R_1R_2R_3$ with $R_i = 1 - q_i$; for equal q, $U = 3q^2 - 2q^3$. For $(0.01, 0.02, 0.05)$, $U = 1.68 \times 10^{-3}$, 14.7 h per year; for $q = 0.02$, $U = 1.184 \times 10^{-3}$, 10.4 h per year.)

Exercise 13.3. *Derivation.* Prove that the service availability of any drawing is multilinear in the unit unavailabilities, and that its posterior expectation equals its value at the posterior means when the unavailabilities are independent a posteriori. State which step fails for a shared unavailability, and show it for a parallel pair. (Each configuration probability is $\prod_i f_{s_i}(q_i)$ with $f_0(q) = q$, $f_1(q) = 1 - q$, of first degree in each q_i, and R_{system} is a sum of such products. Then $\mathbb{E}[\prod_i f_{s_i}(q_i)] = \prod_i \mathbb{E}[f_{s_i}(q_i)]$ by independence and $= \prod_i f_{s_i}(\bar{q}_i)$ by linearity. For a shared q the first equality fails, since the factors are functions of one variable: a parallel pair has $U = q^2$ and $\mathbb{E}[q^2] = \bar{q}^2 + \mathrm{Var}(q) > \bar{q}^2$.)

Exercise 13.4. *Derivation.* (a) Derive the derivative of the k-out-of-n unavailability with respect to an equal unavailability q, and the Birnbaum importance of one unit; evaluate both for 3-out-of-4 at $q = 0.03$. (b) For three units in series with unavailabilities q_1, q_2, q_3, derive the error of the rare-event approximation, and evaluate it for $(0.01, 0.02, 0.05)$. ((a) $U = \Pr(\text{Binomial}(n, q) \ge n - k + 1)$, and the derivative of a binomial upper tail $\Pr(X \ge r)$ is $n\binom{n-1}{r-1}q^{r-1}(1 - q)^{n-r}$, so $\mathrm{d}U/\mathrm{d}q = n\binom{n-1}{k-1}q^{n-k}(1 - q)^{k-1}$ and each unit's importance is $\binom{n-1}{k-1}q^{n-k}(1 - q)^{k-1}$; for 3-out-of-4, $12q(1 - q)^2 = 0.3387$ and $3q(1 - q)^2 = 0.0847$. (b) $1 - \prod_i (1 - q_i) = \sum_i q_i - \sum_{i<j} q_iq_j + q_1q_2q_3$, so the error is $\sum_{i<j} q_iq_j - q_1q_2q_3 = 0.0017 - 0.00001 = 0.00169$, against the exact 0.07831, 2.2 %.)

Exercise 13.5. *Application.* A 2-out-of-3 pump set has a pooled record of three pumps found down in 60 checks, posterior $q \sim \text{Beta}(3, 57)$ shared by the three pumps. Alternatively the same checks are kept per pump, one finding in 20 checks each, with posteriors $\text{Beta}(1, 19)$ for each pump, independent. (a) Compute the expected unavailability of the pump set under both treatments and the plug-in value, and explain the difference. (b) For the shared treatment, whose state is the one logit ξ, write the curvature Λ of the check row at the MAP state, the variance σ_ξ^2 and the slope $g = \mathrm{d}\ln U/\mathrm{d}\xi$; for the separate treatment write the state, the Jacobian of the check rows at the MAP state, $\mathbf{\Sigma}$ and the gradient $\mathbf{g}_{\ln U}$. Compare the two log-scale spreads $g\sigma_\xi$ and $(\mathbf{g}_{\ln U}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}_{\ln U})^{1/2}$. ((a) Shared: $\mathbb{E}[U] = 3\mathbb{E}[q^2] - 2\mathbb{E}[q^3]$ with $\mathbb{E}[q^2] = 12/3660 = 3.279 \times 10^{-3}$ and $\mathbb{E}[q^3] = 60/226\,920 = 2.644 \times 10^{-4}$, so $\mathbb{E}[U] = 9.31 \times 10^{-3}$. Separate: the plug-in value at the common mean 0.05, $3 \times 0.0025 - 2 \times 0.000125 = 7.25 \times 10^{-3}$, exact by multilinearity. Same marginal means, a factor of 1.28: the shared entry makes coincident outages more likely, $\mathbb{E}[q^2]$ against $\bar{q}^2 = 2.5 \times 10^{-3}$. (b) Shared: the state is the number ξ, one row, with curvature $\Lambda = 60 \times 0.05 \times 0.95 = 2.85 = 1.6882^2$, variance $\sigma_\xi^2 = 1/\Lambda = 0.35088 = 60/171$ and slope $g = 6q(1 - q) \cdot q(1 - q)/U = 1.8672$. Separate: $\mathbf{z} = (\xi_1, \xi_2, \xi_3)^{\mathsf{T}}$, three rows, $\mathbf{J} = \mathrm{diag}(0.9747,\ 0.9747,\ 0.9747)$, $\mathbf{\Sigma} = \mathrm{diag}(1.0526,\ 1.0526,\ 1.0526)$, $\mathbf{g}_{\ln U} = (0.6224,\ 0.6224,\ 0.6224)^{\mathsf{T}}$. Both spreads are 1.106. The shared entry is the separate state restricted to the direction $\mathbf{1}$: its slope is $g = \mathbf{1}^{\mathsf{T}}\mathbf{g}_{\ln U}$ and its information $\Lambda = \mathbf{1}^{\mathsf{T}}\mathbf{\Sigma}^{-1}\mathbf{1}$, and with equal entries $g^2/\Lambda = \mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}$. The delta method sees only that, and the difference between the treatments lives in the curvature, the Jensen gap of part (a).)

Exercise 13.6. *Application.* The pump set of the previous exercise, with its shared posterior, is 2-out-of-3 if two pumps are needed and 1-out-of-3 if one pump can carry the load, with

weights 0.6 and 0.4. (a) Compute the reported unavailability under the mixture. (b) Find the weight on 2-out-of-3 below which the mixture report is less than half the 2-out-of-3 report. (c) The posterior standard deviation of the 2-out-of-3 unavailability is 1.04×10^{-2} and its 95 % band $[3.4 \times 10^{-4},\ 3.8 \times 10^{-2}]$; find the weights at which the mixture report differs from the 2-out-of-3 report by less than one standard deviation, and at which it stays inside the band. ((a) 1-out-of-3: $\mathbb{E}[q^3] = 2.64 \times 10^{-4}$; mixture $0.6 \times 9.31 \times 10^{-3} + 0.4 \times 2.64 \times 10^{-4} = 5.69 \times 10^{-3}$. (b) $w\, 9.31 \times 10^{-3} + (1 - w)\, 2.64 \times 10^{-4} < 4.65 \times 10^{-3}$ for $w < 0.485$. (c) The difference is $(1 - w) \times 9.04 \times 10^{-3}$, below one standard deviation for every weight, and the mixture stays inside the band for $w \geq 0.008$: the band of the unavailabilities is so wide that it hides the open structure, while the structure still changes the expectation by up to a factor of 35. The doubt about the drawing must be reported through the mixture, not absorbed into the band.)

Chapter 14

Fault Trees and Minimal Cut Sets

14.1 The combinations that take a service down

The data hall of Chapter 13 has a figure for its loss of cooling, 9.1×10^{-3}, about 80 hours a year (6.7×10^{-3}, about 58 hours, from the exact posterior), computed for its four CRAC units alone. That drawing left out the two pumps that feed the units with chilled water, either of which suffices, and the one control system that sequences them, and this chapter adds both. The owner now asks the facilities engineer where the loss of cooling comes from and what to buy to lower it. The records are three series of availability checks: of a hundred checks of the CRAC units, three found a unit down; of four hundred checks of the pumps, four; of a thousand checks of the controller, one. The first reading lists the combinations of failed units that take the cooling down, with the fractions found down standing for the units' unavailabilities: the controller alone, 1.0×10^{-3}; each of the six pairs of CRAC units, $0.03^2 = 9.0 \times 10^{-4}$; the two pumps together, 1.0×10^{-4}. Their sum, 6.5×10^{-3}, is read as an upper bound on the probability of losing cooling, and the controller, the one unit whose failure alone stops the cooling and the largest entry of the list, is read as the hall's weak point: a second controller is the change to buy.

The engineer who must defend the purchase asks what the list rests on. Each entry is a power of a fraction found down, the fractions come from three, four and one units found down, and so few failures do not fix an unavailability closely. The list does not say whether the controller truly ranks first, how much of the loss each combination carries and how firmly, whether its sum bounds the loss that the records imply, which change removes most of it, or which further record would settle what remains open. The question of this chapter is the one the owner of the hall puts in plain words,

which combinations take the service down, and each one's probability, read off the enumeration and ranked, with the band that the evidence leaves on each.

What changes from Chapter 13 is the reading of its enumeration: the same configurations of up and down units are read for the sets of down units that fail the hall, rather than summed into one number, and every probability built from those sets is a function of a state of three logits with a posterior of its own.

Jargon: fault tree, top event, basic event, gate, OR gate, AND gate, voting gate, coherent structure, cut set, minimal cut set, order of a cut set, rare-event approximation, inclusion–exclusion, Bonferroni bounds, Birnbaum importance, Fussell–Vesely importance, CRAC unit, chilled water

A *fault tree* is a drawing of the logic by which failures of components produce an undesired event of the system, the *top event*. Its leaves are the *basic events*, here the failure (being down) of one unit, and its internal nodes are *gates*: an *OR gate* occurs when any of its inputs occurs, an *AND gate* when all of them occur, and a *voting gate* k_{g}/n when at least k_{g} of its n inputs occur. A structure is *coherent* when failing one more unit can never restore the service. A *cut set* is a set of basic events whose joint occurrence produces the top event, whatever the other units do; it is *minimal* when no proper subset is a cut set, and its *order* $|c|$ is the number of events in it. The *rare-event approximation* sums the probabilities of the minimal cut sets as if they were disjoint. *Inclusion–exclusion* is the exact formula for the probability of a union of events, and the *Bonferroni bounds* are its truncations, which alternate above and below the exact value. The *Birnbaum importance* of a basic event is the derivative of the top event's probability with respect to that event's probability; the *Fussell–Vesely importance* is the fraction of the top event's probability carried by the cut sets that contain the event. Both are importance measures, derived in a later chapter on importance measures and only named here. A *CRAC unit* (computer room air conditioner) cools the air of a data hall with *chilled water*, water cooled by a central plant and circulated by pumps.

Key idea

The minimal cut sets are the smallest down sets of the failed configurations of the 2^n-configuration enumeration. Given the units' unavailabilities, a cut set's probability is the product of its members' unavailabilities; under a posterior shared by its members it is a moment of that posterior, which exceeds the power of the mean and can reorder the expectations. Each cut-set probability is a function of the state, and the gradients of all of them form one matrix $\mathbf{G}$, a row per cut set and a column per state entry: a cut set's band is $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ with $\mathbf{g}$ its row, and to first order the top event's gradient is the sum of the rows, which is the rare-event approximation written in the state. The rare-event sum of the cut-set probabilities bounds the top event from above at every state, and the enumeration gives it exactly; the share of the top event that one group of cut sets carries is itself uncertain, and whether it exceeds one half is a posterior probability, not a fact of the drawing, as is the order of any two cut sets.

14.1.1 From one number to the combinations behind it

Chapter 13 wrote the service unavailability as a sum over configurations. For n units with unavailabilities $\mathbf{q} = (q_1, \dots, q_n)$, each unit is up ($s_i = 1$) or down ($s_i = 0$), the configuration $\mathbf{s} = (s_1, \dots, s_n)$ has the probability $\Pr(\mathbf{s} \mid \mathbf{q}) = \prod_i q_i^{1-s_i}(1-q_i)^{s_i}$ under independence given $\mathbf{q}$, and the structure function $\phi(\mathbf{s})$ equals one when the service is delivered and zero otherwise. The unavailability of the service is

$$U_{\mathrm{sys}}(\mathbf{q}) = \sum_{\mathbf{s}\in\{0,1\}^n} (1 - \phi(\mathbf{s}))\, \Pr(\mathbf{s} \mid \mathbf{q}), \tag{14.1}$$

a sum over the failed configurations. The sum is a single number, and it hides its own composition. Two plants with the same U_{sys} can owe it to quite different combinations: one to a single unit whose failure alone stops the service, the other to many pairs of redundant units failing together. The remedies differ, and so do the records that would narrow the band.

14.1.2 The fault tree as a generator of combinations

The structure function of Chapter 13 was written from the arrangement of the units: in series, in parallel, or so many out of so many. The fault tree states the same function from the other side, as the logic of failure: the top event occurs if this OR that occurs, and that occurs if these AND those occur. The companion volume *Reliability Engineering* develops the fault-tree symbols, the construction of a tree from a system's functions, and the Boolean algebra that reduces a tree to its minimal cut sets; those are taken as known and are not repeated. The tree is used here only as a drawing that, evaluated on a configuration, says whether the top event holds. The minimal cut sets are then read off the enumeration of all configurations, with no algebra, and every probability derived from them is an expectation over the posterior of the unavailabilities, so that it comes with its band. The algebraic route is needed when 2^n configurations are too many to enumerate, and Section 14.3.1 says where that limit lies.

14.1.3 The physical system and its graph (step 0)

The data hall of Chapter 13 is cooled by four CRAC units, C1 to C4, in an N+1 arrangement: the hall is cooled while at least three of the four run, and it loses cooling when two or more are down, the 3-out-of-4 structure of that chapter. The CRAC units take their chilled water from a central plant through two pumps, P1 and P2, either of which can carry the flow, so that the hall loses its chilled water only when both pumps are down. A single control system, K, sequences the pumps and the CRAC units, and its failure takes the cooling down whatever the other units do. The power feed of Chapter 13 belongs to a different top event, the loss of power to the hall, and is left out. Seven basic events remain: the failures of C1 to C4, P1, P2 and K. Chapter 13 left open whether three or two of the four CRAC units must run, weighting the three-unit drawing at 0.7 and the two-unit drawing at 0.3; this chapter takes the three-unit drawing, of weight 0.7.

Figure 14.1 draws the plant as a graph. The gray box on the left is the chilled-water plant, a pinned boundary; the solid green arrows are the chilled-water path, through the two pumps and a common header, the small circle, to the four CRAC units, and from the CRAC units the cooled air into the hall, the gray box on the right, which is cooled when $\phi(\mathbf{s}) = 1$. The green boxes are the units, each up or down, and the frames around the pumps and around the CRAC units mark the free entries of the state: one logit ξ_P shared by the two pumps, one logit ξ_C shared by the four CRAC units, and the controller's own logit ξ_K at the top, whose dotted arrows are its control signals to the two groups. The dashed arrows leave each group for its check record, the evidence attached as a count row.

The state entries. Each unavailability lies in $(0, 1)$, and the graph carries it as a logit, $\xi = \ln(q/(1-q))$, as Chapter 5 carried a failure probability. Units of the same design checked under one record share one entry, as the four CRAC units shared one in case A of Chapter 13, so the state has three entries,

$$\mathbf{z} = \begin{bmatrix} \xi_C \\ \xi_P \\ \xi_K \end{bmatrix}, \qquad q_{C1} = \cdots = q_{C4} = q_C, \qquad q_{P1} = q_{P2} = q_P, \qquad q_K, \tag{14.2}$$

while the enumeration runs over the configurations of the seven units. The two words are kept apart throughout: the *state* is the vector of logits the evidence informs, and a *configuration* is a list of which units are up and which are down. Sharing an entry is a statement about knowledge, not about physics: given q_C the four CRAC units fail independently, but the uncertainty about q_C is common to all four, and Section 14.3.2 shows what that does to a cut set. The cut sets, their probabilities, the top event, the share of a group and every gradient asked for are derived quantities, functions of the state.

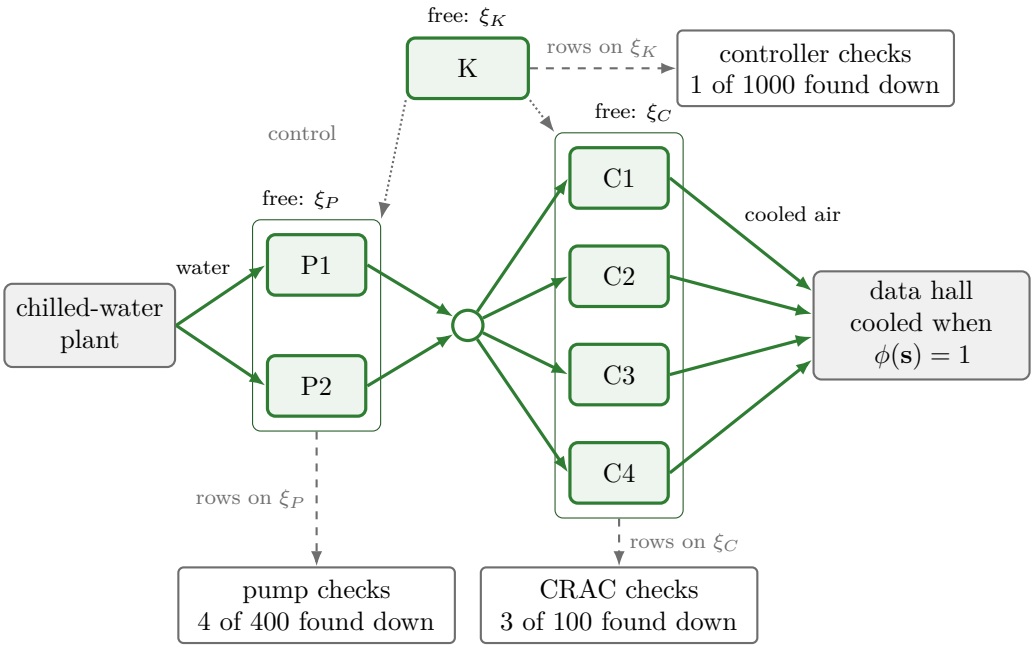


Figure 14.1: Plant graph of the cooling of the data hall. Pinned: the chilled-water plant and the hall (gray). Free: the logits of the unavailabilities, one shared by the two pumps (ξ_P), one shared by the four CRAC units (ξ_C) and one for the controller (ξ_K), carried by the units in the green frames. Solid arrows: the chilled-water path and the cooled air; dotted: the controller's signals. Observed: the availability checks of each group (dashed), with the number of checks that found a unit down.

14.1.4 The fault tree of the top event

Figure 14.2 draws the fault tree. The box at the top is the top event, the loss of hall cooling. Under it an OR gate, the shape with a curved base and a pointed top, combines three inputs. On the left, a voting gate marked 2/4 occurs when two or more of the four CRAC units are down; its four inputs are the basic events C1 to C4, drawn as circles. In the middle, an AND gate, the shape with a flat base and a round top, occurs when both pumps are down. On the right, the controller's failure, a basic event, enters the OR gate directly. Under each circle is the fraction of its group's checks that found a unit down, from Section 14.1.5.

The tree states in a drawing what the arrangement of Section 14.1.3 states in words, and the two must agree; building the tree is the step at which a misunderstanding of the plant is found. The tree also fixes what is not modeled: no basic event stands for a common supply, a shared maintenance crew or a cause that fails two units at once.

14.1.5 The check records

Each group's record is a series of availability checks. A check finds a unit up or down, a Bernoulli trial with probability q of finding it down, the evidence of Chapter 5. The three records are those of Chapter 13's case A and two more kept the same way. Of $n_C = 100$ checks of the CRAC units, $k_C = 3$ found a unit down; of $n_P = 400$ checks of the pumps, $k_P = 4$; of $n_K = 1000$ checks of the controller, $k_K = 1$. The controller has the longest record and the fewest units found down; the CRAC units the shortest record and the largest fraction found down. No other information on the units is used: each record stands alone, without a prior drawn from other halls. The evidence $\mathbf{y}$ is the three records.

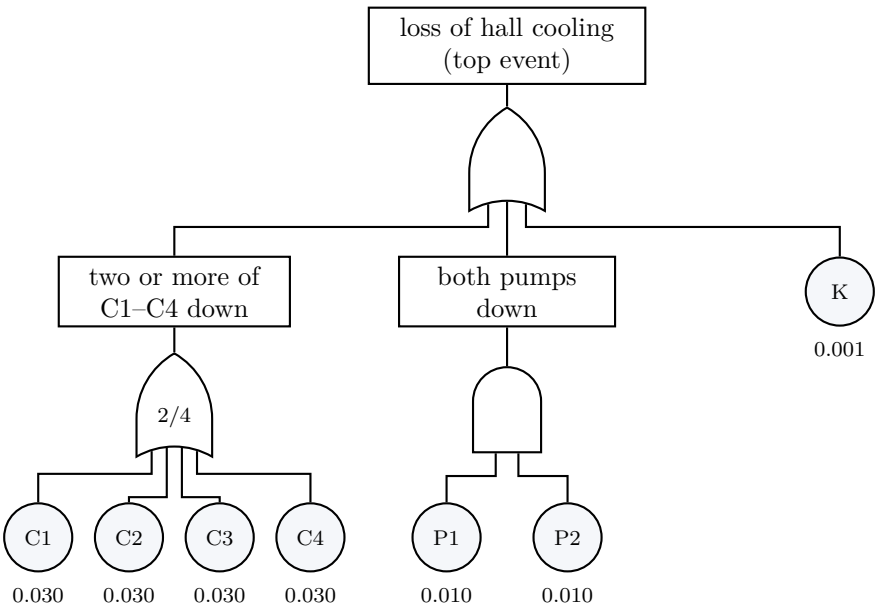


Figure 14.2: Fault tree of the loss of hall cooling: top event $(\geq 2$ of C1–C4 down$) \vee (\mathrm{P1} \wedge \mathrm{P2}) \vee \mathrm{K}$. OR gate at the top, voting gate 2/4 on the CRAC units, AND gate on the pumps, the controller's basic event entering the OR gate directly; basic events as circles, each with the fraction of its group's checks that found a unit down (3/100, 4/400, 1/1000).

14.1.6 The probabilistic questions (step 1)

The owner asks six questions of the tree and the three records, and step 1 writes each as a quantity over the state, with the unknown on the left of the bar and the records on the right. Table 14.1 lists them: the first column is the question in engineering language; the second the general form of its category in the notation of Chapter 2, for a quantity $f(\mathbf{z})$ derived from the state and a level f^* of it; the third its mathematical form for the hall; the fourth its category; and the last the subsection where it is answered. The first two together are the owner's question as stated in Section 14.1. The first needs no posterior at all: the minimal cut sets are a property of the drawing. The second and third are estimation of derived quantities, the cut-set probabilities, their order and the top event; the third also measures the error of an approximation. The fourth is a prediction under a threshold, the posterior probability that a derived share exceeds one half. The fifth is a design question answered without new evidence: a change to the plant itself is valued by the present posterior under each alternative drawing, whereas the next record of step 7 values evidence not yet taken. The sixth is estimation again, of a derived quantity that is a slope: the change of the top event per change of one unit's unavailability. The atlas of Appendix A lists the same forms.

14.1.7 The evidence as rows (step 2)

Step 2 attaches each check record to its entry as a count row. A check of a CRAC unit is a Bernoulli trial on ξ_C, down with probability q_C and up with probability $1 - q_C$, and the hundred checks together are one count row on ξ_C; in the same way the pump record is one row on ξ_P and the controller record one row on ξ_K. Which row acts on which entry is a property of the graph, and it fixes where the Jacobian of the rows can be nonzero, marked $*$:

$$\mathbf{J} = \begin{bmatrix} * & 0 & 0 \\ 0 & * & 0 \\ 0 & 0 & * \end{bmatrix} \begin{matrix} \text{CRAC checks} \\ \text{pump checks} \\ \text{controller checks} \end{matrix}, \qquad \text{columns } \xi_C, \xi_P, \xi_K. \tag{14.3}$$

Table 14.1: The questions asked of the hall's cooling, the general form of their category, their mathematical form for the hall, their category and where each is answered. $\mathbf{y}$ is the three check records; $\mathcal{M}$ the family of minimal cut sets; P_c the probability of cut set c given $\mathbf{q}$; π_{C} the share of the CRAC pairs in the cut-set sum; f_d the top event under drawing d; $\mathbf{G}$ the matrix of the cut-set probabilities' gradients with respect to the state.

Engineering question	General form	Mathematical form	Category	Answered in
Which combinations take the service down?	none: a property of the drawing	$\mathcal{M}$, the minimal down sets of the failed configurations of the enumeration	read-off of the drawing; no posterior	14.3.1
How likely is each combination, and which dominate?	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$, quantiles; $\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y})$	$\mathbb{E}[P_c(\mathbf{q}) \mid \mathbf{y}]$ for $c \in \mathcal{M}$, with bands; $\Pr(P_c > P_{c'} \mid \mathbf{y})$	estimation of derived quantities	14.3.2
What is the probability of losing cooling, and how well does the rare-event sum do?	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$, quantiles	$\mathbb{E}[U_{\mathrm{sys}}(\mathbf{q}) \mid \mathbf{y}]$ by enumeration, against $\mathbb{E}[S_1 \mid \mathbf{y}]$	estimation; error of an approximation	14.3.2
How much of the top event's probability is in the CRAC pairs, and is that conclusion firm?	$\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y})$	π_{C} and $\Pr(\pi_{\mathrm{C}} > \frac{1}{2} \mid \mathbf{y})$	prediction under a threshold	14.3.3
Which single change removes the most?	$\mathbb{E}[f_d(\mathbf{z}) \mid \mathbf{y}]$ for each drawing d	$\mathbb{E}[U_{\mathrm{sys}} \mid \mathbf{y}]$ with a redundant controller and with a fifth CRAC unit	design, by alternative drawings	14.3.4
Which basic event's probability matters most to the top event?	$\Pr(f(\mathbf{z}) \leq f^* \mid \mathbf{y})$, location and quantiles	$f = \partial U_{\mathrm{sys}}/\partial q_i$, with bands; at $\hat{\mathbf{z}}$, $\mathbf{g} = \partial U_{\mathrm{sys}}/\partial \mathbf{z}$ against the columns of $\mathbf{G}$	estimation of a derived quantity	14.3.5

No row acts on two entries, at any state. No row is a physics row either, so no model tolerance σ_p arises. The tree adds no rows: it is the derived quantity, a function of the state through equation (14.1), and every question of Table 14.1 after the first is a question about that function or about its terms.

The modeling choices. Three choices make the rows and the cut-set probabilities read from them, and each buys something and can fail. The units are taken as down independently given their unavailabilities: this makes a configuration's probability a product and a cut set's probability the product of its members' unavailabilities, and it fails under a cause that takes two units down at once, which the next chapter models. Units of one design share one entry of the state: this lets three records inform seven units and makes a CRAC pair's probability a moment of one posterior, and it fails if one CRAC unit is worse than the others, which group totals cannot reveal (Section 14.6.6). Each posterior is its record alone, Beta($k_j,\ n_j - k_j$), the limit of the Beta update of Chapter 5 as the prior's two parameters tend to zero: this lets the records set every answer, and it is fragile where few units were found down, three, four and one, so that any prior moves the means, the controller's most; Section 14.4.3 measures by how much.

The rows of the check records. A record alone is a flat prior on the logit, so each entry has one row, the count of units found down. With k_j units found down in n_j checks of group j and $q_j = 1/(1+\mathrm{e}^{-z_j})$, the objective is the sum of the three rows' negative log-likelihoods,

$$\mathcal{L}(\mathbf{z}) = \sum_{j\in\{C,P,K\}} \mathcal{L}_j(\mathbf{z}), \qquad \mathcal{L}_j(\mathbf{z}) = -k_j \ln q_j - (n_j - k_j)\ln(1-q_j), \tag{14.4}$$

and each row enters Gauss–Newton in the working form of a count row, the binomial-in-the-logit pair of Chapter 4, equation (4.43), written for one entry of a several-entry state in Chapter 13, equation (13.5):

$$\begin{aligned} r(\mathbf{z}) &= \begin{bmatrix} (100\,q_C - 3)/\sqrt{100\,q_C(1-q_C)} \\ (400\,q_P - 4)/\sqrt{400\,q_P(1-q_P)} \\ (1000\,q_K - 1)/\sqrt{1000\,q_K(1-q_K)} \end{bmatrix}, \\ \mathbf{J}(\mathbf{z}) &= \begin{bmatrix} \sqrt{100\,q_C(1-q_C)} & 0 & 0 \\ 0 & \sqrt{400\,q_P(1-q_P)} & 0 \\ 0 & 0 & \sqrt{1000\,q_K(1-q_K)} \end{bmatrix}, \end{aligned} \tag{14.5}$$

three rows on three entries and a 3×3 Jacobian with the pattern (14.3). The gradient of the objective is $\mathbf{J}^{\mathsf{T}} r$, and its curvature is $\mathbf{J}^{\mathsf{T}}\mathbf{J} + \mathbf{S}$ with $\mathbf{S}$ the curvature term of the count rows; the logit is the canonical parameter of the binomial row, so $\mathbf{S} = 0$ at every state and the curvature equals $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ exactly: Newton's method and Gauss–Newton coincide. Each diagonal entry of $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ is $n_j q_j(1-q_j)$, close to $n_j q_j$, so the information a record carries is about its expected number of units found down, not its number of checks, and one further check of group j adds $q_j(1-q_j)$ to it. With the rows normalized, as here, the information matrix is $\mathbf{\Lambda} = \mathbf{J}^{\mathsf{T}}\mathbf{J}$.

14.2 The posterior of the state (step 3)

The six questions of Table 14.1 rest on one drawing and one posterior. Four objects of the fit recur in the specialized forms. The *MAP state* $\hat{\mathbf{z}}$ is the minimum of the objective that the three count rows put on the state; the *information matrix* $\mathbf{\Lambda} = \mathbf{J}^{\mathsf{T}}\mathbf{J}$ is the curvature of that objective at the minimum, diagonal here by the pattern (14.3); the *Laplace covariance* $\mathbf{\Sigma} = \mathbf{\Lambda}^{-1}$ makes the posterior of the state approximately normal, $\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$; and a derived quantity f is carried by its *gradient* $\mathbf{g} = \partial f/\partial \mathbf{z}$ at $\hat{\mathbf{z}}$, with the variance $\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}$. Beside them stands the exact

posterior that this problem admits, three independent Beta posteriors, one per record: their moments give every expectation in closed form, their quantiles the band of any quantity that depends on one entry, and their draws the bands and threshold probabilities of quantities that depend on several. Every answer is computed from $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$, and the exact posterior gives the value beside it. This section derives that posterior and computes it for the hall; Section 14.3 reads every answer from it.

14.2.1 The posterior from the three rows

The fit solves the normal equations of Chapter 3,

$$(\mathbf{J}^{\mathsf{T}}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^{\mathsf{T}}r, \qquad \mathbf{z}_{j+1} = \mathbf{z}_j + \Delta\mathbf{z}, \tag{14.6}$$

all evaluated at $\mathbf{z}_j$, which here are three separate equations, one per entry, because $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ is diagonal. The gradient vanishes where every working residual of equation (14.5) is zero, at $q_j = k_j/n_j$, so the MAP state and the Laplace covariance are

$$\hat{z}_j = \ln\frac{k_j}{n_j - k_j}, \qquad \boldsymbol{\Sigma} = \operatorname{diag}\Big(\frac{1}{n_j\,\hat{q}_j(1-\hat{q}_j)}\Big) = \operatorname{diag}\Big(\frac{1}{k_j} + \frac{1}{n_j - k_j}\Big), \tag{14.7}$$

with zero off-diagonal entries: the three posteriors are independent, each from its own record, a result of the pattern (14.3) and not an assumption. The answer the fit returns is the entry $\hat{z}_j$ with $\sqrt{\Sigma_{jj}}$, and the values of the derived quantities at $\hat{\mathbf{z}}$ are called the plug-in values.

The problem also admits the exact posterior. The objective of equation (14.4) is the negative logarithm of a Beta density written as a density of the logit, and each unavailability has

$$q_j \mid \mathbf{y} \sim \operatorname{Beta}(a_j,\ b_j), \qquad a_j = k_j, \quad b_j = n_j - k_j, \qquad \mathbb{E}[q_j \mid \mathbf{y}] = \frac{k_j}{n_j} = q(\hat{z}_j), \tag{14.8}$$

independently across the three groups, so that the plug-in values are the posterior means. Every expectation of a product of powers of the three unavailabilities is a product of the moments of equation (13.29); a band of a quantity of one entry is a Beta quantile mapped through it; and a band of a quantity of several entries is read from draws of the three posteriors. The exact variance of a logit is the counterpart of Σ_{jj}: the logit of a $\operatorname{Beta}(a,b)$ variable is the difference of the logarithms of two independent gamma variables of shapes a and b, so with ψ' the trigamma function of Chapter 8,

$$\operatorname{Var}(\xi_j \mid \mathbf{y}) = \psi'(a) + \psi'(b), \tag{14.9}$$

which exceeds $1/a + 1/b$ and does so most when a is small.

14.2.2 The posterior of the hall

The hall's inputs are the three check records, three CRAC units found down in 100 checks, four pumps in 400 and one controller in 1000.

Table 14.2 runs the normal equations (14.6) from the state $\mathbf{z}_0 = (-4,\ -4,\ -4)^{\mathsf{T}}$, $q = 0.018$ for every group. Each row is an iteration; the columns give the iteration number, the three entries of the state and the three working residuals of equation (14.5). The CRAC and pump entries settle in four steps; the controller's, whose record holds one unit found down and whose start is far above it, takes six, its residual roughly halving over the first two steps and then falling quadratically.

Table 14.2: Gauss–Newton iterations of the normal equations for the three logits, rows and Jacobian of equation (14.5), from $\mathbf{z}_0 = (-4,\ -4,\ -4)^{\mathsf{T}}$.

Iteration	ξ_C	ξ_P	ξ_K	CRAC r_C	Pumps r_P	Controller r_K
0	−4.0000	−4.0000	−4.0000	−0.9040	1.2018	4.0417
1	−3.3198	−4.4522	−4.9617	0.2669	0.2847	2.2654
2	−3.4652	−4.5856	−5.8239	0.0185	0.0190	1.1360
3	−3.4760	−4.5951	−6.4865	0.0001	0.0001	0.4231
4	−3.4761	−4.5951	−6.8298	0.0000	0.0000	0.0769
5	−3.4761	−4.5951	−6.9039	0.0000	0.0000	0.0029
6	−3.4761	−4.5951	−6.9068	0.0000	0.0000	0.0000

The iteration ends at the MAP state, where every working residual is zero, and the Jacobian and the Laplace covariance there are

$$\hat{\mathbf{z}} = \begin{bmatrix} -3.4761 \\ -4.5951 \\ -6.9068 \end{bmatrix}, \qquad \mathbf{J}(\hat{\mathbf{z}}) = \begin{bmatrix} 1.7059 & 0 & 0 \\ 0 & 1.9900 & 0 \\ 0 & 0 & 0.9995 \end{bmatrix},$$
$$\boldsymbol{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1} = \begin{bmatrix} 2.9100 & 0 & 0 \\ 0 & 3.9600 & 0 \\ 0 & 0 & 0.9990 \end{bmatrix}^{-1} = \begin{bmatrix} 0.3436 & 0 & 0 \\ 0 & 0.2525 & 0 \\ 0 & 0 & 1.0010 \end{bmatrix}. \tag{14.10}$$

The MAP state maps to $\hat{q}_C = 0.030$, $\hat{q}_P = 0.010$ and $\hat{q}_K = 0.001$, the diagonal of $\boldsymbol{\Sigma}$ is $1/k_j + 1/(n_j - k_j)$ by equation (14.7), and the logits are read off with $\sqrt{\Sigma_{jj}} = 0.586$, 0.503 and 1.000.

Table 14.3 gives the exact posteriors of equation (14.8): each row is a group of basic events, the columns give its units, its record, its posterior, the posterior mean and the central 95 % interval. The posterior means are the plug-in values $\bar{q}_C = 0.030$, $\bar{q}_P = 0.010$ and $\bar{q}_K = 0.001$. The interval of q_C spans a factor of eleven, from 0.0063 to 0.071; that width, not the mean, drives most of what follows.

Table 14.3: The basic events of the loss of hall cooling, their check records and the posteriors of their unavailabilities. Units of one group share one unavailability; each posterior is the record alone.

Group	Units	Found down	Posterior	Mean	95 % interval
CRAC units	C1–C4	3 of 100	$q_C \sim \mathrm{Beta}(3,\ 97)$	0.030	[0.0063, 0.0711]
chilled-water pumps	P1, P2	4 of 400	$q_P \sim \mathrm{Beta}(4,\ 396)$	0.010	[0.0027, 0.0218]
control system	K	1 of 1000	$q_K \sim \mathrm{Beta}(1,\ 999)$	0.0010	[0.000025, 0.0037]

The pair $(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of equation (14.10) is what step 4 uses: every answer of the chapter after the first is computed from it. Because this problem also admits the exact posteriors, each answer has an exact value beside it, a function of the Beta parameters of Table 14.3, and the gap between the two is reported with it.

14.3 The answers (step 4)

The questions of Table 14.1 fall into four categories: a read-off of the drawing that needs no posterior (the first), estimation of derived quantities (the second, the third and the sixth), prediction under a threshold (the fourth) and design by alternative drawings (the fifth). The last

five are answered from the posterior of Section 14.2. The questions name the tree of Figure 14.2 and the two alternative trees, and the year of 8760 hours. Every answer from the posterior is computed from the MAP state $\hat{\mathbf{z}}$ and the covariance $\boldsymbol{\Sigma}$ of step 3, and that is the value the chapter reports; the exact Beta posteriors give the value set beside it, with the gap, those without a closed form from 4×10^5 triples (q_C, q_P, q_K) drawn from the three posteriors, with seed 14.

14.3.1 The minimal cut sets

Engineering question. Which combinations take the service down? The owner asks for the sets of failed units whose joint failure stops the cooling, and among them the essential ones.

General form. None of the five forms of Chapter 2 applies, because the answer involves no probability. With $\mathcal{D}(\mathbf{s})$ the set of units down in configuration $\mathbf{s}$ and ϕ the structure function, it is the family of the smallest down sets of the failed configurations,

$$\mathcal{M} = \min_{\subseteq} \{\mathcal{D}(\mathbf{s}) : \phi(\mathbf{s}) = 0\}, \tag{14.11}$$

a property of the drawing, the same under every posterior, and the evidence enters only when the sets are weighted, in the next question.

Specialized form. Two choices enter here. The structure is taken as coherent, failing one more unit never restoring the service: this is what makes the cut sets exactly the down sets of the failed configurations and lets the minimal ones represent the whole tree, and it fails when one failure can undo another, as a tripped interlock that lets a stopped unit run again; the enumeration verifies it (Section 14.4.2). The drawing is taken as the true arrangement of the plant: this is what lets the configurations of seven units be enumerated at all, and it fails if a pump cannot carry the full flow alone or a CRAC unit has its own local controller, which Section 14.6.6 discusses.

The question is equation (14.11) with ϕ the complement of the tree of Figure 14.2 and $\mathbf{s}$ the configurations of the seven units; no state enters.

The gates. Write $x_i = 1 - s_i$ for the indicator that unit i is down, and $\mathbf{x} = (x_1, \ldots, x_n)$. Each gate of the tree is a function of its inputs' indicators with values in $\{0, 1\}$:

$$\text{OR: } 1 - \prod_l (1 - x_l), \qquad \text{AND: } \prod_l x_l, \qquad k_{\mathrm{g}}/n\text{: } \mathbf{1}\Big(\sum_{l=1}^{n} x_l \geq k_{\mathrm{g}}\Big). \tag{14.12}$$

The OR form is one minus the indicator that no input occurs; the AND form is one only when every input is one. Composing the gates of Figure 14.2 from the leaves upward gives the top-event indicator of the hall, written ϕ^{D} for a reason given below,

$$\phi^{\mathrm{D}}(\mathbf{x}) = 1 - (1 - v(\mathbf{x}))(1 - x_{P1}x_{P2})(1 - x_K), \qquad v(\mathbf{x}) = \mathbf{1}(x_{C1} + x_{C2} + x_{C3} + x_{C4} \geq 2). \tag{14.13}$$

The top event holds in configuration $\mathbf{s}$ exactly when the service is not delivered, so the tree evaluated on a configuration is the complement of the structure function of Chapter 13,

$$\phi^{\mathrm{D}}(\mathbf{1} - \mathbf{s}) = 1 - \phi(\mathbf{s}), \tag{14.14}$$

and equation (14.1) is the sum of ϕ^{D} over the enumeration weighted by the configuration probabilities. The tree and the structure function are one object written in the indicators of failure and of success; a function related to ϕ in this way is its *dual* structure function, hence the name ϕ^{D}. For the hall, $\phi(\mathbf{s})$ is the product of three factors, the 3-out-of-4 factor of

the CRAC units, the parallel factor of the pumps and the series factor of the controller, and equation (14.13) is that product written for failure.

Cut sets. For a configuration $\mathbf{s}$, $\mathcal{D}(\mathbf{s}) = \{i : s_i = 0\}$ is the set of its down units. A set c of units is a *cut set* when every configuration with all units of c down is failed:

$$c \text{ is a cut set} \iff \phi^{\mathrm{D}}(\mathbf{x}) = 1 \text{ for every } \mathbf{x} \text{ with } x_i = 1 \text{ for all } i \in c. \tag{14.15}$$

The tree of the hall is coherent: each gate in equation (14.12) is nondecreasing in each input, and so is their composition, so a failed configuration stays failed when more units go down. Two consequences follow, each in one line. If $\mathbf{s}$ is a failed configuration, its down set $\mathcal{D}(\mathbf{s})$ is a cut set, because any configuration with those units down has at least the down units of $\mathbf{s}$ and, by coherence, is failed. Conversely, if c is a cut set, the configuration with exactly the units of c down is failed by definition. The cut sets are therefore exactly the down sets of the failed configurations of the enumeration, and there are as many of them as failed configurations.

Minimal cut sets. A cut set is *minimal* when no proper subset of it is a cut set. Let $\mathcal{M}$ be the family of minimal cut sets. Every cut set contains a minimal one, found by removing units one at a time while the remainder is still a cut set. Then the top event holds exactly when all units of at least one minimal cut set are down: if $\phi^{\mathrm{D}}(\mathbf{x}) = 1$, the down set is a cut set and contains some $c \in \mathcal{M}$, all of whose units are down; if all units of some $c \in \mathcal{M}$ are down, $\phi^{\mathrm{D}}(\mathbf{x}) = 1$ by equation (14.15). In the indicators,

$$\phi^{\mathrm{D}}(\mathbf{x}) = 1 - \prod_{c \in \mathcal{M}} \Big(1 - \prod_{i \in c} x_i\Big), \tag{14.16}$$

an OR gate over AND gates, one AND gate per minimal cut set. This is the minimal-cut-set representation of a coherent structure, a standard result here derived from the enumeration. Every tree of a coherent structure, whatever its gates, is equivalent to this two-level tree, and its minimal cut sets are unique: they are determined by the function ϕ^{D}, not by the way the tree was drawn.

The enumeration. The definitions turn into an algorithm with three steps.

1. *Enumerate.* Visit the 2^n configurations and evaluate the tree, equation (14.13), on each.
2. *Filter.* Keep the failed configurations and record the down set of each; these are all the cut sets.
3. *Remove supersets.* Delete every recorded set that has a proper subset among the recorded sets; what remains is $\mathcal{M}$.

The cost grows as 2^n: thirty basic events give about 10^9 configurations and fifty give about 10^{15}, beyond any enumeration. Larger trees obtain the minimal cut sets from the tree's algebra instead, expanding the gates from the top down into products of basic events and absorbing every product that contains another, as the companion volume's chapter on cut sets describes; that route never visits a configuration. The enumeration has a property the algebra lacks: the same pass that finds the cut sets gives the exact top-event probability, equation (14.1), against which every approximation below is checked.

Building blocks. The gate functions of equation (14.13), the $2^7 = 128$ configurations of the seven units and the superset filter.

The minimal cut sets of the hall. Table 14.4 shows the enumeration of Section 14.3.1 for the hall. Each row is a size $|\mathcal{D}|$ of the down set, the number of units down; the columns give the number of configurations with that many units down, $\binom{7}{|\mathcal{D}|}$, the number of them that are failed, the number of minimal cut sets among their down sets, and the total probability of those failed configurations at the plug-in values, with its share of the plug-in U_{sys}. Of the 128 configurations, 113 are failed. One configuration with a single unit down is failed, the controller's, and its down

set {K} is a minimal cut set of order one. Thirteen of the twenty-one configurations with two units down are failed: the six with two CRAC units down, the one with both pumps down, and the six that pair the controller with another unit; the last six contain {K} and are not minimal, so the second-order minimal cut sets are the six CRAC pairs and $\{\mathrm{P1}, \mathrm{P2}\}$. Every configuration with three or more units down is failed, since at most one CRAC unit and one pump can be down without failing the hall, and every such down set contains a pair or {K}. The hall has eight minimal cut sets:

$$\mathcal{M} = \{\{\mathrm{K}\},\ \{\mathrm{C1}, \mathrm{C2}\},\ \{\mathrm{C1}, \mathrm{C3}\},\ \{\mathrm{C1}, \mathrm{C4}\},\ \{\mathrm{C2}, \mathrm{C3}\},\ \{\mathrm{C2}, \mathrm{C4}\},\ \{\mathrm{C3}, \mathrm{C4}\},\ \{\mathrm{P1}, \mathrm{P2}\}\}. \tag{14.17}$$

Table 14.4: The $2^7 = 128$ configurations of the hall grouped by the number $|\mathcal{D}|$ of units down: their count, the number that are failed and the number of minimal cut sets among their down sets. Probabilities at the plug-in values $q_C = 0.030$, $q_P = 0.010$, $q_K = 0.001$; their total is the plug-in $U_{\mathrm{sys}} = 6.281 \times 10^{-3}$.

Units down $\lvert\mathcal{D}\rvert$	Count	Failed	Minimal cut sets	Probability of the failed ones	Share of U_{sys}
0	1	0	0	0	0
1	7	1	1	8.68×10^{-4}	13.8 %
2	21	13	7	5.19×10^{-3}	82.6 %
3	35	35	0	2.21×10^{-4}	3.5 %
4	35	35	0	3.59×10^{-6}	0.06 %
5 to 7	29	29	0	3.0×10^{-8}	0.0 %
Total	128	113	8	6.281×10^{-3}	100 %

One row of the arithmetic, the configuration with C1 and C2 down and every other unit up: $q_C^2(1-q_C)^2(1-q_P)^2(1-q_K) = 0.0009 \times 0.9409 \times 0.9801 \times 0.999 = 8.29 \times 10^{-4}$. Six such configurations, one per CRAC pair, give 4.97×10^{-3}; the pump pair adds 8.8×10^{-5} and the six configurations pairing the controller with another unit 1.25×10^{-4}, for the 5.19×10^{-3} of the row $|\mathcal{D}| = 2$. Most of U_{sys} sits in configurations with exactly the units of one minimal cut set down, the probability of a configuration falling by a factor of about q_C with each further unit down. For seven units the algorithm evaluates the tree 128 times and makes at most $113 \times 112 = 12\,656$ subset tests.

14.3.2 The cut-set probabilities and the top event

Engineering questions. How likely is each combination, and which dominate? What is the probability of losing cooling, and how well does the rare-event sum do? The first asks for the probability of each minimal cut set with its band, and for their order; the second for the probability of the top event itself, and for the error of the sum of cut-set probabilities that the first reading of Section 14.1 offers in its place.

General form. Both are the estimation form of Chapter 2, equations (2.9) to (2.11), for a quantity $f(\mathbf{z})$ derived from the state: its mean and its central 95 % interval,

$$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}] = \int f(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \qquad \Pr(f(\mathbf{z}) \le f_L \mid \mathbf{y}) = 0.025, \quad \Pr(f(\mathbf{z}) \le f_U \mid \mathbf{y}) = 0.975.$$

The order of two cut sets is a threshold event of the form (2.12), with f the difference of their log probabilities and $f^* = 0$.

Specialized form. The three choices of Section 14.1.7 enter here. The two questions are equations (2.10) and (2.11), with (2.12) for the order, with $f(\mathbf{z})$ the probability of a cut set, the difference of two of their logarithms, and the top event, and $p(\mathbf{z} \mid \mathbf{y})$ the posterior of step 3, the Laplace normal $\mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$ of equation (14.10).

The cut sets. Let E_c be the event that every unit of c is down, whatever the other units do. Its probability is the sum of $\Pr(\mathbf{s} \mid \mathbf{q})$ over the configurations with $s_i = 0$ for $i \in c$. Each unit outside c contributes the factor $q_i + (1 - q_i) = 1$ to that sum, so

$$P_c(\mathbf{q}) = \Pr(E_c \mid \mathbf{q}) = \prod_{i \in c} q_i. \tag{14.18}$$

For the hall, a CRAC pair has $P_c = q_C^2$, the pump pair q_P^2 and the controller q_K. The question asks for $\mathbb{E}[P_c \mid \mathbf{y}]$. When the units of c have independent posteriors, the expectation of the product is the product of the means. When they share one unavailability, as the CRAC units do, $P_c = q^{|c|}$ and its expectation is the moment of order $|c|$ of the posterior. Each is an integral over one entry z_j of the state, against its normal, which the Gauss–Hermite sum of Chapter 13, equation (13.26), evaluates,

$$\mathbb{E}[q_j^h \mid \mathbf{y}] \approx \frac{1}{\sqrt{\pi}} \sum_{m=1}^{M} w_m \, q\big(\hat{z}_j + \sqrt{2\Sigma_{jj}}\, t_m\big)^h, \tag{14.19}$$

and the same sum with any function of q_j in place of the power; h is the order of the moment. The value beside it comes from the exact posterior. For $q \sim \mathrm{Beta}(a, b)$ Chapter 13 derived these moments, equation (13.29):

$$\mathbb{E}[q^{|c|} \mid \mathbf{y}] = \prod_{h=0}^{|c|-1} \frac{a + h}{a + b + h}. \tag{14.20}$$

For a pair, $|c| = 2$, the moment exceeds the square of the mean by the posterior variance, $\mathbb{E}[q^2] - \mathbb{E}[q]^2 = \mathrm{Var}(q \mid \mathbf{y}) \geq 0$, and the ratio of the two is

$$\frac{\mathbb{E}[q^2 \mid \mathbf{y}]}{\mathbb{E}[q \mid \mathbf{y}]^2} = \frac{(a + 1)(a + b)}{a\,(a + b + 1)}, \tag{14.21}$$

which tends to $(a + 1)/a$ as b grows: the gap is set by the number of units found down, not by the number of checks. The cause is the one Chapter 13 found for the whole hall, the Jensen gap of a shared posterior, now seen in one term: if q_C is high, every pair is likely together, and the average of q_C^2 is weighted toward those values. In the state the same coupling is a shared column: all six pairs depend on ξ_C alone. The gap grows with the order of the cut set, each further factor $(a + h)/(a + b + h)$ exceeding the mean.

A ranking computed at point values can therefore misstate even the order of the expectations, whenever the cut sets differ in order and the higher-order ones share a poorly known parameter. And an order of expectations is itself a claim about uncertain quantities: whether one cut set is more probable than another is the event $P_c > P_{c'}$, with a posterior probability of its own, to be reported beside the order.

> **Principle 14.1.** *Under a posterior shared by its members, the expected probability of a cut set of order $|c|$ is the moment of that order of the posterior, which exceeds the power of the mean by a factor that grows with $|c|$ and with the width of the posterior. Cut sets of different orders must be compared by their expectations, not by their plug-in values, and the comparison reported with the posterior probability that it holds.*

The top event. By equation (14.16) the top event is the union $\bigcup_{c \in \mathcal{M}} E_c$. Its indicator is $1 - \prod_c (1 - \mathbf{1}_{E_c})$; multiplying out the product over the $m = |\mathcal{M}|$ factors gives one term

for each nonempty subfamily $\mathcal{F} \subseteq \mathcal{M}$, with sign $(-1)^{|\mathcal{F}|+1}$, and a product of indicators of events of the form "all units of c down" is the indicator that all units of their union are down, $\prod_{c\in\mathcal{F}} \mathbf{1}_{E_c} = \mathbf{1}_{E_{\cup\mathcal{F}}}$. Taking the expectation given $\mathbf{q}$ with equation (14.18),

$$U_{\text{sys}}(\mathbf{q}) = \sum_{\nu=1}^{m} (-1)^{\nu+1} S_\nu, \qquad S_\nu = \sum_{\mathcal{F}\subseteq\mathcal{M},\ |\mathcal{F}|=\nu} \prod_{i\in\cup\mathcal{F}} q_i, \tag{14.22}$$

with ν the number of cut sets taken together, the inclusion–exclusion formula. The joint event of two cut sets is the event of their union: two CRAC pairs that share a unit, such as $\{\text{C1}, \text{C2}\}$ and $\{\text{C1}, \text{C3}\}$, occur together exactly when C1, C2 and C3 are down, with probability q_C^3, not q_C^4. Truncating the alternating sum after an odd number of terms bounds U_{sys} from above and after an even number from below, the Bonferroni inequalities, a standard result for any union of events:

$$S_1 - S_2 \ \le\ U_{\text{sys}} \ \le\ S_1, \qquad U_{\text{sys}} \ \le\ S_1 - S_2 + S_3. \tag{14.23}$$

The first truncation is the rare-event approximation,

$$U_{\text{sys}} \approx S_1 = \sum_{c\in\mathcal{M}} P_c(\mathbf{q}), \qquad 0 \le S_1 - U_{\text{sys}} \le S_2, \tag{14.24}$$

which counts each failed configuration once for every minimal cut set it contains, and whose error is at most the sum over pairs of cut sets of the probability that both occur. The second-order formula $S_1 - S_2$ subtracts those joint occurrences once. Chapter 13 used the same bounds for subsystems in series, whose down events involve disjoint sets of units, so that each joint probability is a product of two subsystem unavailabilities; cut sets of one tree can share units, and then a joint probability is the probability of the union of the two sets, which is larger.

The bounds hold at every $\mathbf{q}$, and so they hold in expectation: $\mathbb{E}[S_1 - S_2 \mid \mathbf{y}] \le \mathbb{E}[U_{\text{sys}} \mid \mathbf{y}] \le \mathbb{E}[S_1 \mid \mathbf{y}]$. Evaluated at the posterior means $\bar{\mathbf{q}}$, the sum $S_1(\bar{\mathbf{q}})$ bounds only the plug-in top event $U_{\text{sys}}(\bar{\mathbf{q}})$, and the expected top event exceeds the plug-in one by the moment gap of equation (14.21) summed over the cut sets that share a posterior. Whether the plug-in sum still covers the expectation is a contest between its overcount of the joint terms and that gap, which the numbers decide.

For the hall the eight minimal cut sets form 28 pairs, of five kinds by the units their unions contain. Twelve pairs of CRAC pairs share one unit, each with a union of three CRAC units and probability q_C^3; three pairs of CRAC pairs are disjoint, q_C^4; six pair a CRAC pair with the controller, $q_C^2 q_K$; six pair a CRAC pair with the pumps, $q_C^2 q_P^2$; one pairs the controller with the pumps, $q_K q_P^2$. Hence

$$S_1 = 6q_C^2 + q_P^2 + q_K, \qquad S_2 = 12q_C^3 + 3q_C^4 + 6q_C^2 q_K + 6q_C^2 q_P^2 + q_K q_P^2. \tag{14.25}$$

The second-order value $S_1 - S_2$ lies at or below the exact value for every tree; that is the Bonferroni bound (14.23), and it does not depend on how the cut sets overlap. What the overlap sets is the size of the gap, and the CRAC part of the tree shows how. The probability that two or more of the four CRAC units are down is, from the binomial tail of Chapter 13,

$$1 - (1-q)^4 - 4q(1-q)^3 = 6q^2 - 8q^3 + 3q^4,$$

expanding $(1-q)^3(1+3q) = 1 - 6q^2 + 8q^3 - 3q^4$. The same probability from inclusion–exclusion over the six pairs has $S_1 = 6q^2$ and $S_2 = 12q^3 + 3q^4$, as above. In S_3, the $\binom{6}{3} = 20$ triples of pairs include four whose union is three units (the three pairs within one set of three) and sixteen whose union is all four, so $S_3 = 4q^3 + 16q^4$; every family of four or more pairs covers all four units, so $S_4 = 15q^4$, $S_5 = 6q^4$ and $S_6 = q^4$. The alternating sum is

$$6q^2 - (12q^3 + 3q^4) + (4q^3 + 16q^4) - 15q^4 + 6q^4 - q^4 = 6q^2 - 8q^3 + 3q^4,$$

the binomial tail, as it must be. The second-order truncation keeps $-12q^3$ where the exact coefficient is $-8q^3$: it subtracts the probability of three CRAC units down three times, once for each pair of pairs within the three, after the rare-event sum had added it three times, and only S_3 adds it back once. Because the pairs share units, the joint terms are of third order in q, and the gap of the second order is the third-order term $4q^3$. Had the cut sets been disjoint, the joint terms, and the gap, would have been of fourth order.

The expectation. The three groups of units enter the hall's top event through independent factors, equation (14.13), and the posteriors of the three shared unavailabilities are independent. Writing $R_{\mathrm{C}}(q_C) = (1-q_C)^3(1+3q_C)$ for the probability that at most one CRAC unit is down,

$$U_{\mathrm{sys}}(\mathbf{q}) = 1 - R_{\mathrm{C}}(q_C)\,(1-q_P^2)\,(1-q_K), \tag{14.26}$$

and the expectation of a product of independent factors is the product of their expectations. The entries of the Laplace normal are independent, $\boldsymbol{\Sigma}$ being diagonal, so on the state route the expectation is three one-entry sums (14.19),

$$\mathbb{E}[U_{\mathrm{sys}} \mid \mathbf{y}] \approx 1 - \mathbb{E}[R_{\mathrm{C}}(q_C)]\,(1-\mathbb{E}[q_P^2])\,(1-\mathbb{E}[q_K]), \tag{14.27}$$

each expectation over the normal of its entry. The value beside it comes from the exact posterior, with the moments of equation (14.20):

$$\mathbb{E}[U_{\mathrm{sys}} \mid \mathbf{y}] = 1 - (1 - 6\mathbb{E}[q_C^2] + 8\mathbb{E}[q_C^3] - 3\mathbb{E}[q_C^4])\,(1-\mathbb{E}[q_P^2])\,(1-\mathbb{E}[q_K]), \tag{14.28}$$

No draw is needed for the mean on either route. The same sums and the same moments give the expectation of every truncation of equation (14.22), term by term, since each term is a product of powers of the three unavailabilities.

The bands and the order. Under the Laplace posterior each quantity has the band of the delta method of Chapter 4: a derived quantity $f(\mathbf{z})$ with gradient $\mathbf{g}$ at $\hat{\mathbf{z}}$ has $\mathrm{Var}(f \mid \mathbf{y}) \approx \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$. A small probability takes the delta method on its logarithm, the rule of Chapter 13, whose gradient is $\mathbf{g}/f$:

$$\mathrm{Var}(\ln f \mid \mathbf{y}) \approx \frac{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}}{f(\hat{\mathbf{z}})^2}, \qquad f \in f(\hat{\mathbf{z}})\,\exp\Big(\pm 1.96\,\frac{\sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}}}{f(\hat{\mathbf{z}})}\Big). \tag{14.29}$$

With $\boldsymbol{\Sigma}$ diagonal the product is a variance budget, one term $g_j^2\Sigma_{jj}$ per entry of the state.

Two incidence matrices carry the drawing. The 8×7 matrix $\mathbf{C}$ has $C_{ci} = 1$ when unit i belongs to cut set c; the 7×3 matrix $\mathbf{A}$ has $A_{ij} = 1$ when unit i takes its unavailability from entry j of the state, so that the units' unavailabilities are $\mathbf{q} = \mathbf{A}\,(q_C, q_P, q_K)^{\mathsf{T}}$. Both are read off the plant graph of Figure 14.1: $\mathbf{A}$ from the frames that place several units under one free entry, $\mathbf{C}$ from the redundant paths of the graph, a cut set being a set of units whose loss leaves the remaining paths unable to carry the water or the cooled air the hall needs, or removes the node that sequences them. With the units in the order C1 to C4, P1, P2, K, and $\mathbf{D} = \mathrm{diag}(q_j(1-q_j))$ the derivative of each shared unavailability with respect to its logit, every gradient with respect to the state passes through the same two factors,

$$\frac{\partial f}{\partial \mathbf{z}} = \frac{\partial f}{\partial \mathbf{q}}\,\mathbf{A}\,\mathbf{D}, \tag{14.30}$$

where $\mathbf{A}$ sums over the units that share an entry and $\mathbf{D}$ converts a change of an unavailability into a change of its logit. Stack the $m = 8$ cut-set probabilities into the vector $\mathbf{P}(\mathbf{z}) = (P_c)_{c\in\mathcal{M}}$. The matrix of their gradients,

$$\mathbf{G} = \frac{\partial \mathbf{P}}{\partial \mathbf{z}}, \qquad G_{cj} = \frac{\partial P_c}{\partial z_j}, \tag{14.31}$$

has a row per minimal cut set and a column per entry of the state, and G_{cj} is the change of P_c per unit change of the logit z_j. It maps a change of the state into changes of the cut sets, $\delta\mathbf{P} \approx \mathbf{G}\,\delta\mathbf{z}$, as $\mathbf{J}$ maps it into changes of the rows, so the Laplace covariance of the cut-set probabilities is $\mathbf{G}\boldsymbol{\Sigma}\mathbf{G}^{\mathsf{T}}$. By equation (14.18), $\partial P_c/\partial q_i = C_{ci}P_c/q_i$, and equation (14.30) gives

$$\mathbf{G} = \operatorname{diag}(\mathbf{P})\,(\mathbf{CA})\,\operatorname{diag}(1-q_j), \qquad \mathbf{CA} = \begin{bmatrix} 2 & 0 & 0 \\ 2 & 0 & 0 \\ 2 & 0 & 0 \\ 2 & 0 & 0 \\ 2 & 0 & 0 \\ 2 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 2 & 0 \end{bmatrix}, \tag{14.32}$$

where the integer matrix $\mathbf{CA}$ counts how many units of each cut set each entry carries, rows in the order of the six CRAC pairs, the controller and the pumps, and columns ξ_C, ξ_P, ξ_K. Its one nonzero entry per row is the drawing's: each minimal cut set draws its units from one group. Divided by P_c, a row of $\mathbf{G}$ is the gradient of $\ln P_c$, the row of $\mathbf{CA}$ times $1-q_j$. Since $\boldsymbol{\Sigma}$ is diagonal, cut sets whose rows have their nonzero entries in different columns are uncorrelated, and the six CRAC rows, which share the column of ξ_C, form a block of equal entries with every correlation equal to one: the shared posterior that made $\mathbb{E}[q_C^2]$ exceed $\mathbb{E}[q_C]^2$ appears in the state as a shared column, and the six pairs rise and fall together.

Cut set c is more probable than c' when $d = \ln P_c - \ln P_{c'}$ is positive, a derived quantity whose gradient $\mathbf{g}_d$ is the difference of the two rows of logarithms, so that under the Laplace posterior

$$\Pr(P_c > P_{c'} \mid \mathbf{y}) \approx \Phi\Big(\frac{d(\hat{\mathbf{z}})}{\sqrt{\mathbf{g}_d^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_d}}\Big), \qquad \mathbf{g}_d = \mathbf{G}^{\mathsf{T}}\Big(\frac{\mathbf{e}_c}{P_c} - \frac{\mathbf{e}_{c'}}{P_{c'}}\Big), \tag{14.33}$$

with $\mathbf{e}_c$ the unit vector of row c. The two terms of $\mathbf{g}_d^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_d$ say which record leaves the order uncertain.

The values beside them come from the exact posterior. Each P_c is an increasing function of one Beta variable, so its quantiles are the Beta quantiles raised to the order of the cut set, the monotone-map rule of Chapter 2; that is the exact band of a cut set. The top event, its truncations and the order of two cut sets from different groups depend on several entries, and their quantiles and probabilities are read from draws of the three Beta posteriors.

Building blocks. The MAP state $\hat{\mathbf{z}}$ and the covariance $\boldsymbol{\Sigma}$ of equation (14.10); the sum (14.19); the matrix $\mathbf{G}$ at $\hat{\mathbf{z}}$; the enumeration of Section 14.3.1 for the plug-in top event; and, for the exact values beside them, the Beta posteriors of the three groups with their moments, equation (14.20), and quantile functions, and draws of the three posteriors.

The cut sets ranked. Table 14.5 evaluates equation (14.19) over the entry ξ_C with five nodes. Each row is a node: the first two columns give the node t_m and its weight $w_m/\sqrt{\pi}$, the third the state, the fourth the unavailability there, and the last three the functions the answers average: the pair's q_C^2, the top-level CRAC factor $1 - R_{\mathrm{C}}$ of the present design, and the same with a fifth unit, for the design question of this step. The middle node is the MAP entry.

For a CRAC pair the sum of the fifth column, unchanged to the digits shown at forty nodes, gives

$$\mathbb{E}[q_C^2 \mid \mathbf{y}] \approx 0.01126 \times 3.32 \times 10^{-5} + 0.22208 \times 1.90 \times 10^{-4} + \cdots \\ + 0.01126 \times 2.01 \times 10^{-2} = 1.660 \times 10^{-3},$$

Table 14.5: The Gauss–Hermite rule (14.19) with five nodes over the Laplace posterior $\xi_C \mid \mathbf{y} \approx \mathcal{N}(-3.4761,\ 0.5862^2)$ of the CRAC units.

Node t_m	Weight $w_m/\sqrt{\pi}$	State $\xi_{C,m}$	q_C	q_C^2	$1 - R_C$	Fifth unit
-2.0202	0.01126	-5.1509	0.00576	3.32×10^{-5}	1.976×10^{-4}	1.90×10^{-6}
-0.9586	0.22208	-4.2708	0.01378	1.90×10^{-4}	1.118×10^{-3}	2.56×10^{-5}
0	0.53333	-3.4761	0.03000	9.00×10^{-4}	5.186×10^{-3}	2.58×10^{-4}
0.9586	0.22208	-2.6814	0.06408	4.11×10^{-3}	2.258×10^{-2}	2.38×10^{-3}
2.0202	0.01126	-1.8013	0.14169	2.01×10^{-2}	9.891×10^{-2}	2.27×10^{-2}
Weighted sum				1.660×10^{-3}	9.145×10^{-3}	9.29×10^{-4}

1.84 times the plug-in 0.03^2. A triple of CRAC units has $\mathbb{E}[q_C^3] \approx 1.054 \times 10^{-4}$, 3.90 times 0.03^3. With forty nodes the pumps have $\mathbb{E}[q_P^2] \approx 1.629 \times 10^{-4}$, 1.63 times 1.00×10^{-4}, and the controller $\mathbb{E}[q_K] \approx 1.644 \times 10^{-3}$, 1.64 times 1.00×10^{-3}, although its cut set has order one: its normal in the logit, from one unit found down, has the widest upper tail of the three. At the plug-in values a CRAC pair (9.0×10^{-4}) ranks below the controller (1.0×10^{-3}); in expectation it ranks above, narrowly (1.660×10^{-3} against 1.644×10^{-3}).

The value beside it comes from the exact posterior. For a CRAC pair, $|c| = 2$ and $\mathrm{Beta}(3, 97)$, equation (14.20) gives

$$\mathbb{E}[q_C^2 \mid \mathbf{y}] = \frac{3 \times 4}{100 \times 101} = 1.188 \times 10^{-3}, \qquad \text{against} \qquad \mathbb{E}[q_C \mid \mathbf{y}]^2 = 0.03^2 = 9.00 \times 10^{-4}, \tag{14.34}$$

a difference of $\mathrm{Var}(q_C \mid \mathbf{y}) = 2.88 \times 10^{-4}$ and, by equation (14.21), the ratio $4 \times 100/(3 \times 101) = 1.320$: the plug-in understates each pair by a quarter of its value. The third moment, which a triple of CRAC units would need, is $(3 \times 4 \times 5)/(100 \times 101 \times 102) = 5.82 \times 10^{-5}$ against $0.03^3 = 2.70 \times 10^{-5}$, a ratio of 2.16. For the pumps, $\mathbb{E}[q_P^2] = (4 \times 5)/(400 \times 401) = 1.247 \times 10^{-4}$ against 1.00×10^{-4}, a ratio of 1.25; the controller's cut set has order one and no exact gap, $\mathbb{E}[q_K] = 1.00 \times 10^{-3}$, so that exactly a pair ranks above it by more (1.19×10^{-3}).

Table 14.6 lists the minimal cut sets. Each row is a cut set or a group; the columns give its order, its probability given $\mathbf{q}$, its expected probability, its plug-in probability, its share of the expected rare-event sum and the central 95 % band of its probability, each on the state route with the exact value in parentheses. One row of the arithmetic, a CRAC pair: $1.660 \times 10^{-3}/1.177 \times 10^{-2} = 14.1\,\%$, and its band by equation (14.29), with the gradient of $\ln P_c$ below, $9.0 \times 10^{-4}\,\mathrm{e}^{\pm 1.96 \times 1.137} = [9.7 \times 10^{-5},\ 8.4 \times 10^{-3}]$; exactly, $1.188 \times 10^{-3}/8.253 \times 10^{-3} = 14.4\,\%$, and the Beta quantiles squared, $[0.0063^2,\ 0.0711^2] = [4.0 \times 10^{-5},\ 5.1 \times 10^{-3}]$.

Figure 14.3 draws the same ranking. The horizontal axis lists the eight minimal cut sets, the six CRAC pairs first, then the controller and the pumps; the vertical axis is the probability of the cut set on a logarithmic scale. Each bar is the state-route expected probability, green for the CRAC pairs, blue for the controller and gold for the pumps; the vertical line through each bar is its 95 % band on the state route, the dashed line beside it the exact band, the circle the exact expected probability, and the cross the plug-in value.

The six CRAC pairs are equal, because the four units share one unavailability, and each bar sits just above the controller's; together they carry 85 % of the sum (86 % exactly). The crosses of the CRAC pairs lie below the controller's, and the bands of a pair and of the controller overlap almost entirely: by equation (14.33), with the numbers worked below, $\Pr(q_C^2 > q_K \mid \mathbf{y}) \approx \Phi(-0.070) = 0.47$, and the six pairs together exceed the controller with probability $\Phi(1.114) = 0.87$. Over the draws of the exact posteriors the two are 0.52 and 0.87: the single order falls on opposite sides of one half on the two routes, and both call it even. Every band spans two orders of magnitude, more than the differences between the bars. The pumps' cut set lies an order of magnitude below the others and contributes 1.4 % (1.5 %).

Table 14.6: The minimal cut sets of the loss of hall cooling: state route, the sum (14.19) and the band (14.29), with the exact value, equation (14.20) and the Beta quantiles, in parentheses. Plug-in at $\bar{q}_C = 0.030$, $\bar{q}_P = 0.010$, $\bar{q}_K = 0.001$; shares of the expected sum $S_1 = 1.177 \times 10^{-2}$ (8.253×10^{-3}).

Cut set	Order	$P_c(\mathbf{q})$	Expected	Plug-in	Share	95 % band
each CRAC pair, $\{Ci, Cj\}$	2	q_C^2	1.660×10^{-3} (1.188×10^{-3})	9.00×10^{-4}	14.1 (14.4) %	$[9.7 \times 10^{-5},\ 8.4 \times 10^{-3}]$ ($[4.0 \times 10^{-5},\ 5.1 \times 10^{-3}]$)
the six CRAC pairs	2	$6q_C^2$	9.96×10^{-3} (7.13×10^{-3})	5.40×10^{-3}	84.6 (86.4) %	$[5.8 \times 10^{-4},\ 5.0 \times 10^{-2}]$ ($[2.4 \times 10^{-4},\ 3.0 \times 10^{-2}]$)
controller, {K}	1	q_K	1.644×10^{-3} (1.000×10^{-3})	1.00×10^{-3}	14.0 (12.1) %	$[1.4 \times 10^{-4},\ 7.1 \times 10^{-3}]$ ($[2.5 \times 10^{-5},\ 3.7 \times 10^{-3}]$)
pumps, {P1, P2}	2	q_P^2	1.629×10^{-4} (1.247×10^{-4})	1.00×10^{-4}	1.4 (1.5) %	$[1.4 \times 10^{-5},\ 7.0 \times 10^{-4}]$ ($[7.5 \times 10^{-6},\ 4.8 \times 10^{-4}]$)
sum S_1			1.177×10^{-2} (8.253×10^{-3})	6.50×10^{-3}	100 %	

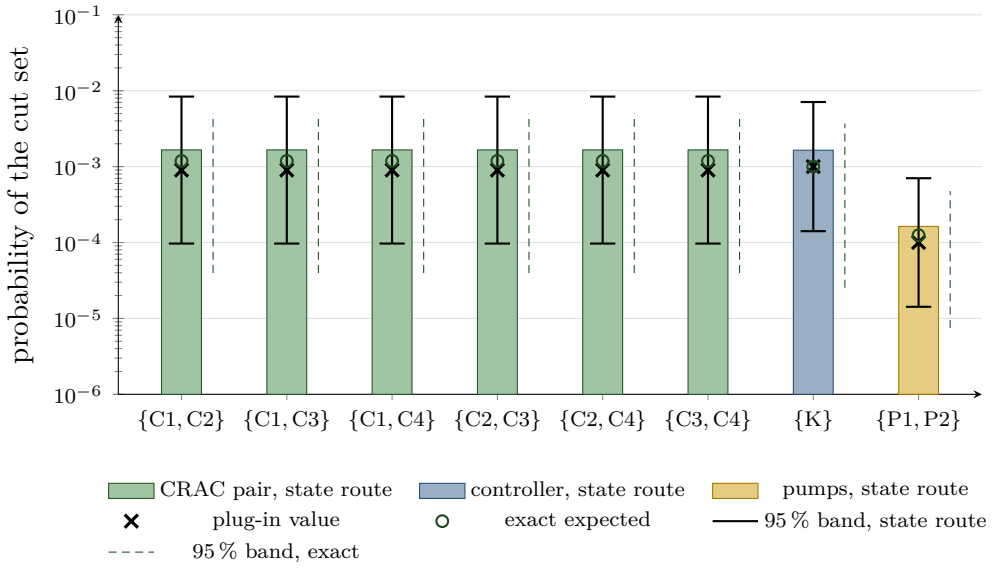


Figure 14.3: The eight minimal cut sets of the loss of hall cooling: state-route expected probability $\mathbb{E}[P_c \mid \mathbf{y}]$ (bars), by the sum (14.19), with its central 95 % band (solid lines), equation (14.29); beside them the exact expected probability (circles) and band (dashed, the Beta quantiles raised to the order of the cut set), with $q_C \sim \text{Beta}(3, 97)$, $q_P \sim \text{Beta}(4, 396)$, $q_K \sim \text{Beta}(1, 999)$; plug-in value at the posterior means (crosses).

At $\hat{\mathbf{z}}$ the matrix of equation (14.32) is

$$\mathbf{G}(\hat{\mathbf{z}}) = \begin{bmatrix} 2(1-q_C)q_C^2 & 0 & 0 \\ 2(1-q_C)q_C^2 & 0 & 0 \\ 2(1-q_C)q_C^2 & 0 & 0 \\ 2(1-q_C)q_C^2 & 0 & 0 \\ 2(1-q_C)q_C^2 & 0 & 0 \\ 2(1-q_C)q_C^2 & 0 & 0 \\ 0 & 0 & (1-q_K)q_K \\ 0 & 2(1-q_P)q_P^2 & 0 \end{bmatrix} = 10^{-3} \begin{bmatrix} 1.746 & 0 & 0 \\ 1.746 & 0 & 0 \\ 1.746 & 0 & 0 \\ 1.746 & 0 & 0 \\ 1.746 & 0 & 0 \\ 1.746 & 0 & 0 \\ 0 & 0 & 0.999 \\ 0 & 0.198 & 0 \end{bmatrix}, \quad (14.35)$$

whose rows divided by P_c, the gradients of $\ln P_c$, are $(1.94,\ 0,\ 0)$ for a CRAC pair, $(0,\ 0,\ 0.999)$ for the controller and $(0,\ 1.98,\ 0)$ for the pumps. For a pair against the controller the gradient of equation (14.33) is $\mathbf{g}_d = (1.94,\ 0,\ -0.999)^{\mathsf{T}}$ and $\mathbf{g}_d^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_d = 1.293 + 0.999 = 2.292$, 56 % from ξ_C and 44 % from ξ_K: the order is uncertain because of both records, and neither alone could settle it. With $d(\hat{\mathbf{z}}) = \ln(9.0 \times 10^{-4}/1.0 \times 10^{-3}) = -0.105$, the standard deviation 1.514 gives the $\Phi(-0.070)$ above, and with $d(\hat{\mathbf{z}}) = \ln 5.4 = 1.686$ for the six pairs together, $\Phi(1.114)$.

The top event and the approximations. The sum of the sixth column of Table 14.5, with the forty-node values $\mathbb{E}[1 - R_{\mathrm{C}}] \approx 9.145 \times 10^{-3}$, $\mathbb{E}[q_P^2] \approx 1.629 \times 10^{-4}$ and $\mathbb{E}[q_K] \approx 1.644 \times 10^{-3}$, gives by equation (14.27)

$$\mathbb{E}[U_{\mathrm{sys}} \mid \mathbf{y}] \approx 1 - 0.990855 \times 0.999837 \times 0.998356 = 1.09 \times 10^{-2},$$

against the plug-in value $U_{\mathrm{sys}}(\bar{\mathbf{q}}) = 6.28 \times 10^{-3}$. The ratio is 1.74. In hours per year of 8760 h the expectation is 96 h and the plug-in 55 h. The value beside it comes from the exact posterior. With the moments $\mathbb{E}[q_C^2] = 1.188 \times 10^{-3}$, $\mathbb{E}[q_C^3] = 5.824 \times 10^{-5}$, $\mathbb{E}[q_C^4] = 3.393 \times 10^{-6}$, $\mathbb{E}[q_P^2] = 1.247 \times 10^{-4}$ and $\mathbb{E}[q_K] = 1.000 \times 10^{-3}$, equation (14.28) gives $\mathbb{E}[R_{\mathrm{C}}] = 1 - 7.129 \times 10^{-3} + 4.659 \times 10^{-4} - 1.018 \times 10^{-5} = 0.993327$ and

$$\mathbb{E}[U_{\mathrm{sys}} \mid \mathbf{y}] = 1 - 0.993327 \times 0.999875 \times 0.999 = 7.79 \times 10^{-3}, \quad (14.36)$$

3.1×10^{-3} below the state route. The ratio to the plug-in value, 1.24, is the shared posterior's gap for the whole tree, set almost entirely by the CRAC pairs; the state route adds the upper tail of the normal in each logit. In hours per year the exact expectation is 68 h.

Table 14.7 compares the methods. Each row is a method; the columns give its value at the plug-in values of $\mathbf{q}$ with its error against the exact plug-in value, its expectation under the Laplace normal, the sum (14.19) term by term, with its error against the state route's top event, and its expectation under the Beta posteriors with its error against the exact expectation (14.36). The state route's expected S_2 is 1.31×10^{-3}, of which the twelve overlapping CRAC pairs carry $12 \times 1.054 \times 10^{-4} = 1.26 \times 10^{-3}$, 97 %; its expected S_3 is 5.64×10^{-4}. The exact expected S_2 is 7.17×10^{-4}, of which the twelve overlapping CRAC pairs carry $12 \times 5.824 \times 10^{-5} = 6.99 \times 10^{-4}$, 97 %; the three disjoint CRAC pairs add 1.0×10^{-5}, the pairs with the controller 7.1×10^{-6}, the pairs with the pumps 8.9×10^{-7} and the controller with the pumps 1.2×10^{-7}. The exact expected S_3 is 2.88×10^{-4}.

On the state route the rare-event sum is 7.6 % above the expected top event, the second-order value 4.3 % below and the third order 0.8 % above, so the two bounds bracket it, $1.05 \times 10^{-2} \le 1.09 \times 10^{-2} \le 1.18 \times 10^{-2}$. Exactly the same three errors are 5.9 %, 3.3 % and 0.4 %, with the bracket $7.54 \times 10^{-3} \le 7.79 \times 10^{-3} \le 8.25 \times 10^{-3}$. At $q_C = 0.03$ the second-order gap of the CRAC part, $4q_C^3$, is 1.1×10^{-4} against its exact 5.19×10^{-3}. The sum at the posterior means, 6.50×10^{-3}, bounds the plug-in top event 6.28×10^{-3} and lies 41 % below the state-route expectation and 17 % below the exact expectation: it bounds neither.

Table 14.7: The probability of losing hall cooling by method. Plug-in: $\mathbf{q}$ at the posterior means; expected on the state route: the sum (14.19) over the Laplace normal, by equation (14.27) and term by term for the truncations; expected exactly: under the three Beta posteriors, by equation (14.28) and its term-by-term analog.

Method	Plug-in	Error	Laplace expected	Error	Exact expected	Error
Exact, enumeration (14.1)	6.281×10^{-3}	–	1.093×10^{-2}	–	7.790×10^{-3}	–
Rare-event sum S_1	6.500×10^{-3}	+3.5 %	1.177×10^{-2}	+7.6 %	8.253×10^{-3}	+5.9 %
Second order $S_1 - S_2$	6.168×10^{-3}	−1.8 %	1.046×10^{-2}	−4.3 %	7.536×10^{-3}	−3.3 %
Third order $S_1 - S_2 + S_3$	6.289×10^{-3}	+0.1 %	1.102×10^{-2}	+0.8 %	7.824×10^{-3}	+0.4 %

The mean is not enough for an unavailability that is known only to within a factor of several, and Figure 14.4 draws its posterior. The horizontal axis is the probability of the top event on a logarithmic scale; the vertical axis is the posterior density per decade, the density of $\log_{10} U_{\text{sys}}$. The dotted blue curve is the Laplace density, the normal of the delta band (14.29), and the green curve the exact density from the draws, whose shaded area holds the exact central 95 %; the vertical lines mark, from left to right, the exact median, the plug-in value, which is the Laplace median, the exact mean, the exact expected rare-event sum and the Laplace mean.

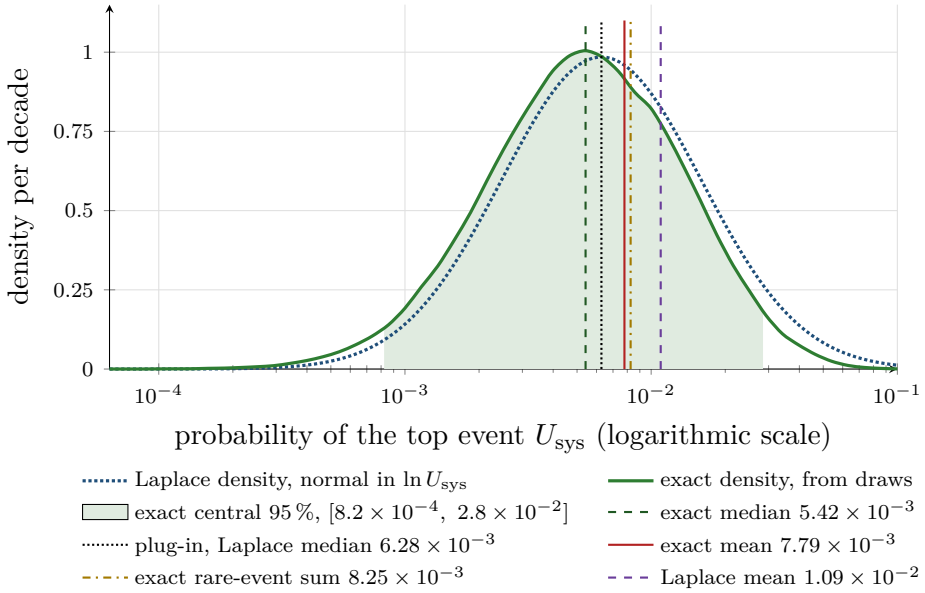


Figure 14.4: Posterior of the probability of losing hall cooling, $U_{\text{sys}} = 1 - (1 - q_C)^3(1 + 3q_C)(1 - q_P^2)(1 - q_K)$. Dotted blue: the Laplace density, $\ln U_{\text{sys}} \approx \mathcal{N}(\ln(6.281 \times 10^{-3}),\ 0.933^2)$, which the answers report; green: the exact density, with $q_C \sim \text{Beta}(3, 97)$, $q_P \sim \text{Beta}(4, 396)$, $q_K \sim \text{Beta}(1, 999)$: kernel density of $\log_{10} U_{\text{sys}}$ over 10^5 of the 4×10^5 draws (seed 14). Quantiles and the mean from all draws and from equation (14.28).

With the gradient $(1.568,\ 0.031,\ 0.158)$ of $\ln U_{\text{sys}}$ at $\hat{\mathbf{z}}$, worked out with the last question of this step, the logarithm has the standard deviation 0.933, and equation (14.29) gives the median 6.28×10^{-3} and the central 95 % band $6.28 \times 10^{-3}\, e^{\pm 1.829} = [1.0 \times 10^{-3},\ 3.9 \times 10^{-2}]$, a factor of 39. In hours per year the median is 55 h, the band $[9,\ 342]$ h and the mean 96 h. The exact density is close to symmetric on the logarithmic scale, with the median 5.42×10^{-3} and the central 95 % band $[8.2 \times 10^{-4},\ 2.8 \times 10^{-2}]$, a factor of 35; its four marked values lie within a fifth of a decade of one another. In hours per year the exact median is 47 h, the band $[7,\ 249]$ h

and the mean 68 h. Against the exact band $[2.4 \times 10^{-4},\ 2.8 \times 10^{-2}]$ of the CRAC units alone in Chapter 13, the pumps and the controller raise the lower end by a factor of about three and a half and leave the upper end; on the state route, against $[5.8 \times 10^{-4},\ 4.6 \times 10^{-2}]$, they raise the lower end by a factor of 1.7 and lower the upper end, because the gradient of $\ln U_{\text{sys}}$ in ξ_C falls from 1.90 to 1.57 when the other terms share the sum.

14.3.3 The share of the CRAC pairs

Engineering question. How much of the top event's probability is in the CRAC pairs, and is that conclusion firm? The owner wants to know where the loss of cooling comes from, which directs maintenance attention and the study of failure causes.

General form. The threshold form of Chapter 2, equation (2.12): the posterior probability of the event $f(\mathbf{z}) \geq f^*$, here with f the share of a group of cut sets and $f^* = \frac{1}{2}$. The probability itself is the answer to "is that conclusion firm"; no assurance level is set by the question.

Specialized form. One choice enters here: the share is defined as the group's part of the rare-event sum rather than of the exact top event. This makes the share a ratio of cut-set probabilities, whose threshold event is a tail of one entry, and it would mislead if the joint terms of the CRAC pairs were large; Section 14.4.2 repeats the question with the exact top event in the denominator. The share of a group $\mathcal{B} \subseteq \mathcal{M}$ of cut sets is

$$\pi_{\mathcal{B}}(\mathbf{q}) = \frac{\sum_{c \in \mathcal{B}} P_c(\mathbf{q})}{S_1(\mathbf{q})}, \qquad \pi_{\text{C}}(\mathbf{q}) = \frac{6q_C^2}{6q_C^2 + q_P^2 + q_K}, \tag{14.37}$$

the second for the six CRAC pairs. The question is equation (2.12) with $f(\mathbf{z}) = \pi_{\text{C}}$, $f^* = \frac{1}{2}$ and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal of step 3. The share is a derived quantity of the state with a posterior of its own, and the Laplace route carries it on its logit, with the gradient built from the rows of $\mathbf{G}$, equation (14.31),

$$\operatorname{logit} \pi_{\mathcal{B}} = \ln S_{\mathcal{B}} - \ln(S_1 - S_{\mathcal{B}}), \qquad \mathbf{g}_\pi = \frac{\mathbf{G}^{\mathsf{T}}\mathbf{1}_{\mathcal{B}}}{S_{\mathcal{B}}} - \frac{\mathbf{G}^{\mathsf{T}}(\mathbf{1} - \mathbf{1}_{\mathcal{B}})}{S_1 - S_{\mathcal{B}}}, \tag{14.38}$$

with $S_{\mathcal{B}} = \sum_{c \in \mathcal{B}} P_c$ and $\mathbf{1}_{\mathcal{B}}$ the indicator vector of the group's rows, and the threshold question becomes $\Pr(\pi_{\mathcal{B}} > \frac{1}{2} \mid \mathbf{y}) \approx \Phi\big(\operatorname{logit} \pi_{\mathcal{B}}(\hat{\mathbf{z}}) / \sqrt{\mathbf{g}_\pi^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g}_\pi}\big)$. The gradient of the share subtracts the column sums of the other cut sets from those of the group, so an entry that carries only the other cut sets enters with a negative sign, and its term $g_{\pi,j}^2 \Sigma_{jj}$ in the variance budget says how much of the share's uncertainty that entry's record leaves. The share of a basic event's cut sets in the top event is the Fussell–Vesely importance of that event, and π_{C} is that importance for the group of CRAC units.

The value beside it comes from the exact posterior. The threshold event $\pi_{\text{C}} > \frac{1}{2}$ is $6q_C^2 > q_P^2 + q_K$, or $q_C > q_C^\dagger$ with the threshold $q_C^\dagger = \sqrt{(q_K + q_P^2)/6}$:

$$\Pr(\pi_{\text{C}} > \tfrac{1}{2} \mid \mathbf{y}) = \mathbb{E}\Big[\, 1 - I_{q_C^\dagger}(a_C,\ b_C) \;\Big|\; \mathbf{y} \Big], \tag{14.39}$$

with $I_x(\cdot, \cdot)$ the regularized incomplete Beta function of Chapter 5 and $\text{Beta}(a_C, b_C)$ the CRAC posterior (Section 14.2.1): the CRAC posterior's tail is exact, and only the expectation over (q_P, q_K) remains, a two-dimensional integral evaluated by draws, from which the exact band of the share also comes.

Building blocks. $\mathbf{g}_\pi$ from the rows of $\mathbf{G}$ and $\boldsymbol{\Sigma}$, for the probability, the band and the budget; and, for the exact values beside them, the incomplete Beta function at $q_C^\dagger$ and draws of (q_P, q_K) for the probability and draws of all three for the band.

The share of the CRAC pairs in the hall. The gradient of the share's logit, equation (14.38) with $\mathcal{B}$ the six CRAC pairs, takes the CRAC rows of $\mathbf{G}$ over $S_{\mathcal{B}} = 5.4 \times 10^{-3}$ and subtracts the other two over 1.1×10^{-3}:

$$\mathbf{g}_\pi = \begin{bmatrix} 10.476/5.4 \\ 0 \\ 0 \end{bmatrix} - \begin{bmatrix} 0 \\ 0.198/1.1 \\ 0.999/1.1 \end{bmatrix} = \begin{bmatrix} 1.940 \\ -0.180 \\ -0.908 \end{bmatrix}, \qquad \mathbf{g}_\pi^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g}_\pi = 1.293 + 0.008 + 0.826 = 2.127.$$

The CRAC entry carries 61 % of the variance of the share's logit, the controller's 39 % and the pumps' 0.4 %. At $\hat{\mathbf{z}}$ the share is 0.831, its logit 1.591, so that $\Pr(\pi_{\mathrm{C}} > \tfrac{1}{2} \mid \mathbf{y}) \approx \Phi(1.591/1.458) = 0.86$, and the band on the logit, mapped back, is $[0.22,\ 0.99]$.

The values beside them come from the exact posterior. The ratio of the exact expected probabilities in Table 14.6 gives the CRAC pairs 86 % of the sum (85 % for the Laplace expectations), a point value of the share (14.37). Figure 14.5 draws the share's posterior. The horizontal axis is the share π_{C} of the six CRAC pairs in the rare-event sum, from zero to one; the vertical axis its posterior density. The dotted blue curve is the Laplace density, the green curve the exact one. The red vertical line marks one half, and the shaded area to its right is the exact probability that the CRAC pairs carry more than half of the sum; the dotted lines mark the ends of the exact central 95 % band.

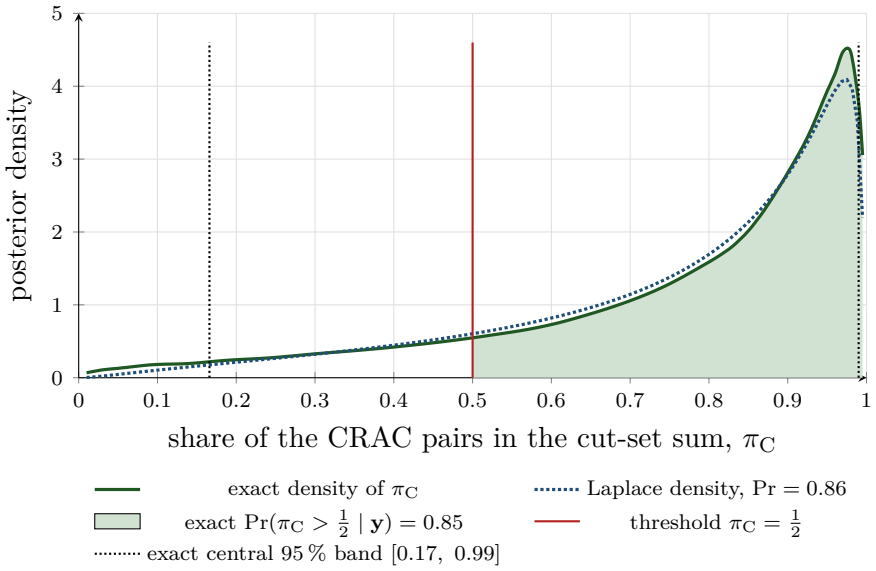


Figure 14.5: Posterior of the share of the six CRAC pairs in the rare-event sum, $\pi_{\mathrm{C}} = 6q_C^2/(6q_C^2 + q_P^2 + q_K)$. Dotted blue: the Laplace density, its logit normal with mean 1.591 and standard deviation 1.458, which the answers report. Green, under the three Beta posteriors: kernel density on the logit of π_{C} over 10^5 draws, mapped back to π_{C} and drawn to 0.995. The shaded area is $\Pr(\pi_{\mathrm{C}} > \tfrac{1}{2} \mid \mathbf{y}) = 0.85$.

The exact density peaks near 0.97 with a long, thin tail toward zero; its median is 0.84, its mean 0.76 and its central band $[0.17,\ 0.99]$, and by equation (14.39) $\Pr(\pi_{\mathrm{C}} > \tfrac{1}{2} \mid \mathbf{y}) = 0.85$, 0.01 below the state route. At the lower end of the band on q_C, 0.0063, the six pairs have $6 \times 0.0063^2 = 2.4 \times 10^{-4}$, a quarter of the controller's mean.

14.3.4 The two design changes

Engineering question. Which single change removes the most? Two changes are proposed, each removing one weakness of the tree.

- *A redundant controller.* A second control system takes over when the first fails, so that the basic event K becomes the AND gate K1 ∧ K2. The cut set {K} of order one is replaced by {K1, K2} of order two.
- *A fifth CRAC unit.* The hall still needs three running units, now out of five, a 3-out-of-5 arrangement, and loses cooling when three or more are down. The vote gate becomes 3/5, and the six CRAC pairs give way to the $\binom{5}{3} = 10$ CRAC triples.

General form. The design form of Chapter 2, equation (2.16), with the action a change of drawing d and u the expected top event under that drawing. A change to the plant brings no new evidence, so the posterior after the action is the present one and the average over outcomes collapses:

$$U(d) = \mathbb{E}\big[f_d(\mathbf{z}) \mid \mathbf{y}\big], \tag{14.40}$$

with f_d the top event of drawing d as a function of the state; the drawings are ranked smallest first. Which of two changes removes more is, draw by draw, the threshold event $f_0 - f_d > f_0 - f_{d'}$, with f_0 the present drawing.

Specialized form. Two choices enter here. The two controllers are of one design, share the posterior $\mathrm{Beta}(a_K, b_K)$ and fail independently given it: this values the redundant controller without a record of its own, through the second moment $\mathbb{E}[q_K^2]$, and it fails if a common cause can take both controllers down, which the next chapter models. The fifth CRAC unit shares the posterior of the other four: this keeps the state at three entries, and it fails if the new unit differs from the old ones.

The question is equation (14.40) with f_d the top event of each drawing and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal of step 3. Each change gives a new tree and a new enumeration, over $2^8 = 256$ configurations (2^9 for both changes together), and the algorithm of Section 14.3.1 returns the new minimal cut sets. The expectation is the sum (14.19) entry by entry, as in equation (14.27): for the redundant controller with $\mathbb{E}[q_K^2]$ in place of $\mathbb{E}[q_K]$; for the fifth CRAC unit, with the CRAC factor R_{C} replaced by the probability that at most two of five units are down, $\sum_{l=0}^{2} \binom{5}{l} q_C^l (1 - q_C)^{5-l}$. The band of each design is equation (14.29) with the gradient of $\ln f_d$ at $\hat{\mathbf{z}}$, and the order of two changes is the contrast of equation (14.33): since f_0 is common, the fifth unit removes more than the second controller exactly when the top event with the second controller exceeds that with the fifth unit, and

$$\Pr(\text{fifth unit removes more} \mid \mathbf{y}) \approx \Phi\Big(\frac{d(\hat{\mathbf{z}})}{\sqrt{\mathbf{g}_d^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g}_d}}\Big), \qquad d = \ln f_{\text{controller}} - \ln f_{\text{fifth}}. \tag{14.41}$$

The value beside it comes from the exact posterior: the same products with the moments (14.20), the CRAC factor of the fifth unit expanded in powers of q_C, and draws of the three posteriors for the medians, the bands and the draw-by-draw comparison.

Where each change acts is read in the columns of $\mathbf{G}$, equation (14.32). The redundant controller acts on the column of ξ_K, whose one nonzero entry is the row of {K}: it replaces that row's cut set and no other, so it can remove at most that cut set's probability, $\mathbb{E}[q_K]$. Because the top event is linear in q_K, equation (14.26), it removes, on either route, the expected coefficient of q_K times the drop of that moment,

$$\mathbb{E}[U_{\mathrm{sys}} \mid \mathbf{y}] - U(\text{redundant controller}) = \mathbb{E}[R_{\mathrm{C}}]\,(1 - \mathbb{E}[q_P^2])\,(\mathbb{E}[q_K] - \mathbb{E}[q_K^2]). \tag{14.42}$$

The fifth CRAC unit acts on the column of ξ_C, whose six nonzero rows are the CRAC pairs, and replaces each pair by triples of the third moment. A change can remove at most what the cut sets of its column carry.

Building blocks. The MAP state $\hat{\mathbf{z}}$ and the covariance $\boldsymbol{\Sigma}$; the sum (14.19) for R_{C}, the CRAC factor of five units, q_P^2, q_K and q_K^2; the gradients of $\ln f_d$ at $\hat{\mathbf{z}}$; the enumerations of the new trees; and, for the exact values beside them, the moments $\mathbb{E}[q_C^h]$ to $h = 5$, $\mathbb{E}[q_P^2]$, $\mathbb{E}[q_K]$ and $\mathbb{E}[q_K^2]$ from equation (14.20), and draws of the three posteriors for the medians, the bands and the draw-by-draw comparison.

The two design changes for the hall. Table 14.8 compares the designs. Each row is a design; the columns give its minimal cut sets, the expected rare-event sum, the expectation (14.40) of the top event, its median and central 95 % band, and the mean and median hours per year without cooling, each on the state route with the exact value in parentheses. On the state route, the redundant controller replaces $\mathbb{E}[q_K] \approx 1.644 \times 10^{-3}$ by $\mathbb{E}[q_K^2] \approx 7.2 \times 10^{-6}$, and the fifth CRAC unit replaces the six pairs by ten triples of 1.054×10^{-4} each. Exactly, the redundant controller replaces $\mathbb{E}[q_K] = 1.0 \times 10^{-3}$ by $\mathbb{E}[q_K^2] = (1 \times 2)/(1000 \times 1001) = 2.0 \times 10^{-6}$, twice the square of the mean; the fifth CRAC unit replaces the six pairs by ten triples of 5.82×10^{-5} each.

Table 14.8: The loss of hall cooling under the present design and three alternatives: state route, the sum (14.19) for the means and equation (14.29) for the medians and bands, with the exact value in parentheses, from the Beta moments and from 4×10^5 draws, seed 14; hours per year are 8760 h times the mean and the median.

Design	Minimal cut sets	Rare-event sum	Mean	Median	95 % band	Hours per year
present	{K}, six CRAC pairs, pumps	1.18×10^{-2} (8.25×10^{-3})	1.09×10^{-2} (7.79×10^{-3})	6.28×10^{-3} (5.42×10^{-3})	$[1.0 \times 10^{-3}, 3.9 \times 10^{-2}]$ ($[8.2 \times 10^{-4}, 2.8 \times 10^{-2}]$)	96; 55 (68; 47)
redundant controller	{K1, K2}, six CRAC pairs, pumps	1.01×10^{-2} (7.26×10^{-3})	9.31×10^{-3} (6.80×10^{-3})	5.29×10^{-3} (4.32×10^{-3})	$[6.2 \times 10^{-4}, 4.5 \times 10^{-2}]$ ($[3.5 \times 10^{-4}, 2.7 \times 10^{-2}]$)	82; 46 (60; 38)
fifth CRAC unit	{K}, ten CRAC triples, pumps	2.86×10^{-3} (1.71×10^{-3})	2.73×10^{-3} (1.66×10^{-3})	1.36×10^{-3} (1.27×10^{-3})	$[2.8 \times 10^{-4}, 6.6 \times 10^{-3}]$ ($[1.9 \times 10^{-4}, 5.3 \times 10^{-3}]$)	24; 12 (14.5; 11)
both changes	{K1, K2}, ten CRAC triples, pumps	1.22×10^{-3} (7.09×10^{-4})	1.10×10^{-3} (6.60×10^{-4})	3.59×10^{-4} (3.37×10^{-4})	$[3.2 \times 10^{-5}, 4.1 \times 10^{-3}]$ ($[3.9 \times 10^{-5}, 3.3 \times 10^{-3}]$)	9.6; 3.1 (5.8; 2.9)

On the state route the redundant controller lowers the expectation from 1.09×10^{-2} to 9.31×10^{-3}, by 15 %, the 1.62×10^{-3} of equation (14.42). The fifth CRAC unit lowers it to 2.73×10^{-3}, by 75 %; the CRAC units' part alone falls to 9.28×10^{-4}, the value of Chapter 13 for 3-out-of-5 cooling. Both changes together give 1.10×10^{-3}. On expected values the fifth unit removes 8.2×10^{-3}, five times the controller's 1.62×10^{-3}, about 72 h of lost cooling per year against 14 h. For the order, equation (14.41) has $d(\hat{\mathbf{z}}) = \ln(5.287 \times 10^{-3}/1.358 \times 10^{-3}) = 1.360$ and $\mathbf{g}_d = (1.320,\ -0.108,\ -0.735)^{\mathsf{T}}$, so $\mathbf{g}_d^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_d = 0.599 + 0.003 + 0.541 = 1.143$, and $\Phi(1.360/1.069) = 0.898$.

The values beside them come from the exact posterior. The redundant controller lowers the expectation from 7.79×10^{-3} to 6.80×10^{-3}, by 13 %, the 9.9×10^{-4} of equation (14.42). The fifth CRAC unit lowers it to 1.66×10^{-3}, by 79 %; the CRAC units' part alone falls to 5.3×10^{-4}. Both changes together give 6.6×10^{-4}. On expected values the fifth unit removes 6.1×10^{-3}, six times the controller's 9.9×10^{-4}, about 54 h of lost cooling per year against 9 h, and taken draw by draw it removes more with posterior probability 0.865. The order and the rough size of the two removals are the same on the two routes.

With the fifth CRAC unit, the design's own rare-event sum is led by the controller, with 57 % (59 % exactly), against 37 % (34 %) for the ten triples and 6 % (7 %) for the pumps. Which of the two changes should be bought depends on their prices and on the cost of an hour without cooling, the subject of a later chapter on redundancy and reliability investment.

14.3.5 The gradients of the top event

Engineering question. Which basic event's probability matters most to the top event? The owner wants to know where a change of a unit's unavailability would move the loss of cooling most.

General form. The estimation form of Chapter 2, equation (2.9), for the derived quantity $f(\mathbf{z}) = \partial U_{\text{sys}}/\partial q_i$, the slope of the top event in one unit's unavailability: a function of the state like any other, summarized by a location and a central band.

Specialized form. No new choice enters. The question is equation (2.9) with $f(\mathbf{z}) = \partial U_{\text{sys}}/\partial q_i$ and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal of step 3. Its value at $\hat{\mathbf{z}}$ is the reported location; its band is equation (14.29) on the logarithm of the slope, or on the logit of a slope near one; and carried to the logits it is the gradient $\mathbf{g} = \partial U_{\text{sys}}/\partial \mathbf{z}$ through which every Laplace band of the top event passes. The value beside the band comes from draws of the exact posterior. Chapter 13 showed that U_{sys} is multilinear in the units' unavailabilities, so that, holding the others fixed, it is a straight line in q_i whose slope is the difference between its values at the two ends of the line, the quantity that chapter wrote $I_{\text{B},i}$,

$$\frac{\partial U_{\text{sys}}}{\partial q_i} = U_{\text{sys}}(\mathbf{q} \mid q_i = 1) - U_{\text{sys}}(\mathbf{q} \mid q_i = 0), \tag{14.43}$$

the probability of losing cooling with unit i certainly down, less the same with unit i certainly up. The enumeration evaluates it for all seven units in one pass, the vector $\nabla_{\mathbf{q}} U_{\text{sys}}$, and equation (14.30) carries it to the state,

$$\mathbf{g} = \frac{\partial U_{\text{sys}}}{\partial \mathbf{z}} = \mathbf{D}\,\mathbf{A}^{\mathsf{T}}\,\nabla_{\mathbf{q}} U_{\text{sys}}. \tag{14.44}$$

For a CRAC unit the slope has a closed form: C1 enters only through the CRAC factor of equation (14.26), which fails with C1 down if any of the other three is down and with C1 up if two or more of them are, a difference of $3q_C(1-q_C)^2$, times the other two factors.

Differentiating the inclusion–exclusion formula (14.22) term by term instead, with $S_1 = \mathbf{1}^{\mathsf{T}}\mathbf{P}$,

$$\mathbf{g} = \mathbf{G}^{\mathsf{T}}\mathbf{1} - \frac{\partial S_2}{\partial \mathbf{z}} + \frac{\partial S_3}{\partial \mathbf{z}} - \cdots, \qquad \mathbf{g} \approx \mathbf{G}^{\mathsf{T}}\mathbf{1}, \tag{14.45}$$

to first order: each cut set adds its row, and the top event's gradient is the vector of column sums of $\mathbf{G}$. This is the rare-event approximation written as a statement about $\mathbf{G}$, and its error is the gradient of the joint occurrences, $\partial S_2/\partial \mathbf{z}$, which sits in the column whose cut sets overlap. Each truncation of equation (14.22) is likewise a function of the state with its own gradient, $\mathbf{G}^{\mathsf{T}}\mathbf{1}$ for S_1, $\mathbf{G}^{\mathsf{T}}\mathbf{1} - \partial S_2/\partial \mathbf{z}$ for $S_1 - S_2$, and so on, each placed in equation (14.29) for its Laplace band. The variance budget of the top event is the same vector,

$$\operatorname{Var}(\ln U_{\text{sys}} \mid \mathbf{y}) \approx \sum_j \Big(\frac{g_j}{U_{\text{sys}}}\Big)^2 \Sigma_{jj}, \tag{14.46}$$

so the entry that carries the top event and the entry that carries its band are read from one vector.

The two gradients answer different questions. An entry of $\nabla_{\mathbf{q}} U_{\text{sys}}$ is the change of the top event per unit change of an unavailability, a sensitivity: a unit that alone stops the service has a slope near one however rarely it fails. An entry of $\mathbf{g}$ is the change per unit change of a logit, that is per relative change of an unavailability, and to first order it is a column sum of $\mathbf{G}$, the probability the cut sets of that column carry: a contribution. A design change acts on the cut

sets of one column (Section 14.3.4), so the ranking for design follows $\mathbf{g}$ and the columns of $\mathbf{G}$. A later chapter on importance measures names these entries: those of $\nabla_{\mathbf{q}} U_{\text{sys}}$ as the Birnbaum importance, those of $\mathbf{g}/U_{\text{sys}}$, divided by $1 - q_j$, as the criticality importance, and the share of S_1 in the nonzero rows of a column of $\mathbf{G}$ as the Fussell–Vesely importance in its rare-event form.

Principle 14.2. *The gradient of the top event with respect to a basic event's unavailability measures sensitivity; the gradient with respect to the state's logits, to first order the column sums of the cut-set gradient matrix, measures contribution. A design change can remove at most what the cut sets of its column carry, so the ranking for design is by the state's gradient and the columns of* $\mathbf{G}$*, with the gradient in the unavailabilities converting a change in an unavailability into a change in the top event.*

Building blocks. The enumeration with one unit pinned down and then up, for $\nabla_{\mathbf{q}} U_{\text{sys}}$; the matrices $\mathbf{A}$ and $\mathbf{D}$ at $\hat{\mathbf{z}}$; $\mathbf{G}$ and its column sums; $\boldsymbol{\Sigma}$ for the bands; and, for the exact bands beside them, draws of the three posteriors.

The gradients of the top event for the hall. Table 14.9 lists the entries of the two gradients of Section 14.3.5. Each row is a basic event; the columns give its posterior mean unavailability, the top event with the unit certainly down and certainly up, the entry of $\nabla_{\mathbf{q}} U_{\text{sys}}$ at the plug-in values, the central 95 % band of that entry on the state route with the exact band in parentheses, and the share of the plug-in rare-event sum carried by the cut sets that contain the unit, its Fussell–Vesely importance in the rare-event form. One row of the arithmetic, a CRAC unit: $0.0883 - 0.0037 = 0.0846$.

Table 14.9: Gradients of the probability of losing cooling with respect to each basic event's unavailability, equation (14.43), at $q_C = 0.030$, $q_P = 0.010$, $q_K = 0.001$, with their central 95 % posterior bands from equation (14.29), on the logit for the controller, and in parentheses from 4×10^5 draws of the exact posterior; last column, the share of the plug-in sum $S_1 = 6.50 \times 10^{-3}$ in the cut sets containing the event.

Basic event	$\bar{q}_i$	U, unit down	U, unit up	$\partial U/\partial q_i$	95 % band	Share of S_1
each CRAC unit, Ci	0.030	0.0883	3.74×10^{-3}	0.0846	[0.030, 0.24] ([0.019, 0.18])	0.415
each pump, Pi	0.010	0.0161	6.18×10^{-3}	0.0099	[0.0037, 0.026] ([0.0027, 0.022])	0.015
controller, K	0.001	1	5.29×10^{-3}	0.995	[0.956, 0.9994] ([0.973, 0.9997])	0.154
shared q_C (four units)				0.338	[0.12, 0.96] ([0.075, 0.73])	0.831
shared q_P (two units)				0.0199	[0.0075, 0.053] ([0.0054, 0.043])	0.015

Summed over the units that share an entry and carried to the state by equation (14.44),

$$\mathbf{A}^{\mathsf{T}} \nabla_{\mathbf{q}} U_{\text{sys}} = \begin{bmatrix} 0.338 \\ 0.0199 \\ 0.995 \end{bmatrix}, \qquad \mathbf{g} = \begin{bmatrix} 0.0291 & 0 & 0 \\ 0 & 0.0099 & 0 \\ 0 & 0 & 0.000999 \end{bmatrix} \begin{bmatrix} 0.338 \\ 0.0199 \\ 0.995 \end{bmatrix} = 10^{-3} \begin{bmatrix} 9.846 \\ 0.197 \\ 0.994 \end{bmatrix},$$

and the column sums of $\mathbf{G}$ and the second-order gradient are

$$\mathbf{G}^{\mathsf{T}} \mathbf{1} = 10^{-3} \begin{bmatrix} 10.476 \\ 0.198 \\ 0.999 \end{bmatrix}, \qquad \frac{\partial (S_1 - S_2)}{\partial \mathbf{z}} = 10^{-3} \begin{bmatrix} 9.512 \\ 0.197 \\ 0.994 \end{bmatrix}.$$

The column sum of ξ_C is 6.4 % high and those of ξ_P and ξ_K are 0.6 % and 0.5 % high, and the second-order gradient is 3.4 % low in the entry of ξ_C, where the joint terms of S_2, almost all $12q_C^3$, sit.

In $\nabla_{\mathbf{q}} U_{\text{sys}}$ the controller leads by a factor of twelve, with the narrow band [0.956, 0.9994] ([0.973, 0.9997] exactly, and the larger entry in every draw); a CRAC unit's entry, 0.085, nearly $3q_C$, has a band of a factor of eight, [0.030, 0.24] (a factor of ten, [0.019, 0.18]). In $\mathbf{g}$ the CRAC entry leads by a factor of ten, 9.85×10^{-3} against 0.99×10^{-3} for the controller and 0.20×10^{-3} for the pumps, and the column of ξ_C carries 83 % of the plug-in sum and 85 % in expectation (86 % exactly). Divided by S_1, the CRAC column sum is $1.612 = 2 \times 0.97 \times 0.831$, the entry 2 of **CA** times $1 - q_C$ times the column's share 0.831.

The variance budget (14.46) of the band of Section 14.3.2 divides the gradient by the top event,

$$\frac{\mathbf{g}}{U_{\text{sys}}} = \frac{10^{-3}}{6.281 \times 10^{-3}} \begin{bmatrix} 9.846 \\ 0.197 \\ 0.994 \end{bmatrix} = \begin{bmatrix} 1.568 \\ 0.031 \\ 0.158 \end{bmatrix}, \qquad \frac{\mathbf{g}^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g}}{U_{\text{sys}}^2} = 0.845 + 0.0002 + 0.025 = 0.870,$$

the three terms $1.568^2 \times 0.3436$, $0.031^2 \times 0.2525$ and $0.158^2 \times 1.001$: the entry of ξ_C carries 97 % of the variance of $\ln U_{\text{sys}}$, the controller's 3 % and the pumps' 0.03 %. The controller's 3 % of that variance compares with its 39 % of the variance of the share's logit (Section 14.3.3): in the top event's gradient ξ_K enters beside the CRAC pairs with the weight 0.158, in the share's against them with -0.908, its cut set being most of the share's denominator.

14.4 The checks (step 5)

The hall has no physics row and no exact constraint, so no limit $\sigma_p \to 0$ of a model tolerance arises and no limit table is needed; the bound $0 < q < 1$ is not a row either, and the logit state handles it. What can fail is the Laplace route, which gives the reported answers, against the exact one, the enumeration and the draws, and the choice of each record alone as the posterior.

14.4.1 The Laplace route against the exact route

The unavailabilities. The spreads of the logits, $\sqrt{\Sigma_{jj}} = 0.586$, 0.503 and 1.000, correspond in the unavailabilities to $q_j(1 - q_j)\sqrt{\Sigma_{jj}} = 0.0171$, 0.0050 and 0.0010, against the exact Beta standard deviations 0.0170, 0.0050 and 0.0010. The exact variances of the logits, equation (14.9), are 0.405, 0.286 and 1.646, against the diagonal (0.344, 0.253, 1.001) of $\boldsymbol{\Sigma}$: the Laplace route understates the controller's most, its record holding a single unit found down.

The cut sets. Equation (14.29) with the pair's row of logarithms (1.94, 0, 0) gives $\mathbf{g}^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g} = 1.94^2 \times 0.3436 = 1.293$ and the band $9.0 \times 10^{-4}\, \mathrm{e}^{\pm 1.96 \times 1.137} = [9.7 \times 10^{-5},\ 8.4 \times 10^{-3}]$, against the exact $[4.0 \times 10^{-5},\ 5.1 \times 10^{-3}]$; the pumps' Laplace band is $[1.4 \times 10^{-5},\ 7.0 \times 10^{-4}]$ against $[7.5 \times 10^{-6},\ 4.8 \times 10^{-4}]$, and the controller's $[1.4 \times 10^{-4},\ 7.1 \times 10^{-3}]$ against $[2.5 \times 10^{-5},\ 3.7 \times 10^{-3}]$. The lower ends lie 1.9 to 2.4 times above the exact ones for the pairs and the pumps, and 5.6 times for the controller, whose logit posterior from a single unit found down is the least normal. For the order of a pair against the controller, equation (14.33) gives $d(\hat{\mathbf{z}}) = \ln(9.0 \times 10^{-4} / 1.0 \times 10^{-3}) = -0.105$ and $\Phi(-0.105/1.514) = \Phi(-0.070) = 0.47$, against 0.52 over the draws; for the six pairs together, $d(\hat{\mathbf{z}}) = \ln 5.4 = 1.686$ and $\Phi(1.114) = 0.87$, as over the draws. Both routes call the single order even.

The top event. Table 14.10 sets the reported band of the top event, and the bands of the three truncations, beside the bands of the same quantities over the draws. Each row is

a quantity; the columns give its plug-in value, the gradient of its logarithm with respect to (ξ_C, ξ_P, ξ_K), the standard deviation of the logarithm, the Laplace band and the band over the draws.

Table 14.10: Laplace bands of the top event and of its truncations, equation (14.29), with the gradients of equation (14.45) and $\mathbf{\Sigma}$ of equation (14.10), against the central 95 % bands of the same quantities over the 4×10^5 draws (seed 14).

Quantity f	Plug-in $f(\hat{\mathbf{z}})$	Gradient of $\ln f$, $\mathbf{g}/f$	sd of $\ln f$	Laplace 95 % band	Draws 95 % band
exact U_{sys}	6.281×10^{-3}	(1.568, 0.031, 0.158)	0.933	$[1.0 \times 10^{-3}$, $3.9 \times 10^{-2}]$	$[8.2 \times 10^{-4}$, $2.8 \times 10^{-2}]$
rare-event sum S_1	6.500×10^{-3}	(1.612, 0.030, 0.154)	0.957	$[1.0 \times 10^{-3}$, $4.2 \times 10^{-2}]$	$[8.3 \times 10^{-4}$, $3.1 \times 10^{-2}]$
second order $S_1 - S_2$	6.168×10^{-3}	(1.542, 0.032, 0.161)	0.919	$[1.0 \times 10^{-3}$, $3.7 \times 10^{-2}]$	$[8.2 \times 10^{-4}$, $2.7 \times 10^{-2}]$
third order $S_1 - S_2 + S_3$	6.289×10^{-3}	(1.571, 0.031, 0.158)	0.934	$[1.0 \times 10^{-3}$, $3.9 \times 10^{-2}]$	$[8.2 \times 10^{-4}$, $2.9 \times 10^{-2}]$

The four gradients are nearly one vector, because the four quantities differ only by the joint terms, and in each the entry of ξ_C carries 97 % of the variance of the logarithm. The bounds hold at every state, so over the draws their quantiles keep the order of equation (14.23), the upper ends $2.7 \times 10^{-2} \le 2.8 \times 10^{-2} \le 3.1 \times 10^{-2}$. The Laplace bands have the width of the draws' bands, standard deviations of 0.92 to 0.96 against 0.909 for $\ln U_{\text{sys}}$ over the draws, and sit too high by a factor of about 1.2: the delta method centers on $\ln U_{\text{sys}}(\hat{\mathbf{z}}) = -5.070$, where the draws' mean of $\ln U_{\text{sys}}$ is -5.245, because the plug-in of a quantity that grows as q_C^2 sits above its median when the posterior of q_C is skewed to the right; the median of $\text{Beta}(3, 97)$ is 0.0269, below its mean 0.030. With the exact variances of the logits in place of the diagonal of $\mathbf{\Sigma}$, the band of the top event widens to $[8.5 \times 10^{-4},\ 4.6 \times 10^{-2}]$ and moves no closer. The delta method on U_{sys} itself does worse: $\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g} = 3.43 \times 10^{-5}$, a standard deviation of 5.9×10^{-3}, gives a normal band from -5.2×10^{-3} to 1.78×10^{-2} that puts 14 % of its probability on negative unavailabilities.

The share. The Laplace route gives the share at $\hat{\mathbf{z}}$, 0.831, the band $[0.22,\ 0.99]$ and $\Pr(\pi_{\text{C}} > \frac{1}{2} \mid \mathbf{y}) \approx \Phi(1.591/1.458) = 0.86$, against the exact $[0.17,\ 0.99]$ and 0.85. Its budget agrees with a decomposition from 2×10^6 draws, the variance of the conditional mean of the share's logit given one entry as a fraction of the total: on the logit 62 % for q_C, 2 % for q_P and 35 % for q_K, against the Laplace 61 %, 0.4 % and 39 %, and on the share itself 66 %, 1 % and 27 %.

What the comparison settles. Where the question is a threshold on a logit, the share above one half or one cut set above another, the Laplace route and the exact one reach the same conclusion. Where the quantity is a small probability under the right-skewed posterior of q_C, the Laplace band has the right width on the logarithmic scale and the wrong location, its lower end too high by a factor of 2.4 for a CRAC pair and 5.6 for the controller, and an expectation over the normal in the logit, a factor of 1.40 on the top event, is too high for the same reason. These are the gaps carried beside each answer in Section 14.3; the variance budget and the preposterior covariance of step 7 are the uses in which the two routes agree most closely.

14.4.2 The enumeration, the bounds and the draws

Table 14.11 collects the checks of the chapter. Each row is a quantity; the columns give its value by the closed form, the cut-set sum or the Laplace route, its value by the full calculation it is

checked against (the enumeration, the exact Beta quantiles, central differences or the draws), and the agreement. Rows six to nine are the comparisons of Section 14.4.1; the paragraphs below take the others.

Table 14.11: The checks of step 5. Laplace bands from equation (14.29) with $\mathbf{\Sigma}$ of equation (14.10); draws 4×10^5, seed 14; bands are central 95 % intervals.

Quantity	Closed form, sum or Laplace	Full calculation	Agreement
coherence of the tree	every gate nondecreasing, equation (14.12)	the 365 successors, one more unit down, of the 113 failed configurations	all 365 failed
plug-in U_{sys}	closed form (14.26), 6.280625×10^{-3}	enumeration of 128 configurations, 6.280625×10^{-3}	every printed digit
$\mathbb{E}[U_{\text{sys}} \mid \mathbf{y}]$	product of moments (14.28), 7.790021×10^{-3}	inclusion–exclusion over 255 subfamilies; draws 7.771×10^{-3}	eight digits; draws 1.6 standard errors low
$\mathbb{E}[U_{\text{sys}} \mid \mathbf{y}]$, state route	product of sums (14.27), 1.09350×10^{-2}	inclusion–exclusion over the same subfamilies with the moments of the sum (14.19), 1.09350×10^{-2}	six digits
bounds (14.23)	$S_1 - S_2 \le U_{\text{sys}} \le S_1$, $U_{\text{sys}} \le S_1 - S_2 + S_3$	each of the 4×10^5 draws	hold in every draw
band of a CRAC pair	Laplace $[9.7 \times 10^{-5},\ 8.4 \times 10^{-3}]$	Beta quantiles squared, $[4.0 \times 10^{-5},\ 5.1 \times 10^{-3}]$	lower end 2.4 times high
band of U_{sys}	Laplace $[1.0 \times 10^{-3},\ 3.9 \times 10^{-2}]$	draws $[8.2 \times 10^{-4},\ 2.8 \times 10^{-2}]$	width right, about 1.2 times high
$\Pr(q_C^2 > q_K \mid \mathbf{y})$	Laplace 0.47	draws 0.52	both even
share π_{C}	Laplace 0.86, band $[0.22,\ 0.99]$	exact 0.85, band $[0.17,\ 0.99]$	threshold to 0.01
$\mathbf{G}$, $\mathbf{g}$, $\mathbf{g}_\pi$	equations (14.35), (14.44), (14.38)	central differences in the logits	$\mathbf{G}$ to six digits; $\mathbf{g}$ and $\mathbf{g}_\pi$ agree

The enumeration. The 128 configuration probabilities at the plug-in values sum to one within 10^{-15}, and the sum over the 113 failed configurations, 6.280625×10^{-3}, equals the closed form (14.26) to every printed digit. The tree is coherent on the enumeration: taking one more unit down in each of the 113 failed configurations gives 365 successors, and every one of them is failed. The full inclusion–exclusion over all $2^8 - 1 = 255$ subfamilies of $\mathcal{M}$, with each term's expectation from equation (14.20), reproduces the expectation (14.28), 7.790021×10^{-3}, to eight digits: the cut sets of equation (14.17) represent the tree completely, as equation (14.16) requires.

The bounds at every draw. On each of the 4×10^5 draws, S_1, S_2 and S_3 from equation (14.22) were set against U_{sys} from the closed form (14.26), and $S_1 - S_2 \le U_{\text{sys}} \le S_1$ and $U_{\text{sys}} \le S_1 - S_2 + S_3$ hold in every draw. The size of the errors moves with the draw. The rare-event sum is 3.1 % high at the median draw and up to 25 % high in the far upper tail of q_C, where the joint terms of the overlapping CRAC pairs, of order q_C^3, are no longer small beside the pairs' q_C^2; the second-order value is 1.6 % low at the median and up to 17 % low; the third order is never more than 5.4 % high. The errors grow with q_C, and the expectation gives weight to the draws with large q_C, which is why the expected values of Table 14.7 carry larger errors, +5.9 % and −3.3 %, than the plug-in values, +3.5 % and −1.8 %.

The draws against the exact mean. The draws give a mean of 7.771×10^{-3} against the exact 7.790×10^{-3}. The standard deviation of U_{sys} over the draws is 7.59×10^{-3}, so the Monte Carlo standard error of the mean is $7.59 \times 10^{-3}/\sqrt{4 \times 10^5} = 1.2 \times 10^{-5}$, and the difference is 1.6 standard errors. The exact mean needs no draws; they are needed for the exact quantiles.

The gradients and the two closed forms. The matrix $\mathbf{G}$ of equation (14.35) agrees with central differences of the eight cut-set probabilities in the logits to six digits, the exact $\mathbf{g}$ with central differences of the closed form (14.26), and $\mathbf{g}_\pi$ with central differences of the share's logit. The closed slope of a CRAC unit, Section 14.3.5, gives $3 \times 0.03 \times 0.97^2 \times 0.9999 \times 0.999 = 0.0846$, the entry of Table 14.9. Equation (14.42) gives the controller change exactly, $0.9933 \times 0.9999 \times 0.000998 = 9.9 \times 10^{-4}$, the difference of the exact means of Table 14.8.

The definition of the share. Defined instead as the probability that two or more CRAC units are down divided by the exact U_{sys}, the share has the band $[0.16,\ 0.99]$ and $\Pr(\text{share} > \frac{1}{2} \mid \mathbf{y}) = 0.85$, the same to two decimals: the choice of denominator does not drive the answer to the fourth question.

14.4.3 Sensitivity to the prior

Each posterior was taken as its record alone, and the controller's record holds a single unit found down. Table 14.12 repeats the analysis under two priors added to every record. Each row is a prior; the columns give the three posteriors, the expected top event, the probability that the CRAC pairs carry more than half of the cut-set sum, the probability that one pair exceeds the controller, the shares of the expected top event that the fifth CRAC unit and the redundant controller remove, and the probability that the fifth unit removes more, each on the state route with the exact value in parentheses. The first row is the record alone; the second adds the uniform prior $\text{Beta}(1, 1)$ to each record; the third the Jeffreys prior $\text{Beta}(\frac{1}{2}, \frac{1}{2})$, the reference prior of Chapter 5. One row of the arithmetic, the uniform prior: the controller's posterior is $\text{Beta}(1 + 1,\ 999 + 1) = \text{Beta}(2, 1000)$, of mean 0.0020, twice the record's, so that $\hat{z}_K = \ln(2/1000)$ and $\Sigma_{KK} = 1/2 + 1/1000$. The exact probabilities are from 4×10^5 draws, seed 14.

Table 14.12: Sensitivity of the hall's answers to the prior on each group's unavailability: state route, by the sums (14.19) and the normal tails at the Laplace posterior of each row, with the exact value in parentheses: expected top event from the Beta moments, the other probabilities from 4×10^5 draws, seed 14, with the share's probability by the exact Beta tail of equation (14.39); removals as fractions of the expected top event of the present design.

Prior	Posteriors q_C, q_P, q_K	$\mathbb{E}[U_{\text{sys}}]$	$\Pr(\pi_{\text{C}} > \frac{1}{2})$	$\Pr(q_C^2 > q_K)$	Fifth unit removes	Second controller removes	Fifth removes more
record alone	$\text{Beta}(3, 97)$, $(4, 396)$, $(1, 999)$	1.09×10^{-2} (7.79×10^{-3})	0.86 (0.85)	0.47 (0.52)	75 (79) %	15 (13) %	0.898 (0.865)
uniform	$\text{Beta}(4, 98)$, $(5, 397)$, $(2, 1000)$	1.60×10^{-2} (1.27×10^{-2})	0.89 (0.87)	0.41 (0.43)	73 (75) %	16 (15) %	0.926 (0.874)
Jeffreys	$\text{Beta}(3.5, 97.5)$, $(4.5, 396.5)$, $(1.5, 999.5)$	1.34×10^{-2} (1.02×10^{-2})	0.88 (0.86)	0.43 (0.46)	74 (76) %	15 (15) %	0.910 (0.864)

The level of the answer moves; the conclusions drawn from it do not. On the state route the expected top event rises by 22 % under the Jeffreys prior and by 46 % under the uniform one, to 117 h and 140 h a year, and exactly by 31 % and 63 %, to 89 h and 112 h a year, because a pseudo-count of one half or one is not small beside three, four and one units found down: the

CRAC units' exact mean rises from 0.030 to 0.035 or 0.039, and with it the second moment of every pair, which carries most of the rise, and the controller's mean rises by half or doubles. The attribution stays at 0.86 to 0.89 (0.85 to 0.87), a single pair against the controller stays an even question, between 0.41 and 0.47 (0.43 and 0.52), and the fifth CRAC unit removes about five times what the redundant controller removes under every prior (five to six times exactly), with posterior probability 0.90 to 0.93 (0.86 to 0.87) of removing more. The report gives these answers with their range over the three priors, which is wide only for the expected loss of cooling.

A dispersion check, whether the four CRAC units share one unavailability or one of them fails more often, would compare the units' own counts with the binomial spread of a common one; the records are kept as group totals, so the check cannot be made from them, and Section 14.6.6 returns to what that leaves open.

14.5 The report (step 6)

Neither half of the first reading survives. The sum of the cut-set probabilities at the records' fractions, 6.5×10^{-3}, was read as an upper bound on the loss of cooling, and it falls 41 % short of the expected probability that the records imply, 1.09×10^{-2} (17 % short of the exact 7.79×10^{-3}): the sum bounds the top event at one value of the unavailabilities, not its expectation over what the records leave open. The list put the controller first and pointed to a second controller; in expectation the six CRAC pairs carry 85 % of the cut-set sum (86 %) and carry more than half of it with posterior probability 0.86 (0.85), a single pair against the controller is an even question, and a redundant controller removes 15 % (13 %) of the expected loss of cooling where a fifth CRAC unit removes 75 % (79 %). The answers below are given in the units in which the questions were asked, from the posterior of the state with the exact value beside each; the expectations differ by up to a factor of 1.4 on so few findings, and no conclusion of the report changes between the two routes.

14.5.1 The answers and their bands

Table 14.13 gives the answers. Each row is a question of Table 14.1, in its order; the second column gives the answer from the posterior of the state, the records alone, the third the exact value, and the last the band or comparison: the central 95 % posterior interval, a posterior probability, or the range over the three priors of Table 14.12 where the prior moves the answer, on the state route with the exact value in parentheses.

The report line. The hall loses its cooling through the control system alone, any two of the four CRAC units, or both chilled-water pumps, about 96 hours a year in expectation and between 9 and 342 hours with 95 % probability (68 hours, and 7 to 249 hours, from the exact posterior), most of it through the CRAC pairs; the figures assume that units fail independently given their unavailabilities.

14.5.2 A ranking that is probable, not certain

The minimal cut sets follow from the drawing and are certain while it stands; their probabilities and their ranking are not. On both routes the controller's cut set exceeds all six pairs together with probability about one in eight and any single pair with probability about one in two. A list in order of plug-in probabilities puts the controller ahead of each pair, where the evidence supports neither order, and hides that the order is uncertain at all. The ranking is therefore reported with the probability that it holds, and a decision that depends on it is checked against the outcomes in which it fails; the choice of change survives that check.

Table 14.13: The answers to the six questions asked of the hall's cooling, from the posterior of the state, with the exact value, from the Beta moments and from draws, beside each. Bands are central 95 % posterior intervals.

Question	Answer	Exact	Band and comparison
Which combinations take the service down?	eight minimal cut sets: {K}, the six CRAC pairs, $\{\mathrm{P1}, \mathrm{P2}\}$	the same; no state route is needed	no cut set of order three or more is minimal; 113 of 128 configurations are failed
How likely is each combination, and which dominate?	each CRAC pair 1.66×10^{-3}; controller 1.64×10^{-3}; pumps 1.63×10^{-4}	1.19×10^{-3}; 1.00×10^{-3}; 1.25×10^{-4}	bands $[9.7 \times 10^{-5},\ 8.4 \times 10^{-3}]$, $[1.4 \times 10^{-4},\ 7.1 \times 10^{-3}]$, $[1.4 \times 10^{-5},\ 7.0 \times 10^{-4}]$ ($[4.0 \times 10^{-5},\ 5.1 \times 10^{-3}]$, $[2.5 \times 10^{-5},\ 3.7 \times 10^{-3}]$, $[7.5 \times 10^{-6},\ 4.8 \times 10^{-4}]$); a pair against the controller is unresolved, $\Pr(q_C^2 > q_K \mid \mathbf{y}) = 0.47$ (0.52), 0.41 to 0.47 over the priors (0.43 to 0.52); the six pairs exceed the controller with probability 0.87 (0.87)
What is the probability of losing cooling, and how well does the rare-event sum do?	mean 1.09×10^{-2} (96 h per year), median 6.28×10^{-3} (55 h)	mean 7.79×10^{-3} (68 h), median 5.42×10^{-3} (47 h)	$[1.0 \times 10^{-3},\ 3.9 \times 10^{-2}]$, $[9,\ 342]$ h per year ($[8.2 \times 10^{-4},\ 2.8 \times 10^{-2}]$, $[7,\ 249]$ h); mean 1.34×10^{-2} to 1.60×10^{-2} under the Jeffreys and uniform priors (1.02×10^{-2} to 1.27×10^{-2}); plug-in sum 6.50×10^{-3}, 41 % (17 %) below the mean; expected rare-event sum +7.6 % (+5.9 %), second order −4.3 % (−3.3 %), third order +0.8 % (+0.4 %)
How much of the top event is in the CRAC pairs, and is that conclusion firm?	0.83 at the MAP state, the median of the share; $\Pr(\pi_{\mathrm{C}} > \frac{1}{2} \mid \mathbf{y}) = 0.86$	0.86 as a ratio of expectations, posterior median 0.84; 0.85	band $[0.22,\ 0.99]$ ($[0.17,\ 0.99]$); 0.86 to 0.89 over the priors (0.85 to 0.87)
Which single change removes the most?	a fifth CRAC unit: mean 2.73×10^{-3}, −75 %	1.66×10^{-3}, −79 %	redundant controller 9.31×10^{-3}, −15 % (6.80×10^{-3}, −13 %); both 1.10×10^{-3} (6.6×10^{-4}); the fifth unit removes more with probability 0.898 (0.865), 0.90 to 0.93 over the priors (0.86 to 0.87)
Which basic event's probability matters most?	by sensitivity the controller, $\partial U / \partial q_K = 0.995$; by contribution the CRAC units, $\mathbf{g} = 10^{-3}(9.85,\ 0.20,\ 0.99)$	the same at the MAP state	controller band $[0.956,\ 0.9994]$ ($[0.973,\ 0.9997]$); each CRAC unit 0.085, band $[0.030,\ 0.24]$ ($[0.019,\ 0.18]$), shared entry 0.34; each pump 0.0099, band $[0.0037,\ 0.026]$ ($[0.0027,\ 0.022]$); the column of ξ_C carries 85 % (86 %) of the expected sum and 97 % of the variance of $\ln U$

Common mistake

Computing the probability of a CRAC pair as the square of the posterior mean, $0.03^2 = 9.0 \times 10^{-4}$, when the four units share one posterior. The expectation is the second moment, 1.66×10^{-3} on the state route, 1.84 times as large, and 1.19×10^{-3} exactly, a third larger, and the error grows with the order of the cut set, to a factor of 3.90 for a triple, 2.16 exactly. With the plug-in values the controller appears as the single largest cut set; in expectation each pair is slightly larger, and on this evidence neither order holds with better than even odds.

14.6 The next measurement (step 7)

The fourth answer is an attribution: it says where the loss of cooling comes from, which is what the operator reports and what directs maintenance attention and the study of failure causes. It does not decide the design. The fifth CRAC unit removes more than the redundant controller with posterior probability 0.898 (0.865 exactly) and five times as much in expectation (six), under every prior tried, and after a hundred further CRAC checks that probability is never below 0.57 (0.60) (Section 14.6.4). The attribution, at 0.86 (0.85), is the less settled answer, and the next measurement is the record that would settle it: more checks of the CRAC units, or more checks of the controller, the second largest term of the sum.

14.6.1 The candidate records

Four candidates are weighed: a hundred or four hundred further checks of the CRAC units, kept as the present record was, and a hundred or a thousand further checks of the controller. Suppose N further checks of the CRAC units find n_{d} of them down. The posterior becomes $\mathrm{Beta}(3 + n_{\mathrm{d}},\ 97 + N - n_{\mathrm{d}})$, and before the checks are made n_{d} has, under the Laplace normal of the entry, the predictive distribution of the sum (14.19), with h the binomial probability of the count, and on the exact route the Beta–binomial predictive distribution of Chapter 5,

$$p(n_{\mathrm{d}} \mid \mathbf{y}) = \binom{N}{n_{\mathrm{d}}} \frac{B(3 + n_{\mathrm{d}},\ 97 + N - n_{\mathrm{d}})}{B(3,\ 97)}, \qquad n_{\mathrm{d}} = 0, \ldots, N. \tag{14.47}$$

For each outcome n_{d} the share has a new posterior and a new band; the preposterior summaries are averages over the outcomes, the design form (2.16) with a sum over the counts. Two of them are useful. The expected width of the share's central 95 % band is the average, weighted by $p(n_{\mathrm{d}} \mid \mathbf{y})$, of the width after each outcome. The probability that the record settles the question is the predictive probability of the outcomes after which $\Pr(\pi_{\mathrm{C}} > \frac{1}{2} \mid \mathbf{y}, n_{\mathrm{d}})$ is at least 0.95 or at most 0.05. A third summary looks natural and is useless: by the law of total probability,

$$\sum_{n_{\mathrm{d}}=0}^{N} p(n_{\mathrm{d}} \mid \mathbf{y})\ \Pr(\pi_{\mathrm{C}} > \tfrac{1}{2} \mid \mathbf{y}, n_{\mathrm{d}}) = \Pr(\pi_{\mathrm{C}} > \tfrac{1}{2} \mid \mathbf{y}), \tag{14.48}$$

so the expected posterior probability of the event equals its present probability for every candidate record. A record cannot be expected to move the probability up or down; it can only be expected to push it toward one end or the other. That law holds for the exact posteriors; the state route after each outcome is a new Laplace fit, and keeps it to about 0.005: from 0.862 at present its averages run from 0.857 to 0.861. A record of controller checks is treated in the same way with the controller's posterior $\mathrm{Beta}(1 + n_{\mathrm{d}},\ 999 + N - n_{\mathrm{d}})$.

14.6.2 The preposterior covariance

In the state the same question has a closed form under the Laplace posterior. A record of N checks of group j is one more count row on entry j, of the form of equation (14.5), whose

information at $\hat{\mathbf{z}}$ is $Nq_j(1-q_j)$. Information adds, so the preposterior covariance of Chapter 2, equation (2.17), is

$$\boldsymbol{\Sigma}_a = \left(\boldsymbol{\Sigma}^{-1} + \boldsymbol{\Lambda}_a\right)^{-1}, \qquad \boldsymbol{\Lambda}_a = Nq_j(1-q_j)\,\mathbf{e}_j\mathbf{e}_j^{\mathsf{T}}, \qquad \mathrm{Var}(\mathrm{logit}\,\pi_{\mathrm{C}} \mid \mathbf{y}, a) \approx \mathbf{g}_\pi^{\mathsf{T}}\boldsymbol{\Sigma}_a\mathbf{g}_\pi, \quad (14.49)$$

with $\mathbf{e}_j$ the unit vector of entry j and $\mathbf{g}_\pi$ the gradient of the share's logit of Section 14.3.3; the outcome of the record does not enter $\boldsymbol{\Sigma}_a$, so no average over outcomes is needed. Chapter 20, in Volume II, ranks candidate rows by the same update in its rank-one form. Each candidate's Jacobian row is $\mathbf{J}_a = \sqrt{Nq_j(1-q_j)}\;\mathbf{e}_j^{\mathsf{T}}$, so that $\boldsymbol{\Lambda}_a = \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a$; at $\hat{\mathbf{z}}$ it is (1.706, 0, 0) for a hundred CRAC checks, (3.412, 0, 0) for four hundred, (0, 0, 0.316) for a hundred controller checks and (0, 0, 0.9995) for a thousand. For the two main candidates,

$$\boldsymbol{\Sigma}_a^{100\ \mathrm{CRAC}} = \begin{bmatrix} 0.172 & 0 & 0 \\ 0 & 0.253 & 0 \\ 0 & 0 & 1.001 \end{bmatrix}, \qquad \boldsymbol{\Sigma}_a^{1000\ \mathrm{K}} = \begin{bmatrix} 0.344 & 0 & 0 \\ 0 & 0.253 & 0 \\ 0 & 0 & 0.500 \end{bmatrix},$$

against the diagonal (0.344, 0.253, 1.001) of $\boldsymbol{\Sigma}$: each record halves the entry it feeds and leaves the others where they were; four hundred CRAC checks take Σ_{CC} to 0.069 and a hundred controller checks take Σ_{KK} to 0.910.

Table 14.14 gives the result for the four candidates. Each row is a candidate record; the columns give the entry it feeds, the information it adds by equation (14.49), the Laplace standard deviation of the share's logit after it, the expected width of the share's band, the probability that it settles the question at 0.95, and the expected posterior probability of $\pi_{\mathrm{C}} > \frac{1}{2}$. The last three columns give the state route, the Laplace normal after each outcome weighted by the Laplace predictive of that outcome, with in parentheses the exact posterior, by 2×10^5 draws of the three posteriors with common random numbers across the outcomes, seed 141, the exact Beta tail of equation (14.39), which equation (14.48) fixes at 0.85 in the last column.

Table 14.14: The preposterior of the share of the CRAC pairs for four candidate records. A CRAC check is one check of one unit. Added information and Laplace standard deviation from equation (14.49); the other columns on the state route, with the exact posterior in parentheses, by 2×10^5 draws. Settled: the posterior probability of $\pi_{\mathrm{C}} > \frac{1}{2}$ after the record is at least 0.95 or at most 0.05.

Record	Entry fed	Added information	Laplace sd of the share's logit	Expected band width	Probability settled	Expected $\Pr(\pi_{\mathrm{C}} > \frac{1}{2})$
present	–	–	1.46	0.77 (0.82)	–	0.86 (0.85)
100 more CRAC checks	ξ_C	2.91	1.22	0.64 (0.66)	0.28 (0.35)	0.86 (0.85)
400 more CRAC checks	ξ_C	11.64	1.05	0.56 (0.54)	0.47 (0.51)	0.86 (0.85)
100 more controller checks	ξ_K	0.10	1.43	0.76 (0.82)	0.00 (0.00)	0.86 (0.85)
1000 more controller checks	ξ_K	1.00	1.31	0.71 (0.78)	0.00 (0.00)	0.86 (0.85)

14.6.3 The ranking of the records

A hundred more CRAC checks, which double the CRAC record, narrow the expected band from 0.77 to 0.64 and settle the question with probability 0.28: if five or more of the hundred find a unit down, the posterior probability that the CRAC pairs dominate passes 0.97. Exactly, they narrow it from 0.82 to 0.66 and settle it with probability 0.35: if four or more of the hundred

find a unit down, the probability passes 0.96. A thousand controller checks, which double the controller's record, narrow the band only to 0.71 (0.78) and settle nothing: even after the most favorable outcome, a thousand checks with no unit found down, the probability is 0.94 (0.92). The reason is not that the share ignores the controller; Section 14.3.3 found that ξ_K carries more than a third of its variance. It is read off equation (14.49). A check adds $q_j(1-q_j)$ to the diagonal entry of $\mathbf{J}^\mathsf{T}\mathbf{J}$ it feeds, Section 14.1.7: 0.0291 per CRAC check against 0.000999 per controller check, twenty-nine times less. A hundred CRAC checks add 2.91, as much as the present CRAC record, and halve Σ_{CC} from 0.344 to 0.172; it takes a thousand controller checks to halve Σ_{KK} from 1.001 to 0.500. The share's logit then loses the squared entry of $\mathbf{g}_\pi$ times the reduction, $1.94^2 \times 0.172 = 0.647$ for the hundred CRAC checks against $0.908^2 \times 0.500 = 0.413$ for the thousand controller checks, because $\mathbf{g}_\pi$ weights ξ_C twice as heavily as ξ_K; the Laplace standard deviations of the table fall in the order of the exact band widths. In the exact posterior the controller's entry moves less than the Laplace route says. The exact variance of its logit, equation (14.9), averaged over the outcomes of a thousand checks, falls from 1.65 to 1.06, by 35 %, where Σ_{KK} falls from 1.00 to 0.50, by 50 %: with probability 0.50 the checks find no unit down, which leaves that variance at 1.65 and only moves the posterior down the scale, raising the probability that the CRAC pairs dominate to 0.92. The CRAC entry follows the Laplace route closely: a hundred checks take its exact variance from 0.405 to 0.220 in expectation, a fall of 46 %, against 50 %. Even four hundred CRAC checks settle the question only about half the time: the share is a ratio of a square of an uncertain probability to a sum that contains another uncertain probability, and a few hundred checks of a probability near 0.03 still record only a handful of units found down.

14.6.4 The recommendation

The reported answer whose band the record is to narrow is the share of the CRAC pairs, band [0.22, 0.99], of width 0.77 ([0.17, 0.99], of width 0.82, exactly). In expectation a hundred CRAC checks remove 0.12 of that width and four hundred remove 0.21 (exactly 0.16 and 0.28); a hundred controller checks remove 0.009 and a thousand 0.06 (exactly 0.006 and 0.04). The next record is therefore a hundred further checks of the CRAC units, kept as the present record was, which double it, remove twice the width a thousand controller checks would (nearly four times exactly) and settle the attribution with probability 0.28 (0.35); four hundred are the choice if the attribution must be settled rather than narrowed, since they settle it about half the time. The record enters as one more count row on ξ_C, and the method returns to step 3 with it. Further controller checks do little for this question, although they remain the record of the controller's own cut set. The same preposterior run on the design question shows how far the design need wait for either record. After a hundred CRAC checks the probability that the fifth unit removes more than the redundant controller reaches 0.95 with probability 0.41 (0.35 exactly) and is never below 0.57 (0.60). After a thousand controller checks the two routes fall on opposite sides: on the state route that probability reaches 0.95 with probability 0.38, the chance of no unit found down, where it stands at 0.985; exactly, a thousand controller checks never take it to 0.95. The state route errs there because it narrows the normal in the logit of q_K at every outcome, where the exact variance of that logit does not fall at all when no unit is found down. The price of the CRAC record, set against the value of the decisions it informs, is the subject of a later chapter on the value of information.

14.6.5 What the analysis showed

The first reading of the hall's records said that the controller is its weak point and that the sum of the cut-set probabilities bounds its loss of cooling; the analysis overturns both. The loss of cooling in this hall lives in its CRAC units. Their six pairs carry most of the expected cut-set sum, and the reason is their record rather than their point value: three units found down in a

hundred checks leave the CRAC unavailability uncertain by a factor of eleven, so that its second moment, not the square of its mean, sets the weight of each pair, and the same width makes the sum at the means fall short of the expected loss. The controller remains the hall's only single point of failure and leads every sensitivity, yet making it redundant buys about a seventh of the expected loss of cooling where a fifth CRAC unit buys three quarters (about an eighth and four fifths from the exact posterior), and that order holds with probability 0.90 (0.865) and under every prior tried. What stays open is the attribution itself, and it is narrowed most cheaply by more checks of the CRAC units, each worth twenty-nine controller checks in information, not by watching the controller longer. The engineer carries back to the site three things: order the CRAC checks, price the fifth unit against a second controller, and report the loss of cooling as a band from a few hours to a few hundred hours a year rather than as the number at the means.

14.6.6 Limits of the model

The minimal cut sets are exact for the drawing and the probabilities are exact for the model; the drawing and the model rest on assumptions that the records do not test.

The drawing. The tree assumes that either pump carries the full flow, that three CRAC units cool the hall, and that the controller is the only unit both groups depend on. A pump that cannot carry the flow alone turns each pump into a cut set of order one; a CRAC unit with its own local controller adds basic events the tree lacks. The drawing is checked against the plant, not against the records. Under Chapter 13's two-unit drawing, of weight 0.3, the CRAC pairs are not cut sets and the controller carries about 83 % of the cut-set sum at the plug-in values (Exercise 14.2).

Independence, shared posteriors and common causes. Every probability of the chapter assumes that the units fail independently given their unavailabilities. The shared posterior already couples them in knowledge: if the CRAC units turn out to be worse than the record suggests, they are all worse together, and that is the source of the gap between $\mathbb{E}[q_C^2]$ and $\mathbb{E}[q_C]^2$. It does not couple them physically. A contaminated chilled-water supply, a control signal that stops both pumps, or a maintenance error repeated on two CRAC units in one visit would fail several units at once, and none of these is a basic event of the tree. Such a cause adds a cut set of order one, the cause itself, whose probability can exceed that of the pair it would otherwise take, and it couples the units' records, since a double failure in the record is then evidence about the cause and not only about each unit. The pumps' cut set, here the smallest at 1.4 % (1.5 % exactly), is the one most exposed, two pumps on one supply and maintained by one crew, and the redundant controller of the design comparison is exposed in the same way.

One unavailability per group. The four CRAC units share one entry because they are of one design under one record. If one of them fails more often, the pairs that contain it are more probable than the others and the six bars of Figure 14.3 are no longer equal; the group totals cannot reveal it, and a record kept per unit is what would test it.

When the rare-event sum and its corrections fail. The rare-event sum is an upper bound with an error of at most S_2, and it is close when the cut sets' probabilities are small and their overlaps few. Both conditions fail in practice. Units with high unavailabilities, standby units checked rarely or units under long maintenance, make the joint occurrence of two cut sets comparable to either alone; Exercise 14.5 has a pair of cut sets that share a unit, where the rare-event sum overstates the top event by 11 %. Many cut sets drawn from one redundant group overlap heavily, as the six CRAC pairs do, and then the second-order value, a lower bound for any tree, falls well below the exact value, here by 4.3 % on the state route and 3.3 %

exactly, so that a value computed as a correction of the rare-event sum is not guaranteed to be better than the sum itself. The Bonferroni bounds always hold and bracket the exact value; when the bracket is too wide for the question, and the tree is small enough, the enumeration ends the matter. For the hall it took 128 configurations.

Common mistake

Treating the second-order inclusion–exclusion value as a refinement that is always closer to the exact top event than the rare-event sum. It is a lower bound. For the hall, in expectation, the rare-event sum is 7.6 % high and the second-order value 4.3 % low on the state route (5.9 % and 3.3 % exactly); with more overlap among the cut sets the second-order error can exceed the first-order one. Only the enumeration, or the full expansion, is exact.

The level and the prior. With three, four and one units found down, the level of the expected loss of cooling depends on the prior by up to 46 % on the state route and 63 % exactly (Section 14.4.3), most of it through the second moment of the CRAC unavailability; the attribution and the design comparison do not. A level that must be quoted to a customer needs a prior argued from the history of similar units, or longer records, the CRAC record first.

The next chapter takes the two chilled-water pumps of this tree and asks whether one cause can take both out at once, freeing a shared cause as an entry of the state with its own evidence.

14.7 What to retain

Definition 14.1 (Minimal cut sets read off the enumeration)**.** For a coherent structure with structure function ϕ over n units, the cut sets are the down sets $\mathcal{D}(\mathbf{s})$ of the failed configurations, $\phi(\mathbf{s}) = 0$, and the minimal cut sets $\mathcal{M}$ are those with no proper subset among them; the top event is $1 - \prod_{c\in\mathcal{M}}(1 - \prod_{i\in c} x_i)$. Given the unavailabilities, $\Pr(E_c \mid \mathbf{q}) = \prod_{i\in c} q_i$; under a posterior $\mathrm{Beta}(a, b)$ shared by the $|c|$ units of c, $\mathbb{E}[P_c \mid \mathbf{y}] = \prod_{h=0}^{|c|-1}(a+h)/(a+b+h)$, and under the Laplace normal of the state it is a Gauss–Hermite sum over the entry. The top event's probability is $\sum_\nu (-1)^{\nu+1} S_\nu$, bounded by $S_1 - S_2 \le U_{\mathrm{sys}} \le S_1$ at every $\mathbf{q}$ and so in expectation. With the unavailabilities carried as logits in the state, each check record is a count row with the Jacobian entry $\sqrt{n_j q_j (1 - q_j)}$, so that $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ holds about the expected number of units found down; the cut sets' gradients form the matrix $\mathbf{G} = \partial\mathbf{P}/\partial\mathbf{z} = \mathrm{diag}(\mathbf{P})\,(\mathbf{C}\mathbf{A})\,\mathrm{diag}(1 - q_j)$, with $\mathbf{C}$ and $\mathbf{A}$ read off the plant graph; the top event's gradient is $\mathbf{g} = \mathbf{D}\mathbf{A}^{\mathsf{T}}\nabla_{\mathbf{q}} U_{\mathrm{sys}} \approx \mathbf{G}^{\mathsf{T}}\mathbf{1}$, and each quantity f has the Laplace band $\mathrm{Var}(\ln f) \approx \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}/f^2$.

> **Principle 14.3.** *Report every band and tail from the Laplace covariance $\boldsymbol{\Sigma}$ and the gradients at $\hat{\mathbf{z}}$, every expectation from a sum over the state, with the exact value beside each where one exists: for a quantity of one entry, such as a single cut set, from that entry's Beta quantiles. The same covariance gives the variance budget across entries, the covariance of the cut sets and the value of a record.*

Three conclusions carry beyond the hall. The enumeration that gives the probability of the top event also gives its minimal cut sets, and for a small tree it gives both exactly; the tree's algebra is needed only when 2^n configurations are too many. Under a shared posterior a cut set's probability is a moment of the posterior, and cut sets of different orders must be compared by their expectations, which can reverse the order seen at point values; the order of the expectations is then a claim with a posterior probability of its own, here about even for a single pair against the controller, and a sum of cut-set probabilities at point values is not a bound on the expected top event. And every band, ranking and importance of the chapter is read from one matrix: the rows of $\mathbf{G}$ give the cut sets' bands and their covariance, the difference of two rows the

probability of their order, and the column sums the top event's gradient in the state, which measures contribution where the gradient in the unavailabilities measures sensitivity; design follows the columns, and the record worth buying is the one that adds information $Nq_j(1-q_j)$ to the entry the share's gradient weights most, which for a probability is a record expected to find failures.

Table 14.15 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 14.15: The terms of fault trees and minimal cut sets and where each is defined.

Term	Expression	Units	Meaning	Defined in
State and covariance	$\mathbf{z}=(\xi_C,\xi_P,\xi_K)^{\mathsf{T}}$, $\mathbf{\Sigma}=(\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1}$	–	Laplace posterior of the logits	14.2.1
Count row	$J_{jj}=\sqrt{n_j q_j(1-q_j)}$	–	information $\approx$ expected number found down	14.1.7
Gates	OR $1-\prod(1-x_l)$; AND $\prod x_l$; k_{g}/n $\mathbf{1}(\sum x_l \geq k_{\mathrm{g}})$	–	Boolean functions of inputs	14.3.1
Dual structure function	$\phi^{\mathrm{D}}(\mathbf{1}-\mathbf{s})=1-\phi(\mathbf{s})$	–	top-event indicator, the tree on a configuration	14.3.1
Cut set	every configuration with c down is failed	–	a down set of a failed configuration	14.3.1
Minimal cut set	no proper subset is a cut set	–	essential combination	14.3.1
Minimal-cut-set form	$\phi^{\mathrm{D}}=1-\prod_c(1-\prod_{i\in c}x_i)$	–	OR over AND gates	14.3.1
Enumeration algorithm	enumerate, filter, remove supersets	–	cost 2^n configurations	14.3.1
Cut-set probability	$P_c=\prod_{i\in c}q_i$	–	given $\mathbf{q}$	14.3.2
Shared-posterior moment	$\mathbb{E}[q^{\lvert c\rvert}]=\prod_{h<\lvert c\rvert}(a+h)/(a+b+h)$	–	exceeds $\mathbb{E}[q]^{\lvert c\rvert}$	14.3.2
Inclusion–exclusion	$U=\sum_\nu(-1)^{\nu+1}S_\nu$	–	exact union	14.3.2
Rare-event sum	$S_1=\sum_c P_c$	–	upper bound	14.3.2
Bonferroni bounds	$S_1-S_2\leq U\leq S_1$	–	truncations alternate	14.3.2
Exact expectation	products of group moments, eq. (14.28)	–	no draws for the mean	14.3.2
Cut-set gradient matrix	$\mathbf{G}=\partial\mathbf{P}/\partial\mathbf{z}=\operatorname{diag}(\mathbf{P})(\mathbf{C}\mathbf{A})\operatorname{diag}(1-q_j)$	–	covariance $\mathbf{G}\mathbf{\Sigma}\mathbf{G}^{\mathsf{T}}$	14.3.2
Laplace band on the log	$\operatorname{Var}(\ln f)\approx\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}/f^2$	–	budget across entries	14.3.2
Share of a group	$\pi_{\mathcal{B}}=\sum_{c\in\mathcal{B}}P_c/S_1$	–	Fussell–Vesely form	14.3.3
Design by drawings	$U(d)=\mathbb{E}[f_d(\mathbf{z})\mid\mathbf{y}]$	–	no new evidence	14.3.4
Gradient in the units	$U(q_i=1)-U(q_i=0)$	–	Birnbaum importance; sensitivity	14.3.5
Gradient in the state	$\mathbf{g}=\mathbf{D}\mathbf{A}^{\mathsf{T}}\nabla_{\mathbf{q}}U\approx\mathbf{G}^{\mathsf{T}}\mathbf{1}$	–	contribution	14.3.5
Preposterior of a share	$\mathbf{g}_\pi^{\mathsf{T}}(\mathbf{\Sigma}^{-1}+\mathbf{\Lambda}_a)^{-1}\mathbf{g}_\pi$; average over $p(n_{\mathrm{d}}\mid\mathbf{y})$	–	expected band width	14.6.2

14.8 Exercises

Exercise 14.1. *Formulation.* A tree has an AND gate at the top whose two inputs are the OR gates $O_1=A\vee B$ and $O_2=B\vee C$. (a) Enumerate the eight configurations of A, B, C, list the failed ones and read off the minimal cut sets. (b) With $q_A=0.1$, $q_B=0.01$ and

$q_C = 0.2$, compute the rare-event sum and the exact probability of the top event. ((a) The failed configurations are the four with B down and the one with A and C down and B up; the minimal cut sets are $\{B\}$ and $\{A, C\}$: the shared input B is a cut set of order one although no gate has it alone. (b) $S_1 = 0.01 + 0.1 \times 0.2 = 0.0300$; exact $q_B + (1 - q_B)q_A q_C = 0.01 + 0.99 \times 0.02 = 0.0298$, an overcount of 0.7 %.)

Exercise 14.2. *Formulation.* The hall's tree is changed so that the hall needs only two of its four CRAC units, a vote gate 3/4. (a) List the minimal cut sets. (b) Compute the rare-event sum and the exact top event at the plug-in values $q_C = 0.03$, $q_P = 0.01$, $q_K = 0.001$. (c) Write the fourth question of Table 14.1 for this tree. (d) Write $\mathbf{G}$ for this tree at $\hat{\mathbf{z}}$, the gradient of $\ln S_1$ by equation (14.45), and its variance budget with $\boldsymbol{\Sigma}$ of equation (14.10); which column carries the band? ((a) The four CRAC triples, {P1, P2} and {K}. (b) $S_1 = 4 \times 0.03^3 + 0.01^2 + 0.001 = 1.208 \times 10^{-3}$; exact $1 - (1 - 4q_C^3(1 - q_C) - q_C^4)(1 - q_P^2)(1 - q_K) = 1.205 \times 10^{-3}$. (c) $\Pr(4q_C^3/(4q_C^3 + q_P^2 + q_K) > \frac{1}{2} \mid \mathbf{y})$; at the plug-in the triples carry 9 % and the controller 83 %. (d) Each triple's row is $(3(1 - q_C)q_C^3,\ 0,\ 0) = (7.86 \times 10^{-5},\ 0,\ 0)$, the pumps' $(0,\ 1.98 \times 10^{-4},\ 0)$ and the controller's $(0,\ 0,\ 9.99 \times 10^{-4})$; $\mathbf{G}^{\mathsf{T}}\mathbf{1}/S_1 = (0.260,\ 0.164,\ 0.827)$ and $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}/S_1^2 = 0.023 + 0.007 + 0.685 = 0.715$. The controller's column carries 96 % of the band, where in the hall's tree the CRAC column carries 97 %.)

Exercise 14.3. *Derivation.* Two minimal cut sets share a unit, $c_1 = \{1, 2\}$ and $c_2 = \{2, 3\}$, with independent unavailabilities q_1, q_2, q_3. (a) Derive the probability of the top event by inclusion–exclusion. (b) Show that the rare-event sum is an upper bound and give its error and its relative error. (c) Evaluate the relative error for $q_1 = q_3 = 0.03$. ((a) $E_{c_1} \cap E_{c_2} = E_{\{1,2,3\}}$, so $\Pr = q_1 q_2 + q_2 q_3 - q_1 q_2 q_3$. (b) $S_1 - \Pr = q_1 q_2 q_3 \geq 0$; relative error $q_1 q_3/(q_1 + q_3 - q_1 q_3)$, free of q_2. (c) $0.0009/0.0591 = 1.5\,\%$.)

Exercise 14.4. *Derivation.* (a) Show that $\mathbb{E}[q^2 \mid \mathbf{y}] \geq \mathbb{E}[q \mid \mathbf{y}]^2$ under any posterior, with equality only when the posterior is a point. (b) With the moments of a Beta posterior from Chapter 13, derive the ratio $\mathbb{E}[q^2]/\mathbb{E}[q]^2$ for $q \sim \text{Beta}(a, b)$ and its limit when b is large. (c) Compute the gap for Beta(3, 97) and its effect on the sum of the six CRAC pairs. ((a) $\mathbb{E}[q^2] - \mathbb{E}[q]^2 = \text{Var}(q \mid \mathbf{y}) \geq 0$, zero only for a point mass. (b) $\mathbb{E}[q^2] = a(a+1)/((a+b)(a+b+1))$ and $\mathbb{E}[q]^2 = a^2/(a+b)^2$, so the ratio is $(a+1)(a+b)/(a(a+b+1))$, which tends to $(a+1)/a$: the gap is set by the number of failures a, not by the number of checks. (c) 1.188×10^{-3} against 9.00×10^{-4}, gap 2.88×10^{-4}, ratio 1.320, against the limit 4/3; for the six pairs 7.13×10^{-3} against 5.40×10^{-3}, a difference of 1.7×10^{-3}, larger than the controller's whole cut set.)

Exercise 14.5. *Application.* A supply is lost when its normal feed F is down together with either of two transfer switches A and B, each of which alone would leave the supply without its alternate path: the top event is $(F \wedge A) \vee (F \wedge B)$. (a) Give the minimal cut sets. (b) With $q_F = 0.02$ and $q_A = q_B = 0.2$, compute the rare-event sum and the exact probability, and the error of treating the two cut sets as disjoint. (c) State why the error is large here and small for the hall. ((a) $\{F, A\}$ and $\{F, B\}$, which share F. (b) $S_1 = 2 \times 0.02 \times 0.2 = 0.0080$; exact $q_F(q_A + q_B - q_A q_B) = 0.02 \times 0.36 = 0.0072$; the sum is 11 % high, by $q_F q_A q_B = 0.0008$. (c) Given F down, the two cut sets need only A or B, whose unavailabilities are large, so their joint occurrence is a fifth of either; in the hall the overlaps need a further unit with unavailability 0.03 or less. Factoring the tree as $F \wedge (A \vee B)$ gives the exact value directly.)

Exercise 14.6. *Application.* Fix the pumps' and the controller's unavailabilities at their posterior means, $q_P = 0.01$ and $q_K = 0.001$, and keep $q_C \sim \text{Beta}(3, 97)$. (a) Show that $\pi_{\text{C}} > \frac{1}{2}$ exactly when q_C exceeds a threshold, and compute $\Pr(\pi_{\text{C}} > \frac{1}{2} \mid \mathbf{y})$. (b) For $N = 100$ further CRAC checks, find the numbers n_{d} of units found down after which the posterior probability is at least 0.95, and the predictive probability of settling the question that way. (c) Explain why the expected posterior probability after the checks equals the present one. ((a) $6q_C^2 > q_P^2 + q_K =$

1.1×10^{-3}, so $q_C > q_C^{\dagger} = \sqrt{1.1 \times 10^{-3}/6} = 0.01354$, and $\Pr = 1 - I_{0.01354}(3, 97) = 0.849$. (b) After n_{d} down in 100, $\Pr = 1 - I_{0.01354}(3 + n_{\mathrm{d}}, 197 - n_{\mathrm{d}})$: 0.494, 0.716, 0.865, 0.945, 0.980 for $n_{\mathrm{d}} = 0$ to 4, so $n_{\mathrm{d}} \geq 4$ settles it, with $\Pr(n_{\mathrm{d}} \geq 4) = 0.346$ under the Beta–binomial$(100, 3, 97)$; no outcome brings it below 0.05. (c) By the law of total probability, equation (14.48), the predictive average of the posterior probability is the prior one, 0.849; a record changes the expected width of the uncertainty, not its expected center.)

Chapter 15

Dependent Failures and Common Causes

15.1 Dependent failures and the common-cause share

Two chilled-water pumps feed the cooling coils of a data hall, and either can carry the full flow alone. The facilities engineer must show that the hall loses its chilled-water supply, both pumps down at once, less than 10^{-4} of the time. The failure log of five years shows nine failures of one pump while the other kept running, and one event in which both pumps stopped within the same shift. The first reading of the log is simple: eleven pump failures in 87 600 pump-hours is a rate of $1.26 \times 10^{-4}\,\mathrm{h}^{-1}$ per pump; with repairs of twelve hours each pump is down 0.15 % of the time, and two independent pumps are down together $(1.50 \times 10^{-3})^2 = 2.3 \times 10^{-6}$ of the time, about a minute a year and forty-four times below the limit, so the pair passes.

The engineer who must sign the cooling design asks something else. The product of two small numbers holds only if the pumps fail independently, and the log contains one event that the product can explain only as a coincidence: two independent failures that happened to fall in the same shift. Whether that event was a coincidence or one cause acting on both pumps, how large a share of the pumps' failures such a cause carries, and what that share does to the pair are statements about unknowns that the single number 2.3×10^{-6} does not contain. The question of this chapter is

how a shared cause of failure is written as one more unknown of the plant, how its share of the failures is estimated from a record of single and simultaneous failures, whether the record supports a shared cause at all against a coincidence of independent failures, whether one observed double failure was a common cause, and what the shared cause does to the unavailability of a redundant pair, with its band.

It is answered with the method of the previous chapters on the same pumps. What changes is the state, which carries a second entry, the logit of the share of each pump's failures that comes from a shared cause, informed by a count row that acts on both entries; and the plant has a second drawing, without the shared cause, that competes with the first. The first reading, 2.3×10^{-6} against a limit of 10^{-4}, looks like a comfortable pass.

Jargon: dependent failure, common-cause failure, β-factor model, common-cause share, superposition, thinning, coincidence window, coincident failures, staggered maintenance, α-factor model, multiple Greek letter model

A *dependent failure* is a failure whose probability depends on the state of another unit; a *common-cause failure* is the failure of two or more units from one shared cause within a short time. The *β-factor model* represents each unit's failures as the sum of an independent part and a common-cause part that fails every unit of the group at once; the *common-cause share* β, written β_{cc} in this book, is the fraction of a unit's failure rate carried by the common-cause part. The *superposition* of point processes is the process of all their events together; *thinning* is the reverse operation, the assignment of each event of a process to one of several sub-processes at random. The *coincidence window* t_{w} is the interval within which two failures recorded on different units cannot be told apart from one simultaneous event; two independent failures that fall within it are *coincident failures*. *Staggered maintenance* schedules the maintenance of redundant units on different days, so that one error is not repeated on both. The *α-factor model* and the *multiple Greek letter model* generalize the β-factor model to groups of more than two units by giving separate shares to failures of two, three or more units together.

Key idea

A shared cause of failure is one more free entry of the state, the common-cause share β_{cc}, informed by one more kind of count row, the simultaneous failures, which acts on both entries of the state. The record compares two drawings of the plant, with and without the shared cause, through their evidences, and the share enters a redundant pair's unavailability at first order, $\beta_{\mathrm{cc}}\lambda/\mu$, where independent failures enter at second order. A share informed by one double failure is poorly known and its posterior is skewed, so every answer is computed from the posterior of the state with the exact value beside it and the gap stated, and the choice between waiting for more record and acting on the share is made by the preposterior.

15.1.1 Independence given the parameters and where it fails

Chapter 13 computed the availability of a service from the 2^n joint states of its units and showed that units which share a posterior are already coupled, even when they fail independently given their parameters, because an unknown common parameter makes their states correlated a posteriori. Chapter 14 read the same enumeration as a list of minimal cut sets and carried the two chilled-water pumps P1 and P2 of the data hall as the cut set $\{\mathrm{P1}, \mathrm{P2}\}$, whose probability was the expectation of q_P^2 under the shared posterior $q_P \sim \mathrm{Beta}(4, 396)$, the two pumps failing independently given q_P.

The coupling of this chapter is of a different kind. It is not a statistical consequence of an unknown parameter but a physical mechanism: one event that takes both units out. No amount of knowledge of the pumps' individual rates removes it. Three examples recur in practice and are named here only as examples, since the model below does not distinguish them: a contaminated or air-bound water supply that damages both pumps drawing from it; a shared control signal or power feed that stops both; and a maintenance error, such as a wrong lubricant or a misset protection relay, repeated on both units. The companion volume *Reliability Engineering* treats the classes of dependent failure, the defenses against them and the parametric models for groups of any size; this chapter derives only the β-factor model of a pair, and derives it as an inference.

15.1.2 A shared cause as one more entry of the state

In the language of the method, the shared cause is a latent source node of the plant graph that feeds both pumps. Its rate is a fraction β_{cc} of each pump's total failure rate λ, and the fraction

is unknown: it is a free entry of the state beside the rate. The evidence acquires a row it did not have before, the count of simultaneous failures, and the count of single failures changes its meaning, since it now counts only the independent part. Whether a shared cause exists at all is then a comparison between two drawings, the graph with the shared node and the graph without it; whether the one observed double failure was a common cause is a posterior over two explanations of one event; and the pair's unavailability is a derived quantity of the two entries. Four of the five categories of Chapter 2 have a question here, which Table 15.2 lists; the fifth, design, is the next measurement.

15.1.3 The physical system and its graph (step 0)

Two identical chilled-water pumps take water from the chilled-water source and deliver it to the distribution header that feeds the hall's cooling coils. Both run continuously, and either can carry the full flow alone: the arrangement is one-out-of-two, and the cooling supply is lost only when both pumps are down at once. Each pump fails alone, from causes that act on it only, at the rate $(1-\beta_{\mathrm{cc}})\lambda$, and both fail together, from a shared cause, at the rate $\beta_{\mathrm{cc}}\lambda$. A failed pump is repaired and returned to service with the repair rate μ; after a common-cause failure both pumps are restored together. The model assumes that the two pumps have the same rates, that the rates are constant over the record, and that the shared cause, when it acts, takes out both pumps within the coincidence window.

Figure 15.1 draws the plant graph. In panel (a) the gray boxes are the chilled-water source and the distribution header, pinned boundaries of the hydraulic path, and the solid green arrows are the two parallel paths through the pumps. The two green boxes are the pumps, free nodes, each failing alone at its independent rate; the box with the red border between them is the shared cause, a latent source node whose red arrows take both pumps out at once. Panel (b) draws the evidence side of the same graph: the two green boxes are the two entries of the state, the log failure rate z_1 and the logit of the share z_2; the two white boxes in the middle are the count rows of the record, the single failures and the double failures, each reached by dashed arrows from both entries because each rate depends on both; and the two white boxes below are the prior rows.

The state entries. The rate must be positive and the share must lie in $(0,1)$, and the graph carries each on the scale that removes its bound,

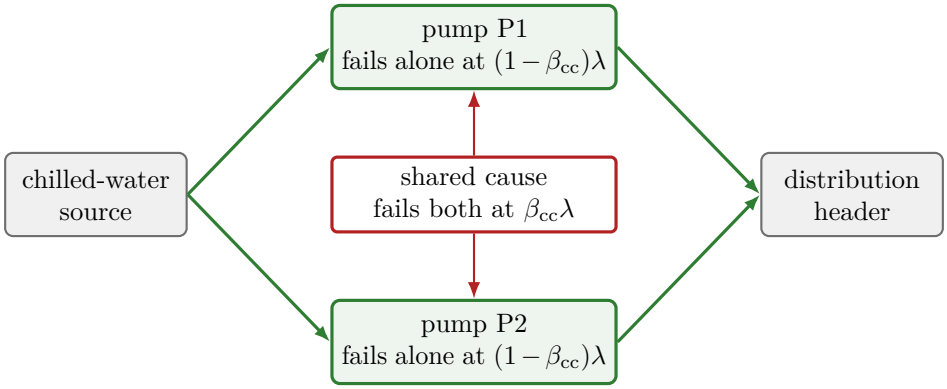
$$\mathbf{z} = \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} = \begin{bmatrix} \ln\lambda \\ \xi \end{bmatrix}, \qquad \xi = \ln\frac{\beta_{\mathrm{cc}}}{1-\beta_{\mathrm{cc}}}, \tag{15.1}$$

with λ in h^{-1}, the log rate of Chapter 8 and the logit of Chapter 5. The inverse maps are $\lambda = \mathrm{e}^{z_1}$ and $\beta_{\mathrm{cc}} = 1/(1+\mathrm{e}^{-\xi})$, with the derivatives $\mathrm{d}\lambda/\mathrm{d}z_1 = \lambda$ and $\mathrm{d}\beta_{\mathrm{cc}}/\mathrm{d}\xi = \beta_{\mathrm{cc}}(1-\beta_{\mathrm{cc}})$ that every change of variable below uses. The share is written β_{cc} rather than β because the book reserves β for the Weibull shape. The repair rate is pinned at $\mu = 1/12\,\mathrm{h}^{-1}$, a mean time to repair of 12 h, the posterior mean repair time of the pump of Chapter 8; its uncertainty is left out so that the chapter isolates the shared cause.

15.1.4 The record, the coincidence window and the priors

Table 15.1 lists the evidence: each row is one item, the middle column its value and the right column the row of the graph it feeds. The record covers five years of continuous operation, $T_{\mathrm{tot}} = 43\,800\,\mathrm{h}$, which is the exposure of each pump and of the shared process; the repair time of the ten outages, of about twelve hours each, a fraction 0.3 % of the record, is not subtracted. In that time the log shows $k_1 = 9$ failures of one pump while the other kept running, counted over both pumps, and $k_2 = 1$ event in which both pumps failed within the coincidence window. The

(a) plant graph of the pump pair



(b) state entries, count rows and priors

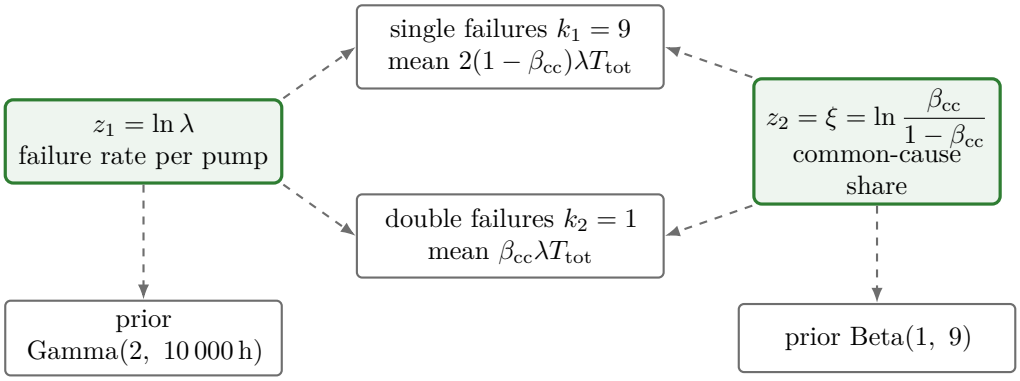


Figure 15.1: The chilled-water pump pair. (a) Plant graph: pinned, the chilled-water source and the distribution header (gray); free, the two pumps, each failing alone at $(1-\beta_{cc})\lambda$ (green), and the latent shared cause, which fails both at $\beta_{cc}\lambda$ (red border and arrows); solid green arrows are the hydraulic paths. (b) The two entries of the state, the count rows of the record of $T_{tot} = 43\,800$ h (dashed arrows from both entries to each row) and the two prior rows.

record is illustrative. It is a different record of the same pumps from the one of Chapter 14: there the unavailability $q_P \sim \text{Beta}(4, 396)$ came from availability checks, which find a pump down for any reason, planned work included; here the rate comes from the failure log, with a repair of twelve hours after each failure. The two records are not combined, and every comparison of this chapter between the shared cause and independence is made within the failure log.

Table 15.1: The record of the chilled-water pump pair and the priors. Counts are events: a double failure is one event in which both pumps fail.

Item	Value	Row of the graph
Record length (both pumps running)	$T_{\text{tot}} = 43\,800$ h (five years)	exposure of every count row
Single-pump failures, either pump	$k_1 = 9$	singles row, mean $2(1-\beta_{\text{cc}})\lambda T_{\text{tot}}$
Double failures, logged times less than t_{w} apart	$k_2 = 1$	doubles row, mean $\beta_{\text{cc}}\lambda T_{\text{tot}}$
Coincidence window (time-stamped log)	$t_{\text{w}} = 8$ h	defines a double; coincidence rate $2\lambda^2 t_{\text{w}}$ under independence
Mean time to repair	12 h, pinned	repair rate $\mu = 1/12\,\text{h}^{-1}$
Prior on the failure rate	Gamma(2, 10 000 h)	prior row on z_1
Prior on the share	Beta(1, 9)	prior row on z_2
Limit on the pair's unavailability	10^{-4} at assurance $\kappa = 0.90$	the threshold question

The coincidence window. A double failure is defined by the record before the record is read. In this chapter the log carries a time stamp for every failure, and a double is a pair of failures, one on each pump, whose logged times t_1 and t_2 satisfy $|t_1 - t_2| < t_{\text{w}}$. The window is at least the resolution of the time stamps, since two failures the log cannot separate must count as one event, and at least the spread with which a shared cause acts, a contaminated supply damaging the second pump some hours after the first. The base case is $t_{\text{w}} = 8$ h, one shift, the time within which a supply fault or an error repeated in one maintenance visit reaches both pumps, with 1 h and 24 h carried as alternatives. The window must not be chosen after looking at the separation of the observed double, since the evidence against independence depends on it directly (Section 15.4.4).

Overlapping outages. A failure of the second pump while the first is under repair, but more than t_{w} after the first failure, is not a double: it is two singles, each pump's own failure, and the time both are down is the term q_{ind}^2 of the pair's unavailability (Section 15.3.4), not the shared term. The record shows none.

The priors. The failure rate per pump has the weak handbook prior $\lambda \sim \text{Gamma}(a_{\lambda 0}, b_{\lambda 0}) = \text{Gamma}(2,\ 10\,000\,\text{h})$, of mean $2 \times 10^{-4}\,\text{h}^{-1}$, written as throughout the book with a shape and a rate parameter, the subscript 0 marking a prior's. The share has the generic prior $\beta_{\text{cc}} \sim \text{Beta}(a_{\beta 0}, b_{\beta 0}) = \text{Beta}(1, 9)$, of mean 0.1, which weighs as much as one common-cause event in ten. The two priors are independent.

15.1.5 The probabilistic questions (step 1)

The operator asks five questions of the record, and step 1 writes each as a probability or an expectation over the state, with the unknown on the left of the bar and the record on the right. Table 15.2 lists them: the first column is the question in engineering language; the second the general form that every question of its category takes, in the notation of Chapter 2 for a

quantity $f(\mathbf{z})$ derived from the state, a level f^* of it, an assurance level κ, the posterior density $p(\mathbf{z} \mid \mathbf{y})$ and the drawings M_j; the third its mathematical form for the pump pair; the fourth its category; and the last the subsection where it is answered. The first is an estimation of a state entry, carried as its logit and reported as the share. The second is a hypothesis comparison between the drawing with the shared cause, M_1, and the drawing of independence, M_0, in which a double failure can only be a coincidence within t_{w}. The third is the diagnosis of one event, a posterior over its possible explanations, whether one cause took the two units out or two independent failures met in the window. The fourth is the estimation of a derived quantity, the pair's unavailability U_{pair}, against the two baselines that independence offers, and the fifth a prediction under a threshold on the same quantity. A sixth question, which helps more, halving the share by staggered maintenance or waiting for more record to pin it, asks nothing new of the present record: it is the next measurement for these five, set against an intervention, and Section 15.6 answers it. The atlas of Appendix A lists the same forms.

Table 15.2: The questions asked of the pump pair's record, the general form of their category, their mathematical form for the pair, their category and where each is answered. $\mathbf{y}$ is the record $(k_1, k_2, T_{\mathrm{tot}})$ with the two priors; M_1 the drawing with the shared cause and M_0 the drawing of independence; $q = \lambda/(\lambda+\mu)$ the unavailability of one pump at its total rate and q_{ind} at its independent rate.

Engineering question	General form	Mathematical form	Category	Answered in
What share of the pumps' failures is a common cause, and how well is it known?	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$, quantiles	$p(\beta_{\mathrm{cc}} \mid \mathbf{y})$, $\beta_{\mathrm{cc}} = 1/(1+\mathrm{e}^{-\xi})$	estimation	15.3.1
Is there a common cause at all, or was the double a coincidence?	$B_{AB} = p(\mathbf{y} \mid M_A)/p(\mathbf{y} \mid M_B)$	$B_{10} = p(\mathbf{y} \mid M_1)/p(\mathbf{y} \mid M_0)$, coincidences within t_{w} under M_0; $\Pr(M_1 \mid \mathbf{y})$	hypothesis comparison	15.3.2
Was the double a common cause, so that one cause took two units out?	$\Pr(M_j \mid \mathbf{y}) \propto p(\mathbf{y} \mid M_j)\ \Pr(M_j)$	$\Pr(\text{common cause} \mid \text{the double}, \mathbf{y})$, over both drawings and within M_1	diagnosis of one event	15.3.3
What is the pair's unavailability with the common cause, and what would it be without?	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$, quantiles	$\mathbb{E}[U_{\mathrm{pair}} \mid \mathbf{y}]$; $\mathbb{E}[q_{\mathrm{ind}}^2 \mid \mathbf{y}]$ and $\mathbb{E}[q^2 \mid \mathbf{y}]$	estimation of a derived quantity	15.3.4
Is the pair's unavailability below 10^{-4}, with assurance 0.90?	$\Pr(f(\mathbf{z}) \ge f^* \mid \mathbf{y}) \ge \kappa$	$\Pr(U_{\mathrm{pair}} < 10^{-4} \mid \mathbf{y}) \ge 0.90$	prediction under a threshold	15.3.4

15.1.6 The evidence as rows (step 2)

Step 2 attaches four rows to the graph of the shared cause. The singles row and the doubles row are Poisson count rows, as in Chapter 6, and the two prior rows carry the gamma prior on the rate and the Beta prior on the share. Unlike the rows of Chapter 8, the two count rows each act on both entries, since both means contain λ and β_{cc}; the prior rows act on one entry each. Which row acts on which entry is a property of the graph, and it fixes where the Jacobian

of the rows can be nonzero, marked $*$:

$$\mathbf{J} = \begin{bmatrix} * & 0 \\ * & * \\ * & * \\ 0 & * \end{bmatrix} \begin{matrix} \text{prior on } \lambda \\ \text{singles row} \\ \text{doubles row,} \\ \text{prior on } \beta_{\text{cc}} \end{matrix} \qquad \mathbf{J}_{M_0} = \begin{bmatrix} * \\ * \\ * \end{bmatrix} \begin{matrix} \text{prior on } \lambda \\ \text{singles row} \\ \text{coincidence row} \end{matrix}, \qquad \text{columns } z_1, \xi. \tag{15.2}$$

The singles and doubles rows share both columns, so the information matrix $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ has an off-diagonal entry and the two entries of the state are correlated a posteriori, a coupling that the factorized records of Chapter 8 did not have. The second question needs a second graph, the same plant without the shared-cause node, whose state has the single entry z_1 and whose doubles row carries the coincidence rate of two independent failures within the window; its three rows all act on that one entry, the right-hand pattern. None of the rows is Gaussian; each enters the objective as its negative log-probability term, equation (4.39), and joins the Jacobian in the working form of Section 4.2.7.

The modeling choices. Four choices make the rows, and each buys something and can fail. The two pumps are taken to share one total failure rate λ and one share β_{cc}, both constant over the five years: this is what lets the record be summarized by k_1, k_2 and T_{tot} alone, and it fails if one pump fails more often than the other or if the rates drift, which the pumps' own counts and the trend tests of Chapter 12 would show (Section 15.6.7). The shared cause is taken to fail both pumps whenever it acts, the β-factor mechanism: this gives the doubles row its mean $\beta_{\text{cc}}\lambda T_{\text{tot}}$, and it fails for a cause that stops one pump and only degrades the other, a mechanism outside the model. The handbook knowledge is carried by the priors Gamma(2, 10 000 h) and Beta$(1, 9)$: they are weak on the rate and worth one common-cause event in ten on the share, and with one double in the record the prior on the share still carries a large part of its information, which Section 15.4.5 measures. And the doubles row of M_1 counts the shared events only, leaving out the rare coincidences of the two independent parts: this makes the share and the pair's unavailability independent of the window, and it is a small omission at practical windows, measured in Section 15.4.4.

The failure processes and the rows of the record. The β-factor model of the pair is three independent homogeneous Poisson processes in the sense of Chapter 6: the independent failures of pump P1, of rate $(1-\beta_{\text{cc}})\lambda$; the independent failures of pump P2, of the same rate; and the shared process, of rate $\beta_{\text{cc}}\lambda$, each of whose events fails both pumps. Let $N_1(t)$, $N_2(t)$ and $N_c(t)$ be their counts in t hours, Poisson with means $(1-\beta_{\text{cc}})\lambda t$, $(1-\beta_{\text{cc}})\lambda t$ and $\beta_{\text{cc}}\lambda t$.

Superposition. The sum $S = X + Y$ of two independent Poisson counts with means a and b is Poisson with mean $a + b$. By the law of total probability over the value of X,

$$\Pr(S = s) = \sum_{j=0}^{s} \frac{a^j \mathrm{e}^{-a}}{j!} \frac{b^{s-j} \mathrm{e}^{-b}}{(s-j)!} = \frac{\mathrm{e}^{-(a+b)}}{s!} \sum_{j=0}^{s} \binom{s}{j} a^j b^{s-j} = \frac{(a+b)^s \, \mathrm{e}^{-(a+b)}}{s!},$$

by the binomial theorem. Because the processes are independent on every interval, the same holds for counts on disjoint intervals, and the superposition of independent Poisson processes is a Poisson process whose rate is the sum of the rates. Three consequences follow. The failures of pump P1 from either cause, $N_1 + N_c$, form a Poisson process of rate $(1-\beta_{\text{cc}})\lambda + \beta_{\text{cc}}\lambda = \lambda$, so λ is the total failure rate of one pump, the quantity a handbook lists, and β_{cc} the fraction of it that comes from the shared cause. The single failures of either pump, $N_1 + N_2$, form a Poisson process of rate $2(1-\beta_{\text{cc}})\lambda$. And all failure events of the pair, singles and doubles together, form a Poisson process of rate

$$\nu = 2(1-\beta_{\text{cc}})\lambda + \beta_{\text{cc}}\lambda = (2-\beta_{\text{cc}})\lambda. \tag{15.3}$$

Thinning. The same model can be built in the other direction: take the stream of all failure events of the pair, of rate ν, and mark each event independently as a double with probability θ and as a single with probability $1-\theta$. Given m events in T_{tot}, the number of doubles is binomial, so the joint probability of k_1 singles and k_2 doubles, with $m = k_1 + k_2$, is

$$\begin{aligned}\Pr(k_1, k_2) &= \frac{(\nu T_{\text{tot}})^m \mathrm{e}^{-\nu T_{\text{tot}}}}{m!} \binom{m}{k_2} \theta^{k_2} (1-\theta)^{k_1} \\ &= \frac{\big((1-\theta)\nu T_{\text{tot}}\big)^{k_1} \mathrm{e}^{-(1-\theta)\nu T_{\text{tot}}}}{k_1!} \times \frac{(\theta \nu T_{\text{tot}})^{k_2} \mathrm{e}^{-\theta \nu T_{\text{tot}}}}{k_2!},\end{aligned}$$

where the second line splits $\mathrm{e}^{-\nu T_{\text{tot}}} = \mathrm{e}^{-(1-\theta)\nu T_{\text{tot}}} \mathrm{e}^{-\theta \nu T_{\text{tot}}}$ and $\nu^m = \nu^{k_1} \nu^{k_2}$. The joint probability is the product of two Poisson probabilities: an independently thinned Poisson process splits into independent Poisson processes with rates $(1-\theta)\nu$ and $\theta\nu$. Matching $\theta\nu$ to the shared rate $\beta_{\text{cc}}\lambda$ identifies the share of failure *events* that are doubles,

$$\theta = \frac{\beta_{\text{cc}}\lambda}{(2-\beta_{\text{cc}})\lambda} = \frac{\beta_{\text{cc}}}{2-\beta_{\text{cc}}}, \tag{15.4}$$

which is not β_{cc}: a single failure involves one pump and a double two, so per event the doubles are about half as frequent as the share of each pump's rate suggests. The distinction returns in Section 15.2.1.

The correlation of the two pumps. The failure counts of the two pumps share the term N_c: $\mathrm{Cov}(N_1 + N_c,\ N_2 + N_c) = \mathrm{Var}(N_c) = \beta_{\text{cc}}\lambda t$, while each has variance λt, so the correlation of the two pumps' failure counts is exactly β_{cc}. The share is therefore also the correlation the shared cause induces between the units, a second reading of the same entry.

The likelihood of the two counts. By either construction, the singles and the doubles in the exposure T_{tot} are independent Poisson counts,

$$k_1 \mid \lambda, \beta_{\text{cc}} \sim \text{Poisson}\big(2(1-\beta_{\text{cc}})\lambda T_{\text{tot}}\big), \qquad k_2 \mid \lambda, \beta_{\text{cc}} \sim \text{Poisson}(\beta_{\text{cc}}\lambda T_{\text{tot}}), \tag{15.5}$$

and the likelihood is the product of the two Poisson probabilities,

$$p(k_1, k_2 \mid \lambda, \beta_{\text{cc}}) = \frac{(2T_{\text{tot}})^{k_1} T_{\text{tot}}^{k_2}}{k_1!\, k_2!}\ (1-\beta_{\text{cc}})^{k_1}\ \beta_{\text{cc}}^{k_2}\ \lambda^m\, \mathrm{e}^{-(2-\beta_{\text{cc}})\lambda T_{\text{tot}}}, \qquad m = k_1 + k_2. \tag{15.6}$$

The rate enters through λ^m, the number of events, and through the exponent with the total event rate $(2-\beta_{\text{cc}})\lambda$; the share enters through the binomial-like factor $(1-\beta_{\text{cc}})^{k_1}\beta_{\text{cc}}^{k_2}$ and again through the exponent. The two entries do not separate, which is the structural difference from the factorized record of Chapter 8. In the state (15.1) the logarithms of the two means are linear in z_1 and not in ξ, and their derivatives form the 2×2 matrix

$$\frac{\partial}{\partial \mathbf{z}^{\mathsf{T}}} \begin{bmatrix} \ln\big(2(1-\beta_{\text{cc}})\lambda T_{\text{tot}}\big) \\ \ln(\beta_{\text{cc}}\lambda T_{\text{tot}}) \end{bmatrix} = \begin{bmatrix} 1 & -\beta_{\text{cc}} \\ 1 & 1-\beta_{\text{cc}} \end{bmatrix} \begin{matrix} \text{singles} \\ \text{doubles'} \end{matrix}$$

whose first column is z_1 and second ξ. Raising the log rate raises both expected counts by the same factor; raising the logit of the share moves expected failures from the singles to the doubles, with entries of opposite sign. The determinant is $(1-\beta_{\text{cc}}) + \beta_{\text{cc}} = 1$ at every state, so the two count rows alone determine both entries: with as many rows as entries, the maximum-likelihood state is the one at which both rows expect exactly the counts observed, $2(1-\beta_{\text{cc}})\lambda T_{\text{tot}} = k_1$ and $\beta_{\text{cc}}\lambda T_{\text{tot}} = k_2$, that is

$$\beta_{\text{cc}}^{\text{ML}} = \frac{2k_2}{k_1 + 2k_2}, \qquad \lambda^{\text{ML}} = \frac{k_2}{\beta_{\text{cc}}^{\text{ML}} T_{\text{tot}}},$$

with the superscript ML marking the maximum-likelihood values, apart from the MAP state with its hats. The thinning form gives the same with the total rate $\nu^{\mathrm{ML}} = m/T_{\mathrm{tot}}$ and the event share $\theta^{\mathrm{ML}} = k_2/m$ in equation (15.4).

The independence drawing and the coincidence rate. Under the drawing M_0 the plant graph has no shared-cause node: each pump fails independently at the rate λ, and a double failure is recorded only when the two independent failures fall within t_{w} of each other in the sense of Section 15.1.4, $|t_1 - t_2| < t_{\mathrm{w}}$ on time-stamped records. The rate of such coincidences follows from the Poisson process. A failure of pump P1 occurs at rate λ; given that it occurs at time t, pump P2 fails in the window $(t - t_{\mathrm{w}},\ t + t_{\mathrm{w}})$ with probability $1 - \mathrm{e}^{-2\lambda t_{\mathrm{w}}} \approx 2\lambda t_{\mathrm{w}}$, since its failures are a Poisson process of rate λ independent of P1's. Each coincident pair is counted once this way, from its P1 failure, so the coincidence rate is

$$c_{\mathrm{w}} = \lambda \times 2\lambda t_{\mathrm{w}} = 2\lambda^2 t_{\mathrm{w}}, \tag{15.7}$$

with the same result if the pair is counted from the first failure of either pump, at rate 2λ, followed by a failure of the other within t_{w} after it, with probability λt_{w}. Some logs record only the shift or the day of a failure, and a double is then two failures in the same bin of length t_{w}: given a failure of P1 in a bin, P2 fails in the same bin with probability λt_{w} on average over the position of the first failure, so the rate is $\lambda^2 t_{\mathrm{w}}$, and a binned log of bin t_{w} is treated as a time-stamped log with window $t_{\mathrm{w}}/2$. The approximation $1 - \mathrm{e}^{-2\lambda t_{\mathrm{w}}} \approx 2\lambda t_{\mathrm{w}}$ needs $\lambda t_{\mathrm{w}} \ll 1$. The derivation also assumes that the second pump is running, and so able to fail, when the first fails; it is under repair with probability q, and the factor $1 - q$ on the coincidence rate is taken as one.

Under M_0 the record is therefore a singles count and a coincidence count, independent to the same order,

$$k_1 \mid \lambda, M_0 \sim \mathrm{Poisson}(2\lambda T_{\mathrm{tot}}), \qquad k_2 \mid \lambda, M_0 \sim \mathrm{Poisson}(2\lambda^2 t_{\mathrm{w}} T_{\mathrm{tot}}), \tag{15.8}$$

the two failures of a coincidence being removed from the singles, a correction of relative order λt_{w}. The state of M_0 has the single entry z_1, with the same gamma prior. The same coincidences can also occur under M_1, between the independent parts of the two pumps, at the rate $(1 - \beta_{\mathrm{cc}})^2 c_{\mathrm{w}}$; the chapter's model of M_1 in equation (15.5) leaves them out, since they are much smaller than the shared rate $\beta_{\mathrm{cc}}\lambda$ at practical windows, and Sections 15.3.3 and 15.4.4 measure what including them changes.

The rows. The posterior is carried by the state $\mathbf{z} = (z_1, \xi)^{\mathsf{T}}$ of equation (15.1), and none of its four rows is Gaussian: each is the negative log-probability term of its record, written in the state, and each takes the working form of Section 4.2.7. The gamma prior on λ becomes, with the change-of-variable factor $\mathrm{d}\lambda/\mathrm{d}z_1 = \lambda$, the term $b_{\lambda 0}\mathrm{e}^{z_1} - a_{\lambda 0} z_1$, a count row of $a_{\lambda 0} = 2$ pseudo-failures in $b_{\lambda 0} = 10\,000$ pseudo-hours, as in Chapter 8. The Beta prior on β_{cc} becomes, with the factor $\mathrm{d}\beta_{\mathrm{cc}}/\mathrm{d}\xi = \beta_{\mathrm{cc}}(1 - \beta_{\mathrm{cc}})$, a density proportional to $\beta_{\mathrm{cc}}^{a_{\beta 0}}(1 - \beta_{\mathrm{cc}})^{b_{\beta 0}}$, the term $-a_{\beta 0}\ln\beta_{\mathrm{cc}} - b_{\beta 0}\ln(1 - \beta_{\mathrm{cc}})$, with the shapes rather than the shapes less one: a binomial row of $a_{\beta 0} = 1$ common-cause events in $a_{\beta 0} + b_{\beta 0} = 10$ pseudo-trials, as in Chapter 5. The singles and doubles rows are the negative logs of the Poisson probabilities (15.5) without their constants. Writing $\bar{k}_0 = b_{\lambda 0}\lambda$, $\bar{k}_1 = 2(1 - \beta_{\mathrm{cc}})\lambda T_{\mathrm{tot}}$ and $\bar{k}_2 = \beta_{\mathrm{cc}}\lambda T_{\mathrm{tot}}$ for the counts the three count rows expect at the state, the four terms are

$$\begin{bmatrix} \mathcal{L}_0 \\ \mathcal{L}_1 \\ \mathcal{L}_2 \\ \mathcal{L}_\beta \end{bmatrix} = \begin{bmatrix} \bar{k}_0 - a_{\lambda 0}\ln\bar{k}_0 \\ \bar{k}_1 - k_1\ln\bar{k}_1 \\ \bar{k}_2 - k_2\ln\bar{k}_2 \\ -a_{\beta 0}\ln\beta_{\mathrm{cc}} - b_{\beta 0}\ln(1 - \beta_{\mathrm{cc}}) \end{bmatrix} \begin{matrix} \text{prior on } \lambda \\ \text{singles row} \\ \text{doubles row'} \\ \text{prior on } \beta_{\mathrm{cc}} \end{matrix}$$

up to constants, and their sum, with the terms collected, is the objective

$$\begin{aligned}\mathcal{L}(\mathbf{z}) = {} & \big(b_{\lambda0} + (2-\beta_{\mathrm{cc}})T_{\mathrm{tot}}\big)\mathrm{e}^{z_1} - (a_{\lambda0}+m)\,z_1 \\ & - (a_{\beta0}+k_2)\ln\beta_{\mathrm{cc}} - (b_{\beta0}+k_1)\ln(1-\beta_{\mathrm{cc}}), \qquad \beta_{\mathrm{cc}} = \frac{1}{1+\mathrm{e}^{-\xi}},\end{aligned} \tag{15.9}$$

with $a_{\lambda0}+m = 12$, $a_{\beta0}+k_2 = 2$ and $b_{\beta0}+k_1 = 18$.

The working residuals and the Jacobian. The three count rows have the log-count Jacobian of Section 15.1.6 with the prior row added,

$$\frac{\partial \ln \bar{\mathbf{k}}}{\partial \mathbf{z}^{\mathsf{T}}} = \begin{bmatrix} 1 & 0 \\ 1 & -\beta_{\mathrm{cc}} \\ 1 & 1-\beta_{\mathrm{cc}} \end{bmatrix} \begin{matrix} \text{prior on } \lambda \\ \text{singles row} \\ \text{doubles row} \end{matrix},$$

and the binomial prior row on the share has the expected count $(a_{\beta0}+b_{\beta0})\beta_{\mathrm{cc}}$, whose derivative in ξ is $(a_{\beta0}+b_{\beta0})\beta_{\mathrm{cc}}(1-\beta_{\mathrm{cc}})$. The working residuals and Jacobian rows of equation (4.43), Poisson for the three count rows and binomial in the logit for the prior on the share, stack into the residual vector and the 4×2 Jacobian

$$\mathbf{r}(\mathbf{z}) = \begin{bmatrix} (\bar{k}_0 - a_{\lambda0})/\sqrt{\bar{k}_0} \\ (\bar{k}_1 - k_1)/\sqrt{\bar{k}_1} \\ (\bar{k}_2 - k_2)/\sqrt{\bar{k}_2} \\ (10\,\beta_{\mathrm{cc}} - a_{\beta0})/\sqrt{10\,\beta_{\mathrm{cc}}(1-\beta_{\mathrm{cc}})} \end{bmatrix}, \qquad \mathbf{J}(\mathbf{z}) = \begin{bmatrix} \sqrt{\bar{k}_0} & 0 \\ \sqrt{\bar{k}_1} & -\beta_{\mathrm{cc}}\sqrt{\bar{k}_1} \\ \sqrt{\bar{k}_2} & (1-\beta_{\mathrm{cc}})\sqrt{\bar{k}_2} \\ 0 & \sqrt{10\,\beta_{\mathrm{cc}}(1-\beta_{\mathrm{cc}})} \end{bmatrix}, \tag{15.10}$$

with $10 = a_{\beta0}+b_{\beta0}$, $\bar{k}_0 = b_{\lambda0}\mathrm{e}^{z_1}$, $\bar{k}_1 = 2(1-\beta_{\mathrm{cc}})T_{\mathrm{tot}}\mathrm{e}^{z_1}$, $\bar{k}_2 = \beta_{\mathrm{cc}}T_{\mathrm{tot}}\mathrm{e}^{z_1}$ and $\beta_{\mathrm{cc}} = 1/(1+\mathrm{e}^{-\xi})$. Each pair reproduces the gradient and the Gauss–Newton curvature of its term, so that the gradient of the objective is $\mathbf{J}^{\mathsf{T}}\mathbf{r}$; the pair is a surrogate for the row, and $\frac{1}{2}\mathbf{r}^{\mathsf{T}}\mathbf{r}$ is not the objective (Chapter 6). With every row normalized in this way the information matrix is $\mathbf{J}^{\mathsf{T}}\mathbf{J}$.

15.2 The posterior of the state (step 3)

In the book's notation, $p(\mathbf{z} \mid \mathbf{y})$ is the posterior density of the state, $f(\mathbf{z})$ a quantity derived from it, f^* a level of that quantity and M_j a drawing of the plant.

Five objects of the fit recur in the specialized forms. The *MAP state* $\hat{\mathbf{z}} = (\hat{z}_1, \hat{\xi})^{\mathsf{T}}$ is the minimum of the objective that the four rows put on the state; the *information matrix* $\mathbf{\Lambda} = \mathbf{J}^{\mathsf{T}}\mathbf{J}$ is built from $\mathbf{J}$, the 4×2 Jacobian of the normalized rows; the *Hessian* $\mathbf{H} = \mathbf{J}^{\mathsf{T}}\mathbf{J} + \mathbf{C}$ adds to it the curvature term of the two count rows whose log means are not linear in the logit; the *Laplace covariance* $\mathbf{\Sigma} = \mathbf{H}^{-1}$, a 2×2 matrix, makes the posterior of the state approximately normal, $\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$; and a derived quantity f is carried by its *gradient* $\mathbf{g} = \partial f/\partial \mathbf{z}$ at $\hat{\mathbf{z}}$, with the variance $\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}$. Every answer of step 4 is computed from these objects, from $\hat{\mathbf{z}}$ and $\mathbf{\Sigma}$ above all. Beside them stands what this problem admits exactly. Given the share, the rate has a gamma posterior, so the joint posterior of the two entries is one integral away from every summary of the share; it is evaluated in that one dimension and, for the questions that need the joint posterior, on a grid in both, and every derived quantity is evaluated by draws from it. That exact posterior gives the value set beside each answer, with the gap. The questions are read from the posterior of the state that the four rows of step 2 give, the second also from the fit of the drawing of independence. This section derives that posterior and computes it for the pump pair; Section 15.3 reads every answer from it.

15.2.1 The exact posterior of the rate and the share

The joint posterior. With the independent priors of Section 15.1.4, densities proportional to $\lambda^{a_{\lambda 0}-1}\mathrm{e}^{-b_{\lambda 0}\lambda}$ and $\beta_{\mathrm{cc}}^{a_{\beta 0}-1}(1-\beta_{\mathrm{cc}})^{b_{\beta 0}-1}$, Bayes' theorem with the likelihood (15.6) gives

$$p(\lambda, \beta_{\mathrm{cc}} \mid \mathbf{y}) \propto \lambda^{a_{\lambda 0}+m-1}\, \mathrm{e}^{-(b_{\lambda 0}+(2-\beta_{\mathrm{cc}})T_{\mathrm{tot}})\lambda}\; \beta_{\mathrm{cc}}^{a_{\beta 0}+k_2-1}(1-\beta_{\mathrm{cc}})^{b_{\beta 0}+k_1-1}. \tag{15.11}$$

For a fixed share the factor in λ is a gamma kernel, so the conditional posterior of the rate is

$$\lambda \mid \beta_{\mathrm{cc}}, \mathbf{y} \sim \mathrm{Gamma}\big(a_{\lambda 0}+m,\ b_{\lambda 0}+(2-\beta_{\mathrm{cc}})T_{\mathrm{tot}}\big) = \mathrm{Gamma}\big(12,\ 10\,000\,\mathrm{h}+(2-\beta_{\mathrm{cc}})\,43\,800\,\mathrm{h}\big), \tag{15.12}$$

the Gamma–Poisson update of Chapter 6 with all ten events in the place of the failures and the exposure multiplied by $2-\beta_{\mathrm{cc}}$, the total event rate of the pair per unit of λ.

The marginal of the share. Integrating λ out of equation (15.11) with the gamma integral $\int_0^\infty \lambda^{a-1}\mathrm{e}^{-b\lambda}\,\mathrm{d}\lambda = \Gamma(a)/b^a$, with the shape $a = a_{\lambda 0}+m$ and the rate parameter $b = b_{\lambda 0}+(2-\beta_{\mathrm{cc}})T_{\mathrm{tot}}$ of the conditional gamma (15.12), leaves

$$\begin{aligned} p(\beta_{\mathrm{cc}} \mid \mathbf{y}) &\propto \beta_{\mathrm{cc}}^{a_{\beta 0}+k_2-1}(1-\beta_{\mathrm{cc}})^{b_{\beta 0}+k_1-1}\,\big(b_{\lambda 0}+(2-\beta_{\mathrm{cc}})T_{\mathrm{tot}}\big)^{-(a_{\lambda 0}+m)} \\ &= \beta_{\mathrm{cc}}\,(1-\beta_{\mathrm{cc}})^{17}\,\big(97\,600-43\,800\,\beta_{\mathrm{cc}}\big)^{-12}, \end{aligned} \tag{15.13}$$

with the exposures in hours. The marginal is known in closed form up to its normalizing constant, a one-dimensional integral. Its moments, quantiles and tail probabilities are one-dimensional quadratures, and the marginal of the rate is the mixture of the gamma distributions (15.12) over it.

The grid and the draws. The exact route for a two-entry state is the posterior evaluated on a grid in both entries: the logarithm of equation (15.11) is evaluated at every point, exponentiated after subtracting its maximum, and normalized with the trapezoidal rule in both directions. Marginals are the trapezoidal sums along one direction, quantiles are read from the cumulative trapezoidal sums, and the expectation of a derived quantity is its trapezoidal sum against the normalized weights. Draws of the posterior are grid cells drawn with probabilities proportional to their weights; the exact construction needs no grid, drawing the share from the marginal (15.13) by the inverse of its cumulative distribution and the rate from the conditional gamma (15.12).

The Beta approximation. Dropping the last factor of equation (15.13) leaves the kernel of a Beta distribution,

$$\beta_{\mathrm{cc}} \mid \mathbf{y} \approx \mathrm{Beta}(a_{\beta 0}+k_2,\ b_{\beta 0}+k_1) = \mathrm{Beta}(2,\ 18), \tag{15.14}$$

the posterior that treats each recorded event as a trial that is a common cause with probability β_{cc}, the Beta–binomial update of Chapter 5 with the doubles as the successes and the singles as the failures. The dropped factor is the coupling through the total event rate $(2-\beta_{\mathrm{cc}})\lambda$. Normalized to one at $\beta_{\mathrm{cc}} = 0$, it is

$$\begin{aligned} \left(\frac{b_{\lambda 0}+2T_{\mathrm{tot}}}{b_{\lambda 0}+(2-\beta_{\mathrm{cc}})T_{\mathrm{tot}}}\right)^{a_{\lambda 0}+m} &= \left(1-\frac{T_{\mathrm{tot}}}{b_{\lambda 0}+2T_{\mathrm{tot}}}\,\beta_{\mathrm{cc}}\right)^{-(a_{\lambda 0}+m)} \approx \exp(s\,\beta_{\mathrm{cc}}), \\ s &= \frac{(a_{\lambda 0}+m)T_{\mathrm{tot}}}{b_{\lambda 0}+2T_{\mathrm{tot}}}, \end{aligned} \tag{15.15}$$

using $-a\ln(1-x) \approx ax$ for small x: the factor tilts the Beta distribution toward larger shares.

The event share behind the approximation. The tilt has a meaning that the thinning construction of Section 15.1.6 makes plain. When the record is long enough that $b_{\lambda 0}$ is negligible beside $2T_{\text{tot}}$, the dropped factor is proportional to $(2-\beta_{\text{cc}})^{-m}$, and the likelihood of the share becomes

$$(1-\beta_{\text{cc}})^{k_1}\beta_{\text{cc}}^{k_2}(2-\beta_{\text{cc}})^{-m} = \left(\frac{1-\beta_{\text{cc}}}{2-\beta_{\text{cc}}}\right)^{k_1}\left(\frac{\beta_{\text{cc}}}{2-\beta_{\text{cc}}}\right)^{k_2} \propto (1-\theta)^{k_1}\theta^{k_2},$$

since $1-\theta = 2(1-\beta_{\text{cc}})/(2-\beta_{\text{cc}})$: the counts are binomial trials of the event share θ of equation (15.4), not of β_{cc}. The Beta approximation (15.14) treats the counts as trials of β_{cc}, and as the record grows its mean tends to k_2/m, the event share $\theta \approx \beta_{\text{cc}}/2$, half the true share. It is close for a short record only while the prior still carries most of the information and the coupling factor is moderate, and its error grows with every year of record (Section 15.6.3).

15.2.2 The Laplace posterior from the four rows

The curvature term. Unlike the rows of Chapter 8, the logarithms of $\bar{k}_1$ and $\bar{k}_2$ are not linear in the state, and the Hessian of a count row is $\mathbf{J}_i^{\mathsf{T}}\mathbf{J}_i$ plus its observed-minus-expected count times the second derivative of its log mean, equation (4.42). Each log mean is z_1 plus a function of ξ alone, $\ln\bar{k}_1 = \ln(2T_{\text{tot}}) + z_1 + \ln(1-\beta_{\text{cc}})$ and $\ln\bar{k}_2 = \ln T_{\text{tot}} + z_1 + \ln\beta_{\text{cc}}$, which is why the log-count Jacobian of Section 15.1.6 has the constant first column; every second derivative that involves z_1 is therefore zero. With $\mathrm{d}\beta_{\text{cc}}/\mathrm{d}\xi = \beta_{\text{cc}}(1-\beta_{\text{cc}})$ the first derivatives in ξ are $\partial\ln\beta_{\text{cc}}/\partial\xi = 1-\beta_{\text{cc}}$ and $\partial\ln(1-\beta_{\text{cc}})/\partial\xi = -\beta_{\text{cc}}$, and differentiating either once more gives

$$\frac{\partial^2\ln\bar{k}_1}{\partial\xi^2} = \frac{\partial^2\ln\bar{k}_2}{\partial\xi^2} = -\beta_{\text{cc}}(1-\beta_{\text{cc}}),$$

so that the Hessian of the objective is $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ plus the curvature term

$$\mathbf{H} = \frac{\partial^2\mathcal{L}}{\partial\mathbf{z}\,\partial\mathbf{z}^{\mathsf{T}}} = \mathbf{J}^{\mathsf{T}}\mathbf{J} + \mathbf{C}, \qquad \mathbf{C} = \begin{bmatrix} 0 & 0 \\ 0 & -\beta_{\text{cc}}(1-\beta_{\text{cc}})\left[(\bar{k}_1 - k_1) + (\bar{k}_2 - k_2)\right] \end{bmatrix}. \tag{15.16}$$

The prior rows contribute nothing to $\mathbf{C}$. The gamma row's log mean, $\ln b_{\lambda 0} + z_1$, is linear. The binomial row is not a Poisson row, and its term $-a_{\beta 0}\ln\beta_{\text{cc}} - b_{\beta 0}\ln(1-\beta_{\text{cc}})$ has the exact second derivative $(a_{\beta 0}+b_{\beta 0})\beta_{\text{cc}}(1-\beta_{\text{cc}}) = 10\,\beta_{\text{cc}}(1-\beta_{\text{cc}})$ in ξ, which is J_{42}^2, its Gauss–Newton curvature: the logit is the canonical link of the binomial row. The matrix $\mathbf{H}$ is the curvature the Laplace covariance inverts, $\boldsymbol{\Sigma} = \mathbf{H}^{-1}$; it differs from the information matrix $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ by $\mathbf{C}$.

The normal equations and the Newton iteration. Each Newton step, the curvature term kept, solves

$$(\mathbf{J}^{\mathsf{T}}\mathbf{J} + \mathbf{C})\,\Delta\mathbf{z} = -\mathbf{J}^{\mathsf{T}}\mathbf{r}, \qquad \mathbf{z} \leftarrow \mathbf{z} + \Delta\mathbf{z}, \tag{15.17}$$

with $\mathbf{J}$, $\mathbf{r}$ and $\mathbf{C}$ evaluated at the current state. The iteration starts at the prior means of λ and β_{cc}, mapped to the two entries, where the two prior rows expect exactly their pseudo-counts and only the count rows pull. The iteration stops where the gradient vanishes,

$$\mathbf{J}^{\mathsf{T}}\mathbf{r} = \begin{bmatrix} (\bar{k}_0 - a_{\lambda 0}) + (\bar{k}_1 - k_1) + (\bar{k}_2 - k_2) \\ -\beta_{\text{cc}}(\bar{k}_1 - k_1) + (1-\beta_{\text{cc}})(\bar{k}_2 - k_2) + (10\,\beta_{\text{cc}} - a_{\beta 0}) \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \tag{15.18}$$

which is not the condition that every residual vanish: the rows can disagree among themselves and only their projections on the two columns must cancel. The first entry says that the three rows carrying z_1 expect, in total, the $a_{\lambda 0} + m$ events they hold, the mean of the conditional gamma of Section 15.2.1; the second weighs the singles row's excess of expected singles by $-\beta_{\text{cc}}$ and the doubles row's shortfall of expected doubles by $1-\beta_{\text{cc}}$, against the prior row's excess of pseudo-events.

The information, the Hessian and the covariance. At the MAP state the Laplace posterior is $\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ with $\boldsymbol{\Sigma} = \mathbf{H}^{-1} = (\mathbf{J}^{\mathsf{T}}\mathbf{J} + \mathbf{C})^{-1}$, and the entries are read off as $\hat{z}_1$ with $\sqrt{\Sigma_{11}}$, $\hat{\xi}$ with $\sqrt{\Sigma_{22}}$, and the correlation $\Sigma_{12}/\sqrt{\Sigma_{11}\Sigma_{22}}$. Each diagonal entry of $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ is the sum of squares down one column of $\mathbf{J}$. The z_1 column gives $\bar{k}_0 + \bar{k}_1 + \bar{k}_2$, the expected counts of the prior, the singles row and the doubles row, which at $\hat{\mathbf{z}}$ equals $a_{\lambda 0} + m$ by the first entry of equation (15.18): the rate per pump is informed by every event. The ξ column gives $\beta_{\mathrm{cc}}^2\bar{k}_1 + (1-\beta_{\mathrm{cc}})^2\bar{k}_2 + 10\,\beta_{\mathrm{cc}}(1-\beta_{\mathrm{cc}})$, in which the singles row enters weighted by β_{cc}^2 and the doubles row by $(1-\beta_{\mathrm{cc}})^2$: the doubles row is the data row that informs the share. The off-diagonal entry is the sum of the products $J_{i1}J_{i2}$ along the rows, $-\beta_{\mathrm{cc}}\bar{k}_1 + (1-\beta_{\mathrm{cc}})\bar{k}_2 = -\beta_{\mathrm{cc}}(1-\beta_{\mathrm{cc}})\lambda T_{\mathrm{tot}}$, negative at every state. The singles row has entries of opposite sign: raising the logit of the share by $\Delta\xi$ and the log rate by $\beta_{\mathrm{cc}}\Delta\xi$ leaves the expected singles unchanged, so the singles row does not resist a larger share that comes with a larger rate, and the negative off-diagonal entry of $\mathbf{H}$ becomes a positive Σ_{12}.

The Beta approximation in the same objects. The Beta approximation of Section 15.2.1 replaces the singles and doubles rows, as far as the share is concerned, by one binomial row of the $m = k_1 + k_2$ events on ξ, with the Jacobian row $\big(0,\ \sqrt{m\,\beta_{\mathrm{cc}}(1-\beta_{\mathrm{cc}})}\big)$: a row with no entry in the z_1 column, which removes the coupling and gives the events the probability β_{cc} in place of the event share θ.

15.2.3 The posterior of the pump pair

The pump pair's inputs are the record of Table 15.1, $k_1 = 9$ singles and $k_2 = 1$ double in $T_{\mathrm{tot}} = 43\,800\,\mathrm{h}$, and the priors Gamma(2, 10 000 h) and Beta(1, 9).

The priors and the maximum-likelihood state. The prior Gamma(2, 10 000 h) has the mean $2 \times 10^{-4}\,\mathrm{h}^{-1}$ and the central 95 % range $[0.24,\ 5.6] \times 10^{-4}\,\mathrm{h}^{-1}$; the prior Beta(1, 9) has the mean 0.1, the median 0.074 and the central 95 % range [0.003, 0.336], and it puts probability 0.63 on shares above 0.05. The record alone is fitted exactly by the maximum-likelihood state of Section 15.1.6,

$$\beta_{\mathrm{cc}}^{\mathrm{ML}} = \frac{2}{11} = 0.182, \qquad \lambda^{\mathrm{ML}} = 1.26 \times 10^{-4}\,\mathrm{h}^{-1},$$

with the total event rate $\nu^{\mathrm{ML}} = 10/T_{\mathrm{tot}} = 2.28 \times 10^{-4}\,\mathrm{h}^{-1}$ and the event share $\theta^{\mathrm{ML}} = 0.1$.

The Newton iteration. Start at the prior means of λ and β_{cc} mapped to the entries, $\mathbf{z}^{(0)} = (\ln(2 \times 10^{-4}),\ \ln(0.1/0.9))^{\mathsf{T}} = (-8.5172,\ -2.1972)^{\mathsf{T}}$. There the two prior rows expect exactly their pseudo-counts and the two count rows do not,

$$\mathbf{r}(\mathbf{z}^{(0)}) = \begin{bmatrix} 0 \\ 1.7044 \\ -0.1325 \\ 0 \end{bmatrix}, \qquad \mathbf{J}(\mathbf{z}^{(0)}) = \begin{bmatrix} 1.4142 & 0 \\ 3.9709 & -0.3971 \\ 0.9359 & 0.8424 \\ 0 & 0.9487 \end{bmatrix},$$

since the singles row expects 15.77 singles against 9 and the doubles row 0.876 doubles against 1. With $\mathbf{C}_{22} = -0.5980$ the first step solves

$$\begin{bmatrix} 18.6440 & -0.7884 \\ -0.7884 & 1.1693 \end{bmatrix} \Delta\mathbf{z} = \begin{bmatrix} -6.6440 \\ 0.7884 \end{bmatrix}, \qquad \Delta\mathbf{z} = \begin{bmatrix} -0.3375 \\ 0.4467 \end{bmatrix};$$

the excess of singles lowers the log rate, and since the singles row's entry in the ξ column is negative, the same excess, with the shortfall of doubles, raises the logit of the share. Table 15.3 lists the iteration: each row is one iterate, and the columns give the two entries, the two entries

of the gradient $\mathbf{J}^{\mathsf{T}}\mathbf{r}$, the three distinct entries of $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ and the curvature term $\mathbf{C}_{22}$. The iteration converges quadratically from the second step, and the gradient is below 10^{-10} at the fourth iterate. A Gauss–Newton iteration, which drops $\mathbf{C}$, reaches the same state, since $\mathbf{C}$ changes the path and not the minimum, but linearly: near the minimum its gradient falls by a factor of about 23 per step and is below 10^{-10} only at the eighth iterate.

Table 15.3: Newton iteration of equation (15.17) for the pump pair from the prior means of λ and β_{cc}, mapped to $\mathbf{z}^{(0)} = (-8.5172,\ -2.1972)^{\mathsf{T}}$, with the gradient $\mathbf{J}^{\mathsf{T}}\mathbf{r}$, the entries of $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ and the curvature term $\mathbf{C}_{22}$ of equation (15.16).

Iteration [–]	z_1 [–]	ξ [–]	$\partial\mathcal{L}/\partial z_1$ [–]	$\partial\mathcal{L}/\partial\xi$ [–]	$(\mathbf{J}^{\mathsf{T}}\mathbf{J})_{11}$ [–]	$(\mathbf{J}^{\mathsf{T}}\mathbf{J})_{12}$ [–]	$(\mathbf{J}^{\mathsf{T}}\mathbf{J})_{22}$ [–]	$\mathbf{C}_{22}$ [–]
0	−8.5172	−2.1972	6.644	−0.788	18.644	−0.788	1.767	−0.598
1	−8.8547	−1.7505	1.004	0.172	13.004	−0.788	2.166	−0.199
2	−8.9392	−1.8716	3.4×10^{-2}	3.3×10^{-3}	12.034	−0.664	1.908	−0.084
3	−8.9422	−1.8745	4.5×10^{-5}	-9.0×10^{-7}	12.000	−0.661	1.902	−0.080
4	−8.9422	−1.8745	8.5×10^{-11}	-5.9×10^{-12}	12.000	−0.661	1.902	−0.080

The MAP state. The iteration stops at

$$\hat{\mathbf{z}} = \begin{bmatrix} \hat{z}_1 \\ \hat{\xi} \end{bmatrix} = \begin{bmatrix} -8.942 \\ -1.874 \end{bmatrix}, \qquad \hat{\lambda} = 1.308 \times 10^{-4}\,\mathrm{h}^{-1}, \quad \hat{\beta}_{\mathrm{cc}} = 0.1330, \tag{15.19}$$

with $\hat{\lambda} T_{\mathrm{tot}} = 5.727$. This is not the joint mode $(1.17 \times 10^{-4}\,\mathrm{h}^{-1},\ 0.075)$ of the density of $(\lambda, \beta_{\mathrm{cc}})$ (Section 15.3.1): the density of (z_1, ξ) is the density of $(\lambda, \beta_{\mathrm{cc}})$ times the change-of-variable factor $\lambda\,\beta_{\mathrm{cc}}(1 - \beta_{\mathrm{cc}})$, so the two densities peak at different places. There the four rows expect $\bar{\mathbf{k}} = (1.308,\ 9.931,\ 0.762)^{\mathsf{T}}$ and $10\hat{\beta}_{\mathrm{cc}} = 1.330$, and the residual vector is $\mathbf{r}(\hat{\mathbf{z}}) = (-0.606,\ 0.295,\ -0.273,\ 0.308)^{\mathsf{T}}$, not zero, while its projections (15.18) vanish: $1.308 + 9.931 + 0.762 = 12$ expected events in the first entry, and in the second the excess of 0.931 expected singles, the shortfall of 0.238 expected doubles and the prior row's excess of 0.330 pseudo-events.

The Jacobian, the Hessian and the covariance. At $\hat{\mathbf{z}}$ the Jacobian (15.10), 4×2, and the matrices of equation (15.16) are

$$\mathbf{J}(\hat{\mathbf{z}}) = \begin{bmatrix} 1.1435 & 0 \\ 3.1513 & -0.4192 \\ 0.8728 & 0.7567 \\ 0 & 1.0739 \end{bmatrix}, \qquad \mathbf{J}^{\mathsf{T}}\mathbf{J} = \begin{bmatrix} 12.000 & -0.661 \\ -0.661 & 1.902 \end{bmatrix}, \qquad \mathbf{C} = \begin{bmatrix} 0 & 0 \\ 0 & -0.080 \end{bmatrix},$$
$$\mathbf{H} = \mathbf{J}^{\mathsf{T}}\mathbf{J} + \mathbf{C} = \begin{bmatrix} 12.000 & -0.661 \\ -0.661 & 1.822 \end{bmatrix}, \qquad \boldsymbol{\Sigma} = \mathbf{H}^{-1} = \begin{bmatrix} 0.0850 & 0.0308 \\ 0.0308 & 0.5601 \end{bmatrix}, \tag{15.20}$$

and the Laplace posterior is $\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$, with $\sqrt{\Sigma_{11}} = 0.292$ in z_1, $\sqrt{\Sigma_{22}} = 0.748$ in ξ and the correlation $\Sigma_{12}/\sqrt{\Sigma_{11}\Sigma_{22}} = 0.14$. The curvature term is small but not negligible: without it $\sqrt{\Sigma_{22}}$ would be 0.732 and the correlation 0.138.

The columns of the Jacobian. The z_1 column of $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ is $12 = 1.31 + 9.93 + 0.76$ from the prior, the singles row and the doubles row. The ξ column is $1.902 = 0 + 0.176 + 0.573 + 1.153$, and with $\mathbf{C}_{22}$ the Hessian entry is $H_{22} = 1.822$, of which the prior Beta$(1, 9)$ supplies 1.15 and the data rows, with the curvature term, 0.67: the share's posterior is still mostly its prior. The

doubles row's ξ entry is 0.757 against 0.873 in its z_1 column; the singles row's is 0.13 times its z_1 entry. The off-diagonal entry is -1.321 from the singles and $+0.661$ from the doubles, -0.661 in all. The eigenvectors of $\boldsymbol{\Sigma}$ are $(0.998,\ -0.064)^{\mathsf{T}}$ and $(0.064,\ 0.998)^{\mathsf{T}}$, with eigenvalues 0.083 and 0.562.

The fit of the independence drawing. The second question also needs the posterior of the state of M_0, the same construction on the single entry z_1 and the three rows of the right-hand pattern of equation (15.2). All three log means are linear in z_1, the coincidence row's with slope 2, so with the expected counts $b_{\lambda 0}\lambda$, $2\lambda T_{\text{tot}}$ and $2\lambda^2 t_{\text{w}} T_{\text{tot}}$ of its three rows the Jacobian rows are $\sqrt{b_{\lambda 0}\lambda}$, $\sqrt{2\lambda T_{\text{tot}}}$ and $2\sqrt{2\lambda^2 t_{\text{w}} T_{\text{tot}}}$, and there is no curvature term. Newton's iteration from the prior mean reaches $\hat{z}_1 = -8.926$ ($\hat{\lambda} = 1.329 \times 10^{-4}\,\text{h}^{-1}$) at $t_{\text{w}} = 8\,\text{h}$, where the rows expect $(1.329,\ 11.646,\ 0.0124)$ against $(2,\ 9,\ 1)$ and

$$\mathbf{r}(\hat{\mathbf{z}}) = \begin{bmatrix} -0.582 \\ 0.775 \\ -8.874 \end{bmatrix}, \qquad \mathbf{J}(\hat{\mathbf{z}}) = \begin{bmatrix} 1.153 \\ 3.413 \\ 0.223 \end{bmatrix}, \qquad \mathbf{H} = \mathbf{J}^{\mathsf{T}}\mathbf{J} = 13.02, \qquad \boldsymbol{\Sigma} = 1/13.02 = 0.0768. \tag{15.21}$$

The coincidence row's residual is large and its Jacobian entry small: the independence drawing cannot bring the expected coincidences near the one observed without contradicting the nine singles, which is the evidence against it. At the other windows of step 4 the same iteration moves $\hat{z}_1$ only from -8.924 at 1 h to -8.935 at 48 h, and $\mathbf{H}$ from 13.00 to 13.15.

The pair $(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of equations (15.19) and (15.20) is what step 4 uses: every answer is computed from it, the comparison of the drawings also from the fit (15.21), and the exact posterior of Section 15.2.1 gives the value beside each answer, with the gap between the two.

15.3 The answers (step 4)

The five questions of Table 15.2 take four general forms of Chapter 2, estimation, hypothesis comparison, diagnosis and prediction under a threshold, and each is answered from the posterior of Section 15.2. The questions name the window $t_{\text{w}} = 8\,\text{h}$ with 1, 4, 24 and 48 h beside it, the repair rate $\mu = 1/12\,\text{h}^{-1}$ and the limit 10^{-4} on the pair's unavailability; the year has 8760 calendar hours. Every answer is computed from the posterior of the state of step 3, the MAP state $\hat{\mathbf{z}}$ and the covariance $\boldsymbol{\Sigma}$, and that is the value the chapter reports; the exact posterior of Section 15.2.1 gives the value set beside it, with the gap.

15.3.1 The common-cause share

Engineering question. What share of the pumps' failures is a common cause, and how well is it known? The question asks for the fraction β_{cc} of each pump's failure rate that comes from the shared cause, with its band; the failure rate per pump λ is estimated with it.

General form. An estimation question asks for the posterior distribution of the quantity, equation (2.9), summarized by a location and by the central interval of equation (2.11):

$$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y}) = \int \mathbf{1}\big[f(\mathbf{z}) \le f^*\big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}. \tag{15.22}$$

Specialized form. The question is equation (15.22) with the state $\mathbf{z} = (z_1, \xi)^{\mathsf{T}}$, the derived quantity $f(\mathbf{z}) = \beta_{\text{cc}} = 1/(1 + \mathrm{e}^{-\xi})$, an increasing function of the second entry alone, and, for the rate per pump, $f(\mathbf{z}) = \lambda = \mathrm{e}^{z_1}$, and $p(\mathbf{z} \mid \mathbf{y})$ the posterior of step 3, the Laplace normal $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of equation (15.20). The marginal of ξ is $\mathcal{N}(\hat{\xi}, \Sigma_{22})$, read off as the second entry of $\hat{\mathbf{z}}$ and the second diagonal entry of $\boldsymbol{\Sigma} = \mathbf{H}^{-1}$, the inverse of the Hessian rather than of $\mathbf{J}^{\mathsf{T}}\mathbf{J}$, since

the count rows are not linear in ξ (Section 15.2.2). The location is the MAP share $\hat{\beta}_{\text{cc}}$, the image of the median of ξ; the central interval of the share is the interval of the entry mapped through the inverse logit,

$$\beta_{\text{cc}} \in \left[\frac{1}{1+\mathrm{e}^{-(\hat{\xi}-1.96\sqrt{\Sigma_{22}})}},\ \frac{1}{1+\mathrm{e}^{-(\hat{\xi}+1.96\sqrt{\Sigma_{22}})}} \right], \tag{15.23}$$

a tail of the share is a normal tail of the entry,

$$\Pr(\beta_{\text{cc}} > \beta^* \mid \mathbf{y}) \approx \Phi\left(\frac{\hat{\xi}-\xi^*}{\sqrt{\Sigma_{22}}} \right), \qquad \xi^* = \ln \frac{\beta^*}{1-\beta^*}, \tag{15.24}$$

with Φ the standard normal cumulative distribution, and the rate per pump has the interval $\mathrm{e}^{\hat{z}_1 \pm 1.96\sqrt{\Sigma_{11}}}$. The correlation $\Sigma_{12}/\sqrt{\Sigma_{11}\Sigma_{22}}$ says how the two entries move together; its sign is read in the singles row of $\mathbf{J}$, whose entries in the two columns have opposite signs. A mean and a variance of a function of the state have no closed form under a normal on a logit, and they are expectations over the state. The Gauss–Hermite sum of Chapter 5, equation (5.33), extends to the two entries by writing the state through the Cholesky factor $\mathbf{L}$ of $\boldsymbol{\Sigma}$, $\mathbf{L}\mathbf{L}^{\mathsf{T}} = \boldsymbol{\Sigma}$: with the nodes t_i and the weights w_i of a rule of Q nodes in each direction,

$$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}] \approx \frac{1}{\pi} \sum_{i=1}^{Q} \sum_{j=1}^{Q} w_i w_j\, f(\mathbf{z}_{ij}), \qquad \mathbf{z}_{ij} = \hat{\mathbf{z}} + \sqrt{2}\,\mathbf{L} \begin{bmatrix} t_i \\ t_j \end{bmatrix}, \qquad \mathbf{L} = \begin{bmatrix} 0.2916 & 0 \\ 0.1057 & 0.7409 \end{bmatrix}, \tag{15.25}$$

which is exact for a polynomial of degree up to $2Q-1$ in each direction; the reported values use $Q = 40$.

The values beside them come from the exact posterior. The rate can be integrated out of the joint posterior in closed form, because given the share its posterior is a gamma distribution; what remains is the marginal of the share, known up to its normalizing constant, a one-dimensional integral (Section 15.2.1). Its mean, quantiles and tail probabilities are one-dimensional quadratures, the marginal of the rate is the mixture of the conditional gammas over it, and the same summaries follow from the joint posterior evaluated on a grid in both entries. A third route counts the events as trials of the share and gives a Beta distribution that drops the coupling of the two entries through the total event rate (Section 15.2.1); it is the common shortcut, and Section 15.4.1 sets it beside the other two.

Building blocks. The MAP state $\hat{\mathbf{z}} = (-8.942,\ -1.874)^{\mathsf{T}}$ and the covariance $\boldsymbol{\Sigma}$ of equation (15.20), with $\sqrt{\Sigma_{11}} = 0.292$ and $\sqrt{\Sigma_{22}} = 0.748$; the inverse logit, the normal tail (15.24) and the sum (15.25); and, for the exact values beside them, the joint posterior, the conditional gamma of the rate and the marginal of the share, with the grid and the one-dimensional quadrature (Section 15.2.1).

The share and the rate per pump. By equation (15.23), the logit interval $-1.8745 \pm 1.96 \times 0.7484 = [-3.3413,\ -0.4077]$ maps to the band of the share, and the MAP share is its median:

$$\text{median } \hat{\beta}_{\text{cc}} = 0.133, \quad \beta_{\text{cc}} \in [0.034,\ 0.399] \text{ with } 95\,\% \text{ posterior probability.} \tag{15.26}$$

By the tail (15.24) with $\xi^* = \ln(0.05/0.95) = -2.9444$, $\Pr(\beta_{\text{cc}} > 0.05 \mid \mathbf{y}) \approx \Phi((-1.8745 + 2.9444)/0.7484) = \Phi(1.430) = 0.924$. The sum (15.25) with $f = \beta_{\text{cc}}$ and with its square gives the mean 0.155 and the standard deviation 0.096; the density of the share under the normal on ξ peaks at 0.088. The rate per pump has the median $\hat{\lambda} = 1.31 \times 10^{-4}\,\text{h}^{-1}$, the mean $1.36 \times 10^{-4}\,\text{h}^{-1}$, by the same sum or by the lognormal mean $\mathrm{e}^{\hat{z}_1 + \Sigma_{11}/2}$, and the central 95 %

interval $\mathrm{e}^{-8.9422 \pm 1.96 \times 0.2916} = [0.74,\ 2.32] \times 10^{-4}\,\mathrm{h}^{-1}$. The entries (z_1, ξ) have the correlation 0.14, and the sum gives 0.133 for $(\lambda, \beta_{\mathrm{cc}})$, both positive: a larger share goes with a larger rate per pump, the sign that Section 15.2.2 reads off the singles row of the Jacobian.

The values beside them come from the exact posterior. The grid of Section 15.2.1 has 1200×1200 points, λ from 10^{-6} to $8 \times 10^{-4}\,\mathrm{h}^{-1}$ in steps of $6.7 \times 10^{-7}\,\mathrm{h}^{-1}$ and β_{cc} from 10^{-4} to 0.999 in steps of 8.3×10^{-4}. Every summary of the share below agrees with the one-dimensional quadrature of equation (15.13) to the digits given:

$$\mathbb{E}[\beta_{\mathrm{cc}} \mid \mathbf{y}] = 0.131, \quad \text{median } 0.116, \quad \beta_{\mathrm{cc}} \in [0.018,\ 0.329] \text{ with } 95\,\% \text{ posterior probability}, \tag{15.27}$$

with mode 0.078, standard deviation 0.082 and $\Pr(\beta_{\mathrm{cc}} > 0.05 \mid \mathbf{y}) = 0.847$. The rate per pump has posterior mean $1.31 \times 10^{-4}\,\mathrm{h}^{-1}$, median $1.27 \times 10^{-4}\,\mathrm{h}^{-1}$ and central 95 % interval $[0.67,\ 2.16] \times 10^{-4}\,\mathrm{h}^{-1}$. The posterior correlation of $(\lambda, \beta_{\mathrm{cc}})$ is 0.140, and that of the state entries (z_1, ξ) is 0.12. The state route lies above the exact one throughout on the share: by 0.017 at the median, 0.024 at the mean, 0.016 and 0.070 at the two ends of the band and 0.077 on the probability above 0.05, because the normal is symmetric in ξ while the exact posterior of ξ has its longer tail toward small shares. On the rate the gap is smaller: $0.04 \times 10^{-4}\,\mathrm{h}^{-1}$ at the median and 0.07 and $0.16 \times 10^{-4}\,\mathrm{h}^{-1}$ at the ends of the band.

Figure 15.2 draws the exact joint posterior. The horizontal axis of the main panel is the failure rate per pump in units of $10^{-4}\,\mathrm{h}^{-1}$ and the vertical axis the common-cause share. The three closed curves are the boundaries of the highest-density regions that hold 50 %, 90 % and 99 % of the posterior probability; the dot is the exact posterior mean, the cross the joint mode and the blue diamond the MAP state of the Laplace posterior. The upper panel draws the marginal density of the rate against the same horizontal axis, green for the exact one and dotted blue for the normal on z_1; the marginal density of the share is drawn in Figure 15.4 with its approximations. Two features carry the argument. The regions are wide in the share and narrow in the rate: the ten events determine the rate per pump to within a factor of about two, while one double leaves the share anywhere from about 0.02 to 0.33. And the regions are skewed upward and tilted slightly to the right, the long upper tail of the share and its positive correlation with the rate.

15.3.2 The evidence of the two drawings

Engineering question. Is there a common cause at all, or was the double a coincidence? The question asks which drawing describes the plant: the graph with the shared-cause node, M_1, or the graph of independence, M_0, in which the double is two independent failures that fell within t_{w} of each other.

General form. A hypothesis comparison is the ratio of the evidences of two complete drawings fitted to the same evidence, the Bayes factor of equation (2.14), each evidence the likelihood averaged over its drawing's prior:

$$B_{AB} = \frac{p(\mathbf{y} \mid M_A)}{p(\mathbf{y} \mid M_B)}, \qquad p(\mathbf{y} \mid M) = \int p(\mathbf{y} \mid \mathbf{z}, M)\, p(\mathbf{z} \mid M)\, \mathrm{d}\mathbf{z}. \tag{15.28}$$

Specialized form. Three choices enter here. The coincidence window is the one fixed with the record (Section 15.1.4): this is what makes the doubles count a datum of both drawings, and the evidence depends on it directly (Section 15.4.4). Under M_0 the probability that a failure of one pump finds a failure of the other within the window is taken as $2\lambda t_{\mathrm{w}}$: this gives the coincidence row its mean and the evidence a closed form, and it needs $\lambda t_{\mathrm{w}} \ll 1$ and the second pump running when the first fails, both checked in Section 15.4.2. And where a probability of

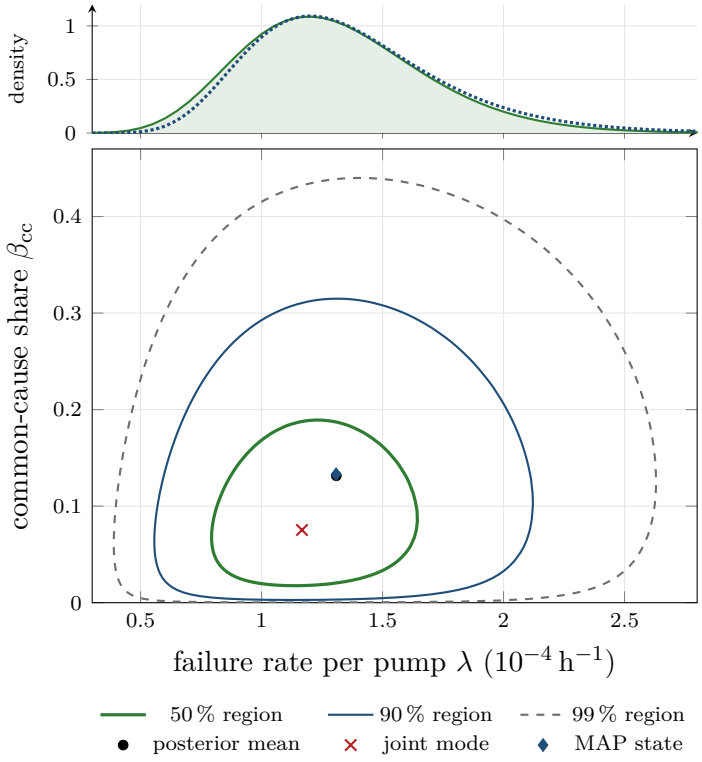


Figure 15.2: The joint posterior of the failure rate per pump and the common-cause share, equation (15.11) with $a_{\lambda 0} = 2$, $b_{\lambda 0} = 10\,000\,\mathrm{h}$, $a_{\beta 0} = 1$, $b_{\beta 0} = 9$, $k_1 = 9$, $k_2 = 1$ and $T_{\mathrm{tot}} = 43\,800\,\mathrm{h}$: boundaries of the highest-density regions holding 50 %, 90 % and 99 % of the probability, the posterior mean ($1.31 \times 10^{-4}\,\mathrm{h}^{-1}$, 0.131), the joint mode ($1.17 \times 10^{-4}\,\mathrm{h}^{-1}$, 0.075) and the MAP state ($1.31 \times 10^{-4}\,\mathrm{h}^{-1}$, 0.133). Upper panel: the marginal density of λ, per $10^{-4}\,\mathrm{h}^{-1}$, exact (green) and the normal $\mathcal{N}(-8.942,\ 0.292^2)$ on z_1 carried to the rate (dotted blue).

M_1 is quoted, the two drawings are given even prior odds: the Bayes factor does not depend on that choice, the posterior probability $\Pr(M_1 \mid \mathbf{y}) = 1/(1 + B_{10}^{-1})$ does.

The question is equation (15.28) with $M_A = M_1$, the state $(z_1, \xi)^\mathsf{T}$ and the four rows of equation (15.2), and $M_B = M_0$, the single entry z_1, a singles row of mean $2\lambda T_{\text{tot}}$ and a coincidence row of mean $c_{\text{w}} T_{\text{tot}}$, where $c_{\text{w}} = 2\lambda^2 t_{\text{w}}$ is the coincidence rate of a time-stamped window (Section 15.1.6). Each integral over the state is taken under the posterior of that drawing's state, the Gaussian integral around its MAP state,

$$\ln p(\mathbf{y} \mid M) \approx -\mathcal{L}(\hat{\mathbf{z}}) + \tfrac{1}{2} \dim(\mathbf{z}) \ln 2\pi - \tfrac{1}{2} \ln \det \mathbf{H}, \tag{15.29}$$

with $\dim(\mathbf{z})$ the number of entries of the drawing's state, $\mathbf{H}$ the Hessian of $\mathcal{L}$ at $\hat{\mathbf{z}}$ and $\mathcal{L}$ including the constants of its rows: for M_1 the MAP state and the Hessian of equations (15.19) and (15.20), and for M_0 the fit (15.21).

The values beside them come from the exact integrals. With $m = k_1 + k_2$ the number of failure events in the record, averaging the likelihood of the two counts of M_1 (Section 15.1.6) over the gamma prior of λ with the gamma integral, and then over the Beta prior of the share,

$$p(\mathbf{y} \mid M_1) = \frac{(2T_{\text{tot}})^{k_1} T_{\text{tot}}^{k_2}}{k_1!\, k_2!} \, \frac{b_{\lambda 0}^{a_{\lambda 0}} \, \Gamma(a_{\lambda 0} + m)}{\Gamma(a_{\lambda 0})} \int_0^1 \frac{(1 - \beta_{\text{cc}})^{k_1} \beta_{\text{cc}}^{k_2}}{\big(b_{\lambda 0} + (2 - \beta_{\text{cc}}) T_{\text{tot}}\big)^{a_{\lambda 0} + m}} \, p(\beta_{\text{cc}}) \, \mathrm{d}\beta_{\text{cc}}, \tag{15.30}$$

a one-dimensional quadrature; the joint posterior on the grid gives the same integral in two dimensions. Under M_0, averaging the likelihood over the gamma prior is a one-dimensional quadrature in λ. Its integrand contains $\exp(-2\lambda T_{\text{tot}} - 2\lambda^2 t_{\text{w}} T_{\text{tot}})$, and the second term of the exponent is the expected number of coincidences, small at practical windows; dropping it leaves a gamma integral,

$$p(\mathbf{y} \mid M_0) \approx \frac{(2T_{\text{tot}})^{k_1} (2t_{\text{w}} T_{\text{tot}})^{k_2}}{k_1!\, k_2!} \, \frac{b_{\lambda 0}^{a_{\lambda 0}} \, \Gamma(a_{\lambda 0} + k_1 + 2k_2)}{\Gamma(a_{\lambda 0}) \, (b_{\lambda 0} + 2T_{\text{tot}})^{a_{\lambda 0} + k_1 + 2k_2}}, \tag{15.31}$$

in which each coincidence contributes two powers of λ, since it is two failures.

The ratio. Dividing equation (15.30) by equation (15.31), the common factors $(2T_{\text{tot}})^{k_1}$, $T_{\text{tot}}^{k_2}$, the factorials and $b_{\lambda 0}^{a_{\lambda 0}} / \Gamma(a_{\lambda 0})$ cancel, and with $a_{\lambda 0} + k_1 + 2k_2 = a_{\lambda 0} + m + k_2$,

$$B_{10} \approx \left(\frac{b_{\lambda 0} + 2T_{\text{tot}}}{2t_{\text{w}}} \right)^{k_2} \frac{\Gamma(a_{\lambda 0} + m)}{\Gamma(a_{\lambda 0} + m + k_2)} \, \mathbb{E}_{p(\beta_{\text{cc}})} \left[\beta_{\text{cc}}^{k_2} (1 - \beta_{\text{cc}})^{k_1} \left(\frac{b_{\lambda 0} + 2T_{\text{tot}}}{b_{\lambda 0} + (2 - \beta_{\text{cc}}) T_{\text{tot}}} \right)^{a_{\lambda 0} + m} \right]. \tag{15.32}$$

For one double, $k_2 = 1$, the gamma ratio is $1/(a_{\lambda 0} + m)$ and the ratio reads

$$B_{10} \approx \frac{E_\beta}{2\tilde{\lambda} t_{\text{w}}}, \qquad \tilde{\lambda} = \frac{a_{\lambda 0} + m}{b_{\lambda 0} + 2T_{\text{tot}}}, \tag{15.33}$$

with E_β the prior expectation in equation (15.32), evaluated by quadrature. The denominator $2\tilde{\lambda} t_{\text{w}}$ is the probability that a failure of one pump finds a failure of the other within the window, the chance explanation of the double; the numerator is the prior average of the probability, of order β_{cc}, that the record contains one shared event among its events, with the tilt of Section 15.2.1. The window enters as the factor ${t_{\text{w}}}^{-k_2}$ of equation (15.32), so that each doubling of the window takes $k_2 \ln 2$ nats from the log Bayes factor.

Building blocks. The MAP state and the Hessian of each drawing, $\hat{\mathbf{z}}$ and $\mathbf{H}$ of equation (15.20) for M_1 and the fit (15.21) for M_0, and the rows' constants; and, for the exact values beside them, the likelihoods of the two drawings (Sections 15.1.6 and 15.1.6), the gamma integral, the prior expectation E_β and the integral over λ by quadrature.

The evidence for a shared cause. By equation (15.29), with $\det \mathbf{H} = 21.43$ and $-\mathcal{L}(\hat{\mathbf{z}}) = -4.982$ for the shared-cause drawing,

$$\ln p(\mathbf{y} \mid M_1) \approx -4.982 + 1.838 - 1.532 = -4.676,$$

and with $\mathbf{H} = 13.02$ and $-\mathcal{L}(\hat{\mathbf{z}}) = -7.517$ for the independence drawing at the window of 8 h,

$$\ln p(\mathbf{y} \mid M_0) \approx -7.517 + 0.919 - 1.283 = -7.881,$$

so that $\ln B_{10} \approx 3.21$, a Bayes factor of 25, and, at even prior odds, $\Pr(M_1 \mid \mathbf{y}) = 0.961$ for the shared-cause drawing. The values beside them come from the exact integrals. The one-dimensional quadrature of equation (15.30) gives $\ln p(\mathbf{y} \mid M_1) = -4.628$, and the two-dimensional quadrature on the grid gives the same. At the window of 8 h the quadrature of the independence drawing gives $\ln p(\mathbf{y} \mid M_0) = -7.875$, so that $\ln B_{10} = 3.25$, a Bayes factor of 26, and $\Pr(M_1 \mid \mathbf{y}) = 0.963$: the state route is lower by 0.05 nats, almost all of it the 0.048 in $\ln p(\mathbf{y} \mid M_1)$, a gap of 0.002 on the probability.

Table 15.4 lists the evidence by window: each row is a coincidence window, and the columns give the log evidence of the independence drawing, the log Bayes factor, the Bayes factor itself, and the posterior probability of the shared-cause drawing if the two drawings were equally probable before the record, each as the Laplace value with the exact one in parentheses. The log evidence of the shared-cause drawing does not depend on the window and is -4.676 (-4.628) throughout.

Table 15.4: The evidence for a shared cause against independence by coincidence window, Laplace with the quadrature in parentheses: $\ln p(\mathbf{y} \mid M_1) = -4.676$ (-4.628); $\ln p(\mathbf{y} \mid M_0)$ one-dimensional in λ, with the coincidence term in the exponent; last column $1/(1 + B_{10}^{-1})$. Windows are $|t_1 - t_2| < t_\mathrm{w}$ on time-stamped records; a log binned in bins of t_w reads the row of $t_\mathrm{w}/2$.

Window t_w (h)	$\ln p(\mathbf{y} \mid M_0)$ [–]	$\ln B_{10}$ [nats]	B_{10} [–]	$\Pr(M_1 \mid \mathbf{y})$, even odds [–]
1	−9.949 (−9.943)	5.27 (5.32)	195 (203)	0.995 (0.995)
4	−8.568 (−8.561)	3.89 (3.93)	49 (51)	0.980 (0.981)
8	−7.881 (−7.875)	3.21 (3.25)	25 (26)	0.961 (0.963)
24	−6.809 (−6.803)	2.13 (2.18)	8.4 (8.8)	0.894 (0.898)
48	−6.155 (−6.149)	1.48 (1.52)	4.4 (4.6)	0.814 (0.821)

15.3.3 The cause of the one double failure

Engineering question. Was the double a common cause, so that one cause took two units out? Within the shared-cause drawing the double could have come from the shared process, or from a coincidence of the two independent parts within the window; under the drawing of independence it can only be a coincidence.

General form. A diagnosis ranks the members of a family of explanations by their posterior probabilities, equation (2.15):

$$\Pr(M_j \mid \mathbf{y}) \propto p(\mathbf{y} \mid M_j)\, \Pr(M_j), \qquad j = 1, \ldots, J. \tag{15.34}$$

Specialized form. One choice enters here: the event question is asked inside M_1 with the coincidences of the independent parts admitted as the second explanation of a double, at the rate $(1 - \beta_\mathrm{cc})^2 c_\mathrm{w}$. Without them the double would be a common cause with certainty and the

question would have no content. The four rows of step 3 leave them out of the doubles row, so the posterior of the state does not know of them, while the exact grid of this question admits them; Section 15.4.4 measures what that difference does to either answer.

The question is equation (15.34) with the family of two members M_j for one recorded event, the shared process and the coincidence of the two independent parts. Under M_1 with coincidences admitted, doubles arrive from two independent Poisson processes, the shared one at rate $\beta_{\mathrm{cc}}\lambda$ and the coincidences of the independent parts at rate $(1-\beta_{\mathrm{cc}})^2 c_{\mathrm{w}} = 2(1-\beta_{\mathrm{cc}})^2\lambda^2 t_{\mathrm{w}}$. Given the state, each explanation's evidence for the one event is the probability that its process produces a double in a short interval, its rate times the interval, and its prior probability is the same for both; by the thinning result of Section 15.1.6 read backward, an event of the superposed process belongs to the first with probability equal to its share of the rate,

$$\Pr(\text{common cause} \mid \text{a double}, \lambda, \beta_{\mathrm{cc}}) = \frac{\beta_{\mathrm{cc}}\lambda}{\beta_{\mathrm{cc}}\lambda + 2(1-\beta_{\mathrm{cc}})^2\lambda^2 t_{\mathrm{w}}}. \tag{15.35}$$

Within M_1 the answer averages it over the posterior of the state, the sum (15.25) with $f(\mathbf{z})$ the probability (15.35) at $\lambda = \mathrm{e}^{z_1}$ and $\beta_{\mathrm{cc}} = 1/(1+\mathrm{e}^{-\xi})$:

$$\Pr(\text{common cause} \mid \text{the double}, \mathbf{y}, M_1) \approx \frac{1}{\pi}\sum_i \sum_j w_i w_j \, \frac{\beta_{\mathrm{cc},ij}\lambda_{ij}}{\beta_{\mathrm{cc},ij}\lambda_{ij} + 2(1-\beta_{\mathrm{cc},ij})^2\lambda_{ij}^2 t_{\mathrm{w}}}. \tag{15.36}$$

The value beside it comes from the exact posterior, the same average as an integral over the joint posterior,

$$\Pr(\text{common cause} \mid \text{the double}, \mathbf{y}, M_1) = \iint \frac{\beta_{\mathrm{cc}}\lambda}{\beta_{\mathrm{cc}}\lambda + 2(1-\beta_{\mathrm{cc}})^2\lambda^2 t_{\mathrm{w}}} \, p(\lambda, \beta_{\mathrm{cc}} \mid \mathbf{y}) \, \mathrm{d}\lambda \, \mathrm{d}\beta_{\mathrm{cc}}, \tag{15.37}$$

evaluated on the grid of the posterior that admits the coincidences in the doubles row; the value at the posterior means, a plug-in, is computed beside it as a check.

The second and third questions sound alike and are answered by different objects. The Bayes factor of Section 15.3.2 asks which drawing describes the plant, a shared-cause mechanism acting over the pumps' life or none, and assigns no event. The event posterior asks, within the shared-cause drawing, which of two explanations produced one recorded event, and it assumes the mechanism. The first is a statement about the model; the second is a statement made with the model. The two combine only through the probabilities of the drawings: under M_0 the double was a coincidence with certainty, so the probability that it was a common cause, with both drawings admitted, is

$$\Pr(\text{common cause} \mid \text{the double}, \mathbf{y}) = \Pr(M_1 \mid \mathbf{y}) \;\Pr(\text{common cause} \mid \text{the double}, \mathbf{y}, M_1), \tag{15.38}$$

the answer to whether one cause took the two units out.

Building blocks. The MAP state $\hat{\mathbf{z}}$, the covariance $\boldsymbol{\Sigma}$ and its factor $\mathbf{L}$ in the sum (15.25); the coincidence rate (Section 15.1.6); $\Pr(M_1 \mid \mathbf{y})$ from the Laplace evidences; and, for the exact values beside them, the grid posterior with the coincidences admitted (Section 15.2.1) and the exact evidences.

The posterior of the one double failure. One row of the arithmetic of equation (15.35), at $t_{\mathrm{w}} = 8\,\mathrm{h}$ and the MAP state $\hat{\lambda} = 1.308\times10^{-4}\,\mathrm{h}^{-1}$, $\hat{\beta}_{\mathrm{cc}} = 0.1330$: $(1-\beta_{\mathrm{cc}})^2 c_{\mathrm{w}} = 2\times0.867^2\times(1.308\times10^{-4})^2\times 8 = 2.06\times10^{-7}\,\mathrm{h}^{-1}$ against $\beta_{\mathrm{cc}}\lambda = 1.74\times10^{-5}\,\mathrm{h}^{-1}$, odds of 85 for the shared process, the probability 0.9883. Table 15.5 evaluates the sum (15.36) with five nodes in each direction. Each row is a node t_i of the first direction, with its weight and the log rate $z_{1,i} = -8.9422 + \sqrt{2}\times0.2916\,t_i$ it fixes; each of the five middle columns is a node t_j of the

second direction, with its weight, and the cell is the probability (15.35) at $\mathbf{z}_{ij}$; the last column averages its row with the weights $w_j/\sqrt{\pi}$. The middle cell is the MAP state. The probability falls in the left column, where the share is small and a coincidence is a real alternative: there ξ_{ij} is near -4, a share near 0.02.

Table 15.5: The Gauss–Hermite sum (15.36) with five nodes in each direction over the Laplace posterior $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ at $t_\mathrm{w} = 8\,\mathrm{h}$. Cells: the probability (15.35) at $\mathbf{z}_{ij}$.

		t_j	-2.0202	-0.9586	0	0.9586	2.0202	
		$w_j/\sqrt{\pi}$	0.01126	0.22208	0.53333	0.22208	0.01126	
t_i	$w_i/\sqrt{\pi}$	$z_{1,i}$						Row average
-2.0202	0.01126	-9.7753	0.9384	0.9794	0.9929	0.9978	0.9995	0.9904
-0.9586	0.22208	-9.3375	0.9203	0.9732	0.9907	0.9972	0.9994	0.9876
0	0.53333	-8.9422	0.9000	0.9660	0.9883	0.9965	0.9993	0.9843
0.9586	0.22208	-8.5469	0.8752	0.9570	0.9853	0.9957	0.9991	0.9802
2.0202	0.01126	-8.1091	0.8419	0.9444	0.9812	0.9946	0.9990	0.9746
Weighted sum of the row averages								0.9841

The weighted sum of the last column is

$$\Pr(\text{common cause} \mid \text{the double}, \mathbf{y}, M_1) \approx 0.984$$

at the window of 8 h; forty nodes change no digit shown. The value beside it comes from the exact posterior. At the posterior means $\lambda = 1.309 \times 10^{-4}\,\mathrm{h}^{-1}$, $\beta_\mathrm{cc} = 0.131$, the plug-in arithmetic gives $(1-\beta_\mathrm{cc})^2 c_\mathrm{w} = 2 \times 0.869^2 \times (1.309 \times 10^{-4})^2 \times 8 = 2.07 \times 10^{-7}\,\mathrm{h}^{-1}$ against $\beta_\mathrm{cc}\lambda = 1.72 \times 10^{-5}\,\mathrm{h}^{-1}$, odds of 83 for the shared process. Averaged over the posterior that admits the coincidences, equation (15.37) gives

$$\Pr(\text{common cause} \mid \text{the double}, \mathbf{y}, M_1) = 0.976$$

at the window of 8 h. The state route is higher by 0.008. Of that gap 0.002 is the model, since the grid without the coincidences in its doubles row, the rows of $\hat{\mathbf{z}}$, gives 0.978, and 0.006 is the normal, which puts less weight on small shares than the exact posterior does. Table 15.6 gives the event's posterior by window: each row is a window, and the columns give the coincidence rate at the posterior means, the posterior that the double was a common cause, and the log Bayes factor of Table 15.4 for comparison, the last two as the Laplace value with the exact one in parentheses.

Table 15.6: The posterior that the observed double failure was a common cause, equation (15.36) over $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$, with equation (15.37) on the grid of the posterior with the coincidences in the doubles row in parentheses, against the log Bayes factor of Table 15.4. Coincidence rate at the posterior means.

Window t_w (h)	$(1-\beta_\mathrm{cc})^2 c_\mathrm{w}$ (h^{-1})	Posterior of a common cause [–]	$\ln B_{10}$ [nats]
1	2.6×10^{-8}	0.998 (0.997)	5.27 (5.32)
8	2.1×10^{-7}	0.984 (0.976)	3.21 (3.25)
24	6.2×10^{-7}	0.955 (0.932)	2.13 (2.18)
48	1.2×10^{-6}	0.916 (0.873)	1.48 (1.52)

Over both drawings at even prior odds, equation (15.38) gives the probability that the double was a common cause, so that one cause took the two pumps out,

$$\Pr(\text{common cause} \mid \text{the double}, \mathbf{y}) \approx 0.961 \times 0.984 = 0.95,$$

and from the exact values beside them

$$\Pr(\text{common cause} \mid \text{the double}, \mathbf{y}) = 0.963 \times 0.976 = 0.94.$$

From 8 to 48 h the log Bayes factor falls by 1.7 nats, a factor of about six, on both routes, and the posterior within M_1 only from 0.984 to 0.916 (0.976 to 0.873 exactly).

15.3.4 The unavailability of the pair and its threshold

Engineering questions. What is the pair's unavailability with the common cause, and what would it be without? Is the pair's unavailability below 10^{-4}, with assurance 0.90? The first asks for the long-run fraction of time both pumps are down, against two baselines: the *independent part* of the same fitted model, the shared events removed, and *independence at the total rate*, every failure of the log treated as an independent failure of one pump. The second asks whether the limit of the cooling design holds with the assurance $\kappa = 0.90$ of Table 15.1.

General form. The first is the estimation of a derived quantity through its moments and quantiles, equations (2.10) and (2.11); the second a prediction under a threshold, equation (2.12):

$$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}] = \int f(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \tag{15.39}$$

$$\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}) = \int \mathbf{1}\big[f(\mathbf{z}) \geq f^*\big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z} \geq \kappa. \tag{15.40}$$

The limit is an upper one, so the requirement is the event $f(\mathbf{z}) = -U_{\text{pair}} \geq -10^{-4}$, accepted if its probability reaches $\kappa = 0.90$.

Specialized form. Two choices enter here. The repair rate is pinned at $1/12\,\mathrm{h}^{-1}$ and taken equal for a single and a common-cause repair: this isolates the share as the unknown that matters, and it fails if a double takes longer to repair because one crew restores two pumps, which a record of common-cause repair times would show. The three processes below are treated as independent two-state processes in their steady state: this gives the pair's unavailability a product form, and it ignores that a pump held down by the shared cause cannot fail on its own, an effect of the order of the overlap term (Section 15.6.7).

The first question is equation (15.39) and the second equation (15.40) with $f(\mathbf{z}) = U_{\text{pair}}$, in the second carried as $-U_{\text{pair}}$ with $f^* = -10^{-4}$, and $p(\mathbf{z} \mid \mathbf{y})$ the posterior of step 3, the Laplace normal $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of equation (15.20); the quantity U_{pair} as a function of the state comes next. The three Poisson processes of Section 15.1.6, each followed by a repair, are three alternating up-and-down processes. The independent process of pump P1 goes down at rate $(1 - \beta_{\text{cc}})\lambda$ and is restored at rate μ; so does the independent process of P2; and the shared process goes down at rate $\beta_{\text{cc}}\lambda$, taking both pumps with it, and is restored at rate μ, both pumps returning together. By the steady state of the two-state model, Section 8.3.3, each is down a fraction of the time equal to its failure rate over the sum of its failure and repair rates:

$$q_{\text{ind}} = \frac{(1 - \beta_{\text{cc}})\lambda}{(1 - \beta_{\text{cc}})\lambda + \mu}, \qquad U_{\text{cc}} = \frac{\beta_{\text{cc}}\lambda}{\beta_{\text{cc}}\lambda + \mu}, \tag{15.41}$$

the independent unavailability of each pump and the unavailability from the shared cause. A pump is up when its own process and the shared process are both up. The pair, in one-out-of-two, is up when the shared process is up and at least one of the two independent processes is up. With the three processes independent, the probability of the pair being up is the product of the probability that the shared process is up and the probability that not both independent processes are down:

$$U_{\text{pair}} = 1 - (1 - U_{\text{cc}})(1 - q_{\text{ind}}^2) = U_{\text{cc}} + q_{\text{ind}}^2 - U_{\text{cc}}\, q_{\text{ind}}^2. \tag{15.42}$$

The first term is the time the shared cause holds both pumps down; the second the time both are down from their own causes; the third removes the overlap counted twice. Each pump's own unavailability, $1-(1-q_{\text{ind}})(1-U_{\text{cc}})$, is close to $\lambda/(\lambda+\mu)$ of Chapter 8, since the two causes together fail it at the rate λ.

The first-order term follows from an expansion in a small ratio. Write $\epsilon=\lambda/\mu$, the ratio of the failure rate to the repair rate. Expanding equation (15.41) in ϵ, $U_{\text{cc}}=\beta_{\text{cc}}\epsilon/(1+\beta_{\text{cc}}\epsilon)=\beta_{\text{cc}}\epsilon-\beta_{\text{cc}}^2\epsilon^2+O(\epsilon^3)$ and $q_{\text{ind}}^2=(1-\beta_{\text{cc}})^2\epsilon^2+O(\epsilon^3)$, while the overlap term is $O(\epsilon^3)$, so

$$U_{\text{pair}}=\beta_{\text{cc}}\,\epsilon+\left[(1-\beta_{\text{cc}})^2-\beta_{\text{cc}}^2\right]\epsilon^2+O(\epsilon^3). \tag{15.43}$$

The common cause enters at first order and the independent failures at second order: whenever β_{cc} exceeds about ϵ, the shared cause sets the pair's unavailability. The standard working form of the β-factor result uses the unavailability of one pump with its total rate, $q=\lambda/(\lambda+\mu)=\epsilon-\epsilon^2+O(\epsilon^3)$:

$$U_{\text{pair}}\approx\beta_{\text{cc}}\,q+\left((1-\beta_{\text{cc}})\,q\right)^2, \tag{15.44}$$

the common-cause share of the unit's unavailability plus the square of its independent share. Its expansion is $\beta_{\text{cc}}\epsilon+[(1-\beta_{\text{cc}})^2-\beta_{\text{cc}}]\epsilon^2+O(\epsilon^3)$, and subtracting it from equation (15.43),

$$U_{\text{pair}}-\left[\beta_{\text{cc}}\,q+\left((1-\beta_{\text{cc}})\,q\right)^2\right]=\beta_{\text{cc}}(1-\beta_{\text{cc}})\,\epsilon^2+O(\epsilon^3), \tag{15.45}$$

a relative error of about $(1-\beta_{\text{cc}})\epsilon$. The independence baseline at the total rate is the row $\beta_{\text{cc}}=0$ of equation (15.42), q^2, and the independent part is its second term, q_{ind}^2.

Principle 15.1. *In a redundant pair the common cause enters the unavailability at first order and the independent failures at second order, $U_{\text{pair}}\approx\beta_{\text{cc}}\lambda/\mu+((1-\beta_{\text{cc}})\lambda/\mu)^2$; once the share exceeds the ratio λ/μ of the failure rate to the repair rate, the shared cause, not the redundancy, sets the pair's unavailability.*

The mean of the quantity under the posterior of the state is the sum (15.25) with $f=U_{\text{pair}}$, and the same sum gives the means of its terms. For the band and the threshold the derived quantity is carried on its logarithm, the scale on which a small probability is approximately linear in the state. Its gradient is assembled from the gradients of the two terms of equation (15.42). For a two-state process with failure rate x and repair rate μ, the unavailability $x/(x+\mu)$ has the derivative $x\mu/(x+\mu)^2$ with respect to $\ln x$, which is the unavailability times the availability; for the shared process, $\mathrm{d}U_{\text{cc}}/\mathrm{d}\ln(\beta_{\text{cc}}\lambda)=U_{\text{cc}}(1-U_{\text{cc}})$, and likewise for q_{ind}. The logs of the shared and the independent failure rates have the gradients of the doubles row and the singles row of the log-count Jacobian (Section 15.1.6), $(1,\ 1-\beta_{\text{cc}})^{\mathsf{T}}$ and $(1,\ -\beta_{\text{cc}})^{\mathsf{T}}$, so that

$$\frac{\partial U_{\text{cc}}}{\partial\mathbf{z}}=U_{\text{cc}}(1-U_{\text{cc}})\begin{bmatrix}1\\1-\beta_{\text{cc}}\end{bmatrix},\qquad\frac{\partial q_{\text{ind}}}{\partial\mathbf{z}}=q_{\text{ind}}(1-q_{\text{ind}})\begin{bmatrix}1\\-\beta_{\text{cc}}\end{bmatrix},$$

and, by equation (15.42),

$$\mathbf{g}=\frac{\partial\ln U_{\text{pair}}}{\partial\mathbf{z}}=\frac{1}{U_{\text{pair}}}\left[(1-q_{\text{ind}}^2)\frac{\partial U_{\text{cc}}}{\partial\mathbf{z}}+2q_{\text{ind}}(1-U_{\text{cc}})\frac{\partial q_{\text{ind}}}{\partial\mathbf{z}}\right], \tag{15.46}$$

evaluated at $\hat{\mathbf{z}}$. When the shared term carries nearly all of the unavailability, the gradient points nearly along the doubles row's direction $(1,\ 1-\beta_{\text{cc}})^{\mathsf{T}}$. The delta method of Chapter 4 then gives the variance budget of the log unavailability and, from the normal on it, the band and the threshold probability,

$$\sigma_{\ln U}^2=\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}=g_1^2\Sigma_{11}+g_2^2\Sigma_{22}+2g_1g_2\Sigma_{12},\qquad\Pr(U_{\text{pair}}<U^*\mid\mathbf{y})\approx\Phi\!\left(\frac{\ln U^*-\ln U_{\text{pair}}(\hat{\mathbf{z}})}{\sigma_{\ln U}}\right), \tag{15.47}$$

with the band $\exp(\ln U_{\text{pair}}(\hat{\mathbf{z}}) \pm 1.96\,\sigma_{\ln U})$ and $U^* = 10^{-4}$. The budget names the entry that limits the answer, which the next measurement needs.

The values beside them come from the exact posterior. The posterior mean of the pair's unavailability is the expectation of equation (15.42) over the joint posterior, evaluated on the grid, and likewise the two baselines. The band and the threshold probability need the distribution of a derived quantity of two entries, which is obtained by draws of the exact posterior: states are drawn from the grid with probabilities proportional to their posterior weights, or by the exact construction, the share from its marginal by the inverse of its cumulative distribution and the rate from its conditional gamma, and equation (15.42) is evaluated at each. By the first-order form (15.43) the mean is close to $\mathbb{E}[\beta_{\text{cc}}\lambda \mid \mathbf{y}]/\mu + \mathbb{E}[q^2_{\text{ind}} \mid \mathbf{y}]$, and the conditional gamma turns the first expectation into a one-dimensional integral, $\mathbb{E}[\beta_{\text{cc}}\lambda \mid \mathbf{y}] = \mathbb{E}\big[(a_{\lambda 0} + m)\beta_{\text{cc}}/(b_{\lambda 0} + (2 - \beta_{\text{cc}})T_{\text{tot}}) \mid \mathbf{y}\big]$, a hand check on the grid.

Building blocks. The MAP state $\hat{\mathbf{z}}$ and the covariance $\boldsymbol{\Sigma}$ of equation (15.20); the three unavailabilities of equations (15.41) and (15.42); the sum (15.25), the gradient (15.46) and the normal cumulative distribution Φ; and, for the exact values beside them, the grid posterior, its draws and the draws of the exact construction (Section 15.2.1).

The unavailability of the pair. For the pumps $\epsilon = \lambda/\mu = 1.57 \times 10^{-3}$ at the posterior mean rate, so any share above about 0.0016 makes the shared cause the larger term of equation (15.43). Table 15.7 evaluates equation (15.42) at the posterior mean rate $\lambda = 1.309 \times 10^{-4}\,\text{h}^{-1}$, within 0.1 % of $\hat{\lambda}$, for six shares. Each row is a share; the columns give the common-cause unavailability, the independent part q^2_{ind}, the pair's unavailability and the pair's downtime in hours per year. One row of the arithmetic, $\beta_{\text{cc}} = 0.10$: $U_{\text{cc}} = 1.309 \times 10^{-5}/(1.309 \times 10^{-5} + 0.08333) = 1.570 \times 10^{-4}$, $q_{\text{ind}} = 1.178 \times 10^{-4}/(1.178 \times 10^{-4} + 0.08333) = 1.411 \times 10^{-3}$, $q^2_{\text{ind}} = 1.99 \times 10^{-6}$, and $U_{\text{pair}} = 1.570 \times 10^{-4} + 1.99 \times 10^{-6} - 3 \times 10^{-10} = 1.590 \times 10^{-4}$.

Table 15.7: The pair's unavailability against the common-cause share at $\lambda = 1.309 \times 10^{-4}\,\text{h}^{-1}$ and $\mu = 1/12\,\text{h}^{-1}$: equation (15.42) and its terms (15.41). 0.1314 is the posterior mean share; the row $\beta_{\text{cc}} = 0$ is independence at the total rate, q^2.

Share β_{cc} [–]	U_{cc} [–]	q^2_{ind} [–]	U_{pair} [–]	Downtime (h per year)
0	0	2.46×10^{-6}	2.46×10^{-6}	0.022
0.05	7.85×10^{-5}	2.22×10^{-6}	8.07×10^{-5}	0.71
0.0657	1.03×10^{-4}	2.15×10^{-6}	1.053×10^{-4}	0.92
0.10	1.570×10^{-4}	1.99×10^{-6}	1.590×10^{-4}	1.39
0.1314	2.063×10^{-4}	1.86×10^{-6}	2.082×10^{-4}	1.82
0.20	3.140×10^{-4}	1.57×10^{-6}	3.156×10^{-4}	2.76

Against the row $\beta_{\text{cc}} = 0$, independence at the same total rate, a share of 0.05 raises the pair's unavailability by a factor of 33, a share of 0.1 by 65 and a share of 0.2 by 128. The independent part is 13 % of the total at a share of 0.01 and under 3 % at every share above 0.05.

The sum (15.25) with $f = U_{\text{pair}}$ gives the mean of the pair's unavailability, and with its terms the means of the shared term and of the independent part:

$$\mathbb{E}[U_{\text{pair}} \mid \mathbf{y}] \approx 2.62 \times 10^{-4}, \qquad \mathbb{E}[q^2_{\text{ind}} \mid \mathbf{y}] \approx 2.1 \times 10^{-6}, \tag{15.48}$$

a downtime of 2.30 h per year, of which $\mathbb{E}[U_{\text{cc}} \mid \mathbf{y}] \approx 2.60 \times 10^{-4}$ is the shared term; five nodes already give 2.62×10^{-4}. The independent part is 0.018 h per year; independence at the total rate is $\mathbb{E}[q^2 \mid \mathbf{y}] \approx 2.9 \times 10^{-6}$, 0.025 h per year. The mean is 114 times the 2.3×10^{-6} of the first reading. The values beside them come from the exact posterior. The posterior mean of the

pair's unavailability, the expectation of equation (15.42) over the joint posterior on the grid, and the posterior mean of its independent part are

$$\mathbb{E}[U_{\text{pair}} \mid \mathbf{y}] = 2.14 \times 10^{-4}, \qquad \mathbb{E}[q_{\text{ind}}^2 \mid \mathbf{y}] = 2.0 \times 10^{-6}, \tag{15.49}$$

a downtime of 1.87 h per year, of which $\mathbb{E}[U_{\text{cc}} \mid \mathbf{y}] = 2.12 \times 10^{-4}$ is the shared term. The independent part is 0.018 h per year; independence at the total rate is $\mathbb{E}[q^2 \mid \mathbf{y}] = 2.7 \times 10^{-6}$, 0.023 h per year. The posterior mean is 94 times the 2.3×10^{-6} of the first reading. The state route is higher by 0.48×10^{-4}, 0.43 h per year: the shared term is nearly linear in the share, and the mean of the share is 0.155 under the normal on ξ against 0.131 exactly.

Figure 15.3 draws the dependence on the share. In the upper panel the horizontal axis is the share and the vertical axis the pair's unavailability on a logarithmic scale; the solid curve is equation (15.42) at the posterior mean rate, the dashed horizontal line the posterior mean of the independent part, 2.0×10^{-6}, the dotted red line the limit 10^{-4} of the fifth question, and the vertical line the exact posterior mean share 0.131. The lower panel draws the posterior density of the share against the same horizontal axis, green for the exact marginal and dotted blue for the normal on ξ. The curve crosses the limit at a share of 0.062, below which the posterior of the state puts $\Phi(-1.118) = 0.13$ of its probability (0.22 exactly).

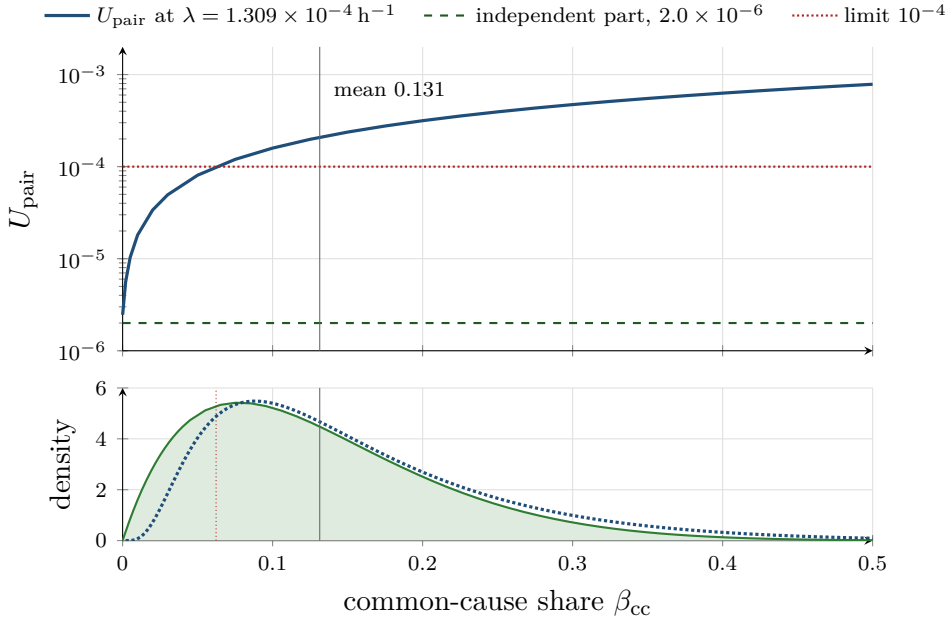


Figure 15.3: Upper panel: the pair's unavailability, equation (15.42), against the common-cause share at $\lambda = 1.309 \times 10^{-4}\,\text{h}^{-1}$ and $\mu = 1/12\,\text{h}^{-1}$, with the posterior mean of the independent part and the limit 10^{-4}, which the curve crosses at $\beta_{\text{cc}} = 0.062$. Lower panel: the posterior density of the share, the marginal of equation (15.13) (green) and the normal $\mathcal{N}(-1.874,\ 0.748^2)$ on ξ carried to the share (dotted blue), with the same vertical line at the mean and a dotted red line at the crossing 0.062.

The band and the threshold. At the MAP state, where $U_{\text{cc}} = 2.087 \times 10^{-4}$, $q_{\text{ind}} = 1.359 \times 10^{-3}$ and $U_{\text{pair}} = 2.105 \times 10^{-4}$, the gradient (15.46) is

$$\begin{aligned} \mathbf{g} = \frac{\partial \ln U_{\text{pair}}}{\partial \mathbf{z}} &= \frac{1}{U_{\text{pair}}}\left[(1 - q_{\text{ind}}^2)\frac{\partial U_{\text{cc}}}{\partial \mathbf{z}} + 2q_{\text{ind}}(1 - U_{\text{cc}})\frac{\partial q_{\text{ind}}}{\partial \mathbf{z}}\right] \\ &= 0.9910 \begin{bmatrix} 1 \\ 0.8670 \end{bmatrix} + 0.0175 \begin{bmatrix} 1 \\ -0.1330 \end{bmatrix} = \begin{bmatrix} g_1 \\ g_2 \end{bmatrix} = \begin{bmatrix} 1.009 \\ 0.857 \end{bmatrix}. \end{aligned}$$

The gradient points almost along the doubles row's direction: the shared term carries 99 % of the unavailability, and the independent term adds a small component along the singles row's direction. Keeping the shared term alone gives the approximate gradient $(1,\ 1-\beta_{\mathrm{cc}})^{\mathsf{T}} = (1,\ 0.867)^{\mathsf{T}}$. With $\boldsymbol{\Sigma}$ of equation (15.20), the delta method of Chapter 4 gives

$$\sigma^2_{\ln U} = \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} = \underbrace{g_1^2\Sigma_{11}}_{0.087} + \underbrace{g_2^2\Sigma_{22}}_{0.411} + \underbrace{2g_1g_2\Sigma_{12}}_{0.053} = 0.551, \qquad \sigma_{\ln U} = 0.742, \tag{15.50}$$

about the center $\ln U_{\mathrm{pair}}(\hat{\mathbf{z}}) = -8.466$, the unavailability 2.11×10^{-4} at the MAP state; the approximate gradient $(1,\ 0.867)^{\mathsf{T}}$ would give 0.559, a standard deviation of 0.748. The variance budget names the entry that limits the answer: three quarters of the variance is the share's term, a sixth the rate's, and the cross term, positive because both entries of $\mathbf{g}$ and Σ_{12} are positive, the sign the singles row put in $\mathbf{J}^{\mathsf{T}}\mathbf{J}$, adds a tenth. By equation (15.47) the band is

$$\begin{gathered} U_{\mathrm{pair}} \in \exp(-8.466 \pm 1.96 \times 0.742) = [4.9 \times 10^{-5},\ 9.0 \times 10^{-4}] \\ \text{with } 95\,\% \text{ posterior probability}, \quad \text{median } 2.11 \times 10^{-4}, \end{gathered} \tag{15.51}$$

and the threshold probability

$$\Pr(U_{\mathrm{pair}} < 10^{-4} \mid \mathbf{y}) \approx \Phi\left(\frac{\ln 10^{-4} + 8.466}{0.742}\right) = \Phi(-1.003) = 0.16. \tag{15.52}$$

In hours per year the band is 0.43 to 7.9 h, the median 1.84 h and the limit 0.88 h. The value 0.16 falls short of the assurance $\kappa = 0.90$: the limit is not demonstrated.

The values beside them come from the exact posterior, by draws of the grid posterior: 2×10^5 grid cells are drawn with probabilities proportional to their posterior weights, by a pseudo-random generator seeded with 15, and equation (15.42) is evaluated at each. The results are

$$U_{\mathrm{pair}} \in [2.5 \times 10^{-5},\ 6.3 \times 10^{-4}] \text{ with } 95\,\% \text{ posterior probability}, \quad \text{median } 1.73 \times 10^{-4}, \tag{15.53}$$

and

$$\Pr(U_{\mathrm{pair}} < 10^{-4} \mid \mathbf{y}) = 0.26. \tag{15.54}$$

In hours per year the band is 0.22 to 5.5 h and the median 1.5 h; the standard deviation of $\ln U_{\mathrm{pair}}$ over the draws is 0.822. The threshold probability 0.26 exceeds the 0.22 below the crossing of Figure 15.3 because the rate is uncertain too. The state route is lower by 0.10 on the threshold, and its band is higher by a factor of about two at the lower end and 1.4 at the upper, because the normal on ξ leaves out the long tail of the exact logit toward small shares. Both values fall short of the assurance 0.90, and the gap does not change the decision.

15.4 The checks (step 5)

Every answer of the chapter has a second route beside the reported one, and with one double in the record the prior on the share carries much of its information. The pump pair has no physics row and no exact constraint, so no limit $\sigma_p \to 0$ of a model tolerance arises and no limit table is needed; the bounds on the rate and the share are not rows either, and the log rate and the logit handle them.

15.4.1 The Laplace and Beta routes for the share

The Laplace marginals. The normal $\xi \sim \mathcal{N}(-1.874,\ 0.748^2)$ of equation (15.20) maps through the inverse logit, equation (15.23), to a median share of 0.133 and a central 95 % interval

[0.034, 0.399], against the exact median 0.116 and interval [0.018, 0.329]; it puts $\Pr(\beta_{cc} > 0.05)$ at 0.924 against the exact 0.847. The exact posterior of ξ has mean -2.115, standard deviation 0.840 and skewness -0.68: it is wider than the Laplace normal and skewed toward small shares, the shape of the logit of a Beta variable with a small first parameter, as the trigamma comparison of Chapter 8 showed for a gamma. The curvature term $\mathbf{C}$ moves the Laplace spread of the logit toward the exact one, from 0.732 to 0.748 against 0.840. The rate is better served: the normal on z_1 gives the median $1.31 \times 10^{-4}\,\mathrm{h}^{-1}$ and the interval $[0.74,\ 2.32] \times 10^{-4}\,\mathrm{h}^{-1}$, against the exact $[0.67,\ 2.16] \times 10^{-4}\,\mathrm{h}^{-1}$, with $\sqrt{\Sigma_{11}} = 0.292$ against the exact 0.298 and the correlation 0.14 against 0.12.

The Beta approximation. The Beta approximation (15.14), Beta(2, 18), has mean 0.100, median 0.087 and central 95 % interval [0.013, 0.260]. The tilt of equation (15.15) has $s = 12 \times 43\,800/97\,600 = 5.39$, a factor of about $\mathrm{e}^{5.39 \times 0.13} = 2.0$ at the posterior mean, and the tilted Beta$(2, 18)\,\mathrm{e}^{5.39\beta_{cc}}$ has mean 0.128 and interval [0.018, 0.320], close to the exact 0.131 and [0.018, 0.329]. On the logit the Beta approximation has the exact spread, standard deviation 0.838 against 0.840, and the whole distribution is shifted toward smaller shares, its mean low by 0.03.

Figure 15.4 compares the four densities of the share. The horizontal axis is the share and the vertical axis the density. The dotted gray curve is the prior Beta$(1, 9)$; the solid green curve the exact marginal (15.13), with the area to the right of 0.05 shaded; the dashed blue curve the Beta approximation (15.14); and the dash-dotted red curve the Laplace normal on ξ, mapped from the logit to the share. The record has moved the prior's mass away from zero, since a double has occurred, but not far: the exact posterior peaks near 0.08 and keeps a long upper tail. The Beta approximation is taller and shifted to the left, the missing tilt. The Laplace density has the right peak but a lower tail that is too thin near zero and an upper tail that is too heavy, the shape of a normal placed on a skewed logit.

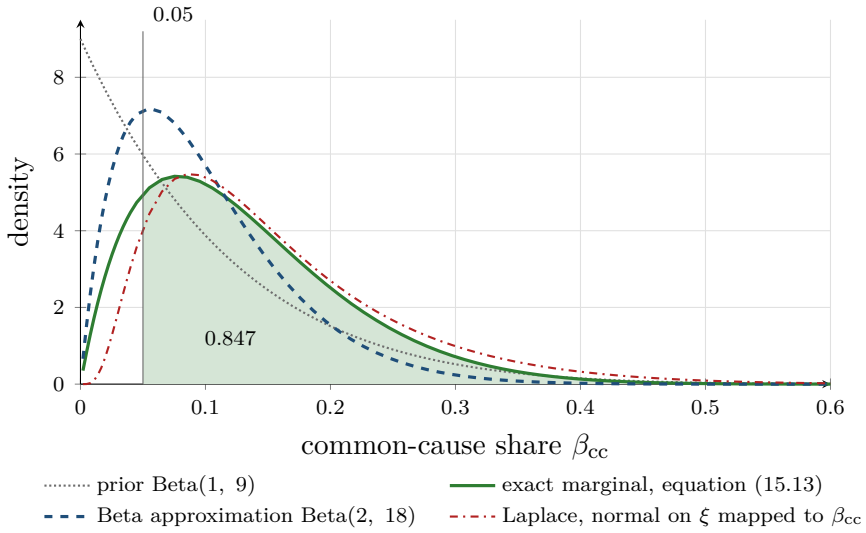


Figure 15.4: Four densities of the common-cause share. Prior Beta$(1, 9)$, density $9(1-\beta_{cc})^8$; exact marginal, equation (15.13) normalized, with the area $\Pr(\beta_{cc} > 0.05 \mid \mathbf{y}) = 0.847$ shaded; the Beta approximation (15.14); and the Laplace normal $\mathcal{N}(-1.874,\ 0.748^2)$ on ξ of equation (15.20), density $\varphi((\xi + 1.874)/0.748)/(0.748\,\beta_{cc}(1-\beta_{cc}))$ with φ the standard normal density.

Table 15.8 sets the routes side by side: each row is a summary, and the columns give it under the prior, the exact posterior on the grid, the Beta approximation and the Laplace normal.

The state route serves the rate closely and the share less well: its band on the share is

Table 15.8: Posterior summaries of the share and the rate per pump by three routes, with the prior. Exact: the grid posterior of equation (15.11), which agrees with the quadrature of equation (15.13); Beta approximation: equation (15.14); Laplace: the normal $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of equation (15.20) mapped back. Rates in $10^{-4}\,\mathrm{h}^{-1}$.

Summary	Prior	Exact	Beta approx.	Laplace
Mean of β_{cc}	0.100	0.131	0.100	0.155
Median of β_{cc}	0.074	0.116	0.087	0.133
Mode of β_{cc}	0	0.078	0.056	0.088
95 % interval of β_{cc}	[0.003, 0.336]	[0.018, 0.329]	[0.013, 0.260]	[0.034, 0.399]
$\Pr(\beta_{\mathrm{cc}} > 0.05)$	0.630	0.847	0.755	0.924
Std. dev. of ξ	–	0.840	0.838	0.748
Mean of λ	2.00	1.31	–	1.36
95 % interval of λ	[0.24, 5.57]	[0.67, 2.16]	–	[0.74, 2.32]

displaced upward, the normal placed on a logit whose exact posterior is skewed toward small shares. It is the route of every answer of the chapter, of the variance budget and of the preposterior, and the gap is stated with each. The Beta approximation is not used for any answer; the tilted version recovers most of its error.

Common mistake

Estimating a common-cause share by counting events, doubles against singles, as if each event were a trial with probability β_{cc}. A single failure involves one unit and a double two, so the counts are trials of the event share $\theta = \beta_{\mathrm{cc}}/(2 - \beta_{\mathrm{cc}})$; the Beta posterior $\mathrm{Beta}(2, 18)$ puts the mean share at 0.100 against the mean 0.155 of the posterior of the state and the exact 0.131, and with a long record it would converge to about half the true share.

15.4.2 The grid, the draws and the closed forms

The grid against the quadrature. The two-dimensional grid and the one-dimensional quadrature of the marginal (15.13) agree on every summary of the share to three decimals and on the evidence of M_1 to four. Several numbers of the chapter are computed on the grid (the posterior mean of the unavailability, the event posterior, the draws of the pair), and the agreement shows that the grid's discretization error causes no discrepancy in any of them at the digits given.

The two sets of draws. A million draws from the exact construction of Section 15.2.1, the share from its marginal by the inverse of its cumulative distribution and the rate from the conditional gamma (15.12), need no grid. They give the band $[2.5 \times 10^{-5},\ 6.3 \times 10^{-4}]$ and the threshold probability 0.258, the same as the 2×10^5 grid draws of Section 15.3.4 to the digits given.

The mean against the plug-in and the first-order form. On the state route the plug-in at the MAP state, 2.11×10^{-4}, is lower than the mean 2.62×10^{-4} of the sum (15.25) by a fifth, the weight the normal on ξ puts on large shares. On the exact posterior the plug-in value at the posterior means, 2.08×10^{-4}, is lower than the posterior mean 2.14×10^{-4} by 2.5 %, because the share and the rate are positively correlated: by the conditional gamma $\mathbb{E}[\beta_{\mathrm{cc}}\lambda \mid \mathbf{y}] = 1.764 \times 10^{-5}\,\mathrm{h}^{-1}$ against $\mathbb{E}[\beta_{\mathrm{cc}}]\mathbb{E}[\lambda] = 1.720 \times 10^{-5}\,\mathrm{h}^{-1}$. The first-order form gives $12 \times 1.764 \times 10^{-5} + 2.0 \times 10^{-6} = 2.14 \times 10^{-4}$, the grid mean. Chapter 13 gives the scale of the effect. There, sharing a posterior coupled units and opened a gap between the plug-in and the predictive answer of tens of percent; here the physical shared cause changes the answer by two

orders of magnitude, and the statistical coupling between the share and the rate is a correction of a few percent on top of it.

The working form. The working form (15.44) with $q = 1.568 \times 10^{-3}$ gives, at the six shares of Table 15.7, 2.46×10^{-6}, 8.06×10^{-5}, 1.052×10^{-4}, 1.588×10^{-4}, 2.079×10^{-4} and 3.152×10^{-4}, within 0.15 % of the exact value in every row: the relative error $(1 - \beta_{cc})\epsilon$ of equation (15.45) is 1.4×10^{-3}, and the working form understates by a seventh of a percent.

The coincidence approximations. The coincidence rate of Section 15.1.6 needs $\lambda t_w \ll 1$, here 10^{-3} at the window of 8 h, and takes the factor $1 - q$ on it as one, with $q \approx 1.6 \times 10^{-3}$ the probability that the second pump is under repair. Overlapping outages, the second pump failing while the first is under repair but more than t_w after it, are expected about $2\lambda^2 \times \text{MTTR} \times T_{tot} = 2 \times (1.31 \times 10^{-4})^2 \times 12 \times 43\,800 = 0.018$ times in the record, and the record shows none.

15.4.3 The delta method on the log unavailability

The band and the threshold probability of Section 15.3.4 come from the delta method on $\ln U_{pair}$, equation (15.50), and this subsection sets them against the exact draws.

Figure 15.5 sets the two routes side by side. The horizontal axis is the pair's unavailability on a logarithmic scale, the vertical axis the density of its base-ten logarithm. The solid green curve is the histogram of the million exact draws of Section 15.4.2, with the area below the limit 10^{-4} shaded; the dashed blue curve is the normal of equation (15.50), the reported route, on the same scale; the vertical red line is the limit. The exact posterior is skewed to the left on this scale, a long tail toward small unavailabilities that comes from the long lower tail of the logit of the share; the normal is narrower and centered to the right of it, and it leaves less probability below the limit.

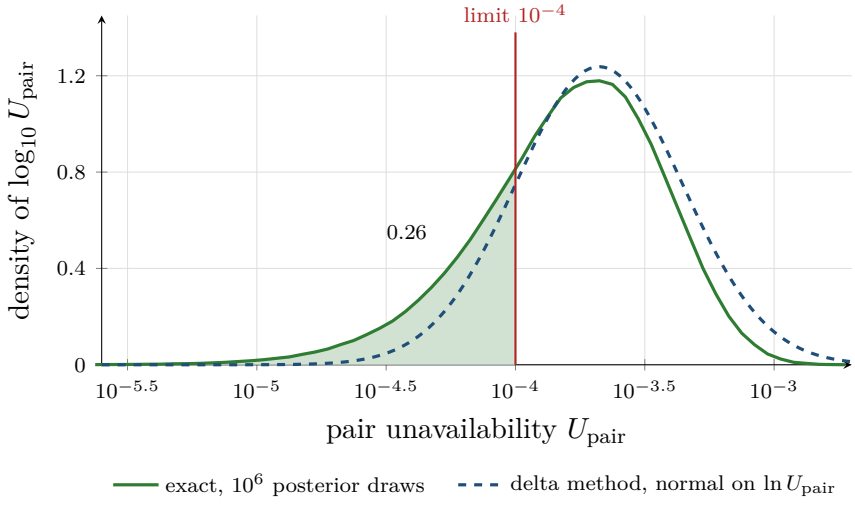


Figure 15.5: The posterior of the pair's unavailability. Exact: histogram of $\log_{10} U_{pair}$ over 10^6 draws of the exact posterior, in bins of 0.05, with $\Pr(U_{pair} < 10^{-4} \mid \mathbf{y}) = 0.26$ shaded. Delta method: $\ln U_{pair} \sim \mathcal{N}(-8.466,\ 0.742^2)$, equation (15.50), drawn as the normal on $\log_{10} U_{pair}$ with mean -3.677 and standard deviation 0.322, which puts 0.16 below the limit.

Against the exact $[2.5 \times 10^{-5},\ 6.3 \times 10^{-4}]$ and 0.26, the delta method's $[4.9 \times 10^{-5},\ 9.0 \times 10^{-4}]$ and 0.16 differ by 0.10 on the threshold and by a factor of two at the lower end of the band. The gap is not in the linearization, since $\ln U_{pair}$ is almost linear in the state, but in the Gaussian

beneath it: the exact standard deviation of $\ln U_{\text{pair}}$ is 0.822, wider than 0.742, and its distribution is skewed toward small values, both inherited from the logit of the share (Section 15.4.1), which carries three quarters of the budget and is informed by one double. Both routes leave the threshold far below the assurance 0.90. With more doubles the logit's posterior becomes more nearly normal and the two routes converge, but slowly: Exercise 15.6 repeats the comparison for a second double, where the delta method gives 0.03 against the exact 0.07.

Table 15.9 collects the band and the threshold: each row is a summary of the pair's unavailability, and the columns give it from the delta method and the sum over the state, the reported route, and from the exact draws.

Table 15.9: The band and the threshold of the pair's unavailability. State route: $\ln U_{\text{pair}} \sim \mathcal{N}(-8.466,\ 0.742^2)$, equation (15.50), and the sum (15.25) for the mean; exact: draws of the grid posterior.

Summary	State route	Exact
Mean, sum (15.25) [–]	2.62×10^{-4}	2.14×10^{-4}
Mean of the normal on $\ln U_{\text{pair}}$ [–]	2.77×10^{-4}	–
Median [–]	2.11×10^{-4}	1.73×10^{-4}
95 % band [–]	$[0.49,\ 9.0] \times 10^{-4}$	$[0.25,\ 6.3] \times 10^{-4}$
Std. dev. of $\ln U_{\text{pair}}$ [–]	0.742	0.822
$\Pr(U_{\text{pair}} < 10^{-4} \mid \mathbf{y})$ [–]	0.16	0.26
Mean downtime, sum (h per year)	2.30	1.87
Mean downtime, normal on $\ln U_{\text{pair}}$ (h per year)	2.43	–

15.4.4 The coincidence window and the evidence

The closed form. With $\tilde{\lambda} = 12/97\,600\,\text{h} = 1.23 \times 10^{-4}\,\text{h}^{-1}$ and $E_\beta = 0.0499$ by quadrature, equation (15.33) gives $B_{10} \approx 0.0499/(2 \times 1.23 \times 10^{-4} \times t_{\text{w}}) = 203\,\text{h}/t_{\text{w}}$, so that

$$\ln B_{10} \approx 5.31 - \ln(t_{\text{w}}/1\,\text{h}), \tag{15.55}$$

5.31 at one hour, 3.93 at four, 3.23 at eight, 2.14 at a day and 1.44 at two days, against the quadratures 5.32, 3.93, 3.25, 2.18 and 1.52 of Table 15.4. At the window of 8 h the closed form (15.31) gives $\ln p(\mathbf{y} \mid M_0) = -7.861$ against the quadrature's -7.875.

Figure 15.6 draws the log Bayes factor against the coincidence window. The horizontal axis is the window on a logarithmic scale from half an hour to two days, the vertical axis the log Bayes factor of the shared-cause drawing against independence, in nats. The solid curve is the quadrature, the dashed line the closed form (15.55), and the three dots mark the windows of one hour, one shift and one day. The two horizontal dotted lines mark 1 and 3 nats, a factor of about 2.7 and of 20.

On the logarithmic axis the curve is a straight line of slope $-k_2 = -1$: doubling the window doubles the chance explanation of the double and takes $\ln 2 = 0.69$ nats from the evidence. One double in five years is evidence of more than 3 nats for a shared cause for windows of a shift or less, and it weakens to about 2 nats for a window of a day and 1.5 for two days. The quadrature lies above the closed form by 0.04 nats at 24 h and by 0.08 at 48 h, and bends away from it at longer windows, as the coincidence term $\exp(-2\lambda^2 t_{\text{w}} T_{\text{tot}})$ that the closed form dropped begins to penalize the independence drawing; the evidence would fall to zero only at a window of about 350 h, two weeks, where a coincidence is no longer improbable. The conclusion that a shared cause acted is therefore robust to the window over the range of practical logs, while its strength is not, and the window has to be stated with the factor.

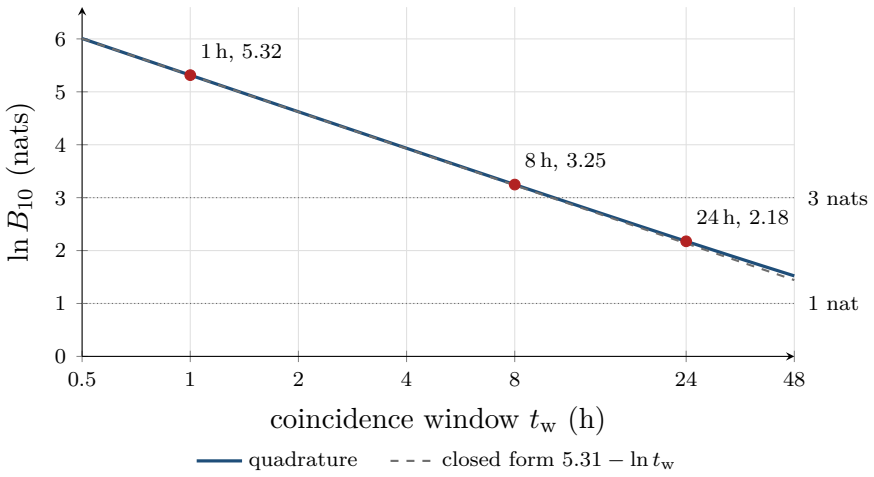


Figure 15.6: The log Bayes factor of the shared-cause drawing against independence as a function of the coincidence window, for the record $k_1 = 9$, $k_2 = 1$, $T_{tot} = 43\,800$ h and the priors Gamma(2, 10 000 h) and Beta(1, 9), windows $|t_1 - t_2| < t_w$. Solid: the quadrature of equation (15.30) against that of the likelihood (15.8) averaged over the prior. Dashed: the closed form $\ln B_{10} = 5.313 - \ln(t_w/1\,\text{h})$, equation (15.55).

The coincidences in the shared-cause drawing. Including the coincidences between the independent parts in the shared-cause drawing, at the rate $(1 - \beta_{cc})^2 c_w$, raises the exact $\ln B_{10}$ to 5.32, 3.26 and 2.22 at the windows of one hour, one shift and one day, changes of at most 0.04, and moves the posterior mean share from 0.131 to 0.130 at 8 h and to 0.127 at 24 h. The exact event posterior of Section 15.3.3 uses that posterior; two simplifications give nearly the same number. The chapter's posterior (15.11), which leaves the coincidences out of the doubles row, gives 0.978 at 8 h against 0.976; replacing the coincidence rate under M_1 by the independence drawing's $2\lambda^2 t_w$, the factor $(1 - \beta_{cc})^2$ dropped, and averaging over the chapter's posterior gives 0.974. At 24 h the two give 0.940 and 0.930 against 0.932, at 48 h 0.893 and 0.875 against 0.873, and at one hour all three are 0.997. The plug-in value at the posterior means, 0.988 at 8 h, is too high, because the draws with small shares, for which a coincidence is a real alternative, pull the average down. The state route, on the rows of $\hat{\mathbf{z}}$ that leave the coincidences out, is compared most directly with the chapter's posterior: 0.984, 0.955, 0.916 and 0.998 at 8, 24, 48 and 1 h against 0.978, 0.940, 0.893 and 0.997.

The Laplace evidences. The Laplace evidences of equation (15.29) give $\ln p(\mathbf{y} \mid M_1) = -4.676$ against the quadrature's -4.628, and $\ln B_{10} = 5.27$, 3.21 and 2.13 at the three windows, within 0.05 of the quadratures. At its own rate, by the fit (15.21), the independence drawing expects 0.0124 coincidences in five years at the 8 h window, so that one observed double is an event of probability $0.0124\,e^{-0.0124} = 0.012$, about one in 82; the Bayes factor weighs that against the probability of one shared event under the prior on the share.

> **Principle 15.2.** *The evidence for a common cause from simultaneous failures is set by the coincidence window: with k_2 doubles, each doubling of the window takes $k_2 \ln 2$ nats from the log Bayes factor. The window is at least the resolution of the record's time stamps and the spread of a shared cause's effects, and it is fixed before the record is read.*

15.4.5 Sensitivity to the prior

The share is informed by one double, and the column of the Jacobian in Section 15.2.3 shows the prior Beta(1, 9) supplying 1.15 of the 1.82 in the Hessian entry of the logit. Table 15.10

repeats the analysis under four alternative priors, on the state route and on the exact one. Each row is a pair of priors; each cell gives the state-route value with the exact one in parentheses. The columns give the posterior mean of the share and its central 95 % interval, from the fit of $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$ under those priors and from the marginal (15.13), the posterior mean of the pair's unavailability and its threshold probability, from the sum (15.25) and the delta method and from a million draws of the exact construction, and the log Bayes factor at the window of 8 h, from the Laplace evidences and by quadrature. The first row is the chapter's; the second a generic share of 0.05 worth twenty events; the third a share of 0.1 worth twenty events, twice as firm as the chapter's; the fourth the uniform prior on the share; and the fifth a rate prior with half the chapter's weight. One row of the arithmetic, the uniform prior: the marginal becomes $\beta_{\mathrm{cc}}^{1}(1-\beta_{\mathrm{cc}})^{9}(97\,600-43\,800\,\beta_{\mathrm{cc}})^{-12}$, of mean 0.256, while the MAP state moves to $\hat{\xi}=-1.014$, of mean share 0.289.

Table 15.10: Sensitivity of the pump pair's answers to the priors; record $k_1 = 9$, $k_2 = 1$ in 43 800 h. Each cell: state route (exact). Share: $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ (exact marginal); unavailability: the sum (15.25) and the delta method (10^6 draws of the exact construction); $\ln B_{10}$ at $t_{\mathrm{w}} = 8$ h: Laplace (quadrature).

Priors	Mean of β_{cc} [–]	95 % interval of β_{cc} [–]	$\mathbb{E}[U_{\mathrm{pair}}]$ [–]	$\Pr(U_{\mathrm{pair}} < 10^{-4})$ [–]	$\ln B_{10}$ [nats]
Gamma(2, 10 000 h), Beta(1, 9)	0.155 (0.131)	[0.034, 0.399] ([0.018, 0.329])	2.62×10^{-4} (2.14×10^{-4})	0.16 (0.26)	3.21 (3.25)
Gamma(2, 10 000 h), Beta(1, 19)	0.097 (0.080)	[0.020, 0.269] ([0.010, 0.210])	1.60×10^{-4} (1.27×10^{-4})	0.38 (0.49)	2.85 (2.89)
Gamma(2, 10 000 h), Beta(2, 18)	0.135 (0.119)	[0.040, 0.311] ([0.027, 0.267])	2.24×10^{-4} (1.91×10^{-4})	0.16 (0.24)	3.36 (3.38)
Gamma(2, 10 000 h), Beta(1, 1)	0.289 (0.256)	[0.074, 0.623] ([0.041, 0.560])	5.27×10^{-4} (4.47×10^{-4})	0.02 (0.07)	2.69 (2.74)
Gamma(1, 5000 h), Beta(1, 9)	0.154 (0.130)	[0.034, 0.397] ([0.018, 0.327])	2.53×10^{-4} (2.05×10^{-4})	0.17 (0.28)	3.22 (3.26)

The rate prior hardly matters: halving its weight moves no answer by more than a few percent, since ten events inform the rate. The share prior matters a great deal. Its posterior mean runs from 0.097 (0.080) under the generic Beta(1, 19) to 0.289 (0.256) under the uniform prior, and the pair's mean unavailability with it, from 1.60 (1.27) to 5.27 (4.47) $\times 10^{-4}$: one double cannot overrule what was assumed about the share before it. Two conclusions survive every prior, on both routes. The pair's unavailability is at least 69 (56) times the 2.3×10^{-6} that independence gives, so the shared cause dominates under every prior examined; and the limit of 10^{-4} holds with probability at most 0.38 (0.49), so it is demonstrated under none. The evidence for a shared cause stays between 2.7 and 3.4 nats. The size of the share, and with it the factor by which the pair falls short, is a property of the prior as much as of the record, which is why the prior is stated with the answer and why a record that adds doubles slowly cannot settle it (Section 15.6).

Principle 15.3. *A share informed by one or two events is mostly prior. The prior on a common-cause share is documented with its source and varied over the defensible range, and a conclusion is reported only when it survives that range, here the dominance of the shared cause and the undemonstrated limit; the size of the share is reported with its prior.*

15.4.6 Redundancy and the checks together

Redundancy and dispersion. The two count rows are as many as the entries of the state, and the maximum-likelihood state reproduces both counts exactly (Section 15.1.6), so the record

has no redundant row and no χ^2 check of the counts against the model is available from it. The residuals at the MAP state, $\mathbf{r}(\hat{\mathbf{z}})$ of Section 15.2.3, have the sum of squares 0.62 and measure only the pull of the two priors against the counts. The checks of the model are the assumptions of Section 15.6.7, each tested as the record grows.

Table 15.11 collects the checks: each row is one quantity, the second and third columns give the second route and its value, and the fourth and fifth the calculation it is set against and its value, the state route where the row names it.

Table 15.11: The checks of step 5: each second route against the calculation it is set against. Grid: the 1200×1200 posterior of equation (15.11); quadrature: one-dimensional; draws: of the grid posterior and of the exact construction; window $t_{\mathrm{w}} = 8\,\mathrm{h}$. All quantities dimensionless except the rate, in $10^{-4}\,\mathrm{h}^{-1}$.

Quantity	Second route	Value	Set against	Value
Mean of β_{cc}	quadrature (15.13)	0.131	grid	0.131
Std. dev. of ξ	grid	0.840	state route	0.748
95 % interval of β_{cc}	grid	[0.018, 0.329]	state route	[0.034, 0.399]
95 % interval of λ	grid	[0.67, 2.16]	state route	[0.74, 2.32]
Mean of β_{cc}	Beta approximation	0.100	grid	0.131
95 % interval of β_{cc}	Beta approximation	[0.013, 0.260]	grid	[0.018, 0.329]
$\mathbb{E}[U_{\mathrm{pair}} \mid \mathbf{y}]$	grid	2.14×10^{-4}	state route	2.62×10^{-4}
$\mathbb{E}[U_{\mathrm{pair}} \mid \mathbf{y}]$	plug-in at the means	2.08×10^{-4}	grid	2.14×10^{-4}
U_{pair} at $\beta_{\mathrm{cc}} = 0.1314$	working form (15.44)	2.079×10^{-4}	equation (15.42)	2.082×10^{-4}
$\Pr(U_{\mathrm{pair}} < 10^{-4} \mid \mathbf{y})$	2×10^5 grid draws	0.26	state route, delta method	0.16
$\Pr(U_{\mathrm{pair}} < 10^{-4} \mid \mathbf{y})$	10^6 exact draws	0.258	2×10^5 grid draws	0.26
$\ln p(\mathbf{y} \mid M_0)$	closed form (15.31)	−7.861	quadrature	−7.875
$\ln p(\mathbf{y} \mid M_0)$	quadrature	−7.875	state route	−7.881
$\ln p(\mathbf{y} \mid M_1)$	quadrature	−4.628	state route	−4.676
$\ln B_{10}$	closed form (15.55)	3.23	quadrature	3.25
$\ln B_{10}$	quadrature	3.25	state route	3.21
Event posterior	coincidences admitted	0.976	state route	0.984
Event posterior	chapter's posterior	0.978	coincidences admitted	0.976
Event posterior	plug-in at the means	0.988	coincidences admitted	0.976

15.5 The report (step 6)

The record overturns the first reading of the log. Two independent pumps with eleven failures in 87 600 pump-hours would be down together 2.3×10^{-6} of the time, forty-four times below the limit; the mean unavailability of the pair is 2.62×10^{-4} (2.14×10^{-4} exactly), 114 (94) times that and more than twice the limit, and the probability that the pair meets the limit of 10^{-4} is 0.16 (0.26). The one double failure is the reason: it is strong evidence that a shared cause acts on both pumps, and a shared cause enters the pair's unavailability at first order, where independent failures enter at second. The answers are computed from the posterior of the state, with the exact grid and its draws beside them, and every prior examined in Section 15.4 leaves the limit undemonstrated on both routes.

15.5.1 The answers and their bands

Table 15.12 gives the answers. Each row is a question of Table 15.2; the second column gives the answer from the posterior of the state, $\mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$ of equation (15.20); the third the exact value; and the last the band of the reported answer with the comparisons that qualify it, the exact band in parentheses and the range over the priors of Table 15.10 among them.

Table 15.12: The answers to the questions asked of the pump pair's record, from the posterior of the state, with the exact value beside each. Bands are central 95 % posterior intervals of the reported route, exact in parentheses; the coincidence window is 8 h unless stated; the naive value is two independent pumps at the rate of eleven failures in 87 600 pump-hours.

Question	Answer (state route)	Exact	Band and comparison
What share of the pumps' failures is a common cause, and how well is it known?	median 0.133, mean 0.155; $\Pr(\beta_{\mathrm{cc}} > 0.05) = 0.92$	mean 0.131, median 0.116; 0.85	[0.034, 0.399] ([0.018, 0.329]); Beta approximation [0.013, 0.260]; mean 0.097 to 0.289 (0.080 to 0.256) over the priors
Is there a common cause at all, or was the double a coincidence?	$B_{10} = 25$ ($\ln B_{10} = 3.21$); $\Pr(M_1 \mid \mathbf{y}) = 0.961$ at even odds	$B_{10} = 26$ ($\ln B_{10} = 3.25$); 0.963	$\ln B_{10} = 5.27$ (5.32) at $t_{\mathrm{w}} = 1$ h, 2.13 (2.18) at 24 h, 3.89 (3.93) for a log binned by shift; 2.7 to 3.4 over the priors
Was the double a common cause, so that one cause took two units out?	0.95 over both drawings; 0.984 within M_1	0.94; 0.976	within M_1, 0.998 (0.997) at 1 h and 0.955 (0.93) at 24 h
What is the pair's unavailability with the common cause, and what would it be without?	mean 2.62×10^{-4} (2.30 h per year), median 2.11×10^{-4}	mean 2.14×10^{-4} (1.87 h per year), median 1.73×10^{-4}	$[4.9 \times 10^{-5},\ 9.0 \times 10^{-4}]$ ($[2.5 \times 10^{-5},\ 6.3 \times 10^{-4}]$); 114 (94) times the naive 2.3×10^{-6}; 90 (80) times independence at the total rate, $\mathbb{E}[q^2] = 2.9 \times 10^{-6}$ (2.7×10^{-6}); 127 (107) times the independent part, $\mathbb{E}[q^2_{\mathrm{ind}}] = 2.1 \times 10^{-6}$ (2.0×10^{-6})
Is the pair's unavailability below 10^{-4}, with assurance 0.90?	$\Pr(U_{\mathrm{pair}} < 10^{-4} \mid \mathbf{y}) = 0.16$; not demonstrated	0.26; not demonstrated	0.02 to 0.38 (0.07 to 0.49) over the priors

The report line. Over five years the two chilled-water pumps had nine single failures and one double failure. Under the β-factor model the share of each pump's failures that comes from a shared cause has the median 0.13, with a 95 % band from 0.03 to 0.40 (from the exact posterior, median 0.12, mean 0.13 and band from 0.02 to 0.33). With a double defined as two failures logged within eight hours of each other, the double failure is strong evidence for a shared cause against a coincidence, a Bayes factor of 25 (26) and a probability of 0.96 for the shared-cause drawing on both routes; over both drawings the double was a common cause, one cause taking the two pumps out, with probability 0.95 (0.94), and within the shared-cause drawing with probability 0.98 on both. The pair's unavailability has mean 2.6×10^{-4} (2.1×10^{-4}), about 2.3 h (1.9 h) of lost cooling supply per year, 114 (94) times the 2.3×10^{-6} of the first reading, which treats the pumps as independent; its band runs from 4.9×10^{-5} to 9.0×10^{-4} (2.5×10^{-5} to 6.3×10^{-4}), and the probability that it is below 10^{-4} is 0.16 (0.26), short of the assurance 0.90. A limit of 10^{-4} is not demonstrated, under the chosen priors or any examined beside them, on either route, and the redundancy does not deliver what independence promises.

Common mistake

Reading the posterior that the observed double was a common cause, 0.984 (0.976 exactly), as the probability that a shared cause exists, or the Bayes factor of 25 (26) as the odds that this particular event was one. The first is computed inside the shared-cause drawing and assumes the mechanism; the second compares the drawings and assigns no event. The probability that the double was a common cause, with both drawings admitted at even prior odds, is 0.95 (0.94).

15.5.2 The limit decision

The decision follows from the last row of Table 15.12. The cooling design asks for an unavailability below 10^{-4} at the assurance 0.90, and the record supports it with probability 0.16 (0.26 exactly) under the chosen priors and at most 0.38 (0.49) under any prior of Table 15.10: the limit is not demonstrated, on either route, and the two routes reach the same decision. The first reading, 2.3×10^{-6}, is not an estimate of the pair at all, since it sets the share of shared failures to zero, the one value the double failure has made improbable.

15.6 The next measurement (step 7)

The report leaves the limit of 10^{-4} undemonstrated, at a posterior probability of 0.16 (0.26 exactly). The logit of the share carries three quarters of the variance of $\ln U_{\text{pair}}$ in the budget of Section 15.3.4 and the whole of its magnitude, so a further record can narrow the share's posterior, while an intervention on the cause can move the share itself.

15.6.1 The candidates

Two courses of action address the share:

- *More of the failure log.* Five, ten or twenty further years of the same log, with forty as a reference, each adding a singles row and a doubles row of exposure $8760\,\text{h}\,t$ for t years: in expectation $2(1-\beta_{\text{cc}})\lambda \times 8760\,\text{h}\,t$ singles and $\beta_{\text{cc}}\lambda \times 8760\,\text{h}\,t$ doubles, 1.99 singles and 0.151 doubles per year at the posterior means.
- *Staggered maintenance now.* Servicing the two pumps on different days removes the repeated maintenance error as a shared cause; it is assumed here, as an illustration, to halve the share.

The intervention is evaluated directly, on the posterior of the state with every node's share halved. By the near-linearity of Table 15.7, halving the share halves the unavailability: at the MAP state, from 2.11×10^{-4} to 1.065×10^{-4}, which is the new median. The sum (15.25) gives the mean 1.33×10^{-4} in place of 2.62×10^{-4}, and the delta method, with the gradient at the halved share $(1.020,\ 0.847)^{\mathsf{T}}$ and $\sigma_{\ln U} = 0.737$, the band $[2.5 \times 10^{-5},\ 4.5 \times 10^{-4}]$, a mean downtime of 1.16 h per year and the probability of meeting the limit

$$\Pr(U_{\text{pair}} < 10^{-4} \mid \mathbf{y}) \approx \Phi\left(\frac{\ln 10^{-4} - \ln(1.065 \times 10^{-4})}{0.737}\right) = \Phi(-0.085) = 0.47,$$

up from 0.16 and still short of the assurance 0.90. The values beside them come from the exact posterior. At the posterior means the unavailability falls from 2.08×10^{-4} to 1.05×10^{-4}. Under the posterior, with every draw's share halved, the mean falls from 2.14 to 1.08×10^{-4}, the median to 0.88×10^{-4}, the band to $[1.3 \times 10^{-5},\ 3.2 \times 10^{-4}]$ and the mean downtime to 0.95 h per year, while the standard deviation of $\ln U_{\text{pair}}$ stays near its present value, 0.807; the probability of meeting the limit rises from 0.26 to 0.57, still short of the assurance 0.90. The gap on the probability stays near 0.10, and the two routes agree that the step is large and that

it stops short of the assurance. The intervention acts at once. The record is evaluated by its preposterior.

15.6.2 The preposterior covariance

A further record of t years adds a singles row and a doubles row of exposure $8760\,\mathrm{h}\,t$. At $\hat{\mathbf{z}}$ they expect $\bar{k}_1' = 1.986\,t$ singles and $\bar{k}_2' = 0.152\,t$ doubles, and by equation (15.10) their Jacobian rows are

$$\mathbf{J}_a = \begin{bmatrix} \sqrt{\bar{k}_1'} & -\beta_{\mathrm{cc}}\sqrt{\bar{k}_1'} \\ \sqrt{\bar{k}_2'} & (1-\beta_{\mathrm{cc}})\sqrt{\bar{k}_2'} \end{bmatrix} = \sqrt{t}\begin{bmatrix} 1.409 & -0.187 \\ 0.390 & 0.338 \end{bmatrix}.$$

A count drawn with its own expectation has residual zero on average, so the curvature term of equation (15.16) averages to zero, and the information the record is expected to add and the preposterior covariance are

$$\boldsymbol{\Lambda}_a = \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a = \hat{\lambda}\,8760\,\mathrm{h}\,t\begin{bmatrix} 2-\beta_{\mathrm{cc}} & -\beta_{\mathrm{cc}}(1-\beta_{\mathrm{cc}}) \\ -\beta_{\mathrm{cc}}(1-\beta_{\mathrm{cc}}) & \beta_{\mathrm{cc}}(1-\beta_{\mathrm{cc}})(1+\beta_{\mathrm{cc}}) \end{bmatrix} = t\begin{bmatrix} 2.138 & -0.132 \\ -0.132 & 0.150 \end{bmatrix},$$

$$\boldsymbol{\Sigma}_a = \left(\boldsymbol{\Sigma}^{-1} + \boldsymbol{\Lambda}_a\right)^{-1}, \tag{15.56}$$

with $\boldsymbol{\Sigma}^{-1} = \mathbf{H}$ of equation (15.20) and t in years: the preposterior covariance of Chapter 2, equation (2.17), with the candidate rows in the place of the new reading. The entries of $\boldsymbol{\Lambda}_a$ carry the finding. Each year adds 2.14 to the information on the log rate and 0.150 to the information on the logit of the share, fourteen times less; of the 0.150 the doubles row supplies 0.114, its expected doubles times $(1-\beta_{\mathrm{cc}})^2$, and the singles row 0.035, its expected singles times β_{cc}^2. The logit gains about one unit of information per expected double, and doubles arrive at 0.15 a year, against the entry $H_{22} = 1.82$ of the present Hessian. That entry is the curvature in ξ at fixed z_1, not the marginal precision of the logit, $1/\Sigma_{22} = 1.785$, which the correlation with z_1 lowers slightly. Halving $\sqrt{\Sigma_{22}}$ needs about four times the present precision and takes 37.8 years of record; halving $\sqrt{\Sigma_{11}}$ takes 17.3 years.

Table 15.13 evaluates equation (15.56) at the horizons of Section 15.6.1: each row is a horizon, and the columns give the two preposterior standard deviations, the band on the share obtained by mapping $\hat{\xi} \pm 1.96\sqrt{(\boldsymbol{\Sigma}_a)_{22}}$ through $\beta_{\mathrm{cc}} = 1/(1+\mathrm{e}^{-\xi})$, its width, and the standard deviation of $\ln U_{\mathrm{pair}}$, $\sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_a\mathbf{g}}$ with $\mathbf{g}$ of Section 15.3.4.

Table 15.13: The Laplace preposterior of the state after further record, equation (15.56), centered at $\hat{\mathbf{z}}$; band: $\hat{\xi} \pm 1.96\sqrt{(\boldsymbol{\Sigma}_a)_{22}}$ mapped to the share; last column with $\mathbf{g} = (1.009,\ 0.857)^{\mathsf{T}}$.

Further record	$\sqrt{(\boldsymbol{\Sigma}_a)_{11}}$ [–]	$\sqrt{(\boldsymbol{\Sigma}_a)_{22}}$ [–]	Band on β_{cc} [–]	Width [–]	$\sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_a\mathbf{g}}$ [–]
none	0.292	0.748	[0.034, 0.399]	0.365	0.742
5 years	0.213	0.633	[0.042, 0.347]	0.304	0.617
10 years	0.176	0.559	[0.049, 0.315]	0.266	0.541
20 years	0.138	0.465	[0.058, 0.276]	0.218	0.448
40 years	0.104	0.366	[0.070, 0.239]	0.170	0.352

The information matrix says why the band narrows slowly: the share is informed at the rate at which doubles occur. Twenty years add 3.0 to the entry $H_{22} = 1.82$, and

$$\mathbf{H} + \boldsymbol{\Lambda}_a\big|_{t=20} = \begin{bmatrix} 54.77 & -3.30 \\ -3.30 & 4.82 \end{bmatrix}, \qquad \boldsymbol{\Sigma}_a\big|_{t=20} = \begin{bmatrix} 0.0191 & 0.0131 \\ 0.0131 & 0.2166 \end{bmatrix},$$

so that the standard deviation of the logit falls only from 0.748 to $\sqrt{0.2166} = 0.465$, while that of the log rate falls from 0.292 to 0.138.

15.6.3 The exact preposterior of the share and of the threshold

The values beside the Laplace preposterior come from the exact posterior. A further record of t years adds, in expectation, $2(1-\beta_{\mathrm{cc}})\lambda \times 8760\,\mathrm{h}\,t$ singles and $\beta_{\mathrm{cc}}\lambda \times 8760\,\mathrm{h}\,t$ doubles, 1.99 singles and 0.151 doubles per year at the posterior means, and the exposure $8760\,\mathrm{h}\,t$. Inserting the expected counts and the longer exposure into the marginal (15.13) gives the band on the share that the longer record is expected to leave. Because the counts are random, the check averages the band over records drawn from the posterior predictive: 4000 records for each horizon, each drawn by taking a state from the posterior and counts from the Poisson rows (15.5), each turned into its own posterior band. Table 15.14 gives the result. Each row is a horizon; the columns give the expected number of further doubles, the band from the expected counts and its width, the average width over the simulated records, and the band the Beta counting of Section 15.2.1 would report with the same expected counts.

Table 15.14: The preposterior of the common-cause share after further record. Expected counts at the posterior means $\lambda = 1.309 \times 10^{-4}\,\mathrm{h}^{-1}$, $\beta_{\mathrm{cc}} = 0.131$; band: central 95 % interval of equation (15.13) with the counts and the exposure added; average width: over 4000 records from the posterior predictive per horizon; Beta counting: $\mathrm{Beta}(2+\mathbb{E}[k_2'],\ 18+\mathbb{E}[k_1'])$.

Further record	Expected doubles	Band on β_{cc}	Width	Average width	Beta counting
none	–	[0.018, 0.329]	0.311	–	[0.013, 0.260]
5 years	0.75	[0.027, 0.295]	0.268	0.264	[0.018, 0.211]
10 years	1.51	[0.035, 0.274]	0.239	0.234	[0.021, 0.185]
20 years	3.01	[0.046, 0.248]	0.202	0.197	[0.027, 0.158]
40 years	6.03	[0.061, 0.221]	0.161	0.155	[0.034, 0.133]

Figure 15.7 draws the band against the length of the further record. The horizontal axis is the further record in years and the vertical axis the share. The shaded region is the band from the expected counts and the solid line its median; the two dotted lines are the band of the Beta counting; the dots mark the ends of the band at 5, 10, 20 and 40 years, the rows of Table 15.14; the dashed horizontal line is the present posterior mean 0.131 and the red line its half, 0.066, the share the intervention would produce. The band narrows slowly: by a seventh in five years, by a third in twenty, and after twenty more years it still runs from 0.046 to 0.25, a range that contains both the present mean and its half. A record of a shared cause that fails a pair once in about seven years cannot tell within a planning horizon whether the share is 0.07 or 0.13; the intervention changes the answer without needing to.

The Beta counting band of the figure is narrower and falls lower with every year, converging toward the event share near 0.07: it would report after twenty years a band of [0.027, 0.158] centered near 0.08, and suggest that the record was settling the share near the intervention's value, which it is not. The exact band is centered near 0.125 throughout.

The Laplace band of Table 15.13 is wider than the exact one at every horizon, by 0.05 to 0.01 in width, the upper tail that the normal on the logit overstates (Section 15.4.1), and it narrows in nearly the same proportion: by a sixth in five years, against a seventh for the exact band, and by two fifths in twenty, against a third. Both routes say the same thing.

The record judged on the pair. The intervention was judged on the pair's unavailability and its threshold probability, and the record must be judged on the same objects. On the state route the record divides the variance of $\ln U_{\mathrm{pair}}$: after t years the posterior of the state has the standard deviation $\sigma_a = \sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_a\mathbf{g}}$, and its center moves with the outcome, a normal about -8.466 of variance $\sigma^2_{\ln U} - \sigma_a^2$, weighted by the Laplace predictive. The threshold probability after a record with center c is $\Phi((\ln 10^{-4} - c)/\sigma_a)$, so it exceeds 0.5 when $c < \ln 10^{-4}$ and 0.9

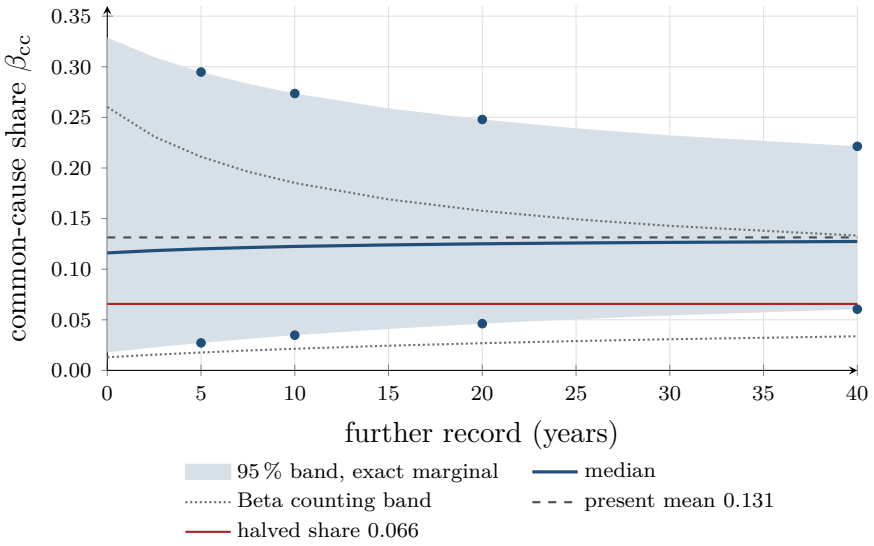


Figure 15.7: The preposterior band on the common-cause share against the length of further record, from the expected counts 1.99 singles and 0.151 doubles per year inserted in equation (15.13) (dots at 5, 10, 20 and 40 years, Table 15.14); the Beta counting band $\mathrm{Beta}(2 + 0.151t,\ 18 + 1.99t)$ for t years; the present posterior mean and its half.

when $c + 1.2816\,\sigma_a < \ln 10^{-4}$, and both fractions are normal tails:

$$\Pr(\text{above } 0.5) = \Phi\left(\frac{-1.003 \times 0.742}{\sqrt{0.742^2 - \sigma_a^2}}\right), \qquad \Pr(\text{above } 0.9) = \Phi\left(\frac{-1.003 \times 0.742 - 1.2816\,\sigma_a}{\sqrt{0.742^2 - \sigma_a^2}}\right). \tag{15.57}$$

Beside it, for each of the 4000 simulated records of each horizon, the posterior of the pair's unavailability is evaluated by 4000 draws of its exact posterior, and its threshold probability and band are recorded. Table 15.15 summarizes both: each row is a course of action; the columns give the expected threshold probability after it, the fraction of records after which the probability would exceed 0.5 and 0.9, and the 95 % band on U_{pair}, about the center 2.11×10^{-4} on the state route and the median over records on the exact one. The expected threshold probability does not move, 0.16 (0.26 exactly) at every horizon, as it must: the expectation of a posterior probability over the records it predicts is the present probability. A record can only redistribute it. After twenty years about one record in ten (one in five exactly) would raise it above one half, about one in a hundred (none exactly) would raise it to 0.9; on the exact route two in five would lower it below 0.1, and the band narrows from a ratio of 18.4 (25) between its ends to about 5.8 (8). On the exact route the records that rise above one half are mostly those with no further double: a third of the records at ten years, with a mean probability of 0.52, and a sixth at twenty, with 0.71, so the fraction above one half falls from 0.21 to 0.19 while each such record says more; the state route, whose center moves smoothly with the outcome, has no such step and rises through 0.071 and 0.10. After forty years 0.13 (0.25 exactly) of the records exceed one half and 0.034 (0.036) reach 0.9. Halving the share raises the probability to 0.47 (0.57) now.

15.6.4 The width each candidate removes

The reported answer is the band on the pair's unavailability, and the two preposteriors measure what each candidate removes from it. On the Laplace route, the last column of Table 15.13, the standard deviation of $\ln U_{\text{pair}}$ falls from 0.742 by 17 % after five years, by 27 % after ten and by 40 % after twenty, and the ratio $\exp(2 \times 1.96\,\sigma_{\ln U})$ between the ends of the band falls from

Table 15.15: The preposterior of the threshold question, $\Pr(U_{\text{pair}} < 10^{-4} \mid \mathbf{y}')$ after further record $\mathbf{y}'$. Each cell: the state route, equation (15.57) with $\boldsymbol{\Sigma}_a$ of equation (15.56), and in parentheses the exact route, over 4000 records per horizon drawn from the posterior predictive, each evaluated by 4000 exact posterior draws; against the halved share now. Exact band: median over records of the central 95 % interval.

Course of action	Expected probability	Records above 0.5	Records above 0.9	Band on U_{pair}
present record	0.16 (0.26)	–	–	$[4.9 \times 10^{-5},\ 9.0 \times 10^{-4}]$ ($[2.5 \times 10^{-5},\ 6.3 \times 10^{-4}]$)
5 more years	0.16 (0.26)	0.035 (0.004)	0.0001 (0)	$[6.3 \times 10^{-5},\ 7.1 \times 10^{-4}]$ ($[2.3 \times 10^{-5},\ 4.7 \times 10^{-4}]$)
10 more years	0.16 (0.26)	0.071 (0.21)	0.002 (0)	$[7.3 \times 10^{-5},\ 6.1 \times 10^{-4}]$ ($[3.8 \times 10^{-5},\ 4.4 \times 10^{-4}]$)
20 more years	0.16 (0.26)	0.10 (0.19)	0.013 (0)	$[8.7 \times 10^{-5},\ 5.1 \times 10^{-4}]$ ($[4.9 \times 10^{-5},\ 3.8 \times 10^{-4}]$)
40 more years	0.16 (0.26)	0.13 (0.25)	0.034 (0.036)	$[1.06 \times 10^{-4},\ 4.2 \times 10^{-4}]$ ($[7.3 \times 10^{-5},\ 3.4 \times 10^{-4}]$)
halved share, now	0.47 (0.57)	–	–	$[2.5 \times 10^{-5},\ 4.5 \times 10^{-4}]$ ($[1.3 \times 10^{-5},\ 3.2 \times 10^{-4}]$)

18.4 to 11.2, 8.3 and 5.8, and to 4.0 after forty years. On the exact draws of Table 15.15 the ratio between the ends of the median band falls from 25 to 20, 12 and 8, and to 4.6 after forty years. The halved share removes almost no width: its band keeps the ratio of 18 on the state route, 25 on the exact, and moves down by a factor of two as a whole. The record narrows the band about its present center, slowly; the intervention moves the center at once. Chapter 20, in Volume II, ranks candidate rows by the same preposterior on graphs whose rows share entries, as the singles and the doubles rows here share both.

15.6.5 The recommendation

The intervention beats waiting for the record: stagger the maintenance of the two pumps now. Halving the share by staggered maintenance halves the pair's unavailability at once, from 2.6 to 1.3×10^{-4} on the mean (2.1 to 1.1×10^{-4} exactly), and raises the probability of meeting the limit of 10^{-4} from 0.16 to 0.47 (0.26 to 0.57), still short of the assurance 0.90 on both routes, so the limit itself calls for a further reduction of the share or a change of the design. Twenty more years of record leave the expected probability at 0.16 (0.26), would raise it above one half in only one case in ten (five) and almost never to 0.9, and narrow the 95 % band on the share only from [0.034, 0.399] to about [0.058, 0.276] ([0.018, 0.329] to about [0.046, 0.248]), which still cannot separate the present share from its half. The gap between the routes does not change the recommendation. The record cannot settle the question within the planning horizon; the intervention can change the answer. The failure log is kept in either case, since its doubles are the record that will show whether the staggering has reduced the share. The price of the intervention against the value of the reduction is a question for the later chapter on the value of information.

15.6.6 What the analysis showed

The first reading of the log, two independent pumps down together 2.3×10^{-6} of the time, said the pair passes the limit with room to spare; the analysis overturns it. The one double failure in five years is not a coincidence the design can set aside: it is strong evidence that a shared cause acts on both pumps, and a shared cause removes the product that made the redundancy look so good, so the pair is down about a hundred times as often as the product promised, 114 times

on the mean (94 exactly), and on average about twice the limit. The second pump protects the hall against each pump's own failures and does almost nothing against the shared cause, which now sets the answer. Waiting does not rescue the design. The share is informed only as fast as doubles occur, about one in seven years, so twenty more years of log would still not tell the present share from its half, and no prior examined demonstrates the limit. What the site can do is act on the cause: staggering the maintenance of the two pumps, so that one error is not repeated on both, moves the answer at once, and the failure log, kept with its time stamps, then shows whether the staggering has reduced the share.

15.6.7 Limits of the model

The derivation is exact under the assumptions of the β-factor model of a pair, and each assumption can be checked as the record grows.

The same rate for both pumps, constant over the record. The pumps are taken to have one total rate and one share; their own single failures, counted separately, test the first. The rates are taken as constant over five years; the trend tests of Chapter 12 apply to the singles count.

The mechanism of the shared cause. The shared cause is taken to fail both pumps within the window whenever it acts. A cause that fails one pump and degrades the other is a different mechanism, outside the β-factor model, and a cause whose effects spread over more than the window is counted as two singles.

The repair rate. The repair rate is pinned, and taken equal for a single and a double repair. One crew restoring two pumps may take longer after a common-cause failure; the record of the double's repair time would inform it.

Independent two-state processes. The pair's unavailability treats the shared and the two independent processes as independent. A pump held down by the shared cause is not running and cannot fail on its own, which changes the unavailability by a term of the order of $U_{\mathrm{cc}} q_{\mathrm{ind}}$ in the overlap, and a shared repair crew couples the repairs of the two pumps; the next chapter's model of the joint states represents both.

The window. The answers that depend on the window, the Bayes factor and the event posterior, are reported for the practical logs; the share and the pair's unavailability do not depend on it.

Beyond two units. For groups of more than two units the shared cause may take out two, three or all units, and the α-factor and multiple Greek letter models give each size of failure group its own share; the companion volume *Reliability Engineering* states them, and each is a set of further entries of the state and further count rows of the same kind as here (Exercise 15.5).

The next chapter replaces the three independent two-state processes of this chapter by a Markov model of the joint states of repairable units, in which a shared repair crew, a common cause and a degraded running state can be represented together.

15.7 What to retain

Definition 15.1 (The β-factor model of a redundant pair)**.** Two identical units with total failure rate λ each, of which a share β_{cc} comes from a shared cause that fails both, are three independent

Poisson processes: two of rate $(1-\beta_{\mathrm{cc}})\lambda$ and one of rate $\beta_{\mathrm{cc}}\lambda$. In an exposure T_{tot} the singles and the doubles are independent counts Poisson$(2(1-\beta_{\mathrm{cc}})\lambda T_{\mathrm{tot}})$ and Poisson$(\beta_{\mathrm{cc}}\lambda T_{\mathrm{tot}})$; with a gamma prior on λ and a Beta prior on the share the rate's conditional posterior is gamma and the share's marginal is $\beta_{\mathrm{cc}}^{a_{\beta 0}+k_2-1}(1-\beta_{\mathrm{cc}})^{b_{\beta 0}+k_1-1}(b_{\lambda 0}+(2-\beta_{\mathrm{cc}})T_{\mathrm{tot}})^{-(a_{\lambda 0}+k_1+k_2)}$ up to a constant. With repair rate μ the pair in one-out-of-two has unavailability $U_{\mathrm{pair}} = 1-(1-U_{\mathrm{cc}})(1-q_{\mathrm{ind}}^2) \approx \beta_{\mathrm{cc}}\lambda/\mu + ((1-\beta_{\mathrm{cc}})\lambda/\mu)^2$.

Five conclusions carry beyond the pumps. A shared cause is an entry of the state with its own count row: in the 4×2 Jacobian the doubles row, $\sqrt{\bar{k}_2}\,(1,\ 1-\beta_{\mathrm{cc}})$, is the data row that informs the logit of the share, the singles row, $\sqrt{\bar{k}_1}\,(1,\ -\beta_{\mathrm{cc}})$, informs the rate and, through its entries of opposite sign, correlates the two, and the counts of singles and doubles inform the share of failure events, $\beta_{\mathrm{cc}}/(2-\beta_{\mathrm{cc}})$, so counting events as trials of the share misreads them by about a factor of two. The evidence for a shared cause from simultaneous failures depends on the coincidence window as $-k_2 \ln t_{\mathrm{w}}$, and the window belongs to the record, not to the analysis. The evidence for a drawing and the posterior of one event are different objects and answer different questions. And a redundant pair's unavailability is first order in the share and second order in the independent failures, so a share of a tenth makes the pair 65 times more unavailable than independence at the same total rate promises, and the posterior of this record puts the pair 114 (94 exactly) times above the first reading, an effect that an intervention on the share changes at once and that a longer record only describes more precisely, since the information a year of record adds to the logit of the share, the $(2,2)$ entry of $\mathbf{\Lambda}_a$, is about the expected number of doubles in that year, 0.15 against 2.14 for the log rate. Finally, an entry of the state informed by one event has a skewed posterior that the Laplace normal does not reproduce: here the delta method gives the threshold probability 0.16 against the exact 0.26, although its linearization was accurate. Every answer of the chapter is computed from the posterior of the state, with the exact value, cheap for two entries, beside it and the gap stated.

Table 15.16 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 15.16: The terms of dependent failures and common causes and where each is defined.

Term	Expression	Units	Meaning	Defined in
Common-cause share	β_{cc}; $\xi = \ln(\beta_{\mathrm{cc}}/(1-\beta_{\mathrm{cc}}))$	–	shared fraction of a unit's rate	15.1.3
Superposition	rates of independent Poisson processes add	h^{-1}	per-pump rate λ	15.1.6
Thinning	marked Poisson process splits into independent ones	–	singles and doubles independent	15.1.6
Event share	$\theta = \beta_{\mathrm{cc}}/(2-\beta_{\mathrm{cc}})$	–	doubles per event	15.1.6
Count rows	Poisson$(2(1-\beta_{\mathrm{cc}})\lambda T_{\mathrm{tot}})$, Poisson$(\beta_{\mathrm{cc}}\lambda T_{\mathrm{tot}})$	–	both act on both entries	15.1.6
Jacobian rows	$\sqrt{\bar{k}_1}\,(1,\ -\beta_{\mathrm{cc}})$, $\sqrt{\bar{k}_2}\,(1,\ 1-\beta_{\mathrm{cc}})$; curvature term in ξ	–	singles inform the rate, doubles the share	15.2.2
Conditional rate posterior	Gamma$(a_{\lambda 0}+m,\ b_{\lambda 0}+(2-\beta_{\mathrm{cc}})T_{\mathrm{tot}})$	h^{-1}	rate given the share	15.2.1
Marginal of the share	equation (15.13)	–	exact up to one integral	15.2.1
Beta approximation	Beta$(a_{\beta 0}+k_2,\ b_{\beta 0}+k_1)$	–	drops the coupling; biased	15.2.1
Coincidence rate	$c_{\mathrm{w}} = 2\lambda^2 t_{\mathrm{w}}$; $\lambda^2 t_{\mathrm{w}}$ for a binned log	h^{-1}	doubles under independence	15.1.6
Bayes factor	$\approx E_\beta/(2\tilde{\lambda}t_{\mathrm{w}})$ for one double	–	shared cause against coincidence	15.3.2
Event posterior	$\beta_{\mathrm{cc}}\lambda/(\beta_{\mathrm{cc}}\lambda + 2(1-\beta_{\mathrm{cc}})^2\lambda^2 t_{\mathrm{w}})$, averaged	–	was this double shared	15.3.3
Pair unavailability	$1-(1-U_{\mathrm{cc}})(1-q_{\mathrm{ind}}^2)$	–	one-out-of-two with shared cause	15.3.4
Working form	$\beta_{\mathrm{cc}}q + ((1-\beta_{\mathrm{cc}})q)^2$, error $\beta_{\mathrm{cc}}(1-\beta_{\mathrm{cc}})\epsilon^2$	–	first order in the share	15.3.4
Delta method on $\ln U$	$\mathbf{g} = (1.009,\ 0.857)^{\mathsf{T}}$, nearly the doubles direction $(1,\ 1-\beta_{\mathrm{cc}})^{\mathsf{T}}$; $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$	–	band and threshold of the pair; a gap with one double	15.3.4
Preposterior band	equation (15.13) with expected counts added	–	width after further record	15.6.3
Preposterior covariance	$\boldsymbol{\Sigma}_a = (\boldsymbol{\Sigma}^{-1}+\boldsymbol{\Lambda}_a)^{-1}$, $\boldsymbol{\Lambda}_a = \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a$	–	0.15 per year on ξ	15.6.2

15.8 Exercises

Exercise 15.1. *Formulation.* Two pumps have a total failure rate $\lambda = 2\times10^{-4}\,\mathrm{h}^{-1}$ each, a common-cause share $\beta_{\mathrm{cc}} = 0.05$ and a repair rate $\mu = 1/12\,\mathrm{h}^{-1}$, all taken as known. (a) Compute the rate of single failures of the pair, the rate of double failures and the total rate of one pump, per hour and per year. (b) Compute the coincidence rate of independent failures within $t_{\mathrm{w}} = 8\,\mathrm{h}$ under the independence drawing, and the probability that a double observed under the shared-cause drawing was a common cause. (c) Compute the pair's unavailability by equation (15.42) and by the working form, and compare it with independence. ((a) Singles $2\times0.95\times2\times10^{-4} = 3.8\times10^{-4}\,\mathrm{h}^{-1}$, 3.33 per year; doubles $10^{-5}\,\mathrm{h}^{-1}$, 0.088 per year; one pump $(1-0.05)\lambda + 0.05\lambda = 2\times10^{-4}\,\mathrm{h}^{-1}$. (b) $c_{\mathrm{w}} = 2\lambda^2 t_{\mathrm{w}} = 6.4\times10^{-7}\,\mathrm{h}^{-1}$, 0.0056 per year ($3.2\times10^{-7}\,\mathrm{h}^{-1}$ if the log is binned by shift); under M_1 the coincidences arrive at $2(0.95\lambda)^2 t_{\mathrm{w}} = 5.78\times10^{-7}\,\mathrm{h}^{-1}$, so $\Pr = 10^{-5}/(10^{-5}+5.78\times10^{-7}) = 0.945$. (c) $U_{\mathrm{cc}} = 1.200\times10^{-4}$, $q_{\mathrm{ind}}^2 = 5.17\times10^{-6}$, $U_{\mathrm{pair}} = 1.252\times10^{-4}$, 1.1 h per year; the working form gives 1.249×10^{-4}, short by $\beta_{\mathrm{cc}}(1-\beta_{\mathrm{cc}})\epsilon^2 = 2.7\times10^{-7}$; independence at the total rate gives

$q^2 = 5.7 \times 10^{-6}$, 22 times less, and the independent part alone $q_{\text{ind}}^2 = 5.2 \times 10^{-6}$, 24 times less.)

Exercise 15.2. *Derivation.* Starting from three independent Poisson processes, two of rate $(1-\beta_{\text{cc}})\lambda$ and one of rate $\beta_{\text{cc}}\lambda$ that fails both units, derive (a) that each unit's failures form a Poisson process of rate λ; (b) that the singles and the doubles in T_{tot} are independent Poisson counts with means $2(1-\beta_{\text{cc}})\lambda T_{\text{tot}}$ and $\beta_{\text{cc}}\lambda T_{\text{tot}}$; (c) that the correlation of the two units' failure counts is β_{cc}. ((a) $N_1 + N_c$ is a sum of independent Poisson counts, and the convolution with the binomial theorem gives Poisson$((1-\beta_{\text{cc}})\lambda t + \beta_{\text{cc}}\lambda t) =$ Poisson(λt). (b) The singles $N_1 + N_2$ are Poisson with mean $2(1-\beta_{\text{cc}})\lambda T_{\text{tot}}$ by the same step, and N_c is independent of them since the processes are; equivalently, thinning the event stream of rate $(2-\beta_{\text{cc}})\lambda$ with probability $\theta = \beta_{\text{cc}}/(2-\beta_{\text{cc}})$ factorizes the joint probability into two Poisson terms. (c) $\text{Cov}(N_1 + N_c, N_2 + N_c) = \text{Var}(N_c) = \beta_{\text{cc}}\lambda t$ and each variance is λt.)

Exercise 15.3. *Derivation.* (a) Derive the pair's unavailability equation (15.42) from the steady state of three two-state processes. (b) Expand it to second order in $\epsilon = \lambda/\mu$ and derive the error of the working form (15.44). (c) Evaluate the error for the chapter's pumps. ((a) Each process is down with probability rate/(rate+μ); the pair is up iff the shared process is up and not both independent ones are down, so $1 - U_{\text{pair}} = (1 - U_{\text{cc}})(1 - q_{\text{ind}}^2)$. (b) $U_{\text{cc}} = \beta_{\text{cc}}\epsilon - \beta_{\text{cc}}^2\epsilon^2 + O(\epsilon^3)$, $q_{\text{ind}}^2 = (1-\beta_{\text{cc}})^2\epsilon^2 + O(\epsilon^3)$; $q = \epsilon - \epsilon^2 + O(\epsilon^3)$ gives the working form $\beta_{\text{cc}}\epsilon + [(1-\beta_{\text{cc}})^2 - \beta_{\text{cc}}]\epsilon^2$; the difference is $\beta_{\text{cc}}(1-\beta_{\text{cc}})\epsilon^2$. (c) With $\epsilon = 1.570 \times 10^{-3}$ and $\beta_{\text{cc}} = 0.131$: 2.8×10^{-7}, a relative error of $(1-\beta_{\text{cc}})\epsilon = 1.4 \times 10^{-3}$, as in Table 15.7.)

Exercise 15.4. *Derivation.* (a) From the joint posterior (15.11), derive the marginal of the share (15.13). (b) Show that dropping its last factor gives the Beta approximation (15.14), and that the factor is approximately $\exp(s\beta_{\text{cc}})$ with $s = (a_{\lambda 0} + m)T_{\text{tot}}/(b_{\lambda 0} + 2T_{\text{tot}})$. (c) Show that for a long record the counts are binomial in $\theta = \beta_{\text{cc}}/(2-\beta_{\text{cc}})$, and state the limit of the Beta approximation's mean. ((a) Integrate λ with $\int \lambda^{a-1}\text{e}^{-b\lambda}\text{d}\lambda = \Gamma(a)/b^a$, $a = a_{\lambda 0} + m$, $b = b_{\lambda 0} + (2-\beta_{\text{cc}})T_{\text{tot}}$. (b) The remaining powers are the Beta kernel; $(1 - T_{\text{tot}}\beta_{\text{cc}}/(b_{\lambda 0} + 2T_{\text{tot}}))^{-a} \approx \exp(aT_{\text{tot}}\beta_{\text{cc}}/(b_{\lambda 0} + 2T_{\text{tot}}))$, $s = 5.39$, and the tilted Beta has mean 0.128 against the exact 0.131 and the untilted 0.100. (c) With $b_{\lambda 0} \ll 2T_{\text{tot}}$ the factor is $\propto (2-\beta_{\text{cc}})^{-m}$ and the likelihood $\propto (1-\theta)^{k_1}\theta^{k_2}$; the Beta mean tends to $k_2/m \to \theta \approx \beta_{\text{cc}}/2$.)

Exercise 15.5. *Application.* Three pumps serve a load in two-out-of-three, each with total failure rate $\lambda = 10^{-4}\,\text{h}^{-1}$, common-cause share $\beta_{\text{cc}} = 0.1$ and repair rate $1/12\,\text{h}^{-1}$. Half of each pump's common-cause rate comes from events that take out it and one other pump, and half from events that take out all three. (a) Write the processes and their rates. (b) Compute the system's unavailability by enumerating the process states and by the rare-event sum, and compare with independence. (c) Compare with a split in which every common-cause event takes out all three. ((a) Three independent processes of rate $0.9\lambda = 9 \times 10^{-5}\,\text{h}^{-1}$; each pump is in two pairs, so $2\lambda_2 + \lambda_3 = \beta_{\text{cc}}\lambda$ with $2\lambda_2 = \lambda_3$: three pair processes of rate $\lambda_2 = 2.5 \times 10^{-6}\,\text{h}^{-1}$ and one triple of rate $\lambda_3 = 5 \times 10^{-6}\,\text{h}^{-1}$. (b) Any pair or triple event takes the system down, as do two independent failures: the rare-event sum $3\lambda_2/\mu + \lambda_3/\mu + 3(0.9\lambda/\mu)^2 = 9.0 \times 10^{-5} + 6.0 \times 10^{-5} + 3.5 \times 10^{-6} = 1.535 \times 10^{-4}$, and the enumeration of the 2^7 process states gives the same, 1.535×10^{-4}; independence gives $3q^2 = 4.3 \times 10^{-6}$, 36 times less. (c) With $\lambda_3 = \beta_{\text{cc}}\lambda$ and no pair events, 1.23×10^{-4}: in two-out-of-three a pair event is as harmful as a triple, and there are three pairs, so the split matters by a quarter.)

Exercise 15.6. *Application.* The same pump pair has a record of $k_1 = 9$ singles and $k_2 = 2$ doubles in 43 800 h, with the same priors, window and repair rate. Answer the five questions and the choice of Section 15.6 in brief, and give the Laplace check: the Jacobian (15.10) at the MAP state, $\mathbf{H}$ and $\boldsymbol{\Sigma}$. ((1) Share: mean 0.187, median 0.175, band [0.044, 0.396], $\Pr(\beta_{\text{cc}} > 0.05) = 0.97$; rate mean $1.46 \times 10^{-4}\,\text{h}^{-1}$. (2) $\ln B_{10} = 7.31$ at 8 h, 11.45 at 1 h, 5.14 at 24 h: two coincidences are far less probable than one, since $k_2 = 2$ puts $(2\lambda t_{\text{w}})^2$ in

the independence evidence. (3) Each double was a common cause with probability 0.988. (4) $\mathbb{E}[U_{\text{pair}} \mid \mathbf{y}] = 3.4{\times}10^{-4}$, against the independent part $2.2{\times}10^{-6}$. (5) Band $[6.5{\times}10^{-5},\ 8.7{\times}10^{-4}]$ and $\Pr(U_{\text{pair}} < 10^{-4} \mid \mathbf{y}) = 0.07$. (6) Halving the share gives a mean of 1.7×10^{-4} and a probability of 0.28; twenty more years narrow the band only to about $[0.087,\ 0.310]$. The second double sharpens the drawing comparison far more than the share. Laplace check: the Newton iteration (15.17) reaches $\hat{\mathbf{z}} = (-8.835,\ -1.453)^{\mathsf{T}}$, $\hat{\lambda} = 1.46 \times 10^{-4}\,\text{h}^{-1}$ and $\hat{\beta}_{\text{cc}} = 0.189$, where the Jacobian has the rows $(1.207,\ 0)$, $(3.215,\ -0.609)$, $(1.099,\ 0.891)$ and $(0,\ 1.239)$, and $\mathbf{H} = \left[\begin{smallmatrix} 13.000 & -0.979 \\ -0.979 & 2.617 \end{smallmatrix}\right]$, $\boldsymbol{\Sigma} = \left[\begin{smallmatrix} 0.0792 & 0.0296 \\ 0.0296 & 0.3932 \end{smallmatrix}\right]$: the second double raises the doubles row's ξ entry from 0.757 to 0.891 and the entry H_{22} from 1.82 to 2.62. The logit's exact standard deviation is 0.675 against the Laplace $\sqrt{\Sigma_{22}} = 0.627$, its skewness -0.49, and the delta method gives $\Pr(U_{\text{pair}} < 10^{-4}) = 0.03$ against the exact 0.07.)

Chapter 16

Repairable and Multistate Systems

16.1 Repair capacity and the states of a redundant pair

Two chilled-water pumps feed the hall's cooling coils side by side, either able to carry the whole flow, and one maintenance crew repairs them. Their five-year failure log holds nine single-pump failures and one event that took both pumps out at once, and the crew keeps a record of its repair durations. The hall's service agreement limits the loss of cooling supply to 10^{-4} of the time, to be shown with an assurance of 0.90, the level Chapter 15 also fixes, and the first reading of the logs is the familiar one: each pump is down a fraction $q = \lambda/(\lambda + \mu)$ of the time, two redundant pumps are both down q^2 of the time, and at the posterior mean rates, 1.31×10^{-4} failures per hour per pump and one repair in 11.25 h, $q^2 = 2.16 \times 10^{-6}$, far below the limit. It is the product of Chapter 15, 2.3×10^{-6}, taken at these rates in place of the raw log's rate and that chapter's pinned twelve-hour repair. If more margin were wanted, a second crew would halve the time a failed pump waits, and with it the downtime.

The engineer responsible for the cooling supply sees what the product leaves out. It treats the two pumps as independent units, each with a crew of its own, while one crew serves both and a pump that fails during the other's repair waits; it ignores the one event of the log that took both pumps out together; and the limit wants 10^{-4} shown with an assurance of 0.90 from ten failure events and a short record of repairs, not from two point rates. Whether a second crew is worth more than a smaller common-cause share depends on what each does to the outages, not to the downtime of one pump. The question of this chapter is

how the long-run fraction of time a repairable pair is out of service is derived from its failure rate, its repair rate, its common-cause share and the size of its repair crew, how well that fraction is known given the records, whether it meets a limit, how soon the pair is whole again after both units are down, how long it runs before its first outage, and which of a second crew or a smaller common-cause share improves it more.

The records are the counts and durations of Chapters 8 and 15; what changes is the derived quantity, which is no longer a formula in the state but an entry of the solution of a linear system whose matrix depends on the state, the stationary distribution of a Markov chain on the pair's states.

Jargon: continuous-time Markov chain, Markov property, transition rate, generator matrix, forward (Kolmogorov) equations, stationary distribution, balance equations, cut, birth–death chain, matrix exponential, time constant, first-passage time, mean down time, repair capacity, implicit differentiation, sensitivity equation, phase-type distribution, multistate system

A *continuous-time Markov chain* is a process $X(t)$ that moves among a finite set of states in continuous time and has the *Markov property*: given its present state, its future is independent of its past. The *transition rate* from state i to state j is the limit of the probability of a jump from i to j in a short interval divided by the length of the interval; the *generator matrix* collects the transition rates, with each diagonal entry minus the sum of the others in its row. The *forward (Kolmogorov) equations* are the differential equations for the state probabilities in terms of the generator. The *stationary distribution* is the distribution of the state that the forward equations leave unchanged, and the long-run fraction of time spent in each state. The *balance equations* state that in stationarity the probability flow out of a set of states equals the flow into it; a *cut* is a division of the states into two sets, across which a balance equation is written (the word has nothing to do with the minimal cut sets of Chapter 14). A *birth–death chain* is a chain whose state changes only by one step up or down. The *matrix exponential* $\mathrm{e}^{\mathbf{A}t}$ is the sum $\sum_{n\geq 0}(\mathbf{A}t)^n/n!$, which solves a linear system of differential equations with constant coefficients; a *time constant* is the reciprocal of the decay rate of one of its modes. The *first-passage time* to a state is the time until the chain first enters it. The *mean down time* is the expected duration of one outage. The *repair capacity* is the number of units that can be under repair at once, here the number of crews. *Implicit differentiation* obtains the derivative of a quantity defined as the solution of an equation by differentiating the equation rather than a formula for the solution. The *sensitivity equation* of a linear differential equation is the differential equation obeyed by the derivative of its solution with respect to a parameter, obtained by differentiating the original equation. A *phase-type distribution* is the distribution of the time a Markov chain takes to pass through a set of auxiliary states, used to represent non-exponential durations. A *multistate system* is a system whose units have more than two states of performance, such as full output, reduced output and none.

Key idea

The state of a repairable pair is the number of its units up, and its probabilities obey the forward equations of a Markov chain whose rates are the failure rate, the repair rate times the number of crews and the common-cause share. The long-run unavailability is an entry of the stationary distribution, a derived quantity of the state entries carried to the posterior like any other. With c crews and each pump of a common-cause outage a separate repair job, a common cause enters the unavailability at first order, $\beta_{\mathrm{cc}}\lambda/(c\mu)$, and the independent double failures only at second.

16.1.1 Dynamics of a shared-cause pair

Chapter 8 derived the availability of one repairable pump from a two-state model: the pump is up or down, fails at the rate λ and is repaired at the rate μ, and its long-run unavailability is $\lambda/(\lambda+\mu)$. Chapter 15 applied that model to three processes at once, the independent failures of each of two pumps and the shared cause, each with its own repair, and combined their steady states into the pair's unavailability. The combination treats the three processes as independent. That is exact only when each failed pump has its own crew and a common-cause repair restores both pumps as one job. When one crew serves both pumps, a pump that fails while the other is under repair waits, and the processes are coupled through the crew.

A model of the pair as a whole removes the approximation. Its state is the number of pumps up, two, one or none, and its transitions are the failures and the completed repairs, at rates set by the failure rate, the common-cause share and the repair capacity. The two-state model of Chapter 8 is the special case of one unit; the three-state chain of the pair is the smallest case in which the repair capacity matters; and the same construction extends to any number of units and states. The companion volume *Reliability Engineering* treats such Markov models in general, with their generator matrices, their Kolmogorov equations and examples of pump systems; the pair is derived here as an inference. For three states the stationary distribution has a closed form, but for a plant of several units, crews and spares it has none, and the derivation works from the linear system itself.

16.1.2 The physical system and its graph (step 0)

The two identical chilled-water pumps of Chapter 15 take water from the chilled-water source and deliver it to the distribution header that feeds the hall's cooling coils. Either pump can carry the full flow, so the cooling supply is lost only when both are down. Both pumps run continuously, sharing the flow: the redundancy is active, and each pump is exposed to failure whenever it is up. (A standby arrangement, with the second pump idle until the first fails, would change the chain only in its first row: the transition from two up to one up would run at the running pump's rate alone, with a probability of failing to start on demand added to the transition from one up to none.) Each pump fails alone at the rate $(1-\beta_{\mathrm{cc}})\lambda$ and both fail together, from a shared cause, at the rate $\beta_{\mathrm{cc}}\lambda$, the β-factor model derived in Chapter 15, with λ the total failure rate of one pump and β_{cc} the common-cause share. A failed pump is repaired by the maintenance crew, which works on one pump at a time; each repair takes an exponentially distributed time with the repair rate μ. With one crew and both pumps down, one pump is repaired first and the other waits; with two crews both are repaired at once.

Figure 16.1 draws the plant graph and the chain. In panel (a) the gray boxes are the chilled-water source and the distribution header, pinned boundaries of the hydraulic path, and the solid green arrows are the two parallel paths through the pumps. The two green boxes are the pumps, free nodes sharing one failure rate, the entry $z_1 = \ln\lambda$. The box with the red border between them is the shared cause, whose red arrows take both pumps out at once and which carries the entry $z_3 = \xi$, the logit of the share. The purple box above is the repair crew, which carries the log repair rate $z_2 = \ln\mu$ and serves every pump inside the dashed purple frame. The two white boxes are the records, attached by dashed arrows: the failure log below, whose rows act on z_1 and z_3, and the repair record at the upper right, whose row acts on z_2. Panel (b) draws the chain of the pair: three states, named by the number of pumps up, and the transitions between them with their rates. The red arrow is the common cause, which takes the pair from two pumps up to none in one jump; the purple arrow is the repair out of the state with no pump up, the only rate that depends on the number of crews.

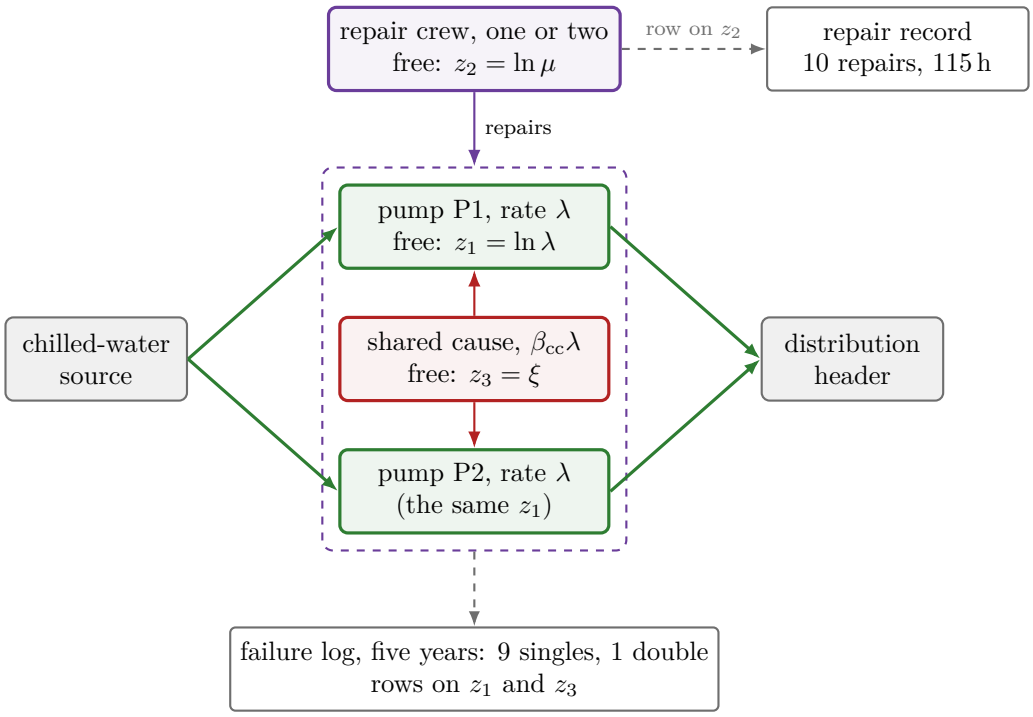
16.1.3 The records and the state

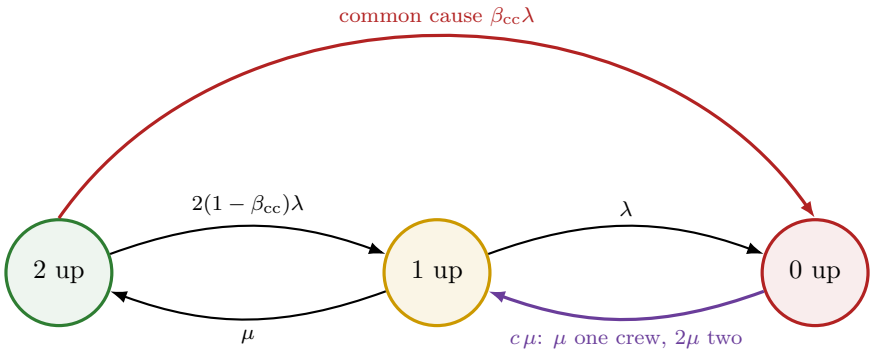
The state entries. The three unknowns are carried on the scales that remove their bounds,

$$z_1 = \ln\lambda, \qquad z_2 = \ln\mu, \qquad z_3 = \xi = \ln\frac{\beta_{\mathrm{cc}}}{1-\beta_{\mathrm{cc}}}, \qquad \mathbf{z} = \begin{bmatrix} z_1 \\ z_2 \\ z_3 \end{bmatrix}, \tag{16.1}$$

the log rates of Chapter 8 and the logit of the share of Chapter 15, with both rates in h^{-1}. The number of crews c is pinned: it is a design choice, one or two, not an unknown; it has nothing to do with the censoring time c_i of the notation. The chain's state $X(t)$, the number of pumps up at time t, is a different object from the state vector $\mathbf{z}$ of the plant graph: $X(t)$ is the random

(a) plant graph of the pump pair with its crew





(b) the chain of the pair

Figure 16.1: The chilled-water pump pair with its repair crew. (a) Plant graph: pinned, the chilled-water source and the distribution header (gray); free, the two pumps sharing the log failure rate z_1 (green), the shared cause carrying the logit of the share z_3 (red border) and the repair crew carrying the log repair rate z_2 (purple), which serves both pumps (dashed frame); observed, the failure log and the repair record (dashed arrows). (b) The chain of the pair: states by the number of pumps up, transitions with their rates; the common cause (red) jumps from two pumps up to none, and the repair out of the state with none up runs at $c\mu$ with c crews (purple).

condition of the pair at one instant, aleatory, and $\mathbf{z}$ holds the rates that govern it, epistemic. The chapter keeps the two apart.

The records. Table 16.1 lists the evidence: each row is one row of the graph, and the columns give the record behind it, the count and exposure it carries, and the entry it acts on.

Table 16.1: The records of the pump pair and the rows they put on the state. A count row of k events in an exposure h; the failure-rate and share rows carry the posteriors of Chapter 15 in that form.

Row	Record behind it	Count and exposure	Acts on
failure-rate row	five-year failure log of Chapter 15: 9 singles, 1 double in 43 800 h, with that chapter's prior; summarized by a gamma with the mean of that chapter's marginal	11 failures in 84 000 h	z_1
prior on μ	the prior Gamma(2, 20 h) of Chapter 8	2 repairs in 20 h	z_2
repair row	10 repair durations of the pumps, totaling 115 h	10 repairs in 115 h	z_2
share row	the same log, by counting events, equation (15.14): the prior Beta(1, 9) and 1 double in 10 events	2 doubles in 20 events	z_3

The failure rate. Chapter 15 derived the joint posterior of the rate and the share from the five-year log; the marginal of the rate there has mean $1.31 \times 10^{-4}\,\mathrm{h}^{-1}$ and interval $[0.67,\ 2.16] \times 10^{-4}\,\mathrm{h}^{-1}$. The pair's failure-rate row summarizes it by the gamma distribution with the same mean and the shape $a_\lambda = 11$, $\lambda \sim \mathrm{Gamma}(11,\ 84\,000\,\mathrm{h})$, a little wider in its tails than the marginal it replaces.

The repair rate. Chapter 15 pinned the repair rate. Here the maintenance log supplies ten repair durations of the pumps, $r = 10$ with total $D = 115\,\mathrm{h}$; the record is illustrative. The durations are not matched one to one with the events of the failure log: they inform the repair rate only, whatever failures they followed. The prior is the Gamma(2, 20 h) of Chapter 8, two pseudo-repairs in twenty pseudo-hours.

The share. The share is carried by the Beta distribution obtained in Chapter 15 by counting the ten events as trials, the one double as a success, $\beta_{\mathrm{cc}} \sim \mathrm{Beta}(1+1,\ 9+9) = \mathrm{Beta}(2,\ 18)$. That chapter showed it to be low: its exact marginal has mean 0.131, against the Beta's 0.100. The Beta form is kept because it makes the three entries independent with standard posteriors, which isolates what the chain and the crew contribute, and it drops the correlation of 0.14 between the rate and the share that Chapter 15 found; Section 16.4.3 evaluates the answers under the exact joint posterior as well.

16.1.4 The probabilistic questions (step 1)

The operator asks six questions of the pair, and step 1 writes each as a probability or an expectation over the state, with the unknown on the left of the bar and the records on the right. Table 16.2 lists them: the first column is the question in engineering language; the second the general form of its category in the notation of Chapter 2; the third its mathematical form for the pair; the fourth its category; and the last the subsection where it is answered. The first three are estimation. The first asks for a derived quantity defined implicitly, the entry π_0 of the stationary distribution, the long-run probability that no pump is up; it is the service availability

of the pair from its components with the repair capacity as a constraint. The second splits the same quantity by changing the drawing, removing the common cause or giving each pump its own crew. The third asks for the mean time to the first outage and the outage frequency. The fourth is a prediction under a threshold. The fifth is also a prediction under a threshold, on the time course after a double outage: the probability that the first passage $T_{0\to 2}$ from none up to both up is at most 24 h, whose distribution function given the rates is written $F_{0\to 2}(t;\mathbf{z})$, averaged over the posterior. The sixth is a design question between two interventions. A seventh, which record would narrow the band most, asks nothing new of the present records; it is the next measurement for these six, Section 16.6. The atlas of Appendix A lists the same forms.

Table 16.2: The questions asked of the pump pair, the general form of their category, their mathematical form for the pair, their category and where each is answered. $\mathbf{y}$ is the failure log, the repair record and the priors; $\pi_0(\mathbf{z})$ the stationary probability that no pump is up, with $U=\pi_0$; $\mathbf{p}(t)$ the state probabilities at time t; $T_{0\to 2}$ the first-passage time from none up to both up and $F_{0\to 2}$ its distribution function given the state; c the number of crews; a an intervention.

Engineering question	General form	Mathematical form	Category	Answered in
What fraction of the time is the pair out in the long run, and how well is that known?	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$, quantiles	$\mathbb{E}[\pi_0(\mathbf{z}) \mid \mathbf{y}]$ with $\boldsymbol{\pi}\mathbf{Q}(\mathbf{z})=\mathbf{0}$; band of $\ln U$	estimation of a derived quantity (implicit)	16.3.3
How much of that is the common cause, and how much the shared crew?	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$ under a changed drawing	$\mathbb{E}[\pi_0 \mid \mathbf{y}]$ with $\beta_{\mathrm{cc}}=0$; $U_{c=1}-q^2$	estimation (decomposition by drawing)	16.3.3
How long until the pair's first outage, and how often do outages occur?	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$, quantiles	$\mathrm{MTTF}_{\mathrm{pair}}=m_2$ with $(-\mathbf{Q}_{TT})\mathbf{m}=\mathbf{1}$; ν_{out}	estimation of derived quantities	16.3.3
Is the pair's unavailability below 10^{-4}, with assurance 0.90?	$\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}) \geq \kappa$	$\Pr(\pi_0(\mathbf{z}) < 10^{-4} \mid \mathbf{y}) \geq 0.90$	prediction under a threshold	16.3.4
After both pumps are down, how soon is the pair back, and how likely is it to be whole again within a day?	$\int p(y_{\mathrm{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	$\mathbf{p}(t)=\mathbf{p}(0)\,\mathrm{e}^{\mathbf{Q}t}$; $\mathbb{E}[F_{0\to 2}(24\,\mathrm{h};\mathbf{z}) \mid \mathbf{y}]$	prediction under a threshold	16.3.5
Which helps more: a second crew, or halving the common-cause share?	$U(a)$, the design form	$\mathbb{E}[\pi_0(\mathbf{z};a) \mid \mathbf{y}]$ with $c=2$ or $\beta_{\mathrm{cc}}/2$; ν_{out}, $\mathrm{MTTF}_{\mathrm{pair}}$	design	16.3.6

16.1.5 The evidence as rows (step 2)

Step 2 attaches the four rows of Table 16.1 to the graph. The failure-rate row is a count row of 11 events in 84 000 h on z_1, the form of the count rows of Chapter 8. The prior on the repair rate and the repair row are gamma-sum rows on z_2, the prior's two pseudo-repairs in twenty pseudo-hours and the record's ten repairs in 115 h. The share row is a binomial row of two doubles in twenty events on z_3, the form of the prior row in the logit of Chapter 5. Which row acts on which entry fixes where the Jacobian of the rows can be nonzero, marked $*$:

$$\mathbf{J}=\begin{bmatrix} * & 0 & 0 \\ 0 & * & 0 \\ 0 & * & 0 \\ 0 & 0 & * \end{bmatrix} \begin{matrix} \text{failure-rate row} \\ \text{prior on } \mu \\ \text{repair row} \\ \text{share row} \end{matrix}, \qquad \text{columns } z_1, z_2, z_3. \tag{16.2}$$

No row acts on two entries. The log's own rows, the singles and doubles count rows of Chapter 15, would act on z_1 and z_3 together, since both expected counts carry the rate and the share; the summary rows drop that shared column. None of the four rows is Gaussian: each enters the objective as its negative log-probability term and joins the Jacobian in the working form of Section 4.2.7, written below. No row is a physics row either. The chain is not attached as evidence: the stationary equations $\boldsymbol{\pi}\mathbf{Q}(\mathbf{z}) = \mathbf{0}$ are solved exactly at every state as the definition of a derived quantity, fixed by the drawing and by the number of crews, so no model tolerance σ_p arises.

The modeling choices. Three modeling choices are made for the whole model, and each buys something and can fail. Every duration of the model, to a failure and to the end of a repair, is taken as exponential: this is what makes the number of pumps up a Markov chain whose generator is built from the three entries, and it fails if the repairs are much more regular than an exponential duration allows, which leaves the long-run unavailability at first order unchanged but changes the wait for the crew and the transient; Section 16.4.4 tests it. The three entries are given independent posteriors, the failure log summarized by a gamma of the rate and the share by the Beta approximation of Chapter 15: this makes $\boldsymbol{\Sigma}$ diagonal and the variance budget three terms with no cross term, and it misleads by the correlation it drops and the share's low mean, which Section 16.4.3 measures. The answers are computed from the posterior of the state, the Laplace normal, with the chain solved at each of its nodes and the delta method on the logarithm: this carries the state to a quantity whose posterior has no closed form, and it misleads where that normal differs from the exact posterior, which draws with the chain solved at each draw measure beside every answer.

The rows. The four rows of equation (16.2) are count rows in the form of equation (8.9), each the negative log of its record written in the state, with the change-of-variable factor of its entry inside: the failure-rate row of 11 events in 84 000 h, the prior on the repair rate of 2 pseudo-repairs in 20 h, the repair row of the record's 10 repairs in 115 h, and the share row of 2 doubles in 20 events, a binomial row in the logit like the prior row of equation (5.15). With $\lambda = \mathrm{e}^{z_1}$, $\mu = \mathrm{e}^{z_2}$ and $\beta_{\mathrm{cc}} = 1/(1+\mathrm{e}^{-z_3})$, the four terms the rows put in the objective $\mathcal{L} = \sum_k \mathcal{L}_k$ are, up to constants,

$$\begin{gathered}\mathcal{L}_1(\mathbf{z}) = 84\,000\,\mathrm{e}^{z_1} - 11\,z_1, \qquad \mathcal{L}_2(\mathbf{z}) = 20\,\mathrm{e}^{z_2} - 2\,z_2, \qquad \mathcal{L}_3(\mathbf{z}) = 115\,\mathrm{e}^{z_2} - 10\,z_2, \\ \mathcal{L}_4(\mathbf{z}) = -2\ln\beta_{\mathrm{cc}} - 18\ln(1-\beta_{\mathrm{cc}}),\end{gathered} \tag{16.3}$$

with the exposures in hours, and in the working form of equation (4.43), Poisson for the first three rows and binomial in the logit for the share row, the residual vector and the 4×3 Jacobian are

$$\mathbf{r}(\mathbf{z}) = \begin{bmatrix} (84\,000\,\lambda - 11)/\sqrt{84\,000\,\lambda} \\ (20\,\mu - 2)/\sqrt{20\,\mu} \\ (115\,\mu - 10)/\sqrt{115\,\mu} \\ (20\,\beta_{\mathrm{cc}} - 2)/\sqrt{20\,\beta_{\mathrm{cc}}(1-\beta_{\mathrm{cc}})} \end{bmatrix}, \qquad \mathbf{J}(\mathbf{z}) = \begin{bmatrix} \sqrt{84\,000\,\lambda} & 0 & 0 \\ 0 & \sqrt{20\,\mu} & 0 \\ 0 & \sqrt{115\,\mu} & 0 \\ 0 & 0 & \sqrt{20\,\beta_{\mathrm{cc}}(1-\beta_{\mathrm{cc}})} \end{bmatrix}, \tag{16.4}$$

rows in the order of Table 16.1 and columns by entry, the nonzero entries where the pattern (16.2) puts them. Each row is a canonical-link row, the log of its expected count linear in its entry or the logit of its probability its entry, so the curvature term of every row is zero, $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ is the exact Hessian of the objective and the Newton step is the Gauss–Newton step. With every row normalized in this way the information matrix is $\mathbf{J}^{\mathsf{T}}\mathbf{J}$.

16.2 The posterior of the state (step 3)

In the book's notation $p(\mathbf{z} \mid \mathbf{y})$ is the posterior density of the state, $f(\mathbf{z})$ a quantity derived from the state and f^* a requirement level on it.

Four objects of the fit recur in the specialized forms. The *MAP state* $\hat{\mathbf{z}} = (\hat{z}_1, \hat{z}_2, \hat{z}_3)^\mathsf{T}$ is the minimum of the objective that the four rows put on the state; the *information matrix* $\mathbf{\Lambda} = \mathbf{J}^\mathsf{T}\mathbf{J}$ is the curvature of that objective at the minimum, with $\mathbf{J}$ the 4×3 Jacobian of the normalized rows; the *Laplace covariance* $\mathbf{\Sigma} = \mathbf{\Lambda}^{-1}$, a 3×3 matrix, makes the posterior of the state approximately normal, $\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$; and a derived quantity f is carried by its *gradient* $\mathbf{g} = \partial f / \partial \mathbf{z}$ at $\hat{\mathbf{z}}$, with the variance $\mathbf{g}^\mathsf{T}\mathbf{\Sigma}\mathbf{g}$. Beside them stands the exact posterior that this problem admits: three independent distributions, $\lambda \mid \mathbf{y} \sim \text{Gamma}(a_\lambda, b_\lambda)$ with $(a_\lambda, b_\lambda) = (11,\ 84\,000\,\text{h})$, $\mu \mid \mathbf{y} \sim \text{Gamma}(a_\mu, b_\mu)$ with $a_\mu = 2 + r$ and $b_\mu = 20\,\text{h} + D$, and $\beta_{\text{cc}} \mid \mathbf{y} \sim \text{Beta}(a_\beta, b_\beta)$ with $(a_\beta, b_\beta) = (2, 18)$.

Every answer of step 4 is computed from the posterior of the state that the four rows of step 2 give, the Laplace normal of the entries, and the exact posterior, the three independent gamma and Beta distributions, gives the value beside it, with the gap. This section derives both and computes them for the pair; Section 16.3 reads every answer from them.

16.2.1 The posterior from the four rows

The normal equations and the posterior of the state. Each step solves $(\mathbf{J}^\mathsf{T}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^\mathsf{T}\mathbf{r}$ at the current state. No row of $\mathbf{J}$ has two nonzero entries, so $\mathbf{J}^\mathsf{T}\mathbf{J}$ is diagonal at every iterate and each entry converges on its own; the two rows of z_2 add, $\mathbf{J}^\mathsf{T}\mathbf{r}$ having the entry $135\,\mu - 12$ and $\mathbf{J}^\mathsf{T}\mathbf{J}$ the entry $135\,\mu$, the prior's and the record's counts and exposures summed. The gradient vanishes where each entry's rows together expect their counts, $84\,000\,\lambda = 11$, $135\,\mu = 12$ and $20\,\beta_{\text{cc}} = 2$, and there the information is the sum of the expected counts of each column, the form $\mathbf{\Lambda} = \text{diag}(a_\lambda,\ a_\mu)$ of equation (8.15) for the two log rates and the binomial information for the logit:

$$\hat{\mathbf{z}} = \begin{bmatrix} \ln(a_\lambda / b_\lambda) \\ \ln(a_\mu / b_\mu) \\ \ln(a_\beta / b_\beta) \end{bmatrix}, \qquad \mathbf{\Lambda} = \mathbf{J}^\mathsf{T}\mathbf{J} = \begin{bmatrix} a_\lambda & 0 & 0 \\ 0 & a_\mu & 0 \\ 0 & 0 & \dfrac{a_\beta b_\beta}{a_\beta + b_\beta} \end{bmatrix}, \qquad \mathbf{\Sigma} = \mathbf{\Lambda}^{-1}, \tag{16.5}$$

with $a_\mu = 2 + r$, $b_\mu = 20\,\text{h} + D$ and $\Sigma_{33} = 1/a_\beta + 1/b_\beta$. The covariance is diagonal because each record informs one column of $\mathbf{J}$, which is the independence of the three posteriors written in the rows, and the MAP state is the posterior mean of each rate and of the share, mapped to the entries.

The exact posterior. The four terms (16.3) are the negative logarithms of gamma and Beta kernels written in the entries, so the posterior also has a closed form. The repair rate's two rows give the gamma update of equation (8.16), $\mu \mid r, D \sim \text{Gamma}(2 + r,\ 20\,\text{h} + D)$; the failure-rate row is the gamma $\text{Gamma}(a_\lambda, b_\lambda)$ it summarizes; the share row is the Beta $\text{Beta}(a_\beta, b_\beta)$; and the three are independent. The normal $\mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$ is the Laplace approximation of that posterior in the entries, whose exact moments in the entries follow by the digamma relations of Section 8.2.2:

$$\mathbb{E}[\mathbf{z} \mid \mathbf{y}] = \begin{bmatrix} \psi(a_\lambda) - \ln b_\lambda \\ \psi(a_\mu) - \ln b_\mu \\ \psi(a_\beta) - \psi(b_\beta) \end{bmatrix}, \qquad \mathbf{\Sigma}_{\text{ex}} = \text{diag}\big(\psi'(a_\lambda),\ \psi'(a_\mu),\ \psi'(a_\beta) + \psi'(b_\beta)\big), \tag{16.6}$$

the exact mean and covariance of the state, the last variance because the logit of a Beta variable is the difference of the logarithms of two independent gamma variables, by the construction of equation (8.18). The posterior mean repair duration is $b_\mu/(a_\mu - 1)$, by equation (8.27), larger than $1/\mathbb{E}[\mu \mid \mathbf{y}] = b_\mu/a_\mu$ because $1/\mu$ is convex.

16.2.2 The posterior of the pair

The pair's inputs are the four rows of Table 16.1. Table 16.3 lists the normal equations solved from $\mathbf{z}^{(0)} = (-9,\ -3,\ -3)^{\mathsf{T}}$: each row is an iterate, and the columns give the three entries and the three entries of the gradient $\mathbf{J}^{\mathsf{T}}\mathbf{r}$. The entry z_1 converges in two steps, because its starting value is close, and the other two in four.

Table 16.3: The normal equations of the rows (16.4) solved from $\mathbf{z}^{(0)} = (-9,\ -3,\ -3)^{\mathsf{T}}$. The gradient $\mathbf{J}^{\mathsf{T}}\mathbf{r}$ is dimensionless; its entries below 10^{-13} are round-off.

Iteration [–]	z_1 [–]	z_2 [–]	z_3 [–]	$(\mathbf{J}^{\mathsf{T}}\mathbf{r})_1$ [–]	$(\mathbf{J}^{\mathsf{T}}\mathbf{r})_2$ [–]	$(\mathbf{J}^{\mathsf{T}}\mathbf{r})_3$ [–]
0	−9.0000	−3.0000	−3.0000	−0.634	−5.279	−1.051
1	−8.9389	−2.2146	−1.8363	2.0×10^{-2}	2.741	0.750
2	−8.9407	−2.4006	−2.1524	1.8×10^{-5}	0.240	8.2×10^{-2}
3	−8.9407	−2.4202	−2.1964	1.4×10^{-11}	2.3×10^{-3}	1.4×10^{-3}
4	−8.9407	−2.4204	−2.1972	$< 10^{-13}$	2.3×10^{-7}	4.5×10^{-7}
5	−8.9407	−2.4204	−2.1972	$< 10^{-13}$	$< 10^{-13}$	$< 10^{-13}$

At $\hat{\mathbf{z}}$ the failure-rate and share rows expect exactly their counts, while the two rows of z_2 pull in opposite directions, the prior expecting $20 \times 0.0889 = 1.78$ repairs against its 2 and the record 10.22 against its 10. The Jacobian, the information and the covariance are

$$\hat{\mathbf{z}} = \begin{bmatrix} \ln(11/84\,000) \\ \ln(12/135) \\ \ln(2/18) \end{bmatrix} = \begin{bmatrix} -8.941 \\ -2.420 \\ -2.197 \end{bmatrix}, \qquad \mathbf{J}(\hat{\mathbf{z}}) = \begin{bmatrix} 3.3166 & 0 & 0 \\ 0 & 1.3333 & 0 \\ 0 & 3.1972 & 0 \\ 0 & 0 & 1.3416 \end{bmatrix},$$
$$\mathbf{\Lambda} = \mathbf{J}^{\mathsf{T}}\mathbf{J} = \begin{bmatrix} 11 & 0 & 0 \\ 0 & 12 & 0 \\ 0 & 0 & 1.8 \end{bmatrix}, \qquad \mathbf{\Sigma} = \mathbf{\Lambda}^{-1} = \begin{bmatrix} 0.0909 & 0 & 0 \\ 0 & 0.0833 & 0 \\ 0 & 0 & 0.5556 \end{bmatrix}, \tag{16.7}$$

the binomial information of the logit being $20 \times 0.1 \times 0.9 = 1.8$, whose inverse is $1/2 + 1/18$. The entries are read off with $\sqrt{\Sigma_{11}} = 0.302$, $\sqrt{\Sigma_{22}} = 0.289$ and $\sqrt{\Sigma_{33}} = 0.745$, and every correlation is zero. The exact posterior of the repair rate is

$$\mu \mid r, D \sim \text{Gamma}(2 + 10,\ 20\,\text{h} + 115\,\text{h}) = \text{Gamma}(12,\ 135\,\text{h}), \tag{16.8}$$

and Table 16.4 summarizes the three exact posteriors: each row is an entry, and the columns give its posterior, its mean and standard deviation, and its central 95 % interval.

Table 16.4: The exact posteriors of the three entries. Gamma distributions with shape and rate parameter, the rate in hours; the repair rate's interval is also given as mean repair durations.

Entry	Posterior	Mean and std. dev.	95 % interval
failure rate λ	Gamma(11, 84 000 h)	1.310×10^{-4}, $0.395 \times 10^{-4}\,\text{h}^{-1}$	$[0.65,\ 2.19] \times 10^{-4}\,\text{h}^{-1}$
repair rate μ	Gamma(12, 135 h)	0.0889, $0.0257\,\text{h}^{-1}$	$[0.046,\ 0.146]\,\text{h}^{-1}$; $[6.9,\ 21.8]$ h
share β_{cc}	Beta(2, 18)	0.100, 0.065	[0.013, 0.260]

The reciprocal of the mean repair rate, $135/12 = 11.25$ h, is not the posterior mean repair duration, $135/11 = 12.27$ h.

The pair $(\hat{\mathbf{z}}, \mathbf{\Sigma})$ of equation (16.7) is what step 4 uses: every answer is computed from it, and the three exact posteriors of Table 16.4 give the value beside each answer, with the gap between the two.

16.3 The answers (step 4)

The six questions of Table 16.2 fall into three categories of Chapter 2 and take four of its general forms: estimation (the first three); prediction under a threshold, as a tail of the posterior (the fourth) and as a tail of the predictive of a future first passage (the fifth); and design (the sixth). Each is answered from the posterior of Section 16.2; the first two subsections derive the chain of the pair at given rates and the gradients of its quantities, and each question then gives its engineering question, the general form of its category, the form it takes for the pair, the quantities of the posterior and of the chain it uses, and the pair's numbers. The questions name the number of crews, one or two; the limit $u^* = 10^{-4}$ at the assurance $\kappa = 0.90$; and the times of the transient. Every answer is computed from the posterior of the state of step 3, the MAP state $\hat{\mathbf{z}}$ and the covariance $\boldsymbol{\Sigma}$, through a Gauss–Hermite sum with the chain solved at each node and the delta method with its gradient, and that is the value the chapter reports. The value beside it comes from 2×10^4 draws from the three independent exact posteriors, with a stationary solve, a first-passage solve and the matrix exponentials at each draw, and the gap is stated with it.

16.3.1 The chain of the pair

Every question below reads the state through one object, the chain of the pair at given rates.

The Markov property. Every duration in the model is exponential: the time to the next independent failure of a running pump, with rate $(1-\beta_{\mathrm{cc}})\lambda$; the time to the next common-cause event, with rate $\beta_{\mathrm{cc}}\lambda$; and the time to the completion of a repair, with rate μ. A repair duration T_r of rate μ is memoryless, since

$$\Pr(T_r > s + t \mid T_r > s) = \frac{\mathrm{e}^{-\mu(s+t)}}{\mathrm{e}^{-\mu s}} = \mathrm{e}^{-\mu t} = \Pr(T_r > t),$$

and the same holds for every exponential duration of the model: a repair that has already lasted s hours has the same remaining duration as a repair just begun, and a pump that has run for s hours fails in the next interval with the same probability as a new one. Given the number of pumps up at time t, the probabilities of every future transition therefore do not depend on how long the pair has been in its state or on how it got there, and $X(t)$ is a Markov chain.

The transition rates. In a short interval of length Δt a single event occurs with a probability proportional to Δt, and two events with a probability of order Δt^2, written $o(\Delta t)$. From two pumps up, either pump fails alone with probability $(1-\beta_{\mathrm{cc}})\lambda\Delta t$, so the pair passes to one up with probability $2(1-\beta_{\mathrm{cc}})\lambda\Delta t$, and the shared cause takes it to none up with probability $\beta_{\mathrm{cc}}\lambda\Delta t$. From one up, the running pump fails, by its own cause or by the shared one, which with one pump already down takes out the other, with probability $\lambda\Delta t$, and the crew completes the repair with probability $\mu\Delta t$. From none up, one crew completes a repair with probability $\mu\Delta t$, and two crews, one per pump, complete one of the two with probability $2\mu\Delta t$, the first of two parallel repairs ending the outage. The last rate is the only place where the number of crews enters the chain. Writing Q_{ij} for the rate from state i to state $j \neq i$, with states named by the number of pumps up,

$$Q_{21} = 2(1-\beta_{\mathrm{cc}})\lambda, \quad Q_{20} = \beta_{\mathrm{cc}}\lambda, \quad Q_{10} = \lambda, \quad Q_{12} = \mu, \quad Q_{01} = c\,\mu, \quad Q_{02} = 0. \tag{16.9}$$

The generator matrix. Each diagonal entry is minus the total rate out of its state, $Q_{ii} = -\sum_{j\neq i} Q_{ij}$, so that every row sums to zero. With the states in the order two, one, none,

$$\mathbf{Q}(\mathbf{z}) = \begin{bmatrix} -(2-\beta_{\mathrm{cc}})\lambda & 2(1-\beta_{\mathrm{cc}})\lambda & \beta_{\mathrm{cc}}\lambda \\ \mu & -(\lambda+\mu) & \lambda \\ 0 & c\,\mu & -c\,\mu \end{bmatrix}, \qquad \mathbf{Q}\,\mathbf{1} = \mathbf{0}, \tag{16.10}$$

with $\mathbf{1}$ the column of ones. The generator is written $\mathbf{Q}$, bold as every matrix of the book; it is not the heat load Q of Chapters 1 to 4, which does not appear in this chapter. The generator depends on the state through $\lambda = \mathrm{e}^{z_1}$, $\mu = \mathrm{e}^{z_2}$ and $\beta_{\mathrm{cc}} = 1/(1+\mathrm{e}^{-z_3})$, and it is linear in each rate: $\mathbf{Q} = \lambda\,\mathbf{Q}_\lambda(\beta_{\mathrm{cc}}) + \mu\,\mathbf{Q}_\mu$, where $\mathbf{Q}_\lambda$ holds the failure terms per unit of λ and $\mathbf{Q}_\mu$ the repair terms per unit of μ. The two-state model of Chapter 8, equation (8.50), has the generator $\left[\begin{smallmatrix} -\lambda & \lambda \\ \mu & -\mu \end{smallmatrix}\right]$ of one pump.

The forward equations. Write $p_j(t) = \Pr(X(t) = j)$ for the probability of state j at time t, and $\mathbf{p}(t) = (p_2(t), p_1(t), p_0(t))$ as a row vector. Chapter 8 derived the forward equations of one pump, equation (8.50), by a balance of probability over Δt; the same balance applies to each state of the pair. The chain is in state j at $t+\Delta t$ if it was in j at t and did not leave, or was in some $i \neq j$ and jumped to j. By the law of total probability over the state at t,

$$\begin{aligned} p_j(t+\Delta t) &= p_j(t)\Big(1 - \sum_{k\neq j} Q_{jk}\Delta t\Big) + \sum_{i\neq j} p_i(t)\,Q_{ij}\Delta t + o(\Delta t) \\ &= p_j(t) + \sum_i p_i(t)\,Q_{ij}\Delta t + o(\Delta t), \end{aligned}$$

using $Q_{jj} = -\sum_{k\neq j} Q_{jk}$ to bring the first term into the sum. Subtract $p_j(t)$, divide by Δt and let it shrink:

$$\frac{\mathrm{d}p_j}{\mathrm{d}t} = \sum_i p_i(t)\,Q_{ij}, \qquad \frac{\mathrm{d}\mathbf{p}}{\mathrm{d}t} = \mathbf{p}(t)\,\mathbf{Q}, \tag{16.11}$$

the forward (Kolmogorov) equations of the chain. Multiplying on the right by $\mathbf{1}$ and using $\mathbf{Q1} = \mathbf{0}$ gives $\mathrm{d}(\mathbf{p1})/\mathrm{d}t = 0$: the probabilities keep summing to one. Written for the column $\mathbf{P}(t) = \mathbf{p}(t)^{\mathsf{T}} = (p_2, p_1, p_0)^{\mathsf{T}}$, the form in which Chapter 8 wrote the generator of one pump, the same equations read

$$\frac{\mathrm{d}\mathbf{P}}{\mathrm{d}t} = \mathbf{Q}^{\mathsf{T}}\mathbf{P}, \tag{16.12}$$

the form in which the sensitivities of Section 16.3.2 are derived. Written out for the pair, the equation of the state with none up is

$$\frac{\mathrm{d}p_0}{\mathrm{d}t} = \beta_{\mathrm{cc}}\lambda\,p_2 + \lambda\,p_1 - c\,\mu\,p_0,$$

probability flowing in from two up by the common cause and from one up by a failure, and out by a repair.

The stationary equations. A distribution $\boldsymbol{\pi} = (\pi_2, \pi_1, \pi_0)$ that the forward equations leave unchanged satisfies

$$\boldsymbol{\pi}\,\mathbf{Q} = \mathbf{0}, \qquad \boldsymbol{\pi}\,\mathbf{1} = 1. \tag{16.13}$$

The three equations $\boldsymbol{\pi}\mathbf{Q} = \mathbf{0}$ are linearly dependent, since the columns of $\mathbf{Q}$ sum to the zero column by $\mathbf{Q1} = \mathbf{0}$, and the normalization supplies the third independent equation. Every state of the pair can reach every other, so the chain is irreducible, and a standard result for finite irreducible chains states that the solution is unique, that it is the limit of $\mathbf{p}(t)$ from any initial

state, and that π_j is the long-run fraction of time spent in state j. The long-run unavailability of the pair is therefore

$$U(\mathbf{z}) = \pi_0(\mathbf{z}), \tag{16.14}$$

the quantity written U_{pair} in Chapter 15; the subscript is dropped, since the pair is the only system here.

A square system. Because one of the three equations $\boldsymbol{\pi}\mathbf{Q} = \mathbf{0}$ is redundant, replace the last column of $\mathbf{Q}$, the one of state none up, by the column of ones, and call the result $\mathbf{A}$. The stationary equations and the normalization together become one square system,

$$\boldsymbol{\pi}\,\mathbf{A}(\mathbf{z}) = \mathbf{e}_0^{\mathsf{T}}, \qquad \mathbf{e}_0 = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}, \tag{16.15}$$

with $\mathbf{e}_0$ the indicator of the last state, none up. The matrix is nonsingular. If a row vector $\mathbf{x}$ has $\mathbf{xA} = \mathbf{0}$, the first two entries of $\mathbf{xQ}$ vanish, since $\mathbf{A}$ shares the first two columns of $\mathbf{Q}$, and the entries of $\mathbf{xQ}$ sum to $\mathbf{xQ1} = 0$, so the third vanishes too: $\mathbf{xQ} = \mathbf{0}$. For an irreducible chain that makes $\mathbf{x}$ a multiple $\alpha\boldsymbol{\pi}$, and the column of ones gives $\mathbf{x1} = \alpha = 0$, so $\mathbf{x} = \mathbf{0}$. One factorization of $\mathbf{A}$ gives $\boldsymbol{\pi}$, at every state and every draw.

The balance equations across a cut. Let $\mathcal{S}$ be a set of states. Summing the stationary equations over $j \in \mathcal{S}$,

$$0 = \sum_{j\in\mathcal{S}} \sum_{i} \pi_i Q_{ij} = \sum_{i\in\mathcal{S}} \pi_i \sum_{j\in\mathcal{S}} Q_{ij} + \sum_{i\notin\mathcal{S}} \pi_i \sum_{j\in\mathcal{S}} Q_{ij},$$

and for $i \in \mathcal{S}$ the zero row sum gives $\sum_{j\in\mathcal{S}} Q_{ij} = -\sum_{j\notin\mathcal{S}} Q_{ij}$. Hence

$$\sum_{i\in\mathcal{S}} \sum_{j\notin\mathcal{S}} \pi_i Q_{ij} = \sum_{i\notin\mathcal{S}} \sum_{j\in\mathcal{S}} \pi_i Q_{ij} : \tag{16.16}$$

in stationarity the probability flow out of any set of states equals the flow into it. Each choice of $\mathcal{S}$ is one balance equation, and a good choice makes each equation involve few unknowns.

The closed form. Take first $\mathcal{S} = \{2\}$: the flow out of two up is $\pi_2(2 - \beta_{\text{cc}})\lambda$, by either kind of failure, and the flow in is $\pi_1\mu$, by a repair. Then take $\mathcal{S} = \{0\}$: the flow out is $\pi_0 c\mu$ and the flow in $\pi_2\beta_{\text{cc}}\lambda + \pi_1\lambda$. With the ratio of the failure rate to the repair rate of Chapter 15,

$$\epsilon = \frac{\lambda}{\mu},$$

the two equations give

$$\pi_1 = (2 - \beta_{\text{cc}})\,\epsilon\,\pi_2, \qquad \pi_0 = \frac{\beta_{\text{cc}}\epsilon + (2 - \beta_{\text{cc}})\epsilon^2}{c}\,\pi_2, \tag{16.17}$$

and the normalization $\pi_2 + \pi_1 + \pi_0 = 1$ completes the solution:

$$U = \pi_0 = \frac{\left[\beta_{\text{cc}}\epsilon + (2 - \beta_{\text{cc}})\epsilon^2\right]/c}{1 + (2 - \beta_{\text{cc}})\epsilon + \left[\beta_{\text{cc}}\epsilon + (2 - \beta_{\text{cc}})\epsilon^2\right]/c}. \tag{16.18}$$

The stationary distribution depends on the two rates only through their ratio ϵ, since multiplying $\mathbf{Q}$ by a constant, which changes the unit of time, leaves $\boldsymbol{\pi}\mathbf{Q} = \mathbf{0}$ unchanged.

The birth–death chain without a common cause. With $\beta_{\mathrm{cc}} = 0$ the chain moves only one step at a time, a birth–death chain, and equation (16.18) becomes, for one crew and for two,

$$U_{c=1} = \frac{2\epsilon^2}{1 + 2\epsilon + 2\epsilon^2}, \qquad U_{c=2} = \frac{\epsilon^2}{1 + 2\epsilon + \epsilon^2} = \left(\frac{\lambda}{\lambda + \mu}\right)^2. \tag{16.19}$$

The second is the square of the unavailability of one pump, equation (8.2): with a crew for each pump and no shared cause the two pumps are independent two-state processes, and the chain reproduces the product q^2, as it must. The first is about twice the second: with one crew, a pump that fails while the other is under repair waits for the crew, and the time both are down doubles at leading order.

The expansion in ϵ. For the pumps ϵ is of order 10^{-3}. Writing equation (16.18) as $N/(1 + (2 - \beta_{\mathrm{cc}})\epsilon + N)$ with N its numerator, expanding $1/(1 + x) = 1 - x + O(x^2)$ and keeping the terms in ϵ and ϵ^2,

$$U = \frac{\beta_{\mathrm{cc}}\,\epsilon + K(\beta_{\mathrm{cc}})\,\epsilon^2}{c} + O(\epsilon^3), \qquad K(\beta_{\mathrm{cc}}) = (2 - \beta_{\mathrm{cc}})(1 - \beta_{\mathrm{cc}}) - \frac{\beta_{\mathrm{cc}}^2}{c}, \tag{16.20}$$

with $K = 2 - 3\beta_{\mathrm{cc}}$ for one crew. The common cause enters at first order and the independent double failures at second, as in equation (15.43) of Chapter 15, whose first-order term $\beta_{\mathrm{cc}}\epsilon$ is the same for one crew. For one crew the chain's second-order coefficient is $2 - 3\beta_{\mathrm{cc}}$ against that chapter's $(1 - \beta_{\mathrm{cc}})^2 - \beta_{\mathrm{cc}}^2 = 1 - 2\beta_{\mathrm{cc}}$; the difference, $1 - \beta_{\mathrm{cc}} = (1 - \beta_{\mathrm{cc}})^2 + \beta_{\mathrm{cc}}(1 - \beta_{\mathrm{cc}})$, is the second failed pump waiting for the crew and the running pump failing by the shared cause. The number of crews divides the whole expansion by c and enters K through the term β_{cc}^2/c.

Outage frequency, mean down time and unavailability. An outage begins each time the chain enters the state with none up. In the long run the chain is in state i a fraction π_i of the time and leaves it for none up at the rate Q_{i0}, so the expected number of outages per hour is

$$\nu_{\mathrm{out}} = \pi_2 Q_{20} + \pi_1 Q_{10} = \pi_2\,\beta_{\mathrm{cc}}\lambda + \pi_1\,\lambda, \tag{16.21}$$

the outage frequency of Chapter 8, equation (8.47), now for the pair; the first term is the common-cause outages and the second the independent doubles. Each outage lasts as long as the chain stays at none up, an exponential time with the rate $c\mu$ of its only exit, so the mean down time is

$$\mathrm{MDT} = \frac{1}{c\,\mu}. \tag{16.22}$$

The balance equation (16.16) across the cut $\mathcal{S} = \{0\}$ reads $\pi_0 c\mu = \pi_2\beta_{\mathrm{cc}}\lambda + \pi_1\lambda$, and by the two definitions it is

$$U = \pi_0 = \nu_{\mathrm{out}} \times \mathrm{MDT}\,: \tag{16.23}$$

the long-run unavailability is the outage frequency times the mean down time. For one unit, Chapter 8 found the same identity, $\nu_{\mathrm{out}}/\mu = U$, below equation (8.47); the companion volume states the relation for any repairable system as the long-run fraction of time down; here it is one balance equation of the chain.

The first-passage time to an outage. The mean time to failure of the pair is the expected time from both pumps up until the first entry into none up. Let m_i be the expected time to reach none up from state i, for the two states that precede it, $i \in \{2, 1\}$. From state i the chain stays an exponential time of mean $1/(-Q_{ii})$ and then jumps to j with probability $Q_{ij}/(-Q_{ii})$;

if j is none up the passage is complete, and otherwise it starts again from j. By the law of total expectation over the first jump,

$$m_i = \frac{1}{-Q_{ii}} + \sum_{j\in\{2,1\},\ j\neq i} \frac{Q_{ij}}{-Q_{ii}}\, m_j.$$

Multiplying by $-Q_{ii}$ and collecting the terms in m, $-Q_{ii}m_i - \sum_{j\neq i} Q_{ij}m_j = 1$, which for the two states together is the linear system

$$(-\mathbf{Q}_{TT})\,\mathbf{m} = \mathbf{1}, \qquad -\mathbf{Q}_{TT} = \begin{bmatrix} (2-\beta_{\mathrm{cc}})\lambda & -2(1-\beta_{\mathrm{cc}})\lambda \\ -\mu & \lambda+\mu \end{bmatrix}, \quad \mathbf{m} = \begin{bmatrix} m_2 \\ m_1 \end{bmatrix}, \tag{16.24}$$

with $\mathbf{Q}_{TT}$ the block of the generator on the states that precede the outage. Solving the two equations by elimination,

$$\mathrm{MTTF}_{\mathrm{pair}} = m_2 = \frac{(3-2\beta_{\mathrm{cc}})\lambda+\mu}{\lambda\big((2-\beta_{\mathrm{cc}})\lambda+\beta_{\mathrm{cc}}\mu\big)} \approx \frac{1}{\beta_{\mathrm{cc}}\lambda + 2(1-\beta_{\mathrm{cc}})^2\lambda\,\epsilon}, \tag{16.25}$$

the approximation to first order in ϵ of the rate $1/m_2$, obtained by dividing numerator and denominator by μ and expanding $1/(1+(3-2\beta_{\mathrm{cc}})\epsilon)$. The pair fails at about the rate $2(1-\beta_{\mathrm{cc}})\lambda\times\epsilon$ at which one pump fails alone and the other fails during its repair, with probability $\lambda/(\lambda+\mu)\approx\epsilon$, plus the rate $\beta_{\mathrm{cc}}\lambda$ of the shared cause, which strikes only while both pumps are up, a fraction $1-2(1-\beta_{\mathrm{cc}})\epsilon$ of the time to first order; the two corrections together give the factor $(1-\beta_{\mathrm{cc}})^2$. Without the common cause the exact form becomes $(3\lambda+\mu)/(2\lambda^2)$. The number of crews does not appear: the time to the first outage is set by what happens before the pair reaches none up, and the crews matter only after. The same system written for the target both up, on the states one and none, gives the mean time from a double outage back to both up.

The transient and its time constants. The forward equation (16.11) is linear with constant coefficients. Its solution from an initial distribution $\mathbf{p}(0)$ is

$$\mathbf{p}(t) = \mathbf{p}(0)\,\mathrm{e}^{\mathbf{Q}t}, \qquad \mathrm{e}^{\mathbf{Q}t} = \sum_{n=0}^{\infty} \frac{(\mathbf{Q}t)^n}{n!}, \tag{16.26}$$

since differentiating the series term by term gives $\mathrm{d}\mathrm{e}^{\mathbf{Q}t}/\mathrm{d}t = \mathrm{e}^{\mathbf{Q}t}\mathbf{Q}$, and $\mathrm{e}^{\mathbf{Q}0}$ is the identity. Row i of $\mathrm{e}^{\mathbf{Q}t}$ is the distribution of the state at t given state i at time zero. For a pair just struck by a double outage, $\mathbf{p}(0) = (0,0,1)$ and $\mathbf{p}(t)$ is the last row. $\mathbf{Q}$ has the eigenvalue 0, with right eigenvector $\mathbf{1}$ and left eigenvector $\boldsymbol{\pi}$, and two further eigenvalues, negative, written $-1/t_1$ and $-1/t_2$. Minus their sum is the trace and their product the sum of the principal 2×2 minors of $\mathbf{Q}$:

$$\frac{1}{t_1}+\frac{1}{t_2} = (2-\beta_{\mathrm{cc}})\lambda+\lambda+(1+c)\mu, \qquad \frac{1}{t_1 t_2} = (2-\beta_{\mathrm{cc}})\lambda^2+\beta_{\mathrm{cc}}\lambda\mu+(2-\beta_{\mathrm{cc}})c\lambda\mu+c\mu^2. \tag{16.27}$$

With distinct eigenvalues, $\mathrm{e}^{\mathbf{Q}t}$ is the sum of three terms, one per eigenvalue, and every state probability has the form

$$p_j(t) = \pi_j + a_j\,\mathrm{e}^{-t/t_1} + b_j\,\mathrm{e}^{-t/t_2},$$

with coefficients fixed by $\mathbf{p}(0)$: the distribution approaches its stationary value through two exponential modes with the time constants t_1 and t_2, the extension to three states of the time constant $1/(\lambda+\mu)$ of equation (8.53). To leading order in ϵ the product and the sum reduce to $c\mu^2$ and $(1+c)\mu$, so that the rates $1/t_k$ are μ and $c\mu$: the time constants are set by the repair, near $1/\mu$ for one crew, and near $1/\mu$ and $1/(2\mu)$ for two.

Principle 16.1. *A repairable system with exponential failure and repair times is a Markov chain on its states of service. Its long-run unavailability is an entry of the stationary distribution, which solves* $\boldsymbol{\pi}\mathbf{Q} = \mathbf{0}$ *and is read from the balance equations across cuts; it equals the outage frequency times the mean down time; the mean time to the first outage solves a second linear system on the states that precede the outage; and the approach to the long run follows the eigenvalues of the generator.*

16.3.2 The gradients of the chain's quantities

The delta method needs the gradient of each derived quantity with respect to the three entries at the MAP state.

The gradient matrix. The object the delta method needs is the derivative of the whole stationary distribution with respect to the whole state, the 3×3 matrix with one row per state of the chain and one column per entry,

$$\boldsymbol{\Pi}_{\mathbf{z}} = \frac{\partial \boldsymbol{\pi}^{\mathsf{T}}}{\partial \mathbf{z}^{\mathsf{T}}} = \begin{bmatrix} \partial\pi_2/\partial z_1 & \partial\pi_2/\partial z_2 & \partial\pi_2/\partial z_3 \\ \partial\pi_1/\partial z_1 & \partial\pi_1/\partial z_2 & \partial\pi_1/\partial z_3 \\ \partial\pi_0/\partial z_1 & \partial\pi_0/\partial z_2 & \partial\pi_0/\partial z_3 \end{bmatrix}, \tag{16.28}$$

whose column i is the change of the three stationary probabilities per unit change of the entry z_i, and whose row j is the gradient of π_j.

Differentiating the system. The solution $\boldsymbol{\pi}(\mathbf{z})$ satisfies equation (16.15) at every state, so the derivative of the left side with respect to each entry is zero. By the product rule, for the entry z_i, $(\partial\boldsymbol{\pi}/\partial z_i)\,\mathbf{A} + \boldsymbol{\pi}\,\partial\mathbf{A}/\partial z_i = \mathbf{0}$, and since the column of ones does not depend on the state, $\partial\mathbf{A}/\partial z_i$ is $\partial\mathbf{Q}/\partial z_i$ with its last column set to zero, written $(\partial\mathbf{Q}/\partial z_i)^{\circ}$. Transposing each of the three equations and setting the three columns side by side gives one matrix equation for $\boldsymbol{\Pi}_{\mathbf{z}}$,

$$\mathbf{A}^{\mathsf{T}}\,\boldsymbol{\Pi}_{\mathbf{z}} = -\mathbf{B}, \qquad \mathbf{B} = \left[\left(\frac{\partial \mathbf{Q}}{\partial z_1}\right)^{\circ\mathsf{T}} \boldsymbol{\pi}^{\mathsf{T}} \quad \left(\frac{\partial \mathbf{Q}}{\partial z_2}\right)^{\circ\mathsf{T}} \boldsymbol{\pi}^{\mathsf{T}} \quad \left(\frac{\partial \mathbf{Q}}{\partial z_3}\right)^{\circ\mathsf{T}} \boldsymbol{\pi}^{\mathsf{T}} \right], \tag{16.29}$$

a linear system with the matrix $\mathbf{A}^{\mathsf{T}}$ of the stationary solve, its factorization reused, and three right-hand sides, one per entry: the whole gradient costs one more solve. The first two entries of column i of $\mathbf{B}$ are the change of the net stationary probability flow into states two and one that a unit change of z_i produces with $\boldsymbol{\pi}$ held fixed, and the solve finds the change of $\boldsymbol{\pi}$ that cancels it. The last row belongs to the normalization, not to a flow: the last row of $\mathbf{A}^{\mathsf{T}}$ is the row of ones and the last entry of each column of $\mathbf{B}$ is zero, because the right side of the normalization does not depend on the state, so $\mathbf{1}^{\mathsf{T}}\boldsymbol{\Pi}_{\mathbf{z}} = \mathbf{0}^{\mathsf{T}}$ and each column of the gradient matrix sums to zero.

The gradient of a stationary quantity. Every stationary quantity of the pair is a function of $\boldsymbol{\pi}$, and by the chain rule its gradient with respect to the state is $\boldsymbol{\Pi}_{\mathbf{z}}^{\mathsf{T}}$ times its gradient with respect to $\boldsymbol{\pi}$. For a quantity linear in the stationary distribution, $f = \mathbf{v}^{\mathsf{T}}\boldsymbol{\pi}^{\mathsf{T}}$ with a weight v_j per state (written $\mathbf{v}$ to keep it apart from the number of crews c), and for the log unavailability,

$$\mathbf{g}_f = \frac{\partial f}{\partial \mathbf{z}} = \boldsymbol{\Pi}_{\mathbf{z}}^{\mathsf{T}}\,\mathbf{v}, \qquad \mathbf{g} = \frac{\partial \ln U}{\partial \mathbf{z}} = \frac{1}{\pi_0}\,\boldsymbol{\Pi}_{\mathbf{z}}^{\mathsf{T}}\,\mathbf{e}_0, \tag{16.30}$$

so that $\mathbf{g}$ is the last row of $\boldsymbol{\Pi}_{\mathbf{z}}$, transposed, divided by π_0. The weights $\mathbf{v} = (1, 1, 0)^{\mathsf{T}}$ give the availability $\pi_2 + \pi_1 = 1 - U$, and $\mathbf{v} = (2, 1, 0)^{\mathsf{T}}$ the expected number of pumps up; one matrix $\boldsymbol{\Pi}_{\mathbf{z}}$ serves them all.

The derivatives of the generator. By the chain rule through the transformed entries and the linearity of $\mathbf{Q}$ in each rate,

$$\frac{\partial \mathbf{Q}}{\partial z_1} = \lambda\, \mathbf{Q}_\lambda, \qquad \frac{\partial \mathbf{Q}}{\partial z_2} = \mu\, \mathbf{Q}_\mu, \qquad \frac{\partial \mathbf{Q}}{\partial z_3} = \beta_{\mathrm{cc}}(1-\beta_{\mathrm{cc}})\,\lambda \begin{bmatrix} 1 & -2 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \tag{16.31}$$

using $\mathrm{d}\lambda/\mathrm{d}z_1 = \lambda$, $\mathrm{d}\mu/\mathrm{d}z_2 = \mu$ and $\mathrm{d}\beta_{\mathrm{cc}}/\mathrm{d}z_3 = \beta_{\mathrm{cc}}(1-\beta_{\mathrm{cc}})$; the last matrix is the derivative of the first row of equation (16.10) with respect to the share. These three matrices are the gradient of the generator, one 3×3 slice per entry.

A null vector of the gradient matrix. The first two slices sum to the generator itself, $\partial \mathbf{Q}/\partial z_1 + \partial \mathbf{Q}/\partial z_2 = \mathbf{Q}$. Applied to the direction $\mathbf{n} = (1, 1, 0)^{\mathsf{T}}$ in the state, which raises both log rates by the same amount and leaves the share, the right side of equation (16.29) is therefore $-\mathbf{B}\mathbf{n} = -(\mathbf{Q}^{\circ})^{\mathsf{T}}\boldsymbol{\pi}^{\mathsf{T}}$, which is zero because $\boldsymbol{\pi}\mathbf{Q} = \mathbf{0}$. Hence

$$\boldsymbol{\Pi}_{\mathbf{z}}\, \mathbf{n} = \mathbf{0}, \qquad \mathbf{n} = \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}: \tag{16.32}$$

the first two columns of $\boldsymbol{\Pi}_{\mathbf{z}}$ are opposite, and the gradient (16.30) of every stationary quantity has $g_2 = -g_1$. Raising both rates by one factor changes only the unit of time, which the long-run fractions do not see; this is the matrix form of the statement that $\boldsymbol{\pi}$ depends on the two rates only through ϵ. For a chain whose rates are not all multiples of λ or μ, such as the degraded chain of Section 16.6.5, the null vector is the direction that raises every log rate together, and the entries of $\mathbf{g}$ on all the log rates sum to zero; entries that are not log rates, such as the logit of the share, carry a zero in $\mathbf{n}$.

The first-passage mean. The mean times $\mathbf{m}$ solve $(-\mathbf{Q}_{TT}(\mathbf{z}))\,\mathbf{m} = \mathbf{1}$, a linear system whose matrix depends on the state and whose right side does not. Write $\mathbf{M}_{\mathbf{z}} = \partial \mathbf{m}/\partial \mathbf{z}^{\mathsf{T}}$ for the 2×3 matrix with one row per state that precedes the outage and one column per entry. Differentiating the system with respect to z_i gives $(-\mathbf{Q}_{TT})\,\partial \mathbf{m}/\partial z_i - (\partial \mathbf{Q}_{TT}/\partial z_i)\,\mathbf{m} = \mathbf{0}$, and the three columns side by side give

$$\begin{aligned} (-\mathbf{Q}_{TT})\,\mathbf{M}_{\mathbf{z}} = \mathbf{C}, \qquad \mathbf{C} &= \left[\frac{\partial \mathbf{Q}_{TT}}{\partial z_1}\,\mathbf{m} \quad \frac{\partial \mathbf{Q}_{TT}}{\partial z_2}\,\mathbf{m} \quad \frac{\partial \mathbf{Q}_{TT}}{\partial z_3}\,\mathbf{m} \right], \\ \mathbf{g}_{\mathrm{MTTF}} = \frac{\partial \ln \mathrm{MTTF}_{\mathrm{pair}}}{\partial \mathbf{z}} &= \frac{1}{m_2}\,\mathbf{M}_{\mathbf{z}}^{\mathsf{T}} \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \end{aligned} \tag{16.33}$$

with $\partial \mathbf{Q}_{TT}/\partial z_i$ the upper-left 2×2 block of the slices (16.31). The direction $\mathbf{n}$ of equation (16.32) is again special: the two log-rate slices sum to $\mathbf{Q}_{TT}$, so $\mathbf{C}\mathbf{n} = \mathbf{Q}_{TT}\mathbf{m} = -\mathbf{1}$ and $\mathbf{M}_{\mathbf{z}}\mathbf{n} = -\mathbf{m}$. Raising both rates by one factor shortens every mean time by that factor, and the entries of $\mathbf{g}_{\mathrm{MTTF}}$ on the two log rates sum to -1.

The sensitivity of the transient. The transient is defined by the forward equation (16.12), a differential equation, and its gradient follows by differentiating that equation. Write $\mathbf{S}(t) = \partial \mathbf{P}(t)/\partial \mathbf{z}^{\mathsf{T}}$ for the 3×3 matrix with one row per state of the chain and one column per entry, the transient counterpart of $\boldsymbol{\Pi}_{\mathbf{z}}$. By the product rule on $\mathrm{d}\mathbf{P}/\mathrm{d}t = \mathbf{Q}^{\mathsf{T}}(\mathbf{z})\,\mathbf{P}$, and since the initial distribution does not depend on the state,

$$\begin{aligned} &\frac{\mathrm{d}}{\mathrm{d}t}\frac{\partial \mathbf{P}}{\partial z_i} = \mathbf{Q}^{\mathsf{T}}\frac{\partial \mathbf{P}}{\partial z_i} + \Big(\frac{\partial \mathbf{Q}}{\partial z_i}\Big)^{\mathsf{T}}\mathbf{P}(t), \qquad \frac{\partial \mathbf{P}}{\partial z_i}(0) = \mathbf{0}, \qquad i = 1, 2, 3, \\ &\frac{\mathrm{d}\mathbf{S}}{\mathrm{d}t} = \mathbf{Q}^{\mathsf{T}}\mathbf{S} + \left[\Big(\frac{\partial \mathbf{Q}}{\partial z_1}\Big)^{\mathsf{T}}\mathbf{P} \quad \Big(\frac{\partial \mathbf{Q}}{\partial z_2}\Big)^{\mathsf{T}}\mathbf{P} \quad \Big(\frac{\partial \mathbf{Q}}{\partial z_3}\Big)^{\mathsf{T}}\mathbf{P} \right], \qquad \mathbf{S}(0) = \mathbf{0}, \end{aligned} \tag{16.34}$$

the sensitivity equation of the chain. Each column obeys the forward equation itself, driven by the source $(\partial\mathbf{Q}/\partial z_i)^{\mathsf{T}}\mathbf{P}(t)$: the probability flow that a unit change of z_i redirects at time t. For the first passage the same equation holds with the generator $\mathbf{Q}'$ of the stopped chain and its slices.

The solution by a block exponential. The equation is linear with constant coefficients, and each column is read from one matrix exponential of twice the size. Let

$$\mathbf{K}_i = \begin{bmatrix} \mathbf{Q}^{\mathsf{T}} & (\partial\mathbf{Q}/\partial z_i)^{\mathsf{T}} \\ \mathbf{0} & \mathbf{Q}^{\mathsf{T}} \end{bmatrix}, \qquad \mathrm{e}^{\mathbf{K}_i t} = \begin{bmatrix} \mathrm{e}^{\mathbf{Q}^{\mathsf{T}} t} & \mathbf{X}_i(t) \\ \mathbf{0} & \mathrm{e}^{\mathbf{Q}^{\mathsf{T}} t} \end{bmatrix}.$$

Every power of a block upper-triangular matrix is block upper-triangular with the powers of the diagonal blocks on its diagonal, so the series $\sum_n (\mathbf{K}_i t)^n/n!$ has the form on the right, with diagonal blocks $\mathrm{e}^{\mathbf{Q}^{\mathsf{T}} t}$ and an upper-right block $\mathbf{X}_i(t)$. The upper-right block of $\mathrm{d}\mathrm{e}^{\mathbf{K}_i t}/\mathrm{d}t = \mathbf{K}_i\,\mathrm{e}^{\mathbf{K}_i t}$ reads $\mathrm{d}\mathbf{X}_i/\mathrm{d}t = \mathbf{Q}^{\mathsf{T}}\mathbf{X}_i + (\partial\mathbf{Q}/\partial z_i)^{\mathsf{T}}\mathrm{e}^{\mathbf{Q}^{\mathsf{T}} t}$, with $\mathbf{X}_i(0) = \mathbf{0}$. Multiplied on the right by $\mathbf{P}(0)$, and with $\mathrm{e}^{\mathbf{Q}^{\mathsf{T}} t}\mathbf{P}(0) = \mathbf{P}(t)$, this is the initial-value problem (16.34) for column i of $\mathbf{S}$, whose solution is unique, so $\mathbf{X}_i(t)\,\mathbf{P}(0)$ is that column. The whole of $\mathbf{S}(t)$ is three 6×6 exponentials, one per entry.

Two relations of the sensitivity. A transient quantity $f = \mathbf{v}^{\mathsf{T}}\mathbf{P}(t)$ has the gradient

$$\mathbf{g}_f(t) = \mathbf{S}(t)^{\mathsf{T}}\mathbf{v}, \qquad \mathbf{1}^{\mathsf{T}}\mathbf{S}(t) = \mathbf{0}^{\mathsf{T}}, \qquad \mathbf{S}(t)\,\mathbf{n} = t\,\mathbf{Q}^{\mathsf{T}}\mathbf{P}(t) = t\,\frac{\mathrm{d}\mathbf{P}}{\mathrm{d}t}, \tag{16.35}$$

and the second and third relations check a computed $\mathbf{S}$. The columns sum to zero because $\mathbf{1}^{\mathsf{T}}\mathbf{P}(t) = 1$ at every state, so its derivative in each entry vanishes. For the third, the direction $\mathbf{n}$ leaves the share and multiplies both rates by e^s, so $\mathbf{Q}(\mathbf{z} + s\mathbf{n}) = \mathrm{e}^s\mathbf{Q}(\mathbf{z})$ and $\mathbf{P}(t;\ \mathbf{z} + s\mathbf{n}) = \mathrm{e}^{\mathrm{e}^s\mathbf{Q}^{\mathsf{T}} t}\mathbf{P}(0) = \mathbf{P}(\mathrm{e}^s t;\ \mathbf{z})$; differentiating at $s = 0$ gives $\mathbf{S}(t)\mathbf{n} = t\,\mathrm{d}\mathbf{P}/\mathrm{d}t$. Unlike the stationary distribution the transient sees the unit of time, so $\mathbf{n}$ is not a null vector of $\mathbf{S}(t)$; its image vanishes only as the transient settles.

16.3.3 The unavailability, its parts and the first outage

Engineering questions. What fraction of the time is the pair out in the long run, and how well is that known? How much of that is the common cause, and how much the shared crew? How long until the pair's first outage, and how often do outages occur? The first asks for the service availability of the pair from its components with the repair capacity as a constraint; the second splits it by drawing; the third asks for the times that a maintenance plan and an outage budget use.

General form. Estimation, the posterior distribution of a derived quantity, equation (2.9), summarized by its location and spread (2.10) and its central interval (2.11); the location reported is the mean,

$$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}] = \int f(\mathbf{z})\,p(\mathbf{z} \mid \mathbf{y})\,\mathrm{d}\mathbf{z}. \tag{16.36}$$

Specialized form. *The long-run unavailability.* The question is equation (16.36) with $\mathbf{z} = (z_1, z_2, z_3)^{\mathsf{T}}$, $p(\mathbf{z} \mid \mathbf{y})$ the posterior of step 3, the Laplace normal $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of equation (16.7), and $f(\mathbf{z}) = U(\mathbf{z}) = \pi_0(\mathbf{z})$, the long-run probability that no pump is up: the last entry of the stationary distribution $\boldsymbol{\pi} = (\pi_2, \pi_1, \pi_0)$ of the chain of the pair, the row vector with $\boldsymbol{\pi}\mathbf{Q}(\mathbf{z}) = \mathbf{0}$ and $\boldsymbol{\pi}\mathbf{1} = 1$, where $\mathbf{Q}(\mathbf{z})$ is the generator whose rates are $\lambda = \mathrm{e}^{z_1}$, $\mu = \mathrm{e}^{z_2}$, $\beta_{\mathrm{cc}} = 1/(1 + \mathrm{e}^{-z_3})$

and the number of crews c (Section 16.3.1). By the law of total probability over the state, as in Section 13.3.1,

$$\Pr(\text{pair out at a random instant} \mid \mathbf{y}) = \int \pi_0(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z} = \mathbb{E}[U(\mathbf{z}) \mid \mathbf{y}], \tag{16.37}$$

the predictive probability, which averages the aleatory condition of the pair and the epistemic uncertainty of its rates together. It is the number an expected loss needs, and 8760 h times it is the expected number of hours per year without cooling supply.

The integral over the state is the Gauss–Hermite sum of Chapter 5, equation (5.33), on the three independent entries: with the nodes t_i and weights w_i of a rule of Q nodes in each direction,

$$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}] \approx \frac{1}{\pi^{3/2}} \sum_i \sum_j \sum_l w_i w_j w_l\, f(\mathbf{z}_{ijl}), \qquad \mathbf{z}_{ijl} = \hat{\mathbf{z}} + \sqrt{2} \begin{bmatrix} \sqrt{\Sigma_{11}}\, t_i \\ \sqrt{\Sigma_{22}}\, t_j \\ \sqrt{\Sigma_{33}}\, t_l \end{bmatrix}, \tag{16.38}$$

with the stationary equations solved at each node; the reported values use $Q = 40$.

The delta method and the variance budget. The unavailability is small and positive and, by its expansion in ϵ, nearly a product of powers of the rates and the share, so its logarithm is nearly linear in the state and the delta method of Chapter 4 belongs on $\ln U$, as Section 13.3.1 established for a service unavailability:

$$\ln U \mid \mathbf{y} \approx \mathcal{N}\big(\ln U(\hat{\mathbf{z}}),\ \sigma^2_{\ln U}\big), \qquad \sigma^2_{\ln U} = \mathbf{g}^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g} = g_1^2 \Sigma_{11} + g_2^2 \Sigma_{22} + g_3^2 \Sigma_{33}, \tag{16.39}$$

with $\mathbf{g} = (g_1, g_2, g_3)^{\mathsf{T}} = \partial \ln U / \partial \mathbf{z}$ at $\hat{\mathbf{z}}$, the last row of the gradient matrix $\partial \boldsymbol{\pi}^{\mathsf{T}} / \partial \mathbf{z}^{\mathsf{T}}$ divided by π_0, which implicit differentiation of the stationary equations supplies (Section 16.3.2). With $\boldsymbol{\Sigma}$ diagonal the variance is a budget of three terms and no cross term, one for the failure log's information on the rate, one for the repair record and one for the share. The delta band is $U(\hat{\mathbf{z}})\, \mathrm{e}^{\pm 1.96\, \sigma_{\ln U}}$. A threshold on U is a normal tail of the same normal.

The values beside them come from the exact posterior. The expectation is not the value at the posterior means. Chapter 13 showed that the expectation of a service availability equals its value at the posterior means when the unit unavailabilities have independent posteriors, because the service is multilinear in them. The pair's unavailability is not multilinear in the rates: to second order in $\epsilon = \lambda/\mu$ it is $U \approx [\beta_{\mathrm{cc}} \epsilon + K(\beta_{\mathrm{cc}})\, \epsilon^2]/c$ with $K(\beta_{\mathrm{cc}}) = (2 - \beta_{\mathrm{cc}})(1 - \beta_{\mathrm{cc}}) - \beta_{\mathrm{cc}}^2/c$ (Section 16.3.1), the product $\beta_{\mathrm{cc}} \lambda \mu^{-1}/c$ plus a term in $\lambda^2 \mu^{-2}$, and μ^{-1} and μ^{-2} are convex. With the three entries independent a posteriori, the expectation of each product is the product of the expectations,

$$\mathbb{E}[U \mid \mathbf{y}] \approx \frac{1}{c} \Big(\mathbb{E}[\beta_{\mathrm{cc}}]\, \mathbb{E}[\lambda]\, \mathbb{E}[\mu^{-1}] + \mathbb{E}[K(\beta_{\mathrm{cc}})]\, \mathbb{E}[\lambda^2]\, \mathbb{E}[\mu^{-2}] \Big), \tag{16.40}$$

$$\mathbb{E}[\lambda] = \frac{a_\lambda}{b_\lambda}, \quad \mathbb{E}[\lambda^2] = \frac{a_\lambda (a_\lambda + 1)}{b_\lambda^2}, \quad \mathbb{E}[\mu^{-1}] = \frac{b_\mu}{a_\mu - 1}, \quad \mathbb{E}[\mu^{-2}] = \frac{b_\mu^2}{(a_\mu - 1)(a_\mu - 2)},$$

the gamma moments of Chapters 6 and 8. The repair enters through the posterior mean repair duration $b_\mu/(a_\mu - 1)$, not through the reciprocal of the mean rate, b_μ/a_μ, and without the common cause the unavailability is second order in ϵ and the gap between the expectation and the plug-in value grows by the factor $\mathbb{E}[\mu^{-2}]/\mathbb{E}[\mu]^{-2}$. The leading terms are a check. The exact route draws the state from its posterior and solves the stationary equations at each draw,

$$\mathbb{E}[U \mid \mathbf{y}] \approx \frac{1}{N} \sum_{n=1}^{N} \pi_0\big(\mathbf{z}^{(n)}\big), \qquad N = 2 \times 10^4,$$

with the draws taken from the three independent exact posteriors by a seeded pseudo-random generator; the same draws give the median, the central 95 % band and every threshold probability as fractions of draws.

The center offset. The Laplace normal centers each entry at its mode, and the exact posterior mean of each transformed entry lies below it (equation (16.6)). With the offset vector $\boldsymbol{\delta} = \mathbb{E}[\mathbf{z} \mid \mathbf{y}] - \hat{\mathbf{z}}$, the same linearization moves the center of $\ln U$ by $\mathbf{g}^{\mathsf{T}}\boldsymbol{\delta}$ and gives it the variance $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_{\text{ex}}\mathbf{g}$:

$$\ln U \mid \mathbf{y} \approx \mathcal{N}\big(\ln U(\hat{\mathbf{z}}) + \mathbf{g}^{\mathsf{T}}\boldsymbol{\delta},\ \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_{\text{ex}}\mathbf{g}\big), \tag{16.41}$$

the normal on the exact moments. Which entry moves the center is read off the products $g_i\delta_i$: the two log rates are offset by nearly the same amount and, since the stationary distribution depends on the two rates only through their ratio, enter with opposite entries of $\mathbf{g}$, so their products nearly cancel; the logit of the share has the large offset of a Beta posterior with a small first parameter, entering with g_3.

The parts of the unavailability. The second question changes the drawing and keeps the posterior. The common-cause part is $\mathbb{E}[U \mid \mathbf{y}] - \mathbb{E}[U_{\beta=0} \mid \mathbf{y}]$, with $U_{\beta=0}(\mathbf{z})$ the stationary probability of none up in the chain with the share set to zero, by the same sum and, beside it, the same draws. The part of the shared crew is the excess of $U_{\beta=0}$ with one crew over the product q^2, $q = \lambda/(\lambda+\mu)$, that two independent pumps would give: with a crew per pump and no shared cause the chain reduces to that product (Section 16.3.1), so the difference is the wait of a pump that fails while the other is under repair.

The first outage and the outage frequency. The mean time to the first outage is $\text{MTTF}_{\text{pair}} = m_2$, the first entry of the mean first-passage times $\mathbf{m}$ to none up, which solve the linear system $(-\mathbf{Q}_{TT})\,\mathbf{m} = \mathbf{1}$ on the block of the generator of the states that precede the outage (Section 16.3.1). It is a derived quantity like U, carried by the sum (16.38) and the delta method on $\ln \text{MTTF}_{\text{pair}}$ with the gradient that differentiating the system gives (Section 16.3.2), and read at each draw beside them. The outage frequency is the long-run rate of entries into none up, $\nu_{\text{out}} = \pi_2\beta_{\text{cc}}\lambda + \pi_1\lambda$, each outage lasting the mean down time $1/(c\mu)$, and the balance of the chain across the state with none up makes their product the unavailability (Section 16.3.1).

Building blocks. The MAP state $\hat{\mathbf{z}} = (-8.941,\ -2.420,\ -2.197)^{\mathsf{T}}$ and the covariance $\boldsymbol{\Sigma}$ of equation (16.7); the generator, the stationary solve, its closed form and the first-passage system (Section 16.3.1); the gradient matrix of the stationary distribution and the gradient of the first passage (Section 16.3.2); the sum (16.38); and, for the exact values beside them, the gamma and Beta parameters with the digamma and trigamma values (Sections 16.2.1 and 16.2.2) and the draws.

The unavailability and its parts. At the MAP state the rates are $\lambda = 1.30952 \times 10^{-4}\,\text{h}^{-1}$, $\mu = 0.0888889\,\text{h}^{-1}$ and $\beta_{\text{cc}} = 0.1$, so that $\epsilon = 1.47321 \times 10^{-3}$. The generator of the pair with one crew is, in h^{-1},

$$\mathbf{Q} = \begin{bmatrix} -2.4881 \times 10^{-4} & 2.3571 \times 10^{-4} & 1.3095 \times 10^{-5} \\ 0.088889 & -0.089020 & 1.3095 \times 10^{-4} \\ 0 & 0.088889 & -0.088889 \end{bmatrix}.$$

Table 16.5 sets out the stationary distribution in three cases. Each column is a case: one crew without a common cause, one crew with it, two crews with it. The upper rows give the transition rates of the generator, the lower rows the stationary distribution by the solve (16.15), the unavailability in hours per year, its expansion (16.20) to second order in ϵ and the mean time to the first outage.

Without the common cause the pair with one crew is out 4.3×10^{-6} of the time, 0.04 h a year. A common-cause share of 0.1 raises it to 1.51×10^{-4}, 1.3 h a year, thirty-five times as much, of which the first-order term $\beta_{\text{cc}}\epsilon = 1.47 \times 10^{-4}$ is 98 %. A second crew halves it to 7.6×10^{-5} and leaves the mean time to the first outage unchanged.

Table 16.5: The generator and the stationary distribution of the pump pair at the MAP state, $\lambda = 1.30952 \times 10^{-4}\,\mathrm{h}^{-1}$, $\mu = 0.0888889\,\mathrm{h}^{-1}$, $\beta_{cc} = 0.1$ (or 0), in three cases. Rates in h^{-1}; stationary distribution by the solve of equation (16.15).

Quantity	Expression	One crew, no common cause	One crew, $\beta_{cc} = 0.1$	Two crews, $\beta_{cc} = 0.1$
Q_{21}, two up to one	$2(1-\beta_{cc})\lambda$	2.619×10^{-4}	2.357×10^{-4}	2.357×10^{-4}
Q_{20}, common cause	$\beta_{cc}\lambda$	0	1.310×10^{-5}	1.310×10^{-5}
Q_{10}, one up to none	λ	1.310×10^{-4}	1.310×10^{-4}	1.310×10^{-4}
Q_{12}, repair	μ	0.0889	0.0889	0.0889
Q_{01}, repair	$c\mu$	0.0889	0.0889	0.1778
π_2, both up		0.997058	0.997058	0.997133
π_1, one up		2.938×10^{-3}	2.791×10^{-3}	2.791×10^{-3}
$U = \pi_0$, none up		4.33×10^{-6}	1.510×10^{-4}	7.55×10^{-5}
Hours per year out	$8760\,\mathrm{h} \times U$	0.038	1.32	0.66
Second-order expansion	eq. (16.20)	4.34×10^{-6}	1.5101×10^{-4}	7.551×10^{-5}
$\mathrm{MTTF}_{\mathrm{pair}}$	eq. (16.25)	297 yr	8.5 yr	8.5 yr

The gradient. At the MAP state, with one crew, the three slices (16.31) of the gradient of the generator are, in h^{-1} per unit of the entry,

$$\frac{\partial \mathbf{Q}}{\partial z_1} = \begin{bmatrix} -2.4881 & 2.3571 & 0.13095 \\ 0 & -1.3095 & 1.3095 \\ 0 & 0 & 0 \end{bmatrix} \times 10^{-4}, \qquad \frac{\partial \mathbf{Q}}{\partial z_3} = \begin{bmatrix} 1.1786 & -2.3571 & 1.1786 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \times 10^{-5},$$

$$\frac{\partial \mathbf{Q}}{\partial z_2} = \begin{bmatrix} 0 & 0 & 0 \\ 0.088889 & -0.088889 & 0 \\ 0 & 0.088889 & -0.088889 \end{bmatrix}.$$

The stationary solve gives $\boldsymbol{\pi} = (0.997058,\ 2.79087 \times 10^{-3},\ 1.50999 \times 10^{-4})$, and the system (16.29) has the matrix and the right-hand sides

$$\mathbf{A}^{\mathsf{T}} = \begin{bmatrix} -2.4881 \times 10^{-4} & 0.088889 & 0 \\ 2.3571 \times 10^{-4} & -0.089020 & 0.088889 \\ 1 & 1 & 1 \end{bmatrix},$$

$$\mathbf{B} = \begin{bmatrix} -2.4808 & 2.4808 & 0.11751 \\ 2.3466 & -2.3466 & -0.23502 \\ 0 & 0 & 0 \end{bmatrix} \times 10^{-4},$$

in h^{-1} except the last row of $\mathbf{A}^{\mathsf{T}}$, the dimensionless normalization, and its solution is

$$\boldsymbol{\Pi}_{\mathbf{z}} = \begin{bmatrix} -2.9373 & 2.9373 & 0.0002 \\ 2.7827 & -2.7827 & -0.13220 \\ 0.15467 & -0.15467 & 0.13200 \end{bmatrix} \times 10^{-3}, \qquad \mathbf{g} = \frac{1}{\pi_0}\,\boldsymbol{\Pi}_{\mathbf{z}}^{\mathsf{T}}\mathbf{e}_0 = \begin{bmatrix} 1.024 \\ -1.024 \\ 0.874 \end{bmatrix}, \tag{16.42}$$

rows by state (two, one, none up) and columns by entry, with $\pi_0 = 1.5100 \times 10^{-4}$. The first column says that raising z_1 moves probability out of both up, almost all of it into one up and a twentieth of it into none up; the third, that raising z_3 barely changes both up, 0.0002 against 0.132, and moves probability from one up straight into none up, because a higher share turns single failures, which pass through one up, into doubles, which skip it.

Table 16.6 carries the gradient matrix into the stationary quantities a report needs. Each row is a quantity $f = \mathbf{v}^{\mathsf{T}}\boldsymbol{\pi}^{\mathsf{T}}$; the columns give its weights $\mathbf{v}$, its value, its gradient $\mathbf{g}_f = \boldsymbol{\Pi}_{\mathbf{z}}^{\mathsf{T}}\mathbf{v}$ by equation (16.30), and the share of $\mathbf{g}_f^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_f$ that each entry carries.

Table 16.6: Stationary quantities of the pair with one crew at the MAP state and their gradients $\mathbf{g}_f = \mathbf{\Pi}_\mathbf{z}^\mathsf{T}\mathbf{v}$ from equation (16.42). Budget: $g_i^2\Sigma_{ii}$ as a percentage of $\mathbf{g}_f^\mathsf{T}\mathbf{\Sigma}\mathbf{g}_f$ with $\mathbf{\Sigma}$ of equation (16.7), entries in the order z_1, z_2, z_3.

Quantity	$\mathbf{v}$ [–]	Value [–]	$\mathbf{g}_f^\mathsf{T}$ [10^{-3}]	Budget [%]
both up, π_2	$(1,0,0)$	0.997058	$(-2.937,\ 2.937,\ 0.000)$	52.2, 47.8, 0.0
one up, π_1	$(0,1,0)$	2.791×10^{-3}	$(2.783,\ -2.783,\ -0.132)$	51.8, 47.5, 0.7
none up, $U=\pi_0$	$(0,0,1)$	1.510×10^{-4}	$(0.155,\ -0.155,\ 0.132)$	15.7, 14.4, 69.9
availability, $1-U$	$(1,1,0)$	0.999849	$(-0.155,\ 0.155,\ -0.132)$	15.7, 14.4, 69.9
pumps up, $2\pi_2+\pi_1$	$(2,1,0)$	1.996907	$(-3.092,\ 3.092,\ -0.132)$	51.9, 47.6, 0.6

The fraction of time on one pump and the expected number of pumps up are set by the rate ratio: their budgets split about evenly between the failure log and the repair record, and the share carries under one percent. The unavailability and the availability have the same gradient up to sign, since $\mathbf{v}=(1,1,0)^\mathsf{T}$ is the column of ones minus $\mathbf{e}_0$ and the columns of $\mathbf{\Pi_z}$ sum to zero, and their budget is seventy percent the share. The same budget holds for $\ln U$, whose gradient $\mathbf{g}$ is that of U divided by U: in the variance (16.39) the terms $g_i^2\Sigma_{ii}$ are 16 % the failure rate, 14 % the repair rate and 70 % the common-cause share. With two crews the gradient is the same to three decimals.

The posterior of the unavailability. With $\mathbf{\Sigma}$ of equation (16.7) and the gradient (16.42),

$$\begin{aligned}\sigma^2_{\ln U} = \mathbf{g}^\mathsf{T}\mathbf{\Sigma}\mathbf{g} &= \begin{bmatrix}1.024 & -1.024 & 0.874\end{bmatrix}\begin{bmatrix}0.0909 & 0 & 0\\ 0 & 0.0833 & 0\\ 0 & 0 & 0.5556\end{bmatrix}\begin{bmatrix}1.024\\ -1.024\\ 0.874\end{bmatrix}\\ &= \underbrace{0.095}_{\text{rate}} + \underbrace{0.087}_{\text{repair}} + \underbrace{0.425}_{\text{share}} = 0.607, \qquad \sigma_{\ln U} = 0.779,\end{aligned} \tag{16.43}$$

the standard deviation of U divided by U under the delta method. The delta-method standard deviations $(\mathbf{g}_f^\mathsf{T}\mathbf{\Sigma}\mathbf{g}_f)^{1/2}$ of the stationary quantities of Table 16.6 are 1.23×10^{-3} for both up, 1.17×10^{-3} for one up, 1.18×10^{-4} for the unavailability and for the availability, and 1.29×10^{-3} for the expected number of pumps up. The band of U is $1.510\times10^{-4}\times\mathrm{e}^{\pm1.96\times0.779} = [3.3\times10^{-5},\ 7.0\times10^{-4}]$, a factor of 21 between its ends; with two crews the center is 7.55×10^{-5} and the band $[1.6\times10^{-5},\ 3.5\times10^{-4}]$.

Table 16.7 evaluates the sum (16.38) for $\mathbb{E}[U\mid\mathbf{y}]$ with five nodes. The share is the entry that moves U most, so each row is a node of z_3: the first three columns give the node, its weight and the entry, the fourth the share there, the fifth U with the rates at their MAP values, and the last the average of U over the 5×5 nodes of the two log rates at that share. The middle row is $\hat{z}_3$; its fifth column is the plug-in value, and the rates' average raises it by a tenth, the convexity in the repair time.

The weighted sum of the last column is $\mathbb{E}[U\mid\mathbf{y}]\approx1.97\times10^{-4}$, 1.73 h of lost cooling supply a year; forty nodes change no digit shown. The value beside it comes from the exact posterior: Table 16.8 gives the posterior of the unavailability by the sum and the delta method, with the 2×10^4 draws in parentheses. Each row is a case; the columns give the plug-in value at the MAP state, the posterior mean, the median, the central 95 % band, the mean hours per year out and the probability that the unavailability is below 10^{-4}.

With one crew the pair is out, in expectation, 1.97×10^{-4} of the time, 1.73 h of lost cooling supply a year, with a 95 % band from 3.3×10^{-5} to 7.0×10^{-4}. The expectation is 30 % above the plug-in value, the median of the normal on $\ln U$, because the normal is skewed toward large

Table 16.7: The Gauss–Hermite sum (16.38) for the unavailability with one crew, five nodes in each direction, over $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of equation (16.7), by the nodes of z_3.

Node t_l	Weight $w_l/\sqrt{\pi}$	$z_{3,l}$	$\beta_{\mathrm{cc},l}$	U at the MAP rates	Average over the rates
-2.0202	0.01126	-4.3267	0.0130	2.345×10^{-5}	2.696×10^{-5}
-0.9586	0.22208	-3.2076	0.0389	6.135×10^{-5}	6.826×10^{-5}
0	0.53333	-2.1972	0.1000	1.510×10^{-4}	1.659×10^{-4}
0.9586	0.22208	-1.1868	0.2338	3.473×10^{-4}	3.798×10^{-4}
2.0202	0.01126	-0.0678	0.4831	7.129×10^{-4}	7.781×10^{-4}
Weighted sum					1.971×10^{-4}

Table 16.8: The posterior of the pair's unavailability: the sum (16.38) and the delta method over $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$, with in parentheses 2×10^4 draws of the posteriors of Table 16.4, a stationary solve at each. Plug-in: at the MAP state (16.7).

Case	Plug-in	$\mathbb{E}[U \mid \mathbf{y}]$	Median	95 % band	Hours per year	$\Pr(U < 10^{-4})$
one crew, no common cause	4.33×10^{-6}	6.1×10^{-6} (6.2×10^{-6})	4.3×10^{-6} (4.3×10^{-6})	$[0.84,\ 22.2] \times 10^{-6}$ ($[0.80,\ 22.9] \times 10^{-6}$)	0.054 (0.054)	1.000 (1.000)
one crew, common cause	1.510×10^{-4}	1.97×10^{-4} (1.66×10^{-4})	1.51×10^{-4} (1.29×10^{-4})	$[0.33,\ 7.0] \times 10^{-4}$ ($[0.19,\ 5.2] \times 10^{-4}$)	1.73 (1.45)	0.30 (0.38)
two crews, common cause	7.55×10^{-5}	9.9×10^{-5} (8.3×10^{-5})	7.6×10^{-5} (6.5×10^{-5})	$[0.16,\ 3.5] \times 10^{-4}$ ($[0.096,\ 2.6] \times 10^{-4}$)	0.86 (0.73)	0.64 (0.71)

unavailabilities when carried to U. By the draws the expectation is 1.66×10^{-4}, 1.45 h a year, with a 95 % band from 1.9×10^{-5} to 5.2×10^{-4}, a factor of 27 between its ends; it is 10 % above the plug-in value 1.51×10^{-4} and above the median 1.29×10^{-4}, the posterior being skewed toward large unavailabilities. The state route is higher by 0.31×10^{-4}, 0.28 h a year, and its band by a factor of 1.7 at the lower end and 1.3 at the upper: the normal on z_3 is centered at $\ln(2/18)$, where the exact logit has its mean lower by 0.242 and a longer tail toward small shares.

The parts. Without the common cause the pair's expected unavailability is 6.1×10^{-6}, so the common cause accounts for 97 % of the expected time out, 1.67 of the 1.73 h a year; by the draws 6.2×10^{-6}, 96 %, 1.40 of the 1.45 h. The wait for the single crew adds $U_{c=1} - q^2 = 4.33 \times 10^{-6} - 2.16 \times 10^{-6} = 2.2 \times 10^{-6}$ at the MAP state and 3.1×10^{-6} in expectation on both routes: one to two percent of the pair's unavailability.

The time to the first outage and the outage frequency.

The mean time to the first outage. At the MAP state equation (16.25) gives, with the numerator $2.8 \times 1.30952 \times 10^{-4} + 0.088889 = 0.089256$ and the denominator $1.30952 \times 10^{-4} \times (1.9 \times 1.30952 \times 10^{-4} + 0.1 \times 0.088889) = 1.19660 \times 10^{-6}$,

$$\mathrm{MTTF}_{\mathrm{pair}} = \frac{0.089256}{1.19660 \times 10^{-6}}\ \mathrm{h} = 74\,591\ \mathrm{h} = 8.5\ \mathrm{yr},$$

against $(3\lambda + \mu)/(2\lambda^2) = 2.603 \times 10^6$ h = 297 yr without the common cause: the common cause shortens the time to the first outage thirty-fivefold. Under the posterior of the state, the plug-in

value is the median; the sum (16.38) gives the mean 11.2 yr, and the delta method, with $\mathbf{g}_{\mathrm{MTTF}}$ below, $\sigma = 0.721$, gives the central band $8.5\,\mathrm{yr} \times \mathrm{e}^{\pm 1.96 \times 0.721} = [2.1,\ 35]\,\mathrm{yr}$ and, with Φ the standard normal cumulative distribution,

$$\Pr(\mathrm{MTTF}_{\mathrm{pair}} > 5\,\mathrm{yr} \mid \mathbf{y}) \approx \Phi\left(\frac{\ln 8.515 - \ln 5}{0.721}\right) = \Phi(0.738) = 0.77.$$

The value beside it comes from the exact posterior. The draws give the median 10.3 yr, the mean 16.0 yr and the central band [3.0, 64] yr, and

$$\Pr(\mathrm{MTTF}_{\mathrm{pair}} > 5\,\mathrm{yr} \mid \mathbf{y}) = 0.85.$$

The plug-in value lies below the exact median because the MTTF is nearly $1/(\beta_{\mathrm{cc}}\lambda)$, a decreasing function of the share, whose median 0.087 is below the mean 0.1 at which the plug-in value is taken; convexity in the share puts the mean above both. The state route is lower by 1.8 yr at the median, by 4.8 yr at the mean and by 0.08 on the probability, because its normal on z_3 gives less weight to small shares.

The gradient of the time to the first outage. The system (16.33) at the MAP state has the matrix, the mean times in hours and the dimensionless right-hand sides

$$-\mathbf{Q}_{TT} = \begin{bmatrix} 2.4881 \times 10^{-4} & -2.3571 \times 10^{-4} \\ -0.088889 & 0.089020 \end{bmatrix} \mathrm{h}^{-1}, \qquad \mathbf{m} = \begin{bmatrix} 74\,591 \\ 74\,492 \end{bmatrix} \mathrm{h},$$

$$\mathbf{C} = \begin{bmatrix} -1.0000 & 0 & -0.8768 \\ -9.7549 & 8.7549 & 0 \end{bmatrix},$$

and its solution and the gradient of $\ln \mathrm{MTTF}_{\mathrm{pair}}$ are

$$\mathbf{M}_{\mathbf{z}} = \begin{bmatrix} -76\,315 & 1725 & -65\,227 \\ -76\,313 & 1820 & -65\,131 \end{bmatrix} \mathrm{h}, \qquad \mathbf{g}_{\mathrm{MTTF}} = \frac{1}{m_2}\,\mathbf{M}_{\mathbf{z}}^{\mathsf{T}} \begin{bmatrix} 1 \\ 0 \end{bmatrix} = \begin{bmatrix} -1.023 \\ 0.023 \\ -0.874 \end{bmatrix}. \tag{16.44}$$

The repair rate's entry is about forty times smaller than the other two: the first outage of the pair is almost always a common-cause event, which no repair prevents, while the failure rate and the share enter with the weight of the unavailability's gradient and the opposite sign. With $\mathbf{\Sigma}$ of equation (16.7), the share carries 82 % of $\mathbf{g}_{\mathrm{MTTF}}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}_{\mathrm{MTTF}}$ and the failure rate the rest.

The outage frequency and the mean down time. By equation (16.21) at the MAP state, $\nu_{\mathrm{out}} = 0.997058 \times 1.30952 \times 10^{-5} + 2.79087 \times 10^{-3} \times 1.30952 \times 10^{-4} = 1.3057 \times 10^{-5} + 0.0365 \times 10^{-5} = 1.3422 \times 10^{-5}\,\mathrm{h}^{-1}$, 0.118 outages a year, of which 0.114, or 97 %, are common-cause events; the sum (16.38) gives the expected frequency 0.147 a year. The mean down time with one crew is $1/\mu = 11.25\,\mathrm{h}$ at the MAP state, 11.73 h by the sum, and 12.27 h on the exact posterior. The reciprocal of the frequency, 8.5 yr, is the mean time between outages, which here nearly equals the mean time to the first one.

Table 16.9 collects the answers to the third question. Each row is a quantity; the columns give its plug-in value at the MAP state with one crew, its posterior summary from the posterior of the state with the draws in parentheses, and its value with two crews.

16.3.4 The requirement on the unavailability

Engineering question. Is the pair's unavailability below 10^{-4}, with assurance 0.90? The service agreement limits the long-run loss of cooling supply to $u^* = 10^{-4}$ of the time and asks that the limit hold with the posterior probability $\kappa = 0.90$.

Table 16.9: The time to the first outage, the outage frequency and the mean down time of the pump pair. Plug-in at the MAP state; posterior: the sum (16.38) and the delta method, with in parentheses the 2×10^4 draws of Table 16.8.

Quantity	Plug-in, one crew	Posterior, one crew	Two crews
$\mathrm{MTTF}_{\mathrm{pair}}$, common cause	8.5 yr	median 8.5 yr (10.3 yr), mean 11.2 yr (16.0 yr), band [2.1, 35] yr ([3.0, 64] yr)	unchanged
$\mathrm{MTTF}_{\mathrm{pair}}$, no common cause	297 yr	median 297 yr (306 yr), band [80, 1100] yr ([88, 1270] yr)	unchanged
$\Pr(\mathrm{MTTF}_{\mathrm{pair}} > 5\,\mathrm{yr} \mid \mathbf{y})$	–	0.77 (0.85)	unchanged
Outage frequency ν_{out}	0.118 per year	mean 0.147 (0.118), median 0.118 (0.097), band [0.029, 0.48] ([0.016, 0.34]) per year	0.118 per year
Common-cause part of ν_{out}	97 %	–	97 %
Mean down time	11.25 h	mean 11.73 h (12.27 h)	5.6 h
$\nu_{\mathrm{out}} \times \mathrm{MDT}$	1.510×10^{-4}	equals U in every draw	7.55×10^{-5}

General form. A prediction under a threshold, equation (2.12):

$$\text{accept if } \Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}) = \int \mathbf{1}\big[f(\mathbf{z}) \geq f^*\big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z} \geq \kappa. \tag{16.45}$$

Specialized form. One choice enters here: the limit is read on the long-run unavailability, the fraction of time out over many years, as the agreement states it, not on the downtime of one year, which stays random even when the rates are known and would make the question a predictive one. The question is equation (16.45) with $f(\mathbf{z}) = -\ln U(\mathbf{z})$, the limit being an upper one, $f^* = -\ln u^*$, the event $U(\mathbf{z}) < u^*$, and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of step 3. Under the delta method of equation (16.39) it is a normal tail,

$$\Pr(U < u^* \mid \mathbf{y}) \approx \Phi\!\left(\frac{\ln u^* - \ln U(\hat{\mathbf{z}})}{\sigma_{\ln U}}\right), \tag{16.46}$$

with one crew and with two. The value beside it comes from the exact posterior, the integral

$$\Pr(U < u^* \mid \mathbf{y}) = \int \mathbf{1}\big[\pi_0(\mathbf{z}) < u^*\big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z},$$

the fraction of the draws below the limit, and on the exact moments the same tail of the normal (16.41).

Building blocks. $U(\hat{\mathbf{z}})$, $\mathbf{g}$ and $\sigma_{\ln U}$ of the previous subsection and the normal distribution Φ; and, for the exact values beside them, the draws and the stationary solve at each, for one crew and for two.

The requirement of 10^{-4}. By equation (16.46), $\Phi((-9.210 + 8.798)/0.779) = \Phi(-0.529) = 0.30$ with one crew and, with the center 7.55×10^{-5}, $\Phi(0.361) = 0.64$ with two. The fraction of draws with the unavailability below 10^{-4} is 0.38 with one crew and 0.71 with two, the last column of Table 16.8, 0.08 and 0.07 higher. Against the assurance $\kappa = 0.90$ the limit is demonstrated by neither, on either route; the plug-in value with one crew, 1.51×10^{-4}, is above the limit, and still the posterior of the state puts 0.30 of its probability below it, and the draws 0.38. With two crews the state route's expectation, 9.9×10^{-5}, and the exact one, 8.3×10^{-5}, are both just below it.

16.3.5 The recovery after a double outage

Engineering question. After both pumps are down, how soon is the pair back, and how likely is it to be whole again within a day? The operator needs the time course of a double outage, the probability of each state at each hour after it, and the probability that both pumps run again within 24 h.

General form. The predictive of a future event, equation (2.13) with y_{new} its indicator: the probability of the event given the state, averaged over the posterior,

$$\Pr(A \mid \mathbf{y}) = \int \Pr(A \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}. \tag{16.47}$$

Specialized form. One choice enters here: the recovery is taken to start at the double outage with the crew at work, $\mathbf{p}(0) = (0, 0, 1)$; this makes the answer a row of a matrix exponential, and a delay before the crew starts would shift every time by that delay.

The question is equation (16.47) with the event A below and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of step 3. Given the state, the probabilities of the three states after the outage are the last row of $\mathrm{e}^{\mathbf{Q}t}$, $\mathbf{p}(t) = \mathbf{p}(0)\,\mathrm{e}^{\mathbf{Q}t}$, approaching the stationary distribution through the time constants of the generator (Section 16.3.1). The event of the question is the first passage, $A = \{T_{0\to 2} \le 24\,\mathrm{h}\}$, with

$$F_{0\to 2}(t; \mathbf{z}) = \Pr(T_{0\to 2} \le t \mid X(0) = 0, \mathbf{z}) = \left[\mathrm{e}^{\mathbf{Q}' t}\right]_{0,2},$$

the entry $(0, 2)$ of the exponential of the generator $\mathbf{Q}'$ whose row of state two is set to zero, so that the chain stops on first reaching both up; it differs from $p_2(t)$ by the paths that return to both up and fail again before t. Under the posterior of the state,

$$\Pr(T_{0\to 2} \le 24\,\mathrm{h} \mid X(0) = 0, \mathbf{y}) = \mathbb{E}\big[F_{0\to 2}(24\,\mathrm{h}; \mathbf{z}) \mid \mathbf{y}\big],$$

the sum (16.38) with a matrix exponential at each node, over the two log rates only, since $F_{0\to 2}$ does not depend on the share. By the delta method the gradient of a transient probability with respect to the state is a row of the sensitivity matrix $\mathbf{S}(t) = \partial \mathbf{P}(t)/\partial \mathbf{z}^{\mathsf{T}}$, which obeys the sensitivity equation of the chain (Section 16.3.2); a probability is bounded, so its band is taken on the logit, $\operatorname{logit} p \pm 1.96\,(\mathbf{g}_p^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g}_p)^{1/2}/(p(1-p))$ mapped back. The value beside it comes from the same expectation by the draws, a matrix exponential at each, and the exact-mean center moves by $\mathbf{g}_p^{\mathsf{T}} \boldsymbol{\delta}$ as for $\ln U$.

Building blocks. The MAP state $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$; the generator at the MAP state and at each node, the matrix exponentials of $\mathbf{Q}$ and $\mathbf{Q}'$ and the eigenvalues of $\mathbf{Q}$ (Section 16.3.1); the sensitivity matrix by the block exponential (Section 16.3.2); and, for the exact values beside them, the draws and $\boldsymbol{\delta}$.

Recovery from a double outage. At the MAP state with one crew the eigenvalues of $\mathbf{Q}$ are 0, -0.08584 and $-0.09232\,\mathrm{h}^{-1}$: their sum is the trace, $-0.17816\,\mathrm{h}^{-1}$, and their product the sum of the principal minors (16.27), $7.9245 \times 10^{-3}\,\mathrm{h}^{-2}$. The time constants are 11.6 h and 10.8 h, on either side of the 11.25 h of one repair at the posterior mean rate; with two crews they are 11.2 h and 5.6 h, near $1/\mu$ and $1/(2\mu)$.

Table 16.10 gives the transient. Each row is a time after the double outage; the columns give the probabilities of two, one and no pumps up with one crew, the probability of both up with two crews, and, for the pair started with both pumps up, the probability of none up as a fraction of its stationary value.

Table 16.10: The transient of the pump pair at the MAP state. From none up: the last row of $e^{\mathbf{Q}t}$; from both up: its entry $(2,0)$ divided by $\pi_0 = 1.510 \times 10^{-4}$, one crew.

Time t (h)	$p_2(t)$, one crew	$p_1(t)$, one crew	$p_0(t)$, one crew	$p_2(t)$, two crews	From both up, $p_0(t)/\pi_0$
6	0.100	0.313	0.587	0.171	0.41
12	0.288	0.367	0.344	0.430	0.65
24	0.628	0.254	0.119	0.776	0.88
48	0.924	0.062	0.014	0.970	0.98
96	0.995	0.004	0.0004	0.997	1.00

Figure 16.2 draws the same functions. In panel (a) the horizontal axis is the time since the double outage and the vertical axis the probability of each state with one crew: the solid green curve is both up, the solid gold curve one up and the solid red curve none up, and the dashed lines of the same colors are their stationary values, 0.9971, 0.0028 and 0.00015, the last two on the axis at this scale. In panel (b) the pair starts with both pumps up; the vertical axis is its probability of being out, in units of 10^{-4}, the dashed line its stationary value and the dotted line 90 % of it.

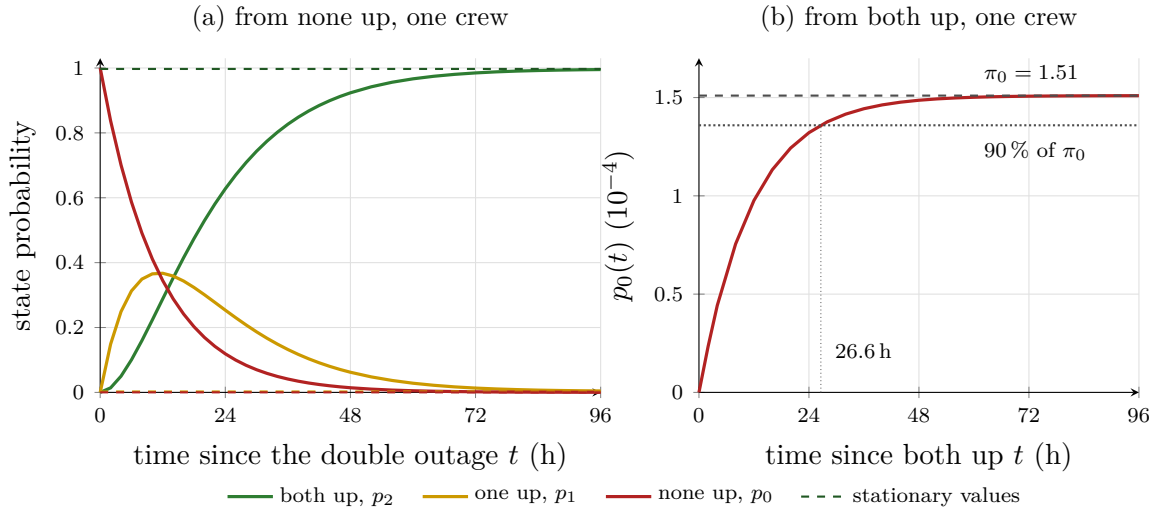


Figure 16.2: The transient of the pump pair at the MAP state, one crew, $\mathbf{p}(t) = \mathbf{p}(0)\, e^{\mathbf{Q}t}$ with $\mathbf{Q}$ of Section 16.3.3; eigenvalues -0.08584 and $-0.09232\,\mathrm{h}^{-1}$, time constants 11.6 and 10.8 h. (a) From none up, the three state probabilities and their stationary values (dashed). (b) From both up, the probability of none up in units of 10^{-4}, its stationary value 1.510 (dashed) and 90 % of it (dotted), reached at 26.6 h.

The recovery at the MAP state. After a double outage with one crew the pair has both pumps up again with probability 0.29 at 12 h, 0.63 at 24 h and 0.92 at 48 h; the median time to both up is 18.9 h, and the mean time, the first-passage mean to both up, is 22.5 h, two repairs in sequence. In panel (a) the probability of one up rises first, to 0.37 at 12 h, as the crew finishes the first pump, and then falls as it finishes the second. The first passage of the question is $F_{0\to 2}(24\,\mathrm{h}; \hat{\mathbf{z}}) = 0.629$, against $p_2(24\,\mathrm{h}) = 0.628$. With two crews the pumps are repaired at once and p_2 at a day rises to 0.78, at two days to 0.97.

The recovery under the posterior. By the sum (16.38) with a matrix exponential at each node,

$$\Pr(T_{0\to 2} \le 24\,\mathrm{h} \mid X(0)=0, \mathbf{y}) = \mathbb{E}\big[F_{0\to 2}(24\,\mathrm{h}; \mathbf{z}) \mid \mathbf{y}\big] \approx 0.625,$$

five nodes giving the same, with the median at the MAP state 0.629 and, with the gradient $\mathbf{g}_F$ below and the standard deviation 0.156, the band on the logit [0.314, 0.862]. The value beside it comes from the exact posterior. By the 2×10^4 draws, a matrix exponential at each,

$$\Pr(T_{0\to 2} \le 24\,\mathrm{h} \mid X(0)=0, \mathbf{y}) = \mathbb{E}\big[F_{0\to 2}(24\,\mathrm{h}; \mathbf{z}) \mid \mathbf{y}\big] = 0.60,$$

with the central 95 % band [0.30, 0.86] on $F_{0\to 2}(24\,\mathrm{h}; \mathbf{z})$ and the median 0.61; $p_2(24\,\mathrm{h})$ gives the same to two digits. The expected probability of being whole again within a day is 0.625 on the state route and 0.60 on the exact one, a gap of 0.025, both below the MAP value 0.63.

The gradient of the recovery. The sensitivity equation (16.34), solved at the MAP state with one crew from none up by the block exponential, gives at one day

$$\mathbf{S}(24\,\mathrm{h}) = \begin{bmatrix} -1.310 \times 10^{-3} & 0.5384 & 2.5 \times 10^{-5} \\ 8.59 \times 10^{-4} & -0.2853 & -7.3 \times 10^{-5} \\ 4.51 \times 10^{-4} & -0.2531 & 4.8 \times 10^{-5} \end{bmatrix},$$

rows by state (two, one, none up) and columns by entry. The repair rate's entry of the gradient of $p_2(24\,\mathrm{h})$, the first row, is 0.538, four hundred times the failure rate's, and the share's is zero to four decimals. For the first passage the stopped chain gives $\mathbf{g}_F = (-0.0003,\ 0.539,\ 0)^{\mathsf{T}}$, the share's entry exactly zero because the share acts only on the row of both up, which $\mathbf{Q}'$ sets to zero. Up to two days the repair rate's entry carries more than 99.98 % of $\mathbf{g}_{p_2}^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g}_{p_2}$, and 97 % at 96 h.

The approach from both up. Started with both pumps up, the pair's probability of being out rises from zero to its stationary value with the same time constants, reaching 88 % of it in a day and 90 % at 26.6 h, panel (b). Any question about a horizon of more than two days may use the stationary distribution; a question about the first day after both pumps return may not.

16.3.6 The second crew and the halved share

Engineering question. Which helps more: a second crew, or halving the common-cause share? A second crew halves the mean down time; staggered maintenance, assumed as in Chapter 15 to halve the share, halves the rate of common-cause events.

General form. The design form, equation (2.16):

$$U(a) = \mathbb{E}_{\tilde{y}_a \mid \mathbf{y}}\Big[\, u\big(p(\mathbf{z} \mid \mathbf{y}, \tilde{y}_a)\big)\Big]. \tag{16.48}$$

Specialized form. Two choices enter here. Each pump, including each of the two pumps of a common-cause outage, is taken as a separate repair job: this is what lets a second crew repair the two pumps of such an outage in parallel and end it at the rate 2μ, and it fails if the shared cause leaves damage that is repaired as one job, as Chapter 15 assumed; Section 16.4.5 gives the answers under that alternative. Staggered maintenance is taken to halve the share: the intervention's effect is an engineering assumption that the record does not test (Section 16.6.5).

The question is equation (16.48) with the actions a the second crew and the halved share. An intervention adds no measurement. It changes the drawing, and the outcome $\tilde{y}_a$ over which the design form averages is absent, so the form reduces to a posterior expectation under the

changed drawing, over the same posterior of the state $\mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$. The quantity u valued in it is here the unavailability itself, and $U(a)$, the design form's value of an action, is the expected unavailability under it, $\mathbb{E}[\pi_0(\mathbf{z}; a) \mid \mathbf{y}]$; it is not to be confused with the unavailability U of one state, nor u with the limit u^*. The second crew sets the rate out of none up to 2μ at every node of the sum (16.38), and beside it at every draw; the halved share replaces β_{cc} by $\beta_{\mathrm{cc}}/2$ in the same way. The outage frequency, the threshold probability and the mean time to the first outage are taken under each action in the same way, the threshold by the delta method with the gradient at the action. To first order each intervention halves the leading term $\beta_{\mathrm{cc}}\epsilon/c$, and the two differ in what they do: by the identity $U = \nu_{\mathrm{out}} \times \mathrm{MDT}$ (Section 16.3.1) the second crew halves the mean down time and leaves the number of outages and the time to the first, while the halved share halves the number of outages and nearly doubles the time to the first. The gradient reads the halved share off the state without a new solve: the action moves the third entry by $\Delta\xi = \ln(\beta_{\mathrm{cc}}/2) - \ln(1 - \beta_{\mathrm{cc}}/2) - \hat{z}_3$, and $\ln U$ moves by about $g_3\,\Delta\xi$. Where each outage carries a cost of its own, the two are priced by their expected losses, Exercise 16.6.

Building blocks. The MAP state $\hat{\mathbf{z}}$ and $\mathbf{\Sigma}$, the sum (16.38) and the delta method with $c = 2$ and with the share halved, the stationary solve and the first-passage system at each node; g_3 and $\hat{z}_3$; and, for the exact values beside them, the draws under the same actions.

The second crew and the halved share. Table 16.11 sets out the two interventions under the posterior, each cell the state route with the draws in parentheses. Each row is a course of action; the columns give the expected unavailability, its band, the threshold probability, the expected outage frequency, the expected hours out per year and the plug-in MTTF.

Table 16.11: The design options under the posterior: the sum (16.38) and the delta method at each action, with in parentheses the 2×10^4 draws of Table 16.8, with every node's and every draw's share halved for the staggered maintenance.

Course of action	$\mathbb{E}[U \mid \mathbf{y}]$	95 % band	$\Pr(U < 10^{-4})$	$\mathbb{E}[\nu_{\mathrm{out}}]$ per year	Hours per year	Pair MTTF
one crew	1.97×10^{-4}	$[0.33,\ 7.0] \times 10^{-4}$	0.30	0.147	1.73	8.5 yr
(present)	(1.66×10^{-4})	$([0.19,\ 5.2] \times 10^{-4})$	(0.38)	(0.118)	(1.45)	
second crew	9.9×10^{-5}	$[0.16,\ 3.5] \times 10^{-4}$	0.64	0.147	0.86	8.5 yr
	(8.3×10^{-5})	$([0.096,\ 2.6] \times 10^{-4})$	(0.71)	(0.118)	(0.73)	
share	1.02×10^{-4}	$[0.17,\ 3.5] \times 10^{-4}$	0.63	0.076	0.89	16.6 yr
halved	(8.6×10^{-5})	$([0.11,\ 2.7] \times 10^{-4})$	(0.70)	(0.061)	(0.75)	
both	5.1×10^{-5}	$[0.086,\ 1.8] \times 10^{-4}$	0.89	0.076	0.45	16.6 yr
	(4.3×10^{-5})	$([0.055,\ 1.4] \times 10^{-4})$	(0.93)	(0.061)	(0.38)	

A second crew and a halved share are worth the same for the unavailability, 9.9 and 10.2×10^{-5} (8.3 and 8.6×10^{-5} by the draws). The two differ in what they do: the second crew leaves the number of outages, 0.147 (0.118) a year, and the time to the first, 8.5 yr, unchanged; the halved share halves the number of outages and nearly doubles the time to the first, to 16.6 yr. Together they reach, by the delta method at $U(\hat{\mathbf{z}}) = 3.88 \times 10^{-5}$ and $\sigma_{\ln U} = 0.771$,

$$\Pr(U < 10^{-4} \mid \mathbf{y}) \approx \Phi\left(\frac{-9.210 + 10.156}{0.771}\right) = \Phi(1.227) = 0.89,$$

and by the draws 0.93. The two values fall on opposite sides of the assurance of 0.90: the state route leaves it unmet by 0.01, the exact draws meet it by 0.03. The gap is the skew of the logit of the share, informed by one double: the normal on z_3 keeps too little of the exact posterior's weight on small shares and too much on large ones. The decision whether both interventions together demonstrate the limit is therefore one the two routes do not make alike.

Figure 16.3 draws the two halves of the design question. Panel (a) is the variance budget of $\ln U$, equation (16.39): one bar per entry, its height the entry's share of the variance in percent. Panel (b) is the expected unavailability under each course of action of Table 16.11, on a logarithmic scale: the dot is the posterior mean, the vertical line the central 95 % band, blue for the state route and green, to its left, for the draws, and the dotted red line the limit 10^{-4}. Each intervention moves the whole band down by a factor of two, while the share carries most of the band's width.

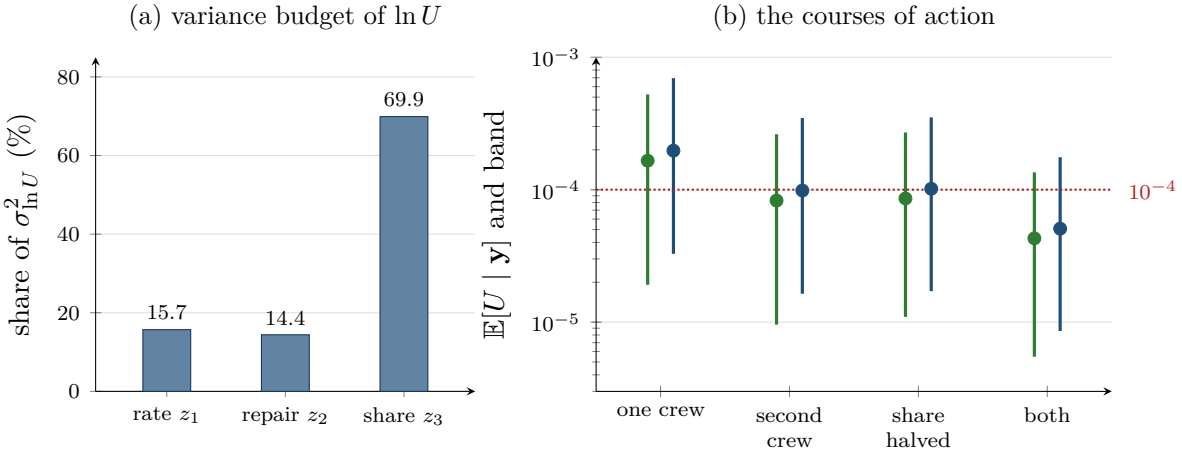


Figure 16.3: The design question. (a) The variance budget of $\ln U$ with one crew, the terms $g_i^2 \Sigma_{ii}$ of equation (16.39) as percentages of their sum. (b) The expected unavailability (dots) and the central 95 % band (lines) under the four courses of action of Table 16.11, blue from the sum and the delta method over $\mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$, green from 2×10^4 draws, with the limit 10^{-4} (dotted).

16.4 The checks (step 5)

The pair has no physics row and no exact constraint attached as evidence, so no limit $\sigma_p \to 0$ of a model tolerance arises: the stationary equations and the forward equations are solved exactly as the definitions of the derived quantities, and what can be checked is what those equations imply, the routes beside the reported one, and the assumptions the chain adds.

16.4.1 The identities of the chain

The stationary solve. One row of the arithmetic of the closed form (16.18), for one crew: $\pi_1/\pi_2 = 1.9{\times}1.47321{\times}10^{-3} = 2.7991{\times}10^{-3}$; $\pi_0/\pi_2 = 0.1{\times}1.47321{\times}10^{-3}{+}1.9{\times}2.17035{\times}10^{-6} = 1.47321 \times 10^{-4} + 4.124 \times 10^{-6} = 1.51445 \times 10^{-4}$; the sum $1 + 2.7991 \times 10^{-3} + 1.51445 \times 10^{-4} = 1.0029506$; and $U = 1.51445 \times 10^{-4}/1.0029506 = 1.5100 \times 10^{-4}$, the value the stationary solve returns to every digit. The expansion (16.20) is within 0.01 % of the exact value with the common cause, for one crew and for two, and within 0.3 % without it, where the leading term is itself of second order and the first neglected term is 2ϵ of it (Table 16.5); the first-order approximation of the MTTF (16.25) is within 0.01 % of its exact form. The balance across none up, equation (16.23), returns $1.3422 \times 10^{-5} \times 11.25 = 1.5100 \times 10^{-4}$, the stationary unavailability to every digit, and it holds in every draw.

The gradients. Each column of $\mathbf{\Pi_z}$ in equation (16.42) sums to zero, the first two are opposite as the null vector (16.32) requires, and central differences of the stationary solve return the same $\mathbf{g}$ to every digit shown. As a scalar remark, differentiating the expansion (16.20) by hand

gives $g_1 \approx 1.024$ and $g_3 \approx 0.874$, within 0.0002 of the entries the solve returns. The reading of the halved share off the gradient is close as well: $\Delta\xi = \ln(0.05/0.95) - \ln(0.1/0.9) = -0.747$, and $g_3\Delta\xi = 0.874 \times (-0.747) = -0.653$ is a factor of 0.52 on U, against the factor 0.514 of the stationary solve at the MAP state with the share halved. The first-passage gradient (16.44) has first two entries summing to -1, the check $\mathbf{M_z n} = -\mathbf{m}$. For the transient, central differences of the matrix exponential agree with $\mathbf{S}(24\,\mathrm{h})$ to 3×10^{-10}, each column sums to zero, and $\mathbf{S}(24\,\mathrm{h})\,\mathbf{n}$ equals $24\,\mathrm{h} \times \mathbf{Q}^{\mathsf{T}}\mathbf{P}(24\,\mathrm{h})$ to round-off, the relations (16.35).

The expectation. The leading terms (16.40) give, with $\mathbb{E}[K(\beta_{\mathrm{cc}})] = 2 - 3 \times 0.1 = 1.7$ for one crew, $0.1 \times 1.30952 \times 10^{-4} \times 12.273 + 1.7 \times 1.8707 \times 10^{-8} \times 165.68 = 1.607 \times 10^{-4} + 0.053 \times 10^{-4} = 1.660 \times 10^{-4}$, the value of the draws to the digits given. The posterior mean repair duration, 12.27 h against the reciprocal 11.25 h of the mean rate, a factor of 1.091, is most of the 10 % gap between the expectation and the plug-in value; without the common cause the gap is the factor $(a_\lambda + 1)/a_\lambda \times \mathbb{E}[\mu^{-2}]/\mathbb{E}[\mu]^{-2} = 1.091 \times 1.309 = 1.43$, which is the ratio of 6.2×10^{-6} to 4.33×10^{-6} in Table 16.8.

16.4.2 The delta method against the draws

The sampling error of the draws. With 2×10^4 draws the standard error of $\mathbb{E}[U \mid \mathbf{y}]$ is 0.6 % of it and that of a threshold probability near 0.38 is 0.003. The interventions are compared on the same draws, so their difference has a much smaller error: the second crew's expected unavailability lies below the halved share's by 3.07×10^{-6}, with a standard error of 2.3×10^{-8}, and the gap between 8.3 and 8.6×10^{-5} is real, not sampling noise.

The unavailability. The band and the threshold probabilities of Sections 16.3.3 and 16.3.4 come from the delta method, equation (16.43), and this paragraph sets them against the draws.

Figure 16.4 draws the two routes for one crew and for two. In each panel the horizontal axis is the unavailability on a logarithmic scale and the vertical axis the density of its base-ten logarithm. The solid green curve is the histogram of the draws in bins of 0.1, with the area below the limit 10^{-4} shaded; the dashed blue curve is the normal of the delta method on the same scale; the vertical red line is the limit. The two panels have the same shape, shifted by $\log_{10} 2 = 0.30$: the second crew divides the unavailability by two in every draw, to first order, and moves a third of the posterior's mass across the limit. In both, the normal is narrower than the histogram, has a thinner lower tail and sits to its right.

Table 16.12 sets the routes side by side for one crew. Each row is a summary; the columns give it from the draws, from the delta method, and from the normal on $\ln U$ with the same gradient but centered at the exact means of the entries (16.6) and with their exact variances, equation (16.41).

Table 16.12: The unavailability with one crew by three routes. Draws: the 2×10^4 draws of Table 16.8. Delta method: $\ln U \sim \mathcal{N}(\ln 1.510 \times 10^{-4},\ 0.779^2)$. Exact moments: the same gradient, the entries centered at $\psi(a) - \ln b$ and $\psi(a_\beta) - \psi(b_\beta)$ with variances ψ', which gives $\ln U \sim \mathcal{N}(\ln 1.217 \times 10^{-4},\ 0.853^2)$.

Summary	Draws	Delta method	Exact moments
Median of U	1.29×10^{-4}	1.51×10^{-4}	1.22×10^{-4}
Std. dev. of $\ln U$	0.846	0.779	0.853
95 % band	$[0.19,\ 5.2] \times 10^{-4}$	$[0.33,\ 7.0] \times 10^{-4}$	$[0.23,\ 6.5] \times 10^{-4}$
$\Pr(U < 10^{-4} \mid \mathbf{y})$	0.38	0.30	0.41

The delta band is shifted upward against the draws, by a factor of about 1.2 at its center and 1.7 at its lower end, and it is narrower. Both effects come from the share's logit. The offset

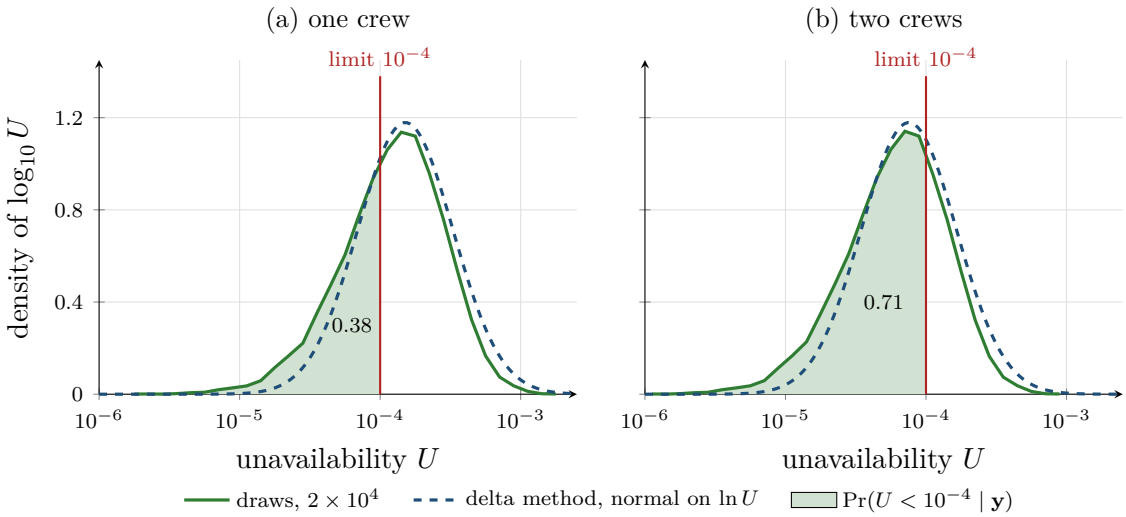


Figure 16.4: The posterior of the pair's unavailability with one crew (a) and two crews (b). Draws: histogram of $\log_{10} U$ over the 2×10^4 draws of Table 16.8, bins of 0.1, with the area below 10^{-4} shaded (0.38 and 0.71). Delta method: $\ln U \sim \mathcal{N}(\ln U(\hat{\mathbf{z}}),\ 0.779^2)$ with $U(\hat{\mathbf{z}}) = 1.510 \times 10^{-4}$ and 7.55×10^{-5}, drawn on $\log_{10} U$ with means -3.821 and -4.122 and standard deviation 0.338; it puts 0.30 and 0.64 below the limit.

vector and the exact covariance of equation (16.6) are

$$\boldsymbol{\delta} = \begin{bmatrix} -0.046 \\ -0.042 \\ -0.242 \end{bmatrix}, \qquad \boldsymbol{\Sigma}_{\text{ex}} = \operatorname{diag}(0.0952,\ 0.0869,\ 0.7021),$$

the third offset being the gap between the Laplace center $\ln(2/18) = -2.197$ of the logit and its exact mean $\psi(2) - \psi(18) = -2.440$. The center of $\ln U$ moves by

$$\mathbf{g}^{\mathsf{T}}\boldsymbol{\delta} = 1.024 \times (-0.046) - 1.024 \times (-0.042) + 0.874 \times (-0.242) = -0.047 + 0.043 - 0.212 = -0.216,$$

a factor of 0.81, from 1.51 to 1.22×10^{-4}: the two log rates' products nearly cancel, as the logit of Chapter 8 cancelled them, and the share's product is the whole shift. The variance $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_{\text{ex}}\mathbf{g} = 0.728$ exceeds the Laplace 0.607 almost entirely through its third term, $0.874^2 \times 0.7021 = 0.537$ against $0.874^2 \times 0.5556 = 0.425$, and widens the normal to 0.853, close to the draws' 0.846. Centered and scaled by the exact moments, the normal comes close to the draws' median, 1.22 against 1.29×10^{-4}, matches their spread and comes within 0.03 of their threshold probability; what remains is the skew of the logit of a Beta variable with a first parameter of two, the same skew that opens the gap on the delta method in Section 15.4.3. The state route stays the route of the reported answers, of the gradient and of the budget, with the draws beside it and the gap stated.

Common mistake

Reporting the threshold probability from the delta method with no exact value beside it, because the logarithm of the unavailability is nearly linear in the state. The linearization is good, the gradient matches central differences of the stationary solve; the Gaussian beneath it has a gap, because the share's logit, which carries 70 % of the budget, has a posterior skewed away from its mode. The delta method gives 0.30 for $\Pr(U < 10^{-4})$ against 0.38 from the draws, and for both interventions 0.89 against 0.93, on opposite sides of the assurance.

The time to the first outage. With the gradient (16.44), $\mathbf{g}_{\mathrm{MTTF}}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_{\mathrm{MTTF}} = 0.0952 + 0.0000 + 0.4248 = 0.520$, a standard deviation of 0.721 on $\ln \mathrm{MTTF}_{\mathrm{pair}}$. Figure 16.5 draws the posterior of the pair's MTTF by both routes. The horizontal axis is the MTTF in years on a logarithmic scale and the vertical axis the density of its base-ten logarithm. The solid green curve is the histogram of the draws in bins of 0.1, with the area above 5 yr shaded; the dashed blue curve is the delta method on $\ln \mathrm{MTTF}_{\mathrm{pair}}$, centered at the plug-in value with the gradient (16.44) and standard deviation 0.721; the red line is 5 yr. The posterior spans more than a decade, skewed toward long lives; the delta normal again sits on the low side of the histogram, at the plug-in value, and puts 0.77 above 5 yr against the draws' 0.85, for the same reason: the share's logit carries 82 % of the variance.

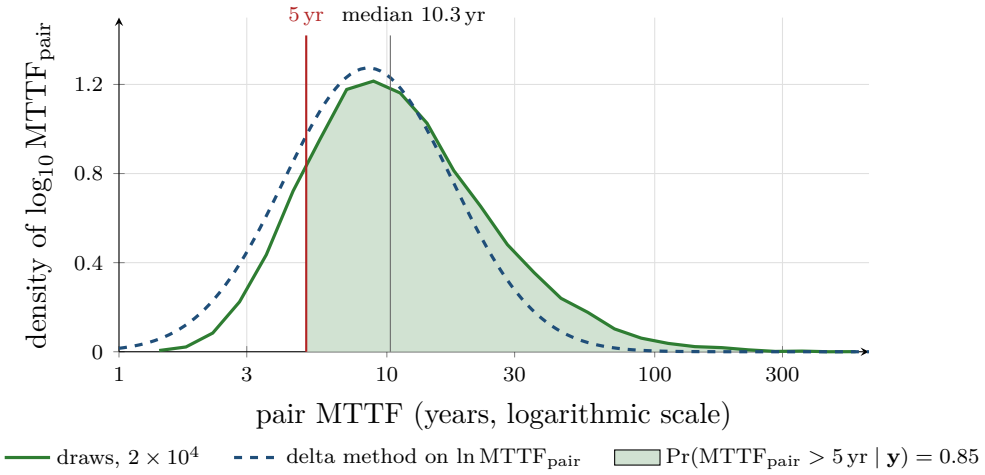


Figure 16.5: The posterior of the pair's mean time to the first outage, one crew. Draws: histogram of $\log_{10} \mathrm{MTTF}_{\mathrm{pair}}$ in years over the 2×10^4 draws, equation (16.25) at each, bins of 0.1; the area above 5 yr is 0.85. Delta method: $\ln \mathrm{MTTF}_{\mathrm{pair}} \sim \mathcal{N}(\ln 8.515\,\mathrm{yr},\ 0.721^2)$, drawn on the base-ten logarithm with mean 0.930 and standard deviation 0.313, which puts 0.77 above 5 yr.

The recovery. Table 16.13 gives the band of the probability of both up after a double outage by both routes. Each row is a time after the outage; the columns give $p_2(t)$ at the MAP state, its gradient $\mathbf{g}_{p_2} = \mathbf{S}(t)^{\mathsf{T}}(1, 0, 0)^{\mathsf{T}}$, its standard deviation $(\mathbf{g}_{p_2}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_{p_2})^{1/2}$, the central 95 % band of the delta method taken on the logit of $p_2(t)$, and the band of the 2×10^4 draws.

Table 16.13: The band of the probability of both pumps up after a double outage, one crew. Gradient: the first row of $\mathbf{S}(t)$ from equation (16.34) at the MAP state, entries in the order z_1, z_2, z_3. Delta band: logit $p_2 \pm 1.96\,(\mathbf{g}_{p_2}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_{p_2})^{1/2}/(p_2(1-p_2))$ mapped back. Draws: the 2×10^4 draws of Table 16.8.

Time t [h]	$p_2(t)$ [–]	$\mathbf{g}_{p_2}^{\mathsf{T}}$ [–]	Std. dev. [–]	Delta band [–]	Draws band [–]
6	0.100	(−0.0001, 0.167, 0.0000)	0.048	[0.038, 0.241]	[0.031, 0.218]
12	0.288	(−0.0004, 0.391, 0.0000)	0.113	[0.121, 0.544]	[0.104, 0.521]
24	0.628	(−0.0013, 0.538, 0.0000)	0.155	[0.314, 0.861]	[0.298, 0.863]
48	0.924	(−0.0026, 0.257, 0.0000)	0.074	[0.607, 0.990]	[0.639, 0.991]
96	0.995	(−0.0029, 0.017, 0.0000)	0.005	[0.963, 0.999]	[0.927, 0.999]

On the logit the delta band comes within 0.04 of the draws' ends at every time, and for $F_{0\to 2}(24\,\mathrm{h}; \mathbf{z})$ it gives [0.314, 0.862] against the draws' [0.30, 0.86]; on p_2 itself the band at

48 h would run past one. The gap between the posterior expectation 0.60 and the MAP value 0.629 is the center offset: with $\boldsymbol{\delta} = (-0.046,\ -0.042,\ -0.242)^{\mathsf{T}}$, the linearized center moves by $\mathbf{g}_F^{\mathsf{T}}\boldsymbol{\delta} = -0.0003 \times (-0.046) + 0.539 \times (-0.042) + 0 \times (-0.242) = -0.023$, to 0.606, against the draws' 0.60; the exact mean of the log repair rate lies below its Laplace center, and only the repair entry of $\mathbf{g}_F$ carries it into the answer. Here the share has no entry, and the delta route works. The standard deviations of the table are close to 0.289 times the repair entry of $\mathbf{g}_{p_2}$, $\sqrt{\Sigma_{22}} = 0.289$: the band of the recovery is the repair record's alone, and $\sqrt{\Sigma_{22}}$ is a factor of $\mathrm{e}^{2\times 1.96\times 0.289} = 3.1$ between the ends of the repair rate's 95 % band.

16.4.3 Sensitivity to the priors and to the posterior used

The priors. Two priors enter the rows: the repair-rate prior Gamma(2, 20 h) and the share's prior Beta(1, 9) inside the share row. Table 16.14 varies each. Each row is a choice of prior; the columns give, each cell the state route with the draws of the same construction in parentheses, the expected unavailability with one crew, its band, the threshold probability with one crew, with a second crew and with the share halved. The repair prior is dropped, which leaves the record's own Gamma(10, 115 h); the share's prior is kept at its mean 0.1 with half its weight, half a pseudo-double in five pseudo-events, and with twice its weight, two pseudo-doubles in twenty pseudo-events, which with the log's one double in ten events give the posteriors Beta(1.5, 13.5) and Beta(3, 27).

Table 16.14: The answers under other priors, one crew unless stated: state route (the sum (16.38) and the delta method at each posterior), with in parentheses 2×10^4 draws of each posterior. The first row is the chapter's choice.

Priors	$\mathbb{E}[U \mid \mathbf{y}]$	95 % band	$\Pr(U < 10^{-4})$	second crew	share halved
Gamma(2, 20 h), Beta(1, 9)	1.97×10^{-4} (1.66×10^{-4})	$[0.33,\ 7.0] \times 10^{-4}$ ($[0.19,\ 5.2] \times 10^{-4}$)	0.30 (0.38)	0.64 (0.71)	0.63 (0.70)
repair prior dropped	2.03×10^{-4} (1.73×10^{-4})	$[0.33,\ 7.3] \times 10^{-4}$ ($[0.19,\ 5.6] \times 10^{-4}$)	0.29 (0.37)	0.63 (0.70)	0.62 (0.68)
share prior of half the weight, posterior Beta(1.5, 13.5)	2.07×10^{-4} (1.65×10^{-4})	$[0.28,\ 8.2] \times 10^{-4}$ ($[0.13,\ 5.6] \times 10^{-4}$)	0.32 (0.41)	0.63 (0.72)	0.62 (0.70)
share prior of twice the weight, posterior Beta(3, 27)	1.87×10^{-4} (1.65×10^{-4})	$[0.40,\ 5.8] \times 10^{-4}$ ($[0.28,\ 4.7] \times 10^{-4}$)	0.27 (0.34)	0.66 (0.72)	0.65 (0.70)

The expected unavailability moves by about 5 % either way on the state route and by under 5 % on the draws, since every prior of the table keeps the share's mean and the repair prior is light against ten repairs. The weight of the share's prior moves the band: half the weight lowers its lower end by a factor of 1.2 (1.5), twice the weight raises it by 1.2 (1.5), and the threshold probability ranges from 0.27 to 0.32 (0.34 to 0.41). No prior of the table brings the limit near the assurance of 0.90, with one crew or with either intervention, on either route, and the two interventions stay worth the same under each.

The posterior used. The share's posterior is the Beta approximation of Chapter 15, whose mean 0.100 is low against the exact 0.131, and the rate and the share are taken as independent. Replacing both by the posterior of the state of that chapter, equation (15.20), with the entry z_2 of equation (16.7), gives by the sum and the delta method with one crew $\mathbb{E}[U \mid \mathbf{y}] = 2.59 \times 10^{-4}$, band $[4.2 \times 10^{-5},\ 9.4 \times 10^{-4}]$ and $\Pr(U < 10^{-4} \mid \mathbf{y}) = 0.19$; with two crews 1.30×10^{-4} and 0.50; with the share halved 1.33×10^{-4} and 0.49; and with both interventions 6.6×10^{-5} and 0.80. Replacing them by the exact joint posterior of that chapter, the share drawn from its marginal

and the rate from its conditional gamma, with the repair rate from equation (16.8), gives with one crew $\mathbb{E}[U \mid \mathbf{y}] = 2.2 \times 10^{-4}$, band $[2.4 \times 10^{-5},\ 7.1 \times 10^{-4}]$ and $\Pr(U < 10^{-4} \mid \mathbf{y}) = 0.27$; with two crews 1.1×10^{-4} and 0.58; with the share halved 1.1×10^{-4} and 0.57; and with both interventions 5.7×10^{-5} and 0.86. Under the posterior of that chapter both interventions fall short of the assurance of 0.90 on both routes. The approximation lowers the answer by about a quarter on both routes and raises the threshold probability by 0.11, more than any prior of Table 16.14. The ranking of the two interventions against each other is not affected, since both act on the same first-order term, but the reported values, the variance budget and the preposterior would all change, the last two because the posterior of the share in that chapter is wider and correlated with the rate. With the exact posterior and one crew, the chain agrees with the three-process construction of Chapter 15, 2.14×10^{-4} and 0.26, to within the effect of the crew and of the uncertain repair rate; on the state route the same comparison is 2.59×10^{-4} and 0.19 against 2.62×10^{-4} and 0.16.

16.4.4 The Markov assumption and the repair-time distribution

The chain rests on exponential holding times. For the failure times the check is the one of Chapter 6, a constant rate over the record, and the trend tests of Chapter 12 apply to the log. For the repair times the check is the coefficient of variation of the recorded durations, which is one for an exponential duration. The four repairs of Chapter 8 had a coefficient of variation of 0.34 (Section 8.4.3), far more regular than exponential; the ten durations behind equation (16.8) must be checked in the same way, and if they are as regular, the exponential repair is the caveat of every result of the chapter.

What a non-exponential repair changes follows from the relations of the chain. The unavailability is the outage frequency times the mean down time, equation (16.23), and at first order in ϵ the frequency is the common-cause rate $\beta_{\mathrm{cc}}\lambda$ and the down time is one repair, so the long-run unavailability depends on the repair distribution, at that order, only through its mean; the second-order term, the wait for the crew, and the transient of Section 16.3.5 depend on its shape, and a regular repair returns the pair more predictably than the exponential curves of Figure 16.2 show. A repair with a coefficient of variation below one is represented in the chain by a sequence of exponential phases, each a further state, a phase-type distribution: k phases of equal rate $k\mu$ give the mean $1/\mu$ and the coefficient of variation $1/\sqrt{k}$, so 0.34 needs about nine phases per repair, and the chain grows but keeps every relation derived for it. The companion volume describes the general multistate and phase-type constructions.

16.4.5 The repair jobs of a common-cause outage

The chain assumes that each pump of a common-cause outage is a separate repair job, so that two crews repair them in parallel and the outage ends with the first completion, at the rate 2μ. If the shared cause instead leaves damage that is repaired as one job, a contaminated supply to be flushed or a protection fault to be found, as Chapter 15 assumed, the outage ends at the rate μ whatever the number of crews, and both pumps return together. The chain then needs two states with none up, one entered by the common cause and left for both up at μ, the other entered from one up by a second failure and left for one up at $c\mu$. The balance equations across the cuts around each of them and around both up give, with $\pi_{0\mathrm{c}}$ and $\pi_{0\mathrm{i}}$ the two states,

$$\pi_1 = 2(1-\beta_{\mathrm{cc}})\,\epsilon\,\pi_2, \qquad \pi_{0\mathrm{c}} = \beta_{\mathrm{cc}}\,\epsilon\,\pi_2, \qquad \pi_{0\mathrm{i}} = \frac{\epsilon}{c}\,\pi_1, \qquad U = \pi_{0\mathrm{c}} + \pi_{0\mathrm{i}},$$

and the number of crews divides only the term of the independent doubles. At the MAP state $U = 1.508 \times 10^{-4}$ with one crew and 1.489×10^{-4} with two, a reduction of 1.3 %; under the posterior of the state, by the sum and the delta method, $\mathbb{E}[U \mid \mathbf{y}] = 1.97 \times 10^{-4}$ and 1.94×10^{-4}, with $\Pr(U < 10^{-4} \mid \mathbf{y}) = 0.30$ and 0.31, while halving the share still gives 1.01×10^{-4} and 0.63;

by the same draws, $\mathbb{E}[U \mid \mathbf{y}] = 1.65 \times 10^{-4}$ and 1.63×10^{-4}, with $\Pr(U < 10^{-4} \mid \mathbf{y}) = 0.38$ and 0.39, while halving the share still gives 8.6×10^{-5} and 0.70. Under one job a second crew is worth one to two percent and the halved share is worth a halving, so the equivalence of the two interventions in Table 16.11 holds only under separate jobs. Which assumption fits the plant is a question about how its common-cause outages are repaired, and the record of their repair times, kept per pump and per job, answers it.

16.5 The report (step 6)

The pair is out of service far more than the first reading said. Two independent pumps would both be down $q^2 = 2.16 \times 10^{-6}$ of the time at the posterior mean rates; with one crew and the shared cause that the log's one double reveals, the pair's expected unavailability is 1.97×10^{-4} (1.66×10^{-4} by the exact draws), 1.73 h (1.45 h) of lost cooling supply a year, and 97 % (96 %) of it is the common cause, while the wait for the single crew adds one to two percent. The probability that the unavailability is below the limit of 10^{-4} is 0.30 (0.38), so the limit is not demonstrated at 0.90. A second crew halves the expected unavailability, to 9.9×10^{-5} (8.3×10^{-5}), only if the two pumps of a common-cause outage are repaired as separate jobs; if they are one job it buys one to two percent. Under either kind of repair it leaves the outages, 0.147 (0.118) a year, and the 8.5 yr to the first unchanged. Halving the share gives nearly the same unavailability, 1.02×10^{-4} (8.6×10^{-5}), under either kind of repair, halves the outages and nearly doubles the time to the first. Neither intervention alone meets the limit at 0.90. Both together reach 0.89 on the state route and 0.93 by the draws, on opposite sides of 0.90, and fall to 0.80 and 0.86 under the joint posterior of the share of Chapter 15.

16.5.1 The answers and their bands

Table 16.15 gives the answers. Each row is a question of Table 16.2; the second column gives the answer from the posterior of the state, $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of equation (16.7); the third the value from the draws of the exact posteriors of Table 16.4; and the last the band of the reported answer, with the band of the draws in parentheses, and the comparisons that qualify it.

16.5.2 The requirement and the design decision

The report line. With one repair crew, the chilled-water pump pair is expected to be out of service 2.0×10^{-4} (1.7×10^{-4} by the exact draws) of the time, about 1.7 h (1.5 h) of lost cooling supply a year, with a 95 % band from 3.3×10^{-5} to 7.0×10^{-4} (1.9×10^{-5} to 5.2×10^{-4}); the probability that its unavailability is below 10^{-4} is 0.30 (0.38). Almost all of it is the common cause, which strikes about once in eight and a half years and takes the whole pair out for one repair, 11 to 12 h; without it the pair would be out 0.05 h a year, and waiting for the single crew adds only one to two percent. After a double outage the pair is whole again within a day with probability 0.63 (0.60), with a band from 0.31 to 0.86 (0.30 to 0.86). The numbers rest on the Beta approximation of the share, which lowers them by about a quarter, and on exponential repair times, which the repair record should be checked against.

The design line. The limit of 10^{-4} at 0.90 is met by neither intervention alone, on either route. Both together reach 0.89 on the state route and 0.93 by the exact draws, on opposite sides of the required 0.90; under the joint posterior of the share of Chapter 15 they reach 0.80 and 0.86, both short of it. Whether both together meet the limit therefore depends on the route and on the posterior used, and the report states the values, not a verdict. Between the two interventions, the halved share is the robust choice on both routes: it gives nearly the same unavailability as the second crew under separate repair jobs and keeps it under one job, and it

Table 16.15: The answers to the questions asked of the pump pair, from the posterior of the state, with the value from 2×10^4 draws of the exact posteriors of Table 16.4 beside each. Bands are central 95 % posterior intervals of the reported route, draws in parentheses.

Question	Answer (state route)	Exact (draws)	Band and comparison
Long-run unavailability, one crew	$\mathbb{E}[U \mid \mathbf{y}] = 1.97 \times 10^{-4}$, 1.73 h per year; median 1.51×10^{-4}	1.66×10^{-4}, 1.45 h per year; median 1.29×10^{-4}	$[3.3 \times 10^{-5},\ 7.0 \times 10^{-4}]$ ($[1.9 \times 10^{-5},\ 5.2 \times 10^{-4}]$); 2.59×10^{-4} (2.2×10^{-4}) under the joint posterior of the share
Common cause and crew	common cause 97 % of the expected time out (without it 6.1×10^{-6}); crew wait 1 to 2 % (3.1×10^{-6})	96 % (without it 6.2×10^{-6}); 3.1×10^{-6}	$q^2 = 2.16 \times 10^{-6}$ for two independent pumps at the posterior mean rates
First outage and frequency	$\text{MTTF}_{\text{pair}}$ median 8.5 yr; 0.147 outages per year, 97 % common cause; $\Pr(\text{MTTF}_{\text{pair}} > 5\,\text{yr}) = 0.77$	median 10.3 yr; 0.118 per year; 0.85	[2.1, 35] yr ([3.0, 64] yr); 297 yr without the common cause; $U = \nu_{\text{out}} \times \text{MDT}$
$\Pr(U < 10^{-4} \mid \mathbf{y})$	0.30 with one crew, 0.64 with two; not demonstrated at 0.90	0.38, 0.71; not demonstrated	0.27 to 0.32 (0.34 to 0.41) over the share's priors; 0.19 (0.27) under the joint posterior
Whole again within a day after a double outage	$\Pr(T_{0\to 2} \le 24\,\text{h} \mid \mathbf{y}) = 0.625$	0.60	[0.314, 0.862] ([0.30, 0.86]); at the MAP state 0.29 at 12 h, 0.63 at 24 h, 0.92 at 48 h, mean 22.5 h; time constants 11.6 and 10.8 h; two crews 0.78 at 24 h (MAP)
Second crew or halved share	second crew 9.9×10^{-5}, Pr = 0.64, 0.147 outages a year; halved share 1.02×10^{-4}, 0.63, 0.076 a year; both 5.1×10^{-5}, 0.89, the limit not met at 0.90	second crew 8.3×10^{-5}, 0.71, 0.118; halved share 8.6×10^{-5}, 0.70, 0.061; both 4.3×10^{-5}, 0.93, the limit met at 0.90	the two routes fall on opposite sides of 0.90 for both together; second crew 1.94×10^{-4} (1.63×10^{-4}) if a common-cause outage is one repair job; halved share unchanged by that; both 0.80 (0.86) under the joint posterior, short of 0.90

halves the number of outages, which a second crew never changes. Which of them is worth its price, and whether either is worth more than the information of a longer record, is a question of cost against the value of the reduction; Exercise 16.6 prices the outages and the hours, and a later chapter, in Volume II, takes up the value of information and investment in general.

16.6 The next measurement (step 7)

The report leaves the limit of 10^{-4} undemonstrated with one crew and with two, under a band whose ends differ by a factor of 21 (27 by the exact draws). A longer record does not lower the unavailability: the expected posterior mean after a record is the present posterior mean. It narrows the band, and the preposterior of Chapter 2 measures how much before the record exists.

16.6.1 The candidate records

Five candidates are ranked: ten more repair durations, which feed the repair row alone; five more years of the failure log counted on the rate alone, which isolates the part of that record that informs z_1; five more years of the log with every row, the singles, the doubles and the repairs that follow them; twenty more years of the same; and, as a floor, rates known exactly, which no record of failures and repairs can pass unless it also informs the share. A candidate record enters as the rows it would add, written in the working form of equation (4.43) at the MAP state with their expected counts in place of the counts not yet seen; its information is $\mathbf{\Lambda}_a = \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a$. The singles and doubles rows are not canonical in ξ: their log means carry $\ln(1-\beta_{\mathrm{cc}})$ and $\ln\beta_{\mathrm{cc}}$, whose second derivative in ξ is $-\beta_{\mathrm{cc}}(1-\beta_{\mathrm{cc}})$, and their curvature term is the observed-minus-expected count times it, equation (15.16). That term is dropped here because a record not yet seen enters with its expected counts, so its residual is zero on average and so is the curvature term.

16.6.2 The preposterior of the log unavailability

Five more years of the failure log, $T'_{\mathrm{tot}} = 43\,800\,\mathrm{h}$, add three rows: the singles, with expected count $\bar{k}_s = 2(1-\beta_{\mathrm{cc}})\lambda T'_{\mathrm{tot}} = 10.32$ and log mean $\ln 2 + \ln(1-\beta_{\mathrm{cc}}) + z_1 + \ln T'_{\mathrm{tot}}$; the doubles, $\bar{k}_d = \beta_{\mathrm{cc}}\lambda T'_{\mathrm{tot}} = 0.574$ with log mean $\ln\beta_{\mathrm{cc}} + z_1 + \ln T'_{\mathrm{tot}}$; and the gamma-sum row of the repairs that follow them, $\bar{k}_r = 2\lambda T'_{\mathrm{tot}} = 11.47$ pump repairs, on z_2. The derivatives of the log means are $(1, 0, -\beta_{\mathrm{cc}})$, $(1, 0, 1-\beta_{\mathrm{cc}})$ and $(0, 1, 0)$, since $\partial\ln(1-\beta_{\mathrm{cc}})/\partial z_3 = -\beta_{\mathrm{cc}}$ and $\partial\ln\beta_{\mathrm{cc}}/\partial z_3 = 1-\beta_{\mathrm{cc}}$, and the 3×3 Jacobian of the candidate rows and its information are

$$\mathbf{J}_a = \begin{bmatrix} \sqrt{\bar{k}_s} & 0 & -\beta_{\mathrm{cc}}\sqrt{\bar{k}_s} \\ \sqrt{\bar{k}_d} & 0 & (1-\beta_{\mathrm{cc}})\sqrt{\bar{k}_d} \\ 0 & \sqrt{\bar{k}_r} & 0 \end{bmatrix} = \begin{bmatrix} 3.2131 & 0 & -0.3213 \\ 0.7573 & 0 & 0.6816 \\ 0 & 3.3870 & 0 \end{bmatrix},$$
$$\mathbf{\Lambda}_a = \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a = \begin{bmatrix} 10.898 & 0 & -0.516 \\ 0 & 11.471 & 0 \\ -0.516 & 0 & 0.568 \end{bmatrix}, \tag{16.49}$$

rows by record (singles, doubles, repairs) and columns by entry. The singles row loads mostly on the rate, and the doubles row carries the largest entry in the share's column, 0.68 against the singles' -0.32; the off-diagonal -0.516 is the two count rows pulling the rate and the share

in opposite directions. The preposterior covariance, equation (2.17), is

$$\boldsymbol{\Sigma}_a = (\boldsymbol{\Sigma}^{-1} + \boldsymbol{\Lambda}_a)^{-1} = \begin{bmatrix} 21.898 & 0 & -0.516 \\ 0 & 23.471 & 0 \\ -0.516 & 0 & 2.368 \end{bmatrix}^{-1} = \begin{bmatrix} 0.04590 & 0 & 0.01001 \\ 0 & 0.04260 & 0 \\ 0.01001 & 0 & 0.42451 \end{bmatrix}, \tag{16.50}$$

with $\boldsymbol{\Sigma}^{-1} = \mathbf{J}^{\mathsf{T}}\mathbf{J} = \mathrm{diag}(11,\ 12,\ 1.8)$ of equation (16.7), and the preposterior standard deviation of $\ln U$ is $(\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_a\mathbf{g})^{1/2} = 0.660$ with the gradient (16.42). Ten more repair durations alone are the one row $\mathbf{J}_a = \sqrt{10}\,(0, 1, 0)$, ten units of precision on z_2, as in Section 8.6.2; five more years counted on the rate alone keep only the information 10.898 in the corner of z_1.

Table 16.16 ranks the records. Each row is a candidate; the columns give the rows it feeds, the information it adds, the preposterior standard deviation of $\ln U$, the width it removes from the present 0.779, and the share's term in the preposterior variance.

Table 16.16: The preposterior of $\ln U$, one crew, $\boldsymbol{\Sigma}_a = (\boldsymbol{\Sigma}^{-1} + \boldsymbol{\Lambda}_a)^{-1}$ and $\sigma_{\ln U,a} = (\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_a\mathbf{g})^{1/2}$ with $\mathbf{g} = (1.024,\ -1.024,\ 0.874)^{\mathsf{T}}$. Added information at the MAP state, $\boldsymbol{\Lambda}_a = \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a$ of the candidate rows; five years by equation (16.49). Width removed: $0.779 - \sigma_{\ln U,a}$. Share term: $g_3^2\,(\boldsymbol{\Sigma}_a)_{33}$.

Record	Rows fed	Added information	$\sigma_{\ln U,a}$	Width removed	Share term
present	–	–	0.779	–	0.425
ten more repairs	repair	10 on z_2	0.753	0.026	0.425
five more years, rate row only	count, on z_1	10.90 on z_1	0.748	0.031	0.425
five more years, every row	count, repair	eq. (16.49)	0.660	0.120	0.324
twenty more years, every row	count, repair	four times the above	0.496	0.283	0.191
rates known exactly	–	infinite on z_1, z_2	0.652	0.128	0.425

16.6.3 The recommendation

The share carries 70 % of the variance of $\ln U$, panel (a) of Figure 16.3, and in the candidate Jacobian (16.49) only the count rows have entries in its column, the half double expected in five years the largest; a repair record has none. Ten more repairs and the rate part of five more years each shrink a term of 15 %, and each removes only 0.03 of the 0.78. The whole of five more years removes 0.12, close to the floor that perfectly known rates would leave; twenty years remove 0.28. The record narrows the band slowly, as Chapter 15 found for the share, and it does not move the expected unavailability.

The next measurement is therefore not a campaign but two ways of keeping the records already kept. The failure log is continued with every event classified as a single or a double within the coincidence window of Chapter 15, since only its doubles reach the share's column; and every repair is recorded per pump and per job, with its duration. The second is the one that can change a decision: the durations settle the exponential assumption of Section 16.4.4, and the record of how a common-cause outage is repaired, as one job or as two, decides whether a second crew is worth a halving of the unavailability or one to two percent of it. The crew's present record already bears on that question without settling it: under separate jobs the log's

nine singles and one double imply eleven pump repairs, under one job ten, and the record holds ten durations; but the durations are not matched to the events and need not cover the same five years, so the count leans toward one job and proves nothing. The choice of the halved share does not wait for either record, since it holds under both kinds of repair and no record of the coming years moves the expected unavailability.

16.6.4 What the analysis showed

The first reading multiplied two pumps' unavailabilities at the posterior mean rates, $q^2 = 2.16 \times 10^{-6}$, and offered a second crew for more margin; the analysis found the pair out 1.97×10^{-4} (1.66×10^{-4} by the exact draws) of the time with one crew, almost all of it from the one kind of event the product leaves out, a shared cause that takes both pumps at once. On the chilled-water plant that changes where attention goes. The crew's wait between two independent failures, the effect the first reading worried about, is one to two percent of the time out. The pair's outages, about one every eight and a half years, are nearly all common-cause events, each lasting one repair, and a limit of 10^{-4} is not met with assurance by the pair as it is. Of the two interventions, the one that acts on the cause, staggering the maintenance of the two pumps, halves the outages as well as the time out, and keeps its value however a common-cause outage is repaired; the second crew matches it only if the two pumps of such an outage are separate repair jobs, which the maintenance records should establish before a crew is hired, and it never reduces the number of outages. Taken together the two reach 0.89 on the state route and 0.93 by the exact draws, on opposite sides of the assurance of 0.90, and 0.80 and 0.86 under the joint posterior of the share of Chapter 15, both short of it. A longer log will narrow the band only slowly, because its width is the share's, informed by a single double failure.

16.6.5 Limits of the model

Exponential repair and constant rates. The chain needs exponential holding times and repairs that return a pump as good as new; Section 16.4.4 gives what a regular repair changes and the phase-type chain that represents it, and an aging pump would need the trend model of Chapter 12 in the rate.

The repair jobs. The value of the second crew rests on separate repair jobs for the two pumps of a common-cause outage; Section 16.4.5 gives the answer under one job, a reduction of one to two percent in place of a halving.

The posterior used. The share's Beta approximation and the independence of the rate and the share lower the answers by about a quarter, Section 16.4.3; under the joint posterior of Chapter 15 the budget and the preposterior would carry a correlation between the rate and the share.

The effect of staggered maintenance. That staggering the maintenance halves the share is an assumption carried over from Chapter 15, not a result of the record; the halved share's numbers are those of the assumption, and a record of the doubles after the change is what would test it.

A degraded state as a fourth state. A pump may not fail outright but degrade, delivering a reduced flow before it fails, and the pair's chain then needs a state for it. Let each pump at full flow degrade at a rate λ_{deg}, let a degraded pump fail at a rate λ_{fd}, larger than λ, and let the crew restore a degraded pump to full flow at a rate μ_{d}, a planned repair that is quicker than a repair after failure; the common cause is left out. The service requires the full flow of at

least one pump, so it is lost whenever no pump is at full flow. The pair then has four states: both pumps at full flow; one degraded with the other at full flow; one down with the other at full flow; and service lost, a state that lumps the three conditions with no pump at full flow (both down, one down and one degraded, both degraded). The transitions are:

- both at full flow to one degraded at $2\lambda_{\mathrm{deg}}$ and to one down at 2λ;
- one degraded back to both at full flow at μ_{d}, to one down at λ_{fd} when the degraded pump fails, and to service lost at $\lambda + \lambda_{\mathrm{deg}}$ when the pump at full flow fails or degrades;
- one down back to both at full flow at μ, and to service lost at $\lambda + \lambda_{\mathrm{deg}}$ when the running pump fails or degrades;
- service lost to one down at $c\mu$, the crew returning one pump to full flow.

The last rate is a simplification of the lumped state: a loss reached by degradation alone could be ended by the quicker restoration at μ_{d}, and treating every loss as a repair after failure errs on the side of a longer outage. The generator is 4×4, its rows sum to zero, and the stationary distribution, the balance equations, the outage frequency, the first passage and the implicit gradient follow as for the pair, with three more entries of the state, $\ln \lambda_{\mathrm{deg}}$, $\ln \lambda_{\mathrm{fd}}$ and $\ln \mu_{\mathrm{d}}$, and their own rows, a record of degradations found at inspection, of failures of degraded pumps and of the durations of planned restorations. The rates are no longer all multiples of λ or μ, so $g_2 = -g_1$ no longer holds; the gradient with respect to the five log rates sums to zero instead. The degraded state also changes the questions: the probability of being in it is a measure of the flow margin, and a reading of the pump's flow, not only its failure record, informs it. Exercise 16.5 works the chain with stated rates.

The close of Volume I. The chapter closes Volume I, whose evidence since Part II has been records of events: demands, failures, repairs, simultaneous outages. Chapter 17, which opens Volume II, returns to the physical model of the plant with which Part I began, a state that carries temperatures, flows and conductances fitted to readings, and asks the questions of reliability of a plant whose physics is part of the graph.

16.7 What to retain

Definition 16.1 (The Markov chain of a repairable pair)**.** A one-out-of-two pair with total failure rate λ per unit, common-cause share β_{cc}, repair rate μ and c crews is a continuous-time Markov chain on the number of units up, with the generator (16.10). Its state probabilities obey $\mathrm{d}\mathbf{p}/\mathrm{d}t = \mathbf{p}\mathbf{Q}$, with solution $\mathbf{p}(0)\mathrm{e}^{\mathbf{Q}t}$; its stationary distribution solves $\boldsymbol{\pi}\mathbf{Q} = \mathbf{0}$, $\boldsymbol{\pi}\mathbf{1} = 1$, and the long-run unavailability $U = \pi_0 \approx [\beta_{\mathrm{cc}}\epsilon + K\epsilon^2]/c$, $K = (2 - \beta_{\mathrm{cc}})(1 - \beta_{\mathrm{cc}}) - \beta_{\mathrm{cc}}^2/c$, with $\epsilon = \lambda/\mu$ and each pump a separate repair job, equals the outage frequency times the mean down time $1/(c\mu)$. The mean time to the first outage solves $(-\mathbf{Q}_{TT})\mathbf{m} = \mathbf{1}$. Under the posterior of the state, U is a derived quantity defined implicitly: the gradient matrix $\boldsymbol{\Pi}_{\mathbf{z}} = \partial\boldsymbol{\pi}^{\mathsf{T}}/\partial\mathbf{z}^{\mathsf{T}}$ solves $\mathbf{A}^{\mathsf{T}}\boldsymbol{\Pi}_{\mathbf{z}} = -\mathbf{B}$, the gradient of any stationary quantity $\mathbf{v}^{\mathsf{T}}\boldsymbol{\pi}^{\mathsf{T}}$ is $\boldsymbol{\Pi}_{\mathbf{z}}^{\mathsf{T}}\mathbf{v}$, and the transient's gradient $\mathbf{S}(t)$ obeys the sensitivity equation $\mathrm{d}\mathbf{S}/\mathrm{d}t = \mathbf{Q}^{\mathsf{T}}\mathbf{S} + [(\partial\mathbf{Q}/\partial z_i)^{\mathsf{T}}\mathbf{P}]_i$ from $\mathbf{S}(0) = \mathbf{0}$.

Four conclusions carry beyond the pump pair. A repair capacity is a constraint that couples units, and the chain of the units' joint states is the model that holds it; the product of unit unavailabilities is right only with a crew per unit and no shared cause. A derived quantity defined as the solution of a linear system has its gradient matrix from the same system, and a transient one from its sensitivity equation, which is what makes the delta method available for models with no closed form; which rate matters is read off a column of that matrix or the entries of $\mathbf{g}$, not off a formula. The unavailability is the outage frequency times the mean down time, and the two design levers act on different factors: more repair capacity shortens outages, while a smaller common cause prevents them. And when most of the variance of the answer is a common-cause share informed by a handful of double failures, the delta method's normal

differs from the draws by amounts that can cross an assurance level, and both values are stated, while a longer record narrows the band slowly and an intervention changes the answer at once.

Table 16.17 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 16.17: The terms of repairable and multistate systems and where each is defined.

Term	Expression	Units	Meaning	Defined in
Transition rates	$Q_{21}=2(1-\beta_{\mathrm{cc}})\lambda$, $Q_{20}=\beta_{\mathrm{cc}}\lambda$, $Q_{10}=\lambda$, $Q_{12}=\mu$, $Q_{01}=c\mu$	h^{-1}	jumps of the pair	16.3.1
Generator matrix	$\mathbf{Q}$, $Q_{ii}=-\sum_{j\neq i}Q_{ij}$, $\mathbf{Q1}=\mathbf{0}$	h^{-1}	rates of the chain	16.3.1
Forward equations	$\mathrm{d}\mathbf{p}/\mathrm{d}t=\mathbf{pQ}$	h^{-1}	balance over Δt	16.3.1
Stationary distribution	$\boldsymbol{\pi}\mathbf{Q}=\mathbf{0}$, $\boldsymbol{\pi}\mathbf{1}=1$	–	long-run fractions	16.3.1
Balance across a cut	flow out of $\mathcal{S}$ = flow into $\mathcal{S}$	h^{-1}	closed forms	16.3.1
Pair unavailability	$U=\pi_0\approx[\beta_{\mathrm{cc}}\epsilon+K\epsilon^2]/c$, $K=(2-\beta_{\mathrm{cc}})(1-\beta_{\mathrm{cc}})-\beta_{\mathrm{cc}}^2/c$	–	second order in ϵ	16.3.1
Outage frequency	$\nu_{\mathrm{out}}=\pi_2\beta_{\mathrm{cc}}\lambda+\pi_1\lambda$	h^{-1}	entries into none up	16.3.1
Mean down time	$\mathrm{MDT}=1/(c\mu)$	h	one outage	16.3.1
Unavailability relation	$U=\nu_{\mathrm{out}}\times\mathrm{MDT}$	–	cut around none up	16.3.1
First-passage mean	$(-\mathbf{Q}_{TT})\mathbf{m}=\mathbf{1}$; $\mathrm{MTTF}_{\mathrm{pair}}=m_2$	h	time to first outage	16.3.1
Transient	$\mathbf{p}(t)=\mathbf{p}(0)\,\mathrm{e}^{\mathbf{Q}t}$	–	after a given state	16.3.1
Time constants	t_k, with $-1/t_k$ the eigenvalues of $\mathbf{Q}$	h	approach to the long run	16.3.1
Predictive unavailability	$\mathbb{E}[\pi_0(\mathbf{z})\mid\mathbf{y}]$	–	by the sum, draws beside it	16.3.3
Rows of the posterior	$\mathcal{L}_k$, $\mathbf{r}(\mathbf{z})$, $\mathbf{J}$ with one nonzero per row; $\boldsymbol{\Lambda}=\mathbf{J}^{\mathsf{T}}\mathbf{J}$, $\boldsymbol{\Sigma}=\boldsymbol{\Lambda}^{-1}$	–	covariance of the entries	16.1.5
Implicit gradient	$\mathbf{A}^{\mathsf{T}}\boldsymbol{\Pi}_{\mathbf{z}}=-\mathbf{B}$; $\mathbf{g}_f=\boldsymbol{\Pi}_{\mathbf{z}}^{\mathsf{T}}\mathbf{v}$	–	gradient of stationary quantities	16.3.2
Transient sensitivity	$\mathrm{d}\mathbf{S}/\mathrm{d}t=\mathbf{Q}^{\mathsf{T}}\mathbf{S}+[(\partial\mathbf{Q}/\partial z_i)^{\mathsf{T}}\mathbf{P}]_i$, $\mathbf{S}(0)=\mathbf{0}$	h^{-1}	gradient of $\mathbf{P}(t)$	16.3.2
Delta method on $\ln U$	$\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$, $\mathbf{g}=\boldsymbol{\Pi}_{\mathbf{z}}^{\mathsf{T}}\mathbf{e}_0/\pi_0$	–	budget of three terms	16.3.3
Center offset	$\boldsymbol{\delta}=\mathbb{E}[\mathbf{z}\mid\mathbf{y}]-\hat{\mathbf{z}}$; center shift $\mathbf{g}^{\mathsf{T}}\boldsymbol{\delta}$	–	exact-moment normal	16.3.3
Phase-type repair	k phases of rate $k\mu$, coefficient of variation $1/\sqrt{k}$	–	non-exponential repair	16.4.4
Preposterior of $\ln U$	$(\boldsymbol{\Sigma}^{-1}+\mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a)^{-1}$, $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_a\mathbf{g}$	–	value of a record	16.6.2

16.8 Exercises

Exercise 16.1. *Formulation.* A cooling service needs two of three identical pumps, each failing at $\lambda=2\times10^{-4}\,\mathrm{h}^{-1}$, repaired one at a time by one crew at $\mu=0.1\,\mathrm{h}^{-1}$, with no common cause, the rates taken as known. (a) Write the states and the generator matrix. (b) Derive the stationary distribution by the balance equations. (c) Compute the service unavailability and the hours per year without service, and write the question whether it is below 10^{-4}, for uncertain rates, as a probability over the state. ((a) States $3,2,1,0$ pumps up; $Q_{32}=3\lambda$, $Q_{21}=2\lambda$, $Q_{10}=\lambda$, $Q_{23}=Q_{12}=Q_{01}=\mu$, diagonal minus the row sums. (b) Cuts between neighbors: $\pi_2=3\epsilon\pi_3$, $\pi_1=2\epsilon\pi_2=6\epsilon^2\pi_3$, $\pi_0=\epsilon\pi_1=6\epsilon^3\pi_3$, with $\epsilon=0.002$ and $\pi_3=1/(1+3\epsilon+6\epsilon^2+6\epsilon^3)=0.99401$. (c) $U=\pi_1+\pi_0=2.39\times10^{-5}$, 0.21 h per year; the

question is $\Pr(\pi_1(\mathbf{z}) + \pi_0(\mathbf{z}) < 10^{-4} \mid \mathbf{y})$ with $\mathbf{z} = (\ln\lambda,\ \ln\mu)^{\mathsf{T}}$.)

Exercise 16.2. *Derivation.* For the two-out-of-three chain of the previous exercise, with $\mathbf{z} = (\ln\lambda,\ \ln\mu)^{\mathsf{T}}$ at $\lambda = 2\times10^{-4}\,\mathrm{h}^{-1}$ and $\mu = 0.1\,\mathrm{h}^{-1}$: (a) write the two slices $\partial\mathbf{Q}/\partial z_i$ and the system for the 4×2 gradient matrix $\mathbf{\Pi_z}$; (b) show that $(1,1)^{\mathsf{T}}$ is a null vector of it; (c) evaluate the gradient of $\ln U$, $U = \pi_1+\pi_0$, and its delta-method standard deviation for $\mathbf{\Sigma} = \mathrm{diag}(0.1,\ 0.1)$, and say why the entries are near ±2. ((a) In the order $3, 2, 1, 0$, $\partial\mathbf{Q}/\partial z_1$ holds the failure rates $3\lambda, 2\lambda, \lambda$ above the diagonal with their negatives on it, and $\partial\mathbf{Q}/\partial z_2$ the repair rates μ below it with theirs; with $\mathbf{A}$ the generator with its last column replaced by ones, $\mathbf{A}^{\mathsf{T}}\mathbf{\Pi_z} = -\mathbf{B}$, column i of $\mathbf{B}$ being $((\partial\mathbf{Q}/\partial z_i)^{\circ})^{\mathsf{T}}\boldsymbol{\pi}^{\mathsf{T}}$. (b) The two slices sum to $\mathbf{Q}$, so $\mathbf{B}(1,1)^{\mathsf{T}} = (\mathbf{Q}^{\circ})^{\mathsf{T}}\boldsymbol{\pi}^{\mathsf{T}} = \mathbf{0}$ and $\mathbf{\Pi_z}(1,1)^{\mathsf{T}} = \mathbf{0}$. (c) $\boldsymbol{\pi} = (0.99401,\ 5.964\times10^{-3},\ 2.386\times10^{-5},\ 4.77\times10^{-8})$, $U = 2.390\times10^{-5}$, $\mathbf{g} = \mathbf{\Pi_z}^{\mathsf{T}}(0,0,1,1)^{\mathsf{T}}/U = (1.996,\ -1.996)^{\mathsf{T}}$ and $(\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g})^{1/2} = 0.893$; the entries are near ±2 because the service is lost at the second failure, so $U \approx 6\epsilon^2$ and $\ln U$ moves twice as fast as $\ln\epsilon = z_1 - z_2$.)

Exercise 16.3. *Derivation.* For the single pump of Chapter 8, with the generator $\left[\begin{smallmatrix}-\lambda & \lambda\\ \mu & -\mu\end{smallmatrix}\right]$, the state $\mathbf{z} = (\ln\lambda,\ \ln\mu)^{\mathsf{T}}$, $\lambda = 0.002\,\mathrm{h}^{-1}$, $\mu = 0.1\,\mathrm{h}^{-1}$ and $\mathbf{\Sigma} = \mathrm{diag}(1/6,\ 1/6)$ of equation (8.22), (a) write the sensitivity equation for the 2×2 matrix $\mathbf{S}(t) = \partial\mathbf{P}/\partial\mathbf{z}^{\mathsf{T}}$ of a pump up at time zero; (b) evaluate the gradient of $P_{\mathrm{down}}(t)$ at the time constant $t_c = 1/(\lambda+\mu)$ and check it against $\mathbf{S}(t)\mathbf{n} = t\,\mathbf{Q}^{\mathsf{T}}\mathbf{P}(t)$ with $\mathbf{n} = (1,1)^{\mathsf{T}}$; (c) give the delta-method standard deviation of $P_{\mathrm{down}}(t_c)$ and say which record carries most of it. ((a) $\mathrm{d}\mathbf{S}/\mathrm{d}t = \mathbf{Q}^{\mathsf{T}}\mathbf{S} + [(\partial\mathbf{Q}/\partial z_1)^{\mathsf{T}}\mathbf{P}\ \ (\partial\mathbf{Q}/\partial z_2)^{\mathsf{T}}\mathbf{P}]$, $\mathbf{S}(0) = \mathbf{0}$, with $\partial\mathbf{Q}/\partial z_1 = \lambda\left[\begin{smallmatrix}-1 & 1\\ 0 & 0\end{smallmatrix}\right]$ and $\partial\mathbf{Q}/\partial z_2 = \mu\left[\begin{smallmatrix}0 & 0\\ 1 & -1\end{smallmatrix}\right]$. (b) At $t_c = 9.804\,\mathrm{h}$, $P_{\mathrm{down}} = 0.012395$ and the second row of $\mathbf{S}$ is $\mathbf{g}^{\mathsf{T}} = (0.012293,\ -0.005080)$; its sum 0.007213 equals $t_c\,\mathrm{d}P_{\mathrm{down}}/\mathrm{d}t = t_c(\lambda P_{\mathrm{up}} - \mu P_{\mathrm{down}})$. (c) $(\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g})^{1/2} = 0.00543$; the failure record carries 85 % of the variance, because shortly after a repair a pump is down only if it has failed again, while the repair rate matters only once a failure has happened.)

Exercise 16.4. *Derivation.* (a) Derive the mean time to failure of one repairable pump as a first-passage mean from the system $(-\mathbf{Q}_{TT})\mathbf{m} = \mathbf{1}$, and check it against $1/\lambda$. (b) Derive the same for the pair without a common cause and one crew, and evaluate it at the chapter's MAP state. ((a) The only state before the outage is up, $-\mathbf{Q}_{TT} = \lambda$, so $m = 1/\lambda$, the MTTF of Chapter 8. (b) $2\lambda m_2 - 2\lambda m_1 = 1$ and $-\mu m_2 + (\lambda+\mu)m_1 = 1$ give $m_2 = (3\lambda+\mu)/(2\lambda^2) = 0.089282/(3.4297\times10^{-8})\,\mathrm{h} = 2.60\times10^6\,\mathrm{h} = 297\,\mathrm{yr}$, about $\mu/(2\lambda^2)$, the reciprocal of the rate $2\lambda\times\epsilon$ at which one pump fails and the other fails during its repair.)

Exercise 16.5. *Application.* In the four-state chain of Section 16.6.5, take $\lambda = 1.31\times10^{-4}\,\mathrm{h}^{-1}$, $\lambda_{\mathrm{deg}} = 5\times10^{-4}\,\mathrm{h}^{-1}$, $\lambda_{\mathrm{fd}} = 2\times10^{-3}\,\mathrm{h}^{-1}$, $\mu_{\mathrm{d}} = 0.25\,\mathrm{h}^{-1}$, $\mu = 1/11.25\,\mathrm{h}^{-1}$, one crew and no common cause. (a) Write the generator. (b) Compute the service unavailability and the mean time to the first loss of service from both pumps at full flow. (c) Compare with the pair without degradation. ((a) States full, degraded, one down, lost: row full $(-(2\lambda_{\mathrm{deg}}+2\lambda),\ 2\lambda_{\mathrm{deg}},\ 2\lambda,\ 0)$; degraded $(\mu_{\mathrm{d}},\ -(\mu_{\mathrm{d}}+\lambda_{\mathrm{fd}}+\lambda+\lambda_{\mathrm{deg}}),\ \lambda_{\mathrm{fd}},\ \lambda+\lambda_{\mathrm{deg}})$; one down $(\mu,\ 0,\ -(\mu+\lambda+\lambda_{\mathrm{deg}}),\ \lambda+\lambda_{\mathrm{deg}})$; lost $(0,\ 0,\ \mu,\ -\mu)$. (b) $\boldsymbol{\pi} = (0.99298,\ 3.93\times10^{-3},\ 3.04\times10^{-3},\ 4.95\times10^{-5})$, so $U = 4.95\times10^{-5}$, 0.43 h per year; the first-passage system on the first three states gives $2.29\times10^5\,\mathrm{h} = 26\,\mathrm{yr}$. (c) Without degradation $U = 4.33\times10^{-6}$ and 297 yr: the degraded state multiplies the unavailability by 11 and divides the time to the first loss by 11, because whenever one pump is down or degraded the service rests on the other alone, which now leaves full flow at $\lambda+\lambda_{\mathrm{deg}} = 6.3\times10^{-4}\,\mathrm{h}^{-1}$, almost five times its failure rate.)

Exercise 16.6. *Application.* For the chapter's pair with one crew, a second crew costs R$ 350 000 per crew-year, staggered maintenance, which halves the share, costs R$ 60 000 per year, each outage costs R$ 1 500 000 at restart and each hour of outage R$ 400 000. Using the posterior expectations of Table 16.11, (a) compute the expected annual cost of the present arrangement and of each option; and (b) choose between the second crew and the staggered maintenance, and state

why the two differ although their unavailabilities are equal. ((a) Present: $0.1177 \times 1.5 + 1.450 \times 0.4 = 0.177 + 0.580 = \text{R\$}\,0.757$ million; second crew: $0.177 + 0.290 + 0.350 = \text{R\$}\,0.817$ million; staggered maintenance: $0.0609 \times 1.5 + 0.752 \times 0.4 + 0.060 = 0.091 + 0.301 + 0.060 = \text{R\$}\,0.452$ million. (b) The staggered maintenance, which saves R$ 0.305 million a year; the second crew saves R$ 0.290 million of outage hours and costs R$ 0.350 million. The crew halves the down time but not the number of outages, so it saves nothing of the cost per outage, while halving the share halves both.)

Appendix A

Atlas of probabilistic questions

This appendix can be read independently of the chapters. Each row is a question a reliability engineer asks, its mathematical form once the plant graph and its state $\mathbf{z}$ are fixed, the object that answers it, and the chapter in which the derivation is made. The rows grow as chapters are added.

Engineering question	Mathematical form	Derives to	Chapter
What is this quantity, and how well is it known?	$p(z_i \mid \mathbf{y})$	the marginal of the posterior: $\hat{z}_i$ and Σ_{ii}	4
What is a quantity that is not an entry of the state?	$p(f(\mathbf{z}) \mid \mathbf{y})$	$f(\hat{\mathbf{z}})$ and $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ by the delta method	4
Will a temperature stay below its limit?	$\Pr(T < T_{\max} \mid \mathbf{y})$	$\Phi\big((T_{\max} - \mu_T)/\sigma_T\big)$	4
What is the failure probability per demand?	$p(p \mid k, n)$	$\mathrm{Beta}(a + k,\ b + n - k)$	5
Is the item at least p^* reliable, with stated assurance?	$\Pr(p \le 1 - p^* \mid k, n)$	the Beta cumulative distribution at $1 - p^*$	5
Will the next m demands succeed?	$\Pr(\text{no failure in } m \mid k, n)$	the Beta–binomial predictive	5
How many tests demonstrate a reliability requirement?	n such that $\Pr(p \le 1 - p^* \mid 0, n) \ge \gamma$	the smallest n satisfying the Beta tail	5
What is the failure rate behind a count over an exposure?	$p(\lambda \mid k, T_{\mathrm{tot}})$	$\mathrm{Gamma}(a + k,\ b + T_{\mathrm{tot}})$	6
Is the rate below a limit, with stated assurance?	$\Pr(\lambda \le \lambda_{\max} \mid k, T_{\mathrm{tot}})$	the gamma cumulative distribution at $\lambda_{\max}$	6
How many failures in the next τ hours?	$\Pr(K_f = j \mid k, T_{\mathrm{tot}})$	the negative-binomial predictive	6

Engineering question	Mathematical form	Derives to	Chapter
Will the item survive a mission of length t?	$\Pr(T > t \mid \mathbf{y})$	the posterior expectation of $R(t)$: $\left(b/(b+t)\right)^a$ under a gamma posterior	6
What lifetime distribution does a record with survivors support?	$p(\ln\eta, \ln\beta \mid \mathbf{y})$	the censored likelihood, its MAP and Hessian	7
How long will a unit that has reached age a last?	$\Pr(T > a+\tau \mid T > a, \mathbf{y})$	the ratio of survivals averaged over the posterior	7
Does the record show wear-out?	$p(\mathbf{y} \mid \beta > 1)/p(\mathbf{y} \mid \beta = 1)$	an evidence ratio on the shape	7
What is the repair rate, and the mean time to repair?	$p(\mu \mid r, D)$	$\text{Gamma}(a+r,\ b+D)$	8
What fraction of the time is the item available, and how well is that known?	$p(A \mid \mathbf{y})$, $A = \mu/(\lambda+\mu)$	a Beta on the exact route; the delta method on the logit	8
Is the availability at least A_{req}?	$\Pr(A \geq A_{\text{req}} \mid \mathbf{y})$	a Beta tail, or a normal tail on the logit	8
Which record tightens the availability more?	$\boldsymbol{\Sigma}_a$ for each candidate record	the preposterior on the logit	8
Will the unit run through a mission of length t?	$\Pr(T > t \mid \mathbf{y}) = \mathbb{E}[R(t;\lambda) \mid \mathbf{y}]$	the Lomax predictive $\left(b/(b+t)\right)^a$; the plug-in understates it	9
How well is the mission reliability itself known?	$p\big(R(t;\lambda) \mid \mathbf{y}\big)$	the quantile map of the gamma posterior	9
Will the unit start and run?	$\mathbb{E}[(1-p)\,\mathrm{e}^{-\lambda t} \mid \mathbf{y}]$	a product of two entries; the budget on $\ln R$	9
How many zero-failure hours demonstrate a mission reliability?	T such that $\Pr(R(t) \geq R^* \mid 0, T) \geq \gamma$	the gamma tail at the required rate	9
Will a unit of age t_0 survive the next period?	$\Pr(T > t_0 + \tau \mid T > t_0, \mathbf{y})$	the ratio of two predictives; the record's own rows	10
How much life is left, on average?	$\mathbb{E}[T - t_0 \mid T > t_0, \mathbf{y}]$	the integral of the conditional survival (restricted horizon)	10
Will any survivor of a fleet fail before the next test?	$\Pr(\text{no failure in } \tau \mid \mathbf{y})$	one expectation over the shared state, not a product	10
How many failures will a fleet have next year?	$p(K_f \mid \mathbf{y})$	the negative binomial; the variance split into Poisson and epistemic parts	11

Engineering question	Mathematical form	Derives to	Chapter
How many spares cover the year with stated probability?	smallest s with $\Pr(K_f \le s \mid \mathbf{y}) \ge q$	a quantile of the predictive	11
Was last year's count consistent with the record before it?	$\Pr(K \ge k \mid \mathbf{y}_{\text{prior years}})$	the leave-one-out predictive tail; an evidence ratio for an alternative rate	11
Is the rate rising?	$\Pr(\gamma > 0 \mid \mathbf{y})$; $p(\mathbf{y} \mid \gamma \text{ free})/p(\mathbf{y} \mid \gamma = 0)$	a normal tail on the trend entry; Savage–Dickey	12
When does the rate cross the limit?	$p(t^* \mid \mathbf{y})$, $t^* = \bar{t} + (\ln \lambda_{\max} - \alpha)/\gamma$	draws or the inversion, not $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ (a pole)	12
What is the service availability of several units under a drawing?	$\mathbb{E}[R_{\text{sys}}(\mathbf{q}) \mid \mathbf{y}]$	the 2^n enumeration; plug-in exact under independent posteriors, a gap under a shared one	13
Which unit matters most to the service?	$\partial R_{\text{sys}}/\partial q_i$	the Birnbaum gradient, the $\mathbf{J}$ row of the band	13
Which combinations take the service down, and how likely is each?	the minimal cut sets; $\mathbb{E}[\prod_{i \in c} q_i \mid \mathbf{y}]$	a filter on the enumeration; ranked with bands	14
How likely is the top event, and how good is the rare-event sum?	$\mathbb{E}[\Pr(\text{top}) \mid \mathbf{y}]$	exact by enumeration; inclusion–exclusion orders	14
What share of the failures is a common cause?	$p(\beta_{\text{cc}} \mid k_1, k_2)$	the grid posterior on (rate, share)	15
Did one cause take two units out?	$p(\mathbf{y} \mid \text{shared cause})/p(\mathbf{y} \mid \text{independence})$; $\Pr(\text{common cause} \mid \text{the double})$	an evidence ratio with a coincidence window; an event posterior	15
What is a redundant pair's unavailability with a common cause?	$\mathbb{E}[U_{\text{pair}} \mid \mathbf{y}]$	the two-state model with the shared process	15
Which of several explanations is the plant?	$\Pr(M_j \mid \mathbf{y})$	the evidences $p(\mathbf{y} \mid M_j)$ and the prior model probabilities	2
What does a measurement reduce?	$\mathbb{E}[\text{posterior spread} \mid a]$	the preposterior $\boldsymbol{\Sigma}_a = (\boldsymbol{\Sigma}^{-1} + \boldsymbol{\Lambda}_a)^{-1}$	2
What action minimizes expected loss?	$\arg\min_a \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]$	the expected loss of each action under the posterior	2

Engineering question	Mathematical form	Derives to	Chapter
What fraction of the time is a repairable system out, with shared repair and a common cause?	$\mathbb{E}[\pi_{\text{down}}(\mathbf{z}) \mid \mathbf{y}]$	the stationary distribution of a Markov chain, $\pi Q = 0$; its gradient by implicit differentiation	16
How soon is the system whole again after an outage?	$p(t) = p(0)\,\mathrm{e}^{Qt}$	the matrix exponential; time constants from the eigenvalues	16
How long until the system's first outage?	$\mathbb{E}[\text{first passage} \mid \mathbf{y}]$	a linear solve on the transient states	16