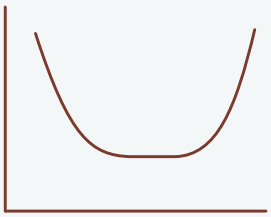
PROBABILISTIC GRAPH MODELING
FOR ENGINEERING DISCIPLINES · VOLUME III



Reliability Engineering by Derivation

Engineering decisions and integrated case studies

222 pages

Ricardo Marquez, Ph.D.

ricardo-marquez.com

Reliability Engineering by Derivation

Volume III

Engineering decisions and integrated case studies

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)

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. Volume I holds Parts I to IV, Chapters 1 to 16: the method itself and its application to records of failures, lifetimes and repairs, to predictions, and to the structure of systems. Volume II holds Parts V to VII, Chapters 17 to 28: the physical model, diagnosis, risk, importance measures, acceptance tests and operating margins. This third volume holds Parts VIII and IX, Chapters 29 to 36: the engineering decisions the posterior supports, and four integrated case studies that run the whole method on one plant each. Chapter numbers run continuously across the three volumes; a reference to a chapter numbered 28 or below points to an earlier volume, whose results this volume builds on and does not repeat.

Every chapter follows the progression fixed in Volume I's preface: step 0, the plant graph; step 1, the question as a probability or expectation over the state; step 2, the graph extended for the question and the evidence attached as rows, with the assumptions they rest on; step 3, the posterior of the state, derived and computed once; step 4, the answers, each computed from that posterior, with the exact value beside it wherever one exists; step 5, evaluation and checks; step 6, the report with its band; step 7, the next measurement. The notation chapter is reproduced here so that this volume can be read with the others closed, and the atlas of questions in the appendix covers all three volumes.

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 v

VIII Information and Engineering Decisions 1

29 Inspection and maintenance intervals 2
29.1 How often to test a battery string . . . 2
29.2 The posterior of the state (step 3) . . . 8
29.3 The answers (step 4) . . . 9
29.4 The checks (step 5) . . . 19
29.5 The report (step 6) . . . 21
29.6 The next measurement (step 7) . . . 23
29.7 What to retain . . . 25
29.8 Exercises . . . 26

30 Repair, replacement and preventive maintenance 28
30.1 When to replace a battery string . . . 28
30.2 The posterior of the state (step 3) . . . 32
30.3 The answers (step 4) . . . 33
30.4 The checks (step 5) . . . 43
30.5 The report (step 6) . . . 45
30.6 The next measurement (step 7) . . . 47
30.7 What to retain . . . 49
30.8 Exercises . . . 50

31 Value of information and sensor placement 52
31.1 Which measurement to buy . . . 52
31.2 The posterior of the state (step 3) . . . 55
31.3 The answers (step 4) . . . 56
31.4 The checks (step 5) . . . 66
31.5 The report (step 6) . . . 68
31.6 The next measurement (step 7) . . . 71
31.7 What to retain . . . 73
31.8 Exercises . . . 73

32 Redundancy, spares and reliability investment 75
32.1 Three requests on one budget . . . 75
32.2 The posterior of the state (step 3) . . . 80
32.3 The answers (step 4) . . . 81

32.4 The checks (step 5) . . . 94
32.5 The report (step 6) . . . 96
32.6 The next measurement (step 7) . . . 98
32.7 What to retain . . . 101
32.8 Exercises . . . 102

IX Integrated Engineering Case Studies 104

33 Industrial pump reliability 105
33.1 Six years of a cooling-water fleet . . . 105
33.2 The posterior of the state (step 3) . . . 109
33.3 The answers (step 4) . . . 112
33.4 The checks (step 5) . . . 121
33.5 The report (step 6) . . . 124
33.6 The next measurement (step 7) . . . 126
33.7 What to retain . . . 129
33.8 Exercises . . . 130

34 Redundant cooling system 131
34.1 Three chillers, two pumps and a request for a fourth chiller . . . 131
34.2 The posterior of the state (step 3) . . . 134
34.3 The answers (step 4) . . . 136
34.4 The checks (step 5) . . . 144
34.5 The report (step 6) . . . 146
34.6 The next measurement (step 7) . . . 148
34.7 What to retain . . . 150
34.8 Exercises . . . 151

35 Data center power and cooling 152
35.1 Six outages a year and a hall at 17.7 degrees . . . 152
35.2 The posterior of the state (step 3) . . . 156
35.3 The answers (step 4) . . . 160
35.4 The checks (step 5) . . . 168
35.5 The report (step 6) . . . 170
35.6 The next measurement (step 7) . . . 172
35.7 What to retain . . . 173
35.8 Exercises . . . 174

36 Integrated reliability assessment 175
36.1 A third hall and a signature due . . . 175
36.2 The posterior of the state (step 3) . . . 178
36.3 The answers (step 4) . . . 180
36.4 The checks (step 5) . . . 188
36.5 The report (step 6) . . . 190
36.6 The next measurement (step 7) . . . 192
36.7 What to retain . . . 194
36.8 Exercises . . . 195

A Atlas of probabilistic questions 196

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).
Physical models and the state (Part V); a row naming a chapter is that chapter's own symbol		
$\mathbf{z}_{\text{full}}$, $\mathbf{r}_{\text{full}}$, $\mathbf{J}_{\text{full}}$	varies	The full state of a plant graph, every free quantity including the temperatures and meter quantities the balances fix; its residual vector and Jacobian (Chapter 17).
$h_i(\mathbf{z})$, σ_p, $\sigma_{p,P}$	varies, K, kW	The balance of node i as a function of the state; the tolerance of a hall's physics row; the tolerance of the header's (Chapters 17 and 18).
$\mathbf{J}_{\text{e}}$, $\mathbf{J}_{\text{p}}$, $\mathbf{J}_{\text{d}}$, $\mathbf{J}_\pi$	varies	The Jacobian of the evidence rows (readings and priors) and of the physics rows (Chapter 17); the design columns, the reading rows alone, and the prior rows, with information $\mathbf{\Lambda}_{\text{prior}} = \mathbf{J}_\pi^{\mathsf{T}} \mathbf{J}_\pi$ (Chapters 19 and 22).
$\mathbf{N}$, $\mathbf{N}_{\text{b}}$	varies	Null-space basis of the physics rows, $\mathbf{J}_{\text{p}}\mathbf{N} = \mathbf{0}$, whose columns are the free directions of the reduced state $\mathbf{z}$, with reduced Jacobian $\mathbf{J} = \mathbf{J}_{\text{e}}\mathbf{N}$; its block on the balanced entries (Chapter 17).

Symbol	Unit	Meaning
$\tilde{Q}_i$	kW	A load profiled out: the least-squares value with the other entries held at the prior point (Chapter 17).
ν, n_r, n_z	–	Redundancy, the number of rows less the number of entries, $\nu = n_r - n_z$: the degrees of freedom of $2\mathcal{L}(\hat{\mathbf{z}})$.
χ^2_ν, $\chi^2_\nu(c)$	–	Chi-square distribution with ν degrees of freedom; noncentral, with noncentrality c (Chapter 18).
$\mathbf{H}$, $\mathbf{S}$	varies	The exact Hessian of the objective, $\mathbf{H} = \mathbf{J}^\mathsf{T}\mathbf{J} + \mathbf{S}$, and its curvature term $\mathbf{S} = \sum_k r_k \nabla^2 r_k$, which Gauss–Newton drops; a subscript names the variable differentiated ($\mathbf{H_q}$, Chapter 25). $\mathbf{H}$ is a Hessian throughout the book.
$\Delta\mathbf{z}$, tol, ϖ	varies, –, –	The Gauss–Newton step, $\mathbf{J}^\mathsf{T}\mathbf{J}\,\Delta\mathbf{z} = -\mathbf{J}^\mathsf{T}\mathbf{r}$; the stopping tolerance on $\max_i \lvert\Delta z_i\rvert/(1+\lvert z_i\rvert)$; the Levenberg–Marquardt damping (Chapter 18). Chapter 23 writes the step $\mathbf{s}$, since there $\mathbf{\Delta}$ is a change between periods.
$\mathbf{j}_k$, $\mathbf{j}_a$	varies	The Jacobian row of reading k as a column, $\mathbf{j}_k = \mathbf{J}_k^\mathsf{T}$; of a candidate row a.
$\mathbf{K}$, K_{kk}, $1 - K_{kk}$	–	The hat matrix $\mathbf{J}(\mathbf{J}^\mathsf{T}\mathbf{J})^{-1}\mathbf{J}^\mathsf{T}$; the leverage of row k, $\mathbf{j}_k^\mathsf{T}\mathbf{\Sigma}\mathbf{j}_k$; its redundancy number, the share of row k the other rows check, summing to ν (Chapters 18, 21 and 22).
r_k^N, r^*, p_fa	–	The normalized residual $r_k(\hat{\mathbf{z}})/\sqrt{1-K_{kk}}$; its test threshold, by Bonferroni $\Phi^{-1}(1 - p_\mathrm{fa}/2n_r)$; the false-alarm probability of the family of tests (Chapters 18 and 22).
s_j, $\mathbf{v}_j$, $\mathbf{V}$, $\mathbf{F}$	varies	The singular values of $\mathbf{J}$ in decreasing order and its right singular vectors, the columns of $\mathbf{V}$; the left factor, $\mathbf{J} = \mathbf{F}\,\mathrm{diag}(s_j)\mathbf{V}^\mathsf{T}$ ($\mathbf{W}$ is the weight matrix). The weakest direction $\mathbf{v}$, of spread $1/s_\mathrm{min}$ (Chapters 17 and 19).
$\mathrm{cond}(\cdot)$	–	Condition number, s_1/s_n for $\mathbf{J}$ and its square for $\mathbf{\Lambda}$ (Chapter 19); κ is the assurance level.
Λ_{ij}, ϱ_i^2, $\sigma_{0,i}$	varies	The entries of the information $\mathbf{\Lambda}$, $1/\Lambda_{ii}$ the conditional variance of entry i; its squared multiple correlation $1 - 1/(\Lambda_{ii}\Sigma_{ii})$; its prior spread (Chapter 19).
$\mathbf{\Lambda}_{-a}$, $\mathbf{\Sigma}_{-a}$, χ_p	varies	The information and covariance with the rows a removed; the radius of the credible ellipse of probability p (Chapter 19).
Propagation, virtual readings and diagnosis (Parts V and VI)		
$\mathrm{sd}(f \mid \mathbf{y})$	varies	The posterior standard deviation of a derived quantity, $\sqrt{\mathbf{g}^\mathsf{T}\mathbf{\Sigma}\mathbf{g}}$ by the delta method; it replaces σ_f from Part V on.
$[\mathbf{g}]_i$, V_i	varies	Entry i of the gradient; the diagonal term $[\mathbf{g}]_i^2\Sigma_{ii}$ of the variance budget (Chapter 20).
$\mathbf{d}$, $m(x)$, $v(x)$	varies	The deviation $\mathbf{z} - \hat{\mathbf{z}}$ (Chapter 20); the conditional mean and variance of the other entries given one entry x under the Laplace normal, $\Sigma_{5\mid1}$ the conditional variance of entry 5 given entry 1 (Chapters 20 and 21).

Symbol	Unit	Meaning
ζ_j, $\bar{\zeta}_j$, N_{eff}	–	The importance weight of draw j and its normalized value; the effective sample size $(\sum_j \bar{\zeta}_j^2)^{-1}$ (Chapter 20).
$\mathbf{y}_-$, $\mathcal{L}_-$, $\hat{\mathbf{z}}_-$, $\boldsymbol{\Sigma}_-$	varies	The evidence, objective, MAP state and covariance of the refit without one reading row; $\tilde{\boldsymbol{\Sigma}}_-$, $\tilde{\boldsymbol{\Lambda}}_-$ the rank-one downdate of $\boldsymbol{\Sigma}$ and $\boldsymbol{\Lambda}$ at $\hat{\mathbf{z}}$ (Chapter 21).
V_{-k}, V_a	varies	The variance of reading k predicted from the other rows, $\mathbf{g}_k^{\mathsf{T}}\boldsymbol{\Sigma}_-\mathbf{g}_k$; of a reading at a place a without one (Chapters 21 and 22).
$\text{sd}(y_k \mid \mathbf{y}_-)$, d_k	varies, –	The predictive spread of a held-out reading, $\sqrt{\sigma_k^2 + V_{-k}}$; the held-out departure $(y_k - f_k(\hat{\mathbf{z}}_-))/\text{sd}(y_k \mid \mathbf{y}_-)$ (Chapter 21).
T_r, $\dot{m}c_p$	°C, kW/K	The return temperature of the chilled-water loop; its heat-capacity flow (Chapters 21, 22 and 24).
b_k, σ_b, $\hat{b}_k$, b_k^*	K or kW	The bias entry of reading k, its prior spread and its estimate; the minimum detectable bias (Chapter 22).
B_{k0}, ψ_k	–	The Bayes factor of a bias on reading k against no bias; the ratio $\sigma_b^2(1-K_{kk})/\sigma_k^2$ of its closed form (Chapter 22).
$\mathcal{F}$, $\lvert\mathcal{F}\rvert$	–	A family of candidate drawings and its size; the posterior over the family is $\Pr(M_j \mid \mathbf{y})$ for $M_j \in \mathcal{F}$ (Chapters 22 and 23).
N, $\bar{y}_k$	d, varies	The number of days in a record and reading k averaged over them, with spread $\sigma_k/\sqrt{N}$ (Chapter 23; bold $\mathbf{N}$ is the null-space basis).
σ_{h}, σ_{w}	varies	The healthy model's prior spread of an entry, from the earlier period; the widened spread of a candidate drawing (Chapter 23).
$\boldsymbol{\Delta}$, Z_i, χ^2_Δ	varies, –, –	The change of the state between two periods, $\hat{\mathbf{z}}_2 - \hat{\mathbf{z}}_1$; its standardized entry; the joint statistic $\hat{\boldsymbol{\Delta}}^{\mathsf{T}}(\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2)^{-1}\hat{\boldsymbol{\Delta}}$ (Chapter 23).
$\Pr(M_d)$, $\ln B_{10}$	–	The prior probability of drawing M_d; the log Bayes factor in nats (Chapter 24).
$\text{KL}(A \,\|\, B)$, ς_d	nat, varies	The Kullback–Leibler divergence of the predictive under A from that under B; the predictive spread of a candidate reading under drawing d (Chapter 24).
c_{f}, Σ_{nn}, U_{c}	varies	The density of a flat prior; the posterior variance of an added entry; the unavailability of the cooling service (Chapter 24).
Risk, importance, acceptance and margins (Part VII)		
$\mathbf{q}$, $\boldsymbol{\Sigma}_{\mathbf{q}}$	–	The unavailabilities of the units as a vector and their posterior covariance (Chapter 25).
t_{yr}, $\mathcal{R}_{\text{yr}}$, L_{yr}	h, R$/yr, R$	The hours of a year, 8760 h; the annual risk $t_{\text{yr}}\mathcal{R}$, the expected loss of a year given $\mathbf{q}$; the loss of one particular year (Chapter 25).
$\mathcal{R}^*$, w_a	R$/yr, decades	A limit on the annual risk; the width of the central 95 % band of the risk after candidate a (Chapter 25).
ν_{loss}, MDT	yr^{-1}, h	The frequency of loss-of-service events; the mean down time of one (Chapter 25).

Symbol	Unit	Meaning
VaR_p, CVaR_p	R$	The p quantile of the predictive year's loss; the tail mean beyond it (Chapter 25).
U_i^+, U_i^-	–	The top event with unit i pinned down ($q_i = 1$) and pinned up ($q_i = 0$) (Chapter 26).
$I_{\mathrm{B},i}$, $\mathbf{I}_\mathrm{B}$	–	Birnbaum importance $U_i^+ - U_i^- = \partial U/\partial q_i$, the probability that unit i is critical; the vector of them (Chapter 26).
$I_{\mathrm{FV},i}$, $I_{\mathrm{CR},i}$	–	Fussell–Vesely importance $(U - U_i^-)/U$, the share a perfect unit removes; criticality importance $q_i I_{\mathrm{B},i}/U$, equal to it for units independent given $\mathbf{q}$ (Chapter 26).
RAW_i, RRW_i	–	Risk achievement worth U_i^+/U and risk reduction worth U/U_i^- (Chapter 26).
V_j, n_j	–, –	Uncertainty importance of shared entry j, its term $(g_j/U)^2\Sigma_{jj}$ in the variance of $\ln U$; the checks of its record (Chapter 26).
G^*, g^*, z_κ	kW/K, –, –	A required conductance and its log-conductance; the standard normal quantile $\Phi^{-1}(\kappa)$ (Chapters 27 and 28).
Q_test, ΔT^*, σ_c	kW, K, K	The held load of a test; the rise at the requirement, Q_test/G^* (the star marks the requirement level); the combined supply spread of sensor and set point (Chapter 27).
$\mathrm{OC}(\mathbf{z}; a)$, $U(a)$	–	The operating characteristic, the probability that test a accepts at a stated truth $\mathbf{z}$; averaged over the posterior, the probability the test ends accepted (Chapter 27).
ε_p, ε_c, π_g	–	The producer's and consumer's risks of a test; the prior's share Σ_{11}/σ_g^2 of the test's posterior (Chapter 27).
m_1^*, Q_cap, Q_κ	K, kW, kW	The margin $T_\mathrm{max} - T_1^*$ at a scenario load; the path's capacity $(T_\mathrm{max} - T_s)\,\mathrm{e}^{g_1}$; the largest load held with assurance κ (Chapter 28).
c_ch, P_pl, P^*, P_hr	–, kW, kW, kW	The chiller's capacity factor; its plate rating; the demand on it at a scenario; its headroom $P_\mathrm{pl}c_\mathrm{ch} - Q_1 - Q_2$ (Chapter 28).
ΔT_s, $(x)^+$, b_P	K, –, kW	The supply's rise when demand exceeds the delivered capacity; the positive part $\max(0, x)$; the bias of the chiller's heat meter (Chapters 23 and 28).
Information and engineering decisions (Part VIII); a row naming a chapter is that chapter's own symbol		
τ, D, d, d_0	yr, yr, –, –	Inspection interval; delay from detectable onset to failure (D is total repair duration in Chapter 8); log-median delay and its prior center (Chapter 29).
s_D, σ_d, s_tot, $\sigma_{d,a}$	–	Spread of the delay across strings; spread of its log-median; their total; the log-median's spread after a record (Chapter 29).
ν_r, p_c, $\bar{p}_c$, ΔC	yr^{-1}, –, –, R$	Replacement rate 1/MTTF; catch probability and its predictive; $C_f - C_r$ (Chapter 29).
EVPI_d, EVPI_{uv}	R$/yr	Expected value of perfect information on the delay and on the lifetime entries (Chapter 29).
T, $W(T)$, $T^*(\mathbf{z})$, T^B	yr	Replacement age; mean cycle length; the best age at a state; the Bayes action (Chapter 30).

Symbol	Unit	Meaning
$\mathcal{R}_{\text{ex}}$, L_c, L_a, $L^*(\mathbf{z})$	R\$/yr	Expected cost on the exact posterior; the condition-based loss on the common basis; age replacement with annual tests; the least loss at a state (Chapter 30).
v_{L}, c_o, C_{x}, $\bar{Q}_\kappa$	R\$/kW-yr, R\$/kW-yr, R\$, kW	Value of admitted load; overshoot price implied by κ; penalty per exceedance; lower-tail mean of the capacity (Chapter 31).
j_a, σ_a, s_a, h	varies, varies, –, nat	A candidate's row and its noise; $1 + j_a^{\mathsf{T}}\boldsymbol{\Sigma}j_a$; differential entropy (Chapter 31).
EVSI_Q, EVSI_L, EVPI_Q, EVPI_L	R\$/yr	Values of a reading and of perfect knowledge on the assurance rule and under the loss (Chapter 31).
c_u, c_k, c_s, C_s	R\$/yr, R\$/yr, R\$/yr, R\$	Annual prices of a unit, a controller and a spare; cost of a failure with no spare (Chapter 32).
U_n, $m_a(\mathbf{z})$, $\mathbf{H}_m$, ϑ	–, R\$/yr, varies, –	Unavailability of n units needing three; marginal saving of an action; its Hessian; the critical fractile $1 - c_s/C_s$ (Chapter 32).
Integrated engineering case studies (Part IX); a row naming a chapter is that chapter's own symbol		
τ, δ_i, d_i, ω_μ, ℓ	–	Population log-spread of the units' rates (τ is the inspection interval in Chapter 29); standardized and centered deviations; $\ln \mu$; $\ln \tau$ (Chapter 33).
$\hat{R}$, N_{ess}	–	A chain's between-against-within spread; its effective sample size (Chapter 33).
c_o, c_{st}, M	R\$/yr	Overhaul price (c_o is the overshoot price in Chapter 31); staggering price; overhaul saving (Chapter 33).
R_{pair}	–	The pair's mission reliability (Chapter 33).
λ_c, β_c, ω_c, ξ_c, ω_p, ξ_p, ξ_K	varies	The chillers' rate and share and their entries; the pumps' entries; the logit of the controller's probability (Chapter 34).
$U_{\text{cc},c}$, q_c, U_C, x_i	–	The chillers' common-cause, independent and group terms; an order-one event's probability (Chapter 34).
D_o, μ_o, ν_o, $S_o(t)$	h, h, yr^{-1}, –	Outage duration, its mean, the outage rate, the duration's survival (Chapter 35).
$t_\times$, T_{lim}	h, °C	Crossing time and the limit it crosses (Chapter 35).
C_{th}, C_{lp}, t_r, $x(t)$	kWh/K, kWh/K, h, K	Hall thermal mass, loop thermal mass, restart time, loop rise (Chapter 35).
$T_{[j]}$, Δt, $\sigma_{\log}$, σ_i, ΔQ_{st}	°C, h, K, K, kW	The step test's temperatures, grid, logger and sensor noise, and shed load (Chapter 35).
Π, Π_s, Π_r, Π_g, C_{cr}, c_a	–, R\$, R\$/yr	Crossing probability per outage and its start, running and generator-up paths; loss per crossing; an action's price (Chapter 35).
x_j, A_j, u_j, m_j, Γ_j, φ_j, D_j	varies	The Kalman filter's and smoother's objects at step j (Chapter 35).
g_3, G_3, b_3, ξ_q	–, kW/K, K, –	Hall 3's path and its conductance; the new sensor's bias; the logit of a cooling unit's unavailability (Chapter 36).

Symbol	Unit	Meaning
Q_a, Q_b, Q_c, σ_h, σ_Q, $\sigma_{\bar{s}}$	kW, K, kW, K	Test loads; the rows' spreads (Chapter 36).
y_c, σ_{ref}, C_T	K, K, R\$/yr	Calibration reading and reference spread; the thermal penalty (Chapter 36).

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).

Physics row, tolerance A balance of the plant graph written as a row with a small standard deviation; in the limit of zero tolerance the balance is exact and the state reduces to its free directions.

Reduced state The free directions left once the balances are exact, the columns of the null-space basis of the physics rows.

Profiling Holding some entries fixed and fitting the others by least squares, to see what the readings say about the fixed ones.

Predictive check Asking whether the readings are probable under the model fitted to them, by the misfit $2\mathcal{L}(\hat{\mathbf{z}})$ against its chi-square law or by a held-out reading against its predictive.

Singular value, weakest direction The scale at which the rows see a direction of the state; the combination of entries with the smallest singular value, which the evidence determines least.

Observability Whether the readings determine an entry at all: an unobservable direction is one the design columns do not see, and only a prior gives it a finite spread.

Condition number The ratio of the largest singular value to the smallest; it bounds the loss of digits in solving with the information matrix.

Virtual sensor The posterior of a quantity no instrument reads, pushed from the state through its function.

Leave-one-out (held-out) check Refitting without a reading and comparing the reading with the predictive of the remaining rows.

Redundancy number The share of a reading the other rows check, $1 - K_{kk}$; zero for a critical measurement, which no other row checks.

Minimum detectable bias The bias on a reading that the residual test flags with a stated probability.

Footprint The pattern a candidate fault leaves on the readings: its column of the Jacobian on the reading rows.

Fault detection and isolation, model-based diagnosis Deciding that the plant has departed from its healthy model, and which component did, by comparing the evidence of candidate drawings generated from the plant graph.

Occam factor The part of a model's evidence that penalizes the prior range its parameters were given and the readings did not use.

Savage–Dickey ratio The Bayes factor of a point value of one entry against the model that frees it, read as the ratio of the posterior to the prior density at that value.

Expected evidence gain The expected log Bayes factor a measurement not yet taken would add between two drawings, the Kullback–Leibler divergence of their predictives.

Value at risk, tail mean A quantile of the predictive loss of a year, and the mean of the loss beyond it.

Importance measures Birnbaum (the slope of the top event in a unit's unavailability), Fussell–Vesely (the share a perfect unit would remove), risk achievement and reduction worths (the factors of an outage and of a perfect unit); each a derived quantity of the state, with a band.

Pivotal decomposition The top event as a straight line in one unit's unavailability, between the unit pinned down and pinned up.

Uncertainty importance A shared entry's term in the variance of the top event: what its record leaves uncertain, as against what its unit carries.

Operating characteristic The probability that a test accepts, as a function of the true value; its complement at a good value is the producer's risk, its value at a bad one the consumer's risk.

Acceptance sampling by variables Accepting on a measured quantity read through a model, rather than on a count of failures (by attributes).

Operating margin, capacity The room between a predicted quantity and its limit at a load; the load at which the limit is reached. Their bands are posteriors, and the load held with an assurance is a tail inverted.

Derating A unit delivering less than its plate rating, estimated as a capacity factor against the plate as a point null.

Delay-time model A failure preceded by a detectable degraded state of random duration; an inspection inside that delay catches the failure.

Age replacement Replacing a unit at a fixed age or at failure, whichever comes first, priced per unit time by the renewal-reward argument.

Plug-in action, Bayes action, regret The action best at the fitted state; the action best in expectation over the posterior; the expected cost of taking the first instead of the second.

Expected value of perfect information The expected cost of the best action now less the expected cost if the state were known before acting: the ceiling on what any record can be worth to that decision.

Expected value of sample information The same gain from one planned record, averaged over its outcomes.

Critical fractile, newsvendor The quantile at which a stocking or admission decision balances its two prices.

Lower-tail mean, overshoot The mean of a capacity below its assurance quantile; the load admitted above the capacity.

Marginal saving The expected cost one more unit of an action removes; a unit is added while it exceeds the unit's price.

Benefit–cost ratio, residual risk The expected loss a change removes over its annual cost; the risk that remains for the next change to act on.

Availability target A required unavailability held with stated assurance, a decision rule that can disagree with the cost minimum.

Hierarchical model, partial pooling, shrinkage Units drawn from a population with an unknown spread; each estimate is pulled toward the population by an amount the record sets.

Funnel The shape of a hierarchical posterior, in which the units' deviations narrow as the population spread falls to zero; a normal at the mode cannot hold it.

Centered and non-centered parametrization The units' deviations written directly, or as standardized deviations multiplied by the spread; the second removes the funnel from a chain's path.

Markov chain Monte Carlo, Metropolis, Gibbs Draws from a posterior by a chain of proposed moves accepted so that the posterior is the chain's stationary distribution.

$\hat{R}$, effective sample size of a chain The between-chain against within-chain spread; the number of independent draws a correlated chain is worth.

Common-cause basic event One event in the fault tree standing for a cause that takes out several units together.

One repair job, separate repair jobs A shared outage repaired as one piece of work, or one piece of work for each unit it took out.

Ride-through, thermal mass A hall's holding its limit through an outage; the stored heat capacity that sets how long.

Sparse information matrix, elimination order, fill-in, Cholesky factor The inverse covariance of a time-indexed state, mostly zeros; the order in which entries are solved out; the new nonzeros that order creates; the triangular factor they land in.

Kalman filter, smoother, augmented state The forward and backward passes of the elimination in time order; a state extended by the constant entries the steps share.

Commissioning test, derated load, redundancy of a row A test a new plant must pass before acceptance; the load accepted below design; the share of a row's information other rows also supply.

Part VIII

Information and Engineering Decisions

Chapter 29

Inspection and Maintenance Intervals

29.1 How often to test a battery string

The twelve battery strings of Chapter 10 are given a capacity test once a year, at R$ 2 000 a string, and the facility engineer has been asked whether once a year is right. A string found degraded at a test is replaced in a planned way for R$ 15 000; a string that fails in service costs R$ 60 000, the replacement and the loss event together. The manufacturer states that a string shows a measurable loss of capacity about half a year before it fails. The quickest reading takes the fit of Chapter 10 as known, a characteristic life of 6.50 years and a shape of 4.73, hence a mean life of 5.95 years, takes the half year as the delay of every string, and tests at that delay, so that no failure is missed: R$ 6 523 per string-year at a test every half year against R$ 8 307 at the present annual test, a saving of R$ 21 400 a year over the twelve strings.

That saving rests on a sharp corner in the cost at exactly half a year, and the corner exists only because one number was given to every string. The manufacturer's half year is the middle of a spread, no string at this site has been timed through its degraded state, and the mean life is one point of a posterior that four failures leave wide. The question of this chapter is

what an inspection interval costs in expectation over everything the record and the manufacturer leave uncertain, which interval minimizes that cost, how firm the choice is, and which unknown it depends on.

What changes from Chapter 10 is the kind of answer: it is an action, chosen by minimizing a posterior expectation. The state gains a third entry, the delay, which carries a prior row and no record row.

Jargon: delay-time model, degraded state, delay, inspection interval, catch probability, cost rate, renewal–reward ratio, partial mean, expected value of perfect information, plug-in

The *delay-time model* says that a failure is preceded by a *degraded state* which a test can detect, entered a random time before the failure; that time is the *delay*. The *inspection interval* is the time between tests of one string. The *catch probability* is the probability that a test falls inside the degraded state, so that the failure is replaced in a planned way. The *cost rate* is the expected cost per string-year; the *renewal–reward ratio* obtains it as the expected cost of one life over the expected length of one life. The *partial mean* of the delay at an interval is the expectation of the delay counted only when it is shorter than the interval. The *expected value of perfect information* on an unknown is the reduction of the expected loss of the best action that would follow from learning that unknown before acting. A *plug-in* value evaluates a quantity at a point estimate of the state.

Key idea

An interval is an action, and its loss is a derived quantity of the state. The expected loss is linear in the catch probability and in the replacement rate, so the two unknowns average apart: the rate over the lifetime posterior, the catch probability over the delay. The expected cost cannot tell the spread of the delay across strings from the spread of what is known about its median, because only their sum enters it; the band on a catch probability and the value of learning the delay depend on the split. The minimizer is not the average of the minimizers.

29.1.1 The delay-time model

A string has a lifetime T, the age at which it would fail in service. It enters its degraded state at $T - D$, where $D > 0$ is its delay. (The symbol D names the delay throughout this chapter; the total repair duration that the notation also writes D does not appear here.) The tests fall on a grid of period τ whose phase relative to the string's failure is taken as uniform on $(0, \tau)$. A test inside the interval $(T - D, T)$ finds the string degraded, and it is replaced; if no test falls inside, the string fails in service.

Figure 29.1 draws the two cases. The horizontal line in each row is the age of one string, the gold segment its degraded state, the red cross the failure it would have and the ticks the tests. In the upper row the delay is longer than the interval's remainder and a tick falls inside the gold segment, so the failure is caught; in the lower row the same interval misses a shorter delay.

The catch probability for a known delay. Measure the time back from the failure to the last test before it, B. With a uniform phase, B is uniform on $(0, \tau)$, and the failure is caught when $B < D$:

$$\Pr(\text{caught} \mid D) = \Pr(B < D) = \min\left(\frac{D}{\tau},\, 1\right). \tag{29.1}$$

A delay at least as long as the interval is always caught; a shorter one is caught in the proportion of the interval it covers.

The cost of one life. A string's life has three costs: the tests, a planned replacement if it is caught, and the failure if it is not. The number of grid points in a life of length T has the mean T/τ under the uniform phase, and with T independent of D and of the phase the expected cost of one life is $C_{\mathrm{i}}\mathrm{MTTF}/\tau + p_{\mathrm{c}}C_{\mathrm{r}} + (1 - p_{\mathrm{c}})C_{\mathrm{f}}$, with p_{c} the catch probability averaged over the delay. Dividing by the expected length of the life, the renewal–reward ratio, gives the cost rate.

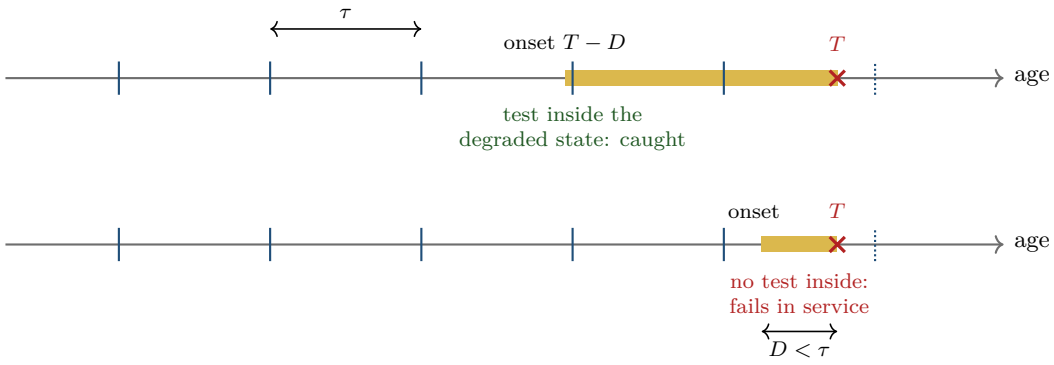


Figure 29.1: The delay-time model. Upper row: a delay longer than the distance to the next test, caught at that test (the dotted tick is not reached). Lower row: a shorter delay falling between two tests, and the failure occurs in service. Blue ticks: the tests, period τ; gold: the degraded state; red cross: the failure. The lengths are illustrative.

With the replacement rate $\nu_{\mathrm{r}} = 1/\mathrm{MTTF}$ and $\Delta C = C_{\mathrm{f}} - C_{\mathrm{r}}$,

$$\text{cost rate} = \frac{C_{\mathrm{i}}}{\tau} + \nu_{\mathrm{r}}(C_{\mathrm{f}} - \Delta C\, p_{\mathrm{c}}). \tag{29.2}$$

The mean life enters only through the failure and replacement term; after the division the inspection term no longer contains it.

29.1.2 The strings and their plant graph (step 0)

The plant is the set of strings of Chapter 10, with the record of its Table 10.1: four failed tests and eight strings in service, 51.9 string-years on test. Its graph, Figure 10.1, has one free node, the string design, with the log-scale $u = \ln \eta$ and the log-shape $v = \ln \beta$. Figure 29.2 adds to it what the question needs. The green box at the left is that node, with the twelve evidence rows reached by a dashed arrow. The second green box is new: the degraded state of the design, carrying the log-median delay d, reached only by the dashed arrow of the manufacturer's figure. The gray boxes are pinned: the spread of the delay across strings and the three costs. The blue box is the action, the interval τ, and the box at the right the loss, computed from the two free nodes, the pinned costs and the action. No row connects the record to the delay.

The record and the delay. The tests of Chapter 10 are inspections, yet they say nothing about the delay. Each row of its record is an age and a failed or passed test; no string was seen degraded and then followed to its failure, and none is recorded as having failed in service between tests. The delay is unobserved: the third column of the Jacobian is zero on every record row.

The costs and the delay figure. Table 29.1 lists the pinned quantities and the manufacturer's figure. Each row is a quantity; the columns give its symbol, value and status. The figure of a median of 0.5 yr with a log-standard deviation of 0.5 is one statement about the delay of a string not yet seen; step 2 separates what in it varies across strings from what is not known.

29.1.3 The probabilistic questions (step 1)

The engineer asks six questions, and step 1 writes each over the state $\mathbf{z} = (u, v, d)^{\mathsf{T}}$, with the unknown on the left of the bar and the evidence on the right. Table 29.2 lists them. The first column is the question in engineering language; the second the general form of its category

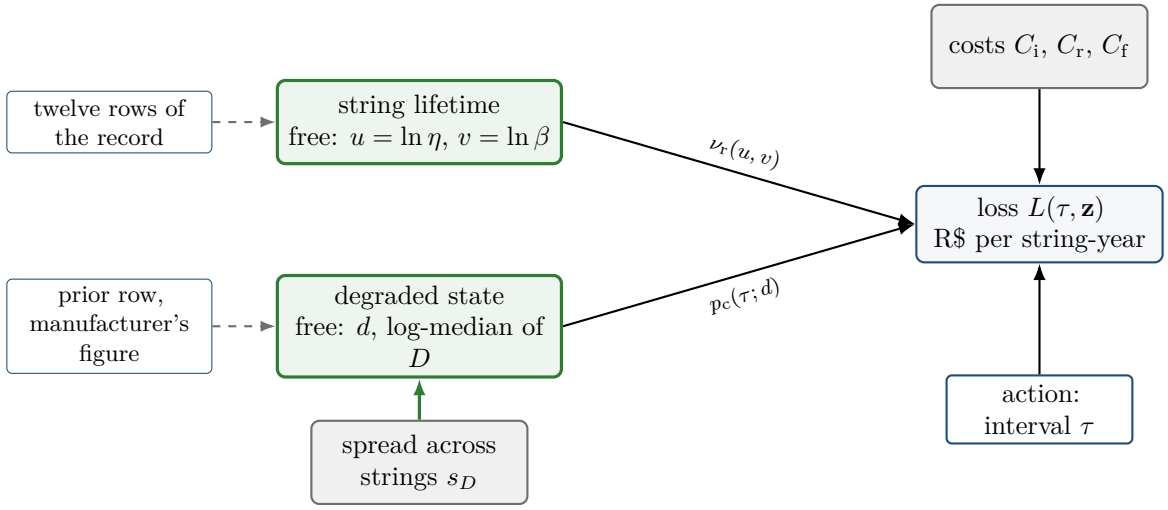


Figure 29.2: The plant graph of the inspection question. Free: the two nodes in green, the lifetime with u and v, the degraded state with d. Pinned: the spread s_D of the delay across strings and the costs (gray). Observed: the twelve rows of the record on the lifetime, and only the manufacturer's prior row on the delay (dashed). The interval τ is an action, not an entry of the state.

Table 29.1: The costs and the delay figure of the strings (illustrative).

Quantity	Symbol	Value	Status
Capacity test of one string	C_i	R$ 2 000	pinned
Planned replacement of a degraded string	C_r	R$ 15 000	pinned
Failure in service, with the loss event	C_f	R$ 60 000	pinned
Delay, median, manufacturer's figure	e^{d_0}	0.5 yr	prior center
Delay, total log-standard deviation	s_tot	0.5	split in step 2

(Chapter 2), with $L(a, \mathbf{z})$ the loss of action a, $f(\mathbf{z})$ a derived quantity and $\mathbf{g}$ its gradient; the third its form for the strings; the fourth its category; the last the subsection that answers it. The interval $\tau^*(\mathbf{z})$ is the interval that would be best if the state were known; τ^* is the interval that is best under the posterior.

Table 29.2: The questions asked about the inspection interval, the general form of their category, their 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) and the prior row on d.

Engineering question	General form	Mathematical form	Category	Answered in
What does an inspection interval cost per string-year, in expectation?	$\mathcal{R}(a) = \int L(a, \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	$\mathcal{R}(\tau) = \mathbb{E}[L(\tau, \mathbf{z}) \mid \mathbf{y}]$	design (expected loss)	29.3.1
How often should the strings be tested?	$a^* = \arg\min_a \mathcal{R}(a)$	$\tau^* = \arg\min_\tau \mathcal{R}(\tau)$	design (the decision)	29.3.2
How sure is the choice?	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$, quantiles	$p(\tau^*(\mathbf{z}) \mid \mathbf{y})$; the band of $L(\tau, \mathbf{z})$; $\Pr(L(a, \mathbf{z}) < L(a', \mathbf{z}) \mid \mathbf{y})$	the posterior of the decision	29.3.3
Which unknown decides the interval?	$\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$, term by term	the terms of $L(\tau^*, \mathbf{z})$ and of $\ln \tau^*(\mathbf{z})$ in (u, v) against d	attribution	29.3.4
What is the probability of catching a failure under a given interval?	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$, quantiles	$\mathbb{E}[\min(D/\tau, 1) \mid \mathbf{y}]$; the band of $p_\mathrm{c}(\tau; d)$	prediction	29.3.5
Would learning the delay change the interval?	$\bar{L}_\mathrm{now} - \bar{L}_\mathrm{after}$	$\min_\tau \mathcal{R}(\tau) - \mathbb{E}_d[\min_\tau \mathcal{R}(\tau \mid d)]$	design (value of information)	29.3.6

29.1.4 The evidence as rows (step 2)

Step 2 keeps the twelve lifetime rows of Chapter 10 and adds one row, the prior row on the delay. Which row acts on which entry fixes where the Jacobian can be nonzero, marked $*$:

$$\mathbf{J} = \begin{bmatrix} * & * & 0 \\ \vdots & \vdots & \vdots \\ * & * & 0 \\ 0 & 0 & * \end{bmatrix} \begin{matrix} \text{string A, failure} \\ \vdots \\ \text{string L, survivor} \\ \text{prior row on } d \end{matrix}, \qquad \text{columns } u, v, d. \tag{29.3}$$

The two blocks share no column, so the covariance is block diagonal, and the record cannot move d: in the language of observability, d is an unobservable direction of the record, and only the prior gives it a finite spread. No row is a physics row, so no tolerance σ_p arises.

The modeling choices. Five choices make the rows and the loss. The lifetime rows are those of Chapter 10, a failed test read as a failure at the test age: this reuses its posterior unchanged, and it mixes two events this chapter separates, since a string that fails its capacity test may be a degraded string caught, which Section 29.4.4 measures by moving the failure ages. The lifetime and the delay are independent given the state: this lets the expected cost factor and the two unknowns average apart, and it fails if a string that wears out early also degrades faster (Section 29.6.5). The delay follows a lognormal distribution across strings: this keeps it positive and gives the catch probability a closed form, and the tail below the interval is the part that matters, which only a record of delays could test. The phase of the test grid is uniform: this gives equation (29.1), and it would fail only for a failure mechanism locked to the

calendar. A test finds every degraded string: this makes the catch a matter of timing alone, and Exercise 29.5 relaxes it.

The aleatory and the epistemic spread of the delay. Two different things are spread in the manufacturer's figure. Strings differ: given the distribution of the delay, one string's D differs from another's, and no measurement of the fleet would remove that variation. That spread is aleatory, written with the log-standard deviation s_D:

$$\ln D \mid d \sim \mathcal{N}(d,\ s_D^2). \tag{29.4}$$

The distribution itself is not known: the log-median d has one true value for this design in this room, and the figure gives it only to within some range. That spread is epistemic, and it is what the prior row carries:

$$r_\pi(d) = \frac{d - d_0}{\sigma_d}, \qquad d_0 = \ln 0.5 = -0.69315. \tag{29.5}$$

The state holds d, the epistemic quantity. It does not hold the delay of any one string: those are aleatory, and the loss integrates them out for each value of d, as it integrates out the lifetime of each string for each value of (u, v). The spread s_D is pinned, a second parameter of the distribution that could be a further entry; its effect is carried by Table 29.11 instead.

Averaging equation (29.4) over the prior row gives the distribution of the delay of a string not yet seen, a sum of two normals in $\ln D$:

$$\ln D \sim \mathcal{N}(d_0,\ s_{\text{tot}}^2), \qquad s_{\text{tot}}^2 = s_D^2 + \sigma_d^2. \tag{29.6}$$

The manufacturer's single log-standard deviation of 0.5 is read as s_{tot}, the spread of a delay about the stated median with both causes in it: that is all a single figure can state. The split is a modeling choice that the figure does not determine. This chapter takes $s_D = 0.4$ and $\sigma_d = 0.3$, so that $0.4^2 + 0.3^2 = 0.5^2$: the median is known to within a factor of $\mathrm{e}^{1.96\times 0.3} = 1.8$, and strings differ from the median by up to a factor of $\mathrm{e}^{1.96\times 0.4} = 2.2$. Table 29.11 carries the other splits, from all aleatory to all epistemic.

The loss. The loss of the interval τ at the state $\mathbf{z}$ is the cost rate (29.2) given the state: the expectation over the lifetime, the delay and the phase of one string, all aleatory, with the epistemic entries held. Its two quantities of the state are the replacement rate, from the Weibull mean of Section 7.1.3,

$$\nu_{\text{r}}(u, v) = \frac{1}{\text{MTTF}(u, v)} = \frac{\mathrm{e}^{-u}}{\Gamma(1 + \mathrm{e}^{-v})}, \tag{29.7}$$

and the catch probability given the log-median, the expectation of equation (29.1) over (29.4),

$$p_{\text{c}}(\tau; d) = \mathbb{E}\Big[\min\Big(\frac{D}{\tau}, 1\Big) \,\Big|\, d\Big] = \frac{\mathcal{M}(\tau; d)}{\tau} + \Pr(D \geq \tau \mid d), \qquad \mathcal{M}(\tau; d) = \mathbb{E}\big[D\,\mathbf{1}[D < \tau] \mid d\big]. \tag{29.8}$$

The partial mean $\mathcal{M}$ comes by completing the square in the lognormal integral: with $x = \ln D$, the factor e^x times the normal density of mean d is $\mathrm{e}^{d + s_D^2/2}$ times the normal density of mean $d + s_D^2$, so

$$\mathcal{M}(\tau; d) = \mathrm{e}^{d + s_D^2/2}\, \Phi\left(\frac{\ln \tau - d - s_D^2}{s_D}\right), \qquad \Pr(D \geq \tau \mid d) = 1 - \Phi\left(\frac{\ln \tau - d}{s_D}\right). \tag{29.9}$$

The loss, with its gradient in the state,

$$L(\tau, \mathbf{z}) = \frac{C_{\text{i}}}{\tau} + \nu_{\text{r}}(u, v)\big[C_{\text{f}} - \Delta C\, p_{\text{c}}(\tau; d)\big], \qquad \mathbf{g}_\tau = \begin{bmatrix} -\nu_{\text{r}}\,[C_{\text{f}} - \Delta C\, p_{\text{c}}] \\ \nu_{\text{r}}\, \psi(1 + \mathrm{e}^{-v})\, \mathrm{e}^{-v}\,[C_{\text{f}} - \Delta C\, p_{\text{c}}] \\ -\nu_{\text{r}}\, \Delta C\, \mathcal{M}(\tau; d)/\tau \end{bmatrix}, \tag{29.10}$$

where ψ is the digamma function and the third entry uses $\partial p_{\mathrm{c}}/\partial d = \mathcal{M}/\tau$: moving the median shifts the density across τ, and the two boundary terms of equation (29.8) cancel. This L is the form $c_a + \sum_e C_e \Pr(e \mid \mathbf{z}, a)$ of equation (2.23), per string-year: the certain cost $c_\tau = C_{\mathrm{i}}/\tau$, and two events, a planned replacement and a failure in service, at the rates $\nu_{\mathrm{r}} p_{\mathrm{c}}$ and $\nu_{\mathrm{r}}(1 - p_{\mathrm{c}})$.

29.2 The posterior of the state (step 3)

All six questions are read from one posterior, that of the three entries, which the two blocks of rows of step 2 give separately. The lifetime block is the posterior of Chapter 10, reused; the delay block is its prior row.

29.2.1 The posterior from the thirteen rows

The rows are the twelve working rows of equation (10.6) and the prior row (29.5),

$$\mathbf{r}(\mathbf{z}) = \begin{bmatrix} \mathbf{r}_{\mathrm{rec}}(u, v) \\ (d - d_0)/\sigma_d \end{bmatrix}, \qquad \mathbf{J} = \begin{bmatrix} \mathbf{J}_{\mathrm{rec}}(u, v) & \mathbf{0} \\ \mathbf{0}^{\mathsf{T}} & 1/\sigma_d \end{bmatrix}, \tag{29.11}$$

a 13×3 Jacobian whose last line is the prior rows' $\mathbf{J}_\pi$. The objective is the sum of the twelve lifetime terms and $\frac{1}{2} r_\pi^2$. Its Hessian is block diagonal: the block of (u, v) is the information of Section 10.2.1, $\mathbf{J}_{\mathrm{rec}}^{\mathsf{T}} \mathbf{J}_{\mathrm{rec}}$ with its named curvature term, and the block of d is $1/\sigma_d^2$, with no curvature term because the row is linear. With unnormalized rows the information would be $\mathbf{J}^{\mathsf{T}} \mathbf{W} \mathbf{J}$. The Newton step of the lifetime block was taken in Table 10.3; the step of d from any start lands on d_0 in one iteration, since its row is linear and alone in its column.

29.2.2 The exact posterior

The lifetime block has the exact route of Section 10.2.2: the marginal $p(v \mid \mathbf{y})$, and $\zeta = \eta^{-\beta}$ given the shape distributed as $\mathrm{Gamma}(k, A(\beta))$. The delay block is linear and Gaussian, so its Laplace posterior *is* its exact posterior, and

$$p(u, v, d \mid \mathbf{y}) = p(u, v \mid \mathbf{y})\, \mathcal{N}(d;\ d_0,\ \sigma_d^2). \tag{29.12}$$

There are therefore two routes for the lifetime and one for the delay. *The state route* takes (u, v, d) as $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$; *the exact route* takes the first factor of equation (29.12) in place of its Gaussian. On a quantity of d alone the two routes are the same posterior.

29.2.3 The posterior of the strings and the delay

The MAP state and the covariance. The lifetime entries and their block are the committed values of equation (10.8), and $\sigma_d = 0.3$ gives the information $1/0.09 = 11.111$ in the third entry:

$$\hat{\mathbf{z}} = \begin{bmatrix} 1.87138 \\ 1.55303 \\ -0.69315 \end{bmatrix}, \quad \boldsymbol{\Lambda} = \begin{bmatrix} 89.331 & 9.009 & 0 \\ 9.009 & 7.491 & 0 \\ 0 & 0 & 11.111 \end{bmatrix}, \quad \boldsymbol{\Sigma} = \begin{bmatrix} 0.01274 & -0.01532 & 0 \\ -0.01532 & 0.15192 & 0 \\ 0 & 0 & 0.09 \end{bmatrix}, \tag{29.13}$$

that is $\hat{\eta} = 6.4973\,\mathrm{yr}$, $\hat{\beta} = 4.7257$ and a median delay of $0.5\,\mathrm{yr}$. The 6.50 and 4.73 of the first reading are the MAP state rounded, and every number below uses the unrounded values. The axes of the lifetime block are $\mathbf{e}_1 = (0.108,\ -0.994)^{\mathsf{T}}$ with $\sqrt{\epsilon_1} = 0.392$ and $\mathbf{e}_2 = (0.994,\ 0.108)^{\mathsf{T}}$ with $\sqrt{\epsilon_2} = 0.105$; the third axis is d itself, with $\sqrt{\Sigma_{dd}} = 0.3$.

The exact posterior. Chapter 10 reports it: the shape has the median 3.64 and the band [1.15, 7.77], against 4.73 and [2.20, 10.14] on the Gaussian, and the scale the median 6.93 yr and the band [5.49, 15.5] yr, against 6.50 yr and [5.21, 8.11] yr (Table 10.15). Bands and tails on both routes below use 2×10^5 draws, of $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ and of equation (29.12) (numpy `default_rng`, seed 28; v by inverting its distribution function on the grid of Section 10.2.3, ζ from its gamma given v).

29.3 The answers (step 4)

The questions of Table 29.2 are answered from $(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of equation (29.13), by a Gauss–Hermite rule, by draws of $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ or by the delta method, and that is the value reported. The exact route of equation (29.12) gives the value beside it, with the gap.

29.3.1 The expected cost of an interval

Engineering question. What does an inspection interval cost per string-year, in expectation?

General form. The expected consequence of an action, equation (2.23) of Chapter 2:

$$\mathcal{R}(a) = \mathbb{E}\big[L(a, \mathbf{z}) \mid \mathbf{y}\big] = \int L(a, \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}. \tag{29.14}$$

Specialized form. The question is equation (29.14) with $a = \tau$, $\mathbf{z} = (u, v, d)^{\mathsf{T}}$, $L(a, \mathbf{z})$ the loss (29.10) and $p(\mathbf{z} \mid \mathbf{y})$ the posterior $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of equation (29.13). One choice enters here: the loss is an expectation per string-year over one long-run life, which is what lets the renewal–reward ratio stand for the cost, and it leaves out that a caught string's life ends early, which Section 29.4.1 measures by simulation. The loss is linear in p_{c}, and $\boldsymbol{\Sigma}$ is block diagonal, so $\nu_{\mathrm{r}}(u, v)$ and $p_{\mathrm{c}}(\tau; d)$ are independent under the posterior and the integral factors:

$$\mathcal{R}(\tau) = \frac{C_{\mathrm{i}}}{\tau} + \bar{\nu}_{\mathrm{r}}\big[C_{\mathrm{f}} - \Delta C\, \bar{p}_{\mathrm{c}}(\tau)\big], \quad \bar{\nu}_{\mathrm{r}} = \mathbb{E}[\nu_{\mathrm{r}}(u, v) \mid \mathbf{y}], \quad \bar{p}_{\mathrm{c}}(\tau) = \mathbb{E}[p_{\mathrm{c}}(\tau; d) \mid \mathbf{y}]. \tag{29.15}$$

The replacement rate is averaged and then its reciprocal is not taken: $1/\bar{\nu}_{\mathrm{r}}$ is not the expected life. The expected life itself is infinite on both routes, since a shallow shape makes $\Gamma(1 + 1/\beta)$ grow without bound faster than either posterior's tail falls. The rate has none of this trouble.

On the state route $\bar{\nu}_{\mathrm{r}}$ is the two-dimensional rule on the axes of the lifetime block, as in equation (10.14), and $\bar{p}_{\mathrm{c}}$ the one-dimensional rule on d:

$$\bar{\nu}_{\mathrm{r}} \approx \sum_m \sum_n \varpi_m \varpi_n\, \nu_{\mathrm{r}}\big(\hat{\mathbf{z}}_{uv} + \sqrt{\epsilon_1}\, x_m \mathbf{e}_1 + \sqrt{\epsilon_2}\, x_n \mathbf{e}_2\big), \qquad \bar{p}_{\mathrm{c}}(\tau) \approx \sum_m \varpi_m\, p_{\mathrm{c}}(\tau;\ d_0 + \sigma_d x_m), \tag{29.16}$$

with $\hat{\mathbf{z}}_{uv}$ the first two entries of $\hat{\mathbf{z}}$. The exact route has one integral for the rate: since $\mathrm{e}^{-u} = \zeta^{1/\beta}$, the moment of the gamma posterior of ζ given the shape is $\Gamma(k + 1/\beta)/(\Gamma(k) A^{1/\beta})$, and

$$\bar{\nu}_{\mathrm{r}}^{\mathrm{exact}} = \int p(v \mid \mathbf{y})\, \frac{\Gamma(k + 1/\beta)}{\Gamma(k)\, A(\beta)^{1/\beta}\, \Gamma(1 + 1/\beta)}\, \mathrm{d}v. \tag{29.17}$$

For $\bar{p}_{\mathrm{c}}$ the exact value needs no quadrature. By equation (29.6) the delay of a string, averaged over d, is lognormal with s_{tot}, so $\bar{p}_{\mathrm{c}}(\tau)$ is equation (29.8) with s_{tot} in place of s_D and $d = d_0$. It is the same posterior as the rule's: the closed form has no separate route here, it checks the quadrature. And it shows at once that $\mathcal{R}(\tau)$ *depends on* s_D *and* σ_d *only through* s_{tot}: the expected cost does not see the split.

Building blocks. $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$ of equation (29.13) with the axes of the lifetime block; the functions (29.7) and (29.8); the five nodes of Table 29.3, and sixty for the reported values; for the values beside them, the marginal $p(v \mid \mathbf{y})$ and $A(\beta)$ of Section 10.2.2.

The rule of five nodes. Table 29.3 evaluates equation (29.16). Its upper block gives ν_{r} at the twenty-five states: the rows are the nodes along the long axis, with the shape there, the columns the nodes along the short axis. The lower block gives the catch probability at five values of the median, at two intervals. Down a column of the upper block the shape falls from 14.38 to 1.55 and the rate changes by a sixth; across a row the scale moves by about a third either way and the rate with it: the mean life hardly depends on the shape.

Table 29.3: The Gauss–Hermite rule (29.16) with five nodes. Upper block: ν_{r} [yr^{-1}] at $(u, v)^{\mathsf{T}} = \hat{\mathbf{z}}_{uv} + 0.392\, x_m \mathbf{e}_1 + 0.105\, x_n \mathbf{e}_2$. Lower block: the median delay and $p_{\mathrm{c}}(\tau; d_0 + 0.3\, x_m)$ at $\tau = 0.5$ and 1 yr. The weights are $\varpi = 0.01126$, 0.22208, 0.53333, 0.22208, 0.01126 on each axis.

x_m	β	$x_n = -2.857$	-1.356	0	1.356	2.857
-2.857	14.38	0.2432	0.2077	0.1801	0.1562	0.1335
-1.356	8.01	0.2337	0.1996	0.1731	0.1501	0.1282
0	4.73	0.2272	0.1940	0.1682	0.1458	0.1245
1.356	2.79	0.2203	0.1882	0.1632	0.1416	0.1209
2.857	1.55	0.2037	0.1745	0.1516	0.1318	0.1128
Weighted sum of the twenty-five cells: $\bar{\nu}_{\mathrm{r}} = 0.16905\,\mathrm{yr}^{-1}$						
x_m		-2.857	-1.356	0	1.356	2.857
Median [yr]		0.212	0.333	0.500	0.751	1.178
$\tau = 0.5$ yr		0.4571	0.6821	0.8733	0.9727	0.9980
$\tau = 1$ yr		0.2299	0.3602	0.5338	0.7446	0.9258
Weighted sums: $\bar{p}_{\mathrm{c}}(0.5) = 0.8496$, $\bar{p}_{\mathrm{c}}(1) = 0.5430$						

The rate and one interval. The sum over the upper block is $\bar{\nu}_{\mathrm{r}} = 0.16905\,\mathrm{yr}^{-1}$, unchanged at sixty nodes, against the plug-in $\nu_{\mathrm{r}}(\hat{\mathbf{z}}) = 1/5.946 = 0.16818\,\mathrm{yr}^{-1}$. The exact integral (29.17) gives $0.15285\,\mathrm{yr}^{-1}$, a gap of 0.0162, or 10 %: the exact posterior keeps long-lived states, with a scale band reaching 15.5 yr, which the Gaussian does not reach. Of the two lower rows, the weighted sums are the closed forms with s_{tot} to all the digits shown. At half a year, equation (29.15) is

$$\mathcal{R}(0.5) = \frac{2\,000}{0.5} + 0.16905\,\big[60\,000 - 45\,000 \times 0.8496\big] = 4\,000 + 3\,680 = \mathrm{R\$}\,7\,680,$$

and on the exact route $4\,000 + 0.15285 \times 21\,768 = \mathrm{R\$}\,7\,327$, a gap of R$ 353. The second term splits into the two events, R$ 2 154 of planned replacement and R$ 1 525 of failure in service.

The expected cost by interval. Table 29.4 gives the cost at five intervals. Each row is an interval; the columns give the fixed delay of half a year at the MAP life; the delay averaged with s_{tot} at the MAP life; the loss at the MAP state, $L(\tau, \hat{\mathbf{z}})$, which holds d at its center and so uses only s_D; the reported $\mathcal{R}(\tau)$; the exact value; and the gap. The last row is the limit of no test, where only the failure term remains, $\bar{\nu}_{\mathrm{r}} C_{\mathrm{f}}$. Figure 29.3 draws the same curves against τ in its left panel, and in its right panel the three terms at half a year for the two delay models.

The dashed curve has its corner at $\tau = 0.5$ yr: below it every failure is caught and only the test cost falls with τ; above it the failures begin to be missed. The corner comes from a delay

Table 29.4: The cost rate by interval [R$/string-year]: fixed delay at the MAP life, equation (29.1) with $D = 0.5$ yr; delay averaged at the MAP life; the loss at $\hat{\mathbf{z}}$; the expected cost $\mathcal{R}(\tau)$ of equation (29.15) on the state and exact routes, and their gap.

τ [yr]	Fixed delay	Delay averaged	$L(\tau, \hat{\mathbf{z}})$	$\mathcal{R}(\tau)$, state	$\mathcal{R}(\tau)$, exact	Gap
0.25	10 523	10 641	10 567	10 655	10 401	254
0.5	6 523	7 661	7 482	7 680	7 327	353
0.75	7 712	7 621	7 590	7 647	7 170	477
1	8 307	7 981	8 051	8 012	7 436	576
2	9 199	8 950	9 041	8 991	8 226	765
no test	10 091	10 091	10 091	10 143	9 171	972

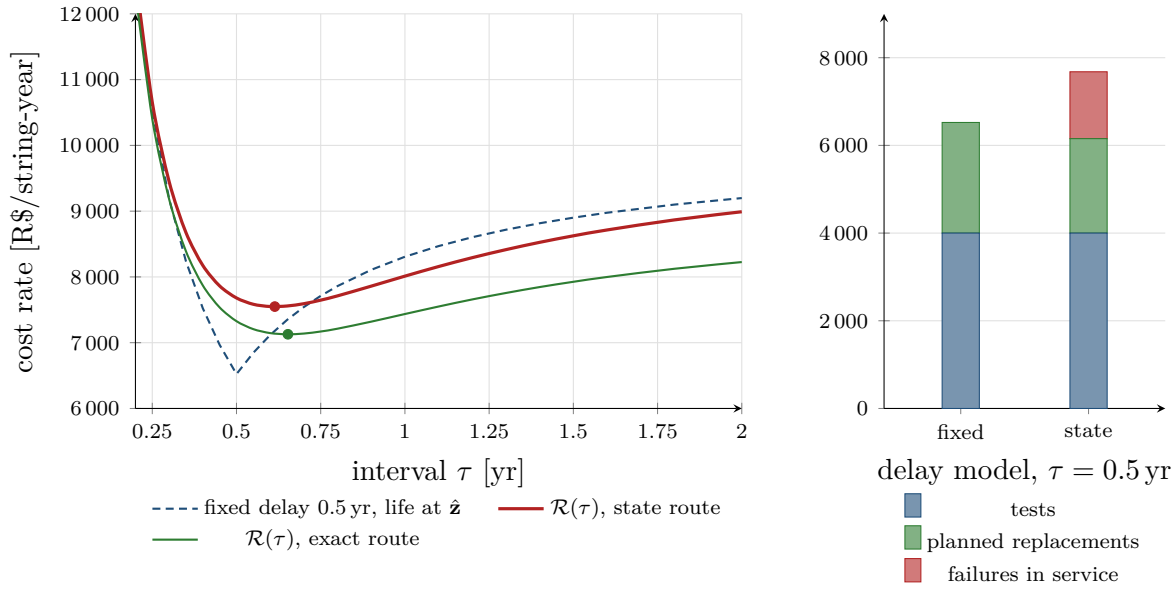


Figure 29.3: Left: the cost rate against the interval, for a fixed delay of 0.5 yr at the MAP life ($C_\mathrm{i}/\tau + \nu_\mathrm{r}(\hat{\mathbf{z}})[C_\mathrm{f} - \Delta C \min(0.5/\tau, 1)]$), and the expected cost $\mathcal{R}(\tau)$ of equation (29.15) with $s_\mathrm{tot} = 0.5$ on the two routes, their minima marked. Right: the three terms of the cost at $\tau = 0.5$ yr, under the fixed delay and on the state route.

with no spread. It is an artefact, and it makes the fixed-delay cost at half a year, R$ 6 523, lower than the expected cost by R$ 1 157, since the right panel's failure term is zero there: a delay shorter than the interval, which half the strings have, costs nothing in that drawing. The expected-cost curves are smooth, with a shallow minimum.

29.3.2 The best interval

Engineering question. How often should the strings be tested?

General form. The action that minimizes the expected consequence, equation (2.23):

$$a^{*} = \arg\min_{a} \mathcal{R}(a). \tag{29.18}$$

Specialized form. The question is equation (29.18) with $a = \tau > 0$ and $\mathcal{R}$ the function (29.15). One choice enters here: the interval is taken as a continuous variable applied to every string alike, which gives a closed form, and it fails if the calendar offers only a few intervals or the strings are not alike, which the comparison below and Exercise 29.6 take up. From equation (29.1), $\partial \min(D/\tau, 1)/\partial\tau = -(D/\tau^{2})\,\mathbf{1}[D < \tau]$, so $\partial \bar{p}_{\mathrm{c}}/\partial\tau = -\bar{\mathcal{M}}(\tau)/\tau^{2}$, with $\bar{\mathcal{M}}$ the partial mean (29.9) at s_{tot} and d_{0}. Then

$$\frac{\mathrm{d}\mathcal{R}}{\mathrm{d}\tau} = \frac{1}{\tau^{2}}\Big[\bar{\nu}_{\mathrm{r}}\,\Delta C\,\bar{\mathcal{M}}(\tau) - C_{\mathrm{i}}\Big]. \tag{29.19}$$

The bracket starts below zero, since $\bar{\mathcal{M}}(0) = 0$, and rises with τ because the partial mean does, so the minimum is unique where it vanishes:

$$\bar{\mathcal{M}}(\tau^{*}) = \frac{C_{\mathrm{i}}}{\Delta C\,\bar{\nu}_{\mathrm{r}}}. \tag{29.20}$$

A test at the optimum buys, in failures it catches, what it costs. The right side is a time, the mean life scaled by the ratio of the test cost to the saving from catching one failure. The left side rises to the whole mean delay $\mathbb{E}[D] = \mathrm{e}^{d_{0} + s_{\mathrm{tot}}^{2}/2}$ as τ grows, so a finite optimum exists only if the mean delay exceeds the right side; otherwise $\mathrm{d}\mathcal{R}/\mathrm{d}\tau < 0$ everywhere, and no test pays. With the partial mean (29.9) and the share π_{D} of the mean delay the right side asks for, the condition inverts in closed form:

$$\tau^{*} = \exp\!\big[d_{0} + s_{\mathrm{tot}}^{2} + s_{\mathrm{tot}}\,\Phi^{-1}(\pi_{D})\big], \qquad \pi_{D} = \frac{C_{\mathrm{i}}}{\Delta C\,\bar{\nu}_{\mathrm{r}}\,\mathbb{E}[D]} < 1. \tag{29.21}$$

The second derivative at the optimum,

$$\left.\frac{\mathrm{d}^{2}\mathcal{R}}{\mathrm{d}\tau^{2}}\right|_{\tau^{*}} = \frac{\bar{\nu}_{\mathrm{r}}\,\Delta C\,f_{D}(\tau^{*})}{\tau^{*}}, \tag{29.22}$$

with f_{D} the density of the delay, follows from $\mathrm{d}\bar{\mathcal{M}}/\mathrm{d}\tau = \tau f_{D}(\tau)$ and gives the flatness of the minimum.

Building blocks. $\bar{\nu}_{\mathrm{r}}$ of Section 29.3.1 on each route, s_{tot} and d_{0} of equation (29.6), and Φ^{-1}.

The optimum. The mean delay is $\mathrm{e}^{-0.69315 + 0.125} = 0.5666\,\mathrm{yr}$. On the state route the right side of equation (29.20) is $2\,000/(45\,000 \times 0.16905) = 0.2629\,\mathrm{yr}$, so $\pi_{D} = 0.4640$, $\Phi^{-1}(\pi_{D}) = -0.0903$ and

$$\tau^{*} = \exp[-0.69315 + 0.25 + 0.5 \times (-0.0903)] = \mathrm{e}^{-0.4883} = 0.614\,\mathrm{yr}, \qquad \mathcal{R}(\tau^{*}) = \mathrm{R\$}\,7\,549.$$

On the exact route the right side is 0.2908 yr, $\pi_D = 0.5132$, $\Phi^{-1}(\pi_D) = 0.0331$ and $\tau^* = \mathrm{e}^{-0.4266} = 0.653$ yr, with $\mathcal{R}(\tau^*) =$ R\$ 7 128. The gap is 0.039 yr, two weeks, in the interval and R\$ 421 in its cost. The two optima, about seven and a half and eight months, are on the same side of every interval the calendar would offer. Taking the other route's optimum costs almost nothing: the exact expected cost at 0.614 yr is R\$ 7 138, R\$ 10 above its minimum, and the state-route cost at 0.653 yr is R\$ 7 559, R\$ 10 above its own. The gap in the cost level does not change the decision.

The shallow minimum. By equation (29.22), with $f_D(0.614) = 1.195\,\mathrm{yr}^{-1}$, the curvature is 14 800 R\$ per string-year per yr^2, so a cost 1 % above the minimum, R\$ 75, lies $\sqrt{2 \times 75/14\,800} = 0.10$ yr away. The curve itself gives the range $[0.52,\ 0.73]$ yr on the state route and $[0.55,\ 0.79]$ yr on the exact one. Half a year is R\$ 131 above the minimum (R\$ 199 exact), nine months R\$ 98 (R\$ 42), and the present annual test R\$ 463 (R\$ 308). Over twelve strings, moving from the annual test to the optimum saves R\$ 5 556 a year on the state route and R\$ 3 696 on the exact one, a quarter and a sixth of the R\$ 21 400 of the first reading.

29.3.3 How firm the choice is

Engineering question. How sure is the choice?

General form. A band of a derived quantity, and the ranking of two actions at the true state, equations (2.11) and (2.24):

$$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y}) = \int \mathbf{1}[f(\mathbf{z}) \le f^*]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \quad \Pr(L(a, \mathbf{z}) < L(a', \mathbf{z}) \mid \mathbf{y}). \tag{29.23}$$

Specialized form. The question is equation (29.23) with $f(\mathbf{z}) = L(\tau, \mathbf{z})$ for the band at each interval, $f(\mathbf{z}) = \tau^*(\mathbf{z})$ for the distribution of the optimum, a and a' two intervals, and $p(\mathbf{z} \mid \mathbf{y})$ the posterior of equation (29.13). The optimum at a known state is the inversion of Section 29.3.2,

$$\tau^*(\mathbf{z}) = \exp[d + s_D^2 + s_D\, \Phi^{-1}(\pi_D(\mathbf{z}))], \qquad \pi_D(\mathbf{z}) = \frac{C_\mathrm{i}}{\Delta C\, \nu_\mathrm{r}(u, v)\, \mathrm{e}^{d + s_D^2/2}}, \tag{29.24}$$

and is unbounded at a state with $\pi_D(\mathbf{z}) \ge 1$, where no test pays. It uses s_D, not s_tot, because at a known state the median is known and only the aleatory spread remains. Hence $\tau^*(\hat{\mathbf{z}}) = 0.580$ yr is not the decision $\tau^* = 0.614$ yr, and neither is the mean of $\tau^*(\mathbf{z})$: this distribution says what the interval would be under each state, and the decision is one interval, the minimizer of the average. The bands are read from draws of the state, and the delta form of equation (2.24), with $\mathbf{g}_\tau$ of equation (29.10), is set beside the ranking.

Building blocks. The draws of Section 29.2.3; $\mathbf{g}_\tau$ and $\boldsymbol{\Sigma}$ of equation (29.13).

The bands and the rankings. Table 29.5 gives the band of the loss at five intervals and the probability that each costs less than the annual test at the true state. Each row is an interval; the columns give $\mathcal{R}(\tau)$ and the central 95 % band of $L(\tau, \mathbf{z})$ on each route, and the probability on each route. Figure 29.4 draws the state-route bands against τ. The horizontal axis is the interval and the vertical the cost rate. The light region is the band with the whole state uncertain, the dark region the band with d held at its center, so that only the lifetime posterior moves; the heavy curve is $\mathcal{R}(\tau)$ on the state route, the thin curve on the exact route. Two rugs at the bottom mark sixty draws of $\tau^*(\mathbf{z})$ on each route.

Table 29.5: The expected cost and the 95 % band of the loss $L(\tau, \mathbf{z})$ [R$/string-year], and the probability that the interval costs less than the annual test, on the state route and, in the last two columns, the exact route. Draws: 2×10^5, seed 28.

	State route			Exact route		
τ [yr]	$\mathcal{R}(\tau)$	Band	$\Pr(L < L_{\tau=1})$	$\mathcal{R}(\tau)$	Band	$\Pr(L < L_{\tau=1})$
0.25	10 655	[10 091, 11 450]	0.000	10 401	[9 060, 11 334]	0.000
0.5	7 680	[6 388, 9 903]	0.807	7 327	[5 363, 9 651]	0.636
0.614	7 549	[5 857, 10 019]	0.931	7 138	[4 820, 9 747]	0.799
0.75	7 647	[5 590, 10 187]	0.971	7 170	[4 471, 9 913]	0.874
1	8 012	[5 662, 10 488]	–	7 436	[4 231, 10 208]	–

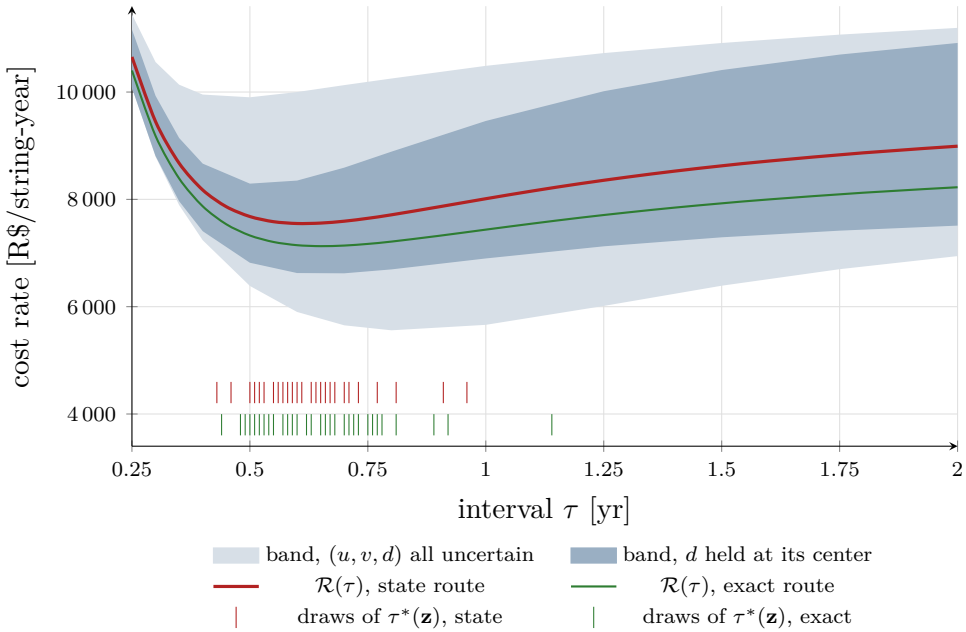


Figure 29.4: The expected cost against the interval with the central 95 % bands of $L(\tau, \mathbf{z})$ on the state route, 2×10^5 draws of $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$, seed 28. Rugs: sixty draws of $\tau^*(\mathbf{z})$, equation (29.24), rounded to 0.01 yr, coincident draws drawn once; on the exact route four of the sixty have no finite optimum and are not drawn.

The bands are wide: at the optimum the loss runs from R$ 5 857 to R$ 10 019, a spread of more than half the cost, and far wider than any difference between two intervals near the minimum. They do not make the choice uncertain in that proportion, since the two losses move together across the posterior. The optimum costs less than the annual test with probability 0.931 on the state route, 0.799 exactly; the delta form of equation (2.24) gives 0.988, because the loss is curved in d, and the draws are the value reported. Half a year beats the annual test with probability 0.807 (0.636). The comparison of nine months with six falls on opposite sides of one half: nine months is cheaper in expectation on both routes, by R$ 33 and R$ 157, but it is the cheaper at the true state with probability 0.452 on the state route and 0.599 on the exact one. The rule is the expected cost, so the two routes agree on the action; the probabilities say how little separates the two intervals.

The distribution of the optimum, equation (29.24), has on the state route the median 0.593 yr and the band [0.459, 0.884] yr, with 1.2 % of the states having no finite optimum; on the exact route the median is 0.632 yr, the lower end 0.472 yr, and 6.9 % of the states have no finite optimum, so the upper end of the band is not finite. Those are the long-lived states: a very long life makes the right side of equation (29.20) exceed the mean delay. The states with $\tau^*(\mathbf{z}) < 1$ yr have the probability 0.986 (0.912).

29.3.4 Which unknown decides the interval

Engineering question. Which unknown decides the interval?

General form. The variance budget of a derived quantity under the delta method (Chapter 4):

$$\mathrm{sd}(f \mid \mathbf{y})^2 = \mathbf{g}^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g} = \sum_i [\mathbf{g}]_i^2 \, \Sigma_{ii} + \sum_{i \neq j} [\mathbf{g}]_i [\mathbf{g}]_j \, \Sigma_{ij}. \tag{29.25}$$

Specialized form. The question is equation (29.25) with two derived quantities: $f = L(\tau^*, \mathbf{z})$, gradient $\mathbf{g}_\tau$ of equation (29.10), and $f = \ln \tau^*(\mathbf{z})$ of equation (29.24), with $\boldsymbol{\Sigma}$ of equation (29.13). Differentiating the logarithm of equation (29.24) with $a_D = s_D \, \pi_D / \phi(\Phi^{-1}(\pi_D))$, ϕ the standard normal density, gives

$$\mathbf{g}_{\ln \tau^*} = \begin{bmatrix} a_D \\ -a_D \, \psi(1 + \mathrm{e}^{-v}) \, \mathrm{e}^{-v} \\ 1 - a_D \end{bmatrix}, \tag{29.26}$$

since $\ln \pi_D$ falls one for one with d and rises one for one with $\ln \mathrm{MTTF}$. The life and the median delay share one unit of sensitivity: a_D is the life's, $1 - a_D$ the delay's. The life enters only through ν_r, and the gradient says what the rate's table said: its v entry is small. The exact route has no gradient at a point; beside the budget stand the variances by draws with one block held at $\hat{\mathbf{z}}$.

Building blocks. $\mathbf{g}_\tau$ and $\mathbf{g}_{\ln \tau^*}$ at $\hat{\mathbf{z}}$, $\boldsymbol{\Sigma}$, and the draws of Section 29.2.3.

The budget. At $\hat{\mathbf{z}}$, $\pi_D = 0.4879$ with $s_D = 0.4$, so $a_D = 0.4894$ and

$$\mathbf{g}_{\ln \tau^*} = \begin{bmatrix} 0.489 \\ 0.028 \\ 0.511 \end{bmatrix}, \qquad \mathbf{g}_{\tau = 0.614} = \begin{bmatrix} -4\,150 \\ -241 \\ -3\,638 \end{bmatrix}.$$

Table 29.6 sets out the terms. Each row is a part of the state; the first columns give the terms of the two variances, the next the value of the loss at the entry's center less and plus one posterior standard deviation, the last the variance by draws with only that block uncertain, on both routes.

Table 29.6: The sensitivity of the cost and of the optimum to the life and to the delay. Terms of $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ for $L(0.614, \mathbf{z})$ [(R$/string-year)^2] and for $\ln \tau^*(\mathbf{z})$ [–] at $\hat{\mathbf{z}}$; the loss one standard deviation either side of the center; and the variances by draws with only one block uncertain, state route with exact in parentheses, $\ln \tau^*$ on the finite draws.

Part of the state	L term	$\ln \tau^*$ term	L at ± 1 sd	Var(L), one block	Var($\ln \tau^*$), one block
u	219 393	0.0031	7 905, 6 966		
v	8 822	0.0001	7 499, 7 320		
u, v cross	−30 645	−0.0004			
Life, in all	197 570	0.0028		2.02×10^5 (6.41×10^5)	0.0029 (0.0164)
Delay, d	1 191 197	0.0235	8 580, 6 494	9.72×10^5 (9.78×10^5)	0.0194 (0.0197)
Share of the delay	86 %	89 %		83 % (60 %)	87 % (55 %)

The delay decides, on the state route: 86 % of the variance of the loss and 89 % of that of the optimum in the budget, 83 % and 87 % by draws, against a lifetime posterior whose variance is nearly all u. With d held the optimum runs over $[0.528,\ 0.651]$ yr, and with the life held over $[0.500,\ 0.855]$ yr. On the exact route the delay's share by draws is 60 % and 55 %, and with d held the band of the optimum has no finite upper end, since 3.9 % of the exact lifetime states call for no test: the exact posterior of the mean life is far wider than the Gaussian, with a band of $[4.97,\ 14.75]$ yr against $[4.82,\ 7.34]$ yr. The delay remains the larger part on both routes, by a margin the state route overstates. Under the fixed delay the optimum sits at the corner for every lifetime state with $\nu_{\mathrm{r}} \Delta C D > C_{\mathrm{i}}$, which says nothing about the life: the fixed-delay cost at half a year, over 4 000 draws (seed 28), has the mean 6 532 and the band [6 055, 7 096] on the state route, about ±8 %, and 6 297 and [5 069, 7 023] on the exact route; it moves with the life alone, and the interval does not move at all.

29.3.5 The probability of catching a failure

Engineering question. What is the probability of catching a failure under a given interval?

General form. A predictive expectation, with the band of a derived quantity, equations (2.13) and (2.11):

$$\Pr(y_{\mathrm{new}} \mid \mathbf{y}) = \int \Pr(y_{\mathrm{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}. \tag{29.27}$$

Specialized form. The question is equation (29.27) with y_{new} the event that the next failure is caught, $\Pr(y_{\mathrm{new}} \mid \mathbf{z}) = p_{\mathrm{c}}(\tau; d)$ of equation (29.8), and the posterior of equation (29.13), of which only the third entry acts: the gradient of p_{c} is $(0,\ 0,\ \mathcal{M}/\tau)^{\mathsf{T}}$, and the life does not enter. Three statements must be kept apart. The predictive $\bar{p}_{\mathrm{c}}(\tau)$ is the probability that one failure is caught, the sum (29.16). The epistemic band is the band of $p_{\mathrm{c}}(\tau; d)$, the long-run fraction of failures caught by this fleet, whose true value is one number not known; since p_{c} rises with d, its quantiles are those of d, $p_{\mathrm{c}}(\tau; d_0 \pm 1.96\, \sigma_d)$. The aleatory range is that of $\min(D/\tau, 1)$ over the central 95 % of strings at d_0, how the catch chance of one string differs from another's given its own delay. The predictive and the epistemic band are on one posterior, for which the state route and the exact route coincide, as Section 29.2.2 says: the closed form with s_{tot} is the value beside the sum, and the gap measures only the quadrature.

Building blocks. The nodes of Table 29.3; $\mathcal{M}$ of equation (29.9); d_0, σ_d and s_D.

The catch probabilities. Table 29.7 gives them at five intervals. Each row is an interval; the columns give the predictive by five nodes, the closed form, the value at d_0, the delta-method spread $(\mathcal{M}/\tau)\sigma_d$, the epistemic band and the aleatory range. At half a year the five-node sum is the weighted sum of Table 29.3, 0.8496. Figure 29.5 draws $p_c(\tau; d)$ against τ for three medians: the horizontal axis is the interval, the vertical the catch probability, and each curve one median delay, at $s_D = 0.4$. The curves are one curve shifted along $\ln \tau$, near one until the interval approaches the median, and then falling as the median over the interval.

Table 29.7: The probability of catching a failure. Predictive by five nodes and in closed form; $p_c(\tau; d_0)$; the delta spread; the 95 % epistemic band of $p_c(\tau; d)$ over the prior row; the range of $\min(D/\tau, 1)$ over the central 95 % of strings at d_0. All entries are dimensionless.

τ [yr]	Five nodes	Closed form	At d_0	sd($p_c \mid \mathbf{y}$)	Epistemic band	Aleatory range
0.25	0.9843	0.9843	0.9941	0.011	[0.909, 1.000]	[0.913, 1]
0.5	0.8496	0.8496	0.8733	0.112	[0.587, 0.989]	[0.457, 1]
0.614	0.7694	0.7694	0.7850	0.144	[0.486, 0.970]	[0.372, 1]
1	0.5430	0.5430	0.5338	0.148	[0.301, 0.831]	[0.228, 1]
2	0.2828	0.2828	0.2708	0.081	[0.150, 0.484]	[0.114, 0.548]

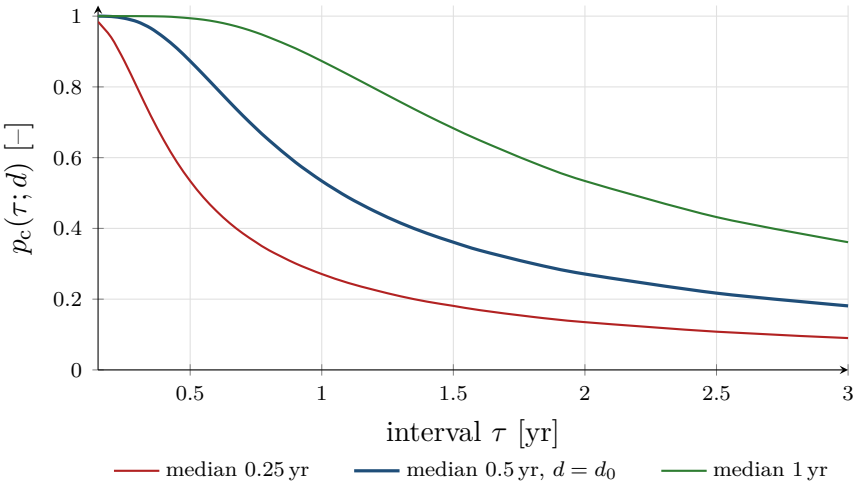


Figure 29.5: The catch probability against the interval for three median delays, equation (29.8) with $s_D = 0.4$.

Under the optimum, one failure in about four is expected to be missed: $\bar{p}_c = 0.769$. The band of the long-run fraction caught, [0.486, 0.970], is where the split of the manufacturer's figure shows: the predictive does not depend on the split, the band does (Table 29.11). The predictive and the value at d_0 differ, by 0.016 at the optimum, because p_c is curved in d; they cross between nine months and a year. The closed form and the five-node sum agree to the four digits shown at every interval.

29.3.6 The value of learning the delay

Engineering question. Would learning the delay change the interval?

General form. The value of information, the difference of two minimum expected losses, equation (2.51):

$$\mathrm{EVSI} = \min_a \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}] - \mathbb{E}_{\tilde{y}_m|\mathbf{y}}\Big[\min_a \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}, \tilde{y}_m]\Big]\,. \tag{29.28}$$

Specialized form. The question is equation (29.28) with $a = \tau$, the loss of equation (29.10), and $\tilde{y}_m$ a measurement that reveals d exactly, which gives the expected value of perfect information on the median. One choice enters here: only the median is learned, not the delay of each string, since the aleatory spread cannot be removed by any record; this is what makes the value depend on the split, and it is set beside the two limits of the split in Section 29.4.3. With d known the best interval is the inversion (29.21) with d for d_0 and s_D for s_{tot}, or no test when $\pi_D \geq 1$, so

$$\mathrm{EVPI}_d = \min_\tau \mathcal{R}(\tau) - \int \min\Big\{\mathcal{R}(\tau_d^* \mid d),\ \bar{\nu}_{\mathrm{r}} C_{\mathrm{f}}\Big\}\, \mathcal{N}(d;\ d_0,\ \sigma_d^2)\, \mathrm{d}d, \tag{29.29}$$

with $\mathcal{R}(\tau \mid d) = C_{\mathrm{i}}/\tau + \bar{\nu}_{\mathrm{r}}[C_{\mathrm{f}} - \Delta C\, p_{\mathrm{c}}(\tau; d)]$, and the integral by the rule in d. The same form with the life revealed, $\nu_{\mathrm{r}}(u, v)$ in place of $\bar{\nu}_{\mathrm{r}}$ and $\bar{p}_{\mathrm{c}}$ in place of p_{c}, gives EVPI_{uv}, read from the draws; both revealed together gives $\mathrm{EVPI}_{\mathbf{z}}$. On the exact route $\bar{\nu}_{\mathrm{r}}$ is that of equation (29.17) and the draws are those of equation (29.12).

Building blocks. The optimum and $\bar{\nu}_{\mathrm{r}}$ of Sections 29.3.1 and 29.3.2; $\tau^*(\mathbf{z})$ of equation (29.24); the draws of Section 29.2.3.

Whether the interval moves. It moves. At $d_0 \pm 1.96\,\sigma_d$, medians of 0.278 yr and 0.900 yr, the best interval is 0.515 yr and 0.826 yr. It is not monotone in the median: by equation (29.26) its slope in d is $1 - a_D$, positive here, but a_D grows as the median falls; the interval reaches its smallest value, 0.497 yr, at a median of 0.317 yr, and rises again, without bound as the median approaches 0.243 yr, below which no test pays (0.268 yr on the exact route; these medians have prior probabilities 0.008 and 0.019).

What the movement is worth. Little. Table 29.8 lists the values. Each row is what is revealed; the columns give the expected cost after it and its value on each route. Learning the median exactly saves R$ 62 per string-year on the state route, 0.8 % of the cost, and R$ 48 on the exact route. The interval moves by a quarter either way, but the curve of Figure 29.3 is shallow, so the saving is small.

Table 29.8: The expected value of perfect information [R$/string-year], with the expected cost of the best interval after the revelation. Row by row, the routes fall on opposite sides on which of the first two is worth more.

Revealed	State, after	State, value	Exact, after	Exact, value
Nothing (the decision now)	7 549	–	7 128	–
Median delay, d	7 487	62	7 081	48
Life, (u, v)	7 538	11	7 050	79
All three entries	7 476	73	6 989	140

The two routes on which is worth learning. They disagree. On the state route the median delay is worth more than five times the life, R$ 62 against R$ 11, as the budget of Table 29.6 suggests. On the exact route the life is worth more, R$ 79 against R$ 48. The state route errs because a normal in (u, v) has almost no long-lived states; the exact route's life

value comes largely from them, the 3.5 % of its lifetime states under which no test pays at all. Section 29.4.4 carries the comparison across the priors on the shape. Both values are stated, and the chapter does not choose between them; on either route the whole of the state is worth at most R$ 140, under 2 % of the cost.

29.4 The checks (step 5)

The answers rest on a renewal–reward ratio, a split of the manufacturer's figure and a flat prior on the shape, and they are computed with quadrature and draws. This section sets each against a fuller calculation or varies it. There is no physics row, so no limit $\sigma_p \to 0$ arises, and the prior row on d cannot be checked against the record, which does not bear on it.

29.4.1 The closed form against a simulated population

A population of 10^4 strings is simulated over their lives (numpy `default_rng`, seed 31): a Weibull lifetime, a lognormal delay and a uniform test phase for each string, the tests counted and the catch decided from the event times, with the onset clipped at installation. Two accountings are kept. The first is that of the closed form: every life runs to T. The second ends a life at the test that catches it, which is what would happen. Table 29.9 gives both at the MAP state. Each row is an interval; the columns give $L(\tau, \hat{\mathbf{z}})$, the first accounting with its standard error, the second accounting and its excess.

Table 29.9: The cost rate by simulation of 10^4 strings at the MAP state against the closed form [R$/string-year], and with each caught life ended at the catching test.

τ [yr]	$L(\tau, \hat{\mathbf{z}})$	Simulated (s.e.)	Life ended at the catch	Excess
0.25	10 567	10 566 (8)	10 930	363
0.5	7 482	7 488 (26)	7 805	323
0.614	7 409	7 396 (32)	7 690	281
1	8 051	8 003 (40)	8 229	178

The first accounting agrees with the closed form within one or two standard errors, and the catch fractions with equation (29.8): 0.8722 against 0.8733 at half a year. The second is dearer by R$ 178 to R$ 363, about 4 % near the optimum: a caught string is replaced before its time, so the same costs fall on fewer string-years. It costs short intervals most, and it moves the optimum of the loss at the MAP state from 0.580 yr to about 0.65 yr, on 4×10^5 simulated strings. This omission moves the interval as much as the gap between the routes does. Averaging over the posterior, 400 states with 10^4 strings each, gives 7 498 with a standard error of 53 against $\mathcal{R}(0.614) = 7\,549$, and for the exact states 7 077 (62) against 7 138; at the annual test 7 953 (60) against 8 012 and 7 356 (76) against 7 436.

29.4.2 The routes and the quadrature

Table 29.10 sets each computed quantity beside an independent calculation. Each row is one check; the columns give the quantity, its value by its route, the other calculation and the verdict.

The second-order form of equation (2.24) reads $L(\tau, \hat{\mathbf{z}}) + \frac{1}{2}\operatorname{tr}(\mathbf{H}_\tau \boldsymbol{\Sigma})$ and overshoots at half a year by R$ 35: its curvature term is R$ 233. The term is nearly all the delay's, the curvature of p_{c} in d, and it is why $\mathcal{R}$ lies above $L(\tau, \hat{\mathbf{z}})$ for short intervals and below it at a year.

Table 29.10: The checks of the routes and of the quadrature.

Quantity	Route	Independent calculation	Verdict
$\bar{\nu}_{\mathrm{r}}$, state	five nodes: 0.16905	sixty nodes: 0.16905; draws: 0.16900	agree
$\bar{\nu}_{\mathrm{r}}$, exact	one integral (29.17): 0.15285	draws: 0.15270	agree within 0.0002
$\bar{p}_{\mathrm{c}}(0.5)$	five nodes: 0.8496	closed form with s_{tot}: 0.8496	agree
$\mathcal{R}(0.5)$, state	factored (29.15): 7 680	draws: 7 679; second order (2.24): 7 715	agree; +35
$\mathcal{R}(0.614)$, state	factored: 7 549	draws: 7 549; second order: 7 568	agree; +19
τ^*, state	inversion (29.21): 0.6137	minimum on a grid of $\mathcal{R}$: 0.6137	agree
EVSI of ten catch rows, annual test	preposterior: 23.1	enumeration of the eleven outcomes: 24.6	within 1.5

29.4.3 The split of the manufacturer's figure

The expected cost and the optimum do not depend on the split, and two of the answers do. Table 29.11 gives them for six splits with $s_D^2 + \sigma_d^2 = 0.25$, from all aleatory to all epistemic. Each row is a split; the columns give the two spreads, the epistemic band of the catch probability at half a year, the best interval at the two ends of the median's band on the state route, and EVPI_d on each route.

Table 29.11: The answers that depend on the split of $s_{\mathrm{tot}} = 0.5$ into the aleatory s_D and the epistemic σ_d. In every row $\tau^* = 0.614\,\mathrm{yr}$ (0.653 yr) and $\mathcal{R}(\tau^*) = \mathrm{R\$\,7\,549}$ (R$ 7 128). A dash marks a median at which no test pays.

s_D	σ_d	Band of $p_{\mathrm{c}}(0.5; d)$	τ^* at $d_0 \mp 1.96\,\sigma_d$ [yr]	EVPI_d, state	EVPI_d, exact
0.500	0	[0.850, 0.850]	0.614, 0.614	0	0
0.450	0.218	[0.670, 0.966]	0.525, 0.751	27	18
0.400	0.300	[0.587, 0.989]	0.515, 0.826	62	48
0.354	0.354	[0.529, 0.997]	0.628, 0.890	105	87
0.300	0.400	[0.477, 1.000]	–, 0.960	167	144
$\to 0$	0.500	[0.375, 1.000]	–, 1.331	703	612

The value of learning the median runs from nothing, when the whole figure is variation among strings, to R$ 703 when the figure is ignorance about a delay every string shares; that last limit is the corner of Figure 29.3, learned. The split is not a detail of the value of information: it is most of it. For $s_D \geq 0.3$ it stays under R$ 170. The manufacturer's data behind the figure, if obtained, would fix the split without a single test at this site.

29.4.4 The prior on the shape and the failure ages

Table 29.12 repeats the decision under the two proper priors on the shape of Section 10.4.4 and, on the state route, under two changes of the failure ages. Each row is a posterior of the life; the columns give the rate, the optimum and its cost, the annual test's excess, and the two values of information. The prior refits are checks and are Chapter 10's; the first change of the ages is its interval fit of Section 10.4.5, read at its point. The second moves the four failure ages half a year later, as if each failed test had caught a degraded string a median delay before its failure.

Table 29.12: The decision under other posteriors of the life, state route with exact in parentheses. Costs in R$/string-year. The interval fit is a plug-in at its point.

Posterior of the life	$\bar{\nu}_{\mathrm{r}}$ [yr^{-1}]	τ^* [yr]	$\mathcal{R}(\tau^*)$	Annual excess	EVPI$_d$	EVPI$_{uv}$
Flat, the chapter	0.1690 (0.1529)	0.614 (0.653)	7 549 (7 128)	463 (308)	62 (48)	11 (79)
$v \sim \mathcal{N}(0, 1^2)$	0.1664 (0.1447)	0.620 (0.677)	7 480 (6 907)	436 (237)	59 (43)	20 (128)
$v \sim \mathcal{N}(0, 0.5^2)$	0.1494 (0.1227)	0.663 (0.767)	7 036 (6 281)	277 (82)	45 (46)	89 (256)
Interval rows, point	0.1723	0.607	7 631			
Failure ages +0.5 yr	0.1620	0.629	7 369	393		

The optimum stays between 0.61 yr and 0.77 yr across all of them, and the annual test stays dearer in every row, though under the narrower prior on the exact route by only R$ 82. Which unknown is worth learning does not hold: the life is worth more than the delay on the exact route under all three priors, and on the state route under the narrowest. The flat prior's tail cannot be blamed alone, since the proper priors, by pulling the shape toward one, add long-lived states rather than remove them.

29.5 The report (step 6)

The corner in the first reading was an artefact: a delay given to every string alike made half a year look like R$ 6 523 per string-year, and over what the record and the manufacturer leave uncertain half a year costs R$ 7 680 (R$ 7 327 on the exact posterior). The best interval is 0.61 yr (0.65 yr), and it beats the annual test by R$ 463 (R$ 308), a quarter (a sixth) of the saving the first reading promised. The answers follow in the order they were asked.

29.5.1 The answers and their bands

Table 29.13 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, the third the exact-route value, the fourth the gap, and the last the band or the qualification.

The report line. On a Weibull lifetime fitted to twelve strings with four failures and a manufacturer's delay figure no record at this site has tested, the expected cost of testing is least at an interval of 0.61 year on the posterior of the state and 0.65 year on the exact posterior of the life, about R$ 7 500 (R$ 7 100) per string-year, within 1 % of the minimum from about six and a half to nine months; at that interval about three failures in four are caught, 0.77, with a long-run fraction known only to lie between 0.49 and 0.97. The annual test costs R$ 463 (R$ 308) per string-year more, and it is the dearer at the true state with probability 0.93 (0.80). The delay contributes most of the spread of the cost; learning it exactly would move the interval by a quarter and save R$ 62 (R$ 48) per string-year, and which of the delay and the life is worth more to learn differs between the two routes.

29.5.2 The inspection decision

Both routes shorten the test from a year to between seven and eight months and neither justifies half a year on cost, so the gap in the interval changes no decision. The gap in the level, R$ 411 at the same interval, changes the budget figure by 5 % and not the choice. A calendar with only half-yearly, nine-month and annual slots gives nine months on both routes, by R$ 33 and by R$ 157 over six, and the state route's margin is smaller than what the renewal accounting of Table 29.9 moves. The saving against the annual test holds under every prior of Table 29.12,

Table 29.13: The report of the inspection question: the answers from $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$, the exact route with its gap, and the bands, in R$/string-year and years. On quantities of d alone the routes are one posterior, and the second value is the closed form.

Engineering question	Answer (state route)	Exact route	Gap	Band, qualification
What does an inspection interval cost per string-year, in expectation?	0.5 yr: 7 680; 0.614 yr: 7 549; 0.75 yr: 7 647; 1 yr: 8 012; no test: 10 143	7 327; 7 138; 7 170; 7 436; 9 171	353; 411; 477; 576; 972	at 0.614 yr [5 857, 10 019] ([4 820, 9 747]); renewal accounting adds about 4 %
How often should the strings be tested?	$\tau^* = 0.614$ yr, R$ 7 549	0.653 yr, R$ 7 128	0.039 yr, 421	within 1 % over [0.52, 0.73] yr ([0.55, 0.79] yr); 0.61 to 0.77 yr under the priors
How sure is the choice?	beats annual with probability 0.931; $\tau^*(\mathbf{z})$ median 0.593 yr	0.799; median 0.632 yr	0.132; 0.039 yr	$\tau^*(\mathbf{z})$ in [0.459, 0.884] yr (lower end 0.472 yr, no finite upper end); no test pays in 1.2 % (6.9 %) of states
Which unknown decides the interval?	the delay: 83 % of the variance of the cost, 87 % of that of $\ln \tau^*$	the delay: 60 %, 55 %	23, 32 points	by draws with one block held (86 %, 89 % in the budget); the state route overstates the margin
What is the probability of catching a failure under a given interval?	0.5 yr: 0.850; 0.614 yr: 0.769; 1 yr: 0.543	closed form: the same	$< 10^{-4}$	epistemic bands [0.587, 0.989], [0.486, 0.970], [0.301, 0.831]; depend on the split
Would learning the delay change the interval?	it moves over [0.515, 0.826] yr; worth R$ 62; the life worth R$ 11	worth R$ 48; the life worth R$ 79	14; 68	opposite on which is worth more; from 0 to R$ 703 with the split

but under the narrowest it falls to R$ 82 on the exact route, and a site that faces any cost of rescheduling tests should read that row. The one statement the gap does change is on what to learn, and it is left in both forms in step 7. Replacement of a string before it degrades, by age, is a different action, and a later chapter treats it: Chapter 30, in Volume III.

Principle 29.1. *An interval is chosen by minimizing the posterior expectation of its loss, and the loss is an expectation over what varies from string to string, at each value of what is not known. The expected cost sees only the sum of those two spreads; what a measurement could be worth depends on their split. Every statement of the decision is reported with its route, and a measurement is valued by what it saves in the decision, not by the width it removes.*

Common mistake

Four errors recur with intervals. Fixing the delay at its median: the corner at half a year gives R$ 6 523 where the expected cost is R$ 7 680 on the state route and R$ 7 327 exactly. Averaging the mean life and inverting: the expected life is infinite on both routes, while the rate averages to 0.16905 (0.15285). Reporting an average or a median of $\tau^*(\mathbf{z})$ as the decision: the median is 0.593 yr (0.632 yr) against the minimizer 0.614 yr (0.653 yr), and with states of no finite optimum the mean does not exist on either route. And adding the aleatory spread into the band of a fleet's catch fraction: at half a year the range over strings, [0.457, 1], is not the band [0.587, 0.989] of what is not known about the fleet.

29.6 The next measurement (step 7)

The value of perfect information of Section 29.3.6 bounds every measurement of the delay. This section prices the ones the site could take, on the state route with the exact-route value beside each. The full treatment of such values is the subject of Chapter 31, in Volume III.

29.6.1 The candidate rows on the delay

The log of how each life ended. Each string that reaches the end of its life under interval τ is either found degraded at a test or fails in service: one trial, with the success probability $p_{\mathrm{c}}(\tau; d)$. Logged as such, n_{e} endings add a binomial count in the working form of equation (4.43), as for the test rows of equation (10.38), with the probability a function of the state and not on its logit:

$$r_a = \frac{n_{\mathrm{e}} p_{\mathrm{c}} - k_{\mathrm{c}}}{\sqrt{n_{\mathrm{e}} p_{\mathrm{c}}(1 - p_{\mathrm{c}})}}, \qquad \mathbf{J}_a = \sqrt{\frac{n_{\mathrm{e}}}{p_{\mathrm{c}}(1 - p_{\mathrm{c}})}}\, \frac{\mathcal{M}(\tau; d_0)}{\tau} \begin{bmatrix} 0 & 0 & 1 \end{bmatrix}, \tag{29.30}$$

with k_{c} the endings found at a test. The row acts on the third column alone, the first evidence from this site to do so. The preposterior of equation (2.17) gives the remaining spread $\sigma_{d,a}^2 = (1/\sigma_d^2 + \mathbf{\Lambda}_{a,dd})^{-1}$, with $\mathbf{\Lambda}_{a,dd} = n_{\mathrm{e}}(\mathcal{M}/\tau)^2/[p_{\mathrm{c}}(1 - p_{\mathrm{c}})]$. Per ending, the information is 0.974 under the annual test, 1.259 at half a year and 1.369 at 0.614 yr.

How often endings occur. The twelve strings end their lives at the rate $12\,\bar{\nu}_{\mathrm{r}} = 2.03$ a year (1.83 exact). Of those, under $\tau = 0.5$ yr, $12\,\bar{\nu}_{\mathrm{r}}\bar{p}_{\mathrm{c}} = 1.72$ a year are found degraded (1.56), and under the annual test 1.10 (1.00). Ten endings, then, take about five years.

The direct reading. A string found degraded, left in service under watch and timed to its failure gives $\ln D$ itself: a Gaussian row $(\ln D - d)/s_D$, with information $1/s_D^2 = 6.25$, about five endings' worth. It costs the failure in service that the test would have prevented, unless the string can be run to its limit off the critical load.

29.6.2 The preposterior and the value

Table 29.14 lists the candidates. Each row is a candidate; the columns give the rows added, the remaining spread $\sigma_{d,a}$, the band of the catch fraction at $\tau = 0.614$ yr that it leaves, and the EVSI of equation (29.28). The EVSI comes from the outcomes of the preposterior mean of d, $\mathcal{N}(d_0, \sigma_d^2 - \sigma_{d,a}^2)$, with the best interval at each outcome by equation (29.21) at $s_D^2 + \sigma_{d,a}^2$. Figure 29.6 draws the same value as a curve: the horizontal axis is the remaining spread, the present 0.3 at the right and perfect knowledge at the left; the vertical axis the reduction of the expected cost of the best interval, on each route; the dashed levels the value of revealing the life.

Table 29.14: The candidate measurements of the delay ranked by the expected value of sample information [R$/string-year], on the state route and the exact route. The band is that of $p_{\mathrm{c}}(0.614; d)$.

Candidate	Rows	$\sigma_{d,a}$	Band of p_{c}	EVSI, state	EVSI, exact
Present record	–	0.300	[0.486, 0.970]	–	–
Five endings logged, annual test	5	0.250	[0.533, 0.954]	14	9
Ten endings logged, annual test	10	0.219	[0.563, 0.941]	23	15
Ten endings logged, $\tau = 0.614$ yr	10	0.201	[0.581, 0.932]	28	19
Twenty endings logged, annual test	20	0.181	[0.602, 0.922]	34	23
Five delays timed directly	5	0.154	[0.630, 0.906]	40	29
Ten delays timed directly	10	0.117	[0.668, 0.882]	49	36
Median known exactly	–	0	0.785	62	48

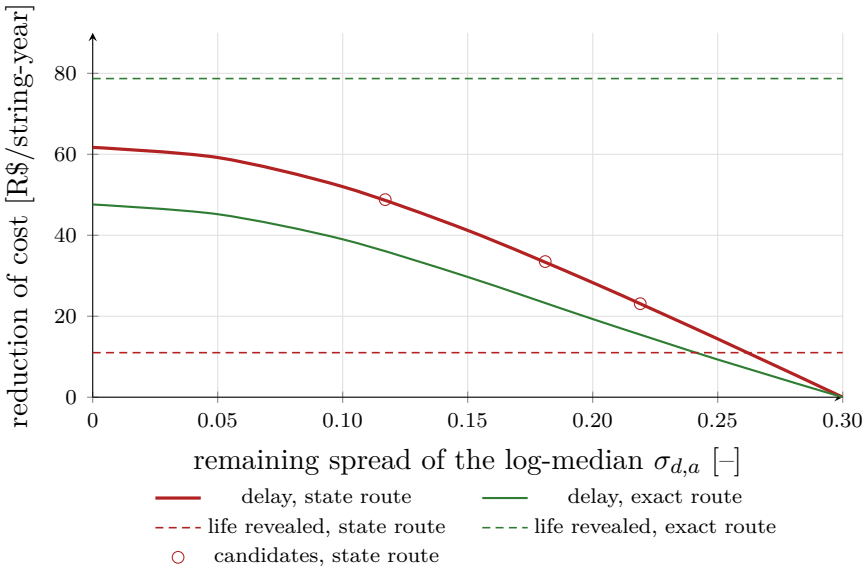


Figure 29.6: The value-of-information curve: the reduction of the expected cost of the best interval as the spread of the log-median falls from its present 0.3, with $s_D = 0.4$. Circles: ten and twenty endings logged at the annual test and ten delays timed. Dashed: EVPI_{uv} on each route.

The curve rises slowly near the present spread and flattens toward perfect knowledge; a measurement that halves the spread collects about two thirds of the perfect value. Ten logged endings narrow the band of the catch fraction by a fifth and are worth R$ 23 per string-year (R$ 15), about R$ 280 a year for the twelve strings (R$ 185) once they are in, five years from now. The enumeration of the eleven binomial outcomes gives 24.6, confirming the Gaussian preposterior within 1.5.

29.6.3 The recommendation

Keep the log. Recording, for each string, whether its life ended at a test or in service costs nothing beyond what the program already pays, it is the first row the delay has ever had, and it needs no change to the interval. Ask the manufacturer what the figure is made of, since the split of Table 29.11 moves the value of the median by more than any measurement at this site could. Do not time delays on strings carrying the critical load: ten of them are worth R$ 49 per string-year at most, R$ 590 a year over the twelve (R$ 430), against a failure in service at R$ 60 000 each. Whether to spend on the life instead is the point where the routes differ: on the state route it is worth R$ 11 and on the exact route R$ 79, and the capacity readings the lifetime record discards, which Chapter 10 also points to, serve the life and the delay together, since a trajectory of readings shows when degradation became measurable. When rows arrive, the fit returns to step 3 with them.

29.6.4 What the analysis showed

The first reading said to test every half year and save more than twenty thousand reais a year; the analysis says to test about every seven or eight months and save four to six thousand, and half a year turns out to be no better than nine months. Nearly all of the difference is that a delay has a spread, which puts half of the strings on the wrong side of the corner; the uncertainty of the life matters much less to the interval, because the mean life is well fixed by four failures on the Gaussian, and on the exact posterior less well than that, but it still moves the optimum by only two weeks. The engineer’s instinct that the delay is the thing nobody

knows is right about the spread of the cost and right that the interval would move if it were known. What is new is how little the movement is worth, because a shallow curve forgives a wrong interval: the delay decides where the minimum is, and the shape of the curve decides that it hardly matters. For the site, the annual test is too long and the right interval is not sharply defined. A longer record is best spent on knowing how lives ended, and the site's own data cannot yet tell whether the delay or the life is the more valuable to learn.

29.6.5 Limits of the model

The derivations are exact under their assumptions, and six of them need checking before a result is used.

One life, one cost per cycle. The renewal–reward ratio over T omits the shortening of a caught life, which adds about 4 % and moves the optimum later (Section 29.4.1), and the replaced string is taken as new. Replacement before degradation, by age, is left out altogether and is the subject of Chapter 30, in Volume III.

The delay's independence. D is independent of T and of the string's age. A delay that shortens as a string ages, or a wear-out in which an early failure is also a quick one, would make the factoring of equation (29.15) fail and would call for an interval that shortens with age.

A perfect test, and a fixed cost. Every test inside the degraded state finds it; Exercise 29.5 shows the catch falling from 0.77 to 0.64 at a detection probability of 0.8. The cost of a failure in service is one number, R$ 60 000; a failure's consequence depends on whether an outage comes while the string is down, and doubling it moves the optimum to 0.42 yr, while halving it to R$ 30 000 removes any finite optimum, since then $\pi_D > 1$. The test cost matters as much: at R$ 4 000 the optimum is 1.33 yr.

The double reading of a failed test. A failed capacity test in the record of Chapter 10 is read as a failure, and here the same test would find a degraded string. Moving the failure ages half a year later changes the optimum from 0.614 yr to 0.629 yr (Table 29.12): the decision barely moves, the reading of the record does.

The delay is a prior. The third entry rests entirely on the manufacturer's figure and the split this chapter chose for it. A figure from another duty or room temperature would shift d_0, on which the optimum depends as Figure 29.5 shows.

The strings as one population. One interval is used for all strings. A fleet of mixed ages is better served by different intervals (Exercise 29.6), and the survivors' own remaining lives of Chapter 10 would then replace the mean life.

The next chapter takes up the action this one left out: replacing a unit before it is found degraded, at an age chosen by the same minimization.

29.7 What to retain

Definition 29.1 (Expected cost rate of an inspection interval)**.** For a lifetime with replacement rate $\nu_{\mathrm{r}}(\mathbf{z}) = 1/\mathrm{MTTF}(\mathbf{z})$, a delay D between detectable degradation and failure, and a test grid of period τ with uniform phase, the loss is $L(\tau, \mathbf{z}) = C_{\mathrm{i}}/\tau + \nu_{\mathrm{r}}[C_{\mathrm{f}} - (C_{\mathrm{f}} - C_{\mathrm{r}})p_{\mathrm{c}}(\tau; d)]$, with $p_{\mathrm{c}} = \mathbb{E}[\min(D/\tau, 1) \mid d]$. The decision minimizes $\mathcal{R}(\tau) = \mathbb{E}[L(\tau, \mathbf{z}) \mid \mathbf{y}]$; its condition is that the partial mean of the delay below τ^* equals $C_{\mathrm{i}}/((C_{\mathrm{f}} - C_{\mathrm{r}})\bar{\nu}_{\mathrm{r}})$, and no test pays when that time exceeds the mean delay.

Four conclusions carry beyond the strings. A decision is one more derived quantity per action, read from the same posterior as every estimate; its band, its ranking against another action and the distribution of what would be best at a known state are three different statements, and the action is none of the last two. A sharp optimum at a parameter's value usually means the parameter was given no spread. What varies from unit to unit is integrated inside the loss, and what is not known is an entry of the state; the expected cost cannot tell them apart, but the band of a fleet's fraction and the value of information can. An entry that has a prior row and no record row is unobservable from the record: its posterior is its prior, its exact and Laplace routes coincide, and the first row that touches its column is worth more than one that narrows a column already covered, if the decision moves. And what a measurement is worth is set by the curve of the expected cost as much as by the spread it removes. Table 29.15 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 29.15: The terms of the inspection decision and where each is defined.

Term	Expression	Units	Meaning	Defined in
Catch, known delay	$\min(D/\tau, 1)$	–	uniform phase	29.1.1
Cost rate	$C_{\mathrm{i}}/\tau + \nu_{\mathrm{r}}(C_{\mathrm{f}} - \Delta C\, p_{\mathrm{c}})$	R$/yr	renewal–reward	29.1.1
Aleatory spread	$\ln D \mid d \sim \mathcal{N}(d, s_D^2)$	–	string to string; in the loss	29.1.4
Epistemic spread	$r_\pi = (d - d_0)/\sigma_d$	–	not known; in the state	29.1.4
Total spread	$s_{\mathrm{tot}}^2 = s_D^2 + \sigma_d^2$	–	all $\mathcal{R}$ sees	29.1.4
Partial mean	$\mathcal{M} = \mathrm{e}^{d + s_D^2/2}\Phi((\ln\tau - d - s_D^2)/s_D)$	yr	delay below τ	29.1.4
Catch probability	$p_{\mathrm{c}} = \mathcal{M}/\tau + \Pr(D \geq \tau \mid d)$	–	given the median	29.1.4
General forms				
Expected consequence	$\mathcal{R}(a) = \int L(a, \mathbf{z}) p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	R$/yr	one action	29.3.1
The decision	$a^* = \arg\min_a \mathcal{R}(a)$	–	minimizer	29.3.2
Ranking	$\Pr(L(a, \mathbf{z}) < L(a', \mathbf{z}) \mid \mathbf{y})$	–	two actions	29.3.3
Variance budget	$\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$, term by term	varies	attribution	29.3.4
Value of information	$\bar{L}_{\mathrm{now}} - \bar{L}_{\mathrm{after}}$	R$/yr	the decision improved	29.3.6
Specialized forms and answers				
Expected cost, factored	$C_{\mathrm{i}}/\tau + \bar{\nu}_{\mathrm{r}}(C_{\mathrm{f}} - \Delta C\, \bar{p}_{\mathrm{c}})$	R$/yr	blocks independent	29.3.1
Rate, exact	$\int p(v \mid \mathbf{y})\, \Gamma(k + 1/\beta)/[\Gamma(k) A^{1/\beta} \Gamma(1 + 1/\beta)]\, \mathrm{d}v$	yr^{-1}	one integral	29.3.1
First-order condition	$\bar{\mathcal{M}}(\tau^*) = C_{\mathrm{i}}/(\Delta C\, \bar{\nu}_{\mathrm{r}})$	yr	test pays for itself	29.3.2
Optimum	$\exp[d_0 + s_{\mathrm{tot}}^2 + s_{\mathrm{tot}}\Phi^{-1}(\pi_D)]$	yr	closed form	29.3.2
Optimum at a state	$\exp[d + s_D^2 + s_D\Phi^{-1}(\pi_D(\mathbf{z}))]$	yr	not the decision	29.3.3
Gradient of the optimum	$(a_D,\ -a_D\psi\, \mathrm{e}^{-v},\ 1 - a_D)^{\mathsf{T}}$	–	life against delay	29.3.4
Perfect information	$\min_\tau \mathcal{R}(\tau) - \mathbb{E}_d[\min_\tau \mathcal{R}(\tau \mid d)]$	R$/yr	bounds any row	29.3.6
Ending row	$\mathbf{J}_a \propto (0,\ 0,\ \mathcal{M}/\tau)$	–	first row on d	29.6.1
Remaining spread	$\sigma_{d,a}^2 = (1/\sigma_d^2 + \boldsymbol{\Lambda}_{a,dd})^{-1}$	–	the preposterior	29.6.1

29.8 Exercises

Exercise 29.1. *Formulation.* A test costs R$ 1 000, a planned replacement R$ 10 000, a failure in service R$ 40 000, the mean life is 8 yr. (a) Compute the cost rate at $\tau = 0.5$ yr with a catch probability of 0.8. (b) Give the catch probability of a fixed delay of 0.3 yr at $\tau = 0.5$ yr and at $\tau = 0.25$ yr. (c) Write the first as an expected loss over a state and say which entries of the state it would need if the life and the delay were uncertain. ((a)

$1\,000/0.5 + (0.8 \times 10\,000 + 0.2 \times 40\,000)/8 = 2\,000 + 2\,000 =$ R\$ 4 000. (b) 0.6 and 1. (c) $\mathcal{R}(0.5) = \mathbb{E}[L(0.5, \mathbf{z}) \mid \mathbf{y}]$ with $\mathbf{z}$ holding the lifetime's entries and the delay distribution's parameters, not the delay of any one string.)

Exercise 29.2. *Formulation.* In the split of the chapter, the catch probability at $\tau = 1\,\text{yr}$ has the predictive 0.543, the epistemic band [0.301, 0.831] and the aleatory range [0.228, 1]. Say in one sentence what each describes and which would narrow under ten logged endings. (The predictive is the chance that one failure is caught; the epistemic band covers the long-run fraction this fleet would catch, and it narrows, to about the $\sigma_{d,a} = 0.219$ of Table 29.14; the aleatory range describes how strings differ given their delays, and no record of the fleet narrows it.)

Exercise 29.3. *Derivation.* Derive equation (29.1) for a uniform phase, then the partial mean (29.9) of a lognormal delay. (The time B back from the failure to the last test is uniform on $(0, \tau)$, so $\Pr(B < D \mid D) = \min(D/\tau, 1)$. With $x = \ln D$, $\mathrm{e}^{x}\mathcal{N}(x; d, s^2) = \mathrm{e}^{d+s^2/2}\mathcal{N}(x; d + s^2, s^2)$ by completing the square; integrating to $\ln \tau$ gives $\mathrm{e}^{d+s^2/2}\Phi((\ln \tau - d - s^2)/s)$.)

Exercise 29.4. *Derivation.* (a) Derive the mean life of a Weibull lifetime. (b) Derive the derivative (29.19) and the condition (29.20). (c) Show that with a fixed delay D the minimum is at $\tau = D$ when $\nu_{\mathrm{r}}\Delta C D > C_{\mathrm{i}}$, and that otherwise no test pays. ((a) $\int_0^\infty \mathrm{e}^{-(t/\eta)^\beta}\,\mathrm{d}t$ with $x = (t/\eta)^\beta$ gives $(\eta/\beta)\Gamma(1/\beta) = \eta\,\Gamma(1 + 1/\beta)$. (b) $\partial \min(D/\tau, 1)/\partial\tau = -(D/\tau^2)\mathbf{1}[D < \tau]$, so $\mathrm{d}\mathcal{R}/\mathrm{d}\tau = -C_{\mathrm{i}}/\tau^2 + \bar{\nu}_{\mathrm{r}}\Delta C \bar{\mathcal{M}}/\tau^2$. (c) For $\tau < D$, $C_{\mathrm{i}}/\tau + \nu_{\mathrm{r}}C_{\mathrm{r}}$ falls with τ; for $\tau > D$ the cost is $\nu_{\mathrm{r}}C_{\mathrm{f}} + (C_{\mathrm{i}} - \nu_{\mathrm{r}}\Delta C D)/\tau$, which rises when the bracket is negative and falls toward no test when it is positive.)

Exercise 29.5. *Application.* A test finds a degraded string with probability $q = 0.8$, independently at each test. Write the catch probability for a known delay, and find the new optimum on the state route of the chapter. (With $D/\tau = m + f$, m an integer and $f \in [0, 1)$, the degraded state holds $m + 1$ tests with probability f and m with probability $1 - f$, so the catch is $1 - (1 - q)^m(1 - qf)$. Averaged over the delay with s_{tot}, $\tau^* = 0.620\,\text{yr}$, $\mathcal{R}(\tau^*) =$ R\$ 8 542 and $\bar{p}_{\mathrm{c}} = 0.634$, against 0.614 yr, R\$ 7 549 and 0.769. The interval hardly moves, because a shorter one catches more only by testing more; the cost rises by R\$ 993.)

Exercise 29.6. *Application.* Two populations of strings, of ages 1 and 5 years, are each to have their own interval, for their remaining lives only. Taking the MAP state and the mean residual life of Chapter 10 in place of the mean life, and $s_{\mathrm{tot}} = 0.5$, find the two intervals. (The mean residual lives at $\hat{\mathbf{z}}$ are 4.95 yr and 1.58 yr. Then $\pi_D = 2\,000 \times 4.95/(45\,000 \times 0.5666) = 0.388$ and 0.124, and equation (29.21) gives 0.557 yr and 0.360 yr, at R\$ 8 357 and R\$ 16 803 per string-year. The older strings need tests more often, since each failure they would have is sooner, and they cost twice as much to hold.)

Chapter 30

Repair, Replacement and Preventive Maintenance

30.1 When to replace a battery string

Chapter 29 shortened the capacity test of the twelve battery strings and left a question standing: whether a string should be replaced before anything is found wrong with it. The facility manager has a supplier's offer to do so at a fixed age, at R$ 15 000 a string, the planned price of that chapter, and a string that fails in service still costs R$ 60 000. The record is the same four failures and eight strings in service, and four of the eight are already past four and a half years.

The quickest reading again takes the fit as known, a characteristic life of 6.50 years and a shape of 4.73. Running to failure then costs R$ 60 000 over a mean life of 5.95 years, R$ 10 091 per string-year; replacing at 3.90 years costs R$ 4 910, so that the age rule saves about R$ 5 180 per string-year, about R$ 62 000 a year over the twelve strings. Both costs are those of one state, and four failures leave the shape wide, as Chapter 10 showed. The question of this chapter is

what each intervention costs in expectation over what the record leaves uncertain, which action minimizes that cost, how far it stands from the action that is best at the fitted state, and what knowing the state would be worth.

What changes from Chapter 29 is that the action set holds interventions of different kinds, compared on one cost basis, and that a state known exactly is priced against all of them.

Jargon: age replacement, run to failure, condition-based replacement, renewal cycle, renewal–reward theorem, plug-in action, Bayes action, regret, expected value of perfect information

Age replacement replaces a string in a planned way at the age T, or at failure if that comes first. *Run to failure* replaces it only at failure. *Condition-based replacement* replaces it when a periodic test finds it degraded, the action of Chapter 29. A *renewal cycle* is the time from one new string to the next. The *renewal–reward theorem* states that the long-run cost per unit time is the expected cost of a cycle over its expected length. The *plug-in action* is the action that minimizes the loss at the point estimate $\hat{\mathbf{z}}$. The *Bayes action* minimizes the posterior expectation of the loss. The *regret* of an action at a state is its loss there less the loss of the best action there. The *expected value of perfect information* is the expectation of the regret of the Bayes action.

Key idea

Three objects carry the name of an optimum and they are different: the plug-in action $T^*(\hat{\mathbf{z}})$, one number computed at one state; the distribution of the optimum $T^*(\mathbf{z})$, what would be best at each state; and the Bayes action, the minimizer of the expected cost, of which each posterior route has its own. The expected value of perfect information is the gap between the expectation of the minimum and the minimum of the expectation. Near the fit it is set by the width of the distribution of the optimum and by the curvature of the loss, on the exact route half of it comes from states whose optimum is very late or absent, and on either route it depends on which actions are on the table.

30.1.1 The renewal cycle and the cost rate of an age

A string is new at age zero, with lifetime T, and is replaced at $\min(T, T)$, after which a new string starts an independent copy of the same cycle. (The symbol T names the replacement age throughout this chapter; the hall temperature that the notation also writes T does not appear here.) Figure 30.1 draws three cycles along the calendar. The horizontal line is time; each segment one cycle, labeled by how it ended; a blue tick a planned replacement at age T and a red cross a failure before T.

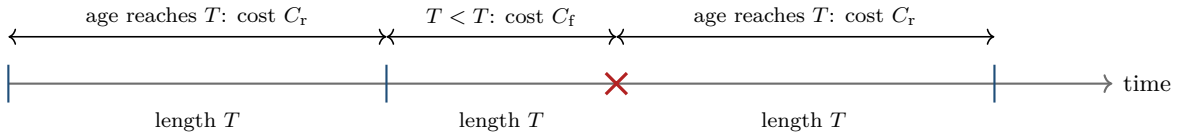


Figure 30.1: Renewal cycles under age replacement. Each cycle ends at $\min(T, T)$, at a planned replacement (blue tick) or at a failure (red cross), and the next string starts new. The lengths are illustrative.

One cycle. Given the lifetime distribution, with the survival function R of equation (7.14) and $F = 1 - R$, a cycle costs C_r with probability $R(T)$ and C_f with probability $F(T)$. Its expected length is the integral of the survival function, since for a nonnegative variable $\mathbb{E}[X] = \int_0^\infty \Pr(X > t)\,\mathrm{d}t$ and $\min(T, T) > t$ means $T > t$ for $t < T$:

$$\mathbb{E}[\text{cost}] = C_\mathrm{r} R(T) + C_\mathrm{f} F(T), \qquad W(T) = \mathbb{E}[\min(T, T)] = \int_0^T R(t)\,\mathrm{d}t. \tag{30.1}$$

The renewal–reward theorem, a standard result treated in the renewal chapters of the companion volume *Reliability Engineering*, says that the cost per unit time over many cycles tends to the ratio of those two expectations. With $\Delta C = C_\mathrm{f} - C_\mathrm{r}$ as in Chapter 29,

$$\text{cost rate} = \frac{C_\mathrm{r} + \Delta C\, F(T)}{W(T)}. \tag{30.2}$$

As $T \to \infty$, $F \to 1$ and $W \to \mathrm{MTTF}$, so run to failure is the limit $C_\mathrm{f}/\mathrm{MTTF}$, the same limit as the no-test row of Table 29.4. As $T \to 0$ the rate grows without bound, as a string replaced almost at once. Between them lies the condition for an optimum.

The condition for an optimum. Setting the derivative of equation (30.2) to zero, with $\mathrm{d}F/\mathrm{d}T = f(T) = \lambda(T)R(T)$ and $\mathrm{d}W/\mathrm{d}T = R(T)$, the factor $R(T)$ cancels and

$$\lambda(T)\, W(T) + R(T) = \frac{C_\mathrm{f}}{\Delta C}, \qquad \text{cost rate at the root} = \Delta C\, \lambda(T). \tag{30.3}$$

The left side equals one at $T = 0$, and its derivative is $\lambda'(T)W(T)$, so it rises exactly while the hazard rises. The right side is above one. A constant or falling hazard therefore gives no root, and the rate falls all the way to run to failure: replacing an unaged string buys nothing. A rising hazard gives a root if the left side reaches the right. For a Weibull lifetime with $\beta > 1$ the hazard grows without bound and the root always exists, however large; with $\beta \leq 1$ none does. The shape question of Chapter 7 thus decides whether age replacement can help at all.

30.1.2 The strings and their plant graph (step 0)

The plant is that of Chapter 29: the strings of Table 10.1 and the degraded state of Figure 29.2. Figure 30.2 redraws it for the wider action set. The two green boxes are the free nodes, the lifetime with $u = \ln \eta$ and $v = \ln \beta$, reached by the twelve rows of the record, and the degraded state with the log-median delay d, reached only by the manufacturer's prior row. The gray boxes are pinned, the spread s_D and the costs. The blue box is the action, one of three kinds, and the box at the right the loss. The arrow from the delay to the loss carries a label because only the third kind uses it.

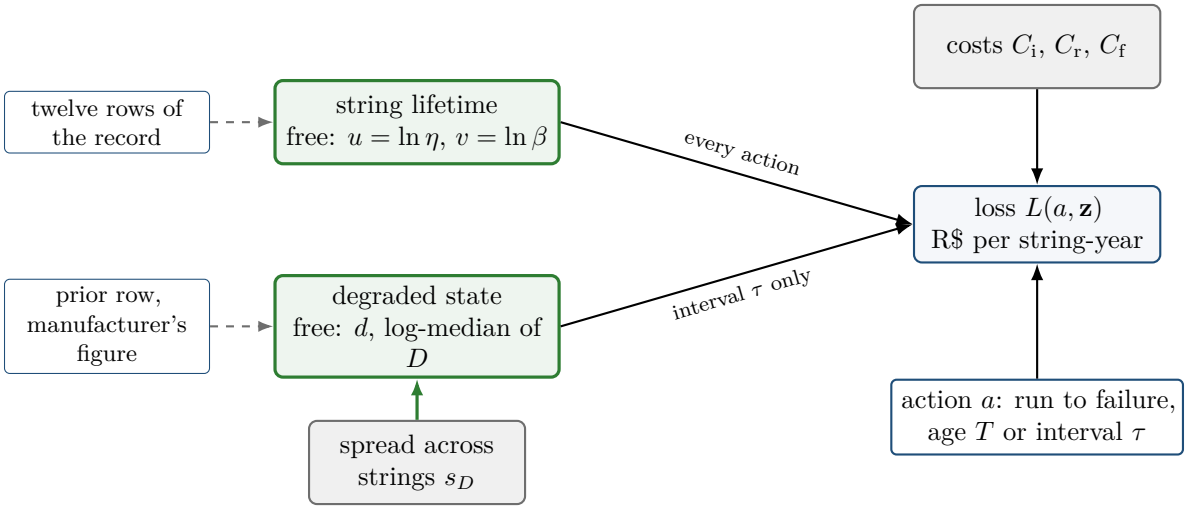


Figure 30.2: The plant graph of the replacement question. Free: the lifetime with u and v, the degraded state with d (green). Pinned: the spread s_D and the costs (gray). Observed: the twelve rows of the record on the lifetime and only the manufacturer's prior row on the delay (dashed). The action is not an entry of the state, and only condition-based replacement reads d.

Table 30.1 lists the pinned costs. Each row is a quantity, with its symbol, value and status. The planned replacement at an age and the planned replacement of a string found degraded are the same work at the same price, so one symbol serves both, and that is the first part of the common cost basis.

Table 30.1: The costs of the three interventions (illustrative), those of Table 29.1, with the delay figure as there.

Quantity	Symbol	Value	Used by
Planned replacement	C_r	R$ 15 000	age, condition
Failure in service, with the loss event	C_f	R$ 60 000	all three
Capacity test of one string	C_i	R$ 2 000	condition
Delay, median and total log-spread	e^{d_0}, s_tot	0.5 yr, 0.5	condition

30.1.3 The probabilistic questions (step 1)

The manager asks six questions, each written over the state $\mathbf{z} = (u, v, d)^{\mathsf{T}}$. Table 30.2 lists them. The first column is the question in engineering language; the second the general form of its category (Chapter 2); the third its form for the strings; the fourth its category; the last the subsection that answers it. Three symbols for an optimal age are kept apart throughout: $T^*(\mathbf{z})$ is the age that would be best if the state were known; $T^*(\hat{\mathbf{z}})$ is that function at the MAP state, the plug-in action; and T^{B} is the Bayes action, written $T^{\mathrm{B}}_{\mathrm{ex}}$ on the exact route.

Table 30.2: The questions asked about the interventions, the general form of their category, their 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) and the prior row on d; $a = \infty$ is run to failure.

Engineering question	General form	Mathematical form	Category	Answered in
What does each intervention cost per string-year?	$\mathcal{R}(a) = \int L(a, \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	$\mathcal{R}(a)$ for $a = \infty$, $a = T$, $a = \tau$	design (expected loss)	30.3.1
At what age should a string be replaced?	$a^* = \arg\min_a \mathcal{R}(a)$	$T^{\mathrm{B}} = \arg\min_T \mathcal{R}(T)$	design (the decision)	30.3.2
Is the age rule better than running to failure, with assurance?	$\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y})$	$\Pr(L(\infty, \mathbf{z}) - L(T^{\mathrm{B}}, \mathbf{z}) > 0 \mid \mathbf{y})$	prediction under a threshold	30.3.3
How different are the plug-in optimum and the Bayes action, and does it matter?	$\mathcal{R}(a') - \min_a \mathcal{R}(a)$; quantiles of $f(\mathbf{z})$	$T^*(\hat{\mathbf{z}})$ against T^{B}; $\mathcal{R}(T^*(\hat{\mathbf{z}})) - \mathcal{R}(T^{\mathrm{B}})$; $p(T^*(\mathbf{z}) \mid \mathbf{y})$	the Bayes action against the plug-in action	30.3.4
How much would knowing the lifetime parameters exactly save?	$\min_a \mathcal{R}(a) - \mathbb{E}[\min_a L(a, \mathbf{z}) \mid \mathbf{y}]$	the same, with $\mathbf{z}$ or a block of it revealed	design (value of information)	30.3.5
Which intervention wins, and by how much?	$\mathcal{R}(a)$; $\Pr(L(a, \mathbf{z}) < L(a', \mathbf{z}) \mid \mathbf{y})$	the three minima, their bands and rankings	design (the comparison)	30.3.6

30.1.4 The evidence as rows (step 2)

Step 2 adds no row. The twelve lifetime rows and the prior row on the delay are those of Chapter 29, with the pattern of equation (29.3): the first two columns of the Jacobian carry the record, the third only the prior row. The new content of this step is the loss of each action as a function of the state.

The modeling choices. Five choices make the losses. The lifetime rows keep the reading of Chapter 10, a failed test taken as a failure at the test age: this reuses its posterior, and the chapter's own interval fit moves it, which step 5 measures. A replacement restores a string to new: this is what makes the cycles independent copies and the renewal–reward ratio available, and it fails if a new string joins an older bank that shortens its life. The age rule is one number applied to every new string: it gives a stationary rate, and it leaves out what to do now with strings already past that age, which the report takes up. The costs do not depend on the age at replacement: this keeps equation (30.2), and Exercise 30.6 relaxes it. A failure's cost is one number: this is the point at which the outage it might meet is averaged away, and the limits show how far it moves the answers.

The losses. With $\Lambda(t;\mathbf{z}) = (t/\eta)^{\beta}$ as in Chapter 10, the age rule has the loss of equation (30.2) at the state, and run to failure its limit:

$$L(T,\mathbf{z}) = \frac{C_{\mathrm{r}} + \Delta C\, F(T;\mathbf{z})}{W(T;\mathbf{z})}, \qquad L(\infty,\mathbf{z}) = C_{\mathrm{f}}\,\nu_{\mathrm{r}}(u,v), \tag{30.4}$$

with ν_{r} the replacement rate of equation (29.7). Neither reads d, so the third entry of their gradients is zero. The gradient in (u,v) uses the survivor-row gradient $\nabla\Lambda(t) = (-\beta\Lambda(t),\ \beta x(t)\Lambda(t))^{\mathsf{T}}$ of Section 10.2.1, with $x(t) = \ln t - u$. Since $\nabla F(T) = R(T)\nabla\Lambda(T)$ and $\nabla W(T) = -\int_0^T R(t)\nabla\Lambda(t)\,\mathrm{d}t$,

$$\mathbf{g}_T = \frac{1}{W(T)}\left[\Delta C\, R(T)\nabla\Lambda(T) + L(T,\mathbf{z})\int_0^T R(t)\nabla\Lambda(t)\,\mathrm{d}t\right], \qquad [\mathbf{g}_T]_d = 0. \tag{30.5}$$

This L is the form $c_a + \sum_e C_e\ \Pr(e \mid \mathbf{z}, a)$ of equation (2.23), per string-year: no certain cost, and two events at the rates R/W and F/W.

Condition-based replacement on one basis. The loss of Chapter 29, equation (29.10), divides by the full life MTTF: a caught string is charged as if it had lived to its failure. An age rule ends its cycle when the string is taken out, so comparing the two requires the same accounting. Measure from the onset $T - D$ to the next test, U, uniform on $(0,\tau)$. The string is caught if $U < D$ and its life ends at the onset plus U; otherwise it runs to T. The cycle then has the length $T - D + \min(U,D)$. The tests before the onset average $(T-D)/\tau$, and a catching test adds one with probability p_{c}. Treating the onset as positive, which the simulation below tests, the renewal–reward ratio is

$$L_{\mathrm{c}}(\tau,\mathbf{z}) = \frac{C_{\mathrm{i}}\big[(\mathrm{MTTF} - \mathbb{E}[D \mid d])/\tau + p_{\mathrm{c}}\big] + p_{\mathrm{c}}C_{\mathrm{r}} + (1-p_{\mathrm{c}})C_{\mathrm{f}}}{\mathrm{MTTF} - \mathbb{E}[D \mid d] + m(\tau;d)}, \tag{30.6}$$

with $p_{\mathrm{c}}(\tau;d)$ and $\mathcal{M}(\tau;d)$ of equations (29.8) and (29.9), and the mean of $\min(U,D)$, by the same completing of the square for the second moment,

$$m(\tau;d) = \mathcal{M}(\tau;d) - \frac{\mathrm{e}^{2d+2s_D^2}}{2\tau}\,\Phi\!\left(\frac{\ln\tau - d - 2s_D^2}{s_D}\right) + \frac{\tau}{2}\,\Pr(D \ge \tau \mid d). \tag{30.7}$$

Of the three interventions, only this loss reads d.

The age rule with the annual test kept. The age rule above has no test in it. The program in place tests every string once a year at C_{i}, and keeping that test beside the age rule is a fourth program, not a free addition. Its cycle ends at the first of the age T, the failure, and a test inside the degraded state, on a test grid of period $\tau = 1\,\mathrm{yr}$ with uniform phase; its loss $L_{\mathrm{a}}(T,\mathbf{z})$ is the expected cost of such a cycle, tests included, over its expected length. It reads d too, and has no closed form: the ratio is integrated by a rule over the lifetime, the delay and the phase at each state.

30.2 The posterior of the state (step 3)

All six questions are read from the posterior of Chapter 29: the rows are unchanged, so the state, its fit and its covariance are unchanged, and the one new computation is the integral over the scale that the age rule requires.

30.2.1 The posterior from the thirteen rows

The residual vector and the 13×3 Jacobian are those of equation (29.11): the working rows of the record in (u, v) and the prior row $(d - d_0)/\sigma_d$. The Hessian is block diagonal, with the (u, v) block $\mathbf{J}_{\mathrm{rec}}^{\mathsf{T}}\mathbf{J}_{\mathrm{rec}}$ plus its named curvature term and the d block $1/\sigma_d^2$. Of the split of the delay figure chosen in Section 29.1.4, $s_D = 0.4$ stays inside the loss and $\sigma_d = 0.3$ is the state entry. The fit, its Newton iteration in Table 10.3, and the information and covariance are those of equation (29.13):

$$\hat{\mathbf{z}} = \begin{bmatrix} 1.87138 \\ 1.55303 \\ -0.69315 \end{bmatrix}, \qquad \boldsymbol{\Sigma} = \begin{bmatrix} 0.01274 & -0.01532 & 0 \\ -0.01532 & 0.15192 & 0 \\ 0 & 0 & 0.09 \end{bmatrix}, \tag{30.8}$$

that is $\hat{\eta} = 6.4973\,\mathrm{yr}$, $\hat{\beta} = 4.7257$, with the long axis of the lifetime block $\mathbf{e}_1 = (0.108,\ -0.994)^{\mathsf{T}}$ at $\sqrt{\epsilon_1} = 0.392$ and the short axis $\mathbf{e}_2 = (0.994,\ 0.108)^{\mathsf{T}}$ at $\sqrt{\epsilon_2} = 0.105$. The MAP state is the joint mode on both routes, since the exact posterior has that mode under the flat prior: there is one plug-in action, not one for each route.

30.2.2 The exact posterior and the two routes

The exact posterior is equation (29.12), $p(u, v \mid \mathbf{y})\,\mathcal{N}(d;\ d_0,\ \sigma_d^2)$, with $\zeta = \eta^{-\beta}$ given the shape distributed as Gamma$(k, A(\beta))$ and the marginal $p(v \mid \mathbf{y})$ of Section 10.2.2. *The state route* takes (u, v, d) as $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$. *The exact route* takes the first factor in place of its Gaussian. Chapter 29 needed one integral for the rate, because the gamma moment is closed. The age rule's loss is not a moment of ζ: it is a ratio with W below. Written on (ζ, β), with $F = 1 - \exp(-\zeta T^{\beta})$ and $W = \int_0^T \exp(-\zeta t^{\beta})\,\mathrm{d}t$, the expected cost needs a second integral,

$$\mathcal{R}_{\mathrm{ex}}(T) = \int p(v \mid \mathbf{y}) \int_0^{\infty} L(T; \zeta, \beta)\, \frac{A^k \zeta^{k-1} \mathrm{e}^{-A\zeta}}{\Gamma(k)}\, \mathrm{d}\zeta\, \mathrm{d}v, \tag{30.9}$$

by the grid of Section 10.2.3 in v and a 48-node gamma rule (Gauss–Laguerre with the factor x^{k-1}) in $A\zeta$.

30.2.3 The posterior of the strings

The optimum as a function of the state. Dividing times by the scale in equation (30.3) leaves one equation in $s = T/\eta$ and the shape:

$$\beta\, s^{\beta-1} \int_0^s \mathrm{e}^{-x^{\beta}}\, \mathrm{d}x + \mathrm{e}^{-s^{\beta}} = \frac{C_{\mathrm{f}}}{\Delta C} = \frac{4}{3}, \qquad T^*(\mathbf{z}) = \eta\, g(\beta), \tag{30.10}$$

with $g(\beta)$ its root, for $\beta > 1$, and no finite optimum for $\beta \le 1$. The pushforward of the posterior through g is therefore a change of the two entries alone. The two routes carry it differently: the shape has the median 4.73 and the band $[2.20,\ 10.14]$ on the state route and 3.64 and $[1.15,\ 7.77]$ on the exact route, as Table 10.15 reports. The probability of a shape at or below one is $\Phi(-1.553/0.390) = 3 \times 10^{-5}$ on the state route and 0.016 on the exact route.

The draws. Bands, tails and the distribution of the optimum use 2×10^5 draws on each route (numpy `default_rng`, seed 28, those of Section 29.2.3, with d drawn separately under seed 29). Expected costs use the sixty-node rule on the axes of the lifetime block and equation (30.9).

30.3 The answers (step 4)

The questions of Table 30.2 are answered from $(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of equation (30.8), by a Gauss–Hermite rule, by draws of $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ or by the delta method, and that is the value reported. The exact route of Section 30.2.2 gives the value beside it, with the gap.

30.3.1 The expected cost of each intervention

Engineering question. What does each intervention cost per string-year?

General form. The expected consequence of an action, equation (2.23) of Chapter 2:

$$\mathcal{R}(a) = \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}] = \int L(a, \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}. \tag{30.11}$$

Specialized form. The question is equation (30.11) with a one of T, ∞ and τ, L the losses (30.4) and (30.6), and $p(\mathbf{z} \mid \mathbf{y})$ the normal of equation (30.8). One choice enters here: the expectation runs over what is not known, and the rate of each state stands for its long run, which holds for a fleet that renews many times and not for one string's next cycle, and which the simulation of Section 30.4.1 tests. For the age rule d integrates out at once, and the rule is that of equation (29.16) on the lifetime block,

$$\mathcal{R}(T) \approx \sum_m \sum_n \varpi_m \varpi_n\, L(T,\ \hat{\mathbf{z}}_{uv} + \sqrt{\epsilon_1}\, x_m \mathbf{e}_1 + \sqrt{\epsilon_2}\, x_n \mathbf{e}_2). \tag{30.12}$$

For run to failure the loss is linear in ν_r, so that $\mathcal{R}(\infty) = C_\mathrm{f} \bar{\nu}_\mathrm{r}$, with the rates of Section 29.3.1. It is not C_f over an expected life. The expected life is infinite on both routes, as Section 29.3.1 found, so that ratio would be zero: a formula that is finite at every state and useless averaged. The R\$ 10 091 of the first reading is $C_\mathrm{f} \nu_\mathrm{r}(\hat{\mathbf{z}})$, a plug-in. For condition-based replacement d integrates by the rule on d and the two blocks by the product rule, since $\mathbf{\Sigma}$ is block diagonal.

Building blocks. $\hat{\mathbf{z}}$ and $\mathbf{\Sigma}$ of equation (30.8) with its axes; the losses (30.4) and (30.6); the rates $\bar{\nu}_\mathrm{r} = 0.16905\,\mathrm{yr}^{-1}$ and $0.15285\,\mathrm{yr}^{-1}$ of Table 29.3 and equation (29.17); the five nodes of Table 30.3, and sixty for the reported values; for the values beside them, equation (30.9).

One state, by hand. At $\hat{\mathbf{z}}$ and $T = 4\,\mathrm{yr}$, $\Lambda = (4/6.4973)^{4.7257} = 0.1010$, so $F = 0.0961$, $R = 0.9039$ and $W = 3.9313\,\mathrm{yr}$, and

$$L(4, \hat{\mathbf{z}}) = \frac{15\,000 + 45\,000 \times 0.0961}{3.9313} = \frac{19\,324}{3.9313} = \mathrm{R\$}\,4\,915.$$

The rule of five nodes. Table 30.3 evaluates equation (30.12) at $T = 4\,\mathrm{yr}$. The rows are the nodes along the long axis, with the shape there, the columns the nodes along the short axis. Down a column the shape falls from 14.38 to 1.55 and the cost rises by half or more than doubles: a shallow hazard spreads failures before the fourth year, where the age rule cannot prevent them, and a steep one concentrates them after. Across a row the scale moves by about a third either way. The left column is high in every row, a short scale making failures before age four common. The weighted sum is R\$ 5 214, and sixty nodes give the same 5 213.78; it stands R\$ 299 above the center cell.

The expected cost by intervention. Table 30.4 gives the cost of each. Each row is an action; the columns give the loss at the MAP state, the expected cost and the central 95 % band of the loss on the state route, the same on the exact route, and the gap, state less exact. The last two rows carry condition-based replacement at each route's best interval on the common basis of equation (30.6); the optimum of that basis is 0.627 yr on the state route and 0.668 yr on the exact route, close to the 0.614 yr and 0.653 yr of Chapter 29. On its own accounting that chapter reports the minima R\$ 7 549 and R\$ 7 128 and, at the one interval $\tau = 0.614\,\mathrm{yr}$, R\$ 7 549 with the band [5 857, 10 019] and R\$ 7 138 with the band [4 820, 9 747]. The difference between the minima is the accounting of a caught life, R\$ 317 and R\$ 268; at the MAP state and

Table 30.3: The Gauss–Hermite rule (30.12) with five nodes: $L(4, \mathbf{z})$ [R$/string-year] at $(u, v)^{\mathsf{T}} = \hat{\mathbf{z}}_{uv} + 0.392\, x_m \mathbf{e}_1 + 0.105\, x_n \mathbf{e}_2$. The weights are $\varpi = 0.01126$, 0.22208, 0.53333, 0.22208, 0.01126 on each axis.

x_m	β	$x_n = -2.857$	-1.356	0	1.356	2.857
−2.857	14.38	7 675	4 235	3 811	3 757	3 751
−1.356	8.01	7 503	4 926	4 124	3 863	3 778
0	4.73	8 219	6 016	4 915	4 328	4 007
1.356	2.79	9 447	7 548	6 322	5 458	4 812
2.857	1.55	11 388	9 757	8 546	7 549	6 656
Weighted sum of the twenty-five cells: $\mathcal{R}(4) =$ R$ 5 214						

$\tau = 0.614$ yr it is R$ 7 716 on equation (30.6) against R$ 7 409, the excess Table 29.9 simulated. The row of the age rule with the annual test is at four years, its band from 3 000 of the draws. Figure 30.3 draws the age-rule rows as curves against T: the horizontal axis is the replacement age, the vertical the cost rate; the dashed curve is the loss at $\hat{\mathbf{z}}$, the heavy curve $\mathcal{R}(T)$ on the state route with its band shaded, the thin curve $\mathcal{R}_{\text{ex}}(T)$, and the two horizontal lines run to failure on the two routes, the asymptotes of those curves.

Table 30.4: The cost of each intervention [R$/string-year]: the loss at $\hat{\mathbf{z}}$; the expected cost and the 95 % band of the loss on the state and exact routes; the gap, state less exact. Draws: 2×10^5, seed 28.

Action	At $\hat{\mathbf{z}}$	$\mathcal{R}$, state	Band	$\mathcal{R}$, exact	Band	Gap
Age 2 yr	7 591	7 777	[7 500, 9 145]	8 169	[7 502, 11 366]	−393
Age 3 yr	5 408	5 679	[5 010, 7 733]	6 122	[5 035, 9 519]	−443
Age 4 yr	4 915	5 214	[3 858, 7 799]	5 576	[3 963, 9 081]	−363
Age 5 yr	5 525	5 870	[3 570, 9 140]	5 861	[3 704, 9 465]	9
Age 6 yr	6 920	7 326	[4 172, 11 658]	6 706	[3 825, 10 490]	621
Age 7 yr	8 523	8 721	[5 321, 12 306]	7 689	[3 928, 11 486]	1 032
Run to failure	10 091	10 143	[8 171, 12 443]	9 171	[4 070, 12 080]	972
Age 4 yr, annual test	6 356	6 525	[5 800, 8 044]	6 745	[5 860, 8 944]	−220
Condition, $\tau = 0.627$ yr		7 866	[6 220, 10 224]			
Condition, $\tau = 0.668$ yr				7 396	[4 753, 9 996]	470

The curves cross. Below five years the exact route is the dearer, by up to R$ 443, and above it the cheaper, by up to R$ 1 161 at eight years. The exact posterior has states of both kinds that the normal lacks: shallow shapes that bring failures before any young replacement age, and long lives in which a late age is seldom reached. Run to failure stands on the other side of the plug-in on each route, R$ 10 143 and R$ 9 171 against R$ 10 091. The band of the loss at four years is wide, over half the cost, and this says what the cost at the true state may be, not how firm a choice between two ages is.

30.3.2 The best age

Engineering question. At what age should a string be replaced?

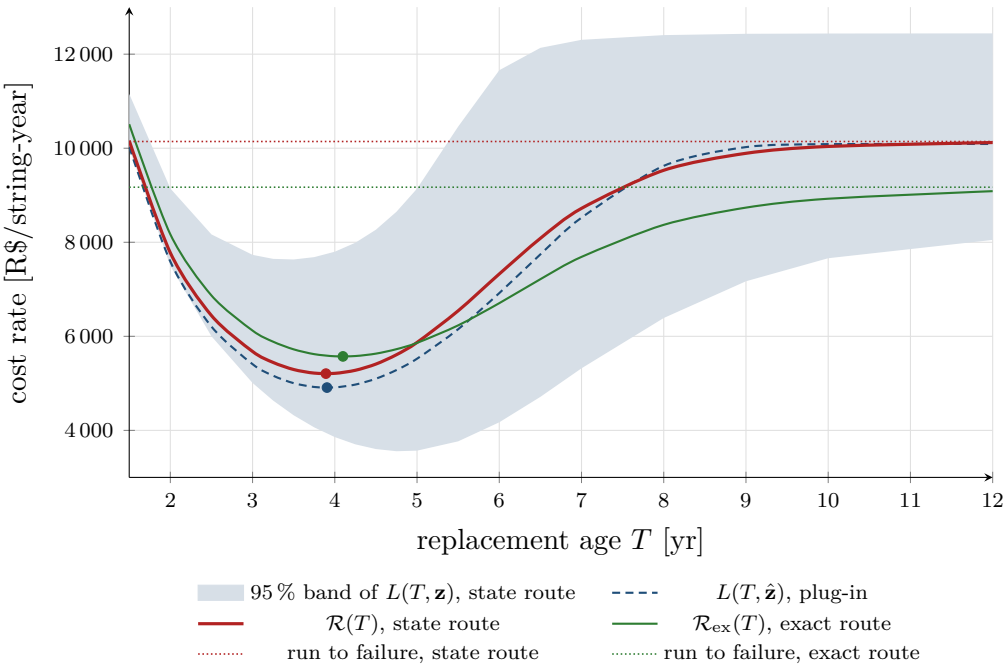


Figure 30.3: The cost rate against the replacement age: the loss at the MAP state, equation (30.4); the expected cost on the two routes, with minima marked; the band of the loss on the state route; and run to failure on the two routes. 2×10^5 draws, seed 28, for the band.

General form. The action that minimizes the expected consequence, equation (2.23):

$$a^* = \arg\min_a \mathcal{R}(a). \tag{30.13}$$

Specialized form. The question is equation (30.13) with $a = T > 0$ and $\mathcal{R}$ of equation (30.12). One choice enters here: the age is a continuous number, which makes the minimum a stationary point, and it fails if replacements can be scheduled only at outages, which the range below covers. The expectation has no root formula of the form of equation (30.3), since that condition holds for one state. Differentiating under the integral gives the stationary condition

$$\mathbb{E}\left[\frac{\partial L(T, \mathbf{z})}{\partial T} \,\middle|\, \mathbf{y}\right] = \mathbb{E}\left[\frac{R(T; \mathbf{z})\left(\Delta C\, \lambda(T; \mathbf{z}) W(T; \mathbf{z}) - C_{\text{r}} - \Delta C\, F(T; \mathbf{z})\right)}{W(T; \mathbf{z})^2} \,\middle|\, \mathbf{y}\right] = 0, \tag{30.14}$$

which weights each state's departure from its own condition by R/W^2, and which is solved on the rule by bracketing. Each state's bracket vanishes at its own $T^*(\mathbf{z})$, and the Bayes action is where the weighted average vanishes.

Building blocks. $\mathcal{R}(T)$ of equation (30.12) on sixty nodes, equation (30.9) beside it.

The optimum. On the state route

$$T^{\text{B}} = 3.892\,\text{yr}, \qquad \mathcal{R}(T^{\text{B}}) = \text{R\$}\,5\,207,$$

and on the exact route $T^{\text{B}}_{\text{ex}} = 4.099\,\text{yr}$ at R$ 5 573. The gap is 0.207 yr, about eleven weeks, and R$ 365. Taking the other route's action costs little: on the exact route 3.892 yr costs R$ 5 590, R$ 17 above its minimum, and on the state route 4.099 yr costs R$ 5 231, R$ 24 above its own. The gap in the action does not change the decision.

The shallow minimum. The curvature at the minimum is 1 117 R$ per string-year per yr^2 on the state route and 780 on the exact one, and the range within 1 % of the minimum is [3.59, 4.20] yr and [3.73, 4.48] yr, which overlap over [3.73, 4.20] yr. Table 30.5 prices the ages a schedule might offer. Each row is an age; the columns give the cost above the minimum on each route. Four years is within R$ 6 on both. Five years costs R$ 663 on the state route, 13 %, and R$ 289 on the exact route, 5 %.

Table 30.5: The expected cost above the minimum at the ages a schedule might offer [R$/string-year].

T [yr]	3	3.5	4	4.5	5	6
State route	472	87	6	203	663	2 119
Exact route	550	151	4	60	289	1 133

30.3.3 The age rule against running to failure

Engineering question. Is the age rule better than running to failure, with assurance?

General form. A tail of a derived quantity under the posterior, equation (2.12):

$$\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}) = \int \mathbf{1}[f(\mathbf{z}) \geq f^*]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}. \tag{30.15}$$

Specialized form. The question is equation (30.15) with $f = L(\infty, \mathbf{z}) - L(T^{\mathrm{B}}, \mathbf{z})$, the saving at the state, and $f^* = 0$, and the posterior of equation (30.8). One choice enters here: the same age is held across the states, since the rule must be chosen before the state is known. Its gradient, from equation (30.5) and the gradient of ν_{r} in equation (29.10), gives the delta form

$$\mathbf{g}_{\mathrm{sav}} = C_{\mathrm{f}} \nu_{\mathrm{r}} \begin{bmatrix} -1 \\ \psi(1 + \mathrm{e}^{-v})\mathrm{e}^{-v} \\ 0 \end{bmatrix} - \mathbf{g}_{T^{\mathrm{B}}},$$

and the draws carry the tail, because the delta form misses what the states below say.

Building blocks. T^{B} of Section 30.3.2, $\mathbf{g}_T$ of equation (30.5), $\boldsymbol{\Sigma}$, and the draws of Section 30.2.3.

The numbers. At $\hat{\mathbf{z}}$ the saving is R$ 5 180 and $\mathbf{g}_{\mathrm{sav}} = (-5\,236,\ 1\,964,\ 0)^{\mathsf{T}}$, so sd $= \sqrt{\mathbf{g}_{\mathrm{sav}}^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g}_{\mathrm{sav}}} =$ R$ 1 118 and the normal tail is $\Phi(4.63)$, above 0.99999. By draws, the saving has the mean R$ 4 936, the median R$ 4 999 and the band [2 380, 7 027], and the probability that the age rule costs less at the true state is 0.999. On the exact route, at its own action, the mean is R$ 3 598, the median R$ 4 134, the band [−2 271, 6 168] and the probability 0.931, a gap of 0.068 and of R$ 1 338 in the mean. The states where the age rule loses are long-lived with a shallow shape: on the exact route their median mean life is 12.6 yr and their median shape 1.27, and three quarters of them have a shape above one, so a rising hazard is not by itself enough for a fixed age of four to win. The rule stands on both routes, and how sure it is differs: one chance in a thousand of losing on one, about seven in a hundred on the other.

30.3.4 The plug-in action, the Bayes action and the optimum at a state

Engineering question. How different are the plug-in optimum and the Bayes action, and does it matter?

General form. The expected cost of one action above the minimum, and the quantiles of a derived quantity, equations (2.23) and (2.11):

$$\mathcal{R}(a') - \min_a \mathcal{R}(a), \qquad \Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y}) = \int \mathbf{1}[f(\mathbf{z}) \le f^*]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}. \tag{30.16}$$

Specialized form. The question is equation (30.16) with $a' = T^*(\hat{\mathbf{z}})$, $\min_a \mathcal{R}(a) = \mathcal{R}(T^{\mathrm{B}})$ and $f(\mathbf{z}) = T^*(\mathbf{z})$ of equation (30.10). No choice enters. The plug-in action reads equation (30.10) at $\hat{\mathbf{z}}$: $g(4.7257) = 0.6010$, so $T^*(\hat{\mathbf{z}}) = 6.4973 \times 0.6010 = 3.905\,\mathrm{yr}$. There equation (30.3) holds with $\lambda = 0.10912\,\mathrm{yr}^{-1}$, $W = 3.8449\,\mathrm{yr}$ and $R = 0.9138$, and $L = 45\,000 \times 0.10912 = \mathrm{R\$}\,4\,910$. The distribution of the optimum comes through the draws; with the gradient of $\ln T^*(\mathbf{z})$,

$$\mathbf{g}_{\ln T^*} = \begin{bmatrix} 1 \\ \partial \ln g / \partial v \\ 0 \end{bmatrix} = \begin{bmatrix} 1 \\ 0.2373 \\ 0 \end{bmatrix}, \qquad \mathrm{sd}(\ln T^* \mid \mathbf{y}) = \sqrt{\mathbf{g}_{\ln T^*}^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g}_{\ln T^*}} = 0.1184, \tag{30.17}$$

the delta form gives $[3.10,\ 4.93]\,\mathrm{yr}$ for the band.

Building blocks. g and the root of equation (30.10); $\mathcal{R}$ on each route; the draws of Section 30.2.3.

The three objects. Table 30.6 sets them apart. Each row is an object; the columns give its age and, where it has one, its cost, on each route. Medians stand beside medians and minimizers beside minimizers. The plug-in action is one number on both routes, and its row gives the plug-in cost and then what that action costs in expectation on each route. Figure 30.4 draws the distribution of the optimum: the horizontal axis is the age, the vertical the probability in each quarter-year bin, and the vertical lines the plug-in action and the Bayes action of each route.

Table 30.6: The three objects called an optimum. Ages in years, costs in R$/string-year; 2×10^5 draws, seed 28, for the distribution.

Object	State route	Exact route
Plug-in action $T^*(\hat{\mathbf{z}})$	3.905	3.905
Its loss at $\hat{\mathbf{z}}$ (one state)	4 910	4 910
Its expected cost $\mathcal{R}(T^*(\hat{\mathbf{z}}))$	5 207.4	5 587.5
Bayes action	$T^{\mathrm{B}} = 3.892$	$T^{\mathrm{B}}_{\mathrm{ex}} = 4.099$
Its expected cost	5 207.3	5 572.5
Regret of the plug-in action	0.1	15.0
Bayes action less plug-in action	-0.013	$+0.194$
Median of $T^*(\mathbf{z})$	3.957	4.170
Band of $T^*(\mathbf{z})$	[3.18, 5.00]	[3.16, 57]
$\Pr(T^*(\mathbf{z}) >$ the route's Bayes action)	0.558	0.545
No finite optimum ($\beta \le 1$)	3×10^{-5}	0.016

The upper end of the exact band is a far-tail quantile of a heavy tail: under other seeds it lies between about 50 yr and 57 yr.

Whether uncertainty pulls the action earlier. It does not, as a general statement. The state route's action lies 0.013 yr below the plug-in action, under a week; the exact route's lies 0.194 yr above it. An expansion on the state route gives the small size of the first and not its

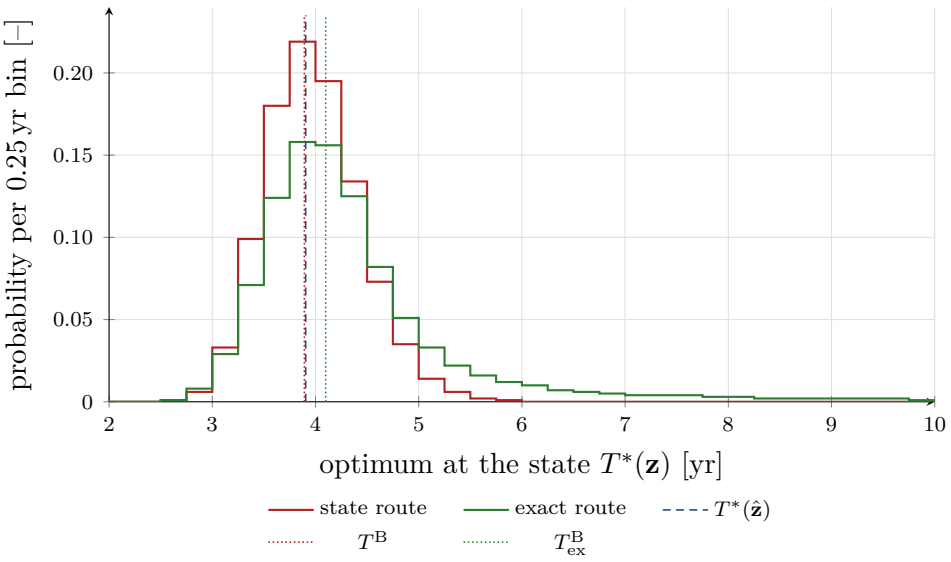


Figure 30.4: The posterior distribution of the optimum at a state, equation (30.10), on the two routes, with the plug-in action and the Bayes actions. 2×10^5 draws, seed 28. On the exact route 5.4 % of the probability lies above 10 yr, of which 1.6 % has no finite optimum; off the axis.

sign: with $\mathbf{H}_T$ the Hessian of $L(T, \mathbf{z})$ in (u, v), the shift to second order is $-\frac{1}{2}\operatorname{tr}(\partial_T \mathbf{H}_T\, \boldsymbol{\Sigma})/\partial_T^2 L$, where $\frac{1}{2}\operatorname{tr}(\partial_T \mathbf{H}_T\, \boldsymbol{\Sigma}) = -9.2$ R$ per string-year per year against the curvature 1 113 at $\hat{\mathbf{z}}$, or $+0.008$ yr, the opposite sign to the actual -0.013 yr. Nor does the shape of the curve argue for an earlier age. The curve is convex near its minimum and steeper on the early side: one year early costs more than one year late, R$ 616 against R$ 516 at $\hat{\mathbf{z}}$, R$ 600 against R$ 543 on the state route and R$ 447 against R$ 352 on the exact route. If anything that asymmetry favors the later side. On the exact route the later action comes from a different posterior, not from curvature. Holding only the states with $T^*(\mathbf{z}) < 7$ yr, 92 % of the mass, the action is 4.016 yr, so the long-lived states supply a little under half of the shift, and the wider mass of later optima supplies the rest. The three priors of Section 30.4.3 repeat the pattern: the state route's action is below the plug-in action in each, the exact route's above it.

Whether it matters. For the choice of age, no, on either route: the regret of the plug-in action is R$ 0.1 and R$ 15. For the cost, yes, on both: the plug-in's R$ 4 910 understates what that same age costs in expectation by R$ 297 on the state route, 6 %, and by R$ 677 on the exact one, 14 %. The action survives the plug-in and its number does not. The median of the optimum is not the action on either route: it stands 0.065 yr and 0.071 yr above the Bayes action, and the mean of $T^*(\mathbf{z})$ does not exist, since states of no finite optimum have positive probability on both.

30.3.5 The value of knowing the state

Engineering question. How much would knowing the lifetime parameters exactly save?

General form. The expected value of sample information, equation (2.51), with a measurement that reveals the state:

$$\mathrm{EVPI} = \min_a \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}] - \mathbb{E}\big[\min_a L(a, \mathbf{z}) \mid \mathbf{y}\big] \geq 0. \tag{30.18}$$

It is nonnegative because for every state $\min_a L(a, \mathbf{z}) \leq L(a', \mathbf{z})$; averaging gives $\mathbb{E}[\min_a L] \leq \mathcal{R}(a')$ for every a', hence for the best.

Specialized form. The question is equation (30.18) with the action set first the ages and run to failure, then all three interventions, and $\mathbf{z}$ revealed whole or by block. One choice enters here: a state revealed now serves every future string, with no cost of learning and no replacement of the strings in service, so this is a ceiling and not a plan. On the age family, equation (30.3) gives the inner minimum in closed form at each state,

$$L^*(\mathbf{z}) = \min_T L(T, \mathbf{z}) = \begin{cases} \Delta C\, \lambda(T^*(\mathbf{z}); \mathbf{z}) = \Delta C\, \dfrac{\beta}{\eta}\, g(\beta)^{\beta-1}, & \beta > 1, \\ C_\mathrm{f} \nu_\mathrm{r}(u, v), & \beta \le 1, \end{cases} \tag{30.19}$$

and $\mathrm{EVPI}_{uv} = \mathcal{R}(T^\mathrm{B}) - \mathbb{E}[L^*(\mathbf{z}) \mid \mathbf{y}]$ by the rule. With all three interventions the inner minimum is taken with L_c as well, and reveals d too; revealing d alone leaves the life averaged in each loss. With the loss nearly quadratic in T around each state's optimum and $T^*(\mathbf{z})$ linearized at $\hat{\mathbf{z}}$, the state route has a local form,

$$\mathrm{EVPI}_{uv} \approx \tfrac{1}{2}\, \frac{\partial^2 L(T, \hat{\mathbf{z}})}{\partial T^2}\Big|_{T^*(\hat{\mathbf{z}})} T^*(\hat{\mathbf{z}})^2\, \mathbf{g}_{\ln T^*}^\mathsf{T} \boldsymbol{\Sigma} \mathbf{g}_{\ln T^*}, \tag{30.20}$$

which shows the two things that set the value near the fit: how wide the optimum is, by the delta width, and how curved the loss. It is no variance of $T^*(\mathbf{z})$, which does not exist.

Building blocks. T^B and $\mathcal{R}$ of Section 30.3.2; g of equation (30.10); L_c of equation (30.6); $\mathbf{g}_{\ln T^*}$ of equation (30.17); the draws of Section 30.2.3.

The values. At the MAP state equation (30.19) reads $45\,000 \times 0.10912 = \mathrm{R\$}\,4\,910$. Over the rule $\mathbb{E}[L^*(\mathbf{z}) \mid \mathbf{y}] = \mathrm{R\$}\,5\,095$, against R\$ 5 284 on the exact route, so that $\mathrm{EVPI}_{uv} = 5\,207 - 5\,095 = \mathrm{R\$}\,112$, 2.2 % of the cost, and on the exact route R\$ 288, 5.2 %. The local form (30.20), with the curvature 1 113 at $\hat{\mathbf{z}}$ and the delta width $3.905 \times 0.1184 = 0.462$ yr, gives R\$ 119. Table 30.7 lists the values. Each row is what is revealed and among which actions; the columns give the expected cost after it and its value on each route. Figure 30.5 draws the cost of the best action against the information state: the horizontal positions are run to failure with no age rule, the action under the posterior, the life known, and the whole state known; the vertical axis is the expected cost of the best action there; the EVPI is the drop from the second position.

Table 30.7: The expected value of perfect information [R$/string-year], with the expected cost of the best action after the revelation, on the two routes. The rows with three interventions use L_c on the common basis.

Revealed, action set	State, after	State, value	Exact, after	Exact, value
Nothing (the decision now)	5 207	–	5 573	–
Life (u, v), ages and run to failure	5 095	112	5 284	288
Life (u, v), three interventions	5 088	120	5 238	334
Median delay d, three interventions	5 207	0	5 571	1
All three entries, three interventions	5 076	131	5 209	364

Whether the uncertainty barely matters. On the state route the value is 2.2 %, and the local form says why: the optimum spreads by under half a year on a curve that forgives half a year. On the exact route the value is more than two and a half times that, and the difference has one source. Of the R\$ 288, R\$ 140 comes from the 4.6 % of states whose optimum is above twelve years or absent, against under R\$ 1 from such states on the state route, and R\$ 75 from the states with an optimum between three and five years against R\$ 92 on the state route. Both

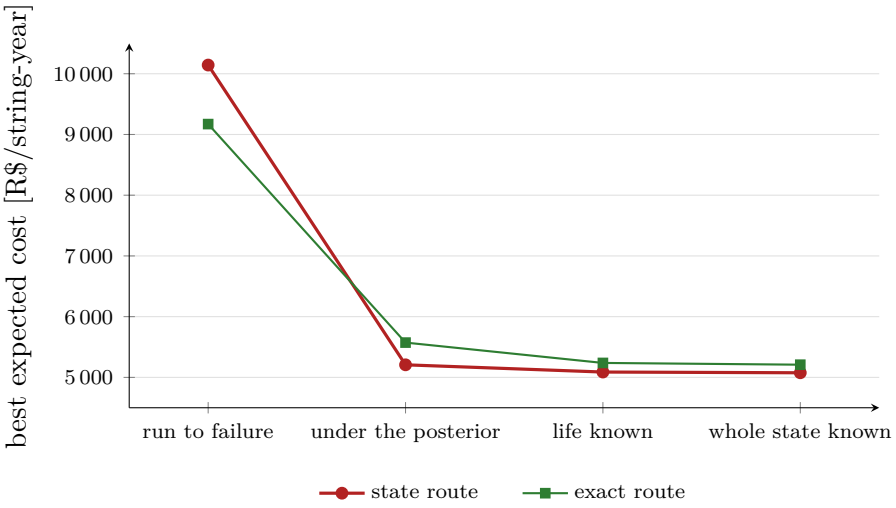


Figure 30.5: The expected cost of the best action against the information state. From the second position on, the drop is the value of perfect information: R$ 131 on the state route and R$ 364 on the exact route for the whole state; the first drop is what adopting the action under the posterior saves against running to failure.

values are stated: 2.2 % and 5.2 %, or over twelve strings about R$ 1 350 and R$ 3 460 a year as a ceiling, and the chapter does not choose between them. What holds on both routes is the comparison with the first figure: the value of perfect information is small against the R$ 3 598 to R$ 4 936 the rule itself saves, and the plug-in's error in the cost is larger than the value on each route, R$ 297 against R$ 112 and R$ 677 against R$ 288.

The action set. The value of learning the delay is R$ 62 (R$ 48) in Table 29.8, where testing is the only intervention, and R$ 0 (R$ 1) here. With the age rule available the median delay almost never changes the action. On the state route no median makes testing cheaper than R$ 5 207: its best cost, at a median of about 2 yr, is R$ 5 360, and it rises at longer medians, since a long delay catches a string early and gives up its life. On the exact route testing wins above a median of about 1.0 yr, of prior probability 0.009. A value of information belongs to a question and to the set of actions it could change.

30.3.6 The three interventions

Engineering question. Which intervention wins, and by how much?

General form. The comparison of expected consequences with the ranking at the true state, equations (2.23) and (2.24):

$$\mathcal{R}(a) - \mathcal{R}(a'), \qquad \Pr(L(a, \mathbf{z}) < L(a', \mathbf{z}) \mid \mathbf{y}). \tag{30.21}$$

Specialized form. The question is equation (30.21) with a and a' the minimizers of the three families, on the common basis of equation (30.6), and the posterior of equation (30.8). One choice enters here: each intervention is taken alone, which is what makes the three comparable, and the annual test kept beside the age rule is priced after them. The third entry of the state acts only through L_c.

Building blocks. The minima of Table 30.4, L_c, and the draws of Section 30.2.3.

The comparison. Table 30.8 gives the three at their optima. Each row is an intervention; the columns give the best parameter, the expected cost, the band and the probability of being the cheapest of the three at the true state, on each route. Figure 30.6 draws the same costs: each horizontal line is a band, each dot an expected cost, the upper of each pair on the state route; the fourth pair is the favorable case of the table after it.

Table 30.8: The three interventions on one cost basis [R$/string-year]: the expected cost at the best parameter, the 95 % band of the loss and the probability of being the cheapest at the true state.

	State route				Exact route			
Intervention	Best	$\mathcal{R}$	Band	Cheapest	Best	$\mathcal{R}$	Band	Cheapest
Age replacement	3.892 yr	5 207	[3 941, 7 740]	0.972	4.099 yr	5 573	[3 905, 9 082]	0.858
Condition-based	0.627 yr	7 866	[6 220, 10 224]	0.028	0.668 yr	7 396	[4 753, 9 996]	0.096
Run to failure		10 143	[8 171, 12 443]	0.000		9 171	[4 070, 12 080]	0.045

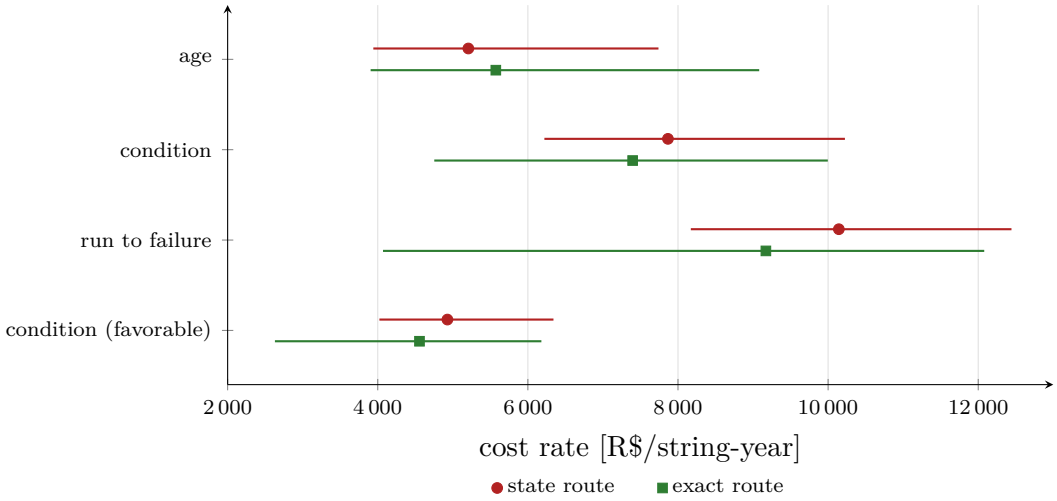


Figure 30.6: The expected cost of each intervention at its best parameter (dots) and the 95 % band of the loss at the true state (bars), from the draws of seed 28. The lowest pair is condition-based replacement with twice the median delay and half the test cost.

The age rule wins on both routes: by R$ 2 659 against condition-based replacement on the state route and by R$ 1 823 on the exact one, and by R$ 4 936 and R$ 3 598 against run to failure. On the accounting of Chapter 29 the margin over testing would be R$ 2 342 and R$ 1 555; the common basis is what widens it. The probability that the age rule is cheaper than condition-based replacement at the true state is 0.972 (0.867). At the MAP state a string replaced at four years is taken out before its failure with probability 0.90 and no test, giving up about two years of mean life; testing pays $2\,000/0.627 = $ R$ 3 190 per string-year before it replaces anything, and catches about three failures in four.

The annual test kept. The rates of Table 30.8 assume that the age rule drops the annual test. Keeping it, L_a of step 2, moves the best age to 4.52 yr (5.02 yr), because a test now catches some of the late failures, at R$ 6 418 (R$ 6 531). The test thus costs a net R$ 1 211 (R$ 958) per string-year, the R$ 2 000 less the failures it prevents, and the age rule alone is the cheaper of the two at the true state with probability 0.978 (0.949) at four years. Age with the test still

beats condition-based replacement, by R$ 1 448 (R$ 865) against the R$ 2 659 (R$ 1 823) of the rule without it: the margin of the table is conditional on dropping the test.

When condition-based replacement wins. Table 30.9 varies the two quantities the chapter on inspection left to the manufacturer and the site. Each row is a case, with the median delay and the test cost multiplied as stated and every spread unchanged; the columns give the best interval and the cost on the common basis on each route, and the margin against the age rule, which the case does not alter, positive when testing wins.

Table 30.9: Condition-based replacement under other delays and test costs, on the common basis [R$/string-year]. The margin is $\mathcal{R}(T^{\mathrm{B}})$, 5 207 on the state route and 5 573 on the exact route, less the cost of the case.

	State route			Exact route		
Case	τ [yr]	$\mathcal{R}$	Margin	τ [yr]	$\mathcal{R}$	Margin
As in the chapter	0.627	7 866	−2 659	0.668	7 396	−1 823
Twice the delay	0.910	6 241	−1 034	0.945	5 818	−245
Half the test cost	0.449	5 918	−711	0.467	5 547	26
Three times the delay	1.207	5 770	−563	1.242	5 316	257
Twice the delay, half the test cost	0.720	4 928	279	0.741	4 556	1 017
A quarter of the test cost	0.356	4 635	572	0.366	4 310	1 263

With twice the delay and half the test cost, condition-based replacement wins on both routes, by R$ 279 and by R$ 1 017; its band is [4 022, 6 341] ([2 630, 6 180]) and it is cheaper than the age rule at the true state with probability 0.564 (0.750). On the accounting of Chapter 29 it would cost R$ 4 388 (R$ 4 101), at 0.706 yr (0.728 yr), and that accounting overstates the margin. Half the test cost alone, or three times the delay alone, falls on opposite sides: the state route keeps the age rule and the exact route takes testing, by a margin of R$ 26 in the first and R$ 257 in the second. Both are stated and neither is chosen. The test cost at which the two tie is R$ 709 on the state route and R$ 1 011 on the exact route at the present delay, and R$ 1 194 and R$ 1 787 at twice the delay. A long delay and a cheap test favor testing as expected, and the wide lifetime spread does too, by raising the age rule's cost while leaving the rate almost where it was, which Section 30.4.3 shows.

30.4 The checks (step 5)

The answers rest on a renewal–reward ratio, a new accounting for the catch, a flat prior on the shape and the reading of a failed test. Each is set against a fuller calculation or varied. There is no physics row, so no limit $\sigma_p \to 0$ arises.

30.4.1 The cost rate against a simulated population

A population of 10^4 strings is simulated at the MAP state (numpy `default_rng`, seed 30), each cycle ended at $\min(T, T)$ and the total cost divided by the total time. Table 30.10 sets it against the formula. Each row is an action; the columns give $L(a, \hat{\mathbf{z}})$, the simulated rate with its standard error, and, for the age rule, the simulated and formula fraction of failures. Below it, 4×10^5 strings with delays and test phases (seed 31) test the common basis of equation (30.6).

All agree within about one and a half standard errors, and the simulated values of the lower block agree within R$ 9, and with the 7 690 of Table 29.9 within about one of its standard errors. Treating the onset as positive costs nothing visible at these delays.

Table 30.10: The cost rate by simulation at the MAP state against the closed forms [R$/string-year], with L_{c} of equation (30.6) in the lower block.

Action	Closed form	Simulated (s.e.)	$F(T;\hat{\mathbf{z}})$	Simulated fraction
Age 3 yr	5 408	5 383 (24)	0.0256	0.0242
Age 3.905 yr	4 910	4 909 (35)	0.0862	0.0862
Age 5 yr	5 525	5 507 (46)	0.2517	0.2504
Run to failure	10 091	10 128 (24)		
Condition, $\tau = 0.5$ yr	7 809	7 811		
Condition, $\tau = 0.614$ yr	7 716	7 725		
Condition, $\tau = 1$ yr	8 286	8 282		
Age 4 yr, annual test, 2×10^6 strings	6 356	6 360 (2)		

30.4.2 The routes and the quadrature

Table 30.11 sets each computed quantity beside an independent calculation. Each row is one check; the columns give the quantity, its value by its route, the other calculation and the verdict.

Table 30.11: The checks of the routes and of the quadrature.

Quantity	Route	Independent calculation	Verdict
$\mathcal{R}(4)$, state	five nodes: 5 213.78	sixty nodes: 5 213.78; draws: 5 213.9 (s.e. 2.4)	agree
$\mathcal{R}(4)$, state	sixty nodes	second order (2.24): 5 224.7	+11
$\mathcal{R}(5)$, state	sixty nodes: 5 870.1	second order: 5 839.7	−30
$\mathcal{R}_{\mathrm{ex}}(4)$	equation (30.9): 5 576.4	draws: 5 574.1 (s.e. 3.1); finer grid: 5 576.4	agree
$\mathcal{R}(\infty)$, state and exact	10 143, 9 171	Table 29.4, no test	agree
$L^*(\hat{\mathbf{z}})$	equation (30.19): 4 910.37	$L(T^*(\hat{\mathbf{z}}),\hat{\mathbf{z}})$ on a grid: 4 910.37	agree
EVPI_{uv}, state	rule: 112.2	draws: 111.6; local form (30.20): 119	agree; +7
EVPI_{uv}, exact	equation (30.9): 288.2	draws: 289.4	agree

The second-order form is good to a fraction of a percent near the optimum and its error changes sign by five years, where the posterior's long axis reaches the steep part of the loss.

30.4.3 The prior on the shape and the failure ages

Table 30.12 repeats the decision under the two proper priors on the shape of Section 10.4.4, refitting each posterior as a check. Each row is a quantity; the columns give it under the flat prior and the two proper priors, the state route with the exact route in parentheses.

The action moves little on the state route, and on the exact route it moves from 4.10 to 5.31 years. The value of knowing the life grows with the weight the prior puts on shallow shapes, to 8.7 % and 12 % of the cost under the narrowest prior, where the statement that the uncertainty barely matters holds on neither route. And the winner changes: under $\mathcal{N}(0, 0.5^2)$ on the exact route condition-based replacement is cheaper than the age rule, by R$ 197. On the state route the age rule keeps it, by R$ 143, and on the accounting of Chapter 29 (R$ 7 036) it would not. This is the wide lifetime spread in numbers: from the flat prior to the narrowest the age rule's cost rises by 37 % (19 %) while condition-based replacement falls by 7 % (13 %). That the age rule wins is thus a statement about a shape near the record's; with the shape's median near two, as under this prior, the ranking is not secure.

Table 30.12: The decision under other priors on $v = \ln \beta$, state route with exact in parentheses. Ages in years, costs in R$/string-year, condition-based on the common basis.

Quantity	Flat	$\mathcal{N}(0, 1^2)$	$\mathcal{N}(0, 0.5^2)$
Plug-in action at the refit $T^*(\hat{\mathbf{z}})$	3.905	3.828	4.409
Bayes action	3.892 (4.099)	3.803 (4.220)	4.077 (5.314)
Age rule $\mathcal{R}$	5 207 (5 573)	5 771 (6 025)	7 155 (6 659)
Condition-based $\mathcal{R}$	7 866 (7 396)	7 790 (7 151)	7 298 (6 462)
Run to failure $\mathcal{R}$	10 143 (9 171)	9 981 (8 679)	8 963 (7 362)
EVPI_{uv}	112 (288)	156 (441)	619 (803)

The reading of a failed test moves the level as well. The interval fit of Section 10.4.5, $\hat{\beta} = 3.93$ and $\hat{\eta} = 6.41$ yr, read at its point, gives $T^* = 3.70$ yr and R$ 5 504, against 3.905 yr and R$ 4 910 at the fit of the chapter on the same plug-in terms.

30.5 The report (step 6)

The fitted Weibull promised R$ 4 910 at the plug-in age, and the expected cost of that same age is R$ 5 207 on the posterior of the state and R$ 5 588 on the exact posterior. The age itself barely moves, 3.89 (4.10) years against 3.90, so the plug-in's action survives and its number does not. The saving against running to failure is R$ 4 936 (R$ 3 598) per string-year, not about R$ 5 180. The answers follow in the order they were asked.

30.5.1 The answers and their bands

Table 30.13 is the report. Each row is a question in the manager's words; the second column gives the answer from the posterior of the state, the third the exact-route value, the fourth the gap, the last the band or the qualification.

The report line. On a Weibull lifetime fitted to twelve strings with four failures, the expected cost of an age rule is least at 3.9 years on the posterior of the state and 4.1 years on the exact posterior, about R$ 5 200 (R$ 5 600) per string-year, within 1 % of the minimum at four years on both. The rule saves R$ 4 936 (R$ 3 598) per string-year against running to failure, with probability 0.999 (0.931) of saving at the true state. It costs R$ 2 659 (R$ 1 823) less than condition-based replacement on the same accounting, and that holds unless the test costs under about R$ 700 (R$ 1 000), or the delay is much longer than the manufacturer states and the test cheaper, on the exact route three times the delay alone sufficing by R$ 257, or the shape is much shallower than the record's fit. Those rates carry no test: keeping the annual test adds R$ 1 211 (R$ 958) per string-year, and it remains cheaper than condition-based replacement by R$ 1 448 (R$ 865). Knowing the whole state exactly would save at most R$ 131 (R$ 364) per string-year, 2.5 % (6.5 %).

30.5.2 The replacement decision

Replace at four years, with the annual test dropped or justified separately. Both routes put it within R$ 6 of their minimum, and it is the one age a scheduler can choose without choosing a route. The gap in the action is two and a half months and changes nothing a calendar can see; the gap in the level, R$ 363 at four years, changes the budget figure by 7 % and not the choice. Where the gap does change a statement is in what follows from a later age, five years costing R$ 663 on one route and R$ 289 on the other, and in the value of information and the winner under cheaper tests, each stated with both values. A rule for new strings does not by

Table 30.13: The report of the replacement question: the answers from $\mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$, the exact route with its gap, and the bands, in R$/string-year and years.

Engineering question	Answer (state route)	Exact route	Gap	Band, qualification
What does each intervention cost per string-year?	age 4 yr: 5 214; run to failure: 10 143; condition, 0.627 yr: 7 866	5 576; 9 171; 0.668 yr: 7 396	−363; 972; 470	age 4 yr [3 858, 7 799] ([3 963, 9 081]); condition 7 549 (7 128) on the accounting of Chapter 29
At what age should a string be replaced?	$T^{\mathrm{B}} = 3.892$ yr, R$ 5 207	4.099 yr, R$ 5 573	0.207 yr, 365	within 1 % over [3.59, 4.20] yr ([3.73, 4.48] yr); 3.80 to 4.08 yr (4.10 to 5.31 yr) under the priors
Is the age rule better than running to failure, with assurance?	probability 0.999; saves R$ 4 936, median R$ 4 999	0.931; R$ 3 598, median R$ 4 134	0.068; 1 338	saving [2 380, 7 027] ([−2 271, 6 168])
How different are the plug-in optimum and the Bayes action, and does it matter?	3.905 yr against 3.892 yr; regret R$ 0.1; plug-in cost low by R$ 297	4.099 yr; regret R$ 15; low by R$ 677	0.207 yr; 15; 380	$T^{*}(\mathbf{z})$ median 3.957 yr (4.170 yr), band [3.18, 5.00] yr ([3.16, 57] yr); action earlier on one route and later on the other
How much would knowing the lifetime parameters exactly save?	R$ 112, 2.2 %; whole state, three interventions: R$ 131	R$ 288, 5.2 %; R$ 364	176; 233	R$ 156 and R$ 619 (R$ 441 and R$ 803) under the priors; the delay alone 0 (R$ 1)
Which intervention wins, and by how much?	age, by R$ 2 659 over condition and R$ 4 936 over run to failure; cheapest with probability 0.972	age, by R$ 1 823 and R$ 3 598; 0.858	836; 1 338; 0.114	condition wins with twice the delay and half the test cost, on both routes; under $\mathcal{N}(0, 0.5^2)$ on the exact route; half the test cost alone, or three times the delay alone, on the exact route only; rates without the annual test, which adds 1 211 (958)

itself settle the strings in service. Strings I, J, K and L, at 4.6 to 6.2 years, are past four years and past both routes' actions. String H, at 3.9 years, is short of four years and of the exact route's action, and at the state route's. The stationary rate does not price the transition: each of them has the conditional survival of Chapter 10 rather than a new string's life, and its order, oldest first, stands. The rate is what the program settles into.

Principle 30.1. *An action is chosen by minimizing the posterior expectation of its loss, and interventions of different kinds are compared only on one accounting of a cycle. The plug-in action may be nearly the Bayes action while the plug-in cost is not its expected cost. The value of perfect information is the expected regret of the Bayes action. Near the fit the width of the optimum and the curvature of the loss set it; far states may supply much of it, as half of it on the exact route here; and it changes with the actions on the table. It is reported with its route, and it bounds what any record is worth for that decision.*

Common mistake

Five errors recur with replacement rules. Reporting the plug-in cost as the cost of the action: R$ 4 910 where the same age costs R$ 5 207 on the state route and R$ 5 588 exactly. Dividing the failure cost by an average life: the expected life is infinite on both routes, while the rate averages to 0.16905 (0.15285). Reporting the median of the optimum as the decision: 3.957 yr (4.170 yr) against the minimizer 3.892 yr (4.099 yr). Saying that uncertainty pulls the action earlier: on this record it holds, by under a week, on the state route alone, and the action moves later by 0.19 yr on the exact route. And comparing an age rule with the inspection cost of a life run to its failure: that accounting flatters condition-based replacement by R$ 317 (R$ 268) as it stands here and overstates its margin in the favorable case.

30.6 The next measurement (step 7)

The values of Table 30.7 bound every measurement of the life for this decision. This section prices the measurements the site could take. The full treatment of such values is the subject of Chapter 31, in Volume III.

30.6.1 The candidate rows on the life

The next tests. The annual test of the eight survivors adds the pass-or-fail rows of equation (10.38), with $\mathbf{\Lambda}_a$ of Section 10.6.1, and two annual tests add the second-year rows weighted by the probability of reaching them; the preposterior of equation (2.17) gives $\mathbf{\Sigma}_a = (\mathbf{\Sigma}^{-1} + \mathbf{\Lambda}_a)^{-1}$ on the lifetime block. The annual test serves this decision through the scale, and its cost is taken up in the recommendation.

A shape figure from the manufacturer's life tests. A figure for the shape is a Gaussian row $(v - y_v)/\sigma_v$, $\mathbf{\Lambda}_a = \mathrm{diag}(0,\ 1/\sigma_v^2)$, acting on the v column alone, which the record covers least. The candidates are spreads of $\sigma_v = 0.2$ and 0.1.

30.6.2 The preposterior and the value

Table 30.14 lists the candidates. Each row is a candidate; the columns give the spreads under $\mathbf{\Sigma}_a$, the spread of $\ln T^*$, $\sqrt{\mathbf{g}_{\ln T^*}^{\mathsf{T}} \mathbf{\Sigma}_a \mathbf{g}_{\ln T^*}}$ of equation (30.17), and the expected value of sample information of equation (2.51) among the three interventions, on the state route and on the exact route. On the state route the outcomes are those of the preposterior mean, $\mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma} - \mathbf{\Sigma}_a)$.

The exact route enumerates the outcomes, the 2^8 pass-or-fail records of one year and the 3^8 of two, or a grid of figures y_v, reweighting equation (30.9) by each outcome's probability at the state. The same enumeration over the state normal is a check on the Gaussian preposterior, 21.9 against 21.2 for one test and 30.4 against 30.5 for two.

Table 30.14: The candidate measurements of the life ranked by the expected value of sample information [R$/string-year], on the state route and the exact route, with the three interventions as the action set.

Candidate	Rows	$\sqrt{\Sigma_{a,uu}}$	$\sqrt{\Sigma_{a,vv}}$	sd of $\ln T^*$	EVSI, state	EVSI, exact
Present record	–	0.113	0.390	0.118	–	–
Annual test of the eight	8	0.089	0.336	0.109	21	44
Two annual tests of the eight	16	0.077	0.291	0.103	30	82
Shape figure, $\sigma_v = 0.2$	1	0.107	0.178	0.109	11	163
Shape figure, $\sigma_v = 0.1$	1	0.106	0.097	0.107	16	184
Life known exactly	–	0	0	0	120	334

The routes rank the candidates oppositely. On the state route the tests are worth more than a shape figure, since $T^*(\mathbf{z})$ is mainly a matter of the scale: the gradient (30.17) has 1 on u and 0.24 on v, and a shape figure that halves $\sqrt{\Sigma_{vv}}$ takes the spread of $\ln T^*$ only to where one annual test takes it. On the exact route the figure is worth more than two years of tests, about R$ 163 against R$ 82, since it removes the shallow and long-lived states that account for half of the value. Both are stated, and the chapter does not choose between them.

30.6.3 The recommendation

Do not wait on any measurement to adopt the rule: waiting gives up the saving of R$ 4 936 (R$ 3 598) per string-year against a ceiling of R$ 120 (R$ 334) on anything learned. The annual test is priced separately. One round of the eight costs R$ 16 000 and is worth R$ 21 (R$ 44) per string-year in this decision, about R$ 250 (R$ 530) a year over the twelve, and kept as a program it costs a net R$ 1 211 (R$ 958) per string-year. This decision does not pay for it: it must be justified on other grounds, such as the assurance of capacity for an outage, which this chapter does not price, or dropped. Ask the manufacturer for its life-test shape: it costs a request, the routes differ on whether it is worth R$ 11 or R$ 163, and it decides whether the shallow states that drive the exact route's value and the winner under narrower priors are live. The value of perfect information is the ceiling on anything further, R$ 120 (R$ 334) per string-year for the life. When rows arrive, the fit returns to step 3 with them. How to spend on such measurements across decisions is taken up in Chapter 31 and how to weigh an investment against them in Chapter 32, both in Volume III.

30.6.4 What the analysis showed

The first reading said to replace at just under four years, at a cost it stated to the real, and to save more than five thousand reais per string-year. The analysis keeps the age and revises the cost and the saving, which are one state's and not the posterior's. Replacing at four years is the conclusion on every route tried under the flat prior, and it wins against both other interventions for this delay and this test cost. What the engineer's instinct expected, that uncertainty about a wearing part argues for earlier replacement, the analysis does not support: on one route there is almost no pull, on the other the pull is later. What the record leaves open is not the age but how sure to be of what follows from it, how much the shallow shapes count, and so what is worth learning, and on those the two routes are not close. For the site, the rule can start now, and whether the annual test stays beside it is a separate decision that this rule does not pay for. A cheaper test could overturn the rule; a measured delay almost certainly would not.

30.6.5 Limits of the model

The derivations are exact under their assumptions, and six of them need checking before a result is used.

A new string is new. A replacement restores the string. A string that joins a bank of aged ones, or a batch fault shared across a replacement generation, would break the independence of the cycles.

One failure cost. Table 30.15 varies the costs; the age depends on them only through their ratio. Each row is a pair of costs; the columns give the Bayes action, its cost and run to failure on each route. With $C_{\mathrm{f}} = \mathrm{R\$}\,22\,500$, a ratio of 1.5, the state route keeps the age rule at 5.68 yr and the exact route's interior minimum at 6.70 yr costs R$ 29 more than run to failure, so the two routes fall on opposite sides.

Table 30.15: The decision under other costs [R$ and R$/string-year]: the Bayes action, its expected cost and run to failure, state route with exact in parentheses.

C_{f}	C_{r}	Action [yr]	$\mathcal{R}$ at the action	Run to failure
60 000	15 000	3.89 (4.10)	5 207 (5 573)	10 143 (9 171)
120 000	15 000	3.22 (3.26)	6 471 (7 422)	20 286 (18 342)
30 000	15 000	4.88 (5.40)	4 079 (4 087)	5 071 (4 586)
22 500	15 000	5.68 (6.70)	3 571 (3 468)	3 804 (3 439)
60 000	30 000	4.88 (5.40)	8 159 (8 174)	10 143 (9 171)

The test beside the age. Only the annual test beside the age rule is computed; an interval chosen jointly with the age is not. Dropping the annual test also gives up the rows of step 7 and the delay log of Chapter 29, and failures then surface only in service.

The double reading of a failed test. The record reads a failed test as a failure. The interval fit raises the plug-in cost by R$ 594 (Section 30.4.3), more than the gap between the routes.

The long run. The rate is a long-run average over many cycles. With twelve strings and a four-year rule the program sees about three replacements a year, enough for a rate over a decade and not for any one year's budget.

One population. The prior on the shape and the comparability of the strings are those of Chapter 10. A string from another room or duty would not share this posterior.

The next chapter takes up what this one bounded: what a given record is worth before it is taken, across the decisions it could change.

30.7 What to retain

Definition 30.1 (Age replacement and its value of information)**.** For a lifetime with survival R and hazard λ, the long-run cost rate of replacing at age T is $L(T, \mathbf{z}) = [C_{\mathrm{r}} + (C_{\mathrm{f}} - C_{\mathrm{r}})F(T)]/\int_0^T R(t)\,\mathrm{d}t$. At a known state a finite optimum satisfies $\lambda(T)W(T) + R(T) = C_{\mathrm{f}}/(C_{\mathrm{f}} - C_{\mathrm{r}})$, which requires a rising hazard. The Bayes action minimizes $\mathcal{R}(T) = \mathbb{E}[L(T, \mathbf{z}) \mid \mathbf{y}]$. The expected value of perfect information is $\min_a \mathcal{R}(a) - \mathbb{E}[\min_a L(a, \mathbf{z}) \mid \mathbf{y}] \geq 0$.

Four conclusions carry beyond the strings. Three objects are called the optimum, and on both routes they differ: the action at $\hat{\mathbf{z}}$, the distribution of what would be best at a state, and the minimizer of the expected loss; the last is the decision. A plug-in action can stand while its plug-in cost falls: on both routes here the regret was negligible and the cost was off by more than the value of perfect information. Whether uncertainty pulls an action earlier, and how much the uncertainty is worth, are route statements that an asymmetry argument does not settle: here the pull was opposite on the two routes and the value differed by more than a factor of two, and the state route's normal has few of the states that drive the exact route's value. And a value of information belongs to an action set: the median delay was worth tens of reais when testing was the only intervention and next to nothing with an age rule beside it, and interventions are comparable only when their cycles are accounted alike. Table 30.16 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 30.16: The terms of the replacement decision and where each is defined.

Term	Expression	Units	Meaning	Defined in
Cycle length	$W(T) = \int_0^T R(t)\,\mathrm{d}t$	yr	mean of $\min(T, T)$	30.1.1
Cost rate of an age	$[C_\mathrm{r} + \Delta C\, F(T)]/W(T)$	R$/yr	renewal–reward	30.1.1
Condition for an optimum	$\lambda W + R = C_\mathrm{f}/\Delta C$	–	needs a rising hazard	30.1.1
Run to failure	$C_\mathrm{f}\nu_\mathrm{r}(u, v)$	R$/yr	the limit $T \to \infty$	30.1.4
Condition, one basis	$L_\mathrm{c}(\tau, \mathbf{z})$, life ended at the catch	R$/yr	comparable with $L(T, \mathbf{z})$	30.1.4
General forms				
Expected consequence	$\mathcal{R}(a) = \int L(a, \mathbf{z}) p(\mathbf{z} \mid \mathbf{y})\,\mathrm{d}\mathbf{z}$	R$/yr	one action	30.3.1
The decision	$a^* = \arg\min_a \mathcal{R}(a)$	–	minimizer	30.3.2
Tail	$\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y})$	–	the saving positive	30.3.3
Regret of an action	$\mathcal{R}(a') - \min_a \mathcal{R}(a)$	R$/yr	plug-in against Bayes	30.3.4
Perfect information	$\min_a \mathcal{R}(a) - \mathbb{E}[\min_a L \mid \mathbf{y}]$	R$/yr	bounds any row	30.3.5
Specialized forms and answers				
Expected cost, exact	$\int p(v \mid \mathbf{y}) \int L\,\mathrm{Gamma}(\zeta; k, A)\,\mathrm{d}\zeta\,\mathrm{d}v$	R$/yr	two integrals	30.2.2
Optimum at a state	$T^*(\mathbf{z}) = \eta\, g(\beta)$	yr	not the decision	30.2.3
Plug-in action	$T^*(\hat{\mathbf{z}})$	yr	one state	30.3.4
Bayes action	$T^\mathrm{B} = \arg\min_T \mathcal{R}(T)$	yr	the decision	30.3.2
Least loss at a state	$L^*(\mathbf{z}) = \Delta C\, \lambda(T^*(\mathbf{z}))$	R$/yr	closed form	30.3.5
Local form of the value	$\frac{1}{2}\,\partial_T^2 L\, T^*(\hat{\mathbf{z}})^2 \mathbf{g}_{\ln T^*}^\mathsf{T} \boldsymbol{\Sigma} \mathbf{g}_{\ln T^*}$	R$/yr	state route, near the fit	30.3.5

30.8 Exercises

Exercise 30.1. *Formulation.* A unit has an exponential lifetime with mean 8 yr, a planned replacement costs R$ 10 000 and a failure R$ 40 000. (a) Compute the cost rate at $T = 2$ yr and $T = 4$ yr. (b) Show why no finite T is optimal. ((a) With $F = 1 - \mathrm{e}^{-T/8}$ and $W = 8F$, the rates are R$ 9 401 and R$ 6 927. (b) The rate is $\lambda[C_\mathrm{f} + C_\mathrm{r}R/F]$, which falls to $\lambda C_\mathrm{f} =$ R$ 5 000 as $T \to \infty$: a constant hazard gives equation (30.3) no root.)

Exercise 30.2. *Formulation.* A Weibull lifetime has $\eta = 10$ yr and $\beta = 3$, with costs of R$ 10 000 and R$ 50 000. Find the best whole-year age on a grid and the saving against running to failure, and say which state entries an expected cost would average over. ($T = 5$ yr at R$ 3 031, 4 yr at R$ 3 170, 6 yr at R$ 3 120; run to failure is $50\,000/(10\,\Gamma(4/3)) =$ R$ 5 599, so the saving is R$ 2 568. The expected cost averages over the scale and the shape, the entries of the lifetime block.)

Exercise 30.3. *Derivation.* Derive the renewal–reward cost rate (30.2) and the condition (30.3), with the cost at the root. (A cycle costs $C_\mathrm{r}+\Delta CF(T)$ on average and has the mean length $\int_0^T R$; the theorem gives the ratio. With $N = C_\mathrm{r}+\Delta CF$, $N'W - NW' = 0$ reads $\Delta C\lambda RW = NR$, so $N/W = \Delta C\lambda$, and $C_\mathrm{r} + \Delta CF = \Delta C\lambda W$ gives $\lambda W + R = C_\mathrm{f}/\Delta C$.)

Exercise 30.4. *Derivation.* (a) Derive EVPI ≥ 0. (b) Derive the local form (30.20). ((a) $\min_a L(a, \mathbf{z}) \leq L(a', \mathbf{z})$ for every $\mathbf{z}$ and a'; taking expectations and then the minimum over a' gives $\mathbb{E}[\min_a L] \leq \min_{a'} \mathcal{R}(a')$. (b) Assume $T^*(\mathbf{z})$ narrowly distributed with finite variance, as under a normal approximation of $\ln T^*$ held near the fit. Near $T^*(\mathbf{z})$, $L(T, \mathbf{z}) \approx L^*(\mathbf{z})+\frac{1}{2}\partial_T^2 L\,(T-T^*(\mathbf{z}))^2$; with the curvature taken constant, the expectation at T is $\mathbb{E}[L^*] + \frac{1}{2}\partial_T^2 L\,[\mathrm{Var}(T^*) + (T - \bar{T}^*)^2]$, minimized at the mean $\bar{T}^*$, and its minimum less $\mathbb{E}[L^*]$ is $\frac{1}{2}\partial_T^2 L\ \mathrm{Var}(T^*)$; the delta method gives $\mathrm{Var}(T^*) \approx T^*(\hat{\mathbf{z}})^2 \mathbf{g}_{\ln T^*}^\mathsf{T} \boldsymbol{\Sigma} \mathbf{g}_{\ln T^*}$.)

Exercise 30.5. *Application.* Under block replacement every string is replaced at the calendar age T regardless, each failure replaced at once at C_f, with the rate $[C_\mathrm{r} + C_\mathrm{f} M(T)]/T$ and M the renewal function, $M(t) = F(t) + \int_0^t M(t - x) f(x)\,\mathrm{d}x$. Find the best T for the strings at the MAP state and compare with age replacement there. (Solving the renewal equation on a grid of 0.005 yr gives $T = 3.73$ yr at R\$ 5 148, against R\$ 4 910 at 3.905 yr under age replacement, and $M(4) = 0.096$, about $F(4)$. Block replacement costs R\$ 238 more because it discards strings just replaced; it is simpler to schedule, a saving not priced here.)

Exercise 30.6. *Application.* A replaced string has a salvage value that depends on its age, $S(T) = 6\,000\,\mathrm{e}^{-T/4}$ reais, received at a planned replacement only. Write the loss and find the action at $\hat{\mathbf{z}}$ and on both routes. ($L = [(C_\mathrm{r} - S(T))R(T) + C_\mathrm{f} F(T)]/W(T)$. At $\hat{\mathbf{z}}$: 3.57 yr at R\$ 4 319; on the state route 3.55 yr at R\$ 4 619; on the exact route 3.68 yr at R\$ 5 048. The salvage pulls the action earlier, by about a third of a year on the state route and about four tenths on the exact one, by paying for life not used.)

Chapter 31

Value of Information and Sensor Placement

31.1 Which measurement to buy

The operator of the two halls of Chapter 17 can add one instrument this year, at about R$ 300 a year with its calibration, and five are on the list: a second temperature sensor in each hall, a supply sensor of 0.1 K, a load meter on hall 2 and a sensor on the return header. Chapter 28 left hall 1 admitted at 60.90 kW with 95 % assurance, and the sales office prices admitted load at R$ 400 per kW-year.

The vendor's argument is the quickest reading. The supply sensor is five times as precise as the temperature sensors on the list, and the supply temperature, known now to 0.28 K, would be known to 0.09 K: no other candidate takes so large a share off an entry of the state. Precision is a property of the instrument, and a share of an entry is a property of the state. Neither names what the plant is to do with the reading. The question of this chapter is

how much a row not yet read is worth, measured against the state, against a named output and against a named decision; which candidate wins under each; and whether any of them is worth its price.

What changes from the step 7 of earlier chapters is that the width removed is now one of three values, and that the reading's outcome, not only its covariance, enters the third.

Jargon: candidate row, preposterior, expected information gain, width removed, expected value of sample information, expected value of perfect information, assurance rule, overshoot, critical fractile, lower-tail mean, action set

A *candidate row* is a reading not yet taken, written as a row of the objective. The *preposterior* is the distribution of the posterior before that reading is known. The *expected information gain* is the expected fall in the entropy of the state, in nats. The *width removed* is the fraction of an output's standard deviation the reading would take away. The *expected value of sample information* is the expected improvement of a decision from the reading, and the *expected value of perfect information* the same with the state revealed. The *assurance rule* admits the largest load whose probability of staying under the limit is κ. *Overshoot* is load admitted beyond the path's capacity. The *critical fractile* is the quantile that minimizes an expected loss linear on each side of a level. The *lower-tail mean* is the mean of a quantity below such a quantile. The *action set* is what the decision may choose among.

Key idea

A sensor has no value in itself. The information gain is a property of the state: it asks how large the virtual variance at the sensor's place is against the sensor's own. The width removed is a property of an output: it multiplies that ratio by the squared correlation between the output and the quantity read. The value of sample information is a property of a decision and of its action set: the same load meter is worth almost nothing to the hall's path and most of all to the chiller's headroom. An assurance rule is not by itself an expected loss, so its gain must be given a loss before it is called a value.

31.1.1 The plant and its candidates (step 0)

The plant is that of Chapters 17 to 19: two halls with loads Q_i and paths of log-conductance g_i to a supply header at T_s, on the reduced state of equation (20.1), with the chiller's capacity factor c_{ch} of Chapter 28 as a sixth entry where its action set needs it. Figure 31.1 draws it. The green boxes are the free nodes and the blue boxes the halls, balance nodes. The solid arrows are the physical paths. The dashed gray arrows are the five readings in hand. The dashed red arrows are the five candidates, none of them yet read; each ends on the node whose function it would read.

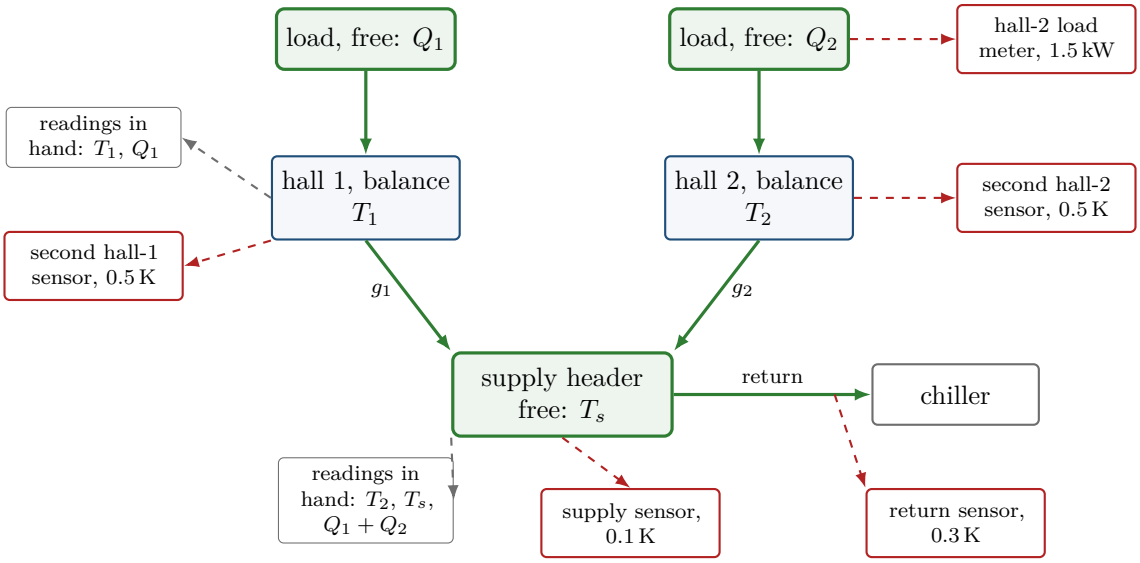


Figure 31.1: The two-hall plant with its five candidate rows. Free (green): the loads, the supply temperature and, labeling the paths, the log-conductances. Balance nodes (blue): the halls. Observed (dashed gray): the five readings in hand. Dashed red: the candidate rows, which are not yet read and enter the objective only when they are.

Table 31.1 lists the candidates. Each row is a candidate, with the function of the state it reads, its noise, and the line $\mathbf{j}_a$ it would add to $\mathbf{J}$ at the MAP state of step 3, whose numbers come only there. All are priced alike at about R$ 300 a year; that figure and the R$ 400 per kW-year are illustrative.

31.1.2 The probabilistic questions (step 1)

The operator asks six questions. Table 31.2 lists them. The first column is the question in engineering language; the second the general form of its category (Chapter 2), with $\tilde{y}_a$ the reading of candidate a; the third its form for the plant; the fourth its category; the last the subsection that answers it. Two outputs are named: the margin of hall 1 at 60 kW, m_1^* of

Table 31.1: The candidate rows: what each reads, its noise, and the entries of the state its line of the Jacobian touches, columns in the order $(g_1, g_2, Q_1, Q_2, T_s)$.

Candidate	Reads $f_a(\mathbf{z})$	σ_a	Pattern of $\mathbf{j}_a$
Second hall-1 sensor	$T_s + Q_1 \mathrm{e}^{-g_1}$	0.5 K	$(*, 0, *, 0, *)$
Second hall-2 sensor	$T_s + Q_2 \mathrm{e}^{-g_2}$	0.5 K	$(0, *, 0, *, *)$
Supply sensor	T_s	0.1 K	$(0, 0, 0, 0, *)$
Hall-2 load meter	Q_2	1.5 kW	$(0, 0, 0, *, 0)$
Return sensor	$T_r = T_s + (Q_1 + Q_2)/\dot{m}c_p$, equation (21.4)	0.3 K	$(0, 0, *, *, *)$

equation (28.11), and the load of hall 2, on which billing rests. One decision is named: the load Q admitted to hall 1.

Table 31.2: The questions asked of the candidates, the general form of their category, their form for the plant, their category and where each is answered. $\mathbf{y}$ is the five readings and the three priors; $Q_\kappa(\tilde{y}_a)$ is the assurance load after the reading.

Engineering question	General form	Mathematical form	Category	Answered in
How much does each candidate reduce the uncertainty of the state?	$U(a) = \mathbb{E}_{\tilde{y}_a \mid \mathbf{y}}[u(p(\mathbf{z} \mid \mathbf{y}, \tilde{y}_a))]$, u the information gained	$\mathrm{EIG}(a) = \frac{1}{2}\ln(\det \boldsymbol{\Sigma} / \det \boldsymbol{\Sigma}_a)$	design (information gain)	31.3.1
How much width does each take off the margin, and off hall 2's load?	the same, u the spread of an output	$\sqrt{\mathbf{g}^\mathsf{T} \boldsymbol{\Sigma}_a \mathbf{g}}$ for m_1^* and Q_2	design (width removed)	31.3.2
Which measurement next?	the candidates ordered by $U(a)$	the ranking under each criterion	design (sensor placement)	31.3.3
What is a candidate worth to the load decision, in money?	$\mathrm{EVSI} = \bar{L}_{\text{now}} - \bar{L}_{\text{after}}$	$v_\mathrm{L}(\mathbb{E}[Q_\kappa(\tilde{y}_a)] - Q_\kappa)$, and under the loss of step 2	design (decision value)	31.3.4
What is the most any measurement could be worth?	$\min_a \mathcal{R}(a) - \mathbb{E}[\min_a L(a, \mathbf{z}) \mid \mathbf{y}]$	the same with the state revealed, on each action set	design (the ceiling)	31.3.5
Why do the rankings differ, and which one to use?	$\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$ term by term; the three criteria compared	V_a/σ_a^2, $\mathrm{corr}(f, f_a)$, the block values	attribution; the comparison	31.3.6

31.1.3 The evidence as rows (step 2)

Step 2 adds no row to the posterior of the present. The eight rows of equation (17.15) stay, with the two chiller rows of equation (28.3). What is new is a row that may be added and a loss.

The candidate row. Each candidate would add

$$r_a(\mathbf{z}) = \frac{f_a(\mathbf{z}) - \tilde{y}_a}{\sigma_a}, \qquad \mathbf{j}_a = \frac{\nabla f_a}{\sigma_a} \ \text{ at } \hat{\mathbf{z}}, \tag{31.1}$$

with the pattern of Table 31.1. The decision's two entries are g_1 and T_s, the pattern of $\mathbf{g}_T$ in equation (28.4). Only the hall-1 sensor's line meets both columns; the hall-2 load meter meets neither and reaches them only through the covariance of the entries.

The modeling choices. Five choices make the values. One reading is taken once, with an error independent of the others: this keeps the update rank one, and it fails for a sensor that repeats a bias, or one averaged for months, which Section 31.6.3 takes up. The value of admitted load is one number, $v_{\mathrm{L}} = \mathrm{R\$}\,400$ per kW-year (illustrative): this makes money proportional to kilowatts, and it fails if the load sells only in blocks, which the answers price as well. The supply is held, as in Chapter 28's path questions: this is the first action set, and it fails whenever the chiller cannot deliver, which the second and third action sets remove. A load is re-decided after the reading, down as well as up: the rule allows no asymmetry, and it fails if a customer once admitted cannot be moved. The value accrues each year while the sensor's cost does: this makes the two comparable per year, and it holds only if the admitted load stays sold.

The decision rule. The load admitted is the assurance load of equation (28.12). This chapter's rule is the inversion of the normal tail of the scenario temperature, the root of equation (28.13), Chapter 28's reported route. The delta inversion on the capacity, 59.21 kW, linearizes another function of the same event and is not used here; the full posterior's quantile of the capacity is the same rule with the Gaussian and the tangent plane removed.

The loss. A rule is not yet a loss. Admitting Q earns $v_{\mathrm{L}}Q$ a year, and a load beyond the capacity $Q_{\mathrm{cap}}(\mathbf{z})$ of equation (28.14) costs something. The loss in which the rule is the best action charges the overshoot in proportion:

$$L(Q, \mathbf{z}) = -v_{\mathrm{L}}Q + c_{\mathrm{o}}\left(Q - Q_{\mathrm{cap}}(\mathbf{z})\right)^{+}, \qquad c_{\mathrm{o}} = \frac{v_{\mathrm{L}}}{1-\kappa} = \mathrm{R\$}\,8\,000 \text{ per kW-year}. \tag{31.2}$$

The value of c_{o} is not a figure from the site. It is what the customer's κ implies, as the fourth answer shows; a site that knows its overshoot cost would have a different κ.

The three action sets. The first holds the supply, the question as posed: the capacity is Q_{cap}. The second takes hall 2 off this chiller and lets the derate raise the supply by equation (28.20), so that the capacity is

$$Q'_{\mathrm{cap}}(\mathbf{z}) = \begin{cases} Q_{\mathrm{cap}}, & Q_{\mathrm{cap}} \le P_{\mathrm{pl}}c_{\mathrm{ch}}, \\ \dfrac{T_{\max} - T_s + P_{\mathrm{pl}}c_{\mathrm{ch}}/\dot{m}c_p}{1/\dot{m}c_p + \mathrm{e}^{-g_1}}, & \text{otherwise}, \end{cases} \tag{31.3}$$

the root of $T_1^{*\prime} = T_{\max}$ of Section 28.3.7. The third keeps hall 2 on this chiller. Chapter 28 found that the steady model then stops at the load where the shortfall begins, so the limit is that onset load,

$$Q_{\mathrm{on}}(\mathbf{z}) = P_{\mathrm{pl}}c_{\mathrm{ch}} - Q_2, \tag{31.4}$$

with the gradient $(0, 0, 0, -1, 0, P_{\mathrm{pl}})^{\mathsf{T}}$ on the six entries. Each has its own assurance load and its own value of knowing.

31.2 The posterior of the state (step 3)

All six questions are read from the present posterior. The rows are unchanged, so the state, its fit and its covariance are those already computed, and the full posterior is reached by the reweighting already derived.

31.2.1 The posterior from the rows

The state is $\mathbf{z} = (g_1, g_2, Q_1, Q_2, T_s)^{\mathsf{T}}$, equation (20.1), with the eight rows of Section 20.2.1, the Jacobian $\mathbf{J} = \mathbf{J}_{\mathrm{e}}\mathbf{N}$, the normal equations iterated in Chapter 18 and the Laplace covariance $\boldsymbol{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1}$. The capacity factor's two rows are linear and alone in their column, so its posterior is the normal of equation (28.8), independent of the hall's. The committed values, Table 18.4 and Table 19.3, are

$$\hat{\mathbf{z}} = \begin{bmatrix} 1.70297 \\ 1.39993 \\ 30.68819 \\ 22.23066 \\ 12.11145 \end{bmatrix}, \quad \boldsymbol{\Sigma} = \begin{bmatrix} 0.011382 & -0.000417 & 0.06358 & -0.06274 & 0.012974 \\ -0.000417 & 0.018765 & -0.07977 & 0.22219 & 0.011989 \\ 0.06358 & -0.07977 & 2.1085 & -1.9490 & 0.002914 \\ -0.06274 & 0.22219 & -1.9490 & 5.5048 & -0.026883 \\ 0.012974 & 0.011989 & 0.002914 & -0.026883 & 0.078433 \end{bmatrix}, \tag{31.5}$$

with $\hat{c}_{\mathrm{ch}} = 0.911536$ and $\Sigma_{66} = 1.0922 \times 10^{-4}$. These reproduce to all digits shown when the fit is rerun with the unrounded hall-1 reading.

31.2.2 The full posterior and the two routes

The state route takes the posterior as $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$, and the output through its tangent plane at the posterior's center. *The full posterior* is $\mathrm{e}^{-\mathcal{L}(\mathbf{z})}/c$, reached by the draws of Chapter 20: 200 000 draws of the Laplace Gaussian (`default_rng(17)`, by the Cholesky factor) reweighted by equations (20.10) and (20.11), with $N_{\mathrm{eff}} = 182\,933$ from equation (20.12), and the capacity factor drawn independently under seed 28. On these draws the rule of step 2 gives

$$Q_{0.95} = 60.90\,\mathrm{kW}\ \text{(state route)}, \qquad Q_{0.95} = 60.99\,\mathrm{kW}\ \text{(full posterior)}, \tag{31.6}$$

Chapter 28's values, $Q'_{0.95} = 58.91\,\mathrm{kW}$ (58.96 kW) with the derate, and on the onset load $Q_{\mathrm{on}}(\hat{\mathbf{z}}) = 54.692 - 22.231 = 32.46\,\mathrm{kW}$ with the spread $\sqrt{60^2\Sigma_{66} + \Sigma_{44}} = 2.429\,\mathrm{kW}$, so that its assurance load is $32.46 - 1.6449 \times 2.429 = 28.47\,\mathrm{kW}$ on both routes. That load is below the present $\hat{Q}_1 = 30.69\,\mathrm{kW}$: with hall 2 kept on this chiller the present load is not assured at 95 %, as the headroom of Chapter 28 already said in other terms.

A reading $\tilde{y}_a$ changes the full posterior by the factor $\mathcal{N}(\tilde{y}_a;\ f_a(\mathbf{z}),\ \sigma_a^2)$. Each outcome therefore needs no new draws: the weights of the same draws are multiplied by it. The figure of merit is N_{eff} at each outcome, reported with the values, and it is the reason no refit is used on this route.

31.3 The answers (step 4)

The questions of Table 31.2 are answered from $(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of equation (31.5), through the rank-one preposterior, and that is the value reported. The full posterior, reweighted for each outcome, gives the value beside it, with the gap.

31.3.1 The information gained about the state

Engineering question. How much does each candidate reduce the uncertainty of the state?

General form. The design form (2.16) with u the information gained. For any posterior this is the expected fall in entropy, which is the mutual information of the state and the reading:

$$\mathrm{EIG}(a) = \mathbb{E}_{\tilde{y}_a|\mathbf{y}}\left[\int p(\mathbf{z} \mid \mathbf{y}, \tilde{y}_a) \ln \frac{p(\mathbf{z} \mid \mathbf{y}, \tilde{y}_a)}{p(\mathbf{z} \mid \mathbf{y})}\, \mathrm{d}\mathbf{z}\right] = h(\tilde{y}_a \mid \mathbf{y}) - h(\tilde{y}_a \mid \mathbf{z}), \tag{31.7}$$

with h the differential entropy. It is defined whenever the posterior is proper, as this one is, and it does not depend on how the state is parameterized.

Specialized form. The question is equation (31.7) with $\tilde{y}_a$ the candidate's reading and $p(\mathbf{z} \mid \mathbf{y})$ the normal of equation (31.5). One choice enters here: the row is taken as linear at $\hat{\mathbf{z}}$, which gives a closed form, and it fails for the two hall sensors to the degree that $Q_i \mathrm{e}^{-g_i}$ curves across the band, which the full value below measures. An entropy of a normal is $\frac{1}{2} \ln \det(2\pi \mathrm{e}\, \boldsymbol{\Sigma})$, so the difference of the two is equation (2.18), with $\boldsymbol{\Sigma}_a$ the preposterior (20.28). The determinant lemma on the same update gives

$$\mathrm{EIG}(a) = \tfrac{1}{2} \ln \frac{\det \boldsymbol{\Sigma}}{\det \boldsymbol{\Sigma}_a} = \tfrac{1}{2} \ln \det(\mathbf{I} + \boldsymbol{\Sigma} \mathbf{j}_a \mathbf{j}_a^{\mathsf{T}}) = \tfrac{1}{2} \ln s_a, \qquad s_a = 1 + \mathbf{j}_a^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{j}_a = 1 + \frac{V_a}{\sigma_a^2}, \quad (31.8)$$

with $V_a = \sigma_a^2 \mathbf{j}_a^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{j}_a$ the virtual variance at the candidate's place, as in equation (21.27). The second form is also the entropy difference of the reading in equation (31.7): the predictive (21.21) has the variance $\sigma_a^2 s_a$ against the noise's σ_a^2. On the full posterior the reading's predictive is a one-dimensional mixture, so the same difference is one quadrature over the weighted draws of f_a.

Building blocks. $\boldsymbol{\Sigma}$ of equation (31.5); the lines $\mathbf{j}_a$; for the values beside them, the reweighted draws of Section 31.2.2.

The numbers. For the hall-1 sensor, $\mathbf{j}_a = (-11.179,\ 0,\ 0.364,\ 0,\ 2)^{\mathsf{T}}$, as in Chapter 20, $\boldsymbol{\Sigma} \mathbf{j}_a = (-0.07814,\ -0.00041,\ 0.06319,\ -0.06236,\ 0.01289)^{\mathsf{T}}$ and $s_a = 1.9223$, so $\mathrm{EIG} = \frac{1}{2} \ln 1.9223 = 0.327$ nat. For the supply sensor $s_a = 1 + 0.078433/0.01 = 8.843$ and $\mathrm{EIG} = 1.090$ nat. Table 31.3 gives all five with the widths of the next question.

31.3.2 The width taken off an output

Engineering question. How much width does each take off the margin, and off hall 2's load?

General form. The design form (2.16) with u the spread of a derived quantity after the outcome:

$$U_f(a) = \mathbb{E}_{\tilde{y}_a \mid \mathbf{y}} \big[\mathrm{sd}(f \mid \mathbf{y}, \tilde{y}_a) \big]. \quad (31.9)$$

Specialized form. The question is equation (31.9) with $f = m_1^*$ at 60 kW and $f = Q_2$, of gradients $\mathbf{g}_T$ of equation (28.9) up to sign and $(0, 0, 0, 1, 0)^{\mathsf{T}}$, and the normal of equation (31.5). One choice enters here: the gradient stays at $\hat{\mathbf{z}}$ under every outcome, which makes the spread independent of the outcome, and it is exact only for a linear row. The result is equation (20.29), derived in Chapter 20. Written with the correlation of the output and the quantity read, $\mathbf{g}^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{j}_a = \mathrm{corr}(f, f_a) \sqrt{\mathbf{g}^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g}}\, \sqrt{V_a}/\sigma_a$, it reads

$$\frac{\mathbf{g}^{\mathsf{T}} \boldsymbol{\Sigma}_a \mathbf{g}}{\mathbf{g}^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g}} = 1 - \mathrm{corr}(f, f_a)^2 \, \frac{V_a}{V_a + \sigma_a^2}, \quad (31.10)$$

so that the share of the variance removed is the squared correlation of the output with the reading itself, the squared correlation with the quantity times the reading's reliability $V_a/(V_a + \sigma_a^2)$.

Building blocks. $\boldsymbol{\Sigma}$, the lines $\mathbf{j}_a$, the two gradients; for the values beside them, the weighted variance of each output under each reweighted outcome, equation (20.11).

Table 31.3: The information gain in nats and the width removed from each output, state route with the full posterior in parentheses. Present spreads: margin at 60 kW 1.074 K (1.071), Q_2 2.346 kW (2.314).

Candidate	EIG [nat]	sd(m_1^*) [K]	Width, m_1^*	sd(Q_2) [kW]	Width, Q_2
Second hall-1 sensor	0.327 (0.325)	0.874 (0.872)	18.7 % (18.6 %)	2.346 (2.314)	0.02 % (0.0 %)
Second hall-2 sensor	0.328 (0.333)	1.074 (1.071)	0.01 % (0.0 %)	2.341 (2.308)	0.24 % (0.27 %)
Supply sensor	1.090 (1.095)	1.053 (1.050)	2.0 % (2.0 %)	2.345 (2.312)	0.07 % (0.1 %)
Hall-2 load meter	0.619 (0.609)	1.048 (1.045)	2.5 % (2.5 %)	1.264 (1.259)	46.1 % (45.6 %)
Return sensor	0.401 (0.400)	1.064 (1.061)	1.0 % (1.0 %)	2.229 (2.204)	5.0 % (4.8 %)

The numbers. Table 31.3 gives the three quantities. Each row is a candidate; the columns give the information gain, the spread of the margin at 60 kW and of hall 2's load after it, and the width each removes, on the state route and, in parentheses, on the full posterior, whose spreads are averaged over the outcomes as in equation (31.9). For the hall-1 sensor, equation (31.10) gives $1 - 0.8401^2 \times 0.4798 = 0.6614$, and $\sqrt{0.6614} \times 1.0744 = 0.874$ K, the number of Table 20.9.

The two routes agree on every entry to within what a table of this kind can use, the largest gap being 0.010 nat on the load meter and 0.5 point of width. The nonlinearity of the two hall rows does not disturb the information gain.

31.3.3 The rankings

Engineering question. Which measurement next?

General form. The candidates ordered by $U(a)$ of the design form (2.16), the largest first for an information gain and the smallest first for a spread.

Specialized form. The question is equation (2.16) with $U(a)$ either the gain of equation (31.8) or the spread of equation (31.10), on each output. No choice enters beyond those already named.

Building blocks. Table 31.3.

The numbers. Figure 31.2 draws the three criteria, each panel in its own order. The horizontal axis is the criterion's own quantity; the red bar is the state route and the green bar the full posterior.

Table 31.4 sets the three orders side by side, with the money of the next question in the last column so that the comparison stands in one place. Three criteria give three winners on both routes: the supply sensor for the state, the hall-1 sensor for the margin, the load meter for hall 2's load. Among the three places for one more temperature sensor, hall 1, hall 2 and the return, Figure 31.5 carries the same comparison on the graph.

Table 31.4: The orders under each criterion, first to last, the same on both routes; hall-2 and hall-1 nearly tie in the first. The last column is the order by money of the next question, on the first action set.

Rank	Information gain	Width, margin	Width, Q_2	EVSI, supply held
1	supply	hall-1	load meter	hall-1
2	load meter	load meter	return	load meter
3	return	supply	hall-2	supply
4	hall-2	return	supply	return
5	hall-1	hall-2	hall-1	hall-2

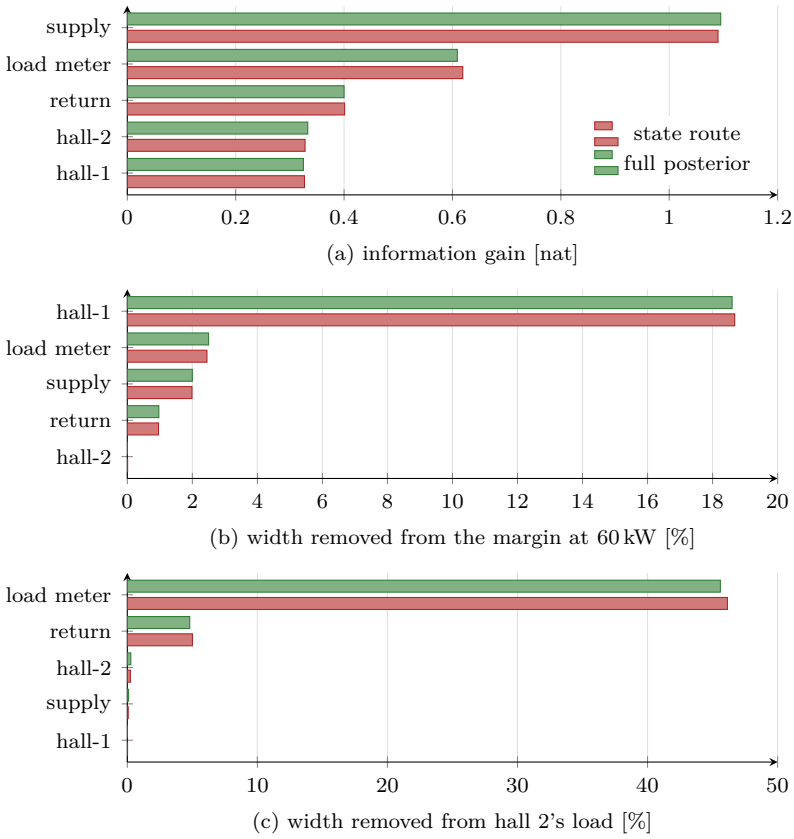


Figure 31.2: The candidates under three criteria, each panel in its own order: information gain, equation (31.8); width removed from the margin and from hall 2's load, equation (31.10).

31.3.4 The value to the load decision

Engineering question. What is a candidate worth to the load decision, in money?

General form. The expected value of sample information, equation (2.51), with the action chosen again after the reading, and the reading weighted by its predictive, equation (2.13):

$$\mathrm{EVSI}(a) = \min_{Q} \mathcal{R}(Q) - \mathbb{E}_{\tilde{y}_a \mid \mathbf{y}}\Big[\min_{Q} \mathbb{E}\big[L(Q, \mathbf{z}) \mid \mathbf{y}, \tilde{y}_a\big]\Big]. \tag{31.11}$$

Specialized form. The question is equation (31.11) with Q the admitted load, L the loss (31.2) and $p(\mathbf{z} \mid \mathbf{y})$ the normal of equation (31.5). One choice enters here: v_{L} is common to all candidates, so it scales every value and cannot change their order, and the cost comparison of the report moves it.

The rule is the action of this loss. Under any posterior with F the distribution function of the capacity,

$$\mathcal{R}(Q) = -v_{\mathrm{L}} Q + c_{\mathrm{o}} \int_0^Q F(q)\,\mathrm{d}q, \qquad \mathcal{R}'(Q) = -v_{\mathrm{L}} + c_{\mathrm{o}} F(Q) = 0 \iff F(Q) = \frac{v_{\mathrm{L}}}{c_{\mathrm{o}}} = 1 - \kappa,$$

which is the critical fractile, and with equation (28.16) it is Q_κ. Substituting, and using $\int_0^Q F = (1-\kappa)Q_\kappa - \mathbb{E}[Q_{\mathrm{cap}};\ Q_{\mathrm{cap}} < Q_\kappa]$ at the root,

$$\min_{Q} \mathcal{R}(Q) = -v_{\mathrm{L}}\, \bar{Q}_\kappa, \qquad \bar{Q}_\kappa = \mathbb{E}[Q_{\mathrm{cap}} \mid Q_{\mathrm{cap}} < Q_\kappa,\ \mathbf{y}] = \frac{1}{1-\kappa} \int_\kappa^1 Q_{\kappa'}\,\mathrm{d}\kappa', \tag{31.12}$$

the lower-tail mean of the capacity. Hence, under the loss,

$$\begin{aligned} \mathrm{EVSI}_L(a) &= v_{\mathrm{L}}\big(\mathbb{E}_{\tilde{y}_a}[\bar{Q}_\kappa(\tilde{y}_a)] - \bar{Q}_\kappa\big) \\ &= v_{\mathrm{L}}\big(\mathbb{E}_{\tilde{y}_a}[Q_\kappa(\tilde{y}_a)] - Q_\kappa\big) + c_{\mathrm{o}}\big(\mathbb{E}[\mathrm{overshoot}]_{\mathrm{now}} - \mathbb{E}[\mathrm{overshoot}]_{\mathrm{after}}\big), \end{aligned} \tag{31.13}$$

with $\mathbb{E}[\mathrm{overshoot}] = (1-\kappa)(Q_\kappa - \bar{Q}_\kappa)$.

The rise in admissible load. The first term of equation (31.13) alone is the value of the rule,

$$\mathrm{EVSI}_Q(a) = v_{\mathrm{L}}\big(\mathbb{E}_{\tilde{y}_a}[Q_\kappa(\tilde{y}_a)] - Q_\kappa\big). \tag{31.14}$$

The proof of Chapter 2 that a value is nonnegative does not apply to it as it stands. That proof keeps the present action available after the reading, at its present expected loss; a rule can make the present load inadmissible after a poor reading. The rise has a different origin. For a capacity normal with mean μ and spread $\mathrm{sd_c}$, linear in the state, $Q_\kappa = \mu - z_\kappa \mathrm{sd_c}$, the mean after the reading averages back to μ and the spread falls to $\mathrm{sd}_{\mathrm{c},a}$, so the rise is $z_\kappa(\mathrm{sd_c} - \mathrm{sd}_{\mathrm{c},a})$: positive for $\kappa > \frac{1}{2}$ and negative for $\kappa < \frac{1}{2}$. The assurance rule's bound tightens toward a center that does not move on average. On the temperature route the map from the center to the load is curved as well, so a term of curvature is added.

What the money measures. Equation (31.14) prices a load that stays at 95 % assurance after every reading. It charges nothing for the 5 %, and it should not be read as if the risk had been paid for. It is a real gain under a stated loss in two senses. If each exceedance costs a fixed sum C_{x}, whatever the load, the expected penalty is $0.05\,C_{\mathrm{x}}$ before and after, so the difference in expected loss between following the rule after the reading and following it now is exactly equation (31.14); the rule is then a policy, not the minimizer. If the cost grows with the overshoot as in equation (31.2), the rule is the minimizer and the true value is equation (31.13), which adds the fall in expected overshoot and is nonnegative on a coherent posterior by Chapter 2's proof; the delta tail of the state route is not such a posterior and can break it, as the action table below shows. In neither is it risk-free income: it is the load admitted at an unchanged probability of exceeding the limit.

The reading on each route. On the state route the reading is summarized by its standardized innovation $\eta_a = (\tilde{y}_a - f_a(\hat{\mathbf{z}}))/(\sigma_a\sqrt{s_a})$, standard normal by the predictive (21.21), and the mean moves by the rank-one update along one line,

$$\mathbf{m}_a(\eta_a) = \hat{\mathbf{z}} + \eta_a \frac{\boldsymbol{\Sigma}\mathbf{j}_a}{\sqrt{s_a}}, \qquad \boldsymbol{\Sigma}_a = \boldsymbol{\Sigma} - \frac{\boldsymbol{\Sigma}\mathbf{j}_a\mathbf{j}_a^{\mathsf{T}}\boldsymbol{\Sigma}}{s_a}, \tag{31.15}$$

so that $\mathbf{m}_a \sim \mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma} - \boldsymbol{\Sigma}_a)$. The load after the reading is the root (28.13) with $\hat{g}_1$, $\hat{T}_s$ replaced by the entries of $\mathbf{m}_a$ and the block of $\boldsymbol{\Sigma}$ by that of $\boldsymbol{\Sigma}_a$, and its expectation is a one-dimensional Gauss–Hermite sum over η_a. On the full posterior the reading's predictive is the mixture of Section 31.3.1; at each outcome the weights are multiplied as in Section 31.2.2 and the load is the weighted 5 % quantile of the capacity. The expectation runs over the 2 000 midpoints in probability of that predictive.

Building blocks. $\hat{\mathbf{z}}$, $\boldsymbol{\Sigma}$ of equation (31.5); $\boldsymbol{\Sigma}\mathbf{j}_a$, s_a of Section 31.3.1; the root (28.13); the present load (31.6); for the values beside them, the reweighted draws.

One candidate worked. For the hall-1 sensor, $\boldsymbol{\Sigma}_a$ has $\Sigma_{a,11} = 0.008206$, $\Sigma_{a,15} = 0.013498$ and $\Sigma_{a,55} = 0.078347$. With the center held, equation (28.13) has $A = 0.977798$, $B = 12.852031$, $C = 165.902748$, so $x = 11.3857\,\mathrm{K}$ and $Q = 62.51\,\mathrm{kW}$, Chapter 28's number. Table 31.5 adds the outcomes: each column is a node of the rule, with the reading it stands for and the load after it. A reading below the predicted 17.701°C says the path is better than thought, and the load rises; the sum is 62.61 kW, which sixty nodes confirm. The rise is 1.708 kW, of which the tightening is 1.606 kW and the curvature 0.102 kW, and it is worth $400 \times 1.708 = \mathrm{R\$}\,683$ a year.

Table 31.5: The second hall-1 sensor by the five-node rule: the standardized innovation, its weight, the reading, and the assurance load after the reading.

η_a	-2.857	-1.356	0	1.356	2.857
Weight	0.01126	0.22208	0.53333	0.22208	0.01126
Reading $\tilde{y}_a$ [°C]	15.720	16.761	17.701	18.641	19.682
$Q_{0.95}(\tilde{y}_a)$ [kW]	73.58	67.54	62.51	57.86	53.11

The value of every candidate. Table 31.6 gives the values on the first action set. Each row is a candidate; the columns give the expected load after the reading, the value of the rule (31.14), the value under the loss (31.13), and the probability that the load moves by more than 2 kW, each on the state route with the full posterior beside it. Figure 31.3 draws the distribution of the load after each reading. The horizontal axis is the load in 1 kW bins, the vertical the probability in each bin, the red curve the state route and the green the full posterior, and the vertical line the present 60.90 kW.

Only the hall-1 sensor moves the load by much, with a reading that also lowers it in about one case in three (0.32 on the state route and 0.31 on the full posterior). The order by money is the same on both routes and under both measures. The full posterior values each candidate higher, by R$ 30 on the hall-1 sensor and by a fifth to a third on the load meter and supply sensor; no gap changes an order or, below, a verdict. The loss adds R$ 108 (R$ 101) to the hall-1 sensor, the fall in expected overshoot, from 0.100 kW (0.111 kW) of load past capacity.

Other action sets. Table 31.7 gives the value of each candidate when the action set changes. Each row is a candidate; the columns give the rule's value with the derate and hall 2 off this chiller, the rule's value with hall 2 kept on it, and the value under the loss when the load sells only in 5 kW blocks, each with the full posterior in parentheses.

Table 31.6: The value of each candidate to the load decision, supply held, $v_{\mathrm{L}} = \mathrm{R\$}\,400$ per kW-year (illustrative): the expected assurance load after the reading, the value of the rule, the value under the loss, the probability the load moves by more than 2 kW, and the median and first percentile of N_{eff} over the outcomes on the full posterior, of 200 000 draws.

	State route				Full posterior				
Candidate	$\mathbb{E}[Q]$	EVSI_Q	EVSI_L	> 2 kW	$\mathbb{E}[Q]$	EVSI_Q	EVSI_L	> 2 kW	N_{eff}
Hall-1 sensor	62.61	683	791	0.61	62.77	713	814	0.58	142 000; 13 000
Load meter	61.13	89	102	0.14	61.26	109	141	0.15	107 000; 9 100
Supply sensor	61.08	71	76	0.16	61.22	94	104	0.10	69 000; 2 900
Return sensor	60.99	35	37	0.04	61.10	46	51	0.02	131 000; 14 000
Hall-2 sensor	60.91	0.2	0.3	0.00	60.99	0.3	0.2	0.00	143 000; 10 000

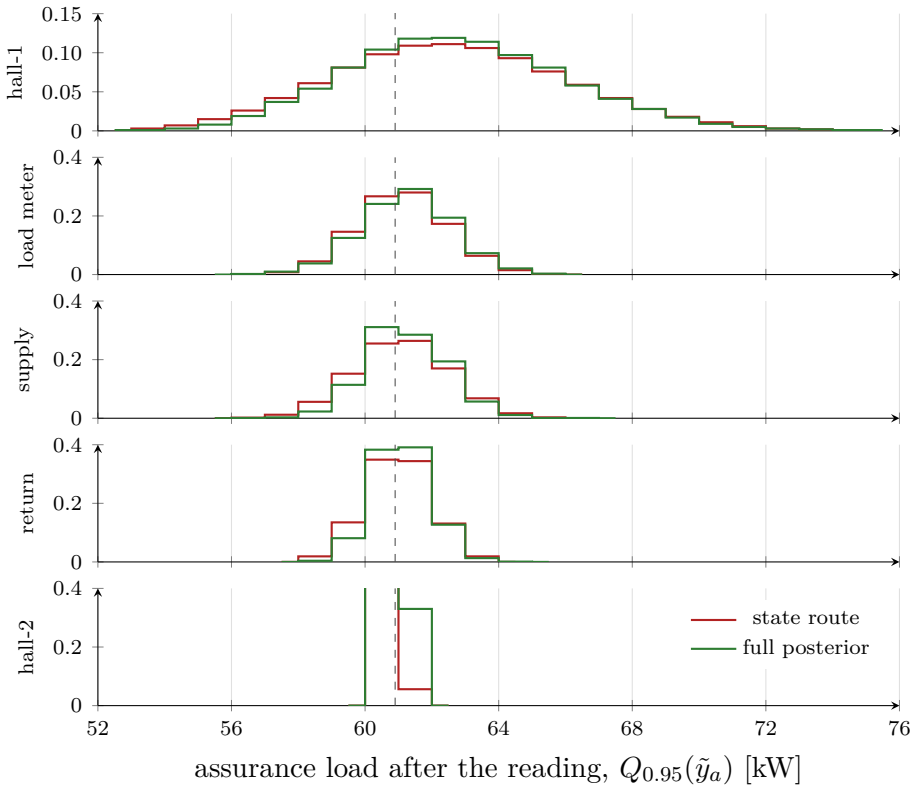


Figure 31.3: The distribution of the assurance load after each candidate's reading, in 1 kW bins, on the state route (red) and the full posterior (green); the dashed line is the present 60.90 kW. The hall-2 sensor's load barely moves: its bars of 0.94 and 0.67 run off the axis, and its split between two bins reflects where a bin edge falls, the load staying within about 0.1 kW.

Table 31.7: The value of each candidate by action set [R$ per year], state route with the full posterior in parentheses. The last column is the Bayes action of the loss (31.2) over blocks of 5 kW; the present action is 60 kW on both routes. The hall-2 sensor's entries are zero within sampling error.

Candidate	Supply held	Derate, hall 2 off	Hall 2 kept	5 kW blocks, loss
Hall-1 sensor	683 (713)	429 (453)	0.3 (1)	664 (710)
Load meter	89 (109)	56 (71)	670 (664)	15 (57)
Supply sensor	71 (94)	38 (54)	1 (3)	−12 (36)
Return sensor	35 (46)	19 (27)	75 (76)	−22 (6)
Hall-2 sensor	0.2 (0.3)	0.1 (−0.5)	4 (3)	−0.02 (0.0)

The derate on the state route. The rule with the derate keeps the delta tail: the center is T_1^* plus the expected shortfall of equation (28.21) over $\dot{m}c_p$, the variance adds $(\Phi(d_X)P_{\mathrm{pl}}/\dot{m}c_p)^2\Sigma_{66}$, and the load is the root in Q by bisection, with $\mathbf{m}_a$ and $\mathbf{\Sigma}_a$ in place of $\hat{\mathbf{z}}$ and $\mathbf{\Sigma}$. At 58.91 kW the shortfall is certain, $d_X = 6.7$, the center is 23.264°C, and each kilowatt adds $0.182 + 0.100$ K, not 0.182. For the hall-1 sensor, with the center held, the load is 59.96 kW; over the outcomes it is 59.98 kW, a rise of 1.071 kW, and $400 \times 1.071 =$ R$ 429. The shortfall's slope is why the same narrowing buys about 0.63 of the kilowatts it buys with the supply held.

With hall 2 kept on this chiller the winner is a different instrument: the chiller's onset load has the gradient of equation (31.4), which meets Q_2 and not the path, so the load meter takes the spread from 2.429 to 1.411 kW and the rise is $1.645 \times 1.018 = 1.674$ kW, while the hall-1 sensor is worth nothing. There is one qualification. The admissible load in that action set stays near the present load: 28.47 kW now and 30.14 kW in expectation after the meter, against $\hat{Q}_1 = 30.69$ kW. Its R$ 670 values all assured load, including load already carried. Counting only assured load above the present, with the present load held at $\hat{Q}_1$, $v_{\mathrm{L}}\mathbb{E}[(Q_\kappa(\tilde{y}_a) - \hat{Q}_1)^+]$, it is R$ 218 (R$ 213). That holding is an assumption: the meter's reading also moves the mean of Q_1, of correlation -0.57 with Q_2, and against the updated mean the figure is R$ 113 (R$ 110). Neither of these two pays at R$ 300. All are stated, since they answer different questions about what the customer pays for.

The negative entries of the last column are not values. Under a coherent posterior equation (31.11) cannot be negative; the full posterior's entries are not. On the state route the proof's step fails, because the delta tail is not a posterior probability averaged over outcomes: at 60.90 kW the average over the hall-1 sensor's readings of the delta probability is 0.943, not 0.950. These entries are small and lie under any verdict, and the state route's error here is one the reader should expect in a coarse action set. Without a loss, the rule alone on blocks does worse: the expected block after the supply sensor is 58.90 kW (59.31), below the present 60 kW, since a poor reading forces the load down a whole block and a good one seldom lifts it.

31.3.5 The most any measurement could be worth

Engineering question. What is the most any measurement could be worth?

General form. The expected value of perfect information, equation (30.18) of Chapter 30. It is the Jensen gap of the minimum, and it bounds every sample value because revealing the state is a reading with no noise:

$$\mathrm{EVPI} = \min_{Q} \mathcal{R}(Q) - \mathbb{E}\big[\min_{Q} L(Q, \mathbf{z}) \mid \mathbf{y}\big] \geq \mathrm{EVSI}(a) \quad \text{for every } a. \tag{31.16}$$

The inequality follows by conditioning: given the reading, the inner minimum on the state is no larger than the minimum of the expectation, and the tower rule removes the reading.

Specialized form. The question is equation (31.16) with the loss (31.2) and the posterior of equation (31.5). One choice enters here: a state revealed serves the load decision only and is obtained free, so this is a ceiling and not a plan. With the state known the overshoot can be avoided exactly, $\min_Q L = -v_{\mathrm{L}} Q_{\mathrm{cap}}(\mathbf{z})$, and the rule's load is the capacity itself, since assurance is then certain. Therefore

$$\mathrm{EVPI}_Q = v_{\mathrm{L}}\big(\mathbb{E}[Q_{\mathrm{cap}} \mid \mathbf{y}] - Q_\kappa\big), \qquad \mathrm{EVPI}_L = v_{\mathrm{L}}\big(\mathbb{E}[Q_{\mathrm{cap}} \mid \mathbf{y}] - \bar{Q}_\kappa\big), \tag{31.17}$$

the first for the rule, the second for the loss. The mean is not the plug-in crossing: under the normal it is closed,

$$\mathbb{E}[Q_{\mathrm{cap}} \mid \mathbf{y}] = \mathrm{e}^{\hat{g}_1 + \Sigma_{11}/2}\big(T_{\max} - \hat{T}_s - \Sigma_{15}\big), \tag{31.18}$$

the normal's moment generating function applied to $(T_{\max} - T_s)\mathrm{e}^{g_1}$. An entry revealed alone keeps the rule, with the conditional covariance and the mean moved along that entry's column of $\mathbf{\Sigma}$.

Building blocks. $\hat{g}_1$, $\hat{T}_s$, Σ_{11}, Σ_{15}; the load (31.6); $\bar{Q}_\kappa$ of equation (31.12), which is 58.90 kW on the state route and 58.76 kW on the full posterior; for the values beside them, the weighted mean of the capacity and of equations (31.3) and (31.4).

The values. Equation (31.18) gives $5.49025 \times 1.005707 \times 12.87553 = 71.09$ kW, the Laplace pushforward's mean of Chapter 28, against the plug-in crossing 70.76 kW and the full posterior's 71.93 kW. So $\mathrm{EVPI}_Q = 400 \times (71.09 - 60.90) =$ R$ 4 076 on the state route and $400 \times (71.93 - 60.99) =$ R$ 4 377 on the full posterior. The state-route figure mixes two distributions of the capacity: 71.09 kW is the true function under the normal, and Q_κ and $\bar{Q}_\kappa$ come from the delta tail, whose own implied mean is 71.52 kW, which would give R$ 4 247; the reported value follows Chapter 28's mean. Table 31.8 lists the values. Each row is what is revealed and on which action set; the columns give the value on each route and the gap. Figure 31.4 draws the ladder: each candidate's value of the rule against the ceiling, on a scale where the ceiling is visible.

Table 31.8: The expected value of perfect information [R$ per year], on the action sets of step 2, with the best candidate's share of each. Revealed entries are on the first action set and the rule.

Revealed, action set and measure	State	Full	Gap	Best candidate's share
State; supply held; rule	4 076	4 377	−301	17 % (16 %), hall-1
State; supply held; loss	4 879	5 268	−389	16 % (15 %), hall-1
State; supply held; 5 kW blocks, loss	3 942	4 336	−394	17 % (16 %), hall-1
State; derate, hall 2 off; rule	2 477	2 650	−173	17 % (17 %), hall-1
State; hall 2 kept; rule	1 598	1 590	8	42 % (42 %), load meter
g_1 alone	3 159	3 429	−270	
T_s alone	81	95	−14	

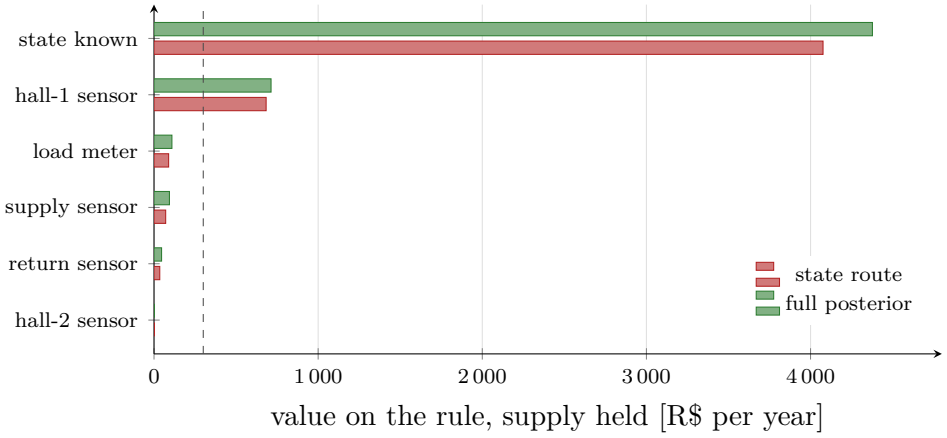


Figure 31.4: The value ladder on the rule with the supply held: each candidate's value and, at the top, the ceiling. The dashed line is the illustrative sensor cost of R$ 300 per year.

The best sensor captures about a sixth of the ceiling on both routes. Most of the ceiling is in the conductance, R$ 3 159 of R$ 4 076.

31.3.6 Why the rankings differ

Engineering question. Why do the rankings differ, and which one to use?

General form. The attribution of the variance a row removes, set beside the criteria of the earlier answers:

$$\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} - \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_{a}\mathbf{g} = \frac{\left(\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{j}_{a}\right)^{2}}{s_{a}}. \tag{31.19}$$

Specialized form. The question is equation (31.19) in the form (31.10), with the gain (31.8) and the value (31.14). No choice enters. Equation (31.8) contains only V_a/σ_a^2; equation (31.10) multiplies a function of it by the squared correlation; and for a decision that is a quantile of one output, the linear rise $z_\kappa(\mathrm{sd_c} - \mathrm{sd_{c,a}})$ of Section 31.3.4 is a width removed in other units, that of the margin near the load decided.

Building blocks. V_a of equation (31.8); $\mathrm{corr}(f, f_a)$ from $\boldsymbol{\Sigma}$; the values of Tables 31.6 and 31.8.

The attribution. Table 31.9 gives the factors. Each row is a candidate; the columns give the ratio of the virtual variance to the sensor's own, the correlation of the quantity read with the margin and with hall 2's load, the share of the margin's variance removed at 60 kW, and the width removed from the margin at the assurance load.

Table 31.9: Why the criteria part, on the state route: the ratio that sets the information gain, the correlations that set the widths, and the width on the margin near the load decided, which sets the money.

Candidate	V_a/σ_a^2	$\mathrm{corr}(m_1^*, f_a)$	$\mathrm{corr}(Q_2, f_a)$	Variance removed, m_1^*	Width at $Q_{0.95}$
Supply sensor	7.843	0.21	−0.04	3.9 %	2.1 %
Load meter	2.447	−0.26	1.00	4.8 %	2.5 %
Return sensor	1.231	0.19	0.42	1.9 %	1.0 %
Hall-2 sensor	0.927	−0.02	0.10	0.01 %	0.01 %
Hall-1 sensor	0.922	−0.84	−0.03	33.9 %	18.6 %

The width on the capacity's own tangent does not set the money: it puts the supply sensor, 2.8 %, above the load meter, 2.4 %, and $z_\kappa$sd on it would give R$ 840, R$ 113 and R$ 128 for the hall-1 sensor, the meter and the supply sensor, against R$ 683, R$ 89 and R$ 71, since that tangent is taken at the plug-in crossing and not at the load decided. The information gain favors the supply sensor because 0.1 K is small against 0.28 K. It says nothing of where that entry stands in any answer, and T_s stands low in both: it contributes $\Sigma_{55} = 0.078\,\mathrm{K}^2$ to a margin's variance of $1.154\,\mathrm{K}^2$, whose largest term is the conductance's, and the value of knowing T_s exactly is only R$ 81. No sensor on T_s, however fine, can exceed that, and the supply sensor's R$ 71 already takes most of it. The hall-1 sensor has the lowest ratio of the five and the one line that reaches g_1, which holds R$ 3 159.

The rule tested on both routes. The rule proposed for practice is to rank by the value of sample information when a decision is named, by the width when a question is named, and by the information gain only when neither is. On this plant, on both routes:

- Three criteria give three winners, on the state route and on the full posterior alike.
- With a named output the width picks the instrument that output needs, and the information gain picks an instrument that takes 2 % off the margin and 0.1 % off hall 2's load.

- With the decision named the money and the width on the margin give the same order. The value adds the verdict and not the order, because this decision is a quantile of one output.
- Money and that width part when the action set changes the output: with hall 2 kept on the chiller the money's winner is the load meter, and the margin's width still points to the hall-1 sensor. Naming the decision is what chooses the output.
- The plant has no case in which nothing is named, so it does not test the last clause. What it shows is what that ranking would cost here: taking the information gain's winner in place of the money's gives up R$ 612 (R$ 619) a year on the first action set. How often the information gain misleads in general is not something one plant can show.

Figure 31.5 carries the question of the three places for one more temperature sensor. The three dashed boxes are the positions, each labeled with its value on the rule and its information gain on the state route. The information gain orders them return, hall-2, hall-1, with hall-2 and hall-1 within 0.001 nat; the money reverses that order.

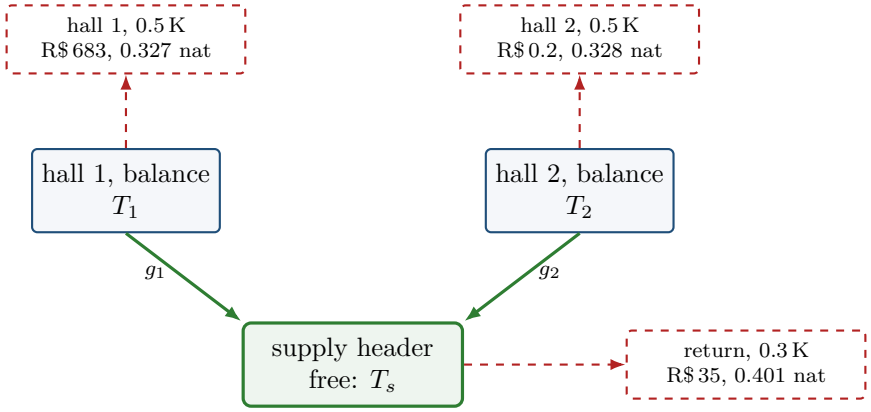


Figure 31.5: Three places for one more temperature sensor, labeled with the value on the rule (supply held, R$ 400 per kW-year, illustrative) and the information gain, on the state route. Free (green): the supply; the log-conductances label the paths. Balance nodes (blue): the halls.

31.4 The checks (step 5)

The answers rest on a rank-one update at a fixed Jacobian, a delta tail after every outcome and a reweighting of fixed draws. Each is set against a fuller calculation, and the two illustrative prices and the priors are varied. No physics row is added; the limit $\sigma_p \to 0$ of the reduced state was checked in Chapter 20.

31.4.1 The routes against each other

Table 31.10 sets each computed quantity beside an independent calculation. Each row is one check; the columns give the quantity, its value by its route, the other calculation and the verdict.

The refit moves the hall-1 sensor's value about two thirds of the way from the state route to the full posterior, and the load meter's the other way; on neither does it cross a verdict. The tail of outcomes is the weak part of the reweighting: over the most extreme readings of the 0.1 K supply sensor N_{eff} falls to about 125, below the tenth of N at which Chapter 20 asks for another sampler; such readings have a small weight in the average, the first percentile being 2 900, and the larger draw sets above measure what they contribute.

Table 31.10: The checks of the preposterior. Amounts in R$ per year unless stated.

Quantity	Route	Independent calculation	Verdict
EIG, all five	$\frac{1}{2}\ln s_a$	$\frac{1}{2}\ln(\det\boldsymbol{\Sigma}/\det\boldsymbol{\Sigma}_a)$	agree to all digits
Q with center held	equation (28.13): 62.51, 61.08, 61.11 kW	Section 28.6.1	agree
$\mathbb{E}[Q]$, hall-1	sixty nodes: 62.612 kW	five nodes: 62.612; 4 000 readings, seed 31: 62.584 (s.e. 0.058)	agree
$\mathbb{E}[Q]$, load meter	61.126 kW	4 000 readings: 61.115 (s.e. 0.022)	agree
$\mathbb{E}[Q]$, supply	61.083 kW	4 000 readings: 61.097 (s.e. 0.023)	agree
$\mathbb{E}[Q]$, return	60.991 kW	4 000 readings: 61.001 (s.e. 0.016)	agree
EVSI_Q against a refit	rank one: 683, 89, 71, 35, 0.2	Gauss–Newton refit on the same 2 000 readings, with $\boldsymbol{\Sigma}$ at the new fit: differences +21, −7.5, +11, +2, +0.3	hall-1 3 % low
Outcomes of the full route	2 000 midpoints of the predictive	tower test over them: mean $\mathbb{E}[Q_{\mathrm{cap}} \mid \tilde{y}]$ 71.928 kW, the present mean; mean $F(60.99\,\mathrm{kW})$ 0.0500, for every candidate	agree
Full posterior, sampling	200 000 draws, seed 17: 713, 109, 94, 46, 0.3	two sets of 10^6 draws, seeds 18 and 19, 400 midpoints: 710, 109, 96, 47, 0.2	within 3
Assurance of the state rule	0.95 on its route	probability under the full posterior after each reading, averaged: 0.9501, 0.9514, 0.9505, 0.9526, 0.9511 (order hall-1, hall-2, supply, meter, return)	holds
$\mathbb{E}[Q_{\mathrm{cap}}]$	equation (31.18): 71.094 kW	2×10^4 draws, seed 31: 70.98 (s.e. 0.05); 2×10^6 draws: 71.094	agree
Identity of the loss	EVSI_L, hall-1: 791	equation (31.13) term by term: $683 + 8\,000\times(0.1005-0.0869) = 792$	agree

31.4.2 The prices, the assurance level and the priors

Table 31.11 repeats the value of the rule on the state route with the priors of Table 28.8, refitting each posterior as a check, and at two other assurance levels. Each row is a case; the columns give the present load, the value of each candidate with the supply held, and the hall-1 sensor's value with the derate and hall 2 off.

Table 31.11: The value on the rule [R$ per year], state route, supply held, under other priors and assurance levels.

Case	Q_κ [kW]	Hall-1	Meter	Supply	Return	Hall-2	Hall-1, derate
As in the chapter	60.90	683	89	71	35	0.2	429
Conductance prior widened, 0.6	61.19	767	104	93	41	0.0	478
Conductance prior narrowed, 0.15	60.44	477	53	22	15	1.4	304
Centered on 4 kW/K	59.79	654	89	66	33	0.2	414
Centered on 6 kW/K	61.92	710	89	76	37	0.2	440
Set-point prior widened, 2 K	60.92	681	89	76	38	0.2	427
$\kappa = 0.90$	62.83	571	75	64	31	0.2	349
$\kappa = 0.99$	57.61	861	110	78	38	0.3	555

The order holds in every case, and so does the verdict at R$ 300 with the supply held. With the derate the margin is thin: at the narrowed prior the hall-1 sensor is worth R$ 304, a net of R$ 4. Its value is largest when the prior on the conductance is weakest.

31.5 The report (step 6)

The most precise instrument on the list, the supply sensor, gains the most information about the state, 1.09 nats, three times the hall-1 sensor's, and is worth R$ 71 (R$ 94) a year to the load decision against a price of about R$ 300. The second hall-1 sensor, which gains the least, is worth R$ 683 (R$ 713) with the supply held. And if hall 2 stays on this chiller, neither of them is worth anything and the load meter is the instrument. The answers follow in the order they were asked.

31.5.1 The answers and their bands

Table 31.12 is the report. Each row is a question in the operator's words; the second column gives the answer from the posterior of the state, the third the full posterior's value, the fourth the gap, the last the qualification.

The report line. With the supply held, a second hall-1 temperature sensor raises the expected assured load of hall 1 from 60.9 kW to 62.6 kW (61.0 to 62.8 kW on the full posterior), at an unchanged 5 % probability of exceeding 25°C, worth R$ 683 (R$ 713) a year at an illustrative R$ 400 per kW-year, about a sixth of what knowing the state would be worth. No other candidate exceeds R$ 109 on that action set. With the derate and hall 2 off this chiller the same sensor is worth R$ 429 (R$ 453). With hall 2 kept on this chiller it is worth nothing, and a hall-2 load meter is worth R$ 670 (R$ 664) if all assured load is paid for and R$ 218 (R$ 213) if only load above the present is, with that load held.

Table 31.12: The report of the candidates: answers from $\mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$ through the rank-one preposterior, the full posterior reweighted for each outcome, and the gap. Money at the illustrative R$ 400 per kW-year.

Engineering question	Answer (state route)	Full posterior	Gap	Qualification
How much does each candidate reduce the uncertainty of the state?	supply 1.090, meter 0.619, return 0.401, hall-2 0.328, hall-1 0.327 nat	1.095, 0.609, 0.400, 0.333, 0.325	at most 0.010	defined on both routes; hall-1 and hall-2 nearly tie
How much width does each take off the margin, and off hall 2's load?	margin: hall-1 18.7 %, meter 2.5 %, supply 2.0 %, return 1.0 %, hall-2 0.01 %; Q_2: meter 46.1 %, return 5.0 %, others under 1 %	18.6 %, 2.5 %, 2.0 %, 1.0 %, 0.0 %; 45.6 %, 4.8 %	at most 0.5 point	0.874 K and 1.264 kW after the winners
Which measurement next?	state: supply; margin: hall-1; Q_2: load meter; decision: hall-1	the same	none	three winners on both routes
What is a candidate worth to the load decision, in money?	on the rule: R$ 683, 89, 71, 35, 0.2; under the loss: R$ 791, 102, 76, 37, 0.3	R$ 713, 109, 94, 46, 0.3; R$ 814, 141, 104, 51, 0.2	−30 to −0.1; −39 to +0.1	hall-1, meter, supply, return, hall-2; a load at unchanged 5 %; derate: hall-1 R$ 429 (R$ 453); hall 2 kept: meter R$ 670 (R$ 664), or R$ 218 (R$ 213) above the present load
What is the most any measurement could be worth?	R$ 4 076 on the rule, R$ 4 879 under the loss	R$ 4 377, R$ 5 268	−301, −389	derate R$ 2 477 (R$ 2 650); hall 2 kept R$ 1 598 (R$ 1 590); g_1 alone R$ 3 159 (R$ 3 429)
Why do the rankings differ, and which one to use?	the gain reads V_a/σ_a^2, the width adds the correlation, the money adds the action set	the same	none	the rule holds on both routes where this plant can test it; the money and the margin's width share an order until the action set changes

31.5.2 The cost comparison

Table 31.13 prices each candidate on one basis: about R$ 300 a year, sensor and calibration, the same for all, which favors the 0.1 K supply sensor, likely the dearest in practice. Each row is a candidate on an action set; the columns give the value on each route, the net value at R$ 300, and the value of admitted load at which it breaks even.

Table 31.13: Cost against value, rule measure, at an illustrative R$ 300 per year for each sensor and R$ 400 per kW-year, state route with the full posterior in parentheses. The break-even price is the value of a kW-year at which the sensor just pays; the break-even cost at R$ 400 is the value itself.

Candidate, action set	Value [R$/yr]	Net [R$/yr]	Break-even [R$/kW-yr]
Hall-1 sensor, supply held	683 (713)	383 (413)	176 (168)
Load meter, supply held	89 (109)	−211 (−191)	1 353 (1 102)
Supply sensor, supply held	71 (94)	−229 (−206)	1 682 (1 274)
Return sensor, supply held	35 (46)	−265 (−254)	3 461 (2 620)
Hall-2 sensor, supply held	0.2 (0.3)	−300 (−300)	–
Hall-1 sensor, derate, hall 2 off	429 (453)	129 (153)	280 (265)
Load meter, hall 2 kept, all assured load	670 (664)	370 (364)	179 (181)
Load meter, hall 2 kept, above the present, held	218 (213)	−82 (−87)	550 (563)

The order within an action set does not depend on either price, since both are common to all candidates; only a different price per instrument could change it, and the supply sensor would have to cost under R$ 71 (R$ 94) to pay. The verdict does depend on them. At R$ 200 per kW-year the hall-1 sensor still pays with the supply held (R$ 342) and no longer with the derate (R$ 214). At R$ 800 no other candidate pays on the first action set, the load meter reaching R$ 177 (R$ 218). The gap between the routes crosses no verdict in the table.

The recommendation. Do not buy the supply sensor, the hall-2 sensor or the return sensor for this decision: they do not pay on any action set at these prices. Which of the other two to buy depends on what the site does with hall 2. Chapter 28 found that hall 1 reaches 60 kW only with hall 2's load off this chiller or with more capacity. If that move is made, buy the second hall-1 sensor: it is worth R$ 429 (R$ 453) against about R$ 300 with the derate, a net of R$ 129 (R$ 153) that a kW-year below R$ 280 (R$ 265) or the narrowed conductance prior would nearly or wholly remove, and R$ 383 (R$ 413) net if capacity is added so that the supply holds. If hall 2 stays, the hall-1 sensor is worth nothing, and the load meter pays only if assurance on load already carried is itself paid for; this chapter states both and does not choose.

Principle 31.1. *A candidate row has three values and they answer three questions: the information gain, of the state; the width removed, of an output; the value of sample information, of a decision on an action set. Rank by what has been named. A rule such as an assurance load earns a money value only with a stated loss: under a fixed cost per exceedance its rise is that gain, at an unchanged probability of exceeding, and under a cost proportional to overshoot the rule is the best action and the value adds the fall in overshoot. A value is reported with its route and action set, and a sensor is bought on a value set against its cost, for the program the site will actually run.*

Common mistake

Five errors recur in sensor placement. Ranking by precision or information gain when an output is named: the supply sensor's 1.09 nats buy 2 % of the margin. Reading a preposterior width as a value: 18.7 % of the margin is not a price until a decision and a value of load are stated. Calling the rise in assured load risk-free: it holds the probability of exceeding at 5 %, at a higher load. Assuming the rule's gain must be nonnegative: on 5 kW blocks the rule alone lowers the expected load after three of the five candidates. And valuing a path sensor when the chiller is the limit: with hall 2 kept on it the hall-1 sensor's R$ 683 is under R$ 1.

31.6 The next measurement (step 7)

The chapter's subject has been the next measurement, so this step asks what comes after the first purchase. A row bought returns the fit to step 3, and the candidates are valued again from $\mathbf{\Sigma}_a$.

31.6.1 The candidates after the first sensor

Table 31.14 values each remaining candidate from the preposterior of the hall-1 sensor, on the rule with the supply held and on the state route. Each row is a candidate; the columns give its value alone, its value after the sensor, and the spread of the margin after it. The last two rows are of another kind, a test row of Chapter 27: hall 1 held at a known load Q_{test} and read once at 0.5 K, with $\mathbf{j}_a = (-Q_{\text{test}}\mathrm{e}^{-\hat{g}_1},\ 0,\ 0,\ 0,\ 1)^{\mathsf{T}}/0.5$, the load pinned and so absent from the row.

Table 31.14: The next candidates on the rule, supply held, state route, R$ 400 per kW-year (illustrative). Test rows are not priced here.

Candidate	Value alone	Value after the hall-1 sensor	sd(m_1^*) after it [K]
Another hall-1 sensor	683	287	0.791
Load meter	89	124	0.838
Supply sensor	71	117	0.842
Return sensor	35	56	0.859
Hall-2 sensor	0.2	0.3	0.874
Test at 40 kW, as if run with the supply held	1 542	1 015	0.589
Test at 50 kW, as if run with the supply held	1 932	1 346	0.502

Two things stand out. After the first sensor no other instrument on the list pays at R$ 300 except possibly another hall-1 sensor, worth R$ 287 on the state route and R$ 300 on the full posterior: that reaches the threshold, and both are stated without choosing. The load meter and the supply sensor are worth more after it than alone: the hall-1 reading confounds Q_1 with g_1, and each of them helps to separate the two. Values do not add; here they are not even smaller in sequence, as Exercise 31.5 shows. And the largest values on the table are not of a sensor. A reading at a known higher load has the lever on g_1 that a reading at the present load lacks, as Chapter 27 derived. The test must itself run on the chiller, though, and a 50 kW test leaves $54.7 - 50 = 4.7$ kW for hall 2 and a 40 kW test 14.7 kW, against 22.2 kW. It therefore requires hall 2 off this chiller, where the derate applies and its value is R$ 956 at 40 kW and R$ 1 190 at 50 kW, not the R$ 1 542 and R$ 1 932 of the table; with hall 2 kept it cannot run, and its row would be worth about nothing to the chiller's limit. These values are on the state route alone, and the test is not priced.

31.6.2 What the analysis showed

The vendor's reading was that the most precise instrument, on the entry it would cut furthest, is the one to buy. The analysis prices that instrument at about a quarter of its cost, and the instrument with the least information gain at more than twice its cost with the supply held, about 1.4 times with the derate and hall 2 off, and nothing if hall 2 stays. For the site the first question is not which sensor but what hall 2 will do, since that decides whether the limit is the path or the chiller, and the path sensor and the load meter trade places accordingly. Most of the four thousand reais a year that knowing the state would be worth lies in the conductance, which a load test with hall 2 off reaches better than another sensor does. The numbers rest on an illustrative price per kilowatt, and the order within an action set does not.

31.6.3 Limits of the model

The derivations are exact under their assumptions, and seven need checking before a value is used.

One reading. A sensor reads every day; the values here are of one reading of σ_a. Averaging shrinks only the independent part of the error, the part a shared calibration bias leaves, and a sensor's lifetime value lies between one reading and the limit of its bias. A one-off instrument would be set against a one-off price.

One decision. The sensors serve other questions, and a value to one decision is a lower bound on a sensor's use. Summing across decisions needs care, since they share the same posterior. How to spend across decisions and against investments is taken up in the next chapter.

The Gaussian preposterior. The reported values use one Jacobian. Section 31.4.1 measures that at 3 % on the hall-1 sensor against a refit and 4 % against the full posterior.

The value per kW. It is stated, not derived, and it assumes the load sells, stays sold, and can be moved down after a poor reading. A load that cannot be moved down makes the rule's downside a real exceedance risk, and the values here are then too high.

The overshoot price. c_{o} is what κ implies, not a measured cost. A real cost of overshoot, in equipment life or in service credits, would replace κ and with it every load in this chapter.

The calibration behind the chiller's limit. The value with hall 2 kept uses the spread of c_{ch} with a calibrated meter. With the meter's bias free, as in Section 28.4.4, that spread is 0.0313; the bias also enters the heat-meter row on $Q_1 + Q_2$, so c_{ch} and Q_2 become correlated, the onset load keeps a spread of 2.44 kW, and after the meter it is 1.81 kW. The load meter's R$ 670 becomes R$ 418 on that refit, and it still pays.

The conductance at load. All the path's values assume the conductance inferred near 31 kW holds at 60 kW, the limit of Chapter 20. A hall-1 sensor at the present load cannot test that, and a load test can.

The next chapter takes up what this one left out: how a value of information competes with an investment in capacity for the same money.

31.7 What to retain

Definition 31.1 (Three values of a candidate row)**.** For a candidate row of Jacobian line $\mathbf{j}_a$ on a posterior $\mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$, with $s_a = 1 + \mathbf{j}_a^{\mathsf{T}} \mathbf{\Sigma} \mathbf{j}_a$: the information gain is $\frac{1}{2} \ln s_a$; the variance of an output of gradient $\mathbf{g}$ falls by $(\mathbf{g}^{\mathsf{T}} \mathbf{\Sigma} \mathbf{j}_a)^2 / s_a$; and the value to a decision is $\min_a \mathcal{R}(a) - \mathbb{E}_{\tilde{y}}[\min_a \mathbb{E}[L \mid \mathbf{y}, \tilde{y}]]$, which on a coherent posterior is nonnegative and bounded by the expected value of perfect information; a route that replaces the posterior's probabilities by a delta tail can break both.

Four conclusions carry beyond the plant. The three criteria measure different things, and on one plant they named three winners on both routes. The gain reads a ratio of variances, the width multiplies it by a squared correlation, and for a decision that is a quantile of one output the money follows that width near the load decided, so that the value's distinct contribution is the verdict and the dependence on the action set. A rule is not a loss: its rise needs a stated loss to be a value, and the assurance load is the critical fractile of a loss linear in the overshoot, whose optimal cost is the lower-tail mean. And values depend on the program: a path sensor is worth nothing where the chiller binds, and sensors can be worth more together than apart. Table 31.15 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 31.15: The terms of the value of a measurement and where each is defined.

Term	Expression	Units	Meaning	Defined in
Candidate row	$r_a = (f_a(\mathbf{z}) - \tilde{y}_a)/\sigma_a$	–	a reading not yet taken	31.1.3
Loss of a load	$-v_{\mathrm{L}} Q + c_{\mathrm{o}}(Q - Q_{\mathrm{cap}})^+$	R$/yr	the rule is its action	31.1.3
Onset load	$Q_{\mathrm{on}} = P_{\mathrm{pl}} c_{\mathrm{ch}} - Q_2$	kW	limit, hall 2 kept	31.1.3
General forms				
Information gain	$h(\tilde{y}_a \mid \mathbf{y}) - h(\tilde{y}_a \mid \mathbf{z})$	nat	of the state	31.3.1
Spread after the reading	$\mathbb{E}_{\tilde{y}_a}[\mathrm{sd}(f \mid \mathbf{y}, \tilde{y}_a)]$	varies	of an output	31.3.2
Sample information	$\min_Q \mathcal{R}(Q) - \mathbb{E}_{\tilde{y}_a}[\min_Q \mathbb{E}[L \mid \mathbf{y}, \tilde{y}_a]]$	R$/yr	of a decision	31.3.4
Perfect information	$\min_Q \mathcal{R}(Q) - \mathbb{E}[\min_Q L \mid \mathbf{y}]$	R$/yr	the ceiling	31.3.5
Specialized forms and answers				
Gain, Gaussian	$\frac{1}{2} \ln s_a$, $s_a = 1 + V_a/\sigma_a^2$	nat	a ratio of variances	31.3.1
Width, Gaussian	$1 - \mathrm{corr}(f, f_a)^2 V_a/(V_a + \sigma_a^2)$	–	variance ratio	31.3.2
Mean after a reading	$\hat{\mathbf{z}} + \eta_a \mathbf{\Sigma} \mathbf{j}_a / \sqrt{s_a}$	varies	along one line	31.3.4
Lower-tail mean	$\bar{Q}_\kappa = (1-\kappa)^{-1} \int_\kappa^1 Q_{\kappa'} \, \mathrm{d}\kappa'$	kW	optimal cost $-v_{\mathrm{L}} \bar{Q}_\kappa$	31.3.4
Value of the rule	$v_{\mathrm{L}}(\mathbb{E}[Q_\kappa(\tilde{y}_a)] - Q_\kappa)$	R$/yr	at unchanged $1 - \kappa$	31.3.4
Ceiling of the rule	$v_{\mathrm{L}}(\mathbb{E}[Q_{\mathrm{cap}}] - Q_\kappa)$	R$/yr	state revealed	31.3.5

31.8 Exercises

Exercise 31.1. *Formulation.* A two-entry posterior has $\mathbf{\Sigma} = \begin{bmatrix} 0.04 & 0.01 \\ 0.01 & 0.09 \end{bmatrix}$ and a candidate has $\mathbf{j}_a = (2,\ 0)^{\mathsf{T}}$. For $f = z_1 + z_2$, compute EIG, the spread before and after, and the width removed. ($\mathbf{\Sigma} \mathbf{j}_a = (0.08,\ 0.02)^{\mathsf{T}}$, $s_a = 1.16$, EIG $= 0.074$ nat. $\mathbf{g}^{\mathsf{T}} \mathbf{\Sigma} \mathbf{g} = 0.15$, $\mathbf{g}^{\mathsf{T}} \mathbf{\Sigma} \mathbf{j}_a = 0.10$, so the variance after is $0.15 - 0.01/1.16 = 0.1414$; spreads 0.387 and 0.376, a width of 2.9 %.)

Exercise 31.2. *Formulation.* A path's capacity is 60 kW with probability 0.3 and 70 kW with 0.7; the load is 60 or 70 kW; the loss is equation (31.2) with $\kappa = 0.95$. A test says "high" or "low" and is right with probability 0.9 in either case. Find the action now, the EVSI and the EVPI. (Now: $\mathcal{R}(60) = -24\,000$ and $\mathcal{R}(70) = -28\,000 + 0.3 \times 80\,000 = -4\,000$, so 60 kW. Pr(high) $= 0.66$, and then Pr(70 | high) $= 0.9545$ and $\mathcal{R}(70) = -24\,364$; after "low" the action

stays. $\bar{L}_{\text{after}} = 0.66 \times (-24\,364) + 0.34 \times (-24\,000) = -24\,240$, so EVSI is R$ 240; known state gives $-26\,800$, so EVPI is R$ 2 800.)

Exercise 31.3. *Derivation.* Derive the determinant form (31.8) and the correlation form (31.10). ($\boldsymbol{\Sigma}_a^{-1} = \boldsymbol{\Sigma}^{-1} + \mathbf{j}_a\mathbf{j}_a^{\mathsf{T}}$, so $\det \boldsymbol{\Sigma} / \det \boldsymbol{\Sigma}_a = \det(\mathbf{I} + \boldsymbol{\Sigma}\mathbf{j}_a\mathbf{j}_a^{\mathsf{T}}) = 1 + \mathbf{j}_a^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{j}_a$. From equation (20.29), with $\mathbf{j}_a = \mathbf{g}_a/\sigma_a$, the removed share is $(\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_a)^2/[\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}\,(\sigma_a^2 + V_a)] = \operatorname{corr}^2 V_a/(V_a + \sigma_a^2)$.)

Exercise 31.4. *Derivation.* (a) For a capacity normal with mean μ and spread $\mathrm{sd_c}$ linear in the state, derive the expected rise of Q_κ after a reading, and say when it is negative. (b) Show that under equation (31.2) the value is $v_{\mathrm{L}}(\mathrm{sd_c} - \mathrm{sd}_{\mathrm{c},a})\phi(z_\kappa)/(1-\kappa)$, and derive EVSI $\leq$ EVPI. ((a) $Q_\kappa(\tilde{y}) = \mu(\tilde{y}) - z_\kappa \mathrm{sd}_{\mathrm{c},a}$ and $\mathbb{E}[\mu(\tilde{y})] = \mu$, so the rise is $z_\kappa(\mathrm{sd_c} - \mathrm{sd}_{\mathrm{c},a})$, negative for $\kappa < \frac{1}{2}$, where $z_\kappa < 0$. (b) The lower-tail mean of a normal is $\mu - \mathrm{sd_c}\phi(z_\kappa)/(1-\kappa)$; averaging over the outcome leaves μ. With the state known the inner minimum in equation (31.11) is taken inside the expectation over the state, which lowers it further: $\mathbb{E}[\min_Q L \mid \mathbf{y}, \tilde{y}] \leq \min_Q \mathbb{E}[L \mid \mathbf{y}, \tilde{y}]$.)

Exercise 31.5. *Application.* Value two rows bought together on the plant, on the state route with the supply held: write the rank-two preposterior and compute the rise for two hall-1 sensors, and for the hall-1 sensor with the load meter. ($\boldsymbol{\Sigma}_a = (\boldsymbol{\Sigma}^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a)^{-1}$, with the mean moved by $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma} - \boldsymbol{\Sigma}_a)$. Two hall-1 sensors: R$ 971 against a sum of R$ 1 366, less than the sum. Hall-1 with the load meter: R$ 807 against R$ 772, more than the sum: the meter removes the confounding with Q_1 that limits the hall-1 reading, through the (Q_1, Q_2) correlation of -0.57. Hall-1 with the supply sensor: R$ 801 against R$ 754. The expected information gain of the meter pair, 0.945 nat, stays just below its sum, 0.946: a value can be superadditive where the gain is not.)

Exercise 31.6. *Application.* Replace the loss by a fixed penalty C_{x} per exceedance, $L = -v_{\mathrm{L}}Q + C_{\mathrm{x}}\mathbf{1}[Q > Q_{\mathrm{cap}}(\mathbf{z})]$. Find the condition for the best load, the C_{x} that the present 60.90 kW implies, the values for larger penalties, and what fails for small ones and for large loads. (The condition is $v_{\mathrm{L}} = C_{\mathrm{x}} f(Q)$, with f the density of the capacity at Q, on the lower side of its peak. On the state route's tail $f(60.90) = 0.0201$ per kW, so $C_{\mathrm{x}} = $ R$ 19 850. The values of the local minimum are hall-1 R$ 699, meter 89, supply 72, return 35, hall-2 0.2; at R$ 40 000 the load is 58.85 kW and they are 807, 103, 77, 38, 0.3; at R$ 80 000, 57.29 kW and 885, 112, 79, 39, 0.3; at R$ 10 000, 64.07 kW and 528, 66, 58, 28, 0.1. The order holds. Since f peaks at 0.0581 per kW, a stationary minimum exists only for C_{x} above R$ 6 889, and it is never global: the penalty is bounded and $-v_{\mathrm{L}}Q$ is not, so past the tail more load always wins, at R$ 19 850 above about 108 kW. A fixed penalty fixes a load only as a local minimum, within a range the site must state.)

Chapter 32

Redundancy, Spares and Reliability Investment

32.1 Three requests on one budget

The facility manager of the hall has three lines to defend in the next budget: how many CRAC units to install, how many spare compressors to hold, and whether to add a second controller. The record is the one the earlier chapters read: three CRAC units found down in a hundred checks, four pumps in four hundred, one controller in a thousand, and twenty-four compressor failures in 116 unit-years. A unit costs R$ 25 000 a year over its life, a controller R$ 40 000, a spare R$ 3 000, and a failure that finds no spare R$ 45 000.

The quickest reading takes the fractions found down as the truth and prices each line there. On the plant of Chapter 25 a fifth unit then removes R$ 810 359 a year, a second controller R$ 173 947, and a sixth unit after both only R$ 2 789 against its R$ 25 000, so the plan stops at five units and a second controller, R$ 332 620 a year all told. A Poisson count at the fitted rate asks for fourteen spares at R$ 46 057. Each of those is the cost of one state, and three findings in a hundred leave the CRAC unavailability far from settled. The question of this chapter is

what each investment costs in expectation over what the records leave uncertain, which whole number of units and spares and which plant changes minimize that cost, how far that stands from the action at the fitted state and in which direction, whether a stated availability target asks for the same thing, and what learning the state would be worth.

What changes from Chapters 29 and 30 is that the action is an integer, so the condition for an optimum is a comparison of one more unit against its price, and that three decisions share one state.

Jargon: redundancy level, annualized cost, marginal saving, newsvendor problem, critical fractile, benefit–cost ratio, residual risk, availability target

The *redundancy level* is the number n of units installed where k are needed. An *annualized cost* spreads a capital cost and its upkeep evenly over the years of its life. The *marginal saving* of an integer action is the loss that one more step of it removes. The *newsvendor problem* chooses a stock before an uncertain demand, against a cost for each item held and a cost for each item short; its optimum is a quantile of the demand, at the *critical fractile*. A *benefit–cost ratio* divides the expected loss a change removes by its annualized cost. The *residual risk* is the expected loss that remains after a change. An *availability target* requires an unavailability below a stated value with a stated assurance.

Key idea

An integer action stops where the expected marginal saving falls below the price of one more step, so the direction in which a plug-in action errs is set by the curvature of the marginal saving in the state, not by the curvature of the loss. That curvature is taken in the coordinate of the route: the logit or log rate on the state route, whose center is the median, and q or λ on the exact route, whose center is the mean. A loss may be convex in the state at every action and the plug-in action may still fall on either side. Convexity of the loss says that the plug-in's cost is too low; it does not say that its action is too small.

32.1.1 The marginal saving of an integer action

Let the action be a whole number a, each step of which has the certain cost c, so that the loss at the state is $L(a, \mathbf{z}) = c\,a + \ell(a, \mathbf{z})$ with ℓ the loss that the action leaves. The saving of one more step at the state is

$$m_a(\mathbf{z}) = \ell(a, \mathbf{z}) - \ell(a+1, \mathbf{z}). \tag{32.1}$$

If the state were known, the step would be taken while $m_a(\mathbf{z}) > c$. Under the posterior the certain cost stays and the saving averages: $\mathcal{R}(a) - \mathcal{R}(a+1) = \mathbb{E}[m_a(\mathbf{z}) \mid \mathbf{y}] - c$. When the expected saving falls as a grows, the Bayes action is the first a at which it no longer exceeds c, and the plug-in action the first at which $m_a(\hat{\mathbf{z}})$ no longer does. The two conditions differ by, to second order,

$$\mathbb{E}[m_a(\mathbf{z}) \mid \mathbf{y}] - m_a(\hat{\mathbf{z}}) \approx \tfrac{1}{2}\,\mathrm{tr}(\mathbf{H}_m \boldsymbol{\Sigma}), \qquad \mathbf{H}_m = \left.\frac{\partial^2 m_a}{\partial \mathbf{z}\,\partial \mathbf{z}^\mathsf{T}}\right|_{\hat{\mathbf{z}}}, \tag{32.2}$$

the first line of equation (2.24) applied to m_a. Equation (32.2) is written in the state, the reported route; on the exact route the same second-order form holds with $\mathbf{H}_m$ and $\boldsymbol{\Sigma}$ in the unavailability or the rate about its mean. Where m_a is convex the Bayes action takes at least as many steps as the plug-in action, and where it is concave at most as many. $\mathbf{H}_m$ is a difference of two Hessians of the loss, and two convex functions may have a difference of either curvature. Both instances of the chapter have the marginal saving in closed form, in step 2.

32.1.2 The hall, the plant, the fleet and their graph (step 0)

Three plants of earlier chapters carry the three lines. *The hall* is case A of Chapter 13, its cooling alone: n CRAC units of which three must run, one shared unavailability, and only the loss of cooling priced. *The plant* is that of Chapter 25: n CRAC units of which three must run, two pumps of which one must run, a controller that must run, or two of which one must, and a schedule that also prices the hours with a unit down. *The fleet* is that of Chapter 11: the compressors, with the coming year's $\tau = 44$ unit-years. The hall and the plant are different objects, and a unit's value differs between them: the hall has no other route to a loss of cooling and no charge for a unit down. Neither is substituted for the other below. The first question is asked of the hall, where availability as a function of n stands clear, and the fourth of the plant, where the changes were priced.

Figure 32.1 draws the whole. The four green boxes are the free nodes, each with its one entry; the dashed arrows reach them from their records. The CRAC node feeds both the hall's enumeration and the plant's, the first through the n units alone and the second with the pump and controller nodes. The compressor node feeds the coming year's count. The blue boxes are the actions, none of which is an entry of the state, and the boxes at the right are the three losses, into which the gray pinned prices enter.

Table 32.1 lists the pinned quantities. Each row is a quantity, with its symbol, value and the decision that uses it. The loss rates are those of Table 25.1; the prices are illustrative, and Section 32.3.1 and Section 32.3.4 state where the decisions turn on them.

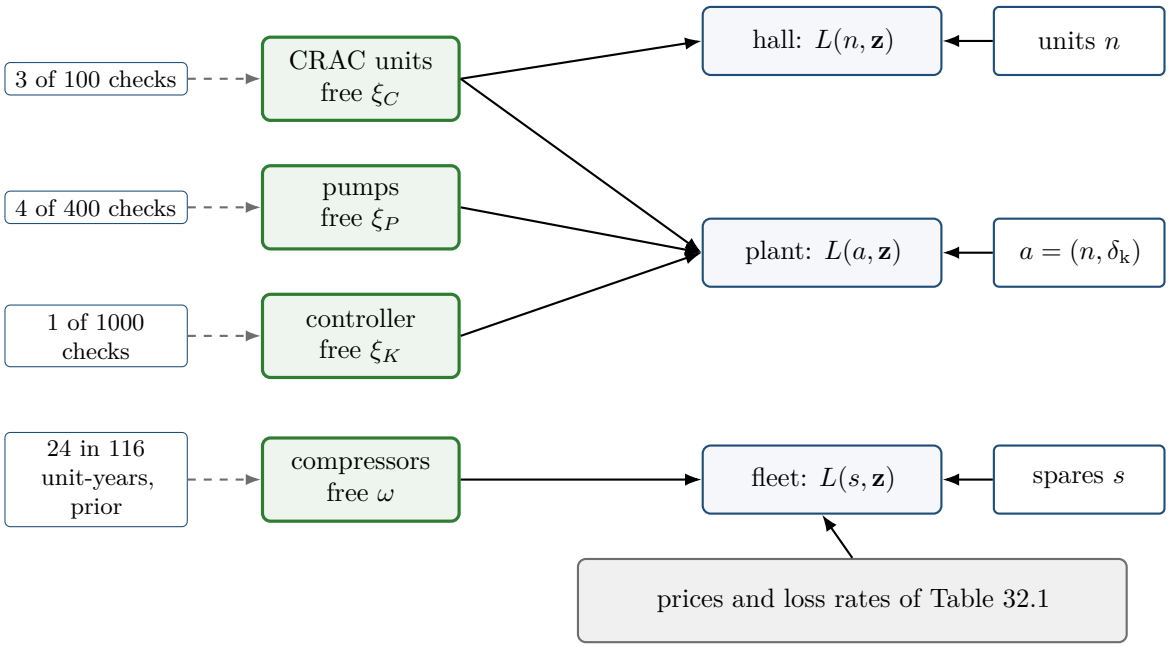


Figure 32.1: The plant graph of the three investments. Free: the four shared entries (green). Pinned: the prices and loss rates (gray), which enter all three losses, the arrow drawn to one of them. Observed: the check records and the yearly counts (dashed). The actions (blue) are not entries of the state. The hall and the plant read the same CRAC entry and are different structures.

Table 32.1: The prices and loss rates of the three decisions (illustrative).

Quantity	Symbol	Value	Used by
A further CRAC unit, annualized over ten years with upkeep	c_{u}	R\$ 25 000 yr^{-1}	hall, plant
A second controller, annualized	c_{k}	R\$ 40 000 yr^{-1}	plant
Loss of cooling	C_{L}	R\$ 20 000 h^{-1}	hall, plant
One CRAC unit down with the cooling kept	C_{C}	R\$ 150 h^{-1}	plant
A pump down with the cooling kept	C_{P}	R\$ 50 h^{-1}	plant
A spare compressor held for the year	c_{s}	R\$ 3 000	fleet
A failure that finds no spare	C_{s}	R\$ 45 000	fleet

32.1.3 The probabilistic questions (step 1)

The manager asks six questions, each written over the state $\mathbf{z} = (\xi_C, \xi_P, \xi_K, \omega)^{\mathsf{T}}$. Table 32.2 lists them. The first column is the question in engineering language; the second the general form of its category (Chapter 2); the third its form here; the fourth its category; the last the subsection that answers it. The superscript B marks a Bayes action and $\hat{n}$, $\hat{s}$ the plug-in actions.

Table 32.2: The questions asked of the three investments, the general form of their category, their form here, their category and where each is answered. $\mathbf{y}$ is the four records; U_n is the unavailability of n units needing three.

Engineering question	General form	Mathematical form	Category	Answered in
Redundancy: how many units? Availability as a function of n against cost	$a^* = \arg\min_a \mathcal{R}(a)$	$\arg\min_n\ (n-3)c_{\mathrm{u}} + C_{\mathrm{L}} t_{\mathrm{yr}} \mathbb{E}[U_n(q_C) \mid \mathbf{y}]$	design (expected loss)	32.3.1
How sure is that number?	$\Pr(L(a,\mathbf{z}) < L(a',\mathbf{z})\ \forall a' \mid \mathbf{y})$; quantiles	the probability that each n is cheapest at the state; the band of the loss	the posterior of the decision	32.3.2
How many spares?	$a^* = \arg\min_a \mathcal{R}(a)$	$\arg\min_s\ c_{\mathrm{s}} s + C_{\mathrm{s}} \mathbb{E}[(K_f - s)^+ \mid \mathbf{y}]$	design (the count predictive)	32.3.3
Which plant change pays, and by how much?	$\mathcal{R}(a') - \mathcal{R}(a)$; $\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y})$	the expected reduction against the annualized cost, the ratio and its band, the residual after each	design (investment)	32.3.4
Why does the plug-in action err, and in which direction?	$\mathcal{R}(a') - \min_a \mathcal{R}(a)$; equation (32.2)	$\hat{n}$, $\hat{s}$ against n^{B}, s^{B}; the curvature of the marginal saving	the Bayes action against the plug-in action	32.3.5
Does an availability target or a cost minimum decide?	$\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}) \geq \kappa$	$\Pr(U_n < 10^{-4} \mid \mathbf{y}) \geq 0.90$ against the cost minimum	two decision rules	32.3.6

32.1.4 The evidence as rows (step 2)

Step 2 adds no new kind of row. The three check records are the binomial rows of Chapter 25 and the fleet's are the four count rows of Chapter 11, in the working form of equation (4.43). Each row acts on one entry:

$$\mathbf{J} = \begin{bmatrix} * & 0 & 0 & 0 \\ 0 & * & 0 & 0 \\ 0 & 0 & * & 0 \\ 0 & 0 & 0 & * \\ \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & * \end{bmatrix} \begin{matrix} \text{CRAC checks} \\ \text{pump checks} \\ \text{controller checks} \\ \text{prior row on } \omega \\ \vdots \\ \text{year 3} \end{matrix}, \qquad \text{columns } \xi_C, \xi_P, \xi_K, \omega, \tag{32.3}$$

so that the covariance is diagonal and each decision averages over its own entries. No physics row arises, so no tolerance σ_p does.

The modeling choices. Five choices make the losses. The same CRAC record serves the hall and the plant: this is what lets one entry carry both, and it holds because Chapters 13 and 25 read the same three findings in a hundred checks. A further unit shares the entry of the others and fails independently given it, as both chapters take it: this values a unit with no record of its own, and it fails for another model or a common cause, which Chapter 15 prices.

The loss is linear in the time without cooling at one rate: this is equation (25.11), and the same R\$ 20 000 h^{-1} is applied from eight percent of the year at three units to a few minutes at six, where neither a shed load nor its restart need cost the same per hour. The prices are constant annual sums and the posterior stands for the life of the purchase: this makes each decision one comparison, and it leaves out discounting and any learning during the ten years. A spare is bought before the year and each failure beyond the store costs C_{s} once: this is the choice of Section 11.3.2, and it leaves out lead times, which Exercise 32.5 takes up.

The loss on the hall. Equation (13.19) with $k = 3$ gives the unavailability, and equation (25.11) with the top event alone gives the annual loss,

$$L(n, \mathbf{z}) = (n-3)c_{\mathrm{u}} + C_{\mathrm{L}} t_{\mathrm{yr}}\, U_n(q_C), \qquad U_n(q) = \sum_{j=0}^{2} \binom{n}{j} (1-q)^j q^{\,n-j}, \tag{32.4}$$

of the form $c_a + \sum_e C_e \Pr(e \mid \mathbf{z}, a)$ of equation (2.23), with $U_5 = 10q^3 - 15q^4 + 6q^5$ and $U_6 = 15q^4 - 24q^5 + 10q^6$. Its marginal saving has a closed form. Let D be the number down among n units. With n units the service is lost when $D \geq n-2$; with one more, by the pivotal decomposition (13.23) on the new unit, $U_{n+1} = q \Pr(D \geq n-2) + (1-q) \Pr(D \geq n-1)$. The difference leaves the configurations in which exactly two of the n run and the new unit is up, the only ones the new unit saves:

$$\begin{aligned} U_n(q) - U_{n+1}(q) &= (1-q) \Pr(D = n-2) = \binom{n}{2} q^{\,n-2} (1-q)^3, \\ m_n(q) &= C_{\mathrm{L}} t_{\mathrm{yr}} \binom{n}{2} q^{\,n-2} (1-q)^3. \end{aligned} \tag{32.5}$$

The ratio $m_{n+1}/m_n = q(n+1)/(n-1)$ is below one for q below $(n-1)/(n+1)$, so the saving falls with n over the states that carry the posterior. Its curvature depends on the coordinate. In q, the exact route's, m_3 is concave below $q = 0.5$; m_4 is convex below $q = 0.155$ and m_5 below 0.276. In the state, with $a_n = n - 2$ and $\mathrm{d}q/\mathrm{d}\xi = q(1-q)$,

$$\frac{\partial^2 m_n}{\partial \xi^2} = m_n \Big[\big(a_n - (a_n + 3)q \big)^2 - (a_n + 3) q (1-q) \Big], \tag{32.6}$$

so that m_3 is convex below $q = 0.10$, m_4 below 0.216 and m_5 below 0.311. Near the record the two coordinates agree from $n = 4$ on and disagree at $n = 3$.

The loss on the plant. The action is $a = (n, \delta_{\mathrm{k}})$ with $\delta_{\mathrm{k}} = 1$ for a second controller. Equations (25.2) and (25.11), with the schedule's rule of a cost per CRAC unit down, which Section 25.3.5 extends to five units and which is extended here to more, a choice named where it enters, give

$$\begin{aligned} L(a, \mathbf{z}) &= (n-4)c_{\mathrm{u}} + \delta_{\mathrm{k}} c_{\mathrm{k}} + \mathcal{R}_{\mathrm{yr}}(\mathbf{q}; a), \\ \mathcal{R}_{\mathrm{yr}}(\mathbf{q}; a) &= t_{\mathrm{yr}} \Big[C_{\mathrm{L}} (1 - R_n Q K) + C_{\mathrm{C}} M_n Q K + C_{\mathrm{P}} R_n \, 2 q_P (1 - q_P) K \Big], \end{aligned} \tag{32.7}$$

with $R_n(q_C) = \Pr(D \leq n-3)$, $M_n(q_C) = \mathbb{E}[D\, \mathbf{1}(D \leq n-3)]$, $Q = 1 - q_P^2$, and $K = 1 - q_K$ or $1 - q_K^2$. The hall's cost counts units above three and the plant's above four, the plant's present design, which changes no comparison.

The loss of a spares level. Given the state the coming year's count is Poisson with mean $\mu(\omega) = \tau \mathrm{e}^{\omega}$, and

$$L(s, \mathbf{z}) = c_{\mathrm{s}} s + C_{\mathrm{s}} \mathbb{E}\big[(K_f - s)^+ \mid \omega\big], \qquad L(s+1, \mathbf{z}) - L(s, \mathbf{z}) = c_{\mathrm{s}} - C_{\mathrm{s}} \Pr(K_f > s \mid \omega). \tag{32.8}$$

The second equation holds because one more spare covers one more failure exactly when the count exceeds s. Its marginal saving therefore is $m_s = C_{\mathrm{s}} \Pr(K_f > s \mid \omega)$. In the mean μ, $\partial \Pr(K_f > s)/\partial \mu = \mu^s \mathrm{e}^{-\mu}/s!$, the Poisson probability of exactly s, whose own derivative is $(s/\mu - 1)$ times it. Hence

$$\frac{\partial^2 \Pr(K_f > s \mid \mu)}{\partial \mu^2} > 0 \iff \mu < s, \qquad \frac{\partial^2 \Pr(K_f > s \mid \omega)}{\partial \omega^2} > 0 \iff \mu < s + 1, \tag{32.9}$$

the first being the exact route's coordinate and the second the state route's, since $\partial/\partial\omega = \mu\, \partial/\partial\mu$. The loss itself is convex in μ at every s, since $(K_f - s)^+$ is convex and the Poisson family is ordered in the convex sense by its mean. The marginal saving is S-shaped, and that is the distinction equation (32.2) turns on.

32.2 The posterior of the state (step 3)

All six questions are read from one posterior, whose four entries come from records already fitted. Nothing new is fitted; the state stacks them.

32.2.1 The posterior from the rows

The rows are those of equation (32.3), each entry with its own Newton iteration: Table 14.2 for the three logits and Table 11.4 for the log rate. With the logit and the log rate both canonical links, the curvature term $\mathbf{C}$ of equation (4.44) is zero and the information is $\mathbf{J}^{\mathsf{T}}\mathbf{J}$; with unnormalized rows it would be $\mathbf{J}^{\mathsf{T}}\mathbf{W}\mathbf{J}$. Equations (25.8) and (11.15) give, at the MAP state, the seven-row Jacobian with the single nonzero entries 1.7059, 1.9900 and 0.9995 in the first three columns and $\sqrt{\mu_k} = (1.4365,\ 2.7255,\ 2.8730,\ 2.8730)$ in the fourth, and

$$\hat{\mathbf{z}} = \begin{bmatrix} -3.4761 \\ -4.5951 \\ -6.9068 \\ -1.5782 \end{bmatrix}, \qquad \mathbf{\Lambda} = \mathrm{diag}(2.9100,\ 3.9600,\ 0.9990,\ 26), \qquad \mathbf{\Sigma} = \mathrm{diag}(0.34364,\ 0.25253,\ 1.00100,\ 0.03846), \tag{32.10}$$

that is $\hat{q}_C = 0.030$, $\hat{q}_P = 0.010$, $\hat{q}_K = 0.001$ and $\hat{\lambda} = 0.2063$ per unit-year, with the spreads $\sqrt{\Sigma_{ii}} = (0.5862,\ 0.5025,\ 1.0005,\ 0.1961)$. The first entry is $\hat{\xi}$ and σ_ξ of equation (13.12).

32.2.2 The exact posterior and the routes

The four entries are independent and each has a conjugate posterior: $q_C \sim \mathrm{Beta}(3, 97)$, $q_P \sim \mathrm{Beta}(4, 396)$, $q_K \sim \mathrm{Beta}(1, 999)$, Table 14.3, and $\lambda \sim \mathrm{Gamma}(26, 126)$, equation (11.11). *The state route* takes $\mathbf{z}$ as $\mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$. *The exact route* takes those four. The MAP state maps to the means of all four, so the plug-in value is one number on both routes and, on the exact route alone, is evaluated at the posterior mean: this is why Jensen's inequality applies there without a correction and why the state route carries a further upward pull, the upper tail of a normal in the logit.

32.2.3 The reported computation on each plant

Each earlier chapter's reported route is kept. On the hall, expectations are the Gauss–Hermite sum (13.26) on ξ_C, with sixty nodes here, and probabilities and bands the delta method on $\ln U_n$, equations (13.33) and (13.36); that reproduces $\mathbb{E}[U_4] = 9.15 \times 10^{-3}$ and 0.790 for five units against 10^{-3}, Table 13.9. On the fleet, the predictive is the sum (11.13) on ω. On the plant, expected risks are the second-order expectation (25.18), by which Table 25.7 reports the reductions, and the same normal summed by quadrature stands beside it where it differs, since Section 25.4.1 found the two apart. Probabilities that a change pays use 2×10^5 draws from one generator (numpy `default_rng`, seed 32), taken in this order: a $2 \times 10^5 \times 3$ standard normal array for (ξ_C, ξ_P, ξ_K) scaled to $\mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$, then the Beta arrays of q_C, q_P, q_K; the exact values on the hall use the moments (13.29) and on the fleet the negative binomial (11.16).

32.3 The answers (step 4)

The questions of Table 32.2 are answered from $(\hat{\mathbf{z}}, \mathbf{\Sigma})$ of equation (32.10), by the routes of Section 32.2.3, and that is the value reported. The exact route gives the value beside it, in parentheses, with the gap.

32.3.1 The number of units

Engineering question. Redundancy: how many units does the hall need, with availability as a function of n priced against cost?

General form. The expected consequence of an action and its minimizer, equation (2.23):

$$\mathcal{R}(a) = \mathbb{E}\big[L(a, \mathbf{z}) \mid \mathbf{y}\big] = \int L(a, \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \qquad a^* = \arg\min_a \mathcal{R}(a). \tag{32.11}$$

Specialized form. The question is equation (32.11) with $a = n$, L the loss (32.4) and $p(\mathbf{z} \mid \mathbf{y})$ the normal of equation (32.10). One choice enters here: the unit is bought once and whole, so that the action set is the integers from three up, and the set is carried to ten units so that the upper end does not decide anything. The loss reads only ξ_C, and

$$\mathcal{R}(n) = (n-3)c_{\mathrm{u}} + C_{\mathrm{L}} t_{\mathrm{yr}} \sum_m \varpi_m\, U_n\big(q(\hat{\xi}_C + \sigma_\xi x_m)\big), \qquad \mathcal{R}(n) - \mathcal{R}(n+1) = \mathbb{E}[m_n(q_C) \mid \mathbf{y}] - c_{\mathrm{u}}, \tag{32.12}$$

with m_n of equation (32.5) and the nodes x_m and weights ϖ_m of the standard normal. On the exact route $\mathbb{E}[U_n]$ is a sum of the moments (13.29). The marginal condition: add the unit while the expected loss it removes exceeds c_{u}.

Building blocks. $\hat{\xi}_C$ and σ_ξ of equation (32.10); U_n and m_n; the five nodes of Table 32.3 and sixty for the reported values; the moments of Beta$(3, 97)$, for the value beside.

The rule of five nodes. Table 32.3 evaluates the sum for the marginal question that matters. Each column is a node, with its weight, q and the saving of the sixth unit. The central node, the plug-in, saves R$ 43 173; the two upper nodes supply R$ 119 402 of the weighted total of R$ 143 408, which sixty nodes make R$ 143 311. The exact value follows from $\mathbb{E}[U_6] = 15\mathbb{E}[q^4] - 24\mathbb{E}[q^5] + 10\mathbb{E}[q^6] = 4.558 \times 10^{-5}$ and $\mathbb{E}[U_5] = 5.329 \times 10^{-4}$, R$ 85 376.

Table 32.3: The sum (32.12) with five nodes for the saving of the sixth unit [R$/yr], $\xi = -3.4761 + 0.5862\,x_m$.

x_m	-2.857	-1.356	0	1.356	2.857
Weight ϖ_m	0.01126	0.22208	0.53333	0.22208	0.01126
q at the node	0.00576	0.01378	0.03000	0.06408	0.14169
$m_5(q)$	329	4 396	43 173	377 917	3 151 331
Weighted	4	976	23 026	83 926	35 476

The expected cost by n. Table 32.4 gives the answer. Each row is a number of units; the columns give the expected unavailability, the hours a year without cooling, the loss and the total with the unit cost, on the state route with the exact value in parentheses. Figure 32.2 draws the same: the horizontal axis is n, the vertical a logarithmic scale of reais per year; the bars are the bands of the loss on the state route, the dotted curve the loss at the plug-in, the rising line the unit cost, and the two heavy curves the total on each route.

Table 32.4: The hall by number of units: state route with the exact value in parentheses. Costs in R$/yr. Gaps are in the text.

n	$\mathbb{E}[U_n]$	Hours	Loss	Cost	Total
3	9.998×10^{-2} (8.649×10^{-2})	876 (758)	17 515 613 (15 153 729)	0	17 515 613 (15 153 729)
4	9.145×10^{-3} (6.673×10^{-3})	80.1 (58.5)	1 602 208 (1 169 103)	25 000	1 627 208 (1 194 103)
5	9.282×10^{-4} (5.329×10^{-4})	8.13 (4.67)	162 629 (93 363)	50 000	212 629 (143 363)
6	1.103×10^{-4} (4.558×10^{-5})	0.97 (0.40)	19 318 (7 986)	75 000	94 318 (82 986)
7	1.551×10^{-5} (4.209×10^{-6})	0.14 (0.04)	2 718 (737)	100 000	102 718 (100 737)
8	2.571×10^{-6} (4.191×10^{-7})	0.02 (0.004)	450 (73)	125 000	125 450 (125 073)

The action. On the state route

$$n^{\mathrm{B}} = 6, \qquad \mathcal{R}(6) = \text{R\$ } 94\,318,$$

and on the exact route $n^{\mathrm{B}} = 6$ at R$ 82 986, a gap of R$ 11 332 and none in the action. The next units either way cost more: five units by R$ 118 311 (R$ 60 376), seven by R$ 8 400 (R$ 17 751). Table 32.5 lists the marginal condition. Each row is one more unit; the columns give the expected loss it removes at the plug-in and on each route, and the loss rate at which it would just pay for itself at $c_{\mathrm{u}} = \text{R\$ } 25\,000$.

The price and the rate. Only the ratio of c_{u} to $C_{\mathrm{L}} t_{\mathrm{yr}}$ decides. By Table 32.5 six units stand for an annual unit cost from R$ 16 601 to R$ 143 311 on the state route and from R$ 7 249 to R$ 85 376 on the exact route, and for a loss rate from R$ 3 489 to R$ 30 119 h^{-1} and from R$ 5 856 to R$ 68 976 h^{-1}. Between those ends the routes give different integers. At the chapter's prices they agree.

The figure at three units. The R$ 15.15 million at three units is the exact value, R$ 17.52 million on the state route, and it is the arithmetic of equation (25.11) with no change of convention: 758 h a year (876 h) at the hourly rate. That rate was stated for short losses, so the

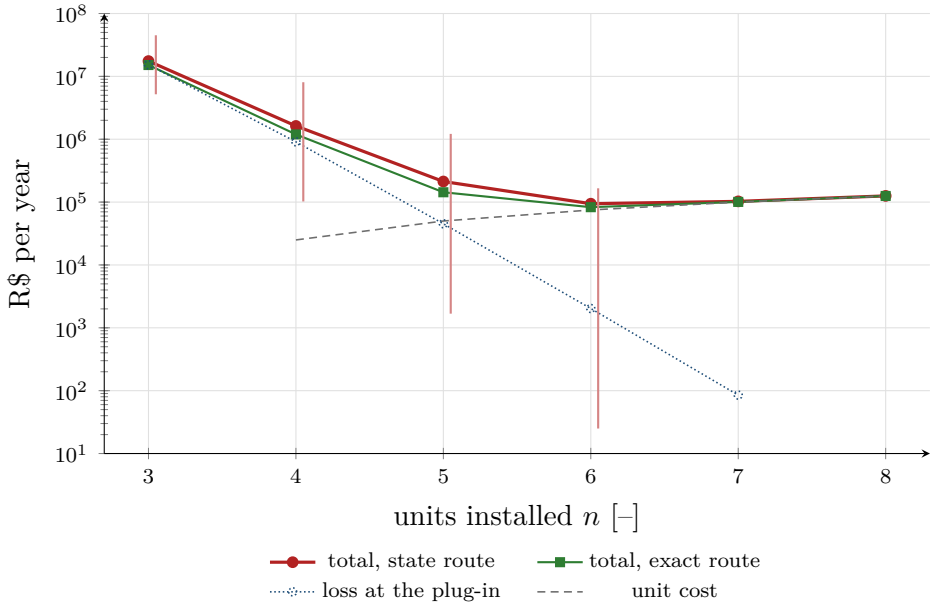


Figure 32.2: The hall's annual loss and total against the number of units, equation (32.4). Vertical bars: the 95 % band of the loss on the state route, delta method on $\ln U_n$, drawn to $n = 6$.

Table 32.5: The expected loss removed by one more unit on the hall [R$/yr], the unit's break-even annual cost at R$ 20 000 h^{-1}, and its break-even loss rate [R$/h] at R$ 25 000 a year.

	Loss removed			Break-even loss rate		
Unit added	Plug-in	State	Exact	Plug-in	State	Exact
3 to 4	14 391 028	15 913 406	13 984 626	35	31	36
4 to 5	863 462	1 439 578	1 075 740	579	347	465
5 to 6	43 173	143 311	85 376	11 581	3 489	5 856
6 to 7	1 943	16 601	7 249	257 362	30 119	68 976
7 to 8	82	2 267	664	6 127 667	220 550	753 103

figure says that three units are not a design, not what such a year would cost. The mean down time of Chapter 25 does not enter any expected loss here; it sets only how a year's hours divide into events.

32.3.2 The firmness of the number

Engineering question. How sure is that number?

General form. The posterior probability of an event on the state and the central interval of a derived quantity, equations (2.12) and (2.11):

$$\begin{gathered}\Pr(\mathbf{z} \in A \mid \mathbf{y}) = \int \mathbf{1}[\mathbf{z} \in A]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \\ \Pr(f(\mathbf{z}) \le f_L \mid \mathbf{y}) = 0.025, \quad \Pr(f(\mathbf{z}) \le f_U \mid \mathbf{y}) = 0.975.\end{gathered} \tag{32.13}$$

Specialized form. The question is equation (32.13) with A the set of states at which n is cheapest and f the loss. One choice enters here: the event compares actions at a state, where the state is treated as known, and so measures what would be best in hindsight and not what to do. Since $m_n(q)$ of equation (32.5) rises with q over the range that matters, A is an interval between the roots of $m_n(q) = c_{\mathrm{u}}$:

$$\Pr(n \text{ cheapest} \mid \mathbf{y}) = \Phi\left(\frac{\xi_n - \hat{\xi}_C}{\sigma_\xi}\right) - \Phi\left(\frac{\xi_{n-1} - \hat{\xi}_C}{\sigma_\xi}\right), \qquad \xi_n = \ln\frac{q_n}{1-q_n}, \quad m_n(q_n) = c_{\mathrm{u}}, \tag{32.14}$$

and on the exact route the same difference of the distribution function of Beta$(3, 97)$.

Building blocks. $\hat{\xi}_C$ and σ_ξ; the roots q_n; the delta band of equation (13.33); the Beta quantiles, for the values beside.

The numbers. The roots are $q_4 = 0.00491$, $q_5 = 0.02487$, $q_6 = 0.05809$ and $q_7 = 0.09851$: six units are best at the state if q_C lies between 2.5 % and 5.8 %. Then $\Phi((-2.7860 + 3.4761)/0.5862) - \Phi((-3.6688 + 3.4761)/0.5862) = 0.8804 - 0.3712 = 0.509$, and the exact value is 0.484, a gap of 0.025. Table 32.6 lists them with the medians and bands of the loss; each row is a number of units. Figure 32.3 draws the probabilities: the horizontal axis is n and each pair of bars the two routes. Six units are the single likeliest number, at about one chance in two; five is next, at 0.370 (0.435). If the action set stopped at six the sixth would collect every state above q_5, 0.629 (0.552), and that overstates it. The band of the loss is wider than the differences: at six units [R$ 25, R$ 165 704] ([R$ 4, R$ 59 667]), at five [R$ 1 680, R$ 1 216 094] ([R$ 433, R$ 563 917]), with medians of R$ 2 028 and R$ 45 201 (R$ 1 321 and R$ 32 810). The expected loss at six units, R$ 19 318 (R$ 7 986), lies far above its median, so that the decision rests on a mean carried by few states.

How far the logit's upper tail drives the answer. The normal in the logit puts 1.46 % of its mass above $q = 0.10$, where Beta$(3, 97)$ puts 0.21 %, and that mass supplies 43 % of $\mathbb{E}[U_6]$ on the state route against 9 % on the exact route; it is why $\mathbb{E}[U_6]$ stands 2.4 times its exact value where $\mathbb{E}[U_4]$ stands 1.37 times. Truncating the normal at $q = 0.10$ brings the totals to R$ 173 029 at five units and R$ 86 128 at six, against the exact R$ 143 363 and R$ 82 986. The tail therefore sets the level and the value of the sixth unit. It does not set the action: truncated anywhere from $q = 0.05$ up, the state route still takes six.

Table 32.6: The firmness of the number on the hall: the probability of being cheapest at the state, the median and the band of the loss [R$/yr], state route with exact in parentheses. The row for eight holds eight or more.

n	Cheapest	Median loss	Band of the loss
4	0.001 (0.013)	908 663 (734 667)	[102 344, 8 067 578] ([41 286, 4 820 619])
5	0.370 (0.435)	45 201 (32 810)	[1 680, 1 216 094] ([433, 563 917])
6	0.509 (0.484)	2 028 (1 321)	[25, 165 704] ([4, 59 667])
7	0.104 (0.066)	85 (50)	[0.3, 21 120] ([0.04, 5 907])
8	0.016 (0.002)	3 (2)	

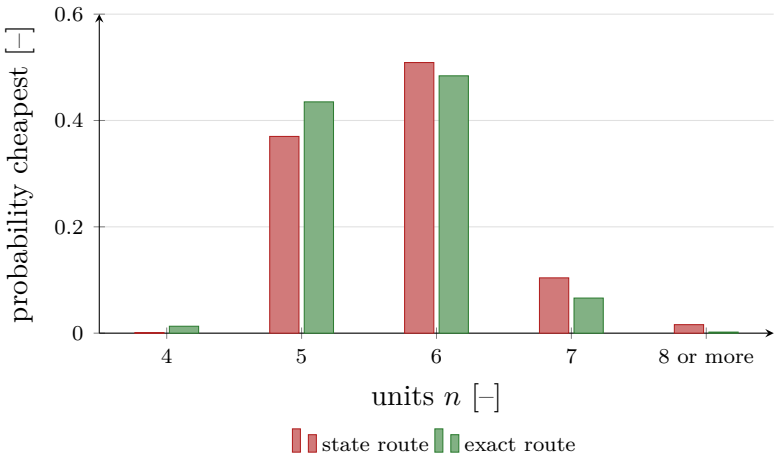


Figure 32.3: The posterior probability that each number of units is the cheapest at the state on the hall, equation (32.14), at $c_{\mathrm{u}} = \mathrm{R\$}\,25\,000$ and $C_{\mathrm{L}} = \mathrm{R\$}\,20\,000\,\mathrm{h}^{-1}$.

32.3.3 The number of spares

Engineering question. How many spares should the store hold?

General form. The expected consequence and its minimizer, equation (32.11), with the event in the future, so that the expectation runs over the predictive of equation (2.13).

Specialized form. The question is equation (32.11) with $a = s$, L the loss (32.8) and the normal of equation (32.10) on ω, the count entering through the predictive (2.13) with $y_{\text{new}} = K_f$. One choice enters here: the two prices are exact and linear, which makes the answer a quantile, and a price rising with the number short would break that. From the difference in equation (32.8), $\mathcal{R}(s+1) - \mathcal{R}(s) = c_{\text{s}} - C_{\text{s}} \Pr(K_f > s \mid \mathbf{y})$, and this rises with s, so that the best level is

$$s^{\text{B}} = \min\{s : \Pr(K_f \le s \mid \mathbf{y}) \ge \vartheta\}, \qquad \vartheta = 1 - \frac{c_{\text{s}}}{C_{\text{s}}} = \frac{C_{\text{s}} - c_{\text{s}}}{C_{\text{s}}}, \tag{32.15}$$

the spares level (11.25) with its service level set by the prices. The net cost of being one short is $C_{\text{s}} - c_{\text{s}}$, since a spare that would have been used also costs its holding; with that as the shortage price the ratio is the familiar shortage price over the sum of the two. The predictive is equation (11.13) on the state route and the negative binomial (11.16) on the exact route, and the plug-in the Poisson (11.18).

Building blocks. $\hat{\omega}$ and σ_ω of equation (32.10); the sum (11.13) on sixty nodes; the shortfall (11.26); the negative binomial, for the values beside.

The numbers. Here $\vartheta = 1 - 3\,000/45\,000 = 0.9333$. Table 32.7 lists the levels. Each row is a level; the columns give the service level and the shortfall on the state route with the exact value in parentheses, the same under the plug-in, and the expected cost under each. Figure 32.4 draws the costs against s. The service level first passes 0.9333 at fifteen on all but the plug-in:

$$s^{\text{B}} = 15, \qquad \mathcal{R}(15) = 45\,000 + 45\,000 \times 0.13022 = \text{R\$}\,50\,860,$$

and on the exact route $s^{\text{B}} = 15$ at R\$ 49 928, the value of Exercise 11.5, a gap of R\$ 932 and none in the action. The service levels of Chapter 11 did differ, sixteen against fifteen at 0.95: the prices put the line at 0.9333, below both routes' value at fifteen, 0.9489 and 0.9553, so the difference does not reach this decision. Holding sixteen costs R\$ 700 (R\$ 986) more. The rate known, the fifteenth spare pays less often than not: it beats fourteen at the true rate with probability 0.377 on the state route (0.352 on the exact route), because it saves little at the likely rates and much at the high ones.

Table 32.7: Spares levels for the coming year: the service level, the shortfall and the expected cost [R$], state route with exact in parentheses, and under the plug-in Poisson.

s	$\Pr(K_f \le s)$	Plug-in	Shortfall	Plug-in	Expected cost	Plug-in
13	0.8806 (0.8912)	0.9221	0.3289 (0.2892)	0.1681	53 802 (52 014)	46 563
14	0.9207 (0.9291)	0.9559	0.2095 (0.1804)	0.0902	51 427 (50 118)	46 057
15	0.9489 (0.9553)	0.9764	0.1302 (0.1095)	0.0461	50 860 (49 928)	47 073
16	0.9679 (0.9726)	0.9880	0.0791 (0.0648)	0.0224	51 560 (50 914)	49 010
17	0.9804 (0.9837)	0.9942	0.0471 (0.0374)	0.0104	53 118 (52 681)	51 470

32.3.4 The plant changes

Engineering question. Which plant change pays, and by how much?

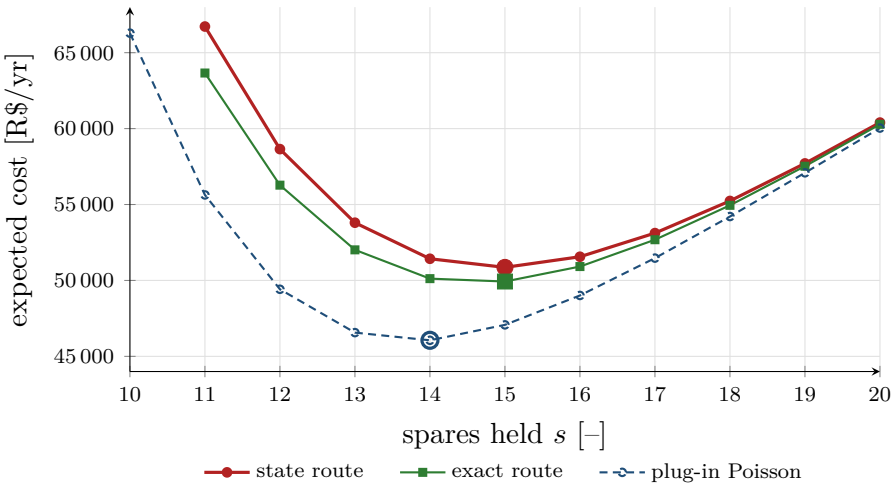


Figure 32.4: The cost of a spares level against s, equation (32.8), under the two predictives and the plug-in Poisson at $\hat{\lambda}\tau = 9.079$, with $c_{\mathrm{s}} = \text{R\$}\,3\,000$ and $C_{\mathrm{s}} = \text{R\$}\,45\,000$; the large marks are the minima.

General form. The comparison of expected consequences with a tail of the difference, equations (2.23) and (2.12):

$$\mathcal{R}(a') - \mathcal{R}(a), \qquad \Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}) = \int \mathbf{1}[f(\mathbf{z}) \geq f^*]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}. \tag{32.16}$$

Specialized form. The question is equation (32.16) with a and a' programs (n, δ_{k}), L the loss (32.7), $f = \Delta\mathcal{R}_{\mathrm{yr}}$ of equation (25.35), f^* the annualized cost, and the normal of equation (32.10) on the first three entries. One choice enters here: the sixth unit takes the schedule's rule of a cost per unit down, C_{C} for up to three of them, though a hall with six units and one down has lost no margin; this keeps Chapter 25's accounting and it charges the sixth unit about R\$ 41 000 a year, which Section 32.4.2 takes off. The ratio of a change is its expected reduction over its cost, and since changes interact it is taken on the residual: the reduction of a change made after others is $\mathcal{R}_{\mathrm{yr}}$ before it less $\mathcal{R}_{\mathrm{yr}}$ after it. The expected risk is equation (25.18), here

$$\mathcal{R}(a) \approx (n-4)c_{\mathrm{u}} + \delta_{\mathrm{k}} c_{\mathrm{k}} + \mathcal{R}_{\mathrm{yr}}(\hat{\mathbf{z}}; a) + \tfrac{1}{2}\operatorname{tr}(\mathbf{H}_a \boldsymbol{\Sigma}). \tag{32.17}$$

The interaction has a sign that can be read off. By equation (32.7) the top event is $1 - R_n Q K$, so a change in K removes $C_{\mathrm{L}} t_{\mathrm{yr}} R_n Q\, \Delta K$: the larger R_n is, the more it removes. Series terms overlap, and a better CRAC group leaves more hours in which the controller alone would take the cooling down. Each of the two changes is thus worth slightly more after the other.

Building blocks. $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$ of equation (32.10); the reductions of Table 25.7; the diagonal of $\mathbf{H}_a$ for the new programs; the draws of Section 32.2.3.

The sixth unit on one line. At $\hat{\mathbf{z}}$, equation (32.7) gives R\$ 264 830 for six units and two controllers, and with $\operatorname{diag}\mathbf{H}_a = (244\,198,\ 76\,403,\ 698)$ and the diagonal of $\boldsymbol{\Sigma}$, the trace term adds $41\,958 + 9\,647 + 349$, so that $\mathcal{R}_{\mathrm{yr}} = \text{R\$}\,316\,785$. Chapter 25's five units and two controllers are R\$ 369 813, so the sixth removes R\$ 53 028; the exact moments give R\$ 275 080 and R\$ 319 145, R\$ 44 065.

The changes. Table 32.8 gives the ratios. Each row is a change, with its cost; the columns give the expected reduction from Chapter 25 or its residual, the central band, the ratio and the probability that the reduction exceeds the cost at the state. Figure 32.5 draws them: the horizontal axis is the annualized cost, the vertical the reduction, both logarithmic; the solid line is break-even, and a point above it pays; each bar is the band on the state route, and each pair of points the two routes.

Table 32.8: The plant changes on Chapter 25's second-order route, with the exact value in parentheses [R$/yr]. The first three reductions and their bands are those of Table 25.7; bands of the residual rows and the probabilities from the draws of seed 32.

Change	Cost	Reduction	Band	Ratio	Pr(> cost)
Fifth unit	25 000	1 282 361 (1 019 408)	[92 061, 7 133 117] ([32 405, 4 105 556])	51.3 (40.8)	0.998 (0.981)
Second controller	40 000	259 912 (173 513)	[24 573, 1 231 329] ([4 353, 636 364])	6.50 (4.34)	0.928 (0.795)
Both	65 000	1 543 959 (1 193 940)	[159 314, 6 091 459] ([110 563, 4 298 730])	23.8 (18.4)	1.000 (0.992)
Controller, after the fifth	40 000	261 598 (174 532)	[24 317, 1 224 914] ([4 543, 645 556])	6.54 (4.36)	0.929 (0.796)
Sixth unit, after both	25 000	53 028 (44 065)	[−14 339, 790 281] ([−14 342, 403 521])	2.12 (1.76)	0.369 (0.300)
The three together	90 000	1 596 988 (1 238 005)		17.7 (13.8)	

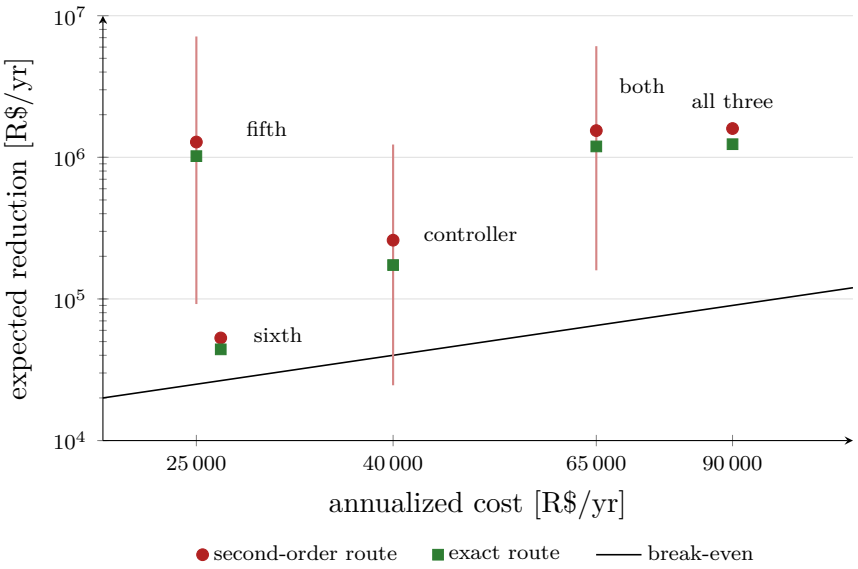


Figure 32.5: The expected reduction against the annualized cost for the plant changes, with the delta bands of Table 25.7. The sixth unit is drawn at its own cost of R$ 25 000, displaced to the right to separate it from the fifth; its band includes negative reductions and is not drawn on the logarithmic axis.

The interaction. Both changes together remove R$ 1 686 more than the two apart, 1 543 959 against 1 542 273 (R$ 1 018 more exactly, 1 193 940 against 1 192 921): the second change acts on the residual and is worth slightly more, not less, there. On the state route the controller goes

from R$ 259 912 to R$ 261 598 after the fifth unit. So close to the sum, the rule of doing every change whose ratio exceeds one, in the order of the ratio, until the budget binds, gives the same order whether it uses first or residual ratios: fifth unit, controller, sixth unit, on both routes. The controller is worth doing after the unit and not instead of it: alone it leaves R$ 1 693 861 (R$ 1 379 572) in the total against R$ 656 411 (R$ 518 677) for the fifth unit alone.

The program. Table 32.9 prices every program. Each row is a pair (n, δ_{k}); the columns give its annual cost, the total at the plug-in, the second-order total, the same normal summed by quadrature, and the exact total, the last three being the expected costs. The minimum on all three posterior routes is six units with two controllers, and at the plug-in it is five units with two. That program is R$ 406 785 (R$ 365 080), of which R$ 316 785 (R$ 275 080) is residual risk, and it stands R$ 28 028 (R$ 19 065) below five units and two controllers, and R$ 69 236 on the summed normal.

Table 32.9: The programs on the plant [R$/yr]: annual cost plus the annual risk at the plug-in, by the second-order route of Chapter 25, by the same normal summed, and on the exact route. The least of each column is in bold.

Program (n, δ_{k})	Cost	Plug-in	Second-order	Summed	Exact
$(4, 0)$, present	0	1 252 736	1 913 773	2 084 301	1 513 085
$(4, 1)$	40 000	1 118 789	1 693 861	1 840 509	1 379 572
$(5, 0)$	25 000	467 377	656 411	739 058	518 677
$(5, 1)$	65 000	**332 620**	434 813	493 020	384 145
$(6, 0)$	50 000	489 591	628 437	669 976	499 656
$(6, 1)$	90 000	354 830	**406 785**	**423 784**	**365 080**
$(7, 1)$	115 000	417 361	470 903	478 587	422 461

Where the routes part. The fifth unit pays with assurance on both routes, 0.998 (0.981). The controller's expected ratio is above one on both, and its probability of paying falls on opposite sides of an assurance of 0.90, 0.928 against 0.795. It pays when q_K exceeds 2.3×10^{-4}: the logit normal around one finding in a thousand gives that 0.93, while Beta$(1, 999)$, with its mode at zero, gives 0.80. Both are stated and neither chosen. The sixth unit is cheaper in expectation on both routes and pays at the state in less than two cases of five on both: it is a purchase against the upper tail of q_C. Its median saving after both changes is R$ 2 553 on the state route and R$ $-$ 4 238 on the exact route, against R$ 53 028 (R$ 44 065) in expectation. The scene's figure for the sixth unit, R$ 2 789, is not far from those medians. The route of the sixth unit's expected saving matters too: R$ 53 028 second-order, R$ 94 235 on the summed normal, R$ 44 065 exact.

The hall and the plant. The same units are worth different sums on the two. On the exact route the fifth removes R$ 1 075 740 on the hall and R$ 1 019 408 on the plant, and the sixth R$ 85 376 and R$ 44 065; the difference is the term $C_{\mathrm{C}} M_n$, which rises by R$ 55 070 and about R$ 41 300. The state-route columns also differ by route, the sum on the hall against the second-order expectation on the plant, and they are not to be compared.

The prices. On the plant the sixth unit stays up to an annual cost of R$ 53 028 second-order and R$ 44 065 exact, and the controller at six units up to R$ 261 652 and R$ 174 576; between them the routes give different programs.

32.3.5 The plug-in action against the Bayes action

Engineering question. Why does the plug-in action err, and in which direction?

General form. The expected cost of one action above the minimum, equation (2.23), and the difference (32.2) between the two stopping conditions:

$$\mathcal{R}(a') - \min_a \mathcal{R}(a), \qquad \mathbb{E}[m_a(\mathbf{z}) \mid \mathbf{y}] - m_a(\hat{\mathbf{z}}). \tag{32.18}$$

Specialized form. The question is equation (32.18) with a' the plug-in action $\hat{n}$, $(\hat{n}, \hat{\delta}_\mathrm{k})$ or $\hat{s}$, the marginal saving m_a that of equation (32.5) or (32.8), and the posterior of equation (32.10). No choice enters. Two separate statements are in play. The first concerns the plug-in's *number*: when $L(a, \cdot)$ is convex in q and the plug-in sits at the mean, $\mathbb{E}[L] \geq L(a, \hat{q})$ by Jensen's inequality, and the plug-in's cost is too low. The second concerns its *action*: by equation (32.2) it is set by the curvature of m_a. For units, m_n of equation (32.5) is, over the states that carry the posterior, concave in q at $n = 3$ on the exact route and convex at $n \geq 4$, and convex in ξ by equation (32.6) at every n on the state route. For spares, $C_\mathrm{s} \Pr(K_f > s)$ turns over at $\mu = s$ on the exact route and $\mu = s + 1$ on the state route by equation (32.9), so that with the fractile above the center of the count the marginal is convex there and the Bayes action is at least the plug-in's, and with the fractile below it the reverse. With a normal approximation this is $s \approx \mu + z_\vartheta \,\mathrm{sd}$, with z_ϑ the normal quantile: a wider predictive moves the quantile away from the mean, up if ϑ lies above one half and down if below.

Building blocks. The plug-in actions from $\hat{\mathbf{z}}$; $\mathcal{R}$ of Sections 32.3.1, 32.3.3 and 32.3.4; the second derivatives (32.6) and (32.9).

The numbers. Table 32.10 sets the three decisions side by side and Figure 32.6 draws them: each horizontal line is one decision's scale, with the plug-in action and the Bayes action of each route marked. Each row of the table is a decision; the columns give the plug-in action and what it claims, what it costs in expectation, the Bayes action, its cost and the regret.

Table 32.10: The plug-in action against the Bayes action [R$/yr], state route with exact in parentheses. The plant's state column is second-order.

Decision	Plug-in	Claimed	Its expected cost	Bayes	Bayes cost	Regret
Hall units	6	77 028	94 318 (82 986)	6 (6)	94 318 (82 986)	0 (0)
Plant program	(5, 1)	332 620	434 813 (384 145)	(6, 1)	406 785 (365 080)	28 028 (19 065)
Spares	14	46 057	51 427 (50 118)	15 (15)	50 860 (49 928)	567 (190)

One step short, or not. On the hall the plug-in takes six units as well: it saves R$ 43 173 for R$ 25 000 and there is no regret. On the plant it stops at five, because the sixth unit's saving there is about R$ 2 789, the same R$ 43 173 less nearly the same rise in the degraded term. So "one unit short" is a statement about the plant under the schedule's rule, not about redundancy. On the hall the plug-in is short of six only where c_u lies between R$ 43 173 and R$ 85 376 exactly (R$ 143 311 on the state route), or the loss rate between the routes' break-even and R$ 11 581 h^{-1}, Table 32.5, and the same holds a step up for the seventh unit. Each is one step, and at these prices its regret is small beside the number that the plug-in misstates: on the plant R$ 28 028 (R$ 19 065) against R$ 102 193 (R$ 51 525) of understatement in its own program, and on the spares R$ 567 (R$ 190) against R$ 5 370 (R$ 4 061).

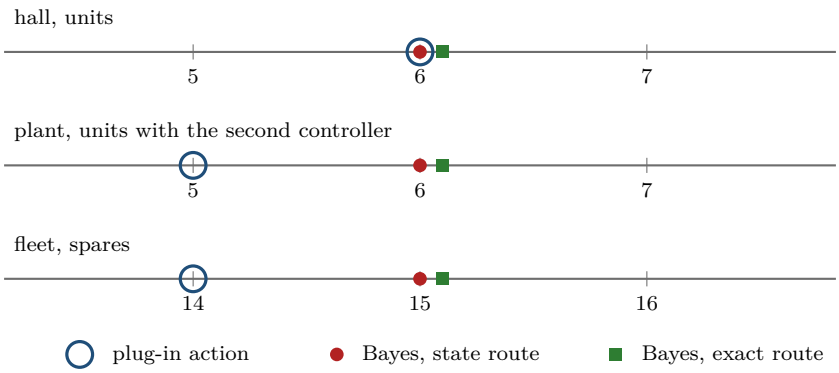


Figure 32.6: The plug-in and Bayes actions for the three decisions. The Bayes marks of a decision belong to the tick to their left, and the square is displaced so that both show. On the hall all three coincide; on the plant and the fleet the plug-in is one step short on both routes.

The direction. Figure 32.7 shows what sets it. The left panel draws the saving of the sixth unit against q, on logarithmic axes, with the cost line, the plug-in at $q = 0.03$, and the two expectations as horizontal lines; the curve is convex, so the expectations stand above the plug-in. The right panel draws the marginal saving of a spare against the mean count, for $s = 14$ and $s = 6$: each curve is S-shaped, with its inflection at $\mu = s$, and the vertical line marks the fitted $\hat{\mu} = 9.08$, in the convex part of the first and in the concave part of the second.

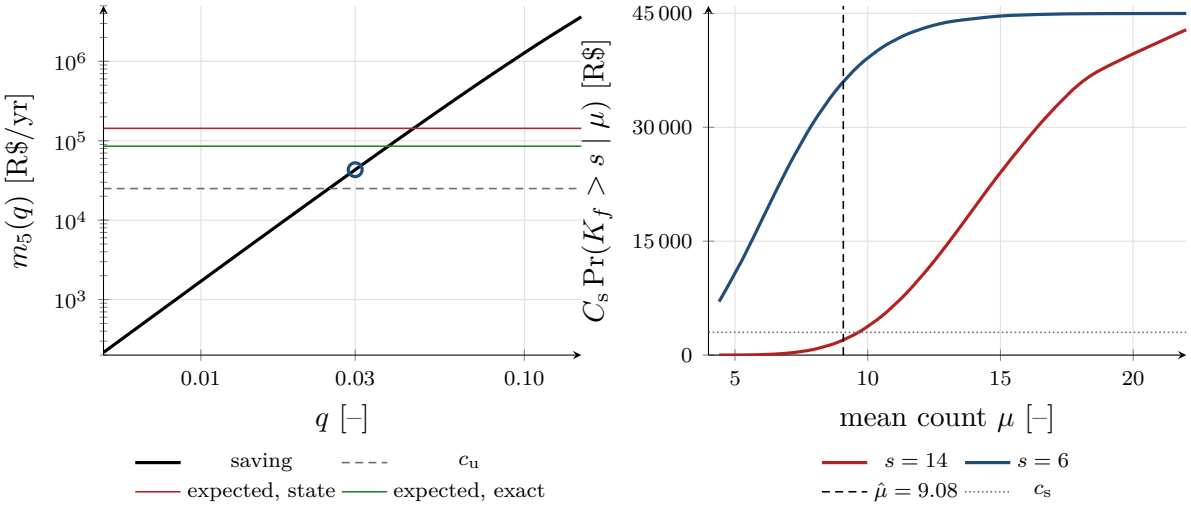


Figure 32.7: What sets the direction. Left: equation (32.5) at $n = 5$, with $C_L = \text{R\$}\,20\,000\,\text{h}^{-1}$; the circle is the plug-in. Right: the marginal saving of a spare, equation (32.8), against the Poisson mean, with $C_s = \text{R\$}\,45\,000$.

The term itself. Equation (32.2) has the right sign in each case and truncates badly in the unit tail. For the sixth unit, $m_5 = 43\,173$ and equation (32.6) with $a_n = 3$ give $\frac{1}{2}\,43\,173 \times 7.778 \times 0.34364 = 57\,696$, in all R\$ 100 869 against the summed R\$ 143 311; in q on the exact route, $43\,173 + 37\,654 = \text{R\$}\,80\,827$ against R\$ 85 376. For the seventh, $1\,943 + 4\,727 = \text{R\$}\,6\,670$ against R\$ 16 601. For the fifteenth spare, $s = 14$, the term is good: $1\,984 + 1\,574 = \text{R\$}\,3\,558$ against the summed R\$ 3 567, and in λ on the exact route $1\,984 + 1\,308 = \text{R\$}\,3\,292$ against R\$ 3 191.

Too few, on what condition. Table 32.11 tests it on the fleet by moving the prices. Each row is a critical fractile; the columns give the ratio $c_\mathrm{s}/C_\mathrm{s}$, the plug-in level, the Bayes level on each route and the regret of the plug-in at $c_\mathrm{s} = \mathrm{R\$\,3\,000}$. The plug-in is short only for a high fractile, and it is above the Bayes level for a low one. Over the whole range it exceeds it only for ϑ below 0.448 on the state route and 0.579 on the exact route, and falls short of it only for ϑ above 0.561 and 0.682. The break comes where the Poisson and predictive distribution functions cross, between $s = 8$ and 9 on the state route and $s = 9$ and 10 on the exact route, as equation (32.9) says, at $\mu < s + 1$ in ω on the state route and $\mu < s$ in λ on the exact route. In every row the loss is convex and the plug-in's number too low. At $\vartheta = 0.20$ that low number comes with too many spares, seven against six.

Table 32.11: The spares level by critical fractile. Levels by equation (32.15); regret of the plug-in level at $c_\mathrm{s} = \mathrm{R\$\,3\,000}$, $C_\mathrm{s} = 3\,000/(1-\vartheta)$ [R$/yr].

ϑ	$c_\mathrm{s}/C_\mathrm{s}$	Plug-in	State	Exact	Regret, state	Regret, exact
0.20	0.80	7	6	6	102	155
0.45	0.55	9	9	8	0	87
0.50	0.50	9	9	9	0	0
0.80	0.20	12	12	12	0	0
0.85	0.15	12	13	13	488	227
0.90	0.10	13	14	14	583	264
0.9333	0.0667	14	15	15	567	190
0.95	0.05	14	16	15	1 822	1 254
0.99	0.01	17	19	18	3 404	1 893
0.999	0.001	20	23	22	13 209	7 620

The same on the units. At $n = 3$ the saving $3q(1-q)^3$ is concave in q below $q = 0.5$, so that on the exact route the fourth unit removes less in expectation than at the plug-in, R$ 13 984 626 against R$ 14 391 028. For a price between the two the plug-in buys a unit that the exact route does not: at a notional R$ 14 million a unit, four against three, with a regret of R$ 15 374. In the logit that saving is convex below $q = 0.10$, which holds 98.5 % of the normal, and the state route's expectation, R$ 15 913 406, is above the plug-in: the routes then disagree in sign. The same reversal holds in the plug-in's number: at three units the loss is concave in q near the record, and the exact expectation is below the plug-in by R$ 145 962, the case of the series noted in Section 13.3.1, while the state route is above it by R$ 2 215 923. From $n = 4$ on, over the states that carry the posterior, both the loss and the marginal saving are convex on both routes, the plug-in's number is low and its action is at most the Bayes action.

Under other prices. Table 32.12 carries the decisions, with the regret just defined. Each row is a set of prices; the columns give the plug-in action, the Bayes action on each route and, for the hall, the regret of the plug-in action on each.

The benefit. The plug-in's reduction for the fifth unit, R$ 810 359, is below R$ 1 282 361 and R$ 1 019 408, by R$ 472 002 and R$ 209 049. For the controller it is below on the state route, R$ 173 947 against R$ 259 912, and above on the exact route by R$ 434, since $q_K - q_K^2$ is concave in q_K and the factor of the other groups is lower in expectation than at the plug-in: there the plug-in overstates, negligibly.

32.3.6 The target and the cost minimum

Engineering question. Does an availability target or a cost minimum decide?

Table 32.12: The actions under other prices, the state route (second-order on the plant) and the exact route, with the regret of the plug-in action on each [R$/yr] for the hall. A dagger marks a row where the routes differ.

Prices	Plug-in	State	Exact	Regret, state	Regret, exact
Hall, number of units					
R$ 25 000, R$ 20 000 h^{-1}	6	6	6	0	0
R$ 10 000, R$ 20 000 h^{-1} †	6	7	6	6 601	0
R$ 50 000, R$ 20 000 h^{-1}	5	6	6	93 311	35 376
R$ 100 000, R$ 20 000 h^{-1} †	5	6	5	43 311	0
R$ 200 000, R$ 20 000 h^{-1}	5	5	5	0	0
R$ 25 000, R$ 5 000 h^{-1} †	5	6	5	10 828	0
R$ 25 000, R$ 10 000 h^{-1}	5	6	6	46 655	17 688
R$ 25 000, R$ 40 000 h^{-1} †	6	7	6	8 202	0
R$ 25 000, R$ 100 000 h^{-1}	6	7	7	58 004	11 245
R$ 14 000 000, R$ 20 000 h^{-1} †	4	4	3	0	15 374
Plant, program					
As in the chapter	(5, 1)	(6, 1)	(6, 1)		
Unit R$ 50 000 †	(5, 1)	(6, 1)	(5, 1)		
Unit R$ 100 000	(5, 1)	(5, 1)	(5, 1)		
Controller R$ 100 000	(5, 1)	(6, 1)	(6, 1)		
Controller R$ 200 000 †	(5, 0)	(6, 1)	(6, 0)		
Controller R$ 300 000	(5, 0)	(6, 0)	(6, 0)		
Loss rate R$ 10 000 h^{-1}	(5, 1)	(5, 1)	(5, 1)		

General form. A tail of a derived quantity against an assurance, equation (2.12):

$$\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}) = \int \mathbf{1}[f(\mathbf{z}) \geq f^*]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \qquad \text{accept if } \Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}) \geq \kappa. \quad (32.19)$$

Specialized form. The question is equation (32.19) with $f = -U_n$, $f^* = -u^*$, $u^* = 10^{-4}$, $\kappa = 0.90$ and the normal of equation (32.10) on ξ_C. One choice enters here: the target is read on the hall's cooling alone, as Chapter 13 reads its own, and it would mean something else on the plant, as below. By equation (13.36),

$$\Pr(U_n < u^* \mid \mathbf{y}) \approx \Phi\left(\frac{\ln u^* - \ln U_n(\hat{q}_C)}{\sigma_{\ln U,n}}\right), \quad (32.20)$$

with $\sigma_{\ln U,n} = 1.114$, 1.680, 2.247 and 2.814 at $n = 4$ to 7; on the exact route $I_{q^*}(3, 97)$ of equation (13.49). The target rule takes the smallest n that passes. It holds no loss rate and no price, and the cost rule holds no u^* and no κ. This chapter derives the cost rule; the target is the form that standards state, as the availability chapters of the companion volume *Reliability Engineering* treat it.

Building blocks. $\hat{\xi}_C$, σ_ξ and the spreads above; U_n at $\hat{q}_C$; the roots q^* and the incomplete Beta function for the values beside.

The numbers. At six units, $\Phi(\ln(10^{-4}/1.157 \times 10^{-5})/2.247) = \Phi(0.960) = 0.831$, and with $q^* = 0.05190$, $I_{q^*}(3, 97) = 0.893$. At seven, 0.971 (0.994). So the rule asks for seven units on both routes, and the cost minimum for six: the two differ by one unit, at R$ 8 400 (R$ 17 751) a year. With 0.893 the exact route misses 0.90 by 0.007, and on neither route does six units pass. With six units the assurance 0.90 holds at $u^* = 2.1 \times 10^{-4}$ (1.1×10^{-4}). The plug-in, $U_6 = 1.16 \times 10^{-5}$, reads six units as passing by a factor of nine. Table 32.13 lists the rule: each row is a number of units, with the plug-in, the spread and the assurance on each route.

Table 32.13: The target $U_n < 10^{-4}$ on the hall: the unavailability at the plug-in, the spread of equation (32.20) and the assurance, state route and exact route.

n	$U_n(\hat{q}_C)$	$\sigma_{\ln U,n}$	State	Exact
4	5.186×10^{-3}	1.114	0.0002	0.008
5	2.580×10^{-4}	1.680	0.286	0.366
6	1.157×10^{-5}	2.247	0.831	0.893
7	4.851×10^{-7}	2.814	0.971	0.994
8	1.938×10^{-8}	3.382	0.994	0.9999

When they part. The target does not move with the prices. The cost minimum is seven only above R$ 30 119 h^{-1} (R$ 68 976 h^{-1}) and five below R$ 3 489 h^{-1} (R$ 5 856 h^{-1}), so a cheap loss rate or a strict assurance sets the rules further apart, and a dear loss rate or a cheap unit makes the cost rule the stricter. Read on the plant, the same target gives another answer: two pumps at $q_P = 0.010$ contribute about 10^{-4} on their own, and even with seven units and two controllers $\Pr(U_{\text{sys}} < 10^{-4})$ is 0.44 (0.54), so no number of CRAC units meets it. A target belongs to a stated boundary of the system.

32.4 The checks (step 5)

The answers rest on a quadrature in a heavy tail, on two earlier chapters' routes and on one prior for each record. Each is set against a fuller calculation or varied. There is no physics row, so no limit $\sigma_p \to 0$ arises.

32.4.1 The routes and the quadrature

Table 32.14 sets each computed quantity beside an independent calculation. Each row is one check; the columns give the quantity, its value, the other calculation and the verdict. Since a tail quantity can look settled before it is, the rule was tried at 5 to 150 nodes against a dense sum, and for every n up to ten it has settled by twenty.

The second-order route at more units. On the plant the second-order expectation stands below the summed normal, by 8 % at the present design, 14 % at $(5, 1)$ and 5 % at $(6, 1)$, and it moves the sixth unit's saving from R$ 94 235 to R$ 53 028. Its truncation is thus not small where the action turns, and here it happens to lie near the exact value. That is no property of the route; the action is the same on all three and the size of the margin is not.

32.4.2 The priors and the drawing

Table 32.15 repeats the decisions under the priors of Tables 13.12 and 11.13, refitting each posterior as a check. Each row is a prior; the columns give the action and its expected cost on the state route with the exact value in parentheses, and for the hall the probability of six and the target at six.

Under the uniform prior the routes give seven and six, seven costing R$ 104 274 on the state route and six R$ 92 000 on the exact one; under the flat prior on the rate they give sixteen and fifteen. The plug-in stays at fourteen spares under all five. On the plant the summed normal and the exact route stay at six units and two controllers under the Jeffreys and uniform priors of Table 25.9, at R$ 471 076 (R$ 411 979) and R$ 519 778 (R$ 459 821), and the sixth unit's margin grows, since those priors raise q_C.

Table 32.14: The checks of the routes and of the quadrature.

Quantity	Route	Independent calculation	Verdict
U_4, U_5, U_6 at 0.03	binomial tail: 5.18643×10^{-3}, 2.57996×10^{-4}, 1.15741×10^{-5}	enumeration of the 2^n configurations: the same	agree
$\mathbb{E}[U_6]$, state	sixty nodes: 1.10264×10^{-4}	five: 1.10392; ten: 1.10265; twenty: 1.10264; dense sum: 1.10264; draws: 1.107 (s.e. 0.014)	agree from ten
$\mathbb{E}[U_8]$, state	sixty nodes: 2.5713×10^{-6}	five: 2.3644; ten: 2.5711; dense sum: 2.5713; draws: 2.63 (s.e. 0.12)	five nodes 8 % low
$\mathbb{E}[U_6]$, exact	moments: 4.5584×10^{-5}	grid: 4.5584; draws: 4.573 (s.e. 0.032)	agree
Pr(6 cheapest)	normal tails: 0.509 (0.484)	draws: 0.510 (0.483)	agree
$\Pr(U_6 < 10^{-4})$	delta: 0.831	normal in ξ through the root: 0.835	0.004
$\mathcal{R}(15)$, state	sum: 50 860	2×10^6 draws: 50 837 (s.e. 22)	agree
$\mathcal{R}(15)$, exact	negative binomial: 49 928	draws: 49 922 (s.e. 20)	agree
$\Pr(K_f \leq 23)$, state	sixty nodes: 0.999325	five: 0.999319; dense sum: 0.999325	agree
Plant, present	equation (32.7): 1 913 773, 1 513 085; summed 2 084 301	Chapter 25: the same three	agree
Plant, $(5, 1)$	equation (32.7): 369 813, 319 145	Table 25.7: the same	agree

Table 32.15: The decisions under other priors, state route with exact in parentheses. Hall probabilities from the normal through the root.

Prior	Action		Cost at six units or s^{B}	Pr(6 cheapest)	$\Pr(U_6 < 10^{-4})$
Hall, q_C					
Vague limit	6	(6)	94 318 (82 986)	0.509 (0.484)	0.835 (0.893)
Jeffreys	6	(6)	99 567 (86 928)	0.575 (0.556)	0.781 (0.839)
Uniform	7	(6)	106 159 (92 000)	0.612 (0.601)	0.718 (0.775)
Unit prior	6	(6)	92 147 (84 117)	0.597 (0.576)	0.818 (0.871)
Fleet, λ					
Handbook, Gamma$(2, 10)$	15	(15)	50 860 (49 928)		
Jeffreys	15	(15)	52 196 (51 031)		
Flat	16	(15)	53 079 (51 843)		
Gamma$(1, 5)$	15	(15)	51 103 (50 102)		
Gamma$(4, 20)$	15	(15)	50 439 (49 624)		

The drawing. Chapter 13 left open whether the hall needs three units or two, at the weights 0.7 and 0.3. With two needed, the cost minimum is five units on both routes, at R$ 57 019 (R$ 52 812). Under the mixture it stays at six, at R$ 88 784 (R$ 80 658). All numbers above assume three.

The degraded term. On the plant, without the term $C_{\mathrm{C}} M_n$, the totals are R$ 209 315 at $(5, 1)$, R$ 133 463 at $(6, 1)$ and R$ 151 795 at $(7, 1)$ (R$ 189 192, R$ 128 831 and R$ 146 583). The Bayes action does not move, and removing the term widens the sixth unit's margin, from R$ 28 028 (R$ 19 065) to R$ 75 852 (R$ 60 361). The plug-in's stop at five does move: with the term taken off the sixth unit saves about R$ 43 000 at the plug-in and it too takes six, as on the hall.

32.5 The report (step 6)

The fitted values promised a plan of five units and a second controller at R$ 332 620 a year, and fourteen spares at R$ 46 057. That plan costs R$ 434 813 (R$ 384 145) and R$ 51 427 (R$ 50 118) in expectation, and the cheapest plan on either posterior has one more unit and one more spare. The direction is a fact of these prices and of the schedule's charge at six units, not of convex losses: on the hall alone the plug-in already takes the right number of units, and on the fleet a low critical fractile would have it hold too many spares.

32.5.1 The answers and their bands

Table 32.16 is the report. Each row is a question in the manager's words; the second column gives the answer from the posterior of the state, the third the exact-route value, the fourth the gap, the last the band or the qualification.

The report line. On the plant with the records in hand, a fifth CRAC unit removes about R$ 1.3 million (R$ 1.0 million) a year for R$ 25 000 and pays with assurance on both routes. A second controller removes R$ 262 000 (R$ 175 000) after it for R$ 40 000; it is worth buying in expectation on both, and it pays with probability 0.93 on the posterior of the state and 0.80 on the exact one, either side of an assurance of 0.90. A sixth unit removes a further R$ 53 000 (R$ 44 000) for R$ 25 000, in expectation only, paying at the true state with probability 0.37 (0.30). With all three the annual risk is R$ 317 000 (R$ 275 000) and the total with prices R$ 407 000 (R$ 365 000), the best action if the decision must be made on the present record; Section 32.6.3 prices taking a hundred checks first. A store of fifteen spares, one more than the fitted count suggests, costs about R$ 51 000 (R$ 50 000) a year. On the hall's cooling alone a target of 10^{-4} at 0.90 would ask for a seventh unit that the costs do not justify. All prices are illustrative.

32.5.2 The investment decision

If the decision must be made on the present record, the best action is six units, the second controller and fifteen spares, at R$ 406 785 plus R$ 50 860 (R$ 365 080 plus R$ 49 928), the minimum on both routes. If a record can be taken first, the plan to compare with it is priced in Section 32.6.3. The steps are not equally firm. The fifth unit does not depend on anything this chapter leaves open. The controller's assurance depends on the route, and its price would have to pass R$ 175 000 for the routes to stop agreeing on it. The sixth unit depends on the price, stopping at R$ 44 065 exactly, and on the schedule. It is a purchase against the upper tail and the first to drop under a budget or under an owner who wants each step to pay with assurance;

Table 32.16: The report of the three investments: the answers from $\mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$, each on its chapter's route, the exact route with its gap, and the bands, in R$/yr.

Engineering question	Answer (state route)	Exact route	Gap	Band, qualification
How many units? (hall)	six, R$ 94 318; loss R$ 19 318, $\mathbb{E}[U_6] = 1.10 \times 10^{-4}$	six, R$ 82 986; R$ 7 986, 4.56×10^{-5}	none in the action; 11 332	six for c_{u} from 16 601 to 143 311 (7 249 to 85 376); uniform prior seven (6)
How sure is that number?	six cheapest at the state 0.509; five 0.370; seven 0.104	0.484; 0.435; 0.066	0.025; 0.065; 0.038	loss at six [25, 165 704] ([4, 59 667]), median 2 028 (1 321)
How many spares?	fifteen, R$ 50 860	fifteen, R$ 49 928	none in the action; 932	fractile 0.9333; service level 0.9489 (0.9553); flat prior sixteen (15)
Which plant change pays? (plant)	fifth unit ratio 51.3; controller 6.50, after the fifth 6.54; both 23.8; sixth after both 2.12; six units and the controller R$ 406 785	40.8; 4.34, 4.36; 18.4; 1.76; R$ 365 080	10.5; 2.16; 5.4; 0.36; 41 705	pays with probability: fifth 0.998 (0.981); controller 0.928 (0.795), opposite sides of 0.90; sixth 0.369 (0.300); interaction +1 686 (+1 018)
Why does the plug-in err?	plant: five against six, regret R$ 28 028; spares: fourteen against fifteen, R$ 567; hall: none	R$ 19 065; R$ 190; none	8 963; 377	set by the curvature of the marginal saving: too few spares only above a fractile near 0.56 (0.68); plug-in's cost understated on every row but three units exact
Target or cost minimum?	target 10^{-4} at 0.90 asks seven (0.971; six 0.831); the cost minimum six	seven (0.994; six 0.893); six	0.023; 0.062	one unit, R$ 8 400 (R$ 17 751) a year; on the plant unreachable by units

such a rule stops at five units on both routes, and for the controller the routes then disagree. Its cost is stated so that the choice can be made: R$ 28 028 (R$ 19 065) a year in expectation.

> **Principle 32.1.** *An integer action stops where the expected marginal saving falls below the price of one more step. Convexity of the loss in the state makes a plug-in's cost too low; the plug-in's action errs in the direction of the curvature of the marginal saving, taken in the coordinate of the route, which can have either sign. Changes on one plant are ranked on the residual, and a report prices the program it recommends, a plan with a measurement in it included.*

Common mistake

Five errors recur with such investments. Saying that a convex loss makes the plug-in action too small: at a critical fractile of 0.20 the plug-in holds seven spares where both routes hold six, with the loss convex throughout. Saying that the plug-in is always a unit short: on the hall it takes the same six units as both routes. Adding two reductions as if the second were diminished: here both together remove R$ 1 686 (R$ 1 018) more than the sum. Moving a unit's value from one plant to another: the sixth unit removes R$ 85 376 on the hall and R$ 44 065 on the plant, exactly. And reading a probability of being cheapest as the decision: the sixth unit on the plant is cheaper in expectation and pays in under two states of five.

32.6 The next measurement (step 7)

The decisions turn on the upper tail of q_C and, for the controller, on the lower end of q_K, so the measurements on offer are more checks and a reading of the load. The value of each is priced in the form of Chapter 30.

32.6.1 The value of knowing the state

Equation (30.18) applies with the state revealed and the actions as stated. Since a value of information belongs to an action set, Table 32.17 takes three of them on the hall and on the plant. Each row is what is revealed and among which actions; the columns give the value on the state route, by the sum over the normal, and on the exact route. The second-order route is not used here: it gives an expectation, and this needs each state's minimum.

On the hall the order of the routes reverses with the action set. With at most six units the exact route's value is the larger; with more on the table the state route's is, since its upper tail would buy seventh and eighth units in states the exact posterior barely holds: 44 % of the state route's value comes from q above 0.10, against 6 % on the exact route. Both are stated. The plant shows the same dependence on the action set from another side: while the fifth unit is the only question q_C is worth almost nothing, since no likely state reverses it, and once a sixth is on the table it is worth over R$ 22 000 on both routes, and q_K is worth five times as much on the exact route as on the state route. On either route and either plant the value is small against what the program removes.

32.6.2 The candidate records and their preposterior

A hundred more checks of the CRAC fleet. This adds a row on ξ_C with $\mathbf{\Lambda}_a = 2.91$, and by equation (2.17) $\Sigma_a = 0.17182$, as in Section 13.6.2. On the state route the outcomes are those of the preposterior mean, $\mathcal{N}(\hat{\xi}_C,\ \Sigma_{11} - \Sigma_a)$; the exact route enumerates the counts, under the Beta–binomial. The probability that six is cheapest averages to its present value, 0.509 (0.484), by the law of total expectation: a planned record spreads it and cannot raise it on average. Table 32.18 shows the spread. Each row is a count of findings; the columns give

Table 32.17: The expected value of perfect information [R$/yr], on the state route and the exact route.

Revealed, action set	State	Exact
Hall		
q_C, four or five units	5	139
q_C, four to six units	5 325	7 269
q_C, three to ten units	16 728	9 815
Plant		
q_C, fifth unit or not	14	223
q_K, fifth unit or not	0	0
q_C, four programs, up to five units	14	223
q_K, four programs, up to five units	885	4 242
The three entries, four programs	899	4 465
q_C, all programs	23 055	22 210
q_K, all programs	884	4 241
The three entries, all programs	23 939	26 450
Fleet		
λ, any level of spares	4 165	3 932

its predictive probability, the hall's action after it and the probability that six is cheapest, on the refit state route with exact in parentheses. Averaged over the counts, the refits give 0.528 against the present 0.509, the drift of successive Laplace fits, not information.

Table 32.18: The hall after a hundred more checks, by the count found: refit of the normal on the state route, with Beta$(3 + k, 197 - k)$ in parentheses.

Findings k	Predictive probability	Action	Pr(6 cheapest)
0	0.088 (0.121)	5 (5)	0.180 (0.125)
1	0.162 (0.186)	6 (5)	0.315 (0.266)
2	0.181 (0.188)	6 (6)	0.478 (0.438)
3	0.161 (0.159)	6 (6)	0.631 (0.601)
5	0.092 (0.084)	6 (6)	0.774 (0.771)
7	0.043 (0.035)	7 (6)	0.675 (0.696)
8	0.029 (0.022)	7 (7)	0.570 (0.598)

So six stays the action with probability 0.682 (0.641) after the hundred checks on the preposterior route, five becomes it with 0.150 (0.307) and seven or more with 0.168 (0.052), the state values from the boundaries -3.9063 and -3.0773 on the preposterior mean of ξ_C.

The value of each record. Table 32.19 lists the candidates. Each row is a record and a decision; the columns give the rows it adds and the expected value of sample information of equation (2.51) on each route, per year in which the decision stands. A reading of the heat load comes from Section 13.6.2: decisive between three needed and two, weights 0.7 and 0.3, it would be worth the mixture's cost less the weighted minima, $88\,784 - (0.7 \times 94\,318 + 0.3 \times 57\,019)$.

32.6.3 The recommendation

Do not wait on a record for the fifth unit, the controller or the store: their gains stand many times above any value in the tables, and a thousand more controller checks would reverse almost nothing. The sixth unit is where a value of information competes with an investment for the same money, the question the previous chapter left to this one. Taking the hundred

Table 32.19: The candidate records ranked by the expected value of sample information [R$/yr]. The state column uses the preposterior mean for a count, the normal summed by quadrature on the plant, and the mixture for the load reading; on the second-order route the plant's first row is 3 669.

Candidate, decision	Rows	State	Exact
100 more CRAC checks, hall units	one on ξ_C	8 680	3 140
400 more CRAC checks, hall units	one on ξ_C	13 665	7 252
Heat-load reading, hall units	the drawing	5 656	6 724
100 more CRAC checks, plant program	one on ξ_C	10 557	13 238
400 more CRAC checks, plant program	one on ξ_C	18 423	18 869
1000 more controller checks, plant program	one on ξ_K	56	0
A further year's record, spares	one on ω	1 026	1 008
Two further years' records, spares	one on ω	1 664	1 585

CRAC checks first gives up one year of the sixth unit's margin and gains the expected value of the record in each later year. Like must be set beside like, each route's margin with its own value: R$ 69 235 against R$ 10 557 a year on the summed normal, 6.6 later years to break even; R$ 28 029 against R$ 3 669 second-order, 7.6 years; R$ 19 066 against R$ 13 238 exact, 1.4 years. The argument assumes that the checks take at most a year and cost next to nothing.

The plan priced. Table 32.20 prices the plan actually open, five units and the controller for a year, the hundred checks in it, and then the program the count gives, against buying six units now and against stopping at five. Each row is a route; the columns give the total of the program with its risk over ten years, the horizon of the annualized price, undiscounted, with one year's deferral and free checks. The plan's total is the first year at $(5, 1)$ and nine years at $\mathbb{E}[\min_a \mathcal{R}(a \mid \tilde{y})]$.

Table 32.20: Ten undiscounted years of the plant program [R$], with the hundred checks free and taken in the first year. The last column is the plan's advantage over buying now, which also bounds what the checks may cost.

Route	Six units now	Stay at five	Checks, then decide	Advantage
Summed normal	4 237 844	4 930 199	4 212 070	25 774
Second-order	4 067 846	4 348 134	4 062 853	4 993
Exact	3 650 798	3 841 455	3 550 726	100 072

The plan is the cheaper over ten years on all three routes, and by sums that differ twentyfold. Over five years it is the dearer on the two state columns, by R$ 27 009 and R$ 13 352, and the cheaper on the exact one, by R$ 33 884. Six units now are the best action only if the decision must be made on the present record. The chapter does not choose a route.

How firm six units are. After the checks five units become the program with probability 0.425 on the preposterior mean of the summed normal, the value reported, against 0.654 exactly; the refit normal gives 0.430 and the second-order route 0.481. By count, the exact program is five units at up to three findings and six from four on, and the refit normal's five at up to two and six from three on, so that at three findings the routes differ. On the hall alone the same checks are worth R$ 8 680 (R$ 3 140) and the heat-load reading R$ 5 656 (R$ 6 724), and on which is worth more the routes differ too; the reading is the one that Chapter 13 put first. When rows arrive, the fit returns to step 3 with them.

32.6.4 What the analysis showed

The first reading said five units, a second controller and fourteen spares, and stated the cost to the real. The analysis adds one unit and one spare and raises the cost, on both routes; but what the engineer expects to follow from an uncertain, convex loss, that the fitted action always falls short, did not survive. On the cooling alone the fitted action was already right, and on the plant it was short because of a charge for units down. What moved the action was the shape of what one more step saves, and the same reasoning that asks for one more spare here would ask for fewer if a spare were dear. The numbers the fit gives are low in every line, and the sixth unit shows how an action can be right in expectation and still lose in most states: it buys protection against a CRAC unavailability that three findings in a hundred cannot exclude. The fifth unit is settled, the sixth is a judgment about that tail, and whether the controller clears an assurance depends on how one finding in a thousand is read.

32.6.5 Limits of the model

The derivations are exact under their assumptions, and five of them need checking before a result is used.

One linear rate and price. The same rate holds at every duration and load, and a unit's price is one annual sum with no discounting. Tables 32.5 and 32.12 show the ranges. A price that fell with the second purchase, or a sixth unit that needs a new pipe, would change the marginal condition.

A unit like the others. New units share q_C and fail independently. A common cause, as in Chapter 15, puts a floor under U_n that further units cannot lower, and would shorten the optimum.

The schedule at more units. The term $C_{\mathrm{C}} M_n$ charges every unit down at every n. Section 32.4.2 shows it moves the plug-in action and the margin, not the Bayes action.

Decisions taken apart. The spares are priced apart from the units, with independent entries. A spare on the shelf shortens a repair and lowers q_C, which would tie the two decisions, and a further unit lowers the cost of a failure that waits. Neither link is in the state.

No lead times and one year. The store is set for a year with no reorder, and the unit decision holds the posterior for a decade, though every year's record would move it.

The next part of the book takes the method through integrated case studies, one plant each.

32.7 What to retain

Definition 32.1 (The marginal condition and the critical fractile)**.** For an integer action a with a price c for each step and the saving $m_a(\mathbf{z}) = \ell(a, \mathbf{z}) - \ell(a+1, \mathbf{z})$, the Bayes action is the first a with $\mathbb{E}[m_a(\mathbf{z}) \mid \mathbf{y}] \le c$ when that expectation falls with a. For a stock against a count, $L(s, \mathbf{z}) = c_{\mathrm{s}} s + C_{\mathrm{s}} \mathbb{E}[(K_f - s)^+ \mid \mathbf{z}]$ and the optimum is the smallest s with $\Pr(K_f \le s \mid \mathbf{y}) \ge 1 - c_{\mathrm{s}}/C_{\mathrm{s}}$.

Four conclusions carry beyond this plant. A unit is worth what its marginal saving is on that plant's structure and schedule, and the closed form $\binom{n}{2} q^{n-2}(1-q)^3$ shows it as the probability that the new unit is critical and up. The direction of a plug-in's error in the action is the sign of the curvature of the marginal saving, while the error in its cost is the sign of the curvature of

the loss; they are separate facts. Changes on one plant are ranked on the residual, and series terms make the residual worth slightly more, not less. And a target and a cost minimum answer different questions, one holding the system's boundary and an assurance fixed and the other its prices. Table 32.21 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 32.21: The terms of the investment decisions and where each is defined.

Term	Expression	Units	Meaning	Defined in
Marginal saving	$m_a(\mathbf{z}) = \ell(a, \mathbf{z}) - \ell(a+1, \mathbf{z})$	R$/yr	one more step	32.1.1
Direction	$\frac{1}{2}\operatorname{tr}(\mathbf{H}_m \boldsymbol{\Sigma})$	R$/yr	Bayes less plug-in	32.1.1
Saving of a unit	$C_{\mathrm{L}} t_{\mathrm{yr}} \binom{n}{2} q^{n-2}(1-q)^3$	R$/yr	the new unit critical and up	32.1.4
Marginal of a spare	$C_{\mathrm{s}} \Pr(K_f > s \mid \omega)$	R$	S-shaped in μ	32.1.4
General forms				
Expected consequence	$\mathcal{R}(a) = \int L(a, \mathbf{z}) p(\mathbf{z} \mid \mathbf{y}) \, \mathrm{d}\mathbf{z}$	R$/yr	one action	32.3.1
Event on the state	$\Pr(\mathbf{z} \in A \mid \mathbf{y})$	–	action cheapest	32.3.2
Tail	$\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y})$	–	pays, or meets a target	32.3.4
Regret of an action	$\mathcal{R}(a') - \min_a \mathcal{R}(a)$	R$/yr	plug-in against Bayes	32.3.5
Specialized forms and answers				
Cost of n units	$(n-3)c_{\mathrm{u}} + C_{\mathrm{L}} t_{\mathrm{yr}} \mathbb{E}[U_n]$	R$/yr	hall	32.3.1
Critical fractile	$\vartheta = 1 - c_{\mathrm{s}}/C_{\mathrm{s}}$	–	the level of the quantile	32.3.3
Program	$a = (n, \delta_{\mathrm{k}})$	–	units and controller	32.3.4
Ratio on the residual	$[\mathcal{R}_{\mathrm{yr}}(\text{before}) - \mathcal{R}_{\mathrm{yr}}(\text{after})]/c_a$	–	order of purchase	32.3.4
Target rule	smallest n with $\Pr(U_n < u^* \mid \mathbf{y}) \geq \kappa$	–	holds no price	32.3.6

32.8 Exercises

Exercise 32.1. *Formulation.* A hall needs three units with a known $q = 0.02$, a unit costs R$ 15 000 a year above three, and a loss of cooling costs R$ 10 000 h^{-1}. Compute the total for $n = 3$ to 6 and name the best. (With $t_{\mathrm{yr}} = 8760$ h: R$ 5 151 581, R$ 219 676, R$ 36 799 and R$ 45 204, so five units.)

Exercise 32.2. *Formulation.* A predictive gives a count 0 to 6 the probabilities 0.1, 0.2, 0.3, 0.2, 0.1, 0.05 and 0.05. A spare costs R$ 120 and a shortage R$ 1 000. Find the optimum and its cost. ($\vartheta = 0.88$; the cumulative sums are 0.1, 0.3, 0.6, 0.8, 0.9, so $s = 4$, at $480 + 1\,000 \times 0.15 =$ R$ 630, against R$ 710 at $s = 3$ and R$ 650 at $s = 5$.)

Exercise 32.3. *Derivation.* Derive the marginal saving of a unit (32.5) and the condition for adding it under a posterior. (Conditioning on the new unit, $U_{n+1} = q \Pr(D \geq n-2) + (1-q)\Pr(D \geq n-1)$; subtracting from $U_n = \Pr(D \geq n-2)$ leaves $(1-q)\Pr(D = n-2) = \binom{n}{2} q^{n-2}(1-q)^3$. The cost is certain, so $\mathcal{R}(n) - \mathcal{R}(n+1) = C_{\mathrm{L}} t_{\mathrm{yr}} \mathbb{E}[U_n - U_{n+1} \mid \mathbf{y}] - c_{\mathrm{u}}$, and the unit is added while this is positive.)

Exercise 32.4. *Derivation.* (a) Derive the optimality condition (32.15). (b) Show that the plug-in's cost at any s is below its expected cost when the plug-in rate is the posterior mean. (c) Show that this does not fix the direction of the plug-in's action. ((a) One more spare adds c_{s} and removes C_{s} when $K_f > s$; the difference rises with s, and it first turns nonnegative where $\Pr(K_f > s) \leq c_{\mathrm{s}}/C_{\mathrm{s}}$. (b) $\mathbb{E}[(K_f - s)^+ \mid \mu]$ has the second derivative $\mu^{s-1}\mathrm{e}^{-\mu}/(s-1)! > 0$ in μ, so Jensen's inequality applies. (c) The action compares $\Pr(K_f > s)$, not the cost, and its second derivative in μ has the sign of $s - \mu$: for s below the mean the expectation lies below the plug-in, and the Bayes level at most the plug-in's.)

Exercise 32.5. *Application.* A store is reordered when it falls to r, with a lead time of half a year, so that $\tau = 22$ unit-years, at the chapter's prices and rate posterior. Find the reorder level on both routes and at the plug-in, and say why they differ from the year's. ($\vartheta = 0.9333$. Cumulative at 7 and 8: state 0.886, 0.939; exact 0.894, 0.944; plug-in 0.910, 0.958. All give $r = 8$. Over a shorter horizon the epistemic share of the variance, equation (11.22), is smaller, and the plug-in is no longer short at these prices.)

Exercise 32.6. *Application.* A budget binds across the plant's three purchases, costing R$ 25 000, R$ 40 000 and R$ 25 000. Find the best set at budgets of R$ 40 000, R$ 50 000, R$ 65 000 and R$ 90 000 on the exact route, and say when the order of the ratios fails. (The sets remove: fifth alone R$ 1 019 408; fifth and sixth R$ 1 063 430; fifth and controller R$ 1 193 940; all three R$ 1 238 005. So the fifth; fifth and sixth; fifth and controller; all three. The order of the ratios, fifth, controller, sixth, skips the controller at R$ 50 000 and takes what still fits, which here is best. It can fail in general, a knapsack whose small high-ratio item blocks a large one, and then the sets are enumerated; and each value must be the set's own, since the ratios interact.)

Part IX

Integrated Engineering Case Studies

Chapter 33

Industrial Pump Reliability

33.1 Six years of a cooling-water fleet

Eight identical pumps supply the cooling water of an industrial plant, and the reliability engineer has six years of their log to take to the budget review: 372 000 operating hours, forty-one failures, forty-one repairs totaling 410 hours. Three requests wait on it. The maintenance planner wants pump 3 overhauled. The workshop proposes to replace the impeller at a fixed age, 8 700 hours. Production wants to know whether two of the pumps will carry a 72-hour campaign.

The quickest reading answers all three. Pump 3 failed nine times in 45 000 hours, $2.0 \times 10^{-4}\,\mathrm{h}^{-1}$, 1.8 times the fleet's $1.10 \times 10^{-4}\,\mathrm{h}^{-1}$, so it is the bad pump. The forty-one times between failures have a mean of 6 531 hours, and a Weibull law fitted to them has the shape 1.35, so the pumps wear out and the workshop's age saves money. Two pumps of that rate almost never fail together in 72 hours, 0.99994. Each answer rests on something the log does not say by itself: that nine against an average of five is more than the scatter of counts, that the time a pump has run since its last failure may be left out of a lifetime fit, and that the two pumps of a pair have nothing in common. The question of this chapter is

what one posterior of the fleet's state says about its rate, whether the pumps differ, whether they wear out, what a pair will do, how often it is down, and whether an age rule pays; which of those answers survive the checks; and what program the numbers support, priced.

The chapter runs the method from step 0 to step 7 as a report and cites the earlier derivations rather than repeating them. What is new is one drawing of the fleet, a population of rates with an unknown spread, whose posterior the normal at the mode cannot represent; there the answers are read from draws, as Section 4.1.3 promised, and the chapter shows what that costs and what else would have served.

Jargon: hierarchical model, partial pooling, shrinkage, funnel, centered and non-centered parametrization, Markov chain Monte Carlo, Metropolis algorithm, Gibbs sampling, potential scale reduction, effective sample size of a chain

A *hierarchical model* gives the units' rates a population with unknown parameters, which are themselves entries of the state. *Partial pooling* is the resulting estimate of a unit between its own record and the population's center, and *shrinkage* is the distance it moves toward that center. A *funnel* is the shape of the joint posterior of a population's spread and its units' deviations: wide at a large spread and narrowing to a neck at a small one. In the *centered parametrization* the deviations themselves are entries of the state; in the *non-centered parametrization* the entries are standardized deviations, multiplied by the spread inside the row. *Markov chain Monte Carlo* draws from a posterior known up to its normalizing constant by a chain whose stationary distribution it is. The *Metropolis algorithm* proposes a move and accepts it with a probability set by the ratio of the posterior at the two states. *Gibbs sampling* draws each block of entries from its conditional. The *potential scale reduction* $\hat{R}$ compares the spread between chains with the spread within them, near one when they agree. The *effective sample size* N_{ess} of a chain is the number of independent draws with the same precision.

Key idea

One posterior answers every question an operator asks of a fleet, and the same record may need different routes through it. A normal at the mode serves a rate with forty-three events behind it and fails a population spread the record cannot separate from zero; the graph, the questions, the rows and the objective are the same either way, and what changes is the computation of steps 3 and 7 and the checks that test it. An honest report of a fleet may say that a difference is not shown, and it says on which route each number stands.

33.1.1 The fleet and its plant graph (step 0)

Figure 33.1 draws the plant. The gray boxes are pinned: the water source and the header. The eight green boxes are the pumps, free nodes in parallel paths, each repaired on failure and each running when it is available. The box with the red border is the shared cause of Chapter 15, drawn on the two pumps that make up a pair for a campaign, and the blue box is the repair crew that serves the fleet. Below them are the entries of the state, each in its green box, and the rows that reach them.

Table 33.1 is the record, which is illustrative. Each row is a pump; the columns give its operating hours, its failures, its running time at each failure since the one before, in hours, in order, and the time it has been running since its last failure when the log closes. The forty-one spells have the mean 6 531 h and the median 4 782 h; with the eight end spells, of 104 229 h, they make up the 372 000 h exactly. They were drawn from one generator (numpy `default_rng`, seed 33), brought onto that mean and median by an order-preserving correction, and assigned to the pumps at random among the assignments that fit each pump's hours. The repairs total 410 h over forty-one jobs. A pump's log starts at a renewal.

33.1.2 The probabilistic questions (step 1)

Table 33.2 lists the six questions in the order they are treated. The first column is the question in engineering language; the second the general form of its category (Chapter 2); the third its form for the fleet; the fourth its category; the last the subsection that answers it. Here λ_i is pump i's rate, λ_{fleet} the fleet's, β the Weibull shape of the spells, T an age and $a = \infty$ run to failure.

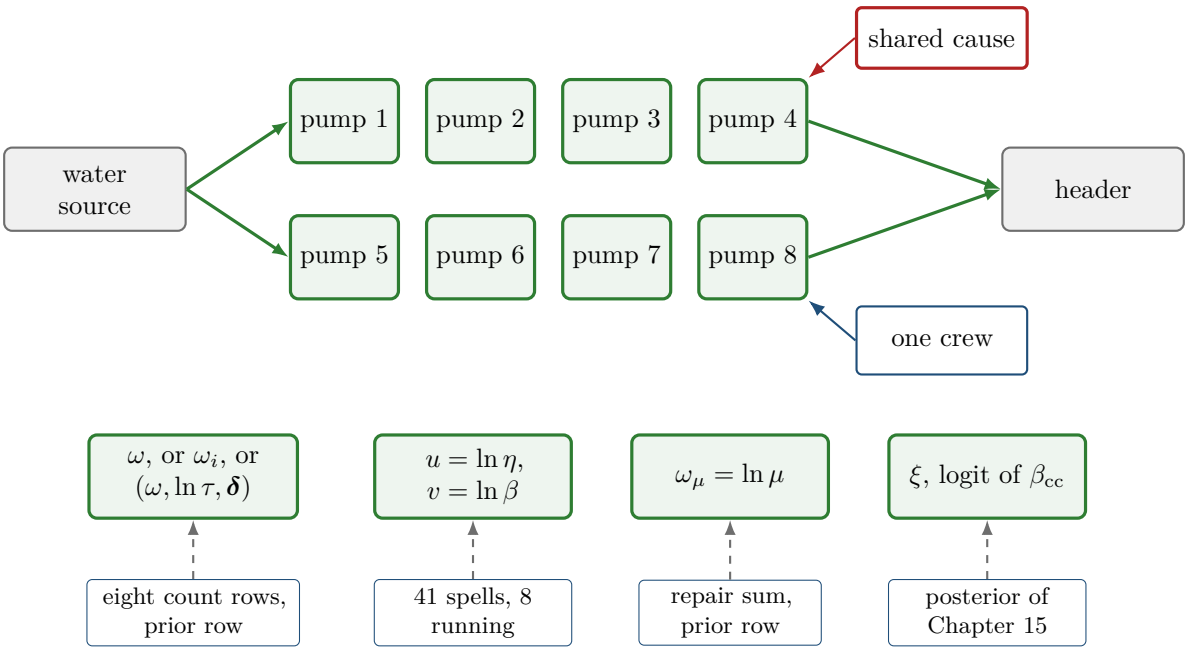


Figure 33.1: The fleet. Above, the plant: pinned, the source and the header (gray); free, the eight pumps (green), all drawing on the source and all repaired by one crew (blue), with a shared cause (red) acting on a pair; the arrows from the cause and the crew stand for one to each pump. Below, the entries of the state (green) and the rows that observe them (dashed). The first box holds one of the three drawings of the rates.

Table 33.1: The record of the eight pumps (illustrative): hours, failures, the spells between failures and the end spell, all in hours.

Pump	Hours h_i	Failures k_i	Spells ending in a failure, in order	End spell
1	48 000	6	11 006, 9 049, 19 214, 2 365, 2 794, 1 777	1 795
2	47 000	4	4 332, 4 539, 13 218, 3 229	21 682
3	45 000	9	4 138, 3 778, 7 052, 4 176, 6 316, 4 003, 1 095, 3 945, 2 815	7 682
4	50 000	5	1 173, 1 450, 10 635, 3 998, 5 925	26 819
5	46 000	3	5 307, 16 126, 13 203	11 364
6	44 000	7	4 782, 11 990, 509, 6 384, 6 768, 1 599, 5 409	6 559
7	49 000	5	14 710, 17 487, 1 107, 7 324, 6 673	1 699
8	43 000	2	1 393, 14 978	26 629
All	372 000	41	sum 267 771	104 229

Table 33.2: The questions asked of the fleet, the general form of their category, their form for the fleet, their category and where each is answered. $\mathbf{y}$ is the record of Table 33.1 with the priors of step 2.

Engineering question	General form	Mathematical form	Category	Answered in
What is the fleet's failure rate, and are the pumps alike?	$p(z_i \mid \mathbf{y})$; $p(\mathbf{y} \mid M_1)/p(\mathbf{y} \mid M_0)$	$p(\omega \mid \mathbf{y})$ pooled; the evidence of one rate, of a population and of eight rates	estimation; hypothesis comparison	33.3.1
Is pump 3 worse than the fleet?	$\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y})$	$\Pr(\lambda_3 > \lambda_{\text{fleet}} \mid \mathbf{y})$; $\Pr(\lambda_3 > 1.5\lambda_{\text{fleet}} \mid \mathbf{y})$	prediction under a threshold	33.3.2
Do the pumps wear out?	$\Pr(z_i > z^* \mid \mathbf{y})$; $p(\mathbf{y} \mid M_1)/p(\mathbf{y} \mid M_0)$	$\Pr(\beta > 1 \mid \mathbf{y})$; the exponential against the Weibull	hypothesis comparison	33.3.3
Will a one-of-two pair carry a 72-hour campaign?	$\int p(y_{\text{new}} \mid \mathbf{z})p(\mathbf{z} \mid \mathbf{y})\,\mathrm{d}\mathbf{z}$	$\mathbb{E}[R_{\text{pair}}(72) \mid \mathbf{y}]$, with and without the common cause	prediction	33.3.4
How often is a pump, and the pair, unavailable?	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$; quantiles	$U = \lambda/(\lambda+\mu)$; U_{pair} with the common cause	estimation of derived quantities	33.3.5
Should the impeller be replaced at an age, and at what age?	$a^* = \arg\min_a \mathcal{R}(a)$; $\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y})$	$\arg\min_T \mathcal{R}(T)$; $\Pr(L(T,\mathbf{z}) < L(\infty,\mathbf{z}) \mid \mathbf{y})$	design	33.3.6

33.1.3 The evidence as rows (step 2)

The modeling choices. Six choices make the rows. The rate of a pump is constant over the six years: this makes hours and failures sufficient, and the trend test of Section 33.4.4 tests it. The failures are read twice, as counts in a rate and as lifetimes in a shape, because the two questions need different summaries: the rows on ω and those on (u, v) are therefore not independent evidence, the state treats them as separate blocks, and no answer here multiplies the two. A repair restores a pump to new, which makes the spells renewal cycles, and it fails if a repair leaves the cause in place. A running pump at the close of the log is a survivor, of the row of equation (7.25): leaving it out is the mistake of Section 7.4.6. The share of common-cause failures is imported from the record of Chapter 15, because this log does not align the pumps' clocks and so carries no count of coincidences: this buys a value for a first-order term, and it fails if the chilled-water pumps of that chapter and these do not share a mechanism, which Section 33.6.4 takes up. And the prices are illustrative.

The rows of the pooled state. The state of the pooled drawing has five entries, $\mathbf{z} = (\omega,\ \omega_\mu,\ \xi,\ u,\ v)^{\mathsf{T}}$, with $\omega = \ln\lambda$ the fleet's log-rate, $\omega_\mu = \ln\mu$ the log repair rate, ξ the logit of the share β_{cc}, and $u = \ln\eta$, $v = \ln\beta$ the Weibull entries. The count rows and the repair row are the terms of Section 6.1.7 and equation (8.16), in the working form (4.43); the spells are the rows (7.29) with their working form (7.33); the prior on the rate is Gamma(2, 20 000 h), of mean $10^{-4}\,\text{h}^{-1}$, and on the repair rate Gamma(2, 20 h); the prior on (u, v) is flat, as in Chapter 7. The share enters as a Gaussian row $(\xi - \hat{\xi}_{15})/\sigma_\xi$ carrying the marginal of equation (15.20), $\hat{\xi}_{15} = -1.8745$, $\sigma^2_\xi = 0.5601$, the route Chapter 15 reports; its correlation with that chapter's rate belongs to other pumps and is not carried. Each row acts on one block:

$$\mathbf{J} = \begin{bmatrix} * & 0 & 0 & 0 & 0 \\ * & 0 & 0 & 0 & 0 \\ 0 & * & 0 & 0 & 0 \\ 0 & 0 & * & 0 & 0 \\ 0 & 0 & 0 & * & * \end{bmatrix} \begin{matrix} \text{prior row on } \omega \\ \text{eight count rows} \\ \text{prior and repair rows,} \\ \text{share row} \\ \text{49 lifetime rows} \end{matrix} \qquad \text{columns } \omega, \omega_\mu, \xi, u, v, \tag{33.1}$$

so that the covariance is block diagonal. With unnormalized rows the information would be $\mathbf{J}^{\mathsf{T}}\mathbf{W}\mathbf{J}$.

The share and Chapter 16's disclosure. Chapter 16 carried the share by the approximation Beta(2, 18), equation (15.14), of mean 0.100, and disclosed that it lowers the answers by about a quarter (Section 16.4.3). There the choice isolated what the chain and the crew contribute. Here the share is the answer to two of the questions, at first order, so the reported value uses the posterior Chapter 15 reports, with its exact marginal (15.13), of mean 0.131, beside it. The Beta approximation answers nothing; Section 33.4.5 states what it would have moved.

Three drawings of the rates. The first entry of the state is replaced under the other drawings. *Pooled*: one rate, as above. *Separate*: eight log-rates ω_i, each with its own count row and its own prior Gamma(2, 20 000 h). *Hierarchical*: a population,

$$\ln \lambda_i = \omega + \tau\delta_i, \qquad \delta_i \sim \mathcal{N}(0,1), \qquad \tau \sim \text{half-normal}(0.5), \tag{33.2}$$

with ω the population's log-median, under the same prior Gamma(2, 20 000 h) as the pooled rate, and τ its spread, which in this chapter alone names that spread and not the interval of Chapter 29. The half-normal of scale 0.5 is a choice: it puts the rates within a factor of about 1.6 of one another at one standard deviation of τ, the tie of Section 6.3.3, where the spread was chosen and not learned, and Section 33.4.5 varies it. The state is $\mathbf{z}_{\mathrm{H}} = (\omega,\ \ell,\ \delta_1, \ldots, \delta_8)^{\mathsf{T}}$ with $\ell = \ln\tau$, the non-centered parametrization. Its count rows have the log mean $\ln h_i + \omega + \mathrm{e}^{\ell}\delta_i$, so that the working Jacobian row of pump i is $\sqrt{\mu_i}\,(1,\ \tau\delta_i,\ 0, \ldots, \tau, \ldots, 0)$, with τ in the column of δ_i, and the curvature term of equation (4.44) does not vanish, since the log mean is a product of two entries. At $\tau = 0$ it is exactly the pooled drawing: every deviation drops out of every row. The centered parametrization takes $d_i = \tau\delta_i$ in place of δ_i, the d_i of Chapter 6, with the row d_i/τ on them.

The losses and the prices. Age replacement has the loss (30.4) of Chapter 30 with $C_{\mathrm{r}} = \text{R\$ } 8\,000$ and $C_{\mathrm{f}} = \text{R\$ } 30\,000$, the mean cycle of equation (30.1), here per pump-hour. Two more prices serve the program of Section 33.5.2. Table 33.3 lists them. Each row is a quantity, with its symbol, value and use. All are illustrative. An overhaul of pump 3 is taken to return it to the population's median rate, with the saving $M(\mathbf{z}) = C_{\mathrm{f}} t_{\mathrm{yr}}(\lambda_3 - \mathrm{e}^{\omega})$ per year against its price.

Table 33.3: The prices of the fleet's decisions (illustrative).

Quantity	Symbol	Value	Used by
Planned impeller replacement	C_{r}	R\$ 8 000	age rule
A failure in service	C_{f}	R\$ 30 000	age rule, overhaul
Overhaul of pump 3, R\$ 40 000 over five years	c_{o}	R\$ 8 000 yr^{-1}	pump 3
Loss of the pair's water supply	C_{L}	R\$ 5 000 h^{-1}	the pair
Staggered maintenance of the pair	c_{st}	R\$ 2 000 yr^{-1}	the pair

33.2 The posterior of the state (step 3)

Every answer is read from the posterior computed here, or, for the population of rates, from the draws of it. The pooled state is fitted once, and the two other drawings of the rates replace its first entry.

33.2.1 The pooled state from the rows

The blocks of pattern (33.1) fit apart. Equation (6.32) with the eight count rows gives $a_n = 43$ and $b_n = 392\,000$ h, so that the gradient vanishes at $\hat{\omega} = \ln(43/392\,000) = -9.11782$ with the information 43. The repair rows give $a_n = 43$ and $b_n = 430$ h, $\hat{\omega}_\mu = \ln 0.1 = -2.30259$, and the same information. Table 33.4 lists the Newton iteration of the lifetime block on the 49 rows, by the gradient and Hessian of equations (7.32) and (7.36), from $v = 0$ and the mean of the 49 times, $u = \ln(372\,000/49)$, where the first gradient entry is $41 - 49 = -8$. Each row is an iterate; the columns give the two entries, the gradient and the Hessian.

Table 33.4: Newton iteration of the lifetime block on the 41 spells and 8 end spells.

Iteration	u	v	$\partial\mathcal{L}/\partial u$	$\partial\mathcal{L}/\partial v$	H_{uu}	H_{uv}	H_{vv}
0	8.93483	0	-8.000	-5.466	49.000	-24.003	67.334
1	9.18081	0.16887	2.092	2.195	54.995	-7.533	71.550
2	9.13795	0.13368	8.5×10^{-2}	9.6×10^{-2}	53.469	-10.491	69.782
3	9.13603	0.13202	1.8×10^{-4}	1.9×10^{-4}	53.389	-10.612	69.695
4	9.13602	0.13201	7.3×10^{-10}	7.9×10^{-10}	53.389	-10.612	69.694

At the minimum, in the lifetime block, $\mathbf{H} = \mathbf{J}^{\mathsf{T}}\mathbf{J} + \mathbf{C}$ with

$$\mathbf{J}^{\mathsf{T}}\mathbf{J} = \begin{bmatrix} 53.389 & -10.612 \\ -10.612 & 28.695 \end{bmatrix}, \qquad \mathbf{C} = \begin{bmatrix} 0 & 0 \\ 0 & 41.000 \end{bmatrix},$$

the curvature term of equation (7.34), not small here. The whole state, with the share's row, is

$$\hat{\mathbf{z}} = \begin{bmatrix} \hat{\omega} \\ \hat{\omega}_\mu \\ \hat{\xi} \\ \hat{u} \\ \hat{v} \end{bmatrix} = \begin{bmatrix} -9.11782 \\ -2.30259 \\ -1.87447 \\ 9.13602 \\ 0.13201 \end{bmatrix}, \qquad \mathbf{\Sigma} = \begin{bmatrix} 0.02326 & 0 & 0 & 0 & 0 \\ 0 & 0.02326 & 0 & 0 & 0 \\ 0 & 0 & 0.5601 & 0 & 0 \\ 0 & 0 & 0 & 0.01932 & 0.00294 \\ 0 & 0 & 0 & 0.00294 & 0.01480 \end{bmatrix}, \quad (33.3)$$

that is $\hat{\lambda} = 1.0969 \times 10^{-4}\,\mathrm{h}^{-1}$, $\hat{\mu} = 0.1\,\mathrm{h}^{-1}$, $\hat{\beta}_{\mathrm{cc}} = 0.1330$, $\hat{\eta} = 9\,284$ h and $\hat{\beta} = 1.141$, with the spreads $\sqrt{\Sigma_{ii}} = (0.1525,\ 0.1525,\ 0.7484,\ 0.1390,\ 0.1216)$. Under the separate drawing each pump has $\hat{\omega}_i = \ln[(2 + k_i)/(20\,000 + h_i)]$ with the variance $1/(2 + k_i)$.

33.2.2 The exact posteriors

Beside the state route stand the conjugate results: $\lambda \sim \mathrm{Gamma}(43,\ 392\,000\,\mathrm{h})$; $\lambda_i \sim \mathrm{Gamma}(2 + k_i,\ 20\,000 + h_i)$ under the separate drawing; $\mu \sim \mathrm{Gamma}(43,\ 430\,\mathrm{h})$; the marginal of the share (15.13), of median 0.116, mean 0.131 and band $[0.018,\ 0.329]$; and, for the lifetime block, the one integral of Section 7.2.2, $\zeta = \eta^{-\beta}$ given the shape distributed as $\mathrm{Gamma}(41, A(\beta))$. Expected costs on it use the rule of equation (30.9). Draws of both routes are 2×10^5 from numpy `default_rng`, seed 33, and expectations on the state route the Gauss–Hermite rule on the axes of $\mathbf{\Sigma}$ with sixty nodes.

33.2.3 The hierarchical posterior and its mode

On the state $\mathbf{z}_{\mathrm{H}}$ the objective is the sum of the eight count terms, the prior term $20\,000\,\mathrm{e}^{\omega} - 2\omega$, the rows $\frac{1}{2}\delta_i^2$, the term $\tau^2/(2{\times}0.5^2)$ and $-\ell$, the last the change-of-variable factor of the logarithm. Newton's iteration with the curvature term reaches a mode at $\hat{\omega} = -9.168$, $\hat{\ell} = -0.274$, that is $\hat{\tau} = 0.761$, and $\hat{\delta}_3 = 0.711$, with $\sqrt{\Sigma_{\ell\ell}} = 0.414$. The eigenvalues of the Hessian run from 1.07 to 46.7, so nothing warns of trouble.

The mode is not where the posterior is. Integrating the eight deviations out by quadrature gives the marginal of τ, whose density is largest at $\tau = 0$: the record does not separate the spread from zero, and for each small τ the deviations stay near their prior, a whole ball of states with little likelihood to gain or lose. The mode sits where one point is highest, which rewards a large τ with the deviations pulled out to fit the counts, and it moves with the coordinate: $\hat{\tau} = 0.620$ on τ itself, and in the centered parametrization there is none, since at $d_i = 0$ the term -7ℓ in the log density grows without bound as $\tau \to 0$. The statement that the mode sits at or near $\tau = 0$ thus holds of the marginal and of the centered density, not of the joint density in the non-centered entries from which a normal would be built.

Figure 33.2 shows the two shapes. In each panel the horizontal axis is $\ell = \ln \tau$; the vertical axis is δ_3 in panel (a) and $d_3 = \tau\delta_3$ in panel (b). The dots are 343 draws of the chain derived below and the closed curve is the 95 % contour of the normal at the mode. In panel (b) the funnel narrows to the left, and panel (a) opens it, the deviations of a small spread being standard normals. The contour covers 26 % of the chain's draws, not 95 %, and it lies to the right of where the draws are.

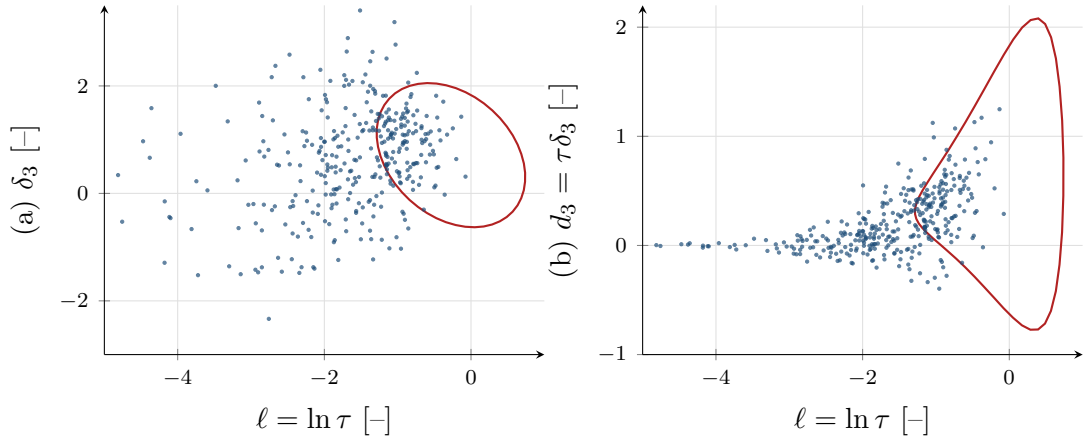


Figure 33.2: The joint posterior of the population's spread and pump 3's deviation: draws of the chain (blue; numpy `default_rng`, seed 33; those below $\ell = -5$ off the axis) and the 95 % contour of the normal at the mode, mapped through $d_3 = \mathrm{e}^{\ell}\delta_3$ in panel (b) (red).

Reweighting from the mode cannot repair this. Draws of the normal, reweighted to the posterior by equation (20.11), have the effective sample size (20.12) of $N_{\text{eff}} = 43$ out of 2×10^5 under seed 33, one draw holding 13 % of the weight, and between 5 and 396 over twenty further seeds; a proposal widened by 1.5 gives 35 to 912, and a Student-t with four degrees of freedom 12 to 1 035. On the pooled drawing the same computation gives 197 962 and on the separate one 10 666 to 89 587. The failure belongs to the proposal centered on the mode and not to reweighting as such: Section 33.4.3 shows a proposal from the priors that works on this record, because eight count rows are weak.

33.2.4 The chain

The target is the normalized product of the rows, $\pi(\mathbf{z}_{\mathrm{H}}) \propto \mathrm{e}^{-\mathcal{L}(\mathbf{z}_{\mathrm{H}})}$. Gibbs sampling would draw each block from its conditional, and where a conditional is a standard distribution that step is exact: on the separate drawing each λ_i given the record is a gamma, and one sweep is an independent draw. Under the population (33.2) no entry has a standard conditional, so the chain uses the Metropolis algorithm on the whole state. From $\mathbf{z}$ it proposes

$$\mathbf{z}' = \mathbf{z} + s\,\mathbf{L}\boldsymbol{\varepsilon}, \qquad \boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \qquad \mathbf{L}\mathbf{L}^{\mathsf{T}} = \boldsymbol{\Sigma}_{\text{pilot}}, \tag{33.4}$$

symmetric in $\mathbf{z}$ and $\mathbf{z}'$, and accepts with the probability

$$\alpha(\mathbf{z}, \mathbf{z}') = \min\bigl(1,\ \exp[\mathcal{L}(\mathbf{z}) - \mathcal{L}(\mathbf{z}')]\bigr), \tag{33.5}$$

keeping $\mathbf{z}$ otherwise. Then $\pi(\mathbf{z})\,\alpha(\mathbf{z}, \mathbf{z}') = \min(\pi(\mathbf{z}), \pi(\mathbf{z}'))$ is symmetric in the two states, so the chain leaves π stationary. Only a difference of objectives enters: the normalizing constant never appears, which is what makes the route possible, and it is also why the draws do not carry the evidence.

Four chains start from dispersed states, with τ at 0.02, 0.1, 0.5 and 1.2, in one generator (numpy `default_rng`, seed 33). Each runs a pilot of 4×10^4 steps that sets $\boldsymbol{\Sigma}_{\text{pilot}}$ and s, both then frozen, and 4×10^5 further steps, of which every twentieth is kept, 8×10^4 draws in all, at the acceptance rate 0.237 and 1.76×10^6 evaluations of the objective against the few that Newton's iteration needs. The chain gives τ the median 0.204 and the band [0.0099, 0.661], against 0.761 and [0.339, 1.713] for the normal at the mode, and it puts probability 0.127 on $\tau < 0.05$, where the normal puts none. The population's median rate e^{ω} has the median $1.063 \times 10^{-4}\,\mathrm{h}^{-1}$ and the band $[0.721,\ 1.490] \times 10^{-4}\,\mathrm{h}^{-1}$. Section 33.4.1 gives the diagnostics.

33.2.5 Which route each answer uses

The book's rule is that the reported value comes from the normal at the mode with the exact value beside it. One drawing departs from it, for the reason just shown: the hierarchical answers are reported from the chain, because the normal misplaces the funnel and reweighting from it cannot correct that, and the normal's value stands beside the chain's with the gap. On this record two other computations also reach the posterior, the nested quadrature and reweighting from the priors; they serve as checks, and the chain is the reported route. Everything else keeps the state route. Table 33.5 says which is which; each row is a kind of answer, and the columns give the reported route and the value beside it.

Table 33.5: The route of each answer. Parentheses in the rest of the chapter hold the second column's other value.

Answer	Reported	Beside it
Pooled and separate rates, their evidences; share, repair rate, shape; mission, unavailability, cost	state route, $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$	exact (gamma, marginal, one integral)
Population spread, partially pooled rates, tails on pump 3	the chain	normal at the mode
Evidence of the hierarchical drawing	Savage–Dickey ratio from the chain	normal at the mode; checked by bridge sampling and quadrature
Value of a record on the hierarchical drawing	simulated outcomes, reweighted draws	rank-one preposterior

33.3 The answers (step 4)

The questions of Table 33.2 are answered from the posterior of Section 33.2, each by the route of Table 33.5. The derivations are those of the chapters cited, and only the fleet's substitutions are shown.

33.3.1 The fleet's rate and the three drawings

Engineering question. What is the fleet's failure rate, and are the pumps alike?

General form. The estimation form of an entry and the ratio of evidences, equations (2.9) and (2.14):

$$p(z_i \mid \mathbf{y}) = \int p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}_{-i}, \qquad B_{10} = \frac{p(\mathbf{y} \mid M_1)}{p(\mathbf{y} \mid M_0)}, \qquad p(\mathbf{y} \mid M) = \int p(\mathbf{y} \mid \mathbf{z}, M)\, p(\mathbf{z} \mid M)\, \mathrm{d}\mathbf{z}. \tag{33.6}$$

Specialized form. The rate is equation (33.6) with $z_i = \omega$ and the pooled normal of equation (33.3). One choice enters: the pumps are comparable if the drawing says so, which is what this question tests and not what it assumes. The comparison is equation (33.6) among the three drawings. Pooled against separate has the Laplace evidence (24.12) on each with the negative binomial (6.68) and (6.69) beside it. The hierarchical drawing nests the pooled at $\tau = 0$ with the priors of condition (24.21), so the Savage–Dickey ratio (24.22) gives their Bayes factor from one density at one point. Its density comes from the draws, without a density estimate at a boundary, by averaging each draw's conditional:

$$p(\tau = 0 \mid \mathbf{y}) \approx \frac{1}{N} \sum_j \frac{\pi_\tau(0)}{\int_0^\infty \pi_\tau(t)\, \exp\!\left[\ell_{\mathrm{c}}(\omega_j, t\boldsymbol{\delta}_j) - \ell_{\mathrm{c}}(\omega_j, \mathbf{0})\right] \mathrm{d}t}, \tag{33.7}$$

with π_τ the half-normal density and ℓ_{c} the summed log-probability of the eight counts.

Building blocks. $\hat{\omega}$ and $\Sigma_{\omega\omega}$ of equation (33.3); the separate fits; the chain of Section 33.2.4, 2×10^4 of its draws in equation (33.7).

The rate. The fleet's rate has the median $\mathrm{e}^{\hat{\omega}} = 1.097 \times 10^{-4}\,\mathrm{h}^{-1}$, 0.96 failures per pump-year, with the band $[0.814,\ 1.479] \times 10^{-4}\,\mathrm{h}^{-1}$. Exactly, the median is $1.088 \times 10^{-4}\,\mathrm{h}^{-1}$ and the band $[0.794,\ 1.448] \times 10^{-4}\,\mathrm{h}^{-1}$, the gap 0.8 % at the center and $0.3 \times 10^{-5}\,\mathrm{h}^{-1}$ at the top. The dispersion statistic of Section 6.4.6 is $X^2 = 7.11$ on seven degrees of freedom, $p = 0.42$, and pump 3's Pearson residual is the largest, -1.81 in the sign convention of that section.

The rates of the pumps. Figure 33.3 draws them. The horizontal axis is the pump; the vertical axis the rate in $10^{-4}\,\mathrm{h}^{-1}$. The gray band is the pooled 95 % band, the blue bars each pump's own record on the separate drawing, and the red bars the same pump partially pooled by the chain. Table 33.6 gives the values, one pump to a row. Each pump is drawn toward the center by an amount the record decides: pump 3's median falls from 1.692 to 1.226, three quarters of the way to the population's 1.063, and pump 8's rises from 0.635 to 0.970.

Table 33.6: The posterior of each pump's rate $[10^{-4}\,\mathrm{h}^{-1}]$: median and band on the separate drawing, state route; the exact gamma; and the chain on the hierarchical drawing.

	Separate, state		Separate, exact		Hierarchical, chain	
Pump	Median	Band	Median	Band	Median	Band
1	1.176	[0.588, 2.353]	1.128	[0.508, 2.121]	1.100	[0.657, 1.800]
2	0.896	[0.402, 1.993]	0.846	[0.329, 1.742]	1.028	[0.543, 1.615]
3	1.692	[0.937, 3.056]	1.641	[0.845, 2.829]	1.226	[0.787, 2.297]
4	1.000	[0.477, 2.098]	0.953	[0.402, 1.866]	1.054	[0.595, 1.663]
5	0.758	[0.315, 1.820]	0.708	[0.246, 1.552]	0.997	[0.492, 1.535]
6	1.406	[0.732, 2.703]	1.355	[0.643, 2.463]	1.150	[0.721, 2.008]
7	1.014	[0.484, 2.128]	0.967	[0.408, 1.893]	1.057	[0.596, 1.667]
8	0.635	[0.238, 1.692]	0.583	[0.173, 1.392]	0.970	[0.438, 1.498]

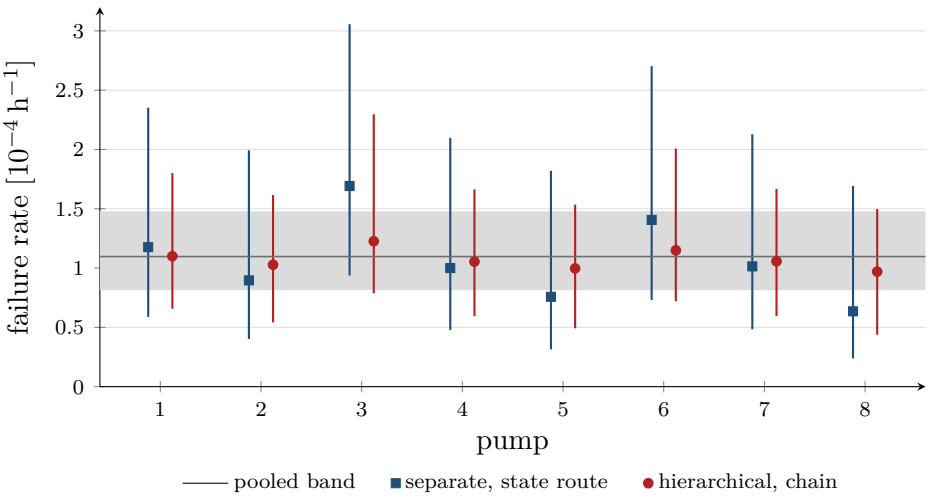


Figure 33.3: The rates of the eight pumps: medians and 95 % bands on the separate drawing (blue, the normal on each ω_i) and on the hierarchical drawing (red, the chain's 8×10^4 draws), against the pooled median and band (gray).

The evidence of the drawings. Table 33.7 sets the three side by side. Each row is a drawing; the columns give its log evidence on the reported route, on the other route and the gap, and then the probability at even prior odds among the three drawings on the reported values, on the exact values and the quadrature, and on the Laplace evidence of all three. Pooled against separate is $\ln B = 1.25$ on the state route and 1.15 exactly, a factor of about three. Pooled against hierarchical is $\ln B = 0.46$ from the chain, a factor of 1.6, and 3.13 from the normal at the mode. The normal thus ranks the hierarchical drawing below the separate one and gives it 0.03; the chain ranks it above and gives it 0.33.

Table 33.7: The evidence of the three drawings, with the Poisson constants included. Reported: Laplace for the pooled and separate drawings, the Savage–Dickey ratio from the chain for the hierarchical one. Probabilities at even prior odds among the three.

Drawing	$\ln p(\mathbf{y} \mid M)$	Other route	Value	Gap	Probability Reported	Exact	Laplace
Pooled, one rate	−18.742	exact	−18.740	−0.002	0.52	0.52	0.75
Hierarchical, a population	−19.205	normal at the mode	−21.869	2.664	0.33	0.32	0.03
Separate, eight rates	−19.992	exact	−19.889	−0.102	0.15	0.16	0.22

So the fleet is not shown to differ, and it is not shown to be alike either: the best drawing holds about half of the probability. The figure is also a statement about the priors, since the separate drawing's cost is the width of its seven extra priors: with the same mean and the weights Gamma(1, 10 000 h), Gamma(4, 40 000 h) and Gamma(0.5, 5 000 h), $\ln B$ of pooled against separate is 2.73, 0.21 and 4.98 exactly. The questions that follow use the pooled rate for a pump of the fleet, as the drawing of largest evidence and for a reason that holds on all three: the fleet's own average rate on the chain, $\sum_i h_i \lambda_i / \sum_i h_i$, has the median 1.089×10^{-4} h^{-1} and the band $[0.797,\ 1.453] \times 10^{-4}$ h^{-1}, within 1 % of the pooled median, so that a rate for an unnamed pump hardly depends on the drawing. Which pump is asked about is the next question.

33.3.2 Pump 3 against the fleet

Engineering question. Is pump 3 worse than the fleet?

General form. A tail of a derived quantity under the posterior, equation (2.12):

$$\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}) = \int \mathbf{1}[f(\mathbf{z}) \geq f^*]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}. \tag{33.8}$$

Specialized form. The question is equation (33.8) with $f = \ln \lambda_3 - \ln \lambda_{\text{fleet}}$ and $f^* = 0$ or $\ln 1.5$. On the hierarchical drawing $\lambda_{\text{fleet}} = \mathrm{e}^{\omega}$, the population's median, so $f = \tau\delta_3$, and the tail is the fraction of the chain's draws above f^*. The first choice is which comparison: pump 3's own posterior against the pooled one, taken as independent, counts pump 3's nine failures on both sides, and it is shown because it is the comparison a report usually makes. The normal at the mode gives its value by draws of $\mathcal{N}(\hat{\mathbf{z}}_{\mathrm{H}}, \boldsymbol{\Sigma}_{\mathrm{H}})$.

Building blocks. The draws of Section 33.2.4; the mode of Section 33.2.3; $\hat{\omega}_3$ with the variance $1/11$.

The tails. Table 33.8 lists them. Each row is a comparison on one route; the columns give the two tails. Within the hierarchical drawing, partial pooling takes the answer from 0.90 and 0.53 to 0.75 and 0.18. Those are conditional on that drawing. Under one shared rate both tails are zero, so averaged over the three drawings at the reported probabilities of Table 33.7 they are 0.384 and 0.133. The chain has f the median 0.102 and the band $[-0.198,\ 0.843]$, so that pump 3's rate stands at 1.15 times the population's and not at the 1.54 of its own record. Nine failures in five years single the pump out and do not demonstrate that it is half again as bad: that tail is 0.18 on the reported route and reaches at most 0.66, on the comparison that takes pump 3 as a different pump from the other seven.

Table 33.8: Is pump 3 worse? The two tails on each comparison and route. The gap between the chain and the normal at the mode is -0.153 and -0.423. Standard error of each chain tail about 0.002.

Comparison and route	$\Pr(\lambda_3 > \lambda_{\text{fleet}})$	$\Pr(\lambda_3 > 1.5\lambda_{\text{fleet}})$
Hierarchical, $\lambda_{\text{fleet}} = \mathrm{e}^{\omega}$, chain (reported)	0.749	0.178
Hierarchical, same, normal at the mode	0.903	0.600
Hierarchical, same, nested quadrature (check)	0.750	0.177
Hierarchical, λ_{fleet} the exposure-weighted mean, chain	0.747	0.137
Own rate against the pooled rate, independent, state route	0.900	0.533
Same, exact gammas	0.876	0.506
Separate, λ_{fleet} the exposure-weighted mean, state route	0.920	0.499
Same, exact gammas	0.926	0.547
Pump 3 against the other seven, state route	0.942	0.658
Same, exact gammas	0.924	0.633
Averaged over the three drawings, reported probabilities	0.384	0.133

The overhaul. The saving $M(\mathbf{z})$ of Table 33.3 has the expectation R\$ 6 196 yr^{-1} on the chain, against R\$ 22 305 yr^{-1} on the normal at the mode and R\$ 16 325 yr^{-1} on the separate drawing, and it is zero on the pooled one. At the price of R\$ 8 000 yr^{-1} the chain and the normal fall on opposite sides, and the separate value, which is another drawing and not a route, falls with the normal. The chain is the reported route here, for the reason of Section 33.2.5, and it says

not to overhaul; the average over the three drawings at the reported probabilities of Table 33.7 is R$ 4 476 yr^{-1} and says the same. A price under about R$ 6 200 would reverse the reported answer, and the saving exceeds R$ 8 000 at the true state with probability 0.30.

33.3.3 Wear-out from the spells

Engineering question. Do the pumps wear out?

General form. A tail of an entry and a ratio of evidences, equations (7.63) and (7.64):

$$\Pr(v > 0 \mid \mathbf{y}) = \int \mathbf{1}[v > 0]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \qquad B_{01} = \frac{p(\mathbf{y} \mid M_0)}{p(\mathbf{y} \mid M_1)}. \tag{33.9}$$

Specialized form. The first form of equation (33.9) is the normal tail (7.65) at $\hat{v}$ and Σ_{vv} of equation (33.3), with the exact one integral beside it. The second has M_0 the exponential and M_1 the Weibull with the prior $v \sim \mathcal{N}(0, 0.5^2)$ of Chapter 7, by the Laplace form (7.69) of the Savage–Dickey ratio, with the exact ratio (7.68) beside it. The family is a choice, and a rising fitted hazard could be a mixture of pumps; Section 33.4.4 tests the fit.

Building blocks. $\hat{v} = 0.13201$, $\Sigma_{vv} = 0.01480$; under the proper prior $\hat{v}_\pi = 0.1246$, $\sqrt{\Sigma_{\pi,vv}} = 0.1188$; the exact marginal of v.

The shape. On the 49 rows the fit is $\hat{\eta} = 9\,284$ h and $\hat{\beta} = 1.141$, with the shape's band [0.899, 1.448] ([0.861, 1.402] exactly, median 1.115). The tail is $\Pr(\beta > 1 \mid \mathbf{y}) = \Phi(0.13201/0.12164) = 0.861$ (0.805 exactly, a gap of 0.056), and the Bayes factor favors the exponential, $B_{01} = 2.43$ (2.78). The maximized log-likelihoods give a likelihood ratio of 1.10, $p = 0.29$. A mild rise is the more likely direction and is far from shown.

What dropping the running pumps does. Fitted to the forty-one spells alone, the same rows give $\hat{\eta} = 7\,147$ h and $\hat{\beta} = 1.352$, $\sqrt{\Sigma_{vv}} = 0.122$, $\Pr(\beta > 1 \mid \mathbf{y}) = 0.993$ (0.983), a likelihood ratio of 5.33, $p = 0.021$, and $B_{01} = 0.25$ (0.45). That is the first reading's wear-out. Figure 33.4 shows why it is wrong. The horizontal axis is the time since the last failure; the vertical axis the survival on a logarithmic scale, on which an exponential is a straight line. The step curve is the Kaplan–Meier estimate of Section 7.4.6 with the end spells as survivors, the solid curve the fit to the 49 rows, the dashed curve the fit to the spells alone, and the dotted line the exponential at the pooled rate. The fit to the spells alone falls below the step curve past 10 000 h and predicts survival to 20 000 h with probability 0.018, while three of the eight end spells have already passed it, at 21 682, 26 629 and 26 819 h. At that fit the eight end spells cost 20.5 nats that the fit never paid.

The constant rate and the Weibull cycles. A renewal process with Weibull cycles has a hazard that rises with age and a long-run rate that is constant, 1/MTTF, so the first and third questions are consistent in form. The test is in the numbers. On the 49 rows 1/MTTF has the median 1.122×10^{-4} h^{-1} and the band $[0.856,\ 1.470] \times 10^{-4}$ h^{-1} (1.108, [0.812, 1.443]), against the pooled 1.097×10^{-4} h^{-1} and $[0.814,\ 1.479] \times 10^{-4}$ h^{-1}: they agree. On the spells alone 1/MTTF $= 1.526 \times 10^{-4}$ h^{-1}, and the mean spell's reciprocal is 1.531×10^{-4} h^{-1}, 39 % above the count's. A mean of completed spells is short of the mean cycle whenever a record closes: a constant rate of the fleet's size, given these counts and hours, would itself leave a mean completed spell of about 7 400 h and 68 600 h in end spells.

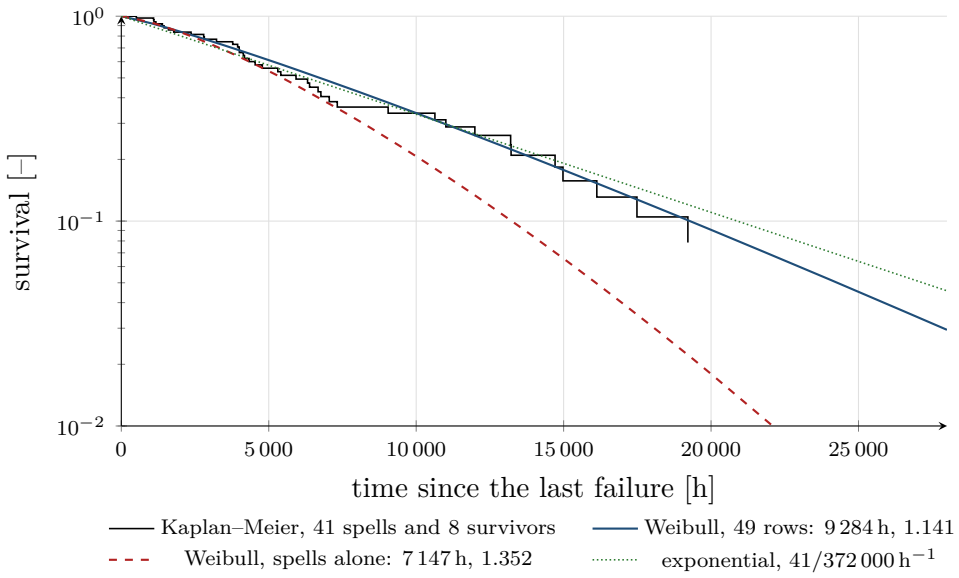


Figure 33.4: The spells against three laws: survival $\exp(-(t/\eta)^{\beta})$ with the parameters of the legend, and the Kaplan–Meier estimate. The exponential is the straight line of a constant hazard.

33.3.4 The campaign

Engineering question. Will a one-of-two pair carry a 72-hour campaign?

General form. The predictive, equation (2.13):

$$p(y_{\text{new}} \mid \mathbf{y}) = \int p(y_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}. \tag{33.10}$$

Specialized form. The question is equation (33.10) with y_{new} the event that the pair lasts $t = 72\,\text{h}$ and $p(y_{\text{new}} \mid \mathbf{z})$ the β-factor form,

$$R_{\text{pair}}(t; \mathbf{z}) = \mathrm{e}^{-\beta_{\text{cc}}\lambda t}\big(2\mathrm{e}^{-(1-\beta_{\text{cc}})\lambda t} - \mathrm{e}^{-2(1-\beta_{\text{cc}})\lambda t}\big), \tag{33.11}$$

averaged over the normal on (ω, ξ) of equation (33.3) by the rule, the two entries independent. Two choices enter. No repair is credited during the campaign: this keeps equation (33.11) closed, and at $\hat{\mathbf{z}}$ a repair at $\mu = 0.1\,\text{h}^{-1}$ raises it only from 0.998904 to 0.998939, since a shared cause leaves nothing to repair onto. And both pumps start the campaign up. The exact value uses Gamma(43, 392 000 h) and the marginal of the share. The value without the common cause is equation (33.11) at $\beta_{\text{cc}} = 0$.

Building blocks. $\hat{\omega}$, $\hat{\xi}$ and their variances; R_{pair}; the sixty-node rule; for the bands, draws of the state.

The numbers. One pump has the predictive 0.99204 (0.99213). Two pumps with no shared cause have 0.999935 (0.999937), a failure probability of 6.5×10^{-5} (6.3×10^{-5}): at $\lambda t = 0.0079$ the averaging over the rate adds a relative 1/43 to a term already second order, so the Jensen and coupling effects are in the sixth decimal. With the share the pair has 0.998714 (0.998917), a failure probability of 1.29×10^{-3} (1.08×10^{-3}), 20 (17) times that of the pair without it. Figure 33.5 draws these. The horizontal axis is the failure probability on a logarithmic scale; each row is a system, a dot the predictive and a bar the 95 % band of the probability at the

state. The common cause is nearly the whole of the pair's mission risk and most of its width, a factor of 10 between the ends of its band against a factor of about 4 from the rate alone. The record leaves that factor to a share read from one double failure.

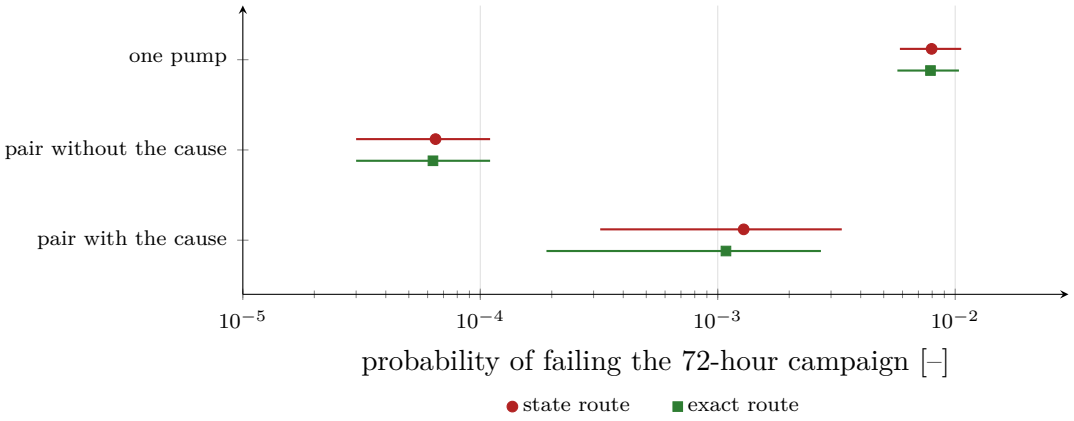


Figure 33.5: The probability of failing the campaign: the predictive (dots) and the 95 % band of equation (33.11) at the state (bars), from 2×10^5 draws, seed 33.

33.3.5 The unavailability of a pump and of the pair

Engineering question. How often is a pump, and the pair, unavailable?

General form. The expectation of a derived quantity and its quantiles, equations (2.10) and (2.11):

$$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}] = \int f(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}. \tag{33.12}$$

Specialized form. The question is equation (33.12) with $f = \lambda/(\lambda + \mu)$ for a pump and $f = U_{\text{pair}}$ of equations (15.41) and (15.42) for the pair. The processes are taken as independent two-state ones, Chapter 15's choice: at $\hat{\mathbf{z}}$ the chain of equation (16.9) with one crew gives 1.478×10^{-4} against 1.468×10^{-4}, so the shared crew adds under one percent. On the state the gradient of $\ln U_{\text{pair}}$ at $\hat{\mathbf{z}}$ in $(\omega, \omega_\mu, \xi)$ is $\mathbf{g} = (1.006,\ -1.006,\ 0.860)^{\mathsf{T}}$, and $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ splits by entry since $\boldsymbol{\Sigma}$ is diagonal there.

Building blocks. $\hat{\omega}$, $\hat{\omega}_\mu$, $\hat{\xi}$ of equation (33.3); the rule in three dimensions; draws for the bands.

The numbers. The repair rate has the median $0.1\,\mathrm{h}^{-1}$, a repair of 10 h ($0.0992\,\mathrm{h}^{-1}$, and a mean repair of 10.24 h). Table 33.9 carries the campaign and the unavailabilities together: each row is a quantity, and the columns give the reported value, the exact value, the gap and the two bands. The pair is down 1.75×10^{-4} (1.49×10^{-4}) of the time, 1.53 h (1.30 h) a year, 133 (113) times the 1.32×10^{-6} of two independent pumps at the total rate. The variance of $\ln U_{\text{pair}}$ at $\hat{\mathbf{z}}$ is 0.461, of which the entry ξ holds 0.414, 90 %; on the exact posterior's variances it holds 92 %.

33.3.6 The age rule

Engineering question. Should the impeller be replaced at an age, and at what age?

Table 33.9: The campaign and the unavailabilities: expected values on the state route and the exact route, the gap, and the 95 % bands of the value at the state.

Quantity	State	Exact	Gap	Band, state (exact)
One pump, $R(72)$	0.99204	0.99213	−0.00009	[0.98940, 0.99415] ([0.98963, 0.99429])
Pair, no common cause	0.999935	0.999937	−0.000002	[0.99989, 0.99997] (the same)
Pair with the common cause	0.998714	0.998917	−0.000203	[0.99667, 0.99968] ([0.99728, 0.99981])
Pump, U	1.121×10^{-3}	1.122×10^{-3}	-1×10^{-6}	$[0.72,\ 1.67] \times 10^{-3}$ ($[0.72,\ 1.68] \times 10^{-3}$)
Pump, hours a year	9.82	9.82	0	
Two independent pumps, U^2	1.32×10^{-6}	1.32×10^{-6}	0	
Independent part of the pair, q_{ind}^2	9.5×10^{-7}	1.00×10^{-6}	-5×10^{-8}	
Pair, U_{pair}	1.752×10^{-4}	1.488×10^{-4}	2.6×10^{-5}	$[0.36,\ 4.78] \times 10^{-4}$ ($[0.20,\ 3.97] \times 10^{-4}$)
Pair, hours a year	1.53	1.30	0.23	

General form. The action that minimizes the expected consequence and a tail of the saving, equations (30.13) and (30.15):

$$\mathcal{R}(a) = \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}], \qquad a^* = \arg\min_a \mathcal{R}(a), \qquad \Pr(L(\infty, \mathbf{z}) - L(T, \mathbf{z}) > 0 \mid \mathbf{y}). \tag{33.13}$$

Specialized form. The question is equation (33.13) with L the loss (30.4) on (u, v), the rule (30.12) on the block of equation (33.3), and equation (30.9) beside it. Two choices enter. Each failure and each planned replacement renews the pump: the renewal–reward ratio needs it, and it is the same reading as the spells' rows. The cost is that of the impeller alone, not of the loss of supply, which the pair's term carries.

Building blocks. $\hat{u}$, $\hat{v}$ and their block; the condition (30.3), which has a root only for a rising hazard; the sixty-node rule and the exact grid.

The cost against the age. Figure 33.6 draws it. The horizontal axis is the age T; the vertical the cost rate. The dashed curve is the loss at $\hat{\mathbf{z}}$, the heavy curve $\mathcal{R}(T)$ on the state route with the band of the loss at the state shaded, the thin curve the exact route, and the two dotted lines run to failure on each route. The curves fall from above and reach those lines from above: on neither route does $\mathcal{R}(T)$ have an interior minimum in the 400 000 h searched. The Bayes action is run to failure, at R$ 3.399 h^{-1} (R$ 3.342 h^{-1}). At $\hat{\mathbf{z}}$ the plug-in action is a root at $T^*(\hat{\mathbf{z}}) = 45\,834$ h, which saves R$ 0.0001 h^{-1} on the plug-in and costs R$ 0.003 h^{-1} (R$ 0.005 h^{-1}) in expectation. The expected cost is within 1 % of the minimum above 20 800 h (24 500 h).

Table 33.10 prices it. Each row is an action; the columns give the expected cost on each route and the probability that it beats run to failure at the true state. The first block is the reported fit; the second is the fit to the spells alone, which is where the first reading stands. There a minimum exists, at 8 884 h (9 996 h), with a saving of R$ 0.082 h^{-1} (R$ 0.049 h^{-1}) and a probability of winning of 0.73 (0.70), a marginal case. On the record with its running pumps the same age is no longer marginal: it costs R$ 0.277 h^{-1} (R$ 0.337 h^{-1}) more than running to failure, and it wins with probability 0.13 (0.08).

The shape at which the rule would pay. With the mean life held at the record's 8 856 h, condition (30.3) gives a saving at the plug-in of 0.6 % at $\beta = 1.25$, 2.2 % at 1.35, 6.0 % at 1.5, 13 % at 1.75 and 19 % at 2. A rule that clearly pays thus needs a shape of about 1.75, with $T^* = 6\,926$ h, and the posterior puts probability 2.2×10^{-4} (2.2×10^{-5}) there. Knowing the

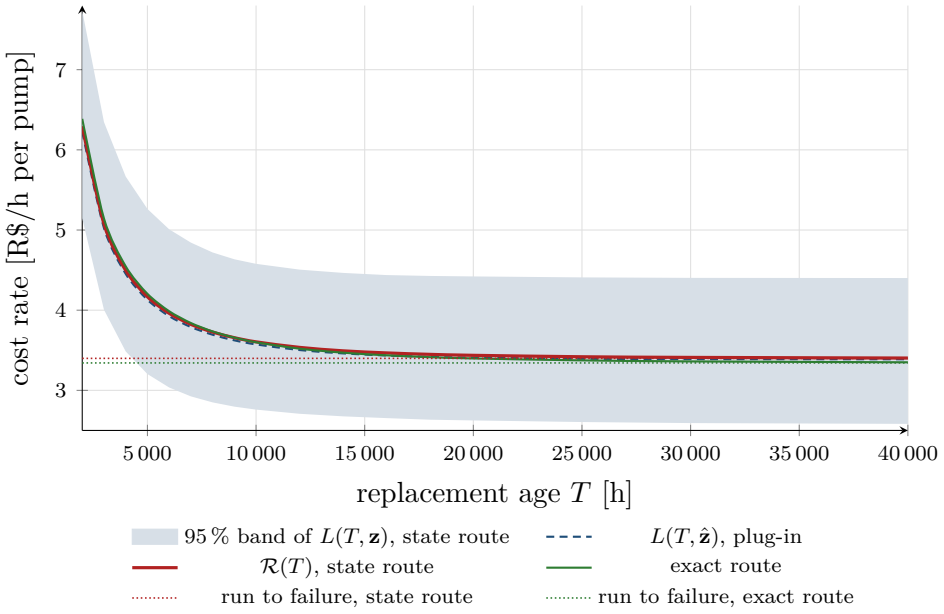


Figure 33.6: The cost rate of an age rule on the 49 rows: the loss at the MAP state, the expected cost on the two routes, the band of the loss, and run to failure. $C_\mathrm{r} = \mathrm{R\$\,8\,000}$, $C_\mathrm{f} = \mathrm{R\$\,30\,000}$ (illustrative).

Table 33.10: The age rule, state route with exact beside: expected cost [R$/h per pump], the saving against run to failure, and the probability that the rule is cheaper at the true state. Costs and savings by the rule on each route, the probability from 2×10^5 draws, seed 33. T^B is the route's Bayes action.

	State route			Exact route		
Action	$\mathcal{R}$	Saving	Probability	$\mathcal{R}$	Saving	Probability
The 49 rows						
Run to failure, the Bayes action	3.399	–	–	3.342	–	–
Age 8 700 h	3.676	−0.277	0.125	3.679	−0.337	0.085
Age 15 000 h	3.479	−0.080	0.283	3.452	−0.109	0.221
Age 30 000 h	3.411	−0.012	0.448	3.361	−0.019	0.376
The spells alone, dropping the running pumps						
Run to failure	4.589	–	–	4.558	–	–
T^B: 8 884 h (9 996 h)	4.507	0.082	0.732	4.509	0.049	0.699

shape and scale exactly would be worth R$ 0.019 h^{-1} (R$ 0.013 h^{-1}), or about R$ 1 350 (R$ 900) a year over the eight pumps, by equation (30.18).

33.4 The checks (step 5)

The answers rest on two routes, a chain, a record read two ways and one prior per entry. Each is tested here. There is no physics row, so no limit $\sigma_p \to 0$ arises.

33.4.1 The chain and its diagnostics

Table 33.11 gives the diagnostics of the chain and of the same algorithm on the centered parametrization, with the nested quadrature as the reference: given (ω, τ) the deviations are independent, so the posterior is a grid in two entries over eight one-dimensional integrals. Each row is a diagnostic or an answer. Figure 33.7 draws ℓ along one chain of each: the horizontal axis is the step in thousands, the vertical the lowest ℓ in each of 400 blocks, so that the excursions into the neck show.

Table 33.11: Four chains, 4×10^5 steps each after the pilot, every twentieth kept, on the two parametrizations, with the quadrature of the same posterior. $\hat{R}$ is the split potential scale reduction.

Quantity	Non-centered	Centered	Quadrature
Acceptance rate	0.237	0.323	
$\hat{R}$ of ℓ; of $\mathbf{1}[\tau < 0.05]$	1.0002; 1.0000	1.0039; 1.0048	
N_{ess} of ℓ	25 607	775	
N_{ess} of $\mathbf{1}[\tau < 0.05]$	37 499	592	
N_{ess} of the deviation of pump 3	37 570	5 601	
Lowest ℓ visited	-11.7	-4.1	
$\Pr(\tau < 0.05)$; range over the four chains	0.127; 0.122 to 0.130	0.095; 0.079 to 0.119	0.128
τ: 2.5 % point, median, 97.5 % point	0.0099, 0.204, 0.661	0.0269, 0.212, 0.675	0.0097, 0.204, 0.664
$\Pr(\lambda_3 > e^{\omega})$; $\Pr(\lambda_3 > 1.5e^{\omega})$	0.749; 0.178	0.762; 0.187	0.750; 0.177

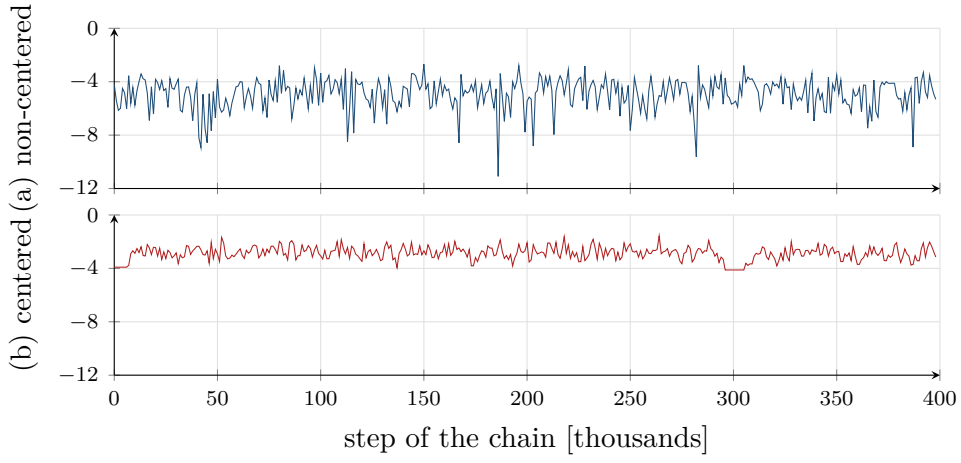


Figure 33.7: The lowest $\ell = \ln \tau$ in each block of 1 000 steps, along the first chain of each parametrization (numpy `default_rng`, seed 33).

The non-centered chain agrees with the quadrature within its standard error on every answer. The centered chain fails in the way Section 4.1.3 warned of: its $\hat{R}$ is close to one, 1.004, yet it

never goes below $\ell = -4.1$, and its sticking is visible in Figure 33.7: the block minimum stays at -3.91 for blocks 0 to 6 and at -4.12 for blocks 296 to 305, the longest stay 8 660 steps without a move against 2 800 for the non-centered chain. At a small spread its deviations must all be small together, a move a single proposal scale cannot make, so the neck is under-visited: the lower quantile of τ is high by a factor of 2.7, $\Pr(\tau < 0.05)$ is low by a quarter and the tail on pump 3 high by 0.010, with a thirty-fold loss of effective sample on ℓ. The trace, the effective sample and the reference expose it, and $\hat{R}$ barely does. That value also depends on tuning: two other settings of the proposal give $\hat{R} = 1.009$ and 1.020, with $N_{\text{ess}} = 520$ and 202, and in both the lowest ℓ stays above -4.3. A chain's diagnostics show disagreement between chains; they cannot show a region none of them entered.

33.4.2 The normal at the mode beside the chain

Table 33.12 sets the normal at the mode beside the chain on the pooled drawing, where both should agree, and on the hierarchical one, where they should not. Each row is an answer. On the pooled drawing the normal's median stands 0.7 % above the chain's, the known offset of a log-rate's median, and the chain equals the exact gamma within its error. On the hierarchical drawing the normal has the spread high by a factor of 3.7, gives no mass to a small spread, and overstates pump 3 and the overhaul.

Table 33.12: The normal at the mode against the chain, on the two drawings. Rates in $10^{-4}\,\text{h}^{-1}$. N_{eff} of the normal as proposal, 2×10^5 draws; N_{ess} of the chain on ω.

Answer	Normal	Chain	Gap
Pooled drawing (exact: 1.088, [0.794, 1.448], $\ln p(\mathbf{y} \mid M) = -18.740$*)*			
Median of the rate	1.097	1.089	0.008
Band of the rate	[0.814, 1.479]	[0.794, 1.447]	0.02, 0.03
N_{eff}; N_{ess}	197 962	176 182	
Hierarchical drawing			
Median of τ	0.761	0.204	0.557
Band of τ	[0.339, 1.713]	[0.0099, 0.661]	
$\Pr(\tau < 0.05)$	0	0.127	−0.127
Median of e^{ω}	1.043	1.063	−0.020
$\Pr(\lambda_3 > \text{e}^{\omega})$	0.903	0.749	0.153
$\Pr(\lambda_3 > 1.5\text{e}^{\omega})$	0.600	0.178	0.423
Overhaul saving $\mathbb{E}[M]$ [R$/yr]	22 305	6 196	16 109
$\ln B$, pooled against hierarchical	3.127	0.465	2.662
N_{eff}; N_{ess} of ω	43	32 484	

33.4.3 The evidence by a second route

The chain's draws do not carry the normalizing constant, so on that route the hierarchical drawing's evidence took a computation of its own; the cost belongs to the chain and not to the drawing. Equation (33.7) gives the posterior density 2.540 at $\tau = 0$ against the prior's 1.596, $\ln B = 0.465$, between 0.453 and 0.479 over the four chains. Bridge sampling between the draws and a normal fitted to them, 2×10^4 of each, gives 0.475; with a Student-t of five degrees of freedom, 0.464. The nested quadrature gives $\ln p(\mathbf{y} \mid M) = -19.2163$ and so 0.476, with the density 2.568. The cost is 2×10^4 one-dimensional integrals against the one formula of the Laplace evidence. The second route agrees and the third is exact, and the Laplace value does not, at 3.127.

Reweighting from the priors. Draw τ and $\boldsymbol{\delta}$ from their priors and ω from the pooled normal widened by 1.3, and weight by the rest of equation (20.11). The effective sample is 12 to 15 % of the draws, above the one tenth of Section 20.2.4. At 2×10^6 draws over three seeds it gives $\ln p(\mathbf{y} \mid M) = -19.216$ to -19.220, the tails 0.749 to 0.750 and 0.174 to 0.176, and the median of τ 0.203, and it carries the evidence directly. At 2×10^5 evaluations its N_{eff} is about 29 000, against the chain's N_{ess} of 25 607 on ℓ from 1.76×10^6. So on this record the chain was not the only way nor the cheapest. The proposal works because the rows are weak: with the same record's counts and hours multiplied by ten, its N_{eff} falls to about 3 of 2×10^5.

33.4.4 The record against its model

The trend. The normalized failure times of Section 6.4.6 sum over the pumps to $S_u = 20.05$, against the mean 20.5 and the standard deviation 1.85 of a constant rate, $z = -0.24$, $p = 0.81$. The record shows no trend; the trend tests of Chapter 12 would apply to a longer one.

The dispersion. The reference chi-square gives $p = 0.417$, and the exact multinomial distribution, by 10^6 draws, 0.417.

The Weibull residuals. The cumulative hazards of the 49 rows at $\hat{\mathbf{z}}$ sum to 41 by the stationary condition, and their Kaplan–Meier estimate stands at most 0.106 from the unit exponential, inside the 0.21 of a 95 % band for forty-one events.

The renewal consistency. Section 33.3.3 set 1/MTTF beside the count's rate: 1.122 against $1.097 \times 10^{-4}\,\text{h}^{-1}$ on the 49 rows and 1.526 against 1.097 on the spells alone. The first passes and the second does not.

33.4.5 The routes, the priors and the share

Table 33.13 varies the prior on τ, by the quadrature. Each row is a prior; the columns give the evidence against the pooled drawing, the median of τ, the two tails and the prior's own tail at 1.5. No prior brings the tail at 1.5 above 0.22 or makes the hierarchical drawing beat the pooled one. The tail 0.178 stands beside a prior 0.134, the record moving its odds by 1.40; at scale 0.25 it goes from 0.05 to 0.09, and at scale 1 from 0.235 down to 0.22. Over half-normal scales from 0.1 to 5, $\ln B$ runs from 0.01 to 2.56. That pump 3 is not demonstrated survives every scale; the number 0.18 and the drawing's probability 0.33 are largely the prior's.

Table 33.13: The hierarchical answers under other priors on τ, by the nested quadrature.

Prior on τ	$\ln B$, pooled against hierarchical	Median τ	$\Pr(\lambda_3 > \text{e}^{\omega})$	$\Pr(\lambda_3 > 1.5\text{e}^{\omega})$	Prior tail
Half-normal, scale 0.25	0.14	0.139	0.71	0.09	0.05
Half-normal, scale 0.5 (reported)	0.48	0.204	0.75	0.18	0.13
Half-normal, scale 1	1.01	0.231	0.77	0.22	0.24
Half-Cauchy, scale 0.25	0.43	0.137	0.71	0.11	0.13
Half-Cauchy, scale 0.5	0.78	0.187	0.74	0.17	0.21

The share. The state route has the share's median 0.133 and mean 0.155; the exact marginal has 0.116 and 0.131. That gap carries the 1.752 against the 1.488×10^{-4} of the pair, 18 %, and it is the largest gap of the chapter on a reported quantity of the pooled state. The approximation Beta(2, 18), had it been used, would give 0.999160 for the campaign and 1.13×10^{-4} for the pair, 24 % below the exact value, which is the quarter that Chapter 16 disclosed. Section 15.4.5 varies the share's own prior.

The rate's prior. It is light against forty-one failures for the rate itself: the three alternatives of Section 33.3.1 move the pooled median by under 1 %, and nothing downstream changes sides, though they move the evidence, as that section shows.

33.5 The report (step 6)

The log does not say what its first reading said. Pump 3 is 1.8 times the fleet on its own count and, within the hierarchical drawing, 1.15 times the fleet partially pooled, with a probability of 0.18 of being half again as bad, 0.13 averaged over the drawings. The wear-out at a probability of 0.99 comes from dropping the eight pumps still running; with them the shape is 1.14, the probability 0.86 (0.81), and an impeller rule at 8 700 h loses money, by about R$ 19 400 (R$ 23 600) a year over the fleet. The record's end spells are an unusually long constructed residual, 104 229 h against 68 600 h expected under a constant rate given the counts ($P = 0.05$) and 62 600 h at the fit ($P = 0.04$), so that reversal is a statement about this record. And the pair's campaign, 0.99994 on the product, is 0.99871 (0.99892) with a shared cause in it. The answers follow in the order they were asked.

33.5.1 The answers and their bands

Table 33.14 is the report. Each row is a question in the engineer's words; the second column gives the answer on the route of Table 33.5, the third the other route's value with what that route is, the fourth the gap, the fifth the band and the qualification, the last the chapter that derives it.

The report line. Over six years the eight cooling-water pumps ran 372 000 hours with forty-one failures, a fleet rate of $1.10 \times 10^{-4}\,\mathrm{h}^{-1}$, about one failure per pump-year, with a 95 % band from 0.81 to 1.48 (0.79 to 1.45). The record does not show the pumps to differ, and it does not show them to be alike: one shared rate holds about half of the probability at even prior odds. If the pumps are a population, pump 3 stands at about 1.15 times the fleet once partially pooled, above the fleet with probability 0.75 and half again above it with probability 0.18; if they share one rate those are zero, and averaged over the drawings they are 0.38 and 0.13. A rising hazard is more likely than not, 0.86 (0.81), with a shape of 1.14, and it is not established; the record slightly prefers a constant hazard. A pump is down about ten hours a year. A pair with a shared cause fails a 72-hour campaign with probability 1.3×10^{-3} (1.1×10^{-3}), and is down together 1.5 h (1.3 h) a year; nearly all of both is the shared cause, whose share comes from another record, with a band spanning a factor of 12 (18). All prices are illustrative.

What changed with the route. The chain changed step 4 for the hierarchical drawing and its checks, and the computation behind steps 3 and 7. It did not change the graph, the questions, the rows or the objective.

33.5.2 The program and its price

Table 33.15 prices the programs on offer. Each row is a program for the fleet and one pair; the columns give its annual cost on each route, made up of the impeller failures and planned replacements of the fleet, the expected loss of the pair's supply at C_{L}, and the staggering at c_{st}.

The recommended program keeps the impellers on run to failure, at R$ 238 183 (R$ 234 222) a year, leaves pump 3 alone, and staggers the maintenance of the pair. The first needs no further record: no age beats it on either route. The second holds on the reported route and at this price, and it depends on which value is taken: R$ 6 196 on the chain says no, R$ 22 305 on the normal at the mode and R$ 16 325 on the separate drawing say yes, as Section 33.3.2 states.

Table 33.14: The report of the fleet: the answer on its route, the other route's value and the gap, the bands, and the chapter behind each. Prices illustrative.

Question	Answer	Other route	Gap	Band, qualification	From
Fleet's rate	median $1.097 \times 10^{-4}\,\mathrm{h}^{-1}$, state	exact $1.088 \times 10^{-4}\,\mathrm{h}^{-1}$	0.009×10^{-4}	[0.814, 1.479] ([0.794, 1.448]); 0.96 per pump-year	6
Are the pumps alike?	not shown either way: $\ln B$ pooled against separate 1.25, state; against hierarchical 0.46, chain; probabilities 0.52, 0.33, 0.15	exact 1.15; normal at the mode 3.13	0.10; −2.66	$X^2 = 7.11$, $p = 0.42$; $\ln B$ from 0.21 to 4.98 over the rate's prior	24
Is pump 3 worse?	within the hierarchical drawing 0.749, and 0.178 at 1.5 times; chain	normal at the mode 0.903, 0.600	−0.153; −0.423	not demonstrated; averaged over the drawings 0.384, 0.133; zero under one rate; against the other seven 0.942, 0.658 (0.924, 0.633); no overhaul at R$ 8 000 yr^{-1}	6
Do they wear out?	$\beta = 1.141$; $\Pr(\beta > 1) = 0.861$; $B_{01} = 2.43$; state	exact 0.805, 2.78	0.056; −0.35	β in [0.899, 1.448] ([0.861, 1.402]); 0.993 (0.983) if the end spells are dropped	7
A 72-hour pair	0.998714 with the cause, 0.999935 without; state	exact 0.998917, 0.999937	−0.000203; −0.000002	[0.99667, 0.99968] ([0.99728, 0.99981]); failure probability 20 (17) times	9, 15
Unavailability	pump 1.12×10^{-3}, 9.8 h a year; pair 1.75×10^{-4}, 1.53 h a year; state	exact 1.12×10^{-3}; 1.49×10^{-4}, 1.30 h	0; 2.6×10^{-5}	pump $[0.72, 1.67] \times 10^{-3}$; pair $[0.36, 4.78] \times 10^{-4}$ ($[0.20, 3.97] \times 10^{-4}$); ξ holds 90 % of the variance	8, 15
Age rule?	no: run to failure, R$ 3.399 h^{-1}; age 8 700 h costs 0.277 more, wins 0.125; state	exact R$ 3.342 h^{-1}; 0.337, 0.085	0.057; −0.060; 0.040	on the spells alone 8 884 h, saves 0.082, wins 0.73 (0.70); clear only at a shape near 1.75	30

Table 33.15: Annual cost of the programs [R$/yr], illustrative prices: eight pumps at 8 760 hours, one pair. The staggering removes half of the share, the assumption of Chapters 15 and 16.

Program	State	Exact	Gap
Run to failure, nothing else	245 857	240 739	5 118
Run to failure, pair staggered (recommended)	244 041	239 503	4 538
Age 8 700 h, pair staggered	263 457	263 083	374

The third has a basis in the numbers and a limit that the numbers cannot lift. The basis: the common-cause term costs R$ 7 632 (R$ 6 473) a year at C_{L}, so that staggering at R$ 2 000 a year pays if it removes more than 26 % (31 %) of the share. If it removes half, the net saving is R$ 1 816 (R$ 1 236) a year, the campaign rises to 0.99932 (0.99943), and staggering pays at the true state with probability 0.75 (0.66). If it removes a quarter it does not pay, by R$ 92 (R$ 382). The limit: nothing in this record, or in that of Chapter 15, says what fraction of the shared cause is a repeated maintenance error and not a supply or a control fault. The half is an assumption of those chapters, and the recommendation is conditional on a fraction above the break-even. It is cheap to try, and the time-stamped log is what will test it.

Principle 33.1. *A case study answers each question on the route the posterior allows and says which. Where the normal at the mode represents the state it is reported, with the exact value beside it; where it does not, as in the funnel of a population's spread, the draws are reported and the normal stands beside them, and the evidence and the next measurement then need their own computations. A record is read with all its rows: a survivor is evidence about a lifetime. A recommendation is the program priced, and a term of that price which no record informs is stated as a condition.*

Common mistake

Five errors recur with a fleet's log. Fitting a lifetime to the times between failures and leaving out the running pumps: here it turns a shape of 1.14 into 1.35, a probability of 0.86 into 0.99, and a losing age rule into a marginal one, on a record whose end spells are long by construction ($P = 0.05$ under a constant rate). Comparing a unit's rate with a fleet rate that contains it, as independent: 0.90 against 0.75 partially pooled, within that drawing. Reading a dispersion test that passes as evidence that the units are alike: $p = 0.42$ here beside a probability of 0.52 for one rate. Building the normal at the joint mode of a population's spread: a spread of 0.76 where the posterior's median is 0.20. And taking $\hat{R}$ near one as convergence: the centered chain has 1.004 and misses the neck.

33.6 The next measurement (step 7)

The pooled state keeps the rank-one preposterior of equation (2.17). On the hierarchical drawing it no longer applies, and the candidates are valued by simulating their outcomes. Two more questions are settled first: what the share is worth, and what a record on pump 3 would have to show.

33.6.1 The candidate records and their price

The route. For a candidate, an outcome is drawn from the predictive by choosing a draw of the chain and simulating the Poisson counts of the new hours at its rates; the posterior after it is the draws reweighted by those counts, with the check that the reweighted effective sample holds. The value is the expected value of sample information (2.51) of the overhaul decision at c_{o}, the gain over the present decision averaged over the outcomes. The cost is 2 000 outcomes over 4×10^4 draws and up to eight rows, against one matrix inversion for the rank-one route beside it, which is equation (2.17) at the mode with the rows $\sqrt{\mu_i}(1,\ \tau\delta_i,\ \tau)$. Three outcomes refitted by the quadrature check the reweighting: for $k = 0$, 3 and 8 failures of pump 3 in two more years, the tails 0.646, 0.786 and 0.958 against 0.644, 0.785 and 0.958, and 0.075, 0.200 and 0.635 against 0.075, 0.198 and 0.633.

The candidates. Table 33.16 lists them. Each row is a record; the columns give, on each route, the value at c_{o}, the probability that the overhaul is the action after it and the probability that the record establishes the pump as above the fleet, $\Pr(\lambda_3 > \mathrm{e}^{\omega}) \geq 0.95$, with the spread of $\ln \lambda_3 - \omega$ after it.

Table 33.16: The candidate records on pump 3: simulated outcomes reweighted, 2 000 outcomes, standard error of the value 55 to 90, against the rank-one preposterior at the mode. Present spread of $\ln \lambda_3 - \omega$: 0.265 on the chain, 0.409 at the mode. Value in R$/yr; the rank-one value takes the delta moment of M throughout, R$ 21 202 at present against the R$ 22 305 of the normal's draws.

	Simulated outcomes				Rank-one preposterior			
Record	Value	Overhaul	Established	Spread	Value	Overhaul	Established	Spread
Two years, fleet	882	0.22	0.012	0.244	456	0.98	0.322	0.375
Two years, pump 3	813	0.20	0.009	0.251	366	0.98	0.285	0.380
Five years, pump 3	1 120	0.25	0.031	0.234	772	0.93	0.411	0.356
Five years, fleet	1 447	0.24	0.058	0.219	1 011	0.93	0.452	0.345

The rank-one route would not have misranked these four: both put them in the same order. What it misstates is everything else. It values each 30 to 55 % lower, the order of the two two-year records being within the standard error of the simulated values, it expects the overhaul after almost every outcome where the simulated route expects it after one in four or five, and it says that two years of log establish the pump with probability 0.32, where the simulated route says 0.012. The ceiling on the value is the perfect-information figure of equation (30.18), R$ 3 006 a year. The five-year fleet row has outcomes with a reweighted N_{eff} as low as 8, a median of 9 239, and those few would need a refit.

What two more years would show. It depends on the population assumed. If pump 3 really does run at its own rate, 1.73 times the others', and two more years of the fleet come in exactly as those rates expect, the hierarchical tails become 0.77 and 0.15; after five years 0.79 and 0.14; after ten 0.84 and 0.14; after twenty 0.91 and 0.18. The first climbs slowly and the second falls, because seven alike pumps pull τ down as fast as pump 3's count pulls it up. That is a property of a normal population, which cannot hold one outlier among seven alike pumps, so on this drawing a record will not demonstrate pump 3 on a planning horizon. Taking pump 3 against the other seven directly, the exact gammas give 0.924 and 0.633 now, 0.949 and 0.656 after two years, and 0.998 and 0.776 after twenty: the pump is soon shown to be above the others, and still not half again above them with assurance. An inspection of pump 3 is a row on a different quantity.

The share. The entry ξ holds 90 % of the variance of $\ln U_{\mathrm{pair}}$, and Section 15.6.5 found that twenty more years of log cannot tell the present share from its half, so it is the least known quantity that matters to the campaign and the pair. It is not the least known in the state: the record took ξ from a prior spread of 1.33 to 0.75 (0.84), and $\ln \tau$ only from 1.11 to 1.08, and that entry touches neither question. And for the staggering decision, at the fraction of one half, knowing the share exactly would be worth only R$ 159 (R$ 270) a year by equation (30.18): the width that dominates the band hardly changes the action.

33.6.2 The recommendation

Do not wait on a record for the impeller rule or the staggering. Keep the log, with a time stamp on every failure, since it costs next to nothing and it is the record that bears on all three open quantities: two years of it are worth about R$ 880 a year in the overhaul decision alone, its

failures would narrow v, whose ceiling is about R$ 1 350 (R$ 900) a year, and its coincidences are the first row this fleet would have on its own share. Inspect pump 3 at its next repair; the verdict against paying for an overhaul at this price is the chain's, and it reverses on the separate drawing's value. Stagger the pair, under the condition stated, and read the doubles in the log for whether the share fell. When rows arrive, the fit returns to step 3.

33.6.3 What the analysis showed

The first reading named a bad pump, a wear-out and a rule, and it cleared the pair. The analysis finds no pump shown to be bad, no wear-out shown, no rule that pays, and a pair whose risk is almost wholly a cause the first reading left out. The two surprises are of different kinds. That partial pooling takes pump 3 from a probability of 0.90 to 0.75 and the half-again tail from 0.53 to 0.18, within that drawing, is a fact about eight small counts and largely about the prior on τ. That the wear-out disappears is a fact about bookkeeping, since the running pumps were in the log all along, and about this record, whose end spells are an unusually long constructed residual, 104 229 h against 68 600 h expected at a constant rate ($P = 0.05$). For the plant, this means spending little: keep the impellers until they fail, look at pump 3 before paying for it, and try the staggering, which is cheap against what a shared cause costs. The engineer who takes this back should expect the next two years of log to confirm a fleet of similar pumps at about one failure a year each, and should not expect that log to settle pump 3.

33.6.4 Limits of the model

The derivations are exact under their assumptions, and six of them need checking before a result is used.

The record's construction. The record is illustrative. The spells were matched to a mean and a median, the hours fixed the end spells, and the conclusions on wear-out and the age rule rest on those end spells, of mean 13 029 h. Those end spells total 104 229 h, against 68 600 h expected at a constant rate given the counts ($P = 0.05$) and 62 600 h at the fit ($P = 0.04$). A record whose running pumps are younger would raise the shape.

A share from another plant. The share is that of a chilled-water pair of another chapter, carried without its correlation with that pair's rate. This log aligns no clocks, so its own coincidences were not read; a time-stamped log would add a count row on ξ, and it could contradict the import.

The fraction of the share that staggering removes. No row informs it. Section 33.5.2 gives the break-even instead of a value.

The double reading of the failures. The counts and the spells are one record, read in two blocks. Their agreement is checked through 1/MTTF, and they are never multiplied, but a joint renewal state would carry the correlation between ω and (u, v).

A repair restores a pump. Pump 3's short spells may be repairs that leave the cause behind, which would be an imperfect renewal and not a high rate; the record cannot tell them apart, and an inspection could.

The other routes that happen to exist. On this plant the deviations are independent given (ω, τ), so nested quadrature reaches the same posterior as the chain, and eight weak count rows leave reweighting from the priors a healthy effective sample. Both serve as checks and both are specific to this plant: shared deviations remove the first, and stronger rows the second. Whether a chain would then do better is not shown here; the chain is the reported route on this record and the other two its checks.

The next chapter takes the method to a redundant cooling system.

33.7 What to retain

Definition 33.1 (The hierarchical drawing and its chain)**.** A fleet of rates drawn from a population has $\ln \lambda_i = \omega + \tau\delta_i$ with $\delta_i \sim \mathcal{N}(0,1)$ and a proper prior on τ; at $\tau = 0$ it is the pooled drawing, so their Bayes factor is $p(\tau = 0 \mid \mathbf{y})/\pi_\tau(0)$. The Metropolis chain on $\mathbf{z}$ with a symmetric proposal accepts with probability $\min(1, \exp[\mathcal{L}(\mathbf{z}) - \mathcal{L}(\mathbf{z}')])$ and has $\mathrm{e}^{-\mathcal{L}}$ as its stationary distribution, without its normalizing constant.

Four conclusions carry beyond the fleet. A method run end to end changes answers that each step alone would pass: the dispersion test, the fit and the plug-in rule each looked sound before the running pumps and the shared cause were put back. A route is chosen by what the posterior looks like, and here one record used three: a normal for forty-three events, the exact forms beside it, and a chain for a spread that the record cannot separate from zero, at the cost of a separate evidence and a simulated preposterior. A chain is checked against something other than itself where that exists, since a chain that never enters a region reports nothing about it. And the width that dominates a band need not be what moves a decision: the share holds nine tenths of the pair's variance and is worth little to the staggering. Table 33.17 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 33.17: The terms of the fleet's report and where each is defined.

Term	Expression	Units	Meaning	Defined in
Hierarchical drawing	$\ln \lambda_i = \omega + \tau\delta_i$	–	a population of rates	33.1.3
Spread	τ, half-normal prior	–	population's log-spread	33.1.3
Non-centered entries	$(\omega, \ln \tau, \delta_1, \ldots, \delta_8)$	–	deviations standardized	33.1.3
Metropolis acceptance	$\min(1, \mathrm{e}^{\mathcal{L}(\mathbf{z}) - \mathcal{L}(\mathbf{z}')})$	–	needs no constant	33.2.4
Potential scale reduction	$\hat{R}$	–	between against within chains	33.4.1
Effective sample of a chain	N_{ess}	–	independent equivalent	33.4.1
General forms				
Entry and evidence	$p(z_i \mid \mathbf{y})$; $p(\mathbf{y} \mid M_1)/p(\mathbf{y} \mid M_0)$	–	estimation, comparison	33.3.1
Tail	$\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y})$	–	pump 3	33.3.2
Predictive	$\int p(y_{\mathrm{new}} \mid \mathbf{z}) p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	–	campaign	33.3.4
Action	$\arg\min_a \mathcal{R}(a)$	–	age rule	33.3.6
Specialized forms and answers				
Savage–Dickey from draws	equation (33.7) over $\pi_\tau(0)$	–	pooled against hierarchical	33.3.1
Pair's mission	$R_{\mathrm{pair}}(t; \mathbf{z})$	–	β-factor pair	33.3.4
Overhaul saving	$M = C_{\mathrm{f}} t_{\mathrm{yr}} (\lambda_3 - \mathrm{e}^{\omega})$	R$/yr	against c_{o}	33.1.3
Break-even of staggering	$c_{\mathrm{st}}/(C_{\mathrm{L}} t_{\mathrm{yr}} \mathbb{E}[U_{\mathrm{cc}}])$	–	fraction of the share	33.5.2

33.8 Exercises

Exercise 33.1. *Formulation.* Five pumps run 20 000 h each with six failures in all, under the prior Gamma(2, 20 000 h). (a) Write the pooled posterior, its mean and its band. (b) Compute one pump's and a one-of-two pair's predictive reliability over 72 h without common cause. ((a) Gamma(8, 120 000 h), mean $6.67 \times 10^{-5}\,\mathrm{h}^{-1}$, band $[2.88,\ 12.0] \times 10^{-5}\,\mathrm{h}^{-1}$. (b) $(b/(b+72))^8 = 0.99521$; $2(b/(b+72))^8 - (b/(b+144))^8 = 0.999974$.)

Exercise 33.2. *Derivation.* Derive the distribution of the dispersion statistic $X^2 = \sum_i (k_i - \hat{\mu}_i)^2/\hat{\mu}_i$ under one shared rate, and say why the exact reference needs no rate. (Given the total K, the counts are multinomial with probabilities h_i/H, free of λ, by dividing the product of Poisson probabilities by the Poisson probability of K. With $\hat{\mu}_i = Kh_i/H$ the multinomial's Pearson statistic tends to chi-square with $I-1$ degrees of freedom, one taken by the fitted rate; for small counts the multinomial is enumerated or drawn.)

Exercise 33.3. *Derivation.* Derive the evidence of one shared rate against I separate rates under the prior Gamma(a, b), and evaluate it for two pumps with two and eight failures in 10 000 h each, $a = 2$, $b = 20\,000$ h. (With the common factor $\prod_i h_i^{k_i}/k_i!$ cancelled, $\ln B = \ln \frac{b^a \Gamma(a+K)}{\Gamma(a)(b+H)^{a+K}} - \sum_i \ln \frac{b^a \Gamma(a+k_i)}{\Gamma(a)(b+h_i)^{a+k_i}}$. Here $\ln B = 0.27$, while $X^2 = 3.6$ has $p = 0.058$: the test nearly rejects where the evidence barely leans the other way.)

Exercise 33.4. *Derivation.* Derive the pair's reliability (33.11) with a β-factor, and evaluate it at $\lambda = 6.67 \times 10^{-5}\,\mathrm{h}^{-1}$, $\beta_{\mathrm{cc}} = 0.15$, $t = 72$ h. (Thinning makes three independent processes, two at $(1-\beta_{\mathrm{cc}})\lambda$ and one at $\beta_{\mathrm{cc}}\lambda$. The pair survives if the shared one has no event and at least one other pump lives: $\mathrm{e}^{-\beta_{\mathrm{cc}}\lambda t}(1 - (1 - \mathrm{e}^{-(1-\beta_{\mathrm{cc}})\lambda t})^2)$, which is the form given. Value 0.99926.)

Exercise 33.5. *Application.* A ninth pump of a different make joins the fleet. (a) State the hierarchical question for its rate and give its posterior predictive from the fleet's chain. (b) If it fails four times in 10 000 h, compute the evidence of one rate for nine pumps' pooled record against two rates, under the chapter's prior. ((a) $\ln \lambda_9 = \omega + \tau\delta_9$ with δ_9 a new standard normal, one draw per draw of the chain: mean $1.122 \times 10^{-4}\,\mathrm{h}^{-1}$, median $1.068 \times 10^{-4}\,\mathrm{h}^{-1}$, band $[0.495,\ 2.098] \times 10^{-4}\,\mathrm{h}^{-1}$, wider than any pump's with a record. A make that is known to differ does not belong to this population, which the question must state. (b) $\ln B = -0.26$ with one rate on 45 failures in 382 000 h against a rate for the 41 and a rate for the 4: four failures do not settle it.)

Exercise 33.6. *Application.* The campaign is extended to a week. Compute one pump, the pair without the common cause and the pair with it, on the state route with the exact value beside each, and say which term grows faster. (One pump 0.98153 (0.98174); the pair without the cause 0.999651 (0.999659); with it 0.996859 (0.997327). The independent part grows as t^2, by about 5.4 here, the shared part about as t, by 2.4, so the cause's factor falls from 20 (17) to about 9 (8).)

Chapter 34

Redundant Cooling System

34.1 Three chillers, two pumps and a request for a fourth chiller

The designer of a cooling plant has a request on the desk: a fourth chiller, at R$ 180 000 a year, annualized and illustrative. The plant runs three chillers of which two are needed, two chilled-water pumps of which one, and one controller. Five years of log show eighteen chiller failures in 131 400 chiller-hours, fifteen alone and three in one simultaneous outage, against one finding of the controller down in 1 000 checks. Eighteen against one is the whole of the case for the chiller.

The quickest reading supports it. Eighteen failures in those hours are $1.37 \times 10^{-4}\,\mathrm{h}^{-1}$; sixteen repairs totaling 290 hours are 18.1 h each; a chiller is down 2.48×10^{-3} of the time, and three independent chillers needing two are down 1.84×10^{-5} of the time. Four chillers needing two are down 6.07×10^{-8}, which is 303 times less. That product leaves out the one event in the log that three independent machines would hardly produce, and it says nothing of what else can stop the plant. The question of this chapter is

> *how available the cooling is and how well that is known, which combinations take it down, which component matters, how soon the plant is back after a double outage, and whether a fourth chiller is worth buying against the alternatives, on one posterior of the plant's state.*

The chapter runs the method from step 0 to step 7 as a report and cites the earlier derivations rather than repeating them. What is new is only that the results of Chapters 13 to 16, 25, 26 and 32 now act on one plant at once, so that a term which looks large in one chapter's plant meets the others'.

Jargon: common-cause basic event, one repair job, separate repair jobs, program

A *common-cause basic event* is the shared cause of a group, carried in the fault tree as an event of its own beside the units' independent failures, so that the events of the tree are independent given the state. A common-cause outage is restored by *one repair job* when the crew clears the cause and the group returns together, and by *separate repair jobs* when the crew returns the units one at a time. A *program* is one combination of the changes on offer, priced as a whole.

Key idea

A redundant group's own redundancy stops paying once a shared cause sets its unavailability, and a component with no redundancy at all can carry more than the group. The questions of importance and of investment are therefore read from the whole tree. Each change is priced against the term it can reach, on one posterior, with the break-even of every price and of every assumption that no record informs.

34.1.1 The plant and its graph (step 0)

Figure 34.1 draws the plant. The gray boxes are pinned: the return and supply headers. The cooling path runs through three stages in series. The three green boxes C1 to C3 are the chillers, of which two must run; the green boxes P1 and P2 the pumps, of which one must run; and the green box K the controller, alone in its stage. The two boxes with red borders are the shared causes, one for the chillers, such as a chilled-water quality event or a control fault that trips all three, and one for the pumps, as in Chapter 15. The blue box is the one crew that repairs the chillers. Below are the six entries of the state, in green, and the rows that reach them.

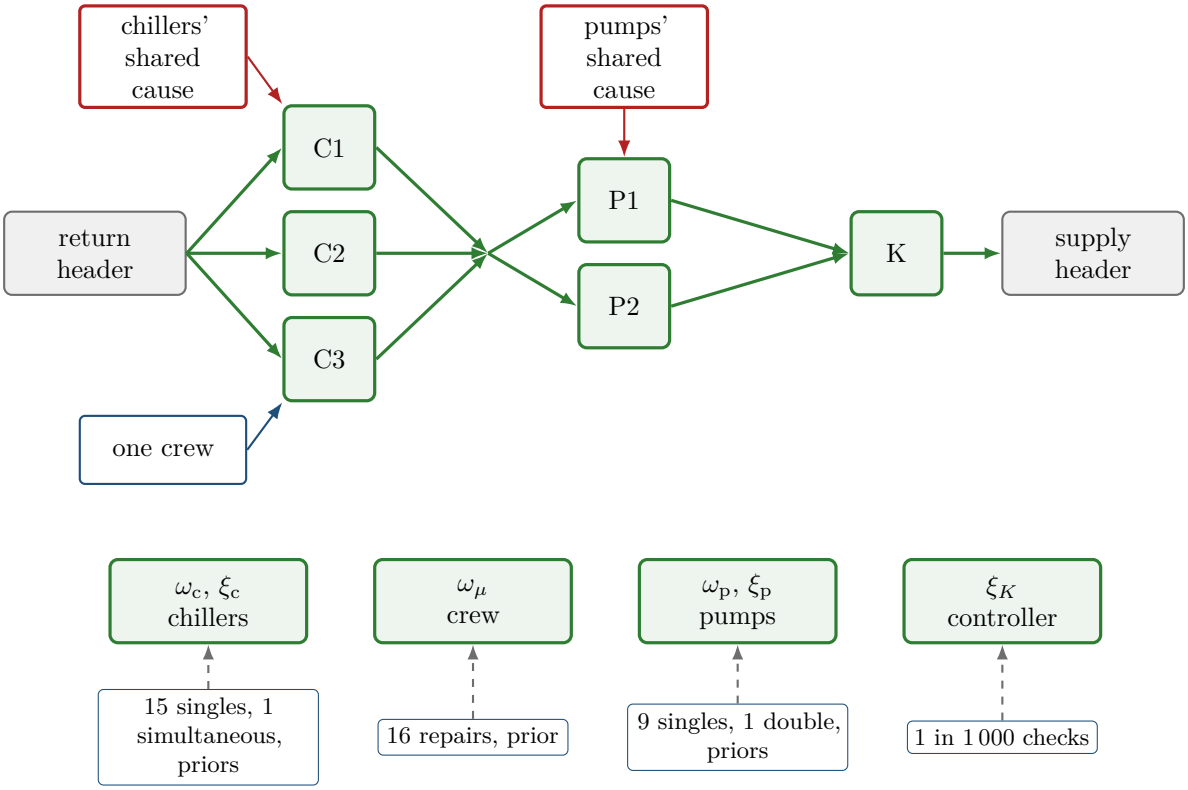


Figure 34.1: The cooling plant. Above, the plant: pinned, the two headers (gray); free, three chillers, two pumps and the controller in three stages in series (green), with a shared cause on each group (red) and the crew of the chillers (blue); each single arrow from a cause or from the crew stands for one to every unit of the group. Below, the entries of the state (green) and the rows that observe them (dashed).

Table 34.1 is the record, which is illustrative. Each row is an item, with its value and the row of the graph it feeds. No record here is drawn from a generator: the counts are stated, not simulated. The pumps' record and posterior are those of Chapter 15, unchanged, with its repair rate pinned at $1/12\,\mathrm{h}^{-1}$. The controller's is the record of Chapter 14.

34.1.2 The probabilistic questions (step 1)

Table 34.2 lists the six questions in the order they are treated. The first column is the question in engineering language; the second the general form of its category (Chapter 2); the third its form for the plant; the fourth its category; the last the subsection that answers it. Here U_{sys} is the plant's top event, the fraction of the time the cooling is lost, and a a change to the plant. The designer also carries an illustrative requirement, a target on the unavailability: cooling lost less than 10^{-4} of the time, the limit Chapter 15 set for the pair, here asked of the whole plant. The last question tests it.

Table 34.1: The records and the priors of the plant (illustrative).

Item	Value	Row of the graph
Chillers, record length	$T_{\mathrm{tot}} = 43\,800\,\mathrm{h}$	exposure of the count rows
Single outages, any chiller	$k_1 = 15$	mean $3(1-\beta_{\mathrm{c}})\lambda_{\mathrm{c}}T_{\mathrm{tot}}$
Simultaneous outages, window 4 h, all three	$k_2 = 1$	mean $\beta_{\mathrm{c}}\lambda_{\mathrm{c}}T_{\mathrm{tot}}$
Repairs of the chillers	16 jobs, 290 h	repair row on ω_μ
Pumps, Chapter 15	9 singles, 1 double, 43 800 h; $\mu = 1/12\,\mathrm{h}^{-1}$, pinned	two count rows
Controller, Chapter 14	1 down in 1 000 checks	binomial row on ξ_K
Priors on rates and on shares	Gamma(2, 10 000 h), Beta(1, 9)	prior rows, as in Chapter 15
Prior on the repair rate	Gamma(2, 20 h)	prior row on ω_μ

Table 34.2: The questions asked of the plant, the general form of their category, their form for the plant, their category and where each is answered. $\mathbf{y}$ is the record of Table 34.1.

Engineering question	General form	Mathematical form	Category	Answered in
What share of the chillers' failures is a common cause?	$p(z_i \mid \mathbf{y})$, quantiles	$p(\beta_{\mathrm{c}} \mid \mathbf{y})$ from the singles and the simultaneous outage	estimation	34.3.1
How often is cooling lost, and how well is that known?	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$, quantiles	$\mathbb{E}[U_{\mathrm{sys}} \mid \mathbf{y}]$, its band; hours per year	estimation of a derived quantity	34.3.2
Which combinations take cooling down, and which dominate?	$\mathbb{E}[P_c \mid \mathbf{y}]$ over the minimal cut sets	the five kinds of cut set and their expected probabilities	the enumeration read off	34.3.3
Which component matters most?	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$, quantiles	Birnbaum and Fussell–Vesely measures with bands	estimation (importance)	34.3.4
After a double chiller outage, how soon is the plant back?	$\int \Pr(A \mid \mathbf{z}) p(\mathbf{z} \mid \mathbf{y})\,\mathrm{d}\mathbf{z}$	the Markov transient with one crew	prediction	34.3.5
Is a fourth chiller worth buying, against the alternatives?	$\mathcal{R}(a') - \mathcal{R}(a)$; $\Pr(f(\mathbf{z}) \ge f^* \mid \mathbf{y})$	expected loss removed against annualized cost, for three options and their programs; $\Pr(U_{\mathrm{sys}} < 10^{-4} \mid \mathbf{y})$	design (investment)	34.3.6

34.1.3 The evidence as rows (step 2)

The modeling choices. Five choices make the rows and the top event. The simultaneous outage is read as a shared cause that trips all three chillers, the β-factor model of Chapter 15 for a group of three: this gives its row the mean $\beta_{\mathrm{c}}\lambda_{\mathrm{c}}T_{\mathrm{tot}}$, and it fails if a shared cause can trip only two, which matters to the fourth chiller and which Section 34.3.6 tests. A common-cause outage of the chillers is restored by one repair job, and no repair waits for another: the record logs 16 repairs for 16 outages, read as jobs, and these two together let each cut set keep the two-state form of Section 8.3.3. The second is not the plant's, which has one crew; it halves the chillers' pairs, a small term that is nevertheless all a fourth chiller reaches. Section 34.3.6 therefore prices the plant's crew beside the two-state form, and Section 34.4.3 tests separate jobs, under which the same log holds 18 jobs. The stages and their causes are independent given the state: this makes the top event a product, and it fails if a control fault that trips the chillers is the controller's own, an overlap that Section 34.6.5 takes up. The controller's record alone sets its posterior, Beta(1, 999), Chapter 14's choice: Section 34.4.5 varies it. And the prices are illustrative.

The rows of the state. The state has six entries,

$$\mathbf{z} = (\omega_{\mathrm{c}},\ \xi_{\mathrm{c}},\ \omega_{\mu},\ \omega_{\mathrm{p}},\ \xi_{\mathrm{p}},\ \xi_{K})^{\mathsf{T}}, \tag{34.1}$$

with $\omega_{\mathrm{c}} = \ln\lambda_{\mathrm{c}}$ the log total rate of one chiller, ξ_{c} the logit of the share β_{c}, $\omega_{\mu} = \ln\mu$ the chillers' log repair rate, ω_{p} and ξ_{p} the same two for a pump and ξ_{K} the logit of the controller's unavailability q_{K}. The chillers' count rows are those of equation (15.10) with three units for two: the singles' expected count is $\bar{k}_1 = 3(1-\beta_{\mathrm{c}})\lambda_{\mathrm{c}}T_{\mathrm{tot}}$ and the doubles' row becomes $\bar{k}_2 = \beta_{\mathrm{c}}\lambda_{\mathrm{c}}T_{\mathrm{tot}}$. The repair row is the gamma sum of equation (8.16), the pumps' rows carry Chapter 15's posterior and the controller's row is the binomial row of equation (14.5). Each row acts on one block:

$$\mathbf{J} = \begin{bmatrix} * & 0 & 0 & 0 & 0 & 0 \\ * & * & 0 & 0 & 0 & 0 \\ * & * & 0 & 0 & 0 & 0 \\ 0 & * & 0 & 0 & 0 & 0 \\ 0 & 0 & * & 0 & 0 & 0 \\ 0 & 0 & 0 & * & * & 0 \\ 0 & 0 & 0 & 0 & 0 & * \end{bmatrix} \begin{matrix} \text{prior on } \lambda_{\mathrm{c}} \\ \text{singles} \\ \text{simultaneous} \\ \text{prior on } \beta_{\mathrm{c}} \\ \text{repairs, prior} \\ \text{pump rows, priors} \\ \text{controller checks} \end{matrix}, \tag{34.2}$$

so that the covariance is block diagonal. The information matrix of these normalized rows is $\mathbf{J}^{\mathsf{T}}\mathbf{J}$; with unnormalized rows it would be $\mathbf{J}^{\mathsf{T}}\mathbf{W}\mathbf{J}$.

The top event and the prices. The two-state forms of equation (15.41), with the chillers' $(\lambda_{\mathrm{c}}, \beta_{\mathrm{c}}, \mu)$, give $U_{\mathrm{cc,c}}$ and q_{c}. Equation (13.19) with a shared cause gives the chillers' term, equation (15.42) the pumps', and the three stages in series the top event:

$$U_{\mathrm{C}} = 1 - (1-U_{\mathrm{cc,c}})(1-3q_{\mathrm{c}}^2+2q_{\mathrm{c}}^3), \qquad U_{\mathrm{sys}}(\mathbf{z}) = 1-(1-q_K)(1-U_{\mathrm{C}})(1-U_{\mathrm{pair}}). \tag{34.3}$$

The loss is $L(a,\mathbf{z}) = c_a + C_{\mathrm{L}}t_{\mathrm{yr}}U_{\mathrm{sys}}(\mathbf{z};a)$, equation (25.12) with one consequence, the loss rate $C_{\mathrm{L}} = \mathrm{R\$}\,25\,000\,\mathrm{h}^{-1}$, so that one unit of unavailability costs $C_{\mathrm{L}}t_{\mathrm{yr}} = \mathrm{R\$}\,219$ million a year. Table 34.3 lists the three changes; each row is a change, with its annual price and what it does to equation (34.3). The prices and the loss rate are illustrative, and Section 34.3.6 gives the break-even of each.

34.2 The posterior of the state (step 3)

Every answer is read from the posterior computed here. By pattern (34.2) the blocks fit apart, and only the chillers' two-entry block needs an iteration.

Table 34.3: The changes on offer (illustrative prices, annualized).

Change	Price [R$/yr]	Effect on the top event
A fourth chiller, two needed	180 000	$3q_{\mathrm{c}}^2 - 2q_{\mathrm{c}}^3$ becomes $4q_{\mathrm{c}}^3 - 3q_{\mathrm{c}}^4$; the shared cause trips all four
Staggered maintenance	30 000	β_{c} becomes $\beta_{\mathrm{c}}/2$, with λ_{c} held; the half is assumed
A redundant controller	40 000	q_K becomes q_K^2, independent given q_K

34.2.1 The posterior from the rows

The chillers' block has the objective (15.9) with the total event rate $(3 - 2\beta_{\mathrm{c}})\lambda_{\mathrm{c}}$, the shape 18, and the share's exponents 2 and 24. Table 34.4 lists the Newton iteration (15.17) from the prior means. Each row is an iterate; the columns give the two entries, the gradient and the Hessian.

Table 34.4: Newton iteration of the chillers' block.

Iteration	ω_{c}	ξ_{c}	$\partial\mathcal{L}/\partial\omega_{\mathrm{c}}$	$\partial\mathcal{L}/\partial\xi_{\mathrm{c}}$	H_{11}	H_{12}	H_{22}
0	−8.51719	−2.19722	8.528	−0.977	26.528	−1.577	1.079
1	−8.81030	−1.72009	1.111	0.265	19.111	−1.683	2.177
2	−8.88417	−1.89900	1.2×10^{-2}	1.1×10^{-2}	18.012	−1.375	1.928
3	−8.88536	−1.90552	-2.0×10^{-5}	2.8×10^{-5}	18.000	−1.367	1.918
4	−8.88536	−1.90553	-1.1×10^{-10}	1.8×10^{-10}	18.000	−1.367	1.918

At the minimum the four rows expect (1.384, 15.831, 0.785) counts and 1.295 pseudo-events, and in the working form of equation (15.10)

$$\mathbf{J} = \begin{bmatrix} 1.1764 & 0 \\ 3.9788 & -0.5152 \\ 0.8860 & 0.7712 \\ 0 & 1.0617 \end{bmatrix}, \quad \mathbf{J}^{\mathsf{T}}\mathbf{J} = \begin{bmatrix} 18.000 & -1.367 \\ -1.367 & 1.987 \end{bmatrix}, \quad \mathbf{C} = \begin{bmatrix} 0 & 0 \\ 0 & -0.069 \end{bmatrix}, \tag{34.4}$$

the curvature term of equation (15.16). The repair rows give the shape 18 and the rate 310 h, and $\Sigma = 1/18$; the controller's row gives $q_K = 10^{-3}$ and the variance $1/(1\,000\, q_K(1 - q_K))$. The whole state, with the pump block of equation (15.20), is

$$\hat{\mathbf{z}} = \begin{bmatrix} \hat{\omega}_{\mathrm{c}} \\ \hat{\xi}_{\mathrm{c}} \\ \hat{\omega}_{\mu} \\ \hat{\omega}_{\mathrm{p}} \\ \hat{\xi}_{\mathrm{p}} \\ \hat{\xi}_{K} \end{bmatrix} = \begin{bmatrix} -8.88536 \\ -1.90553 \\ -2.84620 \\ -8.94217 \\ -1.87447 \\ -6.90675 \end{bmatrix}, \quad \mathbf{\Sigma} = \begin{bmatrix} 0.05873 & 0.04185 & 0 & 0 & 0 & 0 \\ 0.04185 & 0.55120 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0.05556 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.08503 & 0.03083 & 0 \\ 0 & 0 & 0 & 0.03083 & 0.56008 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1.00100 \end{bmatrix}, \tag{34.5}$$

that is $\hat{\lambda}_{\mathrm{c}} = 1.3840 \times 10^{-4}\,\mathrm{h}^{-1}$, $\hat{\beta}_{\mathrm{c}} = 0.1295$, $\hat{\mu} = 0.05806\,\mathrm{h}^{-1}$, $\hat{\lambda}_{\mathrm{p}} = 1.3076 \times 10^{-4}\,\mathrm{h}^{-1}$, $\hat{\beta}_{\mathrm{p}} = 0.1330$ and $\hat{q}_K = 10^{-3}$, with the spreads $\sqrt{\Sigma_{ii}} = (0.2423,\ 0.7424,\ 0.2357,\ 0.2916,\ 0.7484,\ 1.0005)$. The correlation in the chillers' block is 0.233.

34.2.2 The exact posterior and the routes

Beside the state route stand the results that Chapters 14 and 15 derive: λ_{c} given the share is Gamma(18, $10\,000\,\mathrm{h} + (3 - 2\beta_{\mathrm{c}})\,43\,800\,\mathrm{h}$), as in equation (15.12); the marginal of the share is proportional to $\beta_{\mathrm{c}}(1 - \beta_{\mathrm{c}})^{23}(141\,400 - 87\,600\,\beta_{\mathrm{c}})^{-18}$, as in equation (15.13); $\mu \sim$ Gamma(18, 310 h); the pumps' are equation (15.13) itself; and $q_K \sim$ Beta(1, 999). Every answer below is reported on the state route, with the exact value in parentheses and the gap

stated. Expectations on the state route use the Gauss–Hermite rule on the axes of $\mathbf{\Sigma}$, block by block, since a product of independent blocks has the product of their expectations; those on the exact route use quadrature in the same blocks. Bands and tails use 2×10^5 draws on each route from one generator (numpy `default_rng`, seed 34), and the ranges printed for tails are over seeds 34 to 38.

34.3 The answers (step 4)

The questions of Table 34.2 are answered from the posterior of Section 34.2. The derivations are those of the chapters cited, and only the plant's substitutions are shown.

34.3.1 The chillers' common-cause share

Engineering question. What share of the chillers' failures is a common cause?

General form. The estimation form of an entry, equation (2.9):

$$p(z_i \mid \mathbf{y}) = \int p(\mathbf{z} \mid \mathbf{y}) \, \mathrm{d}\mathbf{z}_{-i}. \tag{34.6}$$

Specialized form. The question is equation (34.6) with $z_i = \xi_\mathrm{c}$ and the normal of equation (34.5), mapped through $\beta_\mathrm{c} = 1/(1 + \mathrm{e}^{-\xi_\mathrm{c}})$. No new choice enters.

Building blocks. $\hat{\xi}_\mathrm{c}$ and Σ_{22} of equation (34.5); the exact marginal for the value beside it.

The numbers. Table 34.5 gives the entries that matter below, one quantity to a row; the columns give the median, the mean and the band on each route. The share has the median 0.1295 (0.1126) and the band [0.034, 0.390] ([0.018, 0.317]): one simultaneous outage leaves it known to a factor of about 12 (18) between the ends of the band. A chiller is down 0.25 % of the time on both routes, and the state route's mean repair is 17.7 h (18.2 h).

Table 34.5: The posteriors: median, mean and 95 % band on each route; the band from 2×10^5 draws, seed 34. The pumps' values are those of Chapter 15. $U_\mathrm{unit} = \lambda_\mathrm{c}/(\lambda_\mathrm{c} + \mu)$.

	State route			Exact route		
Quantity	Median	Mean	Band	Median	Mean	Band
λ_c $[10^{-4}\,\mathrm{h}^{-1}]$	1.384	1.425	[0.860, 2.231]	1.357	1.386	[0.813, 2.125]
β_c	0.1295	0.1512	[0.034, 0.390]	0.1126	0.1272	[0.018, 0.317]
μ $[\mathrm{h}^{-1}]$	0.0581	0.0597	[0.0365, 0.0922]	0.0570	0.0581	[0.0343, 0.0878]
U_unit $[10^{-3}]$	2.38	2.52	[1.23, 4.62]	2.38	2.52	[1.22, 4.64]
β_p	0.133	0.155	[0.034, 0.399]	0.116	0.131	[0.018, 0.329]
q_K $[10^{-3}]$	1.00	1.64	[0.14, 7.06]	0.69	1.00	[0.025, 3.70]

34.3.2 The top event and its band

Engineering question. How often is cooling lost, and how well is that known?

General form. The expectation of a derived quantity and its quantiles, equations (2.10) and (2.11):

$$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}] = \int f(\mathbf{z}) \, p(\mathbf{z} \mid \mathbf{y}) \, \mathrm{d}\mathbf{z}, \qquad \Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y}) = \int \mathbf{1}[f(\mathbf{z}) \le f^*] \, p(\mathbf{z} \mid \mathbf{y}) \, \mathrm{d}\mathbf{z}. \tag{34.7}$$

Specialized form. The question is equation (34.7) with $f = U_{\text{sys}}$ of equation (34.3) and the normal of equation (34.5). Since $\boldsymbol{\Sigma}$ is block diagonal, $\mathbb{E}[U_{\text{sys}}] = 1 - (1 - \mathbb{E}[q_K])(1 - \mathbb{E}[U_{\text{C}}])(1 - \mathbb{E}[U_{\text{pair}}])$. The width is read from the gradient of $\ln U_{\text{sys}}$ at $\hat{\mathbf{z}}$,

$$\mathbf{g} = (0.2179,\ 0.1730,\ -0.2179,\ 0.1385,\ 0.1176,\ 0.6520)^{\mathsf{T}}, \qquad \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} = 0.4610. \tag{34.8}$$

Bands are read from draws of the normal, since at a logit spread of one the delta band, $[4.0 \times 10^{-4},\ 5.8 \times 10^{-3}]$, sits low.

Building blocks. $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$ of equation (34.5); the top event (34.3); the rule on each block; $\mathbf{g}$.

The numbers. Table 34.6 gives them. Each row is a term; the columns give the expected value on each route, the gap, and the band on each route. Cooling is lost 2.31×10^{-3} (1.56×10^{-3}) of the time, 20.3 h (13.7 h) a year. The medians are 1.72×10^{-3} (1.29×10^{-3}), a gap of 4.3×10^{-4}, and the value at $\hat{\mathbf{z}}$ is 1.53×10^{-3}. The gap between the means, 48 % of the exact value, is 86 % q_K: a normal on the logit around one finding has the mean 1.64×10^{-3} where Beta(1, 999) has 10^{-3}. The chillers' independent part is 1.5×10^{-5} (1.6×10^{-5}), and their common cause is 96 % (95 %) of their term. In the variance (34.8) the entry ξ_K holds 0.4256, which is 92 %; with each cross term allocated by $g_i(\boldsymbol{\Sigma}\mathbf{g})_i$, ξ_{c} holds 3.9 % and ξ_{p} 1.8 %, and directly, by $g_i^2\Sigma_{ii}$, 3.6 % and 1.7 %. The band of the top event spans a factor of 14 (12), and it is the controller's single finding that makes it so.

Table 34.6: The top event and its terms: expected value on the state route and the exact route, the gap, and the 95 % bands of the value at the state (2×10^5 draws, seed 34).

Term [10^{-4}]	State	Exact	Gap	Band, state	Band, exact
Controller, q_K	16.44	10.00	6.44	[1.40, 70.6]	[0.25, 37.0]
Chillers, U_{C}	4.09	3.49	0.60	[0.71, 12.6]	[0.46, 10.7]
common cause, $U_{\text{cc,c}}$	3.94	3.33	0.61	[0.62, 12.4]	[0.35, 10.4]
independent part	0.152	0.161	−0.009		
Pumps, U_{pair} (Chapter 15)	2.62	2.14	0.49	[0.46, 7.8]	[0.25, 6.3]
Top event, U_{sys}	23.14	15.62	7.52	[5.6, 77.6]	[3.5, 43.2]
Hours a year	20.3	13.7	6.6	[4.9, 68]	[3.0, 38]

34.3.3 The minimal cut sets

Engineering question. Which combinations take cooling down, and which dominate?

General form. The minimal cut sets and the rare-event sum, equation (14.24), under the posterior as equation (34.7):

$$U_{\text{sys}} \approx \sum_c P_c(\mathbf{z}), \qquad \mathbb{E}[P_c \mid \mathbf{y}] = \int P_c(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}. \tag{34.9}$$

Specialized form. The question is equation (34.9) with the tree of Figure 34.2, whose common-cause events are basic events of their own, so that the enumeration over 2^8 configurations returns the cut sets with P_c a product of independent terms given the state, and with the normal of equation (34.5). Figure 34.2 draws the tree: an OR gate at the top, a 2/3 voting gate on the chillers' independent failures, an AND gate on the pumps', and circles for the eight basic events, with the expected probability of each cut set beneath, state route with the exact value in parentheses.

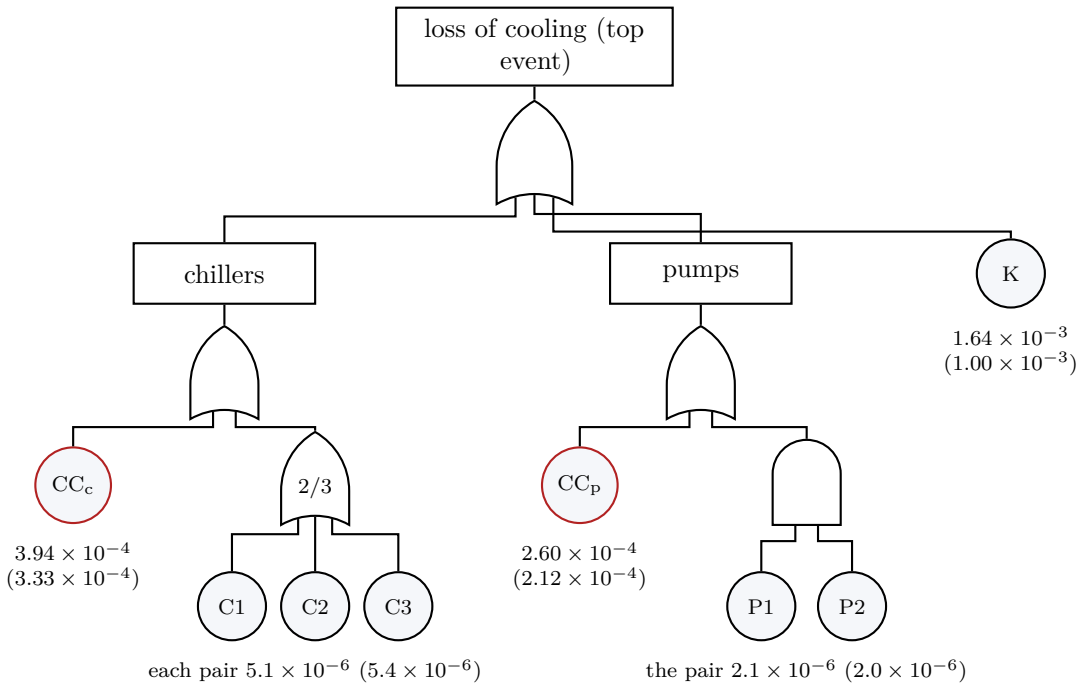


Figure 34.2: The fault tree of the loss of cooling, with the shared causes (red) as basic events and the expected probability of each cut set beneath it, state route with the exact value in parentheses. The vote gate holds the chillers' independent failures.

Building blocks. The enumeration; q_K, $U_{\mathrm{cc,c}}$, q_{c}, U_{cc} and q_{ind} at each node of the rule; $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$.

The numbers. Table 34.7 lists the cut sets. Each row is a kind of cut set; the columns give the value at $\hat{\mathbf{z}}$, the expected probability on each route, and the expected share of the sum on each route. Three cut sets of order one are nearly everything, and the pairs that the first reading priced are 0.7 % (1.2 %) of the top event with the pumps' pair among them. Figure 34.3 draws the shares. The horizontal axis is the share of the top event, each thick bar the expected share of a cut set on one route and the thin line the 95 % band of that share at the state. The bands overlap heavily: the controller's runs from 0.16 to 0.94 (0.04 to 0.92), so the expected share of 0.71 (0.64) is not a well-known one.

Table 34.7: The minimal cut sets: value at $\hat{\mathbf{z}}$, expected probability and expected share of the sum, by route. The pairs of chillers are three cut sets: each probability is for one, the share for the three.

Cut set	At $\hat{\mathbf{z}}$	$\mathbb{E}[P_c]$, state	$\mathbb{E}[P_c]$, exact	Share, state	Share, exact
{K}	1.000×10^{-3}	1.644×10^{-3}	1.000×10^{-3}	0.710	0.640
{CC$_{\mathrm{c}}$}	3.085×10^{-4}	3.941×10^{-4}	3.327×10^{-4}	0.170	0.213
{CC$_{\mathrm{p}}$}	2.087×10^{-4}	2.603×10^{-4}	2.116×10^{-4}	0.112	0.135
{C$_i$, C$_j$}, three	4.29×10^{-6}	5.06×10^{-6}	5.38×10^{-6}	0.007	0.010
{P1, P2}	1.85×10^{-6}	2.07×10^{-6}	2.00×10^{-6}	0.001	0.001
Sum	1.532×10^{-3}	2.315×10^{-3}	1.562×10^{-3}	1	1

34.3.4 The importance of each component

Engineering question. Which component matters most?

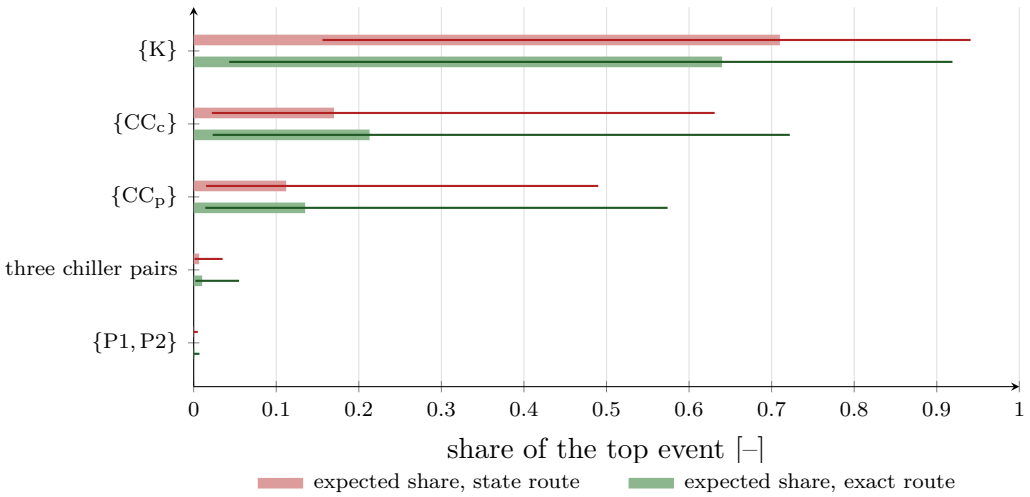


Figure 34.3: The top event's share by cut set: expected probability over the sum of expected probabilities (thick bars), and the 95 % band of the share P_c/U_{sys} at the state (thin lines), 2×10^5 draws, seed 34.

General form. The estimation form of Chapter 2, equation (2.11), for a derived quantity between its quantiles:

$$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y}) = \int \mathbf{1}[f(\mathbf{z}) \le f^*]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}. \tag{34.10}$$

Specialized form. The question is equation (34.10) with $f = I_{\text{B},i}$ of equation (26.7) and $f = I_{\text{FV},i}$ of equation (26.10), over the eight basic events of Figure 34.2, and with the normal of equation (34.5). One choice enters: the shared causes are basic events and not properties of the units, which keeps the events independent given the state, so that the pivotal decomposition (26.6) and the identity (26.11) hold, and the question that identity fails on, pinning a unit whose failures carry a shared cause, does not arise. By equation (34.3), every event of order one has

$$I_{\text{B},i} = \frac{1 - U_{\text{sys}}}{1 - x_i}, \qquad x_i \in \{q_K,\ U_{\text{cc,c}},\ U_{\text{cc}}\}, \tag{34.11}$$

and a chiller's own failure has $I_{\text{B}} = (1 - U_{\text{sys}})\,(2q_{\text{c}} - 2q_{\text{c}}^2)/(1 - 3q_{\text{c}}^2 + 2q_{\text{c}}^3)$.

Building blocks. $\hat{\mathbf{z}}$, the measures at each draw of the normal, and the exact draws beside them.

The numbers. Table 34.8 gives them. Each row is a basic event; the columns give each measure's median and band on the state route, then on the exact route. The controller is first on both measures, and the two say different things about that. By equation (34.11) the Birnbaum measure of every order-one event is within 10^{-3} of one in the median: the controller's 0.99941 against the chillers' cause 0.99871. That measure says that each of them, once down, stops the plant, and its order among them is only the order of their x_i. The Fussell–Vesely measure does the ranking: medians of 0.63, 0.19 and 0.12 (0.58, 0.22 and 0.14). The inequality $I_{\text{B},K} > I_{\text{B,CC}_c}$ holds exactly when $q_K > U_{\text{cc,c}}$, which has probability 0.83 (0.74). The controller is the largest of the three order-one events with probability 0.79 (0.70), a range of 0.791 to 0.794 (0.695 to 0.697) over the seeds, and its Fussell–Vesely measure exceeds one half with probability 0.68 (0.59). The chillers' importance is the common-cause share's: the shared cause has the median 0.19 where one chiller's own failure has 0.0047, 40 (32) times less.

Table 34.8: The importance measures of the basic events: median and band on each route, 2×10^5 draws, seed 34. The three chillers share one row, and so do the two pumps.

	Birnbaum, I_{B}				Fussell–Vesely, I_{FV}			
Event	State	Band	Exact	Band	State	Band	Exact	Band
K	0.99941	[0.99838, 0.99980]	0.99951	[0.99864, 0.99985]	0.630	[0.156, 0.941]	0.578	[0.043, 0.919]
$\mathrm{CC_c}$	0.99871	[0.99266, 0.99968]	0.99907	[0.99606, 0.99984]	0.186	[0.022, 0.630]	0.215	[0.023, 0.722]
$\mathrm{CC_p}$	0.99855	[0.99251, 0.99962]	0.99893	[0.99590, 0.99980]	0.122	[0.015, 0.490]	0.135	[0.014, 0.574]
C_i	0.00400	[0.00204, 0.00774]	0.00412	[0.00212, 0.00798]	0.0047	[0.0006, 0.0232]	0.0068	[0.0012, 0.0366]
P_j	0.00132	[0.00072, 0.00235]	0.00131	[0.00069, 0.00224]	0.0010	[0.0001, 0.0049]	0.0013	[0.0002, 0.0069]

34.3.5 The recovery after a double outage

Engineering question. After a double chiller outage, how soon is the plant back?

General form. The predictive of a future event, equation (2.13), as in equation (16.47):

$$\Pr(A \mid \mathbf{y}) = \int \Pr(A \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}. \tag{34.12}$$

Specialized form. The question is equation (34.12) with A the event that two or more chillers are up at t, $\Pr(A \mid \mathbf{z}) = p_3(t) + p_2(t)$ from the forward equation (16.11) of the four-state chain, and the normal on $(\omega_{\mathrm{c}}, \xi_{\mathrm{c}}, \omega_\mu)$ of equation (34.5). The states are named by the number of chillers up, with one crew: three up goes to two at $3(1-\beta_{\mathrm{c}})\lambda_{\mathrm{c}}$, two to one at $2(1-\beta_{\mathrm{c}})\lambda_{\mathrm{c}}$, one to none at λ_{c}, the shared cause goes from three or two to none at $\beta_{\mathrm{c}}\lambda_{\mathrm{c}}$, and each state below three climbs by one at μ. A double outage starts at one up. Two choices enter. The clock starts with the crew at work: a delay would shift every time by itself. And the outage the record actually holds, all three down, is a different start, computed beside it: with one job, as step 2 chose, it needs one repair and follows the first-passage curve $1 - \mathrm{e}^{-\mu t}$; with separate jobs it needs two, at the rate of Section 34.4.3. The chain climbs one unit at a time, also after a shared cause met during the recovery, where one job would return all three at once; p_0 stays near 10^{-3}, so this moves only the last digits.

Building blocks. $\hat{\omega}_{\mathrm{c}}$, $\hat{\xi}_{\mathrm{c}}$, $\hat{\omega}_\mu$ and their block; the generator and its matrix exponential at each node.

The numbers. Table 34.9 gives the transient, one time to a row; the columns give the expected state probabilities from one up, the tail on the two routes with its band, and the tail from all three down under each repair reading. Figure 34.4 draws it: the horizontal axis is the time since the outage, the vertical a probability, the solid curves the states on the state route, the dashed green curve the tail on the exact route and the dotted curve the tail from all three down with separate jobs. The plant has an even chance of cooling again at 12 h and is not sure of it at two days. The gap is under 0.01 at each time. All of that rests on one repair and so on ω_μ, and a second failure during a repair hardly matters. The outage the record holds, all three down, restored by one job, is nearly the same, 0.505 (0.495) at 12 h. Only the repair reading moves it: with separate jobs it is 0.19 (0.18).

34.3.6 The fourth chiller and the alternatives

Engineering question. Is a fourth chiller worth buying, against the alternatives?

Table 34.9: After an outage, one crew: expected state probabilities from one up on the state route (p_0 stays near 10^{-3}); the probability that two or more are up on each route, with the band of the value at the state; and the same from all three down, with one job or with separate jobs.

	One up, state			One up, Pr(≥ 2 up)			All down, one job		All down, separate	
t [h]	p_3	p_2	p_1	State	Exact	Band, state (exact)	State	Exact	State	Exact
6	0.052	0.246	0.701	0.298	0.292	[0.196, 0.424] ([0.187, 0.410])	0.299	0.292	0.062	0.060
12	0.163	0.341	0.495	0.504	0.495	[0.354, 0.667] ([0.338, 0.651])	0.505	0.495	0.190	0.183
24	0.412	0.335	0.253	0.746	0.737	[0.581, 0.889] ([0.561, 0.877])	0.748	0.739	0.463	0.450
48	0.752	0.176	0.071	0.928	0.923	[0.823, 0.987] ([0.805, 0.984])	0.930	0.925	0.802	0.791

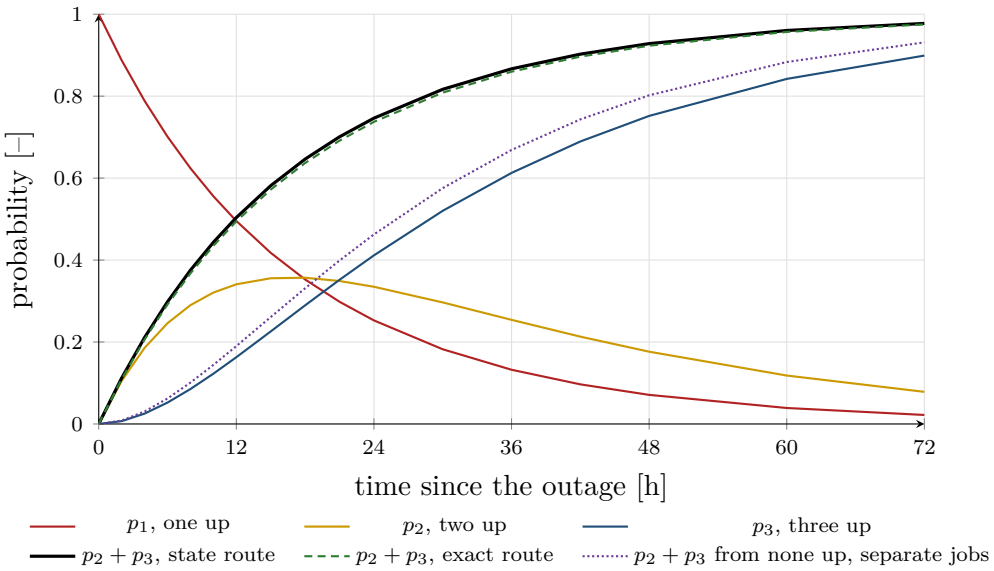


Figure 34.4: The chillers after an outage, expected probabilities with one crew. Solid: from one up on the state route. Dashed: the tail on the exact route. Dotted: the tail on the state route from all three down, with separate jobs at the rate Gamma(20, 310 h). The same start with one job lies within 0.002 of the heavy curve and is not drawn; nor is p_0, near 10^{-3}.

General form. The comparison of expected consequences with a tail of the difference, equation (32.16):

$$\mathcal{R}(a') - \mathcal{R}(a), \qquad \Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}) = \int \mathbf{1}[f(\mathbf{z}) \geq f^*]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}. \tag{34.13}$$

Specialized form. The question is equation (34.13) with a and a' programs of Table 34.3, L the loss of step 2, f the reduction $\Delta\mathcal{R}_{\mathrm{yr}}$ of equation (25.35) and f^* the annualized price, and with the normal of equation (34.5), block by block. Each program is a new top event (34.3) on the same posterior; the reduction of a later change is taken on the residual, as in Section 32.3.4. The redundant controller's cut set has the second moment of equation (14.42), and the two controllers share q_K and are independent given it: this values them without a record of their own, and it fails if they have a shared cause, tested below. The fractions and the jobs that no record informs are carried as stated conditions.

Building blocks. $\hat{\mathbf{z}}$ and $\mathbf{\Sigma}$; the moments $\mathbb{E}[q_K]$, $\mathbb{E}[q_K^2]$; $\mathbb{E}[U_{\mathrm{C}}]$ for each (n, β_{c} or $\beta_{\mathrm{c}}/2$) and each repair reading, by the rule; the draws for bands and tails.

The options. Table 34.10 gives them, one change to a row; the columns give the price and the expected reduction under each reading and route, one route to a column: the two-state form of step 2, the chain with the plant's one crew and one job, and the chain with one crew and separate jobs, at the refitted rate of Section 34.4.3. Figure 34.5 draws the two-state values: the horizontal axis is the annualized price, the vertical the reduction, both logarithmic; the line is break-even, and a point above it pays; each bar is the 95 % band on the state route. On the two-state form the fourth chiller takes the top event from 2.314 to 2.299×10^{-3} (1.562 to 1.546×10^{-3}), 0.65 % (1.03 %), since it reaches only the independent part, 1.5×10^{-5} (1.6×10^{-5}). One crew doubles that part, and with the plant's crew and one job the fourth chiller removes R$ 6 562 (R$ 6 977), 1.3 % (2.0 %). That doubling comes from the crew and not from the jobs: separate jobs give about the same order. Against R$ 180 000 it is a ratio of 0.036 (0.039), it exceeds its price in none of the draws, and the first reading's factor of 303 is the factor on a term worth R$ 4 010 a year at its own number. The redundant controller removes 71 % (64 %), R$ 358 179 (R$ 218 440) against R$ 40 000, a ratio of 8.95 (5.46). Staggering, if it halves the share, removes 8.4 % (10.5 %), R$ 42 428 (R$ 35 834) against R$ 30 000, a ratio of 1.41 (1.19), marginal on the exact route and on neither route certain.

Table 34.10: The expected annual loss removed [R$/yr], illustrative prices, under each reading of the repairs, one route to a column. Separate jobs at the rate Gamma(20, 310 h). Band: two-state form, state route, 2×10^5 draws, seed 34.

		Two-state		One crew, one job		One crew, separate		
Change	Price	State	Exact	State	Exact	State	Exact	Band
Fourth chiller	180 000	3 300	3 509	6 562	6 977	5 258	5 553	[684, 9 852]
Staggering	30 000	42 428	35 834	41 799	35 283	76 449	64 392	[6 760, 132 982]
Controller	40 000	358 179	218 440	358 174	218 436	358 063	218 380	[30 743, 1 535 319]
Staggering, after the controller	30 000	42 498	35 870	41 868	35 318	76 575	64 456	[6 772, 133 186]
Fourth, after both	180 000	3 928	4 058	7 808	8 065	6 262	6 423	[854, 11 633]

The programs. Table 34.11 prices them. Each row is a program; the columns give the price, then the expected price plus residual loss under each reading and route. Table 34.12 gives, on

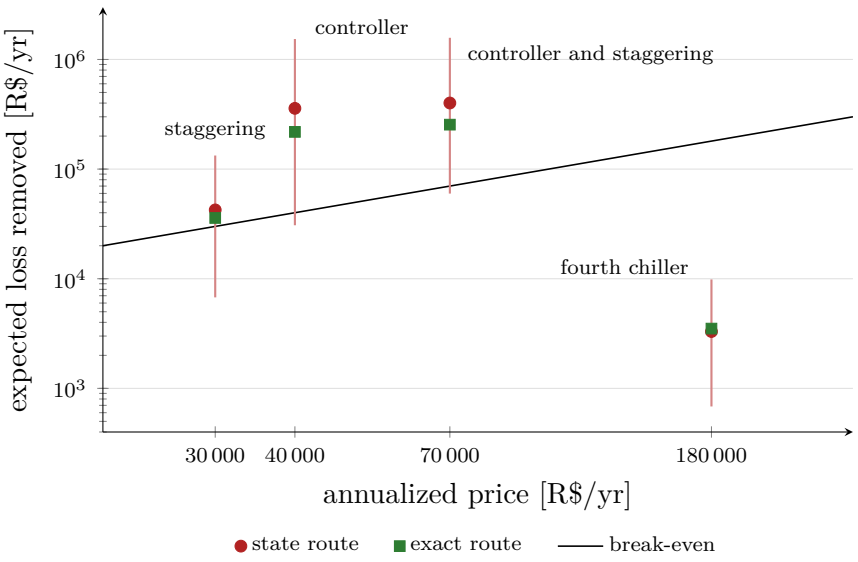


Figure 34.5: The expected loss removed against the annualized price, two-state form, illustrative prices (the plant's crew about doubles the fourth chiller's point and hardly moves the others): both routes, and the 95 % band on the state route (2×10^5 draws, seed 34).

the two-state form, the probabilities at the state that the program beats its price and that the top event is below the 10^{-4} of step 1. The least total on all six columns is the redundant controller with staggering. On the two-state columns it stands at R$ 176 135 (R$ 157 721), of which R$ 106 135 (R$ 87 721) is the residual loss, and it improves on the controller alone by R$ 12 497 (R$ 5 869). After the controller, staggering alone exceeds its price at the state with probability 0.55 (0.47), and that is what a step to staggering rests on, under its assumed half. On the two-state columns, adding the fourth chiller to any program adds about R$ 177 000 to its total.

Table 34.11: The programs [R$/yr], illustrative prices: price plus expected loss, one route to a column, under each reading of the repairs.

		Two-state		One crew, one job		One crew, separate	
Program	Price	State	Exact	State	Exact	State	Exact
None	0	506 812	342 030	510 133	345 562	577 724	401 968
Fourth	180 000	683 511	518 520	683 570	518 585	752 466	576 415
Staggering	30 000	494 384	336 196	498 333	340 280	531 275	367 576
Controller	40 000	188 632	163 590	191 959	167 126	259 661	223 588
Controller, staggering	70 000	176 135	157 721	180 091	161 808	213 086	189 132
Controller, fourth	220 000	365 326	340 077	365 386	340 142	434 394	398 030
All three	250 000	352 207	333 663	352 283	333 744	386 825	362 709

What would break the three verdicts. The fourth chiller is floored because a shared cause trips the fourth with the others, and that is a model. Suppose instead that every shared cause trips exactly two chillers, a drawing the record's one event does not exclude if it was mislogged. Two of three still lose the cooling, while two of four do not, and the fourth chiller then removes R$ 89 042 (R$ 75 944), nearly the whole chillers' term, still under its price; the term alone exceeds R$ 180 000 with probability 0.09 (0.06). At the model of the record it pays only if the loss rate is R$ 1 363 000 h^{-1} (R$ 1 282 000 h^{-1}) or more on the two-state form, R$ 686 000 h^{-1} (R$ 645 000 h^{-1}) with the plant's crew, or, on the two-state form, if a chiller's independent unavailability is 1.68×10^{-2}, eight times the present value. The verdict holds under

Table 34.12: The tails of the programs, two-state form: the loss removed exceeds the price at the state, and $\Pr(U_{\text{sys}} < 10^{-4} \mid \mathbf{y})$ with the range over seeds 34 to 38 where it moves in the digits shown.

	Removed > price		Below 10^{-4}	
Program	State	Exact	State	Exact
None			0	0.0001–0.0002
Fourth	0.000	0.000	0	0.0002–0.0003
Staggering	0.55	0.47	0	0.0004–0.0005
Controller	0.955	0.833	0.0005–0.0007	0.006
Controller, staggering	0.959	0.847	0.003–0.004	0.018–0.019
Controller, fourth	0.50	0.37	0.001	0.009–0.010
All three	0.53	0.39	0.005–0.006	0.027–0.028

these alternatives, and it is a statement about a plant that needs two chillers all year.

The redundant controller's verdict rests on one finding. It pays at the price if $q_K > 1.83{\times}10^{-4}$, and the probability of that is 0.955 (0.833); no requirement here sets an assurance, and the action is the same on both routes in expectation. It stays above its price up to a shared-cause share between the two controllers of 0.89 (0.82), and at the pumps' share of about 0.13 it removes R$ 311 795 (R$ 190 092). Its loss-rate break-even is R$ 2 792 h^{-1} (R$ 4 578 h^{-1}).

Staggering's verdict is the weak one. Its break-even loss rate is R$ 17 677 h^{-1} (R$ 20 930 h^{-1}), not far from the assumed one. Its price requires the staggering to remove 35 % (42 %) of the share under one job and 20 % (23 %) under separate jobs. Removing a quarter, it fails under one job, at R$ 21 221 (R$ 17 922), and passes under separate jobs, at R$ 38 225 (R$ 32 196). The half is an assumption. The causes the plant names, a water-quality event and a control fault, are not maintenance errors, and the log does not say what caused its one outage, so nothing here says the half is available.

The target. Under no priced program, for the target on the unavailability of step 1, does $\Pr(U_{\text{sys}} < 10^{-4} \mid \mathbf{y})$ come near an assurance. On the state route it is below 5×10^{-6} in the draws without the controller change and at most 0.006 with all three changes; on the exact route it runs from 0.0001 with none to 0.028 with all three. So the target is out of reach with the controller change as well as without it. At $\hat{\mathbf{z}}$ the residual after the controller is 5.3×10^{-4}, the two common causes nearly all of it, which must fall by a factor of 5.3.

34.4 The checks (step 5)

The answers rest on a rare-event sum, a normal at the mode, a reading of one repair and the priors. Each is tested here. There is no physics row, so no limit $\sigma_p \to 0$ arises.

34.4.1 The cut-set sum against the enumeration

At $\hat{\mathbf{z}}$ the 2^8 configurations give 244 that lose the cooling, summing to 1.53131×10^{-3}, equal to the product (34.3), against the rare-event sum 1.53193×10^{-3}, a gap of 0.04 %. The same in expectation: 2.3142 against 2.3154×10^{-3} (1.5618 against 1.5624×10^{-3}).

34.4.2 The normal at the mode beside the exact route

Table 34.5 sets the routes side by side. For β_{c} the median stands 15 % above the exact one and the mean 19 %; the upper end of the band is 0.390 against 0.317, the skew of a logit around one event, as in Section 15.4.1. For q_K the mean is 64 % high, and the state route's lower band end

is high by a factor of about 5.5, since Beta(1, 999) has its density's mode at zero while the logit normal has none near it. That gap carries the 48 % on the top event. No action changes between the routes in expectation; the controller's probability of paying is 0.955 (0.833), as Section 34.3.6 states.

34.4.3 The chain against the two-state cut sets

The transient of Section 34.3.5 tends to its stationary value at every start: at $\hat{\mathbf{z}}$, 2 000 h after any start, the chain has $\pi = (0.992869,\ 0.006487,\ 0.000335,\ 0.000309)$ for three to none up, the probability of cooling lost 6.446×10^{-4}, which equals the linear solve. That is not the two-state value 3.214×10^{-4} of equation (34.3). The difference is the two choices of step 2. In the four-state chain a shared-cause outage climbs one chiller at a time, so that it holds the cooling down for two repairs and not one, and one crew serves the pairs. Table 34.13 sets the readings together, one row to a reading; the columns give the value at $\hat{\mathbf{z}}$ and the expected value on each route. A chain with one job, a separate state for the outage that returns to three up at μ, gives 3.342×10^{-4}, 4 % above the two-state form, because one crew doubles the pairs' term, as the pumps' chain did in Chapter 33 by under one percent. That is the plant's own crew, and it is small against the term and not against the pairs, which is why Table 34.10 carries it. Separate jobs also change the log, since three chillers returned one at a time make 18 jobs in the 290 h, with $\mu \sim \text{Gamma}(20,\ 310\,\text{h})$ and $\hat{\omega}_\mu = \ln(20/310)$, $\Sigma = 1/20$. The chillers' term is then 7.34×10^{-4} (6.23×10^{-4}), 1.79 times the two-state value and not twice, as it would be with the rate kept. They raise the top event to 2.638×10^{-3} (1.836×10^{-3}), and the controller's expected share falls to 0.62 (0.54); the ranking of the programs does not move, and staggering is worth about 1.8 times as much. The record's sixteen repairs for sixteen outages favor the one job if they count jobs, and the log does not say whether it counts jobs or outages; one logged outage cannot settle it.

Table 34.13: The chillers' term under each reading [10^{-4}]: at $\hat{\mathbf{z}}$, and expected on each route.

Reading	At $\hat{\mathbf{z}}$	State	Exact
Two-state cut sets, equation (34.3) (reported)	3.214	4.092	3.488
Chain, one crew, one job for the shared cause	3.342	4.244	3.650
Chain, one crew, separate jobs, rate kept from 16 jobs	6.446	8.208	6.999
Chain, one crew, separate jobs, Gamma(20, 310 h)	5.777	7.336	6.229

34.4.4 The shared cause against a coincidence

The simultaneous outage could be a coincidence, the drawing M_0 of Chapter 15, here with the coincidence rate $6\lambda_c^2 t_w$ over the three pairs, generous to M_0 since a coincidence of all three would be far rarer. At $t_w = 4$ h and at $\hat{\mathbf{z}}$ the record expects 0.020 coincidences against one observed. Table 34.14 gives the Bayes factor by equation (15.29), the exact value beside it, one window to a row, and what staggering is worth when that factor enters: the reduction of Table 34.10 times $\Pr(M_1 \mid \mathbf{y})$ at even prior odds, since under M_0 there is no shared cause to stagger. The window was fixed by the log and not after looking at the event, and with its stated 4 h the shared cause is well supported. At 24 h, staggering would stand at R$ 31 219 (R$ 26 650) and the two routes would fall on opposite sides of its price; both are stated.

Table 34.14: The evidence for the chillers' shared cause by window, on each route, and staggering's expected reduction averaged over the two drawings at even prior odds [R$/yr]. One job.

	$\ln B_{10}$		$\Pr(M_1 \mid \mathbf{y})$		Staggering	
Window [h]	State	Exact	State	Exact	State	Exact
1	4.09	4.13	0.984	0.984	41 728	35 268
4	2.72	2.76	0.938	0.940	39 802	33 698
8	2.04	2.09	0.885	0.889	37 566	31 874
24	1.02	1.07	0.736	0.744	31 219	26 650

34.4.5 The priors and the record

Table 34.15 varies the priors. Each row is one change; the columns give the expected top event, the three reductions and the controller's expected share, on each route. The fourth chiller stays near R$ 3 000 on this form and the controller above its price on all of them, even if the single finding had been none, with a Jeffreys prior. The share's prior moves staggering across its price: with Beta(1, 19) it is R$ 25 708 (R$ 21 214), and with Beta(1, 4) R$ 61 970 (R$ 53 242). With Gamma(4, 40 000 h) on the rates it sits at the price on the exact route. The record itself is not unusual under its model: at $\hat{\mathbf{z}}$ the expected simultaneous count is 0.785, with a probability of 0.54 of at least one, and the singles expect 15.8 against 15.

Table 34.15: The answers under other priors [R$/yr for the reductions], two-state form; state route with the exact route beneath, each by quadrature.

Prior changed	Route	$\mathbb{E}[U_{\text{sys}}]$ [10^{-3}]	Fourth	Staggering	Controller	K share
None	state	2.31	3 300	42 428	358 179	0.71
	exact	1.56	3 509	35 834	218 440	0.64
Share Beta(1, 19)	state	2.06	3 514	25 708	358 272	0.80
	exact	1.34	3 673	21 214	218 488	0.75
Share Beta(1, 4)	state	2.62	3 075	61 970	358 071	0.63
	exact	1.83	3 322	53 242	218 381	0.55
Rates Gamma(4, 40 000 h)	state	2.21	2 701	35 690	358 218	0.75
	exact	1.47	2 879	30 080	218 459	0.68
Controller, Jeffreys, Beta(1.5, 999.5)	state	2.76	3 299	42 409	454 906	0.76
	exact	2.06	3 508	35 816	327 169	0.73
Controller, no finding, Beta(0.5, 1 000.5)	state	2.02	3 301	42 441	291 951	0.67
	exact	1.06	3 511	35 852	109 165	0.47

34.5 The report (step 6)

The fourth chiller that looks 303 times better on the independent product removes 1.3 % (2.0 %) of the risk with the plant's one crew, R$ 6 562 (R$ 6 977) a year against a price of R$ 180 000, because the chillers' shared cause is 96 % (95 %) of their term and the fourth chiller goes down with the others. The controller, with one finding against the chillers' eighteen failures, carries 71 % (64 %) of the expected top event. The answers follow in the order they were asked.

34.5.1 The answers and their bands

Table 34.16 is the report. Each row is a question in the engineer's words; the second column gives the answer on the state route, the third the exact value, the fourth the gap, the fifth the

band and the qualification, the last the chapter that derives it.

Table 34.16: The report of the plant: the answer on the state route, the exact value and the gap, the bands, and the chapter behind each. Prices illustrative.

Question	Answer, state route	Exact	Gap	Band, qualification	From
Chillers' common share	median 0.130, mean 0.151	median 0.113, mean 0.127	0.017; 0.024	[0.034, 0.390] ([0.018, 0.317]); $\Pr(M_1) = 0.94$ on both at 4 h	15
Cooling lost	mean 2.31×10^{-3}, 20.3 h a year; median 1.72×10^{-3}	mean 1.56×10^{-3}, 13.7 h; median 1.29×10^{-3}	7.5×10^{-4}; 4.3×10^{-4}	$[5.6,\ 77.6] \times 10^{-4}$ ($[3.5,\ 43.2] \times 10^{-4}$); ξ_K holds 92 % of the variance; R$ 506 812 (R$ 342 030) a year	14, 25
Combinations	{K} 0.71, {CC_c} 0.17, {CC_p} 0.11, pairs 0.01	0.64, 0.21, 0.14, 0.01	0.07; −0.04; −0.02; 0	expected shares; controller's band [0.16, 0.94] ([0.04, 0.92])	14
Component	controller on both; median FV 0.63, I_B 0.9994	median FV 0.58; I_B 0.9995	0.05; −0.0001	first of the three with probability 0.79 (0.70); I_B does not separate order-one events	26
Back after a double	0.50 at 12 h, 0.75 at 24 h, 0.93 at 48 h	0.49, 0.74, 0.92	0.01	[0.35, 0.67] at 12 h ([0.34, 0.65]); from all three down, 0.505 (0.495) at 12 h with one job, 0.19 (0.18) with separate jobs	16
Fourth chiller	no: removes R$ 6 562 with one crew against R$ 180 000	R$ 6 977	R$ −415	two-state R$ 3 300 (R$ 3 509); exceeds its price in no draw; if every cause trips two, R$ 89 042 (R$ 75 944), still no; buy the controller, staggering under its condition	32

The report line. Over five years the plant's record gives a cooling loss of about 20 (14) hours a year in expectation, with a band from 5 to 68 (3 to 38) hours, costing about R$ 510 000 (R$ 340 000) a year at the illustrative loss rate; the first figure is high because a normal on the logit around one finding has a mean of 1.64×10^{-3} for the controller where the exact posterior has 10^{-3}. About seven tenths (two thirds) of it runs through the controller, and that figure rests on one finding in a thousand checks, which gives the band most of its width. Of the chillers' term 96 % (95 %) is a shared cause, whose share of their failures has the median 0.13 (0.11), from one event. A chiller outage leaving one running is even odds of being over in twelve hours. A fourth chiller is not worth its price at any assumption examined within the stated loss rate; a redundant controller is, in expectation on both routes, with a probability of paying of 0.955 (0.833). Staggered maintenance pays if it removes at least 35 % (42 %) of the share. The illustrative requirement of 10^{-4} on the unavailability is not met by any program on offer, with probability at most 0.006 (0.028).

34.5.2 The investment decision

The recommended program is priced in Table 34.11: buy the redundant controller and do not buy the fourth chiller, at R$ 188 632 (R$ 163 590) a year in all on the two-state form, leaving 6.79×10^{-4} (5.64×10^{-4}) of top event, 5.9 h (4.9 h) a year, band $[2.0,\ 16.3] \times 10^{-4}$ ($[1.5,\ 13.6] \times 10^{-4}$). Add staggering on one condition: that the log shows a cause it reaches, a maintenance error repeated on the chillers, carrying at least 35 % (42 %) of the share. The

causes the plant names, water quality and a control fault, are not reached by a schedule. If that holds at the half, the program is R$ 176 135 (R$ 157 721), leaving 4.85×10^{-4} (4.01×10^{-4}), and staggering's marginal probability of paying is 0.55 (0.47). All of the prices are illustrative. The controller pays in expectation down to a loss rate of about R$ 2 800 h^{-1} (R$ 4 600 h^{-1}), up to a shared cause between the two of 0.89 (0.82), and with a probability of 0.955 (0.833) at the state. The fourth chiller stays out up to a loss rate over R$ 640 000 h^{-1}. Staggering also needs a loss rate above R$ 17 677 h^{-1} (R$ 20 930 h^{-1}) and, as Table 34.15 shows, depends on the share's prior.

> **Principle 34.1.** *In a plant in series, each stage is judged against the others. A redundant group is bounded below by its shared cause, so more units buy little once that cause dominates, and a change is valued by the cut sets it can reach and no others. The expected share is reported with its band and on a stated route; a measure that does not separate the leading events, as the Birnbaum measure does not among order-one cut sets, is not the one that ranks them. And what no record informs, a removed fraction or a repair reading, is stated with its break-even and not chosen.*

Common mistake

Four errors recur with a redundant plant. Pricing a unit by the independent product: here a factor of 303 on a term that is 0.7 % (1.0 %) of the risk. Ranking components by counted failures: eighteen for the chillers against one finding for the controller, which carries the larger share. Reading the Birnbaum measure as a ranking among order-one cut sets: 0.99941 against 0.99871 says only which x_i is larger. And taking two-state cut sets with a crew on the plant without asking how the outages are repaired: one crew doubles what a fourth chiller reaches, and separate jobs raise the chillers' term 1.8 times.

34.6 The next measurement (step 7)

The preposterior of equation (2.17) answers what would narrow the top event, and the value of information what would change the program; the two lists differ.

34.6.1 The candidate records and their preposterior

Table 34.17 gives the candidates. Each row is a record, with its rows expected at $\hat{\mathbf{z}}$; the columns give the spread of $\ln U_{\text{sys}}$ after it and the factor between the ends of its 95 % band, by the rank-one form with $\mathbf{g}$ of equation (34.8), on the state route. A thousand more checks of the controller take the factor from 14.3 to 7.1 and ten thousand to 2.9; five more years of the chillers' whole log take it to 14.0. The controller's check record is what narrows the band, as ξ_K's 92 % promised.

34.6.2 The value of knowing more

The band is not the decision. By equation (30.18) over the seven programs of Table 34.11 and the one with staggering and the fourth chiller, each route's own draws on both sides, knowing the whole state exactly is worth R$ 5 860 to R$ 5 940 a year on the state route and R$ 10 680 to R$ 10 750 on the exact route, over seeds 34 to 38. Those ceilings are small against the prices. A fine grid across the kink gives R$ 532 (R$ 3 439) for knowing q_K alone, and each route's quadrature R$ 5 375 (R$ 7 276) for knowing the chillers' block alone once the controller is bought. A thousand more checks have no outcome that would turn down the controller: their value by equation (2.51) is zero on both routes. Ten thousand are worth under R$ 1 (R$ 1 843), with the update on the state route meaning the present normal as the prior and the new row reweighting

Table 34.17: The candidate records: spread of $\ln U_{\text{sys}}$ and the factor across its 95 % band after each, rank-one preposterior at $\hat{\mathbf{z}}$.

Record	Spread	Factor
Present	0.679	14.3
1 000 more controller checks	0.498	7.1
10 000 more controller checks	0.272	2.9
Five more years of the chillers' log	0.674	14.0
The same, with 16 more repairs	0.673	14.0
Five more years of the pumps' log	0.677	14.2
All the logs and 1 000 checks	0.488	6.8

it; a gap that is the difference between a logit normal, with little mass below the break-even q_K, and Beta(1, 999), with 0.167. A refit of the logit normal from the pooled count is not used here, since its expected mean is not the present one, and at a thousand checks, where no action changes, it reports R$ 401. On the decision the chillers' block is the more valuable one to learn.

34.6.3 The recommendation

Do not wait on a record for the controller or against the fourth chiller. Strengthen two records. A time-stamped log of the chillers' outages, with the units tripped and the jobs that restored them: it settles the window and the coincidence of Table 34.14, whether a cause takes two chillers or three, and whether the outage is one job or several, which moves the chillers' term by a factor of 1.8, with the cause of each outage, which decides whether staggering reaches it. And the controller's check record, for the band of the report rather than for the action. After the changes, the pumps' shared cause and the chillers' remain, and the next priced question is that of the pumps, which this chapter's options do not reach. When rows arrive, the fit returns to step 3.

34.6.4 What the analysis showed

The first reading set eighteen chiller failures against one controller finding and a factor of 303 in favor of a fourth chiller. The analysis finds the factor correct and the term it acts on negligible, the chillers' risk almost wholly a cause that the new chiller shares, and the larger part of the plant's risk in the controller, which the log barely mentions. For the plant, this means buying the controller's redundancy first, trying the staggering under its stated condition, and spending nothing on the chiller. The engineer who takes this back should expect the top event to fall by about two thirds, and three quarters if staggering's condition holds, should not promise the 10^{-4} target from these changes, and should say that the size of the controller's share rests on a single finding, while its place at the head of the list survives the priors examined.

34.6.5 Limits of the model

The derivations are exact under their assumptions, and five of them need checking before a result is used.

The repair of a shared-cause outage. Section 34.4.3 gives both readings; the record's sixteen repairs for sixteen outages favor one job and do not prove it.

The multiplicity of the shared cause. The one event is read as tripping all three. If causes of two units exist, the fourth chiller is worth more, by the values of Section 34.3.6, and the record's one event cannot separate them.

The fraction staggering removes. No row informs it. A water-quality event is not reached by a maintenance schedule, and the same holds for a control fault.

The chillers' control fault and the controller. A control fault that trips the chillers may be the controller's own. Then the two terms overlap, the sum double counts a little, and the redundant controller also removes part of the chillers' term, which would raise its value.

The pumps' record, and a steady state. The pumps are carried with Chapter 15's pinned repair and its share, and the controller's check finds it down for any reason, planned work included. Every unavailability is stationary, and a season that needs three chillers is outside the plant described.

The next chapter takes the method to the plant where power and cooling couple.

34.7 What to retain

Definition 34.1 (The floor of a redundant group)**.** A group whose shared cause trips every unit has about the unavailability U_{cc} plus a term that falls with redundancy. A further unit tripped by the same cause removes at most the second term, so U_{cc} is a floor that no number of such units lowers.

Four conclusions carry beyond the plant. A change is valued by the cut sets it reaches, on the whole tree: a factor on a small term is a small reduction. Importance is asked of the measure that can answer: on order-one cut sets the Birnbaum measure is near one for all and the Fussell–Vesely measure ranks them. The entry that holds the variance need not be the one the decision turns on: here ξ_K holds nine tenths of the band and the chillers' block holds more of the value. And a recommendation is the program priced, with the break-even of every number that no record informs. Table 34.18 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 34.18: The terms of the plant's report and where each is defined.

Term	Expression	Units	Meaning	Defined in
Chillers' term	$1-(1-U_{\mathrm{cc,c}})(1-3q_{\mathrm{c}}^2+2q_{\mathrm{c}}^3)$	–	2-out-of-3 with a cause	34.1.3
Top event	$1-(1-q_K)(1-U_{\mathrm{C}})(1-U_{\mathrm{pair}})$	–	three stages in series	34.1.3
Annual loss	$c_a + C_{\mathrm{L}} t_{\mathrm{yr}} U_{\mathrm{sys}}$	R$/yr	price and consequence	34.1.3
General forms				
Entry	$p(z_i \mid \mathbf{y})$	–	the share	34.3.1
Expectation, quantile	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$	–	top event, importance	34.3.2
Cut-set sum	$\sum_c P_c$	–	combinations	34.3.3
Predictive	$\int \Pr(A \mid \mathbf{z}) p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	–	recovery	34.3.5
Comparison	$\mathcal{R}(a') - \mathcal{R}(a)$	R$/yr	options	34.3.6
Specialized forms and answers				
Order-one Birnbaum	$(1-U_{\mathrm{sys}})/(1-x_i)$	–	nearly one	34.3.4
Break-even of staggering	removed fraction with reduction $= c_{\mathrm{st}}$	–	0.35 (0.42)	34.3.6

34.8 Exercises

Exercise 34.1. *Formulation.* (a) Three chillers need two, with $q = 2 \times 10^{-3}$ the total unavailability and $\beta_{\text{cc}} = 0.12$. Compute the working form of the group's unavailability. (b) The top event has the cut sets 10^{-3}, 3×10^{-4}, three of 5×10^{-6}, 2×10^{-4} and 2×10^{-6}. Compute the rare-event sum and the top event in series, taking the cut sets as independent. ((a) $\beta_{\text{cc}} q + 3((1 - \beta_{\text{cc}})q)^2 = 2.40 \times 10^{-4} + 9.3 \times 10^{-6} = 2.49 \times 10^{-4}$. (b) 1.5170×10^{-3} and 1.5164×10^{-3}.)

Exercise 34.2. *Derivation.* Derive the β-factor unavailability of two of three chillers. (Thinning makes four independent up-and-down processes: three at $(1 - \beta_{\text{cc}})\lambda$ and one at $\beta_{\text{cc}}\lambda$. Each is down the fraction of equation (15.41). Cooling holds if the shared process is up and at most one own process is down, $(1 - U_{\text{cc}})(1 - 3q_{\text{ind}}^2 + 2q_{\text{ind}}^3)$. With $\epsilon = \lambda/\mu$ the result is $\beta_{\text{cc}}\epsilon + O(\epsilon^2)$: the shared cause at first order, the group's redundancy at second.)

Exercise 34.3. *Derivation.* Derive the controller's cut-set probability under a redundant controller, independent given q_K and then with a shared share β_K, and its expectation under Beta(1, 999). (Independent given q_K, $P = q_K^2$ with $\mathbb{E}[q_K^2] = 2/(1\,000 \times 1\,001) = 2.0 \times 10^{-6}$. With a share, $P \approx \beta_K q_K + ((1 - \beta_K)q_K)^2$; at $\beta_K = 0.13$, $\mathbb{E}[P] = 1.32 \times 10^{-4}$, and the removal is still R\$ 190 092 a year with the chapter's other terms.)

Exercise 34.4. *Derivation.* Derive the options' ranking as a function of the loss rate, on the exact route with one job. (Each reduction is proportional to C_{L}: ΔU of 9.97×10^{-4}, 1.64×10^{-4} and 1.60×10^{-5}. Each pays above $c_a/(\Delta U t_{\text{yr}})$: R\$ 4 578 h^{-1}, R\$ 20 930 h^{-1} and R\$ 1 282 000 h^{-1}. Their order by ratio is the order of $\Delta U/c_a$ and does not depend on the loss rate; only which ones pay does.)

Exercise 34.5. *Application.* A plant of four chillers needs two, with a shared cause under an α-factor model: of the shared events, half trip exactly two chillers, three tenths three, and two tenths all four. With $U_{\text{cc}} = 3.09 \times 10^{-4}$ and $q_{\text{ind}} = 2.07 \times 10^{-3}$, say which events take cooling down and give the group's unavailability. (An event of two leaves two running, so it loses cooling only with one more unit down; events of three or four lose it at once. $U \approx 0.5U_{\text{cc}} + 0.5U_{\text{cc}} \times 2q_{\text{ind}} + 4q_{\text{ind}}^3 = 1.55 \times 10^{-4}$, half of the β-factor value: the fourth chiller's value depends on the multiplicity, which one logged event does not give.)

Exercise 34.6. *Application.* The budget is R\$ 70 000 a year. Say which options it admits together, what they remove and what they leave, on both routes, one job. (The redundant controller and staggering, exactly at the budget; the fourth chiller is excluded by its price alone. They remove R\$ 400 677 (R\$ 254 309) and leave R\$ 106 135 (R\$ 87 721) of expected loss. If staggering's half is not credited, the same budget leaves R\$ 30 000 unspent: the controller alone leaves R\$ 148 632 (R\$ 123 590).)

Chapter 35

Data Center Power and Cooling Reliability

35.1 Six outages a year and a hall at 17.7 degrees

The grid fails about six times a year at this data center, for two hours on average, and each time the emergency generator of Chapter 9 is asked to start and the chiller of Chapter 28 to restart on its power. The facilities engineer has three proposals on the desk: a chilled-water tank that doubles the mass of the loop, a control change that cuts the chiller's restart from 0.25 to 0.1 h, and a maintenance program on the generator. The limit on hall 1 is 27°C, and a crossing is priced at R$ 150 000, an illustrative figure.

The quickest reading says that none of this matters. Chapter 18's fit puts hall 1 at 17.70°C, 9.30 K under the limit. The chiller delivers 54.7 kW of the 52.9 kW asked of it (Chapter 28). With no cooling at all, the 30.7 kW of the hall into a thermal mass of about 9 kWh/K, the design figure, is a rise of 3.4 K/h, so the hall rides through for 2.7 h, and the average outage is shorter than that. One failed start in forty tests is 2.5 %. The hall rides through the average outage even in the worst case. An average, however, is not the outage that crosses, and that reading has not said which entry the risk hangs on. The question of this chapter is

how often an outage takes hall 1 over its limit, through which path, on which entry of the state that rests, what it costs a year, which proposal is worth its price, and what to measure next.

The chapter runs the method from step 0 to step 7 as a report and cites the earlier derivations. Two pieces are new: the transient of the hall through its thermal mass, and a state indexed by time, some hundreds of entries long, on which the order of elimination promised in Section 4.1.3 decides the cost of the fit.

Jargon: ride-through, thermal mass, sparse information matrix, elimination order, fill-in, Cholesky factor, Kalman filter, Rauch–Tung–Striebel smoother, augmented state

Ride-through is the time a plant holds its limit with its cooling lost, and the *thermal mass* is the stored heat capacity that sets it. An information matrix is *sparse* when most of its entries are zero, which is the case when most rows touch few entries. The *Cholesky factor* is the lower-triangular $\mathbf{L}$ with $\mathbf{L}\mathbf{L}^{\mathsf{T}} = \mathbf{\Lambda}$. Eliminating one entry of the state is one column of that factorization, and the *elimination order* is the sequence in which the entries are taken. *Fill-in* is a nonzero of $\mathbf{L}$ where $\mathbf{\Lambda}$ had a zero. The *Kalman filter* computes the posterior of a time-indexed state from the rows up to each time, and the *Rauch–Tung–Striebel smoother* carries the result backward to give the posterior from all the rows. An *augmented state* adds constant entries to the state of the filter so that they are learned with it.

Key idea

A coupled plant is asked one question through all its subsystems, and the answer is carried by whichever path has a chance of crossing, not by whichever subsystem is worst. Here that is the generator's start, under the drawing in which the circulation stops with the power. The drawing is part of the answer: the same state, with the pumps kept running, moves the largest saving from the generator to the loop. A long time-indexed record changes the computation and not the method, provided the entries are eliminated in time order.

35.1.1 The coupled plant and its graph (step 0)

Figure 35.1 draws the plant. The gray box on the left is the grid, pinned, with its outages, which are aleatory. The chain below it is the coupling: the generator (green), which must start and then run; the chiller, with its capacity factor; the loop, with a pinned mass and a restart; and hall 1, with its conductance, its load and its thermal mass. The dashed boxes along the bottom are the rows that observe each block. Hall 2 acts on the loop only, through its load.

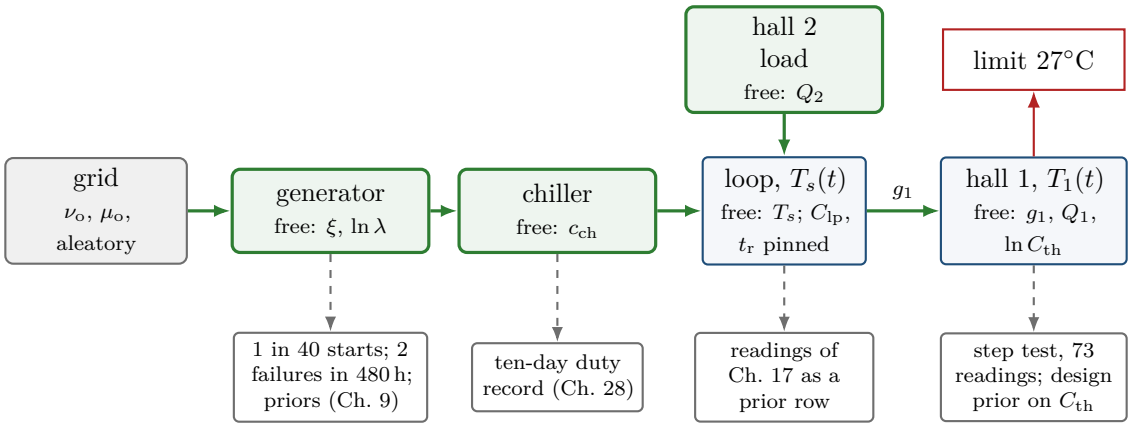


Figure 35.1: The coupled plant. Pinned and aleatory: the grid's outages (gray). Free: the generator's and chiller's entries (green) and the entries of the loop and the hall (blue boxes, whose temperatures in time are functions of them). Observed (dashed): the records. The red box is the event asked about.

An outage has the duration $D_o = \min(E, 12\,\text{h})$ with E exponential of mean $\mu_o = 2\,\text{h}$; there are $\nu_o = 6$ a year. The chiller is the one of Chapter 28, plate $P_{pl} = 60\,\text{kW}$, and its capacity factor is that chapter's: a load-bank result, taken here to hold on generator power, which that record did not test. The loop has the mass $C_{lp} = 20\,\text{kWh/K}$ and the restart delay $t_r = 0.25\,\text{h}$, both pinned and illustrative. A crossing is priced at C_{cr}. Hall 1's thermal mass is new, with the design figure $9.0 \pm 1.5\,\text{kWh/K}$, also illustrative. Table 35.1 lists the three proposals and three other changes with placeholder prices: none is a quotation, and the halvings are assumed, so a later table gives the break-even of each, which is what the conclusions rest on.

35.1.2 The probabilistic questions (step 1)

Table 35.2 lists the six questions in the order they are treated. The columns give the question in engineering language, the general form of its category, its form on this plant, the category and the subsection that answers it. Here $\Pi(\mathbf{z})$ is the probability that one outage crosses the limit given the state, with the outage and the generator's outcome averaged out.

Table 35.1: The actions and their annual prices, placeholders and not quotations. The last three are not among the proposals; Section 35.1.3 says why the last two are here.

Action	What changes in the model	c_a [R$/yr]
Tank	C_{lp} from 20 to 40 kWh/K	25 000
Fast restart	t_{r} from 0.25 to 0.1 h	4 000
Generator maintenance	$p \to p/2$ at every state, an assumed effect	3 000
Running-failure program	$\lambda \to \lambda/2$ at every state, an assumed effect	3 000
Pumps and coil fans on batteries	circulation continues with the generator down	12 000
The same, with the tank	the same, and $C_{\mathrm{lp}} = 40$ kWh/K	37 000

Table 35.2: The questions asked of the coupled plant.

Engineering question	General form	Mathematical form	Category	Answered in
How often does an outage push hall 1 over 27°C?	$\int p(y_{\mathrm{new}} \mid \mathbf{z}) p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	$\mathbb{E}[\Pi(\mathbf{z}) \mid \mathbf{y}]$	prediction under a threshold	35.3.1
Through which path does it cross?	the same, split	the terms of Π: start failure, running failure, generator running	attribution	35.3.2
Which entry of the posterior does the risk hang on?	$\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$	the budget of $\ln \Pi$; $\mathbb{E}[\Pi \mid C_{\mathrm{th}} \in \mathrm{bin}, \mathbf{y}]$	estimation (sensitivity)	35.3.3
What does a year of outages cost, and how well is that known?	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$; quantiles	$\nu_{\mathrm{o}}\Pi$ and $C_{\mathrm{cr}}\nu_{\mathrm{o}}\Pi$	estimation of a derived quantity (risk)	35.3.4
Tank, faster restart, or generator maintenance?	$\mathcal{R}(a') - \mathcal{R}(a)$	$C_{\mathrm{cr}}\nu_{\mathrm{o}}\mathbb{E}[\Pi_0 - \Pi_a \mid \mathbf{y}]$ against c_a	design	35.3.5
What to measure next?	$\boldsymbol{\Sigma}_a$; EVSI(a)	$\boldsymbol{\Sigma}_a$ after n start tests; the step test, computed; their value to the maintenance decision	design (information)	35.3.6

35.1.3 The evidence as rows (step 2)

The modeling choices. Seven choices make the rows. The carried posteriors enter as prior rows, which keeps each chapter's fit and fails if those records share an error with this one. The hall follows a single first-order mass, which is what makes C_{th} one number and fails for a hall with fast air and slow structure, where the first minutes rise faster than the mean. The hall's load is held through the outage: this is what makes the thermal question stand, and it requires batteries that carry the load for hours; with batteries of less than about 2.7 h the load drops first, and the crossing is replaced by a loss of service that this chapter does not price. A start failure or a running failure stops the chiller, the pumps and the coil fans together, so the hall has no path to the loop; this is the drawing the chapter adopts, and the option table tests the alternative because the conclusion on the tank depends on it. The hall's cooling resumes the moment the grid returns; the options' answer tests that as well. The outage's parameters are pinned, as aleatory draws with a known law, and Section 35.4.3 prices what that hides. And the effect of each maintenance program on p or λ is assumed, with a break-even in its place.

The transient. Hall 1 is the node of Chapter 17 with a mass added to its balance:

$$C_{\mathrm{th}}\,\frac{\mathrm{d}T_1}{\mathrm{d}t} = Q_1 - G_1(T_1 - T_s(t)), \qquad G_1 = \mathrm{e}^{g_1}, \qquad T_1^0 = T_s + \frac{Q_1}{G_1}, \tag{35.1}$$

with T_1^0 the steady temperature at the start of the outage, and with the time constant C_{th}/G_1. With no cooling the path is cut, $G_1 \to 0$ in the transient, and the hall rises at Q_1/C_{th}, so that it reaches the limit $T_{\mathrm{lim}} = 27^{\circ}\mathrm{C}$ after the crossing time

$$t_{\times}(\mathbf{z}) = \frac{(T_{\mathrm{lim}} - T_1^0)\,C_{\mathrm{th}}}{Q_1}. \tag{35.2}$$

A no-cooling path crosses if and only if the outage outlasts $t_{\times}$, under the choice that cooling resumes at once, and the duration's survival is

$$S_{\mathrm{o}}(t) = \Pr(D_{\mathrm{o}} > t) = \mathrm{e}^{-t/\mu_{\mathrm{o}}}\,\mathbf{1}[t < 12\,\mathrm{h}]. \tag{35.3}$$

With the generator running, the supply moves. During the restart the chiller removes nothing and the loop warms at the whole demand, $P = Q_1 + Q_2$; afterwards it holds its capacity against the demand, with the headroom $P_{\mathrm{hr}} = P_{\mathrm{pl}}c_{\mathrm{ch}} - P$ of equation (28.19):

$$C_{\mathrm{lp}}\frac{\mathrm{d}x}{\mathrm{d}t} = \begin{cases} P, & t < t_{\mathrm{r}}, \\ -P_{\mathrm{hr}}, & t \geq t_{\mathrm{r}}, \end{cases} \quad x \geq 0, \qquad T_s(t) = T_s + x(t) + \mathbf{1}[t \geq t_{\mathrm{r}}]\frac{(-P_{\mathrm{hr}})^{+}}{\dot{m}c_p}, \tag{35.4}$$

the last term the onset shortfall (28.20), with $\dot{m}c_p = 10\,\mathrm{kW/K}$ pinned. The restart's rise is $x(t_{\mathrm{r}}) = Pt_{\mathrm{r}}/C_{\mathrm{lp}}$, 0.66 K at 52.9 kW. The supply does not settle at a derated value: with positive headroom the rise decays slowly, at $P_{\mathrm{hr}}/C_{\mathrm{lp}}$, and with negative headroom the loop has no steady state and goes on warming, the limit that Chapter 28 named. Equations (35.1) and (35.4) are integrated in steps of three minutes, exactly within each step for a supply linear on it.

The step test. The thermal mass needs a record, and it gets one: a generated illustrative test. At $t = 0$ a load bank of $\Delta Q_{\mathrm{st}} = 10\,\mathrm{kW}$ in hall 1 is shed. A shed is chosen because the chiller has only 1.77 kW of headroom for a rise. The hall temperature is logged every five minutes for six hours, 73 readings, on a test logger of spread $\sigma_{\log} = 0.1\,\mathrm{K}$. The transient is carried on a grid of one minute, $\Delta t = 1/60\,\mathrm{h}$, with $T_{[j]}$ the hall temperature at step $j = 0, \ldots, 360$, each an entry

of the state. The rows are:

$$\mathbf{r}_{\mathrm{st}} = \begin{bmatrix} (T_{[0]} - T_s - Q_1 \mathrm{e}^{-g_1})/\sigma_{\mathrm{i}} \\ \Big(T_{[j+1]} - T_{[j]} - \Delta t\, \mathrm{e}^{-\ln C_{\mathrm{th}}}\big(Q_1 - \Delta Q_{\mathrm{st}} - \mathrm{e}^{g_1}(T_{[j]} - T_s)\big)\Big)/\sigma_p, \quad j = 0, \ldots, 359 \\ (T_{[5i]} - y_{[i]})/\sigma_{\log}, \quad i = 0, \ldots, 72 \\ \mathbf{L}_{\mathrm{c}}^{-1}\big((g_1, Q_1, T_s)^{\mathsf{T}} - (\hat{g}_1, \hat{Q}_1, \hat{T}_s)^{\mathsf{T}}_{\mathrm{carried}}\big) \\ (\ln C_{\mathrm{th}} - \ln 9)/0.1667 \end{bmatrix}, \tag{35.5}$$

438 rows. The first says the hall is at steady state before the shed, with $\sigma_{\mathrm{i}} = 0.1\,\mathrm{K}$; the second are the physics rows, at the tolerance $\sigma_p = 0.002\,\mathrm{K}$ per step; the third the readings; the fourth the carried block of Table 19.3 as one prior row, with $\mathbf{L}_{\mathrm{c}}\mathbf{L}_{\mathrm{c}}^{\mathsf{T}}$ its 3×3 marginal; the last the design figure, with $1.5/9 = 0.1667$. With unnormalized rows the information would be $\mathbf{J}^{\mathsf{T}}\mathbf{W}\mathbf{J}$. The record was generated by these rows from numpy `default_rng`, seed 35, drawing $T_{[0]}$, then the physics steps in order, then the readings, with (g_1, Q_1, T_s) at their carried fit and $C_{\mathrm{th}} = 9.0\,\mathrm{kWh/K}$: the truth sits at the center of the priors, a choice, and Section 35.4.2 varies it. The explicit step shortens the time constant by about half a step, 0.5 %, well inside the posterior of the mass.

The state and the pattern. The state is

$$\mathbf{z} = (T_{[0]}, \ldots, T_{[360]},\ g_1,\ Q_1,\ Q_2,\ T_s,\ \ln C_{\mathrm{th}},\ c_{\mathrm{ch}},\ \xi,\ \ln\lambda)^{\mathsf{T}}, \tag{35.6}$$

369 entries, with g_2 integrated out of the carried block because no question reaches it. The generator's entries are those of equation (9.3), with ξ the logit of the start-failure probability p. The Jacobian's pattern is

$$\mathbf{J} = \begin{bmatrix} \mathbf{B}_{TT} & \mathbf{B}_{Ts} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{B}_{\pi} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & * & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & * \end{bmatrix} \begin{matrix} \text{the steady, physics and reading rows} \\ \text{the prior rows on the shared entries} \\ \text{duty and plate rows, Chapter 28} \\ \text{two records and two priors, Chapter 9} \end{matrix}, \tag{35.7}$$

with columns $T_{[j]}$; $(g_1, Q_1, Q_2, T_s, \ln C_{\mathrm{th}})$; c_{ch}; $(\xi, \ln\lambda)$. $\mathbf{B}_{TT}$ is bidiagonal, each physics row touching two consecutive temperatures, and $\mathbf{B}_{Ts}$ is dense in the shared entries. No row ties the generator to the hall: they meet only in Π.

35.2 The posterior of the state (step 3)

Every answer is read from the normal computed here. The last two blocks of pattern (35.7) are quoted, and the first is fitted, which is where the time index matters.

35.2.1 The carried blocks

The generator's block is equations (9.16) and (9.17): $\hat{\xi} = \ln(2/58) = -3.3673$ with $\Sigma_{\xi\xi} = 0.51724$, and $\widehat{\ln\lambda} = \ln(4/880) = -5.3936$ with variance 0.25, on the record of one failure in forty starts and two failures in 480 h, with the exact posteriors Beta(2, 58) and Gamma(4, 880 h) beside them. The Beta's two and fifty-eight are not a record of two failures in sixty starts: $2 = 1 + 1$ and $58 = 19 + 39$, the prior Beta(1, 19) with the record's one failure and thirty-nine successes. The chiller's is equation (28.7): $\hat{c}_{\mathrm{ch}} = 0.911536$, $\Sigma = 1.0922 \times 10^{-4}$, exact because its rows are linear. The hall's block before the test is Chapter 19's, with $\hat{g}_1 = 1.70297$, $\hat{Q}_1 = 30.688\,\mathrm{kW}$, $\hat{Q}_2 = 22.231\,\mathrm{kW}$, $\hat{T}_s = 12.1115°\mathrm{C}$.

35.2.2 The step test and the structure of the information

Figure 35.2 shows the structure. In each square the axes are the entries of the step-test block, 365 of them, in the elimination order, with a mark at every nonzero. Panel (a) is $\mathbf{\Lambda} = \mathbf{J}^{\mathsf{T}}\mathbf{J}$ in time order with the shared entries last: a tridiagonal band, from the physics rows, and a border, because every physics row touches the shared entries. Panel (b) is the Cholesky factor in that order. Panels (c) and (d) are the same matrix and factor with the shared entries first.

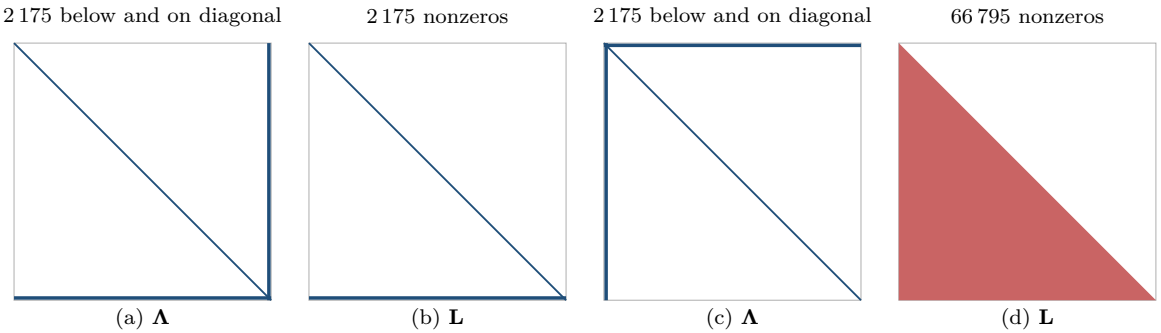


Figure 35.2: The sparsity pattern of the step-test block, 365 entries. Time order, shared entries last: (a) the information matrix, (b) its Cholesky factor, with no fill-in beyond the band and the border. Shared entries first: (c) the same information matrix, (d) its factor, entirely full (red). The band and border are drawn wider than one entry so as to be seen.

Elimination says why. Eliminating an entry is the Schur complement of Chapter 3: it joins all the entries still connected to it. In time order, $T_{[0]}$ is connected to $T_{[1]}$ and the four shared entries of the fit, g_1, Q_1, T_s and $\ln C_{\text{th}}$, which are already connected to each other. Removing it joins nothing new, and the same is true at every step. Each column of $\mathbf{L}$ holds five entries below its diagonal, and the work is fixed per step. If a shared entry is eliminated first, it is connected to every temperature, and removing it joins all of them to one another: the factor is full from then on. The count below takes a column with d entries below its diagonal to cost $d + 1$ operations to scale and $d(d + 1)/2$ to update what remains. Table 35.3 prints it. Each row is a number of time steps; the columns give the size of the block, the nonzeros of $\mathbf{\Lambda}$ on and below its diagonal, for comparison with the triangular factor, and the nonzeros of $\mathbf{L}$ and operations in each order. In time order the two counts are equal: no fill-in.

Table 35.3: The cost of one factorization in the two elimination orders, by exact symbolic elimination; the last entry of the last row is the formula $n^3/6$. The reported record has 360 steps.

			Time order, shared last		Shared first	
Steps	Entries	lower nnz($\mathbf{\Lambda}$)	nnz($\mathbf{L}$)	Operations	nnz($\mathbf{L}$)	Operations
180	185	1 095	1 095	3 815	17 205	1 072 445
360	365	2 175	2 175	7 595	66 795	8 171 255
720	725	4 335	4 335	15 155	263 175	63 776 075
1 440	1 445	8 655	8 655	30 275	1 044 735	about 5.0×10^8

Doubling the steps doubles the cost in time order and multiplies it by nearly eight with the shared entries first: linear against cubic. At the record's size the two differ by a factor of 1 076. The posterior is the same to round-off in either order: the order changes the work and nothing else. On a state of a few hundred entries neither order is out of reach; the difference matters for a year of readings, and it is why a general solver reorders a sparse matrix before factoring it.

35.2.3 The fit of the step-test block

Gauss–Newton on rows (35.5) starts from (g_1, Q_1, T_s) at the carried fit, $\ln C_{\text{th}} = \ln 9$ and the temperatures on the model's path. Table 35.4 lists the shared entries at each iterate, with the objective and the largest entry of the gradient. Each iteration factors the block in time order. The rows are nonlinear in the products of e^{g_1}, $\mathrm{e}^{-\ln C_{\text{th}}}$ and $T_{[j]}$, and the curvature term of a row of small residual is left out, as in Chapter 18; the reweighting below tests it.

Table 35.4: Gauss–Newton on the step-test block, 365 entries and 438 rows, in time order.

Iteration	g_1	Q_1 [kW]	T_s [°C]	$\ln C_{\text{th}}$	$\mathcal{L}$	$\max_i \lvert\mathbf{J}^{\mathsf{T}}\mathbf{r}\rvert_i$
0	1.70297	30.688	12.1115	2.19722	37.292	22.4
1	1.73714	31.004	12.1758	2.25880	32.936	13.0
2	1.73660	30.991	12.1746	2.26164	32.932	8.6×10^{-2}
3	1.73659	30.991	12.1745	2.26180	32.932	1.6×10^{-3}
4	1.73659	30.991	12.1745	2.26181	32.932	1.0×10^{-4}
8	1.73659	30.991	12.1745	2.26181	32.932	2.5×10^{-9}

At the minimum $2\mathcal{L} = 65.86$ on $\nu = 438 - 365 = 73$ degrees of freedom. The thermal mass is $C_{\text{th}} = \mathrm{e}^{2.26181} = 9.60\,\text{kWh/K}$, the time constant 1.69 h. Figure 35.3 shows the record. The horizontal axis is the time since the shed, the vertical the hall temperature; the dots are the readings, and the line and the shaded band are the smoothed $T_{[j]}$ and its 95 % band. The curve falls by the shed's $10/\mathrm{e}^{\hat{g}_1} = 1.76\,\text{K}$ and the reading scatter is many times the band, because 73 readings and the physics rows pool.

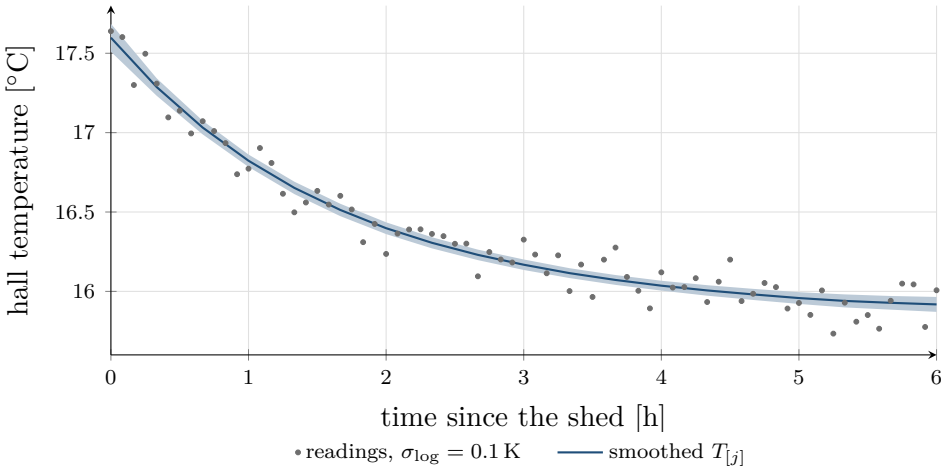


Figure 35.3: The step test (generated, seed 35, illustrative): a shed of 10 kW at $t = 0$, the 73 readings, and the smoothed trajectory with its 95 % band.

35.2.4 The Kalman filter and the smoother

The time-ordered factorization is a filter and a smoother, and this is a statement about the same algebra. Let $\mathbf{x}_j = (T_{[j]},\ g_1,\ Q_1,\ T_s,\ \ln C_{\text{th}})^{\mathsf{T}}$, the augmented state, and write the physics row, linearized at the current iterate, as

$$\mathbf{x}_{j+1} = \mathbf{A}_j \mathbf{x}_j + \mathbf{u}_j + \sigma_p\, \varepsilon_j\, \mathbf{e}_1, \qquad \mathbf{A}_j = \begin{bmatrix} a_j & \mathbf{b}_j^{\mathsf{T}} \\ \mathbf{0} & \mathbf{I} \end{bmatrix}, \tag{35.8}$$

with a_j and $\mathbf{b}_j^{\mathsf{T}}$ the row's entries of $\mathbf{J}$ divided by the first, and $\mathbf{e}_1$ the first unit vector: the shared entries do not change in time, and the row's tolerance enters only the first entry. Let the

rows up to step j give $\mathbf{x}_j \sim \mathcal{N}(\mathbf{m}_j, \mathbf{\Gamma}_j)$. Elimination of $T_{[j]}$ from the rows up to step j and the row to step $j+1$ is the marginal of the joint normal of $(\mathbf{x}_j, T_{[j+1]})$, with the Schur complement of Chapter 3. Written on the covariance rather than the information it is the *prediction*,

$$\mathbf{m}^-_{j+1} = \mathbf{A}_j \mathbf{m}_j + \mathbf{u}_j, \qquad \mathbf{\Gamma}^-_{j+1} = \mathbf{A}_j \mathbf{\Gamma}_j \mathbf{A}_j^\mathsf{T} + \sigma_p^2 \mathbf{e}_1 \mathbf{e}_1^\mathsf{T}. \tag{35.9}$$

Adding a reading at step $j+1$ adds a row, and the rank-one update (4.63) is the *update*,

$$\mathbf{m}_{j+1} = \mathbf{m}^-_{j+1} + \boldsymbol{\phi}_{j+1}(y - \mathbf{e}_1^\mathsf{T}\mathbf{m}^-_{j+1}), \quad \mathbf{\Gamma}_{j+1} = \mathbf{\Gamma}^-_{j+1} - \boldsymbol{\phi}_{j+1}\mathbf{e}_1^\mathsf{T}\mathbf{\Gamma}^-_{j+1}, \quad \boldsymbol{\phi}_{j+1} = \frac{\mathbf{\Gamma}^-_{j+1}\mathbf{e}_1}{\mathbf{e}_1^\mathsf{T}\mathbf{\Gamma}^-_{j+1}\mathbf{e}_1 + \sigma^2_{\log}}. \tag{35.10}$$

That is the Kalman filter on the augmented state, started from the first row and the prior rows. The forward pass of the time-ordered elimination computes it, and each column's work is that of a five-entry matrix, which is the constant in Table 35.3. Its last $\mathbf{m}_{360}$ already holds the posterior of the shared entries, since the last column of the factorization has seen every row. Back substitution recovers each eliminated entry from those after it, which is the conditional of $\mathbf{x}_j$ given $\mathbf{x}_{j+1}$:

$$\mathbf{m}^\mathrm{s}_j = \mathbf{m}_j + \mathbf{D}_j(\mathbf{m}^\mathrm{s}_{j+1} - \mathbf{m}^-_{j+1}), \quad \mathbf{\Gamma}^\mathrm{s}_j = \mathbf{\Gamma}_j + \mathbf{D}_j(\mathbf{\Gamma}^\mathrm{s}_{j+1} - \mathbf{\Gamma}^-_{j+1})\mathbf{D}_j^\mathsf{T}, \quad \mathbf{D}_j = \mathbf{\Gamma}_j \mathbf{A}_j^\mathsf{T} \left(\mathbf{\Gamma}^-_{j+1}\right)^{-1}, \tag{35.11}$$

the Rauch–Tung–Striebel smoother. Repeated at each Gauss–Newton iterate, this pair is the iterated form on a nonlinear row.

The check is numeric, at the first iterate, where the step is large. Equations (35.9) to (35.11) give the same step as the joint normal equations to 1.4×10^{-13} in the temperatures and 1.2×10^{-12} in the shared entries, and the same variances to 2.3×10^{-15} and covariances to 1.4×10^{-12}. At the converged state the smoothed $\ln C_\mathrm{th}$ is 2.261809 and its spread 0.082006 on both routes. The difference between filter and smoother is visible at an inner time: at 3 h the filter has 16.180°C with spread 0.031 K, the smoother 16.168°C with 0.017 K. The filtered spread of $\ln C_\mathrm{th}$ runs 0.119, 0.107, 0.104 and 0.082 at 1, 2, 3 and 6 h.

35.2.5 The posterior of the whole state

Since no row of the test touches Q_2, it follows the fitted block through its carried regression on (g_1, Q_1, T_s). In the order of equation (35.6), leaving out the temperatures, the posterior is

$$\hat{\mathbf{z}} = (1.73659,\ 30.9911,\ 21.9313,\ 12.1745,\ 2.26181,\ 0.91154,\ -3.36730,\ -5.39363)^\mathsf{T},$$

$$\mathbf{\Sigma} = \begin{bmatrix} \mathbf{\Sigma}_\mathrm{h} & \mathbf{0} \\ \mathbf{0} & \mathrm{diag}(1.092 \times 10^{-4},\ 0.51724,\ 0.25) \end{bmatrix},$$

$$\mathbf{\Sigma}_\mathrm{h} = \begin{bmatrix} 0.002487 & 0.024862 & -0.024572 & 0.005191 & 0.002142 \\ 0.024862 & 1.397756 & -1.245965 & -0.147441 & 0.016303 \\ -0.024572 & -1.245965 & 4.809365 & 0.121854 & -0.016103 \\ 0.005191 & -0.147441 & 0.121854 & 0.046602 & 0.003373 \\ 0.002142 & 0.016303 & -0.016103 & 0.003373 & 0.006725 \end{bmatrix}, \tag{35.12}$$

with $\mathbf{\Sigma}_\mathrm{h}$ the block of the first five entries. Table 35.5 reads off the entries, with the state before the step test beside them; each row is an entry, and the columns give the carried or design value and spread, and the value and spread with the test. The hall's steady temperature is $T_1^0 = 17.633$°C with spread 0.087 K, where before the test it had 0.48 K.

35.2.6 The routes and the exact posterior

The reported route is $\mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$. On it ξ and $\ln\lambda$ are averaged by the sixty-node Gauss–Hermite rule and the hall's correlated entries, with g_2 kept for the reweighting, by 10^5 antithetic draws,

Table 35.5: The posterior of the state: before the step test (carried fits and the design figure, the routes of Chapters 9, 19 and 28) and with it (equation (35.12)). The last three rows do not change.

	Before the test		With the test	
Entry	$\hat{z}_i$	$\sqrt{\Sigma_{ii}}$	$\hat{z}_i$	$\sqrt{\Sigma_{ii}}$
g_1	1.70297	0.10669	1.73659	0.04987
Q_1 [kW]	30.688	1.4521	30.991	1.1823
Q_2 [kW]	22.231	2.3462	21.931	2.1930
T_s [°C]	12.1115	0.2801	12.1745	0.2159
$\ln C_{\text{th}}$	2.19722	0.16667	2.26181	0.08201
c_{ch}	0.91154	0.01045	the same	
ξ	−3.36730	0.71920	the same	
$\ln \lambda$	−5.39363	0.50000	the same	

numpy `default_rng`, seed 35. Beside it stand Beta(2, 58), Gamma(4, 880 h), and the hall's exact posterior, reached by reweighting those draws by equation (20.11): the objective at a draw is that of rows (17.15), the design row, and the test's readings with the temperatures integrated out, which for fixed shared entries is the filter's prediction-error sum in one entry, $T_{[j]}$. So each draw costs one pass of a scalar filter. The effective sample is $N_{\text{eff}} = 62\,224$ of 10^5, and 6 664 to 28 800 of 4×10^4 over seeds 1 to 10 and 36 to 45, a heavy tail on a few draws; without the test, 83 044, and 34 823 to 37 193. There C_{th} has the median 9.66 and the band [8.20, 11.42] kWh/K, against the state route's 9.60 and [8.17, 11.28].

35.3 The answers (step 4)

The questions of Table 35.2 are read from the posterior of Section 35.2.5. Parentheses hold the exact route's value, each column of a table is one route, and every number below uses the generated step test.

35.3.1 The crossing probability per outage

Engineering question. How often does an outage push hall 1 over 27°C?

General form. The predictive, equation (2.13):

$$p(y_{\text{new}} \mid \mathbf{y}) = \int p(y_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}. \tag{35.13}$$

Specialized form. The question is equation (35.13) with y_{new} the crossing in one outage and $p(y_{\text{new}} \mid \mathbf{z}) = \Pi(\mathbf{z})$. Aleatory and epistemic parts stay apart, the distinction of Chapter 9: duration, start and running outcomes are averaged inside Π, and the state outside it. Given the state, the outage's three outcomes are exclusive,

$$\begin{aligned} \Pi(\mathbf{z}) &= p\,\Pi_{\text{s}}(\mathbf{z}) + (1-p)\big[\Pi_{\text{r}}(\mathbf{z}) + \Pi_{\text{g}}(\mathbf{z})\big], \qquad \Pi_{\text{s}} = S_{\text{o}}(t_\times), \\ \Pi_{\text{r}} &= \int_0^{12} \lambda\, \mathrm{e}^{-\lambda u}\, S_{\text{o}}(u + t_\times(u))\, \mathrm{d}u, \qquad t_\times(u) = \frac{(T_{\text{lim}} - T_1(u))\, C_{\text{th}}}{Q_1}, \end{aligned} \tag{35.14}$$

with Π_{s} the start failure, by equations (35.2) and (35.3); Π_{r} the running failure at u, from the temperature $T_1(u)$ that equations (35.1) and (35.4) reach by then; and Π_{g} the joint probability, given a start, that the generator survives until the limit is reached and that the outage lasts that long. Since ξ and $\ln \lambda$ are independent of the hall by pattern (35.7), the expectation factors:

$$\mathbb{E}[\Pi \mid \mathbf{y}] = \bar{p}\, \mathbb{E}[\Pi_{\text{s}} \mid \mathbf{y}] + (1-\bar{p})\, \mathbb{E}[\Pi_{\text{r}} + \Pi_{\text{g}} \mid \mathbf{y}], \qquad \bar{p} = \mathbb{E}[\mathrm{e}^{\xi}/(1+\mathrm{e}^{\xi}) \mid \mathbf{y}]. \tag{35.15}$$

For orientation only, ignoring the warming before u and the cap gives the closed form $\Pi \approx (p + (1-p)\lambda/(\lambda + 1/\mu_o))\,\mathrm{e}^{-t_\times/\mu_o}$, by the memoryless duration, which is the analytic route of Section 35.4.1.

Building blocks. $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$ of equation (35.12); $t_\times$; the transient; the two averages.

The numbers. At $\hat{\mathbf{z}}$, $t_\times = (27 - 17.633) \times 9.600/30.991 = 2.90\,\mathrm{h}$, and $\mathrm{e}^{-2.90/2} = 0.234$: about a quarter of the no-cooling outages cross, though the mean outage is shorter than the ride-through, because the exponential law puts that much past it. The rise is $Q_1/C_{\mathrm{th}} = 3.23\,\mathrm{K/h}$. Figure 35.4 draws the three paths at $\hat{\mathbf{z}}$ for a four-hour outage. The horizontal axis is time, the vertical the hall temperature, with the limit in red: the generator running stays near 18°C, the start failure crosses at 2.90 h, and a running failure at 1 h crosses at 3.82 h.

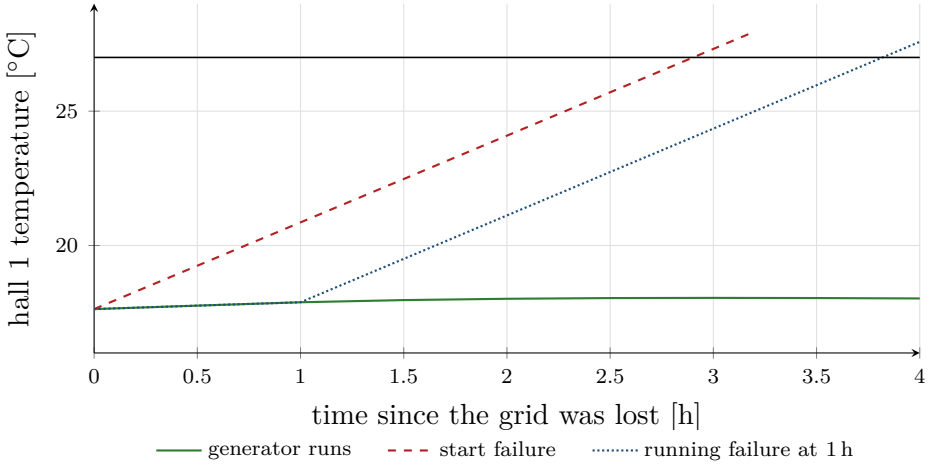


Figure 35.4: Hall 1's temperature at $\hat{\mathbf{z}}$ on the three paths: equations (35.1) and (35.4) with the generator running, and the rise $Q_1/C_{\mathrm{th}} = 3.23\,\mathrm{K/h}$ after the cooling stops. The black line is the limit. The start failure is drawn only to the point past its crossing.

Figure 35.5 carries this to the posterior. In panel (a) the horizontal axis is the thermal mass and the vertical the crossing time, for T_1^0 at its median and at the ends of its band, a narrow bundle; the shaded strip is the mass's band. In panel (b) the same crossing time is the horizontal axis and the survival (35.3) the vertical. The band of $t_\times$, $[2.43,\ 3.46]\,\mathrm{h}$ ($[2.44,\ 3.51]\,\mathrm{h}$), maps into survival between 0.30 and 0.18, so the mass moves the answer by a factor below two.

On the state route the Gauss–Hermite rule gives $\bar{p} = 0.04167$, above the exact mean $2/60 = 0.03333$ because the normal on the logit has a longer upper tail than the Beta; the draws give $\mathbb{E}[\Pi_s] = 0.23484$ (0.23214) and the factorization (35.15) gives

$$\mathbb{E}[\Pi \mid \mathbf{y}] = 0.01215 \quad (0.00982), \tag{35.16}$$

a gap of 0.00233, of which about four fifths is $\bar{p}$. Without the step test it is 0.01338 (0.01085). The band of $\Pi(\mathbf{z})$, from draws of the whole state, is $[0.00370,\ 0.03186]$ ($[0.00254,\ 0.02383]$). The test's construction enters the location of these numbers; Section 35.4.2 gives their range over other records.

35.3.2 The paths

Engineering question. Through which path does it cross?

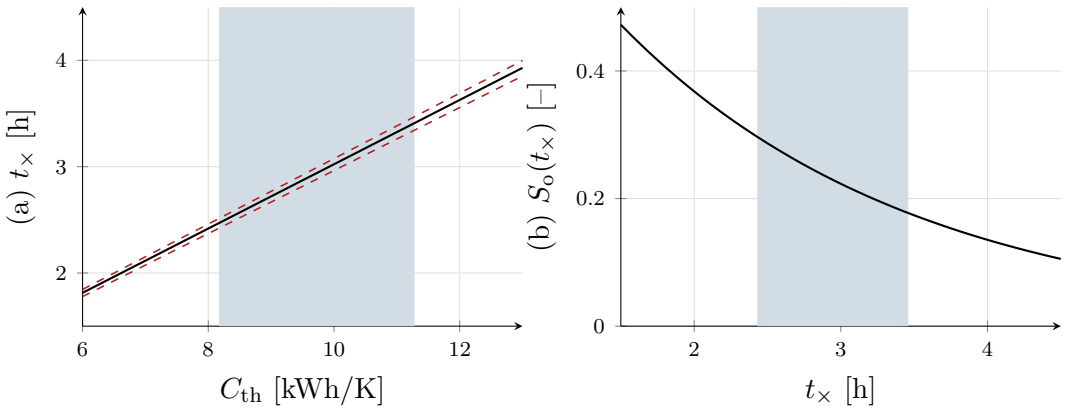


Figure 35.5: (a) The crossing time (35.2) against the thermal mass, at $Q_1 = 30.991\,\text{kW}$ and $T_1^0 = 17.633°\text{C}$ (black) and 17.469, 17.816°C (red), the band of T_1^0, with the state route's band of C_{th} shaded. (b) The duration's survival $\text{e}^{-t/2}$, with the band of $t_\times$ shaded.

General form. The predictive (35.13), split by the exclusive outcomes b of the event, as in the attribution of Chapter 25:

$$p(y_{\text{new}} \mid \mathbf{y}) = \sum_b \int p(y_{\text{new}}, b \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \text{d}\mathbf{z}. \tag{35.17}$$

Specialized form. The question is equation (35.17) with b the three outcomes of equation (35.14), taken term by term. The generator-running term Π_{g} is evaluated from the integrated path on each draw: it is the outage's and the generator's survival to the first time T_1 reaches the limit, if it does. The restart's rise and the headroom are its inputs, the latter by equation (28.19).

Building blocks. The terms of equation (35.15); Π_{g} on the draws; the headroom's spread.

The numbers. Table 35.6 is the split. Each row is a path; the columns give the crossing probability per outage on each route and the gap, with the test and without it. Start failure carries four fifths on both routes. The two terms sum to the total exactly.

Table 35.6: The crossing probability per outage, by path. State route: normal on $\mathbf{z}$; exact: Beta, gamma, reweighted hall.

	With the step test			Before the step test		
Path	State	Exact	Gap	State	Exact	Gap
Start failure	0.009785	0.007738	0.002047	0.010779	0.008548	0.002231
Running failure	0.002367	0.002086	0.000281	0.002603	0.002301	0.000302
Generator running	0	0	0	0	0	0
Total	0.012152	0.009824	0.002328	0.013382	0.010849	0.002533

With the generator running, the hall's maximum over the outage has the median 17.91°C and the 95th percentile 18.32°C (17.91 and 18.31), and the rise over the steady value has the median 0.28 K and the 95th percentile 0.66 K. The largest over 2×10^4 draws of a twelve-hour outage is 22.0°C. The restart raises the supply by 0.66 K, and at $\hat{\mathbf{z}}$ the hall has risen only 0.05 K by its end, since the time constant is 1.69 h; because the positive headroom of 1.77 kW then removes it at only 0.09 K/h, the hall peaks later, 0.42 K up at 3 h. The derate's own contribution is smaller. The headroom is 1.77 kW with spread 2.03 kW, negative with probability 0.19, and

the onset shortfall has the mean 0.02 K and the 95th percentile 0.15 K. A crossing with the generator running needs a sustained shortfall of about 14.4 kW through a twelve-hour outage, hall lag included, 8.0 spreads from the headroom's center and under a duration of probability e^{-6}: under this first-order loop it does not happen, and 0 in the table stands for less than 10^{-13} on the normal, against a margin of 9.4 K.

35.3.3 The entry the risk hangs on

Engineering question. Which entry of the posterior does the risk hang on?

General form. The variance budget of a derived quantity, equations (20.20) and (20.21):

$$\operatorname{Var}[f \mid \mathbf{y}] \approx \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} = \sum_i V_i + \sum_{i \neq j} [\mathbf{g}]_i [\mathbf{g}]_j \Sigma_{ij}, \qquad V_i = [\mathbf{g}]_i^2 \, \Sigma_{ii}. \tag{35.18}$$

Specialized form. The budget is equation (35.18) with $f = \ln \Pi$ of equation (35.14), $\mathbf{g}$ its gradient at $\hat{\mathbf{z}}$, and $\boldsymbol{\Sigma}$ of equation (35.12). The logarithm is chosen because Π spans a decade. Beside the delta form stands the correlation ratio on the draws, $\operatorname{Var}(\mathbb{E}[\ln \Pi \mid z_i])/\operatorname{Var}(\ln \Pi)$, on the exact posterior. The bins are the conditional expectation, $\mathbb{E}[\Pi \mid C_{\mathrm{th}} \in \mathrm{bin}, \mathbf{y}]$, from the same draws.

Building blocks. The gradient, by central differences of the path; $\boldsymbol{\Sigma}$; the weights of Section 35.2.6.

The numbers. At $\hat{\mathbf{z}}$, with the entries in the order of the state, the gradient is

$$\mathbf{g} = (-0.8415,\ 0.0754,\ 0.0016,\ 0.1549,\ -1.4490,\ -0.0833,\ 0.7546,\ 0.2103)^{\mathsf{T}}.$$

The entries of $\mathbf{g}$ for $\ln C_{\mathrm{th}}$ and ξ are close to closed forms: $-t_{\times}/\mu_{\mathrm{o}} = -1.451$ and, by the memoryless form, $p(1-p)/(p + \lambda/(\lambda + 1/\mu_{\mathrm{o}})) = 0.76$. Figure 35.6 and Table 35.7 give the budget. In the figure the horizontal axis is the share of $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$, one bar per entry, before the test and with it; in the table each row is an entry, and the columns give V_i and its share on the delta form, and the correlation ratio on the exact posterior. The total is 0.3226 with the test and 0.3667 before it, the shares summing to 1.025 and 1.030 with the cross terms the remainder. So the start entry holds 0.91 with the test and 0.80 before it. With $\Pi(\hat{\mathbf{z}}) = 0.00991$ the delta band of Π is $[0.00326,\ 0.03018]$, beside the band of the draws.

Table 35.7: The budget of $\ln \Pi$: the delta form on the state route and the correlation ratio on the exact posterior.

	With the step test			Before the step test		
Entry	V_i	Share	Exact ratio	V_i	Share	Exact ratio
ξ (start)	0.2945	0.913	0.868	0.2946	0.804	0.760
$\ln C_{\mathrm{th}}$	0.0141	0.044	0.048	0.0515	0.140	0.147
$\ln \lambda$	0.0111	0.034	0.055	0.0110	0.030	0.047
Q_1	0.0080	0.025	0.003	0.0111	0.030	0.014
g_1	0.0018	0.005	0.009	0.0076	0.021	0.003
T_s	0.0011	0.003	0.011	0.0017	0.005	0.0005

Table 35.8 reads the conditional crossing probabilities. Each row is a bin of the mass or of the conductance, and the columns give the probability of the bin and the crossing probability given that the state lies in it, on each route, before and with the test. A lighter hall crosses more often, by a factor of 1.7 across the mass bins before the test on either route; the conductance

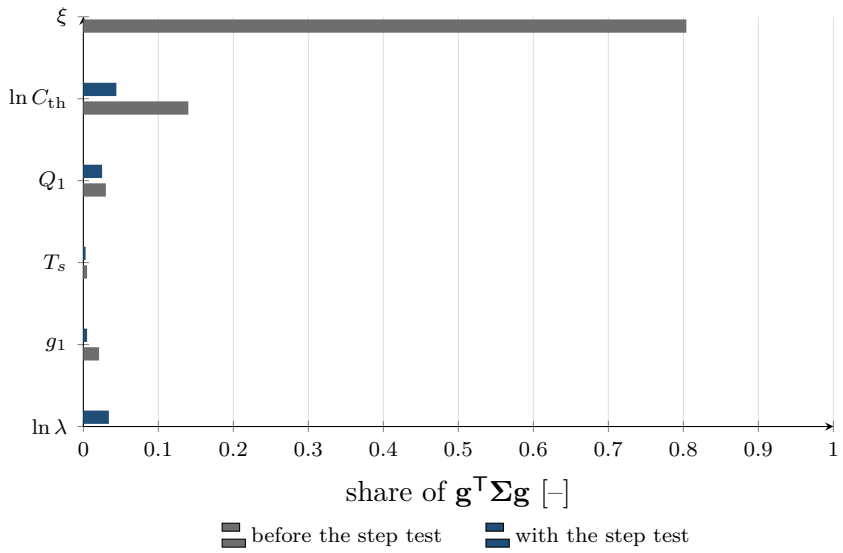


Figure 35.6: The variance budget of $\ln \Pi$ on the state route: $V_i / \mathbf{g}^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g}$ at $\hat{\mathbf{z}}$. The shares of Q_2 and c_{ch} are below 10^{-4}.

moves it by 6 %. The test empties the first mass bin and tightens the conductance. The start entry still sets the width.

Table 35.8: The crossing probability per outage given a bin of the state, $\mathbb{E}[\Pi \mid \mathrm{bin}, \mathbf{y}]$, with the bin's probability on the state route. The mass in kWh/K, the conductance in kW/K. The bins do not exhaust the state.

	Before the step test			With the step test		
Bin	Prob.	State	Exact	Prob.	State	Exact
$C_{\mathrm{th}} \in [6, 8)$	0.232	0.0172	0.0139	0.013	0.0160	0.0131
$C_{\mathrm{th}} \in [8, 10)$	0.497	0.0134	0.0109	0.676	0.0128	0.0105
$C_{\mathrm{th}} \in [10, 12)$	0.220	0.0102	0.0083	0.308	0.0105	0.0086
$G_1 \in [4.5, 5.2)$	0.273	0.0138	0.0112	0.040	0.0135	0.0109
$G_1 \in [5.2, 5.8)$	0.392	0.0133	0.0109	0.626	0.0124	0.0101
$G_1 \in [5.8, 7.0)$	0.291	0.0130	0.0105	0.334	0.0115	0.0093

35.3.4 The annual risk

Engineering question. What does a year of outages cost, and how well is that known?

General form. The expectation of a derived quantity and its quantiles, equations (2.10) and (2.11):

$$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}] = \int f(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}. \tag{35.19}$$

Specialized form. The question is equation (35.19) with $f = \nu_{\mathrm{o}} \Pi(\mathbf{z})$, the crossings per year, and $f = C_{\mathrm{cr}} \nu_{\mathrm{o}} \Pi(\mathbf{z})$ with $C_{\mathrm{cr}} = \mathrm{R\$}\,150\,000$ illustrative. One choice enters: the expected count per year is ν_{o} times the probability per outage whatever the count's law, so no Poisson assumption is used; a price that grew with the hours over the limit would need more.

Building blocks. Equation (35.16) and its band; $\nu_{\mathrm{o}} = 6$; the price.

The numbers. Table 35.9 gives the answer. Each row is a quantity; the columns give the value on each route, the gap and the bands. A crossing about once in fourteen years (seventeen), at about R$ 11 000 a year (R$ 8 800), is a small risk. The band spans a factor of 8.6 (9.4), and it comes mostly from the entry ξ, one failed start in forty tests and a prior.

Table 35.9: The annual risk of crossing. Price illustrative.

Quantity	State	Exact	Gap	Band, state	Band, exact
Crossings a year	0.0729	0.0589	0.0140	[0.0222, 0.1911]	[0.0152, 0.1430]
Risk [R$/yr]	10 937	8 842	2 095	[3 333, 28 671]	[2 287, 21 444]
Before the test, crossings	0.0803	0.0651	0.0152		

35.3.5 The three proposals and two others

Engineering question. Tank, faster restart, or generator maintenance?

General form. The comparison of expected consequences with a tail of the difference, equation (32.16):

$$\mathcal{R}(a') - \mathcal{R}(a), \qquad \Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}) = \int \mathbf{1}[f(\mathbf{z}) \geq f^*]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}. \tag{35.20}$$

Specialized form. The question is equation (35.20) with the loss $L(a, \mathbf{z}) = c_a + C_{\mathrm{cr}}\nu_{\mathrm{o}}\Pi_a(\mathbf{z})$, with Π_a equation (35.14) under action a, the marginal saving (32.1) $m_a(\mathbf{z}) = C_{\mathrm{cr}}\nu_{\mathrm{o}}(\Pi_0 - \Pi_a)(\mathbf{z})$, $f = m_a$ and $f^* = c_a$. Each action is computed on the same draws, so a difference closes, and a small difference is not simulation noise. Two alternative drawings are computed because the conclusion on the tank depends on the choices of step 2. With the pumps and fans on batteries the loop still takes the demand, so that $T_s(t) = T_s + (P/C_{\mathrm{lp}})t$ with the hall on its path, and equation (35.1) gives

$$T_1(t) = T_1^0 + \frac{P}{C_{\mathrm{lp}}}\Big(t - \frac{C_{\mathrm{th}}}{G_1}\big(1 - \mathrm{e}^{-G_1 t/C_{\mathrm{th}}}\big)\Big), \tag{35.21}$$

whose crossing time replaces $t_\times$ in Π_{s} and Π_{r}. And if the chiller also needs its restart when the grid returns, a no-cooling hall keeps rising for t_{r} after the outage, and $S_{\mathrm{o}}(t_\times)$ becomes $S_{\mathrm{o}}(t_\times - t_{\mathrm{r}})$.

Building blocks. The draws and averages of Section 35.2.6; Table 35.1.

The numbers. At $\hat{\mathbf{z}}$ equation (35.21) crosses at 5.15 h with 20 kWh/K and 8.76 h with 40 kWh/K. Figure 35.7 draws the options. The horizontal axis is the crossing probability per outage on a logarithmic scale; each row is an action, with a dot for the expected value on each route and a bar for the state route's band. Table 35.10 gives the numbers: each row is an action, the columns the probability per outage, the saving, the price and the probability that the saving at the state exceeds it, each on both routes.

The numbers split the claim that the cooling options buy nothing into what holds and what does not. Under the drawing of step 2, the tank and the restart act on the generator-running path, which never crosses, and they remove 0.4 and 0.5 % only through the slightly cooler hall that path hands to a running failure. Maintenance removes two fifths and the running program one tenth. That statement holds for this drawing, in which the circulation stops with the power. If the restart is also needed when the grid returns, the fast restart removes 7.7 % of a baseline 13 % higher, still short of its price. And if the pumps and fans stay up, the loop's mass becomes the hall's ride-through: at the pinned $C_{\mathrm{lp}} = 20$ kWh/K the batteries alone remove 68 %, and the

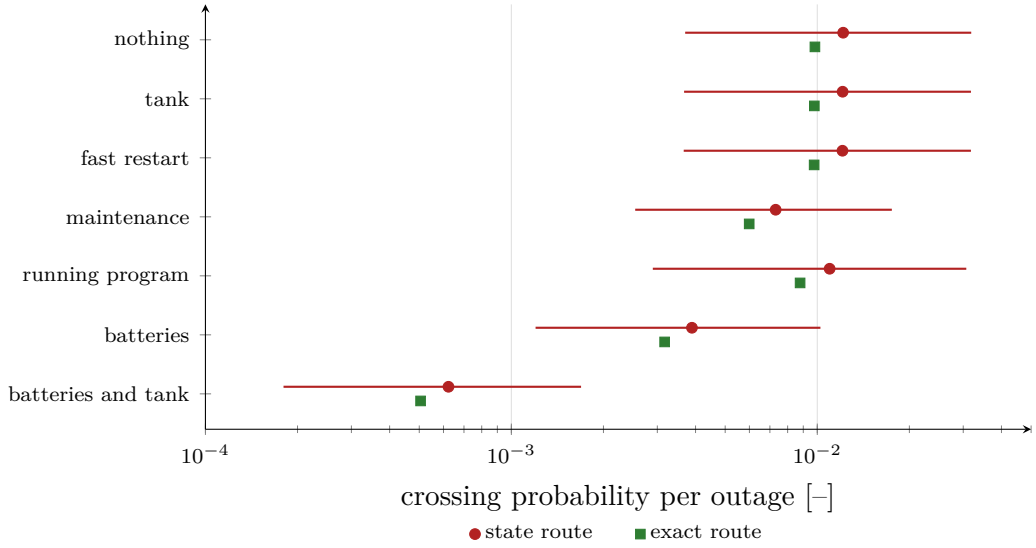


Figure 35.7: The crossing probability per outage under each action: expected values on each route, and the 95 % band of $\Pi_a(\mathbf{z})$ on the state route, seed 35. The bands of the first three rows nearly coincide because the actions hardly change them.

Table 35.10: The actions: crossing probability per outage, annual saving at $C_{\mathrm{cr}}\nu_{\mathrm{o}} = \mathrm{R\$}\,900\,000$, placeholder price, and probability that the saving at the state exceeds it, with the exact route's value in parentheses. The saving is the break-even price. Below the line, the same under the restart at the grid's return, against its own baseline.

	$\mathbb{E}[\Pi_a]$		Removed		Saving [R$/yr]		
Action	State	Exact	State	Exact	State	Exact	c_a; $\Pr(m_a > c_a)$
Nothing	0.012152	0.009824					
Tank	0.012102	0.009780	0.4 %	0.4 %	45	40	25 000; 0 (0)
Fast restart	0.012088	0.009767	0.5 %	0.6 %	57	51	4 000; 0 (0)
Maintenance, $p/2$	0.007311	0.005991	39.8 %	39.0 %	4 357	3 450	3 000; 0.58 (0.48)
Running program, $\lambda/2$	0.010976	0.008787	9.7 %	10.6 %	1 058	934	3 000; 0.01 (0.002)
Batteries on pumps and fans	0.003894	0.003170	68.0 %	67.7 %	7 432	5 989	12 000; 0.13 (0.06)
The same, with the tank	0.000623	0.000505	94.9 %	94.9 %	10 376	8 387	37 000; 0.006 (0.0003)
Restart at return, baseline	0.013773	0.011135					
Fast restart	0.012709	0.010269	7.7 %	7.8 %	958	779	4 000
Tank	0.013717	0.011085	0.4 %	0.4 %	50	45	25 000
Maintenance, $p/2$	0.008287	0.006792	39.8 %	39.0 %	4 937	3 909	3 000

tank then helps, taking that to 95 %. Batteries lengthen the crossing time and do not change the entry: every crossing still needs a generator failure, the start failure is 82 % (80 %) of what remains, a halving still removes 40 % (40 %), and ξ holds 0.94 of that budget. Like must be put with like: the batteries' 68 % is a change of drawing, the 40 % an assumed halving; a perfect start would remove 80 % (78 %), and a maintenance program would need a reduction of 85 % (87 %) to match the batteries. Among the priced options, the largest saving is a cooling measure, though not one of the two proposed, and it depends on the pinned mass: Section 35.4.3 varies it.

The tank's and the restart's verdicts do not depend on the placeholders: they save R$ 45 and R$ 57 a year (R$ 40 and R$ 51), below any plausible price. The maintenance's does. At the placeholder of R$ 3 000 it nets R$ 1 357 a year (R$ 450). Its break-even price is R$ 4 357 (R$ 3 450), so a price between the two values would divide the routes. At that placeholder it pays only if it removes more than 34 % (43 %) of the start-failure probability, against the half assumed, which no record here tests, and only if a crossing costs more than R$ 103 000 (R$ 130 000). At the true state its saving exceeds its price with probability 0.58 (0.48): the expected saving pays and the state may not. The batteries beat the maintenance at any price under R$ 6 075 (R$ 5 539) a year, where those two figures are their saving less the maintenance's net. Their R$ 12 000, like the R$ 3 000, is a placeholder, so this ranking is open until the plant has quotations. Table 35.11 carries the whole annual cost of each program.

Table 35.11: Expected annual cost of each program [R$/yr]: the risk at R$ 900 000 times $\mathbb{E}[\Pi_a]$ plus the placeholder price. The order holds at these placeholders only.

Program	Price	Risk, state	Total, state	Total, exact
Maintenance	3 000	6 580	9 580	8 392
Nothing	0	10 937	10 937	8 842
Running program	3 000	9 878	12 878	10 908
Fast restart	4 000	10 879	14 879	12 790
Batteries on pumps and fans	12 000	3 505	15 505	14 853
Tank	25 000	10 892	35 892	33 802
Batteries and tank	37 000	561	37 561	37 455

35.3.6 What to measure

Engineering question. What to measure next?

General form. The preposterior (2.17) and the value of sample information, equation (31.11), with its ceiling (31.16):

$$\mathrm{EVSI}(a) = \min_{a'} \mathcal{R}(a') - \mathbb{E}_{\tilde{y}_a|\mathbf{y}}\Big[\min_{a'} \mathbb{E}\big[L(a', \mathbf{z}) \mid \mathbf{y}, \tilde{y}_a\big]\Big] \le \mathrm{EVPI}. \tag{35.22}$$

Specialized form. The candidates are equation (35.22) with $\tilde{y}_a$ the outcomes of n more start tests, one a month, and the step test. The value is taken on one decision, maintain or not, at c_a = R$ 3 000, with L of Section 35.3.5; each value is compared on its own route. On the state route n start tests add the rank-one information $n\hat{p}(1-\hat{p})$ to the entry ξ, with $\hat{p} = 1/30$, so that $\Sigma_{a,\xi\xi} = (1/\Sigma_{\xi\xi} + n\hat{p}(1-\hat{p}))^{-1}$; the mean moves as in equation (31.15), and the outer expectation is a Gauss–Hermite sum over the innovation. On the exact route it is the sum over the Beta-binomial outcomes. The step test is not assumed: its $\mathbf{\Sigma}_a$ is the fitted (35.12), set against the state before it, and its value is the same form over 4 000 simulated moves of the shared entries' mean, of covariance $\mathbf{\Sigma}$ before less $\mathbf{\Sigma}_a$, at the saving before the test, seed 35. A

value is per year once the record is complete and is not discounted for the wait: sixty starts take five years, and until then the decision has only what has arrived.

Building blocks. $\Sigma_{\xi\xi}$; the moments of Π_{s} and Π_{r} on the draws; the saving of Section 35.3.5.

The numbers. Table 35.12 is the answer. Each row is a record; the columns give the expected spreads after it, the value to the maintenance decision, and the probability that the record reverses that decision, on each route. The ceiling on the same decision is the perfect-information value, R$ 461 a year on the state route and R$ 685 on the exact route, by draws.

Table 35.12: The candidate records, state route (exact route): the expected posterior spreads, averaged over the outcomes on both routes, the value to the maintenance decision at the placeholder R$ 3 000, and the probability of reversing it.

Record	$\sqrt{\Sigma_{a,\xi\xi}}$	sd of p	sd of Π	Value [R$/yr]	Reversal
Now	0.719	0.0309 (0.0230)	0.00752 (0.00559)		
12 more starts	0.657	0.0278 (0.0208)	0.00683 (0.00512)	43 (86)	0.11 (0.69)
36 more starts	0.569	0.0236 (0.0178)	0.00590 (0.00448)	151 (323)	0.24 (0.38)
60 more starts	0.509	0.0208 (0.0159)	0.00531 (0.00406)	215 (414)	0.29 (0.49)
Before the step test	0.719		0.00888 (0.00665)		
With the step test, this record	0.719		0.00752 (0.00559)	8, before it was run	0.03

The test on the mass was computed rather than assumed. It halves the spread of $\ln C_{\mathrm{th}}$ on this logger, 0.167 to 0.082, and it also moves entries it was not aimed at: T_1^0 goes from a spread of 0.48 to 0.087 K. It takes the budget's second entry from a share of 0.14 to 0.04, and the width of Π falls by a sixth. The start tests are slower: to halve the spread of ξ takes 180 more starts, fifteen years of monthly tests. So the dominant entry is again the one a record pins slowest, as Chapter 15 found of the share. The start tests' value is small for a reason that should not be misread: the plant sits near indifference at the placeholder price, with a net of R$ 450 on the exact route, so the record may well reverse the decision, with probability 0.29 (0.49) after sixty starts, and yet each reversal saves little. The exact value R$ 414 is close to that net. The decision is sensitive to these records, and the stakes of getting it wrong are small.

35.4 The checks (step 5)

The checks test the route, the record, the path model and what the posterior leaves out.

35.4.1 The state route against the exact one and the simulation

Table 35.13 collects the agreement. Each row is a quantity; the columns give the state route, the other route and the gap, and what the other route is. The largest gap is on $\bar{p}$: the normal on the logit gives 0.0417 against the record's exact 0.0333, 25 %, and the crossing probability inherits it, at 24 %. The gap does not change the order of any action or the sign of any net saving at the listed prices, but it does separate two break-even values. The analytic route of Section 35.3.1 sits 0.6 % (0.7 %) below the integrated one, the warming before a running failure being what it leaves out. Simulating everything, duration, start and running outcomes and the transient, on the exact posterior gives 0.00938 for 4×10^4 outages, seed 35, and 0.0089 to 0.0106 over eight other seeds, so that the differences of a few 10^{-4} that a simulation of this size shows are its noise, with standard error about 5×10^{-4}.

Table 35.13: The checks across routes, with the step test.

Quantity	State	Other	Gap	Other route
$\bar{p}$	0.04167	0.03333	0.00834	Beta(2, 58)
$\mathbb{E}[\Pi_{\mathrm{s}}]$	0.23484	0.23214	0.00270	reweighted hall
$\mathbb{E}[\Pi]$	0.012152	0.009824	0.002328	exact
$\mathbb{E}[\Pi]$, memoryless	0.012073	0.009755	0.002318	the closed form on each route
$\mathbb{E}[\Pi]$, simulated		0.00938		4×10^4 outages on the exact posterior
Shares of $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$, sum	1.025			the cross terms are the rest
Peak of T_1, generator up	18.05°C	17.63°C	0.42 K	the steady value at $\hat{\mathbf{z}}$; the rise is the restart's

35.4.2 The generated record against its model

The step test is generated, so its bearing on the conclusions needs stating. It is not unusual under its own model: $2\mathcal{L} = 65.86$ on $\nu = 73$ has $p = 0.71$, and $\ln \hat{C}_{\mathrm{th}}$ lies 0.39 prior spreads, or 0.79 posterior spreads, above the truth of its generation. Over seeds 36 to 55 with the same truth, $\hat{C}_{\mathrm{th}}$ runs from 8.06 to 10.28 kWh/K, $2\mathcal{L}$ from 44.4 to 91.7, and $\mathbb{E}[\Pi]$ from 0.0108 to 0.0152 (0.0087 to 0.0123). Over seeds 36 to 235 with the truth drawn from the priors in place of their center, numpy `default_rng`, seed 35, $\hat{C}_{\mathrm{th}}$ runs from 6.34 to 13.61, $\mathbb{E}[\Pi]$ from 0.0063 to 0.0206 (0.0050 to 0.0168), and the maintenance's saving from R$ 2 259 to R$ 7 403 (R$ 1 760 to R$ 5 895), under the R$ 3 000 placeholder in 2 % (15 %) of the records. So 0.01215 stands for one record, the second significant figure of every crossing probability rests on it, and the comparison of maintenance with nothing can reverse on another record. The tank's and the restart's savings stay at tens of reais on every such record, and the budget's top entry does not change.

35.4.3 What the posterior leaves out

The outage's pinned parameters. The budget's statement that ξ is the most consequential entry is true of the entries the state holds. The mean duration is pinned, and $\partial \ln S_{\mathrm{o}} / \partial \ln \mu_{\mathrm{o}} = t_{\times}/\mu_{\mathrm{o}} = 1.45$ is as strong a gradient as the mass's. On the same draws, $\mu_{\mathrm{o}} = 1.5$, 2 and 3 h give $\mathbb{E}[\Pi_{\mathrm{s}}] = 0.145$, 0.235 and 0.380. If the mean came from a record of n years of six outages, its variance would be about $1/(6n)$, and the term $1.45^2/(6n)$ would be 0.35 at $n = 1$, above the 0.29 of ξ, and 0.07 at $n = 5$. A site without its own outage log should carry that entry before reading ξ as the first concern.

The start prior. The start record is one failure in forty. With priors Beta($\frac{1}{2}, \frac{1}{2}$), Beta(1, 1) and Beta(2, 38) in place of Beta(1, 19), the exact posterior mean is 0.0366, 0.0476 and 0.0375, $\mathbb{E}[\Pi]$ is 0.0106, 0.0131 and 0.0108 against 0.0098, and the maintenance's saving is R$ 3 786, R$ 4 928 and R$ 3 881 against R$ 3 450. Every one of these priors keeps it above the placeholder. The pins are less forgiving: the saving falls to R$ 3 000 at a mean duration of 1.59 h (1.82 h), and at an outage rate of 4.1 (5.2) a year, against the pinned 2 h and 6.

The restart's and tank's claims. Hall conditions change the tank's value. With the hall steady at 25°C, a margin of 2 K, and the rest of the state as it is, the generator-running path crosses with probability 0.0004 and the tank takes it to 5×10^{-6}, while the no-cooling paths, at $t_{\times} = 0.62$ h, hold 0.038 per outage and move by under 1 %. So at 2 K the tank still buys little under this drawing. At a margin of 1 K the path crosses with probability 0.013 and the tank takes it to 0.001, and at 0.5 K from 0.15 to 0.025: that is where it matters, within about the 0.66 K of the restart's rise.

The loop's pinned mass. The batteries' figures rest on $C_{\mathrm{lp}} = 20\,\mathrm{kWh/K}$. With the pumps running, the hall is tied to a loop that warms at P/C_{lp}, which helps only if the loop is heavy enough. The batteries save R$ 3 368 (R$ 2 675) at 12 kWh/K, R$ 4 729 (R$ 3 785) at 14, and pass the halving's saving at about 13.5; at 10 kWh/K they remove 15 % (14 %), and at 8 they raise the risk by 5 % (6 %), since the loop then drives the hall past its isolated ride-through, 2.78 against 2.90 h at $\hat{\mathbf{z}}$. A loop mass is cheap to establish, and the conclusion on batteries waits on it.

The physics rows' tolerance. With $\sigma_p = 0.01\,\mathrm{K}$ per step in place of 0.002, the spread of $\ln C_{\mathrm{th}}$ is 0.113 against 0.082 on the same logger: the test's narrowing rests partly on how first-order the hall really is, which only a test on the hall can say.

35.5 The report (step 6)

The hall does not ride through the outage that matters. The average outage of 2 h is shorter than the 2.90 h a no-cooling hall holds, and still 23 % of those outages cross, so about one outage in eighty (0.012; 0.0098, one in a hundred, exact) takes the hall over 27°C, four fifths of them through a failed start. The derate and the restart, the two faults the chiller's owner watches, do not reach the limit alone. Under the drawing in which a start failure also stops the pumps, the tank and the faster restart remove 0.4 and 0.5 % and halving the start failure removes 40 %. That drawing decides the tank: with the pumps and coil fans on batteries and the loop at its pinned mass, circulation alone removes 68 %, and the tank then matters. The answers follow in the order they were asked. The thermal mass rests on a generated step test.

35.5.1 The answers and their bands

Table 35.14 is the report. Each row is a question in the engineer's words; the columns give the answer on the state route, the exact value, the gap, the band and qualification, and the chapters behind it.

The report line. In about one outage in eighty, 0.012 with a 95 % band from 0.0037 to 0.032 (0.0098, from 0.0025 to 0.024), hall 1 passes 27°C, about 0.07 crossings a year, costing about R$ 11 000 a year (R$ 8 800) at an illustrative R$ 150 000 each. Almost none of that comes from the chiller, whose derate and restart leave the hall under 19°C in nearly every outage. Four fifths of it is a generator that fails to start, and the remainder one that fails while running, in a plant where the pumps stop with the power and the hall's load continues. Halving the start failure, if a program can do that, removes two fifths. The numbers for the thermal mass come from a generated step test, over whose other records the probability per outage on the state route runs from 0.0108 to 0.0152, and from 0.0063 to 0.0206 if the mass's truth differs from its design figure.

35.5.2 The program and its limit

Two verdicts stand on the numbers: decline the tank and the restart as proposed, whose break-evens are tens of reais. The third is a condition and not a program: do the generator maintenance if its quoted price is under R$ 4 357 (R$ 3 450) a year and a reduction of the start-failure probability of more than 34 % (43 %) can be supported at the R$ 3 000 placeholder, keeping the monthly start tests; Table 35.11 is that condition with the placeholders written in. A quotation for batteries on the pumps and coil fans, and a figure for the loop's mass, are the further steps, since its break-even against the maintenance, R$ 6 075 (R$ 5 539) a year at

Table 35.14: The report of the coupled plant. Loss per crossing illustrative, action prices placeholders; the step test generated.

Question	Answer, state route	Exact	Gap	Band, qualification	From
How often?	0.01215 per outage	0.00982	0.00233	[0.0037, 0.0319] ([0.0025, 0.0238]); 0.0108 to 0.0152 over other test records, 0.0063 to 0.0206 with other truths	9, 20
Which path?	start 0.00979, running 0.00237, generator up 0	0.00774, 0.00209, 0	0.00205, 0.00028	hall's maximum with the generator up, median 17.9°C, 95th 18.3°C	25
Which entry?	delta shares: ξ 0.91, $\ln C_{\text{th}}$ 0.04	ratios 0.87, 0.05	–	before the test 0.80, 0.14; among the state's entries, with the mean duration pinned	20
Annual cost?	0.073 crossings, R$ 10 937 a year	0.059, R$ 8 842	0.014, R$ 2 095	[R$ 3 333, R$ 28 671] ([R$ 2 287, R$ 21 444])	25
Tank, restart, or maintenance?	not tank or restart; maintenance conditional, net R$ 1 357 a year; tank removes 0.4 %, restart 0.5 %	net R$ 450; 0.4 %, 0.6 %	R$ 907	placeholder price; pays if the price is under R$ 4 357 (R$ 3 450) and the reduction over 34 % (43 %); batteries on pumps remove 68 % at $C_{\text{lp}} = 20$, tank then matters; ranking open until priced	32
What to measure?	60 starts worth R$ 215 a year, ceiling R$ 461, reversing the decision with probability 0.29	R$ 414, R$ 685, 0.49	−R$ 199, −R$ 224, −0.20	small because the plant is near indifference, not because it is insensitive; keep start tests, the only check on the halving; an outage log first if there is none	31

20 kWh/K, would change which proposal is the right one. If the batteries are bought, the tank's standing changes too, and the question should be asked again with its own price.

Principle 35.1. *A coupled plant's risk lives on the path that can cross, and a cooling engineer's instinct is right about the cooling and may still be spending in the wrong place. Which paths exist is itself a modeling choice: a statement that an option buys nothing holds under the drawing that says so, and it is tested against the drawing that would break it. A time-indexed state is eliminated in time order with its shared entries last, which is the Kalman filter and the smoother with those entries in an augmented state, at a cost linear in the steps.*

Common mistake

Four errors recur with an outage study. Comparing the ride-through with the mean outage: 2.90 h against 2 h reads as safe, and 23 % of the no-cooling outages cross. Reading differences of a simulation as effects: at 4×10^4 outages the standard error is about 5×10^{-4}, ten times what the tank does; the options need common draws or integrals. Calling an option worthless without stating the drawing: the tank removes 0.4 % with the pumps stopped and takes the batteries' 68 % to 95 % with them running. And eliminating shared entries first in a long record: the factor fills, with 66 795 nonzeros against 2 175 and 8.2 million operations against 7 595, and the answer does not change.

35.6 The next measurement (step 7)

The candidates and their values are those of Section 35.3.6. What remains is which record to keep and why, and what the chapter leaves behind.

35.6.1 The record to keep

Keep the monthly start tests, for a reason other than their value of R$ 215 (R$ 414) a year or less to this decision. They are the only row on the assumed effect of the maintenance: without them the halving is never tested. On the state route, 36 tests leave a spread of at least 0.57 on ξ, and more at a lower p, so they will test a program only slowly, and they should be logged as demands with the program's date on them. Keep the outage log with durations if it exists and start it if not, since the pinned mean is worth as much as ξ in a young record and costs nothing. Run the step test once on a real logger, since its cost is a shed and its value lies in replacing the generated values of this chapter. The instrument matters, by refits of the same seed: with the hall's own sensor, $\sigma_T = 0.5\,\mathrm{K}$, the spread of $\ln C_{\mathrm{th}}$ only falls to 0.149; at 0.25 K, to 0.125; at 0.1 K, to the 0.082 reported; and at 0.1 K with a 15 kW shed, to 0.060. When rows arrive, the fit returns to step 3.

35.6.2 What the analysis showed

The first reading set the ride-through against the mean outage and found the hall safe, and the analysis finds about one outage in eighty crossing (one in a hundred on the exact route), all of it with no cooling at all. The risk is small, a crossing in about fourteen years (seventeen). Its causes are not where the cooling proposals looked: the derate and the restart are invisible against a margin of more than nine kelvin. For the plant this means declining the tank and the restart as proposed and pricing the generator maintenance against its break-even. It also means asking one further question, which the numbers forced: why do the pumps have to stop when the generator does, since, with a loop of the pinned mass, keeping them running removes more than a halving and nearly what a perfect start would, and it is what would make the tank useful. An engineer who takes this back should expect the next few years of start tests to leave the start probability nearly as uncertain as it is now, and should not wait for them.

35.6.3 Limits of the model

The derivations are exact under their assumptions, and six of them need checking before a result is used.

The generated test. The step test is generated, at a truth set at the priors' center, on an illustrative logger and tolerance. Section 35.4.2 gives the range. The logger is taken unbiased: a bias would move T_1^0 directly, by the same amount, without moving the shape from which C_{th} comes.

The hall's load through the outage. Crossing assumes the load is carried by batteries for the outage. Batteries that last less than $t_\times$ turn a thermal crossing into a loss of service, and the question to ask is then a different one.

The chiller on generator power. The capacity factor is Chapter 28's, measured with no outage in mind. Nothing here says the generator derates it further; a further derate moves a path that is 9 K from mattering, so it is not a limit on this conclusion unless it reaches the demand by a wide margin.

One lumped mass and a one-way loop. The hall has one time constant, and the loop takes the demand without feeling the hall's lag. The first understates the first minutes' rise. On the batteries' path the second overstates the loop's rise, since the halls' masses would absorb part of the heat, while Π_{r} there starts from the steady hall and drops the restart's warming; these act in opposite directions and are not netted here, and the pinned C_{lp} matters more than either (Section 35.4.3).

The options' effects. Halving p or λ is assumed, and independent of the state. Common cause between the start failure and the outage itself, a storm or a fuel fault, is not drawn: the start tests are not demands under an outage, and a higher conditional p would raise every number here and the maintenance's value with them.

The prices. The R$ 150 000 is illustrative, and the action prices are placeholders with no source, as is the halving. The break-evens are what carry over, and they are conditional on C_{cr}, μ_{o} and ν_{o}.

The next chapter brings the method together in one integrated assessment.

35.7 What to retain

Definition 35.1 (Time-ordered elimination)**.** On a state of time-indexed entries with a band of physics rows and a border of m shared entries, Cholesky factorization in time order with the shared entries last has nnz($\mathbf{L}$) linear in the steps, and its forward pass and back substitution are the Kalman filter and the Rauch–Tung–Striebel smoother on the augmented state ($T_{[j]}$, shared). With the shared entries first the factor is full and the cost is cubic. Both orders give the same posterior.

Four conclusions carry beyond the data center. A coupled question is answered by the path that crosses, so a large margin on one subsystem can remove it from the discussion altogether. An average against a threshold is a mistake that a survival curve corrects in one line: the mean outage was shorter than the ride-through, and a quarter of outages were not. A claim that an option buys nothing belongs to a drawing, and a study that tests the drawing can reverse its own recommendation; here the pumps' power supply, at the pinned loop mass, moves the largest saving among the priced options. And a record's value is computed: the step test halved a spread on one instrument and not on another, and it took one factorization, in the right order, to know that. Table 35.15 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 35.15: The terms of the coupled plant's report and where each is defined.

Term	Expression	Units	Meaning	Defined in
Transient	$C_{\mathrm{th}}\dot{T}_1 = Q_1 - G_1(T_1 - T_s(t))$	kW	hall through its mass	35.1.3
Crossing time	$t_\times = (T_{\mathrm{lim}} - T_1^0)C_{\mathrm{th}}/Q_1$	h	ride-through	35.1.3
Duration survival	$S_{\mathrm{o}}(t) = \mathrm{e}^{-t/\mu_{\mathrm{o}}}\mathbf{1}[t < 12\,\mathrm{h}]$	–	outage outlasts t	35.1.3
Augmented state	$\mathbf{x}_j = (T_{[j]},\ g_1,\ Q_1,\ T_s,\ \ln C_{\mathrm{th}})^{\mathsf{T}}$	varies	shared entries in the filter	35.2.4
Filter	equations (35.9) and (35.10)	varies	forward elimination	35.2.4
Smoother	equation (35.11)	varies	back substitution	35.2.4
General forms				
Predictive	$\int p(y_{\mathrm{new}} \mid \mathbf{z})p(\mathbf{z} \mid \mathbf{y})\,\mathrm{d}\mathbf{z}$	–	crossing per outage	35.3.1
Budget	$\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$, $V_i = [\mathbf{g}]_i^2\Sigma_{ii}$	–	entry the risk hangs on	35.3.3
Expectation	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$	varies	annual risk	35.3.4
Change	$\mathcal{R}(a') - \mathcal{R}(a)$; a tail of it	R$/yr	options	35.3.5
Value	$\mathrm{EVSI}(a) \le \mathrm{EVPI}$	R$/yr	candidates	35.3.6
Specialized forms				
Crossing per outage	$\Pi = p\Pi_{\mathrm{s}} + (1-p)(\Pi_{\mathrm{r}} + \Pi_{\mathrm{g}})$	–	three outcomes	35.3.1
Memoryless form	$(p + (1-p)\lambda/(\lambda + 1/\mu_{\mathrm{o}}))\mathrm{e}^{-t_\times/\mu_{\mathrm{o}}}$	–	analytic route	35.3.1
Circulating rise	$T_1(t)$ of equation (35.21)	K	pumps on batteries	35.3.5

35.8 Exercises

Exercise 35.1. *Formulation.* (a) Compute $t_\times$ for $T_1^0 = 18°\mathrm{C}$, $C_{\mathrm{th}} = 8\,\mathrm{kWh/K}$, $Q_1 = 30\,\mathrm{kW}$. (b) Compute the probability that one outage crosses by a start failure at $p = 0.02$ and $\mu_{\mathrm{o}} = 3\,\mathrm{h}$. ((a) $9 \times 8/30 = 2.4\,\mathrm{h}$. (b) $0.02\,\mathrm{e}^{-0.8} = 0.0090$.)

Exercise 35.2. *Derivation.* Derive the crossing probability given no cooling from the exponential duration, and say what the cap at 12 h changes. (The hall crosses if and only if $D_{\mathrm{o}} > t_\times$; $\Pr(\min(E, 12) > t) = \mathrm{e}^{-t/\mu_{\mathrm{o}}}$ for $t < 12$ and zero after. The cap changes only ride-throughs above 12 h, where the probability is under $\mathrm{e}^{-6} = 0.0025$.)

Exercise 35.3. *Derivation.* Derive the probability that the generator fails during the outage, $\lambda/(\lambda + 1/\mu_{\mathrm{o}})$, and the restart's rise. ($\int_0^\infty \lambda\mathrm{e}^{-\lambda u}\mathrm{e}^{-u/\mu_{\mathrm{o}}}\,\mathrm{d}u = \lambda/(\lambda + 1/\mu_{\mathrm{o}})$, with $1/\mu_{\mathrm{o}} = 1/2$ here; at $\lambda = 4/880\,\mathrm{h}^{-1}$, 0.0090. The rise is $Pt_{\mathrm{r}}/C_{\mathrm{lp}} = 53 \times 0.25/20 = 0.66\,\mathrm{K}$, from the loop's balance.)

Exercise 35.4. *Derivation.* A state of $n_t + 1$ time-indexed entries has a tridiagonal band and m shared entries. Count $\mathrm{nnz}(\mathbf{L})$ in the two orders and verify the second row of Table 35.3. (Time order: n_t columns of $m+2$ nonzeros, one of $m+1$, and $m(m+1)/2$ in the shared triangle; at $n_t = 360$, $m = 4$: $2\,160 + 5 + 10 = 2\,175$. Shared first: full, $n(n+1)/2$ with $n = 365$: 66 795.)

Exercise 35.5. *Application.* A hall runs at 25°C, a margin of 2 K, the rest of the state as in this chapter, on the state route. Recompute the paths, and say whether the tank now matters. ($t_\times$ falls to 0.62 h and $\mathbb{E}[\Pi_{\mathrm{s}}]$ rises to 0.733, so the no-cooling paths hold 0.0380 per outage. The generator-running path crosses with probability 0.0004, and the tank takes it to 5×10^{-6}, for a total change from 0.0384 to 0.0378, about 1.5 %. At 2 K it does not yet matter; at a margin of 1 K the running path has 0.013 and the tank 0.001, and there it does.)

Exercise 35.6. *Application.* A second generator in 1-of-2 joins the plant, with the shared-cause share of Chapter 15, which has the state route's mean 0.155 and the exact mean 0.131, at the p of this chapter. Compute the start-failure probability and the crossing probability per outage, with the running path unchanged. (With both start failures: $\beta_{\mathrm{cc}}p + (1-\beta_{\mathrm{cc}})^2p^2$ in expectation over p, with the share at its mean, 0.0084 (0.0056) against 0.0417 (0.0333); $\mathbb{E}[\Pi] = 0.0044$ (0.0034) against 0.0122 (0.0098), about 65 % removed, against the 40 % of a halving; the running failure is left as one generator's, which understates the gain.)

Chapter 36

Integrated Reliability Assessment

36.1 A third hall and a signature due

A third hall is being commissioned on the loop, and the commissioning engineer must sign the handover. The clause of Chapter 27 asks for a conductance of at least 5.0 kW/K with probability 0.95, the rating of Chapter 28 is 25°C, and the design load is 60 kW. The record on the desk is one afternoon: electric heaters at three loads, with a metered load, a reading of a new and uncalibrated hall sensor and a reading of the supply at each. At the top load the meter shows 60.07 kW and the hall sensor 25.402°C, 0.40 K over the rating.

The contractor has an answer. The line through the lowest and highest points has the slope $(25.402 - 16.572)/(60.07 - 19.88)$, a conductance of 4.55 kW/K, 91 % of the design. Its intercept is 12.20°C while the supply readings average 11.98°C, so the new sensor reads about 0.22 K high, the hall is really at about 25.18°C, and that, says the contractor, is close enough to a rating to sign with a note. Two points do not say how well either number is known, what the note should promise, or what the hall's cooling risks in a year while everyone looks at the sensor. The question of this chapter is

whether the hall meets its design, whether it stays under its limit at the design load, whether the new sensor is lying, how available its cooling is, what it risks in a year, what to measure next and what load to accept.

The chapter runs the method once, from step 0 to step 7, on one record, and cites the derivations. Nothing in the method is new. What is new is only that all its questions are asked of one posterior, with one further fit, the drawing without the bias, for the Bayes factor.

Jargon: commissioning test, derated load, redundancy of a row

A *commissioning test* is the record taken on a new plant before handover, here at loads held by electric heaters. A *derated load* is a load accepted below the design load because the assurance at the design load is not met; it is the assurance load Q_κ of Chapter 28 at a stated level. The *redundancy* of a row is the redundancy number of Chapter 22, one minus the row's leverage, and the redundancies of all the rows sum to the rows less the entries.

Key idea

Seven steps turn a handover into several answers of different standing. A requirement can be decided while a margin is open and a sensor undecided, and the largest money in the report can rest on none of the test's rows. Each answer is read from one posterior with the assumption and the record it holds under, and a next measurement is chosen for a question, not for the plant.

36.1.1 The plant and the record (step 0)

Figure 36.1 draws the plant. The green box on the left is the supply, free, with the prior carried from Chapter 18. The solid edge from it is the path of hall 3, with log-conductance g_3, $G_3 = \mathrm{e}^{g_3}$, to the blue box, the hall, whose temperature is a function of the free entries on the reduced state of Chapter 17. The green boxes along the top are the three test loads; the outlined boxes joined by dashed edges are the rows that observe them; the sensor carries its bias b_3. The separate block at the bottom is the cooling: two CRAC units in 1-of-2 with a shared cause, red, which no row of the test touches.

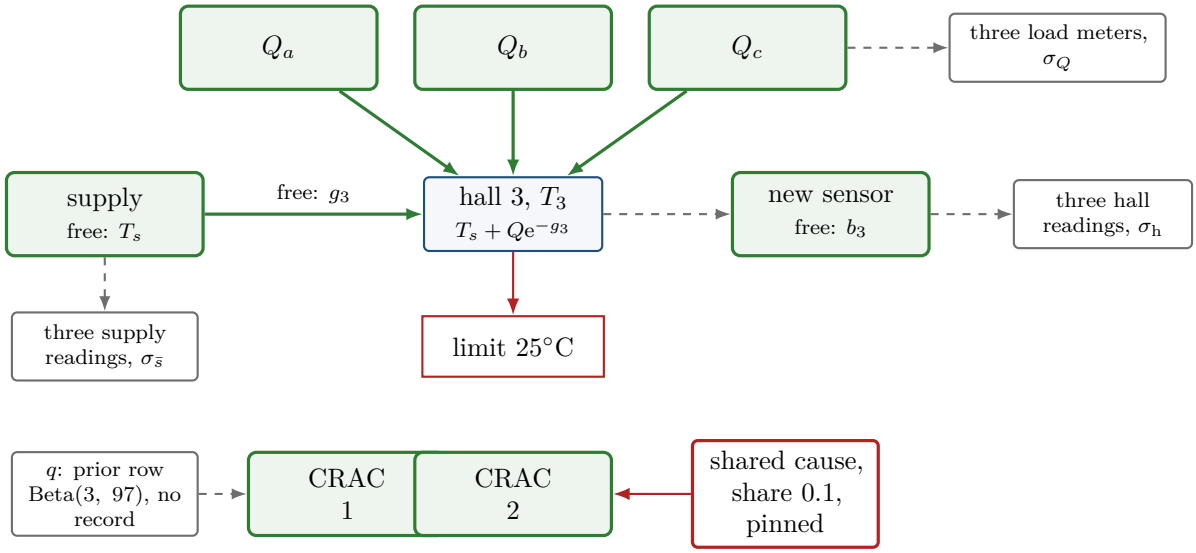


Figure 36.1: Hall 3's commissioning graph. Free: the supply, the path's log-conductance, the three test loads, the new sensor's bias and the CRAC units' shared unavailability (green); the hall's temperature is a function of them (blue). Observed (dashed edges): the records. The red box is the event asked about and the red-bordered box the pinned shared cause. The single arrow from the cause stands for one to each unit.

Table 36.1 is the record. Each column is a load; the rows give the reading of the meter, of the hall sensor and of the supply. Each hall and supply figure is the mean of ten readings. The record is generated and illustrative: numpy `default_rng`, seed 36, drawing the three hall errors, then the three load errors, then the three supply errors, about loads of 20, 40 and 60 kW, from a plant set below its design, a conductance under 5.0 kW/K and a nonzero bias. That construction carries the first answer of this chapter; Section 36.4.2 states the true values, tests the answers against them, and runs other records from the same plant and from plants at the design.

Table 36.1: The commissioning record (generated, seed 36; generator's values, used unrounded).

Row	Load a	Load b	Load c	Spread
Load meter [kW]	19.88159	40.55940	60.06655	$\sigma_Q = 0.5$ kW
Hall sensor [°C]	16.57235	21.05013	25.40187	$\sigma_h = 0.5/\sqrt{10} = 0.158$ K
Supply sensor [°C]	12.03405	12.00055	11.91477	$\sigma_{\bar{s}} = 0.3/\sqrt{10} = 0.095$ K

Two pieces of the plant have no test behind them. A CRAC unit's unavailability q has only the design file's Beta(3, 97), mean 0.03, and the share of its failures that are shared is pinned at $\beta_{cc} = 0.1$. The prices are illustrative: $C_L =$ R$ 15 000 per hour of cooling lost, and a penalty $C_T =$ R$ 50 000 a year if the hall is over 25°C at its accepted load.

36.1.2 The probabilistic questions (step 1)

Each target has a source, and one has none beyond this chapter. The conductance $G^{*} = 5.0\,\mathrm{kW/K}$ and the assurance $\kappa = 0.95$ are the acceptance clause of Chapter 27: 60 kW at a rise of 12 K, the N+1 condition. The limit $T_{\max} = 25^{\circ}\mathrm{C}$ is the rating of Chapter 28. The target of 10^{-3} on the unavailability of a hall's cooling is that of Chapter 24. The level $\kappa = 0.90$ for a derated load is this chapter's and illustrative: no contract sets it, so the answer gives the load at the customer's 0.95 of Chapter 28 beside it. Table 36.2 lists the six questions in the order they are treated, with the general form of the category, the form on this plant, the category and the subsection that answers each.

Table 36.2: The questions of the commissioning.

Engineering question	General form	Mathematical form	Category	Answered in
Does the hall meet its design conductance? Is N+1 true?	$\Pr(f(\mathbf{z}) \geq f^{*} \mid \mathbf{y}) \geq \kappa$	$\Pr(G_3 \geq G^{*} \mid \mathbf{y}) \geq 0.95$; the capacity at 25°C	prediction under a threshold (acceptance)	36.3.1
Will the hall stay under 25°C at 60 kW?	the same	$\Pr(T_3 < T_{\max} \mid \mathbf{y})$ at $Q = 60\,\mathrm{kW}$, the true temperature and not the reading	prediction under a threshold (margin)	36.3.2
Is the new sensor lying?	$p(z_i \mid \mathbf{y})$; $B_{AB} = p(\mathbf{y} \mid M_A)/p(\mathbf{y} \mid M_B)$	$p(b_3 \mid \mathbf{y})$; the bias entry against none	estimation; hypothesis comparison	36.3.3
How available is the hall's cooling?	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$; a tail	$\mathbb{E}[U_{\mathrm{pair}} \mid \mathbf{y}]$, its band, $\Pr(U_{\mathrm{pair}} < 10^{-3} \mid \mathbf{y})$	estimation of a derived quantity	36.3.4
What does the hall risk per year?	$\mathcal{R}(a) = c_a + \sum_e C_e \Pr(e \mid \mathbf{y}, a)$	the cooling-loss and thermal terms	estimation (risk)	36.3.5
What to measure next, and what load to accept?	$U(a) = \mathbb{E}_{\tilde{y}_a \mid \mathbf{y}}[u]$; the threshold inverted	$\boldsymbol{\Sigma}_a$ under each candidate; Q_κ now and after	design	36.3.6

36.1.3 The evidence as rows (step 2)

The modeling choices. Six choices make the rows. The test is taken with one CRAC out of service, as the N+1 clause requires: this makes the record speak to the clause, and with both running the conductance read would be of a different condition. The readings are steady at each load: this drops the hall's thermal mass, and it fails if the holds were short against the time constant, Section 36.6.3. The supply is one entry across the test, which the falling readings of Table 36.1 do not contradict and which Section 36.6.3 takes up. The bias is one constant, the same at all three loads, and the conductance does not depend on load: together these make the test a straight line with an intercept, and either failing would bend it, which three points barely test. The load meters' errors are independent, which is what lets σ_Q fall with the rows; a shared scale error would move the conductance one for one, and Section 36.3.6 says how large it would have to be. The CRAC pair has one q and a pinned share: the first is Chapter 15's form with its rates folded into q, and the second is tested by Table 36.8.

The rows. The state has seven entries,

$$\mathbf{z} = (g_3,\ Q_a,\ Q_b,\ Q_c,\ T_s,\ b_3,\ \xi_q)^{\mathsf{T}}, \qquad \xi_q = \ln \frac{q}{1-q}, \tag{36.1}$$

six for the test and the logit of q. The test's twelve rows, each divided by its spread, are

$$\mathbf{r}(\mathbf{z}) = \begin{bmatrix} (T_s + Q_k \mathrm{e}^{-g_3} + b_3 - y_{T,k})/\sigma_{\mathrm{h}}, \quad k = a,b,c \\ (Q_k - y_{Q,k})/\sigma_Q, \quad k = a,b,c \\ (T_s - y_{s,k})/\sigma_{\bar{s}}, \quad k = a,b,c \\ (g_3 - \ln 5)/0.3 \\ (T_s - 12.11145)/0.2801 \\ b_3/0.5 \end{bmatrix}, \tag{36.2}$$

with the design value as the prior's center, the supply's prior from Chapter 19, and a new sensor's prior spread of 0.5 K; the last is a judgment that Section 36.3.3 varies. The entry ξ_q has one row, the Beta(3, 97) as its negative log-probability term in the logit, in the form of Chapter 5. So there are 13 rows on 7 entries and the redundancy is six. The information matrix of these normalized rows is $\mathbf{J}^{\mathsf{T}}\mathbf{J}$; with unnormalized rows it would be $\mathbf{J}^{\mathsf{T}}\mathbf{W}\mathbf{J}$. The pattern is

$$\mathbf{J} = \begin{bmatrix} * & * & * & * & * & * & 0 \\ 0 & * & * & * & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & * & 0 & 0 \\ * & 0 & 0 & 0 & * & * & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & * \end{bmatrix} \begin{array}{c} \text{hall rows, each on its own load} \\ \text{load rows} \\ \text{supply rows} \\ \text{prior rows, one entry each} \\ \text{the term of } q \end{array} . \tag{36.3}$$

No row ties ξ_q to the test, so the covariance is block diagonal and no test of this hall can move the cooling's answer.

The competing drawing of the third question removes the sensor's bias: the entry b_3, its prior row and its term in the hall rows go, leaving 11 rows on 5 test entries.

36.2 The posterior of the state (step 3)

Every answer below reads from the normal computed here. The objects are those of Chapters 18 and 19; the derivations are not repeated.

36.2.1 The fit of the test

The normal equations $(\mathbf{J}^{\mathsf{T}}\mathbf{J})\Delta\mathbf{z} = -\mathbf{J}^{\mathsf{T}}\mathbf{r}$ on rows (36.2) start from the design and the priors, with the loads at their meters. Table 36.3 lists the three entries that move, with the objective and the largest entry of the gradient.

Table 36.3: Gauss–Newton on the test's six entries.

Iteration	g_3	T_s [°C]	b_3 [K]	$\mathcal{L}$	$\max_i \lvert\mathbf{J}^{\mathsf{T}}\mathbf{r}\rvert_i$
0	1.60944	12.11145	0	54.154	959
1	1.50688	11.98962	0.15352	0.80360	54.4
2	1.51147	11.98959	0.15038	0.67018	0.120
3	1.51148	11.98959	0.15045	0.67018	1.3×10^{-4}
5	1.51148	11.98959	0.15046	0.67018	3.9×10^{-9}

At the fit, $\mathrm{e}^{-\hat{g}_3} = 0.22058$, and the numerical Jacobian of rows (36.2), with the columns in

the order of equation (36.1) and without the last, is

$$\mathbf{J} = \begin{bmatrix} -27.83 & 1.395 & 0 & 0 & 6.325 & 6.325 \\ -56.51 & 0 & 1.395 & 0 & 6.325 & 6.325 \\ -83.82 & 0 & 0 & 1.395 & 6.325 & 6.325 \\ 0 & 2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 10.54 & 0 \\ 0 & 0 & 0 & 0 & 10.54 & 0 \\ 0 & 0 & 0 & 0 & 10.54 & 0 \\ 3.333 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 3.570 & 0 \\ 0 & 0 & 0 & 0 & 0 & 2 \end{bmatrix}. \tag{36.4}$$

The posterior of the state is

$$\hat{\mathbf{z}} = (1.51148,\ 19.9510,\ 40.5051,\ 60.0847,\ 11.98959,\ 0.15046,\ -3.47610)^{\mathsf{T}},$$

$$\boldsymbol{\Sigma} = \begin{bmatrix} \boldsymbol{\Sigma}_{\mathrm{t}} & \mathbf{0} \\ \mathbf{0} & 0.34364 \end{bmatrix},$$

$$\boldsymbol{\Sigma}_{\mathrm{t}} = \begin{bmatrix} 0.000734 & -0.004411 & 0.000529 & 0.005235 & 0.000071 & 0.006133 \\ -0.004411 & 0.220672 & 0.022824 & -0.005444 & -0.000626 & -0.054167 \\ 0.000529 & 0.022824 & 0.194556 & 0.029774 & -0.000149 & -0.012904 \\ 0.005235 & -0.005444 & 0.029774 & 0.231494 & 0.000305 & 0.026402 \\ 0.000071 & -0.000626 & -0.000149 & 0.000305 & 0.002865 & -0.002131 \\ 0.006133 & -0.054167 & -0.012904 & 0.026402 & -0.002131 & 0.065627 \end{bmatrix}. \tag{36.5}$$

The last entry is the term of the Beta(3, 97) in the logit, $\hat{\xi}_q = \ln(3/97)$ with information $3 \times 97/100 = 2.91$. The rest are read off: $\sqrt{\Sigma_{g_3 g_3}} = 0.02710$, $\sqrt{\Sigma_{T_s T_s}} = 0.05352$, $\sqrt{\Sigma_{b_3 b_3}} = 0.25618$, and the correlation of g_3 and b_3 is 0.88352. So $\hat{G}_3 = \mathrm{e}^{1.51148} = 4.5335\,\mathrm{kW/K}$. The curvature term of the nonlinear rows is small: with it the three spreads are 0.02714, 0.05352 and 0.25651.

At the minimum $2\mathcal{L} = 1.3404$. Set against χ^2_6, with six redundant rows, its upper tail is 0.969, which reads as an uncommonly clean record. The checks of Section 36.4.2 show that this reading is wrong as well as flattering.

36.2.2 The observability of the test

The decomposition (19.20) of Jacobian (36.4) says what the test can see. Table 36.4 gives, for each singular value, the right singular vector, one column per entry, and the singular value with the three prior rows removed. Every entry is observed, and none by the prior alone, since the smallest value without the priors is 1.868. The two largest directions are g_3 and T_s. The third is the bias, with $s = 4.69$, which is why its band is wide. The weakest is the three loads moving together with a little of the bias, carried by the meters' own weight $1/0.5 = 2$. The bias and the conductance trade along the line's intercept: raising b_3 lowers the line and flattens it, which raises G_3. The loads span 40 kW, so the slope holds g_3 tight, while the intercept, an extrapolation to no load, is what the bias has to share. A one-load test has no such lever (Exercise 36.5).

36.2.3 The drawing without the bias, and the exact route

Without the bias entry the same iteration converges to $\hat{g}_3 = 1.49750$, $\hat{G}_3 = 4.4705\,\mathrm{kW/K}$, $\hat{T}_s = 11.99439°\mathrm{C}$, with $2\mathcal{L} = 1.6821$, $\sqrt{\Sigma_{g_3 g_3}} = 0.01257$ and $\sqrt{\Sigma_{T_s T_s}} = 0.05289$. The conductance's

Table 36.4: The singular values and right singular vectors of the test's Jacobian, and the singular values of the reading rows alone.

s_j	g_3	Q_a	Q_b	Q_c	T_s	b_3	s_j, no prior rows
105.91	0.990	−0.004	−0.007	−0.010	−0.099	−0.096	105.86
18.99	−0.103	−0.014	−0.003	0.008	−0.994	−0.040	18.65
4.69	−0.086	−0.280	−0.066	0.138	0.052	−0.942	4.30
2.44	0	0.398	−0.816	0.418	0	0	2.44
2.002	0.034	−0.659	0.032	0.691	0	0.292	2.00
2.000	0	0.573	0.573	0.573	0	−0.126	1.87

spread falls by more than half, because the intercept no longer has anything to absorb.

The reported route is $\mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$. Beside it stands the exact posterior, which here has a short form. Given g_3, rows (36.2) are linear in the other five entries, so these integrate in closed form to a conditional normal, leaving $p(g_3 \mid \mathbf{y})$ on a grid of 16 001 points and no draws. The exact route gives g_3 the mean 1.51250, the median 1.51228 and the spread 0.02716, b_3 the mean 0.1561 and the spread 0.2565, and T_s the mean 11.98965 and the spread 0.05352. The exact posterior of q is Beta(3, 97). In step 4, parentheses hold the exact route's value and each column of a table is one route.

36.3 The answers (step 4)

The questions of Table 36.2 are read from the posterior of Section 36.2.1. A value at $\hat{\mathbf{z}}$ is labeled so: it is not a median of the exact route, whose medians stand in the parentheses where they are compared.

36.3.1 The acceptance

Engineering question. Does the hall meet its design conductance? Is N+1 true?

General form. The threshold form with its rule, equation (2.12):

$$\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}) = \int \mathbf{1}\big[f(\mathbf{z}) \geq f^*\big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \qquad \text{accept if} \geq \kappa. \tag{36.6}$$

Specialized form. The question is equation (36.6) with $f = g_3$, $f^* = \ln G^*$ and $\kappa = 0.95$, as in equation (27.8), an entry of the state read directly: $\Phi((\hat{g}_3 - \ln G^*)/\sqrt{\Sigma_{g_3 g_3}})$. The capacity at the limit is the derived quantity (28.14), $Q_{\text{cap}} = (T_{\max} - T_s)\mathrm{e}^{g_3}$, with gradient $(Q_{\text{cap}},\ 0,\ 0,\ 0,\ -\mathrm{e}^{\hat{g}_3},\ 0,\ 0)^{\mathsf{T}}$. No new assumption enters.

Building blocks. $\hat{g}_3$, $\hat{T}_s$ and the (g_3, T_s) block of $\mathbf{\Sigma}$ in equation (36.5).

The numbers. The fit is $\ln G^* - \hat{g}_3 = 0.09795$ below the requirement, 3.61 spreads, so that

$$\Pr(G_3 \geq 5.0 \mid \mathbf{y}) = \Phi(-3.615) = 1.5 \times 10^{-4} \quad (2.4 \times 10^{-4}).$$

The hall is not accepted as designed. At $\hat{\mathbf{z}}$ it delivers $4.5335/5 = 90.7\,\%$ of the design conductance (the exact median, 4.537, is 90.7 %), with the band [4.299, 4.781] ([4.305, 4.788]): the top of the band is 4.4 % short. The N+1 condition, 60 kW at a rise of 12 K, does not hold; at $\hat{\mathbf{z}}$ the rise is 13.23 K. The capacity at 25°C is 58.98 kW at $\hat{\mathbf{z}}$ with spread 1.60 kW (exact median 59.03 kW, spread 1.61 kW), against the 60 kW design. As for the prior: the design value was its

center, and the record moved the fit by 0.098 in the logarithm against a prior spread of 0.3. Section 36.4.3 refits without that row.

Figure 36.2 shows why the conductance is decided and the bias is not. The horizontal axis is g_3 and the vertical b_3; the ellipse is the 95 % region of their block of $\mathbf{\Sigma}$, its long axis the ridge of the correlation 0.88; the dot is $\hat{\mathbf{z}}$ and the cross the design with an unbiased sensor. The region is wide in the bias and narrow in the conductance, and on the conductance alone the design is 3.61 spreads away. Taken together with no bias, as the contractor's design file has it, it is 8.84 spreads away, since it also lies off the ridge. The square is the fit of the drawing with no bias, on the line $b_3 = 0$.

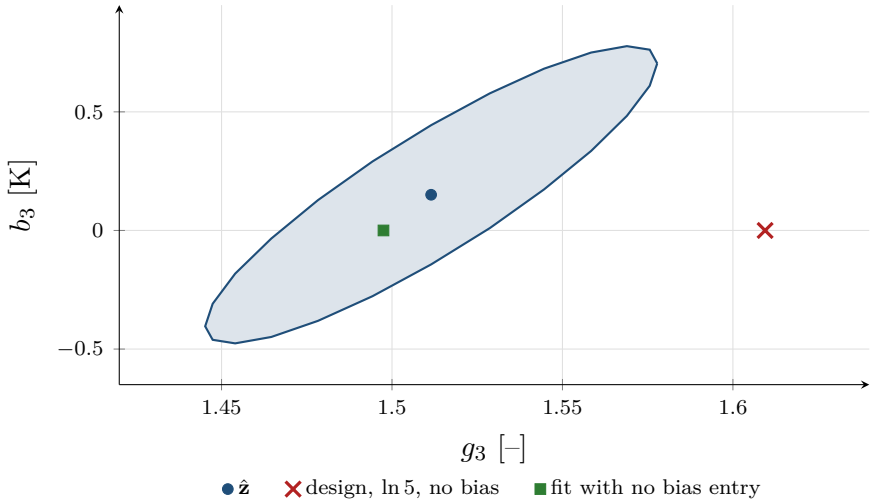


Figure 36.2: The posterior of (g_3, b_3) on the state route: the ellipse of 95 % probability, $\chi^2_2 = 5.99$, of the block of $\mathbf{\Sigma}$ in equation (36.5).

36.3.2 The margin at the design load

Engineering question. Will the hall stay under 25°C at 60 kW?

General form. The threshold form (36.6) on a derived quantity, whose tail on the normal of the state is

$$\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}) \approx \Phi\left(\frac{f(\hat{\mathbf{z}}) - f^*}{\sqrt{\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}}}\right), \qquad \mathbf{g} = \left.\frac{\partial f}{\partial \mathbf{z}}\right|_{\hat{\mathbf{z}}}. \tag{36.7}$$

Specialized form. The question is equation (36.7) with $f = -T_3$, $f^* = -T_{\max}$, and $T_3(\mathbf{z}; Q^*) = T_s + Q^* \mathrm{e}^{-g_3}$, equation (28.1), with Q^* pinned. The event is on the true temperature: b_3 is not in T_3, so a hot reading does not count against the hall nor a cool one for it. The gradient is equation (28.9) on this state, $\mathbf{g}_T = (-Q^* \mathrm{e}^{-\hat{g}_3},\ 0,\ 0,\ 0,\ 1,\ 0,\ 0)^{\mathsf{T}}$, and the tail is $\Phi((T_{\max} - T_3(\hat{\mathbf{z}}))/\sqrt{\mathbf{g}_T^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}_T})$, its sign taken up in $\mathbf{g} = -\mathbf{g}_T$. What the sensor would show is $T_3 + b_3$, with the gradient $\mathbf{g}_T + \mathbf{e}_{b_3}$.

Building blocks. $\hat{\mathbf{z}}$, the (g_3, T_s, b_3) block of $\mathbf{\Sigma}$.

The numbers. At $Q^* = 60\,\mathrm{kW}$, $\mathbf{g}_T = (-13.2349,\ 0,\ 0,\ 0,\ 1,\ 0,\ 0)^{\mathsf{T}}$, and

$$\mathbf{g}_T^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}_T = 0.128623 - 0.001876 + 0.002865 = 0.129612, \qquad \sqrt{\cdot} = 0.3600\,\mathrm{K},$$

in which the terms are those of g_3, of the cross term and of T_s: the conductance holds 99 % of it. Table 36.5 gives the answer around that load. Each row is a load; the columns give the temperature at $\hat{\mathbf{z}}$, its spread, the probability of staying under 25°C on each route and, for the exact route, the median temperature.

Table 36.5: The margin with one CRAC out of service, the test's condition: the true temperature at each load, on the state route and the exact route.

Q^* [kW]	$T_3(\hat{\mathbf{z}})$ [°C]	$\sqrt{\mathbf{g}_T^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_T}$ [K]	$\Pr(T_3 < 25)$, state	Exact	Exact median [°C]
57	24.563	0.342	0.899	0.902	24.552
58	24.783	0.348	0.733	0.742	24.773
60	25.225	0.360	0.266	0.276	25.213
65	26.327	0.390	3.3×10^{-4}	2.8×10^{-4}	26.315

With one unit out of service, the condition of the test, the hall probably fails the rating at the design load, with a probability of holding it of 0.266 (0.276). The odds are about three to one against, so this is not an even chance: the even chance sits at 58.98 kW (59.03 kW), one kilowatt lower. The rating is a looser requirement than the clause, a rise of 13.01 K and not 12, so it needs only 4.61 kW/K, which is inside the conductance's band. That is why the first answer is decided and this one is not. The band of the true temperature is [24.52, 25.93]°C.

The sensor would read 25.375°C with spread 0.169 K at $\hat{\mathbf{z}}$ (25.372 mean, 0.169). That band is less than half of the truth's, by the covariance: the reading's variance adds $\Sigma_{b_3 b_3} + 2\mathbf{g}_T^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{e}_{b_3} = 0.065627 - 0.166600$ to 0.129612, for 0.02864. The bias is large and negatively correlated with the hall's temperature, so it cancels in the row. A tight reading is therefore not a tight temperature: the test has fixed what the sensor would read at 60 kW, close to its own third point, far better than what the hall would do.

Figure 36.3 draws the two. In both panels the horizontal axis is the load and the vertical the temperature, with the right panel a zoom on the top of the left. The dots are the test's readings; the solid blue line is the fitted reading $T_s + Q/G_3 + b_3$ at $\hat{\mathbf{z}}$, and the green line is the true temperature $T_s + Q/G_3$ with its shaded 95 % band, which opens with load. The dashed line is the design, $G = 5$ with no bias on a 12°C supply. The red line is the limit, and the verticals mark 57 and 60 kW. The green line crosses the limit between the two marks, at 59 kW, the band straddles the limit at both, and the design line, which reaches the limit only at 65 kW, is under it at both.

36.3.3 The sensor

Engineering question. Is the new sensor lying?

General form. The marginal of an entry and the ratio of evidences, equation (2.14):

$$p(z_i \mid \mathbf{y}) = \int p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}_{-i}, \qquad B_{10} = \frac{p(\mathbf{y} \mid M_1)}{p(\mathbf{y} \mid M_0)}, \qquad p(\mathbf{y} \mid M) = \int p(\mathbf{y} \mid \mathbf{z}, M)\, p(\mathbf{z} \mid M)\, \mathrm{d}\mathbf{z}. \tag{36.8}$$

Specialized form. The question is equations (36.8) with $z_i = b_3$, M_1 the drawing with the entry and M_0 the drawing without it, nested at $b_3 = 0$. Each evidence is the Laplace form (24.12), with the same rows but one. Since M_0 has one row and one entry fewer, the log ratio splits, in the manner of equation (24.25), into

$$\ln B_{10} = \left(\mathcal{L}_0(\hat{\mathbf{z}}_0) - \mathcal{L}_1(\hat{\mathbf{z}})\right) + \tfrac{1}{2}\ln 2\pi - \tfrac{1}{2}\ln\frac{\det \boldsymbol{\Lambda}_1}{\det \boldsymbol{\Lambda}_0} - \ln\!\left(\sqrt{2\pi}\,\sigma_b\right), \qquad \sigma_b = 0.5\,\mathrm{K}. \tag{36.9}$$

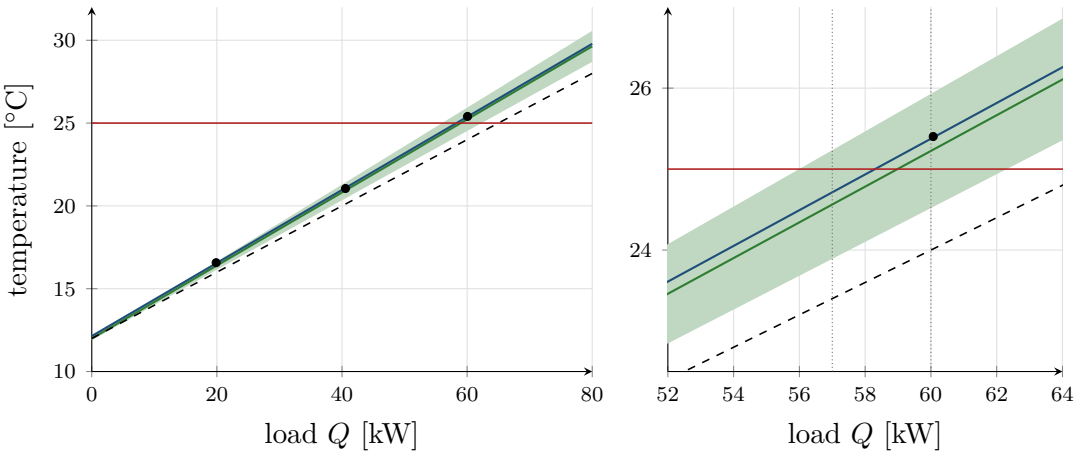


Figure 36.3: The test and its fit, from $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$ of equation (36.5). Dots: the readings of Table 36.1. Blue: $\hat{T}_s + Q\mathrm{e}^{-\hat{g}_3} + \hat{b}_3$. Green: $\hat{T}_s + Q\mathrm{e}^{-\hat{g}_3}$ with $\pm 1.96\sqrt{\mathbf{g}_T^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_T}$. Dashed: $12 + Q/5$. Red: 25°C. Right, a zoom, with dotted verticals at 57 and 60 kW.

The choice of prior spread enters here: it sets what the bias entry is charged, and Section 36.4.3 moves it. Whichever drawing is better supported, one choice remains, and it is an engineering one: the chapter keeps the entry, because the sensor is new and uncalibrated, and the drawing without it asserts a calibration nobody did.

Building blocks. Both fits of Section 36.2.3; their $\mathcal{L}$ and $\boldsymbol{\Lambda}$.

The numbers. The bias is 0.150 ± 0.256 K (0.156 ± 0.257), with the band $[-0.35,\ 0.65]$ K and $\Pr(b_3 > 0 \mid \mathbf{y}) = 0.72$ (0.73). Equation (36.9) is

$$\ln B_{10} = 0.1709 + 0.9189 - 1.3390 - 0.2258 = -0.475 \quad (-0.473),$$

so the improved misfit does not pay for the extra entry and the evidence leans, weakly, to no bias: at even prior odds $\Pr(M_1 \mid \mathbf{y}) = 0.38$. The drawing is undecided. Table 36.6 shows what the drawing costs. Each row is a drawing, with the fitted conductance and bias and two consequences: the probability of holding the rating at 60 kW with one unit out, on each route, and the assurance load at 0.90.

Table 36.6: The two drawings and what depends on them. $Q_{0.90}$ on the state route.

Drawing	$\hat{G}_3$ [kW/K]	$\hat{b}_3$ [K]	$\Pr(T_3 < 25)$, state	Exact	$Q_{0.90}$ [kW]
No bias entry	4.4705		0.0037	0.0036	57.29
Bias entry, kept	4.5335	0.150	0.266	0.276	56.99

Dropping the entry changes the margin sharply. The drawing with no bias reports 4.4705 kW/K with half the spread, and it holds the rating at 60 kW with probability 0.0037 (0.0036); the drawing with it, 0.266 (0.276). On the two drawings mixed at the even-odds weights the probability is 0.104 (0.108). All three fall short of any level in this chapter, so the verdict at 60 kW does not turn on the drawing, while the probability does. That is the plug-in error of Chapter 25 in another form: an entry fixed at a value is a tighter and possibly wrong answer. The assurance load scarcely moves, 56.99 against 57.29 kW, because the higher conductance of one drawing is paid for by the wider band of the other. Nor does the acceptance turn on it: the tail on the conductance is 1.5×10^{-4} with the entry and 3×10^{-19} without. The sensor being undecided matters to the margin and not to the clause.

36.3.4 The cooling's availability

Engineering question. How available is the hall's cooling?

General form. The moments, equation (2.10), the band, equation (2.11), and the threshold form (36.6):

$$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}] = \int f(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}. \tag{36.10}$$

Specialized form. The question is equation (36.10), and equation (36.6) with $f = -U_{\text{pair}}$ and $f^* = -10^{-3}$, where U_{pair} is equation (15.42) with the shared and independent terms written in q:

$$U_{\text{pair}}(\mathbf{z}) = \beta_{\text{cc}} q + \big((1-\beta_{\text{cc}})q\big)^2, \qquad q = \frac{1}{1+\mathrm{e}^{-\xi_q}}, \tag{36.11}$$

which drops the product of the two terms, 9×10^{-6} (5×10^{-6}) in expectation. Three assumptions enter. The share is pinned, which buys one entry and is tested below. The units' independent outages do not wait for one another: under equation (15.41) the second term needs a crew per unit, and with one crew the double outage's term doubles at small q, Section 36.3.4 below. And q is a prior: this answer rests on no record of this plant. On the state route the expectation is a Gauss–Hermite sum over ξ_q, sixty nodes, and since U_{pair} is monotone in ξ_q the band and tail are those of the entry, the tail $\Phi((\xi_q^* - \hat{\xi}_q)/\sqrt{\Sigma_{\xi_q\xi_q}})$. On the exact route they are the moments and the quantiles of the Beta.

Building blocks. $\hat{\xi}_q$, $\Sigma_{\xi_q\xi_q}$ in equation (36.5); $\xi_q^* = \text{logit}\ 0.00930 = -4.668$, the logit of the q at which equation (36.11) is 10^{-3}.

The numbers. Table 36.7 gives the answer: each row is a quantity, with the value on each route and the gap. The tail is $\Phi(-2.034) = 0.021$ (0.065), so the 1-of-2 pair does not meet the target of 10^{-3} at any level in use. The state route's mean sits above the exact one by 22 %, the skew of a normal on a logit that Chapter 35 also saw in the start probability. The gap crosses nothing here: no level separates 0.021 from 0.065.

Table 36.7: The unavailability of the CRAC pair at the pinned share $\beta_{\text{cc}} = 0.1$, from the prior on q.

Quantity	State	Exact	Gap
$\mathbb{E}[U_{\text{pair}}]$	4.84×10^{-3}	3.96×10^{-3}	8.8×10^{-4}
Hours a year, $t_{\text{yr}}\mathbb{E}[U_{\text{pair}}]$	42.4	34.7	7.7
Shared term	3.50×10^{-3}	3.00×10^{-3}	5.0×10^{-4}
Independent term	1.34×10^{-3}	9.6×10^{-4}	3.8×10^{-4}
Share of the mean from the shared term	0.72	0.76	−0.04
Median	3.73×10^{-3}	3.28×10^{-3}	4.5×10^{-4}
95 % band	$[1.05,\ 15.3]\times10^{-3}$	$[0.66,\ 11.2]\times10^{-3}$	
$\Pr(U_{\text{pair}} < 10^{-3})$	0.021	0.065	−0.044

The pinned share. Three quarters of the unavailability is the shared cause at $\beta_{\text{cc}} = 0.1$, and that is what the pinned number says, not what a record found. Table 36.8 moves it. Each row is a share, with the mean, the shared term's part of it and the tail, on each route. The shared term stops being the larger part below $\beta_{\text{cc}} = 0.043$ (0.037), a share of the size that Chapter 15 would not exclude for its pumps at the bottom of its band. Under one crew the double term is doubled: the mean is 6.18×10^{-3} (4.92×10^{-3}), the shared part 0.57 (0.61), the tail 0.016

(0.057), and the crossing moves up to 0.080 (0.069). Even with no shared cause at all, the tail is 0.54 (0.61): on this prior the target's failure does not depend on the shared cause alone.

Table 36.8: The pair under other values of the pinned share, with the exact route in parentheses.

β_{cc}	$\mathbb{E}[U_{\mathrm{pair}}]$	Shared part	$\Pr(U_{\mathrm{pair}} < 10^{-3})$
0	1.66 (1.19) $\times 10^{-3}$	0 (0)	0.537 (0.609)
0.03	2.61 (2.02) $\times 10^{-3}$	0.40 (0.45)	0.249 (0.327)
0.05	3.25 (2.57) $\times 10^{-3}$	0.54 (0.58)	0.127 (0.202)
0.1	4.84 (3.96) $\times 10^{-3}$	0.72 (0.76)	0.021 (0.065)
0.133	5.90 (4.88) $\times 10^{-3}$	0.79 (0.82)	0.007 (0.035)

36.3.5 The risk of a year

Engineering question. What does the hall risk per year?

General form. The expected consequence, equation (2.23):

$$L(a, \mathbf{z}) = c_a + \sum_e C_e \Pr(e \mid \mathbf{z}, a), \qquad \mathcal{R}(a) = \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]. \tag{36.12}$$

Specialized form. The question is equation (36.12) with two events and no price: the cooling term $C_{\mathrm{L}} t_{\mathrm{yr}} U_{\mathrm{pair}}(\mathbf{z})$, by equation (25.12), and the thermal term $C_{\mathrm{T}} \mathbf{1}[T_3(\mathbf{z}; Q^*) > T_{\max}]$. The block structure of pattern (36.3) lets them be added without a joint integral. Both prices are this chapter's illustrative figures, which is an assumption about money rather than about the plant, so any ordering of the terms carries a break-even. The thermal term also assumes a penalty that follows the accepted load, so derating moves it; this is a clause to be agreed, not something the test shows. And its probability is the one with one unit out, charged for the year: this takes one such period in a year to incur the penalty, which the 528 h (456 h) a year with one unit out of equation (36.11)'s independent term make likely. Charged only pro rata to those hours, the term falls to R$ 2 209 (R$ 1 885). The temperature with both units running is not in the state.

Building blocks. $\mathbb{E}[U_{\mathrm{pair}}]$ of Table 36.7, the tail of Section 36.3.2.

The numbers. Table 36.9 gives the terms: each row a term, with its value and band on each route. At a state the thermal term is either nothing or R$ 50 000, so it has a probability and no band of interest. At these prices the cooling term is 17 times the thermal one (14 times), and this rests on the prior, the pinned share and both prices. The ordering would reverse only at a loss rate under about R$ 865 per hour (R$ 1 043), a seventeenth of the one assumed, so the ordering is firm against the prices although each value is not. The cooling term has no capacity state, as Section 36.6.3 discloses.

Table 36.9: The annual risk of the hall at the design load, in R$ a year (prices illustrative).

Term	State	Exact	Band, state	Band, exact
Cooling, $C_{\mathrm{L}} t_{\mathrm{yr}} U_{\mathrm{pair}}$	635 961	520 656	[137 591, 2 009 353]	[86 910, 1 471 625]
Thermal, $C_{\mathrm{T}} \Pr(T_3 > 25)$, one unit out	36 679	36 201		
Total	672 640	556 857		

36.3.6 The next record and the load to accept

Engineering question. What to measure next, and what load to accept?

General form. The design form, equation (2.16), and the threshold inverted, equation (28.12):

$$U(a) = \int u\big(p(\mathbf{z} \mid \mathbf{y}, \tilde{y}_a)\big)\, p(\tilde{y}_a \mid \mathbf{y})\, \mathrm{d}\tilde{y}_a, \qquad Q_\kappa : \ \Pr\big(T_3(\mathbf{z}; Q_\kappa) < T_{\max} \mid \mathbf{y}\big) = \kappa. \tag{36.13}$$

Specialized form. The question is equations (36.13) with u the spread of T_3 at 60 kW and the load Q_κ after the record, on the state route: the root (28.13) with the block of $\mathbf{\Sigma}_a$ and a center that moves. The candidates are four. A calibration of the new sensor against a reference of $\sigma_{\mathrm{ref}} = 0.1$ K is the row $(b_3 - y_{\mathrm{c}})/\sigma_{\mathrm{ref}}$, rank one, so that equation (31.15) gives $\mathbf{m}_a(\eta_a)$ along one line and the expectation over the innovation η_a is a Gauss–Hermite sum. A fourth test at 80 kW is a hall row, a load row and a supply row with a new load entry, on a hall run about 4.6 K over its rating. A second supply sensor of $\sigma_{\bar{s}}$ is one row on T_s. A log of the CRAC units, n morning checks, is a binomial row on ξ_q, and it is the only candidate that reaches the cooling, by pattern (36.3). For the multi-row candidate the mean's move $\mathbf{m}_a \sim \mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma} - \mathbf{\Sigma}_a)$ is drawn, 10^5 draws, numpy `default_rng`, seed 36, with seeds 36 to 40 giving the ranges quoted. A move in Q_κ is valued on the illustrative $v_{\mathrm{L}} = \mathrm{R\$}\,400$ per kW-year of Chapter 31, as EVSI_Q of equation (31.14).

Building blocks. $\mathbf{\Sigma}$, the lines $\mathbf{j}_a$; the root (28.13) with $z_{0.90} = 1.2816$, $a = T_{\max} - \hat{T}_s = 13.0104$, $A = 0.99879$, $B = 13.01030$, $C = 169.2661$.

The load now. The root is $x = 12.5718$ K, so that $Q_{0.90} = 12.5718\,\hat{G}_3 = 56.99$ kW (57.02 kW). At the customer's level, $Q_{0.95} = 56.45$ kW (56.47 kW); at an even chance, 58.98 kW (59.03 kW). The gap between routes crosses nothing, and the two levels differ by half a kilowatt.

The candidates. Table 36.10 is the answer. Each row is a candidate; the columns give its expected spreads, the load with the center held at $\hat{\mathbf{z}}$, the expected load after the record, its 95 % band, the probabilities that the load after the record reaches 60 kW or falls below the present 56.99 kW, and its value. Figure 36.4 draws the first of these. The horizontal axis is the true temperature at 60 kW. Each row is a candidate, with its thick bar the 95 % band after the record with the center held, and its thin bar the 95 % range of where the center itself may land. The red line is the limit. The calibration has the shortest thick bar and the longest thin bar: it is the record that can move the answer most.

Table 36.10: The candidates, on the state route. The row marked known is the ceiling, the bias known exactly. Spreads are expected over the outcomes; the load at $\kappa = 0.90$ in kW. The value is EVSI_Q at R$ 400 per kW-year (illustrative).

Record	sd(g_3)	sd(b_3)	sd(T_3)	Center held	Expected	Band	≥ 60	< 56.99	Value [R$/yr]
Now	0.0271	0.256	0.360	56.99					
Calibration	0.0154	0.093	0.195	57.89	57.91	[55.35, 60.55]	0.059	0.247	365
Test at 80 kW	0.0180	0.217	0.241	57.64	57.65	[55.40, 59.96]	0.023	0.291	261
Second supply	0.0271	0.255	0.360	57.00	57.00	[56.84, 57.15]	0	0.486	1
Bias known	0.0127	0	0.155	58.11	58.13	[55.39, 60.99]	0.098	0.213	456

The calibration is the best of the three test records by the margin's spread, 0.360 to 0.195 K, ahead of the test at 80 kW at 0.241 K and above a hall at its rating. The second supply sensor buys nothing, as T_s is already at 0.054 K. But the load of 57.89 kW is the load after a calibration that lands exactly where it is expected, and a calibration does not. Figure 36.5 draws the load

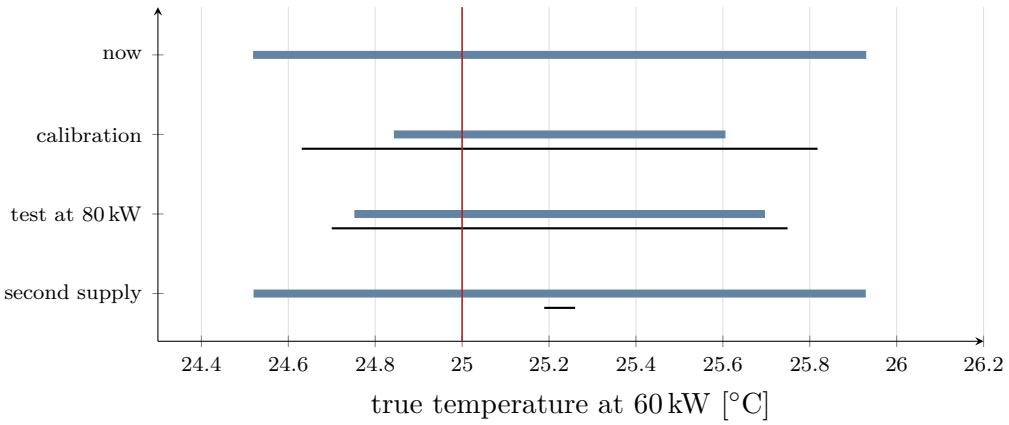


Figure 36.4: The true temperature at 60 kW under the candidates, by $\mathbf{g}_T^{\mathsf{T}}\boldsymbol{\Sigma}_a\mathbf{g}_T$. Thick: the 95 % band about $T_3(\hat{\mathbf{z}}) = 25.225$°C with the center held. Thin: $\pm 1.96\sqrt{\mathbf{g}_T^{\mathsf{T}}(\boldsymbol{\Sigma} - \boldsymbol{\Sigma}_a)\mathbf{g}_T}$, the 95 % range of the center after the record. Red: the limit.

as a function of the outcome. The horizontal axis is the calibration's reading y_{c}, the vertical the load $Q_{0.90}$ after it, with the shaded strip the reading's 95 % predictive band, centered on $\hat{b}_3$ with spread $\sigma_{\mathrm{ref}}\sqrt{s_a} = 0.275$ K, and the dashed lines at the present load and at 60 kW. A reading above 0.58 K returns the design load to the hall, with probability 0.059 (0.062). A reading below −0.04 K lowers the load below the present one, with probability 0.247 (0.240). The exact route's expected load is 57.95 kW, with the band [55.43, 60.58] kW, by quadrature over the reading's predictive. So the calibration decides the margin, in either direction, and the figure of 58 kW is its center and not a promise.

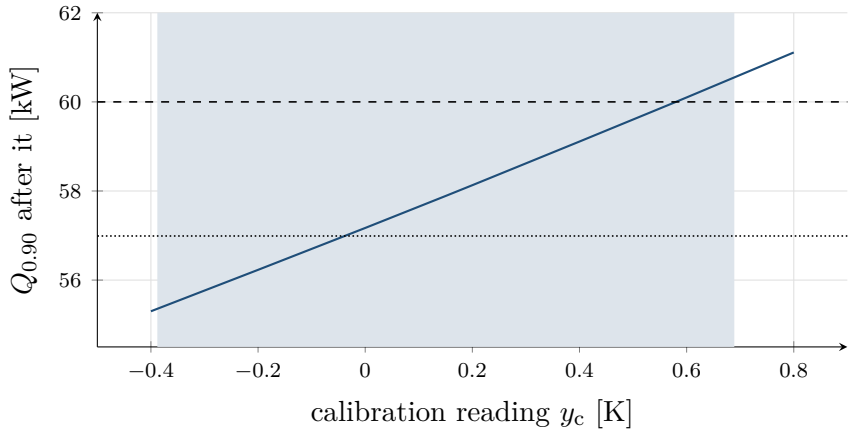


Figure 36.5: The assurance load at $\kappa = 0.90$ after a calibration reading, on the state route by equation (31.15). Shaded: the reading's 95 % predictive band, $0.1505 \pm 1.96 \times 0.275$. Dashed: 60 kW. Dotted: the present 56.99 kW.

The acceptance after a record. None of the candidates reopens the clause, and within this state that holds under every measurement, for a reason that needs no list. The posterior probability after a record averages back to the present one, $\mathbb{E}_{\tilde{y}_a}[\Pr(G_3 \geq G^* \mid \mathbf{y}, \tilde{y}_a)] = \Pr(G_3 \geq G^* \mid \mathbf{y})$, so by Markov's inequality any record reaches κ with probability at most $1.5 \times 10^{-4}/0.95 = 1.6 \times 10^{-4}$ (2.5×10^{-4}). For the calibration it is much less: it would take a reading of 1.67 K, $\eta_a = 5.53$, with probability 1.6×10^{-8}. That statement holds under this model and these rows, so what matters is what would break it. Two things outside the state

could. A shared scale error on the three load meters moves G_3: by refit, the meters would have to read low by 9.4 % to center the fit on the design (true load over metered, 1.104) and by 13.2 % to pass the clause (1.152). And a conductance that rises with load is not in the state; a temperature record at another load, such as the test at 80 kW, could show it. So the statement is: within this state, at most 1.6×10^{-4}; outside it, the meters' scale and a load-dependent conductance.

The cooling's record. The log moves what the test cannot. After $n = 1\,000$ checks, at the present center, $\sqrt{\Sigma_{a,\xi_q\xi_q}}$ falls from 0.586 to 0.177 and the cooling term's band narrows to [R$ 330 000, R$ 737 000]. The share is not in it: the log needs its simultaneous outages counted for that, as in Chapter 15. Its probability of showing the target met at 0.95 is 0.004 (0.006), since on the exact route it would need at most 2 down in 1 000 checks; a refit on the state route at 2 down gives 0.946, so that route needs at most 1, with a predictive probability of 2×10^{-4}, against the 0.004 of its rank-one form with the spread held. So the log narrows the largest number in the report, and it would very probably confirm the failure against the target.

The load to accept. The calculation of the derate is one of prices, so it is stated by its break-even. Moving from 60 kW to $Q_{0.90}$ gives up 3.01 kW and takes the thermal term from R$ 36 679 to R$ 5 000, a saving of R$ 31 679 (R$ 31 201, for 2.98 kW). It pays for any value of admitted load under R$ 10 537 per kW-year (R$ 10 471), twenty-six times the figure of Chapter 31; at $Q_{0.95}$, under R$ 9 638 (R$ 9 542). The derate is thus firm against that price. The level is not: at R$ 400 per kW-year and the R$ 50 000 the expected loss has its minimum at 54.98 kW (54.97 kW), where the probability is 0.996, so 0.90 is a rule chosen by the operator and not the action of these two illustrative prices, under which it would derate further. The level should be set by whoever signs: the load is 56.99 kW at 0.90 and 56.45 kW at 0.95.

36.4 The checks (step 5)

The checks test the routes, the record against its generator, the closed forms against the plant, and the preposterior against a refit.

36.4.1 The routes

Table 36.11 collects the agreement. Each row is a quantity, with the state route, the other route, the gap, and what the other route is. The test's answers agree closely because g_3 is tight: its normal hardly differs from the exact marginal. The cooling's answers do not, because q is a prior: a spread of 0.586 on a logit is wide. No gap crosses a level this chapter uses. The two values of the capacity event at 60 kW are the same event by equation (28.16), with two gradients on the state route and one value on the exact route: its state-route tails, 0.263 and 0.266, differ by what a delta form leaves out.

36.4.2 The generated record against its model

The record was generated at $G_3 = 4.6$ kW/K, $b_3 = 0.4$ K, $T_s = 12.0$°C, loads 20, 40, 60 kW. So the plant has been set below the clause, and the first answer rests on that construction: a test of the method, not a finding about a hall. Against those values the fit is low by 0.54 spreads in g_3, 0.97 in b_3 and 0.19 in T_s, while the drawing without the bias is 2.27 spreads off in g_3. The truth's capacity is 59.8 kW and its temperature at 60 kW is 25.04°C, so the hall truly misses the rating, by little.

The reading of the record as uncommonly clean does not survive. The rows' redundancies, one minus each row's leverage $[\mathbf{J}\boldsymbol{\Sigma}\mathbf{J}^{\mathsf{T}}]_{kk}$, are 3.31 on the nine reading rows and 2.69 on the

Table 36.11: The checks across routes.

Quantity	State	Other	Gap	Other route
$\Pr(G_3 \geq 5)$	1.5×10^{-4}	2.4×10^{-4}	-9×10^{-5}	exact marginal of g_3
$\Pr(T_3 < 25)$ at 60 kW	0.266	0.276	−0.010	exact
$\Pr(Q_{\text{cap}} \geq 60)$	0.263	0.276	−0.013	exact; the same event
$Q_{0.90}$ [kW]	56.99	57.02	−0.03	exact
$\ln B_{10}$	−0.475	−0.473	−0.002	exact evidences by the g_3 grid
$\ln B_{10}$	−0.475	−0.496	0.021	the ratio of densities at $b_3 = 0$, equation (24.22), on the normal
$\mathbb{E}[U_{\text{pair}}]$	4.84×10^{-3}	3.96×10^{-3}	8.8×10^{-4}	Beta moments
U_{pair} at $q = 0.03$	3.729×10^{-3}	3.7268×10^{-3}	2×10^{-6}	enumeration of the pair, equation (15.42) with the product kept

three prior rows, and those prior rows fit almost freely, contributing 0.39 of the 1.34. When the truth is drawn from the priors, $2\mathcal{L}$ does follow χ^2_6: 3 % of 400 records lie below 1.34. But with the plant held fixed, as here, the priors are not noise and 10 % of 400 records lie below it, with a median of 3.31. So the record is not unusual under its own generator, and a very small $2\mathcal{L}$ would be no virtue if it were: it is no evidence for the plant, since it says only that this record's errors are small.

Table 36.12 asks how far the answers move. Over seeds 36 to 235 at the same truth, each row is an answer, with its value on this record, its extreme values and its central 95 % range. The clause's verdict does not move: none of the 200 records accepts. Everything else does. The rating at 60 kW runs from nearly lost to nearly held, and 29.5 % of the records give a probability above one half. The load $Q_{0.90}$ runs from 53.8 to 61.3 kW, and it is at least 60 kW on 3 % of records. Those records would admit 60 kW on a hall whose true capacity is 59.8 kW, an error the 0.90 rule allows. The bias is undecided on this record, and 24 % of records give $\ln B_{10} > 1$ and 4 % above 3: the truth has a bias of 0.4 K, and this record, 0.97 spreads low in it, is one on which the bias is hard to see. The conclusion that the sensor is undecided belongs to this record, and the median $\ln B_{10}$ over the others is 0.04.

Table 36.12: The answers over 200 records from the plant that generated this one (seeds 36 to 235; state route, $\ln B_{10}$ exact).

Answer	This record	Extremes	Central 95 %
$\hat{G}_3$ [kW/K]	4.533	[4.293, 4.897]	[4.371, 4.841]
$\sqrt{\Sigma_{g_3 g_3}}$	0.0271	[0.0265, 0.0287]	[0.0266, 0.0283]
$\Pr(G_3 \geq 5)$	1.5×10^{-4}	$[7 \times 10^{-9}, 0.233]$	$[3 \times 10^{-7}, 0.122]$
$\hat{b}_3$ [K]	0.150	[−0.36, 0.86]	[−0.07, 0.72]
$\ln B_{10}$	−0.47	[−0.67, 5.19]	[−0.65, 3.43]
$\Pr(T_3 < 25)$ at 60 kW	0.266	[0.003, 0.980]	[0.021, 0.938]
$Q_{0.90}$ [kW]	56.99	[53.77, 61.30]	[54.83, 60.42]
$2\mathcal{L}$	1.34	[0.49, 14.47]	[0.72, 9.76]

The verdict on the clause also needs its operating characteristic, since it would look the same from a test that could not accept. With the same design and truth otherwise, the rule of Chapter 27 accepts 1.2 % of records from a hall truly at 5.0 kW/K and 20.8 % from a hall at 5.2 kW/K, over 4 000 seeds each, with sampling errors of 0.2 and 0.6 points. The bound $1 - \kappa$ of equation (27.22) supposes the prior's centers as the truth; there, with no bias, the rate is 3.1 ± 0.3 %. So this test accepts a good hall seldom. A rejection by this rule at a tail near 0.5 would say little; here the tail is 1.5×10^{-4}, and among 4 000 records from a true 4.8 kW/K none is accepted. The test has little power to accept, and it is still strong against a hall this far

below the clause.

36.4.3 The priors of the conductance and the bias

With the prior row on g_3 removed the fit is 4.530 kW/K, against 4.5335: the prior was not holding up the answer in either direction. Table 36.13 refits with the bias's prior spread moved. Each row is a spread, with the log Bayes factor against the drawing with no bias on each route, the fitted conductance and bias, the probability of holding the rating at 60 kW on each route, and the assurance load. The evidence never turns decisive, and the load hardly moves.

Table 36.13: Refits under the bias's prior spread. $Q_{0.90}$ on the state route.

σ_b [K]	$\ln B_{10}$, state	Exact	$\hat{G}_3$ [kW/K]	$\hat{b}_3$ [K]	$\Pr(T_3 < 25)$, state	Exact	$Q_{0.90}$ [kW]
0.1	−0.026	−0.026	4.4789	0.020	0.024	0.024	57.18
0.25	−0.157	−0.155	4.5054	0.084	0.142	0.148	57.02
0.5, kept	−0.475	−0.473	4.5335	0.150	0.266	0.276	56.99
1	−1.013	−1.012	4.5491	0.187	0.326	0.336	57.01
2	−1.659	−1.658	4.5543	0.200	0.344	0.355	57.02

36.4.4 The closed forms and the preposterior

Equation (36.11) against equation (15.42) is in Table 36.11; the one-crew form is in Section 36.3.4. The preposterior of the calibration was tested by refitting: 4 000 draws of the state from $\mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$, a simulated reading at each, and Gauss–Newton on the thirteen rows, numpy `default_rng`, seed 36. The mean of $Q_{0.90}$ is 57.90 kW against the 57.91 kW of the rank-one sum, the probability of reaching 60 kW is 0.054 against 0.059, with 0.053 to 0.059 over seeds 37 to 40 at 1 000 draws, and $\mathrm{sd}(T_3)$ at 60 kW averages 0.1946 K against 0.1945.

36.5 The report (step 6)

The contractor's slope was right and the conclusion drawn from it was not. The hall delivers 90.7 % of its design conductance, at $\hat{\mathbf{z}}$, and the clause is settled against it at 1.5×10^{-4}. The bias he subtracted is undecided, and subtracting it does not reach the rating: with one unit out, the hall holds 25°C at 60 kW with probability 0.27, three to one against. And the largest number in this report has nothing to do with the test.

36.5.1 The answers and their bands

Table 36.14 is the report. Each row is a question in the engineer's words; the columns give the answer on the state route, the exact value, the gap, what the band and qualification are and the chapters behind it.

The report line. Hall 3 is not accepted as designed: its conductance is about 4.53 kW/K, 91 % of the 5.0 in the clause, and within this state no further record has more than 1.6×10^{-4} of changing that; outside it, the meters' scale and a load-dependent conductance could. At the design load with one unit out it probably exceeds 25°C, with a probability of 0.73 on the state route (0.72 exact), and the test cannot say more until the new sensor is calibrated, which may move the admissible load either way, about 58 kW expected, between 55 and 61 kW. Until then the load to accept is 56.45 kW at the customer's 95 % or 57 kW at the illustrative 90 %, for the signatory to choose, on a generated record whose load runs from 54.8 to 60.4 kW over other draws from the same plant. Its largest risk, on illustrative prices, is the cooling, R$ 636 000 a

Table 36.14: The commissioning report. The record generated; the prices illustrative; q a prior; β_{cc} pinned.

Question	Answer, state route	Exact	Gap	Band, qualification	From
Meets its design? N+1?	Not accepted; $\Pr(G_3 \geq 5) = 1.5{\times}10^{-4}$; $\hat{G}_3 = 4.533$ at $\hat{\mathbf{z}}$; capacity 58.98 kW at $\hat{\mathbf{z}}$	2.4×10^{-4}; median 4.537; 59.03 kW	-9×10^{-5}	G_3 in [4.30, 4.78]; within this state any record reopens it with at most 1.6×10^{-4}; outside it, the meters' scale (13 % low) and a load-dependent conductance; the test's truth set below design	27
Under 25°C at 60 kW, one unit out?	0.266; 25.22°C at $\hat{\mathbf{z}}$, spread 0.36 K	0.276; median 25.21°C	−0.010	0.003 to 0.98 over other records; the even chance is at 59 kW	28, 20
Sensor lying?	$b_3 = 0.15 \pm 0.26$ K; $\ln B_{10} = -0.47$, undecided	0.156; −0.47	−0.006; −0.002	−0.03 to −1.66 with the prior (Table 36.13); −0.67 to 5.19 over records; matters to the margin and not the clause	22, 24
Availability?	$\mathbb{E}[U_{\mathrm{pair}}] = 4.84 \times 10^{-3}$, 42 h a year; target tail 0.021	3.96×10^{-3}, 35 h; 0.065	8.8×10^{-4}	band $[1.05,\ 15.3] \times 10^{-3}$; shared part 0.72 (0.76) at the pinned 0.1, not the larger part below 0.043 (0.037); no record	15, 13
Annual risk?	cooling R$ 635 961; thermal, one unit out, R$ 36 679	R$ 520 656; R$ 36 201	R$ 115 305; R$ 478	band of the first [R$ 137 591, R$ 2 009 353]; first larger unless the loss rate is under R$ 865 per hour; no capacity state in the first	25
Next, and what load?	calibrate; $Q_{0.90} = 56.99$ kW, after calibration 57.91 kW expected, band [55.35, 60.55]; $Q_{0.95} = 56.45$ kW	57.02; 57.95; 56.47	−0.03; −0.04; −0.02	reaches 60 kW with probability 0.059 (0.062), falls under the present load with 0.247 (0.240); log the CRAC units for the cooling	31

year (R$ 521 000) against R$ 37 000 for the load, and this rests on a prior on the CRAC units, a pinned shared-cause share and a 1-of-2 form with no capacity state: no test of this hall touches it.

36.5.2 The recommendation and its conditions

Three things stand on the numbers, and two depend on something still to be supplied.

Do not sign the clause: that stands on the tail, with the meters' scale worth checking first. Derate, at the level the signatory chooses: $Q_{0.90}$ pays for any value of admitted load under R$ 10 537 per kW-year and $Q_{0.95}$ under R$ 9 638. Calibrate before fixing that load for good, and read the load again from the new posterior, not from the 58 kW at the center: it is the only candidate that can decide the margin without running the hall over its rating, and it can return 60 kW or take away one more. It is not a precondition of the clause's verdict, which the bias does not affect.

The cooling is not yet a program. Halving the share is worth R$ 209 442 a year (R$ 182 659) at these prices, and this is the break-even price of an assumed halving and not a recommendation. The condition for acting on the shared cause is a record: a CRAC log with its simultaneous outages counted, without which the three quarters is the pinned value speaking.

> **Principle 36.1.** *One record gives answers of different standing, and the report has to carry the standing with the number: decided, open, undecided, or resting on no row at all. A probability near zero and a probability of a quarter from one fit are not the same kind of result, and the second is also the one that moves on another record. What a measurement can do is read from the posterior beforehand: an average of posteriors returns to the present one, which bounds what any record can reopen, and it also says that a calibration moves a load both ways.*

Common mistake

Four errors recur in a commissioning. Reading the sensor's band as the hall's: 0.17 K against 0.36 K at 60 kW, with the bias canceling in the reading's row. Fixing the bias at zero because the evidence leans there: on the state route the probability at 60 kW goes from 0.266 to 0.0037 on a sensor nobody calibrated. Quoting the load after a calibration at the center: 57.9 kW is not what the calibration will give, and a quarter of its outcomes fall below the present load. And reading a small $2\mathcal{L}$ as a quality: with three loose prior rows, 1.34 on six is not rare.

36.6 The next measurement (step 7)

The candidates and their numbers are those of Section 36.3.6. What remains is which to take for which question, and what the chapter leaves behind.

36.6.1 The next record for each question

For the margin and the load, the calibration: it is the only record among the three that both narrows the temperature and can move it across the limit without running the hall over its rating, and when its row arrives the fit returns to step 3 with thirteen test rows. For the clause, nothing within this state; outside it, check the load meters' scale, and a load-dependent conductance if that is in doubt. For the cooling, the CRAC log with simultaneous outages counted, which the calibration cannot replace, since pattern (36.3) has no row between them. These answer different questions and do not compete, so the chapter does not rank them by a common value: a value of R$ 365 a year for the calibration is on a load price, and the log's would be on the share of a program no one has yet proposed.

36.6.2 What the analysis showed

The contractor's two-point line gave nearly the right conductance and the wrong reading of what it meant. Subtracting the sensor's apparent offset seemed to bring the hall near its rating, and the analysis finds the hall below the clause by more than three spreads, probably above the rating, and with a bias the record cannot decide. For the plant it means a hall accepted at a lower load, a sensor to calibrate, and a cooling pair whose unavailability is still only a design file's opinion, which is where most of the money lies. On this generated record the hall's true capacity is just short of the design load, which is why the margin is so hard to decide; an engineer should expect another afternoon's record to read the rating quite differently, and the clause no differently.

36.6.3 Limits of the model

The derivations are exact under their assumptions, and six of them need checking before a result is used.

The generated record. The plant was set below its design, the sensor biased, and the supply on the prior's lower side. The clause's verdict comes from that setting, and Section 36.4.2 gives what moves.

The steady readings. Each hold is taken as settled. Chapter 35's hall 1 has a time constant of 1.69 h; if hall 3's were similar, a short hold would read a lower temperature than the steady one, the fit would see a higher conductance, and the true conductance would be lower still, so the clause's verdict would stand and the margin would be worse. A hold of several time constants should be on the test sheet.

The supply sensor and its one entry. A persistent error of the supply sensor moves T_s and enters b_3 opposite it, and the means of ten readings do not reduce it, nor do they reduce any error shared in time. The hall's heat also reaches the loop, which the fixed T_s leaves out. If the prior from Chapter 18 used the same sensor, its row shares that error with this record.

The straight line. A constant bias and a load-independent conductance make a line. Three points give one degree of freedom against a curve, so the record cannot test this; on the sensor's calibration, a bias that varied with temperature would be seen only at the reference points.

The cooling's form. The pair has no repair process, no capacity and no record. The 1-of-2 form counts one unit as carrying the hall, and the test taken with one unit out says otherwise at the design load: the N+1 clause is not met and the rating is exceeded with probability 0.73. So at 60 kW the form sits on what step 4 refutes; at an accepted load it is consistent only to the assurance of that load, and only for the unit that ran. The cooling term has no capacity state and none is added here. Its ceiling, counting every one-unit hour as cooling lost, is R$ 8.5 million a year (R$ 7.4 million), against R$ 636 000 (R$ 521 000); a hall a fraction of a kelvin over its rating has not lost its cooling, so this is a bound and not an estimate. The crew is taken up in Section 36.3.4.

The prices and the levels. The loss rate, the penalty, the load value and the 0.90 are illustrative, and the penalty's following the accepted load is a clause to be agreed. The tail on the clause and the temperature band do not use them.

36.7 What to retain

Four conclusions carry beyond this hall. A requirement, a margin and a hypothesis read from one fit can have entirely different standing, and the difference is visible only when each has its own number. A record's fit can be reported with a chi-square whose reference is honest only if the prior rows are truly random, and it should be read against the generator or the process that made the record. What no record within a state can reopen follows from an average of posteriors, and what could reopen it lies outside that state and is to be named. And the largest term in a report can rest on a prior, which is a reason to say so and to start the log. Table 36.15 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 36.15: The terms of the commissioning report and where each is defined.

Term	Expression	Units	Meaning	Defined in
Hall row	$(T_s + Q_k \mathrm{e}^{-g_3} + b_3 - y_{T,k})/\sigma_{\mathrm{h}}$	–	reading with a bias	36.1.3
True temperature	$T_3 = T_s + Q^* \mathrm{e}^{-g_3}$	°C	the event's quantity	36.3.2
Reading's gradient	$\mathbf{g}_T + \mathbf{e}_{b_3}$	varies	why the band is tighter	36.3.2
Pair	$\beta_{\mathrm{cc}} q + ((1-\beta_{\mathrm{cc}})q)^2$	–	unavailability	36.3.4
Calibration row	$(b_3 - y_{\mathrm{c}})/\sigma_{\mathrm{ref}}$	–	the next record	36.3.6
Reopening bound, within the state	$\Pr(\text{reach } \kappa) \le \Pr(G_3 \ge G^* \mid \mathbf{y})/\kappa$	–	any record	36.3.6
General forms				
Threshold	$\Pr(f(\mathbf{z}) \ge f^* \mid \mathbf{y}) \ge \kappa$	–	acceptance, margin	36.3.1
Evidence ratio	$B_{10} = p(\mathbf{y} \mid M_1)/p(\mathbf{y} \mid M_0)$	–	the bias	36.3.3
Expectation	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$	varies	availability	36.3.4
Consequence	$\mathcal{R}(a) = \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]$	R$/yr	annual risk	36.3.5
Design	$U(a)$; Q_κ	varies	the next record, the load	36.3.6

36.7.1 The seven steps as one method

This book has asked one thing of every reliability question: write it as a probability or an expectation over the unknowns of a physical model, derive what that requires, compute it, check it and act on it. The plant graph sets out the state; the question stands on the left of the bar with the evidence on the right; the evidence becomes rows; the rows give one posterior; every answer is read from that posterior; the answers are tested against the full calculation; the report carries each with its band; and the next measurement returns a new row to the fit. Those steps did not change between a valve with fifty tests and this hall, with one stated exception of route: Chapter 33's hierarchical drawing reports from a chain, because the normal at the mode misplaces its funnel.

The five categories, estimation, prediction under a threshold, hypothesis comparison, diagnosis and design, are not five subjects. Here estimation read the conductance and the bias from entries of the state, and the year's risk as an expected consequence; prediction under a threshold read the clause and the margin from tails of an entry and of a gradient; hypothesis comparison read the bias from a ratio of evidences, which took a second drawing's fit, and such drawings are what a diagnosis ranks; and design read the preposterior and the assurance load. Each category uses the posterior through the computation it needs.

What the capstone adds to this is that a capacity, a temperature, a sensor, a pair of machines and a year's cost can be asked of one state, as the fleet of Chapter 33, the cooling plant of Chapter 34 and the coupled plant of Chapter 35 were, and that the method then says, with the same care, where its answers are firm and where they rest on a choice. The deliverable was never a number alone. It is the report: the hall is not accepted as designed; until its sensor is

calibrated the load to accept is 56.45 kW at 95 % assurance or 57 kW at 90 %; and its largest risk is a cooling pair on which no record has yet been kept.

36.8 Exercises

Exercise 36.1. *Formulation.* (a) Compute the capacity at 25°C for $T_s = 12$°C and $G_3 = 4.5$ kW/K. (b) Compute the unavailability of the 1-of-2 pair of equation (36.11) at $q = 0.02$ and $\beta_{\mathrm{cc}} = 0.1$, and the shared part of it. ((a) $13 \times 4.5 = 58.5$ kW. (b) $0.002 + 0.018^2 = 2.32 \times 10^{-3}$, of which 0.86 is shared.)

Exercise 36.2. *Derivation.* Derive why the reading's band at 60 kW is tighter than the truth's, from the gradients and the entries of $\boldsymbol{\Sigma}$. ($\mathbf{g}_{\mathrm{read}} = \mathbf{g}_T + \mathbf{e}_{b_3}$, so $\mathrm{Var}_{\mathrm{read}} = \mathbf{g}_T^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_T + 2\mathbf{g}_T^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{e}_{b_3} + \Sigma_{b_3 b_3}$. With $\mathbf{g}_T^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{e}_{b_3} = -13.2349 \times 0.006133 - 0.002131 = -0.08330$: $0.12961 - 0.16660 + 0.06563 = 0.02864$, and $\sqrt{\cdot} = 0.169$. The middle term is negative because g_3 and b_3 are positively correlated and $\mathbf{g}_T$ has a negative entry on g_3.)

Exercise 36.3. *Derivation.* Derive the load meeting a limit with probability κ from the delta band, and check 56.99 kW. (Set $(a - x)/\sqrt{x^2\Sigma_{g_3 g_3} - 2x\Sigma_{g_3 T_s} + \Sigma_{T_s T_s}} = z_\kappa$ with $a = T_{\max} - \hat{T}_s$, $x = Q\mathrm{e}^{-\hat{g}_3}$; square to $Ax^2 - 2Bx + C = 0$, equation (28.13). With $z_{0.90}^2 = 1.6424$: $A = 0.99879$, $B = 13.01030$, $C = 169.2661$, $x = 12.5718$, and $12.5718 \times 4.53345 = 56.99$ kW.)

Exercise 36.4. *Derivation.* Derive the evidence of a bias entry from the two fits' objectives and information matrices, and evaluate it. (Equation (24.12) on each drawing; the row normalizers cancel but for $\sqrt{2\pi}\sigma_b$, one entry adds $\frac{1}{2}\ln 2\pi$, so $\ln B_{10} = (\mathcal{L}_0 - \mathcal{L}_1) + \frac{1}{2}\ln 2\pi - \frac{1}{2}\ln(\det \boldsymbol{\Lambda}_1 / \det \boldsymbol{\Lambda}_0) - \ln(\sqrt{2\pi}\sigma_b) = 0.1709 + 0.9189 - 1.3390 - 0.2258 = -0.475$.)

Exercise 36.5. *Application.* A one-load test takes three holds all at 60 kW. Show that g_3 and $T_s + b_3$ are no longer separated by the reading rows, and say what the priors then do. (The hall rows' g_3 column is proportional to the b_3 column, up to the small change in Q_k, so the reading rows' smallest singular value drops from 1.87 to 0: at the truth of this chapter the singular values with the priors end 2.43, 2.43, 2.01, 2.00, and the last direction is carried by the priors alone. The spread of b_3 stays near its prior, 0.496 K, the correlation is 0.971, and $\sqrt{\Sigma_{g_3 g_3}} = 0.039$, the prior's 0.5 K over the rise, the floor (27.15) with the bias in place of the supply.)

Exercise 36.6. *Application.* The sensor is calibrated before the test, with a certificate of $b_3 = 0.15 \pm 0.1$ K. Refit the record of this chapter: what changes, and what does not. (The prior row on b_3 is replaced by the certificate: $\sqrt{\Sigma_{g_3 g_3}} \to 0.0155$, $\mathrm{sd}(T_3)$ at 60 kW $\to 0.196$ K, $\Pr(T_3 < 25) \to 0.133$ and $Q_{0.90} \to 57.91$ kW; the bias question and its evidence fall away. The clause's tail falls further, to 10^{-10}. The cooling's answers and the risk's first term do not change, since no row of the test touches ξ_q.)

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?	$p(R(t;\lambda) \mid \mathbf{y})$	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
What can these readings say about the plant's state?	$r(\mathbf{z})$, $\mathbf{J}$: the rows and their design columns	the formulation: physical, reading and prior rows over one state vector	17
What are the plant's conductances, loads and supply temperature?	$\hat{\mathbf{z}} = \arg\max p(\mathbf{z} \mid \mathbf{y})$	the objective $\frac{1}{2}\sum r_k^2$ minimized by Gauss–Newton	18
Do the readings agree with the physics and with each other?	$2\mathcal{L}(\hat{\mathbf{z}})$ against χ^2 with rows minus entries degrees of freedom	the normalized residuals and the redundancy check	18
Is the estimate physically admissible?	$\Pr(G_i > 0,\ Q_i > 0,\ T_i > T_s \mid \mathbf{y})$, a joint event	the joint tail under the Laplace normal; distances from the boundary in standard deviations	18
How well is each entry known, and which are correlated?	$\boldsymbol{\Sigma} = \boldsymbol{\Lambda}^{-1}$	the Laplace approximation for n entries; marginals and correlations	19
Can my sensors see this parameter at all?	rank of the design columns; prior spread against posterior spread	observability analysis: the singular values of $\mathbf{J}$ and the weak directions	19
With instruments out, how likely is a load above a limit?	$\Pr(Q_2 > 45\,\text{kW} \mid \mathbf{y}')$ with the meter rows removed	the exact tail as a one-dimensional integral over the conductance; the prior's spread swept	19
What will the hall's temperature be at a future load, and how well is that known?	$p(f(\mathbf{z}) \mid \mathbf{y})$ at a scenario node: location and central interval	the delta method $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ against the full posterior by reweighted draws	20
Will the hall stay under the limit at a future load?	$\Pr(f(\mathbf{z}) < T_{\max} \mid \mathbf{y})$ at a scenario node	the delta method $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ and a normal tail; a check against reweighted draws	20

Engineering question	Mathematical form	Derives to	Chapter
Which entry's uncertainty drives an output?	the budget of $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$	one term per entry plus the cross terms, which may be negative	20
What are a ratio, a total and a conductance, with bands?	$p(f(\mathbf{z}) \mid \mathbf{y})$ for nonlinear f	the delta band against the full posterior; the delta band shifted where f is curved	20
What is the temperature where there is no sensor?	$p(f(\mathbf{z}) \mid \mathbf{y})$ for a node quantity with no reading row	a virtual sensor: the same $\boldsymbol{\Sigma}$ pushed through f; the downdate of a removed row; the predictive of an unseen reading	21
Is the return temperature below the chiller's limit?	$\Pr(T_r(\mathbf{z}) < 18°\mathrm{C} \mid \mathbf{y})$	a derived quantity's tail by the delta method and by reweighted draws	21
Do the readings disagree with the model?	$\Pr(\chi^2_\nu > 2\mathcal{L}(\hat{\mathbf{z}}_0))$, a tail of the predictive	bad-data detection: the misfit against its chi-square law	22
Is this reading lying? Which one?	the normalized residuals $r_k/\sqrt{1-K_{kk}}$; the largest against r^*	bad-data identification: the hat matrix, the redundancy number, the Bonferroni threshold	22
Which sensor is biased, by how much, and how sure is that?	$\Pr(M_k \mid \mathbf{y}) \propto p(\mathbf{y} \mid M_k)\,\Pr(M_k)$ over a family of bias entries; $\hat{b}_k$ and its band	diagnosis by evidence: one freed entry per sensor, the posterior over the family	22
Can a bias on this sensor be detected at all?	$\Pr(\lvert r^{\mathrm{N}}_k\rvert > r^* \mid$ bias $\Delta)$; $1-K_{kk}$	the minimum detectable bias b^*_k from the predictive of the other rows; the critical (unredundant) measurement	22
What does a biased reading do to the answers if it is kept?	$p(f(\mathbf{z}) \mid \mathbf{y})$ under the drawings with and without the bias	the scenario output refitted under each drawing	22
Which row makes a sensor checkable?	$1-K_{kk}$ and b^*_k after a candidate row, from $\boldsymbol{\Sigma}_a$	the redundancy number by the preposterior; a candidate ranked by the bias it would detect	22
Has the plant changed since the last period?	a tail of the predictive under the healthy model	fault detection: the earlier period's marginals as prior rows, the misfit and the held-out departures	23
Which component?	$\Pr(M_j \mid \mathbf{y}) \propto p(\mathbf{y} \mid M_j)\,\Pr(M_j)$ over the candidates the graph generates	fault isolation: a widened prior or an added entry per candidate; the posterior over the family, conditional on the family	23

Engineering question	Mathematical form	Derives to	Chapter
How large is the fault, and how well is it known?	$p(z_i \mid \mathbf{y}, M_j)$ within the leading drawings	fault identification: the freed entry's posterior and its band	23
Which parameter moved between two periods?	$(\hat{\mathbf{z}}_2 - \hat{\mathbf{z}}_1)_i$ against $\sqrt{(\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2)_{ii}}$; the joint statistic	drift attribution: two fits compared entry by entry	23
How long a record would have told?	the answers as functions of the record length N	the averaged rows $\bar{y}_k$ with spread $\sigma_k/\sqrt{N}$; the limits as N grows	23
Is it really parallel? Is this hall really on the loop?	$p(\mathbf{y} \mid M_\mathrm{A})/p(\mathbf{y} \mid M_\mathrm{B})$	model selection: the evidence as a normalizing constant, the Occam factor, the Laplace evidence and what cancels	24
Does the record contradict the handbook value?	$B_{01} = p(z_j^\mathrm{h} \mid \mathbf{y}, M_1)/p(z_j^\mathrm{h} \mid M_1)$	a point null by the Savage–Dickey ratio; the closed-form evidence of a count	24
What availability is reported when the drawing is uncertain?	$\sum_d \Pr(M_d \mid \mathbf{y})\, \mathbb{E}[U_\mathrm{c} \mid M_d, \mathbf{y}]$	the mixture over drawings and its between-drawing variance	24
Which reading settles the structure, before it is installed?	$\sum_d \Pr(M_d \mid \mathbf{y})\, \mathbb{E}[\ln B_{dd'}(\tilde{y}_a) \mid M_d]$	the expected evidence gain: the Kullback–Leibler divergence of the predictives	24
What is the expected loss per year, and how well is it known?	$\mathcal{R}(a) = \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]$, $\mathcal{R}_\mathrm{yr} = t_\mathrm{yr} \sum_\mathbf{s} \Pr(\mathbf{s} \mid \mathbf{q})\, C(\mathbf{s})$	probabilistic risk assessment: the enumeration weighted by consequence, and its band	25
Where does the risk come from?	$\mathbb{E}[f_k(\mathbf{z}) \mid \mathbf{y}]$ for the parts of the risk	the sum split by state class and by minimal cut set	25
Is the annual loss below a limit with a stated assurance?	$\Pr(\mathcal{R}_\mathrm{yr}(\mathbf{q}) < \mathcal{R}^* \mid \mathbf{y}) \geq \kappa$	risk-informed acceptance: a tail of the risk posterior; the plug-in risk against the expected risk	25
What is the loss in a bad year?	the 95th percentile and the tail mean of the predictive year's loss	a quantile of a compound predictive: aleatory and epistemic variance split	25
How much risk does each design change remove?	$\mathcal{R}(a)$ ranked over the actions; $\Pr(L(a, \mathbf{z}) < L(a', \mathbf{z}) \mid \mathbf{y})$	the alternative drawings' risks, draw by draw	25
Which unit's unavailability matters most to the top event?	$I_{\mathrm{B},i} = \partial U/\partial q_i = U_i^+ - U_i^-$	the Birnbaum importance by the pivotal decomposition, with its band	26

Engineering question	Mathematical form	Derives to	Chapter
What share of the top event runs through each unit?	$I_{\mathrm{FV},i} = (U - U_i^-)/U$	the Fussell–Vesely importance, equal to the criticality importance; the cut-set forms bound it	26
How much worse if the unit were down, how much better if it never failed?	U_i^+/U, U/U_i^-	the risk achievement and reduction worths, and their identities with the other two measures	26
Are the rankings firm?	$\Pr(I_{\mathrm{FV},K} > I_{\mathrm{FV},C} \mid \mathbf{y})$	a threshold on the log-ratio of two shares; one exact integral	26
Why does the band of the top event use the same numbers?	$\mathrm{Var}(\ln U \mid \mathbf{y}) = \sum_j (g_j/U)^2 \Sigma_{jj}$	the gradient of the band as the importances summed over the units sharing an entry	26
Does a capacity read through a model meet its requirement?	$\Pr(G_1 \geq G^* \mid \mathbf{y}) \geq \kappa$	the acceptance rule as a normal tail of one entry, with the full posterior beside it; the one-sided bound	27
How strongly does the record favor the requirement over its complement?	posterior odds over prior odds of $G_1 \geq G^*$	the Bayes factor of two halves of one prior, a ratio of tails	27
How well does a test fix a capacity, and what limits it?	$\mathrm{sd}(g_1 \mid \tilde{y}_a)$ as n grows	reading noise over n against the supply's error over the rise	27
How likely is a test to accept a good or a bad path?	$\mathrm{OC}(\mathbf{z}_{\mathrm{t}}; a) = \Pr(\mathrm{accept} \mid \mathbf{z}_{\mathrm{t}}, a)$	the operating characteristic, a preposterior at a stated truth; producer's and consumer's risks	27
How long a test, at what load?	the smallest test meeting both risks	planning by the operating characteristic; the load is the lever	27
What does accepting on a point estimate get wrong?	the operating characteristic at $\kappa = \frac{1}{2}$	the consumer's risk of the point rule: one half at the requirement	27
Will the hall stay under the limit at a scanned load?	$\Pr(T_1^*(Q_1^*) < T_{\max} \mid \mathbf{y})$	the tail by load, on the route the propagation rule selects	28
How much margin is there, and how well is it known?	$p(T_{\max} - T_1^* \mid \mathbf{y})$	the margin's band, its variance quadratic in the load	28
How much load can be taken with an assurance?	$Q_\kappa = \max\{Q_1^* : \Pr(T_1^* < T_{\max} \mid \mathbf{y}) \geq \kappa\}$	the tail inverted: a quantile of the path's capacity	28

Engineering question	Mathematical form	Derives to	Chapter
What is the path's capacity, and is it at least the design load?	$p(Q_{\text{cap}} \mid \mathbf{y})$, $\Pr(Q_{\text{cap}} \geq Q \mid \mathbf{y})$	load–capacity interference: the same event as the temperature tail	28
Is the unit delivering what its plate says?	$p(c_{\text{ch}} \mid \mathbf{y})$; the evidence against $c_{\text{ch}} = 1$	a rank-one update and the Savage–Dickey ratio; a meter bias as the competing drawing	28
How much load does the chiller have in hand?	$p(P_{\text{pl}} c_{\text{ch}} - Q_1 - Q_2 \mid \mathbf{y})$	a linear derived quantity in which the meter bias cancels	28
What does a derate do to the margin at full demand?	$\Pr(T_1^* + (P^* - P_{\text{pl}} c_{\text{ch}})^+ / \dot{m} c_p < T_{\max} \mid \mathbf{y})$	the shortfall's partial expectation propagated through the supply	28
What does an inspection interval cost, and how often should we test?	$\mathcal{R}(\tau) = \mathbb{E}[L(\tau, \mathbf{z}) \mid \mathbf{y}]$, $\tau^* = \arg\min_\tau \mathcal{R}(\tau)$	the delay-time loss, a partial-mean condition for the minimum	
How firm is the choice of action?	$p(\tau^*(\mathbf{z}) \mid \mathbf{y})$; $\Pr(L(a, \mathbf{z}) < L(a', \mathbf{z}) \mid \mathbf{y})$	draws of the state and the variance budget	
Would learning the delay change the interval?	$\mathcal{R}(\tau^*) - \mathbb{E}_d[\min_\tau \mathcal{R}(\tau \mid d)]$	the expected value of perfect information and the preposterior rows	
At what age should a unit be replaced?	$T^{\text{B}} = \arg\min_T \mathcal{R}(T)$	renewal-reward cost rate; a stationary condition that exists only for a rising hazard	
How different are the plug-in and the Bayes actions?	$T^*(\hat{\mathbf{z}})$ against T^{B}; the regret	the expected cost of each action over the posterior	
How much would knowing the state save?	$\min_a \mathcal{R}(a) - \mathbb{E}[\min_a L(a, \mathbf{z}) \mid \mathbf{y}]$	a Jensen gap, nonnegative on a coherent posterior	
How much does a row tell about the state?	$\text{EIG} = \frac{1}{2} \ln(1 + j_a^{\mathsf{T}} \boldsymbol{\Sigma} j_a)$	the entropy of the reading	
How much width does it take off an output?	$\mathbf{g}^{\mathsf{T}} \boldsymbol{\Sigma}_a \mathbf{g}$ against $\mathbf{g}^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g}$	the rank-one update; the squared correlation of output and reading	
What is a reading worth to a decision?	EVSI_a, the value after the reading averaged over its predictive, less the value now	Gauss–Hermite over the reading, or reweighted outcomes	
How many units?	$\arg\min_n \ (n-3)c_u + C_L t_{\text{yr}} \mathbb{E}[U_n \mid \mathbf{y}]$	the first n whose marginal saving falls below the unit's price	

Engineering question	Mathematical form	Derives to	Chapter
How many spares?	$\arg\min_s \; c_s s + C_s \mathbb{E}[(K-s)^+ \mid \mathbf{y}]$	the predictive quantile at the critical fractile $1 - c_s/C_s$	
Which plant change pays?	the loss removed on the residual over c_a; the tail above c_a	a benefit–cost ratio with its band	
Why does the plug-in action err?	$\mathbb{E}[m_a(\mathbf{z}) \mid \mathbf{y}] - m_a(\hat{\mathbf{z}})$	$\frac{1}{2}\operatorname{tr}(\mathbf{H}_m \boldsymbol{\Sigma})$, the curvature of the marginal saving	
Does a target or a cost minimum decide?	$\Pr(U_n < u^* \mid \mathbf{y}) \geq \kappa$ against $\arg\min_n \mathcal{R}(n)$	a delta tail against the marginal condition	
Do the units differ?	three evidences; the population spread τ	Savage–Dickey from draws at $\tau = 0$	
Is one unit worse?	$\Pr(\tau\delta_i > \ln c \mid \mathbf{y})$	partial pooling by the chain, conditional on the drawing	
What is a record worth when the normal fails?	the expected value of sample information	simulated outcomes, reweighted draws	
Which cut sets dominate?	$\mathbb{E}[P_c \mid \mathbf{y}]$ with causes as basic events	shares with bands	
Which order-one event matters?	Fussell–Vesely and Birnbaum importance	Fussell–Vesely ranks; Birnbaum does not separate them	
Is a further unit worth buying?	the risk removed against its price	only the non-shared term is reachable	
How often does an outage cross the limit?	$\mathbb{E}[\Pi(\mathbf{z}) \mid \mathbf{y}]$	the start, running and generator-up paths	
Is an option worth its price?	m_a against a quoted c_a	break-evens of the saving	
How is a long time-indexed fit solved?	the factor of the information matrix in time order	filter and smoother on the augmented state, linear cost	
Accept a new hall?	$\Pr(G_3 \geq G^* \mid \mathbf{y}) \geq \kappa$	the tail of an entry	
Does it hold the limit at a load?	$\Pr(T_3 < T_{\max} \mid \mathbf{y})$	a normal tail through the gradient	
Is the new sensor biased?	$\ln B_{10}$	two fits' objectives and determinants	
What can any record reopen?	$\Pr(\text{reach } \kappa) \leq \Pr(\text{event} \mid \mathbf{y})/\kappa$	the average of posteriors, within the state	
What load after a calibration?	Q_κ under $\boldsymbol{\Sigma}_a$, with its predictive	a rank-one update and its shifted center	