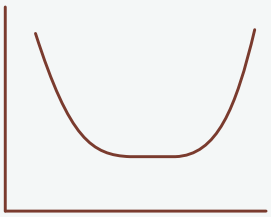
PROBABILISTIC GRAPH MODELING
FOR ENGINEERING DISCIPLINES · VOLUME III



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.

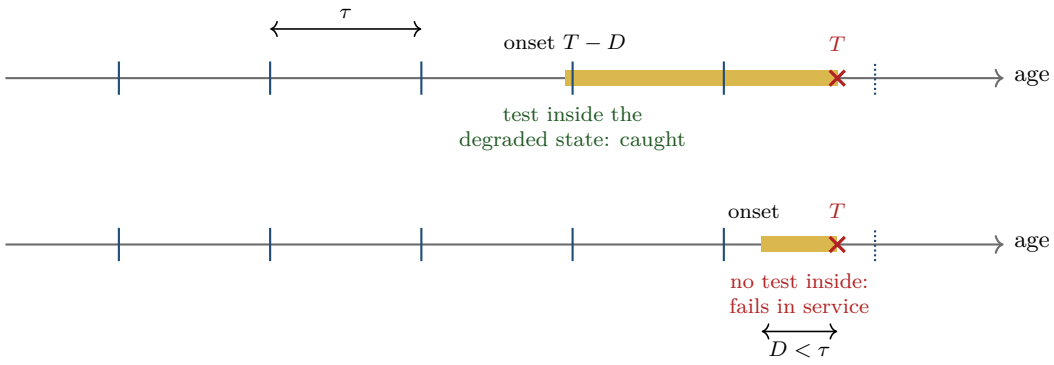


Figure 29.1: The delay-time model. Upper row: a delay longer than the distance to the next test, caught at that test (the dotted tick is not reached). Lower row: a shorter delay falling between two tests, and the failure occurs in service. Blue ticks: the tests, period τ; gold: the degraded state; red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

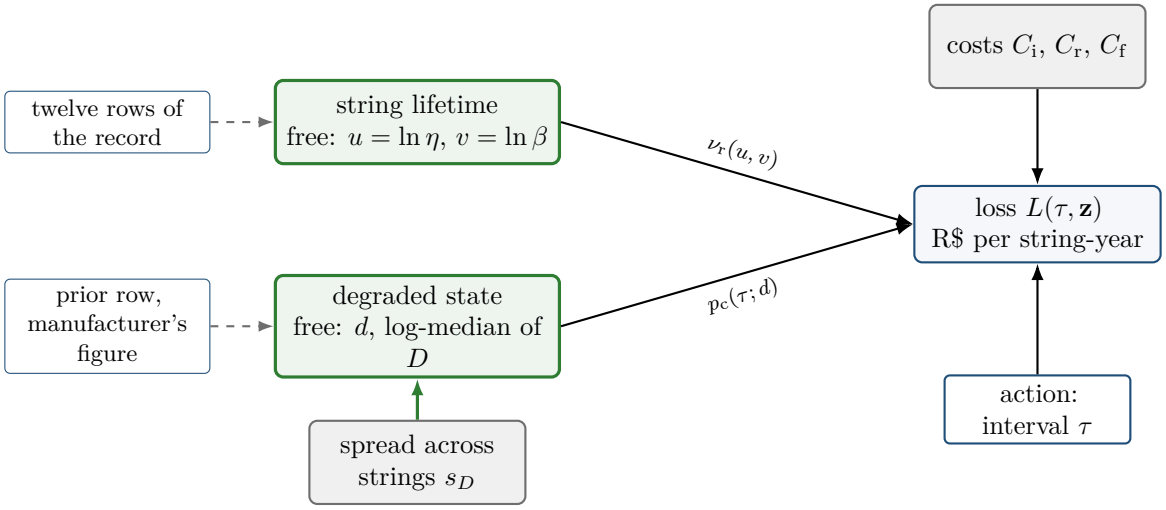


Figure 29.2: The plant graph of the inspection question. Free: the two nodes in green, the lifetime with u and v, the degraded state with d. Pinned: the spread s_D of the delay across strings and the costs (gray). Observed: the twelve rows of the record on the lifetime, and only the manufacturer's prior 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

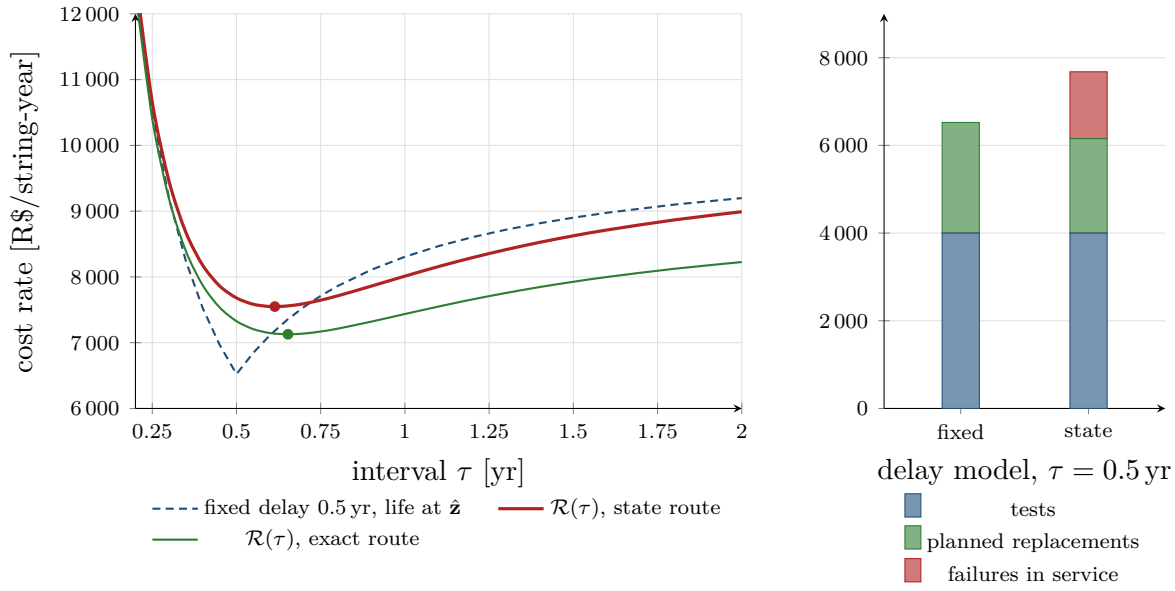


Figure 29.3: Left: the cost rate against the interval, for a fixed delay of 0.5 yr at the MAP life ($C_\mathrm{i}/\tau + \nu_\mathrm{r}(\hat{\mathbf{z}})[C_\mathrm{f} - \Delta C \min(0.5/\tau, 1)]$), and the expected cost $\mathcal{R}(\tau)$ of equation (29.15) with $s_\mathrm{tot} = 0.5$ on the two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

END OF THE PREVIEW

Reliability Engineering by Derivation

These are the opening pages. The complete book runs 222 pages.

The full PDF opens with a token. Write to **rmarquez123@gmail.com**, name the book, and I will send you one.

One token opens every book on the site.

Ricardo Marquez, Ph.D.

ricardo-marquez.com