

Asset Management by Derivation

Remaining life, maintenance decisions and the value of condition monitoring

Draft · 361 pages

Ricardo Marquez, Ph.D.

ricardo-marquez.com

Asset Management by Derivation

Volume II

From the asset manager's questions to probabilistic inference and decisions

Ricardo Marquez

November 1, 2026 (work in progress: chapters are added as they are written)

How to read this book

Volume I of this book holds Parts I to III, Chapters 1 to 13: the method itself (a plant graph, a question as a probability over its state, and the five computations beneath every answer: the posterior, the trend, the evidence, the preposterior and the expected consequence), the condition of the hall's coils, its battery string and a fleet of fifty coils today, and the rates at which the coils' condition changes, read from three years of monthly windows. This second volume holds Parts IV to VI, Chapters 14 to 24. It turns from how fast to when and what: Part IV forecasts the need for attention (the date the coils can no longer hold the hall, the probability of a failure in a mission window, and a fleet's failure rate by age); Part V makes the maintenance decisions as expected consequences over that one posterior (repair or replace, the inspection interval, the trigger for a work order, the order in which one crew visits fifty units, and a warranty claim); and Part VI values the monitoring itself before it is logged (whether the sensors can resolve a stated loss in the time available, which sensor narrows the crossing date most, and how often to log). Chapter numbers run continuously across the two volumes; a reference to a chapter numbered 13 or below points to Volume I, 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; step 3, the posterior of the state, derived and computed once; step 4, the answers, each read from that posterior down to its number; step 5, the checks against the full calculation and against the record's constructed truth; step 6, the report with its band and the entry in the maintenance log; step 7, the next measurement, ranked by the width it removes. The notation chapter below is the one of Volume I with this volume's symbols added, and the atlas of questions in the appendix collects the rows of both volumes.

Contents

How to read this book **ii**

Notation and vocabulary **vi**

IV Forecasting the Need for Attention **1**

14 Remaining useful life **2**
14.1 The dashboard's date and the coils' life 2
14.2 The posterior of the state (step 3) 8
14.3 The answers (step 4) 13
14.4 The checks (step 5) 21
14.5 The report (step 6) 26
14.6 The next measurement (step 7) 28
14.7 What to retain 31
14.8 Exercises 32

15 Failure probability over a mission window **34**
15.1 The work-order screen and the coil's next ninety days 34
15.2 The posterior of the state (step 3) 39
15.3 The answers (step 4) 42
15.4 The checks (step 5) 49
15.5 The report (step 6) 53
15.6 The next measurement (step 7) 54
15.7 What to retain 57
15.8 Exercises 58

16 Age-dependent failure rates across a fleet **59**
16.1 The failure chart and the coil's place on the curve 59
16.2 The posterior of the state (step 3) 65
16.3 The answers (step 4) 68
16.4 The checks (step 5) 79
16.5 The report (step 6) 82
16.6 The next measurement (step 7) 83
16.7 What to retain 86
16.8 Exercises 87

V Maintenance and Replacement Decisions **89**

17 Repair or replace **90**
17.1 The calendar, the dashboard and a new coil 90

17.2 The posterior of the state (step 3) . . . 97
17.3 The answers (step 4) . . . 100
17.4 The checks (step 5) . . . 107
17.5 The report (step 6) . . . 113
17.6 The next measurement (step 7) . . . 115
17.7 What to retain . . . 117
17.8 Exercises . . . 118

18 Inspection intervals **120**
18.1 The calendar's test and the hall's last year . . . 120
18.2 The posterior of the state (step 3) . . . 127
18.3 The answers (step 4) . . . 130
18.4 The checks (step 5) . . . 140
18.5 The report (step 6) . . . 143
18.6 The next measurement (step 7) . . . 145
18.7 What to retain . . . 148
18.8 Exercises . . . 149

19 Condition-based maintenance triggers **151**
19.1 The dashboard's trigger and the cost of an order . . . 151
19.2 The posterior of the state (step 3) . . . 158
19.3 The answers (step 4) . . . 162
19.4 The checks (step 5) . . . 169
19.5 The report (step 6) . . . 173
19.6 The next measurement (step 7) . . . 175
19.7 What to retain . . . 178
19.8 Exercises . . . 179

20 Prioritizing maintenance across a fleet **180**
20.1 The worst-first list and the crew's next visit . . . 180
20.2 The posterior of the state (step 3) . . . 185
20.3 The answers (step 4) . . . 187
20.4 The checks (step 5) . . . 195
20.5 The report (step 6) . . . 197
20.6 The next measurement (step 7) . . . 199
20.7 What to retain . . . 202
20.8 Exercises . . . 203

21 Warranty claims and specification compliance **204**
21.1 The dashboard's claim and the maker's specification . . . 204
21.2 The posterior of the state (step 3) . . . 210
21.3 The answers (step 4) . . . 213
21.4 The checks (step 5) . . . 223
21.5 The report (step 6) . . . 226
21.6 The next measurement (step 7) . . . 228
21.7 What to retain . . . 231
21.8 Exercises . . . 232

VI The Value of Condition Monitoring 234

22 Detectability of degradation 235
22.1 The dashboard's power analysis and the coil's . . . 235
22.2 The posterior of the state (step 3) . . . 241
22.3 The answers (step 4) . . . 244
22.4 The checks (step 5) . . . 253
22.5 The report (step 6) . . . 256
22.6 The next measurement (step 7) . . . 258
22.7 What to retain . . . 261
22.8 Exercises . . . 262

23 Sensor selection for prognostics 264
23.1 The noisiest trace and the crossing date . . . 264
23.2 The posterior of the state (step 3) . . . 271
23.3 The answers (step 4) . . . 275
23.4 The checks (step 5) . . . 283
23.5 The report (step 6) . . . 288
23.6 The next measurement (step 7) . . . 290
23.7 What to retain . . . 291
23.8 Exercises . . . 292

24 Sampling frequency and correlated readings 294
24.1 The proposal to log faster . . . 294
24.2 The posterior of the state (step 3) . . . 301
24.3 The answers (step 4) . . . 304
24.4 The checks (step 5) . . . 314
24.5 The report (step 6) . . . 319
24.6 The next measurement (step 7) . . . 321
24.7 What to retain . . . 324
24.8 Exercises . . . 325

A Atlas of asset-management questions 327

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. Coil subscripts are $u \in \{1, 2\}$, rack subscripts $j \in \{1, 2, 3\}$ and fleet subscripts $i = 0, \ldots, 49$; windows are months, indexed t.

Symbol	Unit	Meaning
The plant graph and its state		
$\mathbf{z}$, z_i	varies	State vector of the plant graph: every free quantity estimated at once; its i-th entry.
$\mathbf{z}_t$, t	varies, –	The state of window t; the window index, $t = 0$ the commissioning test and $t = 1, \ldots, 36$ the months from January of year 1 to December of year 3.
$\mathbf{y}$, y	varies	The evidence: all readings, test results, counts and prior information treated as known; a single reading, subscripted by what it reads.
$r_k(\mathbf{z})$	varies	Row k of the model: a physical relation that must hold ($r_k = 0$), a reading compared with the state, a trend row, or a prior. As a subscript k is this row (and in Chapter 24 a channel's rows); in Chapters 2, 3 and 15, whose bare k is the count of failures, the row is ℓ.
w_k, σ_k	varies	Weight (precision) of row k and the standard deviation it corresponds to, $w_k = 1/\sigma_k^2$.
σ_p	varies	Tolerance of a physics row; an exact constraint is the limit $\sigma_p \to 0$.
The equipment hall of Part I		
T, T_c	°C	Temperature of the hall air; the supply air held by the cooling unit's discharge controller, pinned.
Q, G	kW, kW/K	Heat load of the hall; effective conductance of its air path to the supply air.
g, G_{ref}	–, kW/K	Log-conductance $\ln(G/G_{\text{ref}})$, the entry of the state that carries the conductance; the design conductance on the design sheet, 6 kW/K.
h, h^*, h_F	–	Part I's health index $h = G/G_{\text{ref}} = \mathrm{e}^g$, the conductance against design; the index a requirement needs; the functional limit, at which the asset stops doing its job. From Part II the health index is against commissioning, $h_{u,t}$ below.
μ_g, σ_g	–	Mean and spread of the handbook prior on g, $\mu_g = \ln(7.5/6)$.

Symbol	Unit	Meaning
g_t, $\hat{g}_0$	–	The log-conductance in window t; the commissioned value, the fit of the commissioning test (window 0).
$\sigma_{\hat{g}_0}$, $\hat{G}_0$	–, kW/K	Spread of the commissioning row on g (the commissioning fit carried forward); the commissioned conductance $G_{\text{ref}}\text{e}^{\hat{g}_0}$.
G_{test}, h_{test}	kW/K, –	The conductance and index one condition test reports when its reading is taken as exact, $Q/(y_T - T_c)$.
g^*, g_F	–	A requirement level on g, $\ln h^*$; from Chapter 19, the level on each coil's log-derate at which two equal coils give G^*, $\ln(G^*/(\bar{G}_1 + \bar{G}_2)) = -0.26313$; the functional limit, $\ln h_F$, in Chapter 20 $\ln(5/6) = -0.18232$ on each fleet coil's log-derate against rating.
$T_{\max}$, T^*	°C	A temperature limit: 24°C now and 27.5°C at 60 kW (the owner's request, Chapter 1), 28°C (the commissioning record's acceptance limit, Chapter 4); a threshold on a test temperature.
Q_2, T_2	kW, °C	The requested load and the hall temperature it produces, a scenario node.
ρ, μ_0, σ_0	K/kW	Thermal resistance of the air path, $1/G$; the commissioning record's pre-test estimate of it and its spread (Chapter 4).
b_T	K	Calibration offset of the hall sensor, a nuisance quantity or a bias entry.
Δ_{m}, σ_{m}	kW	A model discrepancy of the physics row and its spread (Chapter 1).
C_{th}, ϑ, ϑ_{th}, ΔT	kJ/K, s or h, s, K	Thermal capacitance of the hall; elapsed time within a record (seconds in Chapter 1's transient, hours into a battery test in Chapter 8, hours within a window in Chapter 24); the hall's time constant C_{th}/G; the departure from steady state.
y_T, σ_T	°C, K	The hall sensor's reading and the standard deviation of its error.
The two-unit hall of Parts II–VI		
$T_s(t)$, $Q_o(t)$	°C, kW	Chilled-water supply temperature and the hall's other heat in window t: known inputs that follow the season.
Q	kW	IT heat load of the hall, an entry of the state read by the hall load meter.
G_u, $\bar{G}_u$	kW/K	Effective conductance of coil u; its rated value, 7 kW/K.
g_u, d_u	–	Log-derate $\ln(G_u/\bar{G}_u)$, the entry that carries coil u's conductance against its rating; the derate factor $d_u = \text{e}^{g_u}$, one at the rating.
$g_{u,t}$, $d_{u,t}$, Q_t	–, –, kW	The log-derate, derate factor and IT load in window t.
$h_{u,t}$	–	Health index of coil u in window t against its commissioned value, $G_{u,t}/G_{u,0} = \text{e}^{g_{u,t}-g_{u,0}}$ (Part II on).
$A_{u,t}$, h_u^{raw}	K, –	Approach of unit u in window t, $T_{d,u} - T_s(t)$; the dashboard's raw-performance index $A_{u,0}/A_{u,t}$, the naive baseline.
$G_{u,\text{test}}$	kW/K	Conductance of coil u with the window's readings taken as exact.

Symbol	Unit	Meaning
σ_W	kW	Standard deviation of a water-side heat meter's reading on a coil's duty.
$g_{u,0}$	–	The commissioned value of g_u, the fit of window 0.
C_u	kW/K	Air capacity rate of unit u, pinned at its fixed fan speed; read by an airflow station only in Chapter 23's candidates.
T_r, T_c, $T_{d,u}$	°C	Return air, cold aisle (the supply air of Part I's hall, now made by two coils), and discharge air of unit u.
T_j, β_j	°C, –	Inlet temperature of rack j; its recirculation fraction.
G^*, f^*	varies	A requirement level: G^* the conductance a duty needs (in Part I, 5.0 kW/K for 24°C at 30 kW and 6.316 kW/K for 27.5°C at 60 kW; from Part II, the total $G_1 + G_2$ of the capacity scenario); f^* a level on any derived quantity. The requirement holds on the safe side of f^*: $f \geq f^*$ for a conductance, $f \leq f^*$ for a temperature.
b_j	varies	Bias of reading j, an entry freed when a sensor is suspected of drifting.
$\check{X}$	as X	The value of X at the scenario node of the commissioned model: today's load and inputs with the commissioned coils ($\check{T}_r$, $\check{D}_u$, $\check{A}_u$; Chapter 6).
$D_{u,t}$	kW	Duty of unit u in window t, the heat its coil passes to the chilled water, $G_{u,t}(T_{r,t} - T_s(t))$.
$\eta_{A,u}$, $\eta_{D,u}$, η_s	K, kW, K	Excess approach, duty shift and supply penalty against the commissioned model at today's conditions (Chapter 6).
$A_{r,t}$, T_c^{sp}	K, °C	Return approach $T_{r,t} - T_s(t)$; a cold-aisle set point.
$G_{\mathrm{tot},t}$, G_t^{app}, $\bar{Q}_o$	kW/K, kW/K, kW	The hall's total conductance $G_{1,t} + G_{2,t}$; its apparent conductance, the metered IT load over the return approach, $y_{Q,t}/(y_{T_r,t} - T_s(t))$, a raw figure; the mean other heat, 8.00 kW (Chapter 13).
ζ_0, s_{season}, s_{load}	K/kW, K, K	The cold aisle's rise per kW at the commissioned coils; the season's and the load's terms of the cold aisle's change since commissioning; from Chapter 13, the season's and the load's terms of a raw figure's change between any two windows.
$b_{d,1}$, σ_b, $\sigma_{\ln h}$	K, K, –	Drift of unit 1's discharge sensor since commissioning (a bias entry); the prior spreads of a freed bias and of a freed change $\ln h_{u,t}$ (Chapter 7). From Chapter 10 the drift is a function of the window, $b_{d,1}(t)$.
$\dot{b}_{d,1}$, t_b, $\sigma_{\dot{b}}$	K/yr, window, K/yr	Drift rate of unit 1's discharge sensor, one entry for the record, with $b_{d,1}(t) = \dot{b}_{d,1}\,(t - t_b + 1)\,\Delta t$ for $t \geq t_b$ and zero before; the drift's first window, $t_b = 20$; the drift rate's prior spread, 2.0 (Chapter 10).
M_{neither}, M_{foul}, M_{drift}, M_{both}	–	The four drawings of Chapter 7: neither term, the derate term, the bias term, both.
K_{kk}, r_k^{N}, n_r	–	Leverage of row k, the diagonal of the hat matrix $\mathbf{K} = \mathbf{J}\boldsymbol{\Sigma}\mathbf{J}^{\mathsf{T}}$, with redundancy number $1 - K_{kk}$; the normalized residual $r_k/\sqrt{1 - K_{kk}}$; the number of rows.

Symbol	Unit	Meaning
ψ_k, $s_{d,1}$, y_b	–, K, K	The ratio in the closed-form Bayes factor of a bias on reading k (Chapter 12's ψ extends it to a step's column); the shift of unit 1's discharge reading the two suspects make; a side-by-side check's reading of a bias.
Time, trend and dates		
Δt	yr	The spacing of two windows, one month.
γ, γ_u, $\dot{\gamma}_u$	yr^{-1}, yr^{-1}, yr^{-2}	Trend (degradation rate) of a parameter; coil u's rate, the trend entry of its log-derate path, read as $100(\mathrm{e}^{\gamma_u}-1)\,\%/\mathrm{yr}$ (Chapter 10), and in a drawing that frees the curvature the rate at the record's center $\bar{t}_\mathrm{y}$; its curvature (acceleration), the change of the rate per year (Chapter 11). Chapter 9's γ_i indexes the fleet's coils instead: its γ_1 is fleet coil 1, not the hall's coil 1.
t_y, $\bar{t}_\mathrm{y}$, $\gamma_u(t_\mathrm{y})$	yr, yr, yr^{-1}	Time since commissioning, $t_\mathrm{y} = t\,\Delta t$ (Chapter 11); the record's center, 1.5 yr (window 18); coil u's rate at time t_y, $\gamma_u + \dot{\gamma}_u(t_\mathrm{y} - \bar{t}_\mathrm{y})$, so that $\gamma_u(3)$ is the rate today.
x_t, $\dot{g}_{u,t}$, $\sigma_{\dot{\gamma}}$	yr, yr^{-1}, yr^{-2}	The curvature's lever on the increment from window t to $t+1$, $(t+\frac{1}{2})\Delta t - \bar{t}_\mathrm{y}$; the path's monthly rate $(g_{u,t+1} - g_{u,t})/\Delta t$; the curvature's prior spread, 0.05 (Chapter 11).
t_knee	window	The knee of a rate, after which it changes by $\dot{\gamma}_u$ per year: 12 or 24 in Chapter 11's sensitivity forms.
Γ	yr^{-1}	The hall's conductance rate, $\ln(G_{\mathrm{tot},36}/G_{\mathrm{tot},0})/(36\,\Delta t)$ (Chapter 13). With an argument, $\Gamma(\cdot)$ is the gamma function (Chapter 2).
α	–	Level entry of a straight-line trend at the record's center $\bar{t}$.
σ_q	–	Spread of a trend row: one monthly window's departure of a parameter from its straight line, $\sigma_q\sqrt{n_\mathrm{w}}$ across n_w windows without a reading; the straight line is the limit $\sigma_q \to 0$.
$r_{q,u,t}$, $\mathbf{D}_t$, $\mathbf{E}$, $\mathbf{F}$	–, varies	Coil u's trend row from window t to $t+1$; the trend rows' blocks of the stacked Jacobian in window t's columns and in the trend entries (γ_1, γ_2) (Chapter 10), and in the curvatures (Chapter 11).
n_w, n_z, t_as	–, –, window	The number of monthly increments a rate is read over, 36 for the whole record, or of unread windows a trend row spans (Chapters 4 and 10); the number of entries of the state (Chapters 3, 4, 7 and 12); the last window of an as-of fit, the record cut at window t_as.
s_u^raw, $\dot{s}_u^\mathrm{raw}$	yr^{-1}, yr^{-2}	Slope of the dashboard's least-squares line through the 37 tiles $h_{u,t}^\mathrm{raw}$ (Chapter 10); the dashboard's bend, twice the quadratic coefficient of the least-squares quadratic through them (Chapter 11).
$\tilde{\mathbf{z}}^0$, $\tilde{\mathbf{z}}^g$	varies	The stacked state with every window's coils at the commissioned values and no drift; with the fitted coils and no drift: the steps of Chapter 10's sequential split of the dashboard's slope.

Symbol	Unit	Meaning
γ_u°	yr^{-1}	Pseudo-true value of γ_u: the drawing's fit to the truth's noise-free readings (step 5, Chapter 10).
M_0, M_1, M_2, M_{12}	–	The constant-rate drawing (Chapter 10's base fit) and the drawings that free coil 1's, coil 2's and both curvatures (Chapter 11). Chapter 12 writes the first M_{none}.
δ, δ_j, t_0	–; –, K or kW; window	Size of a step change in a parameter (with an argument, $\delta(\cdot)$ is the Dirac delta, Chapter 3); a step of kind j (Chapter 12): δ_1, δ_2 a coil's, δ_{12} both coils', δ_{T_r}, $\delta_{T_{d,1}}$, $\delta_{T_{d,2}}$, δ_{T_1} a temperature sensor's, δ_Q the load meter's; the window at which it occurs.
s, σ_δ, Π_s	window, as δ_j, –	A candidate window of a step, 1 to 36; the step's prior spread, 0.15 on a coil, 1.0 K on a temperature sensor, 10 kW on the meter; the location prior, the prior probability of a step's window s (Chapter 12).
M_{none}, $M_{j,s}$, $\mathbf{a}_{j,s}$	–, –, varies	The drawing with no step (Chapter 10's base fit); the drawing with one step of kind j into window s; the step's column in the normalized rows, $-1/\sigma_q$ on a coil's trend row into s and $1/\sigma_k$ on a sensor's readings from s on (Chapter 12).
$B_{j,s}$, $\bar{B}_s$, $\bar{B}$, $B_{\max}$	–	Bayes factor of $M_{j,s}$ against M_{none}; its mean over the eight kinds at window s; its mean over all 288 candidates; the best candidate's (Chapter 12).
ψ, ψ_a	–	A step's prior variance over its variance from the other rows, $\sigma_\delta^2\,\mathbf{a}^\mathsf{T}(\mathbf{I}-\mathbf{K})\mathbf{a}$; the same after a candidate measurement a, $\sigma_\delta^2/\sigma_{\text{dat},a}^2$ (Chapter 12).
D_h, D_γ, $\mathbf{g}_{D_\gamma}$, $\mathbf{g}_{\text{split}}$	–, yr^{-1}, –, mixed	The difference $\ln h_{2,18} - \ln h_{1,18}$ of the two coils' indices (Chapter 5); the difference $\gamma_1 - \gamma_2$ of the two rates; its gradient; the gradient of coil 2's two parts of loss, $36\,\Delta t\,\gamma_2 + \delta_2$ (Chapter 12).
σ_{dat}^2, $\sigma_{\text{dat},a}^2$, σ_{new}^2	–	A step's data-only variance from the present record, with the prior taken out; the same after a candidate measurement a; that of the candidate's rows alone (Chapter 12, step 7).
$\tilde{\delta}$, $\tilde{\delta}_0$, $\tilde{\delta}_{\text{new}}$, $\tilde{\mu}$, $\tilde{v}$	–	The mean of a step's likelihood, the estimate the rows alone give; the present record's; the candidate rows'; the combined estimate's mean and variance under the posterior predictive (Chapter 12, step 7).
t_1, t_2	window	The two windows of a raw figure's change; for the summer drop, July of year 3 (window 31) and today (window 36) (Chapter 13).
ϕ_t, mo_t	rad, –	The season's phase of window t, $2\pi(\mathrm{mo}_t - 0.5)/12$, with $\cos\phi_t = 1$ between January and February; the month index of window t, 0 for January. Window 0 takes zero seasonal regressors (Chapter 13).

Symbol	Unit	Meaning
Y_t, y_t, $\check{Y}_t$, $\check{Y}_t^{\mathrm{m}}$	as the figure	The model value of a raw figure ($T_{r,t}$ or $\ln G_t^{\mathrm{app}}$) at window t; its reading; its value with the coils at the commissioned ones, at the window's own load and inputs; that value at the season's mean inputs, window 0's (Chapter 13).
s_{coils}, s_{err}	K or –	The coils' term and the reading errors' term of a raw figure's change from window t_1 to t_2, beside s_{season} and s_{load}; in bold, each term's column over the windows (Chapter 13).
$\mathbf{X}$, $\mathbf{X}^+$, $\hat{\mathbf{k}}$, k_c, k_s	varies	The design of a least-squares regression on the windows' times (Chapters 2, 11 and 13); its pseudo-inverse $(\mathbf{X}^{\mathsf{T}}\mathbf{X})^{-1}\mathbf{X}^{\mathsf{T}}$; the fitted coefficients; a seasonal regression's cosine and sine coefficients (Chapter 13).
$\mathbf{y}^{\mathrm{raw}}$, $\mathbf{w}$, $\mathbf{w}^{\mathrm{line}}$, $\mathbf{w}^{\mathrm{seas}}$	as the figure, yr^{-1}	The readings of a raw figure in windows 0 to t_{as}; the slope row of a regression, the second row of $\mathbf{X}^+$, without and with an annual harmonic (Chapter 13).
s^{reg}, s^{line}, s^{seas}, $s^{\mathrm{seas}}_{\mathrm{drop}}$	varies	The slope of a regression; of the line and of the seasonal regression through a raw figure, per year; the seasonal regression's reading of the summer drop, its fitted trend over the five months plus the residuals' change (Chapter 13).
$\hat{\nu}_t$, $\hat{\sigma}$, $\mathrm{se}(\cdot)$	as the figure	A regression's residual at window t; the residuals' standard deviation; the practitioner's standard error from the residual scatter (Chapter 13).
$\mathbf{\Sigma}_{31,36}$, $\mathbf{g}_{31,36}$	–	The covariance block of $(g_{1,31}, g_{2,31}, g_{1,36}, g_{2,36})$; the gradient of the coils' term of the summer drop on those entries (Chapter 13).
g_u^c, g_u^s, $\mathbf{g}^{\mathrm{sea}}$, σ_{sea}	–	Cosine and sine amplitudes of coil u's seasonal conductance, on top of its trend; the four amplitudes as one vector; their prior spread, 0.05 (Chapter 13).
M_{base}, M_{sea}	–	The base drawing (Chapter 10's base fit) and the seasonal-conductance drawing (Chapter 13).
t^*, t_c	window	The window at which a parameter's path crosses its requirement level: the remaining useful life is t^* less the present window. From Chapter 14 it is the first passage of $G_{\mathrm{tot},t}$ to G^* on the extended state, read in continuous windows by interpolation between $t_c - 1$ and t_c, with $t_c = \lceil t^* \rceil$ the first window at or below the limit, the date; from Chapter 17, $t_c(\mathbf{z}; \mathcal{P}_{\mathrm{v}})$ is counted from a visit.
t_{sd}	window	A planning date: the October shutdown of year 4 (window 46).
$\mathcal{C}_0$, $\mathcal{C}_1$, j, $M = (\mathcal{C}_l, j)$	–	The two factors of the forecast family (Chapter 14 on, through Chapter 23): the curvature state, $\mathcal{C}_0$ constant rates (Chapter 10's base fit) or $\mathcal{C}_1$ coil 1's curvature $\dot{\gamma}_1$ free (Chapter 11's drawing); the March step's kind at window 15, $j \in \{\mathrm{none}, 1, 2, 12, T_r, T_{d,1}, T_{d,2}, Q, T_1\}$ (Chapter 12's); a drawing of the family, one of 18.

Symbol	Unit	Meaning
$\mathbf{z}_{\text{ext}}$, $\mathbf{g}_t$, t_{hor}	varies, –, window	A drawing's state extended by the forecast windows $t = 37, \ldots, t_{\text{hor}}$, which carry trend rows and no reading; the two log-derates $(g_{1,t}, g_{2,t})$ of window t; the forecast horizon, window 276 (Chapter 14).
$e_{u,t}$, $\mathbf{J}_{\text{f}z}$, $\mathbf{J}_{\text{ff}}$	–, varies	A forecast window's departure from the trend, $\sigma_q\, r_{q,u,t}$; the forecast rows' Jacobian blocks on the record's entries and on the forecast entries, the second square and lower bidiagonal (Chapter 14).
Failures and the fleet		
λ, τ	yr^{-1}, yr	A constant failure rate; the length of a mission window (Chapters 2, 3, 15 and 16).
k, T_{tot}	–, unit-yr	A count of failures or corrective work orders; the exposure behind it.
ω, λ_{ref}	–, yr^{-1}	Log-rate $\ln(\lambda/\lambda_{\text{ref}})$, the entry that carries a rate; its reference, 0.1 per year.
σ_ω, T_f, K_τ	–, yr, –	Posterior spread of the log-rate; time to a unit's next corrective failure; count of corrective work orders in a period τ.
F, H, S	–	Events of a diagnosis: the coil fouled; the condition flag fires; an inspection grade, of value s in $\Pr(S = s)$ (Chapter 3).
F_A, F_B, q_S, q_A, q_B	–	Two findings or replacements; the probabilities of a shared cause and of each one's own wear.
a, b, a_λ, b_λ	–, yr	Shape and rate of a gamma prior on a rate (primed after the record); the context tells them from an action a. From Chapter 15, where a is the age, they are written a_λ, b_λ: the prior is a count row of a_λ failures in b_λ unit-years.
$\lambda(a)$, λ_0, θ_a	yr^{-1}, yr^{-1}, yr^{-1}	Failure rate at age a, $\lambda_0 \text{e}^{\theta_a a}$, the rate of occurrence of failures under minimal repair (Chapter 16); its value at age zero; its log-slope in age, a scalar.
n_a, E_a	–, coil-yr	Failures and exposure of the fleet in age bin a, the ages a to $a + 1$ (Chapter 16).
μ_γ, τ_γ	yr^{-1}	Mean and spread of the fleet's degradation rates, the shared prior of partial pooling. In Chapter 9 $\tau_\bullet$ stands for either of the fleet's spreads, τ_γ or τ_{c}, and $n_{\text{coil}} = 50$ counts the coils.
μ_{c}, τ_{c}	–	Mean and spread of the fleet's commissioning values.
$g_{i,0}$, γ_i, y_{ij}, a_{ij}, σ_{test}, k_i	–, yr^{-1}, –, yr, –, –	Commissioning value and fouling rate of fleet coil i; its j-th annual condition test and the age at it; the test's spread (0.04); its number of tests.
ℓ_{c}, ℓ_γ, $\varepsilon_{\text{c},i}$, $\varepsilon_{\gamma,i}$	–	$\ln \tau_{\text{c}}$, $\ln \tau_\gamma$, the entries that carry the spreads; the standardized deviations of the non-centered parametrization.
$\mathbf{\Omega}$, $\mathbf{\Lambda}_i$, $\mathbf{\Lambda}_\mu$, $\mathbf{\Lambda}_{\text{pr}}$, $\mathbf{\Sigma}_i$	varies	Population covariance $\text{diag}(\tau_{\text{c}}^2, \tau_\gamma^2)$; coil i's information given the spreads and means; the information on the means; their prior information; coil i's covariance with the means integrated (Chapter 9).

Symbol	Unit	Meaning
Ψ, ϖ_m	–	The Gaussian rows' sum of squares at the conditional minimum; the weight of node m of a grid on the state: the spreads' grid (Chapters 9 and 20), the log-rate's (Chapter 15), the age slope's (Chapter 16).
π_i, $\mathcal{K}_i$, $\tilde{\gamma}_0$	–	Fleet position: coil i's percentile in the design's population of rates and its rank among the fifty (1 the fastest); coil 0's standardized rate.
γ_{new}, $\hat{\gamma}_0^{\text{chart}}$	yr^{-1}	The rate of a new coil of the design (the posterior predictive); the fleet chart's slope of coil 0.
$p_\tau(\lambda)$, $\bar{p}_\tau$, R_{pred}	–	The failure probability over a window τ at the rate λ, $1 - \mathrm{e}^{-\lambda\tau}$; the mission failure probability, its predictive mean $\mathbb{E}[p_\tau(\lambda) \mid \mathbf{y}]$; the predictive survival, $R_{\text{pred}}(\tau \mid \mathbf{y}) = 1 - \bar{p}_\tau$ (Chapter 15).
p^*, λ^*, τ_p	–, yr^{-1}, yr	A limit on a failure probability, 5 %; the rate at which p_τ reaches it, $-\ln(1-p^*)/\tau$; the window length at which the unit's own failure probability reaches it, $-\ln(1-p^*)/\lambda$ (Chapter 15).
τ_{sd}	yr	The time from window 36 to the October shutdown, 0.75 yr: the planning date t_{sd} as a window length (Chapter 15).
n_ℓ, E_ℓ, μ_ℓ	–, coil-yr, –	Count and exposure of count row ℓ and its expected count $E_\ell\lambda$: Chapter 15's four rows on the age band, the fourth the prior.
Λ, $\sigma_{\omega,a}$, ΔE	–, –, coil-yr	The information on the log-rate, $a_\lambda + k$, the sum of the rows' expected counts; the log-rate's spread after a candidate log of ΔE further coil-years in the band (Chapter 15).
ω_m, ω_q, λ_q	–, –, yr^{-1}	A node of the grid on the log-rate, weighted by ϖ_m; the q-quantile of the log-rate and of the rate, its image by the quantile map (Chapter 15).
$\lambda_{\text{pred}}(\tau)$	yr^{-1}	The predictive hazard after a time τ without a failure, $a'_\lambda/(b'_\lambda + \tau)$ (Chapter 15).
M_{own}, M_{sh}, $B_{\text{own,sh}}$, $\mathbf{y}_0$, $\mathbf{y}_{\text{sib}}$	–	The drawings in which coil 0's year in the band has its own rate and in which it shares the siblings'; the Bayes factor between them; coil 0's own record and the siblings' (Chapter 15).
$\bar{a}_a$, μ_a, N_f, a_w	yr, –, –, yr	The mid-age of bin a, $a + \frac{1}{2}$; its expected count $E_a\lambda(\bar{a}_a)$; the log's total failures, 32; the failure-weighted mean age $\sum_a n_a\bar{a}_a/N_f$, the age of the band's waist (Chapter 16).
$\mathbf{X}_a$, $\mathbf{k}$, a_{o}	–, –, yr	The log-count Jacobian $\partial \ln \boldsymbol{\mu}_a/\partial \mathbf{z}$, with rows $(1, \bar{a}_a)$; the bins' counts $(n_0, \ldots, n_{11})$; an origin of age at which the level is read, $\ln\lambda(a_{\text{o}})$ (Chapter 16).
$\mathcal{E}(\theta_a)$, Γ_{reg}, θ_m	coil-yr, –, yr^{-1}	The age-weighted exposure $\sum_a E_a\mathrm{e}^{\theta_a\bar{a}_a}$, the rate of the level's gamma conditional $\text{Gamma}(N_f, \mathcal{E}(\theta_a))$; the regularized incomplete gamma function; a node of the grid on the slope, weighted by ϖ_m (Chapter 16).

Symbol	Unit	Meaning
t_{dbl}, t_{lim}, $I_\tau(\theta_a)$	yr, yr, yr	The doubling time of the rate, $\ln 2/\theta_a$, infinite for $\theta_a \le 0$; the span within which a doubling is asked (Chapter 16); the integral of $e^{\theta_a a}$ over ages 3 to $3+\tau$, so that coil 0's expected count over the span is $\lambda_0 I_\tau(\theta_a)$.
s_θ, M_{age}, M_{const}, B_{age}	yr^{-1}, –, –, –	The prior spread of the slope θ_a in the evidence; the drawings of the age curve and of the constant rate, $\theta_a = 0$; the Bayes factor of the first against the second (Chapter 16).
θ_g, n_{ia}, μ_{ia}, $g_i(a)$	–, –, –, –	The log-rate's coefficient on a coil's log-derate, a covariate; coil i's failures and expected count in its year of age a; its line $g_{i,0} + \gamma_i a$ from Chapter 9's fit (Chapter 16).
θ_{aa}, c_{p}, $\theta_{\ln a}$	yr^{-2}, –, –	The quadratic coefficient of a log-rate curved in age; the level and the exponent of the power law in age (Crow–AMSAA), $\ln\lambda = c_{\text{p}} + \theta_{\ln a}\ln a$: the curve-form checks of Chapter 16.
Z_i, ν, $\mu_{i\cdot}$	–	Coil i's gamma frailty, its multiplier on the rate, mean one; the frailty's coefficient of variation, the unit-to-unit spread of failure proneness; coil i's expected total over its years (Chapter 16).
E_b^r, $\mathbf{J}_r$, $\mathbf{\Sigma}_r$	coil-yr, varies, varies	The exposure a candidate record r adds in age bin b; its Jacobian lines; the preposterior covariance $(\mathbf{\Sigma}^{-1} + \mathbf{J}_r^{\mathsf{T}}\mathbf{J}_r)^{-1}$ (Chapter 16).
a_i, t_i^*, $g_i(a)$	yr, yr, –	Fleet coil i's age today; the time after today at which its log-derate reaches g_F, negative (-1 yr) for a coil already past it and ∞ for one whose line does not fall; its log-derate at age a, $g_{i,0} + \gamma_i a$ (Chapter 20).
$\mathbf{x}_i$, $\hat{\mathbf{x}}_i$, $\boldsymbol{\ell}_m$	–	Coil i's pair $(g_{i,0}, \gamma_i)$ and its conditional mean (Chapter 9); node m of the spreads' grid $(\ell_{\text{c}}, \ell_\gamma)$, of weight ϖ_m (Chapters 9 and 20).
$f_{1,i}$, $f_{2,i}(t_{\text{ah}})$, $\mathbf{g}_{1,i}$, $\mathbf{g}_{2,i}(t_{\text{ah}})$	–	Coil i's margin over the functional limit today and t_{ah} years ahead, $g_i(a_i) - g_F$ and $g_i(a_i + t_{\text{ah}}) - g_F$; their gradients on $\mathbf{x}_i$, $(1, a_i)$ and $(1, a_i + t_{\text{ah}})$ (Chapter 20).
$\bar{f}_{1,i}$, $\bar{f}_{2,i}(t_{\text{ah}})$, u_1, u_2, $u_{1,i,m}$	–, –, –, –, –	The margins' means at a node of the grid; their standardized levels $-\bar{f}/\sqrt{\text{Var}(f)}$; u_1 for coil i at node m (Chapter 20).
$P_i(t_{\text{ah}})$, G_F	–, kW/K	Probability that coil i crosses its functional limit after today and by time t_{ah}, $\Pr(0 < t_i^* \le t_{\text{ah}} \mid \mathbf{y})$; the functional limit on Part I's hall, 5.0 kW/K on its design value, 5.8333 kW/K on a 7 kW/K coil (Chapter 20).
The UPS battery string		
$C_{\text{bat},t}$, $R_{\text{bat},t}$	Ah, mΩ	Usable capacity and ohmic resistance of the string in test t; $\bar{C}_{\text{bat}}$, $\bar{R}_{\text{bat}}$ their ratings, 100 Ah and 40 mΩ.
ξ_t, h_C, ΔR	–, –, mΩ	Log-capacity $\ln(C_{\text{bat},t}/\bar{C}_{\text{bat}})$; capacity against commissioning; resistance rise since commissioning.
v_0, Δv, h_{ch}, N	V, V, V, –	A block's loaded curve at the usable-empty point, its rise to full, and the charge curve's offset above it; the number of blocks, 40.
I_d, I_c, SOC	A, A, –	Discharge and recharge currents; state of charge.

Symbol	Unit	Meaning
$V_{d,t}$, $\dot{V}_{d,t}$, $V_{c,t}$, $\dot{V}_{c,t}$	V, V/h	Level and slope of test t's discharge line and of its recharge line.
$T_{\mathrm{run},t}$, V_{cut}, $\mathrm{SOC}_{\mathrm{cut},t}$, $V_{\mathrm{h},t}$	h, V, –, V	Run time to the cutoff; the cutoff on the string, 420 V; the state of charge there; the loaded head $N(v_0 + \Delta v) - V_{\mathrm{cut}} - I_d R_{\mathrm{bat},t}/1000$.
σ_V, $\sigma_{\mathrm{run},t}$, $\mathbf{s}$, $\mathbf{n}$	V, h, mixed	The voltmeter's spread; the run-time row's spread; the two discharge lines as coordinates; the ridge (null direction) of the discharge rows.
Probability		
$\Pr(A)$, $\Pr(A \mid B)$	–	Probability of event A; conditional on event B.
$p(\cdot)$	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, φ its density.
Inference objects		
$\mathcal{L}(\mathbf{z})$	–	The objective, minus the log posterior up to a constant.
$\hat{\mathbf{z}}$	varies	The maximum a posteriori (MAP) state, the minimizer of $\mathcal{L}$.
$\mathbf{J}$	varies	Jacobian of the normalized rows with respect to the state, evaluated at $\hat{\mathbf{z}}$.
$\mathbf{\Lambda}$	varies	Information matrix (Gauss–Newton Hessian of $\mathcal{L}$), $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ for normalized rows.
$\mathbf{\Sigma}$, Σ_{ij}	varies	Posterior covariance of the state under the Laplace approximation, $\mathbf{\Lambda}^{-1}$; its entries. The posterior spread of entry i is written $\sqrt{\Sigma_{ii}}$.
$f(\mathbf{z})$, $\mathbf{g}$	varies	A derived quantity of the state and its gradient $\partial f/\partial \mathbf{z}$ at $\hat{\mathbf{z}}$.
$\mathbf{N}$, $\mathbf{J}_{\mathrm{red}}$	varies	Null-space basis of the physics rows; the Jacobian of the reduced state (Chapter 4).
f_L, f_U	varies	The ends of the credible interval of a derived quantity.
ϱ, $\boldsymbol{\varrho}$, Φ_2	–	A correlation coefficient (Chapters 2 and 3; in Chapter 24 that of two consecutive logged readings' errors); the posterior's correlation matrix (Chapter 7); the distribution function of two standard normals with correlation ϱ.
υ, $\mathbf{\Upsilon}$, S_{in}	–	Weight of a reading; the gain of the correction form; the innovation variance of one reading (Chapter 4).
x_j, w_j	–	Gauss–Hermite nodes and weights.
ϖ_j, N_{eff}, p_{L}	–	Weight of draw j in the reweighted draws; their effective sample size $1/\sum_j \varpi_j^2$; the Laplace normal used as the proposal (Chapter 5).
ℓ_t, ε_q	–	Magnitude of the slope of window t's reading row; the summed departure of unread windows from the line (Chapter 4).
y_{new}, $\tilde{y}_a$	varies	A future reading (the predictive) and the outcome of a candidate measurement a before it is taken (the preposterior).

Symbol	Unit	Meaning
U_{p}, $U(a)$	varies	A quantity of the posterior after $\tilde{y}_a$; its expectation over the outcomes, which ranks the candidates (Chapter 2's design form).
M, M_j, B_{AB}	–	A drawing (model) of the plant; the members of a family of candidates; the Bayes factor $p(\mathbf{y} \mid M_A)/p(\mathbf{y} \mid M_B)$.
M_{clean}, M_{fouled}, M_{bias}, B_{fc}, B_{fb}	–	Chapter 2's drawings of a warm condition test: the air path clean, the air path fouled, the hall sensor with a bias; the Bayes factors of fouled against clean and of fouled against bias.
M_{one}, M_{two}, $B_{\mathrm{two,one}}$	–	Chapter 4's alternatives for the hall's air path, a single thermal resistance and two distinct heat-removal paths (beside the sensor bias M_{bias}); the Bayes factor between them.
$\mathbf{J}_a$, $\mathbf{\Lambda}_a$, $\mathbf{\Sigma}_a$	varies	The Jacobian rows of a candidate measurement, its information $\mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a$, and the preposterior covariance $(\mathbf{\Sigma}^{-1} + \mathbf{\Lambda}_a)^{-1}$.
P_M, $\mathcal{M}$, $\mathcal{M}_k$	–	The posterior probability $\Pr(M \mid \mathbf{y})$ of a drawing, its weight in the mixture; a group of drawings that share one factor, the curvature state or the step's kind (Chapter 14); the groups of one open question, whose partial perfect information Chapter 17 values.
$\hat{\mathbf{z}}_M$, $\mathbf{\Sigma}_M$, $\mathbf{\Sigma}_M^{[\cdot]}$, $\mathbf{\Sigma}_{M,a}$	varies	The MAP state and Laplace covariance of drawing M; its block on the entries the forecast reads (five, $(g_{1,36}, g_{2,36}, \gamma_1, \gamma_2, \dot{\gamma}_1)$, in Chapter 14; seven with the commissioning entries in Chapter 17); its preposterior covariance after candidate a, $(\mathbf{\Sigma}_M^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a)^{-1}$ (Chapter 14 on).
$\mathbf{g}_{t^*}$, $\mathbf{g}_*$, $\mathbf{g}_s$, $\mathbf{g}_G$	varies	The pattern of the crossing's gradient on a drawing's state; its gradient on the five forecast entries at the MAP, behind the delta band; the gradient of $G_{\mathrm{tot},s}$ on those entries; that of $G_{\mathrm{tot},36}$ on today's log-derates (Chapter 14).
$\mathbf{y}_{0:t_{\mathrm{as}}}$, $\tilde{\mathbf{y}}_{37:45}$	varies	The readings of windows 0 to t_{as} with the prior rows, an as-of fit's evidence; the readings of windows 37 to 45 not yet logged, drawn from their predictive (Chapter 19).
$\boldsymbol{\theta}_t$, $\bar{\boldsymbol{\theta}}_t$, $\bar{\boldsymbol{\theta}}_t^-$	mixed	The carried entries $(g_{1,t}, g_{2,t}, \gamma_1, \gamma_2, \dot{\gamma}_1, \beta_1, \dot{b}_{d,1}, \delta)$, all a later window's rows touch; their mean after window t's rows; before them, with the trend applied (Chapter 19).
$\mathbf{\Sigma}_t$, $\mathbf{\Sigma}_t^-$, $\mathbf{T}_t$, $\mathbf{S}_t$	varies	Covariance of $\boldsymbol{\theta}_t$ after and before window t's rows; the trend rows' transition from window t to $t+1$; the innovation covariance $\mathbf{J}_t\mathbf{\Sigma}_t^-\mathbf{J}_t^{\mathsf{T}} + \mathbf{I}$ of a look, with the gain $\mathbf{\Upsilon}_t$ of the correction form (Chapter 19).
$P_M(t)$, $\mathbf{m}$, $\mathbf{V}$, $m_{1\mid2}$, $s_{1\mid2}$	–, –, –, –, –	A drawing's probability after the rows of window t, $P_M(36) = P_M$; mean and covariance of $(g_{1,t_{\mathrm{as}}+1}, g_{2,t_{\mathrm{as}}+1})$ one window ahead; the conditional mean and spread of g_1 given g_2 in the region's integral (Chapter 19).

Symbol	Unit	Meaning
$\mathbf{A}$, $\boldsymbol{\Sigma}_0$	varies	The map from the state to the quantities a block of correlated readings reads, and the prior covariance, in the precision $\boldsymbol{\Sigma}^{-1} = \boldsymbol{\Sigma}_0^{-1} + \mathbf{A}^{\mathsf{T}}(\sigma^2\mathbf{R})^{-1}\mathbf{A}$ (Chapter 24).
Decisions		
a, $L(a, \mathbf{z})$	–, R$	An action (repair, replace, inspect, wait); its loss when the state is $\mathbf{z}$.
$\mathcal{R}(a)$, c_a, C_e	R$	The expected loss of action a, $\mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]$; its fixed cost; the consequence of event e.
c, c^*, c_1^*	varies, –, –	A decision level on a monitored parameter (the work-order trigger), from Chapter 19 a level on the probability $q_{t_{\text{as}}}$; the level of least expected loss with a refit every month, 0.0197; the one-look level $C_{\text{false}}/(C_{\text{false}} + C_{\text{miss}})$, 0.0980 (Chapter 19).
κ	–	Assurance level of an acceptance rule: the posterior probability a requirement must reach (0.95 in Chapter 22's verdict; Chapter 21's break-even p_{be} replaces it with a level the costs set).
EVSI, EVPI, EVPPI	R$	Expected value of sample information: what a measurement is worth before an action. Expected value of perfect information, $\mathcal{R}(a^*) - \mathbb{E}[\min_a L(a, \mathbf{z}) \mid \mathbf{y}]$: what knowing the whole state would be worth, the most any measurement can be (Chapter 17, its general form; Chapter 20 specializes it to the fleet's orders). Expected value of partial perfect information: the same with one open question known instead of the whole state (Chapter 17).
SF, OW, FO, RF, FC	R$, R$, R$, –, R$	Part V's cost schedule, each a stated input: a month with the hall at or below G^* (SHORTFALL, 150 000 per month; Chapters 17–19); the premium for work away from an October outage (OFF_WINDOW, 12 000 a visit; Chapters 17–19); a work order raised while the hall would have held (FALSE_ORDER, 3 000; Chapters 18 and 19); the share of the loss since commissioning that a clean restores (RESTORE_FRACTION, 0.85; Chapter 17, named in words in Chapter 20); a fleet coil crossing its functional limit before its visit (FLEET_CROSS, 40 000; Chapter 20). CLEAN (8 000 a coil), REPLACE (60 000 a coil) and INSPECT (1 500 a crew visit at reference grade) are named in words.
$\mathcal{P}$, $\mathcal{P}^*$, $\mathcal{P}_{\text{v}}$, $\mathcal{P}_{\text{dash}}$, $\mathcal{P}_{\text{cal}}$, $\varnothing$	–	A maintenance plan, the work at each visit of the horizon (one of 256); the best plan of the chosen action a^*; a one-visit plan; the dashboard's and the calendar's schedules; the plan of no work (Chapter 17). An action a is a class of plans, priced by its best plan.
$\mathcal{O}$, t_{v}, $W_t(\mathcal{P})$, $c_{\mathcal{P}}$	window, window, R$, R$	The October outages of the horizon, $\{46, 58, 70, 82, 94\}$; a visit's window; the price of the work plan $\mathcal{P}$ does at window t (CLEAN for each coil cleaned, REPLACE for each replaced); the plan's discounted work with the off-window premium, its c_a (Chapter 17).

Symbol	Unit	Meaning
$\mathrm{df}(t)$, $\mathcal{U}_t$	–, –	The discount factor to today, $1.08^{-(t-36)/12}$; the shortfall event of month t, $\{G_{\mathrm{tot},t} \le G^*\}$, whose consequence is $\mathrm{SF}\,\mathrm{df}(t)$ (Chapter 17).
$d^+_{u,t}$, $g^{\mathcal{P}}_{u,t}(\mathbf{z};\mathcal{P})$, $G_{\mathrm{tot},t}(\mathbf{z};\mathcal{P})$	–, –, kW/K	Coil u's derate factor after a clean at window t, $d_{u,0} - (1-\mathrm{RF})(d_{u,0} - d_{u,t})$; the value of $g_{u,t+1}$ at which the row from window t is zero under plan $\mathcal{P}$; the coils' total conductance once the plan's work is done (Chapter 17).
S_{sd}, $\mathbf{g}_L$	R\$, varies	The expected discounted shortfall from window 37 to the shutdown with no work, set against OW; the pattern of the loss on the extended state, whose gradient is zero wherever it is defined (Chapter 17).
G_{P}, m_{P}, t_{P}	kW/K, kW/K, window	The potential-failure level $G^* + m_{\mathrm{P}}$, the P of the P–F interval, at which a condition test raises the work order, 11.261 kW/K; its margin over the capacity limit, 0.5 kW/K; the first passage of $G_{\mathrm{tot},t}$ to G_{P} (Chapter 18).
Δ_{PF}, $\mathbf{g}_{\mathrm{PF}}$	window, varies	The P–F interval of a path, $t^* - t_{\mathrm{P}} \ge 0$, in months; its gradient on the forecast entries for the delta band (Chapter 18).
Δ_{I}, s_1, $\mathcal{S}$	window, window, –	An inspection interval; the window of a schedule's first test, its phase; the set of test windows, $\mathcal{S}(\Delta_{\mathrm{I}}, s_1) = \{s_1, s_1 + \Delta_{\mathrm{I}}, \ldots\}$ or a dated set (Chapter 18; Chapter 22's $\mathcal{S}$ is likewise a schedule of looks).
$\hat{G}_s$, σ_{I}, $p_{\mathrm{no},s}$	kW/K, kW/K, –	A condition test's estimate of $G_{\mathrm{tot},s}$; its standard deviation, 0.1033 kW/K at the reference grade; the probability, given the path, that the test at window s raises no order, $\Phi((G_{\mathrm{tot},s} - G_{\mathrm{P}})/\sigma_{\mathrm{I}})$ (Chapter 18).
$\mathbf{r}_{\mathrm{I}}$, $\mathbf{z}_{\mathrm{I}}$, $\mathbf{J}_{\mathrm{I}}$, $\mathbf{\Sigma}_{\mathrm{I}}$, $\mathbf{g}_G$	varies	A condition test's rows on its window's four entries $(Q_s, g_{1,s}, g_{2,s}, \beta_1)$; those entries; the rows' Jacobian and covariance; the gradient of $G_1 + G_2$ on them, with $\sigma_{\mathrm{I}}^2 = \mathbf{g}_G^{\mathsf{T}} \mathbf{\Sigma}_{\mathrm{I}} \mathbf{g}_G$ (Chapter 18).
P_{U}, Δ_5, Δ^*	–, window, window	Probability of an undetected crossing: of an interval, averaged over its phases and the drawings, or of a dated set; the longest interval that keeps it, and every shorter one's, at or below 0.05; the interval of least expected cost (Chapter 18).
C_V, N_V, N_{sf}, $s^+(t^*)$	R\$, –, window, window	The price of one test visit reading both coils, two INSPECTs, 3 000; the number of visits of a program; its months at or below G^*; the first test window at or after the crossing (Chapter 18).
$\mathcal{U}$, $q_{\mathcal{U}}$	–	The event a work-order decision turns on, the hall below its limit at a window an order raised now would protect, the shortfall event $\mathcal{U}_t$ of that window; its posterior probability (Chapter 19).
$q_{t_{\mathrm{as}}}$, Δ_{L}	–, window	The trigger's probability at look t_{as}, $\Pr(G_{\mathrm{tot},t_{\mathrm{as}}+\Delta_{\mathrm{L}}} \le G^* \mid \mathbf{y}_{0:t_{\mathrm{as}}})$, $t_{\mathrm{as}} = 36, \ldots, 45$; the lead time from an order to the work, one window (Chapter 19).

Symbol	Unit	Meaning
C_{false}, C_{miss}	R$	The extra cost of an order when the event does not occur, FO + OW, and of waiting when it does, SF − OW (Chapter 19).
$T_{\text{ord}}(c)$, $\mathcal{R}(c)$, $\mathcal{R}_a(c)$	window, R$, R$	The first look at which $q_{t_{\text{as}}} \geq c$, window 46 if none; the rule's expected loss over the records not yet logged; the same with candidate measurement a's rows (Chapter 19).
$\text{OC}(c;\ \cdot)$	–	Operating characteristic of a rule: its probability of firing at one look or more over records drawn at a stated truth, $\text{OC}(c;$ no loss) for the false orders (Chapter 19; Chapter 22 writes $\text{OC}(z_i^{\text{t}}; a)$ for a monitoring plan).
$\mathcal{V}$, v_i, $t_{\text{ah},v}$	–, week, yr	A visiting order of the fifty fleet coils; the week in which it visits coil i; the time of week v after today, $(v-1)/52$ (Chapter 20).
$\mathcal{V}^*$, $\mathcal{V}'$, $\mathcal{V}_{\text{dash}}$, $\mathcal{V}_{\text{now}}$	–	The order of least expected consequence, by linear assignment; another order; the dashboard's worst-looking-first order; the present plan, before a candidate's outcome (Chapter 20).
X_{iv}, v_i^*	–, week	The assignment matrix, one when coil i is visited in week v; the week in which coil i crosses on a draw, $\lceil 52\, t_i^* \rceil$, so that a visit in week v comes too late when $1 \leq v_i^* \leq v-1$ (Chapter 20).
$\mathcal{W}_i(\Delta)$, Δ	R$, yr	Cost of waiting of coil i over a delay, FC $P_i(\Delta)$; the delay (Chapter 20).
γ_{S}	yr^{-1}	The maker's specification as a threshold on a coil's own rate, $\ln(1-0.03) = -0.030459$: faster than specification when $\gamma_u < \gamma_{\text{S}}$ (Chapter 21).
P_u, P_u^0, B_u^{S}	–	Posterior and prior probabilities that coil u is faster than specification; the Bayes factor of the specification, posterior odds of the two halves over their prior odds (Chapter 21).
M_{fast}, M_{within}	–	The claim's drawing with γ_u's prior cut to $\gamma_u < \gamma_{\text{S}}$, and to $\gamma_u \geq \gamma_{\text{S}}$, each renormalized (Chapter 21).
CF, CV, p_{be}	R$, R$, –	Cost of filing a claim, R$ 5 000; value of a granted claim, a replacement coil, R$ 60 000; the break-even probability CF/CV = 1/12 above which filing has the lower expected loss, in place of an assurance level κ (Chapter 21).
$M_{2,15}^{(d)}$, q_d	–	The March step drawing with the curvatures of the coils in $d \in \{0, 1, 2, 12\}$ freed (none, coil 1, coil 2, both); its probability within the reported family of four (Chapter 21).
M_{arg}, M_{claim}	–	A counter-argument's drawing (the maker's alternative account of the record); the claim's drawing (Chapter 21).
$\hat{\gamma}_{u,d}$, $\varsigma_{u,d}$	yr^{-1}	Mean and spread of γ_u under drawing $M_{2,15}^{(d)}$ (Chapter 21).

Symbol	Unit	Meaning
γ_u^{td}, γ_u^{5y}	yr^{-1}	The loss to date, $(g_{u,36} - g_{u,0} - \mathbf{1}[u = 2]\,\delta)/(36\,\Delta t)$, the conceded step removed; the five-year rate, the mean of $\gamma_u(t_{\mathrm{y}})$ over the specification's five years, $\gamma_u + 1.0\,\mathrm{yr} \cdot \dot{\gamma}_u$ (Chapter 21).
$\mathbf{g}_{\mathrm{td},u}$, $\mathbf{g}_{\mathrm{5y},u}$	yr^{-1}	The gradients of the loss to date and of the five-year rate on the state (Chapter 21).
EVPI_u, EVSI_u	R$	A claim's expected value of perfect information, $\mathrm{CF}(1 - P_u)$, the cost of a claim that fails; a candidate measurement's expected value of sample information on coil u's claim (Chapter 21).
ς'_u, $\hat{\gamma}'_u$, P'_u	yr^{-1}, yr^{-1}, –	The spread of γ_u after a candidate measurement; the posterior mean it will leave, $\mathcal{N}(\hat{\gamma}_u, \varsigma_u^2 - \varsigma_u'^2)$ before its outcome; the probability of a breach after it (Chapter 21).
$\mathcal{D}(\Delta_{\mathrm{d}})$, $\mathcal{R}_{\mathrm{proof}}$	–, R$	The discount factor of a payment Δ_{d} years ahead at 8 % a year, Δ_{d} the delay; the expected loss of waiting under a standard of proof, by which a claim is filed only once $P'_u \geq 0.95$ (Chapter 21).
The value of monitoring		
ς_u, $\varsigma_u(t_{\mathrm{as}})$	yr^{-1}	The posterior standard deviation of coil u's rate, the spread of its rate (Chapters 21 and 22); in Chapter 22 the predictive spread of a record to window t_{as}, $\sqrt{[\boldsymbol{\Sigma}_{t_{\mathrm{as}}}]_{\gamma_u\gamma_u}}$, fixed before the record is read and equal to the spread of the rate's estimate over the records the plan may log. Chapters 7 and 11's ς_j is a quantity's standard deviation under drawing M_j, the same object within one drawing.
γ_{D}, $\gamma_{\mathrm{MD},u}(t_{\mathrm{as}})$	yr^{-1}	The stated loss a plan must detect, $\ln 0.95 = -0.05129$; the minimum detectable rate, $(z_\kappa + z_{\mathrm{POD}^*})\,\varsigma_u(t_{\mathrm{as}})$ (Chapter 22).
$\mathrm{OC}(z_i^{\mathrm{t}}; a)$, OC^*, z_i^{t}	–	Operating characteristic of a measurement or plan a: the probability that its outcome meets an acceptance rule when entry i of the state has the stated value z_i^{t}; the level it must reach (Chapter 22; Chapter 19's $\mathrm{OC}(c;\ \cdot)$ is the same object for a trigger).
$\mathrm{POD}_u(t_{\mathrm{as}})$, $\mathrm{PFA}_u(t_{\mathrm{as}})$, POD^*	–	Probability of detection of coil u's loss by the record to window t_{as}, the operating characteristic at $\gamma_u = \gamma_{\mathrm{D}}$; the false-alarm probability, the same at $\gamma_u = 0$; the power asked, 0.9 (Chapter 22).
z_κ, z_{POD^*}	–	The standard normal quantiles $\Phi^{-1}(\kappa) = 1.6449$ and $\Phi^{-1}(\mathrm{POD}^*)$ (Chapter 22).
$T_{90,u}$, $T_{90,u}^B$	window	Record length for a power of 0.9: the first window at which POD_u reaches POD^*; the same under the Bayes-factor rule (Chapter 22).
$\boldsymbol{\Sigma}_{t_{\mathrm{as}}}$, $\mathbf{z}_{t_{\mathrm{as}}}^{\mathrm{D}}$, $\hat{\gamma}_{u,t_{\mathrm{as}}}$	varies, varies, yr^{-1}	Covariance of a record to window t_{as} not yet logged; the hypothetical state at which it is linearized, the coils on the stated loss's line; the rate's estimate from that record (Chapter 22).

Symbol	Unit	Meaning
$\mathbf{J}_{t_{\text{as}}}$, $\mathbf{J}_{\text{pr}}$, $\mathbf{J}_{0:t_{\text{as}}}$, $\mathbf{J}_{r,t_{\text{as}}}$, $\boldsymbol{\varepsilon}_{t_{\text{as}}}$	varies	Jacobian of the record to window t_{as}; its four prior rows' lines; its reading and trend rows' lines; its reading rows' lines alone; the normalized errors of its rows, $\text{Cov}(\boldsymbol{\varepsilon}_{t_{\text{as}}}) = \mathbf{I}$ (Chapter 22).
$\mathcal{S}$, $\text{P}_{\mathcal{S},u}$	window, –	In Chapter 22 the schedule of looks, the windows at which the growing record is refitted (Chapter 18's set of test windows); the probability that the verdict passes at one look of $\mathcal{S}$ or more, the false alarm of repeated looks at $\gamma_u = 0$ (Chapter 22).
σ_γ, B^*, ς_{tr}, $\hat{\gamma}_{\text{tr}}$	yr^{-1}, –, yr^{-1}, yr^{-1}	The trend's prior spread in the Bayes-factor rule, 0.10; the stated level of the factor, 19; the spread and mean of the rate under that prior (Chapter 22).
$z_u^B(t_{\text{as}})$, POD_u^B, PFA_u^B	–	The Bayes-factor rule as a threshold on $-\hat{\gamma}_{u,t_{\text{as}}}/\varsigma_u(t_{\text{as}})$, growing with the record; its probability of detection and false alarm (Chapter 22).
$\varsigma_u^r(t_{\text{as}})$, POD_u^r, PFA_u^r	yr^{-1}, –, –	The estimate's spread from the reading rows' errors alone, for a truth with no wander; the probability of detection and false alarm it gives (Chapter 22).
$\boldsymbol{\Sigma}_{a,t_{\text{as}}}$, $\varsigma_u'(t_{\text{as}})$	varies, yr^{-1}	Preposterior covariance of the record to window t_{as} with candidate instrument a; the predictive spread it leaves (Chapter 22; Chapter 21's ς_u' at window 36).
$\mathcal{B}_a$, $\mathcal{B}_{M,a}$, $\bar{\mathcal{B}}_a$, $\mathcal{B}_{\text{cont}}$	window	The 5–95 % band of the crossing window t^* from the refit at window 45 under candidate a, mixed over the family at today's P_M; drawing M's own band; the own bands weighted, $\sum_M P_M \mathcal{B}_{M,a}$; the band under the record's continuation alone (Chapter 23).
$\Delta\mathcal{B}_a$, $\Delta\bar{\mathcal{B}}_a$	window	The value of candidate a for the date: the windows it takes off the family's band and off the own bands beyond the continuation (Chapter 23).
$F_{t^*}(s)$, $F_{t^*,a}(s)$	–	The distribution function of the crossing window, $\Pr(t^* \le s \mid \mathbf{y})$ (Chapter 14); the family's under candidate a, $\sum_M P_M \Pr(t^* \le s \mid \mathbf{z}_M \sim \mathcal{N}(\hat{\mathbf{z}}_M, \boldsymbol{\Sigma}_{M,a}))$ (Chapter 23).
$\mathbf{J}_{\text{now}}$, $\mathbf{J}_{\text{cont}}$	varies	The lines of the record's rows to window 36 and of every prior and trend row to window 45; those of the existing readings of windows 37 to 45, the continuation, from which each drawing's preposterior $\boldsymbol{\Sigma}_{M,a}$ at the refit is built (Chapter 23).
$\mathbf{z}_{M,45}$	varies	Drawing M's state carried to the refit at window 45, with the water-flow factors ι_1, ι_2 appended (Chapter 23).
ι_u, C_w	–, kW/K	Log of coil u's chilled-water flow against the balancing report's, prior $\mathcal{N}(0, 0.10^2)$, constant from window 37; that report's chilled-water capacity rate per coil, 14.0 (Chapter 23).
$r_{W,u,t}$, $r_{d',t}$, $r_{\Delta,u,t}$, $r_{r',t}$, $r_{c,t}$, $r_{C,u,t}$	–	The candidates' rows in windows 37 to 45: a water-side heat meter on coil u; a second discharge sensor on unit 1; a water-side temperature difference on coil u, $D_{u,t}/(C_w \text{e}^{\iota_u})$; a new return-air sensor; a supply-air sensor; an airflow station on unit u (Chapter 23).

Symbol	Unit	Meaning
$\mathbf{\Sigma}_5$	varies	The covariance block of the five forecast entries $(g_{1,45}, g_{2,45}, \gamma_1, \gamma_2, \dot{\gamma}_1)$ at a drawing's MAP state, on which the crossing's gradient $\mathbf{g}_*$ acts (Chapter 23).
$\mathbf{g}_G$, $\mathbf{g}_{\dot{G}}$, $\mathbf{g}_{\Delta g}$, $\mathbf{g}_{\Delta\gamma}$, $\mathbf{g}_{\dot{\gamma}}$	varies	The gradients on the same five entries of the total's level, its rate of change, the split of the log-derates, the split of the rates and the curvature: the five directions of the date's variance (Chapter 23).
$\dot{G}_{\mathrm{tot},t}$	kW/K/yr	The rate of change of the coils' total conductance at window t (Chapter 23).
$\mathbf{Q}$	–	Rows of a set of directions $\mathbf{Qz}$ of the state taken as known exactly, in $\mathrm{Var}(f \mid \mathbf{Qz})$ (Chapter 23).
$c_a^{\times}$, $a^{\dagger}$	R$, –	The price at which candidate a would match the best permanent sensor $a^{\dagger}$ per real on the own bands, $\Delta\bar{\mathcal{B}}_a\, c_{a^{\dagger}}/\Delta\bar{\mathcal{B}}_{a^{\dagger}}$; c_a is the candidate's price, the fixed cost of the action of buying it (Chapter 23).
p_{46}	–	A predicted record's refit probability of a crossing by the shutdown, $\Pr(t^* \le 46 \mid \mathbf{y})$; the wrong-call chance at the outage is $\mathbb{E}[\min(p_{46}, 1-p_{46})]$ (Chapter 23).
n_{L}, $\Delta\vartheta$, ϑ_w	–, h, h	The logging rate, readings of each channel per window (1 monthly, 30 daily, 720 hourly); the logging interval $\vartheta_w/n_{\mathrm{L}}$; a window's span, 720 h (Chapter 24).
ϱ_1, ϱ	–	Correlation of two reading errors an hour apart, 0.9, the correlation time $-1/\ln\varrho_1 = 9.49$ h; the correlation of two consecutive logged readings' errors, $\varrho_1^{\Delta\vartheta/1\,\mathrm{h}}$ (Chapter 24).
e_i, ϵ_i	as the reading	The AR(1) error of reading i of a channel, $e_i = \varrho e_{i-1} + \epsilon_i$; its new part, independent of every earlier error, variance $\sigma_k^2(1-\varrho^2)$ (Chapter 24).
$y_{k,t,i}$, $\mathbf{y}_{k,t}$, $\bar{y}_{k,t}$, f_k	as the reading	Reading i of channel k in window t; the window's n_{L} readings stacked; their generalized least-squares mean, weights 1 at the ends and $1-\varrho$ inside; the model value the channel reads (Chapter 24).
$r_{k,t,i}$, $\mathbf{r}_{k,t}$, $\tilde{r}_{k,t,i}$, $\mathbf{R}$, $R_{ii'}$	–	The row of reading i; the window's n_{L} rows stacked; the whitened row, the new part of its error over its own standard deviation; the readings' correlation matrix, $\mathrm{Cov}(\mathbf{r}_{k,t}) = \mathbf{R}$, with entries $\varrho^{\lvert i-i'\rvert}$ (Chapter 24).
$\mathbf{j}_{k,t}$	varies	The one Jacobian line the n_{L} rows of a channel's window share, $\sigma_k^{-1}\,\partial f_k/\partial\mathbf{z}$ (Chapter 24).
$n_{\mathrm{eff}}(n_{\mathrm{L}})$, $n_{\mathrm{eff}}^{\mathrm{mean}}$	–	The effective number of readings of a window, $\mathbf{1}^{\mathsf{T}}\mathbf{R}^{-1}\mathbf{1} = (n_{\mathrm{L}} - (n_{\mathrm{L}}-2)\varrho)/(1+\varrho)$, which saturates at $1 + \frac{1}{2}\vartheta_w \ln(1/\varrho_1)$; the same for the simple average (Chapter 24).
f_{sh}, n^{AR}, $\tilde{n}$	–	The part of a reading's variance shared by the whole window; the AR(1) effective number of the remaining part; $n^{\mathrm{AR}}/(1-f_{\mathrm{sh}})$, so that $n_{\mathrm{eff}} = \tilde{n}/(1+f_{\mathrm{sh}}\tilde{n}) < 1/f_{\mathrm{sh}}$ (Chapter 24).

Symbol	Unit	Meaning
$\mathbf{J}_y$, $\mathbf{J}_o$, $\mathbf{J}_{y,\mathrm{f}}$	varies	The record's 185 reading rows; its other 76 rows (priors and trend rows, from window 37 every trend row); the reading rows of the forecast windows 37 to 45 (Chapter 24).
$\mathbf{\Sigma}_{n_\mathrm{L}}$, $\mathbf{\Sigma}^{\mathrm{line}}_{n_\mathrm{L}}$, $U(n_\mathrm{L})$	varies, varies, yr^{-1}	The preposterior of the record logged at rate n_L, $(\mathbf{\Sigma}^{-1} + (n_{\mathrm{eff}}(n_\mathrm{L}) - 1)\mathbf{J}_y^\mathsf{T}\mathbf{J}_y)^{-1}$; the same in the straight-line drawing, the readings' part; the rate's spread $\sqrt{[\mathbf{\Sigma}_{n_\mathrm{L}}]_{\gamma_u\gamma_u}}$, Chapter 2's utility $U(a)$ at that rate (Chapter 24).
n_p, $\mathbf{\Lambda}_Q$, $\mathbf{\Sigma}_\mathrm{h}$	–, kW^{-2}, varies	The weight of the record's past reading rows in year 4's preposterior, 1 as it stands and n_{eff} with the logger's archive read; the future loads' prior precision, $(1000\,\mathrm{kW})^{-2}$ on each; year 4's preposterior at hourly logging, the base of step 7 (Chapter 24).

Vocabulary

Plant graph 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.

State vector The collection of free quantities the fit estimates at once.

Pinned, free, known input A quantity treated as known (a boundary condition, or an input such as the supply temperature read by a precise instrument) against one carried in the state.

Residual, row A relation written so that it equals zero when satisfied: physics rows constrain the state, reading rows compare it with the records, trend rows tie one window's parameter to the next, prior rows carry handbook or commissioning knowledge.

Window The span of records fitted as one steady condition: one representative hour of a month here.

Health index A physical parameter of the model read against its commissioned value; never a score built from raw performance.

Condition monitoring Following the health index through the windows of a record.

Degradation rate The trend of the health index per year.

Remaining useful life The time until the health index crosses the level at which the asset can no longer do its job, as a date with a band.

P–F interval The time from the first detectable change of the health index to its crossing of the functional limit.

Condition-based maintenance Maintenance triggered by the health index crossing a decision level, rather than by the calendar.

Partial pooling Estimating each unit of a fleet with a shared prior whose mean and spread are themselves estimated, so that the fleet informs the unit and the unit informs the fleet.

Change-point detection Locating a step in a parameter's path by the evidence of a step freed at each candidate window.

State of health The capacity and internal resistance of a battery against their values when new.

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

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

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

Evidence, Bayes factor The normalizing constant of one model's posterior; the ratio of two models' evidences.

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

The five computations Posterior, trend, evidence, preposterior and expected consequence: the objects every asset-management question is answered from.

Part IV

Forecasting the Need for Attention

Chapter 14

Remaining Useful Life and Survival to the Next Outage

14.1 The dashboard's crossing date and the coils' remaining life

It is December of year 3, window 36, three years after the two-unit hall was handed over. The October shutdown of year 4, window 46, is the next outage at which the coils can be taken out of service for cleaning, and the planner must decide now whether the work goes into that outage or whether the hall cannot wait for it. The capacity study states the limit as a conductance: the two coils together must keep $G_1 + G_2 \geq 10.761\,\mathrm{kW/K}$ for rack 3 to stay at 27°C on the design day. The condition-monitoring dashboard gives each coil's health-index tile and its straight line, Chapter 10's: $1.0118 - 0.1576$ a year for coil 1 and $1.0154 - 0.1482$ a year for coil 2, with today's tiles at 0.4873 and 0.5947. The planner turns the lines into kW/K with the commissioning test's conductances, 6.7164 and 6.4638 kW/K, and gets $13.3588 - 2.0164\,t_\mathrm{y}\,\mathrm{kW/K}$ with t_y in years. That line reached the limit at $t_\mathrm{y} = 1.2884\,\mathrm{yr}$, window 15.46, in April of year 2, and today it reads 7.3097 kW/K, 3.4513 kW/K below the limit. The conclusion the chart invites is that the hall has been past its design capacity since April of year 2, and that both coils must be cleaned now, not at the October shutdown. With the rated 7 kW/K per coil as the reference the line crosses in August of year 2, window 19.22, and the conclusion is the same.

The chart cannot say whether the tiles' fall is the coils at all: Chapter 10 showed that a tile is a ratio of approaches, which falls with the load, the season and unit 1's drifting sensor. Nor can it say at what date the coils' own conductance reaches the limit, when its rate is uncertain and may itself be changing (Chapter 11), and when part of coil 2's loss may be a single step that does not continue (Chapter 12). And it cannot say how likely the hall is to hold until October, which is a probability, not a date. The question of this chapter is

when the sum of the two coils' conductances first falls below the capacity study's limit, read as the first passage of the path that the condition record's fit carries beyond the record, averaged over the drawings that Chapters 11 and 12 left open, with the probability that it holds until the October shutdown.

It is answered with the method of Chapter 5, on Chapter 10's stacked state extended to the forecast windows by the trend rows of Chapter 4, over the family of Chapters 11 and 12 averaged as in Chapter 7.

Jargon: forecast window, survival to an outage, forecast averaged over drawings, within- and between-drawing width

A *forecast window* is a window of the state past the last reading: it carries the parameters of that window and the trend rows that tie it to its neighbors, and no reading. *Survival to an outage* is the probability that the crossing comes after a stated planning date. A *forecast averaged over drawings* mixes the crossing's distribution under each drawing with the drawing's posterior probability as its weight. Its *within-drawing width* is the band each drawing gives on its own, averaged over the weights, and the *between-drawing width* is what the mixture's band adds to it. Chapter 2 defined the first passage, the plug-in date, the delta band, the band of a date by inversion and the mass that never crosses (Section 2.3.4); the box does not repeat them.

Key idea

A remaining life is a date read from the posterior of the path, not a line drawn through a chart. The physics puts the path on the coils' conductance, the trend rows carry it past the record with the wander they already allow, and the date is the first window at which the path falls below the limit: a quantity of the extended state, with a median and an inversion band but no mean. When the record has not settled how the path bends, the date is averaged over the drawings, and its band carries that indecision.

14.1.1 A date from the posterior of a path

Chapter 2 wrote the crossing of a level as the second line of the trend form (2.21), on a straight line with no wander, and worked it on Part I's hall (Figure 2.2). For a quantity that falls to its level, as a conductance does, the crossing by a date s there is the joint event that the trend is negative and the quantity stands at or below the level at s. That event is the first passage only because a straight line, once below, stays below. The path of this chapter is the trend rows' path, which wanders from month to month about its trend, so a path can dip below the level and come back. The crossing is then defined on the path itself. With $G_{\mathrm{tot},t} = G_{1,t} + G_{2,t}$ the coils' total conductance in window t, Chapter 13's equation (13.1), and G^* the limit, the first window at or below the limit is t_c, and the crossing time t^* is read between that window and the one before it by interpolating G_{tot} linearly:

$$t_c(\mathbf{z}) = \min\{t > 36 : G_{\mathrm{tot},t}(\mathbf{z}) \le G^*\}, \qquad t^* = t_c - 1 + \frac{G_{\mathrm{tot},t_c-1} - G^*}{G_{\mathrm{tot},t_c-1} - G_{\mathrm{tot},t_c}} \in (t_c - 1,\ t_c],$$

$$\{t^* \le s\} = \bigcup_{t=37}^{s} \{G_{\mathrm{tot},t} \le G^*\} \supseteq \{G_{\mathrm{tot},s} \le G^*\} \quad \text{for a window } s. \tag{14.1}$$

The date is the window $t_c = \lceil t^* \rceil$, with its month and year. The union counts a path that has been below the limit at any window up to s, whether or not it is below at s; the threshold at s alone counts only the paths below at s. The remaining useful life is t^* less today's window, $t^* - 36$, in months.

The windows after the record are entries of the state like the windows of the record. Each forecast window carries its two log-derates, tied to the previous window by the same trend row as the record's windows, and it has no reading. The record's entries keep their posterior, and each forecast window adds one month's wander to the trend the drawing carries. The crossing is a function of that extended state, and its distribution is read by the second line of equation (2.21) with the first passage in place of the joint event.

Principle 14.1. *A date is a ratio of a margin to a rate, so it has a median and a band but no mean, and its band is read by inverting the probability of having crossed by each date. A line drawn through raw performance gives a date with no band, on a quantity that is not the coils.*

14.1.2 The two-unit hall carried past its record (step 0)

The graph is Chapter 10's, Figure 10.1: Chapter 5's window graph at each of the 37 windows of the record, with its five physical rows (5.2) at that window's known inputs $T_s(t)$ and $Q_o(t)$; the 72 trend rows between consecutive windows; the two trend entries γ_1 and γ_2; the shared recirculation fraction β_1; and the drift node $\dot{b}_{d,1}$ on unit 1's discharge reading from window 20, equation (10.2). Chapter 11's curvature node $\dot{\gamma}_1$ hangs on coil 1's trend rows in the drawings that free it, and Chapter 12's step node hangs at window 15 in the drawings that carry a step, on the part the drawing names. The modeling choices of step 0 are Chapter 5's paragraph, carried through Chapter 10: steady state within each window, the airflow pinned with no discharge control, the season's quantities as known inputs, and one conductance per coil.

What the forecast adds to the graph. The forecast windows $37, 38, \ldots$ are path nodes $g_{1,t}$, $g_{2,t}$, with trend squares between them and no reading squares. Their IT loads do not enter: no reading row needs them, and the limit is stated at the design load. On every forecast window hangs a requirement node, $G_{\text{tot},t} \geq G^*$, with G^* a stated input, pinned. The shutdown is marked at window 46. One choice of step 0 is new here, and no question block owns it. *The forecast windows carry the record's trend rows unchanged*: the drawing's trend, its rate and, where it is freed, coil 1's curvature, with the wander $\sigma_q = 0.005$ per window. This buys a forecast that is the same fit, carried forward, with no second tool. It fails if the coils' condition changes by a cause the record has not seen, such as a cleaning, a valve restored or a new fouling regime; the forecast assumes no maintenance before the crossing.

The graph. The trend layer of both coils from window 34 on is the content of Figure 14.1. The upper chain is coil 1's path, the lower chain coil 2's; the black squares between neighbors are the trend rows. Windows 34 to 36 are read: each has a block between the chains for its load and five readings, drawn faintly because Figure 10.1 shows them. From window 37 on the windows have no block, and between the chains each carries its requirement node, linked to the pinned limit. Runs of windows are dashed boxes. The trend entries γ_1 and $\dot{\gamma}_1$ feed coil 1's squares from above and γ_2 feeds coil 2's from below, across window 36 as before it.

The record. The chapter reads the whole record as Chapters 10 to 13 read it, Table 10.1: window 0, the commissioning test, and windows 1 to 36, January of year 1 to December of year 3, five readings each. Three stated inputs enter here. The *limit* is the capacity study's $G^* = 10.761\,\text{kW/K}$ on $G_1 + G_2$: rack 3 at 27°C at 150 kW of IT load with the 15.5°C design summer supply. It is 0.8164 of the commissioning test's sum and 0.7686 of the rated 14 kW/K. The season enters the forecast through the limit's design supply, not through the trend, which Chapter 13 showed carries no season. The *planning date* is the October shutdown of year 4, window 46 (Oct Y4), ten windows ahead. And the *commissioning test's conductances*, its readings taken as exact (Chapter 5's $G_{u,\text{test}}$), 6.7164 and 6.4638 kW/K, sum 13.1803 kW/K, are what the dashboard uses to turn its tiles into kW/K.

The maintenance log so far. The two-unit hall's log holds the commissioning test of window 0; Chapter 5's condition assessment and Chapter 6's effort entry at window 18 (Jun Y2); Chapter 7's assessment at window 25 (Jan Y3), which held coil 1's cleaning order and

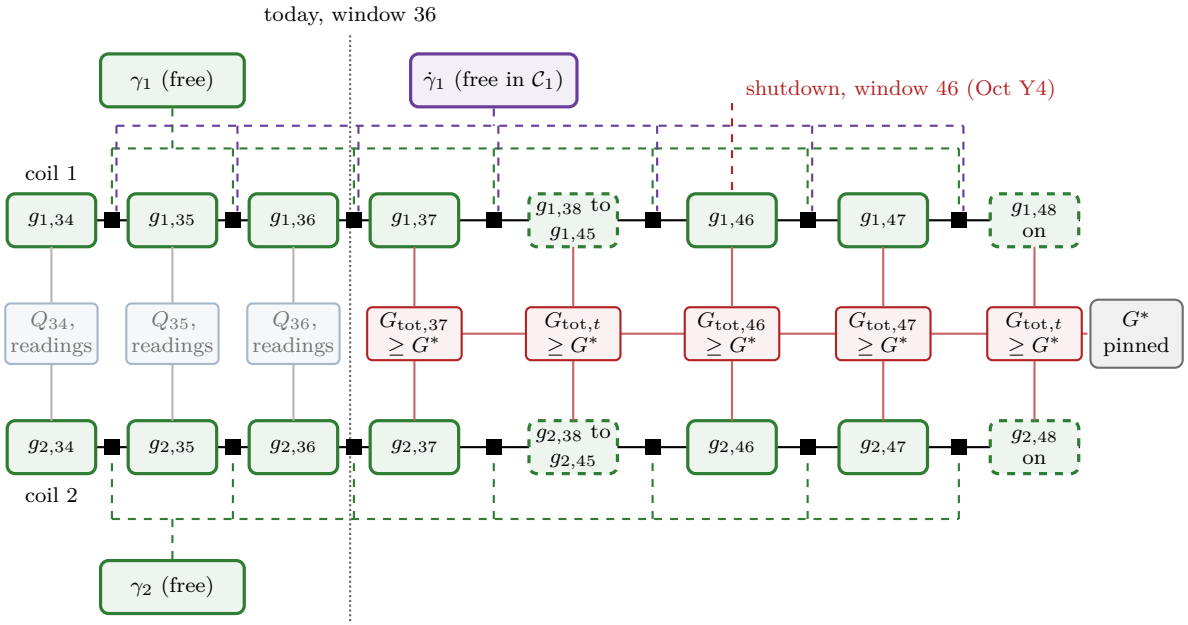


Figure 14.1: The trend layer of the two coils carried past the record, from window 34. Free (green): the log-derates $g_{u,t}$ of each coil, read windows 34–36 and forecast windows from 37 on (runs as dashed boxes); the trend entries γ_1, γ_2; and (purple) coil 1's curvature $\dot{\gamma}_1$ in the drawings that free it. Black squares: the trend rows (11.4), the same on either side of window 36. Faint blocks: the read windows' loads, physical rows and five readings (Figure 10.1); the forecast windows have none. Red: the requirement node $G_{\mathrm{tot},t} \geq G^*$ on every forecast window, linked to the pinned limit $G^* = 10.761\,\mathrm{kW/K}$, a stated input of the capacity study; the shutdown at window 46. Known inputs and pinned quantities of the read windows are not drawn; the step node of Chapter 12, at window 15, lies outside the windows drawn.

asked for a side-by-side check of unit 1's discharge sensor, a check not made; and four entries at window 36. Chapter 10's read the rates on the conductance as constants and confirmed the drift (Section 10.5.2). Chapter 11's left coil 1's acceleration undecided and wrote that the forecast of the capacity limit must carry both drawings, the October shutdown's work order waiting on it (Section 11.5.2). Chapter 12's found a step in March of year 2 likely, coil 2 the likely part, and wrote that the rates for the capacity forecast wait for the logger archive (Section 12.5.2). Chapter 13's found that the season is not aging (Section 13.5.2). This chapter writes the forecast those entries deferred.

14.1.3 The asset manager's questions (step 1)

The asset manager asks five questions of the record, and step 1 writes each as a probability over the extended state of the drawings of the family, with the unknown on the left of the bar and the 185 readings and the priors on the right. Table 14.1 lists them, with the columns of Chapter 5's table: the question in the asset manager's words, the two drill questions first; the general form of its category from Chapter 2; its mathematical form for the hall; its category; the computation of Section 2.3.6 that answers it; and the subsection where it is answered. The second and third questions are the second and third lines of the trend form at one date. What to do about the date, whether to clean at the shutdown, how often to inspect and at what reading to raise the work order, is the question of Part V (Chapters 17 to 19); which sensor most narrows the date is Chapter 23's. Which measurement would narrow the date asks nothing new of the record; it is step 7.

Table 14.1: The questions asked of the two-unit hall's condition record at window 36 (Dec Y3). $\mathbf{y}$ is the 185 readings of windows 0–36 and the prior rows; t^* the first passage (14.1); $G_{\mathrm{tot},t} = G_{1,t} + G_{2,t}$; M runs over the family's drawings, with $P_M = \Pr(M \mid \mathbf{y})$.

Asset manager's question	General form	Mathematical form	Category	Computation	Answered in
When does the conductance cross the level at which the hall can no longer be held?	(2.21), second line; (2.11) by inversion; (7.10)	$\Pr(t^* \le s \mid \mathbf{y}) = \sum_M P_M \Pr(t^* \le s \mid \mathbf{y}, M)$: median and 5–95 % band, as windows and dates; the remaining life $t^* - 36$	estimation	trend	14.3.1
Will it hold until the planned shutdown in October?	(2.21), second line at $s = t_{\mathrm{sd}}$	$\Pr(t^* > 46 \mid \mathbf{y})$, with $\Pr(t^* \le 45 \mid \mathbf{y})$ beside it	prediction under a threshold	trend	14.3.2
How much margin will be left at the shutdown?	(2.21), third line	$G_{\mathrm{tot},46} - G^*$: median and band	estimation	trend	14.3.3
Which open question moves the date, the acceleration or the step?	(2.9), (7.10), (7.11)	the forecast within each factor of the family; the band's width within the drawings against the mixture's	estimation	evidence	14.3.4
Is the hall already past its limit, as the dashboard says?	(2.12)	$\Pr(G_{\mathrm{tot},36} \le G^* \mid \mathbf{y})$	prediction under a threshold	posterior	14.3.5

14.1.4 The evidence as rows (step 2)

Step 2 attaches the record to the graph as the rows of each drawing of a family and extends every drawing by the forecast windows. The rows of the record are Chapter 10's, kept as they are: each window's five readings, unit 1's discharge row carrying the drift (10.2) from window 20, the rating prior on window 0's coils, the survey prior on β_1, the drift prior and the 72 trend rows (10.5) at $\sigma_q = 0.005$, 261 rows on the 115 entries of the stacked state (10.3).

The family. The forecast carries both questions that Part III left open, crossed. The first is coil 1's curvature, Chapter 11's: the curvature state $\mathcal{C}_0$, constant rates, the base fit; or $\mathcal{C}_1$, coil 1's curvature $\dot\gamma_1$ free with the prior $\mathcal{N}(0,\ 0.05^2)\,\mathrm{yr}^{-2}$, its rate at time t_y in years $\gamma_1(t_\mathrm{y}) = \gamma_1 + \dot\gamma_1(t_\mathrm{y} - \bar t_\mathrm{y})$ over the whole record, $\bar t_\mathrm{y} = 1.5\,\mathrm{yr}$, through the lever x_t of the trend rows (11.4). The second is the March step, Chapter 12's named-window family: no step, or one step into window 15 (Mar Y2) of one of the eight kinds of Table 12.2, coil 1 ($j = 1$), coil 2 ($j = 2$), both coils ($j = 12$), the return-air sensor (T_r), unit 1's and unit 2's discharge sensors ($T_{d,1}$, $T_{d,2}$), the load meter (Q) or rack 1's inlet sensor (T_1), with that chapter's priors on the step: 0.15 on a coil, 1.0 K on a temperature sensor, 10 kW on the meter. A drawing of the family is a pair $M = (\mathcal{C}_l, j)$, $l \in \{0, 1\}$ and $j \in \{\text{none}, 1, 2, 12, T_r, T_{d,1}, T_{d,2}, Q, T_1\}$: 18 drawings. $(\mathcal{C}_0, \text{none})$ is Chapter 10's base fit, $(\mathcal{C}_1, \text{none})$ Chapter 11's drawing that frees coil 1's curvature, and $(\mathcal{C}_0, j)$ Chapter 12's step drawing of kind j at window 15. Their priors are the two chapters' own, taken independent:

$$\begin{gathered}\Pr(\mathcal{C}_1) = \Pr(\mathcal{C}_0) = \tfrac{1}{2}, \qquad \Pr(j = \text{none}) = \tfrac{1}{2}, \qquad \Pr(j) = \tfrac{1}{2}\cdot\tfrac{1}{8} = 0.0625 \quad (j \neq \text{none}),\\ \Pr((\mathcal{C}_l, j)) = \Pr(\mathcal{C}_l)\ \Pr(j).\end{gathered} \tag{14.2}$$

The state of a drawing is the stacked state (10.3) with $\dot\gamma_1$ appended under $\mathcal{C}_1$ and the step δ_j under a step, 115 to 117 entries, and its rows are the 261 rows with the curvature's prior row and the step's column and prior row added, 261 to 263 rows: 146 degrees of freedom in every drawing.

The extended state. Each drawing's state is extended by the forecast windows $t = 37, \ldots, t_\mathrm{hor}$, with the horizon $t_\mathrm{hor} = 276$, twenty years past the record:

$$\mathbf{z}_\mathrm{ext} = \left(\mathbf{z}_M,\ \mathbf{g}_{37},\ \mathbf{g}_{38},\ \ldots,\ \mathbf{g}_{t_\mathrm{hor}}\right)^\mathsf{T}, \qquad \mathbf{g}_t = (g_{1,t},\ g_{2,t}), \tag{14.3}$$

with $\mathbf{z}_M$ the drawing's own state. The forecast windows add two entries and two trend rows each, and no reading row. Their trend rows are the rows (11.4) continued past window 36,

$$r_{q,u,t}(\mathbf{z}_\mathrm{ext}) = \frac{g_{u,t+1} - g_{u,t} - \Delta t\left[\gamma_u + \dot\gamma_u\, x_t\right]}{\sigma_q}, \qquad t = 36, \ldots, t_\mathrm{hor} - 1, \quad u = 1, 2, \tag{14.4}$$

with $x_t = (t + \frac{1}{2})\Delta t - \bar t_\mathrm{y}$ the curvature's lever about Chapter 11's center $\bar t_\mathrm{y} = 1.5\,\mathrm{yr}$, $\dot\gamma_2 = 0$ in every drawing and $\dot\gamma_1 = 0$ under $\mathcal{C}_0$. The lever grows with the distance from the record's center: $x_{36}\Delta t = 0.1285\,\mathrm{yr}^2$ on the row from window 36 to 37, 0.1910 on 45 to 46 and 0.2882 on 59 to 60. The quantity asked of the extended state is the total conductance at each window, record or forecast,

$$G_{\mathrm{tot},t}(\mathbf{z}_\mathrm{ext}) = \bar G_1\, \mathrm{e}^{g_{1,t}} + \bar G_2\, \mathrm{e}^{g_{2,t}}, \tag{14.5}$$

with $\bar G_u = 7.0\,\mathrm{kW/K}$ the ratings, and the crossing is its first passage (14.1), read in continuous windows by interpolation and dated by the first window at or below the limit. This is how the record's truth defines its own crossing window.

Which row acts on which entry. The Jacobian of a drawing's record is Chapters 10–12's pattern (10.6), with the curvature's column on coil 1's trend rows (11.5) and a step's column on its rows (Table 12.2). Each forecast row touches $g_{u,t}$, $g_{u,t+1}$, γ_u and, for coil 1 under $\mathcal{C}_1$, $\dot{\gamma}_1$, and nothing else. With $\mathbf{J}_M$ the record's Jacobian, $\mathbf{J}_{\mathrm{f}z}$ the forecast rows' columns on the record's entries and $\mathbf{J}_{\mathrm{ff}}$ their columns on the forecast entries,

$$\mathbf{J}_{\text{ext}} = \begin{bmatrix} \mathbf{J}_M & \mathbf{0} \\ \mathbf{J}_{\mathrm{f}z} & \mathbf{J}_{\mathrm{ff}} \end{bmatrix},$$

$$\frac{\partial r_{q,1,36}}{\partial(g_{1,36},\ g_{1,37},\ \gamma_1,\ \dot{\gamma}_1)} = \Big(-\frac{1}{\sigma_q},\ \frac{1}{\sigma_q},\ -\frac{\Delta t}{\sigma_q},\ -\frac{x_{36}\Delta t}{\sigma_q}\Big) = (-200,\ 200,\ -16.6667,\ -25.6944),$$

$$\mathbf{g}_{t^*}^{\mathsf{T}} = (* \text{ on } g_{1,36},\ g_{2,36},\ \gamma_1,\ \gamma_2,\ \dot{\gamma}_1;\ \ 0 \text{ elsewhere in } \mathbf{z}_M), \tag{14.6}$$

the entries in the columns of γ_1 and $\dot{\gamma}_1$ in yr and yr^2 and the others dimensionless. The upper right block is zero because no record row sees a forecast entry, and $\mathbf{J}_{\mathrm{ff}}$ is square and lower bidiagonal, $+1/\sigma_q$ on its diagonal and $-1/\sigma_q$ below it, one row for each forecast entry. The second line is not a row. It gives the pattern of the crossing's gradient on the drawing's state, with each forecast window written through the record's entries and its own departures from the trend: by the forecast rows (14.4), with $e_{u,t} = \sigma_q\, r_{q,u,t}$,

$$g_{u,s} = g_{u,36} + (s-36)\,\Delta t\,\gamma_u + \Delta t\,\dot{\gamma}_u \sum_{t=36}^{s-1} x_t + \sum_{t=36}^{s-1} e_{u,t}, \qquad s > 36, \tag{14.7}$$

so that the crossing depends on the record through five entries only, $(g_{1,36}, g_{2,36}, \gamma_1, \gamma_2, \dot{\gamma}_1)$, and on the forecast windows' departures. Two facts of the graph are read here, before any number. *The forecast factors*: the forecast entries appear in forecast rows only, linearly, and their block of the Jacobian is square, so the forecast changes nothing in the record's fit, and it needs only the five entries' block of each drawing's covariance (Section 14.2.1 derives it). *The date is a ratio*: on a straight path the crossing is the margin today over the rate of fall, and the rate enters in the denominator, so the date has no mean, and a rate whose band reaches zero would give the date no upper end.

Modeling choices at the rows. Five choices make these rows. *The family is Part III's two open questions crossed, and no other*: this buys a forecast that carries what Part III could not decide, and it fails if the truth's path has a form none of the drawings has, coil 1's acceleration from a knee rather than over the whole record, a second step, or a curvature of coil 2's; the knee and coil 2's curvature are varied in step 5. Coil 2's curvature is left out because Chapter 11's coil-2 curvature was the echo of the step. *The two factors are independent, at even odds each*, equation (14.2): this buys priors that are Chapters 11 and 12's own, and the joint evidence then moves both marginals, which step 3 reports. *The forecast carries the wander*, the choice of step 0, whose effect step 5 measures. *The limit is the capacity study's stated input*: it is never fitted and carries no band of its own, and the plant is not re-chosen. *No maintenance before the crossing*: a cleaning would end the path the rows carry.

14.2 The posterior of the state (step 3)

Every answer of this chapter is read from the posteriors of the 18 drawings of step 2, each extended by its forecast windows, and from the drawings' evidences, which weigh them.

14.2.1 The rows and the posterior of a drawing

The derivation form is Chapter 5's (Section 5.2.1), on each drawing's stacked state, as Chapters 10 to 12 fitted it and not derived again. The residual vector of drawing M stacks its record rows; its Jacobian $\mathbf{J}_M$ has the pattern (10.6) with the curvature's and the step's columns; the MAP state $\hat{\mathbf{z}}_M$ solves the normal equations (5.9) by Gauss–Newton; and the Laplace covariance is $\mathbf{\Sigma}_M = (\mathbf{J}_M^{\mathsf{T}}\mathbf{J}_M)^{-1}$ at $\hat{\mathbf{z}}_M$. With every row normalized the information matrix is $\mathbf{J}^{\mathsf{T}}\mathbf{J}$; with unnormalized rows and their weights it is $\mathbf{J}^{\mathsf{T}}\mathbf{W}\mathbf{J}$. The evidence of each drawing is its Laplace evidence (7.7), and with the priors (14.2) the posterior probability of a drawing is equation (2.47),

$$P_M = \Pr(M \mid \mathbf{y}) = \frac{p(\mathbf{y} \mid M)\ \Pr(M)}{\sum_{M'} p(\mathbf{y} \mid M')\ \Pr(M')}, \tag{14.8}$$

summed over the 18 drawings.

The forecast windows change nothing in the fit. The extended state (14.3) adds the forecast rows to the objective, $\mathcal{L}_{\text{ext}} = \mathcal{L}_M(\mathbf{z}_M) + \frac{1}{2}\sum_{t,u} r_{q,u,t}^2$, and by the pattern (14.6) its information matrix is

$$\mathbf{\Lambda}_{\text{ext}} = \mathbf{J}_{\text{ext}}^{\mathsf{T}}\mathbf{J}_{\text{ext}} = \begin{bmatrix} \mathbf{J}_M^{\mathsf{T}}\mathbf{J}_M + \mathbf{J}_{\text{f}z}^{\mathsf{T}}\mathbf{J}_{\text{f}z} & \mathbf{J}_{\text{f}\bar{z}}^{\mathsf{T}}\mathbf{J}_{\text{ff}} \\ \mathbf{J}_{\text{ff}}^{\mathsf{T}}\mathbf{J}_{\text{f}z} & \mathbf{J}_{\text{ff}}^{\mathsf{T}}\mathbf{J}_{\text{ff}} \end{bmatrix}. \tag{14.9}$$

The covariance of the record's entries is the inverse of the Schur complement of the forecast block (Section 3.5.3). Because $\mathbf{J}_{\text{ff}}$ is square and invertible, $\mathbf{J}_{\text{ff}}(\mathbf{J}_{\text{ff}}^{\mathsf{T}}\mathbf{J}_{\text{ff}})^{-1}\mathbf{J}_{\text{ff}}^{\mathsf{T}}$ is the identity, and

$$\mathbf{J}_M^{\mathsf{T}}\mathbf{J}_M + \mathbf{J}_{\text{f}z}^{\mathsf{T}}\mathbf{J}_{\text{f}z} - \mathbf{J}_{\text{f}z}^{\mathsf{T}}\mathbf{J}_{\text{ff}}\left(\mathbf{J}_{\text{ff}}^{\mathsf{T}}\mathbf{J}_{\text{ff}}\right)^{-1}\mathbf{J}_{\text{ff}}^{\mathsf{T}}\mathbf{J}_{\text{f}z} = \mathbf{J}_M^{\mathsf{T}}\mathbf{J}_M :$$

the record's block of $\mathbf{\Sigma}_{\text{ext}}$ is $\mathbf{\Sigma}_M$. At the MAP state every forecast row is zero, with the forecast windows on the drawing's trend from $\hat{\mathbf{z}}_M$, so $\hat{\mathbf{z}}_M$ is unchanged as well. The forecast rows are linear and Gaussian in the forecast entries, and each forecast window's entries appear only in its own row and the next one, so the posterior of the extended state factors exactly into the record's posterior and one Gaussian step per window and coil,

$$p(\mathbf{z}_{\text{ext}} \mid \mathbf{y}, M) = p(\mathbf{z}_M \mid \mathbf{y}, M) \prod_{t=36}^{t_{\text{hor}}-1} \prod_{u=1}^{2} \mathcal{N}(g_{u,t+1};\ g_{u,t} + \Delta t\,[\gamma_u + \dot{\gamma}_u x_t],\ \sigma_q^2), \tag{14.10}$$

the band of unread windows of equation (4.69), carried past the last reading. Integrating the forecast entries out of the evidence integral gives $\sqrt{2\pi}\,\sigma_q$ per row, which the row's normalizing constant cancels, so every drawing's evidence, and with it its weight (14.8), is Part III's. Under the Laplace normal of $\mathbf{z}_M$ the forecast is then a Gaussian process on the five entries of equation (14.7), and their block of $\mathbf{\Sigma}_M$ is all the forecast needs.

14.2.2 The family's evidence

Table 14.2 gives the 18 drawings. Each row is a drawing; the columns give its numbers of entries and rows, its log evidence (7.7), its log Bayes factor against the base fit, its prior (14.2) and its posterior probability (14.8). The base fit's log evidence, -68.9335, and the factors of the steps under $\mathcal{C}_0$ are Chapter 12's, and the factor of $(\mathcal{C}_1, \text{none})$, $+0.0549$, is Chapter 11's. The leading drawing is $(\mathcal{C}_1, 2)$, coil 1 curving with coil 2 stepped, at 0.3428; then $(\mathcal{C}_0, 2)$ at 0.2061, $(\mathcal{C}_1, T_{d,2})$, coil 1 curving with unit 2's discharge sensor stepped, at 0.1606, and $(\mathcal{C}_1, \text{none})$ at 0.1097. The four leading drawings hold 0.8192 of the family, and the base fit 0.1038.

Table 14.2: The family's 18 drawings: coil 1's curvature state ($\mathcal{C}_0$ constant, $\mathcal{C}_1$ curving) crossed with the step into window 15 (none, or one of Chapter 12's eight kinds), with their Laplace evidences against the base fit $(\mathcal{C}_0, \text{none})$ and their probabilities. Every drawing has 146 degrees of freedom.

Drawing $M = (\mathcal{C}_l, j)$	Entries [–]	Rows [–]	$\ln p(\mathbf{y} \mid M)$ [–]	$\ln B$ against $(\mathcal{C}_0, \text{none})$ [–]	Prior [–]	P_M [–]
$(\mathcal{C}_0, \text{none})$, constant rates, no step	115	261	−68.9335	+0.0000	0.2500	0.1038
$(\mathcal{C}_0, 1)$, constant rates, a step on coil 1	116	262	−70.8015	−1.8680	0.0312	0.0020
$(\mathcal{C}_0, 2)$, constant rates, a step on coil 2	116	262	−66.1687	+2.7648	0.0312	0.2061
$(\mathcal{C}_0, 12)$, constant rates, a step on both coils	116	262	−69.8773	−0.9438	0.0312	0.0051
$(\mathcal{C}_0, T_r)$, constant rates, the return-air sensor stepped	116	262	−71.1863	−2.2528	0.0312	0.0014
$(\mathcal{C}_0, T_{d,1})$, constant rates, unit 1's discharge sensor stepped	116	262	−70.2767	−1.3432	0.0312	0.0034
$(\mathcal{C}_0, T_{d,2})$, constant rates, unit 2's discharge sensor stepped	116	262	−67.8453	+1.0881	0.0312	0.0385
$(\mathcal{C}_0, Q)$, constant rates, the load meter stepped	116	262	−70.5769	−1.6435	0.0312	0.0025
$(\mathcal{C}_0, T_1)$, constant rates, rack 1's inlet sensor stepped	116	262	−70.4612	−1.5277	0.0312	0.0028
$(\mathcal{C}_1, \text{none})$, coil 1 curving, no step	116	262	−68.8785	+0.0549	0.2500	0.1097
$(\mathcal{C}_1, 1)$, coil 1 curving, a step on coil 1	117	263	−70.7194	−1.7860	0.0312	0.0022
$(\mathcal{C}_1, 2)$, coil 1 curving, a step on coil 2	117	263	−65.6599	+3.2736	0.0312	0.3428
$(\mathcal{C}_1, 12)$, coil 1 curving, a step on both coils	117	263	−69.2412	−0.3077	0.0312	0.0095
$(\mathcal{C}_1, T_r)$, coil 1 curving, the return-air sensor stepped	117	263	−70.9338	−2.0004	0.0312	0.0018
$(\mathcal{C}_1, T_{d,1})$, coil 1 curving, unit 1's discharge sensor stepped	117	263	−70.5571	−1.6236	0.0312	0.0026
$(\mathcal{C}_1, T_{d,2})$, coil 1 curving, unit 2's discharge sensor stepped	117	263	−66.4180	+2.5155	0.0312	0.1606
$(\mathcal{C}_1, Q)$, coil 1 curving, the load meter stepped	117	263	−70.5632	−1.6298	0.0312	0.0025
$(\mathcal{C}_1, T_1)$, coil 1 curving, rack 1's inlet sensor stepped	117	263	−70.4888	−1.5553	0.0312	0.0027

The two open questions read jointly. Summed over the family, the probability that coil 1's curvature is in the drawing is $\Pr(\mathcal{C}_1 \mid \mathbf{y}) = 0.6344$, and the probability of a step into window 15 is 0.7864; coil 2 is the stepped part with 0.5489 and unit 2's discharge sensor with 0.1991. Split by the curvature's sign over the family, coil 1 is accelerating with 0.6055, constant with 0.3656 and slowing with 0.0289. These are this chapter's reading of the two questions taken together. Each was asked alone by its own chapter, and those answers stand as written: Chapter 11 gave 0.5392 for coil 1's curvature over its four drawings, with 0.5003 accelerating, 0.4608 constant and 0.0390 slowing (Section 11.3.2), and Chapter 12 gave 0.7160 for a step in March, 0.5637 for coil 2 and 0.1054 for unit 2's sensor (Section 12.3.1). Within $\mathcal{C}_0$ the step family is Chapter 12's, as it must be, 0.7160, 0.5637 and 0.1054; within $\mathcal{C}_1$ the step is more probable, 0.8271, with coil 2 at 0.5403 and unit 2's sensor at 0.2532. The reason runs through the evidences. With coil 2's step freed, coil 1's curvature gains: $\ln B(\mathcal{C}_1 : \mathcal{C}_0) = +0.5088$ against $+0.0549$ without a step, and $+1.4273$ with unit 2's sensor stepped. Unit 2's high discharge readings after March are explained either by coil 2 stepping or by its sensor stepping, and either way the record leaves coil 1 to carry more of year 3's loss through the return air, which a curvature does. The two questions support each other, so neither can be read for the forecast without the other.

14.2.3 The forecast entries of the leading drawings

Table 14.3 gives, for each drawing, the two coils' rates today, $\gamma_u(3) = \gamma_u + \dot{\gamma}_u\,(3 - \bar{t}_{\mathrm{y}})$ at $t_{\mathrm{y}} = 3\,\mathrm{yr}$ (window 36), and the total conductance at window 36, each at the MAP state with its standard deviation by its gradient. Each row is a drawing, in the order of Table 14.2; the last row is the family, the mixture (7.10) of the drawings' Laplace normals, with its median and 5–95 % band. Rates are reported as $100(\mathrm{e}^{\gamma_u(3)} - 1)\,\%/\mathrm{yr}$, equation (10.1).

Table 14.3: The rates today and the total conductance at window 36 under each drawing (MAP and standard deviation), and over the family (median and 5–95 % band).

Drawing	P_M [–]	$\gamma_1(3)$ [%/yr]	sd [yr^{-1}]	$\gamma_2(3)$ [%/yr]	sd [yr^{-1}]	$G_{\mathrm{tot},36}$ [kW/K]	sd [kW/K]
$(\mathcal{C}_0, \text{none})$	0.1038	−4.151	0.01463	−6.062	0.01323	11.7897	0.1193
$(\mathcal{C}_0, 1)$	0.0020	−4.323	0.01846	−6.047	0.01327	11.7845	0.1237
$(\mathcal{C}_0, 2)$	0.2061	−4.303	0.01465	−3.090	0.01673	11.8407	0.1205
$(\mathcal{C}_0, 12)$	0.0051	−3.104	0.01598	−5.190	0.01431	11.8392	0.1235
$(\mathcal{C}_0, T_r)$	0.0014	−4.383	0.01540	−5.973	0.01335	11.7603	0.1322
$(\mathcal{C}_0, T_{d,1})$	0.0034	−4.414	0.01485	−5.910	0.01328	11.7597	0.1217
$(\mathcal{C}_0, T_{d,2})$	0.0385	−6.138	0.01737	−3.920	0.01602	11.8266	0.1195
$(\mathcal{C}_0, Q)$	0.0025	−3.959	0.01632	−6.047	0.01324	11.8269	0.1797
$(\mathcal{C}_0, T_1)$	0.0028	−3.799	0.01503	−6.052	0.01324	11.8307	0.1261
$(\mathcal{C}_1, \text{none})$	0.1097	−8.417	0.03575	−5.867	0.01328	11.6868	0.1388
$(\mathcal{C}_1, 1)$	0.0022	−8.356	0.03598	−5.878	0.01330	11.6895	0.1399
$(\mathcal{C}_1, 2)$	0.3428	−9.453	0.03611	−2.695	0.01686	11.7193	0.1389
$(\mathcal{C}_1, 12)$	0.0095	−8.357	0.03576	−4.769	0.01448	11.7182	0.1399
$(\mathcal{C}_1, T_r)$	0.0018	−9.270	0.03773	−5.699	0.01344	11.6283	0.1557
$(\mathcal{C}_1, T_{d,1})$	0.0026	−7.785	0.03711	−5.805	0.01331	11.6871	0.1388
$(\mathcal{C}_1, T_{d,2})$	0.1606	−13.401	0.04256	−3.060	0.01632	11.6750	0.1372
$(\mathcal{C}_1, Q)$	0.0025	−8.316	0.03780	−5.864	0.01329	11.7000	0.2008
$(\mathcal{C}_1, T_1)$	0.0027	−7.926	0.03607	−5.867	0.01329	11.7288	0.1459
family: median		−7.336		−3.601		11.7489	
5–95 % band		−15.786 to −2.495		−7.049 to −0.447		11.5012 to 11.9827	

Over the family coil 1 is losing 7.34 %/yr today, against Chapter 11's 5.70 over its own four drawings (Section 11.3.3): the curvature's weight is higher in the joint family. Coil 2 is losing 3.60 %/yr, against Chapter 12's 3.09 given the step and 4.14 mixed (Section 12.3.5). Coil 1 is the faster coil today in every drawing with its curvature and in the two constant-rate drawings

that put the step on coil 2's side, $(\mathcal{C}_0, 2)$ and $(\mathcal{C}_0, T_{d,2})$; coil 2 is the faster in the other seven constant-rate drawings. The family puts the total today at 11.7489 kW/K, band 11.5012 to 11.9827: a margin of 0.9879 kW/K above the limit, 9.18 % of G^*. Across the drawings the level today runs only from 11.6283 to 11.8407 kW/K, while coil 1's rate today runs from −3.104 to −13.401 %/yr: the date will be the rates'.

The covariance of the forecast entries. The leading drawing $(\mathcal{C}_1, 2)$ has the MAP values $\hat{g}_{1,36} = -0.16571$, $\hat{g}_{2,36} = -0.19007$, $\hat{\gamma}_1 = -0.04825\,\mathrm{yr}^{-1}$, $\hat{\gamma}_2 = -0.02732\,\mathrm{yr}^{-1}$ and $\hat{\dot{\gamma}}_1 = -0.03403\,\mathrm{yr}^{-2}$, and the block of $\mathbf{\Sigma}_M$ on these five entries, in that order, the trend entries in yr^{-1} and the curvature in yr^{-2}, is

$$\mathbf{\Sigma}^{[5]}_{(\mathcal{C}_1,2)} = \begin{bmatrix} 8.8145 & -3.2768 & 2.7847 & -1.6225 & 3.4891 \\ -3.2768 & 3.2228 & -1.1289 & 1.7298 & -1.0536 \\ 2.7847 & -1.1289 & 2.2372 & -0.63295 & 0.54692 \\ -1.6225 & 1.7298 & -0.63295 & 2.8430 & -0.50099 \\ 3.4891 & -1.0536 & 0.54692 & -0.50099 & 4.0703 \end{bmatrix} \times 10^{-4}, \tag{14.11}$$

with the standard deviations 0.02969, 0.01795, 0.01496, 0.01686 and 0.02018. Table 14.4 sets the four leading drawings side by side: each row is an entry or a correlation, and the columns give each drawing's value. The correlations read by sign, and the rows fix each sign. The two log-derates today are correlated negatively, −0.6148 in the leading drawing: the return air reads their total, so a higher coil 1 leaves less of it to coil 2, and the total, with its standard deviation of 0.1389 kW/K, is known better than either coil. Coil 1's level today is correlated positively with its rate, +0.6271, and with its curvature, +0.5825: with the earlier windows held, a higher value today means a slower fall and less acceleration, so an optimistic level comes with an optimistic trend and the two errors add in the date. Coil 2's level and rate are correlated the same way, +0.5715. The two rates are correlated by −0.2510, Chapter 10's sign, small. Under $(\mathcal{C}_1, T_{d,2})$ the two levels are correlated by −0.8206: with unit 2's discharge readings discounted after March by their step, the coils are separated by the return air and unit 1's reading alone.

Table 14.4: The five forecast entries of the four leading drawings: MAP values, standard deviations $\sqrt{\Sigma_{ii}}$ and the correlations that set the date. $\dot{\gamma}_1$ is absent in $(\mathcal{C}_0, 2)$.

	$(\mathcal{C}_1, 2)$		$(\mathcal{C}_0, 2)$		$(\mathcal{C}_1, T_{d,2})$		$(\mathcal{C}_1, \text{none})$	
	MAP	sd	MAP	sd	MAP	sd	MAP	sd
Entry [unit]	[entry's]	[entry's]	[entry's]	[entry's]	[entry's]	[entry's]	[entry's]	[entry's]
$g_{1,36}$ [–]	−0.16571	0.02969	−0.13737	0.02395	−0.24044	0.04361	−0.15629	0.02933
$g_{2,36}$ [–]	−0.19007	0.01795	−0.19858	0.01736	−0.12604	0.03114	−0.20551	0.01743
γ_1 [yr^{-1}]	−0.04825	0.01496	−0.04398	0.01465	−0.07555	0.01863	−0.04584	0.01490
γ_2 [yr^{-1}]	−0.02732	0.01686	−0.03139	0.01673	−0.03108	0.01632	−0.06047	0.01328
$\dot{\gamma}_1$ [yr^{-2}]	−0.03403	0.02018	–	–	−0.04556	0.02164	−0.02806	0.02002
Correlation [–]								
$g_{1,36}$, $g_{2,36}$		−0.6148		−0.5754		−0.8206		−0.6141
$g_{1,36}$, γ_1		+0.6271		+0.6493		+0.7794		+0.6234
$g_{1,36}$, $\dot{\gamma}_1$		+0.5825		–		+0.6220		+0.5740
$g_{2,36}$, γ_2		+0.5715		+0.5627		+0.7272		+0.5399
γ_1, γ_2		−0.2510		−0.2325		−0.5226		−0.2807

14.2.4 The full posterior by reweighted draws

The Laplace normal of each drawing is the reported posterior. Beside it stands the full posterior of the same rows, by Chapter 5's reweighted draws (Section 5.2.4), run on each drawing's whole state: 20 000 draws per drawing from its Laplace normal, each weighted by equation (5.13) with

the drawing's objective, and the drawings weighted by their importance-sampling evidence, equation (7.12), in place of the Laplace evidence. Because the forecast factors, equation (14.10), each reweighted draw is carried past window 36 by the forecast rows exactly as a draw of the Laplace normal is, and keeps its weight. The effective sample sizes run from 95.2 to 98.9 % of the draws, and step 5 states the Monte Carlo errors.

14.3 The answers (step 4)

Every answer is read from the extended state's Laplace normal under each drawing, the factored posterior (14.10) with $\mathbf{z}_M$ at $\mathcal{N}(\hat{\mathbf{z}}_M, \boldsymbol{\Sigma}_M)$, averaged over the family by the drawings' probabilities (14.8), and that value is the one reported. The full posterior of Section 14.2.4 stands beside each, a median beside a median. A date is a continuous window t^*, and its calendar date is the window $\lceil t^* \rceil$ with its month and year; bands are 5–95 % throughout.

14.3.1 The crossing of the limit

Asset manager's question. When does the conductance cross the level at which the hall can no longer be held? The answer is a date with a band, and the remaining life it leaves.

General form. The crossing time's distribution function is the second line of the trend form, equation (2.21),

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

read by inversion at the levels of the credible interval (2.11), here 0.05, 0.5 and 0.95; over a family of drawings the posterior is the mixture (7.10).

Specialized form. The question is equation (2.21), second line, with $\mathbf{z}$ the extended state (14.3) of each drawing, the event the first passage (14.1) of $G_{\mathrm{tot},t}$ of equation (14.5), $p(\mathbf{z} \mid \mathbf{y})$ the factored posterior (14.10), and the mixture (7.10) over the drawings with the weights (14.8):

$$\begin{aligned} \Pr(t^* \le s \mid \mathbf{y}) &= \sum_M P_M \, \Pr(t^* \le s \mid \mathbf{y}, M), \\ \Pr(t^* \le s \mid \mathbf{y}, M) &= \int \mathbf{1}\Big[\min_{36 < t \le s} G_{\mathrm{tot},t} \le G^*\Big]\, p(\mathbf{z}_{\mathrm{ext}} \mid \mathbf{y}, M)\, \mathrm{d}\mathbf{z}_{\mathrm{ext}}. \end{aligned} \tag{14.12}$$

The median solves the first equation at 0.5 and the band's ends at 0.05 and 0.95. No closed form exists, because the minimum over windows is not a linear function of the state, and the integral is read by draws of the factored posterior, the Laplace route of Section 4.4.7.

1. For each drawing, draw $N_{\mathrm{d}} = 40\,000$ values of the five forecast entries $(g_{1,36}, g_{2,36}, \gamma_1, \gamma_2, \dot{\gamma}_1)$ from their block of $\mathcal{N}(\hat{\mathbf{z}}_M, \boldsymbol{\Sigma}_M)$, through its Cholesky factor ($\dot{\gamma}_1 = 0$ under $\mathcal{C}_0$).
2. Carry each draw window by window by the forecast rows (14.4), adding to each coil's increment a departure $e_{u,t} \sim \mathcal{N}(0,\ \sigma_q^2)$, which is the factor (14.10) drawn, and evaluate $G_{\mathrm{tot},t}$ at each window to $t_{\mathrm{hor}} = 276$.
3. Record the first window t with $G_{\mathrm{tot},t} \le G^*$ and the continuous crossing $t^* = t - 1 + (G_{\mathrm{tot},t-1} - G^*)/(G_{\mathrm{tot},t-1} - G_{\mathrm{tot},t})$; a draw that never crosses by t_{hor} is recorded at infinity.
4. Give each draw of drawing M the weight P_M/N_{d}, and read the distribution function, its quantiles and the mass at infinity from the weighted draws of all 18 drawings, 720 000 in all.
5. State the Monte Carlo standard error of each result by batch means over 20 batches, as in Section 5.2.4.

One choice enters here: the crossing is read as the first passage, not as the threshold at each date. This counts a path that dips below the limit and comes back as having crossed, which is what the hall's design day would see, and it differs from the threshold only by the paths that come back; the margin's question, Section 14.3.3, measures them.

Two other forms stand beside the reported one, each named for what it is. The *plug-in date* $t^*(\hat{\mathbf{z}}_M)$ is the root in continuous time of $G_{\text{tot}} = G^*$ on the path (14.7) at the MAP state with every departure zero, carried between windows as the continuous path $g_u(t_{\text{y}}) = g_{u,36} + \gamma_u(t_{\text{y}} - 3) + \frac{1}{2}\dot{\gamma}_u[(t_{\text{y}} - \bar{t}_{\text{y}})^2 - (3 - \bar{t}_{\text{y}})^2]$, and the *delta band* is

$$t^*(\hat{\mathbf{z}}_M) \pm 1.6449 \sqrt{\mathbf{g}_*^{\mathsf{T}} \boldsymbol{\Sigma}_M \, \mathbf{g}_*}, \qquad \mathbf{g}_* = \frac{\partial t^*}{\partial (g_{1,36},\ g_{2,36},\ \gamma_1,\ \gamma_2,\ \dot{\gamma}_1)}\bigg|_{\hat{\mathbf{z}}_M}, \tag{14.13}$$

with the gradient's pattern the second line of equation (14.6). The *threshold line* is Chapter 2's normal tail of equation (2.22) on the threshold at each date, with $G_{\text{tot},s}$ linearized at each drawing's MAP and the wander's variance added,

$$\Pr(G_{\text{tot},s} \le G^* \mid \mathbf{y}) \approx \sum_M P_M \, \Phi\left(\frac{G^* - G_{\text{tot},s}(\hat{\mathbf{z}}_M)}{\sqrt{\mathbf{g}_s^{\mathsf{T}} \boldsymbol{\Sigma}_M \, \mathbf{g}_s + (s - 36)\, \sigma_q^2 \sum_u G_{u,s}(\hat{\mathbf{z}}_M)^2}} \right), \tag{14.14}$$

with $\mathbf{g}_s$ the gradient of $G_{\text{tot},s}$ on the five entries through equation (14.7). It is the threshold, not the first passage.

Building blocks. Each drawing's MAP state and the block of its covariance on the five forecast entries (Section 14.2.3); the weights (14.8); the forecast rows (14.4) and the factor (14.10); the draws and their weighted quantiles; the gradient $\mathbf{g}_*$ of equation (14.13); the normal distribution function Φ.

The crossing. Table 14.5 gives the family's answer. Each row is a quantity of the crossing, with its value in windows and its calendar date or its years. The median crossing is window 52.55 (Monte Carlo standard error 0.02): the limit is first crossed in window 53, May of year 5. The band runs from window 43.74 to 81.67 (standard errors 0.01 and 0.15), August of year 4 to October of year 7. The remaining useful life from window 36 is 16.55 months at the median, 1.38 years, with a band of 7.74 to 45.67 months, 0.64 to 3.81 years. The probability of having crossed is 0.0604 by window 44, 0.1595 by window 46, 0.4775 by window 52, 0.7290 by window 60 and 0.9002 by window 72, and the mass that never crosses within twenty years is 0.0089. The full posterior of Section 14.2.4 gives a median of 52.63 (standard error 0.03) and a band of 43.74 to 82.03 (standard errors 0.02 and 0.22) beside these.

Each drawing's own date. Table 14.6 gives the answer under each drawing alone. Each row is a drawing; the columns give its probability, its plug-in date and the delta band (14.13), and the first passage's median, band and probability of a crossing by the shutdown. The medians of the four leading drawings run from window 46.87 to 66.49. The most probable, $(\mathcal{C}_1, 2)$, says window 50.21, March of year 5; the constant-rate drawing with the same step, $(\mathcal{C}_0, 2)$, says 66.49, July of year 6; the drawing that blames unit 2's sensor, $(\mathcal{C}_1, T_{d,2})$, says 46.87, November of year 4. No single drawing's date is the answer.

The band within and between the drawings. The date has no variance, so the split of equation (7.11) is made on the band's width instead. The mixture's band is 37.93 windows wide. Each drawing's own band, averaged over the weights, $\sum_M P_M(\text{hi}_M - \text{lo}_M)$, is 27.63 windows: the width if the drawing were known. The family's indecision adds 10.30 windows, 0.272 of the mixture's width, and the mixture's band is 1.373 times a drawing's own. The bands behind

Table 14.5: The crossing of the capacity limit over the family, by the first passage (14.12) on the Laplace normal of each drawing, 40 000 draws per drawing. Dates are the first window at or below the limit.

Quantity [unit]	Value	Date or years
Median crossing t^* [window]	52.55	window 53 (May Y5)
5–95 % band [window]	43.74 to 81.67	Aug Y4 to Oct Y7
2.5–97.5 % band [window]	42.91 to 94.91	Jul Y4 to Nov Y8
Remaining useful life, median [months]	16.55	1.38 yr
Remaining useful life, 5–95 % band [months]	7.74 to 45.67	0.64 to 3.81 yr
Never crossing by window 276 [–]	0.0089	
$\Pr(t^* \le s \mid \mathbf{y})$ at $s = 40, 42, 44$ [–]	0.0003, 0.0095, 0.0604	
at $s = 46, 47, 48, 50$ [–]	0.1595, 0.2183, 0.2771, 0.3856	
at $s = 52, 54, 56, 58, 60$ [–]	0.4775, 0.5551, 0.6218, 0.6798, 0.7290	
at $s = 64, 68, 72, 84$ [–]	0.8078, 0.8625, 0.9002, 0.9565	

Table 14.6: The crossing under each drawing: the plug-in date and the delta band (14.13) beside the first passage's median, band and probability of a crossing by window 46, from 40 000 draws per drawing.

Drawing	P_M [–]	Plug-in [window]	Delta band [window]	Median [window]	Band [window]	$\Pr(t^* \le 46)$ [–]
$(\mathcal{C}_0, \text{none})$	0.1038	57.08	48.86–65.29	57.18	49.20–71.78	0.0051
$(\mathcal{C}_0, 1)$	0.0020	56.63	47.45–65.80	56.72	48.53–73.29	0.0080
$(\mathcal{C}_0, 2)$	0.2061	66.33	50.44–82.23	66.49	53.29–100.40	0.0003
$(\mathcal{C}_0, 12)$	0.0051	63.40	48.96–77.83	63.60	51.70–93.13	0.0011
$(\mathcal{C}_0, T_r)$	0.0014	56.17	47.81–64.52	56.27	48.27–70.73	0.0099
$(\mathcal{C}_0, T_{d,1})$	0.0034	56.20	48.25–64.15	56.33	48.59–70.38	0.0072
$(\mathcal{C}_0, T_{d,2})$	0.0385	58.07	49.66–66.48	58.33	50.04–73.41	0.0025
$(\mathcal{C}_0, Q)$	0.0025	58.31	46.83–69.78	58.41	48.45–78.58	0.0136
$(\mathcal{C}_0, T_1)$	0.0028	58.78	49.33–68.23	58.92	49.96–75.82	0.0030
$(\mathcal{C}_1, \text{none})$	0.1097	48.14	41.18–55.10	48.22	42.94–64.26	0.2877
$(\mathcal{C}_1, 1)$	0.0022	48.18	41.20–55.16	48.30	42.99–64.50	0.2848
$(\mathcal{C}_1, 2)$	0.3428	50.10	41.70–58.49	50.21	43.91–70.34	0.1676
$(\mathcal{C}_1, 12)$	0.0095	49.10	41.62–56.59	49.16	43.50–65.72	0.2197
$(\mathcal{C}_1, T_r)$	0.0018	46.91	40.37–53.46	46.97	41.90–60.87	0.4036
$(\mathcal{C}_1, T_{d,1})$	0.0026	48.91	40.81–57.00	48.95	43.12–76.44	0.2500
$(\mathcal{C}_1, T_{d,2})$	0.1606	46.65	41.40–51.91	46.87	42.46–56.47	0.4002
$(\mathcal{C}_1, Q)$	0.0025	48.38	40.03–56.73	48.46	42.23–68.77	0.3018
$(\mathcal{C}_1, T_1)$	0.0027	49.05	41.28–56.82	49.12	43.35–68.98	0.2298
the family				52.55	43.74–81.67	0.1595

that split are set side by side in Figure 14.2. Each horizontal bar is one drawing's 5–95 % band with its median as a dot, ordered by weight from the top, the bar's thickness growing with P_M; the top bar, in black, is the family's. The dashed red line is the shutdown and the dotted line the stated truth's crossing, a check marker for Section 14.4.6. The bars split into two groups: the drawings with coil 1 curving have medians at or before 50.2, the constant-rate drawings at or after 56.2, and the family's band spans both.

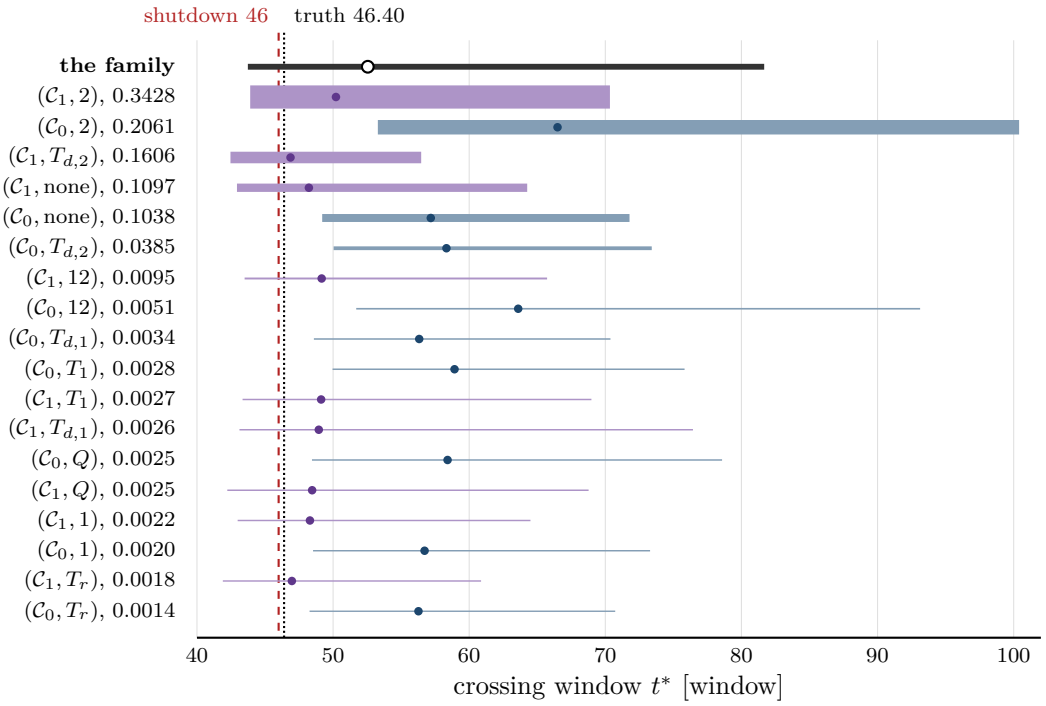


Figure 14.2: Each drawing's crossing band (5–95 %) and median (dot), ordered by its probability P_M, printed beside its name, the bar's thickness growing with P_M; purple: coil 1 curving ($\mathcal{C}_1$), blue: constant rates ($\mathcal{C}_0$). Top, black: the family's band and median (open dot). Dashed red: the shutdown, window 46. Dotted: the stated truth's crossing, window 46.40, a check marker (Section 14.4.6).

The delta band and the threshold line. The delta band is symmetric about the plug-in date, and the date is skewed toward late crossings, so it is wrong as Chapter 2's was. On the leading drawing, with the gradient

$$\mathbf{g}_* = (+71.01,\ +77.21,\ +83.42,\ +90.70,\ +174.13) \text{ windows per unit}$$

on $(g_{1,36}, g_{2,36}, \gamma_1, \gamma_2, \dot{\gamma}_1)$ and a spread of 5.10 windows, it runs from 41.70 to 58.49 about the plug-in 50.10, where the first passage runs from 43.91 to 70.34: it starts before the first passage's band and ends far short of it. Every entry of the gradient is positive, so a higher level today or a slower rate delays the crossing, and each cross term of the variance $\mathbf{g}_*^{\mathsf{T}} \boldsymbol{\Sigma}\, \mathbf{g}_*$ takes the sign of its correlation in Table 14.4. Of the 26.03 windows^2, the curvature's own term is 12.34 and the five diagonal terms sum to 22.60; the positive cross terms add +15.94, led by $g_{1,36}$ with $\dot{\gamma}_1$ (+8.63) and with γ_1 (+3.30), the level's and the trend's errors adding, and the negative ones −12.51, led by $g_{1,36}$ with $g_{2,36}$ (−3.59), $g_{2,36}$ with $\dot{\gamma}_1$ (−2.83) and $g_{1,36}$ with γ_2 (−2.09). The threshold line (14.14), mixed over the family, gives a median of 52.28 and a band of 43.54 to 77.94: close to the first passage at the median and short in the upper tail, where the curvature of G_{tot} in the log-derates matters. Neither is reported.

Figure 14.3 reads the crossing off. In the left panel the horizontal axis is the window and the vertical axis $G_{\mathrm{tot},t}$ in kW/K. The solid line and its shaded band are the family's fitted total over the record, to window 36, and its forecast fan beyond: the median and the 5–95 % band of $G_{\mathrm{tot},t}$ at the windows marked. The dashed red line is the limit, the dashed gray line the shutdown, and the dotted black line the stated truth, a check marker. In the right panel the horizontal axis is the date s and the vertical axis $\Pr(t^* \le s \mid \mathbf{y})$: the thick line is the family, the thin lines the four leading drawings; the dotted lines mark the median and the band's ends, and the dashed gray line the shutdown. The lower edge of the fan meets the limit between windows 43 and 46 and the median between windows 52 and 56, and the fan's upper edge stays above the limit through window 72. In the right panel the family's curve lies between the curving drawings, which rise steeply from window 43, and $(\mathcal{C}_0, 2)$, which barely leaves zero before window 52.

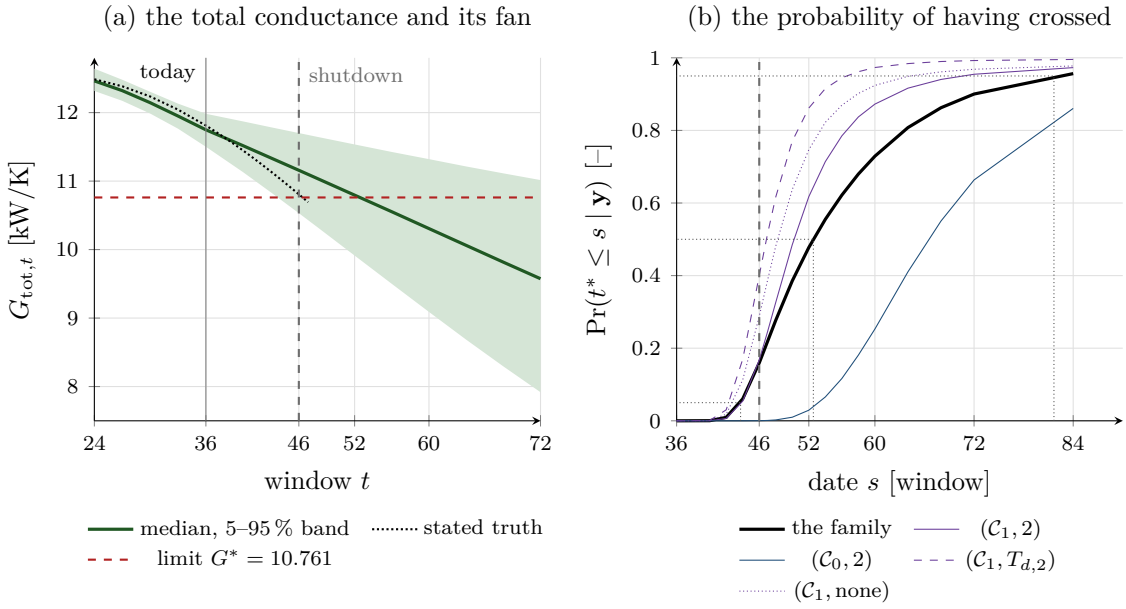


Figure 14.3: The forecast of the capacity limit over the family. (a) The total conductance $G_{\mathrm{tot},t}$, equation (14.5): over the record (to window 36) the mixture of the drawings' Laplace normals, and beyond it the first-passage draws of equation (14.12), median and 5–95 % band at windows 36, 40, 43, 46, 48 and every fourth window to 72; dashed red: the limit; dotted: the stated truth, a check marker, to its crossing at window 47. (b) The distribution function of the crossing time for the family (thick) and the four leading drawings (thin), from 40 000 draws per drawing, at the windows marked by the dotted guides: the band's ends 43.74 and 81.67 and the median 52.55; dashed gray: the shutdown.

14.3.2 Survival to the October shutdown

Asset manager's question. Will it hold until the planned shutdown in October? The answer is a probability at a stated date.

General form. The second line of the trend form, equation (2.21), at one date, the complement of the crossing by it:

$$\Pr(t^* > s \mid \mathbf{y}) = 1 - \int \mathbf{1}\left[t^*(\mathbf{z}) \le s\right] p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}.$$

Specialized form. The question is equation (2.21), second line, at $s = t_{\mathrm{sd}} = 46$, on the mixture (14.12):

$$\Pr(t^* > 46 \mid \mathbf{y}) = 1 - \sum_M P_M \Pr(t^* \le 46 \mid \mathbf{y}, M), \tag{14.15}$$

read from the same draws. One choice enters here: the shutdown's own window counts as one the hall must hold, so a crossing during October itself is a crossing before the work is done. This makes the answer the probability that the planned outage comes in time, and it would overstate the risk if the cleaning were done on the first day of the month; $\Pr(t^* \le 45 \mid \mathbf{y})$, a crossing before the shutdown's month, stands beside it.

Building blocks. The draws and weights of Section 14.3.1; the mixture (14.12) at one date.

The tail at the shutdown. By equation (14.15) the hall holds to the shutdown with probability 0.8405: the limit is crossed in or before the shutdown's month with probability 0.1595 (standard error 0.0008), and before the shutdown's month with 0.1055. The full posterior gives 0.1589 (standard error 0.0011) for the crossing by the shutdown. By December of year 4, window 48, the crossing probability is 0.2771, and by March of year 5, window 51, 0.4335. Under the leading drawings alone it is 0.1676 for $(\mathcal{C}_1, 2)$, 0.0003 for $(\mathcal{C}_0, 2)$, 0.4002 for $(\mathcal{C}_1, T_{d,2})$ and 0.2877 for $(\mathcal{C}_1, \text{none})$. The tail is the curvature's: the drawings with coil 1 curving carry 0.9950 of it, and the three leading ones 0.9608, $(\mathcal{C}_1, T_{d,2})$ with 0.4029 of it, $(\mathcal{C}_1, 2)$ with 0.3600 and $(\mathcal{C}_1, \text{none})$ with 0.1979. The leading drawing's plug-in date, window 50.10, says "after the shutdown" with no tail at all: a date at an estimate cannot answer a survival question. In the report's voice, there is about one chance in six that the hall cannot meet its design day before the shutdown's work is done.

14.3.3 The margin at the shutdown

Asset manager's question. How much margin will be left at the shutdown? The answer is the total conductance at window 46, against the limit.

General form. The third line of the trend form, equation (2.21), the forecast at a time beyond the record:

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

Specialized form. The question is equation (2.21), third line, with $f_t = G_{\mathrm{tot},46}$ of equation (14.5) on the extended state, read for every level f^* as a distribution with its median and band, and the margin $G_{\mathrm{tot},46} - G^*$. The same draws of Section 14.3.1 give $G_{\mathrm{tot},46}$ at window 46. The threshold and the crossing are related by equation (14.1): a path below the limit at 46 has crossed by 46, so

$$\Pr(G_{\mathrm{tot},46} \le G^* \mid \mathbf{y}) = \Pr(t^* \le 46 \mid \mathbf{y}) - \Pr(t^* \le 46,\ G_{\mathrm{tot},46} > G^* \mid \mathbf{y}), \tag{14.16}$$

the last term the paths that crossed and came back above the limit by window 46.

Building blocks. The draws of Section 14.3.1 at window 46; the relation (14.16).

The margin. At window 46 the total has median 11.1581 kW/K and band 10.5319 to 11.7015: a median margin of +0.3972 kW/K over the limit. At window 40 (Apr Y4) it is 11.5170 (11.1307 to 11.8714), a margin of +0.7561, and at window 48 (Dec Y4) 11.0373 (10.3269 to 11.6455), a margin of +0.2763. By equation (14.16), $\Pr(G_{\text{tot},46} \le G^* \mid \mathbf{y}) = 0.1595$, equal to the first passage to four decimals: the paths that crossed and came back hold 0.0001 of the posterior. Ten windows of wander at $\sigma_q = 0.005$ are small against the trend, so the first passage and the threshold differ here by less than the Monte Carlo error; Chapter 2's caution, that the threshold is not the crossing, concerns paths that never fell, and this record has almost none. The full posterior, from a second pass of the reweighted draws of Section 14.2.4 that stores each draw's total at window 46, gives a median of 11.1646 kW/K, a band of 10.5312 to 11.7032 and a margin of +0.4036; that pass's own crossing median and tail, 52.67 and 0.1611, lie within about two Monte Carlo standard errors of the first pass's 52.63 and 0.1589.

14.3.4 The open question that moves the date

Asset manager's question. Which open question moves the date, the acceleration or the step? The answer tells the planner which measurement the date waits for.

General form. An estimation question, equation (2.9), asked under a posterior that is a mixture over drawings, equation (7.10), with its split into the parts within and between the drawings, equation (7.11).

Specialized form. The question is equation (2.9) for $f = t^*$, under the mixture (7.10) restricted to the drawings of one factor, $\mathcal{M}$, the curvature state or the step's kind, and renormalized,

$$\Pr(t^* \le s \mid \mathbf{y}, \mathcal{M}) = \frac{\sum_{M \in \mathcal{M}} P_M \, \Pr(t^* \le s \mid \mathbf{y}, M)}{\sum_{M \in \mathcal{M}} P_M}, \tag{14.17}$$

read by the same draws, with the band's width within the drawings, $\sum_M P_M(\text{hi}_M - \text{lo}_M)$, as the analogue of the first term of equation (7.11) for a quantity with no variance.

Building blocks. The draws of Section 14.3.1, selected by drawing; the weights (14.8).

The forecast by factor. Table 14.7 gives the forecast within each factor. Each row is a set of drawings; the columns give its probability and its median, date, band and tail at the shutdown. The acceleration moves the date; the step barely does. Coil 1's curvature moves the median by −13.12 windows, from 61.93 (Feb Y6) without it to 48.82 (Jan Y5) with it, and the shutdown's tail by +0.2480. Coil 2's step moves the median by +2.18 windows and the tail by −0.0455, against no step. The step matters only through the side it lands on. A step on coil 2 slows coil 2's rate, because part of its loss was one event, while a step on unit 2's sensor leaves coil 2's rate steep and gives coil 1 the loss of year 3; the two kinds move the date in opposite directions, to 54.90 and 48.00, and largely cancel in the mixture. The full posterior, by the second pass of the reweighted draws with the drawings weighted by their importance-sampling evidence, gives the same reading: 62.08 under $\mathcal{C}_0$ and 48.87 under $\mathcal{C}_1$, the curvature moving the median by −13.21 windows and the tail by +0.2504, and coil 2's step moving them by +2.29 and −0.0470. The measurement that decides the acceleration is therefore worth more to the date than the one that decides the step, which step 7 takes up.

14.3.5 The hall against its limit today

Asset manager's question. Is the hall already past its limit, as the dashboard says? The answer is a probability today, at window 36.

Table 14.7: The forecast within each factor of the family, equation (14.17): the curvature states, and the three leading step states, with the full posterior's medians (the family's from the first pass of the reweighted draws, the factors' from the second).

Drawings	Pr [–]	Median [window]	Date [month]	Band [window]	$\Pr(t^* \le 46)$ [–]	Full posterior: median [window]
$\mathcal{C}_0$, constant rates	0.3656	61.93	Feb Y6	50.74–91.36	0.0022	62.08
$\mathcal{C}_1$, coil 1 curving	0.6344	48.82	Jan Y5	43.17–66.07	0.2502	48.87
no step at window 15	0.2136	52.72	May Y5	43.79–70.08	0.1503	52.78
coil 2 stepped at window 15	0.5489	54.90	Jul Y5	44.58–89.58	0.1048	55.07
unit 2's sensor stepped at window 15	0.1991	48.00	Dec Y4	42.68–64.87	0.3232	47.91
the family	1	52.55	May Y5	43.74–81.67	0.1595	52.63

General form. A prediction under a threshold, equation (2.12),

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

here asked of the complement, the requirement failed.

Specialized form. The question is equation (2.12) with $f = G_{\mathrm{tot},36}$ of equation (14.5), $f^* = G^*$ and the event reversed, on each drawing's Laplace normal, mixed. The quantity is an entry's function at a read window, carried by its gradient $\mathbf{g}_G$, with $G_{1,36}$ and $G_{2,36}$ in the columns of $g_{1,36}$ and $g_{2,36}$ and zeros elsewhere (Chapter 5's equation (5.21)):

$$\Pr(G_{\mathrm{tot},36} \le G^* \mid \mathbf{y}) = \sum_M P_M\, \Phi\left(\frac{G^* - G_{\mathrm{tot},36}(\hat{\mathbf{z}}_M)}{\sqrt{\mathbf{g}_G^{\mathsf{T}} \boldsymbol{\Sigma}_M\, \mathbf{g}_G}}\right). \tag{14.18}$$

Building blocks. Each drawing's total today and its standard deviation, Table 14.3; the weights (14.8); Φ.

The hall today. By equation (14.18) the probability that the hall is past its limit today is 3.7×10^{-9}. The leading drawing puts the total at 11.7193 kW/K with standard deviation 0.1389, 6.90 standard deviations above the limit. The full posterior stands beside this answer only as a bound: its 20 000 reweighted draws per drawing cannot resolve a probability of 4×10^{-9}, far below one in 20 000, and the Laplace tail is reported alone. At window 16 (Apr Y2), where the dashboard's line passed the limit, the same drawing puts the coils at 12.7008 kW/K (standard deviation 0.0936). The dashboard's 7.3097 kW/K today is not the coils' conductance: its tile is a ratio of approaches, and Chapter 10 split its slope into the load and the season, the drifting sensor and the coils magnified (Section 10.3.3). The readings taken as exact through the physics, Chapter 10's line on each window's $G_{u,\mathrm{test}}$, put the total today at 11.2553 kW/K and reach the limit at window 43.53 (Aug Y4): the physics alone moves the date from window 15.46 to 43.53 before the probability adds its band.

Figure 14.4 sets the dashboard against the coils. The horizontal axis is the window and the vertical axis the total conductance in kW/K. The gray line is the dashboard's line through its tiles, turned into kW/K, and the gray dots its tiles; the green line and its shaded band are the coils' total over the record, the mixture of the drawings' Laplace normals by their gradients, with its 5–95 % band; the dashed red line is the limit, and the dotted black line the stated truth, a check marker. The dashboard's line falls through the limit at window 15.46, where the coils'

total stands at 12.8332 kW/K at window 15, and today the line reads 7.3097 against the coils' 11.7489 kW/K. The tiles scatter about the line with the load and the season, and the truth lies inside the coils' band at every window drawn.

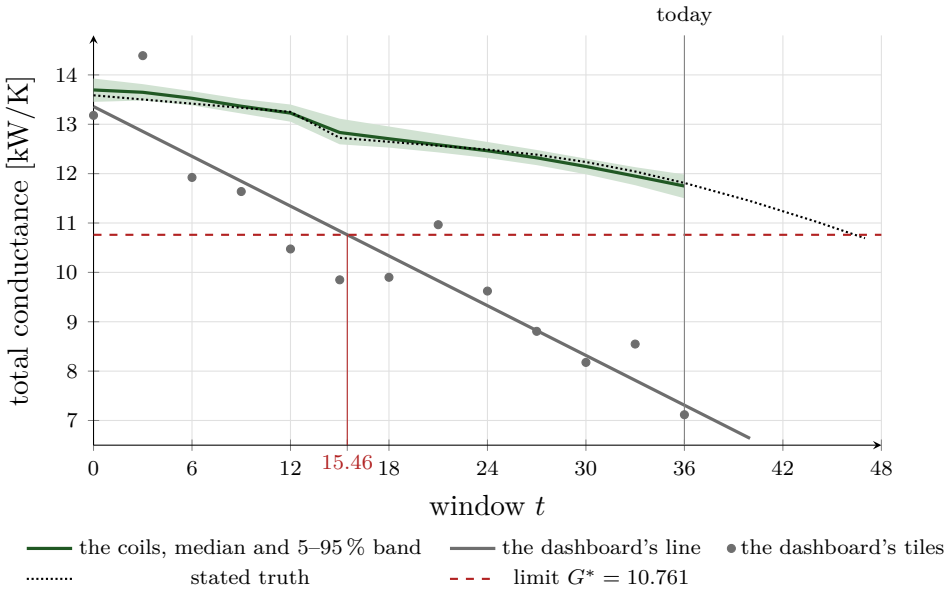


Figure 14.4: The dashboard against the coils over the record. Gray: the dashboard's line through its tiles in kW/K, $13.3588 - 2.0164\,t_{\mathrm{y}}$ with $t_{\mathrm{y}} = t/12$ years, crossing the limit at window 15.46 (red tick), and its tiles turned into kW/K with the commissioning test's conductances, every third window. Green: the coils' total conductance $G_{\mathrm{tot},t}$, the mixture of the drawings' Laplace normals with its 5–95 % band, every third window. Dashed red: the limit. Dotted: the stated truth, a check marker, to its crossing at window 47.

14.4 The checks (step 5)

Each answer of Section 14.3 rests on the Laplace normal of each drawing, on the draws that carry it forward, on the family and its priors, and on a record of three years; this hall lets each be tested.

14.4.1 The full posterior against the Laplace answers

The reweighted draws of Section 14.2.4, 20 000 per drawing over 20 batches, give the family's importance-sampling evidences within 0.0160 of the Laplace evidences, the largest gap that of $(\mathcal{C}_0, T_{d,2})$, and the drawings' weights move by 0.0007 at most, from 0.2061 to 0.2054 for $(\mathcal{C}_0, 2)$. On the full posterior the median crossing is window 52.63 (standard error 0.03), the band runs from 43.74 to 82.03 (standard errors 0.02 and 0.22), and the crossing by the shutdown has probability 0.1589 (standard error 0.0011), against the Laplace route's 52.55, 43.74 to 81.67 and 0.1595. The gaps are +0.07 windows on the median, −0.00 and +0.36 on the band's ends, and −0.0006 on the tail. The median's gap is larger than the two routes' standard errors, 0.03 and 0.02, and small against a band 37.93 windows wide; the upper end's gap is of the size of its standard errors, 0.22 and 0.15, and the tail's lies within them. Under the leading drawings the full posterior gives 50.32 for $(\mathcal{C}_1, 2)$ (band 43.94 to 70.37, tail 0.1629), 66.64 for $(\mathcal{C}_0, 2)$ (53.42 to 101.85, 0.0005), 46.74 for $(\mathcal{C}_1, T_{d,2})$ (42.40 to 55.91, 0.4071) and 48.23 for $(\mathcal{C}_1, \text{none})$ (42.95 to 64.06, 0.2860). The Laplace normal is the posterior of each drawing to the precision of the reported date. The full posterior is not an exact route: it is the same rows without the Gaussian

approximation, to its Monte Carlo error. A second pass of the same draws, which stores each draw's total at window 46 for questions 3 and 4, gives a median of 52.67, a band of 43.74 to 81.77 and a tail of 0.1611: two passes of the full posterior differ by about two of their standard errors, which is the scale on which its gaps from the Laplace route are read.

14.4.2 The routes, the wander and the exact rows

Table 14.11 collects the routes to the date. Each row is a route; the columns give its median, its band, its probability of a crossing by the shutdown and, for the reported route, the probability it gives of a crossing by the truth's window 47. The threshold line (14.14) gives 52.28, 43.54 to 77.94 and 0.1683: it errs in the upper tail. The forecast without the wander, $\sigma_q \to 0$ after window 36 only with the record's fit unchanged, gives 52.72, 44.06 to 79.53 and 0.1489; at the same seed and 10 000 draws per drawing the forecast with the wander gives 52.55, 43.79 to 82.09 and 0.1588. The wander moves the median from 52.72 to 52.55 and the tail from 0.1489 to 0.1588, which is the choice of step 0 measured. The plug-in date and the delta band (14.13) of the leading drawing, 50.10 and 41.70 to 58.49, err at both ends, as Section 14.3.1 showed.

No check of the limit of exact rows, and no check of the curvature term, is new here. The forecast rows (14.4) are linear trend rows and add no physical row, and the fits are Chapters 10 to 12's, whose limits of exact rows hold (Tables 10.6, 11.10 and 12.7).

14.4.3 Sensitivity to the family, its priors and the base fit's choices

Table 14.8 refits every drawing of the family under each variant and forecasts it with 10 000 draws per drawing; the reported family at the same size gives 52.56, 43.73 to 81.30 and 0.1596. Each row is a variant; the columns give the probability of coil 1's curvature, the probability of a step at window 15, and the date's median, band and tail at the shutdown.

Table 14.8: The forecast under other families, priors and base-fit choices. Each variant refits every drawing; 10 000 draws per drawing.

Variant	$\Pr(\mathcal{C}_1)$ [–]	Pr(step) [–]	Median [window]	Band [window]	$\Pr(t^* \le 46)$ [–]
the reported family (2×9)	0.6344	0.7864	52.56	43.73–81.30	0.1596
the core family: $\mathcal{C}_0$, $\mathcal{C}_1$ × (none, coil 2)	0.6194	0.9536	54.66	44.49–88.92	0.1079
coil 2's curvature also (Chapter 11's four × 9)	0.7467	0.7809	52.60	43.76–117.67	0.1553
the step on the coils' kinds only	0.6108	0.8764	54.48	44.45–88.06	0.1123
constant rates only (Chapter 12's family)	0	0.7160	61.88	50.68–91.45	0.0022
coil 1 curving only	1	0.8271	48.83	43.16–65.80	0.2509
the base fit alone (Chapter 10's drawing)	0	0	57.12	49.31–71.26	0.0044
coil 1's curvature from a knee at window 24	0.6978	0.7590	51.48	43.60–87.76	0.1849
curvature prior sd $0.02\,\mathrm{yr}^{-2}$	0.6121	0.7519	56.11	46.29–87.05	0.0424
curvature prior sd $0.10\,\mathrm{yr}^{-2}$	0.5385	0.7847	53.45	43.32–82.89	0.1760
prior $\Pr(\mathcal{C}_1) = 0.25$	0.3665	0.7567	57.08	44.63–86.86	0.0931
$\sigma_q = 0.0025$, record and forecast	0.6778	0.8573	52.72	44.48–77.44	0.1227
$\sigma_q = 0.01$, record and forecast	0.5475	0.5557	51.87	42.34–104.51	0.2206
drift onset window 16	0.8287	0.8698	48.08	42.53–69.64	0.3224
drift onset window 28	0.4320	0.7548	56.73	45.49–88.38	0.0645
no drift entry	0.9485	0.9957	45.85	42.00–54.68	0.5209
rating prior sd 0.05	0.6157	0.7711	52.90	43.80–81.34	0.1520

The median stays between windows 51.48 and 54.66 for every variant of the family's composition that keeps both curvature states: the core family, coil 2's curvature added, the coils' kinds only and the knee. The upper end moves more: coil 2's curvature, which can slow coil 2, takes it to 117.67. The tail at the shutdown moves most with the curvature's weight, from 0.0424 at a tight curvature prior and 0.0931 at prior odds of one to three to 0.2509 with

the curvature certain, and with the drift's onset, 0.3224 at window 16 against 0.0645 at window 28. The onset acts through the coils one at a time. An earlier drift takes less of unit 1's late rise as the sensor's, so coil 1 must lose more in year 3 and accelerate, and coil 2, read through the return air the two share, changes little; this is Chapter 10's ridge between the drift and coil 1, now on the date. Chapter 10's rates move little with the onset, coil 2's between -0.06253 and $-0.05985\,\mathrm{yr}^{-1}$ over onsets 16 to 28 (Table 10.9), but the record does not hold the onsets equal: the base fit's log evidence is -65.8724 at window 28 against -68.9335 at window 20 and -71.1058 at window 16, and the family's -65.4159 against -68.0549 and -69.2455. The record favors onset 28, the onset that lowers the tail to 0.0645; the onset stays at window 20, Chapter 10's choice, as every Part III chapter reads it. Without the drift entry, which the record rejects (Section 10.3.5), all of unit 1's late rise is coil 1's, and the hall crosses before the shutdown with probability 0.5209. The reported family and Chapter 12's constant-rate family give medians of 52.56 and 61.88 and tails of 0.1596 and 0.0022: the forecast cannot be read from either Part III chapter alone.

14.4.4 The forecast as of each window

Table 14.9 refits the family to the record cut at window t_{as}, using nothing later, and starts the forecast there. Each row is an as-of window; the columns give the probability of coil 1's curvature, the probability of coil 2's step at window 15, the date's median, band and tail at the shutdown, the mass that never crosses by window 276, and the plug-in date of the leading drawing. The median and band appear in the left panel of Figure 14.5, the horizontal axis the as-of window and the vertical axis the crossing window, with arrows where the band has no upper end within the horizon, and the tail at the shutdown and $\Pr(\mathcal{C}_1)$ in its right panel, the vertical axis a probability.

Table 14.9: The forecast as of each window t_{as}: the family refitted to the record cut at t_{as}, the forecast started at t_{as}; 10 000 draws per drawing.

t_{as} [window]	$\Pr(\mathcal{C}_1)$ [–]	Pr(coil 2 at 15) [–]	Median [window]	Band [window]	$\Pr(t^* \le 46)$ [–]	Never by 276 [–]	Plug-in, leading drawing [window]
24 (Dec Y2)	0.4349	0.3232	63.46	41.11–beyond 276	0.1306	0.1889	79.25, $(\mathcal{C}_0, 2)$
27 (Mar Y3)	0.4144	0.3511	66.98	44.18–beyond 276	0.0750	0.1645	83.29, $(\mathcal{C}_0, 2)$
30 (Jun Y3)	0.3834	0.4334	65.37	45.48–beyond 276	0.0579	0.1128	78.60, $(\mathcal{C}_0, 2)$
33 (Sep Y3)	0.3818	0.5530	65.82	46.90–beyond 276	0.0364	0.0672	77.86, $(\mathcal{C}_0, 2)$
36 (Dec Y3)	0.6344	0.5489	52.58	43.72–81.82	0.1614	0.0090	50.10, $(\mathcal{C}_1, 2)$

Until September of year 3 the forecast had no upper end, because more than 5 % of the posterior never crossed within twenty years, and its median sat between windows 63.46 and 66.98. The last quarter of year 3 moved the median from 65.82 to 52.58, as coil 1's curvature gained weight, from 0.3818 to 0.6344, and gave the band an upper end. The shutdown's tail fell through year 3, from 0.1306 to 0.0364, and rose to 0.1614 in the last quarter. Three windows of record moved the date by more than a year. Each as-of fit uses nothing later than its window, and none is a verdict time: it is the forecast a planner would have read then.

14.4.5 Fresh records at the stated truth

The record is one draw of the sensors' errors. The check redraws them 40 times at the record's stated truth (Section 14.6.6), refits on each record the core family, $\mathcal{C}_0$ and $\mathcal{C}_1$ crossed with no step and coil 2's step, at even priors, and forecasts with 10 000 draws per drawing. The core family is the reported family's four leading members' factors, and this record's core-family answer, median 54.66 and tail 0.1079, is the like-with-like reference.

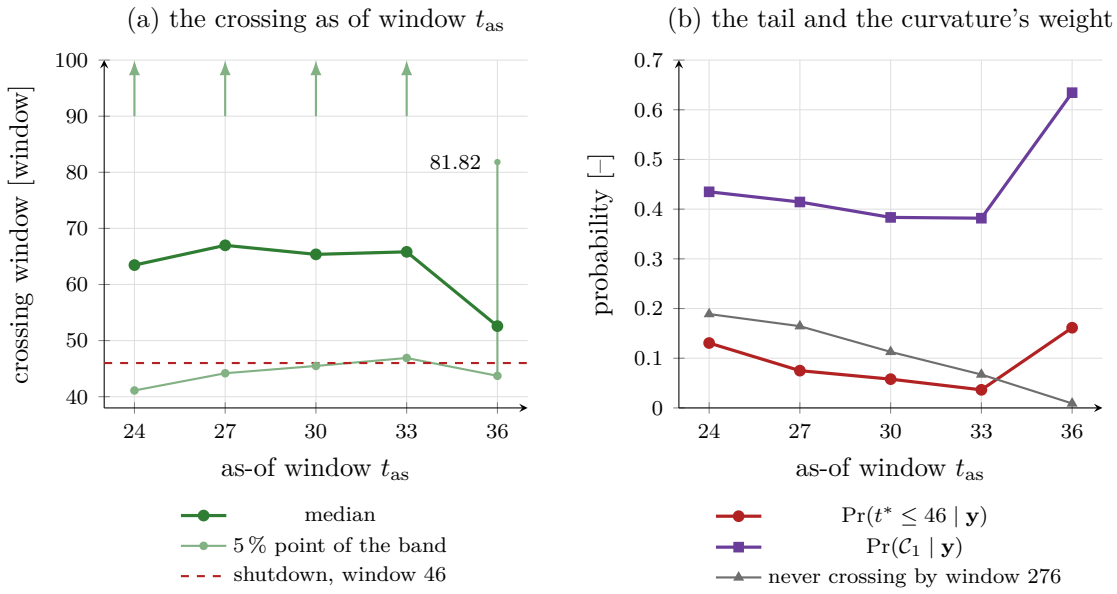


Figure 14.5: The forecast as of each window t_{as}, Table 14.9. (a) The median crossing and the 5 % point of the band, with the 95 % point beyond the horizon (arrows) until window 33 and at 81.82 at window 36; dashed: the shutdown. (b) The probability of a crossing by the shutdown, the probability of coil 1's curvature, and the mass that never crosses by window 276.

The finding. At this truth the forecast is late in two records out of three. Its 5–95 % band covers the truth's window 47 in 0.350 of the records; the truth lies below the band, earlier than its lower end, in the other 0.650 and above it in none. Over the records the median crossing has median 65.99 (5–95 % over the records 48.35 to 89.81) and the tail at the shutdown median 0.0132 (0.0002 to 0.3307). The forecast's probability of a crossing by the truth's window has median 0.0243 over the records, and it is 0.05 or less in 0.650 of them. This record is one of the earlier forecasts: its core-family median, 54.66, sits at the 0.150 quantile of the records' medians, and its tail, 0.1079, at the 0.875 quantile of their tails.

The triage. Table 14.10 tests each candidate cause on its own. Each row is a test; the columns give the probability of coil 1's curvature, the median, the band, the tail at the shutdown, the probability of a crossing by window 47 and, for the fresh records, the share of records whose band covers window 47. The first four rows fit the core family to the truth's noise-free readings, the pseudo-true answer of Chapter 10; the next three are medians over the 40 fresh records.

Here $F_{t^*}(s) = \Pr(t^* \le s \mid \mathbf{y})$ is the crossing's distribution function, and the triage reads the table row by row. *Not the sensor errors*: the noise-free readings forecast the same late date, median 66.03, with the truth's window at $F_{t^*}(47) = 0.0253$. *Not the drift's ridge*: with the drift settled by a side-by-side check reading its true value, the date stays late, 64.48 noise-free, and the fresh records' coverage falls to 0.300. *Not mainly the form*: with the curvature from the truth's own knee at window 24, the noise-free date is 64.14 and the coverage stays at 0.350. *Partly the curvature's prior*: the truth's curvature in year 3, $-0.12\,\mathrm{yr}^{-2}$, is more than twice Chapter 11's prior spread of 0.05; with a spread of 0.12 and the knee the noise-free date moves to 58.91 and the truth's window to $F_{t^*}(47) = 0.2036$. *Mostly the record's length*: three years hold at most one year of an acceleration whose effect on the path grows with the square of time, so that a record of this length reads it with little power (Table 11.14), and the as-of forecasts of Section 14.4.4 moved by more than a year in one quarter. The last row of the table tests this cause directly: the record continued at the truth to window 45, 20 times with fresh errors, and the core family refitted there with nothing later. Nine more windows decide the

Table 14.10: The triage of the late forecast at the stated truth: the core family on the truth's noise-free readings (pseudo-true), on 40 fresh records (medians over the records) under the reported form and its alternatives, and on the record continued at the truth to window 45 (medians over 20 continuations); $F_{t^*}(47) = \Pr(t^* \le 47 \mid \mathbf{y})$.

Test	$\Pr(\mathcal{C}_1)$ [–]	Median [window]	Band [window]	$\Pr(t^* \le 46)$ [–]	$F_{t^*}(47)$ [–]	Covers 47 [–]
pseudo-true, the reported form	0.4252	66.03	48.49–140.44	0.0136	0.0253	
pseudo-true, the curvature from a knee at window 24 (the truth's form)	0.5444	64.14	47.26–beyond 276	0.0245	0.0439	
pseudo-true, the knee with curvature prior sd $0.12\,\mathrm{yr}^{-2}$	0.5537	58.91	43.49–165.09	0.1580	0.2036	
pseudo-true, the side-by-side check reading the true drift, 0.85 K	0.4397	64.48	48.83–120.67	0.0090	0.0186	
fresh records, the reported form	0.3920	65.99		0.0132	0.0243	0.350
fresh records, the knee at window 24	0.5540	65.13		0.0211	0.0378	0.350
fresh records, the side-by-side check made	0.3882	65.78		0.0033	0.0075	0.300
the record continued at the truth to window 45 and refitted there (medians over 20 continuations)	0.9991	47.33		0.1818	0.4206	0.95

acceleration, $\Pr(\mathcal{C}_1)$ median 0.9991, and the band covers the truth's window 47 in 0.95 of the continuations, with the family and its priors unchanged. The finding is classified as a data limit of the three-year record, with a representation part in the curvature's prior. The prior is Chapter 11's and the family is kept as reported: a forecast of this chapter's own with a wider prior would read a different acceleration from the one Part III reports. The remedy is a measurement, the refit at window 45, and step 7 sets it among the candidates.

The side-by-side check on this record. This record with the side-by-side check of unit 1's sensor made at window 36, reading 0.8911 K, the truth's 0.85 K plus a 0.1 K error, and the reported family refitted: the drift reads higher, 0.6195 K/yr (standard deviation 0.0658), $\Pr(\mathcal{C}_1)$ falls to 0.5752 and a step at window 15 to 0.7416, and the median moves later, to 54.50 (Jul Y5), with band 45.25 to 81.97 and tail 0.0777. Settling the sensor gives the late rise back to the sensor and takes it from coil 1. It is the right check for Chapter 7's entry and the wrong one for this date.

14.4.6 The answers against the stated truth

The truth crosses in window 47 (Nov Y4), one month after the shutdown: its continuous crossing is 46.40, and at window 46 it stands at 10.8071 kW/K, 0.0461 above the limit. It sits at the 0.218 quantile of the forecast, between $F_{t^*}(46) = 0.1595$ and $F_{t^*}(47) = 0.2183$, inside the band, and the median is 5.55 windows late. The hall survives the shutdown, as 0.84 of the posterior says, by one month. Under the leading drawings the truth's window is the 0.2449 quantile of $(\mathcal{C}_1, 2)$, 0.5146 of $(\mathcal{C}_1, T_{d,2})$ and 0.3865 of $(\mathcal{C}_1, \text{none})$, and 0.0010 of $(\mathcal{C}_0, 2)$ and 0.0122 of the base fit: the drawings without coil 1's curvature put the truth in their far tail. The true total at window 36 is 11.8113 kW/K, against the family's median 11.7489 and the leading drawing's 11.7193, -0.66 standard deviations from it: the level is right, and the trend is what is late. The true coil 1 is losing 13.929 %/yr today, a rate of $-0.1500\,\mathrm{yr}^{-1}$, and is still accelerating at $-0.12\,\mathrm{yr}^{-2}$; the true coil 2 is losing 1.980 %/yr. The truth's own state at window 36, carried on with its curvature, crosses at 46.40; carried on at today's true rates without further acceleration, at 49.53. The dashboard's date, window 15.46, is 31.54 windows before the truth's crossing.

Table 14.11 gives the truth as its last row, and it is the dotted marker of Figures 14.3, 14.2 and 14.4.

Table 14.11: The routes to the crossing date and the stated truth. Bands are 5–95 %; $F_{t^*}(47)$ is the route's probability of a crossing by the truth's window.

Route	Median [window]	Band [window]	$\Pr(t^* \le 46)$ [–]	$F_{t^*}(47)$ [–]
first passage, Laplace draws (reported)	52.55	43.74–81.67	0.1595	0.2183
first passage, full posterior by reweighted draws	52.63	43.74–82.03	0.1589	–
threshold line (14.14), mixed	52.28	43.54–77.94	0.1683	–
forecast windows without wander	52.72	44.06–79.53	0.1489	–
plug-in and delta band, leading drawing $(\mathcal{C}_1, 2)$	50.10	41.70–58.49	–	–
plug-in and delta band, base fit $(\mathcal{C}_0, \text{none})$	57.08	48.86–65.29	–	–
stated truth	46.40 (window 47, Nov Y4)	–	–	–

14.5 The report (step 6)

The dashboard says the hall passed its capacity limit in April of year 2 and reads 7.3 kW/K today. The coils' own conductance today is 11.75 kW/K, 0.99 above the limit, and the hall is not past it, with probability 4×10^{-9} of the contrary. The limit is reached in May of year 5 at the median, window 53, with a band from August of year 4 to October of year 7: a remaining life of about 17 months, between 8 and 46. The hall holds until the October shutdown with probability 0.84; there is one chance in six that it crosses before or during the shutdown's month. Whether coil 1 is accelerating moves the date by thirteen windows, to February of year 6 if it is not and January of year 5 if it may, while whether coil 2 stepped moves it by two, and the band carries the record's indecision on the first. At the record's stated truth a forecast from three years errs late more often than early, because three years have seen at most one year of coil 1's acceleration; nine more windows decide it, and the refit at window 45 (Sep Y4), before the work order, is the measurement the date waits for.

14.5.1 The answers and their bands

Table 14.12 gives the answers in the order of Table 14.1. Each row is a question; the second column gives the reported value, the first passage on each drawing's Laplace normal mixed over the family, the third the full posterior's value beside it where the draws give one, and the fourth the band: the 5–95 % band of a date or a conductance, and for a probability its range over the priors and the drift's onsets of Table 14.8.

The answers rest on two things the dashboard cannot give. The physics puts the path on the coils' conductance, which is 0.99 kW/K above the limit where the tiles say 3.45 below, and the family carries Part III's two open questions into the date. Within any one drawing the date would be narrower, 27.63 windows of band against 37.93; the rest is what the record has not yet decided about coil 1.

14.5.2 The entry in the maintenance log

The forecast of window 36 is the next entry of the two-unit hall's maintenance log, after Chapters 10 to 13's entries at the same window. It uses nothing later than window 36, and in the values of Table 14.12 it reads:

Table 14.12: The report on the two-unit hall's capacity limit at window 36 (Dec Y3). Reported values from the first passage on the Laplace normal of each drawing, mixed over the family; bands are 5–95 %, or for a probability the range over the variants named.

Asset manager's question	Reported value	Full posterior	Band
When does the conductance cross the level at which the hall can no longer be held?	median $t^* = 52.55$ [window], first crossed in window 53 (May Y5); remaining life 16.55 [months]	median 52.63 [window]	43.74 to 81.67 [window], Aug Y4 to Oct Y7; remaining life 7.74 to 45.67 [months]
Will it hold until the planned shutdown in October?	survival 0.8405 [–]; crossing by window 46 0.1595, before October 0.1055 [–]	crossing by window 46 0.1589 [–]	crossing by window 46: 0.0424 to 0.3224 over the curvature priors and the drift's onsets [–]
How much margin will be left at the shutdown?	$G_{\mathrm{tot},46}$ median 11.1581 [kW/K], margin +0.3972 [kW/K]	median 11.1646, margin +0.4036 [kW/K]	10.5319 to 11.7015 [kW/K]
Which open question moves the date, the acceleration or the step?	the acceleration: −13.12 [window] on the median, +0.2480 [–] on the tail; the step: +2.18 [window], −0.0455 [–]; read jointly, curvature 0.6344 (Ch. 11: 0.5392), step 0.7864 (Ch. 12: 0.7160) [–]	curvature −13.21 [window], +0.2504 [–]; step +2.29 [window], −0.0470 [–]	medians: C_0 band 50.74 to 91.36, C_1 band 43.17 to 66.07 [window]
Is the hall already past its limit, as the dashboard says?	no: $\Pr = 3.7 \times 10^{-9}$ [–]; $G_{\mathrm{tot},36}$ median 11.7489 [kW/K]	not resolved by 20 000 draws per drawing	11.5012 to 11.9827 [kW/K]; 6.90 sd above G^* in $(\mathcal{C}_1, 2)$

Window 36 (Dec Y3), the forecast of the capacity limit ($G_1 + G_2 \geq 10.761$ kW/K, capacity study) from the condition record (windows 0 to 36), over the drawings Part III left open.
Today: 11.75 kW/K (11.50 to 11.98), 0.99 above the limit; not past it.
Crossing: May Y5 at the median (window 53), band Aug Y4 to Oct Y7; remaining life 17 months (8 to 46). Crossed by the October Y4 shutdown: 0.16; holds to the shutdown: 0.84.
The date depends on coil 1's acceleration (Feb Y6 if its rate is constant, Jan Y5 if it may change), little on coil 2's March step. Read jointly: coil 1's curvature 0.6344 (Ch. 11: 0.5392), a March step 0.7864 (Ch. 12: 0.7160).
Dashboard (line through the tiles, crossed Apr Y2): not a conductance; disregarded.
Action: plan the coil cleaning for the October Y4 shutdown; do not defer it to the following year's outage (crossed by Oct Y5, window 58, with probability 0.68).
Caution: three years of record hold at most one year of coil 1's possible acceleration; the last quarter alone moved the median from window 65.82 to 52.58.
Next measurement: continue the record and refit at window 45 (Sep Y4) before the work order is written; the nine windows take 65.5 % off the leading drawing's band and decide the acceleration in 0.545 of the predicted records (Chapter 11). The side-by-side check of unit 1's sensor stays open for Chapter 7's entry; it does not narrow this date.

The action line is the date's direct reading and the refit; whether the cleaning at the shutdown is worth its cost against the risk it removes is the expected consequence of Chapter 17, and the reading at which to raise the order is Chapter 19's.

Common mistake

Answering a survival question with a date. The plug-in date of the leading drawing, window 50.10, says "after the shutdown" with certainty; the first passage of the same drawing gives a crossing by the shutdown with probability 0.1676. A survival question asks for a tail of the date's distribution, which a single date does not have.

14.6 The next measurement (step 7)

The report leaves a band of 37.93 windows on the crossing and the acceleration that sets it undecided. A further measurement is one row or more of the residual vector, now or in windows not yet read, and what it takes off the band within each drawing is known before its value is.

14.6.1 The preposterior of the crossing band

The preposterior is equation (4.65) on each drawing's state extended to window 45 (Sep Y4), as in Chapter 10's step 7: the forecast windows carry their loads at the load's forecast growth and their path on the drawing's MAP trend, linked by the forecast rows (14.4), 144 entries for the leading drawing. A candidate a with Jacobian lines $\mathbf{J}_a$ in the columns of that state gives

$$\mathbf{\Sigma}_{M,a} = \left(\mathbf{\Sigma}_M^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a\right)^{-1}, \tag{14.19}$$

and the crossing band after it is read by the draws of Section 14.3.1 from the block of $\mathbf{\Sigma}_{M,a}$ on the forecast entries, at window 45 for a candidate that reads windows 37 to 45, with each drawing's center and weight held. The width it takes off is one less the ratio of that band's width to the present one. With no new rows the extension leaves the leading drawing's rates where they were, standard deviations 0.01496 and 0.01686, and the band where it was, 43.74 to 81.96 at 10 000 draws per drawing, 38.22 windows wide: the check row of the table. Holding the weights measures what a candidate does within the drawings; what decides between them is the evidence, Section 14.6.3.

The candidates are of three kinds and two limits. A side-by-side check of unit 1's discharge sensor now, ten paired readings against a reference of 0.1 K, reads the drift at window 36. A reference condition test now reads the return air, both discharges and rack 1's inlet at 0.1 K, and a water-side duty test reads each coil's duty at 3.0 kW. The record continued adds the five readings of each of windows 37 to 45, alone or with a permanent instrument from window 37: water-side heat meters on both coils, or a second discharge sensor on unit 1 (0.3 K), the permanent instrument at the place of the question. The two limits are not measurements: the path known exactly at window 36, the level's floor, and the rates known exactly, the trend's floor.

14.6.2 The candidates

Table 14.13 ranks the candidates. Each row is a candidate; the columns give the family's band after it, its width, the share of the width it takes off, and the same share for the leading drawing alone. The two shares are compared as bars in Figure 14.6, ranked by the family's, with the two floors as dashed lines.

Within a drawing the date is mostly the trend's: the rates known take 71.3 % off the leading drawing's band and the level known 46.2 %. Nine more windows of the record already kept take off 65.5 %, nearly as much as knowing the rates, and no instrument added now comes close: a condition test reads the level, not the trend, and takes 34.1 %. Across the family, with the weights held, the continuation still leaves a band of 27.94 windows while the leading drawing's falls to 9.11: what remains is mostly the family's indecision, and only evidence on the acceleration removes it. The permanent instruments add one to two points to the continuation, 28.3 and 29.1 against 26.9 %.

14.6.3 The evidence that decides the drawings

The weights move only with evidence on the drawings themselves, which Chapters 11 and 12 computed and this chapter cites. Chapter 11's record continued to window 45 decides coil 1's curvature, $|\ln B| > 1$, in 0.545 of the predicted records (Section 11.6.3), and Chapter 12's logger

Table 14.13: Candidate measurements and the width each takes off the crossing band, by the preposterior (14.19) with the drawings' centers and weights held; 10 000 draws per drawing.

Candidate	Band [window]	Width [window]	Off, family [%]	Width, $(\mathcal{C}_1, 2)$ [window]	Off, $(\mathcal{C}_1, 2)$ [%]
no new rows (the check)	43.74–81.96	38.22	0.0	26.43	0.0
side-by-side check of unit 1's discharge sensor now, ten pairs, 0.1 K	44.12–79.81	35.69	6.6	22.07	16.5
reference condition test now, 0.1 K	44.56–77.53	32.97	13.7	17.41	34.1
water-side duty test of each coil now, 3.0 kW	43.90–80.71	36.81	3.7	23.63	10.6
the record continued to window 45 (Sep Y4)	45.78–73.71	27.94	26.9	9.11	65.5
with water-side heat meters on both coils, permanent	45.89–73.29	27.40	28.3	8.08	69.4
with a second discharge sensor on unit 1, permanent	45.94–73.03	27.09	29.1	7.75	70.7
the path known exactly at window 36 (the level's floor)	45.12–75.31	30.19	21.0	14.21	46.2
the rates known exactly (the trend's floor)	45.42–71.53	26.10	31.7	7.59	71.3

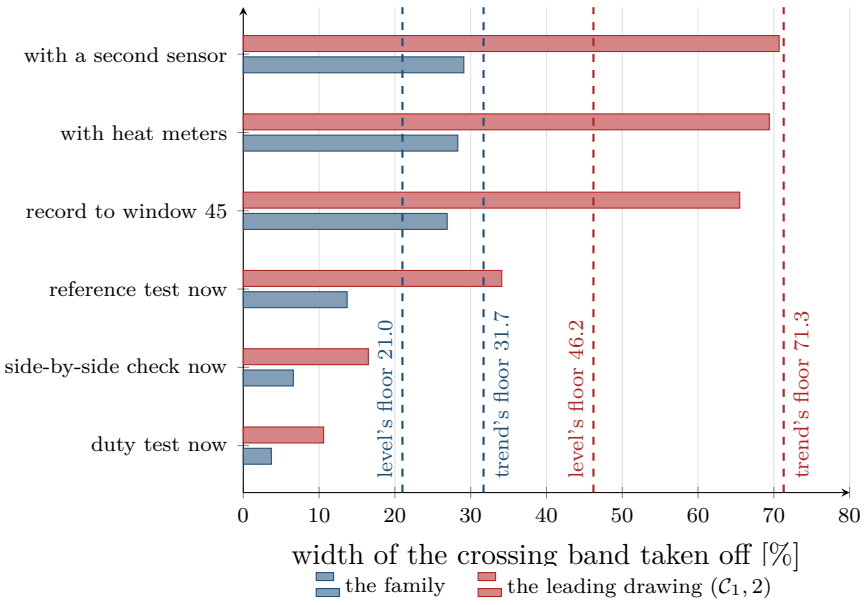


Figure 14.6: The candidates of Table 14.13, ranked by the width they take off the family's crossing band (blue), with the leading drawing's beside each (red), by the preposterior (14.19) with the drawings' centers and weights held. Dashed: the level's floor (the path known exactly at window 36) and the trend's floor (the rates known exactly), for the family (blue) and for the leading drawing (red).

archive of windows 12 to 18 raises the expected $\ln B$ of coil 2's step from +3.03 to +10.70 (Section 12.6.2). By Section 14.3.4 the first is worth more to the date.

The refit at window 45 at the stated truth. Step 5 found the forecast late in two records of three at the record's truth, and its last test, the last row of Table 14.10, is this measurement's check: the record continued at the truth to window 45, 20 times with fresh errors, the core family refitted there using nothing later, and the forecast made from window 45. In every continuation the record decides the acceleration: $\Pr(\mathcal{C}_1)$ has median 0.9991 (5–95 % over the continuations 0.9830 to 1.0000), and $\ln B(\mathcal{C}_1 : \mathcal{C}_0)$ with coil 2's step has median +7.0293, above +1 in all of them. The forecast's median crossing has median 47.33 (45.91 to 49.24), the tail at the shutdown median 0.1818 (0.0299 to 0.5276), and the probability of a crossing by window 47 median 0.4206. The band covers the truth's window in 0.95 of the continuations: the refit restores the forecast's calibration, with the family and its priors as reported. The tail exceeds one half in 0.15 of them, so in some records the refit would advance the work. The side-by-side check made now, by contrast, moved this record's date later, to 54.50 (Section 14.4.5). The continuation is the measurement that corrects the date; the check is not.

14.6.4 The recommendation

Keep the record running and refit at window 45 (Sep Y4), before the shutdown's work order is written: it takes 65.5 % off the leading drawing's band, decides the acceleration in 0.545 of the predicted records (Chapter 11), and costs nothing beyond the record already kept. Plan the cleaning for the October shutdown now, as the log's action says; the refit confirms the plan or advances it. If one window is to be spent now, a reference condition test takes 13.7 % off the family's band, against the duty test's 3.7 %; the side-by-side check, 6.6 %, stays recommended for Chapter 7's reason only. Of the permanent instruments, the second discharge sensor on unit 1 and the heat meters add one to two points to the continuation; their ranking for the date is Chapter 23's question.

14.6.5 What the analysis showed

The dashboard's line said the hall had been past its limit since April of year 2 and must be cleaned now; read on the coils' conductance, the hall stands 0.99 kW/K above the limit today and crosses it in May of year 5 at the median, between August of year 4 and October of year 7. The physics put the path on the coils' conductance, the trend rows carried it past the record with no new tool, the first passage turned it into a date, and the mixture carried Part III's two open questions into the date's band. The probability added what no chart gives: the survival to a stated outage, 0.84, and which open question the date depends on, coil 1's acceleration far more than coil 2's step. It also showed the limit of three years of record, which have seen too little of an acceleration to date it without erring late, and that the remedy is the record itself, continued to window 45. The analysis did not model maintenance before the crossing, a second event in the path, or a knee swept over windows.

14.6.6 Limits of the model

A constructed record. The condition record is generated from a stated truth, with the stated sensor errors, and step 5 compares the answers with that truth (Section 14.4.6). It is not field data, and the agreement is a check of the method on its own model. At that truth a forecast from three years is late in two records of three, coverage 0.350: the record has seen at most one year of an acceleration whose effect grows with the square of time, and Chapter 11's curvature prior is narrower than the truth's curvature. Nine more windows restore the coverage to 0.95. This is the record's limit, with its remedy, not a failure of the drawings.

The family. The family is Part III's two open questions and no other. A form outside it, coil 1's acceleration from a knee or a curvature of coil 2's, is a variant of step 5, not a member, and a second step is not drawn.

The limit as stated. The limit G^* is the capacity study's stated input, with no band of its own; the season enters the forecast only through its design supply.

No maintenance, the record's wander. The forecast assumes no cleaning or repair before the crossing and carries the record's wander, $\sigma_q = 0.005$ per window, unchanged; a new fouling regime would break both.

Chapters 10 to 12's limits. The readings are independent across windows, the drift is linear from window 20, and each window is in steady state with known inputs, as in Chapter 5.

The next chapter. Chapter 15 asks a different question of a date: the probability of a failure, not of a crossing, within a window of service. Whether to clean at the shutdown, how often to inspect and at what reading to raise the order are Chapters 17, 18 and 19's, and Chapter 23 ranks sensors by the width they take off this date.

14.7 What to retain

Definition 14.1 (Remaining useful life from the posterior of a path)**.** The crossing time of a parameter's path at a limit is its first passage, $t^* = \min\{t : G_{\mathrm{tot},t} \le G^*\}$, on the state extended by forecast windows: each forecast window carries its parameters and the record's trend rows and no reading, so the extended posterior factors into the record's posterior and one Gaussian step of the wander per window, and the forecast needs only the block of the covariance on the entries that carry the path forward. The date is $\lceil t^* \rceil$, the remaining useful life t^* less today's window. Being a ratio of a margin to a rate, t^* has no mean; it is reported by its median and the inversion band of its distribution function, with the mass that never crosses within the horizon. Survival to a planning date is the tail of that distribution at the date. When the record leaves the path's form open, the distribution is the mixture over the drawings, and its band's width splits into the part within the drawings and the part between them.

Three conclusions carry beyond the hall. A remaining life is read on the parameter of the physics, and a line through raw performance dates everything that performance depends on. A date is a distribution, so a survival question is answered by its tail, never by a plug-in date or a symmetric delta band. And a forecast from a short record is limited less by its sensors than by how much of the path's bend the record has seen; the most informative measurement is then often the record continued, with the forecast refitted before the decision.

Table 14.14 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 14.14: The terms of the remaining useful life and where each is defined. Units for the two-unit hall.

Term	Expression	Units	Meaning	Defined in
Forecast window	eqs. (14.3), (14.4)	–	trend rows, no readings	14.1.4
First passage	eq. (14.1)	window	the first window at or below the limit	14.1.1
Date and remaining life	$\lceil t^* \rceil$; $t^* - 36$	window, month	the calendar date; the months left	14.1.1
Factored forecast	eq. (14.10)	–	the record's posterior and the wander	14.2.1
Questions and answers				
Crossing over the family	eq. (14.12)	–	first passage, mixed	14.3.1
Inversion band, mass beyond the horizon	levels 0.05, 0.5, 0.95 of eq. (14.12)	window	the date's band; never crossing	14.3.1
Plug-in date, delta band	eq. (14.13)	window	symmetric, no tail; not reported	14.3.1
Threshold line	eq. (14.14)	–	the threshold, not the crossing	14.3.1
Survival to an outage	eq. (14.15)	–	the tail at a planning date	14.3.2
Margin at a date	eq. (14.16)	kW/K	the third line against the second	14.3.3
Within and between drawings	eq. (14.17), $\sum_M P_M(\mathrm{hi}_M - \mathrm{lo}_M)$	window	what the family's indecision adds	14.3.4
As-of forecast	Table 14.9	window	the forecast a planner had then	14.4.4
Coverage over fresh records	Table 14.10	–	calibration at a stated truth	14.4.5
Level's and trend's floors	Table 14.13	%	the band with the level or the rates known	14.6.2

14.8 Exercises

Exercise 14.1. *Application.* Turn the dashboard's two tile lines into kW/K with the commissioning test's conductances and find the date at which they reach G^*. (The line of coil u is $a_u + b_u t_\mathrm{y}$; the conductance is $\sum_u G_{u,\mathrm{test}}(a_u + b_u t_\mathrm{y})$, with $\sum_u G_{u,\mathrm{test}} a_u = 13.3588\,\mathrm{kW/K}$ and $\sum_u G_{u,\mathrm{test}} b_u = -2.0164\,\mathrm{kW/K}$ per yr; it reaches 10.761 at $t_\mathrm{y} = (13.3588 - 10.7610)/2.0164 = 1.2884\,\mathrm{yr}$, window 15.46.)

Exercise 14.2. *Derivation.* From the base fit's MAP at window 36, linearize G_tot in time and give the date at which the margin is spent; then give the plug-in date of the path, and say why the delta band is symmetric and the first passage is not. (With $\hat{g}_{1,36} = -0.13314$, $\hat{g}_{2,36} = -0.21208$, $\hat{\gamma}_1 = -0.04240$ and $\hat{\gamma}_2 = -0.06253\,\mathrm{yr}^{-1}$, $G_{\mathrm{tot},36} = 11.7897\,\mathrm{kW/K}$ and $\mathrm{d}G_\mathrm{tot}/\mathrm{d}t_\mathrm{y} = G_1\gamma_1 + G_2\gamma_2 = -0.6139\,\mathrm{kW/K}$ per yr; the margin $1.0288\,\mathrm{kW/K}$ lasts $1.6758\,\mathrm{yr}$, 20.11 windows. The path's plug-in, with its exponentials, crosses at 57.08. The delta band linearizes t^* and is symmetric about it, 48.86 to 65.29; the first passage keeps the rate in the denominator, which stretches the late side, 49.20 to 71.78.)

Exercise 14.3. *Formulation.* Write "will it hold until December of year 4?" and "will it hold until the shutdown a year later?" in the form of Table 14.1, and read them. (Both are equation (2.21), second line, at $s = 48$ and $s = 58$, by equation (14.15) with the date moved; the hall fails to hold to them with $\Pr(t^* \le 48 \mid \mathbf{y}) = 0.2771$ and $\Pr(t^* \le 58 \mid \mathbf{y}) = 0.6798$.)

Exercise 14.4. *Derivation.* Show that $\Pr(G_{\mathrm{tot},46} \le G^* \mid \mathbf{y}) \le \Pr(t^* \le 46 \mid \mathbf{y})$, name the paths in the difference, and read both. (By equation (14.1) the event $G_{\mathrm{tot},46} \le G^*$ is one term of the

union that makes $t^* \leq 46$, so its probability is at most the union's; the difference is the paths that crossed and stood above the limit again at window 46, equation (14.16). By the draws: 0.1595 and 0.1595, the paths back above 0.0001.)

Exercise 14.5. *Application.* Refit the family as of window 24 and read the forecast; say why its band has no upper end. (Median 63.46, $\Pr(t^* \leq 46 \mid \mathbf{y}) = 0.1306$; the mass that never crosses by window 276 is 0.1889, more than the 0.05 above the band's upper level, so the 0.95 point lies beyond the horizon. At window 36 the same forecast gives 52.58 and 0.1614.)

Chapter 15

Failure Probability over a Mission Window

15.1 The work-order screen and the coil's next ninety days

It is December of year 3, window 36, and the hall's coil 1 is three years old. January to March is the summer quarter, since the hall stands in the southern hemisphere, and the asset manager is deciding whether the leak-repair contingency, a contractor on call and a spare tube bundle on site, should cover coil 1 over the next 90 days. Two numbers are on the screen. The work-order system shows no corrective work order on coil 1 in its three years: its own MTBF is undefined, and its own rate, 0/3, gives a 90-day failure probability of 0. The fleet's reliability report, over the fifty coils of the design, gives the design an MTBF of 9.031 years, 289 coil-years over 32 tube failures, a rate of $0.1107\,\mathrm{yr}^{-1}$, and the report's formula $1 - \mathrm{e}^{-\tau/\mathrm{MTBF}}$ turns it into a 90-day failure probability of 2.69 %, and 7.97 % to the October shutdown. The conclusion the screen invites is that coil 1 has never leaked, so that it sits below its design's 2.7 %, near zero, and can come off the contingency.

The screen cannot say which coils' failures speak for a coil of three years: the whole fleet, aged one to twelve, or its siblings at its own age. Nor can it say what three clean coil-years tell about a rate near 0.1 a year, how well the rate is known, or whether 2.7 % is the probability of a failure in the next 90 days or the failure probability of a coil whose rate is known. And it cannot say whether coil 1's probability is below 5 % with assurance. The question of this chapter is

the probability that coil 1 fails in the next 90 days, given its rate today, read from its own record pooled with its siblings at its age, as the failure probability averaged over the posterior of that rate, with the band of that rate's failure probability and its assurance against 5 %*, set against the plug-in at an estimate.*

It is answered with the method of Chapter 4, on a state of one entry.

Jargon: mission window, mission failure probability, predictive and plug-in failure probability, age band, pooled exposure, concave function, Jensen's inequality, quantile map, assurance of a failure probability, predictive hazard, log-gamma posterior, zero-failure demonstration

A *mission window* is a stated stretch of service ahead, here the next 90 days, and the *mission failure probability* the probability of at least one failure within it. The *predictive failure probability* is that probability given the evidence, the failure probability at each rate averaged over the posterior of the rate; the *plug-in failure probability* is the failure probability at one estimate of the rate. An *age band* is a range of ages whose units are taken to share one rate, and the *pooled exposure* is the coil-years that all the fleet's coils have served at those ages. A function is *concave* when every chord lies below its graph, and *Jensen's inequality* states that the expectation of a concave function is at most the function of the expectation. The *quantile map* carries the quantiles of a quantity to those of a monotone function of it, in the same order for an increasing function and in reverse order for a decreasing one. The *assurance of a failure probability* is the posterior probability that it lies below a stated limit. The *predictive hazard* is the rate at which a unit that has run so far without failure fails next, under the predictive. A *log-gamma posterior* is the posterior of the logarithm of a gamma-distributed rate. A *zero-failure demonstration* is the exposure without a failure that would show a failure probability below a limit with stated assurance.

Key idea

The probability of a failure in a window is the failure probability at each rate averaged over the posterior of the rate, $\bar{p}_\tau = \mathbb{E}[1 - \mathrm{e}^{-\lambda\tau} \mid \mathbf{y}]$. Over a short window it is nearly linear in the rate, so it equals the plug-in at the posterior *mean* rate to second order, and it errs at first order at any other estimate. The band of the rate does not move this number; it moves how well the coil's own failure probability $p_\tau(\lambda)$ is known, which is what an assurance asks.

15.1.1 A failure probability averaged over the rate

For a unit with a constant rate λ of corrective failures, the probability of at least one failure in a window of length τ is

$$p_\tau(\lambda) = \Pr(T_f \le \tau \mid \lambda) = 1 - \mathrm{e}^{-\lambda\tau}, \tag{15.1}$$

with T_f the time to the next failure. The rate is known only through a record, and Chapter 3 averaged the conditional survival over its posterior, equation (3.16). With the event reversed, the probability of a failure in the window is

$$\bar{p}_\tau = \Pr(T_f \le \tau \mid \mathbf{y}) = \int_0^\infty (1 - \mathrm{e}^{-\lambda\tau})\, p(\lambda \mid \mathbf{y})\, \mathrm{d}\lambda = 1 - R_{\mathrm{pred}}(\tau \mid \mathbf{y}), \tag{15.2}$$

and for a gamma posterior the gamma integral of equation (3.17) gives it in closed form: one less the ratio of the posterior's rate parameter to that parameter plus the window, raised to the posterior's shape.

Two objects stand here, as in *Reliability Engineering by Derivation*, Chapter 9. The predictive $\bar{p}_\tau$ is one number, the probability that the coil fails in the next 90 days with everything known taken into account; it is the number to multiply by the consequence of a failure. The failure probability $p_\tau(\lambda)$ of equation (15.1) is a derived quantity of the state, with one true value for the coil and a posterior of its own; its band says how well the coil's failure probability is known, and its tail answers a limit on it. By equation (15.2) the predictive is the mean of that posterior, and the two meet there and nowhere else.

Principle 15.1. *A mission failure probability is a posterior expectation, not a failure probability at an estimate. Since $p_\tau(\lambda)$ is concave in λ, the expectation lies below the plug-in at the posterior mean rate, by at most $\frac{1}{2}\tau^2 \operatorname{Var}(\lambda \mid \mathbf{y})$ to first order.*

15.1.2 Coil 0 and its age band as a graph (step 0)

The hall's coil 1 is coil 0 of the fleet of Chapter 9, fifty coils of the same design at the operator's sites, one to twelve years old. Each coil's tube failures are logged as corrective work orders, each failure repaired and the coil returned to service, and the fleet's failure log counts the failures and the coil-years of exposure in one-year age bins, $[a, a+1)$, for every year each coil has served. The annual condition tests of Chapter 9 are not read here: this chapter reads the failure log by age, and nothing else.

The record. Table 15.1 gives the log. Each column is an age bin; the first two rows are the fleet's failures and coil-years in it, and the last two the part of each that is coil 0's. The two bins of coil 0's age band, ages 2 to 4, are in bold. The fleet holds 32 failures over 289 coil-years. Coil 0 has served three coil-years, one in each of the bins 0–1, 1–2 and 2–3, and has no failure in any of them. The same log is drawn in Figure 15.1: the upper panel gives the coil-years of each bin and the lower one its failures, the band shaded and the bins coil 0 has served marked by dots. The exposure falls with age, because fewer coils are old; the failures scatter between 0 and 7 a bin.

Table 15.1: The fleet's failure log by one-year age bin, at window 36 (Dec Y3). The band of the chapter, ages 2 to 4, is in bold; the last two rows are coil 0's part of each bin.

Age bin [yr]	0–1	1–2	**2–3**	**3–4**	4–5	5–6	6–7	7–8	8–9	9–10	10–11	11–12	Total
Failures [–]	2	4	**3**	**5**	3	2	7	3	1	1	1	0	32
Exposure [coil-yr]	50	47	**38**	**31**	26	22	22	20	14	9	7	3	289
Coil 0's failures [–]	0	0	**0**	–	–	–	–	–	–	–	–	–	0
Coil 0's exposure [coil-yr]	1	1	**1**	–	–	–	–	–	–	–	–	–	3

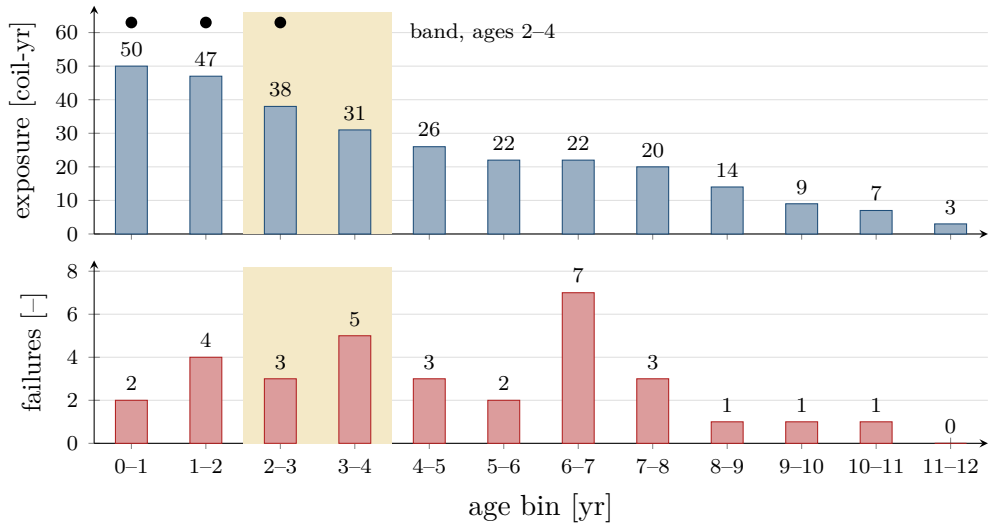


Figure 15.1: The fleet's failure log of Table 15.1: the coil-years of each one-year age bin (upper panel) and its tube failures (lower panel). Shaded: coil 0's age band, ages 2 to 4. Dots: the three bins in which coil 0 has served, each one coil-year without a failure. No line is fitted.

The mission and the planning date. Today is the end of window 36 (Dec Y3). The mission is the next 90 days, windows 37 to 39 (Jan Y4 to Mar Y4),

$$\tau = 90/365\,\text{yr} = 0.24658\,\text{yr},$$

Chapter 2's summer season of 0.2466 yr, over which coil 0 ages from 3.000 to 3.247 years. The planning date is the October shutdown, window 46 (Oct Y4), reached after the nine windows 37 to 45, $\tau_{\text{sd}} = 0.75\,\text{yr}$.

The state. The coil carries one free entry, its rate of tube failures, through its log-rate, as every rate does in this book (Section 2.4.3):

$$\omega = \ln\frac{\lambda}{\lambda_{\text{ref}}}, \qquad \lambda = \lambda_{\text{ref}}\,e^{\omega}, \qquad \lambda_{\text{ref}} = 0.1\,\text{yr}^{-1}. \tag{15.3}$$

The state is the one number ω, and the rate, the MTBF $1/\lambda$, the failure probability $p_\tau(\lambda)$ and every quantity the questions ask for are functions of it.

The graph. The rows hung on the state are drawn in Figure 15.2. Coil 0's node (green) carries ω. Its own year in the band, the year it served at ages 2 to 3, hangs on it as a dashed reading: no failure in one coil-year. The siblings' two age bins of the band and the prior row hang on it as black squares, each labeled with its count and exposure. The dotted arrow on the right is the question's extension of the graph, the event $T_f \le \tau$ over windows 37 to 39, which adds no row.

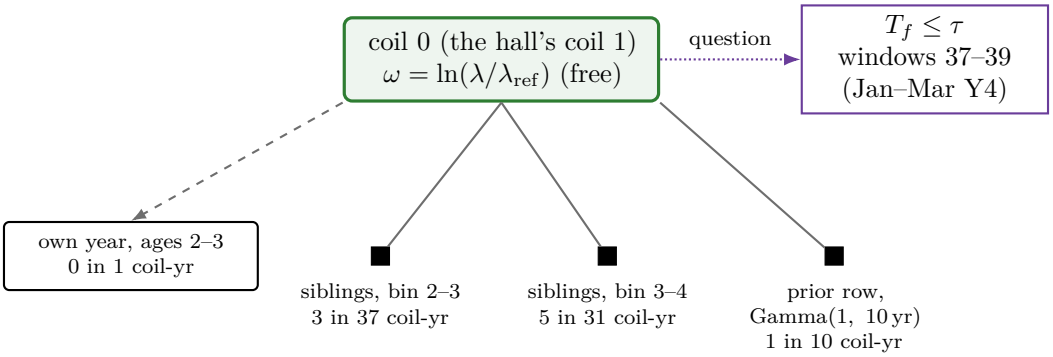


Figure 15.2: The graph of coil 0 in its age band. Free (green): the coil's log-rate ω. Observed (dashed): its own year in the band, no failure in one coil-year. Count rows (black squares): the siblings' two age bins of the band and the gamma prior, each with its count and exposure. Dotted: the question's extension, a failure within windows 37 to 39. The siblings share coil 0's entry by the chapter's choice of a constant rate in the band: within the band the drawing is the fully pooled limit of Chapter 9, a population spread of zero, applied to the coils' failure rates.

Modeling choices at the graph. Four choices that no question owns make this graph. The failures are taken as a Poisson process in the coil's service time, each repaired and returned to service, one fault one count, Chapter 2's choice: this is what lets the log be summarized by counts and exposures, and it fails if a repair leaves the coil weaker or one leak raises several work orders. The rate is taken as constant within coil 0's age band, ages 2 to 4, with today's age at its center, and every coil in the band at the same rate. This is the chapter's central choice, and step 5 tests it (Section 15.4.4). It buys 69 coil-years where coil 0 alone has 3. It fails in two ways: if the rate changes with age within the band, the band's rate is an exposure-weighted average read at the band's mean age, 2.95 years, not at today's; and if coils differ in rate at the

same age, by site, duty or water quality, the band's rate is the fleet's mean and coil 0's own departure is unseen. Whether the rate changes with age is the question Chapter 16 tests; this chapter neither asserts nor uses it. Coil 0's years outside the band, ages 0 to 2, with no failure in two coil-years, are not rows of this state, because they are at other ages. And the rate is taken as independent of the season and of the coil's condition: no row ties the summer duty or the conductance g to the failures.

15.1.3 The asset manager's questions (step 1)

The asset manager asks five questions of the band's record, and step 1 writes each as a probability over the state (15.3), with the unknown on the left of the bar and the counts and the prior on the right. Table 15.2 lists them. The first column is the question in the asset manager's words, the first of them the drill question; the second the general form of its category from Chapter 2; the third its mathematical form for coil 0; the fourth its category; the fifth the computation of Section 2.3.6 that answers it; and the last the subsection where it is answered. "It" is coil 0, the hall's coil 1. The atlas of Appendix A lists the same rows. Which coil the crew visits first is Chapter 20's question, and the cost of a failure and whether to act on it Chapter 17's. Which record would tighten the rate asks nothing new of today's evidence; it is the next measurement, Section 15.6.

Table 15.2: The questions asked of coil 0 at window 36 (Dec Y3). $\mathbf{y}$ is the band's three count rows and the prior row; $\tau = 90$ days, the windows 37 to 39; $p_\tau(\lambda)$ the failure probability at the rate λ, equation (15.1); $\bar{p}_\tau$ its predictive, equation (15.2); $p^* = 0.05$ the limit and $\kappa = 0.95$ the assurance.

Asset manager's question	General form	Mathematical form	Category	Computation	Answered in
Given its rate today, what is the probability it fails in the next 90 days?	the predictive, (2.13), in its event form (2.8)	$\bar{p}_\tau = \Pr(T_f \le \tau \mid \mathbf{y}) = \mathbb{E}[1 - \mathrm{e}^{-\lambda\tau} \mid \mathbf{y}]$, beside the plug-in $1 - \mathrm{e}^{-\hat{\lambda}\tau}$	prediction under a threshold	posterior	15.3.1
What is its failure rate today, and its MTBF?	(2.9), (2.11)	$p(\lambda \mid \mathbf{y})$, $\lambda = \lambda_{\mathrm{ref}}\mathrm{e}^{\omega}$, and $1/\lambda$: medians and 5–95 % bands	estimation	posterior	15.3.2
How well is that probability known, and is it below 5 % with 95 % assurance?	(2.12)	$p(p_\tau(\lambda) \mid \mathbf{y})$, median and band; $\Pr(p_\tau(\lambda) < p^* \mid \mathbf{y}) \ge \kappa$	prediction under a threshold	posterior	15.3.3
Will it run to the October shutdown, and in which window does its failure probability pass 5 %?	the predictive, (2.13); its inverse, (2.9)	$1 - \bar{p}_{\tau_{\mathrm{sd}}}$, $\tau_{\mathrm{sd}} = 0.75\,\mathrm{yr}$; $\tau_p(\lambda) = -\ln(1-p^*)/\lambda$, median and inversion band, as a window	prediction under a threshold	posterior	15.3.4
It has never failed: is its record evidence that it is better than its siblings?	(2.14)	$\ln B$ of its own rate against the shared rate, on its year in the band; Pr(three clean years | the siblings)	hypothesis comparison	evidence	15.3.5

15.1.4 The evidence as rows (step 2)

Step 2 attaches the band's record to the graph as four rows on ω. Each is a count row of Chapter 4, equation (4.42): a count n_ℓ in an exposure E_ℓ, Poisson with the expected count

$$\mu_\ell(\omega) = E_\ell\,\lambda = E_\ell\,\lambda_{\text{ref}}\,\mathrm{e}^{\omega}, \qquad \mathcal{L}_\ell(\omega) = \mu_\ell(\omega) - n_\ell \ln \mu_\ell(\omega), \qquad \ell = 1, \ldots, 4, \tag{15.4}$$

with the term $\mathcal{L}_\ell$ its negative log-probability up to a constant and

$$\begin{array}{lll} \ell = 1: & \text{coil 0, ages 2–3,} & (n_1, E_1) = (0,\ 1 \text{ coil-yr}), \\ \ell = 2: & \text{siblings, bin 2–3,} & (n_2, E_2) = (3,\ 37 \text{ coil-yr}), \\ \ell = 3: & \text{siblings, bin 3–4,} & (n_3, E_3) = (5,\ 31 \text{ coil-yr}), \\ \ell = 4: & \text{prior row,} & (n_4, E_4) = (a_\lambda,\ b_\lambda) = (1,\ 10\,\text{yr}). \end{array}$$

The band holds $k = n_1 + n_2 + n_3 = 8$ failures in $T_{\text{tot}} = E_1 + E_2 + E_3 = 69$ coil-years, coil 0's own 0 in 1. The siblings' bin 2–3 is the fleet's bin less coil 0's year, $38 - 1 = 37$ coil-years.

The prior as a count row. The rate's prior is a gamma distribution, $\lambda \sim \text{Gamma}(a_\lambda, b_\lambda)$, here Gamma(1, 10 yr), its shape a_λ and rate b_λ written with the subscript because a is the age in this part of the book. Carried to ω with the change-of-variable factor $\mathrm{d}\lambda/\mathrm{d}\omega = \lambda$, it is the term $b_\lambda\lambda - a_\lambda \ln\lambda$ of equation (2.38), which is equation (15.4) with $n_4 = a_\lambda$ and $E_4 = b_\lambda$: the count row of one failure in ten coil-years, of mean $a_\lambda/b_\lambda = \lambda_{\text{ref}}$.

Which row acts on which entry. Every row's expected count is proportional to e^{ω}, so every row acts on the one entry, and the Jacobian is a column of four nonzero lines. With $*$ for an entry that can be nonzero,

$$\mathbf{J} = \begin{bmatrix} * \\ * \\ * \\ * \end{bmatrix} \begin{array}{c} \text{coil 0, ages 2–3} \\ \text{siblings, bin 2–3} \\ \text{siblings, bin 3–4'} \\ \text{prior row} \end{array} \qquad \frac{\mathrm{d}p_\tau}{\mathrm{d}\omega} = \lambda\tau\,\mathrm{e}^{-\lambda\tau} \neq 0. \tag{15.5}$$

The second line is not a row: it is the gradient of the first question's quantity, equation (15.1) through $\mathrm{d}\lambda/\mathrm{d}\omega = \lambda$. Two facts are read here, before any number. The siblings' rows and coil 0's own row share one column, which is the chapter's choice of a constant rate in the band written into the graph: they can only add information on one entry, never separate coil 0 from its siblings. And the question's quantity is a function of that same entry, so every answer is a one-dimensional integral over its posterior.

Modeling choices at the rows. Two choices make these rows. The prior Gamma(1, 10 yr) is centered on the book's reference rate and worth one failure in ten coil-years: this keeps the posterior proper whatever the count and lets the band's record dominate, and it misleads if the design's rate lies far from λ_{ref} and the band is short; step 5 varies it (Section 15.4.5). Coil 0's own year is its own row, not merged into the siblings' bin 2–3: this gives the same posterior as the merged bin, because counts in exposures at one rate add, and it lets step 3 show the row's share and the fifth question test it.

15.2 The posterior of the state (step 3)

All five questions of Table 15.2 are read from one posterior, that of the log-rate (15.3) under the four rows of step 2.

15.2.1 The objective and its Newton fit

The objective is the sum of the four terms, equation (4.39),

$$\mathcal{L}(\omega) = \sum_{\ell=1}^{4} \big(\mu_\ell - n_\ell \ln \mu_\ell\big) = (b_\lambda + T_{\text{tot}})\,\lambda - (a_\lambda + k)\ln\lambda + \text{const}, \qquad \lambda = \lambda_{\text{ref}}\,\mathrm{e}^{\omega}, \tag{15.6}$$

and $p(\omega \mid \mathbf{y}) \propto \exp[-\mathcal{L}(\omega)]$. On the log-rate each row's $\ln \mu_\ell = \ln E_\ell + \ln \lambda_{\text{ref}} + \omega$ is linear in the state, with derivative 1 and second derivative 0, so equation (4.42) gives each row's slope and curvature, and they add:

$$\mathcal{L}'_\ell = \mu_\ell - n_\ell, \quad \mathcal{L}''_\ell = \mu_\ell, \qquad \mathcal{L}' = (b_\lambda + T_{\text{tot}})\,\lambda - (a_\lambda + k), \quad \mathcal{L}'' = (b_\lambda + T_{\text{tot}})\,\lambda. \tag{15.7}$$

The curvature term $\mathbf{C}$ of equation (4.44) is zero: the log-rate is the canonical parameter of the Poisson count, and the curvature of each row is its expected count. The Newton step of equation (4.40), in one entry, is

$$\Delta\omega = -\frac{\mathcal{L}'(\omega)}{\mathcal{L}''(\omega)}, \qquad \omega \leftarrow \omega + \Delta\omega, \tag{15.8}$$

and the slope vanishes where $(b_\lambda + T_{\text{tot}})\lambda = a_\lambda + k$. The MAP state, the information and the variance are

$$\hat\lambda = \lambda_{\text{ref}}\,\mathrm{e}^{\hat\omega} = \frac{a_\lambda + k}{b_\lambda + T_{\text{tot}}}, \qquad \Lambda = \mathcal{L}''(\hat\omega) = a_\lambda + k, \qquad \sigma_\omega^2 = \frac{1}{\Lambda} = \frac{1}{a_\lambda + k}, \tag{15.9}$$

and the Laplace normal is $\omega \mid \mathbf{y} \approx \mathcal{N}(\hat\omega, \sigma_\omega^2)$. The information is the sum of the rows' expected counts at the MAP state, $\Lambda = \sum_\ell \mu_\ell$, so each row's share of it is μ_ℓ/Λ: a count row's information is its expected count, whether or not its count is zero.

The rows stacked. Stacked as a column, the four rows enter through the working pair of equation (4.43), here with $\nabla \ln \mu_\ell = 1$,

$$r_\ell = \frac{\mu_\ell - n_\ell}{\sqrt{\mu_\ell}}, \qquad j_\ell = \sqrt{\mu_\ell}, \qquad \mathbf{J}^{\mathsf{T}}\mathbf{r} = \sum_\ell j_\ell r_\ell = \mathcal{L}', \qquad \mathbf{J}^{\mathsf{T}}\mathbf{J} = \sum_\ell j_\ell^2 = \mathcal{L}'', \tag{15.10}$$

with $\mathbf{J}$ the column of the pattern (15.5) and $\mathbf{r}$ the column of the four residuals. The fit uses only the two products, which are the slope and the curvature of equation (15.7); the residuals themselves are read once more, in the residual check of Section 15.4.3.

15.2.2 The fit of the band's record

The iterations. Table 15.3 runs the Newton step (15.8) from the band's maximum-likelihood rate $k/T_{\text{tot}} = 8/69$, where the count rows' slope is zero. Each row is an iteration; the columns are the log-rate, the rate, the slope of the three count rows, that of the prior row, the slope and the curvature of the objective, and the step. The slope of the objective falls from 0.1594 to 1.110×10^{-15} in three steps, and the curvature settles at 9.0000.

The MAP state and its variance. By equation (15.9),

$$\hat\omega = 0.13036, \qquad \hat\lambda = \frac{1+8}{10+69}\,\text{yr}^{-1} = \frac{9}{79}\,\text{yr}^{-1} = 0.113924\,\text{yr}^{-1},$$
$$\Lambda = 9.0000, \qquad \sigma_\omega^2 = 0.11111, \qquad \sigma_\omega = 0.33333.$$

The state has one entry, so there are no cross terms: the budget of the variance is the budget of the information, row by row.

Table 15.3: Newton iterations of the log-rate, equation (15.8), from the band's maximum-likelihood rate 8/69.

j [–]	ω_j [–]	λ [yr^{-1}]	$\sum_{\ell\le 3}\mathcal{L}'_\ell$ [–]	$\mathcal{L}'_4$ [–]	$\mathcal{L}'$ [–]	$\mathcal{L}''$ [–]	$\Delta\omega$ [–]
0	0.14792	0.115942	0.0000	0.1594	1.594×10^{-1}	9.1594	-1.741×10^{-2}
1	0.13052	0.113942	−0.1380	0.1394	1.379×10^{-3}	9.0014	-1.532×10^{-4}
2	0.13036	0.113924	−0.1392	0.1392	1.057×10^{-7}	9.0000	-1.174×10^{-8}
3	0.13036	0.113924	−0.1392	0.1392	1.110×10^{-15}	9.0000	-1.234×10^{-16}

The rows at the MAP state. Table 15.4 gives each row's count, exposure, expected count, working residual and working line, equation (15.10). The products are $\mathbf{J}^\mathsf{T}\mathbf{r}=\sum_\ell(\mu_\ell-n_\ell)=-2.22\times10^{-16}$ and $\mathbf{J}^\mathsf{T}\mathbf{J}=\sum_\ell\mu_\ell=9.0000$. Each row's share of the information, μ_ℓ/Λ, is 46.84 % for the siblings' bin 2–3, 39.24 % for their bin 3–4, 12.66 % for the prior row and 1.27 % for coil 0's own year. A clean coil-year carries $E_1\hat\lambda=0.1139$ of information, as much as a coil-year with a failure: what it lacks is not information but length.

Table 15.4: The four rows at the MAP state, in the working form of equation (15.10), with each row's share of the information $\Lambda=9.0000$.

Row	n_ℓ [–]	E_ℓ [coil-yr]	μ_ℓ [–]	r_ℓ [–]	j_ℓ [–]	share [%]
coil 0, ages 2–3	0	1	0.1139	0.3375	0.3375	1.27
siblings, bin 2–3	3	37	4.2152	0.5919	2.0531	46.84
siblings, bin 3–4	5	31	3.5316	−0.7813	1.8793	39.24
prior row	1	10	1.1392	0.1305	1.0674	12.66

15.2.3 The exact posterior of the log-rate

The objective (15.6) identifies the exact posterior. Written back in λ, without the change-of-variable factor, the negative log posterior is $(b_\lambda+T_{\text{tot}})\lambda-(a_\lambda+k-1)\ln\lambda$, the kernel of a gamma density, so that by equation (2.39)

$$\lambda\mid\mathbf{y}\sim\text{Gamma}(a'_\lambda,b'_\lambda),\qquad a'_\lambda=a_\lambda+k=9,\quad b'_\lambda=b_\lambda+T_{\text{tot}}=79\,\text{yr},\tag{15.11}$$

and the log-rate has the log-gamma density

$$p(\omega\mid\mathbf{y})=\frac{(b'_\lambda\lambda_{\text{ref}})^{a'_\lambda}}{\Gamma(a'_\lambda)}\exp(a'_\lambda\,\omega-b'_\lambda\lambda_{\text{ref}}\,\mathrm{e}^{\omega}),\tag{15.12}$$

which is $\exp[-\mathcal{L}(\omega)]$ of equation (15.6), normalized. Its mode is $\hat\omega$. Its mean and variance are those of the logarithm of a gamma variable, $\mathbb{E}[\omega\mid\mathbf{y}]=(\ln\Gamma)'(a'_\lambda)-\ln(b'_\lambda\lambda_{\text{ref}})$ and $\text{Var}(\omega\mid\mathbf{y})=(\ln\Gamma)''(a'_\lambda)$, the first and second derivatives of the log-gamma function (the digamma and trigamma functions) at the posterior shape.

The grid on the state. The exact posterior of the state is read on a grid of the one entry. With equally spaced nodes ω_m and weights

$$\varpi_m=\frac{\exp[-\mathcal{L}(\omega_m)]}{\sum_{m'}\exp[-\mathcal{L}(\omega_{m'})]},\qquad\mathbb{E}[f(\omega)\mid\mathbf{y}]\approx\sum_m\varpi_m\,f(\omega_m),\tag{15.13}$$

a mean is a weighted sum, a probability the sum of the weights where the event holds, and a quantile is read from the cumulative sum of the weights. The grid has 20 001 nodes, from $14\sigma_\omega$ below the MAP state to $8\sigma_\omega$ above it, wider on the left where the density leans.

The numbers. On ω the grid gives the mean 0.07378, the standard deviation 0.34280 and the skewness -0.3424; the closed forms give $(\ln\Gamma)'(9)-\ln(7.9) = 0.07378$ and $\sqrt{(\ln\Gamma)''(9)} = 0.34280$. The mean lies below the mode $\hat{\omega} = 0.13036$, and the standard deviation exceeds the Laplace $\sigma_\omega = 0.33333$: the log-gamma leans to the left. On λ, the exact posterior (15.11) has the median 0.1097, the mean 0.1139, the standard deviation 0.0380 and the 5–95 % band $[0.0594,\ 0.1827]\,\mathrm{yr}^{-1}$. The Laplace normal on ω makes the rate lognormal, with the median 0.1139, the mean 0.1204, the standard deviation 0.0413 and the band $[0.0658,\ 0.1971]\,\mathrm{yr}^{-1}$.

The step's surprise. The Laplace mode $\hat{\lambda}$ equals the gamma's mean, $0.113924\,\mathrm{yr}^{-1}$ against $0.113924\,\mathrm{yr}^{-1}$, as Chapter 2 found for every gamma posterior. But the Laplace normal is symmetric in ω where the log-gamma leans left, so its mean rate is 0.1204 against the exact 0.1139, 5.71 % high. A short window's failure probability is nearly linear in the mean rate, and that excess passes to the answer.

15.2.4 The route of each answer

The reported answers are read from the exact posterior of the state: the log-gamma (15.12) on the grid (15.13), with the closed form of the gamma posterior (15.11) beside each value. The Laplace normal $\mathcal{N}(\hat{\omega}, \sigma_\omega^2)$ stands beside both, a median beside a median. The Laplace route is not the reported one because of the step's surprise: its mean rate is 5.71 % above the exact posterior's, and a short window's failure probability is nearly linear in the mean rate, so every probability read from it carries that excess. The exact posterior of this one entry is at hand, on the grid and in closed form, and nothing is gained by approximating it. Exact gamma draws, $N_\mathrm{d} = 200\,000$ from a seeded generator, carry the Monte Carlo check of step 5.

15.3 The answers (step 4)

Every answer is read from the posterior of the state of Section 15.2: the exact posterior of ω on the grid (15.13), which is reported, with the closed form of the gamma posterior (15.11) beside it and the Laplace normal $\mathcal{N}(\hat{\omega}, \sigma_\omega^2)$ beside both.

15.3.1 The probability of a failure in the next 90 days

Asset manager's question. Given its rate today, what is the probability it fails in the next 90 days? The answer is the number the contingency is weighed against.

General form. The predictive of Chapter 2, equation (2.13), in its event form (2.8),

$$\Pr(A \mid \mathbf{y}) = \int \Pr(A \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z},$$

for an event A that stays random when the state is known.

Specialized form. The question is equation (2.13) in its event form (2.8) with A the event $T_f \le \tau$, the state the one number ω, and $\Pr(A \mid \omega) = p_\tau(\lambda_\mathrm{ref} \mathrm{e}^{\omega})$ of equation (15.1). On the exact posterior of the state it is the grid sum (15.13),

$$\bar{p}_\tau = \int \left(1 - \mathrm{e}^{-\lambda_\mathrm{ref}\mathrm{e}^{\omega}\tau}\right) p(\omega \mid \mathbf{y})\, \mathrm{d}\omega \approx \sum_m \varpi_m \left(1 - \mathrm{e}^{-\lambda_\mathrm{ref}\mathrm{e}^{\omega_m}\tau}\right), \tag{15.14}$$

and the same integral over the gamma posterior (15.11) is equation (15.2), whose closed form is equation (3.17) with the event reversed,

$$\bar{p}_\tau = 1 - \left(\frac{b'_\lambda}{b'_\lambda + \tau}\right)^{a'_\lambda}. \tag{15.15}$$

On the Laplace route $p(\omega \mid \mathbf{y})$ is $\mathcal{N}(\hat{\omega}, \sigma_\omega^2)$, and equation (15.14) is the five-node Gauss–Hermite sum of equation (2.41) with the integrand reversed, $\bar{p}_\tau^{\mathrm{L}} = \sum_{j=1}^{5} w_j (1 - \mathrm{e}^{-\lambda_{\mathrm{ref}} \mathrm{e}^{\omega_j} \tau})$ at $\omega_j = \hat{\omega} + \sqrt{2}\,\sigma_\omega x_j$. One choice enters here: the window starts at the end of window 36 and the band's rate is held over it, the window lying at ages 3.000 to 3.247, inside the band. This is what makes the conditional probability the exponential form (15.1), and it fails if the rate within the window differs from the band's, by age or by the summer duty.

The plug-in and its gap. A plug-in evaluates equation (15.1) at one estimate of the rate. The failure probability is concave in the rate, $\mathrm{d}^2 p_\tau / \mathrm{d}\lambda^2 = -\tau^2 \mathrm{e}^{-\lambda\tau} < 0$, so by Jensen's inequality $\bar{p}_\tau \le p_\tau(\mathbb{E}[\lambda \mid \mathbf{y}])$ under any posterior, with equality only when the posterior is a single point; a plug-in at another estimate can lie on either side. For the gamma posterior the gap has a bound. With $\chi = \tau / b'_\lambda$, the window in units of the posterior's rate parameter, and the mean rate a'_λ / b'_λ,

$$\ln \frac{1 - \bar{p}_\tau}{1 - p_\tau(\mathbb{E}[\lambda \mid \mathbf{y}])} = -a'_\lambda \ln(1 + \chi) + a'_\lambda \chi = a'_\lambda (\chi - \ln(1 + \chi)).$$

The bracket is zero at $\chi = 0$, and its derivative, $\chi/(1 + \chi)$, lies between 0 and χ, so $0 \le \chi - \ln(1 + \chi) \le \chi^2/2$; and $a'_\lambda \chi^2 / 2 = \frac{1}{2} \tau^2 a'_\lambda / b'^2_\lambda = \frac{1}{2} \tau^2 \operatorname{Var}(\lambda \mid \mathbf{y})$. Therefore

$$0 \le p_\tau(\mathbb{E}[\lambda \mid \mathbf{y}]) - \bar{p}_\tau \le (1 - p_\tau(\mathbb{E}[\lambda \mid \mathbf{y}])) (\mathrm{e}^{\tau^2 \operatorname{Var}(\lambda \mid \mathbf{y})/2} - 1) \approx \tfrac{1}{2}\, \tau^2 \operatorname{Var}(\lambda \mid \mathbf{y}), \tag{15.16}$$

the bound of *Reliability Engineering by Derivation*, Chapter 9, with the event reversed. The gap is second order in the window, while a plug-in at an estimate $\tilde{\lambda}$ other than the mean differs from $\bar{p}_\tau$ by about $\tau(\tilde{\lambda} - \mathbb{E}[\lambda \mid \mathbf{y}])$, first order.

The coordinate of the Laplace route. In the log-rate the same function has the slope and curvature

$$\frac{\mathrm{d}p_\tau}{\mathrm{d}\omega} = \lambda\tau\, \mathrm{e}^{-\lambda\tau}, \qquad \frac{\mathrm{d}^2 p_\tau}{\mathrm{d}\omega^2} = \lambda\tau\, (1 - \lambda\tau)\, \mathrm{e}^{-\lambda\tau},$$

the second positive while $\lambda\tau < 1$: over a window short enough that less than one failure is expected, p_τ is convex in ω. An expansion about $\hat{\omega}$ under the Laplace normal gives

$$\bar{p}_\tau^{\mathrm{L}} \approx p_\tau(\hat{\lambda}) + \tfrac{1}{2} \frac{\mathrm{d}^2 p_\tau}{\mathrm{d}\omega^2}\Big|_{\hat{\omega}} \sigma_\omega^2, \tag{15.17}$$

which lies above the plug-in at $\hat{\lambda}$, and $\hat{\lambda}$ is the gamma's mean (Section 15.2.3). The exact predictive lies below that plug-in, and the cause is the concavity in λ of the bound (15.16), which holds under any posterior. Read in ω, the same fact is two effects of opposite sign: the convexity in ω raises the average, and the log-gamma's lean to the left, its mean of ω below its mode, offsets that rise and more. The Laplace normal keeps the first effect and drops the second. The two routes straddle the plug-in at the mean, and the Laplace route's error is the excess of its mean rate, not the averaging.

The predictive hazard. The closed form (15.15) is the distribution function of the time to the next failure under the predictive, and its hazard is

$$\lambda_{\mathrm{pred}}(\tau) = -\frac{\mathrm{d}}{\mathrm{d}\tau} \ln(1 - \bar{p}_\tau) = \frac{a'_\lambda}{b'_\lambda + \tau} \tag{15.18}$$

(Exercise 15.2). It falls with τ although every rate the posterior admits is constant, because running τ without a failure is one more clean record, which favors the smaller rates: the frailty of a mixture of exponential lifetimes of *Reliability Engineering by Derivation*, Chapter 9, here over one coil's unknown rate. It is an effect of not knowing the rate, not of age.

Building blocks. The exact posterior (15.12) on the grid (15.13), and its closed form with $a'_\lambda = 9$ and $b'_\lambda = 79\,\mathrm{yr}$ (Section 15.2.3); for the Laplace route $\hat{\omega} = 0.13036$ and $\sigma_\omega = 0.33333$ (Section 15.2.2) with the five Gauss–Hermite nodes; the window $\tau = 0.24658\,\mathrm{yr}$; the posterior mean a'_λ / b'_λ and variance $a'_\lambda / b'^2_\lambda$ of the rate for the bound (15.16); the slope and curvature in ω at $\hat{\omega}$ for equation (15.17).

The 90-day probability. The grid sum (15.14) gives $\bar{p}_\tau = 0.027657$, and the closed form (15.15), $1 - (79/79.24658)^9$, gives 0.027657: the probability that coil 0 fails in the next 90 days is 2.77 %. Table 15.5 evaluates the Laplace route. Each row is a node; the columns are the node and its weight, the log-rate and the rate there, the failure probability at that rate and the weighted term. The terms sum to 0.029209, and 200 nodes give the same: the Laplace route lies 0.001551 above the exact value, 5.61 % of it.

Table 15.5: The five-node Gauss–Hermite sum of the Laplace route for the 90-day failure probability, $\hat{\omega} = 0.13036$, $\sigma_\omega = 0.33333$, $\tau = 0.24658\,\mathrm{yr}$.

x_j [–]	w_j [–]	ω_j [–]	λ_j [yr^{-1}]	$p_\tau(\lambda_j)$ [–]	term [–]
−2.02018	0.01126	−0.82196	0.043957	0.010780	0.000121
−0.95857	0.22208	−0.32151	0.072505	0.017719	0.003935
0	0.53333	0.13036	0.113924	0.027700	0.014773
0.95857	0.22208	0.58224	0.179004	0.043178	0.009589
2.02018	0.01126	1.08269	0.295260	0.070217	0.000790
Sum, $\bar{p}_\tau^{\mathrm{L}}$					0.029209

The plug-ins. Table 15.6 sets the plug-in at six estimates of the rate against the exact predictive. Each row is an estimate; the columns give the rate, the failure probability at it, and the exact $\bar{p}_\tau$ less that plug-in. At the posterior mean, which is also the Laplace mode, the plug-in is 0.027700 and the gap −0.000043, inside the bound (15.16): the plug-in exceeds the predictive by 4.25×10^{-5}, against the bound 4.26×10^{-5} and $\frac{1}{2}\tau^2 \operatorname{Var}(\lambda \mid \mathbf{y}) = 4.38 \times 10^{-5}$. At every other estimate the gap is first order: the exact median sets the plug-in 0.000963 low, the band's maximum-likelihood rate 0.000526 high, the fleet dashboard 0.000724 low, and coil 0 alone at zero.

Table 15.6: The plug-in failure probability $1 - \mathrm{e}^{-\lambda\tau}$ over 90 days at six estimates of the rate, against the exact predictive $\bar{p}_\tau = 0.027657$.

Estimate of the rate	λ [yr^{-1}]	$p_\tau(\lambda)$ [–]	$\bar{p}_\tau - p_\tau(\lambda)$ [–]
Posterior mean a'_λ/b'_λ, the Laplace mode	0.1139	0.027700	−0.000043
Exact posterior median	0.1097	0.026695	+0.000963
The band's maximum likelihood $k/T_{\mathrm{tot}} = 8/69$	0.1159	0.028184	−0.000526
The Laplace mean	0.1204	0.029259	−0.001602
The fleet dashboard, 32/289	0.1107	0.026933	+0.000724
Coil 0 alone, 0/3	0	0	+0.027657

The orders of the readings. The first-order reading $\tau\mathbb{E}[\lambda \mid \mathbf{y}] = 0.028091$ sits above the predictive by the curvature it leaves out; on the Laplace mean it is 0.029696. The second-order reading (15.17) is $0.027700 + 0.001475 = 0.029175$, beside the Laplace route's 0.029209 and not beside the exact value. The expected number of failures in the window is $\mathbb{E}[\lambda \mid \mathbf{y}]\,\tau = 0.02809$, and the probability of two or more is 0.000428. The predictive hazard (15.18) is $0.11392\,\mathrm{yr}^{-1}$ today, 0.11357 at the end of the 90 days and 0.11285 at the shutdown.

The densities. The two posteriors of the rate are drawn in Figure 15.3: the horizontal axis is the rate, the vertical axis its posterior density; the solid curve is the exact gamma and the dashed one the Laplace lognormal, the shaded area the exact 5–95 % band, and the vertical lines the rates the plug-ins use and the stated truth, a check marker for Section 15.4.6. The

two curves peak together near $0.10\,\mathrm{yr}^{-1}$; the Laplace curve sits slightly right of the exact one and carries the heavier upper tail, which is its higher mean. Coil 0's own zero is the vertical axis itself, where the exact density is small.

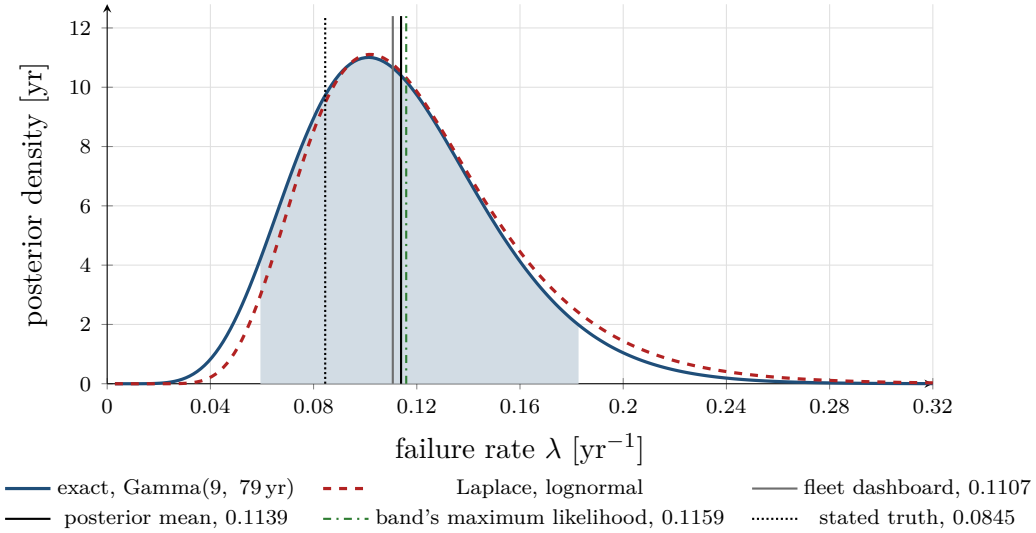


Figure 15.3: The posterior of the rate of coil 0's tube failures: the exact gamma Gamma(9, 79 yr) of equation (15.11) (solid) and the Laplace lognormal, $\ln(\lambda/0.1) \sim \mathcal{N}(0.13036,\ 0.33333^2)$ (dashed). Shaded: the exact 5–95 % band, 0.0594 to $0.1827\,\mathrm{yr}^{-1}$. Vertical lines: the rates of the plug-ins of Table 15.6, except coil 0 alone, whose rate of zero lies on the vertical axis, and the stated truth's rate for the window, a check marker (Section 15.4.6).

The step's surprise. Averaging over the rate changes the 90-day probability by 0.000043; the Laplace normal changes it by 0.001551, 36.5 times as much, and a plug-in at the median instead of the mean by 0.000963. For a window this short the question is which rate and which summary, not whether to average.

15.3.2 The rate today and the MTBF

Asset manager's question. What is its failure rate today, and its MTBF?

General form. An estimation question asks for the posterior distribution of a derived quantity, equation (2.9),

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

summarized by its median and its credible interval, equation (2.11), at the levels 0.05 and 0.95.

Specialized form. The question is equations (2.9) and (2.11) with the state ω and $f(\omega) = \lambda_{\mathrm{ref}} \mathrm{e}^{\omega}$, increasing, for the rate, and $f(\omega) = \mathrm{e}^{-\omega}/\lambda_{\mathrm{ref}}$, decreasing, for the MTBF. The quantiles of both are the quantile map of the quantiles ω_q of the log-rate:

$$\lambda_q = \lambda_{\mathrm{ref}}\, \mathrm{e}^{\omega_q}, \qquad (1/\lambda)_q = 1/\lambda_{1-q}, \tag{15.19}$$

so the MTBF's band is the rate's band reversed. On the exact posterior the ω_q are read from the cumulative weights of the grid (15.13), and they are the logarithms of the gamma quantiles; on the Laplace route $\omega_q = \hat{\omega} + z_q \sigma_\omega$ with $z_{0.95} = -z_{0.05} = 1.6449$. A mean does not pass through the map: the exact mean of the MTBF is $\mathbb{E}[1/\lambda \mid \mathbf{y}] = b'_\lambda/(a'_\lambda - 1)$, against $1/\mathbb{E}[\lambda \mid \mathbf{y}] = b'_\lambda/a'_\lambda$.

Building blocks. The grid's cumulative weights, or the gamma quantiles of Gamma(9, 79 yr); $\hat{\omega}$ and σ_ω for the Laplace route; the map (15.19).

The rate and the MTBF. The rate today has the median $0.1097\,\mathrm{yr}^{-1}$ and the band $[0.0594,\ 0.1827]\,\mathrm{yr}^{-1}$; on the Laplace route the median is 0.1139 and the band [0.0658, 0.1971]. The MTBF has the median 9.11 years and the band [5.47, 16.83] years; on the Laplace route 8.78 and [5.07, 15.19]. The exact mean of the MTBF, $79/8 = 9.88$ years, is not the reciprocal of the mean rate, 8.78 years. The probability that the rate exceeds the dashboard's fleet rate, $0.1107\,\mathrm{yr}^{-1}$, is 0.4894 (Laplace 0.5340): on this record the band's rate sits on the fleet's.

The prior's effect. For the one coil, the prior pulls the mean rate from the band's maximum likelihood, $0.1159\,\mathrm{yr}^{-1}$, to $0.1139\,\mathrm{yr}^{-1}$, toward λ_{ref}: it is one failure in ten coil-years against eight in sixty-nine, a rate below the band's.

15.3.3 How well the probability is known, and the 5 % limit

Asset manager's question. How well is that probability known, and is it below 5 % with 95 % assurance?

General form. A prediction under a threshold is a tail of a derived quantity under the posterior, compared with an assurance level, equation (2.12),

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

Specialized form. The question is equation (2.12) with $f(\omega) = -p_\tau(\lambda_{\mathrm{ref}}\mathrm{e}^{\omega})$ and $f^* = -p^*$: the requirement is $p_\tau(\lambda) < p^*$. The failure probability is increasing in the rate and the rate in ω, so the event is a level on the rate,

$$\Pr\big(p_\tau(\lambda) < p^* \mid \mathbf{y}\big) = \Pr\big(\lambda < \lambda^* \mid \mathbf{y}\big), \qquad \lambda^* = -\frac{\ln(1-p^*)}{\tau}, \tag{15.20}$$

the sum of the grid's weights below $\ln(\lambda^*/\lambda_{\mathrm{ref}})$ on the exact posterior, the gamma distribution function at λ^* in closed form, and $\Phi\big((\ln(\lambda^*/\lambda_{\mathrm{ref}}) - \hat{\omega})/\sigma_\omega\big)$ on the Laplace route. The posterior of $p_\tau(\lambda)$ itself follows by the quantile map (15.19), its quantiles $p_\tau(\lambda_q)$; its mean is the predictive $\bar{p}_\tau$ of equation (15.2), and its spread to first order is the slope of the pattern's second line (15.5) times σ_ω, the delta method of equation (4.53). One choice enters here: the limit $p^* = 5\,\%$ and the assurance $\kappa = 0.95$ are the asset manager's levels for the contingency. They are what turns a band into a verdict, and another policy would set other levels; the answer is given at three more limits below.

Building blocks. The grid's weights or the gamma distribution function, and $\hat{\omega}$, σ_ω for the Laplace route; the level (15.20); the map of the rate's quantiles; the slope $\mathrm{d}p_\tau/\mathrm{d}\omega$ at $\hat{\omega}$.

The band and the assurance. Coil 0's 90-day failure probability has the median 2.67 % and the band [1.45, 4.41] %; on the Laplace route 2.77 % and [1.61, 4.74] %. Its mean is the predictive, 0.027657 (Laplace 0.029209): the two objects meet there. The limit $p^* = 0.05$ is the rate $\lambda^* = 0.2080\,\mathrm{yr}^{-1}$, and $\Pr(p_\tau(\lambda) < 0.05 \mid \mathbf{y}) = \Pr(\lambda < 0.2080 \mid \mathbf{y}) = 0.9827$ (Laplace 0.9646): the 95 % assurance is met. The level shown with 95 % assurance is 4.41 % (Laplace 4.74 %). At other limits the assurance is 0.2052 at 2 % ($\lambda^* = 0.0819$; Laplace 0.1614), 0.6394 at 3 % (0.1235; 0.5959) and 0.9038 at 4 % (0.1656; 0.8689). The delta method gives the standard deviation of $p_\tau(\lambda)$ as $0.027313 \times 0.33333 = 0.009104$.

15.3.4 The shutdown and the window at which 5 % is passed

Asset manager's question. Will it run to the October shutdown, and in which window does its failure probability pass 5 %?

General form. The predictive, equation (2.13), for the first part; for the second, the posterior distribution of a derived quantity, equation (2.9), read by its median and credible interval.

Specialized form. The first part is equation (2.13) in the event form of Section 15.3.1 at the window τ_{sd}: the probability of running to the shutdown is $1 - \bar{p}_{\tau_{\mathrm{sd}}}$, the grid sum (15.14) and the closed form (15.15) with τ_{sd} for τ. The second is equation (2.9) with $f(\omega)$ the window length at which the coil's own failure probability reaches p^*,

$$\tau_p(\lambda) = -\frac{\ln(1-p^*)}{\lambda}, \qquad \ln \tau_p = \ln \frac{-\ln(1-p^*)}{\lambda_{\mathrm{ref}}} - \omega, \tag{15.21}$$

a ratio, decreasing in the rate. Its quantiles are the rate's reversed, $(\tau_p)_q = \tau_p(\lambda_{1-q})$, the median and the inversion band. Its logarithm is linear in ω, so on the Laplace route the delta form is exact: $\ln \tau_p$ is normal with the standard deviation σ_ω, and its band is the Laplace band. A time τ ahead becomes a window by

$$t = 36 + \lceil 12\,\tau \rceil, \tag{15.22}$$

with the window's calendar label. Beside the coil's own $\tau_p(\lambda)$ stands the window at which the predictive itself passes p^*, the solution of $\bar{p}_\tau = p^*$ in equation (15.15), $\tau = b'_\lambda((1-p^*)^{-1/a'_\lambda} - 1)$.

Building blocks. The grid and the closed form at $\tau_{\mathrm{sd}} = 0.75\,\mathrm{yr}$; the rate's quantiles and the inverse window (15.21); $\hat{\omega}$ and σ_ω for the delta form; the window map (15.22).

The shutdown. The probability of a failure before the shutdown is $\bar{p}_{\tau_{\mathrm{sd}}} = 8.15\,\%$, the probability of running to it 0.91848 (Laplace 8.59 %; the dashboard 7.97 %). The coil's own failure probability over the nine windows has the median 7.90 % and the band [4.36, 12.81] % (Laplace 8.19 % and [4.82, 13.74] %).

The window at 5 %. The inverse window τ_p has the median 0.4674 yr, 5.61 months, in window 42 (Jun Y4), and the inversion band [0.2807, 0.8630] yr, from window 40 (Apr Y4) to window 47 (Nov Y4). The delta form on $\ln \tau_p$ gives the Laplace values, the median 0.4502 yr and the band [0.2602, 0.7790] yr, windows 42 and 40 to 46. The predictive passes 5 % at 0.4515 yr, in window 42 (Jun Y4); on the Laplace route at 0.4272 yr, the plug-in at the mean at 0.4502 yr and the dashboard at 0.4632 yr, all in window 42.

Month by month. Table 15.7 gives the predictive failure probability from today to the end of each window of year 4. Each row is a window; the columns are its label and the exact and Laplace values. Figure 15.4 draws the same against the window: the solid curve is the exact predictive, the dashed one the Laplace route, the shaded area the 5–95 % band of the coil's own failure probability, the thin curve the dashboard's plug-in and the horizontal line the 5 % limit; the vertical lines mark the 90 days, at the end of window 39, and the shutdown, at the end of window 45; the circles are the stated truth. The exact predictive crosses the limit in window 42, the band's upper edge in window 40 and its lower edge in window 47.

15.3.5 Coil 0's own clean record

Asset manager's question. It has never failed: is its record evidence that it is better than its siblings?

Table 15.7: The predictive failure probability of coil 0 from today (window 36, Dec Y3) to the end of each window of year 4. The October shutdown falls at the end of window 45 (Sep Y4).

Window (label) [–]	$\bar{p}$, exact [–]	$\bar{p}$, Laplace [–]
37 (Jan Y4)	0.0094	0.0100
38 (Feb Y4)	0.0188	0.0198
39 (Mar Y4)	0.0280	0.0296
40 (Apr Y4)	0.0372	0.0393
41 (May Y4)	0.0462	0.0488
42 (Jun Y4)	0.0552	0.0582
43 (Jul Y4)	0.0641	0.0676
44 (Aug Y4)	0.0728	0.0768
45 (Sep Y4)	0.0815	0.0859
46 (Oct Y4)	0.0901	0.0950
47 (Nov Y4)	0.0986	0.1039
48 (Dec Y4)	0.1070	0.1127

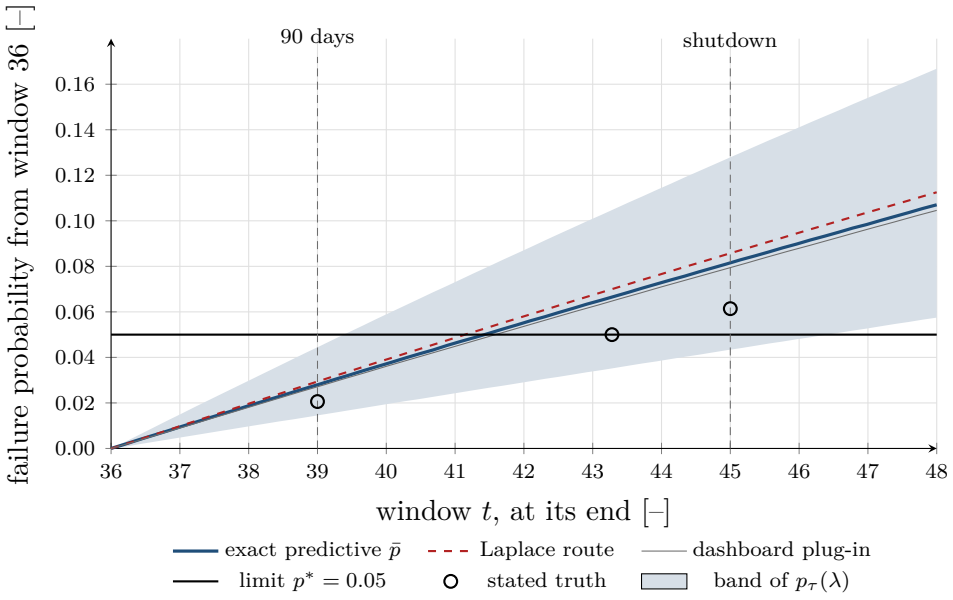


Figure 15.4: The failure probability of coil 0 from today, window 36, to the end of each later window: the exact predictive $1-(79/(79+\tau))^9$ (solid), the five-node Laplace route of Table 15.5 (dashed), the 5–95 % band of the coil's own failure probability $1-\mathrm{e}^{-\lambda\tau}$ between the rates 0.0594 and 0.1827 yr^{-1} (shaded), the dashboard's plug-in at 0.1107 yr^{-1} (thin), and the limit $p^* = 0.05$; $\tau = (t-36)/12$ yr. Dashed verticals: the 90 days (end of window 39) and the October shutdown (end of window 45). Circles: the stated truth, the generator's rate of the age bin 3–4, at the 90 days, at its 5 % window and at the shutdown, check markers (Section 15.4.6).

General form. A hypothesis comparison is the ratio of the evidences of two drawings fitted to the same evidence, the Bayes factor of equation (2.14),

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

Specialized form. The question is equation (2.14) with two drawings of coil 0's year in the band. In M_{sh}, the chapter's drawing, its rate is the siblings'; in M_{own} its rate is its own, with the same prior Gamma(a_λ, b_λ), and the siblings keep theirs. Each evidence is the normalizer of a posterior, the integral of equation (3.14) over the rate. The evidence of M_{sh} is the normalizer of step 3's posterior, all four rows on the one rate, $p(\mathbf{y})$; that of the siblings' rows alone with the prior is $p(\mathbf{y}_{\text{sib}})$, and their ratio is the probability of coil 0's row $\mathbf{y}_0$ given the siblings,

$$p(\mathbf{y}_0 \mid \mathbf{y}_{\text{sib}}, M_{\text{sh}}) = \frac{p(\mathbf{y})}{p(\mathbf{y}_{\text{sib}})}.$$

Under M_{own} the siblings' rows give the same factor $p(\mathbf{y}_{\text{sib}})$ and coil 0's row its own marginal under the prior, $p(\mathbf{y}_0 \mid M_{\text{own}})$, so the factor $p(\mathbf{y}_{\text{sib}})$ cancels from the ratio. For a count of zero in E_1 coil-years under a gamma with shape a_λ and rate b_λ the marginal is $\mathbb{E}[\mathrm{e}^{-\lambda E_1}] = (b_\lambda/(b_\lambda + E_1))^{a_\lambda}$, and under the siblings' posterior the shape and rate are $a_\lambda + 8$ and $b_\lambda + 68$, so

$$B_{\text{own,sh}} = \frac{p(\mathbf{y}_0 \mid M_{\text{own}})}{p(\mathbf{y})/p(\mathbf{y}_{\text{sib}})} = \frac{\big(b_\lambda/(b_\lambda + E_1)\big)^{a_\lambda}}{\big((b_\lambda + 68)/(b_\lambda + 69)\big)^{a_\lambda + 8}}, \tag{15.23}$$

with the Laplace evidence of equation (4.19), each normalizer from its own fit, beside it. One choice enters here: the whole-life version reads coil 0's three clean years against the siblings' bins 0–1, 1–2 and 2–3, each with its own rate from the siblings' log and the same prior, so that it assumes nothing across ages; it asks how often a coil of the design runs its first three years without a failure.

Building blocks. The normalizer of step 3's posterior (Section 15.2.3) and that of the siblings' rows; the gamma–Poisson marginal of a zero count; the Laplace evidence of each from its own fit.

The clean record. The log normalizers are $\ln p(\mathbf{y}) = -4.9944$ for step 3's posterior, $\ln p(\mathbf{y}_{\text{sib}}) = -4.8798$ for the siblings' rows with the prior, and $\ln p(\mathbf{y}_0 \mid M_{\text{own}}) = -0.0953$ for coil 0's row under the prior alone; the Laplace evidence gives -5.0037, -4.8890 and -0.1764. The probability of no failure in coil 0's year is therefore $\exp(-4.9944 + 4.8798) = 0.8917$ at the siblings' rate and 0.9091 at its own rate under the prior alone, and by equation (15.23) $\ln B_{\text{own,sh}} = -0.0953 + 4.9944 - 4.8798 = 0.0193$. The Laplace evidence gives -0.0617; it is worst for coil 0 alone, a zero count, at -0.1764 against the exact -0.0953. Over the whole life, the probability of three clean years given the siblings' bins is 0.8000: four coils in five of the design show coil 0's record. Its own constant rate under the prior gives 0.7692, and $\ln B = -0.0392$. Without coil 0's row the 90-day probability is 0.028006; the row moves it by -0.000349. Read alone, its three years with the prior, Gamma(1, 13 yr), coil 0's rate has the median $0.0533\,\text{yr}^{-1}$ and the band $[0.0039,\ 0.2304]\,\text{yr}^{-1}$, and a 90-day probability of 1.86 %: three coil-years leave the rate anywhere from 0.0039 to 0.2304 a year, and the asset manager's zero lies below that band. The clean record is no evidence either way; it is what the design's coils show.

15.4 The checks (step 5)

Each answer of Section 15.3 rests on a route, a choice of coils and a prior, and each can be set against a calculation that does without it.

15.4.1 The routes and the draws

Table 15.8 collects the routes for the first and third questions. Each row is a route; the columns give the 90-day probability, the median and band of the coil's own failure probability, and its assurance against 5 %. The grid on the exact posterior of the state and the conjugate closed form agree to six decimals. Exact draws from the gamma posterior, $N_{\mathrm{d}} = 200\,000$, give the predictive 0.027655 with a Monte Carlo standard error of 0.000020, the median of $p_\tau(\lambda)$ 0.0267 and its band [0.0146, 0.0440] with standard errors of 0.0000, 0.0000 and 0.0001 at four decimals, the assurance 0.9828 with 0.0003, and the probability of a failure before the shutdown 0.08152 with 0.00006, against the exact 0.08152. The draws resolve the Laplace route's gap, 0.001551 against a standard error of 0.000020; the plug-in's gap, 0.000043, is below what they resolve and is read from the closed form.

Table 15.8: The 90-day failure probability of coil 0 by route, with the band of its own failure probability and its assurance against 5 %, and the stated truth.

Route	$\bar{p}_\tau$ [–]	$p_\tau(\lambda)$ median [band] [–]	$\Pr(p_\tau < 0.05)$ [–]
Exact posterior of the state, grid on ω (reported)	0.027657	0.0267 [0.0145, 0.0441]	0.9827
Conjugate closed form	0.027657	0.0267 [0.0145, 0.0441]	0.9827
Exact draws, $N_{\mathrm{d}} = 200\,000$	0.027655 (se 0.000020)	0.0267 [0.0146, 0.0440]	0.9828 (se 0.0003)
Laplace normal, five nodes (200 the same)	0.029209	0.0277 [0.0161, 0.0474]	0.9646
Second-order reading in ω	0.029175	–	–
Plug-in at the posterior mean	0.027700	–	–
Fleet dashboard (MTBF 9.031 yr)	0.026933	–	–
Coil 0's own history	0	–	–
Stated truth, the generator's bin rate	0.0206	–	–

15.4.2 The limits of a longer record and a shorter window

The chapter has no physical rows, so the limit of exact physical rows has no counterpart here. Two other limits test the answers, and the curvature term is checked numerically. The record multiplied at the same rate, every count and exposure ten, a hundred and a thousand times, takes the Laplace route's excess over the exact value from 1.55×10^{-3} to 1.72×10^{-4}, 1.73×10^{-5} and 1.74×10^{-6}, and the exact value less the plug-in at the mean from -4.25×10^{-5} to -4.88×10^{-6}, -4.96×10^{-7} and -4.96×10^{-8}: both gaps fall as the reciprocal of the exposure, and the predictive becomes the plug-in. The window shortened to a hundredth gives the exact $\bar{p}_\tau/\tau = 0.113906$ against $\mathbb{E}[\lambda \mid \mathbf{y}] = 0.113924$, and the Laplace 0.120412 against its mean rate 0.120432: the Laplace route's excess is its mean rate. The finite-difference curvature of the objective (15.6) at $\hat{\omega}$ is 9.000000 against $\mathbf{J}^{\mathsf{T}}\mathbf{J} = 9.000000$, so $\mathbf{C} = 0$ as Section 15.2.1 stated. Carried on λ itself, the mode moves to $(a'_\lambda - 1)/b'_\lambda = 0.10127$ and the curvature there is 616.35, a coordinate in which a count row's curvature term does not vanish: the log-rate is the coordinate in which the information is the Hessian.

15.4.3 The rows' residuals

The working residuals of Table 15.4 are read by the residual check of Chapter 7. A row's leverage is $K_{ii} = j_\ell^2/\Lambda = \mu_\ell/\Lambda$, its share of the information, and its normalized residual $r_\ell/\sqrt{1 - K_{ii}}$: 0.3397 for coil 0's year, 0.8118 for the siblings' bin 2–3 and -1.0024 for their bin 3–4. The deviance of the three count rows, $\sum_\ell 2\big[\mu_\ell - n_\ell + n_\ell \ln(n_\ell/\mu_\ell)\big]$, is 1.1578 on 2 degrees of freedom, three rows less one entry, with $P = 0.5605$. The record does not reject one rate across the band's two bins; with eight failures it cannot test more.

15.4.4 The age band

The constant rate in the band is the chapter's central choice, and the band is a choice too. Table 15.9 repeats the exact posterior, with the same prior, for six other sets of coil-years. Each row is a band; the columns give its failures and exposure, the mean age of that exposure, the rate's median and band, the 90-day probability on the exact and the Laplace routes, and the standard deviation of the coil's own failure probability; the last line gives each band's assurance against 5 %. The same answers are drawn in Figure 15.5, one row per band: the dot is the exact predictive, the bar the 5–95 % band of $p_\tau(\lambda)$, the dashed line the dashboard's 2.69 %, the dotted line the stated truth and the solid line the 5 % limit.

Table 15.9: The 90-day answer under other choices of the coil-years that speak for coil 0: the exact posterior with the prior Gamma(1, 10 yr). This is how the constant rate in a band can fail.

Band [yr]	k [–]	T_{tot} [coil-yr]	mean age [yr]	rate median [band] [yr^{-1}]	$\bar{p}_\tau$, exact [–]	$\bar{p}_\tau$, Laplace [–]	sd p_τ [–]
3–4	5	31	3.50	0.1383 [0.0637, 0.2564]	0.0353	0.0383	0.0141
2–4 (reported)	8	69	2.95	0.1097 [0.0594, 0.1827]	0.0277	0.0292	0.0091
2–5	11	95	3.37	0.1111 [0.0659, 0.1734]	0.0278	0.0289	0.0079
1–5	15	142	2.75	0.1031 [0.0660, 0.1520]	0.0256	0.0264	0.0063
0–4	14	166	1.80	0.0833 [0.0525, 0.1244]	0.0208	0.0215	0.0053
0–12, the fleet	32	289	3.95	0.1093 [0.0808, 0.1438]	0.0268	0.0272	0.0046
coil 0 alone, 0–3	0	3	1.50	0.0533 [0.0039, 0.2304]	0.0186	0.0300	0.0183
Assurance $\Pr(p_\tau < 0.05)$ [–], in the same order: 0.8526, 0.9827, 0.9917, 0.9992, 1.0000, 1.0000, 0.9331.							

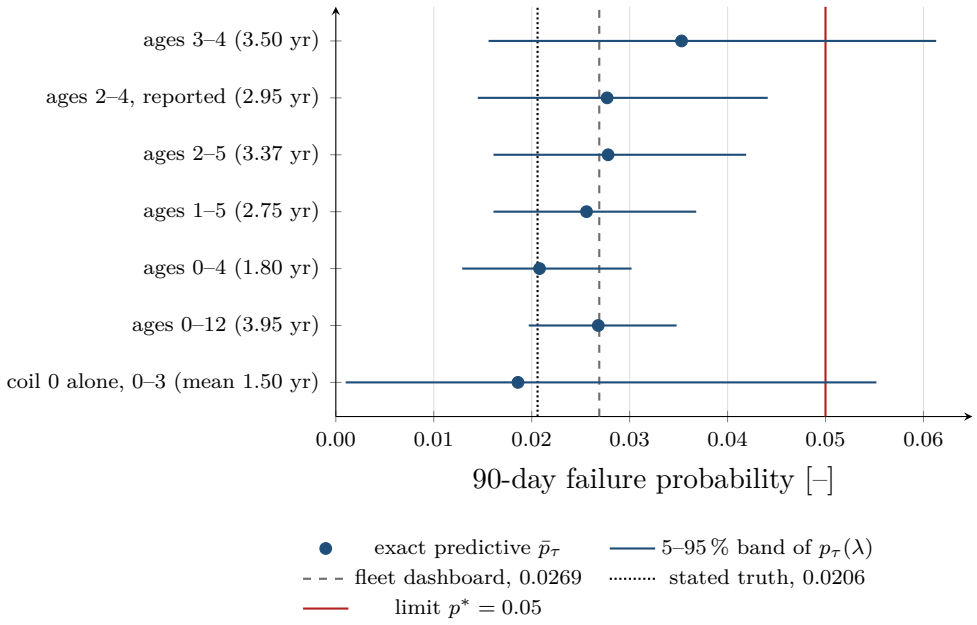


Figure 15.5: The 90-day failure probability of coil 0 under each choice of coil-years of Table 15.9, with the mean age of each band's exposure: the exact predictive (dots) and the 5–95 % band of the coil's own failure probability by the quantile map of the rate (bars), against the dashboard's plug-in (dashed), the stated truth (dotted, a check marker) and the 5 % limit.

The band moves the answer from 2.08 to 3.53 %, more than the reported band's own standard deviation of $p_\tau(\lambda)$, 0.0091. A narrow band lies near the window's ages and is noisy: the bin 3–4 alone gives 3.53 % with assurance 0.8526, below 0.95. A wide one is precise and is read at

another mean age: the bins 0–4 give 2.08 % at a mean age of 1.80 years. Which is right depends on whether the rate changes with age, which is Chapter 16's question; this chapter states the dependence of the answer on the band and does not settle it.

15.4.5 Sensitivity to the prior

Table 15.10 repeats the exact posterior under five other priors. Each row is a prior; the columns give the posterior it makes, the rate's median and band, the 90-day probability, the probability of a failure before the shutdown, and the assurance against 5 %. With eight failures in the band the prior moves the 90-day answer between 2.59 and 3.27 %, and the assurance stays between 0.9419 and 0.9995, met under every prior. The reported prior is among the lower ones because its mean, $0.1\,\mathrm{yr}^{-1}$, lies below the band's maximum likelihood, $0.1159\,\mathrm{yr}^{-1}$: the priors with no weight, Jeffreys and the flat prior, give 2.99 and 3.16 %, and Part I's maintainer prior of Section 2.4.3, Gamma(2, 5 yr), gives 3.27 %.

Table 15.10: The answers under other priors on the rate, on the exact posterior.

Prior	Posterior [–, yr]	rate median [band] [yr^{-1}]	$\bar{p}_\tau$ [–]	$\bar{p}_{\tau_{\mathrm{sd}}}$ [–]	$\Pr(p_\tau < 0.05)$ [–]
Gamma(1, 10 yr), the chapter's	Gamma(9, 79)	0.1097 [0.0594, 0.1827]	0.0277	0.0815	0.9827
Jeffreys, Gamma(1/2, 0)	Gamma(8.5, 69)	0.1184 [0.0628, 0.1999]	0.0299	0.0878	0.9627
Flat in λ, Gamma(1, 0)	Gamma(9, 69)	0.1256 [0.0680, 0.2092]	0.0316	0.0927	0.9479
Part I's maintainer, Gamma(2, 5 yr)	Gamma(10, 74)	0.1307 [0.0733, 0.2122]	0.0327	0.0959	0.9419
Wide, Gamma(1, 2 yr)	Gamma(9, 71)	0.1221 [0.0661, 0.2033]	0.0307	0.0902	0.9578
Tight on λ_{ref}, Gamma(10, 100 yr)	Gamma(18, 169)	0.1045 [0.0688, 0.1509]	0.0259	0.0766	0.9995

15.4.6 The answers against the stated truth

The record is constructed from a stated truth (Section 15.6.5), and the last check sets the answers against it. The truth is the generator of the log, which draws each one-year bin's failures at one rate for that bin; the window and the shutdown both lie in coil 0's bin 3–4. The truth's rate for the window is $0.0845\,\mathrm{yr}^{-1}$, at the 0.2299 quantile of the exact posterior (Laplace 0.1853), inside the band; its 90-day failure probability is 0.0206, inside the band [0.0145, 0.0441] of $p_\tau(\lambda)$, and it is the last row of Table 15.8. Before the shutdown the truth's failure probability is 0.0614, against the reported 8.15 %; its 5 % level, at 0.6069 yr, falls in window 44 (Aug Y4), inside the inversion band of windows 40 to 47, against the reported median of window 42. The truth is the dotted marker of Figures 15.3 and 15.5 and the circles of Figure 15.4.

This record against fresh logs. An answer is one record's draw. Over 20 000 fresh logs of the band's rows drawn from the generator at the same exposures, the truth expects 5.3847 failures where this record has 8, 1.13 Poisson standard deviations above. The reported 90-day probability over the fresh logs has the median 0.0185 and the 5–95 % range [0.0093, 0.0307], and this record's 0.0277 sits at its 0.904 quantile. Their mean, 0.0196, lies just below the truth's 0.0206, well inside the spread of the fresh logs: this record's excess over the truth is mostly its draw. The 5–95 % band of $p_\tau(\lambda)$ covers the truth in a share 0.8819 of the fresh logs, with a standard error of 0.0023.

15.5 The report (step 6)

Coil 1's clean record is not a zero. Its probability of a failure in the next 90 days is 2.77 %, its design's at its age; the band of the rate puts its own failure probability between 1.45 and 4.41 %, below 5 % with assurance 0.98. Its three clean years are what four coils in five of the design show ($P = 0.80$), and its year in the band is 1.3 % of the information: the record is no evidence that coil 1 is better than its siblings ($\ln B = 0.02$). Averaging over the rate changed the answer by 0.00004, and the dashboard's fleet MTBF lands at 2.69 %, 0.07 percentage points away on this record. What the dashboard could not give is the band, the assurance and the ages it speaks for: the age band alone moves the answer between 2.08 and 3.53 %, the question of Chapter 16. Before the October shutdown the probability of a failure is 8.15 %, and the coil's own failure probability passes 5 % in window 42 (Jun Y4), with the band from April to November of year 4.

15.5.1 The answers and their bands

Table 15.11 gives the answers in the order of Table 15.2. Each row is a question; the second column gives the reported value from the exact posterior of the state, the third the Laplace normal's value beside it, and the fourth the band: the 5–95 % band of a rate, a probability or a window, and for an assurance its range over the priors of Table 15.10.

Table 15.11: The report for coil 0, the hall's coil 1, at window 36 (Dec Y3). Reported values from the exact posterior of the state; bands are 5–95 %, or for an assurance the range over the priors of Table 15.10.

Asset manager's question	Reported value	Laplace normal	Band
Given its rate today, what is the probability it fails in the next 90 days?	$\bar{p}_\tau = 0.0277$ [–]	0.0292 [–]	$p_\tau(\lambda)$: 0.0145 to 0.0441 [–]
What is its failure rate today, and its MTBF?	rate median 0.1097 [yr^{-1}]; MTBF median 9.11 [yr]	0.1139 [yr^{-1}]; 8.78 [yr]	rate 0.0594 to 0.1827 [yr^{-1}]; MTBF 5.47 to 16.83 [yr]
How well is that probability known, and is it below 5 % with 95 % assurance?	median 0.0267 [–]; $\Pr(p_\tau < 0.05) = 0.9827$ [–], met; level at 95 % assurance 0.0441 [–]	0.0277; 0.9646; 0.0474 [–]	assurance 0.9419 to 0.9995 over the priors [–]
Will it run to the October shutdown, and in which window does its failure probability pass 5 %?	$\bar{p}_{\tau_{\mathrm{sd}}} = 0.0815$ [–]; 5 % passed in window 42 (Jun Y4)	0.0859 [–]; window 42	$p_{\tau_{\mathrm{sd}}}(\lambda)$: 0.0436 to 0.1281 [–]; window 40 (Apr Y4) to 47 (Nov Y4)
It has never failed: is its record evidence that it is better than its siblings?	$\ln B_{\mathrm{own,sh}} = 0.0193$ [–]; three clean years $P = 0.8000$ [–]	−0.0617 [–]	–

15.5.2 The entry in the maintenance log

The assessment goes in the two-unit hall's maintenance log, under coil 1, at window 36. It uses nothing later than window 36, and in the values of Table 15.11 it reads:

> *Window 36 (Dec Y3), coil 1 (fleet coil 0), age 3; fleet failure log, ages 2–4: 8 failures in 69 coil-years (coil 1: none in its year, none in three).*
> Failure rate today: 0.110 yr^{-1} (5–95 % band 0.059 to 0.183); MTBF 9.1 years (5.5 to 16.8).
> Failure in the next 90 days (windows 37–39, Jan–Mar Y4): 2.77 %; the rate's band gives 1.45 % to 4.41 %; below 5 % with assurance 0.98.
> Before the October shutdown (window 46): 8.15 %; 5 % is passed in window 42 (Jun Y4),

band window 40 (Apr Y4) to window 47 (Nov Y4).
The clean record: no evidence that coil 1 differs from its siblings ($\ln B = 0.02$).
Action: coil 1 stays on the leak-repair contingency through the summer quarter; no change to the shutdown plan.
Next measurement: the fleet's failure log for year 4 (9.0 % off the rate's standard deviation); coil 1's own log adds 0.6 %.

The entry records what the numbers rest on: the band of ages 2 to 4, the constant rate within it and the fleet's log. Whether the contingency is worth its cost against a failure's consequence is not settled by a probability; it is Chapter 17's question.

Common mistake

Reading a clean record as a zero rate. A unit's own zero in a few unit-years is a count row with a small expected count, and it carries almost no information about a rate near one failure in ten unit-years; coil 1's three clean years are what four coils in five of its design show. Its failure probability is read from the coils that share its rate, with the band that their record leaves.

15.6 The next measurement (step 7)

The report leaves the coil's own 90-day failure probability anywhere from 1.45 to 4.41 %, a band set by eight failures in the band. A further record of failures is one more count row on ω, and what it takes off the band is known before its count is.

15.6.1 The preposterior of a longer log

A candidate log of ΔE further coil-years in the band is a count row of exposure ΔE on the one entry. Its information is its expected count, $\Delta E\,\hat{\lambda}$, whatever count it shows, so the preposterior variance of equation (4.65) is, on the Laplace route,

$$\sigma^2_{\omega,a} = \frac{1}{\Lambda + \Delta E\,\hat{\lambda}}. \tag{15.24}$$

The exact route averages over the count the log will show. Under the present posterior the count $\tilde{k}$ in ΔE coil-years has the negative binomial predictive, $\Pr(\tilde{k} \mid \mathbf{y}) = \frac{\Gamma(a'_\lambda+\tilde{k})}{\Gamma(a'_\lambda)\,\tilde{k}!}\left(\frac{b'_\lambda}{b'_\lambda+\Delta E}\right)^{a'_\lambda}\left(\frac{\Delta E}{b'_\lambda+\Delta E}\right)^{\tilde{k}}$, and each count gives the posterior Gamma$(a'_\lambda+\tilde{k},\ b'_\lambda+\Delta E)$, of standard deviation $\sqrt{a'_\lambda+\tilde{k}}/(b'_\lambda+\Delta E)$. The expected standard deviation of the rate after the log is the sum of these over $\Pr(\tilde{k} \mid \mathbf{y})$, and the expected width of the band of $p_\tau(\lambda)$ is formed the same way. Today they are 0.03797 yr^{-1} and 0.02951. Nine coils of the fleet are aged two today and seven aged three, so the fleet's log of next year adds 16 coil-years to the band.

15.6.2 The candidates

Table 15.12 ranks the candidates. Each row is a record the asset manager can collect; the columns give the coil-years it adds to the band, the standard deviation of the log-rate after it by equation (15.24) and the share it takes off, the expected standard deviation of the rate on the exact route and the share it takes off, and the expected width of the band of $p_\tau(\lambda)$. The exact shares are drawn as bars in Figure 15.6, ranked, with the floor of the rate known exactly.

No measurement moves the 90-day answer; each only narrows its band. The expected predictive after any candidate is today's, 0.027657, because today's predictive is the average of the predictives the possible counts would give; with the rate known exactly the coil's 90-day failure probability would still be 2.77 % at the posterior mean, and 0.0206 at the truth's rate.

Table 15.12: Candidate records for coil 0's rate, from the posterior of the state; the shares are taken from the present $\sigma_\omega = 0.33333$ and $\mathrm{sd}(\lambda \mid \mathbf{y}) = 0.03797\,\mathrm{yr}^{-1}$, and the present band width of $p_\tau(\lambda)$ is 0.02951.

Candidate	ΔE [coil-yr]	$\sigma_{\omega,a}$ [–]	off [%]	sd λ, exact [yr^{-1}]	off [%]	width of p_τ band [–]
Coil 1's own work-order log, next year (the permanent instrument at the place of the question)	1	0.33124	0.6	0.03773	0.6	0.02932
The fleet's log, next year	16	0.30397	8.8	0.03455	9.0	0.02690
The fleet's log, next two years	28	0.28642	14.1	0.03252	14.4	0.02535
The fleet's log, next three years	31	0.28249	15.3	0.03206	15.6	0.02500
The maker's field log of the design, as many coil-years in the band	69	0.24354	26.9	0.02757	27.4	0.02155
The hall's condition monitoring of coil 1	0	0.33333	0.0	0.03797	0.0	–
The rate known exactly (the floor)	–	–	–	0	100	0

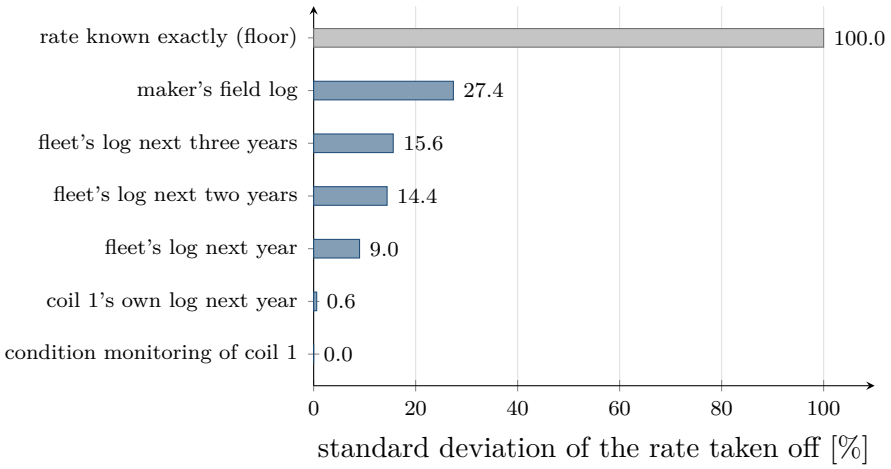


Figure 15.6: The candidate records of Table 15.12, ranked by the share of the rate's posterior standard deviation they take off on the exact route, with the floor (gray), the rate known exactly.

A failure in a window is aleatory, and a record narrows only what is known of the rate. The permanent instrument at the place of the question, coil 1's own work-order log, takes off least: it adds one coil-year a year against the band's 69, while the fleet's next year adds 16, two years 28 and three years 31, because few coils are young enough to enter the band. The condition monitoring of coil 1 takes off nothing, because no row of this model ties the conductance to the failures.

15.6.3 The recommendation

Carry the fleet's year-4 failure log into the band when it closes, which takes 9.0 % off the rate's standard deviation; seek a second source of coil-years in the band, the maker's field log of the design, which would take 27.4 % off, before narrowing the band, which Chapter 16 decides. Coil 1's own log is kept as it always is, and it adds little to this question.

15.6.4 What the analysis showed

The screen offered two readings, coil 1's own zero and the fleet's 2.69 %; read from the coils that share its rate at its age, coil 1 has a 2.77 % probability of a failure in the next 90 days, and its own failure probability lies between 1.45 and 4.41 %. The zero was never a rate: a clean record is a count row whose information is its expected count, and three clean years carry almost none at a rate near 0.1 a year. For a 90-day window the mission probability is a posterior expectation that equals the plug-in at the posterior mean to within 4×10^{-5}, so the error of a quick answer lies in the summary chosen, the route taken and the coils pooled, not in the averaging. The constant rate in an age band is a choice whose answer moves with the band, and this record cannot settle it.

15.6.5 Limits of the model

A constructed record. The failure log is generated from a stated truth, and step 5 compares the answers with that truth (Section 15.4.6). It is not field data. This record's band holds more failures than the truth expects, and the check against fresh logs shows that excess as the record's draw.

A constant rate in the band. The rate is taken as one value across ages 2 to 4, read at the band's mean age of 2.95 years. If the rate changes with age the answer belongs to that mean age, not to the window's; Chapter 16 asks whether it does.

Exchangeable coils. The coils in the band share one rate, with no site, duty or water-quality covariates. Coil 0's own departure from its siblings could be seen only by a record far longer than three years.

A Poisson process in service time. Each failure is repaired and the coil returned to service at the same rate, one leak one count. No row ties the failures to the season, the duty or the coil's condition, although the next 90 days are the summer quarter.

The prior. The prior holds 12.7 % of the information and moves the answer between 2.59 and 3.27 % across the priors of Section 15.4.5.

The next chapter. Chapter 16 reads the failure log across all twelve age bins with a rate that may change with age, and places coil 0 on that curve at age 3.

15.7 What to retain

Definition 15.1 (Mission failure probability)**.** The probability that a unit with a constant failure rate fails within a window τ, given its record, is the predictive $\bar{p}_\tau = \mathbb{E}[1 - \mathrm{e}^{-\lambda\tau} \mid \mathbf{y}]$, the mean of the posterior of the unit's own failure probability $p_\tau(\lambda)$. For a gamma posterior $\mathrm{Gamma}(a'_\lambda, b'_\lambda)$ it is $1 - (b'_\lambda/(b'_\lambda + \tau))^{a'_\lambda}$. It lies below the plug-in at the posterior mean rate by at most $\frac{1}{2}\tau^2 \operatorname{Var}(\lambda \mid \mathbf{y})$ to first order, and differs at first order from a plug-in at any other estimate. The band of $p_\tau(\lambda)$ and its assurance against a limit p^* follow from the rate's quantiles, $\Pr(p_\tau(\lambda) < p^* \mid \mathbf{y}) = \Pr(\lambda < -\ln(1-p^*)/\tau \mid \mathbf{y})$. The rate is read from the coils that share it, a choice of an age band whose answer moves with the band.

Three conclusions carry beyond the coil. A probability of a failure in a window is an average over what is not known of the rate, and for a short window that average is the plug-in at the mean, so the care goes into the rate and its summary. A clean record is evidence like any other, with the information of its expected count, and a short one says little. And the coils that speak for a unit are a choice to be named and tested, because the band of ages changes the answer by more than one standard deviation of the unit's own failure probability.

Table 15.13 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 15.13: The terms of the mission failure probability and where each is defined. Units for coil 0's band.

Term	Expression	Units	Meaning	Defined in
Failure probability at a rate	$p_\tau(\lambda) = 1 - \mathrm{e}^{-\lambda\tau}$, eq. (15.1)	–	a derived quantity of the state	15.1.1
Mission failure probability	$\bar{p}_\tau$, eqs. (15.2), (15.15)	–	the predictive, the mean of $p_\tau(\lambda)$	15.1.1
Prior as a count row	$(n_4, E_4) = (a_\lambda, b_\lambda)$, eq. (15.4)	–, yr	a_λ failures in b_λ unit-years	15.1.4
Rows and posterior				
Clean record's information	$\mu_1 = E_1\hat{\lambda}$, eq. (15.9)	–	its expected count, failure or not	15.2.1
Log-gamma posterior	eq. (15.12)	–	the exact posterior of ω	15.2.3
Grid on the state	eq. (15.13)	–	the reported route	15.2.3
Questions and answers				
Plug-in and its gap	eq. (15.16)	–	Jensen for a concave probability	15.3.1
Second-order reading in ω	eq. (15.17)	–	where the Laplace route lands	15.3.1
Predictive hazard	$\lambda_{\mathrm{pred}}(\tau) = a'_\lambda/(b'_\lambda + \tau)$, eq. (15.18)	yr^{-1}	falls by knowledge, not age	15.3.1
Quantile map	eq. (15.19)	yr^{-1}, yr	bands of the rate and the MTBF	15.3.2
Assurance of a failure probability	eq. (15.20)	–	a level on the rate	15.3.3
Inverse window	$\tau_p(\lambda)$, eq. (15.21)	yr	median and inversion band	15.3.4
Bayes factor of a clean year	eq. (15.23)	–	own rate against the shared rate	15.3.5
Age band	Table 15.9	yr	a choice whose answer moves with it	15.4.4
Preposterior of a longer log	eq. (15.24)	–	narrows the band, leaves $\bar{p}_\tau$	15.6.1

15.8 Exercises

Exercise 15.1. *Formulation.* The hall's two coils share the band's rate. Write "at least one of them fails in the next 90 days" in the form of Table 15.2 and evaluate it. (The predictive, equation (2.13), with $\Pr(A \mid \lambda) = 1 - \mathrm{e}^{-2\lambda\tau}$ for two coils independent given the rate: $1-(b'_\lambda/(b'_\lambda+2\tau))^{a'_\lambda} = 0.054468$; treated as independent at the predictive, $1-(1-\bar{p}_\tau)^2 = 0.054550$.)

Exercise 15.2. *Derivation.* Show that the predictive hazard is $a'_\lambda/(b'_\lambda + \tau)$ and evaluate it today and at the shutdown. (From equation (15.15), $\ln(1-\bar{p}_\tau) = a'_\lambda \ln b'_\lambda - a'_\lambda \ln(b'_\lambda + \tau)$, whose derivative with respect to τ, with the sign changed, is $a'_\lambda/(b'_\lambda + \tau)$: 0.11392 and $0.11285\,\mathrm{yr}^{-1}$.)

Exercise 15.3. *Derivation.* Show that $p_\tau(\mathbb{E}[\lambda \mid \mathbf{y}]) - \bar{p}_\tau \le \frac{1}{2}\tau^2 \operatorname{Var}(\lambda \mid \mathbf{y})$ to first order, and evaluate both sides to the shutdown. (With $\chi = \tau/b'_\lambda$, $\ln[(1-\bar{p}_\tau)/(1-p_\tau(\mathbb{E}\lambda))] = a'_\lambda(\chi - \ln(1+\chi)) \le a'_\lambda\chi^2/2 = \frac{1}{2}\tau^2 \operatorname{Var}(\lambda \mid \mathbf{y})$, and $\mathrm{e}^w - 1 \approx w$: the bound is 4.056×10^{-4}, the gap 3.701×10^{-4}.)

Exercise 15.4. *Application.* Coil 1 fails in January of year 4, one failure in 1/12 coil-year. Update the posterior and give the probability of a failure in February to April. (The row adds one failure and 1/12 coil-year: Gamma(10, 79.0833 yr), mean rate $0.1264\,\mathrm{yr}^{-1}$; by equation (15.15), 0.0307.)

Exercise 15.5. *Application.* How many zero-failure coil-years in the band would show the 90-day failure probability below 2 % with 95 % assurance? (A zero-failure demonstration: by equation (15.20) the limit is $\lambda^* = 0.08193\,\mathrm{yr}^{-1}$, and Gamma(9, $79 + \Delta E$) must put 0.95 below it: $\Delta E = 97.2$ coil-years; today the assurance is 0.2052.)

Chapter 16

Age-Dependent Failure Rates Across a Fleet

16.1 The failure chart and the coil's place on the curve

Coil 1 of the two-unit hall is coil 0 of a fleet of fifty coils of the same model (Chapter 9), one to twelve years old at the end of the record. A tube leak is repaired and the coil returned to service, and the fleet's failure log counts the leaks and the coil-years of service in one-year bins of age: 32 failures in 289 coil-years. The fleet's reliability dashboard plots the log as twelve bars, failures per coil-year by year of age, and draws a straight line through them. The line runs from $0.1102\,\mathrm{yr}^{-1}$ at age 0.5 to $0.1174\,\mathrm{yr}^{-1}$ at age 11.5, a slope of $+0.00065\,\mathrm{yr}^{-1}$ per year of age, $+0.57\,\%$ per year of age: flat. The fleet's one rate is $32/289 = 0.1107\,\mathrm{yr}^{-1}$, a fleet MTBF of 9.03 years, and the line puts coil 0, three years old today (window 36, Dec Y3), at $0.1118\,\mathrm{yr}^{-1}$. The conclusion the chart invites is that these coils do not wear out with age: every coil, coil 0 included, is planned at about 0.11 failures a year, and the old coils need no different treatment from the young.

The chart cannot say how much each bar knows. The four bars from age 8 on hold $11.4\,\%$ of the coil-years and $33.3\,\%$ of the line's weight, and the oldest bar is three coil-years with no failure. Nor can it say whether its tallest bar, ages 6 to 7 at $0.3182\,\mathrm{yr}^{-1}$ (7 failures in 22 coil-years), is wear or chance, or where coil 0, with no failure of its own, sits on any curve. The question of this chapter is

whether the fleet's failure rate rises with age, how fast, and where coil 0 sits on that curve at age 3, read from the failure log's counts by age with each bin weighted by its exposure, against the chart's flat line and against one constant rate.

It is answered with the method of Chapter 4, on a state of two entries.

Jargon: rate of occurrence of failures, minimal repair, exposure in coil-years, age bin, log-linear rate curve, rate ratio per year of age, doubling time, the level integrated out, the band's waist, frailty, covariate

The *rate of occurrence of failures* (ROCOF) of a repairable unit is the expected number of its failures per unit of time at a given age; it is not the hazard of a unit's first failure. *Minimal repair* returns a unit to service at the rate of its age, neither renewed nor harmed. *Exposure in coil-years* is the service time a count of failures was observed over, summed across the coils. An *age bin* gathers the coil-years served between two ages, here one year apart. The *log-linear rate curve* (the Cox–Lewis form) makes the log of the rate a straight line in age. Its *rate ratio per year of age* is the factor by which the rate grows in one year of age, and its *doubling time* the age span over which the rate doubles. *The level integrated out* (the conjugate route) removes the curve's level from the posterior in closed form, leaving its slope alone. *The band's waist* is the age at which the band of the curve is narrowest. *Frailty* is the unit-to-unit spread of failure proneness among units of one design at the same age. A *covariate* is a known quantity of each unit, entered in its rows, that may move its rate.

Key idea

A failure log is count rows in exposure. Weighted by its coil-years, each bin speaks with the information it holds, its expected count μ_a, and the rate's slope in age is an entry of the state. A line through raw bars weighs a bar of three coil-years like a bar of fifty.

16.1.1 Age as the clock of a trend

Chapter 4's trend rows tie a parameter across windows in calendar time (Section 4.6.5), and the first line of Chapter 2's trend form, equation (2.21), asks the sign of the slope. Here the clock is the coil's age, and the parameter is the logarithm of the failure rate,

$$\lambda(a) = \lambda_0 \, \mathrm{e}^{\theta_a a}, \qquad \ln \lambda(a) = \ln \lambda_0 + \theta_a \, a, \tag{16.1}$$

with a the age in years, λ_0 the rate at age zero and θ_a the slope of the log-rate in age, both per year. The straight line in $\ln \lambda$ is the limit $\sigma_q \to 0$ of the trend rows (4.68) over the age bins, and the rows that read it are counts, not readings. Under minimal repair the failures of one coil are a Poisson process in its age with the rate (16.1), so that its expected number of failures between ages a_1 and a_2 is $\int_{a_1}^{a_2} \lambda(a) \, \mathrm{d}a$.

Reliability Engineering by Derivation, Chapter 12, read the same form with the calendar year as the clock, for a fleet observed together over six years. Here the clock is each coil's own age, and the coils enter the log at different ages: a coil eight years old today gave its first year to the bin of ages 0 to 1 eight years ago, and a coil commissioned last year has given that bin its only year so far. The log gathers coil-years by age, not by date, so the young bins are full and the old ones thin.

> **Principle 16.1.** *A fleet's failure rate is a function of age with entries of its own, read from counts weighted by exposure. A unit's position is that function at the unit's age, a derived quantity of the same posterior.*

16.1.2 The fleet's failure log as a graph (step 0)

The graph has one fleet node with two free entries, the log-rate at age zero $\ln \lambda_0$ and its slope in age θ_a. Twelve count readings hang on it, one per age bin, each with its exposure E_a pinned: the coil-years the fleet served between ages a and $a+1$, $a = 0, \ldots, 11$. Coil 0's rate at age 3 is the question's node, a derived quantity of the fleet node with no row of its own.

No physics rows. A failure is an event of the log, not a reading of the hall's physics, and the chapter has no conservation row. The condition of each coil enters only the fifth question of Section 16.1.3, as a known input from Chapter 9's fit.

Modeling choices at the graph. Four choices that no question owns make this graph. The failures are taken as a Poisson process in age with minimal repair: a repaired coil returns to service at the rate of its age. This buys counts that are Poisson in their exposure and independent across coil-years, and it fails if a repair renews the tube bundle, which restarts the clock, or if one leak raises the chance of the next. The coils are taken as alike given their age: before the log is read nothing distinguishes one coil's rate from another's at the same age, the exchangeability of Chapter 9. This lets the coils of one age share one rate, and it fails if some coils are failure-prone, a frailty that Section 16.4.5 frees. The rate is taken as a function of age alone, not of the calendar year, the site or the duty: this keeps the state at two entries, and it fails if one manufacturing year or one site's water fails more, which would enter as a known input of each coil's rows. And the log is taken as complete, with each failure's age: every leak is logged with its coil and the coil's age. This lets the log be read by age at all, and it fails if repairs go unlogged; a register that keeps only each coil's total loses most of what the log says of the slope (Section 16.4.2).

The record. Table 16.1 gives the failure log. Each row is an age bin; the columns give its failures n_a, its exposure E_a, the raw rate n_a/E_a that the dashboard's bar shows, and coil 0's share of the bin. The failures are tube leaks, each repaired and the coil returned to service. Behind the bins are 289 coil-year rows, one per coil per year of age, 30 of them with a failure: 28 coils have no failure, 14 one, 6 two and 2 three. Coil 0 has served three coil-years, one in each of the first three bins, with no failure. The exposure falls with age, from 50 coil-years in the first bin to 3 in the last, because every coil has been through its first year and only three through their twelfth.

Table 16.1: The fleet's failure log by age bin: failures, exposure, the raw rate of the dashboard's bars, and coil 0's own coil-years.

Age bin [yr]	Failures n_a [–]	Exposure E_a [coil-yr]	Raw rate n_a/E_a [yr^{-1}]	Coil 0: failures, coil-yr [–, coil-yr]
0–1	2	50	0.0400	0, 1
1–2	4	47	0.0851	0, 1
2–3	3	38	0.0789	0, 1
3–4	5	31	0.1613	–
4–5	3	26	0.1154	–
5–6	2	22	0.0909	–
6–7	7	22	0.3182	–
7–8	3	20	0.1500	–
8–9	1	14	0.0714	–
9–10	1	9	0.1111	–
10–11	1	7	0.1429	–
11–12	0	3	0.0000	–
Total	32	289	0.1107	0, 3

The graph. The rows are drawn in Figure 16.1. The fleet node at the top (green) carries the two free entries. Below it the count rows are black squares, each labeled with its working residual, and each joined to its bin's reading, a dashed box with the bin's failures and pinned exposure; the first, the fourth and the last bins are drawn in full and the other nine as one

dashed box. To the right, coil 0's rate at age 3 is the question's node (red), joined to the fleet node by a dotted edge: it is computed from the two entries and adds no row.

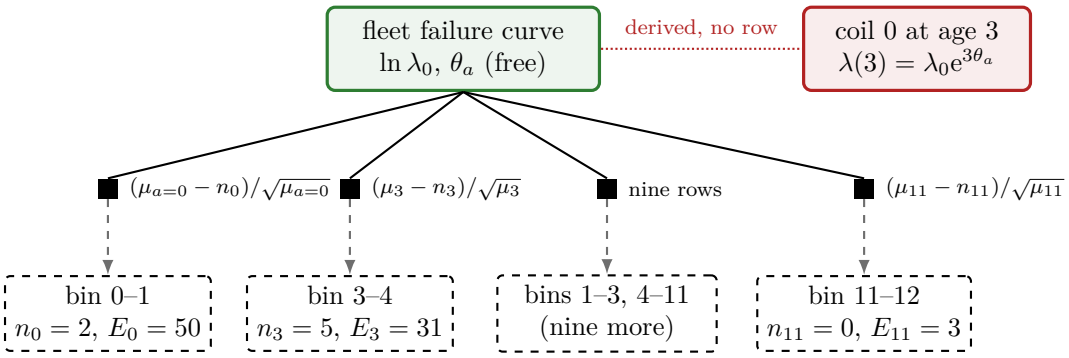


Figure 16.1: The fleet's failure log as rows. Free (green): the log-rate at age zero $\ln \lambda_0$ and the slope θ_a, the two entries of the fleet node. Black squares: the twelve count rows, each with its working residual $(\mu_a - n_a)/\sqrt{\mu_a}$, $\mu_a = E_a \lambda(a + \frac{1}{2})$. Dashed: each bin's reading, its failures n_a with its exposure E_a in coil-years, pinned; the first, fourth and last bins in full, the other nine as one box. Red, dotted: coil 0's rate at age 3, the question's node, a derived quantity that adds no row.

The state. The state of the fleet node is

$$\mathbf{z} = (\ln \lambda_0,\ \theta_a)^{\mathsf{T}}, \tag{16.2}$$

the first entry dimensionless (the logarithm of a rate in yr^{-1}) and the second in yr^{-1}; z_1 and z_2 in that order. Each coil's rate at its own age is a derived quantity of these two entries, and no coil carries an entry of its own until Section 16.4.5 frees one.

16.1.3 The asset manager's questions (step 1)

The asset manager asks five questions of the failure log, and step 1 writes each as a probability over the state (16.2), with the unknown on the left of the bar and the twelve counts on the right. Table 16.2 lists them. The first column is the question in the asset manager's words, the first of them the drill question, in which "this unit" is coil 0 at age 3; the second the general form of its category from Chapter 2; the third its mathematical form for the fleet; the fourth its category; the fifth the computation of Section 2.3.6 that answers it; and the last the subsection where it is answered. The third question is the one Chapter 15 answered under a constant rate, read again on the curve. Which coil the crew visits first is Chapter 20's question, and whether coil 0 is worth restoring is Chapter 17's. Which record would settle the slope asks nothing new of today's evidence; it is the next measurement, Section 16.6.

16.1.4 The evidence as rows (step 2)

Step 2 attaches the log of Table 16.1 to the graph as twelve count rows on the state (16.2). Bin a has the mid-age $\bar{a}_a = a + \frac{1}{2}$, and its count is Poisson in its exposure,

$$n_a \sim \mathrm{Poisson}(\mu_a), \qquad \mu_a = E_a\, \lambda(\bar{a}_a) = E_a \exp(\ln \lambda_0 + \theta_a \bar{a}_a), \qquad a = 0, \ldots, 11. \tag{16.3}$$

Each row enters in the working form of Chapter 4, equation (4.43), with the line $\sqrt{\mu_a}\, \nabla \ln \mu_a^{\mathsf{T}}$:

$$r_a = \frac{\mu_a - n_a}{\sqrt{\mu_a}}, \qquad \mathbf{j}_a = \sqrt{\mu_a}\,(1,\ \bar{a}_a). \tag{16.4}$$

The priors are flat on both entries for the estimates. A normal prior $\theta_a \sim \mathcal{N}(0, s_\theta^2)$ enters only the evidence of the fourth question.

Table 16.2: The questions asked of the fleet's failure log at the end of the record (window 36, Dec Y3). $\mathbf{y}$ is the twelve counts with their exposures; $\lambda(a)$ the curve (16.1); $0.1107\,\mathrm{yr}^{-1}$ the fleet's one rate; $\tau = 0.2466\,\mathrm{yr}$ the next 90 days.

Asset manager's question	General form	Mathematical form	Category	Computation	Answered in
Does the failure rate rise with age, and where is this unit on that curve?	the trend form over age, its first line, (2.21); moments, credible interval, (2.9)–(2.11)	$\Pr(\theta_a > 0 \mid \mathbf{y})$; $p(\lambda(a) \mid \mathbf{y})$ along age; $p(\lambda(3) \mid \mathbf{y})$ and $\Pr(\lambda(3) < 0.1107 \mid \mathbf{y})$	prediction under a threshold; estimation	trend	16.3.1
How fast does it rise, and in how many years does it double?	moments, credible interval; a ratio by inversion, (2.11)	$100(\mathrm{e}^{\theta_a} - 1)$ in %/yr; $t_{\mathrm{dbl}} = \ln 2/\theta_a$, median and inversion band, with $\Pr(\theta_a \le 0 \mid \mathbf{y})$ at $t_{\mathrm{dbl}} = \infty$	estimation	trend	16.3.2
What rate should I plan with for coil 0 over the next year and the next 90 days, and for the oldest coils?	moments; the predictive, (2.13)	$\int_3^4 \lambda(a)\,\mathrm{d}a$; $\mathbb{E}[1 - \mathrm{e}^{-\int_3^{3+\tau} \lambda(a)\,\mathrm{d}a} \mid \mathbf{y}]$; $\lambda(46/12)$; $\lambda(11.5)$	estimation; prediction under a threshold	trend	16.3.3
Is one constant rate good enough, or does the record need the age curve?	$B_{\mathrm{age}} = p(\mathbf{y} \mid M_{\mathrm{age}})/p(\mathbf{y} \mid M_{\mathrm{const}})$, (2.14)	$\ln B$(age curve : constant), $\theta_a \sim \mathcal{N}(0, s_\theta^2)$, over s_θ	hypothesis comparison	evidence	16.3.4
Do the coils in worse condition fail more, once age is counted?	moments; B_{age}, (2.14)	$p(\theta_g \mid \mathbf{y})$ with and without θ_a; $\ln B$ of the age, derate and age-plus-derate drawings	estimation; hypothesis comparison	posterior; evidence	16.3.5

Which row acts on which entry. The log-count Jacobian $\mathbf{X}_a = \partial \ln \boldsymbol{\mu}_a / \partial \mathbf{z}$ has the rows $(1, \bar{a}_a)$, $\bar{a}_a = 0.5, 1.5, \ldots, 11.5$, so every row touches both entries. With $*$ for an entry that can be nonzero, the 12×2 Jacobian and the gradient of the first question's quantity, $\ln \lambda(3)$, are

$$\mathbf{J} = \begin{bmatrix} * & * \\ * & * \\ \vdots & \vdots \\ * & * \end{bmatrix} \begin{matrix} \text{bin 0–1} \\ \text{bin 1–2} \\ \vdots \\ \text{bin 11–12} \end{matrix}, \qquad \mathbf{g}_{\ln \lambda(3)}^{\mathsf{T}} = (1,\ 3). \tag{16.5}$$

The second line is not a row: it is the gradient of $\ln \lambda(3) = \ln \lambda_0 + 3\theta_a$. The pattern is dense, so nothing in the graph separates the level from the slope; their covariance is a matter of the ages the rows sit at. Every answer of step 4 is a function of the same two entries. With the rows in the working form the information is $\mathbf{J}^{\mathsf{T}}\mathbf{J} = \mathbf{X}_a^{\mathsf{T}}\mathbf{W}\mathbf{X}_a$, $\mathbf{W} = \operatorname{diag}(\mu_a)$: with unnormalized rows, the log-counts and their weights, it is $\mathbf{X}_a^{\mathsf{T}}\mathbf{W}\mathbf{X}_a$, the form $\mathbf{J}^{\mathsf{T}}\mathbf{W}\mathbf{J}$ of the notation table.

The rows at the constant rate. At the state of one constant rate, $\mathbf{z} = (\ln 0.1107,\ 0)^{\mathsf{T}}$, every bin's expected count is its exposure times 0.1107. Bin 0–1 has $\mu_{a=0} = 5.5363$, the residual $r_0 = +1.5029$ and the line $\mathbf{j}_0 = (2.3529,\ 1.1765)$; bin 11–12 has $\mu_{11} = 0.3322$, $r_{11} = +0.5764$ and $\mathbf{j}_{11} = (0.5764,\ 6.6280)$. The gradient is $\mathbf{J}^{\mathsf{T}}\mathbf{r} = (0.0000,\ -26.4948)$: the level's column is already balanced, because the constant rate reproduces the total count, and the slope's column says that the late bins hold more failures than a constant rate expects. A negative gradient on θ_a asks for a larger slope.

The score equations. The gradient of the count rows is $\mathbf{J}^{\mathsf{T}}\mathbf{r} = \mathbf{X}_a^{\mathsf{T}}(\boldsymbol{\mu}_a - \mathbf{k})$ by equation (4.42), with $\mathbf{k} = (n_0, \ldots, n_{11})^{\mathsf{T}}$. It vanishes at the MAP state, so the fit satisfies

$$\mathbf{X}_a^{\mathsf{T}}\hat{\boldsymbol{\mu}}_a = \mathbf{X}_a^{\mathsf{T}}\mathbf{k} = \Big(\sum_a n_a,\ \sum_a n_a \bar{a}_a\Big)^{\mathsf{T}} = (N_f,\ N_f a_w)^{\mathsf{T}} = (32,\ 153.0)^{\mathsf{T}}, \tag{16.6}$$

with N_f the log's total failures and a_w the failure-weighted mean age. The fit reproduces the total count and the count-weighted sum of the mid-ages, and the two numbers are all the log contributes to it: they are the sufficient statistic of the curve. Their ratio, $a_w = 4.7812\,\text{yr}$, is later than the exposure-weighted mean age, $3.9533\,\text{yr}$: the failures sit older than the coil-years.

The coil-year rows collapse into the bins. Behind each bin are its coil-years, each a count row of exposure one at the bin's mid-age. The coil-years of one bin share their age, so their terms $\lambda(\bar{a}_a) - n_{ia} \ln \lambda(\bar{a}_a)$ sum to $E_a\lambda(\bar{a}_a) - n_a \ln \lambda(\bar{a}_a)$, the bin's term $\mu_a - n_a \ln \mu_a$ of Section 4.2.7 plus $n_a \ln E_a$, which is free of the state. The 289 coil-year rows and the 12 bins give the same objective up to a constant; Section 16.3.5 shows it numerically.

Modeling choices at the rows. Four choices make these rows. The curve is taken as log-linear in age, the Cox–Lewis form, the Gompertz hazard of a non-repairable unit: this buys two entries on the canonical link of the count, so that the curvature term vanishes and the level integrates out in closed form, and it fails if the rate bends in age, which Section 16.4.4 tests with a quadratic and a power law beside it. The priors are flat on both entries: the record decides, and the posterior is proper because the log has failures in bins of different ages. The evidence's prior on the slope is $s_\theta = 0.10\,\text{yr}^{-1}$, a rate ratio of $\mathrm{e}^{0.10}$ per year of age at one standard deviation; the fourth question gives the evidence as a curve over s_θ, because the evidence depends on it. And the rows are read at the mid-age: a bin's mean count is the integral of the rate over its year, $E_a\lambda_0 \int_a^{a+1} \mathrm{e}^{\theta_a x}\,\mathrm{d}x = \mu_a \sinh(\theta_a/2)/(\theta_a/2)$, and the factor is the same for every bin, so it moves only $\ln \lambda_0$ (Section 16.4.3).

16.2 The posterior of the state (step 3)

All five questions of Table 16.2 are read from one posterior, that of the state (16.2) under the twelve count rows of step 2, derived two ways: the Laplace normal at the MAP state, and the exact posterior with the level integrated out.

16.2.1 The count rows on the two entries

The residual vector stacks the twelve working residuals (16.4) in the order of the pattern (16.5), and the objective is the sum of the count rows' terms of equation (4.42),

$$\mathbf{r}(\mathbf{z}) = \begin{bmatrix} (\mu_{a=0} - n_0)/\sqrt{\mu_{a=0}} \\ \vdots \\ (\mu_{11} - n_{11})/\sqrt{\mu_{11}} \end{bmatrix}, \qquad \mathcal{L}(\mathbf{z}) = \sum_{a=0}^{11} (\mu_a - n_a \ln \mu_a), \tag{16.7}$$

twelve dimensionless rows on two entries, with μ_a of equation (16.3). As Chapter 4 warns, the half square of the working residuals is not the objective; the fit uses only their products. The gradient is $\nabla\mathcal{L} = \mathbf{J}^\mathsf{T}\mathbf{r} = \mathbf{X}_a^\mathsf{T}(\boldsymbol{\mu}_a - \mathbf{k})$, and the Hessian is $\mathbf{H} = \mathbf{J}^\mathsf{T}\mathbf{J} + \mathbf{C}$, equation (4.44), with the curvature term $\mathbf{C} = \sum_a (\mu_a - n_a)\nabla^2 \ln \mu_a$. Each $\ln \mu_a = \ln E_a + \ln \lambda_0 + \theta_a \bar{a}_a$ is linear in the state, the log-rate being the count's canonical parameter, so every $\nabla^2 \ln \mu_a$ is zero and

$$\mathbf{H} = \mathbf{J}^\mathsf{T}\mathbf{J} = \mathbf{X}_a^\mathsf{T}\mathbf{W}\mathbf{X}_a = \begin{bmatrix} \sum_a \mu_a & \sum_a \mu_a \bar{a}_a \\ \sum_a \mu_a \bar{a}_a & \sum_a \mu_a \bar{a}_a^2 \end{bmatrix} \tag{16.8}$$

exactly, at every state. Newton's step and Gauss–Newton's are then the same step, the normal equations

$$(\mathbf{J}^\mathsf{T}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^\mathsf{T}\mathbf{r}(\mathbf{z}), \qquad \mathbf{z}^{(l+1)} = \mathbf{z}^{(l)} + \Delta\mathbf{z}, \tag{16.9}$$

with $\mathbf{J}$ and $\mathbf{r}$ at $\mathbf{z}^{(l)}$, and at the minimum the Laplace covariance is

$$\boldsymbol{\Sigma} = (\mathbf{J}^\mathsf{T}\mathbf{J})^{-1}\big|_{\hat{\mathbf{z}}}, \tag{16.10}$$

the form of equation (4.38) with nothing dropped. At the MAP state the score equations (16.6) fix the first row of the information (16.8): $\sum_a \hat{\mu}_a = N_f$ and $\sum_a \hat{\mu}_a \bar{a}_a = N_f a_w$.

16.2.2 The fit of the failure log

The iterations. Table 16.3 runs the normal equations (16.9) from the constant rate, $\mathbf{z}^{(0)} = (\ln 0.1107,\ 0)^\mathsf{T}$. Each row is an iteration, with the state, the objective and the largest entry of its gradient. The first step takes the slope from zero to 0.097363 and the objective from -2.248286 to -3.437350; the third step reaches the minimum objective to six decimals, and the gradient is at round-off, 2.66×10^{-14}, from the fifth.

Table 16.3: Newton iterations of the failure curve, equation (16.9), from the fleet's one rate. Newton and Gauss–Newton are the same iteration on the canonical link.

Iteration l [–]	$\ln \lambda_0$ [–]	θ_a [yr^{-1}]	Objective $\mathcal{L}(\mathbf{z}^{(l)})$ [–]	$\max\lvert\nabla\mathcal{L}\rvert$ [–]
0	-2.200691	$+0.000000$	-2.248286	2.65×10^{1}
1	-2.585595	$+0.097363$	-3.437350	8.86×10^{0}
2	-2.595702	$+0.090857$	-3.473887	2.63×10^{-1}
3	-2.595345	$+0.090553$	-3.473921	2.78×10^{-4}
4	-2.595344	$+0.090553$	-3.473921	3.37×10^{-10}
5	-2.595344	$+0.090553$	-3.473921	2.66×10^{-14}
6	-2.595344	$+0.090553$	-3.473921	2.66×10^{-14}

The MAP state and the information. The iteration stops at

$$\hat{\mathbf{z}} = \big(\widehat{\ln\lambda_0},\ \hat{\theta}_a\big)^{\mathsf{T}} = (-2.59534,\ 0.09055)^{\mathsf{T}}, \qquad \hat{\lambda}_0 = 0.07462\,\mathrm{yr}^{-1},$$

and there the information (16.8) is

$$\mathbf{J}^{\mathsf{T}}\mathbf{J} = \begin{bmatrix} 32.0000 & 153.0000 \\ 153.0000 & 1041.5510 \end{bmatrix},$$

its first row the sufficient statistic $(N_f, N_f a_w)$ of equation (16.6): the fitted counts sum to 32.000000 and their age-weighted sum is 153.000000. The fitted rates run from $0.0781\,\mathrm{yr}^{-1}$ in the first bin to $0.2114\,\mathrm{yr}^{-1}$ in the last.

The covariance. Its inverse, the Laplace covariance (16.10), is

$$\boldsymbol{\Sigma} = \begin{bmatrix} 0.104988 & -0.015422 \\ -0.015422 & 0.003226 \end{bmatrix},$$

with the spreads $\sqrt{\Sigma_{11}} = 0.32402$ and $\sqrt{\Sigma_{22}} = 0.05679\,\mathrm{yr}^{-1}$ and the correlation -0.8381. The correlation is negative because every coil-year lies at a positive age: the rows see the rate where the coil-years are, around age 4, and a steeper slope must start lower at age zero to pass through the same counts there.

The change of origin. The first entry is the log-rate at age zero, where no coil-year lies, and its spread is the widest of the curve's. Written as the log-rate at another age a_o, the state is $(\ln\lambda(a_\mathrm{o}),\ \theta_a)^{\mathsf{T}} = \mathbf{A}\mathbf{z}$ with $\mathbf{A} = \left[\begin{smallmatrix} 1 & a_\mathrm{o} \\ 0 & 1 \end{smallmatrix}\right]$, the rows of $\mathbf{X}_a$ become $(1,\ \bar{a}_a - a_\mathrm{o})$, and the covariance becomes $\mathbf{A}\boldsymbol{\Sigma}\mathbf{A}^{\mathsf{T}}$: the same fit, with the first entry moved. Table 16.4 gives four origins. Each row is an origin; the columns give the first entry, its spread, and its correlation with the slope, which keeps its own spread 0.05679 at every origin.

Table 16.4: The first entry of the state at four origins of age: the same fit, with $\ln\lambda_0$ replaced by $\ln\lambda(a_\mathrm{o}) = \ln\lambda_0 + a_\mathrm{o}\theta_a$.

Origin a_o [yr]	$\ln\lambda(a_\mathrm{o})$ [–]	sd [–]	Correlation with θ_a [–]
0 (λ_0)	−2.59534	0.32402	−0.8381
3, coil 0's age	−2.32369	0.20368	−0.4967
3.9533, the exposure-weighted mean age	−2.23736	0.18292	−0.2571
4.7812, the failure-weighted mean age a_w	−2.16239	0.17678	−0.0000

The correlation vanishes at the failure-weighted mean age, and there the first entry's spread is the smallest. In $\mathbf{A}\boldsymbol{\Sigma}\mathbf{A}^{\mathsf{T}}$ the covariance of the moved entry with the slope is $\Sigma_{12} + a_\mathrm{o}\Sigma_{22}$, zero at $a_\mathrm{o} = -\Sigma_{12}/\Sigma_{22}$. The information shows which age that is. At origin a_o the off-diagonal entry of equation (16.8) at the MAP state is $\sum_a \hat{\mu}_a(\bar{a}_a - a_\mathrm{o}) = N_f(a_w - a_\mathrm{o})$ by the score equations (16.6), which is zero at $a_\mathrm{o} = a_w$; the information is then diagonal, and the first entry's variance is the reciprocal of its diagonal entry, $1/N_f$:

$$\mathrm{sd}\big(\ln\lambda(a_w) \mid \mathbf{y}\big) = \frac{1}{\sqrt{N_f}} = 0.17678, \qquad a_w = -\frac{\Sigma_{12}}{\Sigma_{22}} = 4.7812\,\mathrm{yr}. \tag{16.11}$$

At the age where the failures sit, the curve is known as well as 32 failures can know one rate, and its spread there does not depend on how the failures are spread in age. With the origin at coil 0's age, coil 0's position is an entry of the state.

16.2.3 The level integrated out

The Laplace normal is the reported posterior. Beside it stands the exact posterior of the same rows, which this state admits in closed form up to one quadrature, because the level integrates out. Write $w = \lambda_0 = \mathrm{e}^{z_1}$. With a flat prior on $z_1 = \ln \lambda_0$ the posterior is the likelihood, and its integral over the level, with the change-of-variable factor $\mathrm{d}z_1 = \mathrm{d}w/w$, is

$$\begin{aligned}\int \prod_a \frac{\mu_a^{n_a} \mathrm{e}^{-\mu_a}}{n_a!}\,\mathrm{d}z_1 &\propto \mathrm{e}^{\theta_a \sum_a n_a \bar{a}_a} \int_0^\infty w^{N_f} \exp\Bigl(-w \sum_a E_a \mathrm{e}^{\theta_a \bar{a}_a}\Bigr) \frac{\mathrm{d}w}{w} \\ &= \Gamma(N_f)\, \mathrm{e}^{\theta_a N_f a_w}\, \mathcal{E}(\theta_a)^{-N_f}, \qquad \mathcal{E}(\theta_a) = \sum_a E_a\, \mathrm{e}^{\theta_a \bar{a}_a},\end{aligned}$$

by the gamma integral, with $\mathcal{E}(\theta_a)$ the age-weighted exposure in coil-years. Its two parts are the exact posterior of the slope and the conditional of the level,

$$p(\theta_a \mid \mathbf{y}) \propto p(\mathbf{y} \mid \theta_a) \propto \mathrm{e}^{\theta_a N_f a_w}\, \mathcal{E}(\theta_a)^{-N_f}, \qquad \lambda_0 \mid \theta_a, \mathbf{y} \sim \mathrm{Gamma}\bigl(N_f,\ \mathcal{E}(\theta_a)\bigr), \tag{16.12}$$

with $N_f = 32$ and $N_f a_w = 153.0$ from the score equations (16.6), and the gamma written by shape and rate. $p(\mathbf{y} \mid \theta_a)$ is the likelihood of the slope with the level integrated out, up to a constant that does not depend on θ_a. The conditional is the gamma update of Chapter 2, equation (2.39), with N_f failures in an exposure of $\mathcal{E}(\theta_a)$ coil-years, in the limit of a flat prior on the log-rate, the gamma prior's shape and rate both taken to zero. At $\theta_a = 0$ the age-weighted exposure is the plain exposure, 289 coil-years, and the conditional is the constant rate's $\mathrm{Gamma}(32, 289)$. Every rate on the curve is a mixture of gammas over the slope: given θ_a, $\lambda(a) = \lambda_0 \mathrm{e}^{\theta_a a}$ is $\mathrm{Gamma}(N_f,\ \mathcal{E}(\theta_a)\mathrm{e}^{-\theta_a a})$, so

$$\Pr(\lambda(a) \le x \mid \mathbf{y}) = \sum_m \varpi_m\, \Gamma_{\mathrm{reg}}\bigl(N_f,\ \mathcal{E}(\theta_m)\, \mathrm{e}^{-\theta_m a}\, x\bigr), \tag{16.13}$$

with θ_m the nodes of a grid in the slope, ϖ_m their weights, the normalized values of $p(\mathbf{y} \mid \theta_m)$, and $\Gamma_{\mathrm{reg}}(k, y)$ the regularized incomplete gamma function. A quantile is the root of equation (16.13) in x.

The grid has 14 001 nodes on $[-0.6,\ 0.8]\,\mathrm{yr}^{-1}$, a step of 0.0001, and the weight within 0.02 of its ends is 2.18×10^{-22}. At the MAP slope the age-weighted exposure is $\mathcal{E}(\hat{\theta}_a) = 428.8384$ coil-years, and $N_f/\mathcal{E} = 0.07462\,\mathrm{yr}^{-1}$ is the mode of the conditional of $\ln \lambda_0$ and the MAP's λ_0: the conditional's density in $\ln \lambda_0$ is proportional to $w^{N_f} \mathrm{e}^{-w\mathcal{E}}$, whose maximum is at $w = N_f/\mathcal{E}$.

16.2.4 The two routes side by side

On the slope the two routes give

Laplace:	median 0.09055, sd 0.05679, band $[-0.00287,\ 0.18397]\,\mathrm{yr}^{-1}$,
exact:	median 0.08949, sd 0.05732, band $[-0.00621,\ 0.18227]\,\mathrm{yr}^{-1}$, mode 0.09055,

and on the rate at age zero, Laplace median 0.07462 with band [0.04379, 0.12715] and exact median 0.07366 with band $[0.04181,\ 0.12238]\,\mathrm{yr}^{-1}$. The exact marginal of the slope has its mode at the MAP's slope, 0.09055, because its score, $N_f a_w - N_f\, \mathcal{E}'(\theta_a)/\mathcal{E}(\theta_a)$, is the MAP's second score equation with the level at its conditional mode; it is slightly wider than the Laplace normal and its median sits a little below the mode. The gap in the rates has another cause. Given the slope, the level's conditional is $\mathrm{Gamma}(N_f,\ \mathcal{E}(\theta_a))$, whose median lies below the mode of its logarithm, the value the MAP reads, by the factor 0.9896 for $N_f = 32$; the slope's spread adds to it wherever the rate is read away from θ_a's center. At $\theta_a = 0$ the Laplace median is 0.11073 and the exact $\mathrm{Gamma}(32,\ 289)$ median $0.10958\,\mathrm{yr}^{-1}$, the gamma's factor alone; at the waist, where the level is uncorrelated with the slope, they are 0.11505 and $0.11208\,\mathrm{yr}^{-1}$.

Every exact median of a rate sits below its Laplace one for this reason. The constant-rate drawing, $\theta_a = 0$, has $\lambda = 0.11073\,\mathrm{yr}^{-1}$ with a spread of its log of $0.17678 = 1/\sqrt{N_f}$, and its exact posterior is Gamma(32, 289). The log-likelihoods at the two MAP states are -19.7784 for the age curve and -21.0040 for the constant rate.

The route of each answer. Table 16.5 says which route each answer uses. Each row is a kind of answer, and the columns give the reported route and the exact route beside it. A derived quantity is carried on the Laplace route by its gradient, $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$, or, when it is not linear in the state, by a Gauss–Hermite rule on the Laplace normal: 60×60 nodes $(x_j, x_{j'})$ with weights $w_j w_{j'}$, mapped to states $\hat{\mathbf{z}} + \mathbf{L}(x_j,\ x_{j'})^\mathsf{T}$ through the Cholesky factor $\mathbf{L}$ of $\boldsymbol{\Sigma}$.

Table 16.5: The route of each answer of the chapter.

Answer	Reported: the Laplace normal	Beside it: the level integrated out
The slope and its sign	the entry $\hat{\theta}_a$, $\sqrt{\Sigma_{22}}$; a normal tail	the grid marginal (16.12)
A rate on the curve, $\lambda(a)$	$\ln\lambda(a) = \mathbf{g}(a)^\mathsf{T}\mathbf{z}$, variance $\mathbf{g}(a)^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}(a)$	the mixture of gammas (16.13)
The rate ratio and the doubling time	quantiles of θ_a, inverted	quantiles of the grid marginal, inverted
The mean rate over the year ahead	the delta method on its log	a mixture of gammas
A window's failure probability	Gauss–Hermite nodes on the Laplace normal	the gamma's transform at each node, weighted
The evidence against a constant rate	Savage–Dickey on the Laplace marginal of θ_a	one-dimensional quadrature of $p(\mathbf{y} \mid \theta_a)$
Condition beside age	the Laplace normal of three entries	the level integrated out; grids in the slopes
The next measurement	the expected information at the MAP	simulated outcomes, the exact route rerun

16.3 The answers (step 4)

Every answer is read from the Laplace normal $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of Section 16.2.2, and that value is the one reported. A derived quantity is carried by its gradient, or by the Gauss–Hermite rule of Section 16.2.4 where it is not linear in the state. The exact route of Section 16.2.3 stands beside each value, a median beside a median.

16.3.1 The curve, its sign and coil 0's place on it

Asset manager's question. Does the failure rate rise with age, and where is this unit on that curve? The first half asks the sign of the slope; the second the rate at coil 0's age, set against the fleet's one rate at which the dashboard plans every coil.

General form. The sign of a trend is the first line of the trend form, equation (2.21),

$$\Pr(\gamma > 0 \mid \mathbf{y}) = \int \mathbf{1}[\gamma > 0]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z},$$

with γ the trend entry, and a quantity along the trend is an estimation question, equation (2.9),

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

summarized by its median and its credible interval, equation (2.11), at the levels 0.05 and 0.95.

Specialized form. The first half is the first line of equation (2.21) with the trend entry the slope $\theta_a = z_2$ and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal, so that

$$\Pr(\theta_a > 0 \mid \mathbf{y}) = \Phi\left(\frac{\hat{\theta}_a}{\sqrt{\Sigma_{22}}}\right), \tag{16.14}$$

a normal tail; on the exact route it is $\sum_m \varpi_m \mathbf{1}[\theta_m > 0]$, the weight of the grid's positive nodes. The second half is equations (2.9) and (2.11) with $f(\mathbf{z}) = \ln \lambda(a) = \mathbf{g}(a)^{\mathsf{T}}\mathbf{z}$, $\mathbf{g}(a) = (1,\ a)^{\mathsf{T}}$, the gradient of the pattern's second line (16.5) at $a = 3$. The quantity is linear in the state, so under the normal its band is exact on the log scale, with

$$\widehat{\ln \lambda}(a) = \hat{z}_1 + a\,\hat{z}_2, \qquad \operatorname{Var}(\ln \lambda(a) \mid \mathbf{y}) = \mathbf{g}(a)^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{g}(a) = \Sigma_{11} + 2a\,\Sigma_{12} + a^2\,\Sigma_{22}, \tag{16.15}$$

a quadratic in a whose minimum is the waist (16.11). Its three terms are the budget of the variance at age a: the level's variance, the slope's variance growing with a^2, and the cross term $2a\,\Sigma_{12}$. The cross term is negative at every positive age, by the negative correlation of step 3: a steeper slope comes with a lower level, and in the rate at any positive age the two errors partly cancel. The slope's term grows with a^2 and the cross term only with a, so far from the waist the slope's term wins. Quantiles pass unchanged through the increasing map $\lambda = \mathrm{e}^{\ln \lambda}$ (Section 4.3.3), so the rate has the median $\mathrm{e}^{\widehat{\ln \lambda}(a)}$ and the band $\mathrm{e}^{\widehat{\ln \lambda}(a) \pm 1.6449\,\mathrm{sd}}$, with sd the square root of the variance (16.15). The comparison with the fleet's one rate is the estimation form at $f^* = \ln 0.1107$, $\Pr(\lambda(3) < 0.1107 \mid \mathbf{y}) = \Phi\big((\ln 0.1107 - \widehat{\ln \lambda}(3))/\mathrm{sd}\big)$. On the exact route the same quantities are read from the mixture (16.13): its value at $x = 0.1107$, and its roots at 0.05, 0.5 and 0.95.

One choice enters here: coil 0's position is the curve at its age. Coil 0's own record has no failure, and read alone, with a flat prior on its log-rate, it has no proper posterior: a likelihood $\mathrm{e}^{-3\lambda}$ does not fall to zero as the rate does. It enters the curve as three of the 289 coil-years. This is what gives coil 0 a rate at all, and it fails if coil 0 is more or less failure-prone than its siblings at the same age, which Section 16.4.5 tests by freeing a multiplier of its own. Whether its record agrees with the curve is a predictive check on the same posterior: coil 0's expected count over its three years is $\mu_{0\cdot} = \lambda_0 \sum_{a=0}^{2} \mathrm{e}^{\theta_a \bar{a}_a}$, and the probability of no failure in them is $\mathbb{E}[\mathrm{e}^{-\mu_{0\cdot}} \mid \mathbf{y}]$, at each node of the exact route the gamma's transform $\big(\mathcal{E}(\theta_m)/(\mathcal{E}(\theta_m) + \sum_{a=0}^{2} \mathrm{e}^{\theta_m \bar{a}_a})\big)^{N_f}$.

Building blocks. The MAP state $\hat{\mathbf{z}}$ and the covariance $\boldsymbol{\Sigma}$ of Section 16.2.2; the gradient $\mathbf{g}(a)$ and the variance (16.15); the normal distribution function Φ; and beside them the grid weights ϖ_m and the mixture (16.13).

The sign. By equation (16.14), $\Pr(\theta_a > 0 \mid \mathbf{y}) = \Phi(0.09055/0.05679) = 0.9446$; the exact route gives 0.9383.

The curve. Table 16.6 gives the curve at eight ages. Each row is an age, the columns the median and the 5–95 % band on each route and the stated truth of the record, a check marker for Section 16.4.7. The rate rises from $0.0781\,\mathrm{yr}^{-1}$ at age 0.5 to $0.2114\,\mathrm{yr}^{-1}$ at age 11.5, and the band is narrow in the middle of the ages and wide at both ends. By equation (16.15) the spread of $\ln \lambda$ is 0.17678 at the waist, $a_w = 4.7812\,\mathrm{yr}$, 0.20368 at age 3, 0.30062 at age 0.5 and 0.42055 at age 11.5. At age 3 the three terms of the budget are $0.104988 - 0.092534 + 0.029030 = 0.041484$: the cross term cancels most of the level's variance. At age 11.5 they are $0.104988 - 0.354715 + 0.426586 = 0.176859$: the slope's term, which grows with the square of the age, carries the band of the oldest coils. Near age zero the level's term is nearly all of it: at age 0.5 the terms are $0.104988 - 0.015422 + 0.000806 = 0.090372$. Coil 0's age is close to the waist; the oldest coils sit far from it.

Table 16.6: The failure curve at eight ages: medians and 5–95 % bands on the Laplace route (reported) and on the exact route, and the stated truth as the record's generator draws it, each bin at its mid-age rate (at age 3, bin 3–4's), a check marker (Section 16.4.7).

Age [yr]	Laplace [yr^{-1}]	band [yr^{-1}]	exact [yr^{-1}]	band [yr^{-1}]	truth [yr^{-1}]
0.5	0.0781	0.0476–0.1280	0.0770	0.0455–0.1231	0.0539
1.5	0.0855	0.0560–0.1304	0.0841	0.0535–0.1254	0.0626
3.0	0.0979	0.0700–0.1369	0.0959	0.0670–0.1318	0.0845
4.5	0.1122	0.0838–0.1502	0.1093	0.0801–0.1449	0.0982
6.5	0.1344	0.0964–0.1874	0.1309	0.0915–0.1795	0.1326
8.5	0.1611	0.1024–0.2534	0.1572	0.0960–0.2402	0.1789
10.5	0.1931	0.1051–0.3548	0.1884	0.0976–0.3338	0.2415
11.5	0.2114	0.1059–0.4222	0.2061	0.0979–0.3961	0.2806

Coil 0 at age 3. At window 36 (Dec Y3) coil 0's rate has median $0.0979\,\mathrm{yr}^{-1}$ and band [0.0700, 0.1369], against 0.0959 and [0.0670, 0.1318] on the exact route. The exact median is 14.3 % below the chart's line at age 3. Against the fleet's one rate, $\Pr(\lambda(3) < 0.1107 \mid \mathbf{y}) = \Phi\big((\ln 0.1107 + 2.32369)/0.20368\big) = 0.7270$, and 0.7646 on the exact route. Coil 0's own record agrees with the curve: on the exact route it expects 0.2590 failures in its three years (0.2571 as the plug-in at the MAP), and the probability of none is 0.7735.

The oldest coils. At age 11.5, $\Pr(\lambda(11.5) > 0.1107 \mid \mathbf{y}) = 0.9379$, and 0.9173 on the exact route.

The record and the fit are set side by side in Figure 16.2. The horizontal axis is the age, the vertical axis the rate. Each bar is a bin's raw rate n_a/E_a, and its width is proportional to the bin's coil-years, a full year for 50; the dashed line is the dashboard's straight line through the bars, the dotted gray line the fleet's one rate, the solid curve the median of the fitted curve with its 5–95 % band shaded, and the thin vertical line coil 0's age. The narrow bars at the right hold few coil-years and scatter widely; the curve passes below the tall bar of ages 6 to 7 and above the empty last bar, which the chart's line weighs as heavily as the first.

The curve is drawn alone in Figure 16.3, with the rate on a logarithmic axis, where the curve is a straight line and its band the image of the normal band of $\ln\lambda$. The shaded band is the Laplace route's, the thin solid lines the exact route's band at the eight ages of Table 16.6; the vertical gray line marks the waist a_w; the red point with its bar is coil 0 at age 3; the dashed line is the fleet's one rate and the dotted curve the stated truth. The band pinches at the waist and opens toward both ends, more toward the old coils, where the log is thin; the exact band lies just below the Laplace band at every age, and coil 0's bar reaches the fleet's one rate.

16.3.2 How fast the rate rises

Asset manager's question. How fast does it rise, and in how many years does it double? The first asks the slope in the units a planner uses, percent per year of age; the second the age span over which the rate doubles.

General form. An estimation question about a derived quantity asks for its posterior distribution, equation (2.9), summarized by its median and its credible interval, equation (2.11),

$$\Pr(f_L \le f(\mathbf{z}) \le f_U \mid \mathbf{y}) = 0.90,$$

at the levels 0.05 and 0.95. A ratio, such as a time, is read by inverting the quantiles of the quantity it is a monotone function of, with its median and its inversion band.

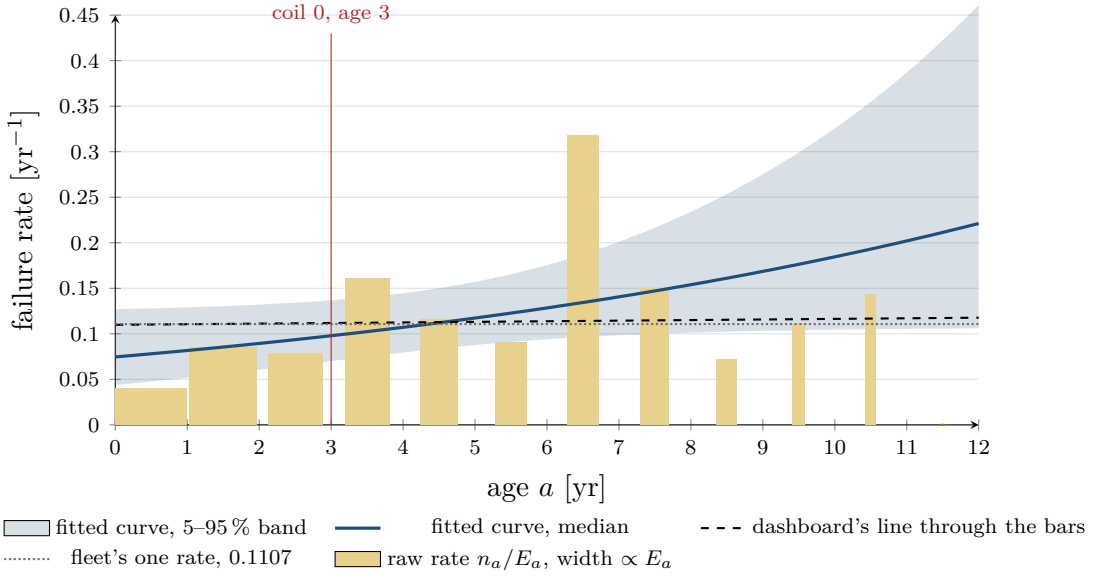


Figure 16.2: The fleet's failure log and the fitted curve. Bars: the raw rate of each age bin, Table 16.1, of width $E_a/50$ yr (the empty last bin, 3 coil-years, is the short base line at age 11.5). Solid and shaded: the curve $\lambda(a) = \exp(-2.59534 + 0.09055a)$ with its 5–95 % band from the variance (16.15). Dashed: the dashboard's straight line through the twelve bars, $0.1099 + 0.00065a$. Dotted: the fleet's one rate.

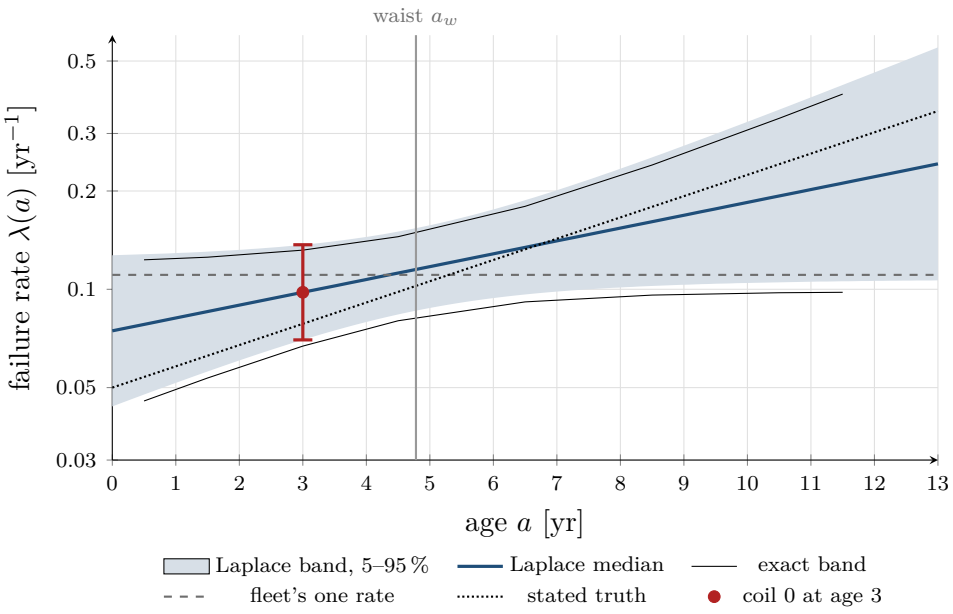


Figure 16.3: The failure curve on a logarithmic axis: the Laplace median $\exp(-2.59534 + 0.09055a)$ with its 5–95 % band (shaded) from equation (16.15), $\Sigma_{11} = 0.104988$, $\Sigma_{12} = -0.015422$, $\Sigma_{22} = 0.003226$; the exact band at the ages of Table 16.6 (thin, joined); the waist $a_w = 4.7812$ yr; coil 0 at age 3 with its band; the fleet's one rate (dashed); and the stated truth $0.05\,\mathrm{e}^{0.15a}$ (dotted), a check marker (Section 16.4.7).

Specialized form. The questions are equations (2.9) and (2.11) with f a function of the slope alone. The rate ratio per year of age is e^{θ_a}, given in percent per year of age as $100(e^{\theta_a} - 1)$, the transform of equation (10.1); it is increasing in θ_a, so its median and band are the transform of the slope's. The doubling time is

$$t_{\text{dbl}} = \frac{\ln 2}{\theta_a}, \quad \theta_a > 0; \qquad t_{\text{dbl}} = \infty, \quad \theta_a \le 0, \tag{16.16}$$

decreasing in θ_a where the slope is positive, with the mass $\Pr(\theta_a \le 0 \mid \mathbf{y})$ at $t_{\text{dbl}} = \infty$: a curve that does not rise never doubles. A quantile of t_{dbl} is the image of the opposite quantile of θ_a. Its median is $\ln 2/\theta_{a,50}$, since less than half the mass is at infinity; its lower end is $\ln 2/\theta_{a,95}$; and its upper end is $\ln 2/\theta_{a,5}$ only if $\theta_{a,5} > 0$, and infinite whenever $\Pr(\theta_a \le 0 \mid \mathbf{y}) > 0.05$. A probability of doubling within t_{lim} years is $\Pr(t_{\text{dbl}} \le t_{\text{lim}} \mid \mathbf{y}) = \Pr(\theta_a \ge \ln 2/t_{\text{lim}} \mid \mathbf{y})$, a normal tail. Beside the inversion stands the delta form, the variance $(\ln 2/\hat{\theta}_a^2)^2\, \Sigma_{22}$ from the gradient $-\ln 2/\theta_a^2$ of equation (4.53).

Building blocks. The slope's entry $\hat{\theta}_a$ and its variance Σ_{22}; the quantiles of the normal; and beside them the quantiles and the tail weights of the grid marginal (16.12).

The rate of rise. The rate rises 9.48 %/yr per year of age, with band $[-0.29,\ 20.20]$ %/yr; the exact route gives 9.36 and $[-0.62,\ 19.99]$.

The doubling time. By equation (16.16) the median is $\ln 2/0.09055 = 7.655$ yr, and the band runs from $\ln 2/0.18397 = 3.768$ yr with no upper end, because $\Pr(\text{never doubles} \mid \mathbf{y}) = 0.0554$ exceeds 0.05. The exact route gives a median of 7.746 yr, a band from 3.803 yr, and a mass at infinity of 0.0617. The doubling time is ten years or less with probability 0.6458 (0.6375 exact) and five years or less with probability 0.1986 (0.1928).

The delta form. The delta form gives 7.655 ± 4.801 yr, a band of $[-0.242,\ 15.552]$ yr. Its lower end is a negative time, and it has no mass at infinity: the linearization of $\ln 2/\theta_a$ about a slope whose band reaches zero cannot carry a function with a pole at zero. The inversion keeps both the order of the times and the curves that never double.

The posterior of the slope is drawn in Figure 16.4. The horizontal axis is the slope θ_a, the vertical axis its density; the solid curve is the exact marginal of equation (16.12), the dashed curve the Laplace normal, and the shaded area under the exact curve left of zero is the mass that never doubles. The marks along the top give the doubling time $t_{\text{dbl}} = \ln 2/\theta_a$ at three slopes, and the dotted line is the stated truth. The two densities nearly coincide; the exact curve's left flank is slightly fuller, which is all that moves its tail to 0.0617.

16.3.3 The rate to plan with

Asset manager's question. What rate should I plan with for coil 0 over the next year and the next 90 days, and for the oldest coils? Chapter 15 answered the 90-day question under a constant rate in coil 0's age band (Section 15.3.1); here it is read again on the curve.

General form. The rate over a span is an estimation question, equation (2.9), summarized by its median and band. The probability of a failure in a window is the tail of the predictive, equation (2.13),

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

with y_{new} the indicator of a failure in the window: the probability given the state, averaged over the posterior.

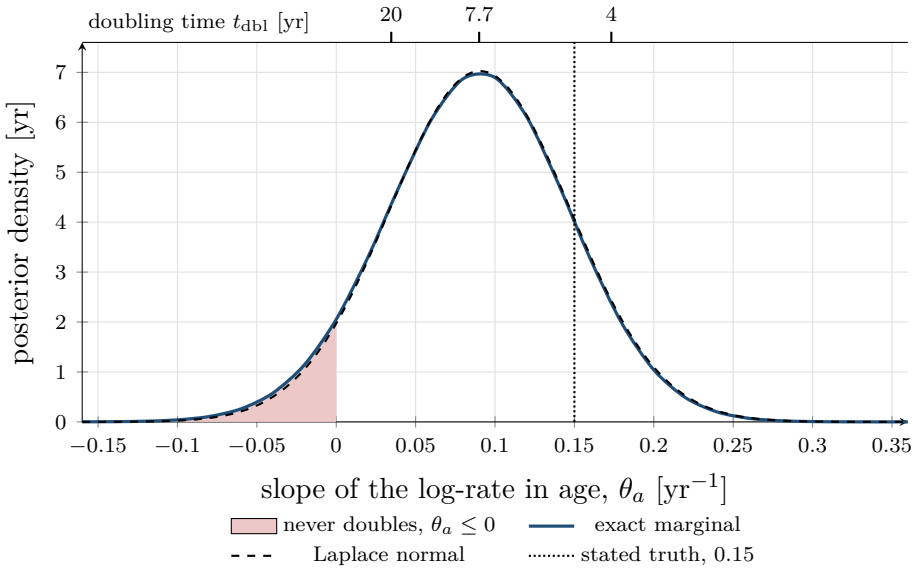


Figure 16.4: The posterior of the slope θ_a: the exact marginal of equation (16.12) (solid) and the Laplace normal $\mathcal{N}(0.09055,\ 0.05679^2)$ (dashed). Shaded: the exact mass at $\theta_a \le 0$, 0.0617, the curves that never double. Top marks: the doubling time $\ln 2/\theta_a$ at 4, 7.7 and 20 years. Dotted: the stated truth, a check marker (Section 16.4.7).

Specialized form. Over a span of ages from 3 to $3+\tau$, coil 0's expected count is the integral of the curve,

$$\int_3^{3+\tau} \lambda(a)\,\mathrm{d}a = \lambda_0\, I_\tau(\theta_a), \qquad I_\tau(\theta_a) = \int_3^{3+\tau} \mathrm{e}^{\theta_a a}\,\mathrm{d}a = \frac{\mathrm{e}^{\theta_a(3+\tau)} - \mathrm{e}^{3\theta_a}}{\theta_a}. \tag{16.17}$$

The first question is equation (2.9) with $f(\mathbf{z}) = \lambda_0 I_1(\theta_a)$, the mean rate over the year of age 3 to 4, windows 37 to 48 (Jan Y4 to Dec Y4). Its logarithm, $\ln\lambda_0 + \ln I_1(\theta_a)$, is nearly linear in the state, and the Laplace route carries it by the delta method, equation (4.53), with the gradient $(1,\ \partial \ln I_1/\partial\theta_a)$ at $\hat{\mathbf{z}}$; on the exact route, given θ_a it is Gamma$(N_f,\ \mathcal{E}(\theta_a)/I_1(\theta_a))$, and its quantiles are those of the mixture (16.13) with $\mathcal{E}(\theta_m)/I_1(\theta_m)$ in place of $\mathcal{E}(\theta_m)\mathrm{e}^{-\theta_m a}$. The second question is equation (2.13) with y_{new} the event of a failure in the next 90 days, τ, windows 37 to 39 (Jan Y4 to Mar Y4). Given the state, the count in the window is Poisson with mean $\lambda_0 I_\tau(\theta_a)$, so

$$\bar{p}_\tau = \mathbb{E}\big[1 - \exp(-\lambda_0 I_\tau(\theta_a)) \mid \mathbf{y}\big]. \tag{16.18}$$

On the Laplace route the expectation is the Gauss–Hermite rule of Section 16.2.4, the integrand evaluated at the 3600 states of the rule. On the exact route the level is averaged in closed form at each node: with $\lambda_0 \mid \theta_m \sim \text{Gamma}(N_f,\ \mathcal{E}(\theta_m))$, the gamma's transform gives $\mathbb{E}[\mathrm{e}^{-\lambda_0 I}] = (\mathcal{E}/(\mathcal{E}+I))^{N_f}$, and

$$\bar{p}_\tau = \sum_m \varpi_m \Big[1 - \Big(\frac{\mathcal{E}(\theta_m)}{\mathcal{E}(\theta_m) + I_\tau(\theta_m)}\Big)^{N_f}\Big],$$

a one-dimensional integral as in Chapter 15, with the age curve inside it. A constant rate with the posterior Gamma(k, E), k failures in E coil-years, gives $1 - (E/(E+\tau))^k$ in the same way. Beside both stands the plug-in $1 - \exp(-\hat{\lambda}(3)\,\tau)$ at the MAP. The rates at the October shutdown of year 4 (window 46, age $46/12 = 3.8333$) and at the oldest bin's mid-age, 11.5, are the curve (16.15) at those ages.

One choice enters here: coil 0 is planned on the fleet's curve at its own age, not at the fleet's one rate nor at a band's rate. This is what lets every bin inform coil 0 while reading the curve

at coil 0's age, and it fails if the curve's form is wrong near age 3 (Section 16.4.4) or if coil 0 differs from its siblings (Section 16.4.5).

Building blocks. The pair $(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$; the span integral (16.17) and the gradient of $\ln \lambda_0 + \ln I_1$; the Gauss–Hermite nodes and weights (x_j, w_j) and the Cholesky factor $\mathbf{L}$; and beside them the grid weights ϖ_m, the age-weighted exposure $\mathcal{E}(\theta_m)$ and the gamma's transform.

The year ahead. Over ages 3 to 4 (windows 37 to 48) the mean rate has median $0.1025\,\mathrm{yr}^{-1}$ and band [0.0749, 0.1403]; the exact route gives 0.1002 and [0.0716, 0.1352].

The next 90 days. Table 16.7 sets the 90-day probability on the curve beside the constant rates a planner could use instead. Each row is a rate model, with its mean rate where it has one and its 90-day failure probability. The curve gives $\bar{p}_\tau = 0.02459$ on the Laplace route and 0.02390 on the exact route, and the plug-in at the MAP gives 0.02385.

Table 16.7: Coil 0's probability of a failure in the next 90 days (windows 37–39, Jan Y4 to Mar Y4, $\tau = 0.2466$ yr) under the age curve and under constant rates. Each constant rate's posterior is Gamma(failures, coil-years), under a flat prior on its log or, for Chapter 15's, under that chapter's prior.

Rate model	Mean rate [yr^{-1}]	90-day failure probability [–]
The age curve, Laplace (reported)	–	**0.02459**
The age curve, exact	–	0.02390
The age curve, plug-in at the MAP, $\hat{\lambda}(3)$	0.0979	0.02385
The fleet's one rate, Gamma(32, 289)	0.1107	0.02692
Bin 2–3, Gamma(3, 38)	0.0789	0.01922
Bin 3–4, Gamma(5, 31)	0.1613	0.03884
Bins 2–4 pooled, Gamma(8, 69)	0.1159	0.02813
Chapter 15's constant rate: bins 2–4 under its prior Gamma(1, 10 yr), posterior Gamma(9, 79 yr)	0.1139	**0.027657**
Stated truth, bin 3–4 as generated	0.0845	0.0206

Against Chapter 15's constant rate in the band of ages 2 to 4, which gives 0.027657 on that chapter's reported exact route, the curve's Laplace 0.02459 is lower by 0.00307, a ratio of 0.8891. Matched route to route, the exact curve's 0.02390 against Chapter 15's exact value is a ratio of 0.8643, and the Laplace curve against that chapter's Laplace route, 0.029209, is 0.842. The band's rate averages ages 2 to 4 and is pulled up by bin 3–4's five failures in 31 coil-years; the curve reads coil 0 at its own age and borrows every bin to do it. The plug-in at the MAP, 0.02385, sits just below both predictives: over 90 days the failure probability is nearly linear in the rate, and averaging over the posterior moves it little.

At the October shutdown. At window 46 (Oct Y4), age 3.8333, coil 0's rate has median $0.1056\,\mathrm{yr}^{-1}$ and band [0.0779, 0.1431]; the exact route gives 0.1031 and [0.0745, 0.1379].

The oldest coils. At age 11.5 the rate has median $0.2114\,\mathrm{yr}^{-1}$ and band [0.1059, 0.4222], 1.91 times the fleet's one rate; the exact route gives 0.2061, [0.0979, 0.3961] and 1.86 times. For fleet coil 2, eleven years old, the 90-day probability on the curve is 0.05290 (0.05079 exact), against 0.02692 at the fleet's one rate.

The block's surprise is how little the curve changes coil 0's plan and how much it changes the oldest coils'. For coil 0 the curve and the constant rates differ little, 0.02459 (Laplace)

against Chapter 15's 0.027657 (exact) and the fleet's 0.02692. The curve, averaged over the fleet's exposure, is the fleet's one rate, 0.110727 against 0.110727, by the first score equation $\sum_a \hat{\mu}_a = N_f$, and coil 0's age lies near the middle of that exposure, whose mean age is 3.9533 years. For an eleven-year-old coil the fleet's one rate gives about half the curve's probability.

16.3.4 The age curve against one constant rate

Asset manager's question. Is one constant rate good enough, or does the record need the age curve? The question is not whether the slope is positive but whether a slope is needed at all.

General form. A hypothesis comparison is the ratio of the evidences of two drawings, equation (2.14),

$$B_{\text{age}} = \frac{p(\mathbf{y} \mid M_{\text{age}})}{p(\mathbf{y} \mid M_{\text{const}})},$$

each evidence the probability of the record with the drawing's state integrated out over its prior.

Specialized form. The question is equation (2.14) with M_{age} the age curve (16.1) under the prior $\theta_a \sim \mathcal{N}(0, s_\theta^2)$ and M_{const} the constant rate, $\theta_a = 0$. Both drawings carry the same flat prior on $\ln \lambda_0$; its improper constant is common to the two evidences and cancels. With the level integrated out, each evidence is the likelihood $p(\mathbf{y} \mid \theta_a)$ of equation (16.12) averaged over the slope's prior, and

$$B_{\text{age}} = \frac{\int p(\mathbf{y} \mid \theta_a)\, \varphi(\theta_a;\ 0,\ s_\theta^2)\, \mathrm{d}\theta_a}{p(\mathbf{y} \mid \theta_a = 0)}, \qquad p(\mathbf{y} \mid \theta_a = 0) \propto 289^{-N_f}, \tag{16.19}$$

with $\varphi(\cdot;\ \mu, \sigma^2)$ the normal density of mean μ and variance σ^2, and the denominator the plain exposure's, since $\mathcal{E}(0) = 289$; numerator and denominator carry the same constant of equation (16.12), which cancels. On the exact route the numerator is a one-dimensional quadrature on the grid of Section 16.2.3. On the Laplace route the drawings are nested at $\theta_a = 0$, and the Savage–Dickey ratio of Chapter 7, equation (7.21), applies: B_{age} is the prior density of the slope at zero over its posterior density there under M_{age}. That posterior is the Laplace marginal $\mathcal{N}(\hat{\theta}_a, \Sigma_{22})$ with the prior applied before its density is read at zero, the correction recorded in *Reliability Engineering by Derivation*, Chapter 12: by the precisions adding (equation (3.29)) it is normal with variance $(1/\Sigma_{22} + 1/s_\theta^2)^{-1}$ and mean $\hat{\theta}_a s_\theta^2/(\Sigma_{22} + s_\theta^2)$, and

$$\ln B_{\text{age}} = \frac{1}{2} \ln \frac{\Sigma_{22}}{\Sigma_{22} + s_\theta^2} + \frac{\hat{\theta}_a^2}{2} \left(\frac{1}{\Sigma_{22}} - \frac{1}{\Sigma_{22} + s_\theta^2} \right). \tag{16.20}$$

The first term is negative, the price of the prior's spread: M_{age} spread its prediction over slopes of size s_θ, most of which the record rules out. The second is positive, the reward for the slope the record shows. As $s_\theta \to 0$ the age curve becomes the constant rate and $\ln B_{\text{age}} \to 0$; as s_θ grows the first term falls without bound and the second approaches $\frac{1}{2}\hat{\theta}_a^2/\Sigma_{22}$, so the evidence has a largest value at a finite spread. The sign probability is read inside M_{age} after its posterior is normalized, and depends on s_θ only through the prior's pull toward zero.

One choice enters here: the spread s_θ of the prior on the slope. The evidence depends on it, and the answer is given as a curve over s_θ with $0.10\,\text{yr}^{-1}$ reported; $\Pr(\theta_a > 0 \mid \mathbf{y})$ under the flat prior does not depend on it.

Building blocks. The slope's entry $\hat{\theta}_a$ and its variance Σ_{22}; the prior's spread s_θ; and beside them the level-integrated likelihood (16.12) on the grid. The likelihood-ratio statistic, twice the difference of the maximized log-likelihoods of step 3, is a remark beside them.

The evidence over the prior's spread. Table 16.8 gives $\ln B_{\text{age}}$ at six spreads. Each row is a spread; the columns give the Savage–Dickey value (16.20), the exact value (16.19), and the probability of a rise under that prior. At the reported spread the reported value, Savage–Dickey on the Laplace marginal, is $\ln B_{\text{age}} = 0.2555$, $B_{\text{age}} = 1.2911$, and at even prior odds the age curve has posterior probability 0.5635; the exact route gives 0.2224, $B_{\text{age}} = 1.2491$ and 0.5554. Over the prior's spread the exact evidence is largest, 0.2760, at $s_\theta = 0.0692\,\text{yr}^{-1}$, and it is negative for every spread above $0.1596\,\text{yr}^{-1}$. The likelihood-ratio statistic is $2(-19.7784 + 21.0040) = 2.4513$ on one degree of freedom, $p = 0.1174$; it is not a probability of either drawing.

Table 16.8: The evidence of the age curve against one constant rate over the spread of the prior on the slope: Savage–Dickey on the Laplace marginal, equation (16.20), and the level-integrated quadrature, equation (16.19), with the probability of a rise under each prior.

s_θ [yr^{-1}]	$\ln B_{\text{age}}$, Laplace (Savage–Dickey) [–]	$\ln B_{\text{age}}$, exact [–]	$\Pr(\theta_a > 0 \mid \mathbf{y})$ under that prior [–]
0.03	0.1542	0.1404	0.7601
0.05	0.2681	0.2449	0.8431
0.10	**0.2555**	**0.2224**	0.9090
0.20	−0.1215	−0.1598	0.9305
0.30	−0.4549	−0.4946	0.9348
1.00	−1.6030	−1.6438	0.9379

The evidence is drawn over the prior's spread in Figure 16.5. The horizontal axis is s_θ on a logarithmic scale, the vertical axis $\ln B_{\text{age}}$; the solid curve is the exact route, the dashed curve Savage–Dickey, the gray line zero, the vertical mark the reported spread, and the dot the exact maximum. The two curves run together, the Laplace one above the exact one at every spread of the table, and neither rises far above zero at any spread.

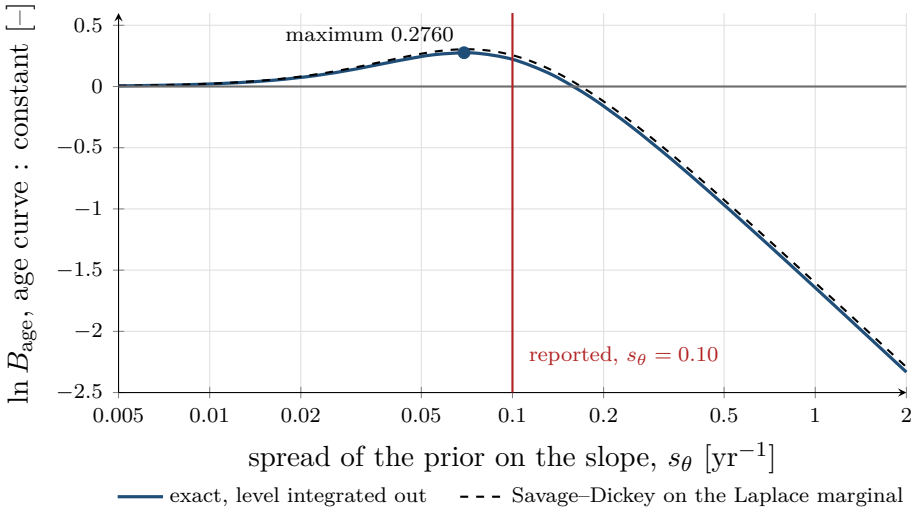


Figure 16.5: The evidence of the age curve against one constant rate over the spread s_θ of the prior on the slope. Solid: the exact quadrature (16.19) at 21 spreads, joined. Dashed: Savage–Dickey, equation (16.20) with $\hat{\theta}_a = 0.09055$ and $\Sigma_{22} = 0.003226$. Gray: no preference. Red: the reported spread; dot: the exact maximum, 0.2760 at $s_\theta = 0.0692$.

The surprise of this block is that the two answers part. The log makes a rise probable, 0.94, and leaves the case for the age curve over a constant rate about even, $B_{\text{age}} = 1.29$ on the Laplace route (1.25 on the exact route). They answer different questions: whether the slope is positive, given that the rate follows a log-linear curve; and whether a slope is needed

at all, against a constant rate that predicted the record nearly as well with one entry fewer. Chapter 15's constant rate is not rejected by this record.

16.3.5 Condition beside age

Asset manager's question. Do the coils in worse condition fail more, once age is counted? Chapter 9 fitted each coil's condition from its annual tests and left the failures alone (Section 9.6.5); the question joins the two.

General form. The effect of condition is an estimation question, equation (2.9), summarized by its median, band and sign, and the need for it a hypothesis comparison, equation (2.14),

$$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y}) = \int \mathbf{1}\left[f(\mathbf{z}) \le f^*\right] p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \qquad B_{AB} = \frac{p(\mathbf{y} \mid M_A)}{p(\mathbf{y} \mid M_B)}.$$

Specialized form. The questions are equations (2.9) and (2.14) on a state of three entries, $(\ln \lambda_0,\ \theta_a,\ \theta_g)^{\mathsf{T}}$, with θ_g the log-rate's coefficient on a coil's derate. The rows return to the 289 coil-years: coil i in its year of age a has

$$n_{ia} \sim \text{Poisson}(\mu_{ia}), \qquad \ln \mu_{ia} = \ln \lambda_0 + \theta_a\, \bar{a}_a + \theta_g\, g_i(\bar{a}_a), \qquad g_i(a) = g_{i,0} + \gamma_i\, a, \qquad (16.21)$$

with $g_i(a)$ coil i's line of equation (9.1), read from Chapter 9's fit at its spreads' medians, 0.03153 for the commissioning values and $0.01199\,\mathrm{yr}^{-1}$ for the rates: the fleet's means are $(-0.00280,\ -0.03022)$ and coil 0's line is $(-0.01292,\ -0.02836)$. The bins no longer suffice, because coils of one age differ in condition. Each row is still on the canonical link, so the curvature term is zero, and the level integrates out as in equation (16.12):

$$p(\theta_a, \theta_g \mid \mathbf{y}) \propto \exp\Big(\theta_a \sum_{ia} n_{ia}\bar{a}_a + \theta_g \sum_{ia} n_{ia} g_i(\bar{a}_a)\Big)\Big(\sum_{ia} \mathrm{e}^{\theta_a \bar{a}_a + \theta_g g_i(\bar{a}_a)}\Big)^{-N_f},$$

read on a grid of 481×521 nodes in (θ_a, θ_g); the drawing with the derate alone, $\theta_a = 0$, is a grid of 1201 nodes in θ_g. The Laplace normal of each drawing, from the normal equations (16.9) on its rows, is reported. Four drawings are compared: the constant rate, age alone, the derate alone, and age and derate together, with the priors $\theta_a \sim \mathcal{N}(0,\ 0.10^2)$ and $\theta_g \sim \mathcal{N}(0,\ 2.5^2)$ in the evidences and the flat level cancelling.

Two choices enter here. The derate is taken as known, Chapter 9's fitted line rather than a reading at each coil-year: this keeps the rows on the curve's form, and it ignores the line's error, which attenuates θ_g toward zero. And the covariate is the coil's condition against its rating, g_i, not its fouling rate γ_i alone, which is the other choice and is not used. The prior's spread 2.5 per unit of g means that a coil at $g = -0.2$, a derate of 0.819, has its rate multiplied by $\mathrm{e}^{-0.2\theta_g}$, a factor of $\mathrm{e}^{0.5}$ at one standard deviation.

Building blocks. Chapter 9's fitted lines as known inputs; the MAP state and covariance of each drawing; the grids of the level-integrated posterior; and the evidences, as in equation (16.19), over the priors.

The covariate. Over the coil-year rows the derate runs from -0.7020 to 0.0264, and its correlation with age is -0.8371: the coils foul as they age. The age drawing on the coil-year rows reproduces the bins' fit of Section 16.2.2, with largest differences of 1.39×10^{-17} in the state and 3.47×10^{-16} in $\mathbf{\Sigma}$, as Section 16.1.4 said.

The derate alone. With the derate and no age the fit gives $\theta_g = -1.9362$ per unit of g, band $[-4.0721,\ 0.1997]$, and $\Pr(\theta_g < 0 \mid \mathbf{y}) = 0.9320$; the exact route gives -1.8286, $[-3.8289,\ 0.4881]$ and 0.9038. Read alone, worse condition means more failures: on the exact median a coil at $g = -0.2$ fails 1.442 times as often as one at its rating.

Age and derate together. With both, $\theta_g = -0.5285$, band $[-4.4666,\ 3.4095]$, and $\Pr(\theta_g < 0 \mid \mathbf{y}) = 0.5874$; exact -0.2759, $[-3.9859,\ 4.0601]$ and 0.5410. The slope in age is 0.0730, band $[-0.0890,\ 0.2350]$, with $\Pr(\theta_a > 0 \mid \mathbf{y}) = 0.7707$; exact 0.0794, $[-0.0869,\ 0.2411]$ and 0.7828. The two slopes are correlated $+0.8177$ in the Laplace normal, positively because the derate falls with age: a larger slope in age is paid for by a less negative coefficient on the derate.

The evidence. Against the constant rate the log-evidences on the exact route, the level-integrated grids, are 0.2224 for age (Section 16.3.4 reports 0.2555 for it on the Laplace route), 0.0035 for the derate and -0.1414 for age and derate together; so age and derate against age alone is -0.3638, and age against the derate 0.2190.

The posteriors of the coefficient are compared in Figure 16.6. In the left panel the horizontal axis is θ_g and the vertical axis its density, under the derate alone and under age and derate together, exact solid and Laplace dashed, with zero marked. In the right panel the horizontal axis is θ_a and the vertical axis θ_g; the ellipses are the contours at one and two standard deviations of the Laplace normal of the two slopes, the gray lines the zeros, and the cross the stated truth. Under the derate alone the density lies mostly below zero; with age beside it, it widens and centers near zero. The tilted ellipses show why: along their long axis the two slopes trade one for the other.

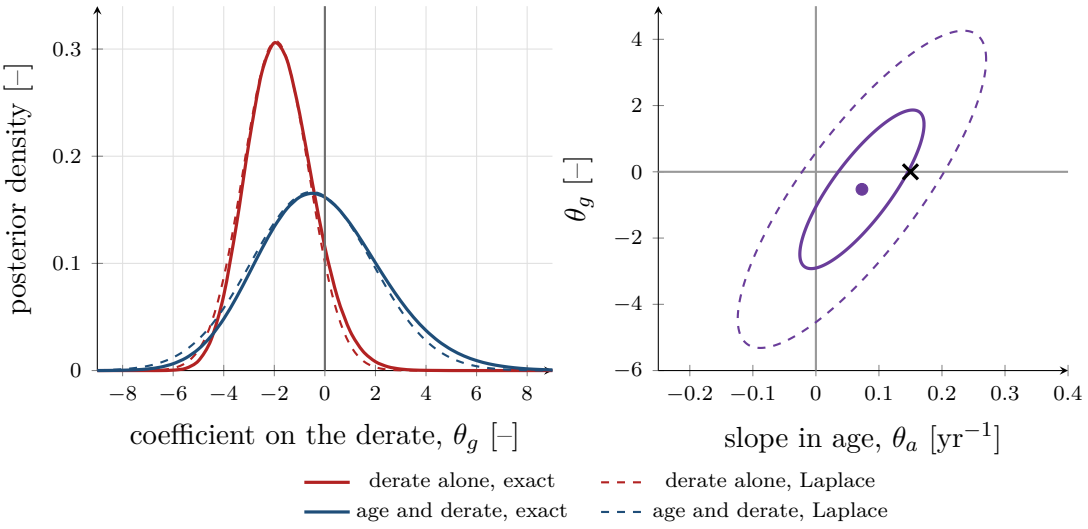


Figure 16.6: The coefficient on the derate. Left: its posterior under the derate alone (red) and under age and derate together (blue), exact (solid, the level and, for the second, the slope in age integrated out) and Laplace (dashed), $\mathcal{N}(-1.93619,\ 1.29852^2)$ and $\mathcal{N}(-0.52855,\ 2.39416^2)$. Right: the Laplace normal of the two slopes under age and derate, means $(0.07302,\ -0.52855)$, standard deviations 0.09851 and 2.39416, correlation 0.8177, contours at one and two standard deviations; dot, the MAP pair; cross, the stated truth $(0.15,\ 0)$, a check marker (Section 16.4.7).

The surprise of this block is that "fouled coils leak more" is the age speaking. Once age is in the rows, the derate's effect has no sign, and this log cannot separate condition from age, because across the fleet they move together. It is the counterpart of the confound of Chapter 13: a regression on condition alone reads age as condition.

16.4 The checks (step 5)

Each answer of Section 16.3 rests on a grid, on the form of the rows and of the curve, and on the coils being alike, and this log lets each be tested against the calculation it stands for.

16.4.1 The grid and the two routes

The exact route has been set beside every answer of step 4; what it rests on is the grid in the slope. From 3501 to 28 001 nodes on $[-0.6,\ 0.8]$ the median of $\lambda(3)$ is 0.095876 and $\ln B_{\text{age}}(s_\theta = 0.10)$ is 0.222422 at every resolution, while the plain sum over the nodes moves $\Pr(\theta_a > 0 \mid \mathbf{y})$ from 0.937855 to 0.938215 and the slope's median from 0.089288 to 0.089463, because it counts each node's weight wholly on one side of the value read. With each node's weight split about it, half on either side, the reading is 0.938266 for the probability and 0.089488 for the median at every resolution, the values quoted in step 4: the corrected grid fixes the tail to four decimals and the median to five. A wider grid, $[-1.5,\ 1.5]$, gives the same $\Pr(\theta_a > 0 \mid \mathbf{y})$, 0.938163, and the same $\ln B_{\text{age}}(s_\theta = 1.0)$, -1.643755. The closed-form integral over the level is itself checked by a grid in both entries, 801×801 nodes over ± 8 standard deviations with flat priors, on which the level is summed rather than integrated: it gives $\Pr(\theta_a > 0 \mid \mathbf{y}) = 0.93776$, a median of $\lambda(3)$ of 0.09588 and a median slope of 0.08949, the one-dimensional route's values within its coarser step. The gaps between the reported Laplace values and the exact ones, 0.9446 against 0.9383 in the probability of a rise and 0.0979 against 0.0959 in coil 0's median rate, are therefore the routes', not the grid's: the gamma conditional's median below the mode of its logarithm, with the slope's spread added (Section 16.2.4). Neither gap would change a plan.

16.4.2 The curvature term and the coils' totals

The chapter has no physics row, so the check of the limit $\sigma_p \to 0$ has no object here; its place is taken by the curvature term of equation (4.44), checked numerically. On the bins the Hessian of the objective by differences of the gradient matches $\mathbf{J}^\mathsf{T}\mathbf{J}$ at the MAP state to 1.78×10^{-8}: $\mathbf{C} = \mathbf{0}$ on the canonical link, as Section 16.2.1 derived.

A register that keeps only each coil's total number of leaks, without their ages, gives fifty rows instead of twelve. Coil i, aged A_i, has the expected total $\mu_{i\cdot} = \lambda_0 \sum_{a<A_i} \mathrm{e}^{\theta_a \bar{a}_a}$, and $\ln \mu_{i\cdot} = \ln \lambda_0 + \ln \sum_{a<A_i} \mathrm{e}^{\theta_a \bar{a}_a}$ is not linear in θ_a, so $\mathbf{C} \neq \mathbf{0}$. At that fit's MAP state, $(-2.17801,\ -0.00577)$, the Hessian is

$$\mathbf{H} = \begin{bmatrix} 32.0 & 124.942 \\ 124.942 & 573.873 \end{bmatrix}, \qquad \mathbf{J}^\mathsf{T}\mathbf{J} = \begin{bmatrix} 32.0 & 124.942 \\ 124.942 & 567.081 \end{bmatrix},$$

and the largest entry of $\mathbf{C} = \mathbf{H} - \mathbf{J}^\mathsf{T}\mathbf{J}$ is 6.7913, in the slope's diagonal. The totals give a spread of 0.10781 on θ_a against 0.05679 from the ages: *the totals keep* 27.8 % *of the information on the slope*, and their MAP slope, -0.00577, barely sees the rise. A register that logs only how many times each coil leaked has lost most of the answer: what tells the slope is the age at which each leak happened.

16.4.3 The bin average and the residuals

The bin average. A bin's mean count is its rate averaged over its year, $\mu_a \sinh(\theta_a/2)/(\theta_a/2)$, against the mid-age row's μ_a. At the MAP slope the factor is 1.000342. It multiplies every bin alike, so the slope is unchanged and $\ln \lambda_0$ moves by -0.000342. The record was drawn at the mid-age rates, so the mid-age row is the truth's own form.

The residuals. The Pearson residuals $(n_a - \hat{\mu}_a)/\sqrt{\hat{\mu}_a}$, the working residuals of equation (16.4) with the sign reversed, give $\chi^2 = 9.4171$ (deviance 8.7998) on $12 - 2 = 10$ degrees of freedom for the age curve, $p = 0.4930$, against 13.1053 (11.2511) on 11 for the constant rate, $p = 0.2865$: neither fit is rejected by its residuals, which is the evidence's verdict in another form. The largest residual is bin 6–7's, +2.351, seven failures against 2.957 expected. The leverages K_{aa}, the diagonal of $\mathbf{J\Sigma J}^{\mathsf{T}}$, run from 0.0889 to 0.3528, the largest in the first bin, which has 50 coil-years at the far end of the age range from the oldest bins.

16.4.4 The shape of the curve

The log-linear curve is a choice of Section 16.1.4; two other shapes stand beside it. A quadratic in age, $\ln\lambda = \ln\lambda_0 + \theta_a a + \theta_{aa} a^2$, has the Laplace MAP $(-3.36258,\ 0.52278,\ -0.04194)$ with spreads $(0.5763,\ 0.24924,\ 0.02378)$ and the maximized log-likelihood -17.9575. With the level integrated out on a grid in (θ_a, θ_{aa}) and flat priors, the curvature has median $-0.04627\,\mathrm{yr}^{-2}$, band $[-0.08850,\ -0.00907]$ and $\Pr(\theta_{aa} < 0 \mid \mathbf{y}) = 0.9809$, and at its MAP the curve peaks at age 6.233 years. With the priors $\theta_a \sim \mathcal{N}(0,\ 0.10^2)$ and $\theta_{aa} \sim \mathcal{N}(0,\ 0.01^2)$ that probability is 0.5714, the curvature's median -0.00138, and $\ln B(\text{quadratic} : \text{log-linear}) = -0.3424$. The likelihood-ratio statistic is 3.6418, $p = 0.0563$. The bump is one record's draw, bin 6–7's seven failures and the shortfall of the old bins: over fresh logs from the truth, which has no curvature, 3.75 % show a Laplace $\Pr(\theta_{aa} < 0 \mid \mathbf{y})$ as high as this record's 0.9611 (Section 16.4.6).

A power law in age, the Crow–AMSAA form, writes the log-rate as $\ln\lambda = c_{\mathrm{p}} + \theta_{\ln a} \ln a$. Its MAP is $(-2.69961,\ 0.42159)$ with spreads $(0.33719,\ 0.21338)$, and its maximized log-likelihood, -18.8156, is 0.9628 above the log-linear curve's on the same two entries. On the exact route its exponent has median 0.4343, band $[0.0967,\ 0.8070]$ and $\Pr(\theta_{\ln a} > 0 \mid \mathbf{y}) = 0.9839$, and coil 0's rate at age 3 has median $0.1037\,\mathrm{yr}^{-1}$ and band $[0.0748,\ 0.1388]$. The verdict on the sign depends on the curve's form: the power law rises steeply over the first years, where the young bins are full, and gives a sharper sign than the log-linear curve's 0.9383. Coil 0's rate moves by less than its band's half-width.

16.4.5 Unit-to-unit spread of the failure rate

The coils were taken as alike given their age. The check frees a gamma frailty: coil i's rate is $Z_i\lambda(a)$, with Z_i gamma of mean one and coefficient of variation ν, so that coil i's total is negative binomial with mean $\mu_{i\cdot}$ of Section 16.4.2 and shape $1/\nu^2$, while the split of each coil's total across its years, a multinomial free of λ_0 and ν, keeps the ages. The posterior is read on a grid of $61 \times 61 \times 101$ nodes in $(\ln\lambda_0, \theta_a, \nu)$ under the prior $\nu \sim$ half-normal(0.5).

- The spread has median $\nu = 0.2359$ and band $[0.0045,\ 0.6655]$, with $\Pr(\nu < 0.2 \mid \mathbf{y}) = 0.4359$. The likelihood is highest at $\nu = 0.0001$, the edge of no spread, and $\ln B(\text{frailty} : \text{none}) = -0.2523$. The per-coil dispersion, $\sum_i (n_i - \hat{\mu}_{i\cdot})^2/\hat{\mu}_{i\cdot}$, is 42.5835 on 48 degrees of freedom. The log shows no spread and cannot rule a modest one out: the posterior, median 0.236 and band $[0.0045,\ 0.666]$, has moved from the prior's median 0.337 and band $[0.031,\ 0.980]$ toward no spread, but not far.
- With the frailty, $\Pr(\theta_a > 0 \mid \mathbf{y}) = 0.9340$ and the slope's median is 0.0885.
- Coil 0's own multiplier, given its three years without a failure and 0.2571 expected at the MAP (a plug-in; averaged over the exact posterior it is 0.2590, Section 16.3.1), is $\mathrm{Gamma}(1/\nu^2,\ 1/\nu^2 + \mu_{0\cdot})$ at each node: its median is 0.9843, its mean 0.9712, and $\Pr(Z_0 < 1 \mid \mathbf{y}) = 0.5655$. Coil 0's own rate at age 3, $Z_0\lambda(3)$, by 200 000 draws of the grid and the gamma (seed 16), has median $0.0914\,\mathrm{yr}^{-1}$ (standard error 0.0001 by batch means) and band $[0.0401,\ 0.1579]$, against 0.0959 for the curve alone on the same grid. Three failure-free years move coil 0 a little below the curve and widen its band; they do not place it.

16.4.6 Fresh failure logs

The answers are one record's draw. The check draws 4000 fresh logs from the stated truth at the same exposures (seed 7) and reads each by the exact route. The truth expects 28.991 failures; the record has 32.

- The slope's median has mean 0.1493 and spread 0.0619 over the logs, and this record's, read on the grid as the logs are, 0.0894, sits at the 16.3rd percentile.
- $\Pr(\theta_a > 0 \mid \mathbf{y})$ has median 0.9938 over the logs; 0.7790 of them exceed 0.95 and 0.5645 exceed 0.99. This record's, read the same way, is 0.9382.
- The 5–95 % band of the slope holds the truth in 0.8955 of the logs (standard error 0.0048), against its nominal 0.90.
- $\ln B_{\text{age}}$ at $s_\theta = 0.10$ has median 1.6405 over the logs; 0.8642 give more than 0 and 0.6460 more than 1. This record's, on the same exact route, is 0.2224.
- The quadratic's Laplace $\Pr(\theta_{aa} < 0 \mid \mathbf{y})$ has median 0.5438, and 0.0375 of the logs reach this record's 0.9611.

This log is a weak draw of a rising truth: most logs from the same truth would have settled the sign and made the case for the curve, and this one did neither, while its band still holds the truth.

16.4.7 The answers against the stated truth

The last check sets the answers against the stated truth of the record (Section 16.6.5). Table 16.9 collects the routes: each row is a route or a variation, the columns the slope's median, its band and the probability of a rise, and coil 0's rate at age 3 with its band. The true slope is $0.15\,\text{yr}^{-1}$, $+1.056$ standard deviations above the exact median and inside the band; the true $\lambda_0 = 0.05$ lies inside $[0.04181,\ 0.12238]$; $(0.05,\ 0.15)$ is the truth of the two entries. The generator draws each bin at its mid-age rate, so the rate coil 0 meets from age 3 to 4, over the next 90 days and over the year ahead, is bin 3–4's, $0.0845\,\text{yr}^{-1}$, inside coil 0's band; its 90-day probability is 0.0206, the value of Chapter 15's step 5. The record fell short of the truth where the lever arm on the slope is: the bins from age 8 hold 3 failures against 6.909 expected. The truth has no condition effect, $\theta_g = 0$, and no spread among the coils, $\nu = 0$.

Table 16.9: The routes to the slope and to coil 0's rate at age 3, and the stated truth. Bands are 5–95 %.

Route	θ_a median [yr^{-1}]	band [yr^{-1}]	$\Pr(\theta_a > 0)$ [–]	$\lambda(3)$ median [yr^{-1}]	band [yr^{-1}]
Laplace normal on $\mathbf{z}$ (reported)	0.09055	−0.00287–0.18397	0.9446	0.0979	0.0700–0.1369
Exact: level integrated, 14 001-node grid in θ_a	0.08949	−0.00621–0.18227	0.9383	0.0959	0.0670–0.1318
Grid in both entries, 801×801, flat priors	0.08949	–	0.93776	0.09588	–
Frailty freed, $\nu \sim$ half-normal(0.5); $\lambda(3)$ coil 0's own	0.0885	–	0.9340	0.0914	0.0401–0.1579
Power law in age, exact; the sign of $\theta_{\ln a}$	–	–	0.9839	0.1037	0.0748–0.1388
Stated truth	0.1500	–	rising	0.0845, bin 3–4	–

16.5 The report (step 6)

The dashboard's flat line is the dashboard's weighting, not the fleet. Weighted by their coil-years, the same bars give a rate that rises 9.48 %/yr per year of age (−0.29 to 20.20), with probability 0.94. Coil 0, at age 3, sits at 0.098 failures a year (0.070 to 0.137), below the fleet's one rate with probability 0.73: a young coil on a rising curve, and its three failure-free years are what that curve expects. The oldest coils run at 0.21 a year, 1.91 times the fleet's one rate, with a band from 0.11 to 0.42, because the log holds few old coil-years. The log makes the rise probable, not proven: against one constant rate the evidence is about even, a Bayes factor of 1.29 (1.25 on the exact route). For coil 0 the difference is small, 0.0246 on the Laplace route against Chapter 15's exact 0.0277 over the next 90 days; for an eleven-year-old coil it is 0.0529 against 0.0269. And "fouled coils leak more" is age: once age is in the rows, condition has no sign of its own in this log.

16.5.1 The answers and their bands

Table 16.10 gives the answers in the order of Table 16.2. Each row is a question; the second column gives the reported value from the Laplace normal, the third the exact route's value beside it, and the fourth the 5–95 % band of the reported value.

Table 16.10: The report for the fleet's failure log and coil 0 at age 3 (window 36, Dec Y3). Reported values from the Laplace normal of the state; bands are 5–95 %.

Asset manager's question	Reported value	Exact route	Band
Does the failure rate rise with age, and where is this unit on that curve?	$\Pr(\theta_a > 0) = 0.9446$ [–]; $\lambda(3) = 0.0979$ [yr^{-1}]; $\Pr(\lambda(3) < 0.1107) = 0.7270$ [–]	0.9383 [–]; 0.0959 [yr^{-1}]; 0.7646 [–]	$\lambda(3)$: 0.0700 to 0.1369 [yr^{-1}]
How fast does it rise, and in how many years does it double?	9.48 [%/yr]; $t_{\mathrm{dbl}} = 7.655$ [yr], never doubles 0.0554 [–]	9.36 [%/yr]; 7.746 [yr], 0.0617 [–]	−0.29 to 20.20 [%/yr]; t_{dbl} from 3.768 [yr], no upper end
What rate should I plan with for coil 0 over the next year and the next 90 days, and for the oldest coils?	year ahead 0.1025 [yr^{-1}]; 90 days 0.02459 [–]; at the shutdown 0.1056, at age 11.5 0.2114 [yr^{-1}]	0.1002 [yr^{-1}]; 0.02390 [–]; 0.1031, 0.2061 [yr^{-1}]	0.0749 to 0.1403; –; 0.0779 to 0.1431; 0.1059 to 0.4222 [yr^{-1}]
Is one constant rate good enough, or does the record need the age curve?	$\ln B_{\mathrm{age}} = 0.2555$, $B_{\mathrm{age}} = 1.2911$ (Savage–Dickey) [–]	0.2224, $B_{\mathrm{age}} = 1.2491$ [–]	over s_θ: at most 0.2760; negative above 0.1596 [yr^{-1}]
Do the coils in worse condition fail more, once age is counted?	with age: $\theta_g = -0.5285$, $\Pr(\theta_g < 0) = 0.5874$; alone: −1.9362, 0.9320 [–]	−0.2759, 0.5410; −1.8286, 0.9038 [–]	−4.4666 to 3.4095; −4.0721 to 0.1997 [–]

The answers rest on twelve count rows and the curve's form. The two routes put every answer on the same side of every comparison the questions ask, and the exact route reads every rate a little lower, because the level's gamma conditional has its median below the mode of its logarithm, and the slope's spread adds to that.

16.5.2 The entry in the fleet register

The fleet's failure assessment is entered in the fleet register under coil 0 at age 3, window 36 (Dec Y3), as Chapter 9's condition assessment was (Section 9.5.2), with the fleet's curve on its first line. It uses nothing later than window 36, and in the values of Table 16.10 it reads:

Coil 0 (the hall's coil 1), age 3, window 36 (Dec Y3); fleet failure log of 50 coils (32 failures in 289 coil-years).
Fleet failure curve: rises 9.5 %/yr per year of age (5–95 % band −0.3 to 20.2); $P(\text{rising}) = 0.94$; against one constant rate, Bayes factor 1.29 (1.25 on the exact route; not decisive).
Coil 0 on the curve: 0.098 failures/yr at age 3 (0.070–0.137); no failure in its three years (0.26 expected).
Year ahead (Jan–Dec Y4): 0.103 failures/yr (0.075–0.140); next 90 days, $P(\text{failure}) = 0.025$.
The dashboard's flat 0.11/yr at every age: not supported for the old coils (age 11.5: 0.21/yr, 0.11–0.42).
Condition beside age: no separate effect resolved.
Action: plan coil 0 at the curve's rate for its age; carry the age curve, not the fleet's one rate, into repair-or-replace and prioritization.
Next measurement: keep the failure log with each failure's age; obtain old coil-years.

The entry records what the band rests on as well as the values: twelve bins of a fleet whose coil-years thin with age, read on a log-linear curve, with coil 0 placed by its age and not by a record of its own. Which coil the crew visits first is not settled by a rate; it needs the expected consequence of waiting that Chapter 20 computes for each coil, and whether coil 0 is worth restoring is Chapter 17's.

Common mistake

Fitting a line through raw failure rates by age. Each bar's rate has a precision set by its exposure, and a line that weighs every bar alike lets the thinnest bins, the old ones, flatten the curve. The count rows weight each bin by its expected count, and the slope is an entry of the state with its band.

16.6 The next measurement (step 7)

The report leaves the slope between −0.29 and 20.20 %/yr and the case for the curve about even. A further record is more count rows, and what it would do to the slope's band and to the two verdicts can be computed before it is kept.

16.6.1 The preposterior by simulated outcomes

A candidate record r adds exposure E_b^r in age bins b, possibly bins beyond the twelfth. Its fixed-weight preposterior is equation (4.65) with the expected information at the MAP state,

$$\boldsymbol{\Sigma}_r = \left(\boldsymbol{\Sigma}^{-1} + \mathbf{J}_r^{\mathsf{T}}\mathbf{J}_r\right)^{-1}, \qquad \mathbf{J}_r^{\mathsf{T}}\mathbf{J}_r = \sum_b E_b^r\, \hat{\lambda}(\bar{a}_b)\,(1,\ \bar{a}_b)^{\mathsf{T}}(1,\ \bar{a}_b), \tag{16.22}$$

each new bin a count row of the form (16.4) with its expected count in place of μ_a. The weights of new count rows depend on the rate they will meet, which the posterior does not fix, so the reported value is the simulated one:

1. Draw a state from the exact posterior: a node θ_m with probability ϖ_m, then λ_0 from $\text{Gamma}(N_f,\ \mathcal{E}(\theta_m))$, equation (16.12).
2. Draw the candidate's counts, $\text{Poisson}(E_b^r \lambda_0 \mathrm{e}^{\theta_m \bar{a}_b})$ in each of its bins.
3. Add them to the log and rerun the exact route: the slope's standard deviation, $\Pr(\theta_a > 0 \mid \mathbf{y})$ and $\ln B_{\text{age}}$ at $s_\theta = 0.10$.
4. Repeat 2000 times (seed 23) and average.

The width taken off is $1 - \text{sd}_r/\text{sd}$ of θ_a, from the exact 0.05732. Each candidate is also asked whether it would settle the two questions the record left open: the probability that the verdict on the sign passes 0.95, and the probability that $\ln B_{\text{age}}$ passes 1.

16.6.2 The candidates

The candidates are of three kinds. The fleet's own failure log, kept for one, three or five more years, is the permanent instrument at the place of the question: each year adds 50 coil-years, every coil one year older, up to age 13 after one year. Coil 0's own next year adds one coil-year at age 3 to 4. An older fleet's log, 30 coil-years at ages 10 to 13, is the maker's field returns or a sister site's coils of this model past age 10. Table 16.11 ranks them. Each row is a candidate; the columns give the slope's standard deviation after it, simulated and with fixed weights, the width it takes off, and the two probabilities of settling. The widths taken off are compared as bars in Figure 16.7, ranked, with the level known exactly as a dashed line.

Table 16.11: Candidate records for the fleet's failure curve, by simulated outcomes from the exact posterior (2000 per candidate); fixed weights by equation (16.22). The width taken off is from the exact $\mathrm{sd}(\theta_a \mid \mathbf{y}) = 0.05732$.

Candidate	sd θ_a, simulated (fixed weights) [yr^{-1}]	width off [%]	Pr(verdict > 0.95) [–]	Pr($\ln B_{\mathrm{age}} > 1$) [–]
The fleet's failure log, one more year (50 coil-yr at ages 1–13)	0.04844 (0.04782)	15.5	0.5690	0.3240
The fleet's log, three more years	0.03604 (0.03544)	37.1	0.7105	0.5675
The fleet's log, five more years	0.02792 (0.02728)	51.3	0.7775	0.6765
Coil 0's own next year (1 coil-yr at age 3–4)	0.05731 (0.05678)	0.0	0.0000	0.0000
An older fleet's log, 30 coil-yr at ages 10–13	0.04300 (0.04233)	**25.0**	0.6375	0.4425
The fleet's log one more year, with those 30 old coil-years	0.03961 (0.03895)	30.9	0.6900	0.5280
The level known exactly, the log as it is	0.03099 (Laplace)	45.9	–	–

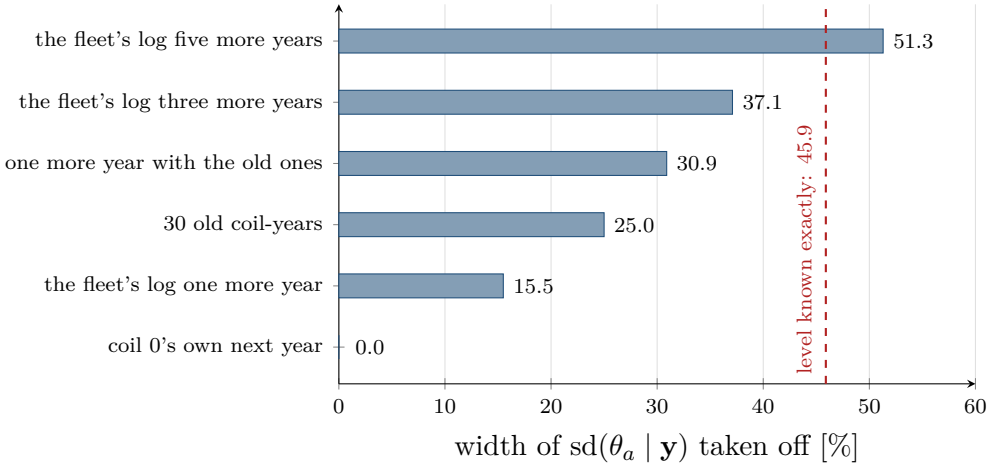


Figure 16.7: The candidate records of Table 16.11, ranked by the width they take off the slope's standard deviation, by simulated outcomes. Dashed: the level λ_0 known exactly with the log as it is, $\mathrm{sd}(\theta_a) = 0.03099$ against the exact 0.05732, 45.9 % off; a record that lengthens the age range informs the slope directly and can pass it.

The step's surprise: thirty old coil-years are worth more than fifty young ones. A year of the whole fleet's log, 50 coil-years at ages 1 to 13, takes 15.5 % off; 30 coil-years at ages 10 to 13 take off 25.0 %. The lever arm on the slope is age, through the information (16.8), whose slope entry weighs each coil-year by its expected count times the square of its age, and the log is short of old coils. Coil 0's own next year says nothing of the curve: it adds one coil-year at

an age the log already holds well. The dashed line is a reference, which a record can pass: with the level known exactly the slope's spread would be the reciprocal root of its own information entry, $1/\sqrt{1041.5510} = 0.03099$, and five more years of the fleet's log pass it, because they add coil-years beyond the twelfth year that inform the slope's own column.

16.6.3 The recommendation

Keep the fleet's failure log, with each failure's age: the totals alone keep 27.8 % of the slope's information (Section 16.4.2). Seek old coil-years: the maker's field returns, or a sister site's log of coils of this model past age 10, which per coil-year take off the most. One more year of the fleet's own log settles the sign with probability 0.57, a verdict above 0.95, and three years with 0.71. Coil 0's own record will not change its place on the curve; the curve will.

16.6.4 What the analysis showed

The dashboard drew a flat 0.11 failures a year through the bars and planned every coil there; weighted by its coil-years the same log rises about 9.5 % a year per year of age, and puts coil 0 at 0.098 at age 3. The difference is not a second record but a second weighting: count rows in exposure weigh each bin by what it can know, and the failure log becomes trend rows over age, with the level integrated out so that the exact posterior is a one-dimensional quadrature beside the Laplace normal. Two questions that a planner tends to merge part here: the probability that the rate rises, 0.94, and the evidence that a rising curve is needed, a Bayes factor of 1.29; the first is inside the curve, the second against the constant rate. Condition and age move together across the fleet, and read alone, condition takes the credit for age. And the record's weak slope is one draw from its truth, at the 16th percentile of the logs the truth would give; its band holds the truth, and what would sharpen it is old coil-years, not more years of the young ones.

16.6.5 Limits of the model

A constructed record. The failure log is generated from a stated truth, and step 5 compares the answers with that truth (Section 16.4.7). It is not field data, and the agreement is a check of the method on its own model, not of the model against a real fleet.

Minimal repair and alike coils. Every repair is taken to return the coil to the rate of its age, the coil-years to be independent, and the coils to be alike at one age. A repair that renews the bundle, a leak that invites the next, or coils that differ in proneness would each move the rows; Section 16.4.5 frees a gamma frailty, which the log neither shows nor rules out.

A log-linear curve. A quadratic and a power law fit the record as well (Section 16.4.4), and the power law gives a sharper sign through its steep start. Coil 0's rate at age 3 moves little under the power law, 0.1037 against 0.0959 on the exact routes; the rates at the oldest ages, where the log is thin, are the ones a change of shape can move.

Age alone. The rate depends on age alone: no site, duty or manufacturing-year covariate. Condition enters only through Chapter 9's fitted line, taken as known, which attenuates its coefficient.

One-year bins. The rows read one-year bins at their mid-ages. A log with each failure's exact age would be read by the same rows, one per coil-year or one per failure, with the same two entries.

The next chapter. Part V turns from forecasts to decisions. Whether restoring coil 0's conductance is worth more than replacing it, Chapter 17, reads the remaining-life band of Chapter 14 and the failure curve here at coil 0's age.

16.7 What to retain

Definition 16.1 (An age-dependent failure rate from count rows)**.** A fleet's failure log is read as count rows in exposure, $n_a \sim \text{Poisson}(E_a \lambda_0 \mathrm{e}^{\theta_a \bar{a}_a})$, on the state $(\ln \lambda_0,\ \theta_a)$, each row in the working form $r_a = (\mu_a - n_a)/\sqrt{\mu_a}$, $\mathbf{j}_a = \sqrt{\mu_a}(1,\ \bar{a}_a)$. On the canonical link the curvature term is zero, the score equations fix the fit to the sufficient statistic $(N_f,\ N_f a_w)$, and the band of $\ln \lambda(a)$ is narrowest at the failure-weighted mean age a_w, with variance $1/N_f$ there. With a flat prior on the level it integrates out: $p(\theta_a \mid \mathbf{y}) \propto \mathrm{e}^{\theta_a N_f a_w} \mathcal{E}(\theta_a)^{-N_f}$ and $\lambda_0 \mid \theta_a, \mathbf{y} \sim \text{Gamma}(N_f,\ \mathcal{E}(\theta_a))$. A unit's position is the curve at its age, a derived quantity of the same posterior.

Three conclusions carry beyond this fleet. A failure rate by age is a trend whose rows are counts: weighted by exposure, the old bins speak softly and the young loudly, and a line through raw rates gets this backward. The probability that a rate rises and the evidence that it needs to rise are different questions, read from the same posterior and the same likelihood, and a record can make the first high while leaving the second near even. And a condition covariate that moves with age takes the credit for age when age is left out of the rows; with age in, a log of this size cannot tell them apart.

Table 16.12 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 16.12: The terms of an age-dependent failure rate and where each is defined. Units for the fleet's failure log.

Term	Expression	Units	Meaning	Defined in
Failure log as count rows	$n_a \sim \text{Poisson}(\mu_a)$, eq. (16.3), working form eq. (16.4)	–, coil-yr	each bin weighted by its expected count	16.1.4
Log-linear rate curve	$\lambda(a) = \lambda_0 \mathrm{e}^{\theta_a a}$, eq. (16.1)	yr^{-1}	the ROCOF under minimal repair	16.1.1
Sufficient statistic	$(N_f,\ N_f a_w)$, eq. (16.6)	–, yr	the total and the age-weighted count	16.1.4
Failure-weighted mean age	$a_w = \sum_a n_a \bar{a}_a / N_f$	yr	where the failures sit	16.1.4
Posterior				
Band's waist	eq. (16.11), $\text{sd} = 1/\sqrt{N_f}$	yr, –	the age the curve is best known	16.2.2
The level integrated out	eq. (16.12)	–	gamma conditional, 1-D quadrature	16.2.3
Age-weighted exposure	$\mathcal{E}(\theta_a) = \sum_a E_a \mathrm{e}^{\theta_a \bar{a}_a}$	coil-yr	the gamma's rate	16.2.3
Questions and answers				
A unit's position	$\ln \lambda(a) = \mathbf{g}(a)^\mathsf{T}\mathbf{z}$, eq. (16.15)	yr^{-1}	the curve at the unit's age	16.3.1
Rate ratio, doubling time	e^{θ_a}; $t_{\text{dbl}} = \ln 2/\theta_a$, eq. (16.16)	–, yr	by inversion, with a mass at infinity	16.3.2
Window failure probability on the curve	eq. (16.18)	–	the predictive with the curve inside	16.3.3
Rise against the need for a rise	eq. (16.14) against eq. (16.19)	–	two questions, one posterior	16.3.4
Confounding of condition by age	eq. (16.21)	–	a covariate correlated with age takes its effect	16.3.5
Value of failure ages	totals against ages	%	27.8 % of the slope's information in the totals	16.4.2
Frailty	$Z_i \sim$ gamma, mean one, coefficient of variation ν	–	unit-to-unit proneness	16.4.5
Preposterior of a record	eq. (16.22); simulated outcomes	–	old coil-years carry the lever arm	16.6.1

16.8 Exercises

Exercise 16.1. *Formulation.* Write the question "will a ten-year-old coil fail more often than once in five years?" in the form of Table 16.2 and evaluate it from the fit of this chapter. (An estimation question at a threshold, equation (2.9) with $f = \ln \lambda(10)$ and $f^* = \ln 0.2$: $\Pr(\lambda(10) > 0.2 \mid \mathbf{y}) = 1 - \Phi\big((\ln 0.2 - \widehat{\ln \lambda}(10))/\text{sd}(\ln \lambda(10))\big)$ from equation (16.15), 0.4079; on the exact route, by the mixture (16.13), 0.3796.)

Exercise 16.2. *Derivation.* Integrate $\ln \lambda_0$ out of the twelve count rows under a flat prior. Show that $\lambda_0 \mid \theta_a, \mathbf{y} \sim \text{Gamma}(N_f,\ \mathcal{E}(\theta_a))$, and that $N_f/\mathcal{E}(\hat{\theta}_a)$ is the MAP's λ_0. (With $w = \lambda_0$ and $\mathrm{d}\ln\lambda_0 = \mathrm{d}w/w$ the posterior density in w is proportional to $\mathrm{e}^{\theta_a N_f a_w} w^{N_f - 1} \mathrm{e}^{-w\mathcal{E}(\theta_a)}$, a gamma kernel of shape N_f and rate $\mathcal{E}$; in $\ln w$ the density is $w^{N_f}\mathrm{e}^{-w\mathcal{E}}$, largest at $w = N_f/\mathcal{E}$, and the first score equation $\sum_a \mu_a = N_f$ is the same condition: $\mathcal{E}(\hat{\theta}_a) = 428.8384$, $N_f/\mathcal{E} = 0.07462$.)

Exercise 16.3. *Derivation.* Show that the score equations are $\mathbf{X}_a^\mathsf{T}\hat{\boldsymbol{\mu}}_a = \mathbf{X}_a^\mathsf{T}\mathbf{k}$, and that the band of $\ln \lambda(a)$ is narrowest at the failure-weighted mean age. (The gradient of $\sum_a(\mu_a - n_a \ln \mu_a)$ is $\sum_a(\mu_a - n_a)\nabla \ln \mu_a = \mathbf{X}_a^\mathsf{T}(\boldsymbol{\mu}_a - \mathbf{k})$; at the MAP the information's off-diagonal entry at

origin a_{o} is $\sum_a \hat{\mu}_a(\bar{a}_a - a_{\mathrm{o}}) = N_f(a_w - a_{\mathrm{o}})$, zero at a_w, where the variance (16.15) is smallest: $\mathbf{X}_a^{\mathsf{T}}\mathbf{k} = (32,\ 153.0)$, $a_w = 4.7812$ yr.)

Exercise 16.4. *Application.* Fleet coil 2 is eleven years old. Give its 90-day failure probability on the curve and at the fleet's one rate. (By equation (16.18) with the span from age 11 to $11 + \tau$: 0.05290 on the Laplace route and 0.05079 on the exact route; at the fleet's one rate, $1 - (289/(289 + \tau))^{32} = 0.02692$.)

Exercise 16.5. *Application.* The oldest bin had logged 2 failures in its 3 coil-years instead of none. Refit by the exact route and give the slope, its sign and the evidence at $s_\theta = 0.10$. (The sufficient statistic gains two failures at the mid-age 11.5; the slope has exact median 0.13007 and band $[0.03938,\ 0.21927]$, $\Pr(\theta_a > 0 \mid \mathbf{y}) = 0.9904$, and $\ln B_{\mathrm{age}} = 1.4015$.)

Part V

Maintenance and Replacement Decisions

Chapter 17

Repair or Replace

17.1 The maintenance calendar, the dashboard and a new coil

It is December of year 3, window 36, and the planner of the two-unit hall has to write the work order for its coils. Chapter 14 dated the crossing of the capacity limit, $G_1 + G_2 \geq 10.761\,\mathrm{kW/K}$: window 53 (May Y5) at the median, with a band from window 44 (Aug Y4) to window 82 (Oct Y7), and about one chance in six, 0.1595, that the hall has crossed by the October shutdown of year 4, window 46. Its maintenance log planned the coil cleaning for that shutdown and left the economics of the work to this chapter. Two other readings are on the desk. The condition-monitoring dashboard still draws its line through the health-index tiles, $13.3588 - 2.0164\,t_{\mathrm{y}}\,\mathrm{kW/K}$, which reached the limit in April of year 2 and reads $7.3097\,\mathrm{kW/K}$ today, 3.4513 below it; its conclusion is to clean both coils now and to keep cleaning them. The maintenance calendar cleans both coils at every October outage, starting with the shutdown, at R$ 16 000 a visit.

Neither reading asks whether coil 1 should be replaced. A new coil costs R$ 60 000 against R$ 8 000 for cleaning one, and a clean restores 85 % of the conductance a coil has lost since commissioning; replacing coil 1 and cleaning coil 2 in one visit costs R$ 68 000. On those prices cleaning is the obvious choice. The chart cannot say how long a clean holds the hall when coil 1's loss may be accelerating, or whether the acceleration survives the clean; what the chance of a crossing before October costs against the R$ 12 000 of working away from an outage; or what a new coil buys over five years when it fouls at the design's rate and not at coil 1's. The question of this chapter is

which work, now and at the October shutdown, and which later cleans minimize the expected cost over five years, priced over Chapter 14's posterior of the two coils' paths: cleaning, replacing coil 1, or waiting, against the dashboard's and the calendar's schedules.

It is answered with the method of Chapter 5, on Chapter 14's extended state with each plan's work written into its forecast windows.

Jargon: repair or replace, planning horizon, discounting, maintenance plan, first move, open-loop plan, shortfall month, break-even cost, insurance action, expected value of perfect information, expected value of partial perfect information

Repair or replace is the choice between restoring a unit's condition and renewing the unit. The *planning horizon* is the span over which the consequences of a choice are counted, here five years; *discounting* weighs a cost by the time until it is paid, so that a real paid later counts for less than one paid today. A *maintenance plan* states the work at each visit over the horizon, and its *first move* is the work at the first visit. An *open-loop plan* is fixed today, with every later visit chosen now; a plan revised as the record grows is not. A *shortfall month* is a month in which the hall stands at or below its capacity limit. The *break-even cost* of a stated price is the value at which the chosen action changes. An *insurance action* has the lower expected loss and is beaten in most outcomes, because it removes the costly tail. The *expected value of perfect information* (EVPI) is the expected loss of the best action now less the expected loss of the best action chosen with the state known: the most any measurement taken before the action can be worth. The *expected value of partial perfect information* (EVPPI) is the same with one open question known instead of the whole state.

Key idea

Repair or replace is a comparison of plans over a horizon, not of two prices. Each plan rewrites the forecast windows of the path: a clean is a step back toward the commissioned conductance, a replacement restarts the coil at its rating on the design's rate. Each plan's loss is its work plus the months the hall spends below the limit, and the remaining-life band enters through those months. Its tail, not its median, decides.

17.1.1 The price of a plan

Chapter 2 wrote the expected consequence of an action as its certain cost plus, for each event it leaves, the event's consequence times its predictive probability, equation (2.23). Here the action is a *plan* $\mathcal{P}$: the work done now, at window 36, and at each October outage of the horizon, $\mathcal{O} = \{46, 58, 70, 82, 94\}$. Let $W_t(\mathcal{P})$ be the price of the work plan $\mathcal{P}$ does at window t, zero where it does nothing: CLEAN for each coil it cleans, REPLACE for each coil it replaces. The cost schedule's SHORTFALL, written SF, is the consequence of a month at or below the limit, and OFF_WINDOW, written OW, the premium of a visit away from an outage. Each window of the horizon carries one event, the hall at or below G^* that month, and every cost is discounted to today. The loss of a plan is

$$\begin{aligned} L(\mathcal{P}, \mathbf{z}) &= c_{\mathcal{P}} + \mathrm{SF} \sum_{t=37}^{96} \mathrm{df}(t)\, \mathbf{1}\big[G_{\mathrm{tot},t}(\mathbf{z}; \mathcal{P}) \le G^*\big], \qquad \mathrm{df}(t) = (1 + 0.08)^{-(t-36)/12}, \\ c_{\mathcal{P}} &= W_{36}(\mathcal{P}) + \mathrm{OW}\, \mathbf{1}\big[W_{36}(\mathcal{P}) > 0\big] + \sum_{t \in \mathcal{O}} \mathrm{df}(t)\, W_t(\mathcal{P}), \end{aligned} \tag{17.1}$$

with $G_{\mathrm{tot},t}(\mathbf{z}; \mathcal{P})$ the coils' total conductance at window t when the plan's work has been done, and $\mathrm{df}(t)$ the discount factor at DISCOUNT, 0.08 a year. It is equation (2.23) with $c_a = c_{\mathcal{P}}$, one event $\mathcal{U}_t = \{G_{\mathrm{tot},t} \le G^*\}$ per window, the shortfall of month t, and the consequence $C_{\mathcal{U}_t} = \mathrm{SF}\, \mathrm{df}(t)$. Its expectation under the posterior is the plan's expected loss,

$$\mathcal{R}(\mathcal{P}) = \mathbb{E}\big[L(\mathcal{P}, \mathbf{z}) \mid \mathbf{y}\big] = c_{\mathcal{P}} + \mathrm{SF} \sum_{t=37}^{96} \mathrm{df}(t)\, \Pr(G_{\mathrm{tot},t}(\mathbf{z}; \mathcal{P}) \le G^* \mid \mathbf{y}), \tag{17.2}$$

the second line of equation (2.23), in which each month's probability is the third line of the trend form (2.21) on the path the plan leaves.

A clean and a replacement act on the forecast windows of the state, so the plan enters the same extended state that Chapter 14 carried past the record; it is not a second model. An *action* is then a class of plans, the work it starts with, and it is priced by its best plan.

Principle 17.1. *Price each action by its best plan over the horizon, under the posterior of the path; never by the cost of one visit, and never by the date at the median.*

17.1.2 The two-unit hall with work on its forecast windows (step 0)

The graph is Chapter 14's, Figure 14.1: the trend layer of both coils carried past window 36 by the trend rows, the trend entries γ_1 and γ_2 and coil 1's curvature $\dot{\gamma}_1$ above it, and the requirement node $G_{\mathrm{tot},t} \geq G^*$ on every forecast window, with the limit $G^* = 10.761\,\mathrm{kW/K}$ the capacity study's. Below that layer it is Chapter 10's graph of the record, Figure 10.1. The modeling choices of the window graph are Chapter 5's paragraph, carried through Chapter 14: steady state within each window, the airflow pinned with no discharge control, the season's quantities as known inputs, and one conductance per coil.

What the decision adds to the graph. A visit node sits at window 36 and at each October outage. Each is a box with its work and its price: nothing, a clean of both coils, or the replacement of one coil with the other cleaned. A clean feeds a step into the trend row that leaves its window, on each coil it touches, toward that coil's commissioning entry $g_{u,0}$, which now enters the question. A replacement cuts coil 1's trend row after the visit and starts a new chain at $g = 0$, the rating, fed by γ_{new}, the rate of a new coil of the design. That rate is Chapter 9's predictive, equation (9.17), drawn as a pinned node: it comes from the fleet's posterior and is not an entry of the hall's. The prices, the requirement nodes and the discount factors feed a loss node. Figure 17.1 draws the layer from window 36 to the end of the horizon. The two chains are coil 1's above and coil 2's below, with the requirement nodes between them; the squares are the trend rows, and the three red squares are those that leave a visit's window. The visit boxes sit above, each linked to its squares on both chains. The commissioning entries stand at the left, and their dashed links run to the squares a clean steps toward them. The new coil's rate enters at the top left, with dashed links to coil 1's squares at the two windows where a replacement can fall.

Modeling choices at the graph. Four choices make the decision's part of the graph, and no question block owns them. *The coil keeps its own trend after a clean*: the clean moves the level only, and the drawing's rate and curvature carry on in calendar time. This buys a clean that is one step in the existing rows, and it fails if the acceleration is the fouling's own and a clean resets it; step 5 runs that variant, and it flips the decision. *A new coil fouls at the design's rate*, Chapter 9's predictive, with the record's wander and no curvature. This gives the new chain a rate the fleet has measured, and it fails if coil 1's acceleration belongs to the site, its water or its air, rather than to the coil; step 5 runs that variant too, and it flips the decision as well. *A plan is fixed today*, open-loop: its later cleans are chosen now and not revised as the record grows. This buys one expected loss per plan, and it overstates the cost of every plan, of a plan that would wait for evidence most; step 7 values the revision. *No maintenance outside the plan, and no failure priced*: a tube leak (Chapter 16) has no stated cost, and it enters only as a sensitivity row of step 5.

The record and the stated inputs. The record is Chapter 14's, windows 0 to 36 (Table 10.1), read through the same 18 drawings. The costs are the cost schedule's stated inputs, and none is re-chosen. Table 17.1 lists those this chapter reads: each row is an input, with its value and

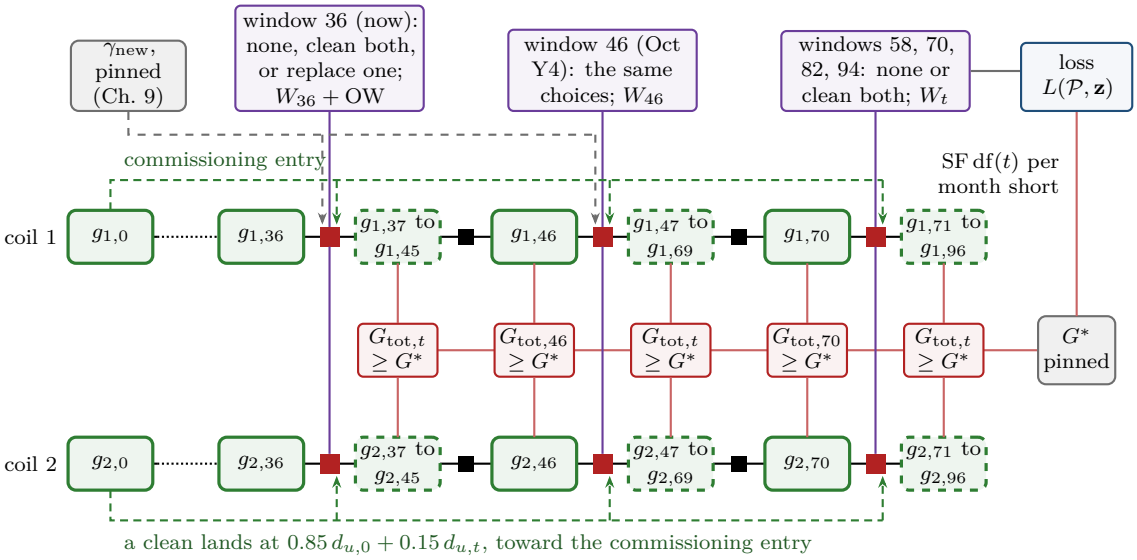


Figure 17.1: The trend layer of Figure 14.1 from window 36 to the end of the horizon, window 96, with the decision drawn on it. Free (green): the coils' log-derates, read at windows 0 and 36 and forecast from window 37 (runs as dashed boxes); windows 1 to 35 and the trend entries γ_1, γ_2, $\dot{\gamma}_1$ that feed every square are not drawn. Black squares: trend rows; red squares: the trend rows that leave a visit's window, which carry the plan's work (purple boxes and links). Dashed green: the commissioning entries, toward which a clean steps. Dashed gray: the new coil's rate γ_{new}, a pinned predictive from the fleet, on coil 1's rows after a replacement now or at the shutdown. Red: the requirement nodes and the pinned limit $G^* = 10.761\,\text{kW/K}$, whose shortfall months, with the visits' prices, feed the loss node.

where it enters. INSPECT, FALSE_ORDER and the fleet's and the warranty's costs do not enter. Today is window 36 (Dec Y3), and the horizon is windows 37 to 96 (Jan Y4 to Dec Y8); the discount factor is 0.9936 for the first month, 0.9379 at the shutdown and 0.6806 at window 96.

The cost schedule names one outage, the October shutdown of year 4. The hall's planned outage is taken to recur each October, as Chapter 14 read the October of year 5 as "the following year's outage"; step 5 sets the later Octobers off-window beside it.

The maintenance log so far. The two-unit hall's log runs from its commissioning test at window 0 to Chapter 14's forecast at window 36, which closed with "plan the coil cleaning for the October Y4 shutdown; do not defer it to the following year's outage" and asked for a refit at window 45 before the work order is written (Section 14.5.2). Whether that cleaning is worth its cost against the risk it removes was left to this chapter.

17.1.3 The asset manager's questions (step 1)

The asset manager asks five questions of the record, and step 1 writes each as an expectation or a probability over the extended state of the family's drawings, with the unknown on the left of the bar and the record's readings and the priors on the right. Table 17.2 lists them in the columns of Chapter 5's table: the question in the asset manager's words, the drill question first; the general form of its category from Chapter 2; its mathematical form for the hall; its category; the computation of Section 2.3.6 that answers it; and the subsection where it is answered. "This unit" in the drill question is coil 1, the coil whose replacement is in question; the state carries both, and step 4 prices replacing either. How often to read the path between visits is

Table 17.1: The stated inputs of the decision, from the cost schedule and the record. Money in reais; the discount factors from today, window 36.

Input [unit]	Value	Where it enters
CLEAN [R$ per coil]	8 000	W_t; both coils 16 000 a visit
REPLACE [R$ per coil]	60 000	W_t; one coil replaced and the other cleaned 68 000, both replaced 120 000
SHORTFALL, SF [R$ per month]	150 000	each month at or below G^*
OFF_WINDOW, OW [R$ per visit]	12 000	a visit away from an outage: here a visit now
RESTORE_FRACTION, RF [–]	0.85	the share of the loss since commissioning a clean restores
DISCOUNT [yr^{-1}]	0.08	df(t): 0.9379, 0.8684, 0.8041, 0.7445, 0.6894 at windows 46, 58, 70, 82, 94
HORIZON [yr]	5	windows 37 to 96
Capacity limit G^* [kW/K]	10.761	the requirement node of every forecast window
October outages $\mathcal{O}$ [window]	46, 58, 70, 82, 94	the visits that pay no OW

Chapter 18's question, at what reading to raise the order Chapter 19's, and which fleet coil the crew visits first Chapter 20's. Which measurement would change the decision asks nothing new of today's evidence; it is step 7.

17.1.4 The evidence as rows (step 2)

Step 2 adds no row to the record. The rows of each drawing are Chapter 14's: the rows of the record on the drawing's state (10.3), with the curvature's and the step's columns where the drawing frees them, extended by the forecast rows (14.4) to the forecast windows of equation (14.3). What the decision adds acts on the forecast rows alone. Write $d_{u,t} = \mathrm{e}^{g_{u,t}}$ for coil u's derate factor in window t, so that $G_{\mathrm{tot},t} = \bar{G}_1 d_{1,t} + \bar{G}_2 d_{2,t}$ by equation (14.5), and $h_{u,t} = d_{u,t}/d_{u,0}$ for its health index against commissioning.

A visit on the forecast rows. A clean of coil u at window t_v restores the fraction RESTORE_FRACTION, written RF, of the conductance lost since commissioning: the health index after the clean is $1-(1-\mathrm{RF})(1-h_{u,t_\mathrm{v}})$, and the derate factor $d_{u,0}$ times it. The clean is a step, of a size the state fixes, on the trend row that leaves the visit's window, Chapter 12's step row (12.1) with its size set by the plan instead of freed. A replacement of coil u at t_v starts a new chain at $g=0$, whose rows carry the new coil's rate γ_new and no curvature. With x_t the curvature's lever of equation (14.4), the forecast rows a plan writes are

$$
\begin{aligned}
&\text{a clean at } t_\mathrm{v}: && r_{q,u,t_\mathrm{v}} = \frac{g_{u,t_\mathrm{v}+1} - \ln d^{+}_{u,t_\mathrm{v}} - \Delta t\,[\gamma_u + \dot{\gamma}_u x_{t_\mathrm{v}}]}{\sigma_q}, \\
& && d^{+}_{u,t_\mathrm{v}} = d_{u,0} - (1-\mathrm{RF})\big(d_{u,0} - d_{u,t_\mathrm{v}}\big), \\
&\text{a replacement at } t_\mathrm{v}: && r_{q,u,t} = \frac{g_{u,t+1} - g_{u,t}\,\mathbf{1}[t > t_\mathrm{v}] - \Delta t\,\gamma_\mathrm{new}}{\sigma_q}, \quad t \geq t_\mathrm{v}, \\
&\text{a clean at } t'_\mathrm{v} > t_\mathrm{v}: && r_{q,u,t'_\mathrm{v}} = \frac{g_{u,t'_\mathrm{v}+1} - \ln d^{+}_{u,t'_\mathrm{v}} - \Delta t\,\gamma_\mathrm{new}}{\sigma_q}, \quad d^{+}_{u,t'_\mathrm{v}} = 1-(1-\mathrm{RF})(1-d_{u,t'_\mathrm{v}}),
\end{aligned}
\tag{17.3}
$$

and every other forecast row is equation (14.4) unchanged. The clean's row departs from the trend row only by the step $\ln(d^{+}_{u,t_\mathrm{v}}/d_{u,t_\mathrm{v}})$, so the coil keeps its own rate and curvature after it, the first choice of step 0. The last line is a clean of a coil replaced earlier, at t_v: the new

Table 17.2: The questions asked of the two-unit hall at window 36 (Dec Y3). $\mathbf{y}$ is the readings of windows 0–36 and the prior rows; L the loss (17.1) of a plan $\mathcal{P}$; an action a is a class of plans; $\mathcal{P}^*$ is the best plan of the chosen action a^*, $\mathcal{P}_{\text{dash}}$ and $\mathcal{P}_{\text{cal}}$ the dashboard's and the calendar's schedules, and $\mathcal{P}_{\text{v}}$ a plan of one visit at window t_{v} and no later work.

Asset manager's question	General form	Mathematical form	Category	Computation	Answered in
Is restoring this unit's conductance worth more than its replacement, given what the remaining life band says?	(2.23)	$\mathcal{R}(a) = \min_{\mathcal{P}\in a} \mathbb{E}[L(\mathcal{P}, \mathbf{z}) \mid \mathbf{y}]$ for $a \in$ {clean now, clean at the shutdown, replace now, replace at the shutdown, wait}; $a^* = \arg\min_a \mathcal{R}(a)$	design	expected consequence	17.3.1
How firm is that choice, and what would knowing the path be worth?	(2.24), second line; perfect information	$\Pr(L(\mathcal{P}^*, \mathbf{z}) < L(\mathcal{P}, \mathbf{z}) \mid \mathbf{y})$ for each action's best plan $\mathcal{P}$; $\text{EVPI} = \mathcal{R}(a^*) - \mathbb{E}[\min_{\mathcal{P}} L(\mathcal{P}, \mathbf{z}) \mid \mathbf{y}]$	design	expected consequence	17.3.2
Should the work be done now, or can it wait for the October shutdown?	(2.21), third line; (2.23)	$\text{SF} \sum_{t=37}^{46} \text{df}(t) \times \Pr(G_{\text{tot},t} \le G^* \mid \mathbf{y})$ with no work, against OW	prediction under a threshold; design	trend; expected consequence	17.3.3
How long does each action hold the hall?	(2.21), second line, from the visit	$t_c(\mathbf{z}; \mathcal{P}_{\text{v}}) = \min\{t > t_{\text{v}} : G_{\text{tot},t}(\mathbf{z}; \mathcal{P}_{\text{v}}) \le G^*\}$: median and band	estimation	trend	17.3.4
What do the dashboard's and the calendar's schedules cost?	(2.23)	$\mathcal{R}(\mathcal{P}_{\text{dash}})$, $\mathcal{R}(\mathcal{P}_{\text{cal}})$ against $\mathcal{R}(a^*)$; $\Pr(L(\mathcal{P}^*, \mathbf{z}) < L(\mathcal{P}_{\text{cal}}, \mathbf{z}) \mid \mathbf{y})$	design	expected consequence	17.3.5

coil's commissioning entry is its rating, $g = 0$, so $d_{u,0} = 1$ and the clean steps toward the rating, and the row carries γ_{new} with no curvature. A visit at window t_{v} changes the windows after it only: the visit's own window is one the hall must hold, as Chapter 14 counted the shutdown's window.

The plans and the actions. A plan fixes the work now and at the shutdown, each one of nothing, a clean of both coils, a replacement of coil 1 with coil 2 cleaned, a replacement of coil 2 with coil 1 cleaned, or the replacement of both, with at most one replacement visit; then any set of the later outages 58, 70, 82 and 94 at which both coils are cleaned. That makes 256 plans. An action is the class its plan falls in. A plan with a replacement is "replace now" or "replace at the shutdown" by its replacement visit, whatever is done before it; a plan without one is "clean now" or "clean at the shutdown" by its first clean, or "wait" if year 4 has no work. The first choice files "clean now, replace in October" under the replacement it leads to: it is the replacement that changes the path.

Which row acts on which entry. The state the loss reads is each drawing's extended state (14.3) together with the new coil's rate. By the carry (14.7) a forecast window depends on the record through $g_{u,36}$, γ_u and $\dot{\gamma}_u$, and the clean's row adds $g_{u,0}$. With the plan's rows in the forecast block, the clean's line (first) and the replacement's line (second) are, at a visit window t_{v},

$$\frac{\partial r_{q,u,t_{\text{v}}}}{\partial(g_{u,0},\, g_{u,t_{\text{v}}},\, g_{u,t_{\text{v}}+1},\, \gamma_u,\, \dot{\gamma}_u)} = \frac{1}{\sigma_q}\Big(-\frac{\text{RF}\, d_{u,0}}{d^{+}_{u,t_{\text{v}}}},\ -\frac{(1-\text{RF})\, d_{u,t_{\text{v}}}}{d^{+}_{u,t_{\text{v}}}},\ 1,\ -\Delta t,\ -x_{t_{\text{v}}}\Delta t\Big),$$

$$\frac{\partial r_{q,u,t_{\text{v}}}}{\partial(g_{u,t_{\text{v}}},\, g_{u,t_{\text{v}}+1},\, \gamma_u,\, \dot{\gamma}_u,\, \gamma_{\text{new}})} = \frac{1}{\sigma_q}(0,\ 1,\ 0,\ 0,\ -\Delta t), \tag{17.4}$$

$$\mathbf{g}_L^{\mathsf{T}} = (* \text{ on } g_{1,36},\, g_{2,36},\, \gamma_1,\, \gamma_2,\, \dot{\gamma}_1,\, g_{1,0},\, g_{2,0};\ * \text{ on } \mathbf{g}_{37}, \ldots, \mathbf{g}_{96},\, \gamma_{\text{new}};\ 0 \text{ elsewhere}),$$

$$\frac{\partial L(\mathcal{P}, \mathbf{z})}{\partial \mathbf{z}} = \mathbf{0} \text{ wherever it is defined,}$$

the curvature's column in yr^2, the rates' in yr and the others dimensionless; the replacement's line in the column of $\dot{\gamma}_u$ is zero for coil 1 and absent for coil 2. The clean's row has a column on the commissioning entry, with weight RF against $1 - \text{RF}$ on today's level, so where a clean lands is mostly the commissioned conductance. The replacement's row has zeros on the old coil's level and trend: it cuts the chain. The last two lines are not rows. The first is the pattern of the loss on the state: the seven record entries the forecast and the clean read, the forecast entries, and γ_{new}, with no reading's entry, no load and no bias. The second is the gradient itself. The loss counts months, so it is constant between the surfaces on which a path touches G^* and jumps across them, and its gradient is zero wherever it is defined. That fact decides how the Laplace posterior is read for this question (step 3), and why its expansion fails (step 5).

Modeling choices at the rows. Three choices make the loss, each named once. *Whole months*: a month is a shortfall if its window stands at or below G^*, SHORTFALL's unit, not pro rata. This keeps one event per window and errs by at most a month at each crossing. *The horizon ends at window 96, with no value left in a coil*: this is the stated HORIZON, and it undervalues a new coil, whose life runs past the horizon; step 5 shows that a four-year horizon flips the decision. *One OFF_WINDOW per visit*, however many coils the visit touches: the premium is the cost of the outage, not of the work.

17.2 The posterior of the state (step 3)

Every answer of this chapter is an expected loss over the posteriors of Chapter 14's 18 drawings, each extended by its forecast windows, with the plan's rows (17.3) in place, and weighted by the drawings' evidences.

17.2.1 The posterior under a plan

The derivation form is Chapter 5's (Sections 5.2.1 and 5.2.2), on each drawing's state as Chapter 14 fitted it and not derived again (Section 14.2.1): the residual vector of drawing M stacks the record's rows, its Jacobian $\mathbf{J}_M$ has the pattern (10.6) with the curvature's and the step's columns, the MAP state $\hat{\mathbf{z}}_M$ solves the normal equations (5.9), and the Laplace covariance is $\mathbf{\Sigma}_M = (\mathbf{J}_M^{\mathsf{T}}\mathbf{J}_M)^{-1}$ at $\hat{\mathbf{z}}_M$. With every row normalized the information matrix is $\mathbf{J}^{\mathsf{T}}\mathbf{J}$; with unnormalized rows and their weights it is $\mathbf{J}^{\mathsf{T}}\mathbf{W}\mathbf{J}$. Each drawing's evidence is its Laplace evidence (7.7), and its probability P_M is equation (14.8) with the family's priors (14.2).

The plan changes nothing in the fit. Under a plan the forecast block of the Jacobian holds the plan's rows (17.3) in place of some of the rows (14.4). Its diagonal is still $1/\sigma_q$, the column of $g_{u,t+1}$ in the row that enters it, and every other entry the clean's and the replacement's lines of equation (17.4) put in the forecast columns lies below the diagonal, on $g_{u,t_{\mathrm{v}}}$; the rest, on $g_{u,0}$ and on the trend entries, fall in the record's columns $\mathbf{J}_{\mathrm{f}z}$. The block $\mathbf{J}_{\mathrm{ff}}$ of equation (14.9) is therefore square, lower triangular and invertible, and Chapter 14's Schur complement leaves the record's block of the covariance at $\mathbf{\Sigma}_M$ and its MAP state at $\hat{\mathbf{z}}_M$. Each forecast window's entries still appear only in the row that enters them and the row that leaves them, so the posterior of the extended state factors under every plan, as equation (14.10) did with no work:

$$p(\mathbf{z}_{\mathrm{ext}},\ \gamma_{\mathrm{new}} \mid \mathbf{y}, M, \mathcal{P}) = p(\mathbf{z}_M \mid \mathbf{y}, M)\ p(\gamma_{\mathrm{new}}) \prod_{t=36}^{95} \prod_{u=1}^{2} \mathcal{N}(g_{u,t+1};\ g_{u,t}^{\mathcal{P}}(\mathbf{z};\mathcal{P}),\ \sigma_q^2), \qquad (17.5)$$

with $g_{u,t}^{\mathcal{P}}(\mathbf{z};\mathcal{P})$ the value of $g_{u,t+1}$ at which the row from window t is zero: $g_{u,t} + \Delta t\,[\gamma_u + \dot{\gamma}_u x_t]$ with no work, $\ln d_{u,t}^{+} + \Delta t\,[\gamma_u + \dot{\gamma}_u x_t]$ after a clean, $g_{u,t}\mathbf{1}[t > t_{\mathrm{v}}] + \Delta t\,\gamma_{\mathrm{new}}$ on a new coil's chain, and $\ln d_{u,t}^{+} + \Delta t\,\gamma_{\mathrm{new}}$ after a clean of the new coil. The new coil's rate has the predictive density $p(\gamma_{\mathrm{new}})$ of equation (9.17), from the fleet's tests, not from this record. Integrating the forecast entries out leaves $\sqrt{2\pi}\,\sigma_q$ per row, which the row's normalizing constant cancels, so every drawing's evidence and weight is Chapter 14's. A plan cannot change what the record says; it changes only the windows the loss reads. The expected loss of a plan is then the mixture (7.10) over the drawings,

$$\mathcal{R}(\mathcal{P}) = \sum_{M} P_M\, \mathbb{E}\big[L(\mathcal{P}, \mathbf{z}) \mid \mathbf{y}, M\big], \qquad (17.6)$$

each term an integral of the loss (17.1) against the factored posterior (17.5).

How the Laplace posterior is read here. For a smooth loss, Chapter 2 read the expected consequence from $\hat{\mathbf{z}}$ and $\mathbf{\Sigma}$ by the expansion (2.24). This loss has no smooth part: by the second part of the pattern (17.4) its gradient and its curvature are zero wherever they are defined, and the expansion keeps only $L(\mathcal{P}, \hat{\mathbf{z}})$, the loss on the most probable path. The Laplace posterior is read instead by draws of each drawing's Laplace normal carried through the factor (17.5), the route of Section 4.4.7 that Chapter 14 used for the first passage. Section 17.4.2 sets the expansion beside it.

17.2.2 The entries the decision reads

The forecast reads five entries of each drawing's state, $(g_{1,36}, g_{2,36}, \gamma_1, \gamma_2, \dot{\gamma}_1)$ by equation (14.7); coil 2's curvature $\dot{\gamma}_2$ is absent in every drawing. The clean's row adds the two commissioning entries $g_{1,0}$ and $g_{2,0}$, which set where a clean lands. In the leading drawing $(\mathcal{C}_1, 2)$, coil 1 curving with coil 2 stepped, the MAP values of the seven entries are $\hat{g}_{1,36} = -0.16571$, $\hat{g}_{2,36} = -0.19007$, $\hat{\gamma}_1 = -0.04825\,\mathrm{yr}^{-1}$, $\hat{\gamma}_2 = -0.02732\,\mathrm{yr}^{-1}$, $\hat{\dot{\gamma}}_1 = -0.03403\,\mathrm{yr}^{-2}$, $\hat{g}_{1,0} = -0.02096$ and $\hat{g}_{2,0} = -0.03298$, and the block of $\mathbf{\Sigma}_M$ on them, in that order, the trend entries in yr^{-1} and the curvature in yr^{-2}, is

$$\mathbf{\Sigma}^{[7]}_{(\mathcal{C}_1,2)} = \begin{bmatrix} 8.8145 & -3.2768 & 2.7847 & -1.6225 & 3.4891 & 0.46039 & 0.90807 \\ -3.2768 & 3.2228 & -1.1289 & 1.7298 & -1.0536 & 0.11002 & -0.83889 \\ 2.7847 & -1.1289 & 2.2372 & -0.63295 & 0.54692 & -0.92700 & 0.59260 \\ -1.6225 & 1.7298 & -0.63295 & 2.8430 & -0.50099 & 0.27635 & -1.3397 \\ 3.4891 & -1.0536 & 0.54692 & -0.50099 & 4.0703 & 1.8483 & 0.0018244 \\ 0.46039 & 0.11002 & -0.92700 & 0.27635 & 1.8483 & 3.2414 & -0.86972 \\ 0.90807 & -0.83889 & 0.59260 & -1.3397 & 0.0018244 & -0.86972 & 2.5022 \end{bmatrix} \times 10^{-4}, \tag{17.7}$$

with the standard deviations 0.02969, 0.01795, 0.01496, 0.01686, 0.02018, 0.01800 and 0.01582. Its upper left 5×5 block is Chapter 14's (14.11), whose cross terms that chapter explained. What is new is the last two rows and columns, the commissioning entries, and their correlations read by sign.

- $g_{u,0}$ *with* γ_u*:* -0.3442 *for coil 1,* -0.5023 *for coil 2.* A higher start needs a steeper line through the same later readings, so the commissioning entry and the rate at the record's center move in opposite directions.
- $g_{1,0}$ *with* $\dot{\gamma}_1$*:* $+0.5089$*.* The rate of coil 1 at time t_y is $\gamma_1 + \dot{\gamma}_1(t_y - \bar{t}_y)$, so a higher curvature makes the early windows fall faster and the late ones slower. With the later windows held, a higher start is reached by a faster early fall, and the curvature carries it: the path bends down less after the record. The two trend terms pull against each other on the rate today, $\gamma_1 + 1.5\,\dot{\gamma}_1$ (standard deviation 0.03611), whose covariance with $g_{1,0}$ is 1.8454×10^{-4}: the γ_1 term contributes -0.5023 times that covariance and the $\dot{\gamma}_1$ term $+1.5023$ times. Net, $g_{1,0}$ and coil 1's rate today are correlated by $+0.2839$. A clean that lands higher, through $g_{1,0}$, therefore comes with a slower fall after it, and the two errors add in the same direction, an optimistic landing point with an optimistic trend, as Chapter 14 found for the level and the trend today.
- $g_{2,0}$ *with* $g_{2,36}$*:* -0.2954*;* $g_{1,0}$ *with* $g_{1,36}$*:* $+0.0861$*.* Coil 2's steeper line from a higher start ends lower today; for coil 1 the curvature's positive term offsets the line's, and the two entries are nearly uncorrelated.
- $g_{1,0}$ *with* $g_{2,0}$*:* -0.3054*.* Window 0's return air reads the sum of the two coils, so their split at commissioning is anticorrelated, as in Chapter 5's fit of window 0.

Where a clean lands. By the clean's row a clean of coil u at window t lands at $d^+_{u,t} = \mathrm{RF}\,d_{u,0} + (1-\mathrm{RF})\,d_{u,t}$, a derived quantity with the gradient $(\mathrm{RF}\,d_{u,0},\ (1-\mathrm{RF})\,d_{u,t})$ on $(g_{u,0}, g_{u,t})$. The weights are 0.85 and 0.15 on two entries whose spreads are of the same order, 0.01800 and 0.02969 for coil 1, so the landing point inherits its spread chiefly from the commissioning entry: the commissioning test, which Chapter 5 found to set the floor of every later health index, also sets how well a clean is known. At the leading drawing's MAP, coil 1 was commissioned at $g = -0.02096$, a derate factor of 0.9793, and stands today at -0.16571, 0.8473; a clean now restores it to -0.04138, 0.9595, or 6.7162 kW/K. Coil 2 was commissioned at -0.03298 (0.9676) and stands at -0.19007 (0.8269); a clean now restores it to -0.05503 (0.9465), 6.6252 kW/K.

The new coil and the old. The new coil's rate, re-read from Chapter 9's predictive (9.17) with that chapter's seeds, has median $-0.03028\,\mathrm{yr}^{-1}$ and band $[-0.05095,\ -0.00935]$, mean -0.03023 and standard deviation 0.01267, the values of Section 9.3.3. Coil 1's own rate today on the leading drawing, $\gamma_1 + \dot{\gamma}_1(3 - \bar{t}_{\mathrm{y}})$, is $-0.09930\,\mathrm{yr}^{-1}$ with standard deviation 0.03611, and coil 2's is -0.02732 with 0.01686. A replacement trades coil 1's rate for the design's slower one, and the new rate does not accelerate.

17.2.3 The routes to the expected loss

The expected loss (17.6) has no closed form, because the loss counts the windows at which a nonlinear function of the path falls below the limit. It is read by draws of the factored posterior (17.5).

1. For each drawing, draw $N_{\mathrm{d}} = 10\,000$ values of the seven entries of equation (17.7) from their block of $\mathcal{N}(\hat{\mathbf{z}}_M, \boldsymbol{\Sigma}_M)$, through its Cholesky factor ($\dot{\gamma}_1 = 0$ under $\mathcal{C}_0$).
2. With each draw, draw the departures $e_{u,t} \sim \mathcal{N}(0,\ \sigma_q^2)$ of windows 37 to 96 for both coils, and one new coil's rate γ_{new} from Chapter 9's predictive draws.
3. For each of the 256 plans, carry each draw window by window by the plan's rows (17.3) and the rows (14.4), evaluate $G_{\mathrm{tot},t}$ at every window of the horizon, and price the loss (17.1). The same draws price every plan, so two plans are compared draw by draw.
4. Give each draw of drawing M the weight P_M/N, and read a plan's expected loss as the weighted mean, the probability that one plan costs less than another as the weighted share of the draws on which it does, and quantiles from the weighted distribution of the loss.
5. State the Monte Carlo standard error of each result by batch means over 20 batches, as in Section 5.2.4.

This is the reported route: the Laplace posterior of each drawing, read by its draws and mixed by the drawings' Laplace evidence. Beside it stands the full posterior of the same rows, by Chapter 5's reweighted draws (Section 5.2.4): 10 000 draws per drawing on its whole state, each weighted by equation (5.13) with the drawing's objective, and the drawings weighted by their importance-sampling evidence, equation (7.12). Because the posterior factors, each reweighted draw is carried past window 36 by the plan's rows exactly as a draw of the Laplace normal is, and keeps its weight. The effective sample sizes run from 95.3 to 98.8 % of the draws, and the importance-sampling evidence moves no drawing's weight by more than 0.0010.

With no work the draws reproduce Chapter 14's forecast: the first window at or below the limit has median 53 (May Y5) and band $[44,\ 82]$, and the probability of a crossing by window 46 is 0.1632 and by window 45 0.1079, against that chapter's 0.1595 and 0.1055 on its own seeds.

17.2.4 Each drawing's best plans

Table 17.3 prices, on each of the five leading drawings alone, the plan that is best over the family for each action: each row is a drawing, with its probability; the next five columns give the expected loss of each action's best plan on that drawing; the last column gives the action best on that drawing alone, with its own best plan where it differs from the family's.

The drawings disagree on the action, and they split on the curvature, not on the step. Every drawing with coil 1 curving replaces now, and on each of them the replacement costs from R$ 92 865.3, its work alone, to R$ 92 913.6, because the new coil holds the hall. Both constant-rate drawings clean at the shutdown: the calendar's five cleans cost little more than their work there, R$ 64 750.3 and R$ 65 526.1 against R$ 64 708.1, and each drawing's own best plan cleans only twice, at the shutdown and at one later October, for R$ 27 420.9 and R$ 28 689.4. Where the drawings differ is how many months a clean leaves below the limit: on the curving drawings the calendar costs from R$ 133 656.7 to R$ 289 639.6, and the largest is the drawing

Table 17.3: Each action's best plan over the family, priced on each leading drawing alone (expected loss in reais, discounted to today), and the action best on that drawing. Probabilities from equation (14.8).

Drawing	P_M [–]	clean now [R$]	clean at the shutdown [R$]	replace now [R$]	replace at the shutdown [R$]	wait [R$]	Best there
$(\mathcal{C}_1, 2)$	0.3428	110 142.9	133 656.7	**92 865.3**	103 688.1	1 090 432.1	replace now
$(\mathcal{C}_0, 2)$	0.2061	92 708.1	64 750.3	92 865.3	103 688.1	142 633.4	clean at the shutdown; own best plan 27 420.9
$(\mathcal{C}_1, T_{d,2})$	0.1606	176 941.5	289 639.6	**92 865.3**	103 688.1	1 624 747.5	replace now
$(\mathcal{C}_1, \text{none})$	0.1097	112 592.1	183 898.4	**92 913.6**	103 766.3	1 343 624.1	replace now
$(\mathcal{C}_0, \text{none})$	0.1038	92 708.1	65 526.1	92 974.3	103 980.7	415 519.9	clean at the shutdown; own best plan 28 689.4

that blames unit 2's sensor for March's step and so gives coil 1 the fastest acceleration. Which coil stepped moves the cost of cleaning; it does not move the action.

17.3 The answers (step 4)

Every answer is an expected loss or a probability over the draws of Section 17.2.3: each drawing's Laplace normal carried through the plan's rows by the factor (17.5) and mixed by the drawings' weights (14.8), and that value is the one reported. The full posterior by reweighted draws stands beside the expected losses of the actions and their comparisons, the months before the shutdown and the life of each visit. Money is in reais, discounted to today; a date is a window with its month and year, and bands are 5–95 %.

17.3.1 The action

Asset manager's question. Is restoring this unit's conductance worth more than its replacement, given what the remaining life band says? The unit is coil 1, and the restoration is a clean.

General form. A design question among actions is the expected consequence of each action and the action that minimizes it, equation (2.23),

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

Specialized form. The question is equation (2.23) with L the loss (17.1) of a plan, $p(\mathbf{z} \mid \mathbf{y})$ the mixture (17.6) of the factored posteriors (17.5), and each action a the class of plans of step 2:

$$\mathcal{R}(a) = \min_{\mathcal{P} \in a} \mathcal{R}(\mathcal{P}) = \min_{\mathcal{P} \in a} \Big\{ c_{\mathcal{P}} + \mathrm{SF} \sum_{t=37}^{96} \mathrm{df}(t)\ \Pr\big(G_{\mathrm{tot},t}(\mathbf{z}; \mathcal{P}) \le G^* \mid \mathbf{y}\big) \Big\}, \qquad a^* = \arg\min_a \mathcal{R}(a), \tag{17.8}$$

for a in {clean now, clean at the shutdown, replace now, replace at the shutdown, wait}, with every plan priced on the same draws. One choice enters here: *an action is priced by the best of its plans.* This lets each action carry the later cleans that suit it, so that cleaning is not charged for a plan that cleans too seldom, and it fails if the later visits of the best plan are not carried out as fixed, which step 7 takes up.

Building blocks. The 256 plans and their rows (17.3); the draws of Section 17.2.3 with the weights P_M/N; the loss (17.1) with the discount factors of Table 17.1; the minimum over each action's plans.

Each action's best plan. Table 17.4 gives each action's expected loss: each row is an action, with its best plan, its expected loss and Monte Carlo standard error, the loss split into the work and the expected shortfall, the expected discounted shortfall months, and the probability of any shortfall month in the horizon. The last row is no work in the horizon at all, the do-nothing counterfactual.

Table 17.4: Each action's best plan and its expected loss over the horizon (windows 37–96), by equation (17.8) on the draws of the Laplace normals mixed over the family. Money in reais discounted to today; "months" are the expected discounted shortfall months.

Action	Its best plan	$\mathcal{R}$ [R$]	se [R$]	work [R$]	shortfall [R$]	months [–]	P(any month) [–]
replace now (a^*)	now: replace coil 1, clean coil 2 (Jan Y4, off the outage); clean both at Oct Y6	**92 884.8**	6.0	92 865.3	19.5	0.0001	0.0001
replace at the shutdown	now: clean both; Oct Y4: replace coil 1, clean coil 2; clean both at Oct Y7	103 731.6	10.8	103 688.1	43.5	0.0003	0.0002
clean now	clean both now and at every October outage (the dashboard's schedule)	114 896.2	538.9	92 708.1	22 188.1	0.1479	0.0737
clean at the shutdown	clean both at every October outage (the calendar)	140 023.2	800.8	64 708.1	75 315.2	0.5021	0.1764
wait	clean both at Oct Y5, Y6, Y7 and Y8	901 025.7	2 790.4	49 702.0	851 323.7	5.6755	0.6812
no work in the horizon	–	4 751 362.5	5 914.9	0.0	4 751 362.5	31.6757	0.9761

The answer is to replace coil 1 now, in January of year 4, off the outage, cleaning coil 2 in the same visit, and to clean both coils at the October outage of year 6. Its expected loss is R$ 92 884.8 (standard error 6.0) over five years, and nearly all of it is the work, R$ 92 865.3; the full posterior gives R$ 92 877.7 (standard error 4.1). The two next-best plans also replace now, with the clean at the October of year 5 (R$ 94 356.0) or of year 7 (R$ 98 634.7). The runner-up action, replace at the shutdown after a clean now, is R$ 103 731.6, and the full posterior gives R$ 103 721.2. The expected losses are compared as bars in Figure 17.2: one row per action, in the table's order, each bar the work (blue) and the expected shortfall (red) stacked, with the expected loss and its standard error printed at its end; the two bars that run off the axis carry their values at its edge. The two replacement bars are almost all blue. The two cleaning bars are shorter in blue and longer in total: the cleans are cheaper, and the months they leave below the limit cost more than the difference.

Which coil to replace is part of the same minimum: the best plan that replaces coil 1 costs R$ 92 884.8, the best that replaces coil 2 R$ 150 738.6, and the best that replaces both R$ 144 891.4. "This unit" of the drill question is coil 1.

Restoring is cheaper per visit, but not per five years. A clean of coil 1 lands near its commissioned conductance and then falls again on coil 1's own trend. On the drawings where that trend accelerates, which hold 0.6344 of the posterior, cleaning every October still leaves

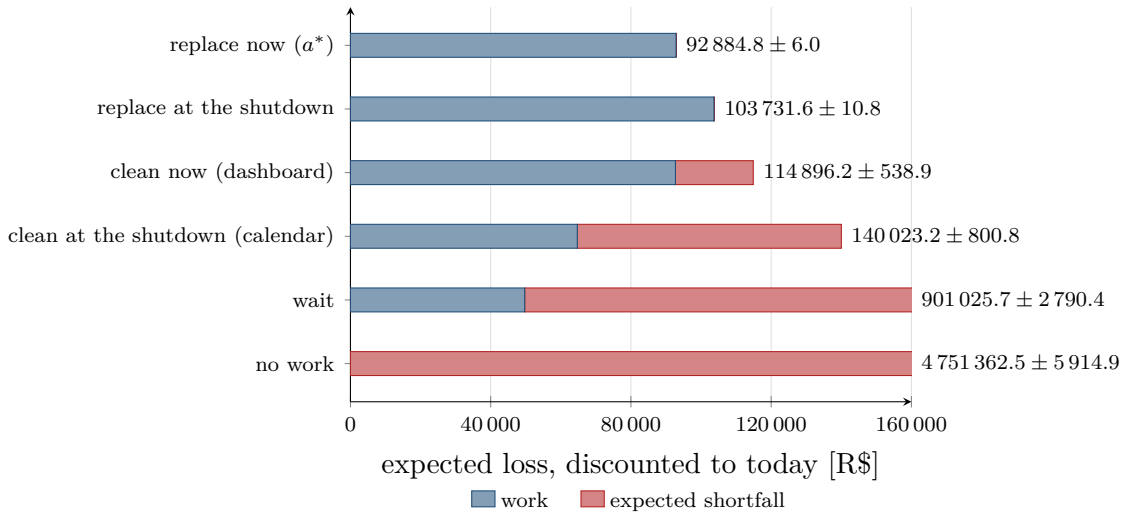


Figure 17.2: Each action's expected loss (Table 17.4), its best plan's work (blue) and expected shortfall (red) stacked, with the expected loss and its Monte Carlo standard error at the end of each bar. Wait and no work run off the axis; their values are printed at its edge. The dashboard's and the calendar's schedules are the best plans of clean now and clean at the shutdown.

months below the limit, and each costs R$ 150 000. Only the replacement takes the acceleration out of the path.

17.3.2 The firmness of the choice and the value of perfect information

Asset manager's question. How firm is that choice, and what would knowing the path be worth?

General form. How firm a ranking is, is the second line of equation (2.24), the probability that one action costs less than another at the true state,

$$\Pr\big(L(a,\mathbf{z}) < L(a',\mathbf{z}) \mid \mathbf{y}\big) = \int \mathbf{1}\big[L(a,\mathbf{z}) < L(a',\mathbf{z})\big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z},$$

whose normal reading there needs the gradient of the difference. What knowing the state would be worth is the expected value of perfect information, the expected loss of the action chosen now less the expected loss of the action chosen with the state known,

$$\mathrm{EVPI} = \mathcal{R}(a^*) - \mathbb{E}\Big[\min_a L(a,\mathbf{z}) \mid \mathbf{y}\Big] \geq 0, \tag{17.9}$$

which holds for any set of actions and any loss. It is nonnegative because $\min_a L(a,\mathbf{z}) \leq L(a^*,\mathbf{z})$ at every state, so the average of the left side is at most $\mathcal{R}(a^*)$. It is the limit of Chapter 2's expected value of sample information (2.52) for a measurement that reveals the whole state, so no measurement taken before the action can be worth more.

Specialized form. The probability is the second line of equation (2.24), read as the integral above because the gradient of a count of months is zero wherever it is defined (equation (17.4)): it is the weighted share of the paired draws of Section 17.2.3 on which $\mathcal{P}^*$ costs less than each action's best plan. Perfect information is equation (17.9) with a running over the 256 plans and the minimum taken on each draw. One choice enters here: *knowing the path* means knowing, before any visit, the draw's whole path, its forecast departures and the new coil's rate. This

makes equation (17.9) a ceiling over every measurement, and it is not itself a measurement; which open question comes near it is step 7's.

Building blocks. The paired draws and their weights; the losses of the five actions' best plans on each draw; the minimum of the 256 losses on each draw.

The chosen plan against the others. Table 17.5 sets the chosen plan against each other action's best plan: each row is an action, with the difference of the expected losses and its standard error on the paired draws, and the probability that the chosen plan costs less or more on the draws.

Table 17.5: The chosen plan against each action's best plan, on the paired draws; the difference is the other action's expected loss less $\mathcal{R}(a^*)$.

a^* against	$\mathcal{R}$ difference [R$]	se [R$]	P(a^* costs less) [–]	P(a^* costs more) [–]
replace at the shutdown	+10 846.7	9.8	0.9999	0.0001
clean now (the dashboard's schedule)	+22 011.4	536.9	0.0737	0.9263
clean at the shutdown (the calendar)	+47 138.4	798.8	0.1764	0.8236
wait	+808 140.9	2 791.0	0.6812	0.3188
no work in the horizon	+4 658 477.6	5 914.9	0.9761	0.0239

Against its runner-up the choice is close to certain, 0.9999: both plans hold the hall, and the runner-up pays for an extra clean. Against the two cleaning schedules it is the other way round. The chosen plan costs more than the calendar in 0.8236 of the outcomes and more than the dashboard's schedule in 0.9263; it costs less only on the draws where cleaning leaves a shortfall month, 0.1764 and 0.0737 of them. The full posterior gives 0.1728 and 0.0740. Replacement is insurance. Its loss is nearly certain, with a standard deviation of R$ 2 270.4 over the draws, against R$ 217 557.3 for the calendar, whose loss is its work, R$ 64 708.1, at the median and R$ 592 880.3 at the 95 % quantile.

The distribution of the loss is plotted in Figure 17.3 as quantile curves: the horizontal axis is the probability level q, the vertical axis the loss at that quantile on a logarithmic scale, one curve for the chosen plan, the runner-up, the dashboard's schedule and the calendar, with points at the levels the draws were read at. To the right of the axis, the dots in the column marked $\mathcal{R}$ are the four expected losses, and the crosses in the column marked "truth" each plan's loss on the record's stated truth, a check marker (Section 17.4.7). The dotted line is the mean loss with perfect information. The two replacement curves are flat; the dashboard's lies just below the chosen plan's to $q = 0.90$ and the calendar's well below it to $q = 0.80$, and both then rise past it.

The value of knowing the path. With the path known before the first visit, the best of the 256 plans on each draw has a mean loss of R$ 35 876.6 (standard error 79.1), and by equation (17.9)

$$\text{EVPI} = 92\,884.8 - 35\,876.6 = \text{R\$ } 57\,008.3,$$

from the unrounded values, 61.4 % of $\mathcal{R}(a^*)$. Most of the chosen plan's cost is insurance against a path the record cannot yet tell apart from a mild one: on the paths that would not accelerate, two cleans would have been enough.

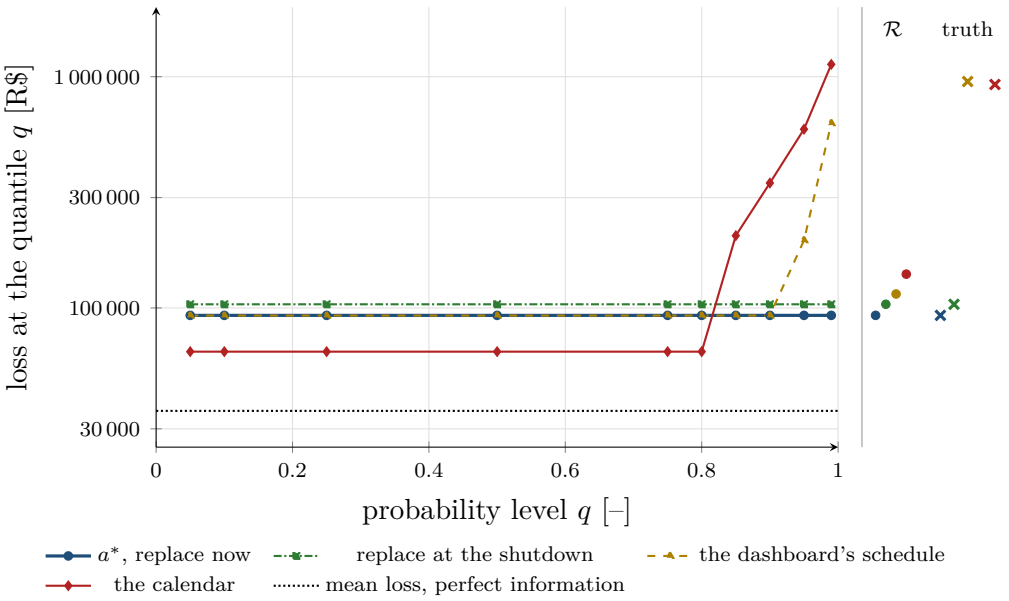


Figure 17.3: The loss of four plans over the draws, as quantile curves read at the levels 0.05, 0.10, 0.25, 0.50, 0.75, 0.80, 0.85, 0.90, 0.95 and 0.99 and joined by straight lines; logarithmic vertical axis. Dots (column $\mathcal{R}$): each plan's expected loss. Crosses (column "truth"): each plan's loss on the stated truth of the record, a check marker (Section 17.4.7). Dotted: the mean loss with perfect information, R$ 35 876.6.

17.3.3 The timing of the work

Asset manager's question. Should the work be done now, or can it wait for the October shutdown?

General form. The forecast of a quantity beyond the record against a level is the third line of the trend form (2.21),

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

and the choice between two times for the work is the expected consequence (2.23) of each.

Specialized form. The months before the shutdown are the third line of equation (2.21) with $f_t = G_{\mathrm{tot},t}$ and $f^* = G^*$ at each window from 37 to 46, with no work, and by equation (2.23) their expected cost is

$$S_{\mathrm{sd}} = \mathrm{SF} \sum_{t=37}^{46} \mathrm{df}(t)\ \Pr(G_{\mathrm{tot},t}(\mathbf{z}; \varnothing) \le G^* \mid \mathbf{y}), \tag{17.10}$$

with $\varnothing$ the plan of no work. A visit now pays OW and takes those months away; waiting for the shutdown pays them and saves OW. The comparison of the full plans is the minimum of equation (17.8) over the plans that start now against those that start at the shutdown. One choice enters here: *each month before the shutdown is priced by its own window's probability*, so a path that dips below the limit and comes back is charged for the months it spends below. This is SHORTFALL's unit, and it differs from the first passage of Chapter 14 only by the paths that come back.

Building blocks. The draws with no work; the probability of each window from 37 to 46; the discount factors; SF and OW; the actions' best plans of Section 17.3.1.

The months before the shutdown. With no work before the October of year 4, the hall is expected below G^* for 0.3734 months of windows 37 to 46, 0.3524 discounted, and by equation (17.10) that is $S_{\text{sd}} = \text{R\$ } 52\,862.5$ of expected shortfall, against $\text{OW} = \text{R\$ } 12\,000$. The probability climbs month by month: 0.0003 at window 40 (Apr Y4), 0.0285 at 43 (Jul Y4), 0.0613 at 44 (Aug Y4), 0.1078 at 45 (Sep Y4) and 0.1632 at the shutdown's own window, 46 (Oct Y4). The full posterior gives 0.3431 discounted months, R\$ 51 472.3, and a crossing by window 46 with probability 0.1588 (standard error 0.0017). Replacing at the shutdown with nothing before it costs R\$ 128 682.5; a clean now that covers the months to October takes it down by R\$ 24 950.9, to the runner-up's R\$ 103 731.6, and replacing now instead takes R\$ 10 846.7 more off. Whatever the action, the work starts now: every action's best plan in Table 17.4 begins with a visit in January, except the two that cannot, clean at the shutdown and wait.

Chapter 14's log planned the cleaning for the shutdown. Priced, that plan is the calendar, the dearest plan short of waiting: the one chance in six of a crossing before October costs R\$ 52 862.5 in expectation, against R\$ 12 000 for working away from the outage. The premium at which the shutdown would win is not that shortfall, because the plans that start at the shutdown differ from a^* in more than the months before October. The best of them, replacing coil 1 at the shutdown with nothing before it and cleaning both at the October of year 7, has an expected shortfall of R\$ 52 994.4 over the whole horizon. Against a^* it also pays the replacement visit later, which discounting makes cheaper by $(68\,000)(1 - 0.9379) = \text{R\$ } 4\,224.2$, and its later clean a year later, cheaper by $16\,000\,(0.8041 - 0.7445) = \text{R\$ } 953.0$, while a^* keeps an expected shortfall of its own, R\$ 19.5. What is left,

$$52\,994.4 - 4\,224.2 - 953.0 - 19.5 = \text{R\$ } 47\,797.6,$$

each term from the unrounded discount factors, is the break-even OFF_WINDOW of step 5, against the stated R\$ 12 000.

17.3.4 The life each action buys

Asset manager's question. How long does each action hold the hall?

General form. The distribution function of a crossing time is the second line of the trend form, equation (2.21),

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

read by inversion at the levels of the credible interval (2.11), 0.05, 0.5 and 0.95.

Specialized form. The question is equation (2.21), second line, with the first passage of Chapter 14, equation (14.12), taken from the visit on the path a one-visit plan $\mathcal{P}_{\text{v}}$ leaves:

$$\Pr(t_c(\mathbf{z}; \mathcal{P}_{\text{v}}) \le s \mid \mathbf{y}) = \sum_{M} P_M \int \mathbf{1}\Big[\min_{t_{\text{v}} < t \le s} G_{\text{tot},t}(\mathbf{z}; \mathcal{P}_{\text{v}}) \le G^*\Big]\, p(\mathbf{z} \mid \mathbf{y}, M, \mathcal{P}_{\text{v}})\, \mathrm{d}\mathbf{z}, \qquad (17.11)$$

with t_c the first window after the visit at or below the limit, as in Table 17.2, and the posterior the factor (17.5). A path that has not crossed by window 96 is recorded as beyond the horizon, so a band end can read "> 96". One choice enters here: *one visit and no later work*. This reads how long a visit holds the hall on its own, the life it buys; it is not a plan anyone would follow, and its expected loss is the drill's integral read literally (step 5).

Building blocks. The draws of Section 17.2.3 carried through one visit's rows (17.3); the first window after the visit at or below G^*; the weighted quantiles.

The next crossing after one visit. Table 17.6 gives, for each one-visit plan, the next crossing's median and band and the probabilities of a crossing by window 70 (Oct Y6) and by the end of the horizon.

Table 17.6: The next crossing after one visit and no later work, by equation (17.11): median and 5–95 % band as windows with their dates, and the probability of a crossing by window 70 (Oct Y6) and by window 96 (Dec Y8).

One visit	Median [window]	Band [window]	By 70 [–]	By 96 [–]
clean both now	75 (Mar Y7)	57 (Sep Y5) to > 96	0.4324	0.7632
clean both at Oct Y4	81 (Sep Y7)	63 (Mar Y6) to > 96	0.2820	0.6718
replace coil 1 now, clean coil 2	> 96	88 (Apr Y8) to > 96	0.0001	0.1425
replace coil 1 at Oct Y4, clean coil 2	> 96	> 96 to > 96	0.0000	0.0436
no work	53 (May Y5)	44 (Aug Y4) to 82 (Oct Y7)	0.8831	0.9761

The distribution functions are drawn in Figure 17.4: the horizontal axis is the window s from 48 to the end of the horizon, the vertical axis the probability that the next crossing has come by s, one curve per one-visit plan, with the October outage of year 6 dashed. The no-work curve rises first and steepest. The two cleans rise from windows 52 and 60 and both cross one half in year 7. The two replacement curves stay below 0.03 until window 84.

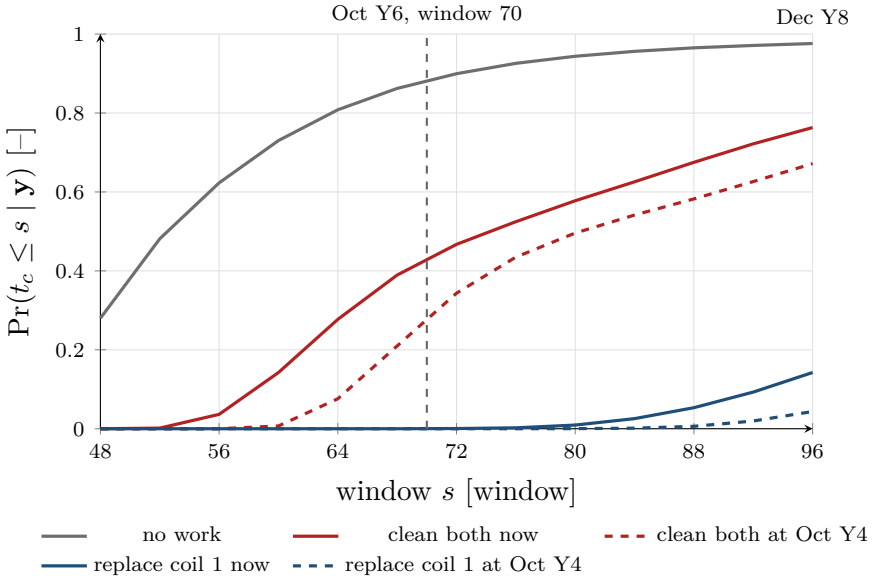


Figure 17.4: The distribution function of the next crossing after one visit and no later work, equation (17.11), read at windows 48, 52, . . . , 96. Red: a clean of both coils now (solid) or at the shutdown (dashed). Blue: coil 1 replaced and coil 2 cleaned, now (solid) or at the shutdown (dashed). Gray: no work. Dashed vertical: the October outage of year 6, window 70; the axis ends at the horizon, window 96.

On the full posterior every median and band end of Table 17.6 falls in the same window, and the probabilities of a crossing by window 96 are 0.7662, 0.6734, 0.1419, 0.0424 and 0.9774 in the table's order. Split by coil 1's curvature, a clean at the shutdown holds the hall to window 72 (Dec Y6) at the median given $\mathcal{C}_1$, with a crossing by window 96 with probability 0.9197, and

past the horizon given $\mathcal{C}_0$, with 0.2416. A replacement at the shutdown holds past the horizon in either case, crossing by window 96 with probability 0.0341 given $\mathcal{C}_1$. This is what "given what the remaining life band says" means for the decision. The band's dependence on the curvature, which moved Chapter 14's date far more than the step did, moves the median life a clean buys from the end of year 6 to beyond the horizon, and with it the action. A replacement buys the horizon either way.

17.3.5 The dashboard's and the calendar's schedules

Asset manager's question. What do the dashboard's and the calendar's schedules cost?

General form. A fixed plan is priced by the expected consequence (2.23), $\mathcal{R}(\mathcal{P}) = \mathbb{E}[L(\mathcal{P}, \mathbf{z}) \mid \mathbf{y}]$.

Specialized form. The question is equation (2.23) for the two plans, each priced by equation (17.2) on the same draws as a^*, with the comparison of Section 17.3.2. The calendar, $\mathcal{P}_{\text{cal}}$, cleans both coils at every October outage, windows 46, 58, 70, 82 and 94. One choice enters here: *the dashboard's conclusion is read as a schedule*, "clean both coils now" followed by a clean of both at every October outage, $\mathcal{P}_{\text{dash}}$. The dashboard says only the first part; the reading gives it the later cleans that make it a five-year plan, and it fails if a planner acting on the chart would clean more often.

Building blocks. The two plans' rows (17.3); the paired draws; $\mathcal{R}(a^*)$ and the comparison of Table 17.5.

The two schedules. The calendar is expected to cost R$ 140 023.2 (standard error 800.8), R$ 47 138.4 more than a^* (standard error 798.8 on the paired draws), 1.507 times as much, and it is cheaper than a^* in 0.8236 of the outcomes. The dashboard's schedule is expected to cost R$ 114 896.2 (538.9), R$ 22 011.4 more (536.9), 1.237 times as much, and it is cheaper in 0.9263 of the outcomes. The full posterior gives R$ 139 254.1 and R$ 115 480.0. Doing no work in the horizon is expected to cost R$ 4 751 362.5, 51.153 times $\mathcal{R}(a^*)$.

Each naive schedule is the best plan of its own action: Table 17.4 shows that the calendar is the best plan that starts at the shutdown, and the dashboard's schedule the best that starts with a clean now. What they get wrong is the action, not the plan: neither considers a replacement. The dashboard reaches the right timing, work now, from the wrong reading, a hall already past its limit, which Chapter 14 found it is not; it still loses R$ 22 011.4 in expectation by not replacing.

17.4 The checks (step 5)

Each answer of Section 17.3 rests on the Laplace normal of each drawing, on the draws that carry it through the plans, on the stated costs, on the two choices of step 0 about where coil 1's acceleration lives, and on the family's weights; this hall lets each be tested.

17.4.1 The full posterior and the Monte Carlo errors

The reweighted draws of Section 17.2.3, 10 000 per drawing with the seeds $2700 + d$, have effective sample sizes from 95.3 to 98.8 %, and the importance-sampling evidence moves no drawing's weight by more than 0.0010; $\Pr(\mathcal{C}_1)$ becomes 0.6358. On the full posterior replace now reads R$ 92 877.7 (standard error 4.1) against the reported R$ 92 884.8 (6.0), replace at the shutdown R$ 103 721.2 (9.3) against R$ 103 731.6 (10.8), clean now R$ 115 480.0 (454.2) against

R$ 114 896.2 (538.9), the calendar R$ 139 254.1 (1 134.1) against R$ 140 023.2 (800.8), and wait R$ 898 849.3 (3 756.3) against R$ 901 025.7 (2 790.4). Each gap is of the size of the two routes' standard errors or smaller, so none is resolved, and no action changes rank. The probability that a^* costs less than the calendar is 0.1728 on the full posterior against 0.1764, and than the dashboard's schedule 0.0740 against 0.0737. The full posterior is not an exact route: it is the same rows without the Gaussian approximation, to its Monte Carlo error.

The Monte Carlo errors are the batch standard errors printed with each expected loss: 6.0 for a^*, 800.8 for the calendar. The two halves of the draws give the same order of the actions, with replace now at R$ 92 880.9 and R$ 92 888.8, replace at the shutdown at R$ 103 732.2 and R$ 103 730.9, clean now at R$ 114 477.0 and R$ 115 315.4, the calendar at R$ 139 467.7 and R$ 140 578.8, and wait at R$ 902 732.9 and R$ 899 318.5.

No check of the limit of exact rows is new here. The plan's rows (17.3) are forecast rows and add no physical row, and the fits are Chapter 14's, whose limits of exact rows hold (Tables 10.6, 11.10 and 12.7).

17.4.2 The expansion at the most probable path

The first line of equation (2.24) reads an expected loss as the loss at the MAP state plus half the trace of its curvature against $\mathbf{\Sigma}$. The loss counts months, so its curvature term is zero wherever it is defined, and the reading is the loss on each drawing's most probable path, with no wander, mixed over the drawings. It gives R$ 92 865.3 for replace now, R$ 103 688.1 for replace at the shutdown, R$ 92 708.1 for clean now, R$ 64 708.1 for the calendar, and for wait R$ 1 115 561.2 on the leading drawing and R$ 864 962.7 mixed. On the most probable paths no plan that does any work crosses the limit, so the expansion ranks the actions' plans by their work alone and puts the calendar first. It is the plug-in action of *Reliability Engineering by Derivation*, Chapter 30, and it fails because the loss is a step function of the path. On the most probable path every plan that does any work stays above the limit, so the months in which the plans differ all lie in the tail of the posterior, on paths that fall faster than the most probable one. The loss at a single path counts none of them; only an average over the posterior does. The draws of the Laplace normal, the reported route, are the Laplace posterior read correctly; the expansion is a reading of it that a count of months does not admit.

17.4.3 The stated costs

Table 17.7 moves each stated cost and re-prices every plan on the same draws. Each row is a variant; the columns give the chosen action, its expected loss, the runner-up and its expected loss.

The break-even value of each price, the value at which replace now stops being the answer, is REPLACE R$ 82 011, above which clean now wins; CLEAN R$ 5 057, below which clean now wins; SHORTFALL R$ 58 674, below which replace at the shutdown, with nothing now and no later clean, wins; and OFF_WINDOW R$ 47 798, above which replace at the shutdown, with nothing now, wins. The SHORTFALL plan costs its work, R$ 63 775.8, plus SHORTFALL times its 0.49591 expected discounted shortfall months, against a^*'s R$ 92 865.3 plus SHORTFALL times 0.00013; at a cheaper month, skipping the visit now and the later clean pays. Further down, at a quarter of the stated SHORTFALL, clean at the shutdown is the answer (Table 17.7). The discount rate does not change the action anywhere from 0 to 0.12. The four break-evens are set against the stated values in Figure 17.5: each row is a price on a logarithmic axis in reais, the dot is its stated value, the bar runs from it to its break-even, over the range in which replace now stays the answer, and the action that wins beyond the break-even is printed past the bar's end. On the logarithmic axis REPLACE's bar is the shortest: its break-even lies nearest its stated value, while SHORTFALL's and OFF_WINDOW's lie far from theirs.

Table 17.7: The stated costs moved one at a time, every plan re-priced on the same draws. Expected losses in reais discounted to today.

Cost moved	a^*	$\mathcal{R}(a^*)$ [R$]	Runner-up	Its $\mathcal{R}$ [R$]
reported	replace now	92 884.8	replace at the shutdown	103 731.6
CLEAN × 0.5 (4 000)	**clean now** (every October)	74 542.2	replace now	82 452.2
CLEAN × 2 (16 000)	replace now	113 750.1	replace at the shutdown	128 474.0
REPLACE × 0.5 (30 000)	replace now	62 884.8	replace at the shutdown	75 595.2
REPLACE × 1.5 (90 000)	**clean now**	114 896.2	replace now	122 884.8
REPLACE × 2 (120 000)	**clean now**	114 896.2	clean at the shutdown	140 023.2
SHORTFALL × 0.25 (37 500)	**clean at the shutdown**	79 051.0	replace at the shutdown	82 372.3
SHORTFALL × 0.5 (75 000)	replace now	92 875.1	replace at the shutdown	98 373.4
SHORTFALL × 2 (300 000)	replace now	92 904.4	replace at the shutdown	103 775.0
OFF_WINDOW 0	replace now	80 884.8	replace at the shutdown	91 731.6
OFF_WINDOW 24 000	replace now	104 884.8	replace at the shutdown	115 731.6
OFF_WINDOW 48 000	**replace at the shutdown** (nothing now)	128 682.5	replace now	128 884.8
DISCOUNT 0	replace now	96 024.2	replace at the shutdown	112 000.0
DISCOUNT 0.04	replace now	94 338.9	replace at the shutdown	107 630.1
DISCOUNT 0.12	replace now	91 623.3	replace at the shutdown	100 272.1
HORIZON 3 yr (windows 37–72)	**clean now**	62 499.1	replace now	80 100.5
HORIZON 4 yr (windows 37–84)	**clean now**	88 158.3	replace now	91 792.0
only Oct Y4 an outage	replace now	102 533.8	replace at the shutdown	104 971.0

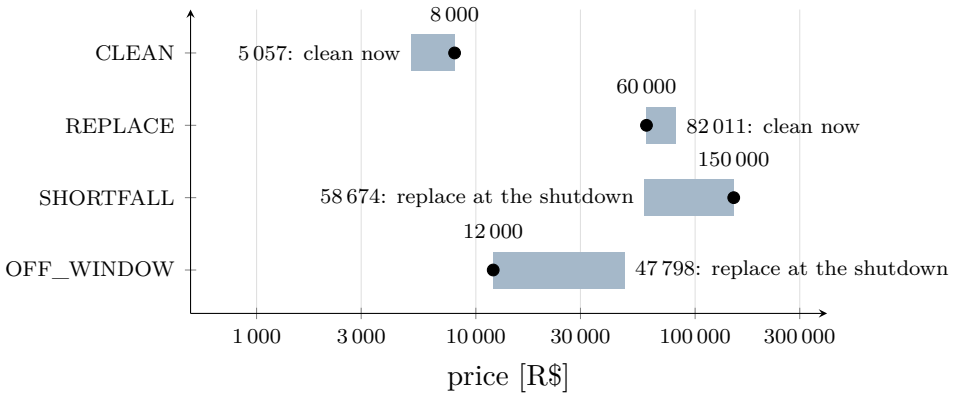


Figure 17.5: The break-even of each stated price. Dot: the stated value, with its number above. Bar: the range from the stated value to the break-even, over which replace now is the answer; past the bar's end, the break-even and the action that wins beyond it. Logarithmic axis. The horizon (four years) and the discount rate (none from 0 to 0.12) are not drawn; Table 17.7 gives them.

The decision is close on REPLACE and on HORIZON. The horizon flips it because the plan gives the new coil no value after window 96, the second choice of step 2: over three or four years a new coil has less time to repay its price than its life will give it. SHORTFALL can be halved and OFF_WINDOW doubled without moving the action (Table 17.7), and CLEAN would have to fall below R$ 5 057. With the later Octobers off-window, the cost schedule's one outage read literally, a^* is unchanged at R$ 102 533.8.

17.4.4 The plan's form and where the acceleration lives

One visit and no later work. The drill's integral, read literally, prices one visit followed by every month below the limit to the horizon. A clean then costs from R$ 1 723 902.0 (at the shutdown) to R$ 2 240 458.7 (now), and replacement wins by default: replace at the shutdown, coil 1 with coil 2 cleaned, at R$ 138 162.0, and replace now at R$ 167 926.6 with both coils replaced; doing nothing costs R$ 4 751 362.5. No planner cleans once in five years, so the plan with its later cleans is the honest comparison, the choice of Section 17.3.1.

The restore fraction. With RF = 0.70 and RF = 0.95 the answer is replace now in both, at R$ 92 973.9 and R$ 92 868.7; clean now is R$ 155 852.8 and R$ 103 792.6. A clean that restored 95 % would still leave coil 1 on its accelerating trend. These variants, and the two below, re-price the 128 plans that replace coil 1 or no coil, since a plan replacing coil 2 or both is never best in step 4.

A clean also ends coil 1's acceleration. If a clean resets the fouling, so that after it coil 1's rate is held at its value at the clean, the answer becomes clean now: R$ 66 533.2, cleaning both coils now and at the Octobers of years 5 and 7, against replace now at R$ 92 884.8.

A new coil fouls on the old coil's own trend. If the cause of the acceleration is the site, the new coil takes the old one's rate and curvature, and the answer becomes clean now at R$ 114 896.2, the dashboard's schedule; replace now rises to R$ 154 421.0 and replace at the shutdown to R$ 155 411.1.

Both are physical questions about where coil 1's acceleration lives: whether it is a property of the fouling layer that a clean removes, and whether it is the coil's or the site's. Before acting, no row of the record can answer either, because it has seen one coil 1 and no clean; after a

replacement, the new coil's windows will bear on the second. A replacement pays only if the acceleration is the coil's and survives a clean. This is the decision's central limit, and the report says so.

17.4.5 The family's open questions and the forecast's lateness

Table 17.8 reweights the drawings on the same draws to answer the decision within each open question of the family: each row is a group of drawings, with its probability, the action best within it, its expected loss and the runner-up.

Table 17.8: The decision within groups of the family's drawings, the drawings reweighted within each group on the same draws. Expected losses in reais discounted to today.

Drawings	P [–]	a^*	$\mathcal{R}(a^*)$ [R$]	Runner-up [R$]
constant rates only ($\mathcal{C}_0$)	0.3656	**clean at the shutdown** (and at Oct Y6)	28 236.2	clean now 41 407.0
coil 1 curving only ($\mathcal{C}_1$): the forecast as late as fresh records show	0.6344	replace now	92 875.7	replace at the shutdown 101 777.0
no step at window 15	0.2136	replace now	92 943.1	clean now 102 923.0
coil 2 stepped at window 15	0.5489	replace now	92 865.3	replace at the shutdown 94 621.8
unit 2's sensor stepped at window 15	0.1991	replace now	92 865.3	replace at the shutdown 95 634.5

Chapter 14's step 5 found that at the record's stated truth the three-year forecast is late in two records of three, and that the refit at window 45 decides coil 1's curvature, with $\Pr(\mathcal{C}_1)$ at a median of 0.9991 over the continuations (Section 14.6.3). A forecast as late as fresh records show is one in which coil 1 is curving, the $\mathcal{C}_1$ row of Table 17.8. There a^* does not flip; it hardens, at R$ 92 875.7 against the runner-up's R$ 101 777.0. It flips on the other side. Holding the drawings within each curvature state and moving $\Pr(\mathcal{C}_1)$, the answer is clean at the shutdown below $\Pr(\mathcal{C}_1) = 0.2437$ and replace now above it. The expected losses of the actions along that line are plotted in Figure 17.6: the horizontal axis is $\Pr(\mathcal{C}_1)$, the vertical axis each action's expected loss, with the switch and today's 0.6344 marked. Replace now is flat, from R$ 92 868.0 to R$ 92 875.7; every cleaning action rises with $\Pr(\mathcal{C}_1)$, and the calendar's line is the steepest. Replace at the shutdown lies below replace now up to $\Pr(\mathcal{C}_1) = 0.2$ and then levels off above it.

The window-45 refit therefore settles which side of 0.2437 the record is on before the October shutdown, and step 7 prices waiting for it. The kind of March's step does not move the action: within every step group the answer is replace now, at R$ 92 865.3 to R$ 92 943.1.

17.4.6 The tube leak

The cost schedule states no cost of a tube leak, and the reported loss prices none. The drill's integral names "failure" beside shortfall, so the leak is set beside the decision on Chapter 16's failure curve, read by its exact route (Section 16.2.3). The expected leak rate is $0.0768\,\mathrm{yr}^{-1}$ for a new coil, 0.0972 at age 3 and 0.1534 at age 8; the probability of a leak in the horizon is 0.4541 for a coil aged 3 today, coil 1, and 0.3725 for a new coil. Over the horizon the expected discounted number of leaks is 0.5006 for coil 1 kept, 0.3860 for coil 1 replaced now and 0.3821 for coil 1 replaced at the shutdown, its old self to October and new after. Priced at REPLACE, a leaking coil replaced, which is the owner's price to state and not the schedule's, a^* is unchanged: replace now at R$ 146 078.1 against replace at the shutdown at R$ 156 690.4. The leak adds $(0.3860 + 0.5006) \times 60\,000 =$ R$ 53 193.3 to a^*'s plan, coil 1 new and coil 2 kept, and

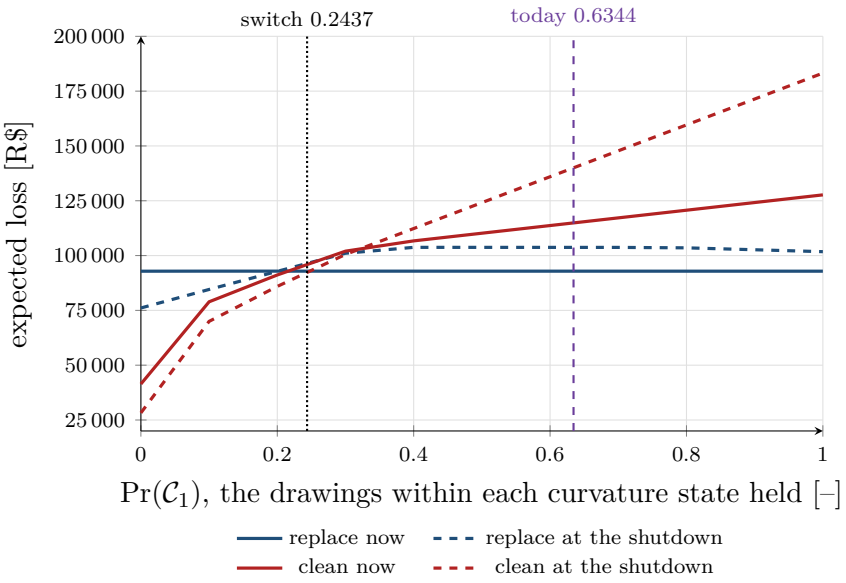


Figure 17.6: Each action's expected loss against the probability of coil 1's curvature, with the drawings within each curvature state held, read at $\Pr(\mathcal{C}_1) = 0, 0.1, \ldots, 1$ and joined by straight lines. Blue: the replacements; red: the cleans; solid now, dashed at the shutdown. Dotted: the switch, 0.2437, below which clean at the shutdown is the answer. Dashed purple: today's 0.6344. Wait lies above the axis throughout, from R$ 227 635.6 at 0 to R$ 1 275 320.5 at 1.

$2 \times 0.5006 \times 60\,000 =$ R$ 60 071.5 to the best clean plan, both coils kept, from the unrounded values. The failure term of the drill's integral leans toward replacement and does not decide it.

17.4.7 The answers against the stated truth

The last check sets the answers against the stated truth of the record (Section 17.6.5). Table 17.9 collects the routes: each row is a route or a variation, the columns the expected loss of each action's best plan, and the last column the action it chooses. The truth's row gives the loss of each of those four plans on the true path.

Table 17.9: The routes to the decision and the stated truth: each action's best plan of Table 17.4, priced by each route, in reais discounted to today. The one-visit row prices one visit and no later work instead.

Route	replace now [R$]	replace at the shutdown [R$]	clean now [R$]	clean at the shutdown [R$]	a^*
Laplace draws, mixed (reported)	92 884.8	103 731.6	114 896.2	140 023.2	replace now
Full posterior, reweighted draws	92 877.7	103 721.2	115 480.0	139 254.1	replace now
Expansion of (2.24) at each drawing's MAP path	92 865.3	103 688.1	92 708.1	**64 708.1**	**clean at the shutdown**
One visit and no later work (the drill's integral read literally; replace now replaces both coils)	167 926.6	138 162.0	2 240 458.7	1 723 902.0	replace at the shutdown
The stated truth (step 5 only)	92 865.3	103 688.1	953 624.7	925 624.7	replace now, of the four

The true coil 1 is losing 13.93 %/yr today and accelerating, and the true total conductance

at the shutdown is 10.8071 kW/K, above the limit. On the truth, a^*'s plan costs R$ 92 865.3, its work, with no shortfall month, and the runner-up's R$ 103 688.1, also with none. The calendar and the dashboard's schedule each leave 8 months below the limit, at R$ 925 624.7 and R$ 953 624.7: the truth's acceleration, in calendar time, outruns an annual clean. Wait leaves 21 months (R$ 2 650 881.0) and no work 50 (R$ 5 998 478.2). In hindsight the best plan is to replace coil 1 at the shutdown with nothing before it, R$ 63 775.8, because the true path crosses after the shutdown, not before (Section 14.4.6); that outcome lies in the part of the posterior that survives to October, which the chosen plan pays to insure against. A single clean at the shutdown would see the truth cross again at window 62; a single replacement at the shutdown would see no crossing by window 96.

17.5 The report (step 6)

The calendar cleans both coils every October and is expected to cost R$ 140 023 over five years; the dashboard's schedule cleans now and every October, at R$ 114 896. Replacing coil 1 now, cleaning coil 2 in the same visit and both coils again at the October outage of year 6 is expected to cost R$ 92 885. A replacement is insurance: it costs more than the calendar in more than four outcomes out of five, but the calendar's bad outcomes are months without the design day, and they cost R$ 150 000 each. Whatever is decided, the work starts in January: the one chance in six of a crossing before October costs more, in expectation, than working away from the outage. The decision rests on coil 1's acceleration being the coil's own and surviving a clean. If the acceleration stops at a clean, or would carry over to a new coil, cleaning now is the answer. It rests also on the stated prices: a new coil above R$ 82 011, or a four-year horizon, makes cleaning the answer.

17.5.1 The answers and their bands

Table 17.10 gives the answers in the order of Table 17.2. Each row is a question; the second column gives the reported value from the draws of the Laplace normals mixed over the family, the third the full posterior's value beside it where it was computed, and the fourth the band, the probability or the break-even that qualifies the value.

The answers rest on two things neither chart gives. The plans are written into the forecast windows of Chapter 14's path, so the remaining-life band enters the price of each action through the months it leaves below the limit. And the posterior carries coil 1's undecided acceleration into the price, which is why the cheapest plan on the most probable path, the calendar, is not the cheapest in expectation.

17.5.2 The entry in the maintenance log

The decision of window 36 is the next entry of the two-unit hall's maintenance log, after Chapter 14's forecast at the same window. It uses nothing later than window 36, and in the values of Table 17.10 it reads:

> *Window 36 (Dec Y3), repair or replace, the two-unit hall's coils, over Chapter 14's forecast (windows 0 to 36), costs from the cost schedule, five-year horizon (windows 37 to 96).*
> Action: replace coil 1 in Jan Y4 (off the outage) and clean coil 2 in the same visit; clean both coils at the Oct Y6 outage. Expected cost R$ 92 885.
> Against: the calendar (clean both each October) R$ 140 023; the dashboard's schedule (clean now and each October) R$ 114 896; replace at the Oct Y4 shutdown after a clean now R$ 103 732; no work in year 4 R$ 901 026.
> Why now: crossing by the shutdown 0.16; expected shortfall to October R$ 52 863 against R$ 12 000 off the outage.
> Why replace: coil 1's acceleration (P 0.63) outruns annual cleaning. If coil 1's rate is

Table 17.10: The report on the two-unit hall's coils at window 36 (Dec Y3). Reported values from the draws of each drawing's Laplace normal, mixed over the family; money in reais discounted to today; bands are 5–95 %.

Asset manager's question	Reported value	Full posterior	Band, probability or break-even
Is restoring this unit's conductance worth more than its replacement, given what the remaining life band says?	no: replace coil 1 now with coil 2 cleaned, clean both at Oct Y6, $\mathcal{R} = 92\,884.8$ [R$]; runner-up 103 731.6 [R$]	92 877.7; 103 721.2 [R$]	se 6.0 [R$]; any shortfall month 0.0001 [–]; break-even REPLACE 82 011 [R$]
How firm is that choice, and what would knowing the path be worth?	beats the runner-up with 0.9999, the dashboard's schedule with 0.0737, the calendar with 0.1764 [–]; EVPI 57 008.3 [R$]	0.9999; 0.0740; 0.1728 [–]	sd of the loss 2 270.4 against the calendar's 217 557.3 [R$]; EVPI 61.4 % of $\mathcal{R}(a^*)$
Should the work be done now, or can it wait for the October shutdown?	now: expected shortfall to October with no work 52 862.5 against OW = 12 000 [R$]	51 472.3 [R$]; crossing by 46 0.1588 [–]	crossing by window 46 0.1632 [–]; break-even OFF_WINDOW 47 798 [R$]
How long does each action hold the hall?	clean at Oct Y4: next crossing window 81 (Sep Y7); replace coil 1 at Oct Y4: beyond window 96	window 81; beyond 96; crossing by 96 0.6734 and 0.0424 [–]	clean: 63 (Mar Y6) to > 96, window 72 (Dec Y6) given $\mathcal{C}_1$; replacement: crossing by 96 0.0436 [–]
What do the dashboard's and the calendar's schedules cost?	calendar 140 023.2, dashboard 114 896.2 [R$]	139 254.1; 115 480.0 [R$]	over a^* by 47 138.4 and 22 011.4 [R$]; cheaper than a^* in 0.8236 and 0.9263 of outcomes [–]

constant, clean at the shutdown and at Oct Y6 (R$ 28 236).
Holds unless: a clean resets coil 1's acceleration; the acceleration belongs to the site; REPLACE above R$ 82 011, CLEAN below R$ 5 057, SHORTFALL below R$ 58 674; the horizon shortened to four years.
Alternative, not taken: clean both coils now, refit at window 45 (Sep Y4), and replace coil 1 at the Oct Y4 shutdown if the refit confirms the curvature; R$ 90 637 on Chapter 11's power to decide the curvature (0.545). That figure is a reading of that power, not a computed preposterior of the refit.
Next measurement: continue the record; from window 37 it reads the new coil, whose rate against the design's bears on whether the acceleration is the site's (how many windows that takes was not computed). Ask the coil's maker and the site's water treatment whether coil 1's acceleration is the coil's.

The entry records what the action rests on as well as its price: the acceleration's weight in the family, the two physical assumptions the record cannot test, and the prices at which it would change. The alternative is recorded because it keeps the replacement decision open until the refit, at an expected cost close to the action's on a reading that is not this family's preposterior (Section 17.6.2).

Common mistake

Comparing repair and replacement by the price of one visit. A clean of both coils costs R$ 16 000, and a replacement visit R$ 68 000 with the R$ 12 000 premium of working now, but the clean leaves coil 1 on its own trend, and the months that trend leaves below the limit are priced at R$ 150 000 each. The comparison is of plans over the horizon, under the posterior of the path.

17.6 The next measurement (step 7)

The report leaves the action resting on coil 1's curvature, which the record has not decided, and on R$ 57 008.3 of value of perfect information. A measurement taken before the action is worth what it lets the plan change, and its value is a preposterior of the expected loss.

17.6.1 Partial perfect information on the family

Perfect information, equation (17.9), knows the whole path. A measurement that reveals one open question of the family before the work, and nothing else, is worth at most its partial perfect information. With the drawings grouped by that question into groups $\mathcal{M}_k$, each with its probability $\Pr(\mathcal{M}_k \mid \mathbf{y})$ and the posterior reweighted within it,

$$\mathrm{EVPPI} = \mathcal{R}(a^*) - \sum_k \Pr(\mathcal{M}_k \mid \mathbf{y}) \min_{\mathcal{P}} \mathcal{R}(\mathcal{P} \mid \mathcal{M}_k), \tag{17.12}$$

the expected loss of the action chosen now less that of the best plan chosen once the group is known. It lies between zero and the EVPI, and like the EVPI it is a limit of Chapter 2's expected value of sample information (2.52): that of a measurement that reveals the group exactly and nothing else. The groups are those of Table 17.8: coil 1's curvature state, $\mathcal{C}_0$ or $\mathcal{C}_1$; the March step's kind; and the drawing itself. The fifth candidate is the measurement the hall already makes: the monthly record continued to window 45 and refitted there, the permanent instrument at the place of the question.

17.6.2 The candidates and the deferred replacement

Table 17.11 ranks the candidates. Each row is a candidate, with its value in reais against a^* and the route that gives it.

Table 17.11: The value of learning each open question before acting, by equation (17.12), against $\mathcal{R}(a^*) =$ R$ 92 884.8; the refit's row is the deferred plan read with Chapter 11's power.

Candidate	Value [R$]	Route
the path known (the ceiling)	57 008.3	equation (17.9)
the drawing known	23 742.7	equation (17.12)
coil 1's curvature state known ($\mathcal{C}_0$ or $\mathcal{C}_1$)	23 641.2	equation (17.12)
the March step's kind known (Chapter 12's logger archive)	0.0	equation (17.12)
the record continued to window 45 and refitted, clean now and act at the shutdown (the permanent instrument, already installed)	2 247.3	a reading of Chapter 11's power 0.545, not a preposterior of this family

The March step's kind is worth nothing to this decision, and coil 1's curvature is worth R$ 23 641.2, nearly all of what knowing the drawing is worth. The value of the curvature depends on what is done while waiting for it.

Clean now, then decide. With the first move a clean of both coils now, and the rest chosen at the shutdown with the curvature known, the plan given $\mathcal{C}_0$ (0.3656) cleans again only at the October of year 6, for R$ 41 407.0, and given $\mathcal{C}_1$ replaces coil 1 at the shutdown, for R$ 101 777.0. Together they cost R$ 79 705.9, R$ 13 179.0 below a^*. That is a bound, since the curvature is known exactly; with nothing learned, the best plan that cleans now and replaces at the shutdown is the runner-up's R$ 103 731.6.

Nothing now, then decide. With no work now and the rest chosen at the shutdown with the curvature known, the plan given $\mathcal{C}_0$ cleans at the shutdown and at the October of year 6 (R$ 28 236.2), and given $\mathcal{C}_1$ replaces coil 1 at the shutdown and cleans at the October of year 7 (R$ 158 880.4). Together they cost R$ 111 117.2, R$ 18 232.4 above a^*: waiting for evidence pays only after a clean that covers the months to October.

The refit at window 45. The record continued to window 45 decides coil 1's curvature, $|\ln B| > 1$, in 0.545 of the predicted records (Chapter 11's step 7, Section 11.6.3). Read with that power, the deferred plan, clean now, refit, and act at the shutdown on what the refit decides, costs

$$0.545 \times 79\,705.9 + 0.455 \times 103\,731.6 = \text{R\$ } 90\,637.5,$$

R$ 2 247.3 below a^*, a decided record taking the bound's plan and an undecided one the runner-up's. A decided record is valued here at the bound for the curvature known exactly, which no record of nine windows reaches, so the reading is optimistic and its margin over a^* is not established; that is why a^* stands as the action. The power at which the deferral costs as much as a^* is 0.451. This is a reading of Chapter 11's power, not a preposterior of this family: a decided record is not a record that knows the curvature exactly, and the continuations were not refitted with this chapter's plans. The one measurement that bears on the decision is the one the hall already makes, and whether it is worth waiting for sits R$ 2 247.3 from a^*, on a power, 0.545, close to the tie, 0.451.

17.6.3 The recommendation

Do the work in January, whichever action is taken: no plan that waits for October is competitive. The computed answer, and the log's action, is to replace coil 1 now, clean coil 2 in the same visit and clean both at the October outage of year 6. The deferred plan, clean both now, refit at window 45 and replace coil 1 at the shutdown only if the refit confirms the curvature, is the recorded alternative. On Chapter 11's power it is R$ 2 247.3 below a^*, which is a reading of that power and not a computed preposterior of the refit, and it keeps the replacement open for the question of where the acceleration lives. Ask the coil's maker and the site's water treatment whether coil 1's acceleration is the coil's: that question, not a sensor, is what flips the decision in step 5. Before acting, no row can answer it; after the replacement, the new coil's windows will, because their rate set against the design's is the record's own test of the site. How many windows that test needs was not computed here.

17.6.4 What the analysis showed

The calendar, at R$ 140 023, and the dashboard's schedule, at R$ 114 896, both clean; replacing coil 1 now is expected to cost R$ 92 885. The analysis wrote each plan into the forecast windows of Chapter 14's path, a clean as a step toward the commissioned conductance and a replacement as a new chain at the rating, so that the remaining-life band entered the decision through the months each plan leaves below the limit. The chosen action loses to the calendar in most outcomes; it wins on the tail, and only the probability saw it, because the expansion at the most probable path put the calendar first. The curvature decides the action and the March step does not: replace if coil 1 accelerates, clean at the shutdown if it does not. The analysis also named the two physical assumptions the record cannot test, whether the acceleration survives a clean and whether it is the coil's, and showed that each one flips the answer.

17.6.5 Limits of the model

A constructed record. The condition record is generated from a stated truth, with the stated sensor errors, and step 5 compares the answers with that truth (Section 17.4.7). It is not

field data, and the agreement is a check of the method on its own model. On the truth, the chosen plan holds the hall, and the calendar leaves eight months below the limit.

Plans fixed today. Every plan is chosen now and not revised; the value of revising it is read only through step 7's bound and its reading of Chapter 11's power.

The horizon. The horizon has no terminal value, so a new coil's life after window 96 is not counted, and a four-year horizon flips the answer.

Where the acceleration lives. The clean keeps coil 1's trend and a new coil takes the design's rate. Both are stated choices, and each flips the decision when reversed.

The costs as stated. A tube leak is not priced, because the cost schedule states no cost for it; a month is a whole month; and the hall has one outage a year, each October.

Chapter 14's limits. The family is Part III's two open questions and no other, the forecast carries the record's wander, and at the truth the three-year forecast is late more often than early.

The next chapter. The plan fixes when the coils are worked on; Chapter 18 fixes how often the path is read between those visits.

17.7 What to retain

Definition 17.1 (Repair or replace as plans over a horizon)**.** A maintenance plan $\mathcal{P}$ is the work at each visit of the horizon, written into the forecast windows of the path's posterior: a clean of coil u at window t is a step on the trend row that leaves t, to $d^{+}_{u,t} = d_{u,0} - (1 - \mathrm{RF})(d_{u,0} - d_{u,t})$, and a replacement is a new chain at $g = 0$ on the design's rate. Its loss is the discounted work plus SF for each discounted month at or below the limit, equation (17.1), and its expected loss is the mixture over the family's drawings, equation (17.6). An action is priced by its best plan, equation (17.8); the probability that it is beaten is read on paired draws; and perfect information, equation (17.9), and partial perfect information, equation (17.12), bound what any measurement before the action is worth.

Three conclusions carry beyond the hall. A repair-or-replace decision is a comparison of plans over a horizon, and the remaining-life band enters it through the months each plan leaves below the limit, its tail more than its median. A loss that counts months has no gradient, so the expected consequence must be read by draws of the posterior; its expansion at the most probable path is the plug-in, and it chooses the cheapest work. And a decision can rest on physical questions that no row of the record touches, here whether a clean ends an acceleration and whether the acceleration is the coil's; the analysis names them and shows what each would change.

Table 17.12 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 17.12: The terms of the repair-or-replace decision and where each is defined. Units for the two-unit hall.

Term	Expression	Units	Meaning	Defined in
Loss of a plan	eq. (17.1)	R$	work plus discounted shortfall months	17.1.1
Discount factor	$\mathrm{df}(t) = 1.08^{-(t-36)/12}$	–	a cost at window t in today's reais	17.1.1
A clean on the forecast rows	eq. (17.3), $d^{+}_{u,t}$	–	a step toward the commissioned conductance	17.1.4
A replacement on the forecast rows	eq. (17.3), γ_{new}	yr^{-1}	a new chain at the rating	17.1.4
The posterior under a plan	eq. (17.5)	–	the record's posterior, the plan in the forecast	17.2.1
Questions and answers				
An action's expected loss	eq. (17.8)	R$	the best of its plans	17.3.1
Insurance action	$\Pr(L(\mathcal{P}^*) < L(\mathcal{P})) < \frac{1}{2}$ with $\mathcal{R}(\mathcal{P}^*) < \mathcal{R}(\mathcal{P})$	–	beaten in most outcomes, cheaper on average	17.3.2
Expected shortfall before an outage	eq. (17.10)	R$	against the off-window premium	17.3.3
Life a visit buys	eq. (17.11)	window	the next crossing after one visit	17.3.4
Expansion of a count loss	$L(\mathcal{P}, \hat{\mathbf{z}})$	R$	the plug-in; picks the cheapest work	17.4.2
Break-even cost	the price at which a^* changes	R$	how far a stated price may move	17.4.3
EVPI and EVPPI	eqs. (17.9), (17.12)	R$	the most a measurement can be worth	17.3.2, 17.6.1
Deferring a decision	clean now, decide at the shutdown	R$	its bound and its reading	17.6.2

17.8 Exercises

Exercise 17.1. *Application.* Price the work of a^*'s plan: coil 1 replaced and coil 2 cleaned in January, off the outage, and both coils cleaned at the October outage of year 6, discounted to today. ($68\,000 + 12\,000 + 16\,000 \times 0.8041 =$ R$\,92\,865.3$ with the discount factor unrounded.)

Exercise 17.2. *Application.* Price the calendar's work over the five October outages. ($16\,000 \times (0.9379 + 0.8684 + 0.8041 + 0.7445 + 0.6894) =$ R$\,64\,708.1$ with the discount factors unrounded.)

Exercise 17.3. *Derivation.* Show that "now or at the shutdown" turns on the expected shortfall before the shutdown against OFF_WINDOW, evaluate it, and find the premium at which the shutdown wins. (With no work, the months of windows 37 to 46 cost equation (17.10): 0.3524 discounted months $\times$ 150 000 = R$\,52\,862.5, against R$\,12\,000. The plans themselves differ by more: the best plan starting at the shutdown has R$\,52\,994.4 of shortfall over the horizon, pays its replacement R$\,4\,224.2 cheaper and its later clean R$\,953.0 cheaper than a^*, and a^* keeps R$\,19.5 of its own, so the break-even premium is $52\,994.4 - 4\,224.2 - 953.0 - 19.5 =$ R$\,47\,797.6$, all from the unrounded values.)

Exercise 17.4. *Derivation.* From the clean's row (17.3), give coil 1's derate factor and log-derate after a clean now at the leading drawing's MAP, and say which entry's spread it inherits. ($d^{+} = 0.9793 - 0.15\,(0.9793 - 0.8473) = 0.9595$, $g = -0.04138$; its gradient puts RF = 0.85 on $g_{1,0}$ and 0.15 on $g_{1,36}$, so it inherits mainly $g_{1,0}$'s.)

Exercise 17.5. *Formulation.* Write perfect information for the set of plans, say why it is

nonnegative, and evaluate it. (Equation (17.9) with a the 256 plans; $\min_{\mathcal{P}} L(\mathcal{P}, \mathbf{z}) \leq L(\mathcal{P}^*, \mathbf{z})$ at every state; EVPI $= 92\,884.8 - 35\,876.6 =$ R\$ 57 008.3 from the unrounded values.)

Chapter 18

Inspection Intervals

18.1 The calendar's condition test and the hall's last year before the limit

It is December of year 3, window 36, and the two-unit hall's coils have a forecast. Chapter 14 put the first crossing of the capacity study's limit, $G_1 + G_2 \geq 10.761\,\mathrm{kW/K}$, in May of year 5 at the median, window 53, with a band from August of year 4 to October of year 7. Until the coils are cleaned, what would show the crossing coming is a condition test: a crew visit that repeats the commissioning test with reference instruments, at R$ 1 500 a coil and R$ 3 000 for the pair. The maintenance plan has one on the calendar every December. The next, December of year 4 (window 48), comes before the forecast's median, window 52.55. The condition-monitoring dashboard cannot help with the timing: its line through the health-index tiles says the hall has been past the limit since April of year 2 (Chapter 14), and it reads 7.31 kW/K today, below any level a test could watch for. The conclusion the calendar invites is to keep the annual test, because December's comes in time. A second planner takes the handbook's P–F interval for air-side fouling of a cooling coil, six months, and the rule of thumb "inspect at half the P–F interval", and proposes a condition test every three months, R$ 12 000 a year, short enough that no crossing is missed.

Neither planner can say how long this hall's P–F interval is. The time from the level at which a test would raise the work order to the capacity limit depends on how fast the coils fall there, and Chapter 14's forecast carries that rate with its band and with coil 1's undecided acceleration. Neither can say how likely a crossing is to slip between two tests, when the test reads the level with an error and the crossing date has a band of three years. And neither can say what an interval costs, when a month below the limit costs R$ 150 000 and a visit R$ 3 000. The question of this chapter is

the probability that the hall crosses the capacity limit before any condition test has raised the work order, for each interval between tests, read as the posterior mean over the forecast's paths of the probability that every test before the crossing read above the potential-failure level; the longest interval that keeps it under 5 %; and the expected cost of each interval under the same posterior.

It is answered with the method of Chapter 5, on Chapter 14's extended state with the condition test's rows added at the test windows, and with Chapter 2's expected consequence.

Jargon: potential failure, P–F interval of a path, undetected crossing, phase of a schedule, dated schedule, inspection cost per year, expected cost of a program, early work order

The *potential failure*, the P of the P–F interval, is here a level of the coils' total conductance at which a condition test raises the work order; the functional limit is the F. The *P–F interval of a path* is the time from its first passage below the potential-failure level to its first passage below the functional limit. A crossing is *undetected* when no test before it raised the order. A schedule tests at a fixed interval from its first test, and the window of that first test is its *phase*; the *dated schedule* puts the first test one interval after today. The *inspection cost per year* is the price of a visit times the visits a year. The *expected cost of a program* is the posterior mean of what its visits and its missed months cost from today to the order. An *early work order* is raised by a test's error while the path is still above the potential-failure level. Chapter 2 defined the expected consequence and Chapter 17 perfect information; Chapter 14 defined the first passage, the forecast window and the family. The box does not repeat them.

Key idea

An inspection interval is an action, and its risk belongs to the paths the forecast carries, not to the date alone. The P–F interval is a margin divided by a rate of fall, so it has a band, and on this hall its short end comes with coil 1's acceleration. A crossing goes undetected when the P–F interval falls between two tests, or when a test's error reads a path near the potential-failure level as above it. Both are integrated over the posterior: the test's error in closed form and the paths by the forecast's draws. The interval that meets a 5 % limit and the interval the costs choose are two answers read from one computation.

18.1.1 Catching a path between two levels

Chapter 14 read the crossing of the limit G^* as the first passage t^* of the coils' total conductance $G_{\mathrm{tot},t} = G_{1,t}+G_{2,t}$, equation (14.1): the first window t_c at or below the limit, with t^* interpolated in $(t_c - 1,\ t_c]$. A condition test does not wait for the limit. It raises the work order when it reads the total at or below a higher level, the potential-failure level $G_{\mathrm{P}} = G^* + m_{\mathrm{P}}$, with m_{P} the margin, so that the coils can be cleaned before the hall runs short. The same first passage at G_{P} gives the window of the potential failure, and the P–F interval of a path is the difference of the two:

$$t_{\mathrm{P}}(\mathbf{z}) = \text{the first passage (14.1) with } G_{\mathrm{P}} \text{ in place of } G^*, \qquad \Delta_{\mathrm{PF}}(\mathbf{z}) = t^*(\mathbf{z}) - t_{\mathrm{P}}(\mathbf{z}) \geq 0, \tag{18.1}$$

in months, one window being one month. It is never negative: every window at or below G^* is at or below G_{P}, so the path passes G_{P} first. On a path that falls steadily at the rate $|\mathrm{d}G_{\mathrm{tot}}/\mathrm{d}t|$ near the two levels, $\Delta_{\mathrm{PF}} = m_{\mathrm{P}}/|\mathrm{d}G_{\mathrm{tot}}/\mathrm{d}t|$: a margin over a rate of fall, a ratio, like the crossing date.

A schedule with interval Δ_{I} and first test at window s_1 tests at the windows

$$\mathcal{S}(\Delta_{\mathrm{I}}, s_1) = \{s_1,\ s_1 + \Delta_{\mathrm{I}},\ s_1 + 2\Delta_{\mathrm{I}},\ \ldots\}, \qquad s_1 \in \{37, \ldots, 36 + \Delta_{\mathrm{I}}\}, \tag{18.2}$$

and s_1 is its phase. With a perfect test and a steady fall a test raises the order exactly when it falls at or after t_{P}, and the crossing is caught when some test falls in $[t_{\mathrm{P}},\ t^*)$, an interval of length Δ_{PF}. The tests form a lattice of spacing Δ_{I}, and with its phase equally likely anywhere within one spacing, the first test at or after t_{P} lies a distance Δ_{s} after it, Δ_{s} uniform on $[0,\ \Delta_{\mathrm{I}})$. The crossing is caught when $\Delta_{\mathrm{s}} < \Delta_{\mathrm{PF}}$:

$$\Pr(\text{caught} \mid \Delta_{\mathrm{PF}}) = \Pr(\Delta_{\mathrm{s}} < \Delta_{\mathrm{PF}}) = \min(\Delta_{\mathrm{PF}}/\Delta_{\mathrm{I}},\ 1). \tag{18.3}$$

This is the catch probability of the delay-time model, *Reliability Engineering by Derivation*, Chapter 29 (inspection and maintenance intervals). It says that an interval no longer than the P–F interval catches every crossing, and that the P–F interval of the path, not of a handbook, is what the interval is set against.

A real test reads the level with an error. Let $\hat{G}_s$ be the test's own estimate of $G_{\mathrm{tot},s}$, normal about the path's value with the test's standard deviation σ_{I}, which step 3 computes from the test's rows. The test at window s raises the order when $\hat{G}_s \le G_{\mathrm{P}}$, so given the path it raises none with probability

$$p_{\mathrm{no},s}(\mathbf{z}) = \Pr(\hat{G}_s > G_{\mathrm{P}} \mid \mathbf{z}) = \Phi\!\left(\frac{G_{\mathrm{tot},s}(\mathbf{z}) - G_{\mathrm{P}}}{\sigma_{\mathrm{I}}}\right). \tag{18.4}$$

The tests' errors are independent given the path, and the crossing is undetected when no test before it raised the order:

$$\Pr(\text{undetected} \mid \mathbf{z}, \mathcal{S}) = \mathbf{1}\big[t^*(\mathbf{z}) \le t_{\mathrm{hor}}\big] \prod_{s \in \mathcal{S},\ s < t^*(\mathbf{z})} p_{\mathrm{no},s}(\mathbf{z}), \tag{18.5}$$

with $t_{\mathrm{hor}} = 276$ Chapter 14's horizon. As $\sigma_{\mathrm{I}} \to 0$ each factor becomes $\mathbf{1}[G_{\mathrm{tot},s} > G_{\mathrm{P}}]$, the product is one only if every test before t^* falls before t_{P}, and equation (18.3) returns for a steady fall. With an error the factors near G_{P} lie between zero and one: a test just after the potential failure can miss it, and a test just before it can raise the order early.

Principle 18.1. *Set an inspection interval from the posterior of the path's P–F interval, not from a handbook figure. Its short end, not its median, sets the interval, and the costs of a missed crossing and of a test set the limit on the probability of missing one.*

18.1.2 The two-unit hall carried past its record, with its tests (step 0)

The graph is Chapter 14's, Figure 14.1: the trend layer of both coils carried past window 36 by forecast windows, each with its two log-derates $g_{1,t}$, $g_{2,t}$, trend squares and no reading squares, the trend entries feeding the squares, and on every forecast window the requirement node $G_{\mathrm{tot},t} \ge G^*$. Behind it stand the 37 read windows of Chapter 10's graph (Figure 10.1), each Chapter 5's window with its five physical rows (5.2). The modeling choices of step 0 are Chapter 5's paragraph, cited through Chapter 14: steady state within each window, the airflow pinned with no discharge control, the season's quantities as known inputs and one conductance per coil, and Chapter 14's own choice that the forecast windows carry the record's trend rows with no maintenance until the crossing.

What the tests add to the graph. On each test window s of a schedule (18.2) the window gets back its load Q_s and the physical rows of Chapter 5's window, with the test's known inputs pinned, the chilled water at $T_s = 14\,^{\circ}\mathrm{C}$ and the other heat at $Q_o = 8\,\mathrm{kW}$, the commissioning test's, and the test's reading squares on them: the return air, both discharges and rack 1's inlet by the crew's reference instruments, and the hall's load meter. Rack 1's inlet row brings in the shared recirculation fraction β_1. On every forecast window a second requirement node, $G_{\mathrm{tot},t} \ge G_{\mathrm{P}}$, stands beside the first, with G_{P} pinned, a stated input. The decision node, the interval Δ_{I} and the first test s_1, says which windows are test windows. The shutdown is marked at window 46.

One choice of step 0. One choice is new here, and no question block owns it: *the work order reads the test's own estimate of* $G_{\mathrm{tot},s}$, not the record's posterior of the level. This buys a rule a crew can apply on the day of the visit, at a stated grade, with no refit of the record. It fails in two ways. The record's posterior of the level at the test window could be sharper than the test itself, when combining the two would read the level better than either; at a test window months past the record, window 46 for one, the forecast's spread of $G_{\mathrm{tot},s}$ is several times the test's, so the record as of window 36 adds little there, and step 3 sets the two side by side. And the program ignores the hall's own monthly readings after window 36: the record goes on being logged, and each new window would narrow the path between two tests. That reading is Chapter 19's: a decision level c on the posterior probability that the hall is below its limit at the end of a one-window lead time, refitted on the hall's own sensors every month and priced by false work orders and missed degradation. Read every month, that rule takes the place of this chapter's test-based order level for orders between tests, and a test of this chapter's program stays as a reference reading of the level.

The graph. Figure 18.1 draws the trend layer from window 36 to window 48 with the four-month schedule's two test windows, 40 and 44. The upper chain is coil 1's path and the lower chain coil 2's, with the black trend squares between neighbors; runs of windows are dashed boxes. Window 36, the last read window, keeps its faint block between the chains. Between the chains on each forecast window stands the red requirement node with its two levels. Above coil 1's chain each test window carries the test's block, its physical rows with five reading squares, linked to both coils of its window, to its own load node and to β_1, and the decision node feeds them from above.

The stated inputs. The *functional limit* is the capacity study's $G_{\mathrm{tot},t} \le G^* = 10.761\,\mathrm{kW/K}$. The *potential-failure level* is a stated input of this chapter, $G_{\mathrm{P}} = G^* + m_{\mathrm{P}}$ with $m_{\mathrm{P}} = 0.5\,\mathrm{kW/K}$, so $G_{\mathrm{P}} = 11.261\,\mathrm{kW/K}$, 4.65 % of G^* above the limit. Today's coils stand about $0.99\,\mathrm{kW/K}$ above G^* (Chapter 14), so the level lies about halfway between today and the limit; step 5 sweeps m_{P}. The *inspection* is the reference condition test, INSPECT in the costs, one per coil: the commissioning test's known inputs, the load as found and read by the hall's meter at 3.0 kW, and the four air temperatures read by reference instruments of 0.1 K with no drift. *One visit reads both coils*, because the limit is on their sum, and costs $C_V =$ R$ 3 000, two INSPECTs. A month below G^* costs SF = R$ 150 000 (SHORTFALL), and an order raised while the path is still above G_{P} costs FO = R$ 3 000 (FALSE_ORDER) in a sensitivity row. The cleaning that follows a detection, CLEAN at R$ 8 000, falls on every path whatever the interval and drops out of the comparison of intervals, and so does OFF_WINDOW, the R$ 12 000 premium of work done away from an outage, because a program carried to the order raises it on every path. Before the October shutdown it does not drop out: an order raised by a test before window 46 brings the work forward off the outage, and the tests before the shutdown carry that premium with the probability that they raise an order. No cost is re-chosen. The *inspection cost per year* of an interval is

$$\frac{12\,C_V}{\Delta_{\mathrm{I}}} \quad [\mathrm{R\$/yr}], \tag{18.6}$$

R$ 36 000 at one month, R$ 18 000 at two, R$ 12 000 at three, R$ 9 000 at four, R$ 7 200 at five, R$ 6 000 at six and R$ 3 000 at twelve. The two schedules of the opening, the calendar's 12 months and the rule of thumb's 3, are the baselines; the handbook's 6-month P–F interval behind the second is the planner's figure for air-side fouling, a stated input of the baseline and not a measurement of this hall. The dashboard's line, $13.3588 - 2.0164\,t_{\mathrm{y}}$ kW/K (Chapter 14), would lose the margin m_{P} in 0.248 yr, 2.98 months, but it already stands below both levels.

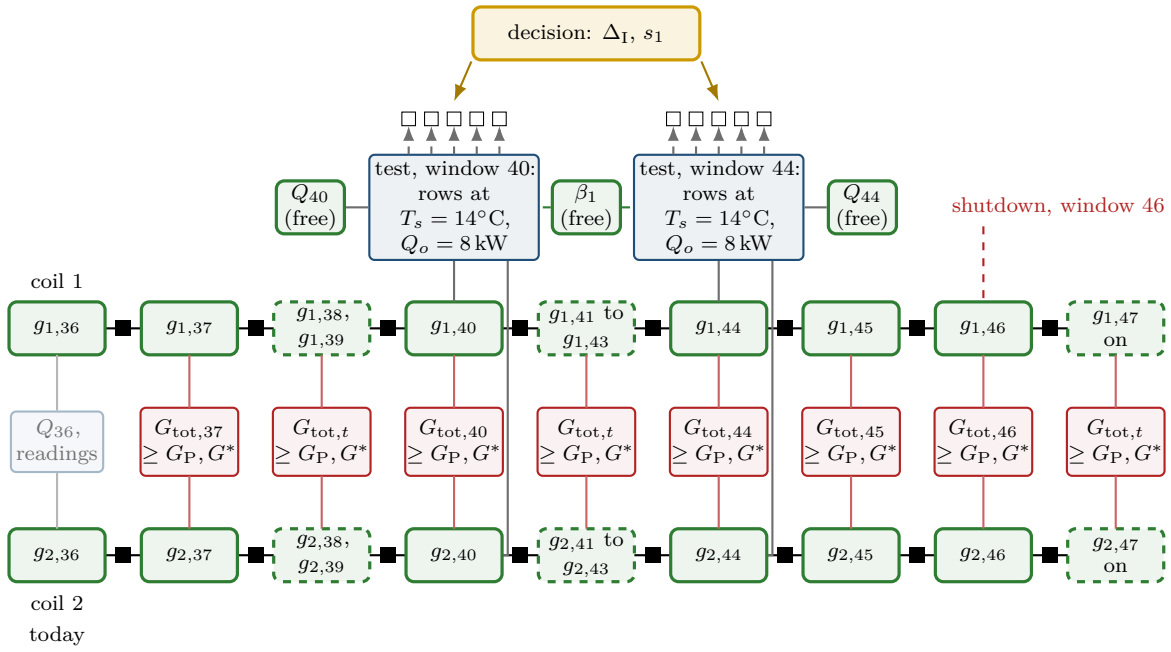


Figure 18.1: The trend layer of the two coils carried past the record, from window 36 to window 48, with the condition tests of the four-month schedule. Free (green): the log-derates $g_{u,t}$ (runs of forecast windows as dashed boxes), each test window's load Q_s, beside its block, and the shared recirculation fraction β_1 (its survey prior not drawn). Black squares: the trend rows (14.4); the trend entries that feed them are those of Figure 14.1, not drawn. Faint: window 36's load, physical rows and five readings. Blue blocks: a test window's physical rows (5.2) at the test's known inputs, $T_s = 14°C$ and $Q_o = 8\,kW$, pinned, with its five reading squares (observed): four temperatures at 0.1 K and the load meter at 3.0 kW. Red: the requirement nodes $G_{tot,t} \geq G_P = 11.261$ and $\geq G^* = 10.761\,kW/K$ on every forecast window, both levels pinned, stated inputs. Gold: the decision node, the interval and the first test. Dashed red: the shutdown at window 46.

The maintenance log so far. The two-unit hall's log holds the commissioning test of window 0; Chapter 5's condition assessment and Chapter 6's entry at window 18 (Jun Y2); Chapter 7's at window 25 (Jan Y3), which asked for a side-by-side check of unit 1's discharge sensor, a check not made; Chapters 10 to 13's entries at window 36; and Chapter 14's forecast at window 36: the crossing in May of year 5, band August of year 4 to October of year 7; the hall holds to the shutdown with 0.84; the cleaning planned for the October shutdown; the refit at window 45 before the work order is written. Then Chapter 17's decision at window 36: replace coil 1 in January of year 4, off the outage, and clean coil 2 in the same visit; its named alternative a deferral, both coils cleaned now and the replacement held for the October shutdown and decided by the refit at window 45; and, if coil 1's rate is constant, a plan that cleans at the shutdown. This chapter's entry follows at window 36: the condition-test program from today.

Which hall the program is for. The program is valued for the hall with no work done before the shutdown, the forecast of Chapter 14 unchanged. That is the plan Chapter 17's step 5 prefers if coil 1's rate is constant, which cleans at the shutdown, and it is also the calendar's plan. Under Chapter 17's action no test is needed before the January work, and the program for the renewed hall after it is not computed here (Section 18.6.4).

18.1.3 The asset manager's questions (step 1)

The asset manager asks five questions of the forecast, and step 1 writes each as a probability or an expectation over the extended state of the family's drawings, with the unknown on the left of the bar and the 185 readings of the record and the priors on the right. Table 18.1 lists them, with the columns of Chapter 5's table: the question in the asset manager's words, the drill question first; the general form of its category from Chapter 2; its mathematical form for the hall; its category; the computation of Section 2.3.6 that answers it; and the subsection where it is answered. The second question is answered first, because the first rests on the band it gives. Whether to clean at the shutdown, now or by replacing the coils is Chapter 17's question; at what posterior probability of a shortfall a monthly refit of the record should raise the order, priced by false work orders and missed degradation, is Chapter 19's; and whether the trend log sees a loss at all is Chapter 22's.

18.1.4 The evidence as rows (step 2)

Step 2 attaches two kinds of rows. The record's rows and the forecast rows are Chapter 14's, kept as they are: for each of the 18 drawings of its family, the 261 to 263 rows of the record, equation (10.5) and the rest of Chapter 10's fit, with coil 1's curvature and the March step where the drawing frees them, and the forecast rows (14.4) carried to the horizon at the wander $\sigma_q = 0.005$ per window. The new rows are the condition test's.

The condition test's rows. A test at window s reads one window of the hall, so its rows are Chapter 5's rows of one window, equation (5.7), at the reference grade and at the test's known inputs. On the test's entries $\mathbf{z}_\mathrm{I} = (Q_s, g_{1,s}, g_{2,s}, \beta_1)$, with the temperatures functions of them by the forward equations (5.3) at $T_s = 14^\circ\mathrm{C}$ and $Q_o = 8\,\mathrm{kW}$,

$$\mathbf{r}_\mathrm{I}(\mathbf{z}_\mathrm{I}) = \begin{bmatrix} (T_r(\mathbf{z}_\mathrm{I}) - y_{T_r,s})/0.1 \\ (T_{d,1}(\mathbf{z}_\mathrm{I}) - y_{T_{d,1},s})/0.1 \\ (T_{d,2}(\mathbf{z}_\mathrm{I}) - y_{T_{d,2},s})/0.1 \\ (Q_s - y_{Q,s})/3.0 \\ (T_1(\mathbf{z}_\mathrm{I}) - y_{T_1,s})/0.1 \\ (\beta_1 - 0.20)/0.08 \end{bmatrix}, \tag{18.7}$$

Table 18.1: The questions asked of the two-unit hall's forecast at window 36 (Dec Y3). $\mathbf{y}$ is the 185 readings of windows 0–36 and the prior rows; t^* and t_{P} the first passages below G^* and G_{P}, equations (14.1) and (18.1); $p_{\mathrm{no},s}$ the probability (18.4) that the test at s raises no order; $\mathcal{S}$ a schedule (18.2); N_V the number of crew visits of a program and N_{sf} its months short, below G^*, both counted from today to the order.

Asset manager's question	General form	Mathematical form	Category	Computation	Answered in
How often should the exchanger be inspected so that the probability of an undetected crossing stays under 5 %?	(2.23) with the indicator of an undetected crossing as the loss; (2.21), second line	$P_{\mathrm{U}}(\Delta_{\mathrm{I}}) = \mathbb{E}\big[\prod_{s\in\mathcal{S},\, s<t^*} p_{\mathrm{no},s} \mid \mathbf{y}\big]$, averaged over the phase s_1; $\Delta_5 = \max\{\Delta_{\mathrm{I}} : P_{\mathrm{U}}(\Delta') \le 0.05\ \forall\, \Delta' \le \Delta_{\mathrm{I}}\}$	design	trend; expected consequence	18.3.2
How long is this hall's P–F interval?	(2.21), second line at two levels; (2.11)	$\Delta_{\mathrm{PF}} = t^* - t_{\mathrm{P}}$: median and 5–95 % band	estimation	trend	18.3.1
What does each interval cost, which is cheapest, and how firm is that choice?	(2.23); (2.24); perfect information, (17.9)	$\mathcal{R}(\Delta_{\mathrm{I}}) = \mathbb{E}[C_V N_V + \mathrm{SF}\, N_{\mathrm{sf}} \mid \mathbf{y}]$; $\Delta^* = \arg\min \mathcal{R}$; $\Pr(L(\Delta^*, \mathbf{z}) < L(\Delta_5, \mathbf{z}) \mid \mathbf{y})$; EVPI	design	expected consequence	18.3.3
What do the calendar and the rule of thumb give?	the same, at the naive actions	P_{U} and $\mathcal{R}$ at $\Delta_{\mathrm{I}} = 12$, and at $\Delta_{\mathrm{I}} = 3$ (half the handbook's 6 months)	design	expected consequence	18.3.4
If the hall is carried to the October shutdown with no work done, which tests before it does it need?	the same, with the program ended at window 46	$P_{\mathrm{U}}(\mathcal{S}) = \mathbb{E}\big[\mathbf{1}(t^* \le 46) \prod_{s\in\mathcal{S},\, s<t^*} p_{\mathrm{no},s} \mid \mathbf{y}\big]$ for dated sets $\mathcal{S} \subset \{37, \ldots, 45\}$ of one, two and three tests; the probability of an order before the shutdown	design	trend; expected consequence	18.3.5

the four temperature rows at the reference instruments' 0.1 K, the meter row at 3.0 kW, and the survey prior on β_1. The reference instruments carry no drift entry, so unit 1's drifting discharge sensor plays no part in the test.

Which row acts on which entry. The test's rows touch only the four entries of their own window. With $*$ for an entry that can be nonzero, columns $(Q_s, g_{1,s}, g_{2,s}, \beta_1)$ and rows in the order of equation (18.7),

$$\mathbf{J}_{\mathrm{I}} = \begin{bmatrix} * & * & * & 0 \\ * & * & * & 0 \\ * & * & * & 0 \\ * & 0 & 0 & 0 \\ * & * & * & * \\ 0 & 0 & 0 & * \end{bmatrix}, \qquad \mathbf{g}_G^{\mathsf{T}} = (0,\ G_{1,s},\ G_{2,s},\ 0), \tag{18.8}$$

$$P_{\mathrm{U}} = P_{\mathrm{U}}\big(G_{\mathrm{tot},s} : s \in \mathcal{S},\ s < t^*\big), \qquad G_{\mathrm{tot},s} = G_{\mathrm{tot},s}\big(g_{1,36}, g_{2,36}, \gamma_1, \gamma_2, \dot{\gamma}_1;\ e_{u,t},\ t < s\big).$$

The first line is the block a test adds; a schedule adds one such block per test window, block-diagonal beside Chapter 14's pattern (14.6) apart from the shared column of β_1. The vector $\mathbf{g}_G$ is not a row: it is the gradient of $G_{\mathrm{tot},s} = \bar{G}_1 \mathrm{e}^{g_{1,s}} + \bar{G}_2 \mathrm{e}^{g_{2,s}}$ on the test's entries. The second line is the pattern of the question's dependence: the probability of an undetected crossing reads the path only at the test windows before the crossing, and by equation (14.7) each forecast window's total is a function of Chapter 14's five forecast entries and of the departures $e_{u,t}$ of the windows before it. Two facts are read here, before any number. The return-air row and the rack-1 row see both coils alike and each discharge row its own coil, so the test reads the sum better than either coil. And the test adds rows only at windows that have no reading in the record, so nothing in the record's fit changes.

Modeling choices at the rows. Six choices make these rows. *The test's grade and inputs are the condition test's*: this buys one σ_{I} for every test, and it would fail if the crew read the hall at the season's own inputs instead, a reference-grade reading of an ordinary window; step 3 gives that variant's spread. *One σ_{I}, taken at the potential-failure level and at the shutdown's load*: this buys the closed form (18.4), and it fails where the level or the load moves the test's spread; the test matters only near G_{P}, which is where the value is taken. *The tests' errors are independent of each other and of the record*: each test is a fresh crew visit with its own instruments. *A test after the crossing finds it*: at G^* the test's reading lies $m_{\mathrm{P}}/\sigma_{\mathrm{I}}$ standard deviations below the order level, so the first test after t^* ends the shortfall. *The phase is averaged*: the first test is taken as equally likely to fall in any of the Δ_{I} windows after today, the parent's uniform phase, which makes the answer a property of the interval as a policy; the dated schedule, with the first test one interval from today, and the best and the worst phase stand beside it. *No lead time*: a test that raises the order before the crossing is taken to prevent the shortfall, however close to the crossing it falls, so the work is done at once. This buys a program read off the paths at the test windows alone, and it fails when the work needs time to arrange: with a lead of one window, a test counts only if it falls before $t^* - 1$, which step 5 tests.

18.2 The posterior of the state (step 3)

Every answer of this chapter is read from two posteriors: the forecast's, Chapter 14's extended state under each drawing of its family, and the condition test's own, the posterior of one test's four entries, which fixes σ_{I}.

18.2.1 The forecast carried past the record

The derivation form is Chapter 5's (Sections 5.2.1 and 5.2.2), on each drawing's stacked state, as Chapter 14 fitted it (Section 14.2.1) and not derived again. Each drawing M has its MAP state $\hat{\mathbf{z}}_M$, its Laplace covariance $\mathbf{\Sigma}_M = (\mathbf{J}_M^{\mathsf{T}}\mathbf{J}_M)^{-1}$ and its weight P_M, equation (14.8); with every row normalized the information matrix is $\mathbf{J}^{\mathsf{T}}\mathbf{J}$, and with unnormalized rows and their weights it is $\mathbf{J}^{\mathsf{T}}\mathbf{W}\mathbf{J}$. The 18 drawings are those of Table 14.2: the leading four are coil 1 curving with coil 2's step, $(\mathcal{C}_1, 2)$, at 0.3428; constant rates with coil 2's step, $(\mathcal{C}_0, 2)$, at 0.2061; coil 1 curving with unit 2's sensor stepped, $(\mathcal{C}_1, T_{d,2})$, at 0.1606; and coil 1 curving with no step, $(\mathcal{C}_1, \text{none})$, at 0.1097. Summed over the family, $\Pr(\mathcal{C}_1 \mid \mathbf{y}) = 0.6344$ and the probability of a step at window 15 is 0.7864. The forecast windows change nothing in the fit, and the posterior of the extended state factors into the record's posterior and one Gaussian step per window and coil, equation (14.10).

The draws. The forecast is read as Chapter 14 read it (Section 14.3.1): 40 000 draws per drawing of the five forecast entries from their block of $\mathcal{N}(\hat{\mathbf{z}}_M, \mathbf{\Sigma}_M)$, each carried window by window by the forecast rows (14.4) with a fresh departure of spread σ_q, on Chapter 14's own seeds, 720 000 draws in all, each weighted $P_M/40\,000$. What this chapter adds is that each draw keeps its total $G_{\text{tot},t}$ at every forecast window up to its crossing, so that the tests' factors (18.4) can be read on it. With the same draws the first passage comes out to Chapter 14's printed digits: median 52.55, window 53 (May Y5); band 43.74 to 81.67; $\Pr(t^* \le 46 \mid \mathbf{y}) = 0.1595$; no crossing by window 276 with 0.0089. The total at window 46 has median 11.16 and band 10.53 to 11.70 kW/K, as in Chapter 14.

18.2.2 The condition test's own posterior

The test's estimate of the level is the MAP state of its own rows, equation (18.7), by the normal equations (5.9) on the four entries $\mathbf{z}_{\text{I}}$,

$$(\mathbf{J}_{\text{I}}^{\mathsf{T}}\mathbf{J}_{\text{I}})\,\Delta\mathbf{z}_{\text{I}} = -\mathbf{J}_{\text{I}}^{\mathsf{T}}\mathbf{r}_{\text{I}}(\mathbf{z}_{\text{I}}), \qquad \mathbf{\Sigma}_{\text{I}} = (\mathbf{J}_{\text{I}}^{\mathsf{T}}\mathbf{J}_{\text{I}})^{-1}, \qquad \sigma_{\text{I}}^2 = \mathbf{g}_G^{\mathsf{T}}\mathbf{\Sigma}_{\text{I}}\,\mathbf{g}_G, \tag{18.9}$$

with $\mathbf{g}_G$ the gradient of the sum on the test's entries from the pattern (18.8). The test's estimate $\hat{G}_s$ is $G_{\text{tot},s}$ at that MAP state, and σ_{I} is its spread by the delta method: the spread that equation (18.4) uses. The Jacobian is formed at the potential-failure level, where the test decides: both coils at the same log-derate, $g_1 = g_2 = -0.21772$, which puts $G_1 + G_2$ at 11.2610 kW/K; the load at window 46's growth, 139.013 kW; and β_1 at the survey's 0.20.

The Jacobian and the covariance as numbers. The normalized Jacobian $\mathbf{J}_{\text{I}}$ is set out in Table 18.2, a row per row of equation (18.7) and a column per entry, and beneath it the test's covariance as standard deviations and the correlation of the two coils. The zeros sit where the pattern (18.8) puts them. Each discharge row has its largest entry, −102.0290, in its own coil's column and −28.5224 in the other's; the return-air and rack-1 rows see the two coils alike, −65.2757 each, because the return air depends on their sum; and only rack 1's row and the survey prior see β_1.

The test's spread of the sum. By equation (18.9), $\sigma_{\text{I}} = 0.1033$ kW/K, and the margin m_{P} is 4.84 of it. The variance of the sum is the two coils' variances times $G_{u,s}^2$ plus twice their covariance times $G_{1,s}G_{2,s}$, and the cross term takes the sign of the correlation, −0.0472, which narrows the sum below the two coils' variances added. The sign comes from the return air, which reads the total and gives the two coils opposite errors, as in Chapter 14's record, where the two log-derates at window 36 are correlated by −0.6148 in the leading drawing. Here the two reference discharges each read their own coil at a third of the record's sensor error, so each

Table 18.2: The condition test's rows (18.7) at the potential-failure level ($g_1 = g_2 = -0.21772$, $Q = 139.013\,\mathrm{kW}$, $\beta_1 = 0.20$, $T_s = 14°\mathrm{C}$, $Q_o = 8\,\mathrm{kW}$): the normalized Jacobian, and beneath it the test's covariance $\mathbf{\Sigma}_\mathrm{I}$ (18.9) as standard deviations and the coils' correlation.

Row	Q_s [kW^{-1}]	$g_{1,s}$ [–]	$g_{2,s}$ [–]	β_1 [–]
Return air T_r, 0.1 K	+0.8880	−65.2757	−65.2757	0
Discharge of unit 1, $T_{d,1}$, 0.1 K	+0.3880	−102.0290	−28.5224	0
Discharge of unit 2, $T_{d,2}$, 0.1 K	+0.3880	−28.5224	−102.0290	0
Load meter Q, 3.0 kW	+0.3333	0	0	0
Rack-1 inlet T_1, 0.1 K	+0.4880	−65.2757	−65.2757	+73.5066
Survey prior on β_1	0	0	0	+12.5000
$\sqrt{\Sigma_{ii}}$ of $\mathbf{\Sigma}_\mathrm{I}$ [entry's unit]	1.8969 [kW]	0.01329	0.01329	0.0155
Correlation of $g_{1,s}$ and $g_{2,s}$ [–]	−0.0472			
Gradient $\mathbf{g}_G$ of $G_1 + G_2$ [kW/K]	0	5.6305	5.6305	0

coil is read nearly on its own and the anticorrelation nearly vanishes. The record's posterior of the level is much wider at a test window: the total at window 46 has a band of 10.53 to 11.70 kW/K, several times σ_I, which is the step-0 choice's premise.

What moves σ_I. By the instruments' grade it is 0.0597 kW/K at 0.05 K, 0.1033 at 0.1, 0.1563 at 0.2 and 0.1956 at 0.3 K, the grade of the hall's own sensors. The hall's own sensors at 0.3 K with unit 1's drifting discharge reading left out give 0.3151. At the reference grade it is 0.0998 at the limit G^*, 0.1068 at $G_\mathrm{P} + 0.5\,\mathrm{kW/K}$, and 0.1286 at the commissioning load of 110 kW, where the smaller temperature differences read the coils less sharply. Read instead at each window's own inputs over windows 37 to 48, a reference-grade reading of an ordinary window rather than a condition test, it runs from 0.1012 to 0.1064.

18.2.3 The tests' errors integrated over a path

Every quantity the questions ask of a schedule is the posterior mean of a function of the path, and given the path the tests' errors integrate in closed form. Fix a draw of the extended state, so that the path $G_{\mathrm{tot},t}$, its crossing t^* and every factor (18.4) are known, and fix a schedule (18.2). The test at s is reached, with no order raised before it, with the probability that every earlier test raised none, $\prod_{s'\in\mathcal{S},\ s'<s} p_{\mathrm{no},s'}$. The visits are the tests reached before the crossing, and, on an undetected crossing, the first test after it, which finds it (step 2's choice); the shortfall runs from t^* to that test, $s^+(t^*) = \min\{s \in \mathcal{S} : s \geq t^*\}$; and an early order is one raised by a test that reads a path still above G_P:

$$\begin{aligned}
\mathbb{E}[N_V \mid \mathbf{z}, \mathcal{S}] &= \sum_{s\in\mathcal{S},\ s<t^*} \prod_{s'\in\mathcal{S},\ s'<s} p_{\mathrm{no},s'} \ + \ \Pr(\text{undetected} \mid \mathbf{z}, \mathcal{S}),\\
\mathbb{E}[N_\mathrm{sf} \mid \mathbf{z}, \mathcal{S}] &= \Pr(\text{undetected} \mid \mathbf{z}, \mathcal{S})\,\big(s^+(t^*) - t^*\big),\\
\Pr(\text{early order} \mid \mathbf{z}, \mathcal{S}) &= \sum_{s\in\mathcal{S},\ s<t^*,\ G_{\mathrm{tot},s}>G_\mathrm{P}} \Big(\prod_{s'\in\mathcal{S},\ s'<s} p_{\mathrm{no},s'}\Big)\,(1 - p_{\mathrm{no},s}),
\end{aligned} \tag{18.10}$$

with N_V the number of visits, N_sf the months below G^*, the undetected probability (18.5), and the sums running to the horizon on a path that does not cross. The program's loss given the path is its visits and its shortfall,

$$L(\mathcal{S}, \mathbf{z}) = C_V\,\mathbb{E}[N_V \mid \mathbf{z}, \mathcal{S}] + \mathrm{SF}\,\mathbb{E}[N_\mathrm{sf} \mid \mathbf{z}, \mathcal{S}], \tag{18.11}$$

from today to the order, or to the horizon. This is the loss of Chapter 2's second line of equation (2.23): the certain cost replaced by the visits' expected number, and the event, a

missed crossing, with its consequence growing with the months it runs. The questions then average over the forecast's draws and, for an interval, over its phases. For the interval Δ_{I},

$$\begin{aligned} P_{\mathrm{U}}(\Delta_{\mathrm{I}}) &= \frac{1}{\Delta_{\mathrm{I}}} \sum_{s_1=37}^{36+\Delta_{\mathrm{I}}} \sum_{M} P_M \, \mathbb{E}\big[\Pr(\text{undetected} \mid \mathbf{z}, \mathcal{S}(\Delta_{\mathrm{I}}, s_1)) \mid \mathbf{y}, M\big], \\ \mathcal{R}(\Delta_{\mathrm{I}}) &= \frac{1}{\Delta_{\mathrm{I}}} \sum_{s_1=37}^{36+\Delta_{\mathrm{I}}} \sum_{M} P_M \, \mathbb{E}\big[L(\mathcal{S}(\Delta_{\mathrm{I}}, s_1), \mathbf{z}) \mid \mathbf{y}, M\big], \end{aligned} \tag{18.12}$$

each expectation read as the weighted mean over the 720 000 draws, with Monte Carlo standard errors by batch means over 20 batches, as in Section 5.2.4. The dated schedule takes the one phase $s_1 = 36 + \Delta_{\mathrm{I}}$, and the best and the worst phase are the smallest and the largest of the Δ_{I} terms. A path's own probability (18.5) needs its totals only at the test windows before t^*, the second line of the pattern (18.8).

18.2.4 The full posterior by reweighted draws

Beside the Laplace normal of each drawing stands the full posterior of the same rows, by Chapter 5's reweighted draws (Section 5.2.4, equation (5.13)), run as Chapter 14 ran it (Section 14.2.4): 20 000 draws per drawing on its whole state, on Chapter 14's seeds, each weighted by the ratio of the drawing's posterior to its Laplace normal, and the drawings weighted by their importance-sampling evidence. Because the forecast factors, each reweighted draw is carried past window 36 as a Laplace draw is, keeps its weight, and takes the tests' factors (18.4) on its path. The effective sample sizes run from 95.2 to 98.9 % of the draws, the drawings' weights move by at most 0.0007, and the first passage reproduces Chapter 14's full posterior, median 52.63 and $\Pr(t^* \le 46 \mid \mathbf{y}) = 0.1589$. Step 5 states the gaps.

18.3 The answers (step 4)

Every answer is read from the forecast's draws of Section 18.2.1, the Laplace normal of each drawing carried past the record and mixed over the family, with the tests' errors integrated in closed form by Section 18.2.3; that value is the one reported, and the full posterior of Section 18.2.4 stands beside it. Intervals are whole windows, that is whole months, and bands are 5–95 % throughout.

18.3.1 The P–F interval of the hall

Asset manager's question. How long is this hall's P–F interval? The answer is a time with a band: from the window at which a test would first raise the order to the window at which the hall runs short.

General form. The distribution function of a crossing time, the second line of the trend form, equation (2.21),

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

read by inversion at the levels of the credible interval (2.11), here 0.05, 0.5 and 0.95.

Specialized form. The question is equation (2.21), second line, taken at two levels on the same draw: with $\mathbf{z}$ the extended state (14.3) of each drawing and $p(\mathbf{z} \mid \mathbf{y})$ its factored posterior (14.10) mixed over the family, the first passage below G_{P} gives t_{P} and the first passage

below G^* gives t^*, both interpolated in the same way, equation (14.1). The quantity is their difference (18.1), a function of the draw, whose distribution is read by equation (2.11) on the weighted draws. Its band is taken over the paths that cross by the horizon, 0.9911 of the posterior, since a path that never crosses has no P–F interval within it. No closed form exists, because each first passage is a minimum over windows. One choice enters here: the two levels are read on the same path, so the interval carries the path's own rate of fall between them and the correlation of its two passages, which two separate quantiles of t_{P} and t^* would lose.

The P–F interval is a margin over a rate of fall, so like the crossing date it is read by its median and band, and its delta form stands beside them as Chapter 14's delta band does beside the date, equation (14.13): the interval on each drawing's MAP path with no wander, $\Delta_{\mathrm{PF}}(\hat{\mathbf{z}}_M)$, with the spread of its gradient on the forecast entries,

$$\Delta_{\mathrm{PF}}(\hat{\mathbf{z}}_M) \pm 1.6449 \sqrt{\mathbf{g}_{\mathrm{PF}}^{\mathsf{T}} \boldsymbol{\Sigma}_M \, \mathbf{g}_{\mathrm{PF}}}, \qquad \mathbf{g}_{\mathrm{PF}} = \left. \frac{\partial \Delta_{\mathrm{PF}}}{\partial (g_{1,36},\ g_{2,36},\ \gamma_1,\ \gamma_2,\ \dot{\gamma}_1)} \right|_{\hat{\mathbf{z}}_M}, \tag{18.13}$$

the gradient taken by central differences of the plug-in path.

Building blocks. The draws of Section 18.2.1 with their weights; the two first passages; the weighted quantiles; each drawing's MAP state and the block of $\boldsymbol{\Sigma}_M$ on the five forecast entries (Section 14.2.3), with the gradient of equation (18.13).

The potential failure and the interval. The potential failure comes at window 44.27 at the median, window 45 (Sep Y4), with a band of 38.67 to 61.74, window 39 (Mar Y4) to window 62 (Feb Y6). It comes before the shutdown with probability $\Pr(t_{\mathrm{P}} \le 46 \mid \mathbf{y}) = 0.6078$, and the hall is already below G_{P} today with probability 0.0004. The P–F interval has a median of 7.93 months and a band of 4.44 to 20.37, with 4.10 to 25.03 at 2.5–97.5 %. It is shorter than 3 months with probability 0.0003, than 4 with 0.0199, than 6 with 0.2581 and than 12 with 0.7732. The full posterior of Section 18.2.4 gives a P–F interval with median 7.94 months and band 4.42 to 20.37 (4.09 to 25.16 at 2.5–97.5 %), shorter than 6 months with probability 0.2593, and a potential failure at window 44.36 at the median. The delta form (18.13) on the leading drawing $(\mathcal{C}_1, 2)$ is $7.09 \pm 1.6449 \times 1.80$ months, 4.13 to 10.05, against that drawing's first passage, median 6.96 and band 4.49 to 14.11; on the base fit it is $10.49 \pm 1.6449 \times 1.75$, 7.62 to 13.37, against 10.26 and 6.48 to 17.59. The symmetric band of the delta form ends far short of the first passage's long upper tail, as the date's did in Chapter 14, and it is not reported. Table 18.3 sets the interval out drawing by drawing. Each row is one of the leading drawings or one state of coil 1's curvature over the family; the columns give its probability, the median and band of its P–F interval, and the median window of its potential failure.

The interval is set by the fall, and the fall by the curvature. The P–F interval runs from 5.70 months under $(\mathcal{C}_1, T_{d,2})$ to 13.97 under $(\mathcal{C}_0, 2)$, and over the family the curvature states split it, 6.46 months with coil 1 curving against 12.11 at constant rates. The interval is a margin over a rate of fall, equation (18.1), and coil 1's curvature is what makes the fall fast at G_{P}: the drawings that free it pass the potential failure earlier, between windows 41.08 and 43.04, and cross the limit sooner after. The paths whose P–F interval is shorter than the handbook's six months, 0.2581 of the posterior, are the paths on which coil 1 accelerates: the probability is 0.4011 with coil 1 curving and 0.0133 at constant rates, and 0.9810 of those short paths' probability lies in the curving drawings.

The distribution is drawn in Figure 18.2. The horizontal axis is the P–F interval in months and the vertical axis the probability per one-month bin. The bars are the family; the two thin step lines are the curvature states, each normalized to its own probability; the dashed line marks the handbook's six months and the dotted line the stated truth's 4.56 months, a check marker

Table 18.3: The P–F interval (18.1) by drawing and by coil 1's curvature state, from the forecast's draws (40 000 per drawing). Bands are 5–95 %; the potential failure's median is a continuous window.

Drawing or group	P_M [–]	P–F median [mo]	band [mo]	t_P median [window]
$(\mathcal{C}_1, 2)$, coil 1 curving, coil 2 stepped	0.3428	6.96	4.49–14.11	43.04
$(\mathcal{C}_0, 2)$, constant rates, coil 2 stepped	0.2061	13.97	7.95–30.48	51.90
$(\mathcal{C}_1, T_{d,2})$, coil 1 curving, unit 2's sensor stepped	0.1606	5.70	3.90–9.66	41.08
$(\mathcal{C}_1, \text{none})$, coil 1 curving, no step	0.1097	6.30	4.17–12.25	41.75
$(\mathcal{C}_0, \text{none})$, the base fit	0.1038	10.26	6.48–17.59	46.61
$\mathcal{C}_0$, constant rates, over the family	0.3656	12.11	7.05–26.13	
$\mathcal{C}_1$, coil 1 curving, over the family	0.6344	6.46	4.20–12.77	
The family	1	7.93	4.44–20.37	44.27

of step 5. The family's mass lies between four and ten months with a long tail to the right. The curving state's mass sits at four to eight months and the constant-rate state's at seven to sixteen, and the left edge of the family's bars, the short P–F intervals, is the curving state's.

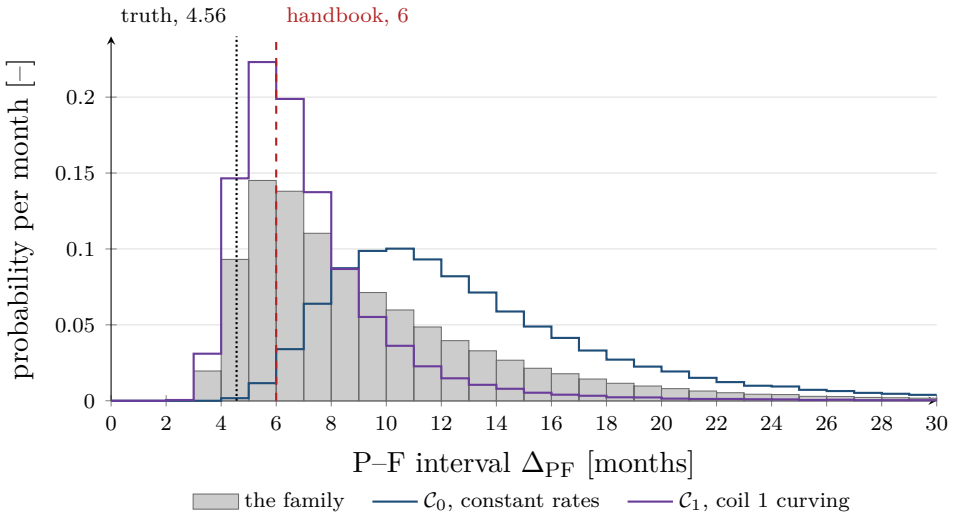


Figure 18.2: The P–F interval (18.1) over the forecast's draws: probability per one-month bin for the family (bars) and for each state of coil 1's curvature (step lines, each normalized to its own probability, 0.3656 and 0.6344); intervals of 30 months or more are not drawn. Dashed: the handbook's 6 months. Dotted: the stated truth's 4.56 months, a check marker (step 5).

18.3.2 The longest interval under five percent

Asset manager's question. How often should the exchanger be inspected so that the probability of an undetected crossing stays under 5 %? The answer is an interval in whole months, the longest that meets the limit.

General form. An expected consequence, equation (2.23),

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

here with the indicator of an undetected crossing as the loss, so that $\mathcal{R}(a)$ is the probability of the event under action a, and the actions those that keep it under the limit.

Specialized form. The question is equation (2.23) with a the interval Δ_{I}, $L(a, \mathbf{z})$ the indicator of an undetected crossing, whose expectation over the tests' errors is the probability (18.5), and $p(\mathbf{z} \mid \mathbf{y})$ the forecast's mixture: that is the first of equations (18.12), $P_{\mathrm{U}}(\Delta_{\mathrm{I}})$, over the tests $s_1, s_1 + \Delta_{\mathrm{I}}, \ldots$ before t^* and averaged over $s_1 \in \{37, \ldots, 36 + \Delta_{\mathrm{I}}\}$. The answer is the longest interval that keeps it and every shorter one under the limit,

$$\Delta_5 = \max\{\Delta_{\mathrm{I}} : P_{\mathrm{U}}(\Delta') \le 0.05 \text{ for every } \Delta' \le \Delta_{\mathrm{I}}\}. \tag{18.14}$$

The interval is a whole number of months, so its band is read from what moves it: the drawings, the curvature states, the full posterior and, in step 5, fresh records.

Building blocks. The draws and weights of Section 18.2.1; σ_{I} of equation (18.9); the factors (18.4) and their product (18.5); the phase average (18.12); the batch-means standard errors.

The probability by interval. It is listed in Table 18.4. Each row is an interval; the columns give P_{U} with its Monte Carlo standard error, the dated schedule's value and the best and worst phases', the expected number of visits and of months below the limit by equation (18.10), the probability of an early order, and the inspection cost per year (18.6). The probability rises from 0.0017 at three months to 0.0101 at four, 0.0302 at five and 0.0621 at six, and to 0.3147 at twelve.

Table 18.4: The probability of an undetected crossing by interval, equation (18.12), the phases averaged (Monte Carlo standard error in parentheses, 0.0000 where below 0.00005), with the dated schedule (first test at window $36 + \Delta_{\mathrm{I}}$) and the best and worst phases beside it; the expected visits and shortfall months by equation (18.10); the inspection cost per year (18.6).

Δ_{I} [mo]	P_{U} (se) [–]	dated [–]	best [–]	worst [–]	visits [–]	shortfall [mo]	early order [–]	cost per year [R$/yr]
1	0.0000 (0.0000)	0.0000	0.0000	0.0000	10.71	0.0000	0.5316	36 000
2	0.0001 (0.0000)	0.0001	0.0001	0.0001	6.15	0.0000	0.3307	18 000
3	0.0017 (0.0000)	0.0017	0.0017	0.0018	4.53	0.0006	0.2383	12 000
4	0.0101 (0.0001)	0.0105	0.0098	0.0105	3.69	0.0053	0.1856	9 000
5	**0.0302** (0.0001)	0.0297	0.0284	0.0320	3.17	0.0220	0.1516	7 200
6	0.0621 (0.0002)	0.0568	0.0568	0.0679	2.82	0.0604	0.1280	6 000
7	0.1022 (0.0002)	0.0899	0.0895	0.1157	2.57	0.1283	0.1106	5 143
8	0.1464 (0.0002)	0.1318	0.1198	0.1719	2.38	0.2293	0.0974	4 500
9	0.1912 (0.0003)	0.1812	0.1509	0.2301	2.23	0.3638	0.0869	4 000
10	0.2347 (0.0003)	0.2369	0.1823	0.2871	2.11	0.5304	0.0785	3 600
11	0.2760 (0.0003)	0.2963	0.2108	0.3442	2.01	0.7266	0.0715	3 273
12	0.3147 (0.0003)	0.3537	0.2370	0.3960	1.92	0.9497	0.0657	3 000
18	0.4941 (0.0003)	0.6110	0.3668	0.6229	1.62	2.7201	0.0441	2 000
24	0.6059 (0.0004)	0.7616	0.4416	0.7659	1.46	4.9629	0.0332	1 500

The answer. A condition test every 5 months keeps the probability of an undetected crossing at $P_{\mathrm{U}} = 0.0302$ (standard error 0.0001); at 6 months it is 0.0621 (standard error 0.0002), and the 5 % level falls at 5.62 months between the two. The dated schedule and the worst phase give the same 5 months. The inspection cost per year is R$ 7 200. The 5-month dated schedule tests at window 41 (May Y4) before the shutdown, and the 4-month one at windows 40 (Apr Y4) and 44 (Aug Y4).

What moves it. Across the 18 drawings, the longest interval under 5 % runs from 4 to 10 months, with a weighted median of 5. It is 5 months on the leading drawing $(\mathcal{C}_1, 2)$, 10 on $(\mathcal{C}_0, 2)$, 4 on $(\mathcal{C}_1, T_{d,2})$, 5 on $(\mathcal{C}_1, \text{none})$ and 7 on the base fit. The open question that sets it is the curvature, as for Chapter 14's date. On the constant-rate drawings $P_{\mathrm{U}}(5) = 0.0027$ and the longest interval is 8 months; on the curving drawings $P_{\mathrm{U}}(5) = 0.0460$ and it is 5, and they carry 0.9671 of the family's $P_{\mathrm{U}}(5)$. The full posterior gives the same 5 months, with gaps in P_{U} of at most +0.0006 (step 5), and over fresh records at the stated truth the interval a record would choose runs from 5 to 10 months (step 5).

Figure 18.3 draws the probability against the interval. The horizontal axis is the interval from 1 to 12 months and the vertical axis P_{U} on a logarithmic scale. The thick line is the family, the thin lines the two curvature states, the dashed line the dated schedule and the dotted line the stated truth's own path (step 5); the horizontal line is the 5 % limit and the vertical one Δ_5. The family's curve crosses the limit between 5 and 6 months, and the curving state's lies just above it and crosses at the same place, 0.0460 at five months against 0.0932 at six, while the constant-rate state's stays under the limit to eight months, 0.0323, and crosses it at nine, 0.0520. The dated schedule stays close to the family's curve to nine months and lies above it from ten, where its first test, fixed one interval from today, comes late. The truth's curve lies above all of them: it is a fast path.

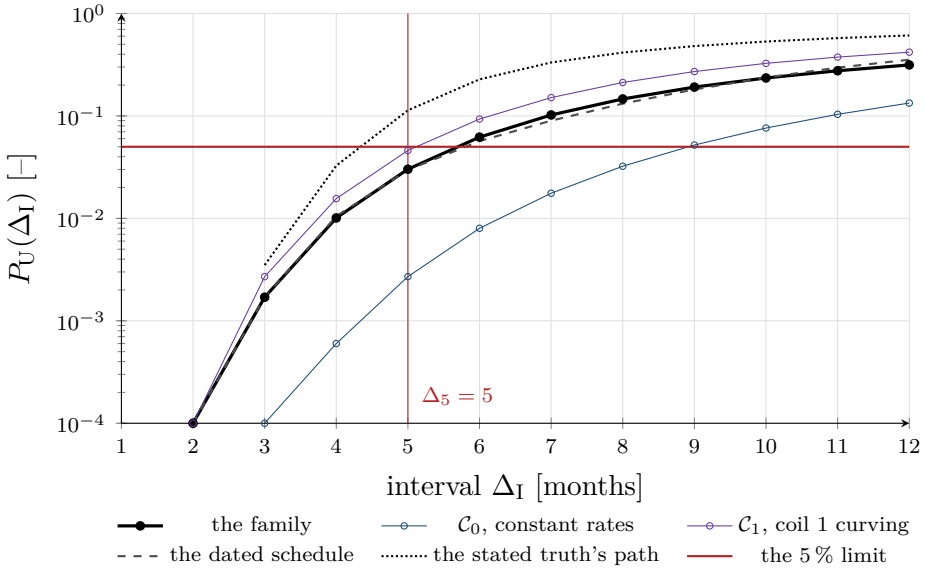


Figure 18.3: The probability of an undetected crossing against the interval, equation (18.12), on a logarithmic axis; values below 0.00005, at one month, are not drawn. Thick: the family, phases averaged; thin: the curvature states $\mathcal{C}_0$ and $\mathcal{C}_1$; dashed: the dated schedule; dotted: the stated truth's path with the tests' errors only, a check marker (step 5). The horizontal line is the 5 % limit and the vertical one $\Delta_5 = 5$ months.

How a path is caught. Two draws of the leading drawing $(\mathcal{C}_1, 2)$ are traced against the window in Figure 18.4, with the two levels as horizontal lines and the dated four-month tests as ticks at windows 40, 44, 48 and 52. The slow path is the draw nearest the family's median P–F interval and median t_{P}; the fast path the draw nearest the family's 5 % P–F interval and the same t_{P}. Both pass the potential failure just after window 44, at 44.25 and 44.23, so the test at window 44 reads each a little above G_{P} and raises no order with probability 0.5430 and 0.5696. The slow path, with a P–F interval of 7.92 months, crosses the limit at 52.17 (window 53), and two more tests fall before its crossing: the test at window 48 raises no order with

probability 0.0100 and the one at 52 with 0.0000. The fast path, with 4.44 months, crosses at 48.67 (window 49), and one test falls between its two passages, at window 48, which raises no order with probability 0.0001. On the dated schedule both are caught, at $P_{\mathrm{U}} = 0.0000$ and 0.0001; over the four phases the slow path is still caught, 0.0000, while the fast path goes undetected with probability 0.0832. In the phase that tests at window 45, a month after its potential failure, it stands at 11.2011 kW/K, only a little below G_{P}, and its next test comes after the crossing. The fast paths are the ones the interval has to be set against.

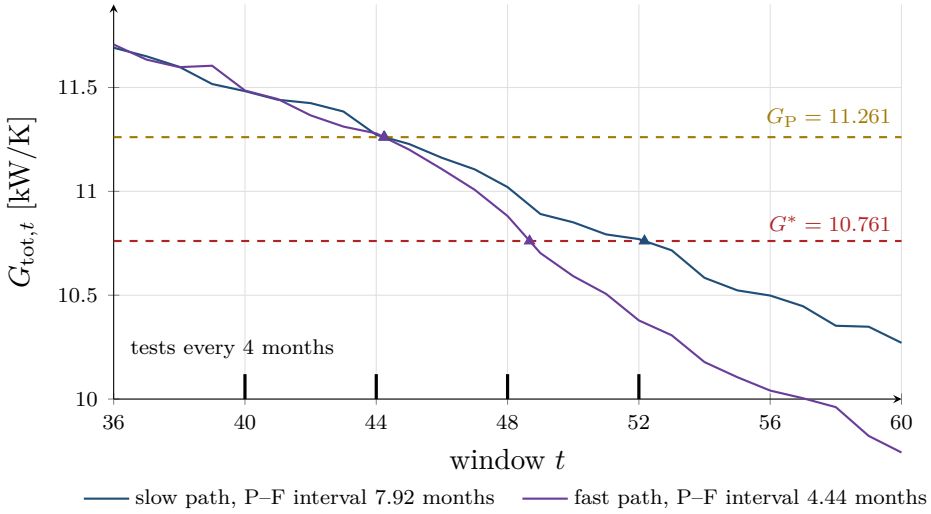


Figure 18.4: Two forecast paths of $G_{\mathrm{tot},t}$ from the leading drawing $(\mathcal{C}_1, 2)$, windows 36–60: the draw nearest the family's median P–F interval and median potential failure (slow) and the draw nearest its 5 % P–F interval with the same potential failure (fast). Horizontal lines: $G_{\mathrm{P}} = 11.261$ and $G^* = 10.761$ kW/K. Ticks: the dated four-month tests, windows 40, 44, 48 and 52. Markers: each path's t_{P} and t^*.

18.3.3 The expected cost of an interval

Asset manager's question. What does each interval cost, which is cheapest, and how firm is that choice? The answer prices the visits and the missed months of each program under the same posterior.

General form. The expected consequence of each action and the action that minimizes it, equation (2.23),

$$\mathcal{R}(a) = \mathbb{E}\big[L(a, \mathbf{z}) \mid \mathbf{y}\big], \qquad a^* = \arg\min_a \mathcal{R}(a),$$

with its Laplace reading, equation (2.24), the probability that the chosen action costs less than another at the true state, and the expected value of perfect information, equation (17.9).

Specialized form. The question is equation (2.23) with $a = \Delta_{\mathrm{I}}$ and $L(a, \mathbf{z})$ the program's loss (18.11): the second of equations (18.12), $\mathcal{R}(\Delta_{\mathrm{I}})$, the posterior mean of $C_V\,\mathbb{E}[N_V \mid \mathbf{z}] + \mathrm{SF}\,\mathbb{E}[N_{\mathrm{sf}} \mid \mathbf{z}]$ averaged over the phases, and $\Delta^* = \arg\min_{\Delta_{\mathrm{I}}} \mathcal{R}(\Delta_{\mathrm{I}})$. The costs it depends on are INSPECT, through C_V, and SHORTFALL; CLEAN and OFF_WINDOW fall on every path and drop out. Two choices enter here. *The costs run undiscounted from today to the order*: this keeps the comparison in today's reais without a horizon of its own, and step 5 discounts them at 8 %/yr. *An early order costs nothing extra*: it is an order the coil needs within months, not a false one, and step 5 prices it at FO. How firm the choice is, the second line of equation (2.24), is read on the draws, $\Pr(L(\Delta^*, \mathbf{z}) < L(\Delta', \mathbf{z}) \mid \mathbf{y})$ as the weighted share of draws on which the

loss of Δ^* is the smaller. Perfect information on the path is equation (17.9) with the interval chosen draw by draw among the candidates, EVPI $= \mathcal{R}(\Delta^*) - \mathbb{E}[\min_{\Delta_\mathrm{I}} L(\Delta_\mathrm{I}, \mathbf{z}) \mid \mathbf{y}]$. The Laplace reading, the first line of equation (2.24), is the loss at the MAP state: the plug-in, each drawing's MAP path with no wander. Its second-order term is not formed, because the loss has a step in the state wherever a test window passes the crossing; the draws are the Laplace route, and the chapter names them so.

Building blocks. The per-path expectations (18.10) and the loss (18.11); the draws and weights; C_V and SF of step 0; the minimum over the candidates on each draw.

The costs by interval. Table 18.5 gives them, in reais from today to the order. Each row is an interval; the columns give $\mathcal{R}$ with its standard error, its two parts, the visits and the shortfall, and the two sensitivity forms, with early orders priced at FO and with the costs discounted. The visits' part falls with the interval and the shortfall's rises, slowly to four months and then fast: R$ 791 at four months, R$ 3 299 at five, R$ 9 066 at six and R$ 142 453 at twelve.

Table 18.5: The expected cost of each interval, equations (18.11) and (18.12), in reais from today to the order (Monte Carlo standard error in parentheses), with its visits and shortfall parts and the two sensitivity forms.

Δ_I [mo]	$\mathcal{R}$ (se) [R$]	visits [R$]	shortfall [R$]	early orders at FO [R$]	discounted 8 %/yr [R$]
1	32 127 (104)	32 127	0	33 722	29 233
2	18 455 (53)	18 453	2	19 447	16 784
3	13 677 (35)	13 588	89	14 392	12 434
4	**11 855** (27)	11 064	791	**12 412**	**10 795**
5	12 812 (27)	9 513	3 299	13 266	11 732
6	17 528 (38)	8 462	9 066	17 912	16 164
7	26 945 (57)	7 703	19 243	27 277	24 952
8	41 519 (78)	7 128	34 391	41 811	38 491
9	61 254 (100)	6 679	54 575	61 514	56 754
10	85 880 (124)	6 317	79 563	86 115	79 467
11	115 009 (149)	6 020	108 989	115 223	106 247
12	148 224 (175)	5 771	142 453	148 422	136 695
18	412 863 (341)	4 855	408 008	412 995	377 193
24	748 826 (517)	4 393	744 433	748 925	679 584

The cheapest interval. The program is cheapest at 4 months, $\mathcal{R}(4) =$ R$ 11 855, R$ 956 below the 5 % interval's R$ 12 812. The costs choose a shorter interval than the 5 % limit does. A month of shortfall costs R$ 150 000 against a visit's R$ 3 000, so the probability of an undetected crossing that the owner can afford is set by the costs, and here it is below 5 %: the cheapest program runs at $P_\mathrm{U} = 0.0101$. Discounted at 8 %/yr, the cheapest is still 4 months, and with early orders priced at FO it is still 4 months, at R$ 12 412. The two costs it depends on move it at their break-evens, where $\mathcal{R}$ of two intervals is equal: with SHORTFALL below R$ 92 782 a month the cheapest is 5 months, and above R$ 539 218 it is 3; with a visit dearer than R$ 4 850 it is 5. The stated R$ 150 000 and R$ 3 000 lie inside both ranges.

How firm it is. $\Pr(L(4, \mathbf{z}) < L(5, \mathbf{z}) \mid \mathbf{y}) = 0.3854$: the interval that is cheaper in expectation is the dearer one on most paths. On most of them the five-month program costs less, by one test fewer, and the expectation is carried by the few paths whose crossing the five-month schedule misses, at R$ 150 000 a month. Against the rule of thumb's three months, $\Pr(L(4) < L(3) \mid \mathbf{y}) = 0.8337$, and against the calendar's twelve, 0.8533.

Perfect information and the plug-in. With the path known, the interval chosen draw by draw would cost $\mathbb{E}[\min L] =$ R$ 8 331, so EVPI = R$ 3 524 against $\mathcal{R}(4) =$ R$ 11 855: knowing the path would buy a well-timed test and drop most of the others. The plug-in, the leading drawing's path at its MAP with no wander, has $t_{\mathrm{P}} = 43.01$, $t^* = 50.10$ and a P–F interval of 7.09 months. On it P_{U} is 0.0005 at four months, 0.0047 at five, 0.0195 at six, 0.0535 at seven and 0.1097 at eight, so the plug-in allows 6 months; mixed over the drawings' plug-ins, $P_{\mathrm{U}}(5) = 0.0105$, the longest interval is again 6 months and the cheapest 5. The plug-in is optimistic by a month, because a single path has none of the short tail of the P–F interval.

The costs are plotted against the interval in Figure 18.5. The horizontal axis is the interval and the left vertical axis the expected cost of the program in reais, stacked as the visits (blue) and the shortfall (red); the right vertical axis carries the inspection cost per year, the dashed line. The cheapest and the 5 % intervals are marked. The stack falls from one month to four as the visits fall, and rises after five as the shortfall overtakes them, while the cost per year keeps falling.

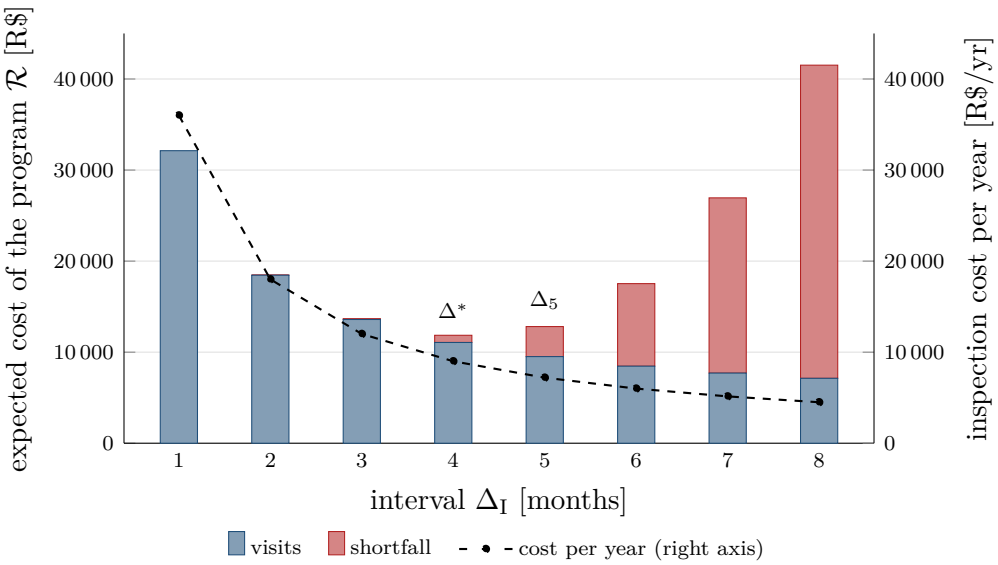


Figure 18.5: The expected cost of the program against the interval, equation (18.12), stacked into the visits (blue) and the shortfall (red), in reais from today to the order (left axis), for 1 to 8 months; at 12 months it is R$ 148 224. Dashed, right axis: the inspection cost per year (18.6). Marked: the cheapest interval $\Delta^* = 4$ and the 5 % interval $\Delta_5 = 5$.

18.3.4 The calendar and the rule of thumb

Asset manager's question. What do the calendar and the rule of thumb give? The naive schedules are priced under the same posterior as the chapter's own.

Specialized form. The question is equation (2.23) again, in the two forms of equations (18.12), at the naive actions: the calendar's $\Delta_{\mathrm{I}} = 12$, and the rule of thumb's half of the handbook's P–F interval, $\Delta_{\mathrm{I}} = 3$. One choice enters here: the handbook's 6 months is the planner's figure for air-side fouling, stated as an input of the baseline and not as a measurement of this hall. The stronger form of the rule, half of this posterior's own median P–F interval, 7.93/2 rounded down to whole months, gives the same 3 months.

Building blocks. The values of Tables 18.4 and 18.5 at the naive intervals; the first passage's distribution at the calendar's first test.

The naive schedules priced. They stand beside the chapter's two answers in Table 18.6. Each row is a schedule; the columns give its interval, its P_{U} phase-averaged and dated, its inspection cost per year, its expected cost and its probability of an early order.

Table 18.6: The naive schedules and the chapter's two answers, priced under the same posterior, equation (18.12). The calendar is a dated schedule, its first test in December of year 4, and its dated value is the one it reports.

Schedule	Δ_{I} [mo]	P_{U}, phases averaged (dated) [–]	cost per year [R$/yr]	$\mathcal{R}$ [R$]	early order [–]
the calendar (annual, in December)	12	0.3147 (**0.3537**)	3 000	148 224	0.0657
half the handbook's P–F interval	3	0.0017 (0.0017)	12 000	13 677	0.2383
half the posterior's median P–F, 7.93/2	3	0.0017 (0.0017)	12 000	13 677	0.2383
the 5 % interval	5	0.0302 (0.0297)	7 200	12 812	0.1516
the cheapest	4	0.0101 (0.0105)	9 000	**11 855**	0.1856

The calendar misses about one crossing in three. The calendar tests every December, so it is the dated annual schedule, and its probability of an undetected crossing is 0.3537; at an annual interval with its phase averaged it would be 0.3147. Its first test is window 48, and the hall has crossed before it with $\Pr(t^* \le 48 \mid \mathbf{y}) = 0.2771$: in more than one path in four the calendar's first look comes after the hall has run short. The annual program's expected cost is R$ 148 224, against the cheapest program's R$ 11 855. The annual test is the wrong schedule for a hall in its last year before the limit.

The rule of thumb is safe here by the luck of its figure. The handbook's 6 months is shorter than this hall's median P–F interval, 7.93 months, so half of it, 3 months, misses with probability 0.0017. Its program costs R$ 13 677 against the four-month program's R$ 11 855. On a hall whose P–F interval were shorter than the handbook's the rule would not be safe, and one path in four of this hall's posterior, 0.2581, is such a hall.

18.3.5 The tests before the October shutdown

Asset manager's question. If the hall is carried to the October shutdown with no work done, which tests before it does it need? The question asks for test windows, not an interval.

Specialized form. Suppose no work is done before the October shutdown and the coils are cleaned there, the plan of Chapter 14's entry, the plan Chapter 17's step 5 prefers if coil 1's rate is constant, and the calendar's. The program then ends at window 46, a crossing by window 46 counts, as in Chapter 14, and the tests can fall at windows 37 to 45. The question is equation (2.23) with the indicator of a crossing by window 46 that no test caught as the loss, for a dated set $\mathcal{S} \subset \{37, \dots, 45\}$:

$$P_{\mathrm{U}}(\mathcal{S}) = \sum_{M} P_M \, \mathbb{E}\Big[\mathbf{1}\big[t^*(\mathbf{z}) \le 46\big] \prod_{s \in \mathcal{S},\ s < t^*} p_{\mathrm{no},s}(\mathbf{z}) \mid \mathbf{y}, M\Big], \tag{18.15}$$

the probability (18.5) with the horizon at the shutdown. Its minimum is found by enumerating every set of one, two and three windows. One choice enters here: a program that ends in ten months is a set of test windows, not a periodic interval. A periodic interval is the wrong form for it, because its dated schedule is not monotone in the interval: phase-averaged, 7 months gives 0.0475, but the dated 7-month schedule's single test falls at window 43 and gives 0.0284, while the dated 6-month schedule's, at window 42, gives 0.0152. The price of a test before the

shutdown is not only its visit: a test that reads at or below G_{P}, or finds a crossing, raises the order before window 46 and brings the work forward off the outage, at the OFF_WINDOW premium of R$ 12 000. Its expected premium is that premium times the probability that some test of $\mathcal{S}$ raises an order, $1 - \mathbb{E}\big[\mathbf{1}(t^* > \max \mathcal{S}) \prod_{s \in \mathcal{S}} p_{\mathrm{no},s} \mid \mathbf{y}\big]$.

Building blocks. The factors (18.4) at windows 37 to 45 on each draw; the indicator of a crossing by window 46; the enumeration of the sets; the OFF_WINDOW premium.

The best dates. With no test the hall crosses before the cleaning with probability 0.1595. The probability of an undetected crossing for each single test window and for the best pair and triple, with their cost, is in Table 18.7. Each row is a set of test windows; the columns give its probability (18.15) and the visits' cost to the shutdown.

Table 18.7: The tests before the October shutdown, equation (18.15): the probability of a crossing by window 46 that no test caught, for one test at each window from 37 to 45 and for the best two and three; with no test it is 0.1595.

Tests [window]	$P_{\mathrm{U}}(\mathcal{S})$ [–]	visits [R$]
none	0.1595	0
37 (Jan Y4)	0.1473	3 000
38 (Feb Y4)	0.1247	3 000
39 (Mar Y4)	0.0891	3 000
40 (Apr Y4)	0.0504	3 000
41 (May Y4)	0.0232	3 000
42 (Jun Y4)	**0.0152**	3 000
43 (Jul Y4)	0.0284	3 000
44 (Aug Y4)	0.0604	3 000
45 (Sep Y4)	0.1055	3 000
the best two: 40 and 43 (Apr and Jul Y4)	0.0018	6 000
the best three: 40, 42 and 44 (Apr, Jun and Aug Y4)	0.0004	9 000

One test in June of year 4, window 42, keeps the hall under 5 % until the cleaning, at 0.0152, for R$ 3 000; the next best is window 41, at 0.0232. Two tests, at windows 40 and 43, take it to 0.0018, and three, at 40, 42 and 44, to 0.0004. An early test misses the late crossings and a late test the early ones. The paths that cross by the shutdown, 0.1595 of the posterior, pass their potential failure at window 39.27 at the median, with a band of 37.32 to 41.05, window 38 (Feb Y4) to window 42 (Jun Y4). The best single date, window 42, sits just past the upper end of that band: a test there comes after the potential failure of nineteen in twenty of those paths, while their crossings, by window 46 at the latest, are still ahead. The full posterior gives 0.0159 for the June test, 0.0018 for the pair and 0.0004 for the triple, with 0.1589 for no test.

The tests' price is the order they bring forward. The single test at window 42 raises an order before the shutdown with probability 0.3286, an expected OFF_WINDOW premium of R$ 3 943 on top of its R$ 3 000; the pair at windows 40 and 43 with 0.4175, R$ 5 010; and the triple with 0.5067, R$ 6 081. Each is more probable than a crossing by the shutdown, 0.1595, so many of those orders fall on paths that would have held to the outage, which a test's reading near G_{P} cannot tell from the paths that would not. Whether to pay the premium on such a reading or to hold the work for the outage is the trigger's question of Chapter 19, on the record's own monthly posterior; this chapter reports the program's probability of a missed crossing and its price.

The single tests are plotted in Figure 18.6. The horizontal axis is the test window and the vertical axis P_{U}; each bar is one test at that window, the dashed line is the probability with no test, the dotted lines the best pair and triple, and the gray line the 5 % limit. The bars fall from window 37 to a minimum at 42 and rise again to 45; only windows 41 to 43 are under the limit.

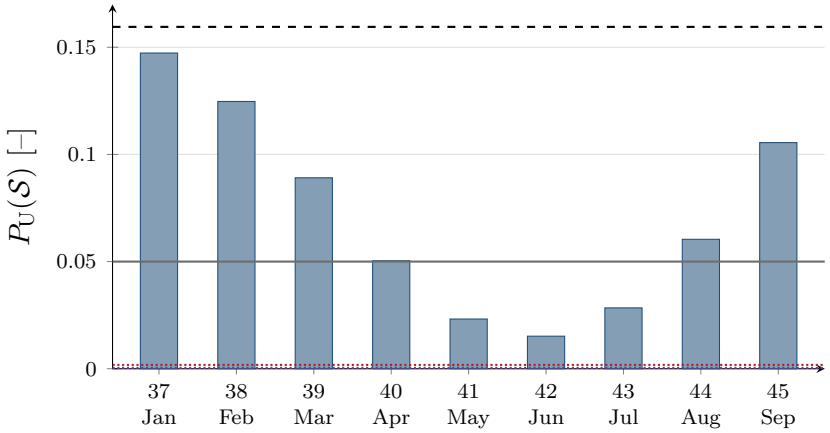


Figure 18.6: The probability of a crossing by the October shutdown that no test caught, equation (18.15), for one test at each window from 37 (Jan Y4) to 45 (Sep Y4). Dashed: no test, 0.1595. Dotted: the best pair (windows 40 and 43, 0.0018) and triple (40, 42 and 44, 0.0004). Gray: the 5 % limit.

18.4 The checks (step 5)

Every answer of step 4 came from the Laplace draws with the tests' errors in closed form; the checks hold it against the other routes, the stated inputs and the stated truth.

18.4.1 The full posterior and the routes

The routes stand side by side in Table 18.8. Each row is a route; the columns give P_{U} at four, five and six months, the longest interval under 5 % and the cheapest interval. The stated truth's own path, with only the tests' errors, is the last row, a check marker.

Table 18.8: The probability of an undetected crossing and the two chosen intervals by route. The reported route is the Laplace draws with 40 000 per drawing; the routes at 10 000 per drawing share the seed 500.

Route	$P_{\mathrm{U}}(4)$ [–]	$P_{\mathrm{U}}(5)$ [–]	$P_{\mathrm{U}}(6)$ [–]	Δ_5 [mo]	Δ^* [mo]
Laplace draws, 40 000 per drawing (reported)	0.0101	0.0302	0.0621	5	4
full posterior by reweighted draws, 20 000 per drawing	0.0103	0.0304	0.0624	5	4
Laplace draws, 10 000 per drawing	0.0102	0.0303	0.0623	5	4
forecast windows without wander, 10 000 per drawing	0.0049	0.0189	0.0460	6	5
plug-in, the leading drawing's MAP path	0.0005	0.0047	0.0195	6	–
the stated truth's path, the tests' errors only	0.0327	0.1135	0.2275	4	–

The full posterior. By the reweighted draws of Section 18.2.4 the gaps in P_{U} against the Laplace draws are +0.0000 at one to three months, +0.0001 at four, +0.0002 at five, +0.0004 at six and +0.0006 at twelve, all in the same direction and at most twice the Laplace route's standard errors of Table 18.4, on half as many draws: the full posterior is at most slightly the more pessimistic. Both give $\Delta_5 = 5$ and $\Delta^* = 4$ months, and $\mathcal{R}(4)$ is R$ 11 883 against R$ 11 855. The chapter calls this route the full posterior. The test's rows are Chapter 5's rows of one window at another grade, whose limit of exact rows and curvature term were checked there (Section 5.4.2), and the forecast is Chapter 14's, so no new check of either is needed.

The wander. Without the wander after window 36 the paths are smoother, fewer of them cross between two tests, and Δ_5 is 6 months and Δ^* 5. The wander is worth a month of interval, in the same direction as for Chapter 14's date: without it the first passage has median 52.72 and $\Pr(t^* \le 46 \mid \mathbf{y}) = 0.1489$, against the reported route's 52.55 and 0.1588 at the same 10 000 draws. The plug-in is the extreme of the same effect, with neither the wander nor the record's spread.

18.4.2 Sensitivity to the levels, the test and the costs

Each stated input and each choice that could flip a decision is varied in Table 18.9, on the 10 000 draws per drawing with seed 500. Each row is a variant; the columns give $P_{\mathrm{U}}(5)$, the longest interval under 5 %, the cheapest interval with its expected cost, and whether a decision flips against the reported row.

Table 18.9: The choices and the stated inputs, 10 000 draws per drawing (seed 500); the rows priced at FO, discounted and at the costs' break-evens use the reported draws. "Flips" compares Δ_5 and Δ^* with the reported row's.

Variant	$P_{\mathrm{U}}(5)$ [–]	Δ_5 [mo]	Δ^* ($\mathcal{R}$) [mo] ([R$])	flips [–]
reported, at 10 000 per drawing	0.0303	5	4 (11 982)	
$m_{\mathrm{P}} = 0.25$ kW/K	0.2793	2	2 (29 210)	yes
$m_{\mathrm{P}} = 0.75$ kW/K	0.0011	8	6 (6 389)	yes
$m_{\mathrm{P}} = 1.00$ kW/K	0.0006	10	6 (4 157)	yes
the test at 0.05 K ($\sigma_{\mathrm{I}} = 0.0597$ kW/K)	0.0195	6	5 (11 340)	yes
the test at 0.2 K ($\sigma_{\mathrm{I}} = 0.1563$ kW/K)	0.0471	5	4 (13 130)	no
a visit at one INSPECT, R$ 1 500	0.0303	5	4 (6 387)	no
early orders priced at FO			4 (12 412)	no
costs discounted at 8 %/yr			4 (10 795)	no
SHORTFALL below R$ 92 782 a month			5	yes
SHORTFALL above R$ 539 218 a month			3	yes
a visit dearer than R$ 4 850			5	yes
a one-window lead time: a test counts only before $t^* - 1$	0.0739	4	3 (14 670)	yes
the core family: $\mathcal{C}_0$, $\mathcal{C}_1$ × (none, coil 2's step)	0.0220	6	4 (12 805)	Δ_5 only
constant rates only, $\mathcal{C}_0$	0.0027	8		yes
coil 1 curving only, $\mathcal{C}_1$	0.0462	5		no

The potential-failure level moves the answer most. With the level half as far above the limit, 0.25 kW/K, the hall needs a test every 2 months; at 0.75 kW/K the 5 % interval is 8 months and at 1.0 kW/K 10 months, with the cheapest at 6 for both. The level is a stated input of this chapter. The costs at another level are not comparable with the reported ones: a higher level ends the program sooner, by cleaning sooner, and the life the earlier cleaning gives up is not priced here.

The test, the costs and the family. A test at 0.05 K allows 6 months and is cheapest at 5; one at 0.2 K keeps both answers. Of the costs, neither a visit at one INSPECT, nor early orders at FO, nor the discount moves the cheapest interval from 4 months; SHORTFALL would have to fall below R$ 92 782 a month, or rise above R$ 539 218, or a visit cost more than R$ 4 850, for it to move (Section 18.3.3). With a one-window lead time, a test counting only before $t^* - 1$, P_{U} is 0.0367 at four months, 0.0739 at five and 0.1190 at six, the longest interval under 5 % is 4 months and the cheapest 3, at R$ 14 670. Both decisions flip, each one month shorter. The core family, the four drawings of Chapter 14's fresh-record check, gives $\Delta_5 = 6$ and keeps $\Delta^* = 4$. The curvature states read alone bound the answer: 8 months at constant rates and 5 with coil 1 curving.

18.4.3 The forecast as late as fresh records show

Chapter 14's step 5 found that at the record's stated truth a three-year forecast is late in about two records out of three (Section 14.4.5): the cause is the record's length, with a part in Chapter 11's curvature prior, and its answer is the refit at window 45. The family is kept as the reported posterior, as the owner ruled, and the finding is carried here as a sensitivity row, in Table 18.10. Each row is a test of the finding; the columns give the longest interval under 5 % and the cheapest, each with the stated truth's own P_{U} at that interval.

Table 18.10: The calibration finding carried as a sensitivity row. The truth's P_{U} is the probability that the truth's own path, with only the tests' errors, goes undetected at the interval chosen. Fresh records: 40 records at the truth, the core family refitted on each, 5 000 draws per drawing.

Test	Δ_5 [mo]	truth's P_{U} at Δ_5 [–]	Δ^* [mo]	truth's P_{U} at Δ^* [–]
the reported family	5	0.1135	4	0.0327
this record, coil 1's curvature from a knee at window 24, prior sd $0.12\,\mathrm{yr}^{-2}$ (18 refits)	3	0.0035	3	0.0035
the truth's own path, the tests' errors only	4	0.0327	–	–
fresh records at the truth: the interval each would choose	median 7, 5–95 % 5 to 10	median 0.3327; above 0.05 in 0.975	median 6	median 0.2275; above 0.05 in 0.875

The drill's answer flips at the truth. The truth's P–F interval is 4.56 months: its potential failure comes at window 41.84 (window 42, Jun Y4) and its crossing at 46.40 (window 47, Nov Y4). That sits at the 0.062 quantile of the posterior's P–F interval: the truth is a fast path, the kind the curvature makes. On it the 5-month interval leaves the crossing undetected with probability 0.1135, its phases ranging from 0.436 to 0.000, while 4 months holds it at 0.0327, the longest interval that does. The dated 5-month schedule, tests at windows 41 and 46, happens to catch it. The cost-chosen 4 months holds at the truth: on this record the decision that flips is the drill's 5 % interval, not the costs'. With coil 1's curvature from the truth's own knee and the wider prior, this record's family crosses earlier, median 44.37 and $\Pr(t^* \le 46 \mid \mathbf{y}) = 0.6354$, and asks for 3 months on both rules, with $P_{\mathrm{U}}(5) = 0.1751$.

Over fresh records both rules choose too long an interval. Refitted on 40 fresh records at the truth, the 5 % rule's interval has median 7 months, from 5 to 10 over the records (4 months once, 5 six times, 6 six, 7 nine, 8 nine, 9 six, 10 once and 12 twice), and it leaves the truth above 5 % in 0.975 of the records; the cheapest interval has median 6 months and leaves the truth above 5 % in 0.875. The records' forecasts are late: the median over the records of

each one's median crossing is window 65.80, and this record is one of the earlier ones: its core family's 6 months sits at the 0.325 quantile of the records' choices, and its core family's cheapest is 4 months.

The window-45 refit settles it before the shutdown. In Chapter 14's 20 continuations of the record at the truth to window 45, refitted there with nothing later, the family puts $G_{\mathrm{tot},45}$ at a median of 10.9566 kW/K against the truth's 10.9199. It reads $\Pr(G_{\mathrm{tot},45} \le G_{\mathrm{P}} \mid \mathbf{y}_{0:45})$ at a median of 0.9962, above 0.5 in every continuation, while $\Pr(G_{\mathrm{tot},45} \le G^{*} \mid \mathbf{y}_{0:45})$ is 0.0427: the record itself shows the potential failure in September of year 4, a month before the shutdown and two before the truth's crossing.

18.4.4 The answers against the stated truth

The record is constructed from a stated truth (Section 18.6.4). The truth's total at windows 36 to 47 is 11.8113, 11.7257, 11.6363, 11.5432, 11.4467, 11.3470, 11.2443, 11.1387, 11.0305, 10.9199, 10.8071 and 10.6923 kW/K: it passes G_{P} in June of year 4 and G^{*} in November. Its own P_{U}, with only the tests' errors, is 0.0035 at three months, 0.0327 at four, 0.1135 at five, 0.2275 at six and 0.6105 at twelve. It is the last row of Table 18.8, the dotted curve of Figure 18.3 and the dotted marker of Figure 18.2. The reported answers place it within their bands: its P–F interval at the 0.062 quantile, inside the band of 4.44 to 20.37 months, and its crossing, as Chapter 14 found, at the forecast's 0.218 quantile. The calendar's first test, in December of year 4, would come a month after the truth's crossing; the June test of the shutdown program falls in the month of its potential failure, and the truth's total there, 11.2443 kW/K, lies just below G_{P}.

18.5 The report (step 6)

The annual condition test misses the crossing about one time in three: there is a 0.28 chance that the hall has crossed the limit before December's test comes. A condition test every 5 months keeps the chance of an undetected crossing under 5 %, at 0.030. At the stated costs the program is cheapest at 4 months, R$ 9 000 a year and R$ 11 855 over the program against R$ 12 812 at 5: a month of shortfall costs R$ 150 000 and a visit R$ 3 000, so the costs ask for more than the 5 % limit does. If no work is done before the October shutdown, one test in June of year 4 carries the hall to the cleaning with a 0.015 chance of a missed crossing; it raises an order before the shutdown, at the R$ 12 000 premium of work off the outage, with probability 0.3286. Under Chapter 17's action, coil 1 replaced in January of year 4, no test is needed before that work. The handbook's rule, every 3 months, misses little here, at 0.0017, but costs R$ 1 822 more than every 4. At the stated truth the hall is a fast path, with a P–F interval of 4.6 months: the 5-month interval would miss it with probability 0.11, while 4 months holds, at 0.03, and the refit at window 45 sees the potential failure before the shutdown. Both answers assume the work follows the order at once: if it needs a month, both flip one month shorter, to 4 months for the 5 % limit and 3 for the costs.

18.5.1 The answers and their bands

The answers follow Table 18.1's order in Table 18.11. Each row is a question; the second column gives the reported value, from the Laplace draws with the tests' errors in closed form, the third the full posterior's value beside it where it was formed, and the fourth the band or what moves the decision.

The answers rest on what the calendar and the handbook cannot give: the P–F interval of this hall's own path, with its band, and the open question that shortens it. Within the

Table 18.11: The report on the condition-test program for the two-unit hall's coils at window 36 (Dec Y3). Reported values from the forecast's Laplace draws, mixed over the family, with the tests' errors in closed form; bands are 5–95 %.

Asset manager's question	Reported value	Full posterior	Band or what moves it
How often should the exchanger be inspected so that the probability of an undetected crossing stays under 5 %?	every 5 [mo]: $P_{\mathrm{U}} = 0.0302$ [–] (se 0.0001); R$ 7 200 a year	5 [mo], $P_{\mathrm{U}} = 0.0304$ [–]	4 to 10 [mo] over the drawings; 8 [mo] at constant rates, 5 with coil 1 curving; at the truth 4 [mo]; with a one-window lead 4 [mo]
How long is this hall's P–F interval?	median 7.93 [mo]; potential failure at window 44.27, window 45 (Sep Y4)	median 7.94, band 4.42 to 20.37 [mo]	4.44 to 20.37 [mo]; 6.46 [mo] with coil 1 curving, 12.11 at constant rates
What does each interval cost, which is cheapest, and how firm is that choice?	cheapest every 4 [mo], $\mathcal{R}$ = R$ 11 855, against R$ 12 812 at 5; $\Pr(L(4) < L(5)) = 0.3854$ [–]; EVPI R$ 3 524	4 [mo], $\mathcal{R}$ = R$ 11 883	4 [mo] with early orders at FO, discounted, or at one INSPECT a visit; 5 [mo] below SHORTFALL R$ 92 782 or above a visit of R$ 4 850; 2 to 6 [mo] over m_{P} from 0.25 to 1.0 [kW/K]
What do the calendar and the rule of thumb give?	the calendar (dated, December): $P_{\mathrm{U}} = 0.3537$ [–], $\mathcal{R}$ = R$ 148 224; every 3 [mo]: 0.0017 [–], R$ 13 677	annual, phases averaged, 0.3153 [–]	crossed before December's test 0.2771 [–]; annual with phases averaged 0.3147 [–]
If the hall is carried to the October shutdown with no work done, which tests before it does it need?	one test at window 42 (Jun Y4), $P_{\mathrm{U}} = 0.0152$ [–], R$ 3 000; an order before the shutdown 0.3286 [–], premium R$ 3 943	window 42: 0.0159; pair 0.0018; triple 0.0004 [–]	windows 40 and 43: 0.0018; 40, 42 and 44: 0.0004 [–]; no test: 0.1595 [–]; none needed under Chapter 17's January action

constant-rate drawings the hall could wait 8 months between tests; coil 1's possible acceleration is what brings the interval to 5, and the costs to 4.

18.5.2 The entry in the maintenance log

The condition-test program is the next entry of the two-unit hall's maintenance log, after Chapter 14's forecast and Chapter 17's decision at the same window. It uses nothing later than window 36, and in the values of Table 18.11 it reads:

> *Window 36 (Dec Y3), the condition-test program for the hall's coils, from Chapter 14's forecast of the capacity limit ($G_1 + G_2 \geq 10.761\,\mathrm{kW/K}$). A work order is raised when a reference condition test reads $G_1 + G_2 \leq 11.261\,\mathrm{kW/K}$ (0.5 kW/K above the limit).*
> P–F interval on this hall: 7.9 months at the median (4.4 to 20.4); shorter when coil 1 accelerates (6.5).
> Undetected crossing: the calendar's December test 0.354; every 5 months 0.030; every 4 months 0.010; every 3 months 0.002.
> Action: under Chapter 17's action (coil 1 replaced in Jan Y4), no condition test before that work; the program for the renewed hall is not set here. If no work is done before the October Y4 shutdown, one reference condition test of both coils in June Y4 (window 42), R$ 3 000 (undetected 0.015; an order before the shutdown, off the outage, probability 0.3286). If the work is not done at the shutdown either, test every 4 months from today, R$ 9 000 a year. The calendar's annual test is withdrawn for these coils until they are worked on.
> Caution: the interval rests on Chapter 14's forecast, which at three years is more often late than early; the window-45 refit re-reads it before the shutdown.
> Next measurement: the refit at window 45 (Chapter 14's), one of the monthly refits of the record (Chapter 19).

The action takes Chapter 17's decision as it stands: its January work needs no test before it. The June test is the program for the hall with no work done before the shutdown, the plan Chapter 17 prefers if coil 1's rate is constant and the calendar's, and the fall-back of 4 months meets the 5 % limit and the costs and holds at the truth.

Common mistake

Setting the interval at a fraction of the P–F interval. A fraction of one summary of the paths gives neither the probability of a missed crossing nor the cost of the program. Half this hall's median P–F interval, 7.93/2, gives 3 months: safe, at $P_{\mathrm{U}} = 0.0017$, but R$ 1 822 dearer than the 4 months the costs choose. Both numbers come from the paths, the probability and the expected cost of each interval, not from the median.

18.6 The next measurement (step 7)

The report leaves the program at a visit every four months and an expected cost of R$ 11 855, most of it the visits themselves. A further measurement either reads the level more sharply at each visit or reads it every window without a visit, and each is priced by the expected cost of its own cheapest program under the same posterior.

18.6.1 The candidates

A candidate changes the test's rows (18.7), and so σ_{I} by equation (18.9), or the visit's price, or both: a crew test at a finer grade keeps C_V and lowers σ_{I}; the hall's own sensors read every window as a level test, with unit 1's drifting discharge reading left out, give $\sigma_{\mathrm{I}} = 0.3151\,\mathrm{kW/K}$ at no visit; and reference-grade air sensors installed and read every window, the permanent instrument at the place of the question, give the reference grade at each window's own inputs, at no visit. Each candidate's P_{U} and $\mathcal{R}$ are equations (18.12) with its σ_{I} and C_V, on the 10 000

draws per drawing with seed 500, where the reported program costs R$ 11 982 at 4 months. Perfect information on the path, equation (17.9), on the reported 40 000 draws per drawing, bounds what any reading of the path is worth to the choice of the condition test's interval.

The candidates are ranked in Table 18.12. Each row is a candidate; the columns give its P_{U} at its own cheapest interval, its longest interval under 5 %, the expected cost of its cheapest program, its probability of an early order there, and its saving against the reported program's cheapest.

Table 18.12: Candidate measurements for the condition-test program, 10 000 draws per drawing (seed 500); the saving is taken from the reported program's R$ 11 982 at 4 months on the same draws. An instrument read every window is evaluated monthly only, with no visit, so it has no longest interval. The perfect-information row is read on the reported draws, 40 000 per drawing.

Candidate	P_{U} at its interval [–]	Δ_5 [mo]	$\mathcal{R}$, cheapest [R$]	early order [–]	saving [R$]
the condition test at 0.05 K	0.0195 (5 mo)	6	11 340	0.0894	641
the hall's own sensors read every window as a level test (0.3 K, unit 1's discharge out)	0.0001 (monthly)	–	3	0.8545	11 979
reference-grade air sensors installed, read every window (0.1 K)	0.0000 (monthly)	–	0	0.5364	11 982
perfect information on the path, over the condition test's intervals					3 524

A finer crew instrument is worth R$ 641 a program. Reading the level every window, without a visit, takes the missed crossing to nearly zero even at the hall's own grade: a monthly look at 0.32 kW/K catches what a four-monthly look at 0.10 kW/K catches, because the P–F interval is longer than three months on all but 0.0003 of the paths. The visit, not the grade, is the cost. The monthly look's price is early orders: the hall's own grade raises an early order in 0.8545 of the programs and the reference grade in 0.5364. A program ends at its first order, so at FO each it costs at most R$ 3 000, and the monthly look stays far ahead of the visits. Perfect information does not bound these instruments, because they change the actions: they read without a visit. Their saving is the installed price below which each pays over this program.

18.6.2 The recommendation

Under Chapter 17's action, coil 1 replaced in January of year 4, no test comes before that work, and the renewed hall's program is read again after it. If no work is done before the October shutdown, test once, in June of year 4; if the work is not done at the shutdown either, test every 4 months from today, in April and August of year 4 before it, and re-read the program at the window-45 refit. Before buying permanent sensors, ask the record itself: a decision level on its own monthly posterior reads every window with no visit, Chapter 19's trigger. Installed reference-grade sensors are worth up to R$ 11 982 over this program as the permanent instrument. The side-by-side check of unit 1's discharge sensor stays open for Chapter 7's reason; with it made, the hall's own discharge reading of unit 1 could join the monthly level test.

18.6.3 What the analysis showed

The calendar said December comes in time, and the handbook said every three months misses nothing; the posterior says the calendar's December test misses the crossing with probability 0.354, that every 3 months is safe but dearer than needed, that 5 months meets the 5 % limit and that the costs ask for 4. The forecast already carried the hall's path, and the P–F interval

is two first passages on it, so the question needed no new fit of the record: the test's rows gave one number, its spread of the coils' sum, and integrating the tests' errors in closed form over the forecast's draws gave the probability of an undetected crossing for any schedule, while the costs read the same paths. The probability added what no handbook gives, the band of this hall's P–F interval and the open question that shortens it, coil 1's acceleration, which is also what makes the truth a fast path. And the same computation showed that the visit, not the grade of the test, is what an inspection program pays for, which points to reading the level every window rather than to a better instrument at each visit. It did not model the life an earlier cleaning gives up, the level G_{P} set against the costs, a lead time beyond one window, the renewed hall after Chapter 17's January work, or what a test and the monthly record tell the forecast.

18.6.4 Limits of the model

A constructed record. The condition record is generated from a stated truth, with the stated sensor errors, and step 5 compares the answers with that truth (Section 18.4.4). It is not field data. At that truth the three-year forecast is late in two records of three (Chapter 14), and the 5 % interval read from it is too long in 0.975 of fresh records; the refit at window 45 corrects it.

The potential-failure level as stated. The level G_{P} is a stated input, and the answer moves with it more than with anything else, from 2 months at 0.25 kW/K to 10 at 1.0. No chapter calibrates it: Chapter 19's trigger is a level on a probability of the hall going short, not on G_{P}.

The hall it is for. The program is valued for the hall with no work done before the shutdown. Under Chapter 17's action, coil 1 replaced in January of year 4, no test is needed before that work, and the program for the renewed hall after it, with a new coil whose path the record has not yet read, is not computed.

No lead time. A test is taken to prevent the shortfall if it raises the order at any time before the crossing, so the work is assumed to follow at once. With a lead of one window the 5 % interval becomes 4 months and the cheapest 3 (Table 18.9); a longer lead would move them further.

No learning from a test or from the monthly record. The program ignores what a test tells the forecast, and it ignores the hall's own monthly readings after window 36, which go on being logged. A test that reads above G_{P} narrows the path, and so does every new window of the record, and the next interval could be read again from the narrowed posterior; that sequential program is not built here; the monthly record is read by Chapter 19's trigger (Section 18.1.2).

The costs priced. The costs are INSPECT and SHORTFALL. The cleaning, the life an earlier cleaning gives up and the disturbance of a visit are not priced, so programs at different levels or ending at different dates are not compared.

The forecast's limits. Chapter 14's limits carry over: the family is Part III's two open questions and no other form, there is no maintenance until the order, and the path wanders as the record's does.

The next chapter. The level at which a monthly refit of the record raises the order, a probability that the hall is below its limit after a one-window lead time, set against the cost of a false work order and of missed degradation, is Chapter 19's. Chapter 22's power curve says the trend log sees a 5 %/yr loss by window 45 with a probability of detection of 0.995, and needs 30 windows on coil 1 for a power of 0.9 (Sections 22.3.1 and 22.3.2): the log sees that the coils are losing, but not when they cross G_{P}, which is this chapter's test.

18.7 What to retain

Definition 18.1 (Probability of an undetected crossing)**.** For a schedule $\mathcal{S}$ of condition tests with spread σ_{I} on the monitored level and a potential-failure level G_{P} above the functional limit, the probability that the crossing comes before any test raised the order is the posterior mean of a product of normal tails over the tests before the crossing, $P_{\mathrm{U}} = \mathbb{E}\big[\prod_{s\in\mathcal{S},\ s<t^*} \Phi\big((G_{\mathrm{tot},s}-G_{\mathrm{P}})/\sigma_{\mathrm{I}}\big) \mid \mathbf{y}\big]$, with the tests' errors integrated in closed form given the path and the paths by the forecast's draws; for an interval it is averaged over the phase. The longest interval keeping it under a limit, and the interval of least expected cost, are read from the same draws. With a perfect test and a steady fall it reduces to the catch probability $\min(\Delta_{\mathrm{PF}}/\Delta_{\mathrm{I}}, 1)$ of the P–F interval $\Delta_{\mathrm{PF}} = t^* - t_{\mathrm{P}}$.

Three conclusions carry beyond the hall. A P–F interval is a margin over a rate of fall, read as two first passages on the forecast's path, so it has a band, and the short end of the band, not the median, sets the interval. The probability of an undetected crossing and the expected cost are two answers of one computation, and the cheaper interval in expectation can be the dearer one on most paths, because the expectation is carried by the few paths a longer interval misses. And a program that ends at a planned outage is a set of dated tests, best placed where the paths that cross before the outage pass their potential failure.

Table 18.13 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 18.13: The terms of the inspection interval and where each is defined. Units for the two-unit hall.

Term	Expression	Units	Meaning	Defined in
Potential-failure level	$G_{\mathrm{P}} = G^{*} + m_{\mathrm{P}}$, a pinned requirement node	kW/K	the level at which a test raises the order	18.1.2
P–F interval	$\Delta_{\mathrm{PF}} = t^{*} - t_{\mathrm{P}}$, eq. (18.1)	mo	two first passages on one path	18.1.1
Catch probability	$\min(\Delta_{\mathrm{PF}}/\Delta_{\mathrm{I}}, 1)$, eq. (18.3)	–	a perfect test, a steady fall, a uniform phase	18.1.1
Phase of a schedule	s_1 in eq. (18.2)	window	averaged for an interval; fixed for a dated schedule	18.1.1
Rows and posterior				
Condition test's rows	eq. (18.7)	–	one window at the reference grade	18.1.4
Test's spread of the sum	$\sigma_{\mathrm{I}}^{2} = \mathbf{g}_{G}^{\mathsf{T}} \boldsymbol{\Sigma}_{\mathrm{I}} \mathbf{g}_{G}$, eq. (18.9)	kW/K	0.1033 at 0.1 K	18.2.2
No order at a test	$p_{\mathrm{no},s} = \Phi((G_{\mathrm{tot},s} - G_{\mathrm{P}})/\sigma_{\mathrm{I}})$, eq. (18.4)	–	given the path	18.1.1
Undetected crossing	product of $p_{\mathrm{no},s}$ before t^{*}, eq. (18.5)	–	given the path	18.1.1
Visits, shortfall, early order	eq. (18.10)	–, mo, –	the tests' errors integrated	18.2.3
Questions and answers				
P_{U} and $\mathcal{R}$ of an interval	eq. (18.12)	–, R$	phase-averaged over the draws	18.2.3
The 5 % interval	Δ_{5}, eq. (18.14)	mo	the drill's answer	18.3.2
The cheapest interval	$\Delta^{*} = \arg\min \mathcal{R}$	mo	cheaper in expectation, dearer on most paths	18.3.3
The plug-in's optimism	the MAP path, no wander	–	a month too long here	18.3.3
The calendar's miss	$\Pr(t^{*} \le 48 \mid \mathbf{y})$	–	crossed before the annual test	18.3.4
Dated tests before an outage	eq. (18.15)	–	a set of windows, not an interval	18.3.5

18.8 Exercises

Exercise 18.1. *Application.* From the condition test's Jacobian of Table 18.2, compute σ_{I} at the reference grade of 0.1 K and at the hall's own 0.3 K, and the margin m_{P} in standard deviations of each. (Divide the four temperature rows by the grade, form $\boldsymbol{\Sigma}_{\mathrm{I}} = (\mathbf{J}_{\mathrm{I}}^{\mathsf{T}}\mathbf{J}_{\mathrm{I}})^{-1}$ and $\sigma_{\mathrm{I}}^{2} = \mathbf{g}_{G}^{\mathsf{T}}\boldsymbol{\Sigma}_{\mathrm{I}}\mathbf{g}_{G}$ with $\mathbf{g}_{G} = (0,\ 5.6305,\ 5.6305,\ 0)$: $\sigma_{\mathrm{I}} = 0.1033$ and $0.1956\,\mathrm{kW/K}$, so m_{P} is 4.84 and 2.56 standard deviations.)

Exercise 18.2. *Derivation.* For a steady fall, a perfect test and a phase uniform within one interval, show that the probability of catching the crossing is $\min(\Delta_{\mathrm{PF}}/\Delta_{\mathrm{I}}, 1)$, and evaluate it for $\Delta_{\mathrm{PF}} = 4$ months at intervals of 3, 6 and 12 months. (With a perfect test an order is raised exactly by a test in $[t_{\mathrm{P}}, t^{*})$; the first test at or after t_{P} lies U after it, U uniform on $[0, \Delta_{\mathrm{I}})$, and the crossing is caught when $U < \Delta_{\mathrm{PF}}$, whose probability is $\min(\Delta_{\mathrm{PF}}/\Delta_{\mathrm{I}}, 1)$: 1.0000, 0.6667 and 0.3333.)

Exercise 18.3. *Formulation.* Write "has the hall crossed before the calendar's December test?" in the form of Table 18.1 and read it; then read the rule of thumb's probability of an undetected crossing. (A prediction under a threshold on the trend, equation (2.21), second line at $s = 48$: $\Pr(t^{*} \le 48 \mid \mathbf{y}) = 0.2771$; the rule's interval of 3 months gives $P_{\mathrm{U}}(3) = 0.0017$.)

Exercise 18.4. *Application.* Give the inspection cost per year at the 5 % interval and at the cheapest interval. (By equation (18.6) with $C_V = \text{R\$ } 3\,000$: at 5 months R\$ 7 200; at 4 months R\$ 9 000.)

Exercise 18.5. *Application.* With the cleaning at the October shutdown, find the best single test before it and the best pair. (By equation (18.15) over the windows 37 to 45: the best single test is window 42, with $P_{\text{U}} = 0.0152$; the best pair is windows 40 and 43, with 0.0018.)

Chapter 19

Condition-Based Maintenance Triggers

19.1 The dashboard's work-order trigger and the cost of an order

The two-unit hall's coils are three years old. Today is December of year 3, window 36, and the maintenance log's entries of this window plan the work on them. The forecast puts the cleaning at the October shutdown of year 4, window 46 (Oct Y4), ten windows ahead; the repair-or-replace decision replaces coil 1 in January of year 4, off the outage, and names a deferral to the shutdown as its alternative; and the inspection program reads both coils in a reference test in June of year 4, window 42. Until work is done the hall runs on its own monthly record, and what is left to settle is the reading of that record that should bring an order forward. The condition-monitoring dashboard has a rule built in for that: it raises a work order when either coil's health-index tile, the commissioning approach over the window's (Chapter 10), falls below a fixed fraction of its commissioned value. The planner set the fraction where the capacity study puts the limit, G^{*} over the sum of the commissioning test's conductances, 0.8164. On the logged record the rule first fired in window 9 (Sep Y1), on coil 1, and the tiles have stood below the fraction in 26 of the 36 windows. Set at the rated fraction, 0.7686, the rule would have fired in window 14 (Feb Y2); at 0.70, in window 16 (Apr Y2); at 0.60, in window 29 (May Y3). Today's tiles read 0.4873 and 0.5947, so the order stands at every look from today on. The planner's conclusion is that the dashboard has been asking for the cleaning for more than two years: raise the order now, away from the outage, at R$ 12 000 more than the shutdown's work would cost.

The chart cannot say whether a tile's fall is the coils at all: Chapter 10 split the tile's slope into the load, the season, unit 1's drifting sensor and the coils. Nor can it say how likely the hall is to fall below its capacity limit before the work could be done, which is a probability of the coils' state, not a reading. It cannot say at what probability an order pays, when an order raised for nothing costs R$ 15 000 with its premium and a month below the limit costs R$ 150 000. And it cannot say how often a rule read every month fires on a hall that did not need the work. The question of this chapter is

the level c such that raising the work order whenever the posterior probability that the hall is below its capacity limit at the end of the lead time reaches c minimizes the expected cost of the monthly looks to the October shutdown, read against the closed form for one look, the levels 0.5, 0.9 *and* 0.95*, and the dashboard's tile threshold, with the false-order and missed-degradation rates the level gives.*

It is answered with the method of Chapter 5, on Chapter 14's family of drawings carried forward window by window by the rows of Chapter 4, and priced by Chapter 2's expected consequence.

Jargon: condition-based maintenance trigger, look, lead time, false work order, missed degradation, one-look level, operating characteristic of a trigger

A *condition-based maintenance trigger* raises a work order when a probability of the asset's state reaches a decision level c (a decision level on the posterior). A *look* is a refit of the record at a window: the posterior updated by that window's rows, with nothing later. The *lead time* is the time between raising an order and the work being done, and the trigger's event is read at its end. A *false work order* is an order raised when the hall would have held to the planned outage without it. *Missed degradation* is a window below the capacity limit before the work is done. The *one-look level* is the probability at which, at a single look, the expected losses of ordering and of waiting are equal. The *operating characteristic of a trigger* is its false-order rate against its miss rate over records not yet logged, as the level runs: the share of records on which it orders for nothing against the share on which the hall goes short before the work.

Key idea

A trigger is a level on a probability of the state, and the costs set it. At one look the level is the cost of a false order over the sum of the costs of the two errors. When the record is refitted every month the level on a one-month probability falls below that, because each look reads only the next window while an order is judged on the whole time to the outage. A conventional confidence of 0.5, 0.9 or 0.95 waits until the hall is nearly short, and pays for the wait in months below the limit.

19.1.1 A level set by the costs

Chapter 2 wrote a choice between actions as the expected consequence (2.23): the loss of each action averaged over the posterior of the state, and the action with the smaller average chosen. A work-order trigger is that choice taken at a look, between two actions, to raise the order now or to wait. The loss of each depends on the state through one event $\mathcal{U}$, the hall below its limit at a window that an order raised now would protect. Write $q_{\mathcal{U}} = \Pr(\mathcal{U} \mid \mathbf{y})$, and write the four losses as $L(\text{order},\mathcal{U})$, $L(\text{order},\neg\mathcal{U})$, $L(\text{wait},\mathcal{U})$ and $L(\text{wait},\neg\mathcal{U})$. Each action is wrong in one state. An order is wrong when $\mathcal{U}$ does not occur, and it then costs C_{false} more than waiting would have; waiting is wrong when $\mathcal{U}$ occurs, and it then costs C_{miss} more than an order would have. The difference of the two expected losses is linear in $q_{\mathcal{U}}$, and it vanishes at one probability:

$$
\begin{aligned}
C_{\text{false}} &= L(\text{order},\neg\mathcal{U}) - L(\text{wait},\neg\mathcal{U}), \qquad C_{\text{miss}} = L(\text{wait},\mathcal{U}) - L(\text{order},\mathcal{U}),\\
\mathcal{R}(\text{order}) - \mathcal{R}(\text{wait}) &= (1 - q_{\mathcal{U}})\, C_{\text{false}} - q_{\mathcal{U}}\, C_{\text{miss}},\\
\mathcal{R}(\text{order}) \le \mathcal{R}(\text{wait}) &\iff q_{\mathcal{U}} \ge c_1^* = \frac{C_{\text{false}}}{C_{\text{false}} + C_{\text{miss}}}.
\end{aligned}
\tag{19.1}
$$

The level needs no fit of its own: the posterior gives $q_{\mathcal{U}}$, and the costs give the level. A conventional level is a statement about the costs made without reading them. A level of 0.5 says that the two errors cost the same; a level of 0.95 says that a false order costs nineteen times a missed month. Two things move the level, and nothing else does. The costs move it. And the looks move it: when the probability a look reads covers only the next window, while the order must come before every later window that could go short, the level at which an order pays on that probability is no longer equation (19.1).

Principle 19.1. *Set a condition-based trigger by the costs of its two errors, read on a probability of the parameter, never on a raw-performance tile and never at a conventional confidence. Set it again when the looks are repeated, because the probability a look reads covers less time than the order must protect.*

19.1.2 The two-unit hall and its looks as a graph (step 0)

The plant graph is Chapter 14's, cited and not re-derived: Chapter 10's stacked record of 37 windows (Figure 10.1), with Chapter 11's curvature node on coil 1's trend rows and Chapter 12's step node at window 15 in the drawings that free them, carried past the record by forecast windows and their trend rows, with the requirement node $G_{\mathrm{tot},t} \geq G^*$ on each forecast window (Figure 14.1). This chapter adds no row to the record. It turns the forecast windows 37 to 45 into windows still to be read, each of which will bring its five readings and its load, and each of which is a look. At each look t_{as} it hangs a question node on window $t_{\mathrm{as}} + \Delta_{\mathrm{L}}$, the end of the lead time Δ_{L}: the probability that the hall is below its limit there,

$$q_{t_{\mathrm{as}}} = \Pr(G_{\mathrm{tot},t_{\mathrm{as}}+\Delta_{\mathrm{L}}} \leq G^* \mid \mathbf{y}_{0:t_{\mathrm{as}}}), \qquad t_{\mathrm{as}} = 36, \ldots, 45, \tag{19.2}$$

with $\mathbf{y}_{0:t_{\mathrm{as}}}$ the readings of windows 0 to t_{as} and the prior rows, and $G_{\mathrm{tot},t}$ the coils' total conductance (14.5). Beside each question node it hangs one decision node, *order*, fed by the level c and the costs.

The stated inputs. The record is Chapter 14's, windows 0 to 36 (Table 10.1). Four stated inputs enter here, and none is a finding. *The looks* are a refit every month from today, windows 36 (Dec Y3) to 45 (Sep Y4), ten looks; each adds its window's five readings, its load entry and its two trend rows, and nothing later. *The lead time* is $\Delta_{\mathrm{L}} = 1$ window: an order raised at the look of window t_{as} has both coils cleaned before window $t_{\mathrm{as}} + 1$ is read. *The October shutdown*, window 46 (Oct Y4), does the work in the outage if no order has fired, and the hall must hold through window 46, Chapter 14's convention (Section 14.3.2). *The costs* are Part V's schedule, each a stated input: a false work order, FO = R$ 3 000; the premium for work away from the outage, OW = R$ 12 000; a month with the hall below its limit, SF = R$ 150 000; cleaning, R$ 8 000 a coil; and an inspection, one crew visit with reference instruments, R$ 1 500, which step 7 prices; Chapter 18's test of both coils' conductance is two such visits, R$ 3 000. An order raised before the shutdown pays OW; it is a false order, and pays FO as well, when the hall would have held through window 46 without it; and every window at or below G^* before the work is done pays SF. The cleaning is paid once on every branch, since the coils are cleaned once by window 46 whatever the rule does, so it cancels from every comparison, and the share of the loss a cleaning restores does not enter for the same reason. For the coil clause of step 5 two per-coil limits are stated. The first is the capacity study's share, $g^* = \ln(G^*/(\bar{G}_1 + \bar{G}_2)) = -0.26313$ on each coil's log-derate, a derate factor of 0.7686: the requirement level on g of the notation table, applied to g_u with the two coils sharing the limit equally. The second is Part I's functional limit (Chapter 2, Section 2.3.4), $g_F = \ln(5/6) = -0.18232$, the hall of Part I at 5 kW/K on its 6 kW/K design, read here against each coil's rating.

The maintenance log so far. The two-unit hall's log holds the commissioning test of window 0; Chapter 5's and Chapter 6's entries at window 18 (Jun Y2); Chapter 7's at window 25 (Jan Y3), whose side-by-side check of unit 1's discharge sensor is still open; and the entries of Chapters 10 to 13 at window 36. Three more entries precede this one at window 36, in this order. Chapter 14's forecast put the crossing in May of year 5 at the median, with probability 0.16 by the October shutdown, and wrote: plan the coil cleaning for the October Y4 shutdown, and refit at window 45 before the work order is written. Chapter 17's decision replaces coil 1 in January of year 4, off the outage, and names as its alternative a deferral: the replacement held for the October shutdown and decided by the refit at window 45. Chapter 18's program reads both coils in a reference condition test in June of year 4, window 42, and raises an order if that test reads $G_1 + G_2 \leq 11.261$ kW/K. This chapter's entry is at window 36 as well. It values its trigger for the hall with no work done before the shutdown, the plan Chapter 17's step 5 prefers when coil 1's rate is constant (it cleans at the shutdown), and the calendar's; and it sets the

rule by which the monthly refits, Chapter 14's at window 45 among them, turn into a work order. Read on the hall's own sensors every month, the rule takes the place of Chapter 18's test-based order level for orders between tests; the June test stays as a reference reading.

Modeling choices at the graph. Four choices that no question block owns make the looks. *The looks read Chapter 14's family unchanged*: each window's rows are added to each drawing, and each drawing's weight is moved by its evidence for them. This buys a trigger on the same posterior the forecast read, and it fails where that forecast fails; step 5 prices the forecast's calibration finding (Section 14.4.5). *Each look is the drawing's Laplace normal updated by the new window's rows, linearized on the drawing's MAP path*, Chapter 4's preposterior applied one window at a time. This buys ten looks on each of many predicted records without a Gauss–Newton refit of 18 drawings at each, and it fails where the rows bend between the MAP path and a record's own path; step 5 refits records. *No maintenance but the order, and one order cleans both coils*: this keeps one decision per look, and it fails for a plan that works on one coil, such as the replacement of coil 1 that Chapter 17 weighs. *The rule is set today*, at window 36, and applied unchanged at every later look: this makes the rule's cost a quantity of today's posterior, and it fails if the costs or the lead time change before October.

The graph. The trend layer of both coils from window 36 is drawn in Figure 19.1. The upper chain is coil 1's path and the lower chain coil 2's, with the trend rows as black squares between neighbors. Window 36 is read, and its faint block holds its load and readings. Windows 37 to 45 are still to be read: each carries a dashed reading block on the left, its load and five readings, and on the right the red question node of the look before it, linked by dotted lines to the two coils of its window. Below each question node a purple decision node, *order*, is fed along one line by the level and the costs. Window 46 carries the last question node, q_{45}, and the shutdown; window 47 is a forecast window with nothing attached. The run of windows 39 to 45 is one dashed column.

19.1.3 The asset manager's questions (step 1)

The asset manager asks four questions of the trigger, and step 1 writes each as an expectation or a probability over the state and the readings not yet logged, with the unknown on the left of the bar and the record to window 36 on the right. Table 19.1 lists them, with the columns of Chapter 5's table: the question in the asset manager's words, the drill question first; the general form of its category from Chapter 2; its mathematical form for the hall; its category; the computation of Section 2.3.6 that answers it; and the subsection where it is answered. T_{ord} is the first look at which the order fires, $L(c, \mathbf{z}, \tilde{\mathbf{y}}_{37:45})$ the loss of the rule at level c on the path $\mathbf{z}$ and the readings $\tilde{\mathbf{y}}_{37:45}$ of windows not yet read, and EVPI Chapter 17's expected value of perfect information (17.9). Not asked here: whether to clean or replace at the shutdown, or whether to move the work to another outage (Chapter 17); how often to inspect (Chapter 18); when the hall crosses (Chapter 14); or whether the trend log can see a given loss (Chapter 22). Which measurement would lower the trigger's expected cost asks nothing new of today's evidence; it is step 7.

19.1.4 The evidence as rows (step 2)

Step 2 attaches nothing new to the record. The state and the rows of each drawing are Chapter 14's: the stacked state (10.3) with coil 1's curvature and the step at window 15 where the drawing frees them, 115 to 117 entries under 261 to 263 rows, and the family's prior (14.2) (Table 14.2). What step 2 adds are the rows of the windows still to be read, which each look will attach, and the loss that prices a rule.

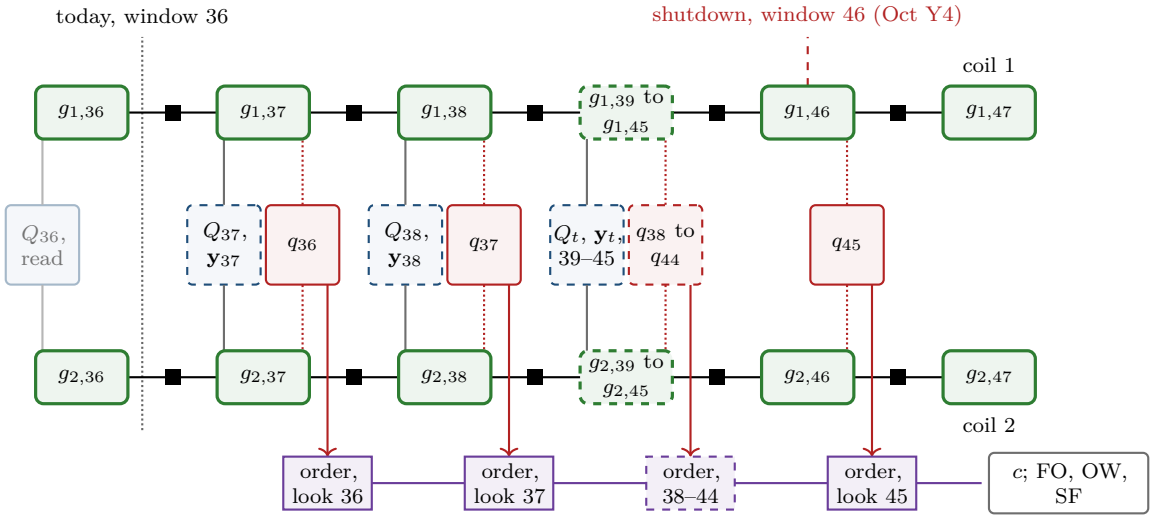


Figure 19.1: The trend layer of the two coils with the looks to the October shutdown, from window 36. Free (green): the log-derates $g_{u,t}$ of each coil, read at window 36 and still to be read at windows 37 to 45 (the run 39 to 45 as dashed boxes), forecast at 46 and 47. Black squares: the trend rows (14.4), fed by the trend entries γ_1, γ_2 and, in the drawings that free it, $\dot{\gamma}_1$, as in Figure 14.1 (not drawn). Faint block: window 36's load and readings, observed. Dashed blocks: the load and the five readings of windows 37 to 45, observed at their own looks and not yet read. Red, dotted: the question node $q_{t_{\text{as}}}$ of equation (19.2) on window $t_{\text{as}} + 1$, reading that window's two coils. Purple: the decision node of each look, fed by the level c and the stated costs FO, OW and SF (white). The known inputs, the pinned limit $G^* = 10.761\,\text{kW/K}$ and the read windows' physical rows are as in Figure 14.1. The decisions are coupled only through the first look whose order fires: after it, the coils are cleaned and the later nodes do not act.

Table 19.1: The questions asked of the trigger at window 36 (Dec Y3). $\mathbf{y}_{0:36}$ is the record to window 36 and the prior rows; $\tilde{\mathbf{y}}_{37:45}$ the readings of the looks to come; $q_{t_{\text{as}}}$ the trigger's probability (19.2); T_{ord} the first look with $q_{t_{\text{as}}} \geq c$; c' any other rule.

Asset manager's question	General form	Mathematical form	Category	Computation	Answered in
At what measured value should the work order fire?	(2.23), (2.24); (2.12)	at one look, $\mathcal{R}(\text{order}) = \text{OW} + \text{FO}\,(1 - q_{45})$, $\mathcal{R}(\text{wait}) = \text{SF}\,q_{45}$; order when $q_{45} \geq c_1^*$	design	expected consequence	19.3.1
With a refit every month, which level costs least to the shutdown, and what do 0.5, 0.9 and 0.95 cost?	(2.16) with the rule's loss; (2.17)	$c^* = \arg\min_c \mathcal{R}(c)$, $\mathcal{R}(c) = \mathbb{E}[L(c, \mathbf{z}, \tilde{\mathbf{y}}_{37:45}) \mid \mathbf{y}_{0:36}]$; $\mathcal{R}(c)$ at c^*, 0.5, 0.9, 0.95; $\Pr(L(c^*) < L(c'))$; EVPI	design	expected consequence; preposterior	19.3.2
When would this record fire?	(2.16), its first look	$\Pr(T_{\text{ord}} \leq s \mid \mathbf{y}_{0:36})$ over the predicted records; $q_{t_{\text{as}}}$ on the as-of fits of the logged record	prediction under a threshold	preposterior	19.3.3
How often does the rule fire on a coil that is not degrading?	(2.16) at a truth with no loss	$\Pr(\exists\, t_{\text{as}} \leq 45 : q_{t_{\text{as}}} \geq c^* \mid \text{no loss})$, from today and on a restarted record; the dashboard's beside it	design	preposterior	19.3.4

The entries a later window touches. A window t after 36 brings five reading rows, on the return air, the two discharge temperatures, the hall's load meter and rack 1's inlet, the rows of Chapter 5's window (5.7) at the window's known inputs, with unit 1's discharge carrying the drift (10.2); a load entry Q_t; and the two trend rows from window $t-1$, equation (14.4). By the forward equations (5.3) and the trend rows, these rows touch eight entries of the drawing's state and the window's own load,

$$\boldsymbol{\theta}_t = \left(g_{1,t},\ g_{2,t},\ \gamma_1,\ \gamma_2,\ \dot{\gamma}_1,\ \beta_1,\ \dot{b}_{d,1},\ \delta\right)^{\mathsf{T}}, \qquad Q_t, \tag{19.3}$$

with $\dot{\gamma}_1$ and δ present only in the drawings that free them. The step δ enters a later window's rows only under a sensor step, where the stepped reading carries it in every window from 15 on; under a coil step it touches no later row. No other entry of the stacked state is read again. A look therefore carries each drawing's Laplace marginal of $\boldsymbol{\theta}_{36}$ forward, the block of $\boldsymbol{\Sigma}_M$ on the eight entries, and the rest of the state stays where window 36 left it.

A look's rows. The load of a later window has no reading but the meter's. Its prior row is centered on the record's growth, 120 kW from window 1 rising 4 %/yr, with a standard deviation of 100 kW, flat in effect beside the meter's 3.0 kW. On the leading drawing, $(\mathcal{C}_1, 2)$, coil 1 curving with coil 2 stepped, the rows of window 37 (Jan Y4) are linearized on its MAP path, where $g_{1,37} = -0.17411$, $g_{2,37} = -0.19234$ and the load is 134.984 kW. The slopes of the five readings on $(\boldsymbol{\theta}_{37}, Q_{37})$, in K or kW per unit entry, with the rows in the order return air, unit 1's

discharge, unit 2's discharge, load meter and rack-1 inlet, are

$$\frac{\partial \hat{\mathbf{y}}_{37}}{\partial(\boldsymbol{\theta}_{37}, Q_{37})} = \begin{bmatrix} -6.3145 & -6.2004 & 0 & 0 & 0 & 0 & 0 & 0 & +0.0858 \\ -9.9613 & -2.5537 & 0 & 0 & 0 & 0 & +1.5000 & 0 & +0.0353 \\ -2.6678 & -9.8471 & 0 & 0 & 0 & 0 & 0 & 0 & +0.0362 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ -6.3145 & -6.2004 & 0 & 0 & 0 & +7.2941 & 0 & 0 & +0.0463 \end{bmatrix}, \tag{19.4}$$

and the Jacobian of the normalized rows divides each line by its sensor's standard deviation,

$$\mathbf{J}_{37} = \mathbf{W}^{1/2} \frac{\partial \hat{\mathbf{y}}_{37}}{\partial(\boldsymbol{\theta}_{37}, Q_{37})}, \qquad \mathbf{W} = \operatorname{diag}(0.3^{-2},\ 0.3^{-2},\ 0.3^{-2},\ 3.0^{-2},\ 0.3^{-2}),$$

with $\mathbf{W}$ the readings' weights: with normalized rows the information is $\mathbf{J}^{\mathsf{T}}\mathbf{J}$, with unnormalized rows and their weights $\mathbf{J}^{\mathsf{T}}\mathbf{W}\mathbf{J}$. Unit 1's discharge has slope $+1.5000$ on the drift rate: the drift's lever at window 37, eighteen windows after its onset, is 1.5 yr. The trend rows from window 36 to 37 have the lines of equation (14.6), $(-200,\ +200,\ -16.6667,\ -25.6944)$ on $(g_{1,36}, g_{1,37}, \gamma_1, \dot{\gamma}_1)$ for coil 1 and $(-200,\ +200,\ -16.6667)$ on $(g_{2,36}, g_{2,37}, \gamma_2)$ for coil 2. Under a sensor-step drawing the stepped reading's line has $1/\sigma_k$ in the column of δ, the step's column of Chapter 12.

Which row acts on which entry. With $*$ for an entry that can be nonzero, $\circ$ for the step's entry on the stepped reading under a sensor-step drawing, and the columns $(g_{1,t-1}, g_{2,t-1}, g_{1,t}, g_{2,t}, \gamma_1, \gamma_2, \dot{\gamma}_1, \beta_1, \dot{b}_{d,1}, \delta, Q_t)$, the rows a look attaches at window t have the pattern

$$\begin{bmatrix} * & 0 & * & 0 & * & 0 & * & 0 & 0 & 0 & 0 \\ 0 & * & 0 & * & 0 & * & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & * & * & 0 & 0 & 0 & 0 & 0 & \circ & * \\ 0 & 0 & * & * & 0 & 0 & 0 & 0 & * & \circ & * \\ 0 & 0 & * & * & 0 & 0 & 0 & 0 & 0 & \circ & * \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \circ & * \\ 0 & 0 & * & * & 0 & 0 & 0 & * & 0 & \circ & * \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & * \end{bmatrix} \begin{matrix} \text{trend row, coil 1} \\ \text{trend row, coil 2} \\ \text{return air } y_{T_r} \\ \text{discharge } y_{T_{d,1}} \\ \text{discharge } y_{T_{d,2}} \\ \text{load meter } y_Q \\ \text{rack-1 inlet } y_{T_1} \\ \text{load prior on } Q_t \end{matrix}, \tag{19.5}$$

$$\mathbf{g}_{G,t}^{\mathsf{T}} = (0,\ 0,\ \bar{G}_1 \mathrm{e}^{g_{1,t}},\ \bar{G}_2 \mathrm{e}^{g_{2,t}},\ 0,\ 0,\ 0,\ 0,\ 0,\ 0,\ 0).$$

The second line is not a row. It is the gradient of the quantity q_{t-1} reads, $G_{\mathrm{tot},t}$, on the same columns: the question reads $(g_{1,t}, g_{2,t})$ and nothing else, with $\partial G_{\mathrm{tot}}/\partial g_u = \bar{G}_u \mathrm{e}^{g_u}$, 5.8815 and 5.7752 kW/K on the leading drawing's MAP path at window 37. Two facts of the graph are read here, before any number. The readings see the coils only through $g_{1,t}$ and $g_{2,t}$ of their own window, and the earlier window's coils only through the trend rows, so a look needs only the eight entries of equation (19.3), carried from one window to the next. And the question shares its two columns with every reading, so each window's readings move the next month's probability directly.

The loss of a rule. A rule at level c orders at the first look whose probability reaches it,

$$T_{\mathrm{ord}}(c) = \min\{t_{\mathrm{as}} \in \{36, \dots, 45\} : q_{t_{\mathrm{as}}} \geq c\}, \qquad T_{\mathrm{ord}} = 46 \text{ if no look reaches } c,$$

$$\begin{aligned} L(c, \mathbf{z}, \tilde{\mathbf{y}}_{37:45}) = \mathrm{SF} \sum_{t=37}^{T_{\mathrm{ord}}} \mathbf{1}\big[G_{\mathrm{tot},t} \leq G^*\big] + \mathrm{OW}\, \mathbf{1}[T_{\mathrm{ord}} \leq 45] \\ + \mathrm{FO}\, \mathbf{1}[T_{\mathrm{ord}} \leq 45] \prod_{t=37}^{46} \mathbf{1}\big[G_{\mathrm{tot},t} > G^*\big], \end{aligned} \tag{19.6}$$

with the cleaning cancelled. An order at look t_{as} protects every window from $t_{\mathrm{as}}+1$ on, so the windows before the work are 37 to T_{ord}, and with no order they run to window 46, the shutdown's own. The rule is a function of the readings through the $q_{t_{\mathrm{as}}}$, and the loss is a function of the path through the indicators: the loss of the rule depends on both the state and the readings not yet logged.

Modeling choices at the rows. Four choices make these rows. *The trigger's quantity is the hall below its limit at the end of the lead time*, equation (19.2): it is what an order at look t_{as} prevents and waiting cannot. It is the third line of the trend form (2.21) at window $t_{\mathrm{as}}+1$, and it differs from the first passage by that window only through paths that dipped below the limit and came back, which hold 0.0001 of the posterior by window 46 (Section 14.3.3). It fails if a crossing that recovers still costs, which no cost here states. *The coil clause is left out of the reported quantity.* With the capacity study's share, $g_1 > g^*$ and $g_2 > g^*$ give $G_{\mathrm{tot}} > G^*$, so "either coil below its share, or the hall below its limit" is "either coil below its share". That clause fires when coil 1 alone passes its share while coil 2 still holds the hall, an event no cost prices; step 5 sets it beside the hall's quantity. *The lead time is one window*: it fixes which window an order protects, and step 5 gives two. *The load of a later window enters with the record's growth as its center and a* 100 kW *spread*: every drawing shares it, so it cancels from their weights, and it fails only if the meter fails.

19.2 The posterior of the state (step 3)

Every answer of this chapter is read from the posteriors of Chapter 14's 18 drawings at window 36, carried to each later look by the rows of step 2, and from the drawings' evidences, which weigh them.

19.2.1 The family's posterior at window 36

The derivation form is Chapter 5's (Sections 5.2.1 and 5.2.2), on each drawing's stacked state, as Chapter 14 fitted it (Section 14.2.1) and not derived again: the MAP state $\hat{\mathbf{z}}_M$ by Gauss–Newton on the normal equations (5.9), the Laplace covariance $\mathbf{\Sigma}_M = (\mathbf{J}_M^{\mathsf{T}}\mathbf{J}_M)^{-1}$ at $\hat{\mathbf{z}}_M$, the Laplace evidence (7.7), and the drawing's weight (14.8). The refit reproduces Chapter 14's family to four decimals: coil 1's curvature with $\Pr(\mathcal{C}_1 \mid \mathbf{y}) = 0.6344$, a step at window 15 with 0.7864, and a step on coil 2 with 0.5489. The six leading drawings are $(\mathcal{C}_1, 2)$ at 0.3428, $(\mathcal{C}_0, 2)$ at 0.2061, $(\mathcal{C}_1, T_{d,2})$ at 0.1606, $(\mathcal{C}_1, \text{none})$ at 0.1097, $(\mathcal{C}_0, \text{none})$ at 0.1038 and $(\mathcal{C}_0, T_{d,2})$ at 0.0385, and the other twelve hold 0.0385 together.

The block the looks carry. The posterior of the eight entries (19.3) at window 36 is given in Table 19.2 for the leading drawing: each row is an entry, with its MAP value, its spread $\sqrt{\Sigma_{ii}}$ and its correlations with the other seven. Its covariance, each entry Σ_{ij} in the product of the units of entries i and j (the two levels, β_1 and δ dimensionless, the rates in yr^{-1}, the curvature in yr^{-2}, the drift rate in K/yr),

$$\mathbf{\Sigma}_{\boldsymbol{\theta},36} = 10^{-4} \times \begin{bmatrix} 8.8145 & -3.2768 & 2.7847 & -1.6225 & 3.4891 & 0.90048 & 44.453 & 0.68269 \\ -3.2768 & 3.2228 & -1.1289 & 1.7298 & -1.0536 & -0.19156 & -17.191 & -1.1278 \\ 2.7847 & -1.1289 & 2.2372 & -0.63295 & 0.54692 & 0.25493 & 14.664 & 0.17733 \\ -1.6225 & 1.7298 & -0.63295 & 2.8430 & -0.50099 & -0.14925 & -8.2205 & -2.4593 \\ 3.4891 & -1.0536 & 0.54692 & -0.50099 & 4.0703 & 0.24176 & 13.695 & 0.44755 \\ 0.90048 & -0.19156 & 0.25493 & -0.14925 & 0.24176 & 0.81127 & 5.8064 & 0.039035 \\ 44.453 & -17.191 & 14.664 & -8.2205 & 13.695 & 5.8064 & 326.97 & 2.4519 \\ 0.68269 & -1.1278 & 0.17733 & -2.4593 & 0.44755 & 0.039035 & 2.4519 & 5.5723 \end{bmatrix}, \tag{19.7}$$

is the block of $\mathbf{\Sigma}_M$ on those entries, in the order of equation (19.3).

Table 19.2: The posterior of the eight entries a look carries, at window 36 on the leading drawing $(\mathcal{C}_1, 2)$ (weight 0.3428): MAP value, spread and correlations.

Entry [unit]	MAP [entry's]	$\sqrt{\Sigma_{ii}}$ [entry's]	$g_{1,36}$ [–]	$g_{2,36}$ [–]	γ_1 [–]	γ_2 [–]	$\dot{\gamma}_1$ [–]	β_1 [–]	$\dot{b}_{d,1}$ [–]	δ [–]
$g_{1,36}$ [–]	−0.16571	0.02969	1	−0.6148	+0.6271	−0.3241	+0.5825	+0.3367	+0.8280	+0.0974
$g_{2,36}$ [–]	−0.19007	0.01795	−0.6148	1	−0.4204	+0.5715	−0.2909	−0.1185	−0.5296	−0.2661
γ_1 [yr^{-1}]	−0.04825	0.01496	+0.6271	−0.4204	1	−0.2510	+0.1812	+0.1892	+0.5422	+0.0502
γ_2 [yr^{-1}]	−0.02732	0.01686	−0.3241	+0.5715	−0.2510	1	−0.1473	−0.0983	−0.2696	−0.6179
$\dot{\gamma}_1$ [yr^{-2}]	−0.03403	0.02018	+0.5825	−0.2909	+0.1812	−0.1473	1	+0.1330	+0.3754	+0.0940
β_1 [–]	+0.21049	0.00901	+0.3367	−0.1185	+0.1892	−0.0983	+0.1330	1	+0.3565	+0.0184
$\dot{b}_{d,1}$ [K/yr]	+0.55732	0.18082	+0.8280	−0.5296	+0.5422	−0.2696	+0.3754	+0.3565	1	+0.0574
δ [–]	−0.07511	0.02361	+0.0974	−0.2661	+0.0502	−0.6179	+0.0940	+0.0184	+0.0574	1

The correlations read by sign, and the rows fix each sign. The largest is coil 1's level against the drift rate, +0.8280: a faster drift explains more of unit 1's late rise in discharge temperature, so the same readings leave coil 1 standing higher today. It is Chapter 7's ridge, read on the level, and it is why a check of unit 1's discharge sensor is a candidate in step 7. The two coils' levels are correlated by −0.6148: the return air reads their sum, so a higher coil 1 leaves less to coil 2, and the sum, which is what $q_{t_{\mathrm{as}}}$ reads, is known better than the split. Each coil's level is correlated positively with its own rate, +0.6271 for coil 1 and +0.5715 for coil 2, and coil 1's with its curvature, +0.5825: with the earlier windows held, a higher level today comes with a slower fall to it, so an error in the level and an error in the rate add in the hall's next month rather than cancel. Coil 2's rate is correlated with the step by −0.6179: a larger March step leaves less of coil 2's loss for its rate to carry, Chapter 12's split.

19.2.2 A look as rows added to the posterior

The quantity each look owes is $q_{t_{\mathrm{as}}}$ of equation (19.2). It is a mixture over the drawings, Chapter 7's form (7.10) with each drawing's probability after the readings to window t_{as}:

$$q_{t_{\mathrm{as}}} = \sum_{M} P_M(t_{\mathrm{as}})\ \Pr\!\left(G_{\mathrm{tot},t_{\mathrm{as}}+1} \le G^* \mid \mathbf{y}_{0:t_{\mathrm{as}}}, M\right). \tag{19.8}$$

Each term needs three things: the drawing's posterior of the carried entries after window t_{as}'s rows, its weight after them, and the region $G_{\mathrm{tot},t_{\mathrm{as}}+1} \le G^*$ under the drawing's prediction one window ahead. Each is derived in turn, from the posterior at window 36 of Section 19.2.1.

The drawing's posterior after a window. Write $\bar{\boldsymbol{\theta}}_t$ and $\mathbf{\Sigma}_t$ for the drawing's mean and covariance of $\boldsymbol{\theta}_t$ after the rows of windows 0 to t, with $\bar{\boldsymbol{\theta}}_{36} = \hat{\boldsymbol{\theta}}_M$, the drawing's MAP values of the eight entries, and $\mathbf{\Sigma}_{36} = \mathbf{\Sigma}_{\boldsymbol{\theta},36}$, its block (19.7). The trend rows from $t-1$ to t are linear and Gaussian in the path, so before window t's readings the carried entries are normal with the trend applied and one month's wander added to each coil,

$$\bar{\boldsymbol{\theta}}_t^{-} = \mathbf{T}_{t-1}\bar{\boldsymbol{\theta}}_{t-1}, \qquad \mathbf{\Sigma}_t^{-} = \mathbf{T}_{t-1}\mathbf{\Sigma}_{t-1}\mathbf{T}_{t-1}^{\mathsf{T}} + \sigma_q^2\ \mathrm{diag}(1,\ 1,\ 0,\ \ldots,\ 0), \tag{19.9}$$

with $\mathbf{T}_{t-1}$, the trend rows' transition, the identity except for the trend rows' terms, Δt in the row of $g_{1,t}$ and the column of γ_1, $\Delta t\, x_{t-1}$ in that row and the column of $\dot{\gamma}_1$, and Δt in the row of $g_{2,t}$ and the column of γ_2, by equation (14.7). The window's load is appended with its prior row of step 2. The readings of window t are then $n_r = 5$ normalized rows $\mathbf{r}_t$ on $(\boldsymbol{\theta}_t, Q_t)$, linearized on the drawing's MAP path, with Jacobian $\mathbf{J}_t$ as in equation (19.4). Adding them to

the normal equations of the record is the preposterior of Chapter 4 (4.65) with the readings' values now known, and the correction form (4.66) with five rows in place of one:

$$\begin{aligned}
\mathbf{S}_t &= \mathbf{J}_t\,\boldsymbol{\Sigma}_t^{-}\,\mathbf{J}_t^{\mathsf{T}} + \mathbf{I},\\
\boldsymbol{\Upsilon}_t &= \boldsymbol{\Sigma}_t^{-}\,\mathbf{J}_t^{\mathsf{T}}\,\mathbf{S}_t^{-1},\\
\bar{\boldsymbol{\theta}}_t &= \bar{\boldsymbol{\theta}}_t^{-} - \boldsymbol{\Upsilon}_t\,\mathbf{r}_t(\bar{\boldsymbol{\theta}}_t^{-}),\\
\boldsymbol{\Sigma}_t &= \left((\boldsymbol{\Sigma}_t^{-})^{-1} + \mathbf{J}_t^{\mathsf{T}}\mathbf{J}_t\right)^{-1} = \boldsymbol{\Sigma}_t^{-} - \boldsymbol{\Upsilon}_t\,\mathbf{S}_t\,\boldsymbol{\Upsilon}_t^{\mathsf{T}},
\end{aligned} \tag{19.10}$$

with $\boldsymbol{\Sigma}_t^{-}$ here the covariance with the load appended, $\mathbf{r}_t(\bar{\boldsymbol{\theta}}_t^{-})$ the five rows at the prediction, the model's readings less the readings over each sensor's standard deviation, and $\boldsymbol{\Upsilon}_t$ the gain of the correction form. The load is then dropped: no later row reads it.

The drawing's weight after a window. The weight of a drawing is its evidence times its prior, normalized, equation (14.8). Adding window t's rows multiplies each drawing's evidence by the predictive density of those rows given the readings before them, Chapter 7's probability (7.9) with the evidence ratio of the new rows in place of the whole record's:

$$P_M(t) = \frac{P_M(t-1)\,\mathcal{N}(\mathbf{r}_t(\bar{\boldsymbol{\theta}}_t^{-});\ \mathbf{0},\ \mathbf{S}_t)}{\sum_{M'} P_{M'}(t-1)\,\mathcal{N}(\mathbf{r}_t(\bar{\boldsymbol{\theta}}_t^{-});\ \mathbf{0},\ \mathbf{S}_t)_{M'}}, \qquad P_M(36) = P_M, \tag{19.11}$$

with each drawing's own prediction and $\mathbf{S}_t$ in its own term. The sensors' standard deviations, which normalize the rows, are common to every drawing and cancel. A drawing whose prediction of the window's readings is poor loses weight; under a sensor-step drawing the stepped reading's prediction carries δ.

The region one window ahead. The question reads $(g_{1,t_{\mathrm{as}}+1}, g_{2,t_{\mathrm{as}}+1})$ alone, by the gradient of equation (19.5). One more trend row by equation (19.9) makes the pair normal. Its mean $\mathbf{m}$ is the first two entries of $\mathbf{T}_{t_{\mathrm{as}}}\bar{\boldsymbol{\theta}}_{t_{\mathrm{as}}}$, and its 2×2 covariance $\mathbf{V}$ is the leading block of $\mathbf{T}_{t_{\mathrm{as}}}\boldsymbol{\Sigma}_{t_{\mathrm{as}}}\mathbf{T}_{t_{\mathrm{as}}}^{\mathsf{T}}$ with σ_q^2 added to each diagonal entry. The region $\bar{G}_1\mathrm{e}^{g_1} + \bar{G}_2\mathrm{e}^{g_2} \le G^*$ is curved, so its probability is not a normal tail of a linear quantity; it is the integral over g_2 of the tail of g_1 given g_2,

$$\Pr(G_{\mathrm{tot},t_{\mathrm{as}}+1} \le G^* \mid \mathbf{y}_{0:t_{\mathrm{as}}}, M) = \int \Phi\left(\frac{\ln(G^*/\bar{G}_1 - \mathrm{e}^{g_2}) - m_{1|2}(g_2)}{s_{1|2}}\right)\mathcal{N}(g_2;\ m_2,\ V_{22})\,\mathrm{d}g_2, \tag{19.12}$$

with $\bar{G}_1 = \bar{G}_2$, the integrand set to one where $\mathrm{e}^{g_2} \ge G^*/\bar{G}_2$, $m_{1|2}(g_2) = m_1 + (V_{12}/V_{22})(g_2 - m_2)$ and $s_{1|2}^2 = V_{11} - V_{12}^2/V_{22}$, Chapter 3's conditional normal. The integral is taken by 48-node Gauss–Hermite quadrature in g_2.

What moves with the record. The rows are linearized on the drawing's MAP path, not on the record, so $\mathbf{J}_t$, $\mathbf{S}_t$, $\boldsymbol{\Upsilon}_t$, $\boldsymbol{\Sigma}_t$ and $\mathbf{V}$ do not depend on the readings: they are computed once per drawing and window. Only the means $\bar{\boldsymbol{\theta}}_t$ and the weights $P_M(t)$ move with the readings, through the rows' values at the prediction. A look on any record is therefore equations (19.10) and (19.11) applied to its readings, and equation (19.8) with (19.12) read at the result.

19.2.3 The next month at each look

Table 19.3 gives, for the leading drawing, the spread of the next window's two coils and of the hall's total at each look, from the covariances of equation (19.10) carried one window ahead. Each row is a look t_{as}; the columns give the spreads of $g_{1,t_{\mathrm{as}}+1}$ and $g_{2,t_{\mathrm{as}}+1}$, their correlation, and the spread of $G_{\mathrm{tot},t_{\mathrm{as}}+1}$ by the gradient (19.5). The spread of the next month narrows only

a little over nine more windows, from 0.1555 to 0.1306 kW/K: each look adds five readings and one month of wander. What moves $q_{t_{\mathrm{as}}}$ from look to look is the mean. On the leading drawing's MAP path the total falls from 11.657 kW/K at window 37 to 11.593, 11.528, 11.463, 11.397, 11.330, 11.262, 11.193, 11.124 and 11.054 kW/K at window 46, against the limit of 10.761.

Table 19.3: The posterior of the next window's coils at each look, the leading drawing $(\mathcal{C}_1, 2)$. The spreads do not depend on the readings.

Look t_{as} [–]	Date	$\mathrm{sd}(g_{1,t_{\mathrm{as}}+1})$ [–]	$\mathrm{sd}(g_{2,t_{\mathrm{as}}+1})$ [–]	correlation [–]	$\mathrm{sd}(G_{\mathrm{tot},t_{\mathrm{as}}+1})$ [kW/K]
36	Dec Y3	0.03243	0.01944	−0.5789	0.1555
37	Jan Y4	0.03192	0.01918	−0.5819	0.1514
38	Feb Y4	0.03144	0.01894	−0.5831	0.1477
39	Mar Y4	0.03104	0.01872	−0.5830	0.1445
40	Apr Y4	0.03072	0.01854	−0.5820	0.1418
41	May Y4	0.03046	0.01839	−0.5805	0.1395
42	Jun Y4	0.03024	0.01825	−0.5788	0.1374
43	Jul Y4	0.03003	0.01812	−0.5771	0.1353
44	Aug Y4	0.02979	0.01797	−0.5755	0.1331
45	Sep Y4	0.02949	0.01781	−0.5741	0.1306

Today's look. At window 36 no row is added, and equations (19.8) and (19.12) give

$$q_{36} = \Pr\big(G_{\mathrm{tot},37} \le G^* \mid \mathbf{y}_{0:36}\big) = 2.222 \times 10^{-8},$$

2.623×10^{-10} on the leading drawing and 8.273×10^{-6} on the drawing that gives the largest value. Today's posterior puts the hall's next month far above its limit.

19.2.4 The looks over predicted records

The answers of step 4 ask how a rule fares over the readings still to come, an expectation over records not yet logged. The design form (2.16) takes that expectation under the predictive of the readings, and it is evaluated here by predicted records drawn from the posterior of Section 19.2.1:

1. draw a drawing by its weight P_M, and its eight entries $\boldsymbol{\theta}_{36}$ from its Laplace normal, through the Cholesky factor of $\boldsymbol{\Sigma}_{\boldsymbol{\theta},36}$ (Section 4.4.7);
2. carry the path to window 47 by the drawing's trend rows, with an independent $\mathcal{N}(0,\ \sigma_q^2)$ wander for each coil and window, $\sigma_q = 0.005$;
3. draw the loads of windows 37 to 45 at the record's growth with its 4.0 kW scatter, and the five readings of each window at the path, the load and the drawing's sensors, with the sensors' errors, rounded to 0.01 as the record is;
4. apply the looks, equations (19.10) to (19.12), to the record's readings, giving $q_{36}, \ldots, q_{45}$;
5. for each level c, find $T_{\mathrm{ord}}(c)$ and the loss (19.6) on the record's own path, and average over the records.

The records are 20 000, and every Monte Carlo standard error is by batch means over 20 batches of 1 000. Each record's loss is a step function of c, which changes only where c passes the record's running maximum of $q_{t_{\mathrm{as}}}$, so one pass of the looks prices every level. The full posterior stands beside this route by drawing step 1 from each drawing's reweighted draws (Section 5.2.4, equation (5.13)), 10 000 per drawing, with the drawings weighted by their importance-sampling evidence (7.12) in place of their Laplace evidence; steps 2 to 5 are unchanged. The route of this section is called the looks' route; it is the Laplace normal updated by linearized rows, and step 5 refits records beside it.

19.3 The answers (step 4)

Every answer is read from the looks' route of Section 19.2.4 over the posterior of Section 19.2.1, and that value is the one reported. The full posterior's value stands beside each, and step 5 tests the agreement.

19.3.1 One look

Asset manager's question. At what measured value should the work order fire? It is asked first at one look, the look of window 45 (Sep Y4), the last before the shutdown and Chapter 14's refit.

General form. An expected consequence compares actions by their posterior expected loss, equation (2.23),

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

with the Laplace reading (2.24) beside it; a rule that acts when a probability reaches a level is the threshold form (2.12) with the level in place of the assurance level κ.

Specialized form. The question is equation (2.23) with $a \in \{\text{order at 45},\ \text{wait}\}$, $\mathbf{y} = \mathbf{y}_{0:45}$, and the event $\mathcal{U} = \{G_{\text{tot},46} \le G^*\}$, the shutdown's own month, which is all an order at look 45 can still protect. Its losses are those of Table 19.4: each row is an action and each column a state. One choice enters here: an order when $\mathcal{U}$ occurs pays the premium OW, since the work had to be done, and waiting when $\mathcal{U}$ occurs pays one month's SF, since the shutdown does the work at window 46. This makes the decision turn on one probability, and it fails if a month below the limit costs more than SF. By the second line of equation (2.23), with $\Pr(\mathcal{U} \mid \mathbf{y}_{0:45}) = q_{45}$ from the looks (19.8),

$$\mathcal{R}(\text{order}) = \text{OW} + \text{FO}\,(1 - q_{45}), \qquad \mathcal{R}(\text{wait}) = \text{SF}\, q_{45}, \tag{19.13}$$

two straight lines in q_{45} that cross at the level of equation (19.1) with $C_{\text{false}} = \text{FO} + \text{OW}$ and $C_{\text{miss}} = \text{SF} - \text{OW}$,

$$c_1^* = \frac{\text{FO} + \text{OW}}{\text{FO} + \text{SF}}.$$

A missed month is charged net of the premium the shutdown saves: waiting when the hall goes short avoids OW and pays SF.

Table 19.4: The loss of each action at the look of window 45 (Sep Y4), in reais, the cleaning cancelled; $\mathcal{U} = \{G_{\text{tot},46} \le G^*\}$.

Action	$\mathcal{U}$, the hall short in October [R$]	not $\mathcal{U}$ [R$]
order at look 45	OW = 12 000	OW + FO = 15 000
wait for the shutdown	SF = 150 000	0

The Laplace reading. Equation (2.24) expands the loss about the MAP state. The loss is a step in $G_{\text{tot},46}$, flat on either side of G^*, so its gradient and curvature vanish at the MAP state and the expansion returns the plug-in, as Chapter 18 found for its test schedule (Section 18.3.3). On the leading drawing's MAP path, with the readings at their predictions, $G_{\text{tot},46} = 11.054\,\text{kW/K}$, above the limit, so the plug-in waits at R$ 0, as if the hall surely held; the looks on that path give $q_{45} = 0.0118$. The reported value is the expectation over the looks, equation (19.13), not the expansion.

Building blocks. The probability q_{45} by equations (19.8) to (19.12); the costs of step 0; the lines (19.13); the predicted records of Section 19.2.4 for the cost of a single look.

The level. With the stated costs, $C_{\text{false}} = \text{R\$ } 15\,000$ and $C_{\text{miss}} = \text{R\$ } 138\,000$, and

$$c_1^* = \frac{3\,000 + 12\,000}{3\,000 + 150\,000} = \frac{15\,000}{153\,000} = 0.098039.$$

The predicted records confirm it. A rule that looks only at window 45 has its lowest expected loss at 0.1082, R\$ 35 323.5 (standard error 1 006.4), and every level from 0.0135 to 0.2725 lies within one standard error of it: the loss is flat near its minimum. At 0.0980 it is R\$ 35 345.4; at 0.5, R\$ 38 802.8; at 0.9, R\$ 47 183.7; at 0.95, R\$ 49 053.9. Every one of these is large because a single look leaves windows 37 to 45 unwatched, and their shortfall is paid at every level. The level rests on q_{45} being a probability that means what it says, and over the predicted records it does. Of the records with q_{45} in $[0.05, 0.15)$, mean 0.0918, 0.0933 are short in October; in $[0.15, 0.50)$, mean 0.3019, the share is 0.3128; in $[0.50, 1.00)$, mean 0.8162, it is 0.8382; and below 0.05, mean 0.0040, it is 0.0039.

The surprise of the answer. The level the costs set is one in ten. A rule that waits for "probably short", 0.5, or "confidently short", 0.95, takes a R\$ 150 000 risk to save R\$ 15 000.

19.3.2 Monthly looks

Asset manager's question. With a refit every month, which level costs least to the shutdown, and what do 0.5, 0.9 and 0.95 cost? The rule is read at every look from today to window 45.

General form. A design question is the preposterior expectation (2.16), with U_{p} the loss of the action taken once the outcome is known; for a rule applied to readings not yet logged, the expected loss of the rule is the loss averaged over the state and over the readings' predictive,

$$\mathcal{R}(c) = \iint L(c, \mathbf{z}, \tilde{\mathbf{y}})\, p(\tilde{\mathbf{y}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\tilde{\mathbf{y}}\, \mathrm{d}\mathbf{z}, \qquad c^* = \arg\min_c \mathcal{R}(c),$$

with the expected value of perfect information (17.9) beside it.

Specialized form. The question is equation (2.16) with L the rule's loss (19.6), $\tilde{\mathbf{y}} = \tilde{\mathbf{y}}_{37:45}$, $\mathbf{y} = \mathbf{y}_{0:36}$ and $p(\mathbf{z} \mid \mathbf{y})$ the family of Section 19.2.1, evaluated by the predicted records of Section 19.2.4 with the rule read at every look 36 to 45. One choice enters here: the level reported is the minimum of the expected loss over a log grid of 332 levels from 0.0001 to 0.999, with the named levels added, and the levels whose expected loss lies within one standard error of it are stated beside it. This makes the level one number with an honest range, and it overstates the precision if the range is dropped. The probability that c^* is the better rule than a level c' is read on the same records, as the share in which its loss is lower, equal or higher; it is the second line of equation (2.24) computed without the Laplace approximation, which a step loss defeats. Perfect knowledge of the path orders at the look before the first short window when there is one and never otherwise, so its expected loss is OW times the probability that the hall needs the work before the shutdown, and $\text{EVPI} = \mathcal{R}(c^*) - \text{OW}\ \Pr(\exists\, t \le 46 : G_{\text{tot},t} \le G^* \mid \mathbf{y}_{0:36})$ by equation (17.9). The full posterior stands beside the level: the predicted records drawn from the reweighted draws of Section 19.2.4, with the looks kept on the Laplace normals.

Building blocks. The looks (19.10) to (19.12) on each predicted record; the loss (19.6); the records' paths, which give the windows short; the batch means of Section 19.2.4.

The level. Table 19.5 gives the rules. Each row is a rule; the columns give its level, its expected loss with the standard error, and the shares of the predicted records in which it orders before the shutdown, raises a false order, and lets a window go short, with the expected number of months short. Over the records the hall goes short by window 46 in 0.1616 of them, against Chapter 14's first passage of 0.1595, by window 45 in 0.1061, and at window 37 in none. Read every month, the level that minimizes the expected loss is $c^* = 0.0197$, at R$ 6 094.9 (standard error 82.9), and every level from 0.0117 to 0.0255 lies within one standard error. Its expected loss is R$ 4 553.4 of premiums, R$ 885.0 of shortfall and R$ 656.5 of false orders. Its two error rates are a false order in 0.2188 of the records, or 0.2610 of those that hold through window 46, and a window short in 0.0055, or 0.0340 of those that need the work. On the full posterior, with the predicted records drawn from the reweighted draws and the looks kept on the Laplace normals, the hall is short by window 46 in 0.1628 of the records, and the level is 0.0152, at R$ 6 111.0 (standard error 110.3), inside the reported level's range; the reported level costs R$ 6 170.7 there. Questions 3 and 4 were not run on the full posterior; their route keeps the Laplace looks and the Laplace records.

Table 19.5: The rules for the coil cleaning with a refit every month (looks 36–45), priced over 20 000 predicted records by the loss of equation (19.6); standard errors by batch means in parentheses.

Rule	level [–]	$\mathcal{R}$ [R$]	order before the shutdown [–]	false orders [–]	a window short [–]	months short [–]
the level at its minimum, c^*	**0.0197**	**6 094.9** (82.9)	0.3795	0.2188	0.0055	0.0059
the one-look level c_1^*	0.0980	7 794.9 (171.9)	0.2778	0.1219	0.0245	0.0273
the level 0.5	0.5	22 396.0 (504.3)	0.1452	0.0244	0.1005	0.1372
the level 0.9	0.9	41 716.2 (1 011.4)	0.0585	0.0014	0.1541	0.2734
the level 0.95	0.95	45 453.9 (1 168.0)	0.0440	0.0005	0.1585	0.2995
the dashboard (fires today)	–	14 515.2	1	0.8384	0	0
the calendar (the shutdown, no watch)	–	55 042.5	0	0	0.1616	0.3669
perfect knowledge of the path	–	1 939.2	0.1616	0	0	0

The naive levels. The conventional levels cost more, and the paired differences on the same records resolve each: 0.5 costs R$ 16 301.1 more than c^* (standard error of the difference 496.8), 0.9 costs R$ 35 621.2 more (999.1) and 0.95 costs R$ 39 358.9 more (1 156.5). At 0.95 the rule misses 0.9811 of the records that need the work: by the time the hall is confidently short next month, it is often short this month. The dashboard, priced under the same posterior, fires today and costs R$ 14 515.2, the premium on every record and a false order on the 0.8384 that would hold through the shutdown. The calendar, the shutdown's work with no watch, costs R$ 55 042.5, all of it shortfall.

How firm the choice is. Against 0.5, the rule at c^* has the lower loss on 0.0995 of the records, the same on 0.7060 and a higher one on 0.1945; against the dashboard, 0.6196, 0.3750 and 0.0055; against the calendar, 0.1606, 0.6206 and 0.2188; against the one-look level, 0.0207, 0.8824 and 0.0969. The rule at c^* is better in expectation, not on most records: on most records every rule pays nothing. Where it loses, it has paid R$ 15 000 for a false order; where it wins, it has avoided R$ 150 000 a month. It saves R$ 48 947.6 in expectation against the calendar and R$ 8 420.2 against the dashboard. Perfect knowledge of the path would cost R$ 1 939.2, so the expected value of perfect information is EVPI = R$ 4 155.8, the most any measurement before the shutdown could save.

How the looks move the level. One look sets the level at 0.0980 by the closed form, 0.1082 on the records; monthly looks set it at 0.0197; quarterly looks at windows 36, 39, 42 and 45, on the same one-window quantity, set it at 0.0009 (R$ 8 888.9, standard error 82.3; within one standard error from 0.0005 to 0.0016). A look reads only the next window, but the order it raises is judged against the whole time to the shutdown. At an early look a $q_{t_{\text{as}}}$ of 0.02 already stands for a much larger chance that the work will be needed before October: on the records that need the work, the orders at c^* come a median of 3.0 windows before the first short window. The sparser the looks, the more months each one must cover, and the lower the level on a one-window quantity must go. Chapter 22 shows what repeated looks do to a rule whose level is fixed: refitting every month raises its false alarm from 0.05 to 0.37 (Section 22.3.4). Here the level is not fixed. The costs set it for the looks, and the false-order rate they accept, 0.2188, is the price of a miss rate of 0.0055.

The expected loss against the level is drawn in Figure 19.2. The horizontal axis is the level c on a logarithmic scale and the vertical axis the expected loss in reais; the three curves are the monthly, quarterly and single looks, each with its minimum marked; the vertical gray lines mark the one-look level and the levels 0.5, 0.9 and 0.95; the dashed horizontal lines are the dashboard and the calendar, and the dotted one perfect knowledge. The monthly curve lies below the other two at every level, and to the right of its minimum it rises steeply: each step toward a conventional level adds shortfall months.

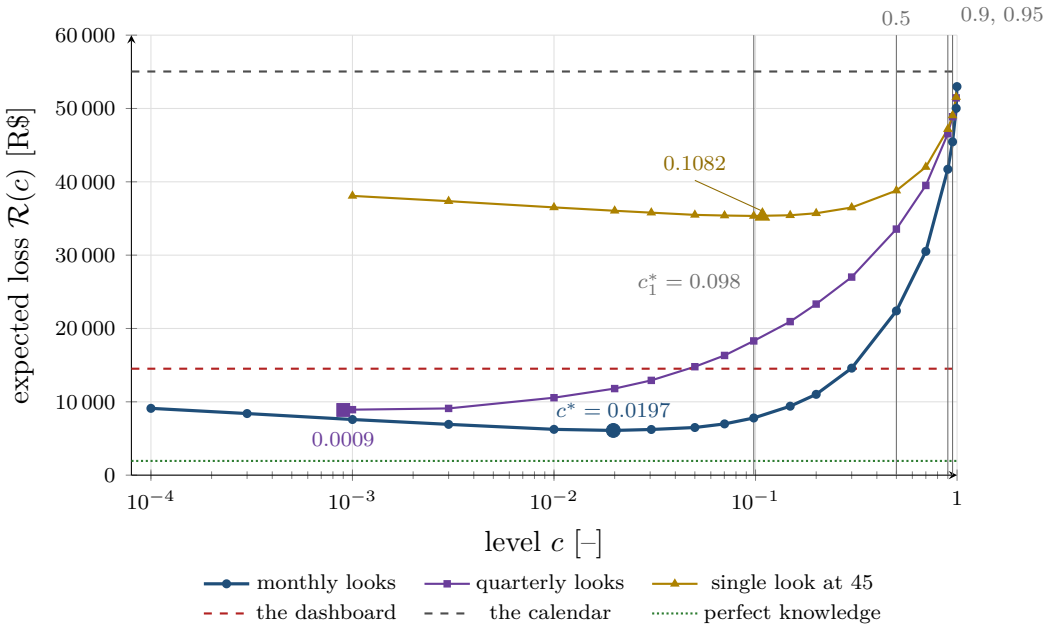


Figure 19.2: The expected loss of the rule against its level, equation (19.6) averaged over the predicted records, for monthly looks (36–45), quarterly looks (36, 39, 42, 45) and a single look at window 45, each with its minimum marked. Vertical gray lines: the one-look level $c_1^* = 0.098$ of equation (19.13) and the levels 0.5, 0.9 and 0.95. Dashed: the dashboard's order today and the calendar's work at the shutdown with no watch. Dotted: perfect knowledge of the path.

The two error rates of the monthly rule are traced against each other in Figure 19.3, the trigger's operating characteristic. The horizontal axis is the share of the records with a window short before the work, the miss rate, and the vertical axis the share with a false order; the curve runs over the levels from 0.0001 at the top left to 0.999 at the bottom right, with the named levels marked, and the dashboard and the calendar are the two single points at the ends of the axes. As the level falls from 0.5 to c^*, the miss rate drops from 0.1005 to 0.0055 while the false orders rise from 0.0244 to 0.2188. Below c^* the false orders keep rising, to 0.4775 at 0.0001, to

remove misses that are already rare, 0.0001 there. The dashboard stands above the whole curve: it pays more false orders, 0.8384, than any level, for a miss rate that the lowest levels nearly reach. The level the costs choose is the point of the curve at which a further cut in misses costs more in false orders and premiums than the months it saves.

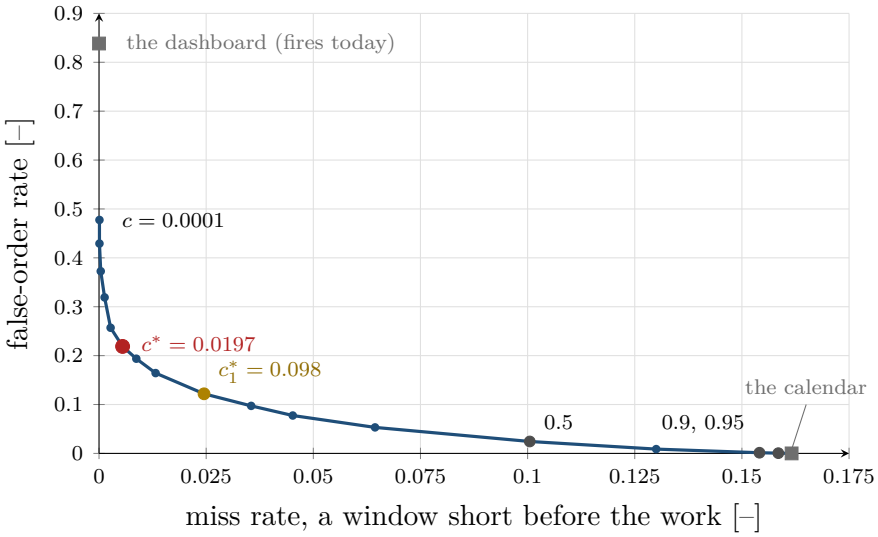


Figure 19.3: The operating characteristic of the monthly trigger: the false-order rate against the miss rate (a window at or below G^* before the work is done) over the predicted records, as the level runs from 0.0001 to 0.999; the named levels are marked. Squares: the dashboard's order today and the calendar's work at the shutdown. No lines of equal expected loss are drawn: the loss (19.6) also depends on the rate of orders, which pay the premium whether or not they are false, and on the number of months short, so a pair of error rates does not fix it.

19.3.3 The month of the order

Asset manager's question. When would this record fire? The asset manager wants the month in which to expect the order, and whether the record as logged has already reached it.

General form. The design form (2.16) with U_{p} an indicator, here of the rule having fired by a look s: the probability, over the readings' predictive, of a first look at or before s,

$$\Pr(T_{\mathrm{ord}} \le s \mid \mathbf{y}) = \int \mathbf{1}\left[T_{\mathrm{ord}}(\tilde{\mathbf{y}}) \le s\right] p(\tilde{\mathbf{y}} \mid \mathbf{y})\, \mathrm{d}\tilde{\mathbf{y}}. \tag{19.14}$$

Specialized form. The question is equation (19.14) with T_{ord} the first look of equation (19.6) at $c = c^*$, $\mathbf{y} = \mathbf{y}_{0:36}$ and the predictive of $\tilde{\mathbf{y}}_{37:45}$ by the predicted records of Section 19.2.4. For the record as logged, the same rule is read at earlier windows on the family refitted to window t_{as} with nothing later, the as-of fits of Chapter 10 (Section 10.4.5), and $q_{t_{\mathrm{as}}}$ follows by equations (19.8) and (19.12). One choice enters here: the month at which a record fires is reported as a distribution over records, not as the date on one. This is what an entry in the log can plan around, and it fails to say which month this record will give, which depends on readings not yet logged.

Building blocks. The rule at c^* on the predicted records of Section 19.2.4; the as-of fits and the look (19.12) at each as-of window.

The predicted records. Table 19.6 gives the distribution. Each row is a look; the columns give the share of the predicted records whose first order comes at that look, the cumulative share at c^* and at c_1^*, the share of the records whose hall has been short at some window up to the one the look protects, and, as the check marker of step 5, the cumulative share of the continuations at the stated truth that have fired at c^*. An order before the shutdown comes in 0.3795 of the records: by window 40 (Apr Y4) in 0.0225, by 42 (Jun Y4) in 0.1434 and by 44 (Aug Y4) in 0.3073. Among the records that fire, the median look is window 43 (Jul Y4), and the 5–95 % range is windows 40 (Apr Y4) to 45 (Sep Y4). Among the records that need the work, the order comes before the first short window in 0.9660, a median of 3.0 windows ahead of it. The look at which a record fires is one record's draw, and its spread over the predicted records is that range: on the leading drawing's MAP path, with the readings at their predictions, $q_{t_{\text{as}}}$ reaches 0.0006 at look 43, 0.0029 at 44 and 0.0118 at 45, and never c^*.

Table 19.6: The month of the order over the predicted records: the share whose first order comes at each look and the cumulative shares at $c^* = 0.0197$ and $c_1^* = 0.0980$; the share whose hall has been short by the window the look protects; and, as a check, the cumulative share of the 400 continuations at the stated truth (step 5).

Look t_{as} [–]	Date	first order at t_{as}, c^* [–]	by t_{as}, c^* [–]	by t_{as}, c_1^* [–]	short by $t_{\text{as}}+1$ [–]	truth, by t_{as}, c^* [–]
36	Dec Y3	0.0000	0.0000	0.0000	0.0000	0.0000
37	Jan Y4	0.0000	0.0000	0.0000	0.0000	0.0000
38	Feb Y4	0.0001	0.0001	0.0000	0.0001	0.0000
39	Mar Y4	0.0024	0.0024	0.0003	0.0003	0.0000
40	Apr Y4	0.0200	0.0225	0.0052	0.0021	0.0050
41	May Y4	0.0481	0.0707	0.0254	0.0083	0.0175
42	Jun Y4	0.0727	0.1434	0.0688	0.0278	0.0725
43	Jul Y4	0.0829	0.2263	0.1331	0.0607	0.2625
44	Aug Y4	0.0809	0.3073	0.2068	0.1061	0.6300
45	Sep Y4	0.0722	0.3795	0.2778	0.1616	0.9200

The orders run ahead of the shortfall at every look. By look 44 the rule at c^* has ordered in 0.3073 of the records, while the hall has been short by window 45 in 0.1061; the one-look level has ordered in 0.2068. The level is low because an order must come before the first short window, and most of the records that will need the work are still above the limit at the look that must act.

The record as logged. The logged record would not have fired. On the family refitted to each window, $q_{24} = 2.590 \times 10^{-15}$ (Dec Y2), $q_{30} = 2.341 \times 10^{-12}$ (Jun Y3), $q_{33} = 1.697 \times 10^{-11}$ (Sep Y3) and $q_{36} = 2.222 \times 10^{-8}$ (Dec Y3). The dashboard fired at window 9.

The cumulative shares are drawn in Figure 19.4. The horizontal axis is the window, a look t_{as} for the orders and a window t for the shortfall; the vertical axis is the share of the records. The solid blue line is the share that has ordered by each look at c^*, the gold line the same at c_1^*, and the red line the share whose hall has been short by each window; the dotted lines are the continuations at the stated truth, at c^* and at c_1^*, the check markers of Section 19.4.6. The blue line rises three to four windows before the red one, which is the lead the order needs. The dotted lines rise later and higher: on the truth's continuations coil 1's acceleration shows only in the last looks, and then the orders come in most of them.

19.3.4 A coil that is not degrading

Asset manager's question. How often does the rule fire on a coil that is not degrading? The asset manager asks how often an order from this rule would be raised for nothing, and how

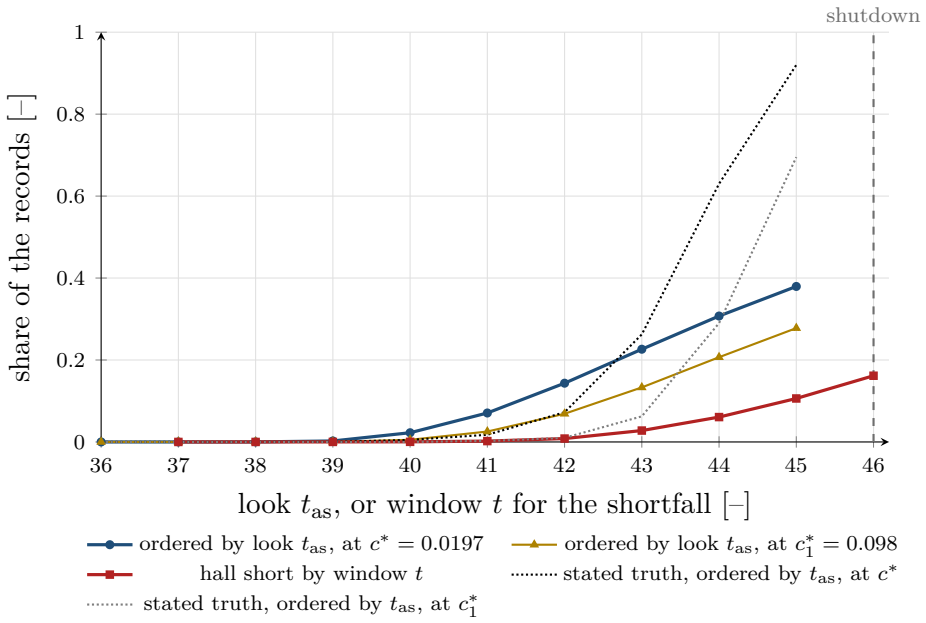


Figure 19.4: When the order fires. Solid: the cumulative share of the 20 000 predicted records with an order by each look, at c^* (blue) and c_1^* (gold), and the cumulative share whose hall has been at or below G^* by each window (red). Dotted: the 400 continuations at the stated truth, which holds through the shutdown, a check marker (Section 19.4.6). Dashed: the October shutdown, window 46.

that compares with the dashboard's.

General form. The design form (2.16), with U_{p} the indicator that the rule fires at one look or more and the outcomes weighted by their distribution at a stated truth rather than by their predictive over the posterior,

$$\text{OC}(c;\ \text{no loss}) = \int \mathbf{1}\Big[\max_{t_{\text{as}} \in \{36,\ldots,45\}} q_{t_{\text{as}}}(\tilde{\mathbf{y}}) \geq c\Big]\, p(\tilde{\mathbf{y}} \mid \text{no loss})\, \mathrm{d}\tilde{\mathbf{y}}, \tag{19.15}$$

the operating characteristic of the schedule of looks. Chapter 22 shows the same form for a rule with a fixed level, equation (22.15). with $p(\tilde{\mathbf{y}} \mid \text{no loss})$ the readings of records whose coils do not lose conductance.

Specialized form. The question is equation (19.15) in two worlds without loss. In the first, *the coils stop losing conductance today*: predicted records as in Section 19.2.4 whose paths from window 36 wander with no trend, read by the looks as reported. In the second, *a record restarted at a new commissioning test*: windows 0 to 45 with the coils held at the truth's commissioned values, $g_1 = -0.02$ and $g_2 = -0.04$, the record's loads, inputs and sensors, and unit 1's drift either absent or present from window 20; the dashboard's tile rule at 0.8164 is read on 4 000 such records, and the trigger on ten of them with the base fit refitted at windows 6, 12, 24, 36 and 45. One choice enters here: the second world reads the trigger on Chapter 10's base fit alone, since a restarted record has no March step to weigh. This keeps the check to the trigger's quantity, and it fails to test the family's own false orders on a new hall; the first world tests those.

Building blocks. The looks (19.10) to (19.12); the records without loss; the tile $A_{u,0}/A_{u,t}$ of the dashboard; the base fit of Chapter 10 at each refit.

The false orders. Table 19.7 gives them. Each row is a rule in one of the two worlds; the columns give the share of the records with an order before the shutdown, every one of them false, and the first window at which the rule fires. When the coils stop losing today, the trigger orders in 0.00575 of the records at c^*, 0.00110 at c_1^* and 0.00005 at 0.5: the largest $q_{t_{\text{as}}}$ over the looks has a median of 2.650×10^{-7} over the records, a 99 % point of 8.912×10^{-3} and a largest value of 0.6168. The dashboard's tiles stand below 0.8164 today, so it orders in every one of them. On a restarted record with sound sensors, the dashboard's tile falls below 0.8164 by window 45 in 0.9950 of the 4 000 records, at a median first window of 2, and by window 12 in 0.8870; with unit 1's drift, in 1.0000 of them. The trigger, refitted at the five windows on ten of those records, never comes near its level: its largest $q_{t_{\text{as}}}$ is 2.812×10^{-31}, and the hall's margin at the next window is never less than 10.5 standard deviations, with a median of 18.1.

Table 19.7: Work orders on coils that are not degrading, every one of them false: the coils stop losing today (predicted records, the looks as reported), and a record restarted at a new commissioning test (the dashboard on 4 000 records; the trigger on ten, the base fit refitted at windows 6, 12, 24, 36 and 45).

Rule and world	orders by window 45 [–]	first window, median [–]
The coils stop losing today		
the trigger at $c^* = 0.0197$	0.00575	–
the trigger at $c_1^* = 0.0980$	0.00110	–
the trigger at 0.5	0.00005	–
the dashboard at 0.8164	1.0000	36
A record restarted at a new commissioning test		
the dashboard, sound sensors	0.9950	2
the dashboard, unit 1's sensor drifting from window 20	1.0000	2
the trigger, base fit (largest $q_{t_{\text{as}}}$ 2.812×10^{-31})	none of ten	–

The false orders accumulate by window in Figure 19.5. The horizontal axis is the window from the restarted commissioning test and the vertical axis, on a logarithmic scale, the share of the records with an order by that window. The solid and dashed red lines are the dashboard's tile rule on the restarted record, with sound sensors and with unit 1's drift; the blue line is the trigger at c^* on the records whose coils stop losing today, from its first nonzero look. The dashboard's two lines nearly coincide: they start near one half at window 1 and pass 0.88 by window 12. The trigger's line stays below 0.006, more than two decades lower, and the restarted record's refits give it nothing to draw.

The surprise of the answer. The dashboard's rule is a false-order generator. Its tile moves with the load's growth and the season, so on a hall whose coils do not foul at all it fires within the first months. Read on the parameter, the trigger fires on such a hall in 0.00575 of the records from today, and on a new hall not once.

19.4 The checks (step 5)

Each answer of Section 19.3 rests on the looks' route, on the costs and on the stated lead time, and on a forecast whose calibration Chapter 14 questioned; each is tested here.

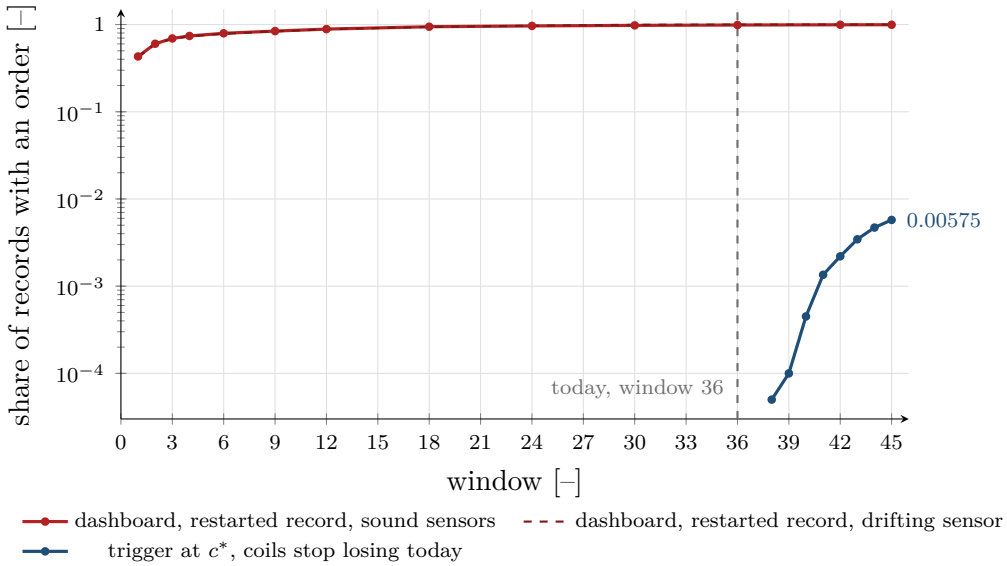


Figure 19.5: Work orders on coils that are not degrading, cumulative by window, on a logarithmic scale. Red: the dashboard's tile rule at 0.8164 on 4 000 records restarted at a new commissioning test, with sound sensors (solid) and with unit 1's drift from window 20 (dashed). Blue: the trigger at $c^* = 0.0197$ on the predicted records whose coils stop losing at window 36, from its first nonzero look; on the restarted records, refitted at windows 6, 12, 24, 36 and 45, it never fires and has no point to draw.

19.4.1 The looks against refits

The looks update each drawing's Laplace normal by rows linearized on its MAP path. The check refits six predicted records in full: every drawing fitted again by Gauss–Newton on windows 0 to 45, its weight from its Laplace evidence, and q_{45} from the refitted posteriors. On the six records, q_{45} by refit and by the looks is 0.6556 against 0.6277, 0.4778 against 0.4467, 0.9988 against 0.9992, and below 0.00005 for both on the other three. The leading drawing's weight at window 45 is 0.5461 against 0.5305, 0.6829 against 0.6414, 0.2957 against 0.2910, 0.5331 against 0.5731, 0.1553 against 0.1481 and 0.4157 against 0.4118. The largest gap is 0.0311 in q_{45} and 0.0556 in a weight. The linearization errs by up to three points of probability, beside a level whose one-standard-error range is 0.0117 to 0.0255: the decision rests on records whose $q_{t_{\text{as}}}$ passes the level by more than that, and the route is the looks' route, not a refit.

19.4.2 The full posterior and the Monte Carlo errors

The full posterior of Section 19.2.4, with the predicted records drawn from each drawing's 10 000 reweighted draws and the drawings weighted by their importance-sampling evidence, moves no drawing's weight by more than 0.0009. On those records the hall is short by window 46 in 0.1628, against 0.1616; the level that minimizes the expected loss is 0.0152, at R$ 6 111.0 (standard error 110.3), with the one-standard-error range 0.0104 to 0.0203: each route's range contains the other route's level. The reported level costs R$ 6 170.7 there, and the levels 0.5, 0.9 and 0.95 cost R$ 23 278.3, R$ 42 325.9 and R$ 45 931.7. The Monte Carlo errors are batch means over 20 batches of 1 000 records, on paired differences: against the reported level, the one-look level costs R$ 1 700.0 more (standard error 149.3), 0.5 R$ 16 301.1 more (496.8), 0.9 R$ 35 621.2 more (999.1) and 0.95 R$ 39 358.9 more (1 156.5). Each difference is many standard errors from zero, and the level's own range is the one stated with it. The chapter adds no physical row: the looks' readings are Chapter 5's rows on the reduced state at later windows,

whose limit $\sigma_p \to 0$ Chapter 10 checked on the record's full state (Section 10.4.2).

Table 19.8 collects the routes and the variants of the next three subsections. Each row is a route or a variant; the columns give the level that minimizes the expected loss on it, with its one-standard-error range, that loss, the loss of the reported level, and the losses of the levels 0.5 and 0.95.

Table 19.8: The routes to the trigger's level and the variants of step 5, in reais; standard errors in parentheses, the one-standard-error range of the level in brackets.

Route or variant	c^* [–]	$\mathcal{R}(c^*)$ [R$]	$\mathcal{R}$ at the reported c^* [R$]	$\mathcal{R}(0.5)$ [R$]	$\mathcal{R}(0.95)$ [R$]
the looks over the Laplace normals (reported)	0.0197 [0.0117, 0.0255]	6 094.9 (82.9)	6 094.9	22 396.0	45 453.9
the full posterior, reweighted draws	0.0152 [0.0104, 0.0203]	6 111.0 (110.3)	6 170.7	23 278.3	45 931.7
the plug-in on the MAP path	none fires	0	0	0	0
a lead time of two windows	0.0241 [0.0131, 0.0331]	6 908.9 (123.2)	–	–	–
the coil clause added	0.5608	7 712.1	14 515.2	–	–
the forecast as late as fresh records show	0.0020 [0.0013, 0.0029]	12 747.6 (53.4)	14 479.2	72 182.4	129 872.6
the stated truth, 400 continuations	–	–	13 800.0	1 650.0	0

19.4.3 The costs that could flip the level

The level rests on three stated costs, FALSE_ORDER, OFF_WINDOW and SHORTFALL, and each is moved here alone, with the level chosen again on the same records; none is chosen again in the answer. Table 19.9 gives the variants. Each row is a variant; the columns give the one-look level (19.13), the monthly level and its expected loss, the expected loss of the reported level under the variant, and those of 0.5 and 0.95.

Table 19.9: The costs moved one at a time, the monthly level chosen again on the same predicted records, in reais.

Variant	c_1^* [–]	c^* [–]	$\mathcal{R}(c^*)$ [R$]	$\mathcal{R}$ at the reported c^* [R$]	$\mathcal{R}(0.5)$ [R$]	$\mathcal{R}(0.95)$ [R$]
as stated	0.0980	0.0197	6 094.9	6 094.9	22 396.0	45 453.9
FALSE_ORDER ×0.5	0.0891	0.0197	5 766.7	5 766.7	22 359.5	45 453.2
FALSE_ORDER ×2	0.1154	0.0215	6 744.0	6 751.5	22 469.1	45 455.4
FALSE_ORDER ×5	0.1636	0.0455	8 451.8	8 721.1	22 688.2	45 459.9
OFF_WINDOW ×0	0.0196	0.0057	1 124.1	1 541.5	20 653.0	44 926.5
OFF_WINDOW ×0.5	0.0588	0.0117	3 706.7	3 818.2	21 524.5	45 190.2
OFF_WINDOW ×2	0.1765	0.0455	10 364.5	10 648.4	24 139.0	45 981.3
SHORTFALL ×0.1	0.8333	0.2805	3 694.3	5 298.4	3 874.1	5 021.4
SHORTFALL ×0.5	0.1923	0.0500	5 435.4	5 652.4	12 106.0	22 991.4
SHORTFALL ×2	0.0495	0.0080	6 645.9	6 979.9	42 976.1	90 378.9

The level stays far below the conventional ones under every variant but a tenfold cut in SHORTFALL, which raises it to 0.2805. The question that matters is what it would take to bring the level up to the first naive level, 0.5, with one cost moved at a time. FALSE_ORDER cannot, within ×0.01 to ×100 of its value, though the closed form for one look reaches 0.5 at ×42.00; OFF_WINDOW would have to rise ×17.78 (closed form ×6.12); SHORTFALL would have to fall to ×0.06 (closed form ×0.18). SHORTFALL is the cost the level rests on.

FALSE_ORDER barely moves the monthly level, 0.0215 when doubled, because the premium for work away from the outage, not the false order, is most of what an order costs. The cleaning is paid once on every branch, and no value of it moves the expected loss of one level against another.

19.4.4 The lead time and the coil clause

Two windows of lead. With an order taking two windows to be done, the looks run from 36 to 44 and read $q_{t_{\mathrm{as}}} = \Pr(G_{\mathrm{tot},t_{\mathrm{as}}+2} \le G^* \mid \mathbf{y}_{0:t_{\mathrm{as}}})$. The level is 0.0241 (one-standard-error range 0.0131 to 0.0331), at R$ 6 908.9 (standard error 123.2), with an order before the shutdown in 0.4067 of the records, a false order in 0.2464 and a window short in 0.0078. The level belongs to the lead time and moves with it, but little.

The coil clause. The per-coil clause a coil-by-coil trigger adds reads either coil below its share, the probability $\Pr(\min_u g_{u,t_{\mathrm{as}}+1} \le g^* \mid \mathbf{y}_{0:t_{\mathrm{as}}})$ at each look, and its event contains the hall's. Today it reads 0.0689, against the hall's 2.222×10^{-8}, so at the reported level it fires today and costs what the dashboard costs, R$ 14 515.2. Chosen again for it, its level is 0.5608, at R$ 7 712.1, with an order before the shutdown in 0.4757 of the records, a false order in 0.3156 and a window short in 0.0063: at its own best level it costs more than the hall's quantity at its own, R$ 6 094.9. The clause orders on coil 1 passing its share while coil 2 still carries the hall, and no cost in the schedule pays for that. With Part I's functional limit g_F as the coil's limit, coil 2 is already past it on the leading drawing, $\Pr(g_{2,36} \le g_F \mid \mathbf{y}) = 0.6669$, and coil 1 is past it with 0.2879, so that clause fires today as well. The hall's quantity is the one the costs price.

19.4.5 A forecast as late as fresh records show

Chapter 14 found that its forecast is late on most fresh records at the stated truth (Section 14.4.5), and its family stays the reported posterior. The variant here draws predicted records whose coils follow the stated truth's trend from today, coil 1 accelerating and coil 2 at 2 %/yr, from the posterior's level at window 36, with the wander, and reads them by the looks as reported. The hall is then short by window 46 in 0.5327 of the records, against the posterior's 0.1616. The level that minimizes the loss there is 0.0020, at R$ 12 747.6 (standard error 53.4), with the one-standard-error range 0.0013 to 0.0029. The reported level costs R$ 14 479.2 there, with a false order in 0.3152 of the records and a window short in 0.0223; the one-look level costs R$ 24 176.0, 0.5 costs R$ 72 182.4 and 0.95 costs R$ 129 872.6; the dashboard R$ 13 401.8 and the calendar R$ 143 017.5.

The level moves, and at the reported level the decision flips. A late forecast pushes the level lower, so the reported level errs toward waiting, and there the dashboard's order today, R$ 13 401.8, costs less than the reported rule, R$ 14 479.2. Only the level chosen again for this world, 0.0020 at R$ 12 747.6, beats ordering now, and every naive level is far worse than either. What settles which world the hall is in before the October shutdown is the refit at window 45 that Chapter 14 recommends, which is the last of this rule's looks: in this world the looks' posterior puts coil 1's curvature at a median of 0.9936 by window 45, against 0.6098 on the reported records, and $q_{t_{\mathrm{as}}}$ follows it.

19.4.6 The answers against the stated truth

The last check sets the rule against the stated truth of the record (Section 19.6.5), continued to window 45 in 400 continuations with fresh loads and sensor errors. The truth's total conductance runs 11.7257, 11.6363, 11.5432, 11.4467, 11.3470, 11.2443, 11.1387, 11.0305, 10.9199, 10.8071 and 10.6923 kW/K at windows 37 to 47: it holds through the shutdown by 0.0461 kW/K and is first below the limit at window 47. Every order before the shutdown is therefore false on

the truth, and the right action is the calendar's. At c^* the rule orders before the shutdown in 0.9200 of the continuations, at a median first look of 44 (Aug Y4), and costs R$ 13 800.0; at c_1^*, in 0.6950, for R$ 10 425.0; at 0.5, in 0.1100, for R$ 1 650.0; at 0.95, in none. At the truth q_{45} has a median of 0.1825 over the continuations, with 5–95 % range 0.0095 to 0.6881, and the looks put coil 1's curvature at a median of 0.9987 by window 45. The dashboard orders today, a false order at R$ 15 000.

A truth that misses the limit by one month is the case the trigger is priced to be wrong on. Its posterior at window 45 puts the October window short with a median probability of 0.1825, and a month short costs R$ 150 000 against a false order's R$ 15 000, so the order is the right bet on that evidence. A rule at 0.95 wins on this truth by waiting, and it loses R$ 39 358.9 in expectation.

19.5 The report (step 6)

The dashboard's tile rule has been asking for the cleaning since September of year 1, and it is not a condition reading: on a new hall with coils that do not foul at all it fires by window 45 in 0.9950 of records, because the tile moves with the load and the season. Raise the work order when the probability that the hall is below its capacity limit next month reaches 2 %, 0.0197. At a single look the level would be one in ten, 0.098; read every month it falls, because each look's probability covers only the next month while the order is judged on the whole time to the shutdown, so a small probability for next month already stands for a larger chance that the work is needed before October. The level comes from the costs: a false order costs R$ 15 000 with its premium, and a month short R$ 150 000. At 0.0197 the rule raises a false order in 0.2188 of the predicted records and lets the hall go short in 0.0055; waiting for 50 % costs R$ 16 301 more in expectation, and waiting for 95 % R$ 39 359 more. The rule fires before the shutdown in 0.3795 of the records, most often in July of year 4, three months before the hall would go short; on this record it has not fired. If the forecast is as late as fresh records show, the reported level loses to an order today, and only a level lowered to 0.0020 beats it; the refit at window 45 settles which forecast holds before the October shutdown.

19.5.1 The answers and their bands

Table 19.10 gives the answers in the order of Table 19.1. Each row is a question; the second column gives the reported value from the looks over the Laplace normals, the third the full posterior's beside it, and the fourth the band: the one-standard-error range of a level, a standard error, or the error rates the level gives.

The answers rest on the family's posterior and on SHORTFALL. The full posterior's level, 0.0152, lies inside the reported level's range, and the reported level costs R$ 6 170.7 there. The costs could move the level to a conventional one only if a month short cost about ×0.06 of its stated value.

19.5.2 The entry in the maintenance log

The trigger of window 36 is the next entry of the two-unit hall's maintenance log, after Chapter 14's forecast, Chapter 17's decision and Chapter 18's test program at the same window. It uses nothing later than window 36, and in the values of Table 19.10 it reads:

> *Window 36 (Dec Y3), condition-based trigger for the coils' work, from the condition record (windows 0 to 36) and Chapter 14's forecast family; costs from the costs schedule (false order* R$ 3 000*, off-window premium* R$ 12 000*, a month below the capacity limit* R$ 150 000*). Valued for the hall with no work done before the shutdown, the plan Chapter 17's step 5 prefers when coil 1's rate is constant (it cleans at the shutdown), and the calendar's.*

Table 19.10: The report for the coil cleaning's trigger at window 36 (Dec Y3). Reported values from the looks over the Laplace normals of the family; the full posterior from reweighted draws; costs in reais.

Asset manager's question	Reported value	Full posterior	Band
At what measured value should the work order fire?	at one look, $q_{45} \geq c_1^* = 0.0980$ [–]	the closed form, from the costs alone; the single look not run on the full posterior	on the records, 0.1082, within one se 0.0135 to 0.2725 [–]
With a refit every month, which level costs least to the shutdown, and what do 0.5, 0.9 and 0.95 cost?	$c^* = 0.0197$ [–], $\mathcal{R} = 6\,094.9$ [R$]; 0.5: 22 396.0, 0.9: 41 716.2, 0.95: 45 453.9 [R$]; EVPI 4 155.8 [R$]	$c^* = 0.0152$ [–], $\mathcal{R} = 6\,111.0$ [R$]	c^* within one se 0.0117 to 0.0255 [–]; se 82.9 [R$]; false orders 0.2188, a window short 0.0055 [–]
When would this record fire?	before the shutdown with probability 0.3795 [–]; median window 43 (Jul Y4)	not run; the Laplace looks and records	windows 40 (Apr Y4) to 45 (Sep Y4), 5–95 %; not fired on the logged record, $q_{36} = 2.2 \times 10^{-8}$ [–]
How often does the rule fire on a coil that is not degrading?	0.00575 from today [–]; never on a new hall (ten refits)	not run; the Laplace looks and records	the dashboard: 1.0000 from today, 0.9950 on a new hall [–]

Rule: at each monthly refit (windows 36 to 45), raise the off-window order when $P(G_1 + G_2$ below 10.761 kW/K next month$) \geq 0.02$. Otherwise the work stays at the October shutdown (window 46). Read on the hall's own sensors, the rule replaces Chapter 18's test-based order level ($G_1 + G_2 \leq 11.261$ kW/K) for orders between tests; the June Y4 test (window 42) stays as a reference reading.
Expected cost to the shutdown R$ 6 095 (false orders 0.22, a month short 0.006); the level 0.5 would cost R$ 22 396, 0.95 R$ 45 454.
Today: $P = 2 \times 10^{-8}$; no order. Expected to fire before the shutdown with probability 0.38, most likely Jul Y4.
Not used: the dashboard's tile rule (fired Sep Y1; fires on an unfouled hall in 0.995 of records).
Caution: if coil 1 is accelerating as fast as fresh records suggest, the rule at 0.02 costs more than an order today (R$ 14 479 against R$ 13 402), and only a level of 0.002 beats the order (R$ 12 748); the refit at window 45 settles which forecast holds before the work order is written.
Next measurement: a second discharge sensor on unit 1 (R$ 900) about breaks even on this decision (R$ 968 saved); a reference reading of the hall's sensors saves R$ 986 against R$ 1 500, and heat meters R$ 796 against R$ 24 000. The side-by-side check of unit 1's discharge sensor stays open for Chapter 7's entry.

The level is rounded to 0.02 in the log, inside its one-standard-error range. The rule brings forward the shutdown's work on a hall with no work done before it; whether that work is a cleaning or a replacement, and whether it waits for the shutdown at all, is Chapter 17's decision, and if coil 1 is replaced in January as that decision recommends, the rule is set again on the new coil's record. Between Chapter 18's tests the rule raises the order from the monthly record; the test itself remains a reference reading.

Common mistake

Setting a trigger at a conventional confidence. A level of 0.95 says that a false order costs nineteen times a missed month; here a missed month, net of the premium, costs about nine times a false order with its premium, 138 000 against 15 000, and the level the costs set for monthly looks is 0.0197. At 0.95 the rule misses 0.9811 of the records that need the work.

19.6 The next measurement (step 7)

The report's rule costs R$ 6 094.9 in expectation, against R$ 1 939.2 with the path known. A further measurement is one row or more added to the looks, and what it saves is known before its value is.

19.6.1 The preposterior of a candidate row

A candidate measurement a adds rows to the looks, at a stated window or at every window from 37. Its preposterior covariance at each look is the update (19.10) with the candidate's lines appended to $\mathbf{J}_t$, Chapter 4's form (4.65), and its readings on each predicted record are drawn from that record's path with the instrument's error. The level is chosen again for each candidate, and the candidate's value is the expected loss it saves at its own level, Chapter 2's expected value of sample information (2.52),

$$\mathrm{EVSI}(a) = \mathcal{R}(c^{*}) - \min_{c} \mathcal{R}_{a}(c), \tag{19.16}$$

with $\mathcal{R}_a$ the expected loss of the rule with the candidate's rows, on the same records. The candidates are a side-by-side check of unit 1's discharge sensor now, at 0.1 K, which reads the drift and whose price the schedule does not state; a reference condition test at window 40 (Apr Y4) or 44 (Aug Y4), one crew visit that reads the return air, both discharges and rack 1's inlet at 0.1 K, priced at one INSPECT, R$ 1 500, as Chapters 23 and 24 price it; water-side heat meters on both coils at 3.0 kW, permanent from window 37, at Chapter 23's R$ 12 000 a coil (Table 23.1); and a second discharge sensor on unit 1 at 0.3 K, permanent from window 37, at Chapter 23's R$ 900, the permanent instrument at the place of the question, where the drift's ridge with coil 1's level lies.

19.6.2 The candidates

Table 19.11 ranks them. Each row is a candidate; the columns give the level chosen again for it, its expected loss, the expected loss saved with its standard error, its price, its two error rates, and the spread of coil 1's rate at window 45 on the leading drawing. The savings are compared as bars in Figure 19.6, with the prices that fall on its axis and the expected value of perfect information marked.

The rule is worth far more than anything it could be given. It saves R$ 48 947.6 against the calendar and R$ 8 420.2 against the dashboard, while the best candidate, both permanent instruments, adds R$ 1 403.1 at a price of R$ 24 900, and no single candidate adds R$ 1 000. The reference condition test saves R$ 985.5 at window 44 against its R$ 1 500, and the heat meters R$ 795.9 against R$ 24 000. The second discharge sensor on unit 1, the permanent instrument at the place of the question, saves R$ 967.6 (standard error 103.1) against its R$ 900: it about breaks even on this decision alone. It saves about what a reference test saves because both break the ridge between the drift and coil 1's level, the correlation +0.8280 of Table 19.2. Every candidate lowers the false orders more than the misses, so the level moves little, from 0.0197 to 0.0271. The sensor's value over the hall's life, beyond this one decision to the shutdown, is Chapter 23's question.

19.6.3 The recommendation

Adopt the rule at 0.02 for the looks to the shutdown, and read it at every monthly refit, Chapter 14's at window 45 included. For the rule's sake alone, no measurement clearly pays: the second discharge sensor on unit 1 about breaks even at R$ 900, a reference reading saves less than its R$ 1 500, and the heat meters save a small fraction of their R$ 24 000. The second discharge sensor is the one to fit if Chapter 23 finds it pays over the hall's life, since this decision

Table 19.11: Candidate measurements for the looks to the shutdown, each with the level chosen again and its saving by equation (19.16), in reais; standard errors of the paired saving in parentheses. Prices: the reference reading is one INSPECT; the instruments are priced by Chapter 23 (Table 23.1); the side-by-side check has no stated price (–).

Candidate	c^* [–]	$\mathcal{R}$ [R$]	saved [R$]	price [R$]	false orders [–]	a window short [–]	sd γ_1 at 45 [yr^{-1}]
the existing looks	0.0197	6 094.9	–	–	0.2188	0.0055	0.01259
the side-by-side check of unit 1's discharge sensor now (0.1 K)	0.0271	5 419.9	675.0 (89.3)	–	0.1726	0.0056	0.01203
a reference condition test at window 40 (Apr Y4)	0.0197	5 114.6	980.4 (74.7)	1 500	0.1787	0.0033	0.01191
a reference condition test at window 44 (Aug Y4)	0.0203	5 109.4	985.5 (44.7)	1 500	0.1635	0.0046	0.01192
water-side heat meters on both coils (3.0 kW), permanent from 37	0.0215	5 299.1	795.9 (100.8)	24 000	0.1777	0.0046	0.01192
a second discharge sensor on unit 1 (0.3 K), permanent from 37	0.0221	5 127.3	967.6 (103.1)	900	0.1671	0.0044	0.01189
both permanent instruments	0.0200	4 691.9	1 403.1 (92.7)	24 900	0.1550	0.0028	0.01177
perfect knowledge of the path (EVPI)	–	1 939.2	4 155.8	–	0	0	–

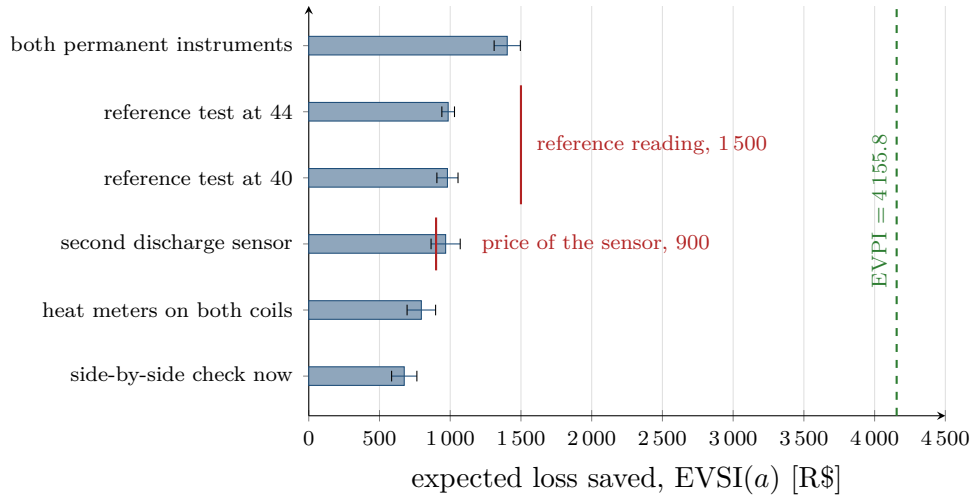


Figure 19.6: The expected loss each candidate saves for the trigger, equation (19.16), with the level chosen again for each; error bars: one standard error of the paired saving. Red: the price of a reference condition test, one INSPECT, beside the two tests, and Chapter 23's price of the second discharge sensor beside its bar; the heat meters' R$ 24 000 lies far off the axis. Dashed: the expected value of perfect information, the most any measurement could save.

already returns its price. The side-by-side check of unit 1's sensor stays open for Chapter 7's reason. If permanent instruments are fitted, choose the level again with them: 0.0221 with the second discharge sensor, 0.0200 with both.

19.6.4 What the analysis showed

The dashboard's tile rule fired in September of year 1, has stood since in most windows, and fires on an unfouled hall in 0.9950 of records; the trigger on the hall's posterior, at a level of 2 %, has not fired, and fires on a hall that stops fouling today in 0.00575 of records. The physics puts the trigger on the coils' conductance and the hall's limit rather than on a ratio of approaches, and that alone removes the load and the season from the alarm. The probability, priced by the costs, sets the level, and the level is not a confidence: one in ten at a single look, about one in fifty when the record is read every month, because a look reads only the next month while the order is judged on the whole time to the shutdown, so a small probability for next month already stands for a larger chance that the work will be needed before October. A conventional confidence costs R$ 22 396.0 at 0.5 and R$ 45 453.9 at 0.95, against R$ 6 094.9. What the analysis did not settle is whether coil 1 is accelerating as fast as fresh records suggest; that moves the level, and the refit at window 45 settles it.

19.6.5 Limits of the model

A constructed record. The condition record is generated from a stated truth, with the stated sensor errors, and step 5 compares the rule with that truth (Section 19.4.6). It is not field data. On the truth, which holds through the shutdown by one month, the rule's orders are false in 0.9200 of the continuations: the case it is priced to be wrong on.

The looks' route. Each look is the Laplace normal updated by rows linearized on the drawing's MAP path. Refits agree to within 0.0311 in probability; a record whose path bends far from the MAP path would need refits at every look.

The horizon and the work. The horizon ends at the shutdown, and one order cleans both coils. The cost of the next cleaning cycle coming earlier after an early order is not priced, nor is a coil cleaned alone, nor is the discounting of costs within the ten months.

The lead time and the looks. The lead time and the look interval are stated inputs, one window each. The level belongs to them and must be chosen again if either changes: quarterly looks set it at 0.0009, and two windows of lead at 0.0241.

The costs. The level rests on FALSE_ORDER, OFF_WINDOW and SHORTFALL, and above all on SHORTFALL, which would have to fall to $\times 0.06$ of its value to bring the level to 0.5. The cleaning cancels.

Chapter 14's and Part III's limits. The family holds Part III's two open questions only; the drift is linear from window 20; the readings are independent across windows; and each window is in steady state with known inputs, Chapter 5's choices.

The next chapter. With one crew and fifty coils, the order raised here competes for the crew's week, and Chapter 20's ranking prices the wait coil by coil.

19.7 What to retain

Definition 19.1 (Condition-based maintenance trigger)**.** A trigger raises the work order at the first look t_{as} at which the posterior probability of the event an order would prevent, $q_{t_{\text{as}}} = \Pr(G_{\text{tot},t_{\text{as}}+\Delta_{\text{L}}} \le G^* \mid \mathbf{y}_{0:t_{\text{as}}})$, reaches a level c. At one look the level is $c_1^* = C_{\text{false}}/(C_{\text{false}}+C_{\text{miss}})$, the probability at which the expected losses of ordering and waiting cross. With repeated looks it is the minimizer of the rule's expected loss over records not yet logged, the design form of Chapter 2, read by applying the looks to predicted records; each look adds a window's rows to every drawing's posterior and moves the drawings' weights by the rows' predictive density. The level belongs to the costs, the lead time and the look interval.

Three conclusions carry beyond the hall. A trigger is a level on a probability of the parameter, and the parameter is what keeps the load and the season out of the alarm; a raw-performance tile has no level at which it stops firing on a hall that does not foul. The level is set by the costs of the two errors and by how often the record is read, never by a conventional confidence, which is a statement about costs made without reading them. And a trigger is judged in expectation: on most records every rule pays nothing, and the trigger's worth lies in the few on which a month below the limit is avoided.

Table 19.12 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 19.12: The terms of the condition-based maintenance trigger and where each is defined. Units for the two-unit hall.

Term	Expression	Units	Meaning	Defined in
Trigger's probability	$q_{t_{\text{as}}}$, eq. (19.2)	–	the hall below its limit at the end of the lead time	19.1.2
One-look level	$c_1^* = C_{\text{false}}/(C_{\text{false}} + C_{\text{miss}})$, eq. (19.1)	–	the two expected losses cross	19.1.1
False order and missed month	$C_{\text{false}} = \text{FO} + \text{OW}$, $C_{\text{miss}} = \text{SF} - \text{OW}$	R$	the rule's two errors	19.3.1
Loss of a rule	eq. (19.6)	R$	premiums, false orders, months short	19.1.4
Rows and posterior				
Carried entries	$\boldsymbol{\theta}_t$, eq. (19.3)	mixed	all a later window's rows touch	19.1.4
A look	eq. (19.10)	–	the window's rows added by the gain	19.2.2
A drawing's weight after a look	eq. (19.11)	–	moved by the rows' predictive density	19.2.2
The region one window ahead	eq. (19.12)	–	a curved region by quadrature	19.2.2
Questions and answers				
Monthly level	$c^* = \arg\min_c \mathcal{R}(c)$	–	the looks re-set the level	19.3.2
Operating characteristic of a trigger	false orders against misses as c runs	–	the two error rates traded	19.3.2
First look	T_{ord}, eq. (19.14)	window	one record's draw	19.3.3
False orders at no loss	eq. (19.15)	–	the dashboard's against the trigger's	19.3.4
Plug-in of a step loss	$L(a, \hat{\mathbf{z}})$	R$	the hall surely holds; misleading	19.3.1
Coil clause	$\Pr(\min_u g_{u,t_{\text{as}}+1} \le g^*)$	–	contains the hall's event; unpriced	19.4.4
Value of a candidate	eq. (19.16)	R$	saved at its own level	19.6.1

19.8 Exercises

Exercise 19.1. *Formulation.* Write the one-look rule with FALSE_ORDER doubled in the form of Table 19.4 and find its level. (The order's losses become 12 000 and 18 000 and waiting's 150 000 and 0, so by equation (19.1) $c_1^* = (6\,000 + 12\,000)/(6\,000 + 150\,000) = 0.115385$.)

Exercise 19.2. *Derivation.* Show that the one-look level is $C_{\text{false}}/(C_{\text{false}} + C_{\text{miss}})$ with $C_{\text{false}} = \text{FO} + \text{OW}$ and $C_{\text{miss}} = \text{SF} - \text{OW}$, and find it when the order's work could be done away from the outage at no premium. (By equation (19.13), $\mathcal{R}(\text{order}) - \mathcal{R}(\text{wait}) = \text{OW} + \text{FO}(1 - q_{\mathcal{U}}) - \text{SF}\, q_{\mathcal{U}} = (1 - q_{\mathcal{U}})(\text{FO} + \text{OW}) - q_{\mathcal{U}}(\text{SF} - \text{OW})$, which vanishes at $q_{\mathcal{U}} = (\text{FO} + \text{OW})/(\text{FO} + \text{SF})$; with $\text{OW} = 0$, $3\,000/153\,000 = 0.019608$.)

Exercise 19.3. *Application.* Evaluate $\mathcal{R}(\text{order})$ and $\mathcal{R}(\text{wait})$ at the look of window 45 for $q_{45} = 0.05$ and $q_{45} = 0.2$, and decide each. (By equation (19.13), at 0.05: $12\,000 + 3\,000 \times 0.95 = 14\,850.0$ against $150\,000 \times 0.05 = 7\,500.0$, wait; at 0.2: 14 400.0 against 30 000.0, order.)

Exercise 19.4. *Application.* Price the dashboard's order under the posterior. (It fires at today's look, so it pays OW on every record and FO on the records that would hold through window 46, whose share is $1 - 0.1616$: $12\,000 + 3\,000 \times 0.8384 = 14\,515.2$.)

Exercise 19.5. *Derivation.* Show that perfect knowledge of the path costs OFF_WINDOW times the probability that the hall needs the work before the shutdown, and evaluate it. (With the path known, the order is raised at the look before the first short window when there is one, which pays OW and no shortfall, since window 37 is short in none of the records, and never otherwise, which pays nothing; so $\mathbb{E}[L] = \text{OW}\ \Pr(\exists\, t \le 46 : G_{\text{tot},t} \le G^* \mid \mathbf{y}_{0:36}) = 12\,000 \times 0.1616 = 1\,939.2$.)

Chapter 20

Prioritizing Maintenance Across a Fleet

20.1 The worst-first list and the crew's next visit

The fifty coils of Chapter 9 stand at the operator's sites, one to twelve years old, each tested once a year by repeating its commissioning test. Today is the end of the record for every coil, window 36 (Dec Y3), and coil 0, three years old, is the hall's coil 1. The fleet's condition-monitoring dashboard lists the coils by their latest annual condition test, lowest first: coil 15, twelve years old, at −0.7394, then coil 44 at −0.4361, coil 21 at −0.4207, coil 28 at −0.4084 and coil 4 at −0.3801. Coil 15's reading is a derate of 0.4774, 52.3 % below its rating, and 23 of the 50 coils read at or below the functional limit of Part I, $g_F = -0.18232$. The planner has one crew, which can reach one site a week. The conclusion the dashboard invites is to send the crew down the list, worst first, starting with coil 15 this week.

A planner who reads the fleet chart's own lines instead, each coil's straight line through its own tests, sees 23 coils past the limit by their lines and two crossings ahead inside the year: coil 33 in 0.35 yr and coil 34 in 0.62 yr. Neither list says which of the coils reading past the limit have already crossed it, so that a visit can no longer prevent their crossing; which coil still above the limit is about to cross it, and when; or how far one test, read to a standard deviation of 0.04, places a coil today. The question of this chapter is

which coil the one crew should visit next, and in what order it should visit the fifty, read as the expected consequence of each order over the fleet's posterior from Chapter 9, against the dashboard's worst-looking-first order.

It is answered with the method of Chapter 4, on Chapter 9's state of 104 entries.

Jargon: visiting order, the crew's cycle, cost of waiting, sunk consequence, backlog, linear assignment, index rule

A *visiting order* (a schedule) assigns each unit of a fleet to one slot of the crew's time. *The crew's cycle* is the span the order fills, here fifty weeks at one coil a week. The *cost of waiting* of a unit over a delay is the expected consequence of the events a visit at the start of the delay would prevent and a visit at its end would not. A *sunk consequence* is the consequence of an event whose probability no action changes; it cancels from every comparison of actions. The *backlog* is the set of units already past their functional limit, whose crossing is behind them. A *linear assignment* chooses a one-to-one matching of units to slots that minimizes a sum of costs, one cost per unit at its slot. An *index rule* (a greedy ranking) visits the units in decreasing order of one number per unit. Chapter 17 defined the expected value of perfect information (EVPI); the box does not repeat it.

Key idea

A visit is worth what it prevents: the chance that the coil crosses its functional limit after today and before the crew arrives, times the consequence of that crossing. A coil already past the limit has crossed, and its visit prevents nothing the cost schedule prices. The expected consequence of an order is then a sum over the coils, one term per coil at its week, and the best of the 50! orders is a linear assignment, solved exactly.

20.1.1 The value of a visit

Chapter 2 wrote the expected consequence of an action as its certain cost plus, for each event it leaves, the consequence of the event times the event's predictive probability, equation (2.23). Here the action is an order $\mathcal{V}$ of the fifty visits, in which coil i is visited in week v_i. The order costs nothing for certain: the crew is paid whatever it visits. Each coil carries one event, *crossed its functional limit after today and before its visit*, and the cost schedule gives that event the consequence FLEET_CROSS, written FC, the site's share of a shortfall and an emergency clean. A visit is a clean that holds the coil above the limit for the rest of the cycle, which step 5 checks, so after its visit a coil cannot cross.

The event is a crossing, the first passage of a falling path of Chapter 2's trend form, the second line of equation (2.21), taken between two dates instead of from the start of time: today and the visit. Write t_i^* for the time after today at which coil i's log-derate falls to g_F, with two conventions that step 0 displays: a coil already at or below the limit has a negative t_i^*, its crossing behind it, and a coil above the limit whose line does not fall has $t_i^* = \infty$. The second line of the trend form gives $\Pr(t_i^* \le \Delta \mid \mathbf{y})$, the probability that the coil has crossed by a delay Δ, and it counts a coil already past the limit as crossed. Waiting is priced instead by the increment over the delay,

$$\Pr(0 < t_i^* \le \Delta \mid \mathbf{y}) = \Pr(t_i^* \le \Delta \mid \mathbf{y}) - \Pr(t_i^* \le 0 \mid \mathbf{y}),$$

and by the first convention the term subtracted, $\Pr(t_i^* \le 0 \mid \mathbf{y})$, is the coil's probability of being at or below the limit today. No order changes that probability, so it enters every order's expected consequence alike and drops out of every comparison between them. A list ordered by how low the coils read mixes the two parts, and the larger part, for the lowest coils, is the one no visit can change.

> **Principle 20.1.** *Rank units by what waiting would cost: the probability of a crossing that a visit can still prevent, times its consequence. Never rank them by how bad they look, which mixes the crossings still ahead with the ones already behind.*

20.1.2 The fleet as a graph (step 0)

The graph is Chapter 9's (Section 9.1.2, Figure 9.2): a fleet node carrying the design's means μ_c and μ_γ and the logarithms ℓ_c and ℓ_γ of its spreads, and fifty coil nodes, each carrying its commissioning value $g_{i,0}$ and its fouling rate γ_i, the pair $\mathbf{x}_i = (g_{i,0},\ \gamma_i)^\mathsf{T}$, with its annual tests hanging on it and two tie rows to the fleet node. The state is Chapter 9's (9.2), 104 entries, and the chapter adds no entry and no row to it.

The question nodes. Each coil's path is the straight line (9.1). With a_i the coil's age today, its log-derate t_{ah} years from now and the time at which it reaches the functional limit are

$$g_i(a_i + t_{\text{ah}}) = g_{i,0} + \gamma_i\,(a_i + t_{\text{ah}}), \qquad t_i^* = \begin{cases} \dfrac{g_F - g_i(a_i)}{\gamma_i} > 0, & g_i(a_i) > g_F,\ \gamma_i < 0, \\ \infty, & g_i(a_i) > g_F,\ \gamma_i \ge 0, \\ -1\,\text{yr}, & g_i(a_i) \le g_F, \end{cases} \tag{20.1}$$

a derived quantity of the coil's two entries. The first case is the coil above the limit and falling, which reaches it at that ratio; the second is the coil above the limit whose line does not fall, which never reaches it; the third is the coil already at or below it, whose crossing is behind it. The value -1 of the third case is a convention: only its sign enters any event of the chapter, and the ratio of the first case would be positive there for a rising line, which would wrongly count a coil climbing back through the limit as a crossing ahead. The fifty crossing times are the question nodes of the graph. The decision node is the order $\mathcal{V}$, fed by the crew's one visit a week and by the consequence FC. Figure 20.1 draws them: the fleet node at the top, three coil nodes and one dashed box for the other coils, each with its tests (dashed) and its tie rows (solid) to the fleet node; below each coil its question node t_i^* (red, with a dotted edge, since it is computed from the coil's entries and read by no sensor); and at the bottom the decision node, a box fed by every question node, by the crew and by FC.

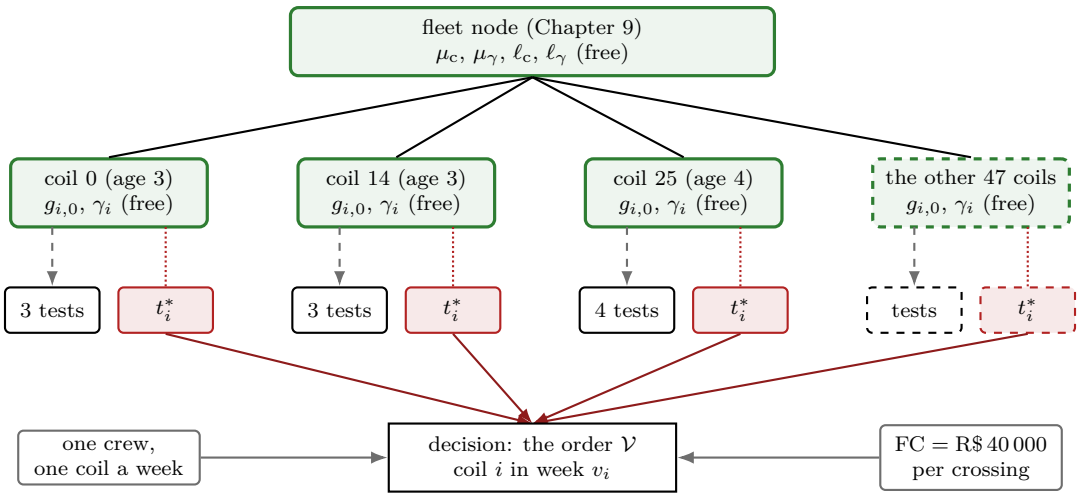


Figure 20.1: The fleet of Chapter 9 with this chapter's question and decision nodes. Free (green): the fleet node's four entries and each coil's commissioning value and rate (three coils drawn, the other 47 as a dashed box). Solid black: the tie rows; dashed gray: each coil's annual condition tests. Red, on dotted edges: each coil's crossing time t_i^* of equation (20.1), a derived quantity. Box: the decision node, the order of visits, fed by the crew and the consequence of a crossing. No row is added; the order couples the coils only through the weeks.

The stated inputs. The record is Chapter 9's: 192 annual condition tests of g, each with standard deviation 0.04, of fifty coils of the hall's design, rated 7 kW/K. The functional limit follows Part I's convention, against rating: the hall of Part I stops doing its job when its conductance falls to $G_F = Q/(T_{\max} - T_c) = 30/(24 - 18) = 5.0$ kW/K on its design value $G_{\text{ref}} = 6$ kW/K, so $h_F = 0.8333$ and $g_F = \ln h_F = -0.18232$ on each coil's log-derate against its rating, 5.8333 kW/K on a 7 kW/K coil; step 2 names it as a choice. The crew visits one coil a week (CREW_PER_WEEK $= 1$), so an order of the fifty coils fills fifty weeks, and week v is visited at

$$t_{\text{ah},v} = \frac{v - 1}{52} \text{ yr}, \qquad v = 1, \ldots, 50, \tag{20.2}$$

after today: week 1 today, in window 37 (Jan Y4), and week 50 in window 48 (Dec Y4). The consequence of a coil crossing after today and before its visit is FC = R$ 40 000. A condition test, INSPECT, costs R$ 1 500 and prices the candidates of step 7, and the clean restores RESTORE_FRACTION $= 0.85$ of the loss since commissioning, which a check of step 5 uses.

Modeling choices at the graph. Five choices that no question owns make this graph. Each coil runs on Chapter 9's straight line, $\sigma_q \to 0$, from today on: this makes the crossing time (20.1) a ratio of the coil's two entries, with a bivariate normal probability under each component of the posterior, and it fails if a coil's fouling departs from its line, by how much Chapter 9's process-noise check (Section 9.4.4) shows for the rates. Today is the end of the record for every coil: this puts every coil's age on one clock, and it fails if the annual tests were taken on different dates, since the record states only whole years of age. The plan is set once: no test is read and no plan revised during the cycle. This buys one assignment, and it fails as a description of a planner who replans each month; step 7 values only tests read before the plan is set, so their values are lower bounds. Each coil's crossing is charged once: a coil past its limit today has crossed, and its consequence is the same under every order, a sunk consequence. The record holds no cleaning history, so the coils past the limit are the record's backlog, and they do not compete for the crew's weeks; a charge for each month below the limit would make them compete, and the cost schedule states none. And a tube leak does not depend on the order. That a clean leaves a coil's leak rate unchanged is a choice, and it rests on Chapter 16, which read condition beside age and resolved no separate effect of the derate once age is counted: it buys a leak risk that is the same under every order and cancels from every comparison, which step 5 shows, and it fails if fouled coils do leak more at the same age, an effect the fleet's failure log cannot separate from age.

20.1.3 The asset manager's questions (step 1)

The asset manager asks four questions of the fleet's posterior, and step 1 writes each as a probability or an expectation over the state (9.2), with the unknown on the left of the bar and the 192 tests and the priors on the right. Table 20.1 lists them in the columns of Chapter 5: the question in the asset manager's words, the first the drill question; the general form of its category from Chapter 2; its mathematical form for the fleet; the category; the computation of Section 2.3.6 that answers it; and the subsection where it is answered. The third question reads Chapter 9's posterior through a threshold. Which test to read before the plan is set asks nothing new of today's evidence; it is the next measurement, Section 20.6. Whether the hall's own coil 1 is cleaned now or at the October shutdown is Chapter 17's question, read on the hall's own path, not here.

20.1.4 The evidence as rows (step 2)

The rows are Chapter 9's, equation (9.3): the 192 test rows, the 100 tie rows that hold each coil's two entries to the design's means with the spreads as their tolerances, and the four priors of the fleet node, 296 rows on the 104 entries. The chapter adds none: every answer is a quantity of the posterior those rows give.

Which row acts on which entry. The Jacobian has Chapter 9's arrow pattern (9.5), which is not recounted: each coil's test block touches only its own $\mathbf{x}_i$, and the coils meet only in the four columns of the fleet node. The second line of the pattern is the question's. Each coil's margin above the limit today and at a time t_{ah} ahead,

$$f_{1,i} = g_i(a_i) - g_F = g_{i,0} + \gamma_i\, a_i - g_F, \qquad f_{2,i}(t_{\mathrm{ah}}) = g_i(a_i + t_{\mathrm{ah}}) - g_F = f_{1,i} + \gamma_i\, t_{\mathrm{ah}},$$
$$\mathbf{g}_{1,i} = \frac{\partial f_{1,i}}{\partial \mathbf{x}_i} = \begin{bmatrix} 1 \\ a_i \end{bmatrix}, \qquad \mathbf{g}_{2,i}(t_{\mathrm{ah}}) = \frac{\partial f_{2,i}}{\partial \mathbf{x}_i} = \begin{bmatrix} 1 \\ a_i + t_{\mathrm{ah}} \end{bmatrix}, \tag{20.3}$$

has its gradient on coil i's two entries alone, every other entry of the 104 zero. Coil 0's gradients are $(1,\ 3)$ and $(1,\ 3 + t_{\mathrm{ah}})$, coil 15's $(1,\ 12)$ and $(1,\ 12 + t_{\mathrm{ah}})$. The fifty questions are block

Table 20.1: The questions asked of the fleet at the end of the record. $\mathbf{y}$ is the 192 tests and the four prior rows; $P_i(t_{\text{ah}}) = \Pr(0 < t_i^* \le t_{\text{ah}} \mid \mathbf{y})$ is coil i's probability of crossing after today and by t_{ah}; v_i its week in the order $\mathcal{V}$ and t_{ah,v_i} the visit time (20.2); $\mathcal{V}_{\text{dash}}$ the dashboard's worst-looking-first order; FC = R$ 40 000.

Asset manager's question	General form	Mathematical form	Category	Computation	Answered in
Fifty units, one crew; which one does the crew visit next?	expected consequence over actions, (2.23); how firm a ranking is, (2.24), second line	$\mathcal{V}^* = \arg\min_{\mathcal{V}} \text{FC} \sum_i P_i(t_{\text{ah},v_i})$ and its week 1; $\Pr(L(\mathcal{V}^*, \mathbf{z}) < L(\mathcal{V}', \mathbf{z}) \mid \mathbf{y})$	design	expected consequence	20.3.1
What does it cost to make each coil wait, and for how long does that hold?	expected consequence of one action, (2.23); first passage, (2.21), second line	$\mathcal{W}_i(\Delta) = \text{FC}\, P_i(\Delta)$ for Δ of 1, 13 and 50 weeks	prediction under a threshold; design	trend; expected consequence	20.3.2
Which coils are already past their functional limit?	$\Pr(f(\mathbf{z}) \ge f^* \mid \mathbf{y})$, (2.12)	$\Pr(g_i(a_i) \le g_F \mid \mathbf{y})$	prediction under a threshold	posterior	20.3.3
What does the year's plan cost against the dashboard's order, and what would knowing every coil's condition be worth?	expected consequence, (2.23); perfect information, (17.9)	$\mathcal{R}(\mathcal{V}^*)$, $\mathcal{R}(\mathcal{V}_{\text{dash}})$, $\Pr(L(\mathcal{V}^*, \mathbf{z}) < L(\mathcal{V}_{\text{dash}}, \mathbf{z}) \mid \mathbf{y})$; $\text{EVPI} = \mathcal{R}(\mathcal{V}^*) - \mathbb{E}[\min_{\mathcal{V}} L(\mathcal{V}, \mathbf{z}) \mid \mathbf{y}]$	estimation; design	expected consequence	20.3.4

diagonal and touch no fleet entry: given the fleet node, they are independent. An order of visits couples them only through the weeks, never through the state.

The loss of an order. An order $\mathcal{V}$ assigns each coil a week v_i, one coil to each week, one of $50! = 3.041 \times 10^{64}$ orders. Its loss at the state $\mathbf{z}$ counts the coils that cross before their visit,

$$L(\mathcal{V}, \mathbf{z}) = \text{FC} \sum_{i=0}^{49} \mathbf{1}\big[0 < t_i^* \le t_{\text{ah},v_i}\big], \tag{20.4}$$

with t_i^* the crossing time (20.1) and t_{ah,v_i} the visit time (20.2). Because $f_{1,i}$ and $f_{2,i}$ are linear in $\mathbf{x}_i$, a straight line crosses g_F in $(0, t_{\text{ah}}]$ exactly when $f_{1,i} > 0$ and $f_{2,i}(t_{\text{ah}}) \le 0$: then $\gamma_i\, t_{\text{ah}} \le -f_{1,i} < 0$, so the line falls and passes the limit once, and a line that crosses in $(0, t_{\text{ah}}]$ stands above the limit today and at or below it at t_{ah}. The indicator of equation (20.4) is therefore that joint event of the two margins.

Modeling choices at the rows. Three choices enter here. The functional limit is $g_F = \ln(5/6)$ against rating, Part I's convention: this gives every coil of the design one limit, the level at which the hall of Part I could no longer hold its temperature, and it fails if a site's duty differs from that hall's; step 5 shows that the stated limit can change the first visit. The visit times are $t_{\text{ah},v} = (v - 1)/52$: the coil in week 1 is visited today and cannot cross first, and a visit later in its week would add a few days to each wait. And FC is the only cost that enters: it multiplies every term of the loss (20.4), so it sets the size of the expected consequence and never the order, which step 5 checks.

20.2 The posterior of the state (step 3)

All four questions of Table 20.1 are read from one posterior, Chapter 9's posterior of the fleet, recomputed from the same rows and the same grid, with its spreads' medians $\tau_{\mathrm{c}} = 0.03153$ and $\tau_{\gamma} = 0.01199\,\mathrm{yr}^{-1}$ and coil 0's rate $-0.02861\,\mathrm{yr}^{-1}$, band $[-0.04471,\ -0.01207]\,\mathrm{yr}^{-1}$, as Chapter 9 reported them.

20.2.1 The expected consequence down to one coil's crossing

The quantity owed is the expected consequence of an order, the posterior expectation of the loss (20.4). The expectation of a sum is the sum of the expectations, and the expectation of an indicator is the probability of its event, so

$$\mathcal{R}(\mathcal{V}) = \mathbb{E}\big[L(\mathcal{V}, \mathbf{z}) \mid \mathbf{y}\big] = \mathrm{FC} \sum_{i=0}^{49} \Pr\big(0 < t_i^{*} \le t_{\mathrm{ah},v_i} \mid \mathbf{y}\big) = \mathrm{FC} \sum_{i=0}^{49} P_i(t_{\mathrm{ah},v_i}), \tag{20.5}$$

one term per coil at its week, each a crossing probability $P_i(t_{\mathrm{ah}})$ of one coil. Each term is derived down to Chapter 9's posterior in turn.

The mixture over the spreads. The posterior of the 102 entries other than the spreads is the mixture (9.11) of exact normal conditionals over the nodes $\boldsymbol{\ell}_m$ of the spreads' grid, with the weights ϖ_m. A probability of the state is the same mixture of the probabilities under each component,

$$P_i(t_{\mathrm{ah}}) = \sum_m \varpi_m\, P_i(t_{\mathrm{ah}} \mid \boldsymbol{\ell}_m), \qquad \Pr\big(g_i(a_i) \le g_F \mid \mathbf{y}\big) = \sum_m \varpi_m\, \Pr\big(f_{1,i} \le 0 \mid \boldsymbol{\ell}_m, \mathbf{y}\big). \tag{20.6}$$

One coil at one node. At node m, with the means $\boldsymbol{\mu}$ integrated through equation (9.8), coil i's pair $\mathbf{x}_i$ is normal with Chapter 9's mean $\hat{\mathbf{x}}_i = \boldsymbol{\Lambda}_i^{-1}(\mathbf{q}_i + \boldsymbol{\Omega}^{-1}\hat{\boldsymbol{\mu}})$ and covariance $\boldsymbol{\Sigma}_i$, the conditional (9.9) widened by the uncertainty of the means. Write its entries

$$\boldsymbol{\Sigma}_i = \begin{bmatrix} \Sigma_{gg} & \Sigma_{g\gamma} \\ \Sigma_{g\gamma} & \Sigma_{\gamma\gamma} \end{bmatrix},$$

in the order $(g_{i,0},\ \gamma_i)$, all at node m. The two margins of equation (20.3) are linear in $\mathbf{x}_i$, so the pair $(f_{1,i},\ f_{2,i}(t_{\mathrm{ah}}))$ is bivariate normal, its means the gradients against $\hat{\mathbf{x}}_i$ and its covariance the gradients through $\boldsymbol{\Sigma}_i$. Since $\mathbf{g}_{2,i}(t_{\mathrm{ah}}) = \mathbf{g}_{1,i} + t_{\mathrm{ah}}\,(0,\ 1)^{\mathsf{T}}$, every entry of that covariance follows from $\mathrm{Var}(f_{1,i})$ and the covariance of the margin today with the rate:

$$\begin{aligned}
\bar{f}_{1,i} &= \mathbf{g}_{1,i}^{\mathsf{T}}\hat{\mathbf{x}}_i - g_F, \qquad \bar{f}_{2,i}(t_{\mathrm{ah}}) = \bar{f}_{1,i} + t_{\mathrm{ah}}\,\hat{\gamma}_i, \\
\mathrm{Var}(f_{1,i}) &= \mathbf{g}_{1,i}^{\mathsf{T}}\boldsymbol{\Sigma}_i\,\mathbf{g}_{1,i} = \Sigma_{gg} + 2a_i\Sigma_{g\gamma} + a_i^2\Sigma_{\gamma\gamma}, \qquad \mathrm{Cov}(f_{1,i}, \gamma_i) = \Sigma_{g\gamma} + a_i\Sigma_{\gamma\gamma}, \\
\mathrm{Var}(f_{2,i}(t_{\mathrm{ah}})) &= \mathbf{g}_{2,i}^{\mathsf{T}}\boldsymbol{\Sigma}_i\,\mathbf{g}_{2,i} = \mathrm{Var}(f_{1,i}) + 2t_{\mathrm{ah}}\,\mathrm{Cov}(f_{1,i}, \gamma_i) + t_{\mathrm{ah}}^2\Sigma_{\gamma\gamma}, \\
\mathrm{Cov}(f_{1,i}, f_{2,i}(t_{\mathrm{ah}})) &= \mathbf{g}_{1,i}^{\mathsf{T}}\boldsymbol{\Sigma}_i\,\mathbf{g}_{2,i} = \mathrm{Var}(f_{1,i}) + t_{\mathrm{ah}}\,\mathrm{Cov}(f_{1,i}, \gamma_i),
\end{aligned} \tag{20.7}$$

with $\hat{\gamma}_i$ the second entry of $\hat{\mathbf{x}}_i$, each moment conditional on $\boldsymbol{\ell}_m$ and $\mathbf{y}$.

The crossing at one node. The event of the loss (20.4) is $f_{1,i} > 0$ and $f_{2,i}(t_{\mathrm{ah}}) \le 0$. Its probability is that of $f_{2,i}(t_{\mathrm{ah}}) \le 0$ less that of both margins at or below zero, so with the standardized levels $u_1 = -\bar{f}_{1,i}/\sqrt{\mathrm{Var}(f_{1,i})}$ and $u_2 = -\bar{f}_{2,i}(t_{\mathrm{ah}})/\sqrt{\mathrm{Var}(f_{2,i}(t_{\mathrm{ah}}))}$ and the correlation $\varrho = \mathrm{Cov}(f_{1,i}, f_{2,i}(t_{\mathrm{ah}}))/\sqrt{\mathrm{Var}(f_{1,i})\,\mathrm{Var}(f_{2,i}(t_{\mathrm{ah}}))}$ of the two margins,

$$P_i(t_{\mathrm{ah}} \mid \boldsymbol{\ell}_m) = \Phi(u_2) - \Phi_2(u_1, u_2; \varrho), \qquad \Pr\big(f_{1,i} \le 0 \mid \boldsymbol{\ell}_m, \mathbf{y}\big) = \Phi(u_1), \tag{20.8}$$

with Φ_2 the distribution function of two standard normals of correlation ϱ of equation (2.22), evaluated by Owen's T function, a standard identity. The first is the increment of Section 20.1.1, the second line of equation (2.21) at t_{ah} less its value today; the second is the coil's probability of being past the limit. Both are closed forms in the node's $\hat{\mathbf{x}}_i$ and $\boldsymbol{\Sigma}_i$, and the mixture (20.6) sums them over the nodes with their weights.

20.2.2 The routes

Chapter 9 showed that the joint normal of this state has no mode in one parametrization and the wrong one in the other (Section 9.2.3), and it reported the mixture over the spreads instead. This chapter reads the same posterior by the same routes, listed in Table 20.2: each row is a route, with its nodes and the quantities it gives.

- *The mixture over the spreads' grid* is the reported route. Equations (20.8) and (20.6) are summed in closed form over the 19 985 nodes, of the grid's 86 000, that hold all but 10^{-10} of its mass, with the weights renormalized.
- *The Laplace route* is Chapter 9's normal on the spreads' marginal, with its mode at $\tau_{\text{c}} = 0.03356$ and $\tau_\gamma = 0.01165$ and the 15×15 Gauss–Hermite nodes on its axes in place of the grid's: the book's Laplace order applied where it can work, to the two entries the linear rows do not fix.
- *The plug-in spreads* (empirical Bayes) put all the weight on one node, the grid node at the marginal's mode, $\tau_{\text{c}} = 0.03300$ and $\tau_\gamma = 0.01162$, beside the Laplace normal's mode at 0.03356 and 0.01165, which is fitted off the grid.
- *The exact draws by composition* are Chapter 9's (Section 9.2.5), 200 000 draws from a generator seeded with 9, which reproduce coil 0's median rate $-0.02863\,\text{yr}^{-1}$. Each draw $\mathbf{z}^{(j)}$ gives every coil's crossing time (20.1), and with it the week $v_i^* = \lceil 52\, t_i^* \rceil$ in which a coil ahead of the limit crosses: a coil visited in week v is crossed first exactly when $1 \le v_i^* \le v - 1$. The draws give every joint quantity of the fifty coils: the loss of an order on each draw, the probability that one order beats another, and the best order with the state known.

Table 20.2: The routes to the chapter's answers, all on Chapter 9's posterior.

Route	Nodes or draws	Gives
Mixture over the spreads' grid (reported)	19 985 of 86 000 grid nodes, renormalized	$P_i(t_{\text{ah}})$, Pr(past) and $\mathcal{R}(\mathcal{V})$ in closed form
Normal on the spreads' marginal (Laplace)	15×15 Gauss–Hermite nodes	the same, beside the reported values
Spreads fixed at the grid node of the marginal's mode (plug-in)	one node	the same, without the spreads' uncertainty
Composition draws (exact draws)	200 000 draws, seed 9	the loss of an order on each draw; the probability one order beats another; perfect information

The expansion of equation (2.24) at the most probable state is not among the routes. The loss (20.4) is a step in the state, and Section 20.4.3 shows what that expansion returns.

20.2.3 Two coils' blocks as numbers

The covariance of one coil's block is shown at the plug-in spreads for the two coils the answers turn on. Coil 25, four years old, has the mean $(g_{25,0},\ \gamma_{25}) = (-0.03187,\ -0.04136\,\text{yr}^{-1})$ and

$$\boldsymbol{\Sigma}_{25} = \begin{bmatrix} 5.6154 \times 10^{-4} & -1.3843 \times 10^{-4} \\ -1.3843 \times 10^{-4} & 7.2952 \times 10^{-5} \end{bmatrix},$$

standard deviations 0.02370 and 0.00854 and correlation −0.6839. By equation (20.7) its margin today has the mean $\bar{f}_1 = -0.01500$ and the variance 6.2135×10^{-4}, a standard deviation of 0.02493, so it is past the limit with probability 0.7263 at this node, and $\mathrm{Cov}(f_1, \gamma) = \Sigma_{g\gamma} + 4\,\Sigma_{\gamma\gamma} = 1.5338 \times 10^{-4}$. The margin's variance grows with the time ahead as

$$\mathrm{Var}(f_2(t_{\mathrm{ah}})) = 6.2135 \times 10^{-4} + 2t_{\mathrm{ah}}\,(1.5338 \times 10^{-4}) + t_{\mathrm{ah}}^2\,(7.2952 \times 10^{-5}),$$

6.2728×10^{-4} over one week ($t_{\mathrm{ah}} = 0.0192$ yr), with the mean −0.01579 and a correlation of 0.999990 between the two margins, and 9.8376×10^{-4} over the cycle ($t_{\mathrm{ah}} = 0.9615$ yr), with the mean −0.05477 and a correlation of 0.983373.

Coil 14, three years old, has the mean $(-0.01650,\ -0.04195\,\mathrm{yr}^{-1})$ and

$$\boldsymbol{\Sigma}_{14} = \begin{bmatrix} 5.6138 \times 10^{-4} & -1.3709 \times 10^{-4} \\ -1.3709 \times 10^{-4} & 9.6939 \times 10^{-5} \end{bmatrix},$$

standard deviations 0.02369 and 0.00985 and correlation −0.5877. Its margin today has the mean 0.03998 and the standard deviation 0.02472, past the limit with probability 0.0529 at this node, and $\mathrm{Cov}(f_1, \gamma) = 1.5372 \times 10^{-4}$. Over the cycle the margin's variance is 9.9652×10^{-4} and its mean −0.00035: coil 14's margin reaches zero at the end of the cycle.

The signs. The covariance $\Sigma_{g\gamma}$ is negative for both coils, for Chapter 9's reason: no coil is tested at age 0, so a steeper line must start higher to pass through the same tests. The cross term of the budget, $\mathrm{Cov}(f_1, \gamma) = \Sigma_{g\gamma} + a_i\Sigma_{\gamma\gamma}$, is nevertheless positive, because at today's age the term $a_i\Sigma_{\gamma\gamma}$ outweighs $\Sigma_{g\gamma}$: the tests lie before today, so a faster-fouling line also ends lower today. A coil low today is then also falling fast, and the variance of the margin ahead grows faster than the rate's term $t_{\mathrm{ah}}^2\Sigma_{\gamma\gamma}$ alone would make it. Over one week the two margins are correlated 0.999990, so a week's crossing probability is the small difference of two nearly equal tails. The closed form (20.8) keeps it exactly; the draws estimate it with the standard error that step 5 states.

20.2.4 The crossing quantities coil by coil

Table 20.3 gives the question's quantities for the 24 coils that carry any cost or any doubt, ordered by the probability of a crossing over the whole cycle. Each row is a coil; the columns give its age and its latest test, the posterior median of g today with its 5–95 % band, its probability of being past the limit, its crossing probabilities $P_i(t_{\mathrm{ah}})$ over one week, thirteen weeks and the cycle, the Laplace values beside the reported ones for two of them, and, from the draws, the median crossing time over the draws in which the coil is not yet past. The other 26 coils have a probability of being past above 0.98 or below 0.02 and a cycle's crossing probability below 0.001. Read down the columns, the coils most likely to cross in the cycle (14, 33, 1) stand above the limit today, while the coil most likely to cross in the next week, coil 25, is probably past already. All fifty coils together expect 2.9929 crossings in the cycle if none is visited, and the 17 coils with a cycle's probability of at least 0.01 hold 2.9687 of them.

20.3 The answers (step 4)

Every answer is read from the posterior of step 3: the mixture over the spreads, equation (20.6), is the reported value, with the Laplace route beside it, and the composition draws give every quantity that involves several coils at once.

20.3.1 The next visit

Asset manager's question. Fifty units, one crew; which one does the crew visit next? The question asks for the first week of the best order, and how firm that choice is.

Table 20.3: The question's quantities for the 24 coils that carry any cost or doubt, from the mixture over the spreads (reported), with the Laplace route beside two columns; $P_i(t_{\mathrm{ah}}) = \Pr(0 < t_i^* \le t_{\mathrm{ah}} \mid \mathbf{y})$, the cost of waiting over t_{ah} divided by FC. The last column is from the draws, over the draws with the coil not yet past. Bands are 5–95 %.

Coil [–]	Age [yr]	Latest test [–]	g today, median [–]	band [–]	P(past) [–]	Laplace [–]	P_i, 1 wk [–]	Laplace [–]	P_i, 13 wk [–]	P_i, cycle [–]	median t^*, if ahead [yr]
14	3	−0.2099	−0.1430	−0.1851, −0.1019	0.0619	0.0591	0.00497	0.00473	0.0842	0.4561	0.988
33	4	−0.1308	−0.1420	−0.1836, −0.1008	0.0552	0.0528	0.00395	0.00383	0.0655	0.3689	1.193
1	5	−0.1639	−0.1462	−0.1871, −0.1052	0.0732	0.0728	0.00414	0.00416	0.0645	0.3309	1.311
38	7	−0.1166	−0.1717	−0.2109, −0.1322	0.3280	0.3270	0.00714	0.00708	0.0928	0.3241	1.005
46	3	−0.1284	−0.1363	−0.1777, −0.0952	0.0337	0.0349	0.00269	0.00268	0.0474	0.3081	1.391
25	4	−0.2266	−0.1980	−0.2402, −0.1566	0.7324	0.7257	0.00947	0.00947	0.1031	0.2295	0.352
2	11	−0.1273	−0.1841	−0.2203, −0.1478	0.5316	0.5349	0.00556	0.00553	0.0694	0.2272	1.002
49	5	−0.2408	−0.1983	−0.2396, −0.1574	0.7388	0.7338	0.00865	0.00874	0.0954	0.2197	0.377
16	3	−0.1045	−0.1214	−0.1626, −0.0802	0.0076	0.0079	0.00077	0.00077	0.0156	0.1563	1.986
40	4	−0.1347	−0.1109	−0.1522, −0.0696	0.0023	0.0022	0.00028	0.00027	0.0065	0.0929	2.283
31	8	−0.1221	−0.1293	−0.1679, −0.0905	0.0118	0.0127	0.00067	0.00072	0.0112	0.0789	3.149
36	9	−0.2419	−0.2190	−0.2565, −0.1815	0.9461	0.9467	0.00173	0.00170	0.0186	0.0430	0.405
0	3	−0.0639	−0.0979	−0.1389, −0.0567	0.0004	0.0004	0.00006	0.00005	0.0015	0.0370	2.951
29	8	−0.2714	−0.2230	−0.2615, −0.1846	0.9594	0.9621	0.00172	0.00164	0.0175	0.0358	0.300
34	2	−0.1273	−0.0887	−0.1300, −0.0482	0.0001	0.0002	0.00002	0.00002	0.0008	0.0337	2.605
41	11	−0.2336	−0.2282	−0.2644, −0.1920	0.9814	0.9820	0.00062	0.00060	0.0066	0.0151	0.398
19	3	−0.1151	−0.0801	−0.1211, −0.0390	0.0000	0.0000	0.00001	0.00000	0.0002	0.0116	3.363
47	2	−0.0683	−0.0702	−0.1110, −0.0296	0.0000	0.0000	0.00000	0.00000	0.0000	0.0051	3.619
32	8	−0.2316	−0.2425	−0.2808, −0.2041	0.9950	0.9947	0.00021	0.00021	0.0021	0.0043	0.346
12	10	−0.1848	−0.2407	−0.2774, −0.2039	0.9954	0.9953	0.00017	0.00017	0.0018	0.0038	0.344
11	2	−0.0525	−0.0645	−0.1052, −0.0239	0.0000	0.0000	0.00000	0.00000	0.0000	0.0026	3.978
10	5	−0.1066	−0.0739	−0.1150, −0.0321	0.0000	0.0000	0.00000	0.00000	0.0000	0.0023	4.953
37	2	−0.0180	−0.0634	−0.1043, −0.0227	0.0000	0.0000	0.00000	0.00000	0.0000	0.0018	4.355
43	2	−0.0384	−0.0595	−0.1002, −0.0188	0.0000	0.0000	0.00000	0.00000	0.0000	0.0014	4.325

General form. A choice among actions is the expected consequence of each and the action that minimizes it, equation (2.23),

$$\mathcal{R}(a) = \mathbb{E}\big[L(a, \mathbf{z}) \mid \mathbf{y}\big], \qquad a^* = \arg\min_a \mathcal{R}(a),$$

and how firm the choice is, the probability that the chosen action costs less than another at the true state, is the second line of equation (2.24), $\Pr(L(a, \mathbf{z}) < L(a', \mathbf{z}) \mid \mathbf{y})$.

Specialized form. The question is equation (2.23) with the action a the order $\mathcal{V}$, no certain cost, the events e the fifty crossings before a visit, each with $C_e = \mathrm{FC}$ and the predictive probability $P_i(t_{\mathrm{ah},v_i})$ of equation (20.6): that is the expected consequence (20.5). With $X_{iv} = 1$ when coil i is visited in week v and 0 otherwise,

$$\mathcal{R}(\mathcal{V}) = \mathrm{FC} \sum_{i=0}^{49} \sum_{v=1}^{50} X_{iv}\, P_i(t_{\mathrm{ah},v}), \qquad \sum_v X_{iv} = 1, \quad \sum_i X_{iv} = 1, \tag{20.9}$$

linear in the assignment matrix X, so its minimum over all 50! orders is a linear assignment problem, solved exactly by the standard algorithm for it (the Hungarian method) on the 50×50 table of FC $P_i(t_{\mathrm{ah},v})$. The answer is the coil in week 1 of the minimizing order $\mathcal{V}^*$. One choice enters here: the answer is read as the first week of the best order for the whole cycle, not as the top of a ranking of the coils; the drill's index rule, a ranking by the cost of waiting, is set beside it in Section 20.3.2. The firmness of the first choice comes from the draws, by setting the best order against the same order with its week-1 coil swapped with another coil, and reading on each draw which of the two losses (20.4) is smaller. The expansion in the second line of equation (2.24) cannot give it: the gradient of a step loss is zero wherever it exists (Section 20.4.3).

Building blocks. The crossing probabilities $P_i(t_{\mathrm{ah},v})$ of equations (20.8) and (20.6) on the fifty visit times (20.2); the assignment (20.9); the composition draws and the loss (20.4) on each draw.

The plan. The crew visits coil 25 first. The plan's weeks 1 to 12 are coils 25, 49, 38, 14, 2, 33, 1, 46, 16, 36, 29 and 31, and its expected consequence over the cycle is $\mathcal{R}(\mathcal{V}^*) = \mathrm{R\$}\,7\,907.4$. Every route puts coil 25 in week 1: the Laplace route's own best order costs R$ 7 822.8 under it, and the reported plan R$ 7 828.1; the plug-in spreads give R$ 7 727.9; the draws' own best order costs R$ 7 902.4 under the draws, and the reported plan R$ 7 907.6. Figure 20.2 draws where the cost sits. The horizontal axis is the plan's week, with its window below; each bar is the expected consequence that week's coil carries, $\mathrm{FC}\,P_i(t_{\mathrm{ah},v})$, labeled with the coil, and the row beneath the axis gives the coil the dashboard's order puts in the same week. The bars start at zero in week 1, since a coil visited today cannot cross first, and rise to R$ 1 077 in week 7. Weeks 1 to 8 hold 0.6637 of $\mathcal{R}(\mathcal{V}^*)$ and weeks 1 to 14 hold 0.9541; every coil after week 37 carries under R$ 1, and 8 coils have a cycle's crossing probability below 10^{-6}, so their weeks are interchangeable at no cost. The order matters for its first fourteen weeks or so, windows 37 to 39 (Jan Y4 to Mar Y4).

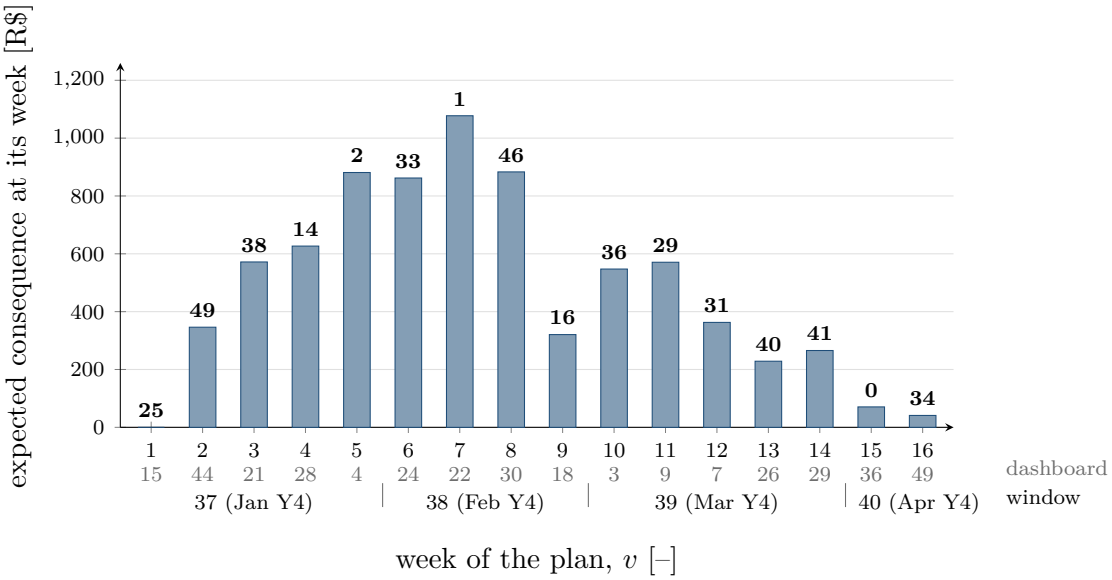


Figure 20.2: The plan's first sixteen weeks: the expected consequence each week's coil carries, $\mathrm{FC}\,P_i(t_{\mathrm{ah},v})$ from the mixture, with the plan's coil above each bar. Beneath the week numbers: the coil the dashboard's worst-looking-first order visits in the same week (gray), and the window of the week. Week 1 carries nothing: its coil is visited today.

How firm the first choice is. Table 20.4 sets the plan against the plan with another coil in week 1 and coil 25 moved to that coil's week. Each row is the coil put first; the columns give that order's expected consequence and its excess over the plan, and, over the draws, the probabilities that the plan's loss is smaller, equal and larger, with the first of these among the draws on which the two losses differ.

The crew goes first to a coil that has probably already crossed: coil 25 is past the limit with probability 0.7324. It also has the largest one-week crossing probability of the fifty, $P_{25}(1\,\mathrm{wk}) = 0.00947$, a cost of waiting of R$ 378.7: what remains of its probability is concentrated in the next weeks. Against its runner-up, coil 49, the choice is close to a coin flip on the draws where the two orders differ (0.5285), and the two agree on 0.9824 of the draws; the choice is

Table 20.4: The plan against the plan with week 1 swapped, from the mixture ($\mathcal{R}$) and the composition draws (the probabilities).

Coil in week 1 instead (coil 25 moves to)	$\mathcal{R}$ [R$]	extra [R$]	P(plan better) [–]	P(equal) [–]	P(plan worse) [–]	P(better), if they differ [–]
coil 49 (week 2)	7 940.0	+32.6	0.0093	0.9824	0.0083	0.5285
coil 38 (week 3)	8 082.5	+175.1	0.0182	0.9678	0.0140	0.5647
coil 14 (week 4)	8 385.0	+477.6	0.0270	0.9580	0.0150	0.6435
coil 15, the dashboard's first (week 46)	16 758.9	+8 851.5	0.2214	0.7786	0.0000	1.0000

worth R$ 32.6. Against the dashboard's coil 15 in week 1 the plan is never worse, on none of the 200 000 draws, and it is better on 0.2214 of them. Coil 0, the hall's coil 1, goes in week 15: it is past with probability 0.000378, its cycle's crossing probability is 0.037014, and it carries R$ 70.56 at its week. This is Chapter 9's register entry kept: coil 0's place is set by the expected consequence of waiting, not by its slope.

20.3.2 The cost of waiting

Asset manager's question. What does it cost to make each coil wait, and for how long does that hold? The drill's computation ranks the units by this cost, and the question asks whether that ranking finds the plan's first visit.

General form. The expected consequence of one action, equation (2.23), here the action of not visiting a unit before a delay Δ, whose event is the first passage of a falling path, the second line of equation (2.21), taken between today and the end of the delay.

Specialized form. Equation (2.23) with the single event "coil i crosses in $(0, \Delta]$", of consequence FC and probability $P_i(\Delta)$ from equation (20.6), gives the cost of waiting

$$\mathcal{W}_i(\Delta) = \mathrm{FC}\, P_i(\Delta) = \mathrm{FC}\, \Pr(0 < t_i^* \le \Delta \mid \mathbf{y}), \tag{20.10}$$

the second line of equation (2.21) at Δ less its value today. An index rule visits the coils in decreasing $\mathcal{W}_i(\Delta)$, and its order is scored by the expected consequence (20.5). One choice enters here: the delay Δ of the ranking. Three are shown: the crew's next slot, one week; a quarter, thirteen weeks; and the cycle, fifty weeks. A delay matched to the next slot prices what a week of waiting costs each coil; a long delay prices the whole year, and it fails when the coils whose risk is now and the coils whose risk is late in the year compete for the first weeks.

Building blocks. The crossing probabilities of equation (20.8), summed by equation (20.6), at Δ of 1, 13 and 50 weeks and at every visit time; the expected consequence (20.5) of each index order.

The index orders. Table 20.5 gives, for each delay, the eight coils with the largest cost of waiting with their costs, the index order's week 1, its expected consequence, and its excess over the plan. Ranked over one week the index starts with coil 25 at R$ 378.7, then coils 49, 38, 2 and 14, and its order costs R$ 7 993.5, R$ 86.1 more than the plan. Ranked over a quarter it also starts with coil 25 and costs R$ 54.7 more. Ranked over the cycle it starts with coil 14, at R$ 18 243.1, and costs R$ 2 035.9 more. The Laplace route beside the mixture gives the one-week probabilities of the plan's first five coils as 0.00947 against 0.00947 for coil 25, 0.00874 against

Table 20.5: The index rule at three delays: the coils ranked by the cost of waiting (20.10), from the mixture, and the expected consequence of the index order.

Delay	First eight coils by $\mathcal{W}_i(\Delta)$, coil: $\mathcal{W}$ [R$]	index week 1 [–]	its $\mathcal{R}$ [R$]	over the plan [R$]
one week	25: 378.7, 49: 346.2, 38: 285.7, 2: 222.3, 14: 198.8, 1: 165.6, 33: 157.8, 46: 107.5	25	7 993.5	+86.1
13 weeks	25: 4 124.2, 49: 3 817.2, 38: 3 713.7, 14: 3 367.9, 2: 2 776.7, 33: 2 621.2, 1: 2 580.6, 46: 1 894.7	25	7 962.1	+54.7
the cycle	14: 18 243.1, 33: 14 756.2, 1: 13 236.7, 38: 12 965.7, 46: 12 322.4, 25: 9 181.0, 2: 9 086.5, 49: 8 788.8	14	9 943.3	+2 035.9

0.00865 for coil 49, 0.00708 against 0.00714 for coil 38, 0.00473 against 0.00497 for coil 14 and 0.00553 against 0.00556 for coil 2.

The cost of waiting per week of delay runs in opposite directions for the two leaders. Coil 25's falls from R$ 378.7 in the first week to R$ 368.1, R$ 357.5 and on to R$ 306.2 in the eighth: its crossing is now or not this year. Coil 14's rises from R$ 198.8 to R$ 208.9, R$ 219.1 and on to R$ 269.7: its crossing is most likely late in the year, with a median of 0.988 yr over the draws in which it is still ahead. The crossing probabilities over the cycle are plotted in Figure 20.3. The horizontal axis is the time ahead t_{ah}, in weeks, and the vertical axis $P_i(t_{\text{ah}})$ for six coils: the mixture as solid lines and the Laplace route as dashed lines of the same color; the dotted verticals mark the weeks in which the stated truth crosses (Section 20.4.6). The slope of a curve at the start is the coil's cost of waiting for the next week, and its height at the right edge its cost over the cycle, and the two rank the coils differently: coil 25's curve is the steepest at the start and flattens, coil 14's starts lower and ends highest. The dashed curves stay within 0.01 of the solid ones, the widest, 9.91×10^{-3}, near the end of coil 14's curve.

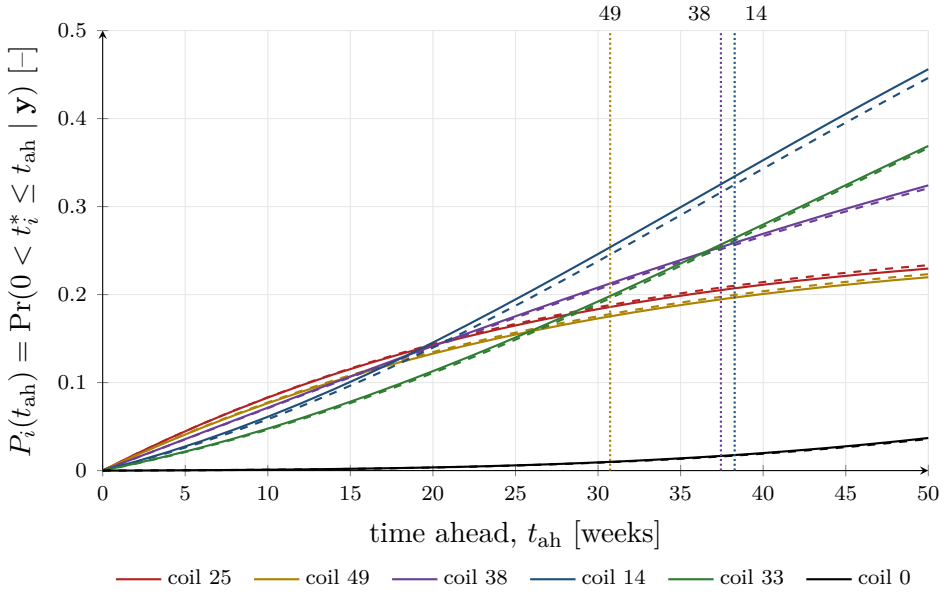


Figure 20.3: The crossing probability $P_i(t_{\text{ah}})$ of six coils over the cycle, equations (20.8) and (20.6): the mixture (solid) and the Laplace route (dashed), drawn every two weeks. The slope at $t_{\text{ah}} = 0$ ranks the coils for the next week, the height at $t_{\text{ah}} = 50$ for the cycle. Dotted: the weeks in which the stated truth crosses for coils 49, 38 and 14, check markers (Section 20.4.6).

The drill's ranking is right when its delay is the crew's next slot. Ranked over a week or a quarter, the index finds the plan's first visit and comes within R$ 86.1 of the exact order. Ranked over the whole cycle, it sends the crew to coil 14, the coil most likely to cross this year,

and costs R$ 2 035.9 more. The plan lets coil 14 wait three weeks, at R$ 627 in its week 4, while coils 25, 49 and 38, whose risk is now, go first.

20.3.3 The coils past the limit

Asset manager's question. Which coils are already past their functional limit? The question asks which of the coils the dashboard reads below the limit have crossed, and which reading above it may have.

General form. A prediction under a threshold is the posterior probability of an event on a derived quantity, equation (2.12),

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

Specialized form. The question is equation (2.12) in its own direction, with $f(\mathbf{z}) = g_i(a_i)$, the log-derate today, and $f^* = g_F$: the requirement is that the coil stands at or above its limit, and the answer is its complement, as Chapter 2 writes an exceedance,

$$\Pr(g_i(a_i) \leq g_F \mid \mathbf{y}) = 1 - \Pr(g_i(a_i) \geq g_F \mid \mathbf{y}),$$

the two events sharing only the boundary, which has probability zero. The event is $f_{1,i} \leq 0$ of equation (20.3). Under each node it is the second closed form of equation (20.8), and over the posterior the second sum of equation (20.6),

$$\Pr(g_i(a_i) \leq g_F \mid \mathbf{y}) = \sum_m \varpi_m\, \Phi(u_{1,i,m}), \qquad u_{1,i,m} = -\frac{\bar{f}_{1,i}}{\sqrt{\mathrm{Var}(f_{1,i})}}\bigg|_{\ell_m}. \tag{20.11}$$

A coil past its limit costs the same under every order: it is the backlog of the modeling choice of step 0, set beside the ranking and not in it. The expected number of coils past is the sum of the fifty probabilities.

Building blocks. The margin's mean and variance of equation (20.7) at each node; the normal distribution function; the weights of the mixture (20.6).

The backlog. By equation (20.11), 23 coils are past the limit with probability above 0.5, 19 above 0.95, and 23 are past with probability below 0.05. The expected number past is 22.4533, and the Laplace route gives 22.4440. The dashboard and the posterior differ on two coils. Coil 14 reads -0.2099, past the limit, yet is past with probability 0.0619 (Laplace 0.0591): its pooled g today is -0.1430, band $[-0.1851,\ -0.1019]$, because its last test is low against a line the fleet pulls in, the chart's slope $-0.0911\,\mathrm{yr}^{-1}$ against the pooled rate $-0.04256\,\mathrm{yr}^{-1}$ of Chapter 9. Coil 2 reads -0.1273, above the limit, yet is past with probability 0.5316 (Laplace 0.5349). The worst-looking coil, coil 15, is past with probability 1.000000, with a median g today of -0.7307, and its cycle's crossing probability is 0: its crossing is behind it.

Each coil's latest test is set against its probability of being past in Figure 20.4. The horizontal axis is the latest test, the vertical axis $\Pr(g_i(a_i) \leq g_F \mid \mathbf{y})$; the dashed vertical line is g_F and the dashed horizontal line the probability 0.5, so that a coil the dashboard and the posterior agree on sits in the lower right (above the limit, ahead) or the upper left (below it, past). The plan's first eight coils are red, the rest gray; coils 14 and 2 are labeled, and the two doubtful coils that are truly past, a check marker of Section 20.4.6, carry a dotted ring. The coils fall in a band that climbs from zero to one as the latest test crosses the limit, and the band is wide: the latest test of a coil is one reading, and its line, pulled by the fleet, can sit well away from it.

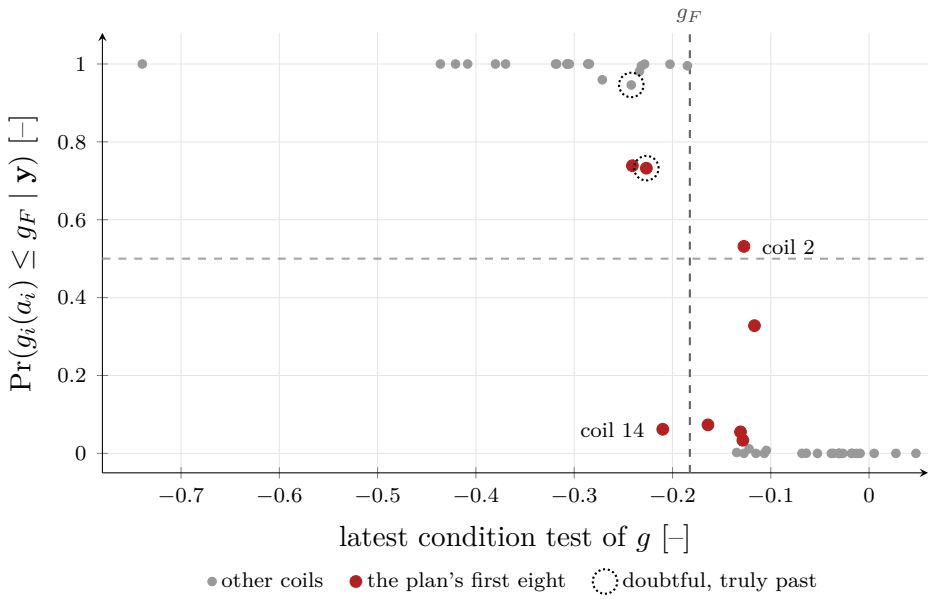


Figure 20.4: Each coil's latest condition test against its posterior probability of being past the functional limit, equation (20.11), from the mixture. Dashed: the limit $g_F = -0.18232$ and the probability 0.5. Red: the plan's first eight coils; gray: the other 42. Dotted rings: the doubtful coils ($0.05 < P < 0.95$) that the stated truth puts past the limit, check markers (Section 20.4.6).

The dashboard's first twelve coils, 15, 44, 21, 28, 4, 24, 22, 30, 18, 3, 9 and 7, are all past the limit with probability 1.000. The worst-looking-first order spends the crew's first twelve weeks where a visit can prevent no crossing.

20.3.4 The year's plan against the naive orders

Asset manager's question. What does the year's plan cost against the dashboard's order, and what would knowing every coil's condition be worth? The question sets each naive order's expected consequence beside the plan's, and asks how much of the plan's own expected consequence is the price of not knowing the coils' state.

General form. The expected consequence of each action, equation (2.23), and the expected value of perfect information in Chapter 17's general form (17.9): the expected loss of the action chosen now less the expected loss of the best action chosen with the state known, never negative, and the most any measurement taken before the action can be worth.

Specialized form. Each order is scored by equation (2.23) in its form (20.5), under the same posterior and on every route. For perfect information the actions are the orders, $a = \mathcal{V}$ and $a^* = \mathcal{V}^*$, and equation (17.9) becomes the fleet's form,

$$\text{EVPI} = \mathcal{R}(\mathcal{V}^*) - \mathbb{E}\Big[\min_{\mathcal{V}} L(\mathcal{V}, \mathbf{z}) \mid \mathbf{y}\Big], \qquad \mathcal{V}^* = \arg\min_{\mathcal{V}} \mathcal{R}(\mathcal{V}). \tag{20.12}$$

Given the state, every crossing week v_i^* of Section 20.2.2 is known, and the best order is a set of one-week jobs with deadlines, coil i due by week v_i^*. The number of coils due by week v is the size of the set $\{i : 1 \le v_i^* \le v\}$, and the least number crossed is the standard result for such jobs,

$$\min_{\mathcal{V}} L(\mathcal{V}, \mathbf{z}) = \text{FC} \max_{v} \Big(\big|\{i : 1 \le v_i^* \le v\}\big| - v \Big)^+, \tag{20.13}$$

the largest number of coils due by any week beyond the weeks there are, met by visiting the coils in the order of their deadlines. Its average over the composition draws is the second term of equation (20.12), and the reported plan's loss averaged over the same draws, $\mathcal{R}(\mathcal{V}^*)$ on the draws, the first, with the standard error by batch means. One choice enters here: the naive baselines. The headline is the dashboard's worst-looking first. Three stronger orders stand beside it: the dashboard with the coils reading past the limit set aside (moved to the end of the cycle), the chart's crossing dates, earliest first, with the coils past by their lines set aside, and the calendar, oldest coil first.

Building blocks. The expected consequence (20.5) of each order on the mixture, the Laplace route and the draws; the loss of each order on each draw; the perfect-information loss (20.13) on each draw.

The orders. Table 20.6 gives each order's week 1 and its expected consequence on the three routes, and, over the draws, the probabilities that the plan's loss is smaller than the order's, equal to it and larger. The dashboard's order costs R$ 53 476.7, R$ 45 569.3 more than the plan, 6.76 times, and the plan is better on 0.7005 of the draws and worse on 9 of the 200 000, which the table's 0.0000 rounds away. Setting aside the coils reading past the limit takes the dashboard to R$ 42 234.9; the chart's dates cost R$ 48 866.4 and the calendar R$ 58 585.3. No visit at all costs R$ 119 716.4.

Table 20.6: The year's plan against the naive orders: the expected consequence (20.5) on three routes, and the comparison of the losses on the composition draws.

Order	week 1 [–]	$\mathcal{R}$, mixture [R$]	Laplace [R$]	draws [R$]	P(plan <) [–]	P(=) [–]	P(plan >) [–]
the plan (assignment on the mixture)	25	**7 907.4**	7 828.1	7 907.6	–	–	–
the dashboard, worst-looking first	15	**53 476.7**	53 038.4	53 371.4	0.7005	0.2995	0.0000
the dashboard, coils reading past set aside	1	42 234.9	41 972.0	42 185.4	0.6384	0.3395	0.0221
the chart's dates, past ones set aside	33	48 866.4	48 414.5	48 857.2	0.6784	0.3065	0.0151
the calendar, oldest first	4	58 585.3	58 003.7	58 552.0	0.7472	0.2502	0.0026
no visit in the cycle	–	119 716.4	119 036.6	–	–	–	–

The loss over the draws, and perfect information. Over the draws the plan loses nothing with probability 0.8194, one crossing with 0.1645, two with 0.0153 and three or more with 0.0008; its mean is R$ 7 907.6 and its standard deviation R$ 17 642.6. The dashboard's order loses nothing with probability 0.2418. The two distributions are drawn in Figure 20.5: the horizontal axis is the number of coils that cross before their visit, the vertical axis its probability over the draws, the blue bars the plan's and the red the dashboard's; the dashed lines mark each order's expected consequence in crossings, $\mathcal{R}$/FC, and the dotted line the mean loss with perfect information. With every coil's path known, equation (20.13) averages R$ 45.8, and

$$\mathrm{EVPI} = \mathcal{R}(\mathcal{V}^*) - \mathbb{E}\Big[\min_{\mathcal{V}} L(\mathcal{V}, \mathbf{z}) \mid \mathbf{y}\Big] = \text{R\$ } 7\,907.6 - \text{R\$ } 45.8 = \text{R\$ } 7\,861.8$$

by equation (20.12), with $\mathcal{R}(\mathcal{V}^*)$ the plan's mean loss over the same draws and a standard error of R$ 35.9, 99.4 % of the plan's expected consequence. Perfect information leaves a loss on only 0.0011 of the draws, where the crossings are so crowded that one crew cannot meet them all.

Of the plan's R$ 7 907.6, R$ 7 861.8 is the price of not knowing where the coils are. The crew has time for every coil at risk, but it does not know which ones are at risk.

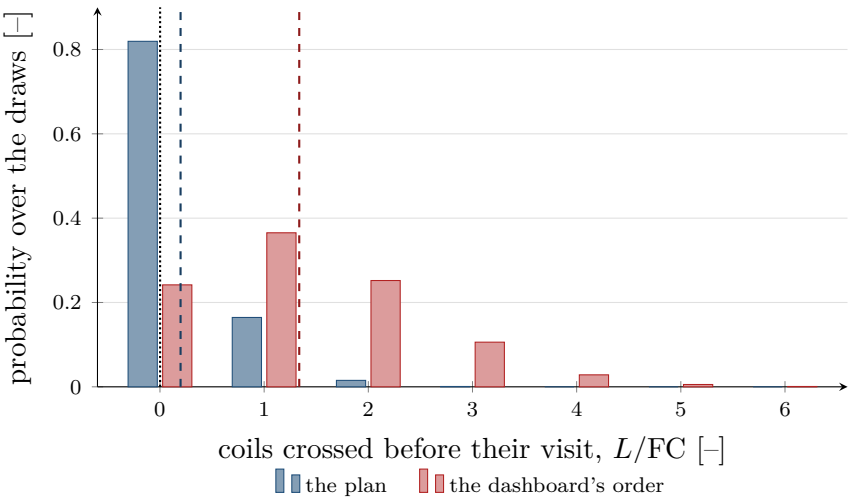


Figure 20.5: The loss of the plan and of the dashboard's order over the 200 000 composition draws, in coils crossed before their visit. Dashed: each order's expected consequence divided by FC, 7 907.4/40 000 for the plan and 53 476.7/40 000 for the dashboard. Dotted: the mean loss with every coil's path known, equation (20.13), 45.8/40 000.

20.4 The checks (step 5)

Each answer of Section 20.3 rests on draws, on routes over the spreads, on the stated limit and costs, on the clean and on the straight lines of step 0, and each can be set against the calculation it stands for.

20.4.1 The draws against the closed form

The plan's mean loss on the composition draws is R$ 7 907.6, with a standard error of R$ 35.5 by batch means, against the closed form's R$ 7 907.4, and the dashboard's is R$ 53 371.4 (standard error 92.5) against R$ 53 476.7: both gaps lie within about one and a half standard errors, the plan's well within one. The one-week crossing probabilities of the plan's first five coils from the draws are 0.00941 (standard error 0.00017) against the closed form's 0.00947 for coil 25, 0.00841 (0.00024) against 0.00865 for coil 49, 0.00704 (0.00016) against 0.00714 for coil 38, 0.00490 (0.00012) against 0.00497 for coil 14, and 0.00543 (0.00020) against 0.00556 for coil 2, each within about one standard error; their probabilities of being past agree to within two, for example 0.7320 (0.0007) against 0.7324 for coil 25. The best order chosen on the draws themselves, scored on fresh draws from a generator seeded with 19, costs R$ 7 937.8, against R$ 7 930.8 for the reported plan on the same fresh draws: the optimization does not feed on the draws' noise.

20.4.2 The routes over the spreads

Table 20.7 collects the routes. Each row is a route or a reading; the columns give the plan's week 1 and its expected consequence, the dashboard's, the expected number of coils past, and coil 14's cycle's crossing probability. The last row is the stated truth (Section 20.4.6). The largest gap between the Laplace route and the mixture is 9.91×10^{-3} in any $P_i(t_{\mathrm{ah}})$ and 6.66×10^{-3} in a probability of being past, and every route's week 1 is coil 25. The plug-in spreads, which drop the spreads' uncertainty from each coil's variance, give the smallest expected consequences, R$ 7 727.9 for the plan; the plan's first visit does not move.

Table 20.7: The routes to the year's plan, the expansion of equation (2.24), and the stated truth.

Route or reading	week 1 [–]	$\mathcal{R}(\mathcal{V}^*)$ [R$]	dashboard [R$]	expected number past [–]	P_{14}, cycle [–]
mixture over the grid (reported)	25	7 907.4	53 476.7	22.4533	0.4561
normal on the spreads' marginal (Laplace)	25	7 822.8	53 038.4	22.4440	0.4462
spreads at the grid node of the mode (plug-in)	25	7 727.9	52 846.6	22.4289	0.4515
composition draws, 200 000	25	7 902.4	53 371.4	–	–
expansion of (2.24) at the most probable lines	–	0	40 000	23	–
stated truth	–	plan loses 0	0	21	crosses in week 39

20.4.3 The curvature term and the plug-in

The chapter has no physics rows, so the check of the limit of exact rows is replaced by the check of the curvature term of equation (2.24), the book's Laplace reading of an expected consequence. The loss (20.4) is a step in the state: it changes only where a crossing time passes a visit time. Its second differences at the most probable lines, the means at the marginal's mode, are 0.0 for each of the plan's first eight coils, so the curvature term vanishes and the expansion returns the plug-in, $L(\mathcal{V}^*, \hat{\mathbf{z}})$. On the most probable lines 23 coils are past the limit and two cross inside the cycle, coil 14 at 0.953 yr (window 48, Dec Y4) and coil 38 at 0.486 yr (window 42, Jun Y4), and the plan meets both: $L(\mathcal{V}^*, \hat{\mathbf{z}}) = \text{R\$}\,0$, against $\mathcal{R}(\mathcal{V}^*) = \text{R\$}\,7\,907.4$. The plug-in says the plan costs nothing, and it would say the same of any order that meets two crossings; the dashboard's plug-in loss is R$ 40 000. A loss made of indicators has no curvature to expand, and its expectation is carried by the mass on either side of each step. The Laplace reading of the expected consequence is therefore the expectation of the loss under the Laplace route, not its expansion, which is why that route is defined on the spreads in Section 20.2.2.

20.4.4 The stated limit and the stated costs

Table 20.8 repeats the plan at two other functional limits. Each row is a limit; the columns give its value of g, the number of coils past it with probability above 0.5, the plan's week 1 and weeks 1 to 6, and the expected consequences of the plan, the dashboard and no visit. The stated limit changes the answer. At $\ln 0.80$ the first visit is coil 29, and at the capacity study's per-coil share, -0.26313, it is coil 9, because the limit decides which coils are near it: a lower limit moves the backlog's edge down the fleet to coils that read low today and are still ahead of it.

Table 20.8: The plan under three functional limits, from the mixture.

Limit	g [–]	P(past) > 0.5 [–]	week 1 [–]	weeks 1–6 [–]	$\mathcal{R}(\mathcal{V}^*)$ [R$]	dashboard [R$]	no visit [R$]
ln(5/6), Part I's convention (the chapter's)	−0.18232	23	**25**	25 49 38 14 2 33	7 907.4	53 476.7	119 716.4
$\ln 0.80$	−0.22314	18	**29**	29 25 49 36 41 32	9 359.1	44 069.3	120 946.9
the capacity study's per-coil share	−0.26313	13	**9**	9 3 42 23 32 30	15 160.8	46 219.9	149 185.5

The costs do not change it. With FC halved or doubled the order is the same, at R$ 3 953.7 and R$ 15 814.7: FC multiplies every term of equation (20.9). CREW_PER_WEEK sets the cycle and the visit times, and no other stated cost enters the order.

20.4.5 The clean and the failure curve

The clean. A visit restores RESTORE_FRACTION = 0.85 of the coil's loss since commissioning. On the draws in which each of the plan's first eight coils is still ahead of the limit, the restored g has a median from −0.0481 (coil 25) to −0.0183 (coil 1), the median time to cross again is 3.56 yr (coil 14) to 10.46 yr (coil 2), and the probability of crossing again within a year is 0.0000 for each. A visited coil stays above the limit through the cycle, as Section 20.1.1 assumed.

The failure curve. A tube leak is the failure of Chapter 16, read on that chapter's exact route at each coil's age. The probability of a leak in the 0.9615-yr cycle is 0.1005 for coil 25 (age 4), 0.1094 for coil 49 (age 5), 0.1303 for coil 38 (age 7), 0.0926 for coil 14 (age 3), 0.1882 for coil 2 (age 11), 0.0926 for coil 0 and 0.2065 for coil 15 (age 12), and the fleet expects 6.2082 coils with a leak in the cycle. Under the choice of step 0, that a clean leaves the leak rate unchanged because Chapter 16 resolved no separate effect of condition beside age, the leaks are the same under every order and cancel from every difference of two orders' expected consequences. This is the one place where the failure, not only the crossing, enters, and it changes nothing; it answers the "No failures" paragraph of Chapter 9's Limits.

20.4.6 The answers against the stated truth

The last check sets the answers against the stated truth of the record (Section 20.6.5), the last row of Table 20.7. In the truth 21 coils are past the limit today, and three cross inside the cycle: coil 49 at 0.591 yr (week 31, window 44, Aug Y4), coil 38 at 0.720 yr (week 38, window 45, Sep Y4) and coil 14 at 0.736 yr (week 39, window 45, Sep Y4). The plan visits them in weeks 2, 3 and 4; the dashboard's order visits them in weeks 16, 31 and 21, also before their crossings. On this record's truth the plan, the dashboard and the calendar all lose R$ 0, and the chart's dates lose R$ 40 000. The dashboard's order met every crossing by chance: under the posterior it loses nothing with probability 0.2418. Coil 25, the plan's first visit, is truly past (posterior probability 0.732), so on the truth its visit prevented nothing; coil 36 is truly past at 0.946. Coils 2 (0.532), 38 (0.328) and 49 (0.739) are truly ahead, as are coils 1, 14 and 33. The fleet's record is Chapter 9's and is not drawn afresh here; its spread over fresh tests is Chapter 9's question.

20.5 The report (step 6)

The dashboard sends the crew first to the coil that looks worst, and every one of its first twelve coils has already crossed its limit. Their crossings are behind them, visiting them first prevents nothing, and the year's order is expected to cost R$ 53 477, against R$ 7 907 for the plan. The crew should go to coil 25 this week. It has probably crossed already (0.73), yet its probability of crossing in the next week, 0.00947 (R$ 378.7), is the largest of the fifty. Then the plan visits coils 49, 38, 14, 2, 33, 1 and 46; after about fourteen weeks the order no longer matters. Coil 14 reads past the limit on the dashboard, but is above it with probability 0.94 and is the coil most likely to cross this year (0.46); it can wait three weeks. Nearly all of the plan's expected cost, 99.4 %, is the price of not knowing exactly where the coils are, and no condition test available before the plan is set is worth its price.

20.5.1 The answers and their bands

Table 20.9 gives the answers in the order of Table 20.1. Each row is a question; the second column gives the reported value, from the mixture over the spreads; the third the Laplace route beside it; and the fourth the band, the probability or the standard error that qualifies it.

Table 20.9: The report for the fleet at the end of the record (window 36, Dec Y3). Reported values from the mixture over the spreads; probabilities over orders from the composition draws.

Asset manager's question	Reported value	Laplace	Band or probability
Fifty units, one crew; which one does the crew visit next?	coil 25 [–]; plan 25, 49, 38, 14, 2, 33, 1, 46 [–]; $\mathcal{R}(\mathcal{V}^*) = 7\,907.4$ [R$]	coil 25 [–]; 7 828.1 [R$]	better than coil 49 first on 0.5285 of the draws where they differ, than coil 15 first on 1.0000 [–]
What does it cost to make each coil wait, and for how long does that hold?	coil 25 over one week, $\mathcal{W} = 378.7$ [R$] ($P = 0.00947$ [–]); coil 14 over the cycle, 18 243.1 [R$] ($P = 0.4561$ [–])	0.00947; 0.4462 [–]	index over one week +86.1 [R$] on the plan; over the cycle +2 035.9 [R$]
Which coils are already past their functional limit?	23 with $P > 0.5$ [–]; expected 22.4533 [–]; coil 14 0.0619, coil 2 0.5316 [–]	22.4440; 0.0591, 0.5349 [–]	19 above 0.95, 23 below 0.05 [–]
What does the year's plan cost against the dashboard's order, and what would knowing every coil's condition be worth?	dashboard 53 476.7 against 7 907.4 [R$]; EVPI 7 861.8 [R$], 99.4 % of the plan's	53 038.4 against 7 828.1 [R$]	plan better on 0.7005 of the draws, worse on 9 of 200 000 [–]; EVPI standard error 35.9 [R$]

The answers rest on Chapter 9's posterior of the fleet and on the stated functional limit: the routes over the spreads agree on the first visit, the costs only scale the plan, and the limit is the one input that moves it.

20.5.2 The entry in the fleet register

The fleet's plan is entered in the fleet register at the end of the record, window 36 (Dec Y3), beside Chapter 9's assessment and Chapter 16's failure curve. It uses nothing later than window 36, and in the values of Table 20.9 it reads:

> *Fleet of 50 coils, end of record (window 36, Dec Y3); Chapter 9's assessment (192 tests); functional limit* $g_F = \ln(5/6)$ *against rating.*
> Past the limit: 23 coils with $P > 0.5$ (expected 22.5); backlog, not ranked. Coil 15 (worst reading) past with $P = 1.00$.
> Crew plan, windows 37 to 48 (Jan–Dec Y4), one coil a week: 25, 49, 38, 14, 2, 33, 1, 46, then the rest; expected consequence R$ 7 907 (no crossing with $P = 0.82$).
> Next visit: coil 25 (P(past) = 0.73; the largest one-week crossing probability of the fleet, 0.0095, R$ 379).
> The dashboard's worst-first order: R$ 53 477; not used.
> Coil 0 (the hall's coil 1): week 15; P(cross in the year) = 0.037.
> Next measurement: none before the crew starts (no test pays its price); keep the annual round.

The entry records what the plan rests on as well as the plan: the fleet's straight lines, one functional limit for every site, and a cost schedule that charges a crossing once. Whether the hall's own coil 1 is cleaned now or at the October shutdown is not settled by its week in the fleet's plan; it is Chapter 17's decision, on the hall's own path.

Common mistake

Sending the crew to the worst-looking unit first. A unit that reads far past its limit has already crossed it, and its consequence is sunk; the crew's week is worth what it prevents, the crossing probability over the wait of the units still ahead of the limit.

20.6 The next measurement (step 7)

The report leaves the plan's expected consequence at R$ 7 907.4, nearly all of it the price of not knowing the coils' state. A condition test read before the plan is set is one more row on a coil's two entries, and what it would be worth to the plan is known before its reading is.

20.6.1 The preposterior by reweighted draws

A candidate is a set of condition tests read today, each a row $(g_i(a_i) - y_i)/\sigma_y$ on coil i's two entries with the gradient $\mathbf{g}_{1,i}$ of equation (20.3). Its value is the design form (2.16) in which the quantity averaged over the outcomes is the expected loss of the best order after the outcome, the value of information (2.52),

$$\mathrm{EVSI} = \mathbb{E}_{\tilde{y}|\mathbf{y}}\Big[\mathcal{R}(\mathcal{V}_{\mathrm{now}} \mid \mathbf{y}, \tilde{y}) - \min_{\mathcal{V}} \mathcal{R}(\mathcal{V} \mid \mathbf{y}, \tilde{y})\Big], \tag{20.14}$$

with $\mathcal{V}_{\mathrm{now}}$ the present plan, since the average of $\mathcal{R}(\mathcal{V}_{\mathrm{now}} \mid \mathbf{y}, \tilde{y})$ over the outcomes returns its present expected consequence. It is computed by simulated outcomes.

1. Draw 500 outcomes: for each, a state from the fresh composition draws (seed 19), and readings of the tested coils at that state with their errors.
2. Reweight the 200 000 composition draws (seed 9) coil by coil, as in Section 5.2.4: each tested coil's crossing probabilities $P_i(t_{\mathrm{ah},v})$ are recomputed from the draws weighted by the likelihood of its own reading only. The 22 coils with a cycle's crossing probability of at least 0.002 are updated by a test, and the rest are held.
3. Score the present plan on the updated table, reassign the best order by equation (20.9), and take the difference.
4. Average the differences over the outcomes, with the standard error of the mean, and set beside it the check that the average of $\mathcal{R}(\mathcal{V}_{\mathrm{now}} \mid \mathbf{y}, \tilde{y})$ returns the present plan's expected consequence on the draws.

Two choices enter here. Each coil is updated by its own reading alone: a reading of one coil also moves the fleet node, and through it every other coil, and the procedure leaves that coupling out. This keeps the update to one coil's weights at a time, and it errs least where the fleet's fifty coils already fix the fleet node, which a few more tests barely move. And the present plan scored on the draws is the draws' own best order, $\mathcal{V}_{\mathrm{now}}$, of expected consequence R$ 7 902.4 on the draws, beside the R$ 7 907.4 of the reported plan $\mathcal{V}^*$ on the mixture, so that the outcomes and the plan they are set against come from one sample. A test is priced at INSPECT, R$ 1 500, the stated cost of one condition test of a coil of this model; the cost schedule describes it as a crew visit, and the chapter assumes, as its choice, that a technician's visit with reference instruments reads it without taking one of the crew's weeks. The ceiling of every candidate is perfect information, equation (20.12), which on the draws' own plan is R$ 7 856.6; it differs from the R$ 7 861.8 of Section 20.3.4 by the gap between that plan and the reported one under the draws, R$ 7 902.4 against R$ 7 907.6.

20.6.2 The candidates

Table 20.10 lists the candidates. Each row is a candidate read before the plan is set; the columns give its value (20.14) with its standard error, its price, the share of outcomes in which the plan's week 1 stays coil 25, and the check of step 4 of the procedure, the average over the outcomes of the present plan's expected consequence, with its standard error, against R$ 7 902.4; every check returns that value to within its standard error or about one and a quarter of it, and the least effective sample size over the outcomes is 485, for the reference tests, and 3 886 or more for the others. The candidates are a condition test of coil 25 alone, condition tests of the eight coils with the largest cycle's crossing probability (14, 33, 1, 38, 46, 25, 2 and 49), reference tests

of those eight at a standard deviation of 0.02, a permanent monitor on those eight read for its first window, at the standard deviation 0.0348 of one window of Chapter 5's fit, and the fleet's annual round brought forward, all fifty coils tested now, the permanent instrument at the place of the question. The values are compared with the prices in Figure 20.6: the horizontal axis is in reais on a logarithmic scale, each row a candidate, the dot its value with a bar of one standard error, the open square its price, and the dashed line the value of perfect information.

Table 20.10: Condition tests read before the plan is set, by the preposterior (20.14) over 500 simulated outcomes. The check is the mean over the outcomes of $\mathcal{R}(\mathcal{V}^* \mid \mathbf{y}, \tilde{y})$, with its standard error, against R$ 7 902.4.

Candidate	EVSI [R$]	se [R$]	price [R$]	week 1 stays coil 25 [–]	check (se) [R$]
a condition test of coil 25 (the plan's week 1), sd 0.04	150.2	14.3	1 500	0.644	7 902.4 (0.0)
condition tests of the eight coils with the largest cycle's crossing probability, sd 0.04	1 146.5	33.6	12 000	0.466	7 867.1 (66.5)
reference tests of those eight, sd 0.02	2 771.6	71.9	12 000	0.304	7 862.1 (113.3)
a permanent monitor on those eight, its first window, sd 0.0348	1 407.2	39.8	not stated	0.424	7 866.2 (74.5)
the fleet's annual round brought forward, all fifty tested now, sd 0.04	1 541.7	38.6	75 000	0.466	7 793.5 (90.1)
perfect information (the ceiling), equation (20.12)	7 861.8	35.9	–	–	–

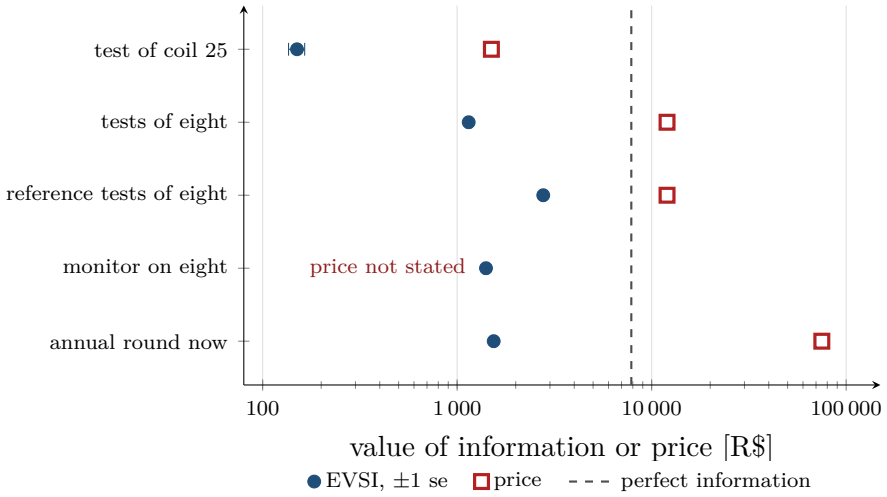


Figure 20.6: The candidate condition tests of Table 20.10: the value of each by the preposterior (20.14) (dots, with bars of one standard error) against its price (open squares), on a logarithmic axis. Dashed: the value of perfect information, R$ 7 861.8, the most any measurement before the plan can be worth.

Perfect information is worth R$ 7 862, yet no candidate recovers more than R$ 2 771.6 of it, and none pays its price. A test reads today's level, while a coil's crossing date also rests on its rate, which one more reading barely moves: the gradient $\mathbf{g}_{1,i} = (1,\ a_i)$ of the test row is that of the margin today, not of the margin at the end of the cycle. The candidate closest to its price is the eight reference tests, R$ 2 771.6 for R$ 12 000. A test of coil 25 alone leaves the plan's cost unchanged: coil 25 is in week 1 and costs nothing whatever its reading, so its test is worth only the re-planning it allows. A reading of -0.22 or lower says the coil has probably crossed

and moves coil 49 to week 1 (Exercise 20.5).

20.6.3 The recommendation

Send the crew to coil 25 now, without a test first. Keep the fleet's annual condition round as it is: its value to this one decision, R$ 1 541.7, is below its cost, and its value across the year's re-planning is not counted here, so that value is a lower bound. If one purchase is to be made before the plan, the reference tests of the eight coils at risk come closest, and still fall short of their price by more than three quarters.

20.6.4 What the analysis showed

The dashboard's worst-looking-first order is expected to cost R$ 53 477 over the year, against R$ 7 907 for the plan, because its first twelve coils have already crossed their limit and a crossing that has already happened is sunk. Ranking the coils by what waiting would cost, the crossing a visit can still prevent, separates the backlog from the coils at risk; the best order over the whole cycle is then an exact linear assignment, and the drill's index over the next week gives the same first visit. Nearly all of the plan's expected cost is uncertainty, 99.4 % of it the value of perfect information, yet no test available before the plan pays: a test reads where a coil stands, not how fast it falls. The stated functional limit, not the costs, decides which coil is first. And on the record's truth, the plan's first four weeks hold all three coils that truly cross this year.

20.6.5 Limits of the model

A constructed record. The fleet record is generated from a stated truth, with the stated test errors, and step 5 compares the answers with that truth (Section 20.4.6). It is not field data, and the agreement is a check of the method on its own model, not of the model against a real fleet.

Straight lines. The coils follow Chapter 9's straight lines, with no process noise and no acceleration. The hall's own coil 1 is read here only as a fleet member, on its fleet line.

The plan set once. No test is read and no plan revised during the cycle, so the values of step 7 are lower bounds on what a re-planning crew would draw from them.

A sunk backlog. The fleet record holds no cleaning history, so 21 coils truly sit past a limit that the cost schedule charges once. The schedule has no charge for each month below the limit; with one, the backlog would compete for the crew's weeks, and that cost is the owner's to state.

One limit, fixed slots. One functional limit serves every site, and the stated limit decides the first visit (Section 20.4.4). Visits are one a week on fixed slots, with no travel between sites and no crew absence.

Leaks. Tube leaks are taken as independent of the order, on the choice that a clean leaves the leak rate unchanged, which the failure log of Chapter 16 can neither confirm nor refute once age is counted. They would enter only through a stated cost of a leak that a visit prevents, and the schedule states none.

The next chapter. Chapter 21 asks of one coil whether its loss is faster than its specification, with the evidence for each answer, and whether to claim.

20.7 What to retain

Definition 20.1 (Maintenance prioritization by the cost of waiting)**.** The order of visits to a fleet minimizes the expected consequence $\mathcal{R}(\mathcal{V}) = \mathrm{FC}\sum_i P_i(t_{\mathrm{ah},v_i})$, with $P_i(t_{\mathrm{ah}}) = \Pr(0 < t_i^* \le t_{\mathrm{ah}} \mid \mathbf{y})$ the probability that unit i crosses its functional limit after today and by its visit. Under a normal component of the posterior, with the margins f_1 today and f_2 at t_{ah} linear in the unit's entries, $P_i(t_{\mathrm{ah}}) = \Phi(u_2) - \Phi_2(u_1, u_2; \varrho)$ and the probability of being past is $\Phi(u_1)$; a mixture sums them with its weights. A unit already past its limit is a sunk consequence, the same under every order. The expected consequence is linear in the assignment of units to slots, so its exact minimum is a linear assignment; the expected value of perfect information, $\mathcal{R}(\mathcal{V}^*) - \mathbb{E}[\min_{\mathcal{V}} L(\mathcal{V}, \mathbf{z}) \mid \mathbf{y}]$, bounds what any measurement before the plan can be worth.

Three conclusions carry beyond the fleet. A visit is worth what it can still prevent, so a ranking by condition alone mixes the units at risk with the units whose consequence is already incurred. A loss built from indicators of events is a step in the state, its expansion at the most probable state returns the plug-in, and its expectation must be taken over the posterior. And the value of perfect information can be large while every available measurement is worth less than its price, when the measurements read the quantity that matters least to the decision.

Table 20.11 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 20.11: The terms of maintenance prioritization and where each is defined. Units for the fleet of coils.

Term	Expression	Units	Meaning	Defined in
Value of a visit	$\Pr(0 < t_i^* \le t_{\mathrm{ah}} \mid \mathbf{y})$, the increment of eq. (2.21)	–	the crossing a visit can still prevent	20.1.1
Backlog	$\Pr(g_i(a_i) \le g_F \mid \mathbf{y})$ near one, eq. (20.11)	–	a sunk consequence, the same under every order	20.1.2
Loss of an order	eq. (20.4)	R$	a sum over the units, one term at each unit's week	20.1.4
Rows and posterior				
Question's gradient pattern	eq. (20.3)	–	block diagonal; no fleet entry	20.1.4
Crossing in $(0, t_{\mathrm{ah}}]$	$\Phi(u_2) - \Phi_2(u_1, u_2; \varrho)$, eq. (20.8)	–	per node, in closed form	20.2.1
Variance budget of the margin	eq. (20.7)	–	$\mathrm{Cov}(f_1, \gamma) > 0$ widens it	20.2.3
Questions and answers				
Exact optimum by linear assignment	eq. (20.9)	R$	the best of the 50! orders	20.3.1
Probability the first choice is better	on the draws, beside eq. (2.24)	–	how firm week 1 is	20.3.1
Index rule and its delay	$\mathcal{W}_i(\Delta)$, eq. (20.10)	R$	exact only when Δ is the next slot	20.3.2
EVPI	eq. (20.12), (20.13)	R$	the price of not knowing the state	20.3.4
Step loss and its expansion	$L(\mathcal{V}^*, \hat{\mathbf{z}})$ against $\mathcal{R}(\mathcal{V}^*)$	R$	the expansion returns the plug-in	20.4.3
The stated limit as the lever	Table 20.8	–	moves the first visit; the costs do not	20.4.4
Leaks that cancel	the failure curve of Chapter 16	–	the same under every order	20.4.5
EVSI against EVPI	eq. (20.14)	R$	a test reads the level, not the rate	20.6.1

20.8 Exercises

Exercise 20.1. *Formulation.* Write "has coil 14, which reads -0.2099, already crossed its limit?" in the form of Table 20.1, and evaluate it from the fit of this chapter, with its cost of waiting over the cycle. (A prediction under a threshold, equation (2.12) with $f = g_{14}(3)$ and $f^* = g_F$, its complement read by equation (20.11): $\Pr(g_{14}(3) \le g_F \mid \mathbf{y}) = 1 - \Pr(g_{14}(3) \ge g_F \mid \mathbf{y}) = 0.0619$; over the cycle $\mathcal{W}_{14}/\mathrm{FC} = P_{14}(0.9615) = 0.4561$.)

Exercise 20.2. *Derivation.* For one coil with $(g_{i,0},\ \gamma_i)$ normal, show that $\Pr(0 < t_i^* \le t_{\mathrm{ah}}) = \Phi(u_2) - \Phi_2(u_1, u_2; \varrho)$, and evaluate it for coil 25 at the plug-in spreads over one week. (A line crosses in $(0, t_{\mathrm{ah}}]$ exactly when $f_1 > 0$ and $f_2 \le 0$, so the probability is $\Pr(f_2 \le 0) - \Pr(f_1 \le 0,\ f_2 \le 0)$, with the moments of equation (20.7). For coil 25, f_1 has mean -0.01500 and standard deviation 0.02493, f_2 over one week mean -0.01579 and standard deviation 0.02505, and $\varrho = 0.99998966$; $\Pr(\text{past}) = \Phi(u_1) = 0.7263$ and $\Pr(0 < t^* \le 1\,\mathrm{wk}) = 0.00954$.)

Exercise 20.3. *Derivation.* Show that FLEET_CROSS cannot change the best order, and give $\mathcal{R}(\mathcal{V}^*)$ when it doubles. (In equation (20.9) FC multiplies every entry of the assignment's cost table, so every order's expected consequence is scaled by the same factor and the minimizing order is unchanged: R\$ 15 814.7, the same order.)

Exercise 20.4. *Application.* Rank the coils by their cost of waiting over a quarter. Which coil comes first, and in which week does the plan visit it? (Coil 25, $\mathcal{W}_{25}(13\,\mathrm{wk}) = $ R\$ 4 124.2; the plan visits it in week 1.)

Exercise 20.5. *Application.* A condition test of coil 25 today, standard deviation 0.04, reads -0.20, or else -0.24. Give its probability of being past and the plan's week 1 for each reading. (The composition draws reweighted by the reading's likelihood: at -0.20, $\Pr(\text{past}) = 0.7751$ and week 1 stays coil 25; at -0.24, $\Pr(\text{past}) = 0.9017$, week 1 is coil 49, and $\mathcal{R}$ falls to R\$ 7 667.0.)

Chapter 21

Warranty Claims and Specification Compliance

21.1 The dashboard's warranty claim and the maker's specification

The hall's two coils are three years old. Today is December of year 3, window 36, early summer in the southern hemisphere where the hall stands. The coils' maker warrants that each coil's own conductance loses no more than 3 % a year over its first five years, and the asset manager is asked whether to claim. The condition-monitoring dashboard answers from its trend chart: the health-index tiles of Chapter 10, each coil's commissioning approach over the window's, with their least-squares line. Coil 1's tile reads 0.4873 today, and its line falls by 15.76 % of the commissioned tile a year, 5.25 times the specification; coil 2's tile reads 0.5947, and its line falls by 14.82 % a year, 4.94 times. The chart's own test of a line steeper than 3 % a year gives "$P = 1.000000$" for both coils, with slope standard errors of 0.0152 and 0.0151 per year on 35 degrees of freedom. The planner's conclusion is that both coils are out of specification five times over, with certainty: file both claims now, at R$ 5 000 each, and expect two replacement coils, an expected gain of R$ 110 000.

The chart cannot say whether either coil's own conductance, the quantity the specification names, is falling faster than 3 % a year, or how much of a tile's slope is the load, the season, the other coil and the sensors. It cannot answer what the maker will answer: that coil 2's loss in March of year 2 came with a chilled-water pump change, not with the coil; that unit 1's discharge sensor drifts; that the record ends in summer. And it cannot say how sure is sure enough to claim, when filing costs R$ 5 000 and a granted claim returns a coil worth R$ 60 000. The question of this chapter is

whether each coil's own degradation rate, with the March step not counted, is faster than the specification's 3 %/yr*, read as the posterior odds of the specification's two halves; whether to file each claim, read as the expected consequence of filing at* R$ 5 000 *against* R$ 60 000*; and what the maker's counter-arguments do to each.*

It is answered with the method of Chapter 5, on Part III's posterior: Chapter 12's step drawing with Chapter 11's curvature family.

Jargon: warranty claim, specification and its envelope, faster than and within specification, Bayes factor of a specification, break-even probability of a claim, grant rule, standard of proof, counter-argument, conceded event, loss to date, five-year rate, operating characteristic of a claim rule

A *warranty claim* asks the maker to replace a unit that has broken a promise of the sale. The maker's *specification* here bounds the loss of a coil's own conductance per year over its first five years, and its *envelope* is the lowest health index the bound allows at each age. A coil is *faster than specification* when its degradation rate is beyond the bound and *within specification* otherwise: two halves of the range of one entry of the state. The *Bayes factor of a specification* compares the two halves as two drawings of the same plant. The *break-even probability of a claim* is the probability of a grant above which filing lowers the expected consequence. A *grant rule* states on what the maker settles a claim; under a *standard of proof* the maker grants only on evidence that reaches a stated probability. A *counter-argument* is the maker's alternative account of the record, read as another drawing of the plant. A *conceded event* is an event the claim keeps out of the coil's rate whatever its evidence. The *loss to date* is the coil's change of log-conductance since commissioning, events removed, per year of service; the *five-year rate* is the mean of its rate over the specification's five years. The *operating characteristic of a claim rule* is the share of records, at a stated truth, in which the rule files.

Key idea

A specification is a threshold on one entry of the state, the coil's own degradation rate. Whether the coil breaches it is the posterior tail beyond the threshold, and the Bayes factor of the breach is that tail's odds over the prior's. Whether to claim is a second question with a different threshold: a claim pays when its probability of being granted exceeds the cost of filing over the value of winning, here one in twelve. The maker's counter-arguments are other drawings of the same graph; each moves the tail, and a claim stands when no drawing the record supports moves it below the break-even.

21.1.1 A specification as a threshold on the state

Chapter 10 read each coil's degradation rate as the trend entry γ_u of its log-derate path, in yr^{-1}; Chapter 12 added a step entry on coil 2's path at the pump change of March of year 2, and Chapter 11 a curvature entry on each coil's rate. The maker's specification speaks of the coil's own conductance, a parameter of the physics, not of a performance figure, so on the state it is a threshold on one entry. A loss of 3 % a year multiplies the conductance by 0.97 each year, so

$$\gamma_{\mathrm{S}} = \ln(1 - 0.03) = -0.030459\,\mathrm{yr}^{-1}, \qquad \text{faster than specification} \iff \gamma_u < \gamma_{\mathrm{S}}, \tag{21.1}$$

which in the form of equation (10.1) is $100(\mathrm{e}^{\gamma_{\mathrm{S}}} - 1) = -3.0000\,\%/\mathrm{yr}$. Its envelope on the coil's own health index against commissioning is $\mathrm{e}^{\gamma_{\mathrm{S}} t_{\mathrm{y}}} = 0.97^{t_{\mathrm{y}}}$ at an age of t_{y} years: 0.91267 at three years, a loss of 8.73 %, and 0.85873 at five, a loss of 14.13 %.

The two halves of the threshold, $\gamma_u < \gamma_{\mathrm{S}}$ and $\gamma_u \geq \gamma_{\mathrm{S}}$, are two drawings of one plant that differ only in the prior of one entry, cut at γ_{S}. Chapter 2's comparison form, equation (2.14), compares them, and *Reliability Engineering by Derivation*, Chapter 27 (reliability demonstration and acceptance tests), shows that for two halves of one entry it needs no new fit: the Bayes factor is the posterior odds of the halves over their prior odds. Whether to claim is a second question. Filing costs a fixed sum, and a grant returns a coil; by Chapter 2's expected consequence the claim pays when the probability of a grant times the coil's value exceeds the cost of filing. The threshold of that question is set by the two costs, and it is not the specification's.

Principle 21.1. *Claim when the probability of winning times the value won exceeds the cost of claiming. How much confidence a claim needs is for the costs to set, not for a convention.*

21.1.2 The two-unit hall and its warranty as a graph (step 0)

The plant graph is Chapter 10's, cited and not re-derived (Figure 10.1): 37 copies of Chapter 5's window, window 0 (commissioning) and windows 1 to 36, each with its free load Q_t and the two log-derates $g_{1,t}$, $g_{2,t}$ under the five physical rows (5.2) at the window's known inputs $T_s(t)$ and $Q_o(t)$; one trend row per coil between consecutive windows, linked to the trend entries γ_1 and γ_2; the shared recirculation fraction β_1; and the drift node $\dot{b}_{d,1}$ on unit 1's discharge reading from window 20, equation (10.2). Part III added two kinds of node to it. Chapter 12's step δ sits on coil 2's trend row into window 15, the pump change, equation (12.1), and Chapter 11's curvature entries $\dot{\gamma}_1$, $\dot{\gamma}_2$ enter every trend row of their coil through the lever x_t of equation (11.2), with the rate at time t_y written $\gamma_u(t_\mathrm{y}) = \gamma_u + \dot{\gamma}_u(t_\mathrm{y} - \bar{t}_\mathrm{y})$ about the record's center $\bar{t}_\mathrm{y} = 1.5\,\mathrm{yr}$, window 18. This chapter adds no row. It adds two question nodes, one per coil, each reading its coil's trend entry against γ_S, and one decision node per coil, *file*, fed by the two costs.

The warranty's stated inputs. The specification is the record's stated input: the loss of the coil's own conductance no faster than 3 %/yr over the first five years, measured on the coil, not on the hall. The first five years end at window 60 (Dec Y5), the last date at which a claim may be filed. Two costs of Part V's stated schedule enter the decision: filing a claim, the evidence pack and the engineer's time, costs CF = R$ 5 000 whatever the outcome, and a granted claim returns a replacement coil at no charge, CV = R$ 60 000, the price of replacing a coil. Two more enter only step 7: the annual discount rate of 8 %, with the factors $\mathcal{D}(1) = 0.9259$ and $\mathcal{D}(2) = 0.8573$ for a payment one and two years ahead, and one condition test of a coil, R$ 1 500. None is re-chosen here.

The maintenance log so far. The two-unit hall's log holds the commissioning test (window 0); the condition assessment and the effort entry of window 18 (Chapters 5 and 6); Chapter 7's assessment of window 25, with its side-by-side check of unit 1's discharge sensor asked for and still open; and at window 36 the rates (Chapter 10), the acceleration left undecided (Chapter 11), the March step probable, with an inspection of coil 2's water side asked for (Chapter 12), and no seasonal change of the coils (Chapter 13). This chapter's entry is also at window 36. The inspection of coil 2's water side is not yet on record.

Modeling choices at the graph. Four choices that no question block owns make this graph. *The claim reads Part III's drawings unchanged*: the trend rows at $\sigma_q = 0.005$, the drift linear from window 20 with Chapter 7's prior, the rating and survey priors, and Chapter 5's steady state, pinned airflow and known inputs. This buys a claim on a record the earlier chapters have already checked, and it fails where they fail; Section 21.3.3 and Section 21.4.3 vary each choice the maker could contest. *The maker settles a claim on the coil's own conductance*, as the specification's clause says: the claim is granted exactly when the coil is truly faster than specification, so that the probability of a grant is the posterior probability of $\gamma_u < \gamma_\mathrm{S}$. This is the grant rule, and it fails if the maker grants on the evidence pack instead, the standard of proof whose answer Section 21.3.2 gives beside it. *Each coil is its own claim*, filed at its own cost. *Today is window 36*, and the claim uses nothing later.

The graph. The graph reduced to the claim is drawn in Figure 21.1. The two coils' paths run as strips of log-derate nodes, five of the 37 windows drawn, with the trend squares between

consecutive windows and the window blocks between the strips. Above them sit the shared fraction β_1, coil 1's trend entry and curvature, and the drift node on unit 1's discharge readings; below, coil 2's trend entry, its step into window 15 and its curvature. On the right stand the two question nodes and the two decision nodes with the costs that feed them.

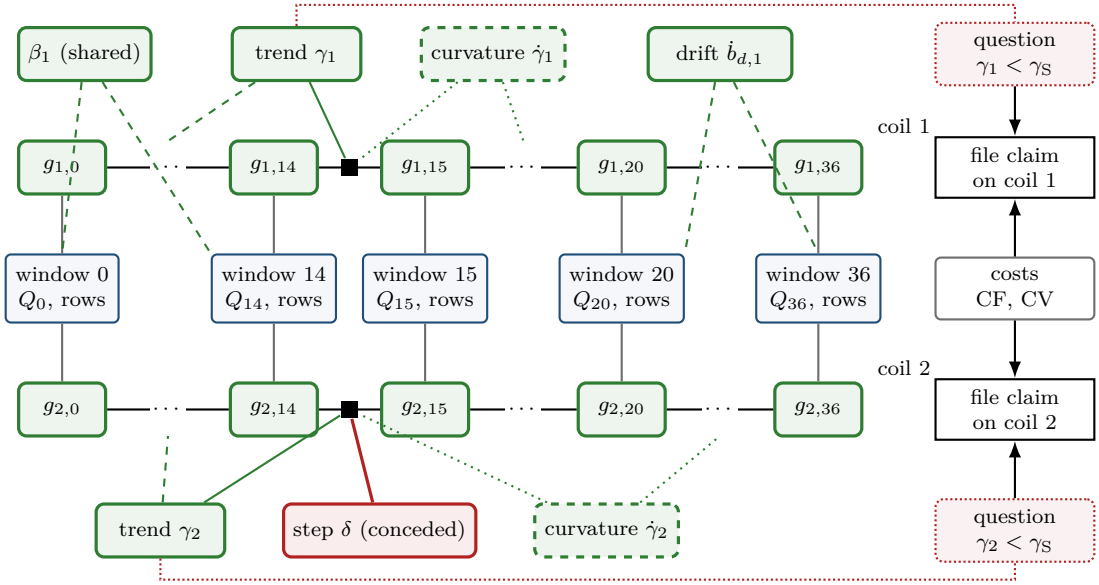


Figure 21.1: The record's graph reduced to the claim, five of the 37 windows. Free (green): each window's log-derates $g_{u,t}$ and load Q_t (inside each window's block, with Chapter 5's physical rows and five readings), the trend entries γ_1, γ_2, the shared fraction β_1 and the drift rate $\dot{b}_{d,1}$ on unit 1's discharge readings from window 20. Dashed green boxes: the curvature entries of Chapter 11, freed in some drawings of the family, linked (dotted) to every trend square of their coil. Black squares: the trend rows between consecutive windows, the runs between the drawn windows marked by dots. Red: the step δ on coil 2's square into window 15, conceded by the claim. Red dotted: the two question nodes, each reading its coil's trend entry against the specification; boxes: the two decisions, fed by the costs CF and CV. No row is added to Part III's; the two questions share the return air through the correlation of γ_1 and γ_2.

21.1.3 The asset manager's questions (step 1)

The asset manager asks four questions of the record, and step 1 writes each as a probability or an expectation over the state of the claim's drawings, with the unknown on the left of the bar and the 185 readings and the priors on the right. Table 21.1 lists them. The first column is the question in the asset manager's words, the first of them the drill question; the second the general form of its category from Chapter 2; the third its mathematical form for the hall; the fourth its category; the fifth the computation of Section 2.3.6 that answers it; and the last the subsection where it is answered. "This unit" in the drill question is each coil in turn. When the capacity limit is crossed is Chapter 14's question, whether to clean or replace either coil Chapter 17's, and whether the monitoring can resolve a loss of 3 %/yr in the time left Chapter 22's. Which measurement would change the claim asks nothing new of the record; it is step 7.

21.1.4 The evidence as rows (step 2)

Step 2 attaches the record to the graph as Part III's rows and adds none. The base is Chapter 10's stacked state (10.3) of 115 entries under 261 rows: the 185 reading rows (10.4), the rating priors on window 0's coils, the survey prior on β_1, the 72 trend rows (10.5) at $\sigma_q = 0.005$ and the drift

Table 21.1: The questions asked of the two-unit hall's condition record at window 36 about the maker's specification. $\mathbf{y}$ is the 185 readings of windows 0–36 and the prior rows; γ_{S} the specification's rate (21.1); P_u and P_u^0 the posterior and prior probabilities that coil u is faster than specification; M_{arg} a counter-argument's drawing and M_{claim} the claim's; γ_u^{td} and $\gamma_u^{5\mathrm{y}}$ the loss to date and the five-year rate.

Asset manager's question	General form	Mathematical form	Category	Computation	Answered in
Has this unit degraded faster than its specification allows, with enough confidence to make the claim?	(2.14) on the two halves of one entry; (2.12)	$P_u = \Pr(\gamma_u < \gamma_{\mathrm{S}} \mid \mathbf{y})$, mixed over the reported family; $B_u^{\mathrm{S}} = \dfrac{P_u/(1-P_u)}{P_u^0/(1-P_u^0)}$	hypothesis comparison	trend	21.3.1
Should the claim be filed, at R\$ 5 000 to win R\$ 60 000?	(2.23), (2.24)	$\mathcal{R}(\text{file}) = \mathrm{CF} - \mathrm{CV}\,P_u$, $\mathcal{R}(\text{not}) = 0$; a^*; $\Pr(a^* \text{ better} \mid \mathbf{y})$; EVPI	design	expected consequence	21.3.2
What do the maker's arguments do to the claim: the March step, the drifting sensor, the summer?	(2.14), each argument a drawing; (2.12) under each	$\Pr(\gamma_u < \gamma_{\mathrm{S}} \mid \mathbf{y}, M_{\mathrm{arg}})$; $\ln B(M_{\mathrm{arg}} : M_{\mathrm{claim}})$	hypothesis comparison	evidence	21.3.3
Read as the loss so far, or over the specification's five years, is it still faster?	(2.21) (line 1), (2.12)	$\Pr(\gamma_u^{\mathrm{td}} < \gamma_{\mathrm{S}} \mid \mathbf{y})$, $\Pr(\gamma_u^{5\mathrm{y}} < \gamma_{\mathrm{S}} \mid \mathbf{y})$	prediction under a threshold	trend	21.3.4

prior. Chapter 12's step adds the entry δ, with the prior row $\delta/0.15$; Chapter 11's curvature adds an entry $\dot{\gamma}_u$ for each coil a drawing frees, with the prior row $\dot{\gamma}_u/0.05$, in yr^{-2}.

The state and the trend rows. The claim reads four drawings, the step drawing with each of Chapter 11's four curvature drawings, written $M^{(d)}_{2,15}$ with $d \in \{0, 1, 2, 12\}$ the coils whose curvature is freed (none, coil 1, coil 2, both). The state of $M^{(d)}_{2,15}$ and its trend rows are

$$\begin{aligned} \mathbf{z}^{(d)} &= \left(\mathbf{z}_{\text{none}},\ \delta,\ \dot{\gamma}_u\ (u \in d)\right)^{\mathsf{T}}, \\ r_{q,u,t}(\mathbf{z}^{(d)}) &= \frac{g_{u,t+1} - g_{u,t} - (\gamma_u + \dot{\gamma}_u\, x_t)\,\Delta t - \delta\,\mathbf{1}[u = 2,\ t+1 = 15]}{\sigma_q}, \qquad t = 0, \ldots, 35, \end{aligned} \tag{21.2}$$

with $\mathbf{z}_{\text{none}}$ the stacked state (10.3), $\dot{\gamma}_u = 0$ for a coil the drawing does not free, and x_t the lever of equation (11.2): 116 entries on 262 rows for $M^{(0)}_{2,15}$, the step drawing of Chapter 12, 117 for $M^{(1)}_{2,15}$ and $M^{(2)}_{2,15}$, and 118 for $M^{(12)}_{2,15}$, each with 146 degrees of freedom. The reading rows and the other priors are those of the base fit.

The family and its prior. The reported family crosses Chapter 12's step drawing with Chapter 11's four curvature drawings at even odds. Beside it stands the record's whole family, which does not concede the step: Chapter 12's nine drawings at the named window, no step (M_{none}) or a step of one of the eight kinds j of Table 12.2 into window 15, crossed with the same four, 36 drawings,

$$\Pr(M^{(d)}_{2,15}) = \tfrac{1}{4}\ \text{(reported)}, \qquad \Pr(M^{(d)}_{j,15}) = \tfrac{1}{4} \cdot \begin{cases} \tfrac{1}{2}, & j = \text{none}, \\ \tfrac{1}{16}, & \text{each of the eight kinds} \end{cases} \ \text{(whole)}. \tag{21.3}$$

Which row acts on which entry. The Jacobian of $M^{(d)}_{2,15}$ is Chapter 12's pattern (12.3), the base pattern (10.6) with the step's column, plus one column for each freed curvature, whose nonzeros sit in its coil's 36 trend rows, $x_t \Delta t/\sigma_q$ each, and in its prior row, Chapter 11's pattern (11.5). The questions' quantities are linear in the state, with the gradients

$$\mathbf{g}^{\mathsf{T}}_{\gamma_u} = \mathbf{e}^{\mathsf{T}}_{\gamma_u}, \qquad \mathbf{g}^{\mathsf{T}}_{\text{td},u} = \frac{\mathbf{e}^{\mathsf{T}}_{g_{u,36}} - \mathbf{e}^{\mathsf{T}}_{g_{u,0}} - \mathbf{1}[u = 2]\,\mathbf{e}^{\mathsf{T}}_{\delta}}{36\,\Delta t}, \qquad \mathbf{g}^{\mathsf{T}}_{5\text{y},u} = \mathbf{e}^{\mathsf{T}}_{\gamma_u} + (2.5\,\mathrm{yr} - \bar{t}_{\text{y}})\,\mathbf{e}^{\mathsf{T}}_{\dot{\gamma}_u}, \tag{21.4}$$

the last only in a drawing that frees $\dot{\gamma}_u$; these are not rows. Two facts of the graph are read here, before any number. The trend entry, the specification's quantity, has no column in a reading row: it is read from the path through the trend rows alone, as in Chapter 10, and its tail carries the wander of the path. And the loss to date reaches across the record, one entry at window 0, one at window 36 and, for coil 2, the step: the conceded step enters the claim through that one entry.

The loss of each action. Under the grant rule of step 0 a filed claim costs CF and returns CV exactly when the coil is faster than specification, and not filing costs nothing,

$$L(\text{file}, \mathbf{z}) = \text{CF} - \text{CV}\,\mathbf{1}\big[\gamma_u < \gamma_{\text{S}}\big], \qquad L(\text{not}, \mathbf{z}) = 0, \tag{21.5}$$

in reais, negative for a gain: a step in the one entry γ_u.

Modeling choices at the rows. Four choices make these rows. *The specification's quantity is the trend entry* γ_u, the coil's mean degradation rate over its service to date; in a curvature drawing it is the rate at the record's center, which is the record's mean rate, because the levers x_t sum to zero. This buys one entry per coil, already in the state, and it fails if the maker reads "over the first five years" as the mean over all five; the loss to date and the five-year rate are question 4, Section 21.3.4. *The step is conceded*: the reported family is Chapter 12's step drawing, so that coil 2's rate is its fouling with the March step taken out, whatever the step's evidence. This buys a claim the maker's strongest argument cannot shake, and it costs coil 2's claim whatever the step carried; the whole family stands beside it. *No prior row on the rates*, Part III's choice: the trend entries keep a flat prior, the limit of a wide normal centered at γ_{S}, so the two halves of the specification carry even prior odds; a proper prior centered at γ_{S} keeps them even, and the fleet's prior does not (Section 21.4.3). *CF and CV are the only costs that enter* a decision taken today; the discount enters only when a decision waits, in step 7.

21.2 The posterior of the state (step 3)

Every answer of this chapter is read from the posteriors of the four drawings of the reported family, mixed by their evidences, with the step drawing $M_{2,15}^{(0)}$ shown in full.

21.2.1 The fit of each drawing and the family's weights

The derivation form is Chapter 5's, on Chapter 10's stacked state with Part III's entries added: the state and the trend rows of equation (21.2), the residual vector stacking the reading rows, the trend rows and the priors, the Jacobian of the pattern of step 2, 262×116 for the step drawing, evaluated at each iterate because the reading rows are nonlinear in the log-derates, the normal equations (5.9) solved by Gauss–Newton from the base MAP state with the new entries at zero, and the Laplace covariance $\mathbf{\Sigma}^{(d)} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1}$ at the minimum, equation (5.10). With every row normalized the information is $\mathbf{J}^{\mathsf{T}}\mathbf{J}$; with unnormalized rows and their weights it would be $\mathbf{J}^{\mathsf{T}}\mathbf{W}\mathbf{J}$. The evidence of each drawing is its Laplace evidence, equation (7.7), and with the even prior (21.3) the family's posterior weights are the normalized evidences of equation (7.9),

$$q_d = \Pr(M_{2,15}^{(d)} \mid \mathbf{y}) = \frac{p(\mathbf{y} \mid M_{2,15}^{(d)})}{\sum_{d'} p(\mathbf{y} \mid M_{2,15}^{(d')})}, \qquad d \in \{0, 1, 2, 12\}. \tag{21.6}$$

Table 21.2 gives the four fits. Each row is a drawing; the columns give its number of entries, its log evidence, its log Bayes factor against $M_{2,15}^{(0)}$, its weight q_d, and the MAP value and posterior spread of each coil's trend entry. The drawing that frees coil 1's curvature carries 0.4381 of the weight, the step drawing 0.2634. Coil 1's rate is -0.04398 to $-0.04829\,\mathrm{yr}^{-1}$ across the four, and coil 2's -0.02732 to $-0.03139\,\mathrm{yr}^{-1}$, with spreads of 0.01465 to 0.01700.

Table 21.2: The reported family: Chapter 12's step drawing with Chapter 11's four curvature drawings, at even prior odds. Laplace evidences (7.7); weights by equation (21.6); trend entries at the MAP state with their posterior spreads.

Drawing	entries [–]	$\ln p(\mathbf{y})$ [–]	$\ln B$ vs $M_{2,15}^{(0)}$ [–]	weight q_d [–]	$\hat{\gamma}_1$ [yr^{-1}]	sd [yr^{-1}]	$\hat{\gamma}_2$ [yr^{-1}]	sd [yr^{-1}]
$M_{2,15}^{(0)}$, constant rates	116	−66.1687	0	0.2634	−0.04398	0.01465	−0.03139	0.01673
$M_{2,15}^{(1)}$, coil 1 curves	117	−65.6599	+0.5088	0.4381	−0.04825	0.01496	−0.02732	0.01686
$M_{2,15}^{(2)}$, coil 2 curves	117	−66.9790	−0.8103	0.1171	−0.04376	0.01465	−0.03008	0.01697
$M_{2,15}^{(12)}$, both curve	118	−66.5412	−0.3725	0.1815	−0.04829	0.01502	−0.02739	0.01700

The curvatures within the family. Under $M_{2,15}^{(1)}$ coil 1's curvature is $\dot{\gamma}_1 = -0.03403\,\mathrm{yr}^{-2}$ with spread 0.02018, and $\Pr(\dot{\gamma}_1 < 0) = 0.9542$; under $M_{2,15}^{(12)}$ it is -0.03422 (0.02119), 0.9469. Coil 2's is -0.00945 (0.01980) under $M_{2,15}^{(2)}$ and $+0.00061$ (0.02071) under $M_{2,15}^{(12)}$. Over the family, coil 1's curvature is freed with probability 0.6195: freed and negative, an accelerating loss, with 0.5898; pinned at zero, a constant rate, with 0.3805; freed and positive with 0.0297. The step moves Chapter 11's split, which without it is 0.5003, 0.4608 and 0.0390, with the weights 0.3284, 0.3470, 0.1323 and 0.1923 that this fit reproduces. As Chapter 11 found, the record leaves coil 1's acceleration undecided. Under the step drawing the step itself is $\delta = -0.07139$ (0.02349), a loss of 6.89 % of coil 2's conductance, with $\ln B(\text{step} : \text{none}) = +2.7648$: Chapter 12's numbers, reproduced.

The record's whole family. Over the 36 drawings of the whole family, with the prior (21.3), a step at window 15 of any kind has probability 0.7809, coil 2's step 0.5421, and coil 1's curvature is freed with 0.6393. The two largest weights are those of coil 2's step with coil 1's curvature, 0.23748, and with constant rates, 0.14278; the drawings without a step carry 0.07195, 0.07601, 0.02899 and 0.04212. Without a step, coil 2's rate is -0.06004 to $-0.06253\,\mathrm{yr}^{-1}$; with it, -0.02732 to -0.03139. Chapter 12's own named-window family, the nine with constant rates, gives a step at window 15 the probability 0.7160, as that chapter reported.

21.2.2 The step drawing's state

The step drawing $M_{2,15}^{(0)}$ has 116 entries, 262 rows and 146 degrees of freedom; at its MAP state $2\mathcal{L}(\hat{\mathbf{z}}) = 145.7166$ and its log evidence is -66.1687. Table 21.3 gives the entries the claim reads, with their posterior spreads: each coil's log-derate at commissioning and today, coil 2's either side of the step, the two trend entries, the step and the drift rate.

Table 21.3: The entries of the step drawing $M_{2,15}^{(0)}$ that the claim reads, at the MAP state, with their posterior spreads.

Entry [unit]	$\hat{z}_i$ [entry's]	$\sqrt{\Sigma_{ii}}$ [entry's]	Meaning
$g_{1,0}$ [–]	-0.00542	0.01543	coil 1's log-derate at commissioning
$g_{1,36}$ [–]	-0.13737	0.02395	coil 1's log-derate today
$g_{2,0}$ [–]	-0.03303	0.01586	coil 2's log-derate at commissioning
$g_{2,14}$ [–]	-0.06756	0.01407	coil 2's log-derate before the step
$g_{2,15}$ [–]	-0.14164	0.01506	coil 2's log-derate after the step
$g_{2,36}$ [–]	-0.19858	0.01736	coil 2's log-derate today
γ_1 [yr^{-1}]	-0.04398	0.01465	coil 1's degradation rate
γ_2 [yr^{-1}]	-0.03139	0.01673	coil 2's degradation rate, the step out
δ [–]	-0.07139	0.02349	the step into window 15
$\dot{b}_{d,1}$ [K/yr]	$+0.66823$	0.16809	unit 1's discharge drift rate

The covariance. The block of $\boldsymbol{\Sigma}$ on these ten entries, in the order of Table 21.3, is, multiplied by 10^4 and written $10^4\,\boldsymbol{\Sigma}_{\text{claim}}$ (the last row and column carry K/yr),

$$\begin{bmatrix}
2.3801 & -1.0976 & -0.86665 & -0.16704 & -0.18149 & 0.58938 & -1.1592 & 0.50414 & -0.056397 & -5.6888 \\
-1.0976 & 5.7363 & 0.90674 & -0.28783 & -0.071710 & -2.3925 & 2.2780 & -1.2051 & 0.31619 & 32.618 \\
-0.86665 & 0.90674 & 2.5139 & 0.056452 & 0.62886 & -0.85615 & 0.59113 & -1.3514 & 0.68427 & 5.0836 \\
-0.16704 & -0.28783 & 0.056452 & 1.9807 & -0.32438 & 0.56691 & -0.040262 & 0.96443 & -2.3828 & -1.6699 \\
-0.18149 & -0.071710 & 0.62886 & -0.32438 & 2.2689 & -0.33140 & 0.036593 & -1.2173 & 2.6917 & -1.1466 \\
0.58938 & -2.3925 & -0.85615 & 0.56691 & -0.33140 & 3.0142 & -0.99395 & 1.6346 & -1.0334 & -13.960 \\
-1.1592 & 2.2780 & 0.59113 & -0.040262 & 0.036593 & -0.99395 & 2.1457 & -0.56976 & 0.12420 & 12.769 \\
0.50414 & -1.2051 & -1.3514 & 0.96443 & -1.2173 & 1.6346 & -0.56976 & 2.7995 & -2.4124 & -6.7083 \\
-0.056397 & 0.31619 & 0.68427 & -2.3828 & 2.6917 & -1.0334 & 0.12420 & -2.4124 & 5.5195 & 1.0811 \\
-5.6888 & 32.618 & 5.0836 & -1.6699 & -1.1466 & -13.960 & 12.769 & -6.7083 & 1.0811 & 282.55
\end{bmatrix}.$$

The correlations of the trend entries read by sign, and the rows fix each sign. γ_1 and γ_2 are correlated by -0.2325: the return air sees the hall's total conductance, so a faster loss on one coil at the same readings is paid for by a slower loss on the other, and the two claims trade. γ_2 and δ are correlated by -0.6137: a larger March loss leaves less of coil 2's fall to its fouling, which is why conceding the step costs coil 2's claim. γ_1 and $\dot{b}_{d,1}$ are correlated by $+0.5186$: a faster drift of unit 1's discharge sensor explains more of the sensor's late rise, which leaves less of it to coil 1, Chapter 7's ridge; γ_2 and $\dot{b}_{d,1}$ by -0.2385, through the return air. Each rate is correlated negatively with its coil's commissioning entry, -0.5130 for coil 1 and -0.5094 for coil 2, and positively with its entry today, $+0.6493$ for coil 1 and $+0.5627$ for coil 2: a rate is a change from window 0 to window 36.

21.2.3 The loss to date and the trend entry

The specification says "loss … no faster than 3 % a year", and two quantities of the state read it. The trend entry γ_u is the rate the coil's path fouls at. The loss to date is what the path did, its change since commissioning per year of service with the conceded step taken out, the second gradient of the pattern (21.4),

$$\gamma_u^{\mathrm{td}} = \frac{g_{u,36} - g_{u,0} - \mathbf{1}[u = 2]\,\delta}{36\,\Delta t}. \tag{21.7}$$

With γ_u free, the residuals of the coil's 36 trend rows sum to zero at the MAP state, the telescoped sum of equation (10.7) with the step's increment removed, so the two share their mean: -0.04398 and $-0.03139\,\mathrm{yr}^{-1}$ on the step drawing, a gap below 10^{-12}. Their variances differ. Given the path, the trend rows are the only rows on γ_u; they put its conditional mean at the path's γ_u^{td} and leave the conditional variance at the wander floor of equation (10.12), so that by the law of total variance, equation (3.34),

$$\mathrm{Var}(\gamma_u \mid \mathbf{y}) = \mathrm{Var}(\gamma_u^{\mathrm{td}} \mid \mathbf{y}) + \frac{\sigma_q^2}{36\,\Delta t^2}. \tag{21.8}$$

On the step drawing the variances are 2.1457×10^{-4} against 1.1457×10^{-4} for coil 1, and 2.7995×10^{-4} against 1.7995×10^{-4} for coil 2, a difference of 1.0000×10^{-4} in each: the floor $\sigma_q/(\sqrt{36}\,\Delta t) = 0.01000\,\mathrm{yr}^{-1}$, squared. The loss to date is the coil's record; the trend entry adds the month-to-month wander a rate of fouling carries, and its tail is the wider of the two.

21.2.4 The Bayes factor of a specification

The first question asks for the Bayes factor of "faster than specification" against "within",

$$B_u^{\mathrm{S}} = \frac{p(\mathbf{y} \mid M_{\mathrm{fast}})}{p(\mathbf{y} \mid M_{\mathrm{within}})},$$

the comparison form (2.14) with M_{fast} the claim's drawing with the prior of γ_u cut to $\gamma_u < \gamma_{\mathrm{S}}$ and renormalized, and M_{within} the same drawing cut to $\gamma_u \geq \gamma_{\mathrm{S}}$. Write $P_u^0 = \Pr(\gamma_u < \gamma_{\mathrm{S}})$ for the prior mass of the first half. The cut prior is $p(\mathbf{z})\,\mathbf{1}[\gamma_u < \gamma_{\mathrm{S}}]/P_u^0$, so the evidence of M_{fast} is the integral of the likelihood against it, and with $p(\mathbf{y} \mid \mathbf{z})\,p(\mathbf{z}) = p(\mathbf{z} \mid \mathbf{y})\,p(\mathbf{y})$, $p(\mathbf{y})$ the evidence of the uncut drawing, the integral is a tail of the posterior,

$$p(\mathbf{y} \mid M_{\mathrm{fast}}) = \frac{1}{P_u^0} \int \mathbf{1}[\gamma_u < \gamma_{\mathrm{S}}]\,p(\mathbf{y} \mid \mathbf{z})\,p(\mathbf{z})\,\mathrm{d}\mathbf{z} = p(\mathbf{y})\,\frac{\Pr(\gamma_u < \gamma_{\mathrm{S}} \mid \mathbf{y})}{P_u^0}.$$

The same steps on the other half give $p(\mathbf{y} \mid M_{\mathrm{within}}) = p(\mathbf{y})\,(1 - P_u)/(1 - P_u^0)$, with $P_u = \Pr(\gamma_u < \gamma_{\mathrm{S}} \mid \mathbf{y})$. Dividing, $p(\mathbf{y})$ cancels,

$$B_u^{\mathrm{S}} = \frac{P_u/(1 - P_u)}{P_u^0/(1 - P_u^0)}, \qquad P_u \geq \kappa \iff B_u^{\mathrm{S}} \geq \frac{\kappa/(1 - \kappa)}{P_u^0/(1 - P_u^0)}, \tag{21.9}$$

the posterior odds of the two halves over their prior odds, the acceptance form of *Reliability Engineering by Derivation*, Chapter 27; the second relation is the assurance rule of equation (2.12) written on the factor. With no prior row on γ_u, the limit of a wide prior centered at γ_{S}, or with any prior centered there, $P_u^0 = \frac{1}{2}$: the factor is the posterior odds, and an assurance level of $\kappa = 0.95$ asks for $B_u^{\mathrm{S}} \geq 19$. No new fit is needed: under each drawing P_u is the normal tail of the Laplace normal's γ_u, and over a family it is their mixture.

21.2.5 The full posterior by reweighted draws

Beside the Laplace normals stands the full posterior of each drawing, computed by the reweighted draws of Section 5.2.4, with the drawing's Laplace normal as the proposal and the weights of equation (5.13). Each drawing takes 20 000 draws, in chunks of 5 000, from a seeded generator, and keeps the claim's six quantities, γ_u, γ_u^{td} and γ_u^{5y} for each coil, with the weights. The same draws give each drawing's evidence by equation (7.12), and the drawings are mixed by those evidences in place of the Laplace ones. The Monte Carlo standard errors are by batch means over ten batches. Section 21.4.1 states the effective sample sizes and the errors.

21.3 The answers (step 4)

Every answer is read from the Laplace normals of the reported family, mixed by the weights (21.6) as in equation (7.10), and that value is the one reported. The full posterior of Section 21.2.5 stands beside it, and the record's whole family beside that.

21.3.1 The breach of specification, coil by coil

Asset manager's question. Has this unit degraded faster than its specification allows, with enough confidence to make the claim? It is asked of each coil in turn; "enough confidence to make the claim" is the second question's, Section 21.3.2.

General form. A prediction under a threshold is the posterior probability that an event holds, equation (2.12),

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

and a hypothesis comparison the Bayes factor of two drawings, equation (2.14), here the two halves of one entry's range.

Specialized form. The question is equation (2.12) with $f(\mathbf{z}) = -\gamma_u$, $f^* = -\gamma_{\mathrm{S}}$ of equation (21.1), the gradient $\mathbf{g}_{\gamma_u}$ of the pattern (21.4), and $p(\mathbf{z} \mid \mathbf{y})$ the mixture (7.10) of the four Laplace normals of Section 21.2.1. Under drawing d the trend entry is normal, with mean $\hat{\gamma}_{u,d}$ and spread $\varsigma_{u,d}$ read off $\boldsymbol{\Sigma}^{(d)}$, so the tail is a weighted sum of normal tails, and the Bayes factor of the specification is equation (21.9) with even prior odds,

$$P_u = \sum_d q_d\, \Phi\!\left(\frac{\gamma_{\mathrm{S}} - \hat{\gamma}_{u,d}}{\varsigma_{u,d}}\right), \qquad B_u^{\mathrm{S}} = \frac{P_u}{1 - P_u}. \tag{21.10}$$

The rate's median and 5–95 % band are the quantiles of the same mixture, the roots of $\sum_d q_d \Phi((\gamma - \hat{\gamma}_{u,d})/\varsigma_{u,d}) = 0.05, 0.5, 0.95$, the form of equation (9.16), carried to %/yr by the increasing map of equation (10.1). The mean and variance of the mixture are equation (7.11), the within-drawing and the between-drawing parts. One choice enters here: the tail is read from each drawing's Laplace normal, which makes it a closed form in $(\hat{\gamma}_{u,d}, \varsigma_{u,d})$, and it is accurate only if each drawing's posterior of γ_u is close to normal; the full posterior beside it tests that.

Building blocks. The four drawings' MAP values and spreads of Table 21.2; their weights q_d; the tail (21.10); the factor (21.9); the mixture quantiles and the map (10.1).

The step drawing first. On $M_{2,15}^{(0)}$ alone, with the entries of Table 21.3, coil 1's rate lies $(\gamma_{\mathrm{S}} - \hat{\gamma}_1)/\varsigma_1 = (-0.030459 + 0.04398)/0.01465 = +0.9233$ standard deviations beyond the specification, so $P_1 = \Phi(0.9233) = 0.8221$ and $B_1^{\mathrm{S}} = 4.6200$ ($\ln B = +1.5304$). Coil 2's lies $(-0.030459 + 0.03139)/0.01673 = +0.0557$ beyond it, $P_2 = 0.5222$ and $B_2^{\mathrm{S}} = 1.0929$ ($\ln B = +0.0888$): on the step drawing the record cannot tell coil 2's rate from the specification's. Over the four drawings P_1 is 0.8221, 0.8829, 0.8181 and 0.8824, and P_2 is 0.5222, 0.4262, 0.4910 and 0.4283: the drawings that free coil 1's curvature raise coil 1's tail and lower coil 2's.

The family. Table 21.4 collects the answer. Each row is a coil and a reading of the record: the reported family, the full posterior beside it, the whole family that does not concede the step, the families of Chapters 12 and 11 alone, the step drawing, and the base fit with no step. The columns give the median rate and its band in %/yr, the probability of the breach and its Bayes factor.

Table 21.4: The breach of specification, coil by coil. Rates as $100(\mathrm{e}^{\gamma} - 1)$, negative for a loss; bands 5–95 %; P_u by equation (21.10) and B_u^{S} by equation (21.9) with even prior odds. The full posterior's Monte Carlo standard errors are in parentheses.

Coil	Reading	median [%/yr]	band [%/yr]	P_u [–]	B_u^{S} [–]
1	**reported**: the step conceded, the curvature family	**−4.552**	**−6.883** to **−2.173**	**0.8592**	**6.10**
1	full posterior of the reported family	–	–	0.8567 (0.0014)	–
1	the record's whole family	−4.832	−8.330 to −2.208	0.8723	6.83
1	Chapter 12's family alone	−4.383	−7.042 to −1.937	0.8237	4.67
1	Chapter 11's family alone	−4.352	−6.677 to −1.981	0.8269	4.78
1	the step drawing $M_{2,15}^{(0)}$	−4.303	–	0.8221	4.62
1	the base fit, no step	–	–	0.7928	3.83
2	**reported**: the step conceded, the curvature family	**−2.832**	**−5.505** to **−0.082**	**0.4595**	**0.85**
2	full posterior of the reported family	–	–	0.4655 (0.0017)	–
2	the record's whole family	−3.616	−7.056 to −0.444	0.6147	1.60
2	Chapter 12's family alone	−4.137	−7.361 to −0.831	0.6983	2.32
2	the step drawing $M_{2,15}^{(0)}$	−3.090	–	0.5222	1.09
2	the base fit, no step	–	–	0.9923	129.33

Coil 1 is probably faster than specification: its rate has median −4.552 %/yr, band −6.883 to −2.173, $P_1 = 0.8592$ and $B_1^{\mathrm{S}} = 6.10$. Coil 2, with the March step not counted, is a coin flip leaning within: median −2.832 %/yr, band −5.505 to −0.082, $P_2 = 0.4595$ and $B_2^{\mathrm{S}} = 0.85$. Neither reaches the assurance rule's $B \geq 19$ of equation (21.9) at $\kappa = 0.95$. Coil 1 is faster than 5 %/yr with probability 0.3770 and slower than 1 %/yr with 0.0073; coil 2 with 0.0917 and 0.1354. The mixture's variance is nearly all within the drawings, 2.2058×10^{-4} within and 4.4568×10^{-6} between for coil 1 (mean $-0.04661\,\mathrm{yr}^{-1}$), 2.8442×10^{-4} and 3.2692×10^{-6} for coil 2 (mean $-0.02873\,\mathrm{yr}^{-1}$): the curvature family barely moves the record's mean rate. Not conceding the step raises coil 2's probability to 0.6147, and dropping the step altogether, the base fit, to 0.9923; the step is what stands between coil 2 and a strong claim. The full posterior moves the two probabilities by −0.0025 and +0.0060, which Section 21.4.1 resolves.

The coils together. The two claims are not independent. On the step drawing, through the correlation −0.2325 of γ_1 and γ_2, the probability that both coils are faster than specification

is 0.4053, against 0.4293 if they were independent, and that neither is, 0.0610. A coil 1 that accelerates takes late loss from coil 2 through the return air, which is why coil 2's P falls from 0.5222 under $M^{(0)}_{2,15}$ to 0.4262 under $M^{(1)}_{2,15}$.

The posterior densities of the two rates are plotted in Figure 21.2, coil 1 above and coil 2 below. The horizontal axis is the rate in %/yr, the vertical axis its density; the solid curve is the reported mixture, its area beyond the specification shaded, and the dashed curve the whole family's. The vertical line at −3 %/yr is the specification, the dotted line the stated truth (Section 21.4.5), and the arrow at the left edge points to the dashboard's slope, far off the axis. Coil 1's density lies mostly beyond the specification. Coil 2's straddles it, and its whole family adds a shoulder toward −6 %/yr, the drawings without a step.

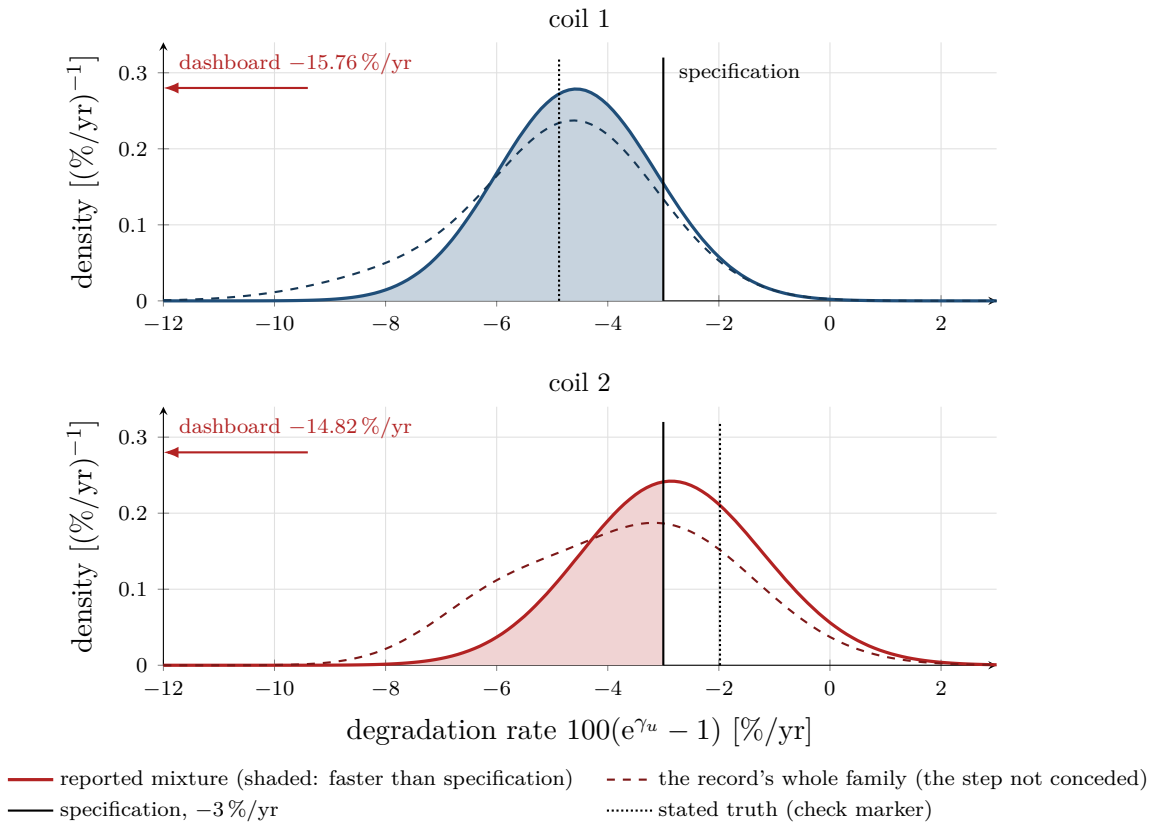


Figure 21.2: The posterior density of each coil's degradation rate in %/yr, by the change of variables from the trend entry: solid, the reported mixture of the four normals of Table 21.2 with the weights q_d, its tail beyond the specification shaded ($P_1 = 0.8592$, $P_2 = 0.4595$); dashed, the mixture over the record's whole family of 36 drawings. Solid vertical: the specification. Dotted: the stated truth of each coil's mean rate over the record, −4.877 and −1.980 %/yr, a check marker (Section 21.4.5). Arrows: the dashboard's slopes, off the axis.

The surprise of the answer. The dashboard's 14.82 % a year on coil 2 is, on the coil's own conductance with the pump change taken out, 2.83 %/yr at the median: within specification.

21.3.2 The decision to file

Asset manager's question. Should the claim be filed, at R$ 5 000 to win R$ 60 000? It is asked of each coil.

General form. An expected consequence compares actions by their posterior expected loss, equation (2.23),

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

with the Laplace reading of equation (2.24) and the probability that the chosen action is the better one.

Specialized form. The question is equation (2.23) with the two actions and their losses (21.5), and $p(\mathbf{z} \mid \mathbf{y})$ the reported mixture. The loss of filing is a constant less CV on the event of equation (21.10), so its expectation is that probability, weighted:

$$\mathcal{R}(\text{file}) = \text{CF} - \text{CV}\, P_u, \qquad \mathcal{R}(\text{not}) = 0, \qquad \text{file} \iff P_u > p_{\text{be}} = \frac{\text{CF}}{\text{CV}} = 0.08333. \tag{21.11}$$

The break-even p_{be}, one in twelve, replaces the assurance level κ of equation (2.12): the costs set it. Filing has the lower loss at the true state exactly when the coil is faster than specification, so $\Pr(a^* \text{ better} \mid \mathbf{y}) = P_u$ when filing is chosen. With the coil's rate known, the claim would be filed only on a coil that is faster, and the expected loss would be $(\text{CF} - \text{CV})\, P_u$; the expected value of perfect information, Chapter 17's form (17.9), is the difference,

$$\text{EVPI}_u = \mathcal{R}(\text{file}) - (\text{CF} - \text{CV})\, P_u = \text{CF}\,(1 - P_u), \tag{21.12}$$

the cost of a claim that fails. The decision flips where $\mathcal{R}(\text{file}) = 0$, at $\text{CF} = \text{CV}\, P_u$ or $\text{CV} = \text{CF}/P_u$.

The Laplace reading. Equation (2.24) expands the loss about the MAP state. The loss (21.5) is a step in γ_u, flat on either side of γ_{S}, so at $\hat{\mathbf{z}}$ its gradient and its curvature are zero: the expansion returns the plug-in $L(\text{file}, \hat{\mathbf{z}})$, and the ranking line divides by a zero spread. The reported value is therefore the expectation of the loss over the Laplace normals, equation (21.11), not the expansion, as in Chapter 20.

Building blocks. The probabilities P_u of Section 21.3.1; the costs CF and CV of step 0; the rule (21.11); the value (21.12).

The decision for each coil. Table 21.5 gives it. Each row is a coil; the columns give the probability of a grant, the expected consequence of filing by the Laplace route and by the full posterior, the chosen action and the probability that it is the better one, the expected value of perfect information, and the cost of filing or the value of a grant at which the decision would flip, with its multiple of the stated cost.

Table 21.5: The decision to file, by equations (21.11) and (21.12), in reais; negative is a gain. The full posterior's Monte Carlo standard errors are in parentheses; the flip points are the stated cost at which $\mathcal{R}(\text{file}) = 0$, with its multiple of CF = R$ 5 000 or CV = R$ 60 000.

Coil	P_u [–]	$\mathcal{R}$(file) [R$]	full posterior [R$]	action [–]	$\Pr(a^* \text{ better})$ [–]	EVPI [R$]	flips at CF [R$]	flips at CV [R$]
1	0.8592	**−46 551.1**	−46 399.1 (83.8)	**file**	0.8592	704.1	51 551 (×10.31)	5 819 (×0.097)
2	0.4595	**−22 569.0**	−22 931.6 (100.4)	**file**	0.4595	2 702.6	27 569 (×5.51)	10 882 (×0.181)

The plan is to file both claims now, with an expected consequence of R$ − 69 120.1. The record's whole family gives R$ − 47 336.5 and R$ − 31 879.2, with values of perfect information of R$ 638.6 and R$ 1 926.7: the same actions. The expansion of equation (2.24) gives R$ −55 000 for both coils, the plug-in at the MAP state, where both rates lie beyond γ_{S}: it prices both claims as sure.

The surprise of the decision. The plan files a claim on coil 2 that is more likely to fail, 0.5405, than to succeed. It pays because a failure costs R$ 5 000 and a success returns R$ 60 000: the claim needs one chance in twelve, and coil 2 has nearly one in two.

The costs that could flip it. Both costs are stated inputs, and each is moved here, not re-chosen. The cost of filing would have to rise 5.51 times, or the value of a grant fall to 0.181 of its stated value, before coil 2's claim stops paying; coil 1's has 10.31 and 0.097. With the cost of filing doubled the expected consequences are R$ $-$ 41 551.1 and R$ $-$ 17 569.0, and with the value of a grant halved R$ $-$ 20 775.6 and R$ $-$ 8 784.5: both still file. The discount rate does not enter a decision taken today.

The naive claims, priced on the same posterior. Table 21.6 sets the plan beside the rules a dashboard or a convention would apply. Each row is a rule; the columns give its action on each coil, its expected consequence under the reported posterior, and its difference from the plan.

Table 21.6: The plan against the naive claims, each priced under the reported posterior by equation (21.11), in reais; negative is a gain.

Rule	coil 1 [–]	coil 2 [–]	$\mathcal{R}$ [R$]	against the plan [R$]
the plan, the expected consequence (21.11)	file	file	−69 120.1	–
the dashboard's claim, on the chart's slope (its own reading −110 000)	file	file	−69 120.1	0
the chart's line at 95 % ($P = 1.000000$ for each)	file	file	−69 120.1	0
the 95 % rule on the coil's own posterior	not	not	0	+69 120.1
the standard of proof, a grant only on evidence at $P \geq 0.95$	not now	not now	0 now	see step 7

The dashboard reaches the plan's action for the wrong reason. Its claim is the same two filings, but its stated certainty overstates their value by R$ 40 879.9, and its evidence, a tile on the hall's approach, is not the coil's conductance that the specification names. The rule the drill question's wording suggests, "enough confidence" read as 95 % on the coil's own posterior, files neither claim and gives up R$ 69 120.1. Under a standard of proof the maker grants only a claim whose evidence reaches 0.95; at 0.8592 and 0.4595 neither is grantable now, filing would cost R$ 5 000 each for nothing, and the question becomes when the evidence could reach that level, which is step 7's.

The decision is drawn in Figure 21.3. The horizontal axis is the probability of a grant, the vertical axis the expected consequence of filing in thousands of reais; the line is equation (21.11), $\mathrm{CF} - \mathrm{CV}\,P$, which crosses zero at the break-even; the two vertical lines mark the break-even 1/12 and the 95 % level; the two large marks are the coils' reported probabilities; and the two strips at the top hold each coil's probability under every reading of Table 21.7. Every mark of both strips lies to the right of the break-even, and most of coil 2's lie left of 0.95 by far.

21.3.3 The maker's counter-arguments

Asset manager's question. What do the maker's arguments do to the claim: the March step, the drifting sensor, the summer? Each is an account of the record in which the coil is not at fault.

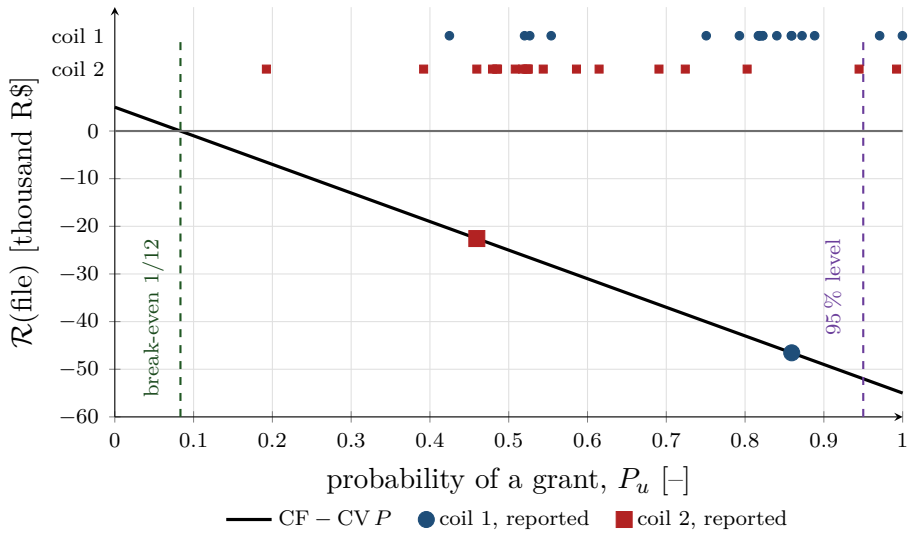


Figure 21.3: The expected consequence of filing against the probability of a grant, equation (21.11) with CF = R$ 5 000 and CV = R$ 60 000; filing pays below the gray zero line, to the right of the break-even $p_{\text{be}} = 1/12$. Large marks: each coil's reported probability and expected consequence. Strips at the top: each coil's probability under every reading of Table 21.7, the claim's own and the base fit's included.

General form. A hypothesis comparison is the Bayes factor of two drawings fitted to the same evidence, equation (2.14),

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

and under each drawing the breach is the threshold form (2.12).

Specialized form. The question is equation (2.14) with $M_A = M_{\text{arg}}$, the step drawing $M^{(0)}_{2,15}$ with the argument's entries or choices, and M_B the drawing it modifies, and under each M_{arg} the tail $\Pr(\gamma_u < \gamma_{\text{S}} \mid \mathbf{y}, M_{\text{arg}})$ of equation (21.10) on that drawing alone. Each argument is refitted by Gauss–Newton from the step drawing's MAP state, and its evidence is equation (7.7). Three kinds of argument are read. *The step*: whether it is conceded, at which window, with which prior, on which part. *The drift*: its onset and its prior, and a second drift on unit 2's discharge sensor, $\dot{b}_{d,2}$ from window 20 with the same prior. *The season*: Chapter 13's seasonal conductance on each coil, the trend rows (13.9) with the amplitudes g^c_u, g^s_u, and the claim read as of a winter window, before the record's last summer. The choice that enters here is the set of arguments: the ones a maker can raise from the record and the work orders. An argument outside the graph, such as a dispute of the specification's measure, is not tested and is named in the limits.

Building blocks. The step drawing's fit of Section 21.2.2; each argument's refit, its evidence and its trend entries; the tail (21.10); the break-even p_{be} of equation (21.11).

The arguments. Table 21.7 gives them. Each row is an argument or a reading; the columns give the log Bayes factor of the argument's drawing against the claim's step drawing $M^{(0)}_{2,15}$, one baseline for every row, and each coil's rate on it and its probability of a breach. The family rows and the readings as of an earlier window have no single drawing on the same record, and no factor. For a row that only widens a prior, the factor is mostly that prior's Occam factor, not evidence against the account.

Table 21.7: The maker's counter-arguments, each a drawing of the same graph refitted on the record; P_u the probability of a breach under it, equation (21.10). The rates are the trend entries in %/yr, medians over the family in the first two rows. $\ln B$: the argument's drawing against the claim's step drawing $M_{2,15}^{(0)}$.

Argument or reading	$\ln B$ vs $M_{2,15}^{(0)}$ [–]	rate 1 [%/yr]	P_1 [–]	rate 2 [%/yr]	P_2 [–]
the claim, the reported family	–	−4.552	0.8592	−2.832	0.4595
The step					
not conceded, the record's whole family	–	−4.832	0.8723	−3.616	0.6147
no step at all, the base fit	−2.7648	–	0.7928	–	0.9923
conceded at window 15, $M_{2,15}^{(0)}$	0	−4.303	0.8221	−3.090	0.5222
a month earlier, window 14	−1.8274	−4.276	0.8169	−3.799	0.6907
a month later, window 16	−0.3929	−4.283	0.8185	−3.179	0.5438
its prior wider, 0.30	−0.5975	−4.306	0.8225	−3.034	0.5083
on both coils, the pump serving both	−3.7086	−3.104	0.5267	−5.190	0.9447
on unit 2's discharge sensor, not the coil	−1.6767	−6.138	0.9708	−3.920	0.7241
The drift					
onset at window 16	−1.9652	−4.531	0.8587	−2.936	0.4843
onset at window 24	+1.3126	−4.375	0.8403	−3.071	0.5175
onset at window 28	+2.8232	−4.627	0.8883	−2.917	0.4794
drift prior 4.0 K/yr	−0.6474	−4.288	0.8191	−3.098	0.5241
drift prior 100 K/yr, all but free	−3.8509	−4.283	0.8181	−3.101	0.5248
unit 2's discharge sensor drifts too	−1.4782	−3.089	0.5201	−4.963	0.8026
no drift entry at all	−5.7497	−7.149	0.9997	−1.633	0.1924
The season					
a seasonal conductance on each coil	−4.3224	−3.968	0.7507	−2.526	0.3920
the claim as of window 30 (Jun Y3, winter)	–	−3.224	0.5539	−2.929	0.4858
the claim as of window 31 (Jul Y3, winter)	–	−2.693	0.4247	−3.417	0.5860

The step. The step is the maker's strongest argument, and the claim concedes it: conceding it takes coil 2 from 0.6147, the whole family's, to 0.4595. Against no step, the step on coil 2 has $\ln B = +2.7648$ at window 15, $+0.9374$ at window 14 and $+2.3719$ at window 16: the record favors March, with April close. Placing it in February or April, or widening its prior to 0.30, moves coil 2 between 0.5083 and 0.6907. A step on both coils at the pump change would cut coil 1 to 0.5267, but the record weighs against it, $\ln B = -3.7086$ against the claim's drawing, and coil 2 would rise to 0.9447. A step on unit 2's discharge sensor in place of the coil, -1.6767, raises both claims, to 0.9708 and 0.7241.

The drift. The drift is already carried. Its onset from window 16 to 28 moves coil 1 between 0.8403 and 0.8883 and coil 2 between 0.4794 and 0.5175; the record prefers the later onsets ($\ln B = +1.3126$ at window 24 and $+2.8232$ at window 28, against -1.9652 at window 16), and at window 28 coil 1 rises and coil 2 falls, both on the same side of the break-even as before. The prior's width moves neither coil, because the readings fix the drift, $\dot{b}_{d,1} = +0.6682$ K/yr (0.1681); its factors, -0.6474 at 4.0 K/yr and -3.8509 at 100 K/yr, are the wider prior's Occam factors. The maker's stronger form, that if one discharge sensor drifted so did the other, is rejected by the record, $\ln B = -1.4782$, with $\dot{b}_{d,2} = -0.2650$ K/yr (0.2242), -0.3754 K at window 36; and granted, it would trade the claims, coil 1 to 0.5201 and coil 2 to 0.8026. Without a drift entry at all, coil 1's claim would be near certain, 0.9997, and coil 2's would fall to 0.1924, with $\ln B = -5.7497$: the drift entry is the maker's argument on coil 1 and the claimant's on coil 2.

The season. The season is in the known inputs. A seasonal conductance on the coils is rejected, $\ln B = -4.3224$; its summer amplitudes, the cosine terms, are $g_1^c = +0.0007$ (0.0086) and $g_2^c = +0.0009$ (0.0091), with the sine terms $g_1^s = +0.0173$ (0.0087) and $g_2^s = +0.0025$ (0.0097); granted, it leaves 0.7507 and 0.3920. Read as of June or July of year 3, before the last summer, coil 1's claim falls to 0.5539 and 0.4247. What carries coil 1's claim is the second half of year 3, which the seasonal drawing says is not the season: coil 1's late loss, Chapter 11's undecided acceleration.

The decision under every argument. The lowest probability of a breach over every reading of Table 21.7 is 0.4247 for coil 1 and 0.1924 for coil 2, both above the break-even 0.08333: the decision to file holds under every argument.

The surprise of the arguments. The arguments that move loss between the coils trade the two claims: the step's window and extent, the drift and a second drift, and the wander σ_q of Section 21.4.3 each take from one coil's conductance what they give to the other's, because the coils share the return air. The season does not trade: a seasonal conductance lowers both claims, to 0.7507 and 0.3920, and so does reading the claim as of window 30, to 0.5539 and 0.4858. What keeps every argument above the break-even is that the break-even is low, one chance in twelve, not the trade between the coils. Filing both claims hedges only against the arguments that move loss between the coils; against the season, both claims fall together.

The arguments are drawn as a forest in Figure 21.4. Each row is a row of Table 21.7, the claim's at the top, with the stated truth as a check row at the bottom; the horizontal axis is the probability of a breach; circles are coil 1 and squares coil 2, drawn hollow where the argument's drawing has $\ln B < -1$ against the claim's; the dashed vertical lines mark the break-even and 0.5. No mark lies left of the break-even. The hollow rows include the arguments that move the claims most, in opposite directions for a step on both coils or a second drift.

What the arguments do to the dashboard's slope. The dashboard's slope is a quantity of the same state, and Chapter 10's sequential split, equation (10.16), takes it apart on the step

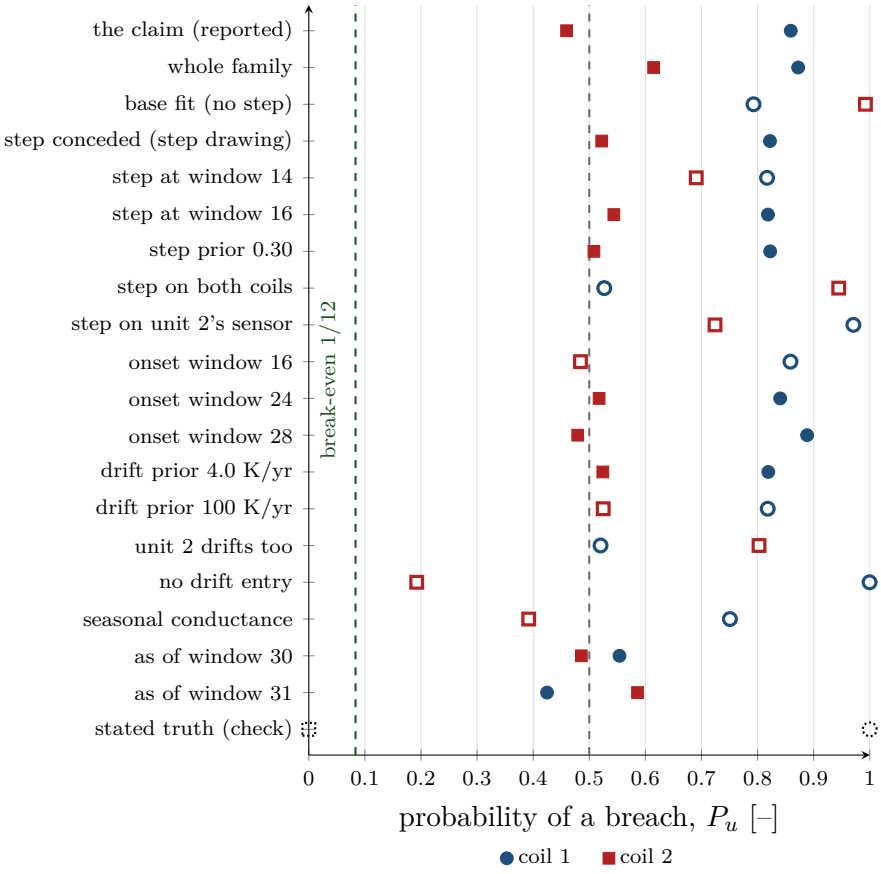


Figure 21.4: The probability of a breach for each coil under each reading of Table 21.7 (circles: coil 1; squares: coil 2). Hollow: a drawing with $\ln B < -1$ against the claim's step drawing $M_{2,15}^{(0)}$, the column of Table 21.7 (for the drift prior of 100 K/yr, the wide prior's Occam factor). Dashed: the break-even $p_{\text{be}} = 1/12$ and 0.5. Bottom row, dotted: the stated truth, coil 1 faster than specification and coil 2 within, a check marker (Section 21.4.5).

drawing's state, with the step as one more term. Table 21.8 gives the split. Each row is a coil; the columns give the slope of the line through the readings' tiles and its five parts: the load and the season, the two coils' fouling, coil 2's March step, unit 1's drift and the reading errors, each in per year with its share of the slope.

Table 21.8: The dashboard's slope on the step drawing's state, split by equation (10.16) with the step as a term; per year of the commissioned tile, shares of the slope in parentheses.

Coil	slope of the readings' tiles [yr^{-1}]	load and season [yr^{-1}]	the coils' fouling [yr^{-1}]	coil 2's March step [yr^{-1}]	unit 1's drift [yr^{-1}]	reading errors [yr^{-1}]
1	−0.1576	−0.0240 (15.2 %)	−0.0823 (52.2 %)	−0.0105 (6.6 %)	−0.0374 (23.7 %)	−0.0035 (2.2 %)
2	−0.1482	−0.0240 (16.2 %)	−0.0715 (48.3 %)	−0.0431 (29.1 %)	0	−0.0095 (6.4 %)

The same split is drawn as stacked bars in Figure 21.5. Each bar is a coil's slope, running left from zero, its segments the five parts in the order of the table; the diamond on each bar is the coil's own degradation rate $\hat{\gamma}_u$ on the step drawing, on the same axis. Granted every argument, the chart still shows 8.23 % and 7.15 % a year, the coils' fouling part, but that is a tile on the approach, 1.87 and 2.28 times each coil's own rate of -0.04398 and $-0.03139\,\mathrm{yr}^{-1}$: the approach magnifies its coil's loss and carries part of the other's, Chapter 5's finding. Only the fit turns the tile into the coil's conductance.

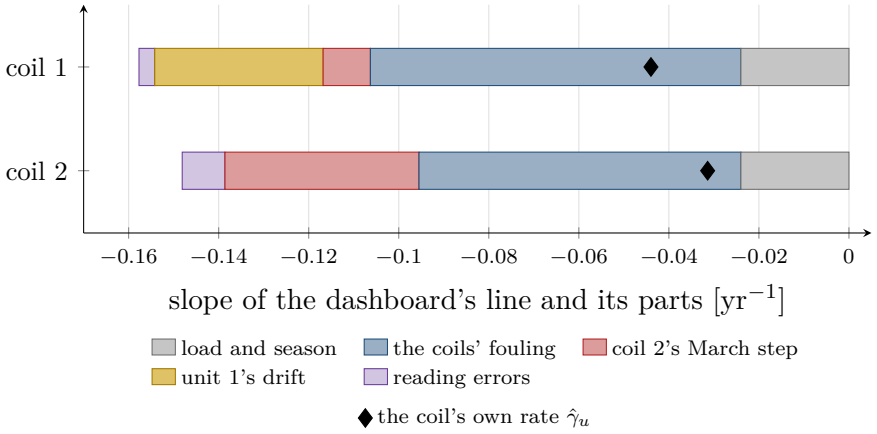


Figure 21.5: The dashboard's slope per coil on the step drawing's state, split by equation (10.16) (Table 21.8): stacked from zero leftward, the load and the season, the coils' fouling, coil 2's March step, unit 1's drift and the reading errors. Diamonds: each coil's own degradation rate on the same axis, -0.04398 and $-0.03139\,\mathrm{yr}^{-1}$.

21.3.4 The loss to date and the specification's five years

Asset manager's question. Read as the loss so far, or over the specification's five years, is it still faster? The specification's words "over the first five years" can be read three ways, and the reported reading is one of them.

General form. The sign of a trend is the threshold form on its trend entry, the first line of equation (2.21),

$$\Pr(\gamma > 0 \mid \mathbf{y}) = \int \mathbf{1}[\gamma > 0]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z},$$

here with the level γ_S in place of zero and the event reversed, a tail of equation (2.12).

Specialized form. The question is the first line of equation (2.21) with the trend entry replaced by two linear quantities of the state, the loss to date (21.7) and the five-year rate, the mean of Chapter 11's rate $\gamma_u(t_\mathrm{y})$ over the first five years,

$$\gamma_u^{5\mathrm{y}} = \frac{1}{5}\int_0^5 \gamma_u(t_\mathrm{y})\,\mathrm{d}t_\mathrm{y} = \gamma_u + \dot\gamma_u\,(2.5\,\mathrm{yr} - \bar t_\mathrm{y}) = \gamma_u + 1.0\,\mathrm{yr}\cdot\dot\gamma_u, \tag{21.13}$$

with $\bar t_\mathrm{y} = 1.5\,\mathrm{yr}$, a forecast in a curvature drawing and equal to γ_u in the others. Their gradients are the second and third of the pattern (21.4); under each drawing each quantity is normal with variance $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}^{(d)}\mathbf{g}$, and the tail, its Bayes factor and the decision are equations (21.10), (21.9) and (21.11) with the quantity in place of γ_u. The choice that enters here is the reading itself: all three are the specification's words taken differently, and the reported reading is the trend entry, chosen at step 2. The coil's own health index today with the step taken out is $\mathrm{e}^{3\,\mathrm{yr}\cdot\gamma_u^{\mathrm{td}}}$, set against the envelope 0.97^3.

Building blocks. The four drawings' fits; the gradients of the pattern (21.4); the variance relation (21.8) for the loss to date; the tail (21.10); the decision (21.11).

The three readings. Table 21.9 sets them side by side. Each row is a coil and a reading; the columns give the median rate and its band, the probability of a breach, its Bayes factor and the expected consequence of filing.

Table 21.9: The specification read three ways, over the reported family: the trend entry (the reported reading), the loss to date (21.7) and the five-year rate (21.13), each in %/yr by the map (10.1); bands 5–95 %.

Coil	Reading	median [%/yr]	band [%/yr]	P [–]	B [–]	$\mathcal{R}$(file) [R$]
1	the trend entry (reported)	−4.552	−6.883 to −2.173	0.8592	6.10	−46 551.1
1	the loss to date	−4.550	−6.299 to −2.788	0.9261	12.54	−50 567.9
1	the five-year rate	−6.162	−11.388 to −2.487	0.9134	10.55	−49 804.5
2	the trend entry (reported)	−2.832	−5.505 to −0.082	0.4595	0.85	−22 569.0
2	the loss to date	−2.833	−4.996 to −0.616	0.4500	0.82	−21 997.0
2	the five-year rate	−2.909	−6.060 to +0.232	0.4803	0.92	−23 815.4

The loss to date shares the trend entry's median and has the narrower band, by the wander of equation (21.8), so its tail moves away from $\frac{1}{2}$: coil 1's breach rises to 0.9261 and coil 2's falls to 0.4500. On the coil's own health index today, the step taken out, coil 1 stands at 0.8696 (band 0.8227 to 0.9187), below the specification's three-year envelope 0.9127, and coil 2 at 0.9174 (0.8575 to 0.9816), its median inside the envelope. The five-year rate separates the drawings: coil 1's is −7.899 %/yr under $M_{2,15}^{(1)}$, $P = 0.9716$, against −4.303 under $M_{2,15}^{(0)}$, 0.8221; coil 2's runs from −2.642 ($M_{2,15}^{(12)}$, 0.4416) to −3.876 ($M_{2,15}^{(2)}$, 0.6479). Coil 1's possible acceleration strengthens a five-year claim, but only as a forecast on a curvature the record does not decide, and the reported claim does not need it. Every reading files both claims.

21.4 The checks (step 5)

Each answer of Section 21.3 rests on the Laplace normals of four drawings, on Part III's priors and choices, and on one record; the checks set each against the calculation or the variation it stands for.

21.4.1 The full posterior against the Laplace answers

The reweighted draws of Section 21.2.5 keep 98.8, 98.7, 98.7 and 98.5 % of their 20 000 draws as effective samples on $M_{2,15}^{(0)}$, $M_{2,15}^{(1)}$, $M_{2,15}^{(2)}$ and $M_{2,15}^{(12)}$, with largest weights of 9.11×10^{-5}, 1.17×10^{-4}, 1.01×10^{-4} and 1.19×10^{-4}: no draw dominates. The evidences by equation (7.12) lie 0.0067 to 0.0102 below the Laplace ones, −66.1777 against −66.1687 on the step drawing, and the family's weights by those evidences are 0.2632, 0.4388, 0.1170 and 0.1811, against 0.2634, 0.4381, 0.1171 and 0.1815; no weight moves by more than 0.0007. Drawing by drawing, coil 1's probability of a breach is 0.8221 (0.0031), 0.8811 (0.0023), 0.8123 (0.0034) and 0.8763 (0.0020), against the Laplace 0.8221, 0.8829, 0.8181 and 0.8824, and coil 2's 0.5248 (0.0036), 0.4366 (0.0027), 0.4855 (0.0045) and 0.4365 (0.0027), against 0.5222, 0.4262, 0.4910 and 0.4283; the median of γ_2 on the step drawing is −0.03153 against the MAP value −0.03139.

Over the reported family, coil 1's probability is 0.8567 (standard error 0.0014) against the Laplace 0.8592, a gap of −0.0025, 1.8 standard errors, which the draws do not resolve. Coil 2's is 0.4655 (0.0017) against 0.4595, a gap of +0.0060, 3.6 standard errors: resolved, and small: the draws raise coil 2's tail in every drawing, most in the two that free coil 1's curvature. The expected consequences of filing are R$ − 46 399.1 (83.8) and R$ − 22 931.6 (100.4), and the decision does not move. The loss to date gives 0.9254 and 0.4559, against 0.9261 and 0.4500, and the five-year rate 0.9138 and 0.4876, against 0.9134 and 0.4803. Over the whole family, with effective sample sizes from 63.7 to 98.8 %, the probabilities are 0.8678 (0.0009) and 0.6179 (0.0010), against 0.8723 and 0.6147, and a step at window 15 has 0.7803 against 0.7809. The route is named the full posterior, the stated model without the Gaussian approximation up to its Monte Carlo error; the Laplace values are reported.

21.4.2 The curvature term

The chapter adds no physical row, so the limit $\sigma_p \to 0$ is not checked again here: Chapter 12 fitted this drawing's full state, 301 entries, with the physical rows at a finite tolerance and showed it reproduces the reduced state (Section 12.4.2). What remains is the curvature term $\mathbf{C}$ of equation (4.44) that the Gauss–Newton covariance (5.10) drops. The full Hessian of the step drawing's objective at $\hat{\mathbf{z}}$ gives the Newton covariance, in which the spread of γ_1 is 0.01457 against 0.01465, a ratio of 0.9946, and that of γ_2 0.01668 against 0.01673, 0.9968; the probabilities of a breach on it are 0.8234 and 0.5223, against 0.8221 and 0.5222. The two log-determinants differ by 0.02, so the Laplace evidence is unchanged at the precision the weights are reported. The Gauss–Newton covariance is the one reported.

21.4.3 Sensitivity to the priors and the choices

Table 21.10 refits the step drawing $M_{2,15}^{(0)}$ under other priors and choices; the last row varies the curvature prior of the whole reported family. Each row is a variant; the columns give each coil's rate with its spread, its probability of a breach and its Bayes factor, the factor taken against the variant's own prior odds.

Prior effects coil by coil. A prior row centered at the specification pulls each rate, taken alone, toward −3 %/yr. That costs coil 1, which sits beyond it (0.8221 falls to 0.7997 and 0.7531), and helps coil 2, which sits at it (0.5222 rises to 0.5335 and 0.5428). Coil 2's rate does not move toward −3 %/yr but away from it, from −3.090 to −3.119 and −3.116: the prior row on γ_1 pulls coil 1's rate toward the specification, a slower loss, and through the correlation −0.2325 of γ_1 and γ_2 the return air hands that loss to coil 2, which outweighs the pull of coil 2's own prior row. The narrower posterior at nearly the same mean explains only part of coil 2's rise from 0.5222 to 0.5428; the rest is that shift of its mean, through coil 1. The fleet's prior is Chapter 9's posterior predictive of a coil of this design, equation (9.17), with the fleet's mean

Table 21.10: Priors and choices, each refitted on the step drawing $M^{(0)}_{2,15}$ unless said; rates as $100(\mathrm{e}^{\gamma}-1)$ with the spread of γ_u in parentheses; B by equation (21.9) against the variant's prior odds. Every variant files both claims.

Variant	rate 1 [%/yr] (sd [yr^{-1}])	P_1 [–]	B_1 [–]	rate 2 [%/yr] (sd [yr^{-1}])	P_2 [–]	B_2 [–]
the step drawing, no prior row on the rates	−4.303 (0.01465)	0.8221	4.620	−3.090 (0.01673)	0.5222	1.093
a prior row $\mathcal{N}(\gamma_{\mathrm{S}},\ 0.030^2)$ on each rate	−4.062 (0.01309)	0.7997	3.993	−3.119 (0.01456)	0.5335	1.144
a prior row $\mathcal{N}(\gamma_{\mathrm{S}},\ 0.015^2)$ on each rate	−3.687 (0.01039)	0.7531	3.050	−3.116 (0.01111)	0.5428	1.187
the fleet's prior $\mathcal{N}(-0.03022,\ 0.01234^2)$	−3.550 (0.00936)	0.7283	2.764	−3.090 (0.00988)	0.5376	1.199
the rating prior's spread 0.05	−4.266 (0.01446)	0.8182	4.501	−3.220 (0.01653)	0.5547	1.246
$\sigma_q = 0.0025$	−3.950 (0.01119)	0.8104	4.275	−3.424 (0.01355)	0.6266	1.678
$\sigma_q = 0.01$	−4.856 (0.02340)	0.7955	3.889	−2.541 (0.02509)	0.4254	0.740
the straight line, $\sigma_q \to 0$	−3.773 (0.00970)	0.7954	3.887	−3.596 (0.01219)	0.6935	2.263
the reported family, curvature prior 0.02 / 0.10 yr^{-2}	–	0.8426 / 0.8567	–	–	0.4857 / 0.4633	–

rate μ_γ at median $-0.03022\,\mathrm{yr}^{-1}$ (spread 0.00290) and its spread τ_γ at median $0.01199\,\mathrm{yr}^{-1}$: before any record it puts a coil of the design faster than specification with probability 0.4923. The specification sits at the design's mean rate, −2.98 %/yr. The fleet's prior pulls coil 1 to −3.550 %/yr, $B = 2.764$, and leaves coil 2 where it is, $B = 1.199$; it double-counts coil 1 a little, since the fleet's coil 0 is this hall's coil 1. A stiffer path ($\sigma_q = 0.0025$, or the straight line) moves rate from coil 1 to coil 2, and a looser one (0.01) the other way, the trade between the coils that Chapter 12 reported. No variant moves either probability below 0.4254, and every variant files both claims.

21.4.4 Fresh records and the claim rule's operating characteristic

A probability of a breach is one record's draw. The check refits the step drawing and the base fit on 100 fresh records at the stated truth, the same windows and inputs with fresh sensor errors, and reads the claim rule "file when $P_u > 1/12$" on each. Table 21.11 gives the result, the rule's operating characteristic. Each row is a coil and a drawing; the columns give the median of the probability of a breach over the records with its 5–95 % spread, the share of records that file, the share that reach 0.95, and this record's place among them.

Table 21.11: The claim rule over 100 fresh records at the stated truth (coil 1 truly faster than specification, coil 2 truly within), on the step drawing and on the base fit without a step.

Coil	Drawing	P_u, median [–]	5–95 % [–]	filed, $P_u > 1/12$ [–]	$P_u \geq 0.95$ [–]	this record's quantile [–]
1	the step drawing	0.7698	0.3889 to 0.9511	1.00	0.06	0.61
1	the base fit	0.7366	–	1.00	0.04	–
2	the step drawing	0.3337	0.0619 to 0.7752	0.94	0.01	0.74
2	the base fit	0.9567	–	1.00	0.55	–

On the coil that is truly faster the rule files in every record, and the 95 % rule in 0.06 of them. On the coil that is truly within it files in 0.94 of the records, by design: the low threshold accepts cheap failures. Without the step, coil 2 would read a median of 0.9567 and reach 0.95 in 0.55 of the records; the concession is what keeps a coil within specification from reading as a

breach. Over the records coil 2's rate on the step drawing has median −2.300 %/yr (spread −4.221 to −0.483), and the step's $\ln B$ against no step has median +2.423. At this truth the rule costs R$ 4 700 a record on coil 2 and gains R$ 55 000 on coil 1. This record's probabilities sit at the 0.61 and 0.74 quantiles of their spreads: an ordinary draw for coil 1 and a somewhat high one for coil 2.

21.4.5 The answers against the stated truth

The last check sets the answers against the stated truth of the record (Section 21.6.5). Coil 1's true mean rate over the record, with no step, is $-0.05000\,\mathrm{yr}^{-1}$, −4.877 %/yr, faster than specification; its true rate is −2.955 %/yr to window 24 and −13.929 %/yr at window 36, and over the five years −12.890 %/yr. The step drawing's posterior lies +0.41 standard deviations from the true mean rate, and the drawing fitted to the truth's noise-free readings, its pseudo-true value, gives −3.931 %/yr, and −4.225 under $M_{2,15}^{(1)}$: a straight line through a path that bends late reads less than its mean. Coil 2's true fouling rate is $-0.02000\,\mathrm{yr}^{-1}$, −1.980 %/yr, within specification, with a true step of −0.0700 into window 15, −6.76 % of its conductance. The posterior lies −0.68 standard deviations from it, and the pseudo-true value is −2.533 %/yr (−2.264 under $M_{2,15}^{(1)}$): the drawing's straight coil 1 hands part of coil 1's year-3 loss to coil 2 through the return air, Chapter 12's finding. On the truth, coil 1's claim is granted and coil 2's refused, so the plan nets R$ − 50 000; the dashboard's claim nets the same, and the 95 % rule R$ 0.

Table 21.12 collects the routes and the truth. Each row is a route or a check; the columns give each coil's probability of a breach, its expected consequence of filing and the decision.

Table 21.12: The routes to the claim and the stated truth. Probabilities of a breach; expected consequences of filing in reais, negative a gain; the full posterior's Monte Carlo standard errors in parentheses.

Route or check	P_1 [–]	P_2 [–]	$\mathcal{R}$(file), coil 1 [R$]	$\mathcal{R}$(file), coil 2 [R$]	decision [–]
Laplace normals over the reported family (reported)	0.8592	0.4595	−46 551.1	−22 569.0	file, file
full posterior, reweighted draws	0.8567 (0.0014)	0.4655 (0.0017)	−46 399.1	−22 931.6	file, file
the expansion of equation (2.24)	–	–	−55 000	−55 000	file, file
Newton covariance, step drawing	0.8234	0.5223	–	–	–
the record's whole family, Laplace; full posterior	0.8723; 0.8678	0.6147; 0.6179	−47 336.5	−31 879.2	file, file
stated truth: coil 1 at −4.877, coil 2 at −1.980 %/yr	faster	within	−55 000	+5 000	granted, refused

21.5 The report (step 6)

The dashboard shows both coils losing about 15 % a year, five times the warranty's 3 %. On each coil's own conductance, with the March pump change taken out, coil 1 is losing 4.6 %/yr (2.2 to 6.9) and is faster than specification with probability 0.86; coil 2 is losing 2.8 %/yr (0.1 to 5.5), and the record cannot say on which side of 3 % it lies (0.46). File both claims. A claim costs R$ 5 000 and wins R$ 60 000, so it pays from one chance in twelve: coil 2's claim is more likely to fail than to succeed, and it is still worth R$ 22 569 in expectation, while a rule of 95 % confidence would file neither and give up R$ 69 120. The maker's arguments move the probabilities, and none moves either below one in twelve, because one in twelve is low; the

arguments that move loss between the coils trade the two claims, and the season lowers both. No measurement is worth taking before filing: waiting a year costs more than it can tell.

21.5.1 The answers and their bands

Table 21.13 gives the answers in the order of Table 21.1. Each row is a question; the second column gives the reported value, from the Laplace normals of the reported family; the third the full posterior's value beside it; and the fourth the band: the 5–95 % band of a rate, the flip points of a decision, and for a probability its range over the arguments of Table 21.7.

Table 21.13: The report for the two-unit hall at window 36 (Dec Y3), the warranty claims against the maker's specification, from the condition record of windows 0–36 with coil 2's March step conceded. Reported values from the Laplace normals of the reported family; bands 5–95 %.

Asset manager's question	Reported value	Full posterior	Band
Has this unit degraded faster than its specification allows, with enough confidence to make the claim?	coil 1: $P = 0.8592$, $B = 6.10$ [–], rate −4.552 [%/yr]; coil 2: $P = 0.4595$, $B = 0.85$ [–], rate −2.832 [%/yr]	0.8567; 0.4655 [–]	rates −6.883 to −2.173 and −5.505 to −0.082 [%/yr]
Should the claim be filed, at R$ 5 000 to win R$ 60 000?	file both; $\mathcal{R}$(file) = −46 551.1 and −22 569.0 [R$], plan −69 120.1 [R$]; break-even 1/12 [–]	−46 399.1; −22 931.6 [R$]	flips at CF ×10.31 and ×5.51, or CV ×0.097 and ×0.181 [–]
What do the maker's arguments do to the claim: the March step, the drifting sensor, the summer?	the decision holds under every argument; lowest P 0.4247 and 0.1924 [–]; a second drift ($\ln B = -1.4782$) and a seasonal conductance (−4.3224) rejected [–]	–	coil 1: 0.4247 to 0.9997; coil 2: 0.1924 to 0.9923 [–]
Read as the loss so far, or over the specification's five years, is it still faster?	loss to date: $P = 0.9261$ and 0.4500 [–]; five-year rate: 0.9134 and 0.4803 [–]; every reading files both	0.9254, 0.4559; 0.9138, 0.4876 [–]	loss to date −6.299 to −2.788 and −4.996 to −0.616; five-year −11.388 to −2.487 and −6.060 to +0.232 [%/yr]

The answers rest on the claim's reading of Part III's posterior: the trend entry as the specification's quantity, the step conceded, the curvature family carried. Each of the three is a choice the maker could contest, and step 5 and Section 21.3.3 show that none of them changes the decision, while each moves the probabilities.

21.5.2 The entry in the maintenance log

The two-unit hall's maintenance log holds the entries listed in Section 21.1.2. This chapter's entry is also at window 36. It uses nothing later than window 36, and in the values of Table 21.13 it reads:

> *Window 36 (Dec Y3), warranty claims against the coil maker's specification (loss of the coil's own conductance no faster than* 3 %/yr *over the first five years), from the condition record (windows 0–36); coil 2's step at the pump change (Mar Y2) conceded.*
> Coil 1: −4.6 %/yr (−6.9 to −2.2); faster than specification with $P = 0.86$ (Bayes factor 6.1). Claim filed: expected consequence R$ − 46 551.
> Coil 2: −2.8 %/yr (−5.5 to −0.1) with the March step taken out; $P = 0.46$ (Bayes factor 0.85). Claim filed: expected consequence R$ − 22 569; break-even $P = 0.083$.
> Evidence pack: each coil's own conductance from the whole record; the March step excluded; unit 1's discharge drift carried; the season in the known inputs. The maker's arguments

tested: none takes either claim below the break-even.
Not used: the dashboard's 15.8 and 14.8 %/yr (a tile on the hall's approach, not the coil's conductance).
Next measurement: none before filing. The water-side inspection of coil 2 (window 36's change-point entry) stays open for the hall, not for the claim. The side-by-side check of unit 1's discharge sensor stays open.

The entry names what the claim concedes and what it carries, so that the maker's answer can be read against it. It does not name the cause of the March step, which the condition record cannot show and the inspection has not yet found.

Common mistake

Claiming on the dashboard's slope, or not claiming without 95 %. A tile's slope is the load, the season, the other coil and the sensors as well as the coil, and the specification names the coil's own conductance. And how confident a claim must be is set by what it costs to file against what it returns, not by a convention: here one chance in twelve.

21.6 The next measurement (step 7)

The report leaves coil 1 probably faster than specification and coil 2 on the threshold, and it files both. A measurement is worth taking for the claim only if its outcome could change that decision, by enough to repay what it costs or how long it delays the filing.

21.6.1 The preposterior of a candidate

The preposterior is read on the step drawing $M^{(0)}_{2,15}$, one normal per coil. A candidate a is a set of rows with Jacobian lines $\mathbf{J}_a$ in the columns of the state; its preposterior covariance is equation (5.22), $\boldsymbol{\Sigma}_a = (\boldsymbol{\Sigma}^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a)^{-1}$, and with it the spread of each trend entry after the measurement, $\varsigma'_u = (\mathbf{e}_{\gamma_u}^{\mathsf{T}}\boldsymbol{\Sigma}_a\,\mathbf{e}_{\gamma_u})^{1/2}$. Before its outcome is known, the posterior mean that the measurement will leave is normal about today's, $\hat{\gamma}'_u \sim \mathcal{N}(\hat{\gamma}_u,\ \varsigma_u^2 - \varsigma_u'^2)$, the preposterior of a linear row, and the probability of a breach after it is $P'_u = \Phi((\gamma_{\mathrm{S}} - \hat{\gamma}'_u)/\varsigma'_u)$. With filing the action today, the measurement's value is what it saves by the claims it would stop, the expected value of sample information,

$$\mathrm{EVSI}_u = \mathbb{E}\big[\max(0,\ \mathrm{CF} - \mathrm{CV}\,P'_u)\big], \qquad \Pr(\mathrm{flip}) = \Pr(P'_u < p_{\mathrm{be}}),$$
$$\mathcal{R}(\mathrm{wait}) = \mathcal{D}(\Delta_{\mathrm{d}})\,\big[\mathcal{R}(\mathrm{now}) - \mathrm{EVSI}_u\big], \qquad \mathcal{R}_{\mathrm{proof}}(\mathrm{wait}) = \mathcal{D}(\Delta_{\mathrm{d}})\,\Pr(P'_u \geq 0.95)\,(\mathrm{CF} - \mathrm{CV}), \tag{21.14}$$

the expectation over the normal of $\hat{\gamma}'_u$ by Gauss–Hermite quadrature with 80 nodes. A candidate that takes time delays the filing by Δ_{d} years and discounts its result by $\mathcal{D}(\Delta_{\mathrm{d}})$; its cost of waiting is $\mathcal{R}(\mathrm{wait}) - \mathcal{R}(\mathrm{now})$. The last expression is the value of waiting under a standard of proof, by which a claim is filed only once the evidence reaches 0.95 and is then granted. On the step drawing $\mathcal{R}(\mathrm{now})$ is R$ $-44\,323.9$ for coil 1 and R$ $-26\,331.8$ for coil 2, and a year later, discounted and with no new information, R$ $-41\,040.6$ and R$ $-24\,381.3$.

The candidates are of four kinds. More record: a fourth year of monthly windows, 37 to 48, or the record continued to the specification's end, window 60, each read at the extrapolated state, and then a decision. The permanent instrument at the place of the question: water-side heat meters on both coils, of standard deviation 3.0 kW as in Chapter 5, read with a fourth year. The logger's archive: four more steady hours in each window from December of year 1 to June of year 2, windows 12 to 18, where the step sits, or one more hour of every window. And a reference condition test now, with reference sensors of 0.1 K on the return air, both discharges and rack 1's inlet, free of drift, and its own load on the meter, at the price of one condition test. A water-side inspection of coil 2 asks a different question, the cause of the step.

21.6.2 The candidates

Table 21.14 gives the candidates. Each row is a candidate; the columns give each coil's spread of γ_u after it, with the share of today's that it removes, its value to each claim by equation (21.14), its cost of waiting, and the probability that the evidence would reach 0.95; the candidate's price, where it is stated, follows its name.

Table 21.14: The candidate measurements for the claim, from the preposterior on the step drawing, equation (21.14); today's spreads are 0.01465 and 0.01673 yr^{-1}. Values in reais; "1 / 2" gives coil 1 and coil 2.

Candidate	ς_1' [yr^{-1}] (cut [%])	ς_2' [yr^{-1}] (cut [%])	EVSI, 1 / 2 [R$]	waiting cost, 1 / 2 [R$]	$\Pr(P' \geq 0.95)$, 1 / 2 [–]
a fourth year (windows 37–48), then decide; a year's delay	0.01131 (22.8)	0.01234 (26.2)	1.4 / 118.6	3 281.9 / 1 840.7	0.2924 / 0.0432
the record to the specification's end (to window 60); two years' delay	0.00938 (36.0)	0.01005 (39.9)	20.7 / 418.2	6 305.6 / 3 398.0	0.4328 / 0.1218
a fourth year with water-side heat meters on both coils, permanent; a year's delay, meters not priced	0.01017 (30.6)	0.01171 (30.0)	8.2 / 197.3	3 275.7 / 1 767.8	0.3806 / 0.0627
the logger's archive, windows 12–18, four more hours each; not priced	0.01445 (1.3)	0.01467 (12.3)	0.0 / 4.7	–	0.0000 / 0.0020
the logger's archive, one more hour of every window; not priced	0.01282 (12.5)	0.01419 (15.2)	0.0 / 14.4	–	0.1434 / 0.0057
a reference condition test now, R$ 1 500	0.01199 (18.1)	0.01495 (10.6)	0.2 / 2.4	–	0.2307 / 0.0008
a water-side inspection of coil 2; not priced	–	–	0 / 0	–	–
perfect information on each coil (EVPI)	–	–	704.1 / 2 702.6	–	–

The candidates' values are set against their costs in Figure 21.6. Each row is a candidate; the horizontal axis is in reais on a logarithmic scale; the filled marks are its value to each coil's claim and the open marks its cost, the cost of waiting for the candidates that delay the filing and the price of the condition test; the dashed lines are the values of perfect information. Every value lies to the left of its candidate's cost, most of them by more than a factor of ten, and below the line of perfect information.

The surprise of the step. The evidence is weak and the decision is not close, so no measurement pays. The best candidate, two more years of record, is worth R$ 418.2 to coil 2's claim and costs R$ 3 398.0 in delay; the reference condition test is worth R$ 2.6 to both claims together against its R$ 1 500; the largest value of any candidate to one claim is R$ 418.2. The decision flips after a candidate with probability 0.0009 and 0.0555 for a fourth year, 0.0093 and 0.1338 for the record to window 60, 0.0044 and 0.0761 with the heat meters, and 0.0020 or less for coil 2 after the archive's four hours or the reference test. The value of a measurement to a decision lives in the outcomes that flip it, and a claim that pays from one chance in twelve has few.

The water-side inspection. It is worth nothing to the claim, and provably so. It would settle what stepped in March, not how fast the coil fouls: across every drawing of the record's

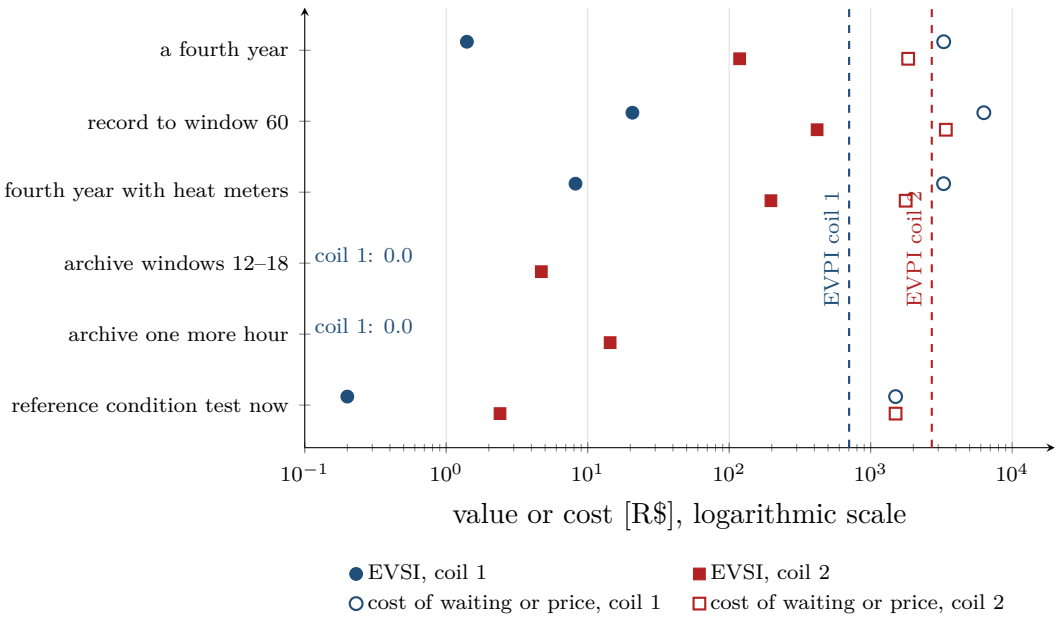


Figure 21.6: The value of each candidate to each claim (filled, by equation (21.14)) against its cost (open): the cost of waiting a year or two for the record candidates, the price of one condition test for the reference test. The archive candidates have no stated price, and their value to coil 1's claim is 0.0. Dashed: the expected value of perfect information on each coil, R$ 704.1 and R$ 2 702.6. The water-side inspection, worth R$ 0 to the claim, is not drawn.

whole family coil 2's probability of a breach lies between 0.4262 and 0.9923, all above the break-even, so no finding about the step can flip the decision. It settles the cause for the hall, the change-point entry's question, not the rate for the claim.

Under a standard of proof. If the maker grants only on evidence at 0.95, neither claim is grantable now, and waiting has a value. Continuing the record to window 60 gives coil 1 a 0.4328 chance of reaching 0.95, $\mathcal{R}_{\text{proof}}(\text{wait}) = \text{R\$}\ -20\,409.1$, and coil 2 0.1218, R$ −5 742.6; a fourth year alone gives R$ −14 890.0 and R$ −2 201.0, and a fourth year with permanent heat meters 0.3806 and 0.0627, R$ −19 380.8 and R$ −3 191.9.

21.6.3 The recommendation

File both claims now, with the evidence pack of the log entry. Take no measurement for the claim; the record keeps running for the reasons of Chapters 14 and 17, not this one. If the maker answers with a standard of proof, read the claim again at window 48 and at window 60, before the specification's five years end. The permanent heat meters add little to a fourth year for the claim, 0.3806 against 0.2924 on coil 1; what they are worth to the crossing date is Chapter 23's question.

21.6.4 What the analysis showed

The dashboard said both coils were losing 15.76 and 14.82 % a year, five times the specification, with certainty; on each coil's own conductance they are losing 4.6 and 2.8 %/yr, and coil 2 is within specification at the median. The physics takes the load, the season, the other coil and the sensor out of the tile, the step entry takes out the pump change, and the curvature family carries coil 1's undecided acceleration. What the probability adds is what the chart cannot give: a Bayes factor of the specification, 6.1 and 0.85, and a claim threshold set by the costs, one in

twelve, not by a convention. The maker's arguments are drawings of the same graph, and the record weighs against the ones that would hurt most: a step on both coils cuts coil 1 to 0.5267, about as far as a second drift (0.5201), and a fit with no drift entry cuts coil 2 to 0.1924, the lowest of all. None crosses the break-even, because the break-even is low; the arguments that move loss between the coils trade the two claims, and the season lowers both. The analysis names what it did not settle: coil 2's side of the specification: even the record continued to the specification's end would show it faster at 0.95 in only 0.12 of its outcomes.

21.6.5 Limits of the model

A constructed record. The condition record is generated from a stated truth, with the stated sensor errors, and step 5 compares the answers with that truth (Section 21.4.5). It is not field data, and the agreement is a check of the method on its own model, not of the model against a real hall.

The grant rule. The maker is taken to settle on the coil's own conductance, read as the truth. A maker who settles on the evidence pack is the standard-of-proof row, and a maker who disputes the specification's measure, or the conductance as its reading, is outside the model.

The specification's quantity. The specification is read as the trend entry over the service to date. The five-year reading is a forecast on the curvature family, and the loss to date leaves out the wander; both are given, and neither changes the decision.

The conceded step and the curvature. The step is taken at window 15, permanent and single, with Chapter 12's limits, and the curvature as linear in time over the whole record, with Chapter 11's.

The costs. Only the cost of filing and the value of a grant enter. The relationship with the maker, the coil's downtime during a replacement and the claim's effect on the capacity forecast of Chapter 14 are not priced.

Part III's limits. The wander's σ_q is chosen, the drift is linear from its onset, the readings are independent across windows, and Chapter 5's steady state and known inputs hold in every window.

The next chapter. Three years of monthly windows know a coil's rate to a standard deviation of 0.01465 and $0.01673\,\mathrm{yr}^{-1}$, against a specification of $-0.030459\,\mathrm{yr}^{-1}$. Whether the monitoring can resolve a loss of a few percent a year in the time an owner has is the question of Chapter 22 (Section 22.3.1).

21.7 What to retain

Definition 21.1 (Specification compliance and a warranty claim)**.** A specification on a degradation rate is a threshold on one entry of the state, $\gamma_u < \gamma_{\mathrm{S}}$ with $\gamma_{\mathrm{S}} = \ln(1 - x)$ for a loss of x a year, read on the coil's own parameter, never on a performance tile. Its probability is the posterior tail P_u beyond the threshold, mixed over the drawings the record leaves open, and its Bayes factor is the tail's odds over the prior's, equation (21.9), even prior odds when the rate has no prior row or a prior centered on the threshold. A claim that costs CF to file and returns CV when granted has the expected consequence $\mathrm{CF} - \mathrm{CV}\,P_u$ and is filed when $P_u > \mathrm{CF}/\mathrm{CV}$; the break-even replaces an assurance level, and the value of perfect information is $\mathrm{CF}(1 - P_u)$.

Three conclusions carry beyond the hall. Compliance with a specification is a question about a parameter of the physics, and the fit that separates the parameter from the load, the season, the other unit and the sensors is what makes the question answerable at all. The confidence a decision needs is the costs' to set, and a cheap claim with a large return pays at a probability far below any assurance level. And an adversary's arguments are other drawings of the same plant: each can be fitted, weighed by its evidence and carried to the decision; where units share a reading, an argument that moves loss between them trades their claims, while one that lowers both, such as the season, does not.

Table 21.15 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 21.15: The terms of a warranty claim against a specification and where each is defined. Units for the two-unit hall.

Term	Expression	Units	Meaning	Defined in
Specification as a threshold	$\gamma_u < \gamma_S = \ln 0.97$, eq. (21.1)	yr^{-1}	a breach is a tail of one entry	21.1.1
Loss of each action	eq. (21.5)	R$	a step in the state	21.1.4
Loss to date against the rate	eqs. (21.7), (21.8)	yr^{-1}	the same mean; the wander between them	21.2.3
Bayes factor of a specification	eq. (21.9)	–	posterior odds over prior odds	21.2.4
Questions and answers				
Probability of a breach	eq. (21.10)	–	mixed over the reported family	21.3.1
Break-even probability	$p_{\mathrm{be}} = \mathrm{CF}/\mathrm{CV}$, eq. (21.11)	–	replaces an assurance level	21.3.2
The plug-in's failure	eq. (2.24) on a step loss	R$	returns $L(\mathrm{file}, \hat{\mathbf{z}})$	21.3.2
EVPI of a claim	$\mathrm{CF}(1 - P_u)$, eq. (21.12)	R$	the cost of a claim that fails	21.3.2
Conceded event	coil 2's step kept out of γ_2	–	the claim excludes it	21.1.4
Counter-argument	a drawing M_{arg}, refitted	–	its $\ln B$ and its tail	21.3.3
Claims that trade	$\mathrm{Cov}(\gamma_1, \gamma_2 \mid \mathbf{y}) < 0$	–	a shared reading; an argument that moves loss between the coils trades them	21.3.3
Five-year rate	eq. (21.13)	yr^{-1}	a forecast on the curvature	21.3.4
Operating characteristic	share of fresh records that file	–	power and false claims	21.4.4
EVSI against waiting	eq. (21.14)	R$	a measurement must repay its delay	21.6.1
A measurement that cannot flip	every drawing's $P_u > p_{\mathrm{be}}$	R$	worth nothing to the decision	21.6.2

21.8 Exercises

Exercise 21.1. *Formulation.* Write the question "has coil 1's own health index, the step taken out, fallen below the specification's three-year envelope 0.97^3?" in the form of Table 21.1, and evaluate it for both coils. (A prediction under a threshold, equation (2.12): $\mathrm{e}^{3\,\mathrm{yr}\cdot\gamma_u^{\mathrm{td}}} < 0.97^3$ is $\gamma_u^{\mathrm{td}} < \gamma_S$, with γ_u^{td} of equation (21.7); over the reported family $\Pr(\gamma_1^{\mathrm{td}} < \gamma_S \mid \mathbf{y}) = 0.9261$ and $\Pr(\gamma_2^{\mathrm{td}} < \gamma_S \mid \mathbf{y}) = 0.4500$.)

Exercise 21.2. *Derivation.* Show that the Bayes factor of the halves $\gamma < \gamma_S$ and $\gamma \geq \gamma_S$ is the posterior odds over the prior odds, and evaluate it for coil 2 on the step drawing, with no prior row and with the fleet's prior. (Each cut prior is $p(\mathbf{z})\mathbf{1}[\cdot]/P^0$, so each evidence is $p(\mathbf{y})$ times the

posterior mass of its half over its prior mass, and $p(\mathbf{y})$ cancels, equation (21.9). With no prior row, $z = +0.0557$, $P = 0.5222$ and $B = 1.0929$; with the fleet's prior, $P = 0.5376$, prior odds 0.9695 and $B = 1.1992$.)

Exercise 21.3. *Derivation.* Show that a claim pays exactly when $P > \mathrm{CF}/\mathrm{CV}$, and find the claim value below which coil 2's claim stops paying. ($\mathcal{R}(\text{file}) - \mathcal{R}(\text{not}) = \mathrm{CF} - \mathrm{CV}\,P < 0 \iff P > \mathrm{CF}/\mathrm{CV}$; with $P_2 = 0.4595$ the claim value at which it vanishes is $\mathrm{CF}/P_2 = \text{R\$}\,10\,882$, and coil 1's is R$ 5 819.)

Exercise 21.4. *Application.* Refit the step drawing with the fleet's population prior on each rate, and say why coil 1's Bayes factor falls and coil 2's barely moves. (The prior $\mathcal{N}(-0.03022,\ 0.01234^2)$ is centered at the design's mean rate, almost at the specification, and it pulls each rate toward it: coil 1, beyond the specification, is pulled in, and coil 2, at it, stays. Coil 1: $P = 0.7283$, $B = 2.7644$; coil 2: $P = 0.5376$, $B = 1.1992$; prior probability of a breach 0.4923.)

Exercise 21.5. *Application.* Read the claim as of window 30 (Jun Y3) and decide each claim. (Refit the step drawing on windows 0–30: coil 1's $P = 0.5539$, $\mathcal{R}(\text{file}) = \text{R\$}\ -28\,237.0$; coil 2's $P = 0.4858$, $\mathcal{R}(\text{file}) = \text{R\$}\ -24\,145.9$: both filed. As of window 31, coil 1's 0.4247, R$ −20 479.2, and coil 2's 0.5860, R$ −30 158.7.)

Part VI

The Value of Condition Monitoring

Chapter 22

Detectability of Degradation

22.1 The dashboard's power analysis and the coil's

The two-unit hall is three years past its handover. Today is window 36 (Dec Y3), early summer in the southern hemisphere where the hall stands, and the asset manager is writing the condition-monitoring plan on which the work orders of the October shutdown of year 4, and any later warranty claim, will rest: the trend log as it stands, five sensors read for one representative hour a month. The plan must say whether that log can see a coil losing 5 % of its conductance a year in the time left before the shutdown, and how long a record that takes, because a coil's record starts again at every new commissioning test. The planner runs the power analysis the dashboard supports. The least-squares line through each coil's 37 health-index tiles scatters about them with a residual standard deviation of 0.0820, and the slope of such a line over windows 0 to t_{as} has a standard error of that scatter over the root of the sum of the squared times about their mean, the residuals taken as independent: 0.0152 per year on a record to window 36 and 0.0109 on one to window 45 (Sep Y4). Reading a 5 %/yr loss as the tile falling 0.05 a year, the planner's one-sided test at the 95 % level concludes "losing" with probability 0.951 on the first record and 0.998 on the second, and a record of 34 windows (Oct Y3) reaches a power of 0.9. The conclusion the chart invites is that the existing sensors are more than enough: the trend log sees a 5 %/yr loss in under three years, well before the shutdown.

The chart cannot say whether its test means anything on a coil that is not losing, because the tile also moves with the load's growth and the season. It cannot say whether 5 % of conductance is 5 % of the tile, when the approach magnifies a coil's change and mixes in the other coil's. Nor can it say how much of the tiles' scatter is reading error and how much the path's own wander from month to month, which a longer record averages in a different way; what unit 1's drifting discharge sensor costs coil 1's record; or what becomes of the test when the planner refits every month and acts on the first verdict. The question of this chapter is

the probability that the hall's trend log, read by Chapter 10's fit, concludes that a coil is losing conductance by the shutdown when its true loss is 5 %/yr*, as a curve over the record's length and the sensors' grade, with the record length a stated power needs and the smallest loss the record resolves.*

It is answered with the method of Chapter 4, on Chapter 10's stacked state, evaluated before the record is logged.

Jargon: probability of detection, POD curve, record length for a stated power, minimum detectable rate, false-alarm probability of a trend test, repeated looks, sensor grade

The *probability of detection* (POD) of a monitoring plan is the probability that the record it will log concludes that a coil is losing, when the coil's true loss is a stated one: the power of the trend test, the operating characteristic of the plan at that loss. A *POD curve* plots it against the record's length or against the loss. The *record length for a stated power* is the shortest record whose probability of detection reaches that power; for a power of 0.9 it is written T_{90}, the counterpart for a rate of the crack size a_{90} of nondestructive testing. The *minimum detectable rate* is the smallest loss that a record of given length resolves at the stated power. The *false-alarm probability* of a trend test, its size, is the probability that it concludes "losing" on a coil that is not losing. *Repeated looks* read the refits of a growing record, one per month, as a sequence of verdicts, each refit a look. The *sensor grade* is the standard deviation of a sensor's reading error.

Key idea

A monitoring plan is judged before it is logged by two numbers read together: how often it concludes when the coil is losing, its probability of detection, and how often when it is not, its false alarm. Read on the coil's own parameter, both come from one quantity, the predictive spread of the rate $\varsigma_u(t_{\mathrm{as}})$, which the rows fix before a reading is taken. A test on raw performance has a probability of detection but no meaningful false alarm, because the load and the season move the raw figure with no loss at all.

22.1.1 Detection as a preposterior

Chapter 10 read each coil's degradation rate as the trend entry γ_u of a record already logged, with the band its posterior gives. Chapter 11 asked a different question of fresh records drawn at the stated truth, how often its evidence test on the curvature would pass a stated level, and found that power low: 0.080 of fresh records reach $\ln B > 1$ (Section 11.4.5). That was the power of one test at one point, one truth and one record length. A monitoring plan needs the power as a curve, over the record's length and the sensors' grade, for the quantity the plan is meant to see, the rate; and it needs it before the record is logged.

Chapter 2's design form, equation (2.16), values a candidate measurement a before its outcome is known by the expected value of a quantity U_{p} of the posterior the outcome would produce. With U_{p} the indicator that an acceptance rule (2.12) is met, and the outcomes weighted by their distribution given a stated value z_i^{t} of one entry of the state rather than by their predictive, the form becomes the operating characteristic of the measurement, as *Reliability Engineering by Derivation*, Chapter 27, derives it for an acceptance test:

$$\begin{aligned} \mathrm{OC}(z_i^{\mathrm{t}}; a) &= \int \mathbf{1}\left[\Pr(f(\mathbf{z}) \geq f^* \mid \tilde{\mathbf{y}}_a) \geq \kappa\right] p(\tilde{\mathbf{y}}_a \mid z_i = z_i^{\mathrm{t}})\, \mathrm{d}\tilde{\mathbf{y}}_a, \\ p(\tilde{\mathbf{y}}_a \mid z_i = z_i^{\mathrm{t}}) &= \int p(\tilde{\mathbf{y}}_a \mid \mathbf{z})\, p(\mathbf{z}_{-i} \mid z_i = z_i^{\mathrm{t}})\, \mathrm{d}\mathbf{z}_{-i}, \end{aligned} \tag{22.1}$$

with $\tilde{\mathbf{y}}_a$ the outcome of a before it is taken and $\mathbf{z}_{-i}$ the other entries of the state, drawn from their prior rows. Averaged over a distribution of z_i^{t} the operating characteristic returns the design form; here the stated value is an input of the question, not a belief.

For a monitoring plan the measurement is the record itself: a is a record of length t_{as}, windows 0 to t_{as}, and $\tilde{\mathbf{y}}_{0:t_{\mathrm{as}}}$ its outcome. The event of the rule is a loss, $f(\mathbf{z}) = -\gamma_u$ and $f^* = 0$, and the stated value is the rate to detect, $\gamma_u = \gamma_{\mathrm{D}}$. The probability of detection is

$$\mathrm{POD}_u(t_{\mathrm{as}}) = \int \mathbf{1}\left[\Pr(\gamma_u < 0 \mid \tilde{\mathbf{y}}_{0:t_{\mathrm{as}}}) \geq \kappa\right] p(\tilde{\mathbf{y}}_{0:t_{\mathrm{as}}} \mid \gamma_u = \gamma_{\mathrm{D}})\, \mathrm{d}\tilde{\mathbf{y}}_{0:t_{\mathrm{as}}}, \tag{22.2}$$

and the false-alarm probability $\mathrm{PFA}_u(t_{\mathrm{as}})$ is the same integral at $\gamma_u = 0$. The two are one form at two values of the rate, so a plan that states one and not the other has stated half of its operating characteristic.

> **Principle 22.1.** *Judge a monitoring plan by its probability of detection and its false-alarm probability together, both read on the parameter. A test whose false alarm is not stated has no power worth quoting.*

22.1.2 The two-unit hall through a record not yet logged (step 0)

The plant graph is Chapter 10's (Figure 10.1), cited and not derived again: Chapter 5's window copied for every window of the record, each copy with its free load Q_t and log-derates $g_{1,t}$, $g_{2,t}$ under the five physical rows (5.2) at the window's known inputs $T_s(t)$ and $Q_o(t)$; one trend row per coil between consecutive windows, equation (10.5), linked to the trend entries γ_1 and γ_2; the recirculation fraction β_1, shared by every window; and the drift node $\dot{b}_{d,1}$ on unit 1's discharge reading from window 20, equation (10.2). The modeling choices of the graph are Chapter 5's paragraph, carried through Chapter 10: steady state within each window, the airflow pinned at the units' fixed fan speed with nothing controlling the discharge air, the chilled-water supply and the other heat as known inputs, and one effective conductance per coil.

What the question adds. The windows run to a free last window t_{as}, Chapter 10's as-of window. Windows 0 to 36 are the record as logged, and windows 37 to t_{as} are windows not yet read; for a record restarted at a new commissioning test, after a replacement or a cleaned coil's re-test, every window from its window 0 is a window not yet read. Above each coil's trend entry stands one question node, the verdict $\mathbf{1}[\Pr(\gamma_u < 0 \mid \mathbf{y}_{0:t_{\mathrm{as}}}) \geq \kappa]$. No row is added: the rule is a stated input of the question, not evidence about the plant.

The stated inputs of the question. Four inputs state the question, and none is a quantity of the fit. The loss to detect is 5 % of conductance a year on each coil, the trend entry

$$\gamma_{\mathrm{D}} = \ln 0.95 = -0.05129\,\mathrm{yr}^{-1}. \tag{22.3}$$

The rule concludes that coil u is losing when $\Pr(\gamma_u < 0 \mid \mathbf{y}_{0:t_{\mathrm{as}}}) \geq \kappa$, with $\kappa = 0.95$, read once, at the last window of the record. The power asked is $\mathrm{POD}^* = 0.9$. And the time available is set by the October shutdown of year 4, window 46: the last window read before it is window 45 (Sep Y4), the refit that Chapter 11's step 7 and Chapter 14's entry ask for before the work order is written. One question adds a stricter rule: a Bayes factor of at least 19 for a coil's trend against no trend, with the trend's prior $\mathcal{N}(0,\ 0.10^2)$ in yr^{-1}.

The maintenance log so far. The two-unit hall's log holds the commissioning test (window 0); the entries of Chapters 5 and 6 at window 18 (Jun Y2); Chapter 7's at window 25 (Jan Y3), with its side-by-side check of unit 1's discharge sensor asked for and still open; and at window 36 the rates (Chapter 10), the acceleration left undecided (Chapter 11), coil 2's step of March of year 2 (Chapter 12), the season (Chapter 13), the forecast of the capacity limit with its refit at window 45 (Chapter 14), coil 1's mission assessment (Chapter 15) and the warranty claims (Chapter 21). This chapter's entry is dated window 36 as well, and it asks what the trend log itself can see. The record as logged enters step 5 only: the curve is a property of the record's design, not of its readings.

The modeling choice at the graph. One choice belongs to this chapter's graph: *the future windows repeat the record's design*, with the same known inputs month by month, the same load growth, 110 kW at commissioning and 120 kW in January of year 1 rising 4 % a year, the same five sensors and one representative hour a month. This buys a curve over t_{as} that any coil of this design on this hall can read, the present coils and a restarted record alike. It fails if the logging or the load changes; how often to log is Chapter 24's question.

The graph. Coil 1's trend layer of the graph, carried to the free last window, is drawn in Figure 22.1. The middle row is coil 1's path: its log-derate in window 0, window 1, window 20, window 36, window 37 and the last window t_{as}, with dots for the runs of windows between them, solid for the windows logged and dashed for those not yet read, as the key at the top left says. The black squares are trend rows. Below each path node is its window's block, with its load, physical rows and five readings; the blocks of windows not yet read are dashed. Above the path stands the trend entry γ_1, linked to every trend row, and above it the verdict node with the rule's input box beside it. The shared fraction β_1 and the drift rate sit below the blocks.

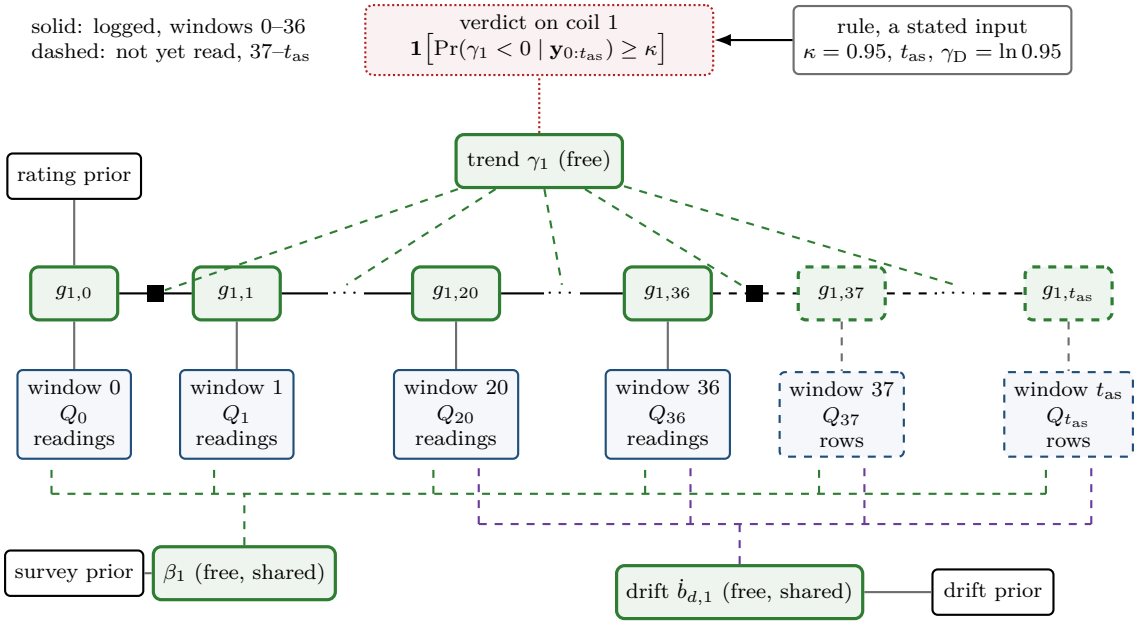


Figure 22.1: Coil 1's trend layer of the condition record's graph (Figure 10.1), carried to a free last window t_{as}. Free (green): the log-derate $g_{1,t}$ of every window, in solid boxes for the windows logged (0 to 36) and dashed boxes for windows not yet read (37 to t_{as}, or every window of a restarted record); each window's load Q_t, inside its block with Chapter 5's physical rows and five readings; the trend entry γ_1, linked to every trend row (black squares; dots for the runs of windows between the drawn ones); the shared fraction β_1, linked to every window's rack-1 reading; and the drift rate $\dot{b}_{d,1}$ on unit 1's discharge reading from window 20 (purple). Pinned in every window, not drawn: the ratings $\bar{G}_u$ and the capacity rates C_u. Known inputs, not drawn: $T_s(t)$ and $Q_o(t)$ of every window, repeated by month. Observed: the five readings of each window logged; a window not yet read carries the rows it will have, with outcomes not yet known. Priors (solid gray): the rating prior on window 0's coil, the survey prior and the drift prior. Red dotted: the verdict node on γ_1, fed by the stated rule (white box); it adds no row. Coil 2's layer is the same without the drift node.

22.1.3 The asset manager's questions (step 1)

The asset manager asks five questions of the plan, and step 1 writes each as a probability over the outcome of a record not yet logged, with the stated rate on the right of the bar. Table 22.1 lists them. The first column is the question in the asset manager's words, the first of them the drill question; the second the general form of its category from Chapter 2; the third its mathematical form for the hall; the fourth its category; the fifth the computation of Section 2.3.6 that answers it; and the last the subsection where it is answered. Both coils are answered. Whether the coils are losing today is Chapter 10's answer, read on the record; whether coil 1's rate changes, Chapter 11's; whether a loss exceeds the maker's 3 %/yr, Chapter 21's; which sensor most narrows the crossing date, Chapter 23's; and how often to log, Chapter 24's. Which instrument shortens the record a stated power needs asks nothing new of the plan; it is step 7, Section 22.6.

Table 22.1: The questions asked of the hall's condition-monitoring plan at window 36 (Dec Y3). $\tilde{\mathbf{y}}_{0:t_{\text{as}}}$ is the outcome of a record to window t_{as}, not yet logged; γ_{D} the stated loss (22.3); $\kappa = 0.95$ the rule's assurance and $\text{POD}^* = 0.9$ the power asked; window 45 (Sep Y4) the last window read before the October shutdown.

Asset manager's question	General form	Mathematical form	Category	Computation	Answered in
Can the existing sensors resolve a 5 % per year loss in the time I have?	(2.16) with the indicator of the rule (2.12) as the utility, conditioned on the rate: (22.1)	$\text{POD}_u(45) = \Pr\big(\Pr(\gamma_u < 0 \mid \tilde{\mathbf{y}}_{0:45}) \geq \kappa \mid \gamma_u = \gamma_{\text{D}}\big)$, and its curve over t_{as}	design	preposterior	22.3.1
How long a record does a power of 0.9 need, and what would better sensors save?	the same, solved for the record	$T_{90,u} = \min\{t_{\text{as}} : \text{POD}_u(t_{\text{as}}) \geq \text{POD}^*\}$, over the temperature sensors' grade σ_T and over (κ, POD^*)	design	preposterior	22.3.2
What is the smallest loss the record resolves by the shutdown?	the same, solved for the stated rate	the smallest $\lvert\gamma\rvert$ with $\text{POD}_u(45) \geq \text{POD}^*$ at $\gamma_u = \gamma$, in %/yr	design	preposterior	22.3.3
If I refit every month and act on the first verdict, when does it come, and how often is it false?	the same, with the rule read at every look	$\Pr\big(\exists\, t_{\text{as}} \leq 45 : \Pr(\gamma_u < 0 \mid \tilde{\mathbf{y}}_{0:t_{\text{as}}}) \geq \kappa \mid \gamma_u\big)$ at γ_{D} and at 0; the first such t_{as}	design	preposterior	22.3.4
Does a stricter rule, a Bayes factor of 19 against no trend, change the answer?	(2.16) with $\mathbf{1}[B \geq 19]$ as the utility, B by (2.14)	$\Pr\big(B(\text{trend} : \text{none}) \geq 19 \mid \gamma_u = \gamma_{\text{D}}\big)$ on $\tilde{\mathbf{y}}_{0:t_{\text{as}}}$; $T^B_{90,u}$; its value at $\gamma_u = 0$	design; hypothesis comparison	preposterior; evidence	22.3.5

22.1.4 The evidence as rows (step 2)

Step 2 attaches the candidate record to the graph as rows. They are Chapter 10's rows, cited and carried to the last window t_{as}: each window's five reading rows, equation (5.7) with unit 1's discharge row (10.4), at 0.3 K for the four temperatures and 3.0 kW for the load meter; the rating prior $g_{u,0} \sim \mathcal{N}(0,\ 0.15^2)$ on each commissioning entry; the survey prior $\beta_1 \sim \mathcal{N}(0.20,\ 0.08^2)$; the drift prior, a row $\dot{b}_{d,1}/\sigma_{\dot{b}}$ with $\sigma_{\dot{b}} = 2.0$ K/yr; and the trend rows (10.5) at $\sigma_q = 0.005$ between consecutive windows. The residual vector is equation (10.8) with its last window 36 replaced by t_{as}, on the stacked state (10.3) carried to window t_{as}. For the record to window 45

the state has 142 entries, the 46 windows' $(Q_t, g_{1,t}, g_{2,t})$, β_1, γ_1, γ_2 and $\dot{b}_{d,1}$, under 324 rows: 230 readings, two rating priors, one survey prior, 90 trend rows and one drift prior, 182 degrees of freedom.

The hypothetical state. A record not yet logged has no readings to fit, and its rows are evaluated at a stated point instead, the hypothetical state $\mathbf{z}^{\mathrm{D}}_{t_{\mathrm{as}}}$: both coils on the 5 %/yr line from their rating, $g_{u,t} = \gamma_{\mathrm{D}}\, t\, \Delta t$, so that $g_{u,45} = -0.19235$; the loads on the record's growth, 110 kW at commissioning and 138.560 kW at window 45; β_1 at the survey's 0.20; the trend entries at γ_{D}; and no drift. The readings do not enter: the Jacobian of these rows depends on the state at which it is evaluated, not on the values read. With the columns $(Q_{45}, g_{1,45}, g_{2,45}, \beta_1, \dot{b}_{d,1}, g_{1,44}, \gamma_1, \gamma_2)$ and the rows window 45's five readings, coil 1's trend row from window 44 to 45, and the drift prior, the Jacobian's lines at the hypothetical state are

$$\begin{bmatrix} 0.2886 & -20.8420 & -20.8420 & 0 & 0 & 0 & 0 & 0 \\ 0.1219 & -32.8785 & -8.8055 & 0 & 7.2222 & 0 & 0 & 0 \\ 0.1219 & -8.8055 & -32.8785 & 0 & 0 & 0 & 0 & 0 \\ 0.3333 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.1553 & -20.8420 & -20.8420 & 24.0731 & 0 & 0 & 0 & 0 \\ 0 & 200 & 0 & 0 & 0 & -200 & -16.6667 & 0 \\ 0 & 0 & 0 & 0 & 0.5 & 0 & 0 & 0 \end{bmatrix} \begin{matrix} y_{T_r} \\ y_{T_{d,1}} \\ y_{T_{d,2}} \\ y_Q \\ y_{T_1} \\ \text{trend 44–45} \\ \text{drift prior} \end{matrix}, \tag{22.4}$$

the load column in kW^{-1}, the trend entries' columns in yr, the drift's in yr/K and the others dimensionless. They are Chapter 10's window lines, equation (10.9), evaluated on the line instead of at a fit: each discharge row has its largest entry in its own coil's column, the return-air row sees both coils alike, and the rack-1 row repeats it in the log-derate columns. Unit 1's discharge row carries $7.2222 = (45 - 19)\,\Delta t/0.3$ in the drift's column, the drift's 26th window.

Which row acts on which entry. The pattern is Chapter 10's, equation (10.6), carried to window t_{as}: of the $324 \times 142 = 46\,008$ places of the Jacobian to window 45, 944 are nonzero, 2.05 %. The column of γ_u holds 45 entries, each $-\Delta t/\sigma_q = -16.6667$, all in coil u's trend rows; no reading row touches it. The second line of the pattern, the gradients of the questions' quantities, is

$$\mathbf{g}_{\gamma_1} = \mathbf{e}_{\gamma_1}, \qquad \mathbf{g}_{\gamma_2} = \mathbf{e}_{\gamma_2}, \tag{22.5}$$

the unit vectors of the two trend entries. Every question of Table 22.1 reads one entry of the state, so each answer is a function of one number of the record's covariance, the variance of γ_u.

Modeling choices at the rows. Four choices make these rows, each named once. *The rate is conditioned on, and the rest is drawn from its priors*: the outcome distribution of equation (22.2) fixes γ_u and draws $g_{u,0}$, β_1 and $\dot{b}_{d,1}$ from the rating, survey and drift priors, with the path wandering about its line by σ_q a month; the loads are entries with no prior row, and their values do not enter. This buys a probability of detection that the model itself predicts, in closed form, and it fails both ways. A path smoother than σ_q says is detected sooner, which Section 22.4.2 sets beside it; a rougher one later, and the evidence of Chapter 10 leans that way, peaking at $\sigma_q = 0.0125$ (Section 10.4.4): at $\sigma_q = 0.01$ the probability of detection at window 45 falls to 0.8177 and 0.8394 and the record a power of 0.9 needs grows to 56 and 54 windows (Section 22.4.4). *The rows are linearized at the* 5 %/yr *line*: this buys one Jacobian per record length, and it fails where the rows bend between the line and a drawn record; step 5 refits simulated records and moves the point to the 3 %/yr line and to an unaged second coil. *Chapter 10's rows are unchanged*: $\sigma_q = 0.005$, the drift from window 20 with Chapter 7's prior, the rating and survey priors. The record's own drift entry is part of "the existing sensors" as they stand, and step 7 sets a sound sensor beside it. *The sensors' grade is one number*, σ_T, for the four temperature sensors together: their reading rows are rescaled by $0.3/\sigma_T$, which keeps the record's design and changes only how well each window is read. The load meter is varied separately, in step 5.

22.2 The posterior of the state (step 3)

Every question of Table 22.1 is read from one object, the posterior covariance that a record to window t_{as} will have, computed from the rows of step 2 at the hypothetical state before any of the record's later windows is read.

22.2.1 The preposterior of a record not yet logged

For linear rows the posterior covariance does not depend on the values read, so the posterior that a record to window t_{as} will give is known, up to its center, before the record is logged. The derivation is Chapter 5's (Sections 5.2.1 and 5.2.2), applied by Chapter 10 to the stacked state: at the hypothetical state $\mathbf{z}^{\text{D}}_{t_{\text{as}}}$ the normal equations are $(\mathbf{J}^{\mathsf{T}}_{t_{\text{as}}}\mathbf{J}_{t_{\text{as}}})\,\Delta\mathbf{z} = -\mathbf{J}^{\mathsf{T}}_{t_{\text{as}}}\mathbf{r}(\mathbf{z}^{\text{D}}_{t_{\text{as}}})$, with $\mathbf{J}_{t_{\text{as}}}$ the Jacobian of the record's normalized rows, and the Laplace covariance is

$$\boldsymbol{\Sigma}_{t_{\text{as}}} = \left(\mathbf{J}^{\mathsf{T}}_{t_{\text{as}}}\mathbf{J}_{t_{\text{as}}}\right)^{-1}\Big|_{\mathbf{z}^{\text{D}}_{t_{\text{as}}}}, \qquad \mathbf{J}^{\mathsf{T}}_{t_{\text{as}}}\mathbf{J}_{t_{\text{as}}} = \mathbf{J}^{\mathsf{T}}_{\text{pr}}\mathbf{J}_{\text{pr}} + \mathbf{J}^{\mathsf{T}}_{0:t_{\text{as}}}\mathbf{J}_{0:t_{\text{as}}}, \tag{22.6}$$

with $\mathbf{J}_{\text{pr}}$ the lines of the four prior rows and $\mathbf{J}_{0:t_{\text{as}}}$ those of every reading row and trend row of windows 0 to t_{as}. This is Chapter 2's preposterior (2.17), with $\boldsymbol{\Sigma}^{-1} = \mathbf{J}^{\mathsf{T}}_{\text{pr}}\mathbf{J}_{\text{pr}}$ the information of the prior rows and $\boldsymbol{\Lambda}_a = \mathbf{J}^{\mathsf{T}}_{0:t_{\text{as}}}\mathbf{J}_{0:t_{\text{as}}}$ that of the record. With every row normalized the information matrix is $\mathbf{J}^{\mathsf{T}}\mathbf{J}$; with unnormalized rows and their weights it is $\mathbf{J}^{\mathsf{T}}\mathbf{W}\mathbf{J}$. The reading rows are nonlinear in the state through $\mathrm{e}^{g_{u,t}}$ and $\beta_1(T_r - T_c)$, so $\boldsymbol{\Sigma}_{t_{\text{as}}}$ depends on where they are linearized, the choice of step 2, which step 5 moves.

The covariance at window 45. The block of $\boldsymbol{\Sigma}_{45}$ on the eight entries $g_{1,0}$, $g_{2,0}$, $g_{1,45}$, $g_{2,45}$, β_1, γ_1, γ_2 and $\dot{b}_{d,1}$, rows and columns in that order, is

$$\boldsymbol{\Sigma}_{45} = 10^{-4} \times \begin{bmatrix} 2.2430 & -0.74089 & -0.82161 & 0.41822 & 0.020639 & -0.81722 & 0.30910 & -2.4795 \\ -0.74089 & 2.2407 & 0.69317 & -0.53452 & 0.18250 & 0.38242 & -0.74006 & 2.4346 \\ -0.82161 & 0.69317 & 5.4633 & -2.1054 & 0.70448 & 1.6760 & -0.74628 & 18.743 \\ 0.41822 & -0.53452 & -2.1054 & 2.4498 & -0.16342 & -0.67296 & 0.79581 & -7.6061 \\ 0.020639 & 0.18250 & 0.70448 & -0.16342 & 0.66036 & 0.18236 & -0.092248 & 3.2074 \\ -0.81722 & 0.38242 & 1.6760 & -0.67296 & 0.18236 & 1.4649 & -0.28143 & 5.6593 \\ 0.30910 & -0.74006 & -0.74628 & 0.79581 & -0.092248 & -0.28143 & 1.2096 & -2.6775 \\ -2.4795 & 2.4346 & 18.743 & -7.6061 & 3.2074 & 5.6593 & -2.6775 & 97.375 \end{bmatrix},$$

with the trend entries' rows and columns in yr^{-1} and the drift's in K/yr. The correlations read by sign, and the rows fix each sign. The rate of coil 1 is correlated with the drift rate by $+0.4739$, and the drift rate with coil 1's last log-derate by $+0.8126$: a faster drift explains more of the late rise of unit 1's discharge reading, so coil 1 stands higher at the end and needs less loss, Chapter 7's ridge read on the rate. That ridge is what makes coil 1's rate the less well known of the two on any record that runs past window 20. Each rate is correlated with its coil's commissioning entry by -0.4508 (coil 1) and -0.4495 (coil 2): a higher commissioned value with the later windows held needs a steeper fall. Each is correlated with its coil's last log-derate by $+0.5924$ and $+0.4623$: a higher value at the end with the start held needs a gentler one, the rate as the telescoped path of equation (10.7). The two rates are correlated by -0.2114 and the two last log-derates by -0.5755: the return air sees the hall's total, so a coil that loses more leaves less to the other. The signs, and nearly the sizes, are those of the record's own posterior at window 36 (Section 10.2.3). The diagonal holds the two variances the answers use, 1.4649×10^{-4} for γ_1 and 1.2096×10^{-4} for γ_2.

22.2.2 The predictive spread of the rate

The answers use one number per coil and record length, the rate's posterior standard deviation on a record to window t_{as},

$$\varsigma_u(t_{\text{as}}) = \sqrt{\mathbf{g}^{\mathsf{T}}_{\gamma_u}\boldsymbol{\Sigma}_{t_{\text{as}}}\,\mathbf{g}_{\gamma_u}} = \sqrt{[\boldsymbol{\Sigma}_{t_{\text{as}}}]_{\gamma_u\gamma_u}}, \tag{22.7}$$

with the gradient (22.5): Chapter 21's spread of a rate's posterior, here fixed before the record is read. Section 22.2.3 shows that it is also the spread of the rate's estimate over the records the plan may log, which is why it is called the predictive spread. At window 45, $\varsigma_1(45) = 0.01210$ and $\varsigma_2(45) = 0.01100\,\mathrm{yr}^{-1}$, the roots of the diagonal above.

Table 22.2 gives it at seven record lengths. Each row is a last window; the columns give each coil's predictive spread, the same rows' spread on the straight line, $\sigma_q \to 0$, the wander floor of equation (10.12) for t_{as} monthly increments, $\sigma_q/(\Delta t\sqrt{t_{\mathrm{as}}})$, and the share of each variance that is not the straight line's.

Table 22.2: The predictive spread of each coil's rate on a record to window t_{as}, at the hypothetical state, with the straight line's spread on the same windows, the wander floor, and the share of the variance that the wander adds to the straight line's.

t_{as} [window]	$\varsigma_1(t_{\mathrm{as}})$ [yr^{-1}]	$\varsigma_2(t_{\mathrm{as}})$ [yr^{-1}]	line, coil 1 [yr^{-1}]	line, coil 2 [yr^{-1}]	floor [yr^{-1}]	wander share 1 [–]	wander share 2 [–]
6 (Jun Y1)	0.09155	0.09155	0.08794	0.08794	0.02449	0.077	0.077
12 (Dec Y1)	0.03832	0.03832	0.03366	0.03366	0.01732	0.229	0.229
24 (Dec Y2)	0.02014	0.01838	0.01482	0.01287	0.01225	0.459	0.509
36 (Dec Y3)	0.01480	0.01309	0.00994	0.00749	0.01000	0.549	0.672
45 (Sep Y4)	0.01210	0.01100	0.00741	0.00546	0.00894	0.625	0.754
48 (Dec Y4)	0.01140	0.01046	0.00674	0.00495	0.00866	0.650	0.776
72 (Dec Y6)	0.00818	0.00789	0.00352	0.00265	0.00707	0.815	0.888

The variance has two parts. The straight line's spread is what the readings' errors give a path with no wander, the readings' part; the rest is what the wander adds to it, and the floor, the spread of a path known at every window, is the part of that rest no reading removes (Chapter 10's two splits, Section 10.3.1). Early in a record the sensors set the spread, and late in it the path's wander does: the wander's share of coil 1's variance is 0.229 at one year, 0.549 at three and 0.815 at six. The two parts fall at different speeds. The readings' part falls by Chapter 2's square law (2.20), nearly as $t_{\mathrm{as}}^{-3/2}$ from three years on: the straight line's spread times $t_{\mathrm{as}}^{3/2}$ is 2.1473 at window 36, 2.2378 at 45 and 2.1514 at 72 for coil 1, after 1.3991 at 12. The floor falls only as $t_{\mathrm{as}}^{-1/2}$, and the total approaches it: $\varsigma_1(t_{\mathrm{as}})\,t_{\mathrm{as}}^{1/2}$ is 0.13275 at window 12, 0.08881 at 36 and 0.06943 at 72, against the floor's constant $\sigma_q/\Delta t = 0.0600$.

The two coils are equal on every record that ends before window 20, where the drift entry first shares unit 1's discharge reading, and they part after it, coil 1 the wider, by the ridge of Section 22.2.1. And the spread does not depend on the readings. At window 36 the hypothetical record gives 0.01480 and 0.01309, and the base fit of the record itself 0.01463 and 0.01323 (Section 10.2.3); the as-of fits of the record at windows 12, 24 and 36 give 0.03778, 0.01994 and 0.01463 for coil 1, against the hypothetical 0.03832, 0.02014 and 0.01480. The rows differ only through the state at which they are evaluated.

22.2.3 The estimate before the record

The rule reads the record's posterior. Under the Laplace normal of the record to window t_{as}, with the estimate $\hat{\gamma}_{u,t_{\mathrm{as}}}$ its MAP value, the posterior probability of a loss is equation (10.11), $\Pr(\gamma_u < 0 \mid \mathbf{y}_{0:t_{\mathrm{as}}}) = \Phi(-\hat{\gamma}_{u,t_{\mathrm{as}}}/\varsigma_u(t_{\mathrm{as}}))$, so the verdict is

$$\mathbf{1}\big[\Phi(-\hat{\gamma}_{u,t_{\mathrm{as}}}/\varsigma_u(t_{\mathrm{as}})) \geq \kappa\big] = \mathbf{1}\big[\hat{\gamma}_{u,t_{\mathrm{as}}} \leq -z_\kappa\,\varsigma_u(t_{\mathrm{as}})\big], \qquad z_\kappa = \Phi^{-1}(\kappa) = 1.6449. \tag{22.8}$$

What the probability of detection needs is the distribution of $\hat{\gamma}_{u,t_{\mathrm{as}}}$ over the records the plan may log, at the stated rate. In the linear Gaussian form one Gauss–Newton step from the hypothetical state reaches the MAP state, $\hat{\mathbf{z}}_{t_{\mathrm{as}}} = \mathbf{z}^{\mathrm{D}}_{t_{\mathrm{as}}} - \boldsymbol{\Sigma}_{t_{\mathrm{as}}}\mathbf{J}_{t_{\mathrm{as}}}^{\mathsf{T}}\mathbf{r}(\mathbf{z}^{\mathrm{D}}_{t_{\mathrm{as}}})$. At the state $\mathbf{z}$ that generates the record, each normalized row is a standard normal: a reading row by its sensor's

error over its standard deviation, a trend row by the path's wander over σ_q, and a prior row by the draw of its entry from the prior. Writing $\boldsymbol{\varepsilon}_{t_{\mathrm{as}}}$ for that vector of standard normals, the rows at any point $\boldsymbol{\zeta}$ are $\mathbf{r}(\boldsymbol{\zeta}) = \mathbf{J}_{t_{\mathrm{as}}}(\boldsymbol{\zeta} - \mathbf{z}) - \boldsymbol{\varepsilon}_{t_{\mathrm{as}}}$, and

$$\begin{aligned}\hat{\mathbf{z}}_{t_{\mathrm{as}}} &= \mathbf{z}^{\mathrm{D}}_{t_{\mathrm{as}}} - \boldsymbol{\Sigma}_{t_{\mathrm{as}}}\mathbf{J}^{\mathsf{T}}_{t_{\mathrm{as}}}\left(\mathbf{J}_{t_{\mathrm{as}}}(\mathbf{z}^{\mathrm{D}}_{t_{\mathrm{as}}} - \mathbf{z}) - \boldsymbol{\varepsilon}_{t_{\mathrm{as}}}\right)\\ &= \mathbf{z} + \boldsymbol{\Sigma}_{t_{\mathrm{as}}}\mathbf{J}^{\mathsf{T}}_{t_{\mathrm{as}}}\boldsymbol{\varepsilon}_{t_{\mathrm{as}}},\end{aligned}$$

because $\boldsymbol{\Sigma}_{t_{\mathrm{as}}}\mathbf{J}^{\mathsf{T}}_{t_{\mathrm{as}}}\mathbf{J}_{t_{\mathrm{as}}}$ is the identity. The estimate's error is a fixed linear function of the rows' errors,

$$\hat{\gamma}_{u,t_{\mathrm{as}}} - \gamma_u = \mathbf{g}^{\mathsf{T}}_{\gamma_u}\boldsymbol{\Sigma}_{t_{\mathrm{as}}}\mathbf{J}^{\mathsf{T}}_{t_{\mathrm{as}}}\boldsymbol{\varepsilon}_{t_{\mathrm{as}}}, \qquad \hat{\gamma}_{u,t_{\mathrm{as}}} \mid \gamma_u \sim \mathcal{N}\big(\gamma_u,\ \varsigma_u(t_{\mathrm{as}})^2\big), \tag{22.9}$$

since $\mathrm{Cov}(\boldsymbol{\varepsilon}_{t_{\mathrm{as}}}) = \mathbf{I}$ and $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_{t_{\mathrm{as}}}\mathbf{J}^{\mathsf{T}}_{t_{\mathrm{as}}}\mathbf{J}_{t_{\mathrm{as}}}\boldsymbol{\Sigma}_{t_{\mathrm{as}}}\mathbf{g} = \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_{t_{\mathrm{as}}}\mathbf{g}$. Neither the mean nor the spread depends on the entries the rows do not hold a prior on, the loads: they are entries of $\mathbf{z}$ and cancel. With the verdict (22.8), $\Pr(\hat{\gamma}_{u,t_{\mathrm{as}}} \le -z_\kappa\varsigma_u(t_{\mathrm{as}}) \mid \gamma_u) = \Phi\big((-z_\kappa\varsigma_u(t_{\mathrm{as}}) - \gamma_u)/\varsigma_u(t_{\mathrm{as}})\big)$, which at $\gamma_u = \gamma_{\mathrm{D}} < 0$ and at $\gamma_u = 0$ gives

$$\mathrm{POD}_u(t_{\mathrm{as}}) = \Phi\!\left(\frac{|\gamma_{\mathrm{D}}|}{\varsigma_u(t_{\mathrm{as}})} - z_\kappa\right), \qquad \mathrm{PFA}_u(t_{\mathrm{as}}) = \Phi(-z_\kappa) = 1 - \kappa \quad \text{at every } t_{\mathrm{as}}. \tag{22.10}$$

The predictive spread is the one quantity of the record's design that the probability of detection uses, and the false alarm of a single look is the rule's own, whatever the record.

The as-of estimates as one sequence. A monthly reader refits the growing record at every window and reads a verdict each time. The as-of estimates $\hat{\gamma}_{u,1}, \ldots, \hat{\gamma}_{u,45}$ are then one Gaussian vector, and equation (22.9) gives its covariance. Records cut at $t_{\mathrm{as}} < t'_{\mathrm{as}}$ are nested: every row of the shorter is a row of the longer, with the same error. Both estimates are unbiased linear functions of those errors, and the longer record's has the smaller variance of the two, the property of the normal equations by which the estimate from all the rows is the most precise unbiased linear one. For such an estimate $\hat{\gamma}'$ and any other unbiased linear one $\tilde{\gamma}$, adding a multiple of $\tilde{\gamma} - \hat{\gamma}'$ to $\hat{\gamma}'$ gives another unbiased linear estimate, whose variance is $\mathrm{Var}(\hat{\gamma}')$ plus twice the multiple times $\mathrm{Cov}(\hat{\gamma}',\ \tilde{\gamma} - \hat{\gamma}')$ plus the multiple's square times $\mathrm{Var}(\tilde{\gamma} - \hat{\gamma}')$. That variance is least at a multiple of zero only if the covariance vanishes. So $\mathrm{Cov}(\hat{\gamma}_{u,t'_{\mathrm{as}}},\ \hat{\gamma}_{u,t_{\mathrm{as}}} - \hat{\gamma}_{u,t'_{\mathrm{as}}}) = 0$, and

$$\mathrm{Cov}\big(\hat{\gamma}_{u,t_{\mathrm{as}}},\ \hat{\gamma}_{u,t'_{\mathrm{as}}} \mid \gamma_u\big) = \varsigma_u(t'_{\mathrm{as}})^2, \qquad t_{\mathrm{as}} \le t'_{\mathrm{as}}. \tag{22.11}$$

Computed from the two records' linear weights on their common rows, $\mathbf{g}^{\mathsf{T}}_{\gamma_1}\boldsymbol{\Sigma}_{t_{\mathrm{as}}}\mathbf{J}^{\mathsf{T}}_{t_{\mathrm{as}}}$ and $\mathbf{g}^{\mathsf{T}}_{\gamma_1}\boldsymbol{\Sigma}_{t'_{\mathrm{as}}}\mathbf{J}^{\mathsf{T}}_{t'_{\mathrm{as}}}$, the covariance of coil 1's estimates at windows $(12, 45)$, $(24, 45)$ and $(36, 45)$ is 1.4649×10^{-4} each time, equal to $\varsigma_1(45)^2 = 1.4649 \times 10^{-4}$. The sequence is a random walk run backward from the last window: $\hat{\gamma}_{u,45} \sim \mathcal{N}(\gamma_u,\ \varsigma_u(45)^2)$, and each earlier estimate is the later one plus an independent normal, $\hat{\gamma}_{u,t_{\mathrm{as}}} = \hat{\gamma}_{u,t_{\mathrm{as}}+1} + \mathcal{N}(0,\ \varsigma_u(t_{\mathrm{as}})^2 - \varsigma_u(t_{\mathrm{as}}+1)^2)$. Repeated looks are then a first passage of that walk below the moving boundary $-z_\kappa\varsigma_u(t_{\mathrm{as}})$.

22.2.4 Simulated records refitted

The closed form (22.10) rests on the rows linearized at the hypothetical state, and on each record judged with the spread of the plan instead of its own. Beside it stands the full route, the simulated tests of *Reliability Engineering by Derivation*, Chapter 27, here simulated records:

1. draw a record of 45 windows at the stated rate on both coils: the commissioning entries, β_1 and the drift rate from their prior rows; each coil's path from its commissioning value by the rate plus a wander of σ_q a month; the loads about the record's growth with its month-to-month scatter; and the five readings of each window with their sensors' errors, rounded as the record's are;

2. refit the record cut at each look, windows $6, 12, \ldots, 42$ and 45, by Gauss–Newton with Chapter 10's rows unchanged;
3. judge each cut by its own posterior: concluded when $\Phi(-\hat{\gamma}_{u,t_{\mathrm{as}}}/\sqrt{[\boldsymbol{\Sigma}]_{\gamma_u\gamma_u}}) \geq \kappa$, with the spread of that refit;
4. read the share of records concluded as the probability of detection, with the Monte Carlo standard error $\sqrt{\hat{P}(1-\hat{P})/200}$ of a share $\hat{P}$ from 200 records; the same at $\gamma_u = 0$ gives the false alarm.

The result is the preposterior of the stated model up to its Monte Carlo error, without the linearization; it is called the simulated records, not the exact answer. Two hundred records are drawn per arm, at a 5 %/yr loss and at no loss, and step 4 sets their shares beside each closed form.

22.3 The answers (step 4)

Every answer is read from the predictive spread $\varsigma_u(t_{\mathrm{as}})$ of Section 22.2.2, through equation (22.10) or the form the question asks of it, and that value is the one reported; the simulated records of Section 22.2.4 stand beside it. A loss is reported as $100(1 - \mathrm{e}^{\gamma})$ %/yr, the rate (10.1) with its sign turned.

22.3.1 Detection by the shutdown

Asset manager's question. Can the existing sensors resolve a 5 % per year loss in the time I have? It is asked of each coil, with the time the refit at window 45 before the October shutdown.

General form. A design question about a rule asks for the design form (2.16) with the indicator of the acceptance rule (2.12) as the utility, its outcomes weighted at a stated value of one entry: the operating characteristic (22.1),

$$\mathrm{OC}(z_i^{\mathrm{t}}; a) = \int \mathbf{1}\big[\Pr(f(\mathbf{z}) \geq f^* \mid \tilde{\mathbf{y}}_a) \geq \kappa\big]\, p(\tilde{\mathbf{y}}_a \mid z_i = z_i^{\mathrm{t}})\, \mathrm{d}\tilde{\mathbf{y}}_a.$$

Specialized form. The question is the design form (2.16) as the operating characteristic (22.1), with a the record to window 45, $f(\mathbf{z}) = -\gamma_u$ and $f^* = 0$, $\kappa = 0.95$, and $z_i = \gamma_u$ at $z_i^{\mathrm{t}} = \gamma_{\mathrm{D}}$ of equation (22.3): equation (22.2) at $t_{\mathrm{as}} = 45$. Inside the indicator the record's posterior is the Laplace normal of step 3, so the indicator is the verdict (22.8); over the outcomes the estimate is normal about the stated rate with the predictive spread, equation (22.9); and the integral is equation (22.10), read on the gradient $\mathbf{e}_{\gamma_u}$ of equation (22.5). One choice enters here: the verdict is read once, at the shutdown's refit. This is what makes the false alarm the rule's own, $1 - \kappa$, and it fails if the verdict is read at every refit, which is the question of Section 22.3.4.

Building blocks. The predictive spread $\varsigma_u(45)$ of Table 22.2, from $\boldsymbol{\Sigma}_{45}$ of equation (22.6); the stated rate γ_{D}; the quantile $z_\kappa = 1.6449$; Φ.

The probability of detection by the shutdown. By equation (22.10), for coil 1 with $\varsigma_1(45) = 0.01210\,\mathrm{yr}^{-1}$,

$$\mathrm{POD}_1(45) = \Phi(4.2380 - 1.6449) = \Phi(2.5932) = 0.995,$$

with $|\gamma_{\mathrm{D}}|/\varsigma_1(45) = 4.2380$ and a miss of 0.00475; the simulated records conclude in 0.995 of 200 (standard error 0.005). For coil 2, with $\varsigma_2(45) = 0.01100$, $|\gamma_{\mathrm{D}}|/\varsigma_2(45) = 4.6639$ and

$$\mathrm{POD}_2(45) = \Phi(4.6639 - 1.6449) = \Phi(3.0190) = 0.999,$$

a miss of 0.00127; the simulated records conclude in all 200. On a record to window 36 the same form gives $\Phi(1.8205) = 0.966$ and $\Phi(2.2728) = 0.988$. The false alarm of the single look is 0.05 by construction; the simulated records at no loss conclude at window 45 in 0.035 (standard error 0.013) and 0.050 (0.015). The answer is yes: read on the conductance, the existing sensors resolve a 5 %/yr loss by the shutdown on both coils, and the single look at window 45 misses a coil-1 record with probability 0.00475.

The curve. Table 22.3 gives the probability of detection by equation (22.10) at nine record lengths. Each row is a last window; the columns give each coil's value. Both coils start equal, 0.3797 at a year, and part from window 20 on: 0.8163 and 0.8741 at two years, 0.9657 and 0.9885 at three, 0.9952 and 0.9987 at window 45.

Table 22.3: The probability of detection of a 5 %/yr loss on a record to window t_{as}, equation (22.10), with the rule $\Pr(\gamma_u < 0 \mid \mathbf{y}_{0:t_{\text{as}}}) \geq 0.95$ read once.

t_{as} [window]	$\text{POD}_1(t_{\text{as}})$ [–]	$\text{POD}_2(t_{\text{as}})$ [–]
6 (Jun Y1)	0.1391	0.1391
12 (Dec Y1)	0.3797	0.3797
18 (Jun Y2)	0.6842	0.6842
24 (Dec Y2)	0.8163	0.8741
30 (Jun Y3)	0.9073	0.9581
36 (Dec Y3)	0.9657	0.9885
45 (Sep Y4)	0.9952	0.9987
54 (Jun Y5)	0.9996	0.9999
72 (Dec Y6)	1.0000	1.0000

The dashboard's curve. The planner's test of the opening scene has a curve of its own: on the tiles' line, with the residuals taken as independent, its probability of concluding at a 5 %/yr loss read as 0.05 of the tile a year is 0.0844 at window 6, 0.1686 at 12, 0.5738 at 24, 0.9509 at 36 and 0.9983 at 45 for coil 1, and nearly the same for coil 2. Nor is 5 % of conductance 5 % of the tile: at the hypothetical state, both coils losing 5 %/yr, the tiles' line falls 0.1197 a year, 2.39 times the planner's 0.05. And the test has no meaningful false alarm. At the hypothetical state with no loss on either coil, with sound sensors and no reading error, the tiles' line over windows 0 to 45 falls by 0.0326 a year on each coil, the load's growth and the season. Judged against the planner's own standard error at window 45, 0.01093 a year, in a single look at that window, the planner's test concludes "losing" with probability 0.9098 for coil 1 and 0.9102 for coil 2. Its power is a power to see the load.

The curves are drawn in Figure 22.2. The horizontal axis is the last window of the record and the vertical axis the probability of detection; the solid lines are the two coils at the existing grade, 0.3 K, the dashed lines the same at 0.1 K (upper) and 1.0 K (lower), and the dotted line the planner's test; the gray horizontal line is the power asked, 0.9, and the gray vertical line window 45. At the existing grade both coils cross 0.9 before the shutdown; at 1.0 K neither does. The planner's curve rises later than the existing grade's through the first three years; it passes coil 1's between windows 36 and 45, 0.9983 against 0.9952 at window 45, and never passes coil 2's, 0.9987 there. And it would rise much the same way on a coil not losing at all.

22.3.2 The record length a power of 0.9 needs

Asset manager's question. How long a record does a power of 0.9 need, and what would better sensors save? The first part sets how long a restarted record must run before it can carry a work order; the second asks what a better grade of sensor is worth in windows.

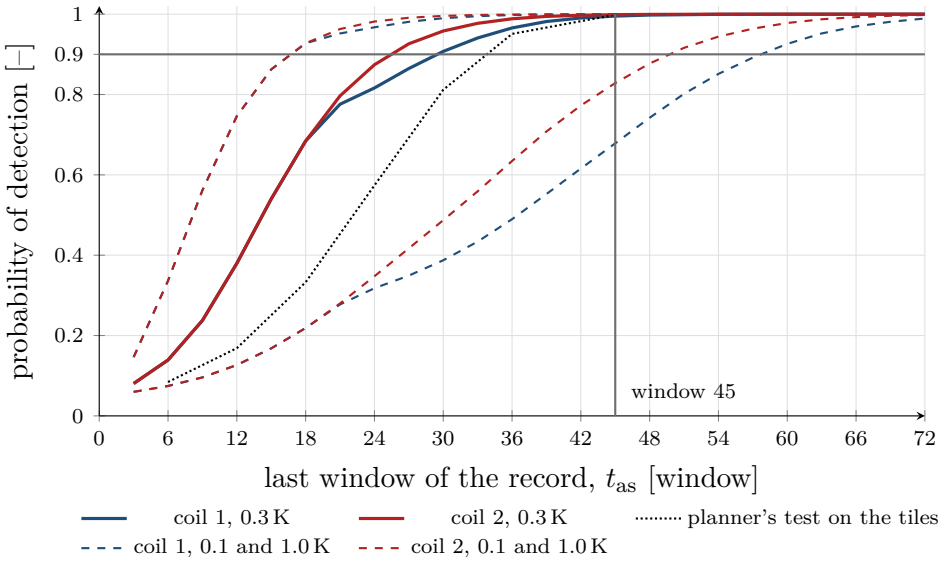


Figure 22.2: The probability of detection of a 5 %/yr loss against the record's length, equation (22.10): $\Phi(|\gamma_{\text{D}}|/\varsigma_u(t_{\text{as}}) - 1.6449)$ with $\varsigma_u(t_{\text{as}})$ from the rows at the hypothetical state. Solid: the existing sensors, 0.3 K. Dashed: the four temperature sensors at 0.1 K (the upper pair) and 1.0 K (the lower pair). Dotted: the planner's test on the tiles' line, which concludes at window 45 on unaged coils with probability 0.91. Gray: the power asked, 0.9, and window 45, the last read before the shutdown.

General form. The design form (2.16) with the acceptance rule's indicator, read as the operating characteristic (22.1) and solved for the design: the smallest measurement a in a family whose characteristic at the stated value reaches a stated power,

$$a^* = \min\{a : \text{OC}(z_i^{\text{t}}; a) \geq \text{OC}^*\}.$$

Specialized form. The question is the design form (2.16) as the operating characteristic (22.1), over the family of records to window t_{as}, with the stated value and rule of Section 22.3.1 and $\text{OC}^* = \text{POD}^* = 0.9$. By equation (22.10) the probability of detection increases as the predictive spread falls, so

$$T_{90,u} = \min\{t_{\text{as}} : \text{POD}_u(t_{\text{as}}) \geq \text{POD}^*\} = \min\Big\{t_{\text{as}} : \varsigma_u(t_{\text{as}}) \leq \frac{|\gamma_{\text{D}}|}{z_\kappa + z_{\text{POD}^*}}\Big\}, \tag{22.12}$$

with $z_{\text{POD}^*} = \Phi^{-1}(\text{POD}^*)$: the first window at which the spread falls below $0.05129/2.9264 = 0.01753\,\text{yr}^{-1}$. The sensors' grade enters through $\boldsymbol{\Sigma}_{t_{\text{as}}}$: the four temperature sensors' reading rows are rescaled by $0.3/\sigma_T$, the fourth choice of step 2, and $\varsigma_u(t_{\text{as}})$ is recomputed by equation (22.6) at each grade. With the path known at every window, the floor (10.12) for t_{as} increments is the spread, $\varsigma = \sigma_q/(\Delta t\sqrt{t_{\text{as}}})$, and the condition becomes a bound no sensor passes,

$$t_{\text{as}} \geq \left(\frac{(z_\kappa + z_{\text{POD}^*})\,\sigma_q}{\Delta t\,|\gamma_{\text{D}}|}\right)^2, \tag{22.13}$$

which Exercise 22.2 derives.

Building blocks. The predictive spread $\varsigma_u(t_{\text{as}})$ at each record length and grade, from equation (22.6); the quantiles $z_\kappa = 1.6449$ and $z_{\text{POD}^*} = 1.2816$, summing to 2.9264; the stated rate; the wander σ_q and the spacing Δt, $\sigma_q/\Delta t = 0.0600\,\text{yr}^{-1}$.

The record length at the existing grade. By equation (22.12) the existing sensors need 30 windows on coil 1, to window 30 (Jun Y3), where its probability of detection passes from 0.8941 at window 29 to 0.9073; and 26 windows on coil 2, to window 26 (Feb Y3), from 0.8942 at window 25 to 0.9116. The two coils differ by the drift entry, which shares coil 1's discharge reading; without it both need 25 windows (Section 22.4.4). For a record restarted at a new commissioning test the same 30 and 26 windows hold only if unit 1's sensor drifts again from the new record's window 20, as the existing sensors do in this record; with a sound sensor the restarted record needs 25 windows on each coil. The simulated records agree at the crossing: at window 30 coil 1's conclude in 0.900 (standard error 0.021), and coil 2's in 0.895 (0.022) at window 24 and 0.970 (0.012) at window 30. By equation (22.13) the wander's bound is $(2.9264 \times 0.0600/0.05129)^2 = 11.72$ windows: no grade of sensor resolves a 5 %/yr loss with power 0.9 in under a year.

The sensors' grade. Table 22.4 repeats the computation at six grades of the four temperature sensors. Each row is a grade; the columns give each coil's record length for a power of 0.9, its predictive spread at window 45 and its probability of detection there.

Table 22.4: The record length for a power of 0.9, equation (22.12), and the predictive spread and probability of detection at window 45, against the grade σ_T of the four temperature sensors; the load meter at 3.0 kW.

Grade σ_T [K]	$T_{90,1}$ [window]	$T_{90,2}$ [window]	$\varsigma_1(45)$ [yr^{-1}]	$\varsigma_2(45)$ [yr^{-1}]	$\mathrm{POD}_1(45)$ [–]	$\mathrm{POD}_2(45)$ [–]
0.05	14 (Feb Y2)	14 (Feb Y2)	0.00935	0.00917	0.9999	1.0000
0.1	17 (May Y2)	17 (May Y2)	0.00985	0.00947	0.9998	0.9999
0.2	24 (Dec Y2)	22 (Oct Y2)	0.01087	0.01015	0.9989	0.9997
0.3 (existing)	30 (Jun Y3)	26 (Feb Y3)	0.01210	0.01100	0.9952	0.9987
0.5	40 (Apr Y4)	34 (Oct Y3)	0.01519	0.01314	0.9585	0.9880
1.0	58 (Oct Y5)	50 (Feb Y5)	0.02433	0.01978	0.6785	0.8284

The record lengths are plotted against the grade in Figure 22.3. The horizontal axis is the grade on a logarithmic scale and the vertical axis the record length for a power of 0.9; the solid lines are the two coils, the dashed line the wander's bound, 11.72 windows, and the dotted line window 45. Both curves fall steeply from 1.0 to 0.2 K and flatten toward the bound below 0.1 K; coil 1's lies above coil 2's at every grade from 0.2 K up. The step's surprise is the size of the grade's effect. The floor suggests that the wander sets the band, but the readings dominate the spread for the first eighteen months, and the two parts are about even at 24 (Table 22.2): 0.1 K sensors cut the record from 30 windows to 17, and 1.0 K sensors would not see a 5 %/yr loss by the shutdown, 0.6785 on coil 1. The 0.05 K sensors need 14 windows, against the bound's 11.72.

The rule as an input. The rule and the power asked are stated inputs, and Table 22.5 moves both. Each row is an assurance κ; each column a power asked; each entry the two coils' record lengths. They run from 18 windows on each coil, at $\kappa = 0.90$ and a power of 0.8, to 42 and 37 windows at $\kappa = 0.99$ and a power of 0.95; at the reported 0.95 and 0.9 they are 30 and 26.

22.3.3 The smallest loss resolved

Asset manager's question. What is the smallest loss the record resolves by the shutdown? The maker's specification of 3 %/yr is one loss of interest, read here only as a loss to resolve from zero.

General form. The design form (2.16) as the operating characteristic (22.1), solved for the stated value: the smallest departure of the entry from its null value that the measurement a

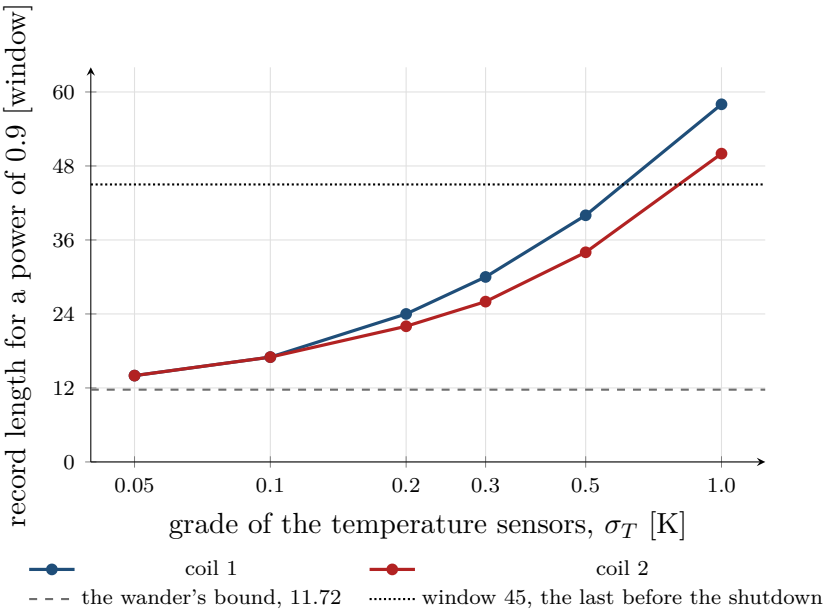


Figure 22.3: The record length for a power of 0.9 against the grade of the four temperature sensors, equation (22.12), the load meter at 3.0 kW (Table 22.4). Dashed: the bound of a path known at every window, equation (22.13). Dotted: window 45.

Table 22.5: The record length for a stated power under the rule $\Pr(\gamma_u < 0 \mid \mathbf{y}_{0:t_{\text{as}}}) \geq \kappa$, at the existing grade, coil 1 / coil 2.

Assurance κ [–]	power 0.80 [window]	power 0.90 [window]	power 0.95 [window]
0.90	18 / 18	25 / 22	30 / 26
0.95 (reported)	23 / 22	30 / 26	35 / 30
0.99	33 / 28	38 / 33	42 / 37

resolves at a stated power,

$$z_i^*(a) = \min\{|z_i^{\mathrm{t}}| : \mathrm{OC}(z_i^{\mathrm{t}}; a) \geq \mathrm{OC}^*\}.$$

Specialized form. The question is the design form (2.16) as the operating characteristic (22.1), with the record to window t_{as}, the rule of Section 22.3.1 and $\mathrm{OC}^* = \mathrm{POD}^*$, solved for the stated rate. By equation (22.10) at $\gamma_u = \gamma$, $\Phi(|\gamma|/\varsigma_u(t_{\mathrm{as}}) - z_\kappa) \geq \mathrm{POD}^*$ holds when $|\gamma| \geq (z_\kappa + z_{\mathrm{POD}^*})\,\varsigma_u(t_{\mathrm{as}})$, so the minimum detectable rate is

$$\gamma_{\mathrm{MD},u}(t_{\mathrm{as}}) = (z_\kappa + z_{\mathrm{POD}^*})\,\varsigma_u(t_{\mathrm{as}}) = 2.9264\,\varsigma_u(t_{\mathrm{as}}), \qquad \text{read as } 100(1 - \mathrm{e}^{-\gamma_{\mathrm{MD},u}(t_{\mathrm{as}})})\ \%/\mathrm{yr}. \tag{22.14}$$

One choice enters here: the rows stay linearized at the 5 %/yr line while the stated loss moves. This keeps one spread per record, and it fails if the spread depends on where the coils are; the rows linearized on the 3 %/yr line test it below.

Building blocks. The predictive spread $\varsigma_u(t_{\mathrm{as}})$; the two quantiles of Section 22.3.2; the map to %/yr.

The smallest loss. Table 22.6 gives the minimum detectable loss by equation (22.14). Its upper part runs over the record length at the existing grade, each row a last window; its lower part runs over the grade at window 45. A record of a year resolves only losses faster than 10.608 %/yr on either coil; by the shutdown, window 45, the record resolves with power 0.9 any loss faster than 3.480 %/yr on coil 1 and 3.167 %/yr on coil 2. A loss at the maker's 3 %/yr, $\gamma = -0.03046\,\mathrm{yr}^{-1}$, is resolved from zero at window 45 with probability 0.8083 on coil 1 and 0.8696 on coil 2; whether a loss exceeds 3 %/yr is Chapter 21's test, not this one. With the rows linearized on the 3 %/yr line instead, the spreads at window 45 are 0.01215 and 0.01098 and the minimum detectable losses 3.492 and 3.163 %/yr, against 3.480 and 3.167.

Table 22.6: The minimum detectable loss at a power of 0.9 under the rule $\Pr(\gamma_u < 0 \mid \mathbf{y}_{0:t_{\mathrm{as}}}) \geq 0.95$, equation (22.14): over the record length at the existing grade, and over the grade at window 45, where the loss is given in %/yr only.

Record and grade	$\gamma_{\mathrm{MD},1}$ [yr^{-1}]	coil 1 [%/yr]	$\gamma_{\mathrm{MD},2}$ [yr^{-1}]	coil 2 [%/yr]
The existing grade, 0.3 K				
$t_{\mathrm{as}} = 12$ (Dec Y1)	0.11214	10.608	0.11214	10.608
$t_{\mathrm{as}} = 24$ (Dec Y2)	0.05895	5.725	0.05379	5.236
$t_{\mathrm{as}} = 36$ (Dec Y3)	0.04332	4.239	0.03831	3.759
$t_{\mathrm{as}} = 45$ (Sep Y4)	0.03542	3.480	0.03218	3.167
$t_{\mathrm{as}} = 60$ (Dec Y5)	0.02758	2.720	0.02609	2.575
$t_{\mathrm{as}} = 72$ (Dec Y6)	0.02395	2.366	0.02309	2.283
At window 45, by the grade σ_T				
0.05 K	not given	2.699	not given	2.649
0.1 K	not given	2.842	not given	2.735
0.2 K	not given	3.131	not given	2.928
0.5 K	not given	4.347	not given	3.773
1.0 K	not given	6.872	not given	5.625

The minimum detectable loss is drawn against the record's length in Figure 22.4. The horizontal axis is the last window and the vertical axis the loss in %/yr; the solid lines are the two coils at the existing grade, the dashed lines at 0.1 K, the gray horizontal line the 5 %/yr of the question and the gray vertical line window 45. Every curve falls steeply in the first two years and slowly after them, as the predictive spread does. Each curve crosses 5 %/yr at its record

length for a power of 0.9, windows 30 and 26 at the existing grade and 17 at 0.1 K, because the two questions read the same condition on $\varsigma_u(t_{\mathrm{as}})$.

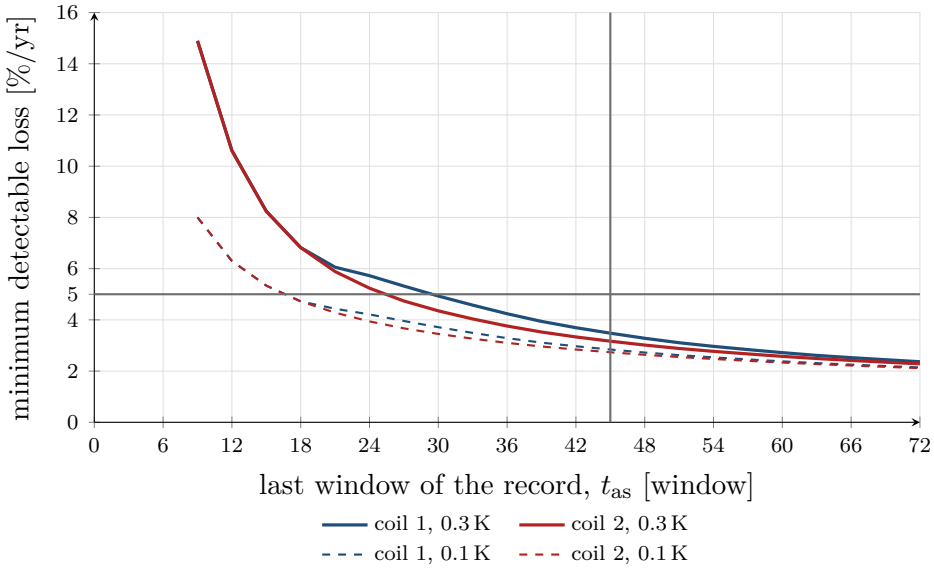


Figure 22.4: The minimum detectable loss at a power of 0.9 against the record's length, equation (22.14): $100(1 - \mathrm{e}^{-2.9264\,\varsigma_u(t_{\mathrm{as}})})$ with $\varsigma_u(t_{\mathrm{as}})$ from the rows at the hypothetical state, from window 9 on. Solid: the existing sensors. Dashed: the four temperature sensors at 0.1 K. Gray: the 5 %/yr of the question and window 45.

22.3.4 Monthly refits and the first verdict

Asset manager's question. If I refit every month and act on the first verdict, when does it come, and how often is it false? The dashboard redraws its line every month, and a planner who reads it each time acts on the first month that crosses.

General form. The design form (2.16) as an operating characteristic (22.1), with the rule read at every look of a schedule $\mathcal{S}$ of growing measurements $a_{t_{\mathrm{as}}}$ and the indicator that it is met at one look or more:

$$\mathrm{OC}_{\mathcal{S}}(z_i^{\mathrm{t}}) = \int \mathbf{1}\Big[\max_{t_{\mathrm{as}}\in\mathcal{S}} \mathbf{1}\big[\Pr(f(\mathbf{z}) \geq f^* \mid \tilde{\mathbf{y}}_{a_{t_{\mathrm{as}}}}) \geq \kappa\big] = 1\Big]\, p(\tilde{\mathbf{y}} \mid z_i = z_i^{\mathrm{t}})\, \mathrm{d}\tilde{\mathbf{y}}. \tag{22.15}$$

Specialized form. The question is equation (22.15) with $a_{t_{\mathrm{as}}}$ the record to window t_{as}, the rule of Section 22.3.1, and $\mathcal{S}$ the schedule of looks up to window 45: every window from 1 (monthly), every third (quarterly), every sixth with 45 (semiannual), windows 12, 24, 36 and 45 (annual), or window 45 alone (the single look). The verdict at look t_{as} is equation (22.8), and by equations (22.9) and (22.11) the as-of estimates are one Gaussian vector with $\mathrm{Cov}(\hat{\gamma}_{u,t_{\mathrm{as}}}, \hat{\gamma}_{u,t'_{\mathrm{as}}}) = \varsigma_u(\max(t_{\mathrm{as}}, t'_{\mathrm{as}}))^2$, so the integral is the probability of a first passage,

$$\mathrm{P}_{\mathcal{S},u}(\gamma_u) = \Pr\big(\exists\, t_{\mathrm{as}} \in \mathcal{S} : \hat{\gamma}_{u,t_{\mathrm{as}}} \leq -z_\kappa\, \varsigma_u(t_{\mathrm{as}}) \mid \gamma_u\big), \tag{22.16}$$

at $\gamma_u = \gamma_{\mathrm{D}}$ for detection and at $\gamma_u = 0$ for the false alarm, with the first window that passes as the time of the verdict. One choice enters here: each look is judged with the plan's spread at its window, as the single look is. This makes the sequence one covariance, computed before any record, and it fails if a record's own refits scatter about the plan's spread; the simulated records at the semiannual looks test it.

Building blocks. The predictive spreads $\varsigma_u(1), \ldots, \varsigma_u(45)$ of Section 22.2.2, linearized at the 5 %/yr line for detection and at no loss for the false alarm ($\varsigma_1(45) = 0.01221$, $\varsigma_2(45) = 0.01096$ there); the covariance (22.11); draws of the sequence by the walk run backward: draw $\hat{\gamma}_{u,45} \sim \mathcal{N}(\gamma_u,\ \varsigma_u(45)^2)$, then for $t_{\text{as}} = 44, \ldots, 1$ set $\hat{\gamma}_{u,t_{\text{as}}} = \hat{\gamma}_{u,t_{\text{as}}+1} + \mathcal{N}(0,\ \varsigma_u(t_{\text{as}})^2 - \varsigma_u(t_{\text{as}}+1)^2)$, independent of the rest, and read the verdict (22.8) at every look of each schedule; 400 000 draws per coil and rate, in chunks of 50 000, with a Monte Carlo standard error of at most 0.0008 on every probability.

The schedules. Table 22.7 gives equation (22.16) for the five schedules. Each row is a schedule; the columns give the probability of a verdict by window 45 at a 5 %/yr loss and at no loss, for each coil. The single look reproduces the closed form, 0.9952 against 0.9952 for coil 1 and 0.9988 against 0.9987 for coil 2, and its false alarm is 0.0499 and 0.0495.

Table 22.7: The probability of a verdict by window 45 under five schedules of looks, equation (22.16), from 400 000 draws of the as-of sequence per coil and rate (Monte Carlo standard errors at most 0.0008).

Schedule of looks	5 %/yr, coil 1 [–]	5 %/yr, coil 2 [–]	no loss, coil 1 [–]	no loss, coil 2 [–]
the single look at window 45 (reported rule)	0.9952	0.9988	0.0499	0.0495
annual: windows 12, 24, 36, 45	0.9966	0.9992	0.1320	0.1327
semiannual: windows 6, 12, ..., 42, 45	0.9973	0.9994	0.1902	0.1942
quarterly: windows 3, 6, ..., 45	0.9978	0.9995	0.2559	0.2623
monthly: windows 1 to 45	0.9984	0.9996	0.3682	0.3750

The surprise of the chapter is in the last row. Refitting every month and acting on the first verdict raises the probability of detection from 0.9952 to 0.9984 on coil 1 and raises the false alarm from 0.0499 to 0.3682: a coil that is not losing at all is written up as losing by the shutdown with probability 0.3682. Most of those false verdicts come early. At no loss the monthly reader's first false verdict falls at a median window of 7 (Jul Y1), with a 5–95 % range of windows 1 (Jan Y1) to 37 (Jan Y4), and 0.2571 of the 0.3682 is spent by window 12. At a 5 %/yr loss the monthly refits first conclude at a median window of 11 (Nov Y1), with a 5–95 % range of windows 1 to 28 (Jan Y1 to Apr Y3) for coil 1 and 1 to 25 (Jan Y1 to Jan Y3) for coil 2; by window 12 they have concluded in 0.5940 and 0.5953, against the single look's 0.3797 at window 12. In the first year a verdict comes on a coil not losing with probability 0.2571, and on a coil losing 5 %/yr with probability 0.5940: the monthly verdict of the first year tells the two apart poorly. The first verdict is one record's draw, and its spread is the ranges just given.

The cumulative probability of a verdict is drawn against the window in Figure 22.5, for coil 1. The horizontal axis is the window of the look and the vertical axis the probability that a verdict has been written by then; the solid line is the monthly reader at a 5 %/yr loss and the dashed line the monthly reader at no loss; the gray solid line is the single look's probability of detection for a record to that window, equation (22.10), and the gray dashed line its constant false alarm, 0.05. The dotted line marks window 13, where the record as logged first reached the verdict, a check marker for Section 22.4.5. The monthly curve at no loss rises steeply in the first months, where the spread falls fastest and each new look shares least with the last, by equation (22.11), and flattens as the looks come to share most of their rows.

The simulated records check the sequence at the semiannual looks: refitted at windows $6, 12, \ldots, 42$ and 45 and each judged by its own posterior, they conclude at some look by window 45 in 0.180 of the records at no loss (standard error 0.027) for coil 1 and 0.215 (0.029) for coil 2, against the sequence's 0.1902 and 0.1942, and in 0.995 and 1.000 at a 5 %/yr loss, against 0.9973 and 0.9994.

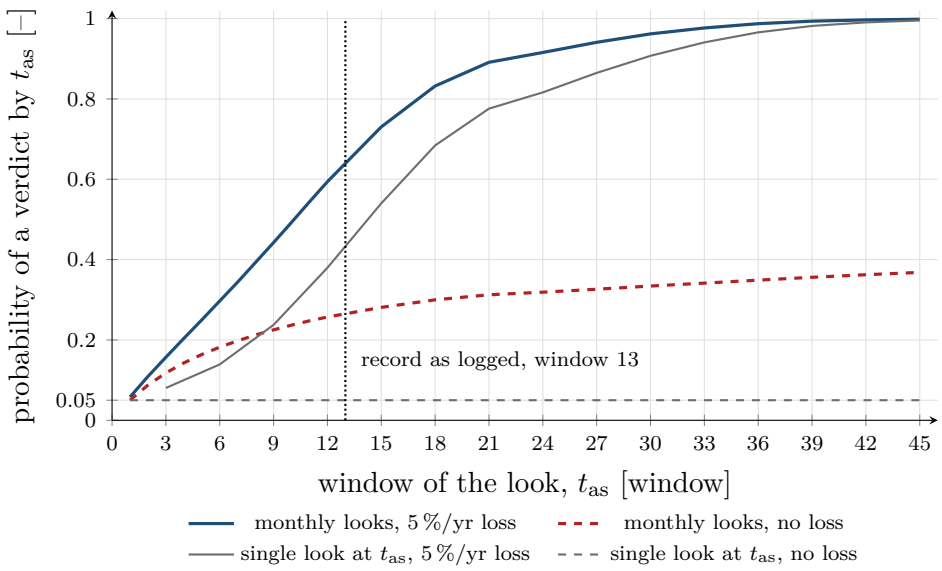


Figure 22.5: Coil 1's probability of a verdict by window t_{as} under monthly looks, equation (22.16), from 400 000 draws of the as-of sequence at a 5 %/yr loss and at no loss, against the single look at window t_{as}, equation (22.10). Dotted: the window at which the record as logged first reached the verdict, a check marker (Section 22.4.5).

22.3.5 The Bayes-factor rule

Asset manager's question. Does a stricter rule, a Bayes factor of 19 against no trend, change the answer? A work order raised on a trend log is sometimes asked to show that the record needs the trend at all, not only that the trend's sign is likely.

General form. The design form (2.16) with the indicator that a Bayes factor (2.14) passes a stated level B^* as the utility, read as an operating characteristic at a stated value of one entry,

$$\mathrm{OC}^{B}(z_i^{\mathrm{t}}; a) = \int \mathbf{1}\left[B_{AB}(\tilde{\mathbf{y}}_a) \geq B^*\right] p(\tilde{\mathbf{y}}_a \mid z_i = z_i^{\mathrm{t}})\, \mathrm{d}\tilde{\mathbf{y}}_a. \tag{22.17}$$

Specialized form. The question is equation (22.17) with a the record to window t_{as}, $z_i = \gamma_u$ at γ_{D} or at 0, $B^* = 19$, M_A the drawing with a trend entry of prior $\mathcal{N}(0,\ \sigma_\gamma^2)$, $\sigma_\gamma = 0.10\,\mathrm{yr}^{-1}$, and M_B the drawing that pins $\gamma_u = 0$. The two are nested, and by the Savage–Dickey ratio (7.21) the factor needs only the larger drawing's posterior at the pinned value. On the record's Laplace marginal $\mathcal{N}(\hat{\gamma}_u,\ \varsigma_u^2)$ of the base fit, whose trend entry has no prior row, the trend's prior gives the posterior $\mathcal{N}(\hat{\gamma}_{\mathrm{tr}},\ \varsigma_{\mathrm{tr}}^2)$ of M_A by precisions that add,

$$\varsigma_{\mathrm{tr}}^{-2} = \varsigma_u^{-2} + \sigma_\gamma^{-2}, \qquad \hat{\gamma}_{\mathrm{tr}} = \frac{\varsigma_{\mathrm{tr}}^2}{\varsigma_u^2}\,\hat{\gamma}_u,$$

and the ratio of its density at zero to the prior's, $B(\text{none} : \text{trend}) = (\sigma_\gamma/\varsigma_{\mathrm{tr}})\,\exp(-\hat{\gamma}_{\mathrm{tr}}^2/(2\varsigma_{\mathrm{tr}}^2))$, gives

$$\ln B(\text{trend} : \text{none}) = \tfrac{1}{2}\ln\frac{\varsigma_{\mathrm{tr}}^2}{\sigma_\gamma^2} + \frac{\hat{\gamma}_{\mathrm{tr}}^2}{2\varsigma_{\mathrm{tr}}^2}. \tag{22.18}$$

The rule $\ln B \geq \ln 19$ with $\hat{\gamma}_u < 0$ is a threshold on the same standardized estimate as the sign rule, $\hat{\gamma}_{\mathrm{tr}}/\varsigma_{\mathrm{tr}} = (\varsigma_{\mathrm{tr}}/\varsigma_u)\,\hat{\gamma}_u/\varsigma_u$:

$$\begin{aligned} \frac{-\hat{\gamma}_{u,t_{\mathrm{as}}}}{\varsigma_u(t_{\mathrm{as}})} \geq z_u^B(t_{\mathrm{as}}) &= \frac{\varsigma_u(t_{\mathrm{as}})}{\varsigma_{\mathrm{tr}}}\sqrt{2\ln 19 + \ln\frac{\sigma_\gamma^2}{\varsigma_{\mathrm{tr}}^2}}, \\ \mathrm{POD}_u^B(t_{\mathrm{as}}) = \Phi\!\left(\frac{|\gamma_{\mathrm{D}}|}{\varsigma_u(t_{\mathrm{as}})} - z_u^B(t_{\mathrm{as}})\right), &\qquad \mathrm{PFA}_u^B(t_{\mathrm{as}}) = \Phi(-z_u^B(t_{\mathrm{as}})), \end{aligned} \tag{22.19}$$

by equation (22.9). The threshold grows as the record lengthens, because $\ln(\sigma_\gamma^2/\varsigma_{\mathrm{tr}}^2)$ grows as the spread falls, and it grows with the prior spread: Chapter 11's Lindley's effect (Section 11.3.2), here read on the rate. One choice enters here: the trend's prior $\mathcal{N}(0,\ 0.10^2)$ on γ_u alone, with the level 19, a posterior probability of 0.95 at even prior odds, Chapter 10's decision level for the drift. This buys the factor from the same predictive spread, and it fails as the prior is widened or narrowed; the answer is given over σ_γ from 0.03 to 0.30.

Building blocks. The predictive spread $\varsigma_u(t_{\mathrm{as}})$; the prior spread σ_γ; the threshold and the tails of equation (22.19); for repeated looks, the sequence of Section 22.3.4 judged by equation (22.18) at each look.

The stricter rule. Table 22.8 gives equation (22.19) at windows 36 and 45. Each row is a coil and a window; the columns give the threshold on the standardized estimate, the probability of detection, the simulated records' share at window 45, and the false alarm of the single look.

Table 22.8: The Bayes-factor rule, $B(\text{trend} : \text{none}) \geq 19$ with the trend's prior $\mathcal{N}(0,\ 0.10^2)$, by equation (22.19); the sign rule's threshold is 1.6449.

Coil [–]	Window [window]	threshold $z_u^B(t_{\mathrm{as}})$ [–]	$\mathrm{POD}_u^B(t_{\mathrm{as}})$ [–]	simulated [–]	$\mathrm{PFA}_u^B(t_{\mathrm{as}})$ [–]
1	36 (Dec Y3)	3.1535	0.62242	–	0.00081
2	36 (Dec Y3)	3.1848	0.76818	–	0.00072
1	45 (Sep Y4)	3.2055	0.84909	0.860 (se 0.025)	0.00067
2	45 (Sep Y4)	3.2312	0.92402	0.915 (se 0.020)	0.00062

The stricter rule changes the answer for coil 1. Its threshold at window 45 is 3.2055 on coil 1, against the sign rule's 1.6449, so by the shutdown it detects a 5 %/yr loss with probability 0.849 on coil 1 and 0.924 on coil 2, and a power of 0.9 needs 49 windows (Jan Y5) on coil 1 and 43 (Jul Y4) on coil 2, against the sign rule's 30 and 26: coil 1 not by the shutdown. Over the trend's prior spread the record length moves with Lindley's effect: at $\sigma_\gamma = 0.03\,\mathrm{yr}^{-1}$ the threshold at window 45 is 3.0222 on coil 1, the probability of detection 0.8880, and the record lengths 46 (Oct Y4) and 41 (May Y4) windows; at 0.30 they are 3.5116, 0.7662, and 52 (Apr Y5) and 48 (Dec Y4). Under monthly looks the Bayes-factor rule concludes by window 45 at a 5 %/yr loss in 0.8942 of draws on coil 1 and 0.9499 on coil 2, at a median first window of 27 (Mar Y3) and 25 (Jan Y3), and it raises a false alarm by window 45 in 0.0106 and 0.0105, against 0.0007 and 0.0006 for its single look. The two rules buy different things: the sign rule a short record at a false alarm of 0.05 per look, the Bayes factor a false alarm near 0.01 even under monthly looks, at the cost of 49 and 43 windows of record for a power of 0.9 in place of 30 and 26.

22.4 The checks (step 5)

Each answer of Section 22.3 rests on the rows linearized at the hypothetical state, on a wander the truth may not have, and on the choices of step 2, and each can be tested against the

computation it stands for.

22.4.1 The simulated records against the closed form

The simulated records of Section 22.2.4, 200 at a 5 %/yr loss (seeds 22000 to 22199) and 200 at no loss (seeds 23000 to 23199), were refitted at windows $6, 12, \dots, 42$ and 45 and set beside every answer of step 4. The Monte Carlo standard error of a share near 0.8 from 200 records is 0.029. At a 5 %/yr loss coil 1's records conclude in 0.135 at window 6 (standard error 0.024), 0.330 at 12 (0.033), 0.795 at 24 (0.029), 0.965 at 36 (0.013) and 0.995 at 45 (0.005), against the closed form's 0.1391, 0.3797, 0.8163, 0.9657 and 0.9952; the largest gap is the one at window 12. Coil 2's conclude in 0.375 at 12, 0.895 at 24 (0.022) and 1.000 at 45, against 0.3797, 0.8741 and 0.9987. At no loss the shares at the eight looks run from 0.025 to 0.080 on coil 1 and from 0.040 to 0.080 on coil 2, about the closed form's 0.0500, with standard errors from 0.011 to 0.019. The records' own posterior standard deviations of γ_1 at window 45 scatter from 0.01131 to 0.01387, mean 0.01218, about the linearized 0.01210: the drawn commissioning entries, β_1 and drift move the rows. Their estimates scatter with standard deviation 0.01141 about a mean of -0.05218. For coil 2 the posterior standard deviations run from 0.01039 to 0.01220, mean 0.01104, against 0.01100, and the estimates scatter by 0.01151. The linearization at the 5 %/yr line holds across the records the plan may log.

22.4.2 A smooth path

The reported probability of detection draws the path's wander from the trend rows, the first choice of step 2. A truth whose path is a smooth line has no wander, and for it the trend rows' errors are zero; with the nuisance entries fixed as well (the coils at their rating, β_1 at 0.20, the drift at the stated truth's 0.6 K/yr) only the reading rows' errors remain in equation (22.9). With $\mathbf{J}_{r,t_{\mathrm{as}}}$ the reading rows' lines, the estimate's spread over the records is then

$$\begin{aligned} \varsigma_u^r(t_{\mathrm{as}})^2 &= \mathbf{g}_{\gamma_u}^{\mathsf{T}} \boldsymbol{\Sigma}_{t_{\mathrm{as}}} \mathbf{J}_{r,t_{\mathrm{as}}}^{\mathsf{T}} \mathbf{J}_{r,t_{\mathrm{as}}} \boldsymbol{\Sigma}_{t_{\mathrm{as}}} \mathbf{g}_{\gamma_u}, \\ \mathrm{POD}_u^r(t_{\mathrm{as}}) = \Phi\!\left(\frac{|\gamma_{\mathrm{D}}| - z_\kappa \varsigma_u(t_{\mathrm{as}})}{\varsigma_u^r(t_{\mathrm{as}})} \right), &\qquad \mathrm{PFA}_u^r(t_{\mathrm{as}}) = \Phi\!\left(-\frac{z_\kappa \varsigma_u(t_{\mathrm{as}})}{\varsigma_u^r(t_{\mathrm{as}})} \right), \end{aligned} \tag{22.20}$$

with the fit, and so its verdict, still carrying σ_q. At two years the readings' spread is 0.01455 and 0.01282, against the posterior's 0.02014 and 0.01838: the probability of detection is 0.8940 and 0.9498, against the reported 0.8163 and 0.8741, and the false alarm 0.01138 and 0.00917, against 0.05. At window 45 the readings' spread is 0.00757 and 0.00567, the probability of detection 1.0000 on both coils, and the false alarm 0.00428 and 0.00071. Two hundred records on the smooth path (seeds 24000 to 24199), refitted, conclude at window 24 in 0.910 (standard error 0.020) and 0.935 (0.017), and at windows 36 and 45 in all 200; their estimates scatter at window 24 with standard deviation 0.01394 for coil 1, against 0.01455. The reported probability of detection is conservative for a smooth path from the second year on, as Chapter 10's bands were, which covered the pseudo-true value in 0.990 of fresh records against the 0.90 they claim (Section 10.4.6). At a year the order is the other way, 0.3626 against 0.3797: the rule's threshold $z_\kappa \varsigma_u(12)$ lies beyond the stated loss, and a narrower scatter about the loss crosses it less often.

22.4.3 The limit of exact rows and the curvature term

The rows run on the reduced state, which equation (4.48) makes the limit $\sigma_p \to 0$ of the full state. The check builds the chapter's own full state for the record to window 45 at the hypothetical state: 372 entries, the 142 of the reduced state and each window's five temperatures, under 230 physical rows (5.2) at a common tolerance beside the readings, priors and trend rows. At $\sigma_p = 0.1$ the predictive spreads are 0.01215 and 0.01100, and the probabilities of detection

0.99501 and 0.99873; at 0.01 and 0.001 the spreads are 0.01210 and 0.01100, those of the reduced state to the fifth decimal, with probabilities 0.99524 and 0.99525 for coil 1 against the reduced state's 0.99525, and 0.99873 for coil 2 in every row.

The covariance (22.6) is the Gauss–Newton form, without the curvature term $\mathbf{C}$ of equation (4.44). On one simulated record of the 5 %/yr arm (seed 22000), fitted to window 45, the full Hessian gives Newton standard deviations of the rates of 0.01225 and 0.01093, against Gauss–Newton's 0.01223 and 0.01093, ratios of 1.0010 and 0.9998; the largest entry of $\mathbf{C}$ is 6.33×10^{-4} of the largest entry of $\mathbf{J}^{\mathsf{T}}\mathbf{J}$. The Gauss–Newton covariance is the one reported.

22.4.4 Sensitivity to the choices

Table 22.9 repeats the closed form under other choices. Each row is a variant of the rows or of the point at which they are linearized; the columns give each coil's predictive spread at window 45, its probability of detection there, and its record length for a power of 0.9.

Table 22.9: The probability of detection at window 45 and the record length for a power of 0.9 under other choices of the rows and of the linearization, coil 1 / coil 2.

Variant	$\varsigma_u(45)$ [yr^{-1}]	$\mathrm{POD}_u(45)$ [–]	$T_{90,u}$ [window]
the base fit's rows (reported)	0.01210 / 0.01100	0.9952 / 0.9987	30 / 26
$\sigma_q = 0.0025$	0.00887 / 0.00730	1.0000 / 1.0000	21 / 21
$\sigma_q = 0.01$	0.02010 / 0.01945	0.8177 / 0.8394	56 / 54
the straight line, $\sigma_q \to 0$	0.00741 / 0.00546	1.0000 / 1.0000	19 / 19
no drift entry (a sound discharge sensor on unit 1)	0.01066 / 0.01066	0.9992 / 0.9992	25 / 25
the drift carried from window 24	0.01205 / 0.01098	0.9955 / 0.9988	27 / 25
rating prior sd 0.05	0.01199 / 0.01090	0.9958 / 0.9989	29 / 25
rating prior sd 0.30	0.01211 / 0.01101	0.9952 / 0.9987	30 / 26
survey sd 0.04 on β_1	0.01210 / 0.01100	0.9953 / 0.9987	30 / 26
the load meter at 1.0 kW	0.01177 / 0.01092	0.9967 / 0.9989	29 / 25
the load meter at 6.0 kW	0.01289 / 0.01113	0.9903 / 0.9985	32 / 27
linearized with the other coil not losing	0.01243 / 0.01082	0.9935 / 0.9990	31 / 25
linearized on the 3 %/yr line	0.01215 / 0.01098	0.9950 / 0.9988	30 / 26

The priors and the point of linearization move the record length by at most one window on either coil. The wander moves it most: half the reported σ_q gives 21 windows on each coil and twice it 56 and 54, and the straight line 19. The curve is conditional on Chapter 10's σ_q, as that chapter's wander floor was, and Chapter 10 kept 0.005 against the evidence's lean toward a larger value (Section 10.4.4); a rougher true path would need a longer record. Coil by coil, with the reason: the drift entry costs coil 1 five windows, 30 against 25 with a sound sensor, by the ridge of Section 22.2.1, and coil 2 one window, 26 against 25, through the return air's shared total, the correlation -0.2114 of the two rates. Carried from window 24 instead of 20, the drift shares fewer of coil 1's windows, and its record length falls to 27. The load meter acts on coil 1 more than on coil 2, 29 and 32 windows at 1.0 and 6.0 kW against 25 and 27. With the other coil linearized as unaged, the two coils' shares of the return air change at the point where the rows are read, and the record lengths move one window each way, 31 and 25.

22.4.5 The record as logged and the stated truth

The last check sets the plan against the record as logged and its stated truth (Section 22.6.5). The base fit of Chapter 10, cut at each window t_{as} from 1 to 36, gives coil 1 a probability of a loss of 0.0592, 0.5787, 0.8917 and 0.9487 at windows 3, 6, 9 and 12, then 0.9140 at 18, 0.9066 at 24, 0.9670 at 30 and 0.9981 at 36. It first reaches 0.95 at window 13 (Jan Y2), falls back below it, and holds above it only from window 33 (Sep Y3). Coil 2's reads 0.0724, 0.1042, 0.3697 and

0.7878 at windows 3 to 12, 0.9994 at 18 and 1.0000 from 24; it first reaches 0.95 at window 16 (Apr Y2) and holds. The record shows the finding of Section 22.3.4 once: a monthly reader would have written coil 1's verdict in January of year 2 and seen it lapse, and lapse again after its return, before it held from September of year 3. The as-of spreads at windows 12, 24 and 36, 0.03778, 0.01994 and 0.01463 for coil 1, lie close to the hypothetical 0.03832, 0.02014 and 0.01480 (Section 22.2.2).

The stated truth's net change over windows 0 to 36 is a loss of 4.877 %/yr on coil 1 ($-0.05000\,\mathrm{yr}^{-1}$) and 4.241 %/yr on coil 2 ($-0.04333\,\mathrm{yr}^{-1}$), close to the chapter's 5 %/yr. At a constant loss of those sizes the monthly refits first conclude at a median window of 11 for coil 1 (Nov Y1; 5–95 % range windows 1, Jan Y1, to 29, May Y3) and 12 for coil 2 (Dec Y1; windows 1, Jan Y1, to 30, Jun Y3), and this record's first windows, 13 and 16, sit at the 0.608 and 0.675 quantiles of those distributions. The truth's paths are not lines: coil 1's trend runs at $-0.03\,\mathrm{yr}^{-1}$ to window 24 and then accelerates, the question of Chapter 11, and coil 2's path steps down by 0.07 into window 15, Chapter 12's; unit 1's discharge sensor drifts 0.6 K/yr from window 20. The comparison is a check of scale, not a test of the curve.

Table 22.10 collects the routes. Each row is a route or a reference; the columns give each coil's probability of detection at windows 45 and 24, the false alarm at window 45, and the record length for a power of 0.9 or, for the record as logged, the window of its first verdict.

Table 22.10: The routes to the probability of detection of a 5 %/yr loss, and the record as logged with its stated truth. Pairs are coil 1 / coil 2; simulated shares are from 200 records, with standard errors in Section 22.4.1.

Route	$\mathrm{POD}_1(45)$ [–]	$\mathrm{POD}_2(45)$ [–]	$\mathrm{POD}_1(24)$ [–]	$\mathrm{POD}_2(24)$ [–]	false alarm at 45 [–]	T_{90} or first verdict [window]
closed form on the base fit's rows (reported)	0.9952	0.9987	0.8163	0.8741	0.0500 / 0.0500	30 / 26
simulated records, refitted	0.995	1.000	0.795	0.895	0.035 / 0.050	–
smooth path, closed form	1.0000	1.0000	0.8940	0.9498	0.0043 / 0.0007	–
smooth path, simulated	1.000	1.000	0.910	0.935	–	–
full state, $\sigma_p = 0.1$	0.99501	0.99873	–	–	–	–
full state, $\sigma_p = 0.001$	0.99525	0.99873	–	–	–	–
monthly looks, the as-of sequence	0.9984	0.9996	–	–	0.3682 / 0.3750	–
the record as logged, as-of verdicts	–	–	–	–	–	13 / 16
stated truth, a net loss of 4.877 / 4.241 %/yr: monthly first verdict, median	–	–	–	–	–	11 / 12

22.5 The report (step 6)

The dashboard's power analysis says that a 5 %/yr loss shows in 34 windows; the same test says that a coil which is not losing at all is losing, with probability 0.91 by the shutdown, because the load's growth and the season move the tile with no fouling. Read on the conductance, the existing sensors resolve a 5 %/yr loss by the shutdown with probability 0.995 on coil 1 and 0.999 on coil 2, at a false alarm of 0.05, provided the verdict is read once and the path wanders no more than Chapter 10's σ_q allows (0.818 and 0.839 at twice that wander). A power of 0.9 needs 30 windows of record on coil 1 and 26 on coil 2; better sensors shorten that more than the wander floor suggests, to 17 windows at 0.1 K, but no sensor brings it below the bound of 11.72 windows. And refitting every month and acting on the first verdict raises the false alarm to 0.37, for a probability of detection of 0.9984 in place of 0.9952.

22.5.1 The answers and their bands

Table 22.11 gives the answers in the order of Table 22.1. Each row is a question; the second column gives the reported value from the closed form on the base fit's rows, the third the simulated records' value beside it, or for the smallest loss the rows linearized on the 3 %/yr line, and the fourth the stated rule with the range of the answer over the inputs and choices of step 5.

Table 22.11: The report on the hall's condition-monitoring plan at window 36 (Dec Y3), for a record to window 45 (Sep Y4) at the record's design. Reported values from the closed form (22.10) on the base fit's rows at the hypothetical state; simulated shares from 200 records per arm.

Asset manager's question	Reported value	Simulated records or second route	Rule and range
Can the existing sensors resolve a 5 % per year loss in the time I have?	yes: POD(45) 0.9952 and 0.9987 [–]; today, window 36, 0.9657 and 0.9885 [–]	0.995 and 1.000 [–]	Pr(loss) $\geq$ 0.95 read once; false alarm 0.05 [–]; POD(45) 0.818 to 1.000 over σ_q 0.0025 to 0.01 [–]
How long a record does a power of 0.9 need, and what would better sensors save?	30 (Jun Y3) and 26 (Feb Y3) [window]; at 0.1 K, 17 and 17 [window]	coil 1 at window 30: 0.900 (se 0.021); coil 2 at 24 and 30: 0.895 and 0.970 [–]	18 to 42 and 18 to 37 over κ and the power; 21 to 56 over σ_q [window]
What is the smallest loss the record resolves by the shutdown?	3.480 and 3.167 [%/yr]; 3 %/yr resolved with POD 0.8083 and 0.8696 [–]	linearized on the 3 %/yr line: 3.492 and 3.163 [%/yr]	power 0.9; 2.842 and 2.735 at 0.1 K [%/yr]
If I refit every month and act on the first verdict, when does it come, and how often is it false?	detection 0.9984 and 0.9996, false alarm 0.3682 and 0.3750 [–]; first verdict at a median window 11 (Nov Y1)	semiannual false alarm 0.180 and 0.215 (sequence 0.1902, 0.1942) [–]	the rule at every look; annual looks, false alarm 0.1320 and 0.1327 [–]
Does a stricter rule, a Bayes factor of 19 against no trend, change the answer?	POD^B(45) 0.849 and 0.924 [–]; record for 0.9: 49 (Jan Y5) and 43 (Jul Y4) [window]; false alarm 0.0007 and 0.0006 [–]	0.860 and 0.915 [–]	$B \geq 19$, prior $\mathcal{N}(0,\ 0.10^2)$; coil 1's record 46 to 52 over σ_γ 0.03 to 0.30 [window]

The answers rest on the predictive spread of the base fit's rows and on one choice above all, the wander σ_q: it sets the floor that the spread approaches, and a rougher or smoother path than 0.005 a month moves the record a power needs more than any prior does (Section 22.4.4). For a smooth path the reported probability of detection is conservative from the second year on.

22.5.2 The entry in the maintenance log

The plan is the next entry of the two-unit hall's maintenance log, at window 36 beside the entries of Section 22.1.2. It reads the record's design and uses the record to window 36 only for the check of Section 22.4.5, nothing later. In the values of Table 22.11 it reads:

> *Window 36 (Dec Y3), whether the trend log can see a* 5 %/yr *loss of conductance by the October shutdown (window 45, the last read), from the base fit's rows at the record's design.*
> Rule: a coil is "losing" when $P(\text{rate} < 0) \geq 0.95$, read once, at the shutdown's refit.
> Coil 1: probability of detection 0.995; record needed for 0.9: 30 windows; smallest loss resolved by the shutdown 3.5 %/yr. Coil 2: 0.999; 26 windows; 3.2 %/yr. False alarm 0.05 each. Conditional on the path's wander, $\sigma_q = 0.005$ a month: at 0.01, 0.82 and 0.84, and 56 and 54 windows.
> Monthly refits acted on at the first verdict: false alarm 0.37. Not to be used for work orders.
> A Bayes factor of 19 against no trend: probability of detection 0.85 and 0.92 by the shutdown, 49 and 43 windows for 0.9, false alarm 0.001.

> Action: no change of sensors for the present coils; the shutdown's refit is the single look. A record restarted at a new commissioning test needs 30 windows (coil 1) and 26 (coil 2) before it can carry a work order at this power if unit 1's sensor drifts again from the new record's window 20; 25 and 25 with a sound sensor.
> Next measurement: if a restarted record must decide within two years, 0.1 K sensors on the return air, both discharges and rack 1 (17 windows).

The entry records the rule beside the numbers, because the probability of detection belongs to the rule as much as to the sensors. At what reading a work order is raised, against the costs of false and missed orders, is Chapter 19's question; the false alarm of repeated looks is what that trigger must be calibrated against.

Common mistake

Quoting a monitoring plan's power without its false alarm, or reading the first monthly verdict as the plan's verdict. A test on the dashboard's tiles detects a loss, and also the load; a monthly reader who acts on the first crossing has a false alarm of 0.37 by the shutdown, against the single look's 0.05. State both numbers, on the parameter, for the schedule of looks the plan will actually use.

22.6 The next measurement (step 7)

The report leaves the present coils' detection saturated by the shutdown, 0.995 and 0.999, and a restarted record needing 30 and 26 windows before it can carry a work order. A permanent instrument added at a commissioning test adds rows to every window of the record, and what it takes off the record a power needs is known before any of them is read.

22.6.1 The preposterior of a candidate instrument

A candidate a is a set of rows in every window from commissioning to the last, with Jacobian lines $\mathbf{J}_{a,t_{\mathrm{as}}}$ in the columns of the record's state at the hypothetical state; it is a design alternative for a record that starts at a commissioning test. Its preposterior covariance is equation (2.17) on the record's covariance (22.6),

$$\boldsymbol{\Sigma}_{a,t_{\mathrm{as}}} = \left(\mathbf{J}_{t_{\mathrm{as}}}^{\mathsf{T}}\mathbf{J}_{t_{\mathrm{as}}} + \mathbf{J}_{a,t_{\mathrm{as}}}^{\mathsf{T}}\mathbf{J}_{a,t_{\mathrm{as}}}\right)^{-1}, \qquad \varsigma_u'(t_{\mathrm{as}}) = \sqrt{\mathbf{g}_{\gamma_u}^{\mathsf{T}}\boldsymbol{\Sigma}_{a,t_{\mathrm{as}}}\,\mathbf{g}_{\gamma_u}}, \tag{22.21}$$

and the candidate's record length for a power of 0.9 is equation (22.12) on $\varsigma_u'(t_{\mathrm{as}})$. The candidates are ranked by that record length, because at window 45 the existing probability of detection is already saturated; beside it stand the width taken off the predictive spread at window 45, $1-\varsigma_u'(45)/\varsigma_u(45)$, and the minimum detectable loss there, equation (22.14).

The candidates open to the asset manager are of four kinds. A sound discharge sensor on unit 1, the present one recalibrated or replaced so that the record carries no drift entry. A second discharge sensor on unit 1 at 0.3 K, permanent, a row on $T_{d,1}$ in every window beside the drifting one. Water-side heat meters on both coils, $\sigma_W = 3.0\,\mathrm{kW}$, permanent, each a row on its coil's duty $D_{u,t} = G_{u,t}(T_r - T_s(t))$ in every window: the permanent instrument at the place of the question, the coil's own heat transfer. And a better grade, 0.1 K sensors on the return air, both discharges and rack 1, alone or with the heat meters. The last line is not a measurement: the path known exactly at every window, the floor of equation (10.12), $\sigma_q/(\Delta t\sqrt{45}) = 0.00894\,\mathrm{yr}^{-1}$ at window 45.

22.6.2 The candidates

Table 22.12 ranks the candidates. Each row is an instrument carried from commissioning; the columns give each coil's record length for a power of 0.9, the width taken off each coil's predictive

spread at window 45, the minimum detectable loss there, and the probability of detection at two years. The record lengths are compared as bars in Figure 22.6, two per candidate, with the floor's 12 windows dashed.

Table 22.12: Candidate instruments, each carried from a commissioning test, by the pre-posterior (22.21) at the hypothetical state; pairs are coil 1 / coil 2, the widths taken off $\varsigma_1(45) = 0.01210$ and $\varsigma_2(45) = 0.01100\,\mathrm{yr}^{-1}$.

Candidate	$T_{90,1}$ [window]	$T_{90,2}$ [window]	off $\varsigma(45)$ [%]	min. detectable at 45 [%/yr]	POD at 24 [–]
the existing sensors	30 (Jun Y3)	26 (Feb Y3)	–	3.480 / 3.167	0.8163 / 0.8741
a sound discharge sensor on unit 1 (no drift entry)	25 (Jan Y3)	25 (Jan Y3)	11.9 / 3.1	3.071 / 3.071	0.8859 / 0.8859
a second discharge sensor on unit 1 (0.3 K), permanent	23 (Nov Y2)	25 (Jan Y3)	14.5 / 3.5	2.981 / 3.059	0.9220 / 0.8927
water-side heat meters on both coils (3.0 kW), permanent	24 (Dec Y2)	22 (Oct Y2)	13.7 / 7.4	3.012 / 2.936	0.9143 / 0.9355
0.1 K sensors on return air, both discharges and rack 1	17 (May Y2)	17 (May Y2)	18.6 / 13.9	2.842 / 2.735	0.9673 / 0.9817
heat meters with 0.1 K sensors	17 (May Y2)	17 (May Y2)	20.4 / 14.3	2.781 / 2.721	0.9731 / 0.9832
the path known exactly at every window (the floor)	12 (Dec Y1)	12 (Dec Y1)	26.1 / 18.7	2.583 / 2.583	0.9945 / 0.9945

The ranking follows each candidate's rows. A sound discharge sensor on unit 1 removes the drift entry, and with it the ridge of Section 22.2.1: it takes 11.9 % off coil 1's spread and 3.1 % off coil 2's, and brings both coils to 25 windows. A second discharge sensor on unit 1 reads the same temperature without the drift, so it separates the drift from coil 1 in every window: coil 1's record falls to 23 windows and coil 2's stays at 25. The heat meters read each coil's duty, the split of the heat between the coils that the return air cannot see (Chapter 5, Section 5.6.3): coil 1's record falls to 24 windows and coil 2's to 22.

The step's surprise is the grade. The permanent instrument at the place of the question, a heat meter on each coil's water side, cuts coil 1's record from 30 windows to 24 and coil 2's from 26 to 22; regrading the four temperature sensors to 0.1 K cuts them to 17 and 17. For the first eighteen months the readings, not the wander, set the spread, and the two are about even at 24 (Table 22.2); every temperature row carries both coils, so a better grade narrows both rates at once. Meters added to 0.1 K sensors take a little more off the spread at window 45, 20.4 and 14.3 % against 18.6 and 13.9 %, and nothing more off the record length, 17 windows. No instrument passes the floor's 12 windows.

22.6.3 The recommendation

For the present coils, nothing: the shutdown's single refit at window 45 is the measurement, at probabilities of detection of 0.995 and 0.999. For a record restarted at a new commissioning test that must decide inside two years, 0.1 K sensors on the four temperature points: 17 windows for a power of 0.9, and a probability of detection of 0.9673 and 0.9817 at two years. What the same sensors are worth for the date the capacity limit is crossed is Chapter 23's question, and how often to log is Chapter 24's. Whatever the sensors, the verdict is read at planned looks, the shutdown's refit, and never at the first monthly crossing.

22.6.4 What the analysis showed

The dashboard's power analysis said that the trend log sees a 5 %/yr loss in 34 windows; the same test concludes "losing" on unaged coils with probability 0.91 by the shutdown, because its tile moves with the load and the season. Read on the conductance, a power of 0.9 needs

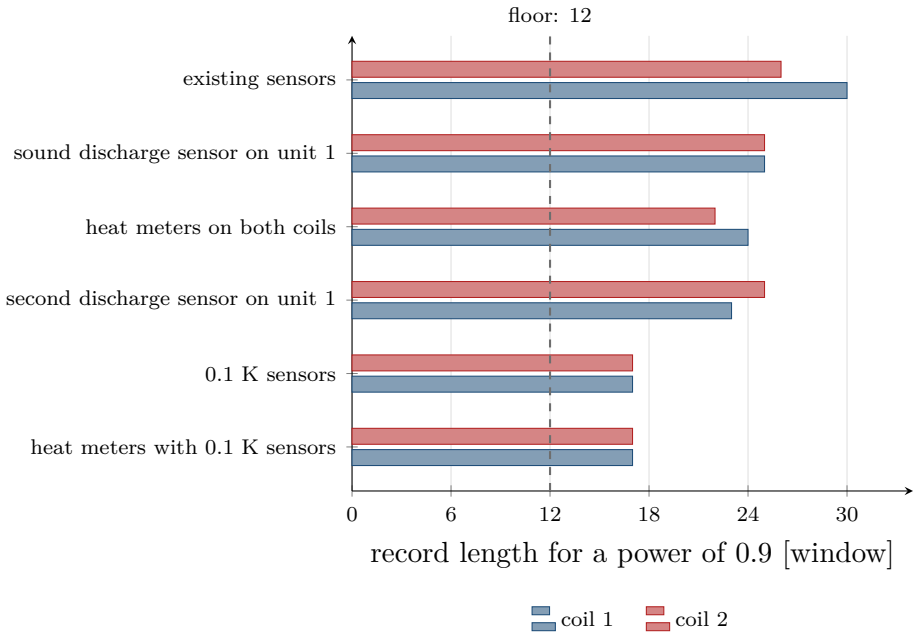


Figure 22.6: The candidate instruments of Table 22.12, each carried from a commissioning test, by the record length they need for a power of 0.9 on each coil, by the preposterior (22.21). Dashed: the record length of a path known exactly at every window, the least any instrument can reach.

30 windows on coil 1 and 26 on coil 2, at a false alarm of 0.05. The physics takes the load, the season and the other coil out of every window before the test is made, and the drift entry, the record's account of unit 1's sensor, costs coil 1 five windows. The probability adds what a power figure on raw performance cannot give: the false alarm beside the power, both from one predictive spread fixed before a reading is taken, and the effect of looking every month, which raises the detection only from 0.9952 to 0.9984 and the false alarm from 0.05 to 0.37. A plan is a rule and a schedule as well as a set of sensors. What the analysis did not test is a change of the logging rate, readings correlated within a window, and a coil whose rate changes during the record.

22.6.5 Limits of the model

A constructed record. The condition record is generated from a stated truth, with the stated sensor errors, and step 5 compares the plan with that truth (Section 22.4.5). It is not field data. The curve is a property of the record's design, and the record itself enters only as a check.

The model's own predictive. The probability of detection is the predictive of the base fit's rows, conditional on $\sigma_q = 0.005$, Chapter 10's choice: a smoother true path is detected sooner, a rougher one later, with record lengths of 21 to 56 windows over σ_q from 0.0025 to 0.01.

Linearized at the stated loss. The rows are read at the 5 %/yr line from rating. The simulated records and the linearizations moved to the 3 %/yr line and to an unaged second coil agree with it within their errors and one window.

A constant true rate. The stated loss is a constant rate. An accelerating or stepped path changes when a verdict comes, as the record as logged shows; a curvature is Chapter 11's entry,

and its power on fresh records was found low there.

One grade for four sensors. The four temperature sensors are regraded together and the load meter separately; a grade per sensor is closer to Chapter 23's ranking. No instrument is priced.

Chapter 10's limits. The readings are independent across windows, one hour a month, the drift is linear from window 20, and each window is in steady state at known inputs.

The next chapters. Chapter 23 asks which added measurement most narrows the date the capacity limit is crossed, rather than the rate, on Chapter 14's path, with the heat meters and the 0.1 K sensors among its candidates; and Chapter 24 asks whether more readings per window lower the readings' part of $\varsigma_u(t_{\mathrm{as}})$, which sets the first eighteen months, once their correlation is carried.

22.7 What to retain

Definition 22.1 (Probability of detection of a monitoring plan)**.** The probability of detection of a monitoring plan is the design form with the indicator of its rule as the utility, its outcomes weighted at a stated value of the monitored entry: $\mathrm{POD}_u(t_{\mathrm{as}}) = \int \mathbf{1}[\Pr(\gamma_u < 0 \mid \tilde{\mathbf{y}}_{0:t_{\mathrm{as}}}) \geq \kappa]\, p(\tilde{\mathbf{y}}_{0:t_{\mathrm{as}}} \mid \gamma_u = \gamma_{\mathrm{D}})\, \mathrm{d}\tilde{\mathbf{y}}_{0:t_{\mathrm{as}}}$, and its false-alarm probability is the same integral at no loss. For linear Gaussian rows both are read from one number fixed before the record is logged, the predictive spread $\varsigma_u(t_{\mathrm{as}}) = \sqrt{[\mathbf{\Sigma}_{t_{\mathrm{as}}}]_{\gamma_u\gamma_u}}$ of the record's preposterior covariance $\mathbf{\Sigma}_{t_{\mathrm{as}}} = (\mathbf{J}_{t_{\mathrm{as}}}^{\mathsf{T}}\mathbf{J}_{t_{\mathrm{as}}})^{-1}$: the estimate is $\mathcal{N}(\gamma_u,\ \varsigma_u(t_{\mathrm{as}})^2)$, so that $\mathrm{POD}_u(t_{\mathrm{as}}) = \Phi(|\gamma_{\mathrm{D}}|/\varsigma_u(t_{\mathrm{as}}) - z_\kappa)$ and the false alarm of a single look is $1-\kappa$. The record length for a power POD^* is the first window with $\varsigma_u(t_{\mathrm{as}}) \leq |\gamma_{\mathrm{D}}|/(z_\kappa + z_{\mathrm{POD}^*})$, bounded below by the wander; the minimum detectable rate is $(z_\kappa + z_{\mathrm{POD}^*})\,\varsigma_u(t_{\mathrm{as}})$. Nested records give as-of estimates with $\mathrm{Cov}(\hat{\gamma}_{t_{\mathrm{as}}}, \hat{\gamma}_{t'_{\mathrm{as}}}) = \varsigma(t'_{\mathrm{as}})^2$ for $t_{\mathrm{as}} \leq t'_{\mathrm{as}}$, and the false alarm of repeated looks is the first passage of that sequence.

Three conclusions carry beyond the hall. A monitoring plan is judged before it is logged, on the parameter, by two numbers read together, its probability of detection and its false alarm, and both come from the predictive spread of the rate, which the rows fix before any reading is taken; a power figure on raw performance has no meaningful false alarm, because the raw figure moves with the load and the season. The spread has a readings' part, which sets the first eighteen months of a record and is about even with the wander at 24, and which a better grade of sensor removes, and a wander's part, which sets the later years and which only a longer record lowers; the record a stated power needs is bounded below by the wander whatever the sensors. And a plan is a rule and a schedule of looks as well as a set of sensors: reading every refit and acting on the first verdict raises the false alarm by the shutdown from 0.05 to 0.37 while the detection moves from 0.9952 to 0.9984, so the looks are planned, and the rule is stated for the schedule actually used.

Table 22.13 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 22.13: The terms of the detectability of degradation and where each is defined. Units for the two-unit hall.

Term	Expression	Units	Meaning	Defined in
Operating characteristic of a plan	eq. (22.1)	–	the design form at a stated value of one entry	22.1.1
Probability of detection; false alarm	$\mathrm{POD}_u(t_{\mathrm{as}})$, eq. (22.2); the same at no loss	–	the power of the trend test and its size	22.1.1
Hypothetical state	$\mathbf{z}^{\mathrm{D}}_{t_{\mathrm{as}}}$, the coils on the stated loss's line	mixed	where a record not yet logged is linearized	22.1.4
The preposterior of the record				
Covariance of a record not yet logged	$\boldsymbol{\Sigma}_{t_{\mathrm{as}}}$, eq. (22.6)	mixed	known before the readings	22.2.1
Predictive spread of the rate	$\varsigma_u(t_{\mathrm{as}})$, eq. (22.7)	yr^{-1}	one number per coil and record	22.2.2
Readings' part and wander's part	Table 22.2	–	the sensors early, the wander late	22.2.2
The estimate before the record	eq. (22.9)	yr^{-1}	$\mathcal{N}(\gamma_u, \varsigma_u(t_{\mathrm{as}})^2)$	22.2.3
Closed form	$\Phi(\lvert\gamma_{\mathrm{D}}\rvert/\varsigma_u - z_\kappa)$; $1-\kappa$, eq. (22.10)	–	detection and false alarm of one look	22.2.3
As-of sequence	$\mathrm{Cov} = \varsigma(t'_{\mathrm{as}})^2$, eq. (22.11)	yr^{-2}	a random walk run backward	22.2.3
Questions and answers				
Record length for a stated power	$T_{90,u}$, eq. (22.12); the bound, eq. (22.13)	window	no sensor passes the wander's bound	22.3.2
Minimum detectable rate	$(z_\kappa + z_{\mathrm{POD}^*})\varsigma_u(t_{\mathrm{as}})$, eq. (22.14)	yr^{-1}, %/yr	the smallest loss resolved	22.3.3
False alarm of repeated looks	eq. (22.16)	–	a first passage of the sequence	22.3.4
Bayes-factor rule	$z^B_u(t_{\mathrm{as}})$, eq. (22.19)	–	a stricter threshold, growing with the record	22.3.5
Dashboard's power analysis	slope over its standard error on the tiles	–	a power with no size	22.3.1
Preposterior of a candidate instrument	eq. (22.21)	–	the record length an instrument saves	22.6.1

22.8 Exercises

Exercise 22.1. *Formulation.* Write coil 2's probability of detection for a record of 24 windows with $\kappa = 0.90$ in the form of Table 22.1 and evaluate it. (A design question, equation (22.2) with $t_{\mathrm{as}} = 24$ and $\kappa = 0.90$: $\Pr(\Pr(\gamma_2 < 0 \mid \tilde{\mathbf{y}}_{0:24}) \geq 0.90 \mid \gamma_2 = \gamma_{\mathrm{D}})$, which by equation (22.10) with $\varsigma_2(24) = 0.01838$ is $\Phi(2.7908 - 1.2816) = 0.9344$.)

Exercise 22.2. *Derivation.* Show that a path known exactly at every window gives $\varsigma = \sigma_q/(\Delta t\sqrt{t_{\mathrm{as}}})$, and hence the bound (22.13); evaluate it at $\kappa = 0.95$ and 0.99 with a power of 0.9. (With every $g_{u,t}$ known, only the t_{as} trend rows act on γ_u, each with the line $-\Delta t/\sigma_q$, so the information is $t_{\mathrm{as}}\Delta t^2/\sigma_q^2$ and $\varsigma = \sigma_q/(\Delta t\sqrt{t_{\mathrm{as}}})$, the floor (10.12); equation (22.12) asks $\varsigma \leq \lvert\gamma_{\mathrm{D}}\rvert/(z_\kappa + z_{\mathrm{POD}^*})$, that is $\sqrt{t_{\mathrm{as}}} \geq (z_\kappa + z_{\mathrm{POD}^*})\sigma_q/(\Delta t\lvert\gamma_{\mathrm{D}}\rvert)$: $(2.9264 \times 0.0600/0.05129)^2 = 11.72$ windows at $\kappa = 0.95$, and 17.81 at 0.99.)

Exercise 22.3. *Derivation.* From equation (22.11), show that the standardized as-of estimates at windows 24 and 45 have correlation $\varsigma(45)/\varsigma(24)$, and compute coil 1's false alarm with two looks, at windows 24 and 45. (With $\mathrm{Cov}(\hat{\gamma}_{24}, \hat{\gamma}_{45}) = \varsigma(45)^2$, the correlation of $\hat{\gamma}_{24}/\varsigma(24)$ and $\hat{\gamma}_{45}/\varsigma(45)$ is $\varsigma(45)^2/(\varsigma(24)\varsigma(45)) = \varsigma(45)/\varsigma(24) = 0.01210/0.02014 = 0.6008$; at no loss

the two looks stay silent when both standardized estimates exceed $-z_\kappa$, so the false alarm is $1 - \Phi_2(z_\kappa, z_\kappa; 0.6008) = 0.0844$.)

Exercise 22.4. *Application.* Compute coil 1's minimum detectable loss for a record of 24 windows. (By equation (22.14), $2.9264 \times 0.02014 = 0.05895\,\mathrm{yr}^{-1}$, a loss of $5.725\,\%/\mathrm{yr}$.)

Exercise 22.5. *Application.* Find the record length for a power of 0.9 under $\kappa = 0.99$ on both coils. (By equation (22.12) with $z_\kappa = \Phi^{-1}(0.99)$, the first windows at which the predictive spread falls below $|\gamma_{\mathrm{D}}|/(z_\kappa + 1.2816)$: 38 windows (Feb Y4) on coil 1 and 33 (Sep Y3) on coil 2.)

Chapter 23

Sensor Selection for Prognostics

23.1 The dashboard's noisiest trace and the crossing date

It is December of year 3, window 36, and the two-unit hall's maintenance log carries Chapter 14's forecast: the coils' total conductance crosses the capacity study's limit in May of year 5 at the median, between August of year 4 and October of year 7, and the record is to be refitted at window 45 (Sep Y4), before the work order of the October shutdown is written. The planner has a small budget to firm up that date before the refit, and the plant engineer turns to the dashboard's five traces to ask which one is read least well. Against its own straight line through windows 0 to 36, the return-air trace scatters with a residual standard deviation of 1.327 K on a 0.3 K sensor, 4.42 times its grade; rack 1's inlet scatters at 4.37 times its grade, unit 1's discharge at 4.06, unit 2's at 3.92 and the load meter at 2.03. The return air is the noisiest trace for its sensor, and it is also more of the sensor already there: a return-air sensor at each unit costs R$ 1 800. Three return-air readings averaged cut the trace's noise by $1/\sqrt{3}$, so, by the planner's arithmetic, the crossing band should narrow by 42.3 %. The conclusion the dashboard invites is to buy the two return-air sensors and to expect a band about 40 % narrower at the refit.

The dashboard cannot say whether the return air's scatter is reading error at all, when the season and the load move the return air with no error in it. It cannot say which entries of the state the crossing date is a function of, or which sensors read those entries rather than the hall's total. Nor can it say what the refit at window 45 already takes off the band before anything is bought, how much of the band is the record's indecision about coil 1's curvature, which no sensor installed now can settle in nine windows, or what a narrower band is worth per real spent. The question of this chapter is

for each instrument that could be installed now and logged to the refit at window 45, how much it narrows the crossing date's 5–95 % *band and the band of the margin at the shutdown beyond what the continued record already does, against the floor at which the rates are known exactly, and at what price.*

It is answered with the design form of Chapter 2 on Chapter 14's extended state, over Chapter 14's family of drawings, before any of the readings is logged.

Jargon: sensor selection for prognostics, value of a sensor for a date, the drawing's own band and the family's indecision, the trend's floor at the refit, load-free reading of a coil, water-side temperature difference without flow metering, value per real

Sensor selection for prognostics (instrumentation design for a remaining-life forecast) ranks the instruments that could be added to a plant by what they would do to a forecast date, before any of their readings exists. The *value of a sensor for a date* is the width it takes off the date's band beyond what the record's own continuation takes off. A family's band, mixed over its drawings, has two parts: *the drawing's own band*, Chapter 14's within-drawing width, and *the family's indecision*, its between-drawing width, which only evidence on the drawings removes. *The trend's floor at the refit* is the band left when the rates are known exactly at that refit. A *load-free reading* of a coil reads its conductance through temperatures alone, with no load meter in it. A *water-side temperature difference without flow metering* reads a coil's duty up to an unknown constant factor, the water flow. The *value per real* is the width a candidate takes off per R$ 1 000 of its price.

Key idea

A sensor is worth what it takes off the answer the asset manager needs, here the band of a crossing date, and not what it takes off the noise of its own trace. The date is a ratio of the coils' margin to their rate, and within a drawing it is set by the total conductance's level and rate. The hall reads its total by two routes. The return air reads $(Q + Q_o)/(G_1 + G_2)$ and needs the load meter, whose error stays however many return-air sensors are added. The units' discharge drops over the approach read G_u/C_u with no load in them, and unit 1's drop runs through the drifting sensor. What frees that second route of the drift therefore narrows the date, and more of the sensor on the first route does not.

23.1.1 Valuing an instrument by a date

Chapter 2's design form, equation (2.16), values a candidate measurement a before its outcome is known by the expected value of a quantity U_{p} of the posterior the outcome would produce. For a forecast the quantity is the width of the date's band: with t^* the first passage of Chapter 14, equation (14.1), and $\Pr(t^* \le s \mid \mathbf{y}, \tilde{\mathbf{y}}_a)$ its distribution function after the outcome $\tilde{\mathbf{y}}_a$, the band's ends are its inverses at 0.05 and 0.95, and U_{p} is their difference. For rows that are linear in the state the posterior covariance after the outcome is the preposterior (2.17) and does not depend on the values read; the outcome moves only the posterior's center. A forecast from a family of drawings mixes the drawings by their probabilities, Chapter 7's mixture (7.10). Held at today's probabilities P_M and today's MAP centers $\hat{\mathbf{z}}_M$, the candidate's band is

$$
\begin{aligned}
F_{t^*,a}(s) &= \sum_M P_M \, \Pr(t^* \le s \mid \mathbf{z}_M \sim \mathcal{N}(\hat{\mathbf{z}}_M, \boldsymbol{\Sigma}_{M,a})), \qquad \mathcal{B}_a = F^{-1}_{t^*,a}(0.95) - F^{-1}_{t^*,a}(0.05), \\
\bar{\mathcal{B}}_a &= \sum_M P_M \, \mathcal{B}_{M,a}, \qquad \Delta\mathcal{B}_a = \mathcal{B}_{\mathrm{cont}} - \mathcal{B}_a, \qquad \Delta\bar{\mathcal{B}}_a = \bar{\mathcal{B}}_{\mathrm{cont}} - \bar{\mathcal{B}}_a,
\end{aligned}
\tag{23.1}
$$

with $\boldsymbol{\Sigma}_{M,a}$ drawing M's preposterior covariance with the candidate's rows, $\mathcal{B}_a$ the family's band, $\mathcal{B}_{M,a}$ drawing M's own band, $\bar{\mathcal{B}}_a$ the own bands weighted, and $\Delta\mathcal{B}_a$ and $\Delta\bar{\mathcal{B}}_a$ the windows candidate a takes off the family's and the own band of the record's continuation, the plan the asset manager already has.

The family's band narrows under a candidate in two ways: within each drawing, by its own band, and by moving the drawings' probabilities. Equation (23.1) holds the probabilities and the centers, as Chapter 14's step 7 did, and so measures the first way only. Predicted records,

refitted and reweighted, measure both (Section 23.2.5). Chapter 14 found the refit at window 45 to be the measurement; this chapter ranks what to add to it.

> **Principle 23.1.** *Rank an instrument by the width it takes off the answer the decision needs, beyond what the record's own continuation takes off, against the floor of the answer's own physics, and per real spent. A sensor's precision on its own trace is not its value.*

23.1.2 The two-unit hall carried to the refit (step 0)

The plant graph is Chapter 14's, Figure 14.1, cited and not derived again: Chapter 10's window graph at each of the record's windows with its five physical rows (5.2) at the window's known inputs $T_s(t)$ and $Q_o(t)$; the trend rows of both coils, with coil 1's curvature $\dot{\gamma}_1$ in the drawings that free it and Chapter 12's step at window 15 in the drawings that carry one; the shared fraction β_1; the drift node $\dot{b}_{d,1}$ on unit 1's discharge reading from window 20; and the forecast windows, each with its requirement node $G_{\mathrm{tot},t} \geq G^*$. The modeling choices of step 0 are Chapter 5's paragraph, carried through Chapter 10: steady state within each window, the airflow pinned at the units' fixed fan speed, the season's quantities as known inputs, and one conductance per coil.

What the question adds. Windows 37 to 45 (Jan Y4 to Sep Y4) are read windows again: the record continues, and each window has its block with its load and the five existing readings. On those windows each candidate hangs its reading squares on the node it reads. A water-side heat meter's square hangs on coil u's duty $D_{u,t} = G_{u,t}(T_r - T_s(t))$, a node fixed by $g_{u,t}$ and the return air. A second discharge-air sensor's square sits on $T_{d,1}$ beside the existing one, with no drift edge. A water-side temperature difference hangs on $D_{u,t}$ through a new free node, the log ι_u of coil u's water flow against the balancing report's. The return-air squares sit on T_r and the supply-air square on the cold aisle's T_c. An airflow station's square sits on the capacity rate C_u, which is pinned and has no edge into the state.

The stated inputs of the question. The plan the candidates are valued against is Chapter 14's: the record continues to window 45 (Sep Y4) and the family is refitted there, before the work order for the October shutdown at window 46 (Oct Y4). Every candidate is installed at window 36 and logged with the record, one representative hour a month in windows 37 to 45, its new points riding the dashboard's monthly hour at R$ 0 a year, a stated input of this chapter; the continuation itself is candidate zero, at R$ 0. Table 23.1 lists the candidates. Each row is an instrument; the columns give what it reads, its grade and its price, installed and connected, one-off. The plant's cost schedule prices no sensor, so the prices are stated inputs of this chapter; the one comparator that is not permanent, a reference condition test at window 45, is the schedule's inspection, R$ 1 500: one crew visit that reads the hall's four temperature points at reference grade. The water-side temperature difference has no flow meter: each coil's chilled-water capacity rate is the balancing report's $C_w = 14.0\,\mathrm{kW/K}$ (3.34 L/s), right to 10 %, which the state carries as the entry $\iota_u \sim \mathcal{N}(0,\ 0.10^2)$, constant from window 37.

The maintenance log so far. The two-unit hall's log holds the commissioning test (window 0); the entries of Chapters 5 and 6 at window 18 (Jun Y2); Chapter 7's at window 25 (Jan Y3), which found the drift and asked for a side-by-side check of unit 1's discharge sensor, a check not made (Section 7.5.2); and Part III's entries at window 36. Chapter 14's entry at window 36 logged the crossing (May Y5, Aug Y4 to Oct Y7), survival to the shutdown 0.84, and the instruction to refit at window 45 before the work order (Section 14.5.2). After it come, at the same window, the entries of Chapters 17 to 22: the cleaning at the shutdown, the inspection

Table 23.1: The candidate instruments, installed at window 36 (Dec Y3) and logged one hour a month with the record in windows 37 to 45, and the one-off comparator at window 45. Prices are stated inputs, installed and connected; the reference test is the cost schedule's inspection.

Candidate	Reads	Grade	Price [R$]
the record continued to window 45 (no new sensor)	T_r, $T_{d,1}$, $T_{d,2}$, Q, T_1, as now	the existing five	0
water-side heat meter, coil 1	$D_1 = G_1(T_r - T_s)$	3.0 kW	12 000
water-side heat meter, coil 2	D_2	3.0 kW	12 000
water-side heat meters, both coils	D_1, D_2	3.0 kW each	24 000
second discharge-air sensor, unit 1	$T_{d,1}$, without the drift	0.3 K	900
water-side temperature difference, coil 1	$D_1/(C_w e^{\iota_1})$	0.1 K, flow to 10 %	2 400
water-side temperature difference, coil 2	$D_2/(C_w e^{\iota_2})$	0.1 K, flow to 10 %	2 400
water-side temperature differences, both coils	both	0.1 K, flow to 10 %	4 800
a return-air sensor at each unit (two)	T_r, twice	0.3 K each	1 800
a supply-air sensor (mixed supply)	$T_c = (T_{d,1} + T_{d,2})/2$	0.3 K	1 200
an airflow station on each unit	C_1, C_2 (pinned)	2 % of flow	28 000
comparator: a reference condition test at window 45 (one-off)	T_r, $T_{d,1}$, $T_{d,2}$, T_1 at window 45	0.1 K, four points	1 500

interval, the work-order trigger, the crew's visiting order, the warranty claims and the trend log's power. This chapter's entry is dated window 36 as well, and it asks which instrument to add before the refit.

The modeling choice at the graph. One choice is new here, and no question block owns it: *the candidates are installed at window 36 and logged like the record*, one representative hour a month, their errors independent and at their grade. This buys candidates valued on the record's own design, beside the readings the record already has. It fails if a new sensor is installed out of calibration, as the existing unit-1 sensor came to be read; how often to log is Chapter 24's question.

The graph. Window 45's block of the graph, with every candidate's square, is drawn in Figure 23.1. Down the middle runs Chapter 5's window: the load and the other heat feed the return air, the return air feeds the two coils, each coil passes its duty to the chilled water and sends its discharge air to the cold aisle, and the cold aisle feeds rack 1's inlet. Unit 1's candidates stand in a column on the left and unit 2's on the right, each square linked to the node it reads. Black boxes are the existing readings. Heavy blue boxes are candidates that read a coil's own share, the split between the coils; heavy red boxes read the hall's total through the return air and so through the load meter; the heavy gold box, the supply air, reads the total through the cold aisle's drop, which needs no meter; the gray dashed box reads a pinned quantity. The nodes the physical rows fix are thin-rimmed light boxes.

23.1.3 The asset manager's questions (step 1)

The asset manager asks five questions of the candidates, and step 1 writes each as a quantity of the posterior that the record continued to window 45, with the candidate's rows, would produce. Table 23.2 lists them. The first column is the question in the asset manager's words, the first of them the drill question; the second the general form of its category from Chapter 2; the third its mathematical form for the hall; the fourth its category; the fifth the computation of Section 2.3.6 that answers it; and the last the subsection where it is answered. Whether to clean

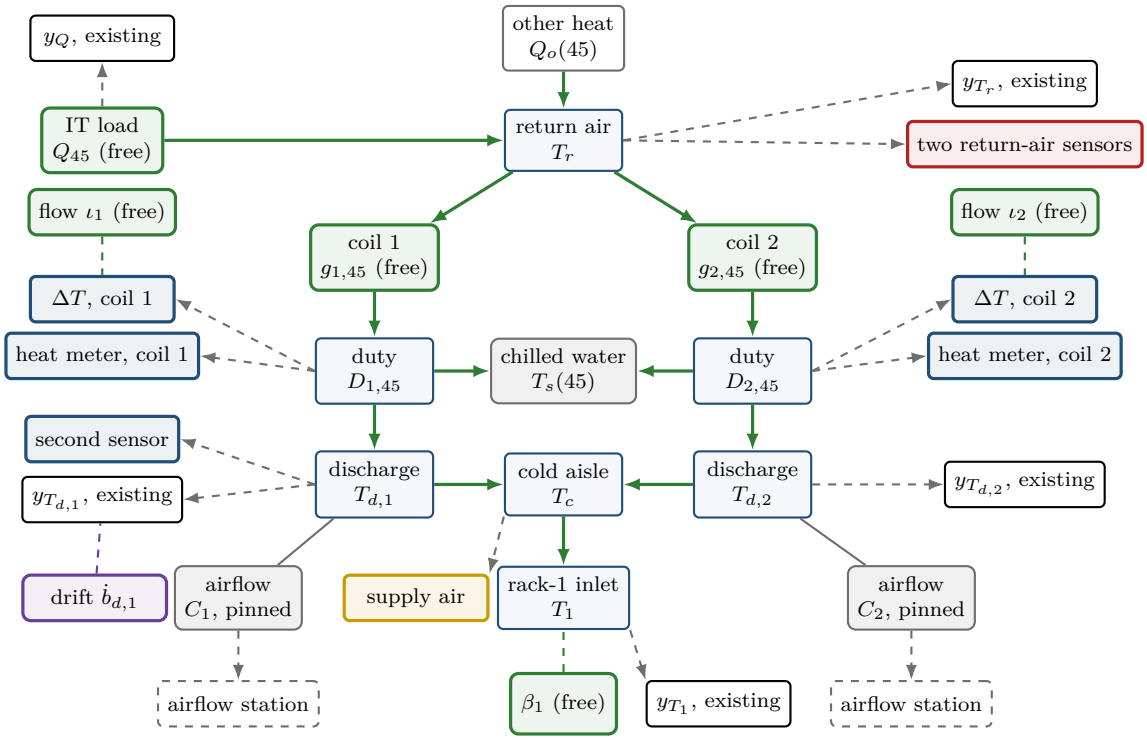


Figure 23.1: Window 45 (Sep Y4) of the two-unit hall's graph with every candidate's reading square. Free (green): the IT load Q_{45}, the log-derates $g_{1,45}$ and $g_{2,45}$ (tied to window 44 and to the trend entries by the trend rows of Figure 14.1, not drawn), the shared fraction β_1, and the flow factors ι_1, ι_2 of the water-side temperature differences, each with the balancing report's prior $\mathcal{N}(0,\ 0.10^2)$; purple: the drift rate $\dot{b}_{d,1}$, which acts on unit 1's existing discharge reading only. Thin-rimmed light boxes: the duties and temperatures fixed by the physical rows (5.2). Pinned (gray): the chilled-water supply, a known input, and the airflows C_1, C_2; the other heat Q_o is a known input (white). Observed (dashed arrows): black boxes, the five existing readings; heavy blue boxes, the candidates that read a coil's share (heat meters, temperature differences, the second discharge sensor on unit 1); the heavy red box, the return-air pair, which reads the hall's total through the load meter; the heavy gold box, the supply-air sensor, which reads the total through the cold aisle's drop with no meter, and unit 1's discharge without the drift; gray dashed, the airflow stations on the pinned airflows. The reference test at window 45 reads T_r, $T_{d,1}$, $T_{d,2}$ and T_1 once at 0.1 K, without the drift; it is not drawn.

at the shutdown is Chapter 17's question, how often to inspect Chapter 18's, whether the trend log resolves a loss at all Chapter 22's, and how often to log Chapter 24's. What measurement remains after the instruments asks nothing new of the candidates; it is step 7, Section 23.6.

Table 23.2: The questions asked of the candidate instruments at window 36 (Dec Y3), valued at the refit at window 45 (Sep Y4). t^* is the first passage of $G_{\mathrm{tot},t}$ below G^*; P_M the drawings' probabilities, held; $\mathbf{\Sigma}_{M,a}$ drawing M's preposterior covariance with the continuation and candidate a; c_a the candidate's price.

Asset manager's question	General form	Mathematical form	Category	Computation	Answered in
Which added measurement most tightens the remaining-life band?	(2.16) with U_{p} the band's width; (2.17); (7.10)	$\mathcal{B}_a$ of (23.1), the band of t^* from window 45 under each $\mathbf{\Sigma}_{M,a}$ mixed at the held P_M; $\Delta\mathcal{B}_a$; the own bands $\bar{\mathcal{B}}_a$ beside it	design	preposterior	23.3.1
Which one makes the call at the shutdown surest?	(2.16); (2.21), third line	the 5–95 % band of $G_{\mathrm{tot},46}$ under each $\mathbf{\Sigma}_{M,a}$	design	preposterior; trend	23.3.2
Which is worth its price?	(2.16) with the price c_a beside the width	$\Delta\bar{\mathcal{B}}_a/c_a$ per R$ 1 000; the price at which a candidate matches the best	design	preposterior	23.3.3
Why does a sensor on unit 1 narrow the date and two more return-air sensors do not?	(2.17); (2.9) with a direction known	the delta form's $\mathrm{Var}\, t^*$ under the continuation, conditioned on each direction known; the spread of $G_{\mathrm{tot},45}$ under each candidate; one window's conductance by each route	estimation	preposterior	23.3.4
How much of the band can any sensor reach, and what is left?	(2.17); (7.10); (2.14)	the floors (rates known; rates and path known); $\mathcal{B}_a$ against $\bar{\mathcal{B}}_a$; the share of predicted records that decide coil 1's curvature	design; hypothesis comparison	preposterior; evidence	23.3.5

23.1.4 The evidence as rows (step 2)

Step 2 attaches the record, its continuation and each candidate to the graph as rows, on every drawing of Chapter 14's family carried to window 45.

The family. The drawings are pairs $M = (\mathcal{C}_c, j)$ of coil 1's curvature state and a March step, Chapter 14's family (Section 14.1.4). Six of its eighteen are read here: $\mathcal{C}_0$ and $\mathcal{C}_1$ crossed with no step ($j = \text{none}$), coil 2 stepped at window 15 ($j = 2$) and unit 2's discharge sensor stepped there ($j = T_{d,2}$), at Chapter 14's prior (14.2) restricted to them. Their posterior probabilities P_M are equation (14.8) over the six.

The state. Each drawing's stacked state, Chapter 10's equation (10.3) with $\dot{\gamma}_1$ under $\mathcal{C}_1$ and the step δ_j under a step, is carried to window 45, and two flow entries are appended for the water-side temperature differences:

$$\mathbf{z}_{M,45} = (Q_0,\ g_{1,0},\ g_{2,0},\ \ldots,\ Q_{45},\ g_{1,45},\ g_{2,45},\ \beta_1,\ \gamma_1,\ \gamma_2,\ \dot{b}_{d,1},\ \dot{\gamma}_1,\ \delta_j,\ \iota_1,\ \iota_2)^{\mathsf{T}}, \qquad (23.2)$$

146 entries on the leading drawing $(\mathcal{C}_1, 2)$: the 46 windows' $(Q_t, g_{1,t}, g_{2,t})$ and eight more. Its base rows number 328: the 230 readings of windows 0 to 45, the two rating priors, the survey prior, the 90 trend rows (14.4), the drift prior, the curvature's and the step's priors and the two flow priors $\iota_u/0.10$. The rows of windows 37 to 45 are the continuation's: the existing five readings of each, equation (5.7) with unit 1's discharge row (10.4).

The candidate rows. Each candidate adds rows in windows 37 to 45, with the temperatures and duties the forward equations (5.3) make of the window's entries:

$$\begin{aligned}
r_{W,u,t} &= \big(D_{u,t}(\mathbf{z}) - y_{W_u,t}\big)/\sigma_W, \quad \sigma_W = 3.0\,\text{kW} && \text{heat meter on coil } u,\\
r_{d',t} &= \big(T_{d,1}(\mathbf{z}_t) - y_{T'_{d,1},t}\big)/0.3 && \text{second discharge sensor, no drift,}\\
r_{\Delta,u,t} &= \big(D_{u,t}(\mathbf{z})/(C_w \mathrm{e}^{\iota_u}) - y_{\Delta_u,t}\big)/0.1 && \text{water-side difference on coil } u,\\
r_{r',t} &= \big(T_r(\mathbf{z}_t) - y_{T'_r,t}\big)/0.3 && \text{each new return-air sensor,}\\
r_{c,t} &= \big(T_c(\mathbf{z}_t) - y_{T_c,t}\big)/0.3 && \text{supply-air sensor,}\\
r_{C,u,t} &= \big(C_u - y_{C_u,t}\big)/(0.02\,C_u) && \text{airflow station on unit } u,
\end{aligned} \tag{23.3}$$

$t = 37, \ldots, 45$, with $C_w = 14.0\,\text{kW/K}$. The reference test reads T_r, $T_{d,1}$ without the drift, $T_{d,2}$ and T_1 once, at window 45, each at 0.1 K. The candidates' channels add 76 rows in all, and a candidate uses its own rows only. The readings are not yet logged, so the rows are evaluated at predicted readings: each drawing's MAP path carried from window 36 by its trend (Chapter 14's extension), with the loads at the record's growth, 138.560 kW at window 45. For linear rows the Jacobian, not the readings, sets the preposterior covariance.

The candidate rows at window 45. Table 23.3 gives the lines of the normalized rows at window 45 on the leading drawing $(\mathcal{C}_1, 2)$, at its MAP path. Each row is a reading, existing or candidate; the columns are the entries Q_{45}, $g_{1,45}$, $g_{2,45}$, β_1, $\dot{b}_{d,1}$, ι_1 and ι_2, with the load column in kW^{-1}, the drift's in yr/K and the others dimensionless. At that path the duties are $D_{1,45} = 70.804$ and $D_{2,45} = 73.634\,\text{kW}$, the water-side differences 5.057 and 5.260 K, and the return air 25.924°C over a supply of 12.939°C.

Table 23.3: The lines of the normalized reading rows at window 45 (Sep Y4) on the leading drawing $(\mathcal{C}_1, 2)$, at its MAP path carried by its trend. The load column is in kW^{-1}, the drift's in yr/K, the others dimensionless.

Row	Q_{45} [kW^{-1}]	$g_{1,45}$ [–]	$g_{2,45}$ [–]	β_1 [–]	$\dot{b}_{d,1}$ [yr/K]	ι_1 [–]	ι_2 [–]
return air (existing, and each new one)	0.2997	−21.2169	−22.0651	0	0	0	0
unit 1's discharge (existing)	0.1363	−33.2488	−10.0332	0	7.2222	0	0
unit 2's discharge (existing)	0.1297	−9.1850	−34.0970	0	0	0	0
load meter (existing)	0.3333	0	0	0	0	0	0
rack 1's inlet (existing)	0.1681	−21.2169	−22.0651	24.0731	0	0	0
heat meter, coil 1	0.1634	12.0319	−12.0319	0	0	0	0
heat meter, coil 2	0.1699	−12.0319	12.0319	0	0	0	0
second discharge sensor, unit 1	0.1363	−33.2488	−10.0332	0	0	0	0
water-side difference, coil 1	0.3501	25.7827	−25.7827	0	0	−50.5742	0
water-side difference, coil 2	0.3641	−25.7827	25.7827	0	0	0	−52.5960
supply air	0.1330	−21.2169	−22.0651	0	0	0	0
airflow station, each unit	0	0	0	0	0	0	0

The lines are read by what each reading sees. The return air sees the hall's total: both coils' columns carry nearly the same slope, and the load enters with the largest positive one of the temperatures. A heat meter reads the split: with the return air free, a higher $g_{1,45}$ takes

duty from coil 2, so its line is +12.0319 on its own coil and −12.0319 on the other. The second discharge sensor's line is the existing sensor's with a zero in the drift's column, where the existing line carries $7.2222 = (45-19)\,\Delta t/0.3$. The water-side difference carries the heat meter's line scaled to its grade, and in its flow column $-D_{1,45}/(C_w \times 0.1) = -50.5742$ (Exercise 23.3): a constant unknown flow absorbs the reading's level and leaves its change. The supply air has the return air's slopes on the coils with a smaller load slope, because the air balance (5.4) gives $T_c = T_r - (Q + Q_o)/(C_1 + C_2)$. The airflow station's line is zero: it reads the airflow, which the model pins (Chapter 5's choice), so its rows add the same constant to every drawing's objective and move no answer in any record.

Which row acts on which entry. The pattern is Chapter 14's, equation (14.6), carried to window 45: each candidate row touches Q_t, $g_{1,t}$ and $g_{2,t}$ of its own window and, for a water-side difference, ι_u; only the existing unit-1 discharge row touches the drift; no candidate row touches a trend entry. The second line of the pattern, the gradient of the date from the refit, has the form of Chapter 14's with window 36 moved to 45,

$$\mathbf{g}_*^{\mathsf{T}} = (* \text{ on } g_{1,45},\ g_{2,45},\ \gamma_1,\ \gamma_2,\ \dot{\gamma}_1;\ \ 0 \text{ elsewhere in } \mathbf{z}_{M,45}), \tag{23.4}$$

so a candidate reaches the date only through the covariance of five entries, and it reaches them only through the windows' log-derates and the drift, which share columns with the trend entries in the trend rows and in unit 1's discharge row.

Modeling choices at the rows. Five choices make these rows, each named once. *The drawings' probabilities and centers are held at today's*: this buys a closed-form preposterior covariance for every candidate, independent of the outcome, and it leaves out a sensor's power to decide between the drawings, which the predicted records of Section 23.2.5 measure. *Six drawings of eighteen*: they keep the predicted records affordable, and step 5 repeats the ranking on all eighteen. *The water flow constant and right to* 10 %: this buys a water-side difference that reads the duty's change, and it fails if the flow is rebalanced between windows 37 and 45. *The airflow pinned*, Chapter 5's choice, not chosen again here: it makes the airflow station worth nothing by construction, and step 5 frees the airflow to show what the station would read then. *The prices are stated inputs*, never fitted: the value per real moves inversely with them, and step 4 gives the price at which each candidate would match the best.

23.2 The posterior of the state (step 3)

Every question of Table 23.2 is read from two objects: the drawings' probabilities today, and each drawing's preposterior covariance at the refit under each candidate, computed from the rows of step 2 before any of windows 37 to 45 is read.

23.2.1 The family today

The eighteen drawings of Chapter 14, refitted on the record to window 36, reproduce that chapter's evidences: the leading drawing $(\mathcal{C}_1, 2)$ has $\ln p(\mathbf{y} \mid M) = -65.6599$, as there. The six read here hold 0.9615 of the eighteen's posterior. Renormalized over the six by equation (14.8), their probabilities are 0.3565 for $(\mathcal{C}_1, 2)$, 0.2143 for $(\mathcal{C}_0, 2)$, 0.1670 for $(\mathcal{C}_1, T_{d,2})$, 0.1141 for $(\mathcal{C}_1, \text{none})$, 0.1080 for $(\mathcal{C}_0, \text{none})$ and 0.0401 for $(\mathcal{C}_0, T_{d,2})$; jointly coil 1's curvature has $\Pr(\mathcal{C}_1 \mid \mathbf{y}) = 0.6376$, coil 2's step 0.5708, unit 2's sensor step 0.2071 and no step 0.2221.

Today's forecast over the six, by Chapter 14's draws with 40 000 per drawing, has its median crossing at window 52.56 (window 53, May Y5) and a band from 43.74 to 81.84, 38.10 windows wide, with $\Pr(t^* \le 46 \mid \mathbf{y}) = 0.1607$; the eighteen give 52.55, 43.74 to 81.67 and 0.1595. Each drawing's own band, weighted, is 27.49 windows, so the family's indecision adds 10.61.

23.2.2 The preposterior of a drawing at the refit

The derivation form is Chapter 5's (Sections 5.2.1 and 5.2.2), applied by Chapter 14 to each drawing's state carried past the record (Section 14.6.1), and not derived again. For drawing M and candidate a, with $\mathbf{J}_{\text{now}}$ the lines of the record's rows to window 36 and of every prior and trend row to window 45, $\mathbf{J}_{\text{cont}}$ those of the existing readings of windows 37 to 45, and $\mathbf{J}_a$ those of the candidate's rows (23.3), all evaluated at the drawing's predicted state, the preposterior covariance of the state (23.2) is

$$\boldsymbol{\Sigma}_{M,a} = \left(\mathbf{J}_{\text{now}}^{\mathsf{T}}\mathbf{J}_{\text{now}} + \mathbf{J}_{\text{cont}}^{\mathsf{T}}\mathbf{J}_{\text{cont}} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a\right)^{-1}. \tag{23.5}$$

This is Chapter 2's preposterior (2.17) twice over: $\boldsymbol{\Sigma}^{-1} = \mathbf{J}_{\text{now}}^{\mathsf{T}}\mathbf{J}_{\text{now}}$ is the information of the record carried to window 45 with no new reading, and the continuation's information $\boldsymbol{\Lambda}_{\text{cont}} = \mathbf{J}_{\text{cont}}^{\mathsf{T}}\mathbf{J}_{\text{cont}}$ and the candidate's $\boldsymbol{\Lambda}_a = \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a$ add to it. With every row normalized the information matrix is $\mathbf{J}^{\mathsf{T}}\mathbf{J}$; with unnormalized rows and their weights it is $\mathbf{J}^{\mathsf{T}}\mathbf{W}\mathbf{J}$. The loads of windows 37 to 45 have no row of their own besides the meter's, so a negligible row holds them, which matters only when no reading is logged there; they do not enter the date. The airflow station's rows have zero lines, so $\boldsymbol{\Sigma}_{M,\text{air}} = \boldsymbol{\Sigma}_{M,\text{cont}}$ exactly.

The band of each drawing is read from $\boldsymbol{\Sigma}_{M,a}$ by Chapter 14's draws (Section 14.3.1), started at the refit:

1. draw 10 000 values of the path entries $g_{u,36}, \ldots, g_{u,45}$ of both coils and of the trend entries γ_1, γ_2 and, under $\mathcal{C}_1$, $\dot{\gamma}_1$ from their block of $\mathcal{N}(\hat{\mathbf{z}}_M, \boldsymbol{\Sigma}_{M,a})$, the center the drawing's predicted state;
2. record a crossing in windows 37 to 45 where the drawn path has one, and otherwise carry the draw from window 45 by the forecast rows (14.4) with the wander $\sigma_q = 0.005$, to the first passage (14.1) and the total $G_{\text{tot},46}$ at the shutdown;
3. give each draw of drawing M the weight $P_M/10\,000$, and read the family's band and each drawing's own band from the weighted draws, equation (23.1);
4. state the Monte Carlo error of each band by batch means over 20 batches.

These draws share the standard normals of the path and the trend entries across candidates, but each draw's wander after window 45 is drawn only while the draw has not crossed, so once draws cross at different windows under different candidates their wander no longer lines up, and a candidate's difference from the continuation carries about the error of a band. The *paired run* removes that: it draws one array of wander for every draw at every window to the horizon, shared by every candidate, carries every draw through all of it, and repeats the whole route over five seeds. A candidate's addition and the difference between two candidates' additions are then read as their mean over the seeds, with the standard deviation over the seeds as the Monte Carlo error of one run. The floors condition the same draws. *The trend's floor at the refit* fixes γ_1, γ_2 and $\dot{\gamma}_1$ at the drawing's center and draws the rest from the conditional normal, the Schur complement (3.31) with no reading error; the second floor fixes the path to window 45 as well, which leaves only the wander after it.

23.2.3 The covariance at the refit

On the leading drawing $(\mathcal{C}_1, 2)$ under the continuation, the center at the refit is $\hat{g}_{1,45} = -0.24976$ and $\hat{g}_{2,45} = -0.21056$, $\hat{\gamma}_1 = -0.04825$ and $\hat{\gamma}_2 = -0.02732\,\text{yr}^{-1}$, and $\hat{\dot{\gamma}}_1 = -0.03403\,\text{yr}^{-2}$: the conductances $G_{1,45} = 5.4529$ and $G_{2,45} = 5.6709\,\text{kW/K}$, the total $11.1238\,\text{kW/K}$, and the rates at window 45, $\gamma_u(3.75) = \gamma_u + \dot{\gamma}_u(3.75 - \bar{t}_{\text{y}})$, -0.12482 and $-0.02732\,\text{yr}^{-1}$. The plug-in crossing from window 45 is 50.10. Table 23.4 gives the block of $\boldsymbol{\Sigma}_{M,\text{cont}}$ on the five entries the date reads, equation (23.4), with the correlations below it.

The rows fix each sign. The two log-derates at the refit are correlated by -0.6128: the return air reads the hall's total, so a higher coil 1 leaves less to coil 2. Coil 1's log-derate

Table 23.4: The block of the leading drawing's preposterior covariance $\mathbf{\Sigma}_{M,\mathrm{cont}}$ at the refit on the five entries the date reads, and their correlations. Rows and columns in the order $(g_{1,45}, g_{2,45}, \gamma_1, \gamma_2, \dot{\gamma}_1)$; the trend entries in yr^{-1} and the curvature in yr^{-2}.

	$g_{1,45}$	$g_{2,45}$	γ_1	γ_2	$\dot{\gamma}_1$
Covariance [mixed]					
$g_{1,45}$	7.4836×10^{-4}	-2.7678×10^{-4}	1.1838×10^{-4}	-1.0731×10^{-4}	1.8537×10^{-4}
$g_{2,45}$	-2.7678×10^{-4}	2.7261×10^{-4}	-5.3302×10^{-5}	1.0980×10^{-4}	-5.6505×10^{-5}
γ_1	1.1838×10^{-4}	-5.3302×10^{-5}	1.5846×10^{-4}	-2.4747×10^{-5}	-4.6085×10^{-5}
γ_2	-1.0731×10^{-4}	1.0980×10^{-4}	-2.4747×10^{-5}	1.7374×10^{-4}	-2.1077×10^{-5}
$\dot{\gamma}_1$	1.8537×10^{-4}	-5.6505×10^{-5}	-4.6085×10^{-5}	-2.1077×10^{-5}	1.7222×10^{-4}
Standard deviation [mixed]					
	0.02736	0.01651	0.01259	0.01318	0.01312
Correlation [–]					
$g_{1,45}$	1	−0.6128	0.3438	−0.2976	0.5163
$g_{2,45}$	−0.6128	1	−0.2565	0.5045	−0.2608
γ_1	0.3438	−0.2565	1	−0.1491	−0.2790
γ_2	−0.2976	0.5045	−0.1491	1	−0.1218
$\dot{\gamma}_1$	0.5163	−0.2608	−0.2790	−0.1218	1

and its curvature are correlated by +0.5163: with the earlier windows held, a higher coil 1 now needs less acceleration to have reached it. Coil 1's rate and its curvature are correlated by −0.2790: with window 45's level held, a steeper constant rate needs less curvature. Each coil's rate is correlated with its own log-derate positively, +0.3438 and +0.5045, the rate as the telescoped path read from its end. Two correlations reach beyond the five entries. The drift rate, whose spread is 0.1026 K/yr, is correlated with $g_{1,45}$ by +0.8108, with the split $g_{1,45} - g_{2,45}$ by +0.7787 and with the total's level $G_{\mathrm{tot},45}$ by +0.6064: a faster drift explains more of unit 1's warm discharge, leaving coil 1 cleaner, Chapter 7's ridge, and it reaches the total because unit 1's drop is one of the two drops that read it. And the total's rate of change, $\dot{G}_{\mathrm{tot},45} = G_{1,45}\gamma_1(3.75) + G_{2,45}\gamma_2$, is correlated with $\dot{\gamma}_1$ by +0.8190 and with the total's level by +0.4959: the date's two leading quantities are nearly one direction of the state.

23.2.4 What each candidate does to the entries the date reads

Each candidate changes the block of Table 23.4, and with it the spread of the quantities the date reads. Five directions of the block carry them, each with its gradient on $(g_{1,45}, g_{2,45}, \gamma_1, \gamma_2, \dot{\gamma}_1)$: the total's level $G_{\mathrm{tot},45} = G_{1,45} + G_{2,45}$; its rate of change $\dot{G}_{\mathrm{tot},45}$; the split $g_{1,45} - g_{2,45}$; the split of the rates, $\gamma_1(3.75) - \gamma_2$; and the curvature,

$$\mathbf{g}_G = (G_{1,45},\ G_{2,45},\ 0,\ 0,\ 0), \qquad \mathbf{g}_{\dot{G}} = \big(G_{1,45}\gamma_1(3.75),\ G_{2,45}\gamma_2,\ G_{1,45},\ G_{2,45},\ 2.25\,G_{1,45}\big),$$
$$\mathbf{g}_{\Delta g} = (1,\ -1,\ 0,\ 0,\ 0), \qquad \mathbf{g}_{\Delta\gamma} = (0,\ 0,\ 1,\ -1,\ 2.25), \qquad \mathbf{g}_{\dot{\gamma}} = (0,\ 0,\ 0,\ 0,\ 1),$$

with $2.25\,\mathrm{yr} = 45\,\Delta t - \bar{t}_{\mathrm{y}}$, the curvature's lever at the refit. The date itself is read by the delta form of equation (14.13) from the refit, with gradient

$$\mathbf{g}_* = \left.\frac{\partial t^*}{\partial(g_{1,45},\ g_{2,45},\ \gamma_1,\ \gamma_2,\ \dot{\gamma}_1)}\right|_{\hat{\mathbf{z}}_M}, \tag{23.6}$$
$$\mathbf{g}_* = (+71.01,\ +77.21,\ +30.16,\ +32.79,\ +74.27) \text{ windows per unit},$$

on the leading drawing: a higher conductance at the refit or a gentler rate moves the date later.

Table 23.5 gives these spreads on the leading drawing under each candidate. Each row is a candidate, all but the first with the continuation; the columns give the posterior standard

Table 23.5: What each candidate does to the quantities the date reads on the leading drawing $(\mathcal{C}_1, 2)$ at the refit, by the preposterior (23.5): posterior standard deviations, and the delta form's standard deviation of the date from window 45.

Candidate (with the continuation)	$G_{\mathrm{tot},45}$ [kW/K]	$\dot{G}_{\mathrm{tot},45}$ [kW/K/yr]	split [–]	rates' split [yr^{-1}]	$\dot{\gamma}_1$ [yr^{-2}]	$\dot{b}_{d,1}$ [K/yr]	date [window]
no new rows (window 36's record carried)	0.2984	0.2395	0.07784	0.05627	0.02018	0.1808	5.339
the record continued (no sensor)	0.1179	0.1512	0.03968	0.03378	0.01312	0.1026	2.216
heat meter, coil 1	0.0995	0.1447	0.02720	0.02929	0.01232	0.0785	1.908
heat meter, coil 2	0.1148	0.1499	0.03019	0.03093	0.01280	0.0866	2.163
heat meters, both coils	0.0993	0.1447	0.02270	0.02835	0.01226	0.0726	1.906
second discharge sensor, unit 1	0.0916	0.1431	0.02471	0.02862	0.01220	0.0666	1.790
water-side difference, coil 1	0.1105	0.1455	0.03631	0.03147	0.01243	0.0992	2.062
water-side difference, coil 2	0.1176	0.1510	0.03792	0.03306	0.01303	0.1012	2.212
water-side differences, both	0.1104	0.1451	0.03382	0.03037	0.01226	0.0974	2.059
return-air sensor at each unit	0.1104	0.1497	0.03730	0.03295	0.01298	0.0925	2.104
supply-air sensor	0.1010	0.1472	0.03681	0.03255	0.01283	0.0911	1.958
reference test at window 45 (one-off)	0.0766	0.1379	0.01727	0.02604	0.01167	0.0691	1.536

deviations of the five directions and of the drift rate, and the delta form's standard deviation of the date, $\sqrt{\mathbf{g}_*^{\mathsf{T}} \boldsymbol{\Sigma}_5 \mathbf{g}_*}$, with $\boldsymbol{\Sigma}_5$ the block of the five entries.

The continuation alone takes the date's standard deviation on this drawing from 5.339 to 2.216 windows, and the total's level from 0.2984 to 0.1179 kW/K; the best permanent sensor, the second discharge sensor on unit 1, takes the date to 1.790 and the level to 0.0916. Coil 2's instruments leave coil 1's curvature and the drift almost where the continuation left them: the heat meter on coil 2 takes the drift's spread to 0.0866 against the coil-1 meter's 0.0785, and the water-side difference on coil 2 to 0.1012.

23.2.5 Predicted records refitted

Equation (23.1) holds each drawing's probability and center. Beside it stands the preposterior simulated over the records the continuation may log, called the predicted records, not an exact answer:

1. draw a drawing by its probability P_M and its state at window 36 from its Laplace normal;
2. carry the path to window 45 by the drawing's trend rows with the wander, and read the existing five sensors and every candidate's channels in windows 37 to 45 at their grades, with one stream of errors per channel, the same for every candidate, and the flow factors ι_u drawn from their prior;
3. refit every one of the six drawings on the record to window 45 with the continuation's rows and each candidate's, by Chapter 10's Gauss–Newton, and reweight them by their Laplace evidence (7.7);
4. forecast from window 45 with 2 000 draws per drawing, the same normals for every candidate, and read the family's band, the band of the drawing that generated the record, the refit's $\Pr(t^* \le 46 \mid \mathbf{y}_{0:45})$ and the Bayes factors of the curvature and of the two steps.

Thirty records are drawn. Each candidate's addition is paired with the continuation's on the same record, and its Monte Carlo standard error is that of the paired differences over the 30.

23.3 The answers (step 4)

Every answer is read from the preposterior covariances $\mathbf{\Sigma}_{M,a}$ of equation (23.5) through the band (23.1), with the drawings' probabilities and centers held at today's, by the draws of Section 23.2.2, and that value is the one reported; the predicted records of Section 23.2.5 stand beside it where they answer differently. Bands are 5–95 %, and dates are continuous windows.

23.3.1 The measurement that tightens the band most

Asset manager's question. Which added measurement most tightens the remaining-life band? It is asked of every instrument that could be installed now and logged to the refit at window 45.

General form. A design question asks for the design form (2.16), the expected value over a candidate's outcomes of a quantity U_{p} of the posterior each outcome would produce,

$$U(a) = \int U_{\mathrm{p}}\big(p(\mathbf{z} \mid \mathbf{y}, \tilde{y}_a)\big)\, p(\tilde{y}_a \mid \mathbf{y})\, \mathrm{d}\tilde{y}_a,$$

with the preposterior covariance (2.17) for linear rows and, over a family, the mixture (7.10).

Specialized form. The question is equation (2.16) with a the continuation plus a candidate's rows (23.3), $\tilde{y}_a$ their readings in windows 37 to 45, and U_{p} the width of the band of t^* under the posterior those readings would give, the mixture (7.10) over the six drawings. For linear rows that posterior has each drawing's covariance $\mathbf{\Sigma}_{M,a}$ of equation (23.5) whatever the readings, and with the centers and probabilities held at $\hat{\mathbf{z}}_M$ and P_M the integral over the outcomes is the band at those centers, $\mathcal{B}_a$ of equation (23.1). The candidates are ranked by $\Delta\mathcal{B}_a$, the windows they take off the continuation's band, with the own bands' $\Delta\bar{\mathcal{B}}_a$ beside it, and by their reach, $\Delta\mathcal{B}_a/(\mathcal{B}_{\mathrm{cont}} - \mathcal{B}_{\mathrm{floor}})$, the share of the distance to the trend's floor $\mathcal{B}_{\mathrm{floor}}$ they cover. The choice that enters here is the first of step 2, the probabilities held, and Section 23.3.5 reads what the predicted records say of it.

Building blocks. Each drawing's preposterior covariance (23.5) for each candidate and its MAP center; the probabilities P_M of Section 23.2.1; the draws of Section 23.2.2, with the floors; the weighted quantiles.

The ranking. Table 23.6 gives each candidate's band. Each row is a candidate or a floor; the columns give the family's band and its width, the share of today's band of 38.10 windows it takes off, the windows it adds to the continuation and its reach, and the same for the own bands weighted. The continuation alone takes 27.0 % off the band, from 38.10 to 27.83 windows, and the own bands from 27.49 to 13.72. No instrument installed now adds more than 0.70 windows to it. Of the permanent sensors, the second discharge-air sensor on unit 1 tightens the band most: +0.70 windows on the family's band and +1.30 on the own bands, 23.4 and 27.5 % of the way to the trend's floor. Close behind come both heat meters, +0.61 and +1.00, and the supply-air sensor, +0.60 and +0.83. The one-off reference test at the refit adds more than any permanent sensor, +1.30 and +2.33.

The leading three lie close on the family's band, and the paired run of Section 23.2.2 reads their differences. Over its five seeds the continuation's family band is 28.00 windows (standard deviation over the seeds 0.13) and the own bands 13.78 (0.10); the second discharge sensor adds +0.64 (0.08) and +1.35 (0.07), both heat meters +0.50 (0.08) and +1.01 (0.03), and the supply air +0.47 (0.05) and +0.92 (0.04). Paired draw by draw, the second sensor's addition to the family's band exceeds the heat meters' by +0.14 windows (0.08) and the supply air's

by +0.17 (0.06), while the heat meters and the supply air differ by +0.02 (0.07). On the own bands the leads are wider: +0.34 (0.07) over the heat meters and +0.43 (0.06) over the supply air. Ranked on the family's band, as reported, the second discharge sensor is first and the heat meters and the supply air are tied behind it; the coil-1 heat meter follows, the second sensor leading it by +0.23 (0.04).

Table 23.6: The crossing band at the refit under each candidate, the drawings' probabilities and centers held, by equation (23.1); 10 000 draws per drawing. "Off today" is against today's 38.10 windows; "adds" is taken off the continuation's band; "reach" is that as a share of the continuation's distance to the trend's floor.

Candidate	Band [window]	Width [window]	Off today [%]	Adds [window]	Reach [%]	Own [window]	Own adds [window]	Own reach [%]
no new rows (the check)	43.76–81.34	37.58	1.4	–	–	27.22	–	–
the record continued to window 45	45.78–73.61	27.83	27.0	–	–	13.72	–	–
heat meter, coil 1	45.87–73.16	27.28	28.4	+0.55	18.3	12.76	+0.97	20.5
heat meter, coil 2	45.77–73.55	27.77	27.1	+0.06	1.9	13.58	+0.14	3.0
heat meters, both coils	45.89–73.11	27.22	28.6	+0.61	20.5	12.73	+1.00	21.1
second discharge sensor, unit 1	45.94–73.07	27.13	28.8	+0.70	23.4	12.42	+1.30	27.5
water-side difference, coil 1	45.85–73.48	27.63	27.5	+0.20	6.8	13.27	+0.45	9.5
water-side difference, coil 2	45.78–73.57	27.79	27.1	+0.04	1.2	13.63	+0.09	1.9
water-side differences, both	45.84–73.48	27.63	27.5	+0.20	6.5	13.19	+0.53	11.3
return-air sensor at each unit	45.85–73.50	27.65	27.4	+0.18	6.1	13.45	+0.27	5.7
supply-air sensor	45.93–73.16	27.23	28.5	+0.60	20.0	12.89	+0.83	17.6
airflow station on each unit	45.78–73.61	27.83	27.0	0.00	0	13.72	0.00	0
reference test at window 45 (one-off)	46.07–72.60	26.53	30.4	+1.30	43.3	11.39	+2.33	49.5
floor: the rates known exactly	45.99–70.82	24.84	34.8	+2.99	100	9.00	+4.72	100
floor: the rates and the path to window 45 known	46.41–70.21	23.80	37.5	+4.03	–	7.04	+6.68	–

Ranked by the family's band, the permanent sensors run: the second discharge sensor +0.70, both heat meters +0.61, the supply air +0.60, the coil-1 heat meter +0.55, the water-side difference on coil 1 +0.20, the pair of differences +0.20, the return-air pair +0.18, the coil-2 heat meter +0.06 and the difference on coil 2 +0.04. The own bands give the same order except that the coil-1 heat meter, +0.97, passes the supply air, +0.83, and the pair of differences, +0.53, passes the coil-1 difference, +0.45. The check row, no new rows from window 45, gives 37.58 windows against today's 38.10 from 40 000 draws per drawing and another seed, 0.52 windows apart; Chapter 14's same check on the eighteen gave 38.22 against 37.93, apart the other way. Both gaps are the Monte Carlo error of a band's slow upper end at these draws, not a property of the rows; for scale, the check row's difference from the continuation carries a batch-means standard error of 0.28 windows.

The chapter's surprise is the dashboard's pick (Exercise 23.1). The planner's 42.3 % from two more return-air sensors comes to 0.18 windows, 0.7 % of the continuation's band, and 0.27 windows, 2.0 %, of the own bands. Read on the continuation's bands, the planner's claim would have taken 11.76 windows off the family's 27.83 and 5.80 off the own 13.72. The noisiest trace is not the one that limits the date. The step's surprise is how much of the band is not

any drawing's: at the refit the family's band is 27.83 windows and the own bands are 13.72, so the family's indecision about coil 1's curvature is 14.11 windows, 50.7 % of the band. No instrument's narrowing within the drawings shows through it undiminished, and on the own bands every sensor is worth close to twice what it is worth on the family's: +1.30 against +0.70 for the second discharge sensor.

The windows each candidate adds are compared as bars in Figure 23.2. Each pair of bars is a candidate, ranked by the family's addition, the blue bar the family's and the red bar the own bands'; the dashed lines are the trend's floor, +2.99 and +4.72; the pair at the top is the planner's claim, 42.3 % of each of the continuation's bands. Every bar is short beside the planner's; the second discharge sensor's bars are the longest of the permanent sensors', the heat meters' and the supply air's stand close together behind it, and the return-air pair is among the shortest.

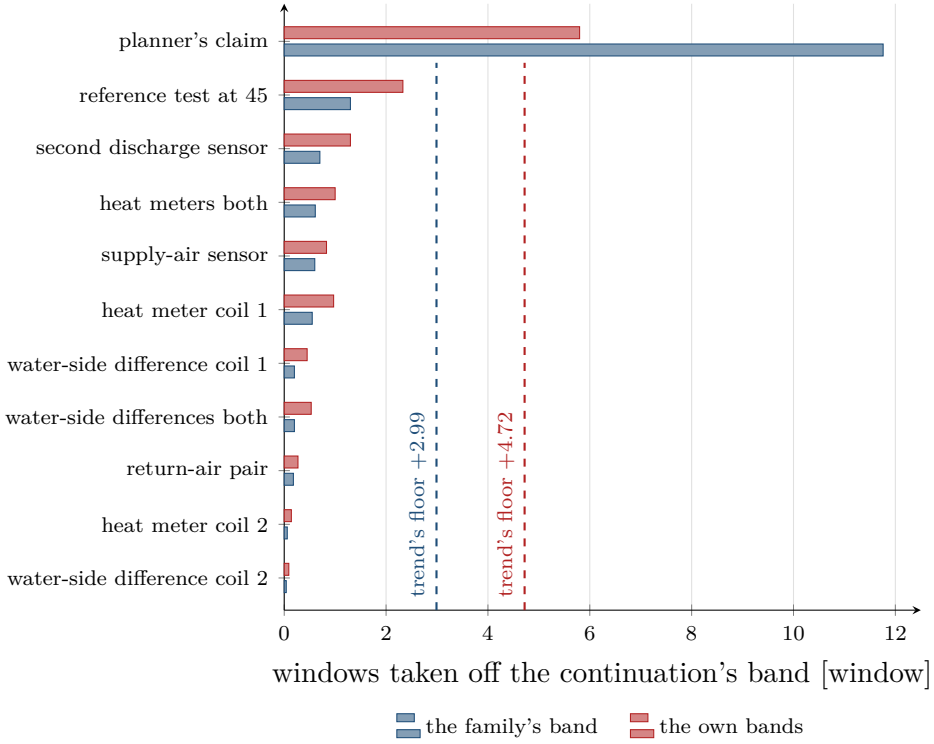


Figure 23.2: The windows each candidate takes off the continuation's crossing band at the refit, the drawings' probabilities and centers held (Table 23.6): blue, the family's band; red, the own bands weighted. Dashed: the trend's floor, the rates known exactly. Top pair: the planner's claim, 42.3 % of each of the continuation's bands, the arithmetic of three return-air readings averaged; the return-air pair is the third pair from the bottom.

23.3.2 The call at the shutdown

Asset manager's question. Which one makes the call at the shutdown surest? The call is whether the coils hold past the October shutdown or must be cleaned in it.

General form. The design form (2.16) with U_{p} the width of the band of a quantity read beyond the record, the third line of the trend form (2.21),

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

read by inversion at 0.05 and 0.95.

Specialized form. The question is equation (2.16) with the candidates of Section 23.3.1 and U_{p} the width of the 5–95 % band of $f_t = G_{\mathrm{tot},46}$, the third line of equation (2.21) at the shutdown, read on the same draws from each $\mathbf{\Sigma}_{M,a}$ of equation (23.5), family and own. Chapter 14 found the shutdown's tail $\Pr(t^* \le 46)$ equal to $\Pr(G_{\mathrm{tot},46} \le G^*)$ to four decimals (Section 14.3.3), so the margin's band is the sharpness of the tail. One choice enters here: the call is read on the margin's band and not on the held tail, because the tail at held centers is the tail of one record's center, not an expectation over records.

Building blocks. The draws of Section 23.2.2, with $G_{\mathrm{tot},46}$ recorded per draw; the probabilities P_M; the weighted quantiles.

The margin's band. Under the continuation the total at the shutdown has a median of 11.1021 kW/K, a margin of 0.3411 over G^*, with a band 0.8271 kW/K wide on the family and 0.4123 on the own bands. The own band under each candidate is 0.3409 with the second discharge sensor, 0.3623 with both heat meters, 0.3628 with the supply air, 0.3641 with the coil-1 heat meter, 0.3904 with the return-air pair, 0.3914 and 0.3921 with the pair of water-side differences and the coil-1 difference, and 0.4057 and 0.4112 with coil 2's meter and difference. The reference test at the refit gives 0.2929, the rates known 0.3389, and the rates and the path known 0.1302. Ranked by what they take off the own band: the reference test, +0.1194 kW/K; the second discharge sensor +0.0713, both heat meters +0.0500, the supply air +0.0494, the coil-1 meter +0.0482, the return-air pair +0.0218, the pair of differences +0.0209, the coil-1 difference +0.0201, the coil-2 meter +0.0066 and the coil-2 difference +0.0011.

The second discharge sensor alone reaches the margin's floor with the rates known, 0.3409 against 0.3389: one window after the refit the margin is set mostly by the level at window 45, and the level is what unit 1's second sensor restores. The held tail at the centers, $\Pr(t^* \le 46)$ of 0.0638 under the continuation and 0.0538 with the second discharge sensor, is not an expectation over the records, and it is not reported as one; the chance of a wrong hold-or-cross call over the predicted records is read in Section 23.4.2.

23.3.3 The value per real

Asset manager's question. Which is worth its price? A candidate's width is set beside its stated price.

General form. The design form (2.16) with the candidate's price c_a, the fixed cost of the action of installing it, beside its utility: the ratio $U(a)/c_a$ against the same ratio of the best candidate $a^\dagger$.

Specialized form. The question is equation (2.16) with the widths of Section 23.3.1: the windows taken off per R$ 1 000, $1000\,\Delta\mathcal{B}_a/c_a$ and $1000\,\Delta\bar{\mathcal{B}}_a/c_a$, and the margin's own band per R$ 1 000; and the price at which candidate a would match the best permanent sensor $a^\dagger$ per real on the own bands,

$$c_a^{\times} = \Delta\bar{\mathcal{B}}_a \, \frac{c_{a^\dagger}}{\Delta\bar{\mathcal{B}}_{a^\dagger}}. \tag{23.7}$$

The prices are the stated inputs of Table 23.1, the fifth choice of step 2, and the matching price shows how far each would have to move.

Building blocks. The additions $\Delta\mathcal{B}_a$ and $\Delta\bar{\mathcal{B}}_a$ of Table 23.6 and the margin's own bands of Section 23.3.2; the prices c_a.

The value per real. Table 23.7 gives the value of each candidate per R$ 1 000. Each row is a candidate, ranked by the own bands' value; the columns give its price and the windows it takes off per R$ 1 000, family and own, and the margin's own band per R$ 1 000; its values carry the Monte Carlo error of the additions, about 0.08 windows on the second sensor's own addition, so they are quoted here to two figures. The second discharge sensor is the best permanent instrument per real, about 1.4 windows per R$ 1 000 on the own bands (1.4402 on the reported draws, 1.50 with a standard deviation of 0.08 over the paired run's seeds), against the supply air's 0.69, the next, and the heat meters' 0.042, a factor of about 35 (34.7 reported, 35.6 with 1.7 over the seeds; Exercise 23.4). By equation (23.7), to match it per real both heat meters would have to cost about R$ 690 instead of R$ 24 000 (R$ 691 reported, R$ 676 with R$ 33 over the seeds), the coil-1 meter about R$ 670, the supply air R$ 580 and the pair of water-side differences R$ 370. The one-off reference test is about as good per real, 1.6, and would match at about R$ 1 600; it does not stay.

Table 23.7: The value of each candidate per R$ 1 000 of its stated price: the windows taken off the continuation's band, family and own, and the width taken off the margin's own band at the shutdown.

Candidate	Price [R$]	Family [window/R$ 1 000]	Own [window/R$ 1 000]	Margin, own [kW/K/R$ 1 000]
reference test at window 45 (one-off)	1 500	0.8647	1.5558	0.079572
second discharge sensor, unit 1	900	0.7784	1.4402	0.079254
supply-air sensor	1 200	0.4997	0.6917	0.041207
water-side difference, coil 1	2 400	0.0850	0.1860	0.008396
return-air sensor at each unit	1 800	0.1022	0.1504	0.012122
water-side differences, both	4 800	0.0407	0.1114	0.004356
heat meter, coil 1	12 000	0.0457	0.0805	0.004015
heat meters, both coils	24 000	0.0255	0.0415	0.002083
water-side difference, coil 2	2 400	0.0155	0.0370	0.000455
heat meter, coil 2	12 000	0.0047	0.0120	0.000546
airflow station on each unit	28 000	0	0	0

The widths are set against the prices in Figure 23.3. The horizontal axis is the price on a logarithmic scale and the vertical axis the windows taken off the own bands; each point is a candidate, and the dashed line is equal value per real through the second discharge sensor, 1.4402 windows per R$ 1 000. Every permanent candidate lies below the line, the heat meters by more than an order of magnitude in price; the reference test lies just above it.

23.3.4 The coils' drops and the hall's total

Asset manager's question. Why does a sensor on unit 1 narrow the date and two more return-air sensors do not? The answer is asked of the state, before any reading.

General form. An estimation question about a derived quantity under a posterior, equation (2.9), read through the preposterior (2.17): the variance $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ of f by the delta form, and the variance left when a set of directions $\mathbf{Q}\mathbf{z}$ of the state is known exactly, the Gaussian conditional (3.31) with no reading error,

$$\operatorname{Var}(f \mid \mathbf{Q}\mathbf{z}) = \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} - \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{Q}^{\mathsf{T}}\left(\mathbf{Q}\boldsymbol{\Sigma}\mathbf{Q}^{\mathsf{T}}\right)^{-1}\mathbf{Q}\boldsymbol{\Sigma}\mathbf{g}.$$

Specialized form. The question is equation (2.9) with $f = t^*$ from the refit on the leading drawing, $\boldsymbol{\Sigma} = \boldsymbol{\Sigma}_5$ the block of Table 23.4, $\mathbf{g} = \mathbf{g}_*$ of equation (23.6), and $\mathbf{Q}$ the rows of the directions of Section 23.2.4: the total's level, its rate, both together, the split with the rates' split, the curvature, and the three trend entries. The share each direction removes is

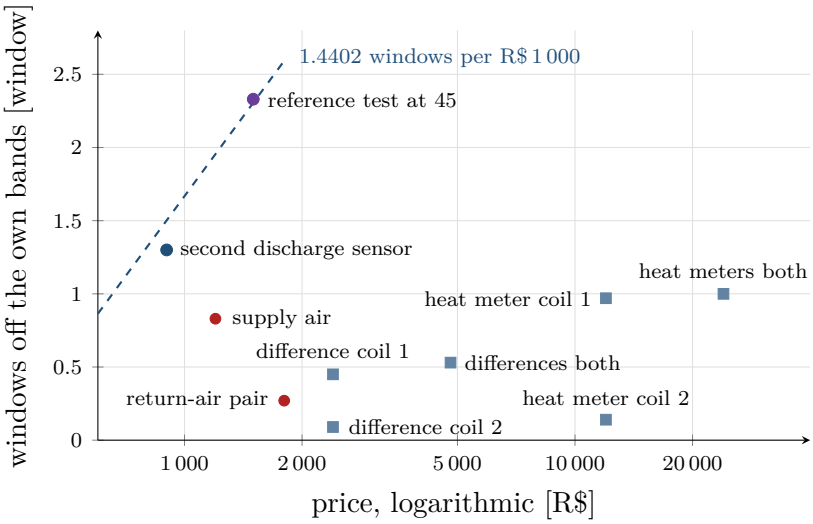


Figure 23.3: The windows each candidate takes off the own bands at the refit against its stated price (Table 23.7). Circles: the second discharge sensor on unit 1 (blue), the reference test at window 45 (purple), the total sensors (red); squares: the water-side instruments. Dashed: equal value per real through the second discharge sensor, 1.4402 windows per R$ 1 000, drawn to R$ 1 800. The airflow stations, R$ 28 000 for nothing, are not drawn.

$1 - \mathrm{Var}(t^* \mid \mathbf{Qz})/\mathbf{g}_*^{\mathsf{T}}\boldsymbol{\Sigma}_5\mathbf{g}_*$. Beside it, one window's conductance is read by two routes from the forward equations (5.3): through the hall's total sensors, $G_{\mathrm{tot}} = (Q + Q_o)/(T_r - T_s)$, and through a unit's discharge drop over the approach,

$$\frac{T_r - T_{d,u}}{T_r - T_s} = \frac{G_u}{C_u}, \tag{23.8}$$

which holds no load.

Building blocks. The block $\boldsymbol{\Sigma}_5$ of Table 23.4 and the gradient (23.6); the directions' gradients; the spreads of Table 23.5; the grades of the load meter, 3.0 kW, and of the temperature sensors, 0.3 K.

The date's variance by direction. Table 23.8 gives the budget. Each row is a direction known exactly; the columns give the date's variance left and the share removed, against the continuation's 4.909 windows^2, a standard deviation of 2.216 windows. The total's level known removes 87.3 %, its rate 59.7 %, and both together all but 0.001 windows^2, 100.0 %: within a drawing the date is a function of the state at the refit only through the total's level and rate. The budget is the delta form's on the five entries of the state at the refit; the wander after window 45 lies outside it, and that wander is about half the own band: the second floor, with every entry to window 45 known, still leaves 7.04 of the continuation's 13.72 windows (Table 23.6). The split removes 37.9 %, coil 1's curvature 45.0 % and the three trend entries 68.5 %, each only through its correlation with the total's level and rate. The shares overlap, because the directions are correlated by the signs of Table 23.4, and they do not add.

The routes to one window's conductance. If the date reads the total, why does more of a sensor of the total not read it better? At window 45 on the leading drawing's center, $Q + Q_o = 144.44$ kW and $T_r - T_s = 12.985$ K. The route through the return air and the load meter reads $G_{\mathrm{tot},45}$ to a relative standard deviation of 0.0311, the meter contributing 0.0208 and the return air 0.0231 (Exercise 23.2); three return-air readings bring it only to 0.0247,

Table 23.8: The date's delta variance from the refit on the leading drawing under the continuation, $\mathbf{g}_*^{\mathsf{T}}\boldsymbol{\Sigma}_5\mathbf{g}_* = 4.909$ windows^2, and what is left when a direction of the state is known exactly.

Direction known exactly	Variance left [window^2]	Share removed [%]
the total's level $G_{\mathrm{tot},45}$	0.626	87.3
the total's rate $\dot{G}_{\mathrm{tot},45}$	1.976	59.7
the total's level and rate	0.001	100.0
the split, $g_{1,45} - g_{2,45}$ and $\gamma_1(3.75) - \gamma_2$	3.050	37.9
coil 1's curvature $\dot{\gamma}_1$	2.699	45.0
the trend entries γ_1, γ_2, $\dot{\gamma}_1$	1.545	68.5

because the load meter's 2.1 % stays. The second route is the units' drops over the approach, equation (23.8), 0.5453 for coil 1 and 0.5671 for coil 2, which carry no load. At 0.3 K one drop reads its G_u/C_u to a relative standard deviation of 0.0465 on unit 1 and 0.0444 on unit 2, and the two drops together, $(2T_r - T_{d,1} - T_{d,2})/(T_r - T_s) = G_1/C_1 + G_2/C_2$, with the shared return air counted once, read the total to 0.0347: close to the meter route's 0.0311, and with no meter in it. The two routes are of like precision, and the date uses both; what sets this one apart is that unit 1's drop runs through the drifting sensor: the drift is correlated with the total's level by +0.6064. A third route opens with the supply-air sensor. By the air balance (5.4) the cold aisle's drop over the approach, $(T_r - T_c)/(T_r - T_s) = G_{\mathrm{tot}}/(C_1 + C_2)$, reads the total with no meter and no drift in it: to 0.0454 at 0.3 K, coarser than the meter route, and to 0.0231 at 0.1 K, finer than either. With the load meter's rows of windows 37 to 45 made exact, the continuation's spread of $G_{\mathrm{tot},45}$ falls from 0.1179 to 0.0913 kW/K (the date's to 1.792 windows), and with the return-air pair added it falls to 0.0705 (1.490). The return-air sensors would be worth buying if the load were known, and it is not.

The mechanism. The second discharge sensor gives unit 1's drop back without the drift. Its line is the existing sensor's with a zero in the drift's column (Table 23.3), so the two readings of $T_{d,1}$ differ by the drift alone and separate $\dot{b}_{d,1}$ from $g_{1,45}$: the drift's spread falls from 0.1026 to 0.0666 K/yr and the total's level from 0.1179 to 0.0916 kW/K, the largest cut of any permanent sensor (Table 23.5). The heat meters read the split, ± 12.0319 on the two log-derates, and through it free the drift as well, to 0.0726. The supply-air sensor reads $(T_{d,1} + T_{d,2})/2$ with no drift in it, so it too separates the drift, to 0.0911, and it adds the cold aisle's load-free route to the total; its line has the return air's slopes on the coils, but it does not lean on the meter. The crossing date is a function of the coils' total, read by two routes of like precision, and the instruments that narrow it are those that free the load-free route of the drift on unit 1's sensor or add a route that has neither the meter nor the drift in it.

The budget and the two spreads are drawn in Figure 23.4. The left panel gives the share of the date's variance each direction removes when known. The right panel places each candidate by its spread of $G_{\mathrm{tot},45}$ (horizontal) against its spread of the drift rate (vertical), on the leading drawing with the continuation; the marker's size grows with the date's delta standard deviation. The instruments that separate the drift, the second discharge sensor, the heat meters and the supply air, fall at the lower left; the return-air pair and coil 2's instruments stay near the continuation, at the upper right.

23.3.5 The reach of any instrument

Asset manager's question. How much of the band can any sensor reach, and what is left? The answer sets the instruments against the band that no instrument installed now can touch.

General form. The preposterior (2.17) at its limit, the rows of an entry made exact, with the band read over the family's mixture (7.10); and the evidence between drawings, the Bayes

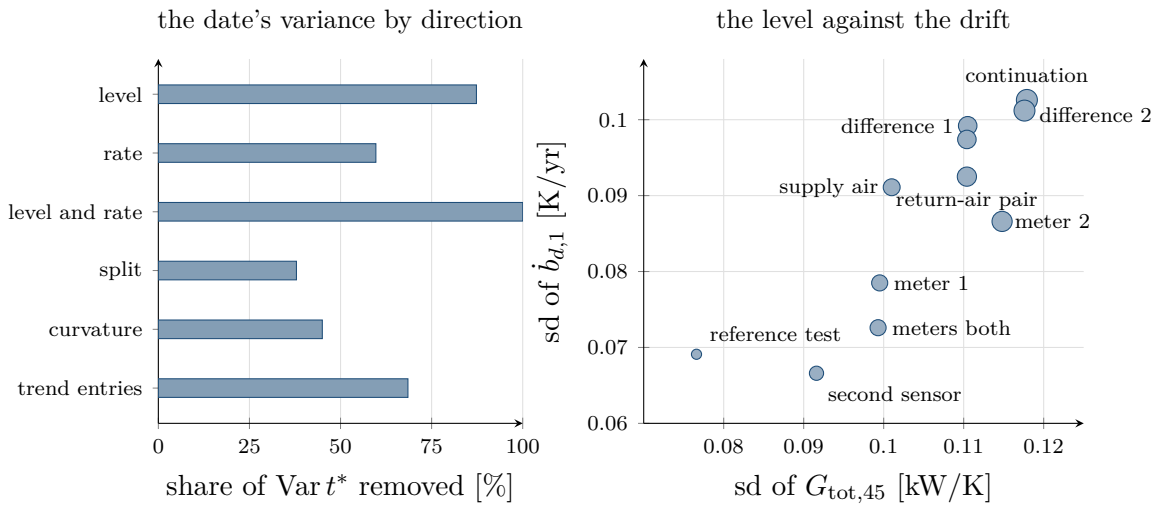


Figure 23.4: Why unit 1's second sensor narrows the date. Left: the share of the date's delta variance from the refit removed when a direction of the state is known exactly (Table 23.8), the leading drawing under the continuation. Right: each candidate's spread of the total's level $G_{\text{tot},45}$ against its spread of the drift rate $\dot{b}_{d,1}$, with the continuation (Table 23.5); the marker grows with the date's delta standard deviation, 1.536 to 2.216 windows. The pair of water-side differences, at $(0.1104, 0.0974)$, sits under the label of the coil-1 difference.

factor of equation (2.14), predicted over the records a measurement may log.

Specialized form. The question is equation (2.17) with the trend entries, and then the path to window 45 as well, known exactly on each drawing: the two floors of Section 23.2.2, read on the family's band and the own bands of equation (23.1). The family's band against the own bands splits the floor's band into the drawings' part and the indecision. The second part moves only with equation (2.14) for $\mathcal{C}_1$ against $\mathcal{C}_0$, $\ln B(\mathcal{C}_1 : \mathcal{C}_0)$, computed on each predicted record of Section 23.2.5: the curvature is decided on a record when $|\ln B| > 1$, and a candidate's effect on it is the mean change of $|\ln B|$ against the continuation on the same records.

Building blocks. The floors of Table 23.6; the bands $\mathcal{B}_a$ and $\bar{\mathcal{B}}_a$; the predicted records' refits and their Laplace evidences, equation (7.7).

The floors. The trend's floor at the refit is 24.84 windows on the family and 9.00 on the own bands. Knowing every entry to window 45 as well leaves 23.80 and 7.04, only the wander after the refit. The continuation already takes the family from 38.10 to 27.83 windows, so any instrument works within the 2.99 windows left above the trend's floor, and the best permanent sensor covers 23.4 % of them. What is left is 23.80 windows even with the path known exactly, because the band carries the drawings: only evidence on coil 1's curvature narrows it.

The evidence on the predicted records. With the probabilities moving, the refit at window 45 decides coil 1's curvature in 0.57 of the predicted records, which is Chapter 11's 0.545 (Section 11.6.3) by another route. The second discharge sensor makes that 0.60, both heat meters 0.67, the reference test 0.70, the return-air pair 0.60 and the supply air 0.53, each a share of 30 records. The mean change of $|\ln B(\mathcal{C}_1 : \mathcal{C}_0)|$ against the continuation is +0.531 with the reference test, +0.435 with the second discharge sensor and +0.381 with the heat meters, against +0.077 with the supply air and +0.039 with the return-air pair. Coil 2's step against unit 2's sensor step is decided in 0.700 of the records with the continuation; the heat meters

raise that to 0.967, the reference test to 0.933 and the second discharge sensor to 0.867, while the return-air pair leaves it at 0.700 and the supply air takes it to 0.667. What separates coil 1 moves the drawings; what reads only the hall's total, by either route, does not. In a line: buy the sensor that separates unit 1 from its drift, and let the record decide the curvature.

23.4 The checks (step 5)

Each answer of Section 23.3 rests on the probabilities held, on six drawings of eighteen, on the stated grades and on the pinned airflow, and each can be set against the computation it stands for.

23.4.1 The eighteen drawings and the rows

On all eighteen drawings, with the probabilities held, the route reproduces Chapter 14's step 7 to the printed digit: no new rows 38.22 windows, the continuation 27.94, with both heat meters 27.40, and with a second discharge sensor on unit 1 27.09 (Table 14.13). This is the one place the chapter repeats that chapter's rows, and it repeats them as a check. The continuation's own bands over the eighteen are 13.88 windows against the six's 13.72, and the candidates' additions to them run: the reference test +2.23, the second discharge sensor +1.30, both heat meters +1.02, the coil-1 meter +1.01, the supply air +0.90, the pair of water-side differences +0.43, the return-air pair +0.33, the coil-1 difference +0.30, and coil 2's meter and difference +0.12 and +0.07. The first six places are the six drawings' order; the return-air pair and the coil-1 difference trade the seventh and eighth.

No new check of the limit of exact rows or of the curvature term is needed: the candidate rows (23.3) are reading rows on Chapter 10's physics, whose checks hold (Section 10.4.2), and the predicted records' refits use the base fit's Gauss–Newton.

23.4.2 Predicted records with the probabilities moving

The thirty predicted records of Section 23.2.5 were drawn from the six drawings as $(\mathcal{C}_0, \text{none})$ 4, $(\mathcal{C}_0, 2)$ 10, $(\mathcal{C}_0, T_{d,2})$ 1, $(\mathcal{C}_1, \text{none})$ 1, $(\mathcal{C}_1, 2)$ 10 and $(\mathcal{C}_1, T_{d,2})$ 4, and each refit took a median of 7 Gauss–Newton iterations, at most 26. The law of total probability holds on them: the refit's expected $\Pr(t^* \le 46 \mid \mathbf{y}_{0:45})$ is 0.1503 (standard error 0.0522) against today's 0.1607, and the expected $\Pr(\mathcal{C}_1 \mid \mathbf{y}_{0:45})$ is 0.6277 (0.0552) against today's 0.6376. Bands are read at most to the horizon, window 276, and one record's continuation band reaches it, so the mean width is driven by that one record; its median over the records is 18.53 windows.

Table 23.9 gives the records' values. Each row is a candidate, and the columns give the family's band at the refit, its mean with the standard error over the records; the windows the candidate adds to the continuation, paired record by record; the addition to the band of the drawing that generated the record; the share of records whose refit decides coil 1's curvature and the mean change of $|\ln B(\mathcal{C}_1 : \mathcal{C}_0)|$; and the expected chance of a wrong hold-or-cross call at the shutdown, $\mathbb{E}[\min(p_{46},\ 1 - p_{46})]$ with $p_{46} = \Pr(t^* \le 46 \mid \mathbf{y}_{0:45})$ the refit's tail, and what the candidate takes off it.

At 30 records the moving probabilities do not rank the permanent sensors. Every paired addition to the family's band is positive, and each is two to six times the held route's: +1.55 for the second discharge sensor (standard error 1.26), +1.79 for the heat meters (1.11), +1.88 for the supply air (1.09) and +1.01 for the return-air pair (0.72). The differences between them are within the larger of the two standard errors in every pair; only the gaps of the heat meters and the supply air over the return-air pair exceed the smaller one. Separating them would take many more records, which is why the reported route holds the probabilities (the first choice of step 2). One answer the records do resolve. The expected chance of a wrong hold-or-cross call

Table 23.9: The predicted records with the probabilities moving: 30 records, every drawing refitted at window 45 with each candidate's rows and reweighted; 2 000 draws per drawing. Standard errors in parentheses; the family band's is 6.90 to 7.08 windows, 3.72 with the reference test; additions paired against the continuation on the same records.

Candidate	Family band [window]	Adds [window]	Own adds [window]	Decided [–]	$\Delta\lvert\ln B\rvert$ [–]	Wrong call [–]	Off it [–]
continuation	28.54	–	–	0.57	–	0.0716	–
second sensor, unit 1	26.99	+1.55 (1.26)	−0.71 (1.39)	0.60	+0.435	0.0417	+0.0299 (0.0145)
heat meters, both	26.75	+1.79 (1.11)	+0.78 (0.68)	0.67	+0.381	0.0695	+0.0021 (0.0087)
supply-air sensor	26.67	+1.88 (1.09)	+1.39 (0.69)	0.53	+0.077	0.0726	−0.0010 (0.0124)
return-air pair	27.53	+1.01 (0.72)	−0.50 (0.57)	0.60	+0.039	0.0662	+0.0054 (0.0045)
reference test at 45	21.89	+6.65 (4.10)	+2.32 (1.28)	0.70	+0.531	0.0554	+0.0162 (0.0202)

is 0.0716 under the continuation, and the second discharge sensor takes 0.0299 off it, standard error 0.0145, the only candidate more than twice its standard error; the heat meters take 0.0021 (0.0087), the return-air pair 0.0054 (0.0045), the supply air −0.0010 (0.0124) and the reference test 0.0162 (0.0202). The records' own-drawing bands, 21.09 windows under the continuation, are wider than the held route's 13.72: a drawn record with a slower path has a far wider band, because the date is a ratio, while the held route reads each drawing at its MAP center. Within each route the candidates are compared like with like.

The paired additions and the decided shares are drawn in Figure 23.5. The left panel gives each candidate's paired addition to the family's band with a bar of one standard error either side; the right panel gives the share of records whose refit decides coil 1's curvature, with the continuation's 0.57 as the first bar. The error bars of the four permanent sensors overlap the others' points, and the decided shares move by a few records out of thirty, most for the instruments that separate coil 1.

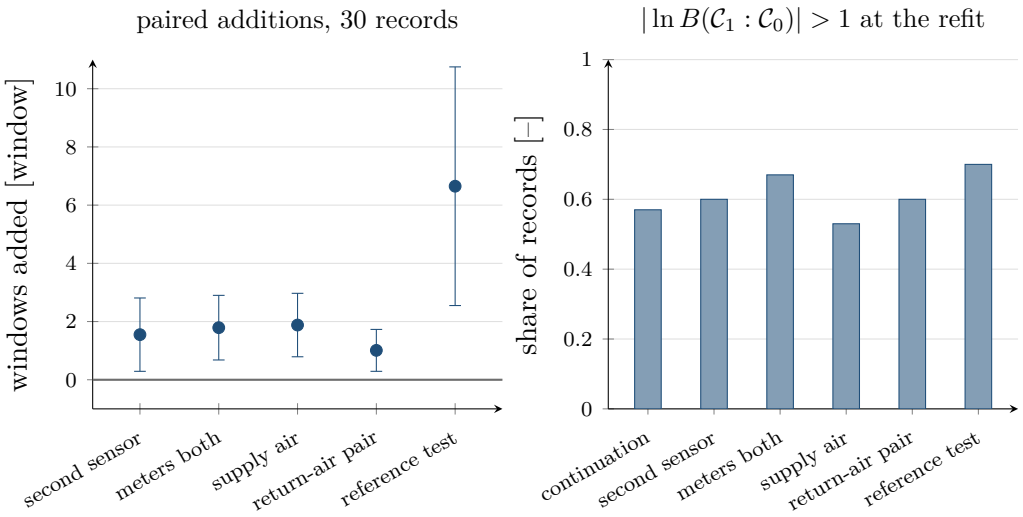


Figure 23.5: The predicted records with the probabilities moving (Table 23.9). Left: each candidate's paired addition to the continuation's family band at the refit, with one standard error either side over the 30 records. Right: the share of records whose refit at window 45 decides coil 1's curvature, $|\ln B(\mathcal{C}_1 : \mathcal{C}_0)| > 1$, under the continuation and with each candidate.

23.4.3 The grades and the pairs

The grades are stated inputs, and Table 23.10 moves them, on the reported route. Each row is a candidate at another grade, or a pair of candidates; the columns give the family's band and its addition to the continuation's 27.83 windows, the own bands and their addition to the continuation's 13.72, and the margin's own band at the shutdown.

Table 23.10: The candidates at other grades, and in pairs, on the reported route (the drawings' probabilities and centers held); additions to the continuation's 27.83 (family) and 13.72 (own) windows.

Variant	Family [window]	Adds [window]	Own [window]	Own adds [window]	$G_{\text{tot},46}$ own [kW/K]
heat meters, both coils, at 1.5 kW	26.94	+0.89	12.12	+1.60	0.3341
heat meters, both coils, at 6.0 kW	27.56	+0.27	13.31	+0.41	0.3889
second discharge sensor at 0.1 K	26.79	+1.04	11.78	+1.94	0.3108
second discharge sensor at 0.5 K	27.12	+0.71	12.84	+0.88	0.3648
water-side differences at 0.05 K	27.39	+0.43	12.88	+0.84	0.3811
water-side differences at 0.2 K	27.63	+0.20	13.49	+0.23	0.4022
water-side differences, flow known to 3 %	27.45	+0.38	13.04	+0.69	0.3784
water-side differences, flow known to 30 %	27.53	+0.30	13.32	+0.40	0.3930
return-air pair at 0.1 K	27.29	+0.54	13.01	+0.71	0.3736
supply-air sensor at 0.1 K	26.41	+1.42	11.40	+2.32	0.2814
second discharge sensor and supply-air sensor	26.78	+1.05	12.10	+1.62	0.3230
second discharge sensor and water-side differences	26.92	+0.91	12.11	+1.61	0.3303
second discharge sensor and heat meters	26.77	+1.06	12.12	+1.60	0.3234
second discharge sensor and the reference test at 45	26.44	+1.39	11.22	+2.50	0.2795

At the stated grades the second discharge sensor leads; regraded, two others pass it. Heat meters at 1.5 kW add +0.89 and +1.60 windows, against the 3.0 kW meters' +0.61 and +1.00, and the 0.3 K second sensor's +0.70 and +1.30. A supply-air sensor at 0.1 K adds +1.42 and +2.32, more than the second discharge sensor at the same grade, +1.04 and +1.94. Both read unit 1's drift-free discharge, the supply air through the mixed cold aisle; the supply air also reads the total by the cold aisle's drop, with no meter, and at 0.1 K that route reads one window's total to 0.0231, finer than the meter route's 0.0311 (Section 23.3.4). A coarser second sensor, 0.5 K, keeps nearly its family value, +0.71, and loses on the own bands, +0.88. The water-side differences gain from a finer sensor, +0.43 and +0.84 at 0.05 K, and from a better-known flow, +0.38 and +0.69 at 3 %, without reaching the second sensor; the return-air pair at 0.1 K, +0.54 and +0.71, does not reach it either.

The pairs add little to the second discharge sensor alone: with the supply air, the water-side differences or the heat meters the family's addition is +1.05, +0.91 and +1.06 against its +0.70, and the own bands' +1.62, +1.61 and +1.60 against +1.30, because the drift, once separated, is separated. Only the reference test at the refit adds clearly more, +1.39 and +2.50, by reading the level at window 45 at 0.1 K.

23.4.4 The airflow freed

The airflow station is worth nothing on the reported route because the airflow is pinned. The check frees it: each unit's capacity rate becomes $C_u\,\mathrm{e}^{\iota_u^{\text{air}}}$, with ι_u^{air} a constant entry of prior $\mathcal{N}(0,\ 0.05^2)$, and the six drawings are refitted on the record to window 36. The probabilities move little, the leading drawing's to 0.3769 against 0.3565, and the record itself reads the two airflows to standard deviations of 0.0343 and 0.0359 on the leading drawing, against the prior's

0.05. With the airflow free, no new rows give a band of 36.26 windows on the family and 26.10 own, and the continuation 26.96 and 12.87. Airflow stations on both units take that to 26.76 and 12.64, adding +0.19 and +0.23; the second discharge sensor takes it to 26.29 and 11.81, both heat meters to 26.19 and 12.27, the water-side differences to 26.60 and 12.51, and the supply air to 26.31 and 12.16. Freed, the airflow station stands below every candidate run with the airflow freed; its +0.19 and +0.23 are not nothing, and they exceed several of the pinned route's additions, the return-air pair's and coil 2's instruments' among them. With the airflow free the heat meters edge past the second discharge sensor on the family's band, and the second sensor keeps the lead on the own bands.

23.4.5 The stated truth

The record is constructed from a stated truth (Section 23.6.5), and the last check reads the candidates on it. Chapter 14's 20 continuations at the truth to window 45, with the same readings, are read with each candidate's channels at the truth, the flow factors zero, and the six drawings are refitted at window 45. Table 23.11 gives the result. Each row is a candidate; the columns give the medians over the continuations of the band's ends and width, the paired addition with its standard error, the median crossing, the tail at the shutdown and the probability of a crossing by window 47, the share of continuations whose band holds window 47, and the median of $\Pr(\mathcal{C}_1 \mid \mathbf{y}_{0:45})$.

Table 23.11: The candidates at the stated truth: Chapter 14's 20 continuations to window 45, each candidate read at the truth, the six drawings refitted there. Medians over the continuations, except the additions (means, with standard errors) and the coverage (a share).

Candidate	Band [window]	Width [window]	Adds [window]	Median [window]	$\Pr(t^* \le 46)$ [–]	$\Pr(t^* \le 47)$ [–]	Covers 47 [–]	$\Pr(\mathcal{C}_1)$ [–]
continuation	45.00–50.62	5.58	–	47.15	0.2201	0.4618	0.95	0.9991
heat meter, coil 1	45.31–50.25	4.95	+0.78 (0.17)	47.19	0.1772	0.4436	0.95	0.9998
heat meter, coil 2	45.18–50.61	5.43	+0.11 (0.10)	47.29	0.1925	0.4269	0.95	0.9990
heat meters, both	45.36–50.29	4.96	+0.82 (0.18)	47.24	0.1693	0.4325	1.00	0.9997
second discharge sensor	45.33–50.03	4.62	+1.05 (0.20)	47.10	0.1828	0.4680	0.95	0.9997
water-side difference, coil 1	45.19–50.35	5.23	+0.71 (0.16)	47.21	0.1876	0.4448	0.95	0.9997
water-side difference, coil 2	44.95–50.41	5.46	+0.08 (0.03)	47.05	0.2334	0.4848	0.95	0.9992
water-side differences, both	45.17–50.44	5.32	+0.69 (0.17)	47.23	0.1922	0.4493	0.95	0.9997
return-air pair	45.24–50.65	5.34	+0.18 (0.11)	47.29	0.1799	0.4257	0.95	0.9992
supply-air sensor	45.19–49.98	4.75	+0.68 (0.19)	47.00	0.2134	0.5015	0.95	0.9995
reference test at 45	45.57–49.53	3.96	+1.91 (0.25)	47.14	0.1326	0.4532	1.00	0.9999

The continuation alone gives a band 5.58 windows wide about a median of 47.15, which holds the truth's window 47 in 0.95 of the continuations, with the curvature decided, $\Pr(\mathcal{C}_1) = 0.9991$: Chapter 14's correction of its late forecast, on the six drawings. The additions to it rank the reference test first, +1.91 windows (standard error 0.25), then the second discharge sensor, +1.05 (0.20), both heat meters, +0.82 (0.18), the coil-1 meter +0.78, the coil-1 difference +0.71, the pair of differences +0.69 and the supply air +0.68, with the return-air pair at +0.18 and coil 2's meter and difference at +0.11 and +0.08. The second discharge sensor's band is 4.62 windows against the continuation's 5.58. At the truth the own bands' leading three hold: second sensor, heat meters, coil-1 meter, with the return-air pair near the bottom. The supply air does not hold its place: third of the permanent sensors on the family's band of the reported route, it falls at the truth to sixth of them, seventh with the reference test, behind the coil-1 meter and both water-side difference candidates. This is the sensitivity row Chapter 14's calibration

finding asks for (Section 14.4.5): if the forecast is as late as fresh records show, the refit corrects the date, and the instrument to add does not change. The decision does not flip.

The bands at the truth are drawn in Figure 23.6. Each row is a candidate; the bar runs between the medians of the band's two ends over the continuations, the dot is the median crossing, and the dotted line is the truth's crossing window, 47 (Nov Y4). Every band holds the truth's window; the instruments that separate coil 1 shorten the band at both ends, and every median lies between 47.00 and 47.29.

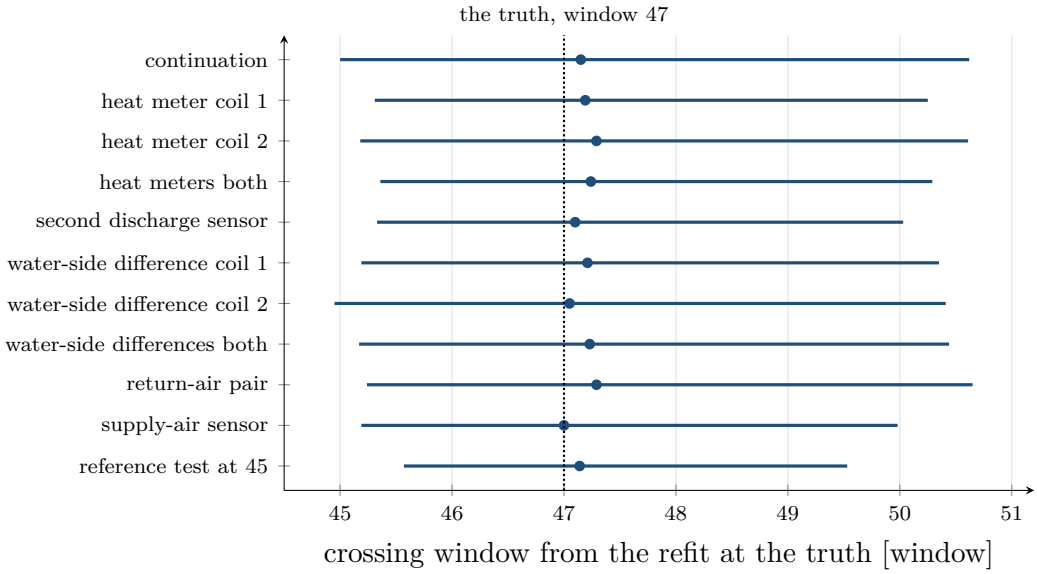


Figure 23.6: The crossing band from the refit at the stated truth under each candidate (Table 23.11): each bar runs between the medians, over Chapter 14's 20 continuations, of the band's 5 % and 95 % ends, and the dot is the median crossing. Dotted: the truth's crossing window, 47 (Nov Y4).

Table 23.12 collects the routes. Each row is a route; the columns give the continuation's family band and own band, the same two with the second discharge sensor, and the top three permanent sensors the route ranks.

Table 23.12: The routes to the candidates' value at the refit: the continuation's band and the second discharge sensor's, family and own, and the leading permanent sensors; td1b the second discharge sensor, hm12 both heat meters, hm1 the coil-1 meter, tc the supply air.

Route	cont., family [window]	cont., own [window]	td1b, family [window]	td1b, own [window]	Top three [–]
six drawings, probabilities held (reported)	27.83	13.72	27.13	12.42	family td1b, then hm12 and tc tied (paired run); own td1b, hm12, hm1
eighteen drawings, probabilities held	27.94	13.88	27.09	12.58	own td1b, hm12, hm1
predicted records, probabilities moving, 30 records (means)	28.54	21.09	26.99	21.80	not resolved
the airflow freed, probabilities held	26.96	12.87	26.29	11.81	family hm12, td1b, tc; own td1b, tc, hm12
the stated truth continued, 20 continuations (medians)	5.58	–	4.62	–	td1b, hm12, hm1

23.5 The report (step 6)

Two more return-air sensors, the dashboard's pick, would narrow the crossing band at the refit by 0.7 %, not 42 %: the return air is the noisiest trace, but the hall's total cannot be read through it without the load meter's error, and the date is read through the coils. The record kept running to the refit at window 45 is the measurement: it takes 27 % off the band, from 38.1 to 27.8 windows, at no cost. Of the instruments, the second discharge-air sensor on unit 1 adds the most, 0.7 windows on the family's band and 1.3 on the drawings' own, at the lowest price, R$ 900, because it gives back unit 1's drop without the drift; it is also the only sensor that measurably lowers the chance of a wrong call at the shutdown on the predicted records. Heat meters do nearly as much on the family's band for R$ 24 000. Coil 2's instruments and the airflow stations do almost nothing for this date. And half the remaining band is the record's indecision about coil 1's acceleration, which no instrument installed now settles, as Chapter 14 found; the refit settles it in more than half the predicted records.

23.5.1 The answers and their bands

Table 23.13 gives the answers in the order of Table 23.2. Each row is a question; the second column gives the reported value, the probabilities held over six drawings; the third the eighteen drawings' and the predicted records' values beside it; and the fourth the price or the range over the choices of step 5.

The answers rest on the probabilities held, which make every candidate's covariance independent of its outcome, and on the stated prices and grades. On the family's band, read by the paired run, the second discharge sensor leads and the heat meters and the supply air tie behind it; on the own bands the leads are wider; the predicted records, with the probabilities moving, do not rank them. The choice of the second discharge sensor rests on the own bands, the margin, the price, the wrong-call chance and the truth, which agree.

23.5.2 The entry in the maintenance log

The ranking is the next entry of the two-unit hall's maintenance log, at window 36 beside the entries of Section 23.1.2. It uses nothing later than window 36; the truth enters the checks only. In the values of Table 23.13 it reads:

> *Window 36 (Dec Y3), which instrument to add before the refit at window 45 (Sep Y4), valued on the crossing date of the capacity limit ($G_1 + G_2 \geq 10.761\,\mathrm{kW/K}$) over Chapter 14's drawings.*
> The refit itself: the band 38.1 → 27.8 windows; no cost; keep it.
> Install now: a second discharge-air sensor on unit 1 (0.3 K, R$ 900), logged with the record. It adds 0.7 windows to the family's band and 1.3 to the drawings' own, and it takes 0.03 off the chance of a wrong call at the shutdown.
> At the refit: a reference condition test (R$ 1 500); with the new sensor, 1.4 and 2.5 windows.
> Not for this date: return-air sensors (0.2 windows), heat meters (0.6 windows for R$ 24 000), coil 2's instruments, airflow stations (nothing while the airflow is pinned).
> Caution: half the band is coil 1's undecided acceleration; the refit decides it in about half the predicted records; no instrument does.
> The side-by-side check of unit 1's sensor (Chapter 7's open entry) is superseded for this date by the permanent second sensor, which reads the drift every month.

The entry names what not to buy beside what to buy, because the dashboard's pick is the instrument the log is most likely to be asked for. Whether the cleaning at the shutdown is worth its cost is Chapter 17's question, and how often to log the new sensor is Chapter 24's.

Table 23.13: The report on the candidate instruments at window 36 (Dec Y3), valued at the refit at window 45 (Sep Y4). Reported values on the six drawings with their probabilities and centers held, 10 000 draws per drawing; predicted records from 30 records.

Asset manager's question	Reported value	Eighteen drawings; predicted records	Price and range
Which added measurement most tightens the remaining-life band?	the continuation: 38.10 → 27.83 [window], 27.0 % off [–]; added to it, the second discharge sensor +0.70 (own +1.30) [window], first; heat meters +0.61 (+1.00) and supply air +0.60 (+0.83), tied; the return-air pair +0.18 (+0.27); paired run: second sensor over heat meters +0.14 (0.08), over supply air +0.17 (0.06) [window]	eighteen: own +1.30, +1.02, +0.90 [window]; records: +1.55, +1.79, +1.88 (standard errors 1.09 to 1.26), not ranked	R$ 0; R$ 900, R$ 24 000, R$ 1 200; at the truth +1.05, +0.82, +0.68 [window]
Which one makes the call at the shutdown surest?	margin's own band 0.4123 → 0.3409 kW/K with the second discharge sensor, at the rates-known floor 0.3389; median margin 0.3411 kW/K	records: wrong-call chance 0.0716, the second sensor takes off 0.0299 (0.0145) [–]	reference test at 45: 0.2929 kW/K
Which is worth its price?	the second discharge sensor, about 1.4 windows per R$ 1 000 (own; Monte Carlo error 0.08) [window/R$ 1 000]; heat meters 0.042, a factor of about 35	–	heat meters would match at about R$ 690; the reference test at 45, 1.6, one-off
Why does a sensor on unit 1 narrow the date and two more return-air sensors do not?	within a drawing the date reads the total's level (87.3 %) and rate; one window's total from the return air and meter to 0.0311, three return-air readings 0.0247 [–]; the second sensor separates the drift, 0.1026 → 0.0666 K/yr	–	with the meter exact, the return-air pair 0.0705 kW/K against 0.1179
How much of the band can any sensor reach, and what is left?	trend's floor 24.84 [window], 2.99 below the continuation; the best permanent sensor reaches 23.4 %; the indecision 14.11 of 27.83 [window]	records: the curvature decided in 0.57 [–], 0.60 with the second sensor	path known too: 23.80 [window]

Common mistake

Choosing an instrument by the scatter of its own trace. The return air scatters at 4.42 times its grade, and averaging three of it would cut that trace's noise by 42.3 %; it takes 0.7 % off the crossing band. Value an instrument by the width it takes off the answer the decision reads, beyond what the record's continuation already takes off, and read that answer through the rows the instrument adds.

23.6 The next measurement (step 7)

The report leaves a band of 27.83 windows at the refit, of which the instruments can take at most the 2.99 above the trend's floor, and half of which is the family's indecision. What remains is set by what the evidence between the drawings can do by the refit.

23.6.1 What remains after the instruments

The family's indecision is 14.11 windows of the continuation's 27.83, and no instrument installed now reaches it with the probabilities held. The measurement that does is time, the record itself. Chapter 11's record continued to window 48 (Dec Y4) decides coil 1's curvature in 0.668 of predicted records, against 0.545 at window 45 (Section 11.6.3); here the refit at window 45 decides it in 0.57 of the predicted records, and with the second discharge sensor the expected $|\ln B(\mathcal{C}_1 : \mathcal{C}_0)|$ rises by 0.435 over the continuation's. The record to window 48 is past the work order, so for this decision the refit at window 45 is the last evidence on the curvature before the shutdown.

23.6.2 The candidates for the refit

The candidates for the refit itself are the pairs and grades of Table 23.10, by the preposterior (23.5) with the probabilities held. The second discharge sensor with a reference test at window 45 adds 1.39 windows to the family's band and 2.50 to the own, the largest pair, for R$ 900 and R$ 1 500. A second discharge sensor of reference grade, 0.1 K, adds 1.04 and 1.94 alone: a better instrument at the same place, worth choosing if one is available at a like price. The permanent instrument at the place of the question is the second discharge sensor on unit 1, because unit 1's drop is the quantity the date reads worst; it is the recommendation, not one candidate among the step's.

23.6.3 The recommendation

Install the second discharge-air sensor on unit 1 now, at 0.3 K or better, and log it with the record. Keep the record running, and refit at window 45 with a reference condition test before the shutdown's work order is written. Leave the heat meters, the return-air pair, the water-side differences and the airflow stations out of this decision: their value elsewhere, the coils' split for a warranty claim (Chapter 21) or the airflow behind its pin, is not this date's.

23.6.4 What the analysis showed

The dashboard pointed at the noisiest trace and promised 42.3 %; two more return-air sensors take 0.18 windows, 0.7 %, off the crossing band at the refit. The physics put the date on the coils' total, which the hall reads through its two drops, one of them through a drifting sensor, and only through the load meter's error by the return air. The probability turned each candidate's rows into a band before a reading was logged, and the mixture showed that half the band at the refit is the drawings' indecision, beyond any instrument. The price turned the widths into a choice: R$ 900 against R$ 24 000 for most of the same narrowing. The analysis did not rank the

sensors with the probabilities moving to a precision that separates them, which 30 records do not, and it did not value any sensor for another question, such as the coils' split or a warranty claim.

23.6.5 Limits of the model

A constructed record. The condition record is generated from a stated truth, with the stated sensor errors, and step 5 reads the candidates on that truth (Section 23.4.5). It is not field data. At that truth the own bands' leading three hold, and the instrument to add is the same.

The probabilities held. The reported route holds the drawings' probabilities and centers; the predicted records show that moving them adds value to every candidate but, at 30 records, does not rank the sensors.

Six drawings of eighteen. The six hold 0.9615 of the family, and the eighteen give the same leading order.

The instruments as modeled. The airflow is pinned, and the station is worth little even when it is freed. The water flow behind the temperature differences is constant over nine windows. New sensors are installed in calibration and logged one hour a month with independent errors. Prices are stated inputs, not quotes.

Chapters 10 to 14's limits. The readings are independent across windows, the drift is linear from window 20, each window is in steady state with known inputs, and no maintenance comes before the crossing.

The next chapter. Chapter 24 asks whether reading the existing sensors more often in each window narrows the same quantities: the logging rate, not a new instrument.

23.7 What to retain

Definition 23.1 (Value of a sensor for a date)**.** The value of a candidate instrument a for a forecast date is the width it takes off the date's band, beyond what the record's own continuation takes off, before any of its readings is logged: $\Delta\mathcal{B}_a = \mathcal{B}_{\text{cont}} - \mathcal{B}_a$, with $\mathcal{B}_a$ the 5–95 % band of the first passage t^* under each drawing's preposterior covariance $\mathbf{\Sigma}_{M,a} = (\mathbf{J}_{\text{now}}^{\mathsf{T}}\mathbf{J}_{\text{now}} + \mathbf{J}_{\text{cont}}^{\mathsf{T}}\mathbf{J}_{\text{cont}} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a)^{-1}$, mixed over the family at its probabilities. For linear rows the covariance does not depend on the readings; holding the drawings' probabilities and centers then makes the value a closed-form function of the candidate's Jacobian lines, its own-band part $\Delta\bar{\mathcal{B}}_a$ the instrument's effect within the drawings, and the remainder the family's indecision, which only evidence between the drawings removes. It is read against the floor of the rates known exactly and per real, $\Delta\bar{\mathcal{B}}_a/c_a$.

Three conclusions carry beyond the hall. A sensor is valued by the answer the decision reads, not by its own trace: the noisiest trace may read a quantity the answer barely depends on, and a trace's noise averaged away is not a band narrowed. The answer's own physics decides which entries matter: a crossing date reads the total's level and rate, the total is read best by load-free routes through the units, and the instrument that removes a corruption on one of those routes, a drift, is worth more than the instrument of the highest grade on another. And the record's own continuation comes first: an instrument's value is what it adds to the continuation, in the

windows left above the floor, and a band that is mostly the family's indecision is narrowed by evidence between the drawings, which time supplies and an instrument mostly does not.

Table 23.14 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 23.14: The terms of sensor selection for prognostics and where each is defined. Units for the two-unit hall.

Term	Expression	Units	Meaning	Defined in
Value of a sensor for a date	$\Delta\mathcal{B}_a$, eq. (23.1)	window	width taken off beyond the continuation	23.1.1
Family's band, own bands	$\mathcal{B}_a$; $\bar{\mathcal{B}}_a = \sum_M P_M \mathcal{B}_{M,a}$	window	the mixture; within the drawings	23.1.1
Family's indecision	$\mathcal{B}_a - \bar{\mathcal{B}}_a$	window	beyond any instrument	23.3.1
Preposterior at the refit	$\mathbf{\Sigma}_{M,a}$, eq. (23.5)	mixed	the record, its continuation, the candidate	23.2.2
Trend's floor at the refit	the rates known exactly	window	the least an instrument can reach	23.2.2
Reading the rows				
Date's variance by direction	$\mathrm{Var}(t^* \mid \mathbf{Qz})$, Table 23.8	window^2	the state at the refit, through the total's level and rate; the wander after it outside	23.3.4
Load-free reading of a coil	$(T_r - T_{d,u})/(T_r - T_s) = G_u/C_u$, eq. (23.8)	–	the conductance without the meter; unit 1's through the drift	23.3.4
The return-air route	$(Q + Q_o)/(T_r - T_s)$	–	carries the load meter's error	23.3.4
Water-side difference without flow	$D_u/(C_w \mathrm{e}^{\iota u})$, eq. (23.3)	K	reads the duty's change	23.1.4
Instrument on a pinned quantity	a zero Jacobian line	–	moves no answer	23.1.4
Choosing				
Value per real; matching price	$\Delta\bar{\mathcal{B}}_a/c_a$; $c_a^{\times}$, eq. (23.7)	window per R$ 1 000	the price beside the width	23.3.3
Wrong-call chance at an outage	$\mathbb{E}[\min(p_{46}, 1 - p_{46})]$	–	the risk of the hold-or-cross call	23.4.2
Predicted records	refit and reweighted, Section 23.2.5	–	the probabilities moving; need many records	23.2.5

23.8 Exercises

Exercise 23.1. *Application.* Give the planner's arithmetic for three return-air readings averaged, and read what the return-air pair adds to the continuation's band at the refit. (Averaging three readings divides the trace's noise by $\sqrt{3}$, so the planner expects $1 - 1/\sqrt{3} = 0.4226$ of the band to go; by Table 23.6 the pair takes 0.18 windows off the family's 27.83, 0.7 %, and 0.27 off the own 13.72, 2.0 %.)

Exercise 23.2. *Derivation.* Show that one window's total conductance read from the return air and the load meter, $(Q + Q_o)/(T_r - T_s)$, and a unit's drop over the approach, equation (23.8),

with $a = T_r - T_{d,u}$ and $b = T_r - T_s$, have the relative standard deviations

$$\sqrt{\Big(\frac{\sigma_Q}{Q + Q_o}\Big)^2 + \Big(\frac{\sigma_T}{T_r - T_s}\Big)^2} \qquad \text{and} \qquad \sqrt{\frac{2\sigma_T^2}{a^2} + \frac{\sigma_T^2}{b^2} - \frac{2\sigma_T^2}{ab}},$$

the second with no load in it; evaluate both at window 45. (The relative error of a ratio is the difference of the relative errors, with Q_o and T_s known inputs, giving the first form; for the drop, the return air's error e_r enters both a and b and the discharge's e_d only a, so the relative error is $e_r(1/a - 1/b) - e_d/a$, whose variance is the second form. At $Q + Q_o = 144.44\,\mathrm{kW}$, $T_r - T_s = 12.985\,\mathrm{K}$ and $0.3\,\mathrm{K}$ sensors: 0.0311 for the total, 0.0247 with three return-air readings, and 0.0465 for unit 1's drop and 0.0444 for unit 2's. The two drops together read the total to 0.0347.)

Exercise 23.3. *Derivation.* Write the water-side difference's row on coil 1 with the unknown constant flow factor ι_1, give its slopes at window 45, and say why it reads the duty's change and not its level. (The row is $(D_{1,45}/(C_w \mathrm{e}^{\iota_1}) - y)/0.1$; its slopes are those of the duty over $C_w \times 0.1$, $+25.7827$ on $g_{1,45}$ and -25.7827 on $g_{2,45}$, and on the flow $-D_{1,45}/(C_w \times 0.1) = -50.5742$. A constant ι_1 scales every window's reading alike, so the reading's level is shared with the flow's prior of 10 % while its changes from window to window are read at the sensor's grade.)

Exercise 23.4. *Application.* Compare the second discharge sensor and both heat meters per real on the own bands, and find the heat meters' matching price. (By Table 23.7, 1.4402 and 0.0415 windows per R$ 1 000, a ratio of 34.7; by equation (23.7) the heat meters would match at R$ 691. To two figures, given the Monte Carlo error of about 0.08 on the second sensor's addition: 1.4, 0.042, about 35 and about R$ 690.)

Exercise 23.5. *Formulation.* Write "how much of the date's variance within a drawing does the total's level carry?" as a conditional variance, read it, and say why the shares of the directions do not add. (An estimation question, equation (2.9), with $f = t^*$ from the refit and the level $\mathbf{Qz} = G_{\mathrm{tot},45}$ known: $1 - \mathrm{Var}(t^* \mid \mathbf{Qz})/\mathbf{g}_*^{\mathsf{T}} \boldsymbol{\Sigma}_5 \mathbf{g}_*$, which is 87.3 % for the level, 37.9 % for the split, and 100.0 % for the level and rate together. The directions are correlated, so each share counts what it shares with the others.)

Chapter 24

Sampling Frequency and Correlated Readings

24.1 The proposal to log faster and what a logged hour holds

The hall is three years past its handover. Today is December of year 3, window 36, the last window of the condition record. The record holds one logged hour a month for each of the hall's five channels, taken from the dashboard: the return air, the two discharge temperatures and rack 1's inlet, each read to a standard deviation of 0.3 K, and the hall's load meter, read to 3.0 kW. The building-management system keeps a trend log of the same five channels and can export it at any rate, down to one minute. The proposal on the asset manager's desk reads: hourly logging gives 720 readings a month instead of one, so the trend will be $\sqrt{720} = 26.8$ times sharper. Read as independent readings on a straight line, three years of hourly logging would know coil 1's degradation rate to ± 0.058 %/yr, the half-width of its 5–95 % band, where today's band is 4.613 %/yr wide; and a one-minute log would be sharper still. The conclusion the proposal invites is to log as fast as the system allows, and the rate will be known almost exactly.

The proposal cannot say how many independent readings an hour of logging holds, when one hour's error carries into the next: a sensor's own error and the hour's departure from the month's steady condition both persist from one hour to the following one. Nor can it say which part of the rate's band any logging can touch, when Chapter 10 found about half of each rate's variance to be the path's own wander from month to month; or whether the gain lies in the rate at all, or in today's health index. The question of this chapter is

how much each logging rate narrows each coil's degradation rate and today's health index, with the readings' autocorrelation carried in the rows, and which rate the logger should be set to.

It is answered by the design form of Chapter 2 on the stacked state of Chapter 10, in which the logged readings of each window are rows with a full correlation matrix.

Jargon: sampling frequency and logging interval, autocorrelated reading errors, correlation time, effective number of readings, prewhitening, saturation, error shared by a window

The *sampling frequency* of a condition record (its logging rate) is the number n_{L} of readings of each channel logged in one window; the *logging interval* is the time between two of them. Reading errors are *autocorrelated* when the error of one reading carries into the next; in the first-order autoregressive form, AR(1), the correlation of two errors falls by one fixed factor for every hour between them, and the *correlation time* is the lag over which it falls by the factor e. The *effective number of readings* of a window is the number of independent readings whose mean would be as precise as the window's logged readings taken together. *Prewhitening* rewrites rows of correlated readings so that their errors are independent (the whitened rows): each keeps only the new part of its reading's error. *Saturation* is the limit of the effective number as the logging interval goes to zero over a fixed span. An *error shared by a window* is a part of the reading error common to every reading of the window, which no logging rate averages away.

Key idea

A window's logged readings add information only as far as their errors are independent. For readings of one steady value the whitened rows collapse to one row, whose standard deviation is the reading's over the square root of an effective number, and that number saturates at about one reading per two correlation times of the window's span. The rate's band then stops at the wander floor, which no logging touches. Today's health index has no such floor: it keeps gaining until the effective number saturates.

24.1.1 Readings that remember

Chapter 3 read two condition tests by one sensor whose errors shared a calibration offset (Table 3.8): the pair added the precision of its own average, less than two independent readings would, and treated as independent it reported a band narrower than the evidence supports. Chapter 10's step 7 asked for a second logged hour of every window from the logger's archive and gave it an error independent of the first. Here a window's readings are many, logged at a fixed interval over the month, and each reading's error remembers the one before it. In the AR(1) form the error of reading i of a channel with standard deviation σ_k is

$$e_i = \varrho\, e_{i-1} + \epsilon_i, \qquad \epsilon_i \sim \mathcal{N}\big(0,\ \sigma_k^2(1-\varrho^2)\big), \qquad \mathrm{Cov}(e_i, e_{i'}) = \sigma_k^2\, \varrho^{|i-i'|}, \tag{24.1}$$

with ϱ the correlation of two consecutive readings' errors and ϵ_i the new part of reading i's error, independent of every earlier error. Every reading keeps the marginal standard deviation σ_k, whatever the interval. What changes with the interval is the share of each error that is new, $1-\varrho^2$ of its variance: near one when the readings are far apart, near zero when they follow each other faster than the error forgets.

> **Principle 24.1.** *A reading adds the information of its error's new part only. Logged faster than the error forgets, the new part per reading shrinks to nothing, and the effective number of readings stops growing however many readings are taken.*

24.1.2 The two-unit hall through its logged record (step 0)

The plant graph is Chapter 10's stacked record, cited and not re-derived (Figure 10.1): 37 copies of Chapter 5's window, window 0 (commissioning) and windows 1 to 36, each with its free load Q_t and the two log-derates $g_{1,t}$, $g_{2,t}$ under the five physical rows (5.2) at the window's

known inputs $T_s(t)$ and $Q_o(t)$; one trend row per coil between consecutive windows, linked to the trend entries γ_1 and γ_2; the shared recirculation fraction β_1; and the drift node $\dot{b}_{d,1}$ on unit 1's discharge reading from window 20, equation (10.2). Its state is the stacked state (10.3), 115 entries, and its fit is Part III's base fit: the straight line with wander at $\sigma_q = 0.005$, the drift from window 20, the rating, survey and drift priors. This chapter adds no entry to that state and no new kind of row.

What logging adds to the graph. Each reading edge of a window carries n_L logged readings in place of one, all of them readings of the same window's state. Their errors form a chain, each error node linked to the one before it by the correlation ϱ of equation (24.1). The chain does not cross from one window to the next, and the five channels of a window have a chain each.

The logging's stated inputs. Three inputs of this chapter are stated, not estimated, because no logged record exists to estimate them from.

- *A logging rate.* n_L readings of each channel per window, equally spaced over the window's span $\vartheta_w = 720\,\mathrm{h}$ (30 days), at the logging interval $\Delta\vartheta = \vartheta_w/n_\mathrm{L}$. Monthly is $n_\mathrm{L} = 1$, the record as it stands; weekly $n_\mathrm{L} = 4$; daily $n_\mathrm{L} = 30$; every six hours $n_\mathrm{L} = 120$; hourly $n_\mathrm{L} = 720$; every 15 minutes $n_\mathrm{L} = 2880$; every minute $n_\mathrm{L} = 43\,200$.
- *The readings' errors.* Each logged reading keeps the record's standard deviation: 0.3 K on the return air, the two discharges and rack 1's inlet, and 3.0 kW on the meter. In time the errors are AR(1), equation (24.1), with the correlation $\varrho_1 = 0.9$ between readings an hour apart and ϱ_1 raised to the lag in hours at any other lag: 0.5314 at 6 h, 0.2824 at 12 h, 0.0798 at a day and 2.05×10^{-8} at a week. The correlation time, $-1/\ln\varrho_1$ in hours, is 9.49 h. The error holds the sensor's own error and the hour's departure from the window's steady condition. The errors of the five channels are independent of each other.
- *The cost of logging.* Monthly costs R$ 0 a year, the dashboard's hour, which is read already. Any finer rate costs R$ 1 500 a year: the historian's export of the five channels into the fit, and its upkeep. This price is the chapter's stated input and not an entry of the plant's cost schedule. One entry of that schedule costs the same: a crew visit with reference instruments, R$ 1 500.

Two planning dates are the record's: the refit at window 45 (Sep Y4), before the October shutdown at window 46 (Oct Y4), which Chapter 11's entry asks for (Section 11.5.2).

The maintenance log so far. The two-unit hall's log holds the commissioning test (window 0); the condition assessment and the effort entry of window 18 (Chapters 5 and 6); Chapter 7's assessment of window 25, which asked for a side-by-side check of unit 1's discharge sensor (Section 7.5.2); Part III's entries at window 36, the rates, the acceleration, the step and the season; and the entries the earlier chapters of this volume have added at the same window, among them the forecast of the crossing with its refit at window 45 (Section 14.5.2) and the power of the monthly trend log (Chapter 22). The side-by-side check is still open: Chapters 7, 10 and 11 kept it so. This chapter's entry is also at window 36.

Modeling choices at the graph. The choices that no question block owns are Chapter 5's paragraph, cited through Chapter 10: steady state within each window, the airflow pinned at the units' fixed fan speed with nothing controlling the discharge air, the chilled-water supply and the other heat as known inputs, and one effective conductance per coil. The logging adds one choice, named here once: *one steady condition per window.* The window's state is its month's representative condition, and every logged reading of the window reads it; the hours' departures from it, of the load, the supply and the daily cycle, are part of each reading's error, which is why that error is correlated over hours. This buys a window with one state whatever n_L is, so that the stacked state keeps its 115 entries at every logging rate. It fails where a departure

moves all five channels together, as an hour's load departure does: that departure is then a quantity of the state, not an error, and it needs an entry of its own for every logged hour, the form of Chapter 10's archive candidate, which gave each archived hour its own load entry.

The graph. One window of Figure 10.1 is opened up in Figure 24.1. The top row holds window t (green) between its two neighbors (dashed), tied to them by the trend rows (black squares), which are linked to the trend entries above. Below window t is one of its five reading edges, the return air, opened into its n_{L} logged readings; under each reading is its error, and the errors form a chain, each arrow labeled ϱ. The brace marks what step 2 makes of the chain: whitened rows, collapsed into one row. Under each neighbor is its own chain of errors, with no link to window t's.

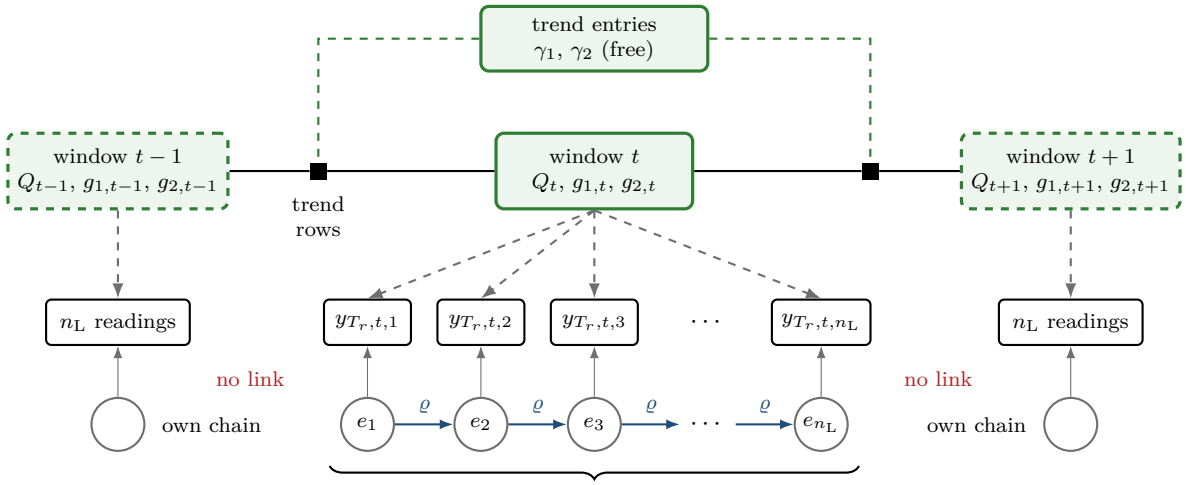


Figure 24.1: One window of the condition record, Figure 10.1, opened up for a logging rate of n_{L} readings per window. Free (green): window t's IT load Q_t and log-derates $g_{1,t}$, $g_{2,t}$, its neighbors' (dashed boxes), and the trend entries γ_1, γ_2, linked to the trend rows (black squares). Observed (dashed arrows): the n_{L} logged readings of window t's return air, all of the one window's state; the other four channels of the window, not drawn, have the same form. Under each reading is its error (gray circle); the errors form an AR(1) chain, each link carrying the correlation ϱ of two consecutive readings, equation (24.1). Brace: the chain whitened and collapsed into one row (step 2). The neighbors' readings have chains of their own, with no link to window t's. Pinned in every window, not drawn: the known inputs $T_s(t)$ and $Q_o(t)$, the ratings $\bar{G}_u = 7.0\,\mathrm{kW/K}$ and the capacity rates $C_u = 10.0\,\mathrm{kW/K}$; also not drawn, the shared fraction β_1 and the drift rate $\dot{b}_{d,1}$ of Figure 10.1.

24.1.3 The asset manager's questions (step 1)

The asset manager asks six questions of the record, and step 1 writes each as a design question over the stacked state (10.3): a spread of the posterior that a record logged at a stated rate would give, before any of it is logged. Table 24.1 lists them. The first column is the question in the asset manager's words, the first of them the drill question; the second the general form of its category from Chapter 2; the third its mathematical form for the hall; the fourth its category; the fifth the computation of Section 2.3.6 that answers it; and the last the subsection where it is answered. Every question is asked of both coils. $\mathbf{\Sigma}_{n_{\mathrm{L}}}$ is the posterior covariance of the record logged at n_{L} readings a window, and n_{eff} the effective number of readings, both derived

in step 2 and step 3.

Table 24.1: The questions asked of the two-unit hall's condition record at window 36 (Dec Y3) about its logging rate. $\mathbf{\Sigma}_{n_\mathrm{L}}$ is the preposterior covariance of windows 0–36 logged at n_L readings per window; $\mathbf{\Sigma}^{\mathrm{line}}_{n_\mathrm{L}}$ the same on the straight line; $\mathbf{\Sigma}_a$ that of the record extended to window 45 with a candidate; $u = 1, 2$.

Asset manager's question	General form	Mathematical form	Category	Computation	Answered in
Hourly, daily, or monthly?	(2.16) with U_p the spread; (2.17)	$\sqrt{[\mathbf{\Sigma}_{n_\mathrm{L}}]_{\gamma_u\gamma_u}}$ for windows 0–36 logged at n_L readings, against the wander floor $\sigma_q/(\sqrt{36}\,\Delta t)$; the band's width in %/yr	design	preposterior; trend	24.3.1
How many independent readings does a month of logging hold?	(2.16); (3.29) with the readings' covariance full	$n_{\mathrm{eff}}(n_\mathrm{L}) = \mathbf{1}^\mathsf{T}\mathbf{R}^{-1}\mathbf{1}$ for the AR(1) correlation $\mathbf{R}$, and its limit as $n_\mathrm{L} \to \infty$	design	preposterior	24.3.2
From which rate on does logging stop narrowing the rate's band, and why?	(2.16); (4.74)	the readings' share $[\mathbf{\Sigma}^{\mathrm{line}}_{n_\mathrm{L}}]_{\gamma_u\gamma_u}/[\mathbf{\Sigma}_{n_\mathrm{L}}]_{\gamma_u\gamma_u}$ over n_L; the first n_L with a share of 10 % or less, and of 5 %	design	preposterior; trend	24.3.3
Does a faster rate help today's health index and the capacity margin?	(2.16)	$\sqrt{\mathbf{g}^\mathsf{T}\mathbf{\Sigma}_{n_\mathrm{L}}\mathbf{g}}$ for $\ln h_{u,36} = g_{u,36} - g_{u,0}$ and for $G_{1,36} + G_{2,36}$	design	preposterior	24.3.4
What does the band look like if the logged readings are treated as independent?	(2.16); (3.29), the covariance diagonal against full	the same spreads with $n_{\mathrm{eff}} = n_\mathrm{L}$, against the rows' own; their ratio	design	preposterior	24.3.5
What should the logger be set to for year 4, against a crew visit at the same price?	(2.16), a candidate's cost a loss of its own	$\sqrt{[\mathbf{\Sigma}_a]_{\gamma_u\gamma_u}}$ and $\sqrt{\mathbf{g}^\mathsf{T}\mathbf{\Sigma}_a\mathbf{g}}$ for $\ln h_{u,45}$ and $G_{1,45} + G_{2,45}$, windows 37–45 logged at each rate, beside the crew visits	design	preposterior	24.3.6

Three questions near these are not asked here. Whether a record logged at a given rate decides an acceleration or a step is a question of power, Chapter 22's computation; step 5 reports what the simulated records show of it and stops there. Which permanent instrument narrows the crossing date is Chapter 23's. And step 7 values the instruments on the rate and the level only, once the logging is set.

24.1.4 The evidence as rows (step 2)

Step 2 takes Chapter 10's rows as they stand and changes only what a window's reading rows become when each channel is logged n_L times. The base fit has 261 rows on 115 entries: 185 reading rows, five per window, each Chapter 5's row (5.7) with unit 1's discharge carrying the drift, equation (10.4); two rating priors and the survey prior; 72 trend rows, equation (10.5); and the drift prior.

The rows of a window logged n_L times. Write k for one of a window's five channels, σ_k for its standard deviation and $f_k(\mathbf{z}_t)$ for its model value, the quantity of the window's entries it

reads: a temperature of equation (5.3) at the window's inputs, with the drift of equation (10.4) added to unit 1's discharge, or the load Q_t for the meter. Logged n_{L} times, channel k of window t gives n_{L} rows on the one window's entries, with errors correlated as in equation (24.1),

$$r_{k,t,i}(\mathbf{z}) = \frac{f_k(\mathbf{z}_t) - y_{k,t,i}}{\sigma_k}, \quad i = 1, \dots, n_{\mathrm{L}}, \qquad \mathrm{Cov}(\mathbf{r}_{k,t}) = \mathbf{R}, \quad R_{ii'} = \varrho^{|i-i'|}, \quad \varrho = \varrho_1^{\Delta\vartheta/1\,\mathrm{h}}, \tag{24.2}$$

with $\mathbf{r}_{k,t}$ the n_{L} rows stacked and their covariance taken at the state that made the readings. Here $\mathbf{R}$ is the readings' correlation matrix, not Chapter 3's covariance of the readings, which for these rows is $\sigma_k^2\mathbf{R}$; that chapter's precision (3.29) holds for them with the covariance $\sigma_k^2\mathbf{R}$ full, and the rows are whitened by a square root of $\mathbf{R}^{-1}$. For the AR(1) correlation the inverse is tridiagonal,

$$\mathbf{R}^{-1} = \frac{1}{1-\varrho^2}\begin{bmatrix} 1 & -\varrho & & & \\ -\varrho & 1+\varrho^2 & -\varrho & & \\ & \ddots & \ddots & \ddots & \\ & & -\varrho & 1+\varrho^2 & -\varrho \\ & & & -\varrho & 1 \end{bmatrix}, \tag{24.3}$$

which the product $\mathbf{R}\,\mathbf{R}^{-1}$ checks row by row: $\varrho^{|i-i'|}$ times $-\varrho$, $1+\varrho^2$, $-\varrho$ sums to zero off the diagonal and to $1-\varrho^2$ on it. It factors as $\mathbf{R}^{-1} = (\mathbf{R}^{-1/2})^{\mathsf{T}}\mathbf{R}^{-1/2}$ with $\mathbf{R}^{-1/2}$ lower bidiagonal, first row $(1, 0, \dots)$ and row $i \geq 2$ holding $-\varrho/\sqrt{1-\varrho^2}$ and $1/\sqrt{1-\varrho^2}$ in columns $i-1$ and i. The whitened rows $\tilde{\mathbf{r}}_{k,t} = \mathbf{R}^{-1/2}\mathbf{r}_{k,t}$ are

$$\tilde{r}_{k,t,1} = r_{k,t,1}, \qquad \tilde{r}_{k,t,i} = \frac{r_{k,t,i} - \varrho\, r_{k,t,i-1}}{\sqrt{1-\varrho^2}}, \quad i = 2, \dots, n_{\mathrm{L}}, \tag{24.4}$$

and at the state that made the readings each is the new part of its error over its own standard deviation, $-\epsilon_i/(\sigma_k\sqrt{1-\varrho^2})$ by equation (24.1): independent rows of unit variance, as every normalized row of the book is.

The effective number of readings. The n_{L} rows read one value, so they share one Jacobian line, $\mathbf{j}_{k,t}^{\mathsf{T}} = \sigma_k^{-1}\,\partial f_k/\partial \mathbf{z}$, Chapter 10's line of that reading. By equation (24.4) the first whitened row keeps the line and every later one carries it scaled by $(1-\varrho)/\sqrt{1-\varrho^2} = \sqrt{(1-\varrho)/(1+\varrho)}$, so the whitened rows add the information

$$\mathbf{j}_{k,t}\Big[1 + (n_{\mathrm{L}} - 1)\,\frac{1-\varrho}{1+\varrho}\Big]\mathbf{j}_{k,t}^{\mathsf{T}} = n_{\mathrm{eff}}\,\mathbf{j}_{k,t}\mathbf{j}_{k,t}^{\mathsf{T}}, \qquad n_{\mathrm{eff}} = \mathbf{1}^{\mathsf{T}}\mathbf{R}^{-1}\mathbf{1} = \frac{n_{\mathrm{L}} - (n_{\mathrm{L}} - 2)\,\varrho}{1+\varrho}, \tag{24.5}$$

the information of n_{eff} independent readings. It is not the effective sample size N_{eff} of Chapter 5's reweighted draws, though both count what a set of correlated or weighted values is worth. It is also the sum of the entries of (24.3): $1/(1+\varrho)$ for each end row and $(1-\varrho)/(1+\varrho)$ for each interior one. At $\varrho = 0$ it is n_{L}, at $\varrho \to 1$ it is one.

The collapsed row. The n_{L} rows enter the objective as $\frac{1}{2}\mathbf{r}_{k,t}^{\mathsf{T}}\mathbf{R}^{-1}\mathbf{r}_{k,t}$, the sum of squares of the whitened rows. With $\mathbf{r}_{k,t} = (f_k\mathbf{1} - \mathbf{y}_{k,t})/\sigma_k$ and $f_k = f_k(\mathbf{z}_t)$, completing the square in f_k gives

$$\begin{aligned} \mathbf{r}_{k,t}^{\mathsf{T}}\mathbf{R}^{-1}\mathbf{r}_{k,t} &= \frac{n_{\mathrm{eff}}\big(f_k(\mathbf{z}_t) - \bar{y}_{k,t}\big)^2}{\sigma_k^2} + \frac{\mathbf{y}_{k,t}^{\mathsf{T}}\mathbf{R}^{-1}\mathbf{y}_{k,t} - n_{\mathrm{eff}}\,\bar{y}_{k,t}^2}{\sigma_k^2}, \\ \bar{y}_{k,t} &= \frac{\mathbf{1}^{\mathsf{T}}\mathbf{R}^{-1}\mathbf{y}_{k,t}}{n_{\mathrm{eff}}} = \frac{y_{k,t,1} + (1-\varrho)\sum_{i=2}^{n_{\mathrm{L}}-1} y_{k,t,i} + y_{k,t,n_{\mathrm{L}}}}{n_{\mathrm{L}} - (n_{\mathrm{L}} - 2)\,\varrho}, \end{aligned} \tag{24.6}$$

the second term free of the state. The window's n_{L} rows of channel k are therefore one row, $(f_k(\mathbf{z}_t) - \bar{y}_{k,t})/(\sigma_k/\sqrt{n_{\mathrm{eff}}})$, on the generalized least-squares mean $\bar{y}_{k,t}$ of the logged readings,

whose standard deviation is $\sigma_k/\sqrt{n_{\text{eff}}}$. The mean weights the two end readings by one and every interior reading by $1-\varrho$, because an interior reading's error shares ϱ with each neighbor. It is sufficient for the window's value, so the collapse loses nothing the n_{L} rows held.

Saturation. At a fixed span the correlation at the interval is $\varrho = \varrho_1^{\vartheta_w/n_{\text{L}}}$, and as n_{L} grows $\varrho \to 1$ with $n_{\text{L}}(1-\varrho) \to \vartheta_w \ln(1/\varrho_1)$. Written as $n_{\text{eff}} = (n_{\text{L}}(1-\varrho)+2\varrho)/(1+\varrho)$, equation (24.5) has the limit

$$n_{\text{eff}} \to 1 + \tfrac{1}{2}\,\vartheta_w \ln(1/\varrho_1), \qquad n_{\text{L}} \to \infty, \tag{24.7}$$

with ϑ_w in hours: one reading for the window, and one more for every two correlation times of its span, $\frac{1}{2}\ln(1/\varrho_1)$ per hour. No logging rate takes a window beyond it.

The rows of the logged record. The logged record is Chapter 10's 261 rows with every reading row's standard deviation divided by $\sqrt{n_{\text{eff}}(n_{\text{L}})}$. Split the Jacobian into its 185 reading rows $\mathbf{J}_y$, Chapter 3's name for the reading block, and its 76 other rows $\mathbf{J}_o$: the two rating priors, the survey prior, the 72 trend rows and the drift prior. Then

$$\mathbf{J}(n_{\text{L}}) = \begin{bmatrix} \sqrt{n_{\text{eff}}(n_{\text{L}})}\,\mathbf{J}_y \\ \mathbf{J}_o \end{bmatrix}, \tag{24.8}$$

261 rows on 146 degrees of freedom at every rate; logged raw, before the collapse, the record would have $185\,n_{\text{L}} + 76$ rows.

Which row acts on which entry. The pattern is Chapter 10's, equation (10.6), unchanged. Logged raw, the rows of one window repeat one line per channel n_{L} times; collapsed, they scale each window's reading lines, in its own three columns, β_1's and $\dot{b}_{d,1}$'s, by $\sqrt{n_{\text{eff}}}$, and add no nonzero. On the second line of the pattern, the questions' gradients are

$$\mathbf{g}_{\gamma_u} = \mathbf{e}_{\gamma_u}, \qquad \mathbf{g}_{\ln h_u} = \mathbf{e}_{g_{u,36}} - \mathbf{e}_{g_{u,0}}, \qquad \mathbf{g}_G = \bar{G}_1 \mathrm{e}^{g_{1,36}}\,\mathbf{e}_{g_{1,36}} + \bar{G}_2 \mathrm{e}^{g_{2,36}}\,\mathbf{e}_{g_{2,36}}, \tag{24.9}$$

with $\mathbf{e}_x$ the unit vector of entry x of equation (10.3): the rate, the health index today against commissioning, and the hall's total conductance today, $G_{1,36} + G_{2,36}$, whose gradient sits on today's two log-derates.

Two facts of the graph. Two facts are read here, before any number. The first: the logging rate scales only the reading rows. The trend rows, with their lines $\pm 1/\sigma_q = \pm 200$ in the columns of $g_{u,t}$ and $g_{u,t+1}$ and $-\Delta t/\sigma_q = -16.6667$ in the column of γ_u, and the four priors are the base fit's at every rate, and no reading row sees γ_u. No logging rate reaches the trend rows' own spread σ_q. The second: n_{eff} is bounded at a fixed span, by equation (24.7). Whatever the logging adds to the reading rows, it adds at most the information of about 39 independent readings a window.

Modeling choices at the rows. Five choices make these rows, each named once. *AR(1) errors with $\varrho_1 = 0.9$ per hour, at the record's marginal standard deviation*: this buys the closed-form whitening (24.4) and the effective number (24.5), and it fails if part of the error forgets slowly, as a calibration wander over weeks or a bias at a fixed hour of the day would; a part f_{sh} of the variance shared by the whole window caps n_{eff} at $1/f_{\text{sh}}$, which step 5 evaluates. *The same correlation on all five channels, independent across channels*: this keeps the collapse channel by channel, and it fails where one cause moves several channels together, the failure of step 0's choice. *Readings independent across windows*, Chapter 10's choice carried to logged windows, which that chapter left to this one for errors that persist across months: under the AR(1) the error a month away is negligible, and logged contiguously, two consecutive windows'

means correlate by 0.0031 at daily logging and 0.0232 at hourly, against 1.13×10^{-33} for two monthly readings; it fails for an error that persists across months, which step 5 carries at a stated share and correlation. *The wander* $\sigma_q = 0.005$, Chapter 10's choice, not re-chosen: the floor the answers meet is σ_q's, and step 5 varies it. *The preposterior on the base fit's Jacobian*: the rows are linear in the readings, so the band does not depend on the readings a logging rate would make (equation (2.17)), and a width is reported about the base fit's MAP state, never as a forecast of where the rate would land; the curvature term this neglects is checked on a simulated logged record in step 5.

24.2 The posterior of the state (step 3)

All six questions of Table 24.1 are read from one posterior and its preposteriors: Part III's base fit of the record as it stands, and the covariance the same rows give when every window is logged at n_{L} readings.

24.2.1 The base fit and the preposterior of a logging rate

The posterior of the record as it stands is Chapter 10's base fit, cited: the MAP state $\hat{\mathbf{z}}$ and the Laplace covariance $\boldsymbol{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1}$ of Section 10.2.2, with $\mathbf{J}$ the 261×115 Jacobian at $\hat{\mathbf{z}}$ and $2\mathcal{L}(\hat{\mathbf{z}}) = 154.9755$ on 146 degrees of freedom. Its rates are $\hat{\gamma}_1 = -0.04240$ with standard deviation 0.01463 and $\hat{\gamma}_2 = -0.06253$ with $0.01323\,\mathrm{yr}^{-1}$, and its drift rate $\hat{\hat{b}}_{d,1} = 0.6837$ with $0.1681\,\mathrm{K/yr}$. The normal equations and the iterations are that chapter's and are not repeated. The information of the normalized rows is $\mathbf{J}^{\mathsf{T}}\mathbf{J}$; with unnormalized rows and their weights it is $\mathbf{J}^{\mathsf{T}}\mathbf{W}\mathbf{J}$.

A logging rate adds no entry, only weight on the reading rows. With the logged record's Jacobian (24.8) at the base fit's MAP state, its covariance is

$$\boldsymbol{\Sigma}_{n_{\mathrm{L}}} = \left(\mathbf{J}_o^{\mathsf{T}}\mathbf{J}_o + n_{\mathrm{eff}}(n_{\mathrm{L}})\,\mathbf{J}_y^{\mathsf{T}}\mathbf{J}_y\right)^{-1} = \left(\boldsymbol{\Sigma}^{-1} + \left(n_{\mathrm{eff}}(n_{\mathrm{L}}) - 1\right)\mathbf{J}_y^{\mathsf{T}}\mathbf{J}_y\right)^{-1}, \tag{24.10}$$

the second form because $\boldsymbol{\Sigma}^{-1} = \mathbf{J}_o^{\mathsf{T}}\mathbf{J}_o + \mathbf{J}_y^{\mathsf{T}}\mathbf{J}_y$. It is equation (2.17) with the candidate rows $\mathbf{J}_a = \sqrt{n_{\mathrm{eff}} - 1}\,\mathbf{J}_y$: a logged record is the record as it stands plus $n_{\mathrm{eff}} - 1$ more of every reading row, the same windows read better. At $n_{\mathrm{L}} = 1$ it is $\boldsymbol{\Sigma}$. The same rows in the limit $\sigma_q \to 0$, the straight line of equation (4.74), give $\boldsymbol{\Sigma}_{n_{\mathrm{L}}}^{\mathrm{line}}$, whose base fit has $\hat{\gamma}_1 = -0.03861$ with standard deviation 0.00970 and $\hat{\gamma}_2 = -0.06637$ with $0.00777\,\mathrm{yr}^{-1}$: the spread the readings' errors alone give a path that has no wander.

A band independent of the readings. The rows are nonlinear in the state, through $\mathrm{e}^{g_{u,t}}$ in the physics, but every reading enters its row linearly. At a fixed linearization point $\boldsymbol{\Sigma}_{n_{\mathrm{L}}}$ holds no reading, so it is known before any reading is logged: this is what makes it a preposterior, and it is why only the width is reported. The center of a logged record's band would be its own fit, which the base fit's MAP state does not forecast; the chapter holds the center at the base fit's MAP state only to state a width in the reader's unit, %/yr. A logged record would also move the linearization point, and its curvature term would enter its covariance; step 5 checks both on a simulated logged record.

24.2.2 The rows of a logged window

Table 24.2 shows the reading lines of one window, window 25 (Jan Y3), in five of its columns. Each pair of rows is one channel: its line in the base fit, logged monthly, and the collapsed line of the same channel logged daily, $\sqrt{n_{\mathrm{eff}}(30)} = \sqrt{25.7153} = 5.0710$ times longer. The window's other columns are zero in every reading row, by the pattern (10.6).

Table 24.2: The Jacobian lines of window 25's five reading rows, monthly (the base fit) and collapsed from a daily log, in the columns of Q_{25}, $g_{1,25}$, $g_{2,25}$, β_1 and $\dot{b}_{d,1}$; the daily line is the monthly one times $\sqrt{25.7153} = 5.0710$.

Reading	Logging	Q_{25} [kW^{-1}]	$g_{1,25}$ [–]	$g_{2,25}$ [–]	β_1 [–]	$\dot{b}_{d,1}$ [yr/K]
return air y_{T_r}	monthly	0.2682	−19.8850	−18.3188	0	0
	daily	1.3603	−100.8375	−92.8953	0	0
discharge of unit 1 $y_{T_{d,1}}$	monthly	0.0947	−31.7336	−6.4702	0	1.6667
	daily	0.4804	−160.9220	−32.8107	0	8.4517
discharge of unit 2 $y_{T_{d,2}}$	monthly	0.1084	−8.0364	−30.1674	0	0
	daily	0.5497	−40.7529	−152.9798	0	0
load meter y_Q	monthly	0.3333	0	0	0	0
	daily	1.6903	0	0	0	0
rack-1 inlet y_{T_1}	monthly	0.1370	−19.8850	−18.3188	23.7371	0
	daily	0.6949	−100.8375	−92.8953	120.3714	0

Every entry of a reading line scales by the same factor, so the lines keep their directions: the discharge of unit 1 still sees coil 1 far more strongly than coil 2 (−31.7336 against −6.4702 monthly) and the drift in its own column, and the return air still sees the two coils nearly alike. What changes is the lines' length against the trend rows', whose ±200 and −16.6667 do not move. In information the readings of a window weigh 25.7153 times more against the same trend rows, and that, not a new direction, is all a logging rate does to the fit.

24.2.3 The covariance at three logging rates

Table 24.3 shows $\mathbf{\Sigma}_{n_L}$ of equation (24.10) as numbers, on the seven entries the questions use: the two rates, the drift rate, today's log-derates and the commissioning log-derates. Its first rows are the standard deviations at monthly logging (the base fit), daily and hourly; its last rows the correlations of the pairs the answers depend on.

Table 24.3: The standard deviations and correlations of seven entries of the stacked state under $\mathbf{\Sigma}_{n_L}$, equation (24.10), for windows 0–36 logged monthly (the base fit), daily and hourly.

Quantity [unit]	Monthly	Daily	Hourly
sd(γ_1), sd(γ_2) [yr^{-1}]	0.01463, 0.01323	0.01042, 0.01034	0.01030, 0.01025
sd($\dot{b}_{d,1}$) [K/yr]	0.16812	0.03717	0.03039
sd($g_{1,36}$), sd($g_{2,36}$) [–]	0.02386, 0.01690	0.00680, 0.00547	0.00574, 0.00470
sd($g_{1,0}$), sd($g_{2,0}$) [–]	0.01543, 0.01558	0.00534, 0.00541	0.00465, 0.00471
corr(γ_1, γ_2) [–]	−0.2676	−0.0328	−0.0245
corr($\dot{b}_{d,1}, \gamma_1$), corr($\dot{b}_{d,1}, \gamma_2$) [–]	+0.5176, −0.2830	+0.1585, −0.0804	+0.1309, −0.0662
corr($g_{1,0}, \gamma_1$), corr($g_{2,0}, \gamma_2$) [–]	−0.5128, −0.5115	−0.1767, −0.1826	−0.1537, −0.1588
corr($g_{1,36}, \gamma_1$) [–]	+0.6479	+0.2223	+0.1881
corr($\dot{b}_{d,1}, g_{1,36}$), corr($\dot{b}_{d,1}, g_{2,36}$) [–]	+0.8104, −0.4894	+0.7200, −0.4103	+0.7033, −0.3950
corr($g_{1,36}, g_{2,36}$), corr($g_{1,0}, g_{2,0}$) [–]	−0.5823, −0.3578	−0.5319, −0.3269	−0.5317, −0.3399

The correlations that fade. Every entry narrows, the log-derates and the drift most: today's two log-derates from 0.02386 and 0.01690 at monthly logging to 0.00574 and 0.00470 at hourly, and the drift rate from 0.16812 to 0.03039 K/yr. The rates narrow least, and their correlations with everything else fade. The reason is the part of a rate's variance that is the wander floor of equation (10.12): it comes from the trend rows alone and is correlated with no other entry, and

as the readings sharpen it becomes nearly all of the rate's variance. The signs are those of the base fit, and each is fixed by a row. γ_1 against γ_2 is negative, -0.2676, because the return air sees the sum of the two coils at both ends of the record, so a steeper loss of one is paid for by a gentler loss of the other; at daily logging it is -0.0328. A commissioning entry against its coil's rate is negative, -0.5128 and -0.5115, because by equation (10.7) the rate is the path's end-to-end change, so a higher start at the same end is a steeper loss; at daily logging it is -0.1767 and -0.1826. The drift rate against coil 1's rate is positive, $+0.5176$: a larger drift explains more of unit 1's warming discharge reading as the sensor and leaves less of it to the coil, a gentler loss; at daily logging it is $+0.1585$, and the drift no longer trades with coil 1's slope.

The correlations that stay. Two correlations within window 36 barely move. The drift rate against coil 1's log-derate today goes from $+0.8104$ to $+0.7200$ at daily and $+0.7033$ at hourly, and coil 1's log-derate against coil 2's today from -0.5823 to -0.5319 and -0.5317. These are the within-window ridges of Chapters 5 and 7: unit 1's discharge reading sees the drift and coil 1 through one line, a warmer reading being either a larger drift or a lower conductance, and the return air sees the coils' sum. A logged reading of the window has the same line as the monthly one (Table 24.2), so logging narrows both entries of each pair and keeps the ridge's direction. Only a row with another line, such as a second sensor at the place of the question, turns it; step 7 takes that up.

24.2.4 The posterior of simulated logged records

No logged record exists, so the full posterior of one cannot be computed by reweighted draws as Chapter 5 computed the record's: there is no record to weight them by. Its stand-in is the posterior of simulated logged records at the stated truth of the record, set beside every answer of step 4 as a spread beside a spread. Two routes are used, each named as such and neither called exact.

Simulated windows. For the effective number, 20 000 windows of hourly AR(1) errors of unit variance are drawn from a seeded generator (seed 24), and each window's generalized least-squares mean, equation (24.6), and its simple mean are formed, at hourly and at daily logging. The effective number is the reciprocal of the means' variance, with its standard error from the spread of a variance over 20 000 draws.

Simulated records. For the spreads of the state, the procedure is the following.

1. Draw 100 records at the stated truth for each of three rates, monthly, daily and hourly (seeds 24100, 24200 and 24300). Each window's five channels are logged n_{L} times with AR(1) errors at the rate's interval, from the record's noise-free readings, and each reading is rounded to 0.01 as the record is.
2. Collapse each channel's n_{L} readings in each window to the generalized least-squares mean of equation (24.6); at hourly logging, form also the simple mean, for the treatment of the readings as independent.
3. Fit each record by the base drawing with every reading row's standard deviation $\sigma_k/\sqrt{n_{\mathrm{eff}}}$ (with $\sigma_k/\sqrt{720}$ for the simple means treated as independent), by Gauss–Newton on the normal equations of Section 10.2.2, and read its MAP state and Laplace covariance.
4. Report the mean posterior standard deviation over the 100 records beside each preposterior spread, the scatter of the 100 estimates, and the share of records whose 5–95 % band covers the drawing's pseudo-true value γ_u°, the base drawing fitted to the truth's noise-free readings at the same rate's weights. The Monte Carlo standard error of a coverage near 0.9 over 100 records is 0.030.

The result is called the posterior of a simulated daily (or monthly, or hourly) record. The plant's own record draws independent reading errors; the AR(1) is this chapter's stated input, and the simulated records carry it.

24.3 The answers (step 4)

Every answer is a spread read from a preposterior of step 3, $\mathbf{\Sigma}_{n_\mathrm{L}}$ of equation (24.10) or its extension to window 45, and that value is the one reported. Beside it stands the mean posterior standard deviation of the 100 simulated records logged at that rate (Section 24.2.4), a spread beside a spread. Rates are read as $100(\mathrm{e}^{\gamma}-1)\,\%/\mathrm{yr}$, equation (10.1), and bands are 5–95 %.

24.3.1 The rate's band at each logging rate

Asset manager's question. Hourly, daily, or monthly? It is asked first of each coil's degradation rate, the quantity the proposal promises to sharpen.

General form. A design question asks for the expected value of a quantity of the posterior a candidate would produce, equation (2.16); here the quantity is the spread of a derived quantity f after the outcome $\tilde{y}_a$,

$$U(a) = \int \mathrm{sd}\big(f(\mathbf{z}) \mid \mathbf{y}, \tilde{y}_a\big)\, p(\tilde{y}_a \mid \mathbf{y})\, \mathrm{d}\tilde{y}_a,$$

and the candidates are ranked by it, the smallest first. For a Gaussian posterior and linear Gaussian candidate rows the spread is $\sqrt{\mathbf{g}^\mathsf{T}\mathbf{\Sigma}_a\mathbf{g}}$ with $\mathbf{\Sigma}_a$ of equation (2.17), the same for every outcome.

Specialized form. The question is equation (2.16) with the candidate a the record of windows 0–36 logged at n_L readings a window, $\tilde{y}_a$ its logged readings, $f = \gamma_u$ with the gradient $\mathbf{g}_{\gamma_u} = \mathbf{e}_{\gamma_u}$ of equation (24.9), and $\mathbf{\Sigma}_a = \mathbf{\Sigma}_{n_\mathrm{L}}$ of equation (24.10). The spread does not depend on $\tilde{y}_a$, so the expectation is the spread itself,

$$U(n_\mathrm{L}) = \sqrt{[\mathbf{\Sigma}_{n_\mathrm{L}}]_{\gamma_u\gamma_u}}, \qquad \Big[100\big(\mathrm{e}^{\hat{\gamma}_u - 1.6449\,U(n_\mathrm{L})} - 1\big),\ 100\big(\mathrm{e}^{\hat{\gamma}_u + 1.6449\,U(n_\mathrm{L})} - 1\big)\Big]\ \%/\mathrm{yr}, \quad (24.11)$$

the band in the reader's unit about the base fit's $\hat{\gamma}_u$, whose width is what the logging changes. One choice enters here: *the archive counterfactual*, the same 37 windows as if they had been logged at n_L readings from the start. This buys a comparison of the logging rates at the record's own length, against Chapter 10's floor; it is not a plan for the logger, which can only change the windows still to come, Section 24.3.6. The floor is unchanged by the logging. By equation (10.13), which holds for $\mathbf{\Sigma}_{n_\mathrm{L}}$ as for $\mathbf{\Sigma}$ because no reading row sees γ_u and the trend rows are the base fit's,

$$U(n_\mathrm{L}) \geq \frac{\sigma_q}{\sqrt{36}\,\Delta t},$$

with equality only for every window's log-derates known exactly, the limit $n_\mathrm{eff} \to \infty$.

Building blocks. The base fit's $\mathbf{\Sigma}$ and its reading and other rows $\mathbf{J}_y$, $\mathbf{J}_o$ (Section 24.2.1); the effective number (24.5) at $\varrho = \varrho_1^{\Delta\vartheta/1\,\mathrm{h}}$; the preposterior (24.10); the band (24.11) with the map (10.1); the floor (10.12).

Table 24.4: The degradation rate's spread at each logging rate, windows 0–36 logged at that rate, by the preposterior (24.10). The band's width is that of equation (24.11) about the base fit's MAP rates.

Rate	n_{L} [–]	n_{eff} [–]	sd γ_1 [yr^{-1}]	width 1 [%/yr]	sd γ_2 [yr^{-1}]	width 2 [%/yr]	sd $\dot{b}_{d,1}$ [K/yr]
monthly (as now)	1	1.0000	0.01463	4.613	0.01323	4.089	0.1681
weekly	4	4.0000	0.01168	3.684	0.01122	3.466	0.0895
daily	30	25.7153	0.01042	3.285	0.01034	3.195	0.0372
every 6 hours	120	37.4092	0.01031	3.252	0.01026	3.170	0.0310
hourly	720	38.8421	0.01030	3.249	0.01025	3.168	0.0304
every 15 minutes	2880	38.9144	0.01030	3.249	0.01025	3.168	0.0304
every minute	43 200	38.9289	0.01030	3.249	0.01025	3.168	0.0304
readings exact	–	–	0.01000	–	0.01000	–	–

The rate at each logging rate. Table 24.4 gives the spreads. Each row is a logging rate, from the record as it stands to one reading a minute, with its readings per window n_{L} and effective number; the columns give each coil's standard deviation of the rate, the width of its band in %/yr, and the drift rate's standard deviation. The last row is the limit of readings known exactly.

Daily logging takes 28.8 % off coil 1's rate standard deviation and 21.9 % off coil 2's, which is 91.0 % and 89.5 % of all that any logging can take, down to the floor of $0.01000\,\mathrm{yr}^{-1}$. Hourly takes a further 1.10 % and 0.84 % off daily, and every 15 minutes or every minute changes nothing in the printed digits. About the base fit's MAP rates, coil 1's band goes from $[-6.430,\ -1.817]$ to $[-5.780,\ -2.495]$ %/yr at daily logging, and coil 2's from $[-8.084,\ -3.995]$ to $[-7.646,\ -4.451]$: widths about today's MAP state, whose centers a logged record would move, as step 5 shows. The drift rate's standard deviation falls from 0.1681 to 0.0372 K/yr at daily logging: given the onset at window 20, the logged record decides the drift's size without the side-by-side check, because seventeen windows of well-read readings tell unit 1's discharge from coil 1 through the other channels, Chapter 7's principle. The onset itself the logging does not place, and the size depends on it: over the onsets Chapter 10 tried, the drift rate runs from 0.4738 to 1.5370 K/yr (Section 10.4.4). Beside the reported values, the mean posterior standard deviation over 100 simulated daily records is 0.01044 and 0.01034, and over hourly records 0.01032 and 0.01025.

The rate's spread is plotted against the logging interval in Figure 24.2. The horizontal axis is the interval in hours on a logarithmic scale, from a month at the left to a minute at the right; the vertical axis is the standard deviation of the rate. The solid curves are each coil's rate in the base drawing with its wander, the dashed curves the straight line's, the spread the readings' errors alone give a path without wander, and the dotted line the wander floor; monthly, daily and hourly are marked. The solid curves fall steeply to about a day and are flat from there; the dashed ones keep falling to a few hours and then flatten too.

A log 720 times denser than today's narrows the rate's band by less than a third, and daily logging already does nine-tenths of what any logging can. The rest of the band, between the solid curves and the floor, is small by a day; below them the floor is the path's own wander.

24.3.2 The effective number of readings in a month

Asset manager's question. How many independent readings does a month of logging hold? The question asks what the proposal's 720 readings are worth, before the rate.

General form. A block of readings of one quantity with a full covariance adds to the posterior's information what Chapter 3's precision (3.29) gives it, with the readings' covariance

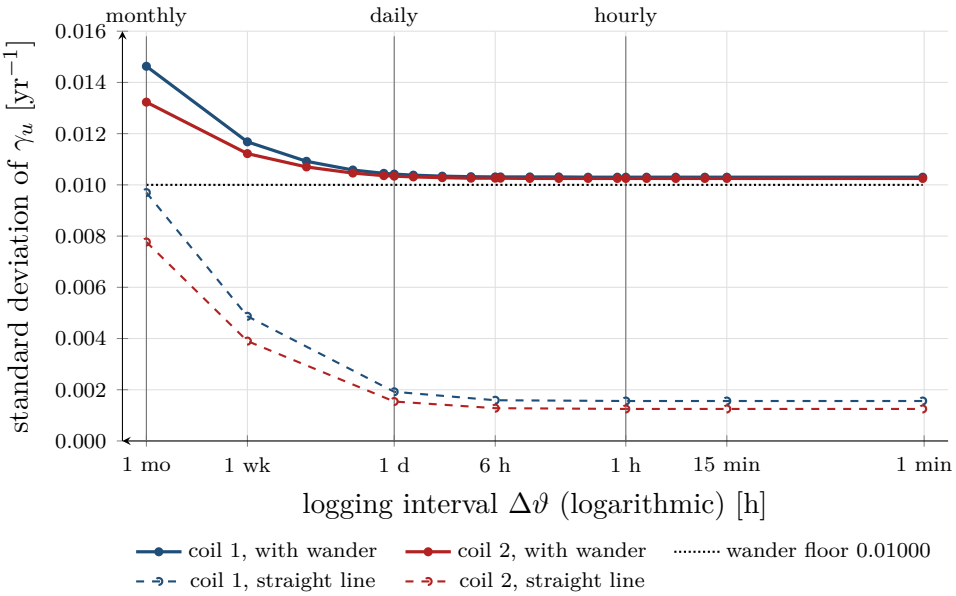


Figure 24.2: The standard deviation of each coil's degradation rate against the logging interval, windows 0–36 logged at that interval, by the preposterior (24.10). Solid: the base drawing with its wander. Dashed: the straight line, $\sigma_q \to 0$, the readings' errors alone. Dotted: the wander floor $\sigma_q/(\sqrt{36}\,\Delta t)$ of equation (10.12). Thin gray lines: monthly, daily and hourly logging.

written as a standard deviation and a correlation matrix,

$$\boldsymbol{\Sigma}^{-1} = \boldsymbol{\Sigma}_0^{-1} + \mathbf{A}^{\mathsf{T}}(\sigma^2\mathbf{R})^{-1}\mathbf{A}, \tag{24.12}$$

with $\mathbf{A}$ the map from the state to the quantities read; the design form (2.16) ranks candidates by what it leaves.

Specialized form. The question is equation (24.12) for one channel k of one window, its n_{L} logged readings of the one value $f_k(\mathbf{z}_t)$: $\mathbf{A}$ is the column of ones times the line $\mathbf{j}_{k,t}^{\mathsf{T}}\sigma_k$, $\sigma = \sigma_k$, and $\mathbf{R}$ is the AR(1) correlation of equation (24.2) at $\varrho = \varrho_1^{\Delta\vartheta/1\,\mathrm{h}}$. The block adds $(\mathbf{1}^{\mathsf{T}}\mathbf{R}^{-1}\mathbf{1})\,\mathbf{j}_{k,t}\mathbf{j}_{k,t}^{\mathsf{T}} = n_{\mathrm{eff}}\,\mathbf{j}_{k,t}\mathbf{j}_{k,t}^{\mathsf{T}}$, equation (24.5), and the readings alone read the window's value to $\sigma_k/\sqrt{n_{\mathrm{eff}}}$. The answer is read from the state as every answer is: the value the channel reads in window t is the derived quantity $f_k(\mathbf{z}_t)$, with gradient $\mathbf{g}_{f_k} = \sigma_k\,\mathbf{j}_{k,t}$, and its spread under the logged record is $\sqrt{\mathbf{g}_{f_k}^{\mathsf{T}}\boldsymbol{\Sigma}_{n_{\mathrm{L}}}\mathbf{g}_{f_k}}$ by equation (24.10), narrower than the readings' own because the window's other channels, its neighbors through the trend rows and the physics read the same value. For comparison, a logger that averages its readings with equal weights gives a mean of variance $\sigma_k^2\,\mathbf{1}^{\mathsf{T}}\mathbf{R}\,\mathbf{1}/n_{\mathrm{L}}^2$, an effective number

$$n_{\mathrm{eff}}^{\mathrm{mean}} = \frac{n_{\mathrm{L}}^2}{\mathbf{1}^{\mathsf{T}}\mathbf{R}\,\mathbf{1}} = \frac{n_{\mathrm{L}}^2}{n_{\mathrm{L}} + 2\sum_{i=1}^{n_{\mathrm{L}}-1}(n_{\mathrm{L}}-i)\,\varrho^i}, \tag{24.13}$$

provided the fit is told that variance. One choice enters here: *the readings equally spaced over the whole span.* This buys the closed forms (24.5) and (24.7), and it fails for a logger that records only some hours, such as the working day: the span the readings cover, not their number, then sets the limit.

Building blocks. The correlation $\mathbf{R}$ and its inverse (24.3); the effective number (24.5) and its limit (24.7); the weights of the mean (24.6); the simple average (24.13); the preposterior (24.10) and the gradient of the channel's value.

The effective number at each rate. Table 24.5 gives it. Each row is a logging rate; the columns give the interval, the correlation of consecutive readings' errors at that interval, the effective number of the whitened rows and of the simple average, the readings logged per effective reading, and the standard deviation of a window's return air, $0.3/\sqrt{n_{\text{eff}}}$.

Table 24.5: The effective number of readings in one window of 720 h at each logging rate, AR(1) errors at $\varrho_1 = 0.9$ an hour apart: equation (24.5) for the whitened rows, equation (24.13) for the simple average.

Rate	interval [h]	ϱ [–]	n_{eff}, whitened [–]	$n_{\text{eff}}^{\text{mean}}$, average [–]	$n_{\text{L}}/n_{\text{eff}}$ [–]	sd of a window's return air [K]
monthly	720	0	1.0000	1.0000	1.00	0.30000
weekly	180	0	4.0000	4.0000	1.00	0.15000
daily	24	0.07977	25.7153	25.7051	1.17	0.05916
every 6 hours	6	0.53144	37.4092	37.1740	3.21	0.04905
hourly	1	0.90000	38.8421	38.4000	18.54	0.04814
every 15 minutes	0.25	0.97400	38.9144	38.4342	74.01	0.04809
every minute	0.017	0.99825	38.9289	38.4365	1109.72	0.04808

The limit as the interval goes to zero is, by equation (24.7), $1 + 720\ln(1/0.9)/2 = 38.9298$: a window holds one effective reading per two correlation times of 9.49 h, 0.05268 per hour of its span, and hourly logging already gives 38.8421 of the limit's 38.9298. At hourly logging the whitened mean weights the two end readings by 1.0 and each of the 718 interior readings by 0.10; the weights sum to 73.80, which is $n_{\text{eff}}(1+\varrho)$ by equation (24.5). The simple average loses little, 38.4000 against 38.8421 at hourly logging: a logger that averages over the window is nearly as good as the whitened rows, provided the fit is told the average's true standard deviation, $0.3/\sqrt{38.4}$ and not $0.3/\sqrt{720}$ on a temperature. Read from the state, window 36's return air has the spread 0.18209 K under $\mathbf{\Sigma}_{n_{\text{L}}}$ at monthly logging, 0.03686 at daily and 0.03007 at hourly, against the readings' own 0.30000, 0.05916 and 0.04814 of Table 24.5; its load has 2.2332, 0.5048 and 0.4182 kW against 3.0000, 0.5916 and 0.4814. The state reads each value better than its own readings alone do, and the gain from logging carries through to it. Chapter 10's archive hour is re-valued by the same form: two adjacent hours are worth $2/(1+0.9) = 1.0526$ readings, and two hours 12 h apart $2/(1+0.2824) = 1.5595$, where that chapter counted two. Beside the closed forms, the 20 000 simulated windows of Section 24.2.4 give 38.642 (standard error 0.386) for the whitened rows at hourly logging and 38.321 (0.383) for the simple average, and 25.653 and 25.654 (both 0.257) at daily logging, each within one standard error of its closed form.

The effective number is plotted in Figure 24.3. The horizontal axis is the logging interval, as in Figure 24.2; the vertical axis is the effective number of readings in one window. The solid curve is the whitened rows', the dashed curve with open marks the simple average's, the gray dashed curve the count n_{L} the proposal uses, and the dotted line the limit. The count leaves the top of the axis near half a day, while both effective numbers turn flat under the limit.

24.3.3 Where the rate's band stops narrowing

Asset manager's question. From which rate on does logging stop narrowing the rate's band, and why?

General form. The design form, equation (2.16), with the spread of the rate as its quantity, read in two drawings of the same rows: the base drawing with its wander, and the straight line of equation (4.74), in which only the readings' errors leave the rate unknown.

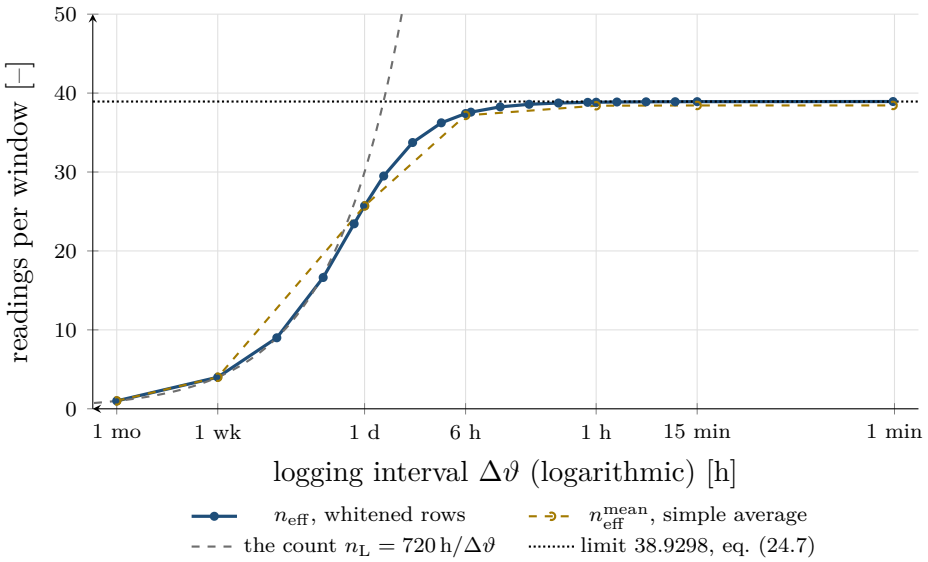


Figure 24.3: The effective number of readings in one window of 720 h against the logging interval, AR(1) errors with $\varrho_1 = 0.9$ an hour apart. Solid: the whitened rows, equation (24.5). Dashed with open marks: the simple average, equation (24.13). Gray dashed: the count of readings n_{L}, the effective number the readings would have if independent; it leaves the axis at 50. Dotted: the limit (24.7).

Specialized form. The question is equation (2.16) with $f = \gamma_u$ and the two preposteriors $\boldsymbol{\Sigma}_{n_{\text{L}}}$ and $\boldsymbol{\Sigma}_{n_{\text{L}}}^{\text{line}}$ of Section 24.2.1, combined into the readings' share,

$$\frac{[\boldsymbol{\Sigma}_{n_{\text{L}}}^{\text{line}}]_{\gamma_u\gamma_u}}{[\boldsymbol{\Sigma}_{n_{\text{L}}}]_{\gamma_u\gamma_u}}, \tag{24.14}$$

read over a fine grid of rates for the first n_{L} at which it falls to 10 % and to 5 %. One choice enters here: *the straight line as the readings' part*. It buys a measure, at every rate, of how much of the band the readings' errors would leave on a path without wander, and it is not a split of the variance: the two drawings' variances do not add, as Chapter 10 noted beside its budget (10.13). The reason the band stops is that budget: the rate's variance is the floor's, $\sigma_q^2/(36\,\Delta t^2)$, plus the variance of the path's two ends over three years, and only the second term is a reading's to narrow. That second term over the rate's variance is a second measure, the path-ends share, and it is a split: the two terms add. The two measures differ because, in the drawing with wander, each end of the path is read mostly by its own window and its neighbors, where the straight line reads it from every window. The thresholds of the question are read on the readings' share (24.14), and the path-ends share is set beside them.

Building blocks. The two preposteriors (24.10) on the base drawing and on the straight line; the share (24.14); the budget (10.13) and its path-ends share; the floor (10.12).

The share and its thresholds. The readings' share of coil 1's rate variance is 44.0 % at monthly logging, 17.4 % at weekly, 3.4 % at daily and 2.3 % at hourly; coil 2's is 34.4, 12.1, 2.2 and 1.5 %. At the saturation limit, $n_{\text{eff}} = 38.93$, it is 2.30 and 1.49 %. The share first falls to 10 % at $n_{\text{L}} = 8$ readings a window for coil 1, a reading every 90.0 h with $n_{\text{eff}} = 8.00$ and a rate standard deviation of 0.01101, and at $n_{\text{L}} = 5$ for coil 2, every 144.0 h, with 0.01104. It falls to 5 % at $n_{\text{L}} = 18$ for coil 1, every 40.0 h ($n_{\text{eff}} = 17.50$, 0.01056), and at $n_{\text{L}} = 11$ for coil 2, every 65.5 h (10.98, 0.01061). With the readings exact, the rate's standard deviation

is $0.01000 = \sigma_q/(\sqrt{36}\,\Delta t)$ for both coils, Chapter 10's wander floor. The dashed curves of Figure 24.2 are the numerator of the share.

The path-ends share of the budget (10.13) is larger at every rate: 53.3 and 42.9 % at monthly logging, 7.9 and 6.4 % at daily, 5.8 and 4.9 % at hourly, and 5.8 and 4.8 % at the saturation limit. On it the share first falls to 10 % at $n_{\mathrm{L}} = 19$ for coil 1, every 37.9 h, and at $n_{\mathrm{L}} = 14$ for coil 2, every 51.4 h; it falls to 5 % for coil 2 at $n_{\mathrm{L}} = 117$, every 6.2 h, and for coil 1 never, at any rate.

A reading every 90.0 h, about twice a week, brings coil 1's readings' share under a tenth, and a reading every 37.9 h its path-ends share; once the logging is daily, the rate's band is the path's month-to-month wander, set by σ_q, not by the sensors, on either measure. What the readings leave at saturation is 2.3 and 1.5 % of the rate's variance on the straight line's measure and 5.8 and 4.8 % on the budget's, what the window's 39 effective readings cannot remove from the path's two ends. The floor is the chosen wander's, and step 5 moves it with σ_q.

24.3.4 Today's health index and the capacity margin

Asset manager's question. Does a faster rate help today's health index and the capacity margin?

General form. The design form, equation (2.16), with the spread of a derived quantity after the outcome as its quantity.

Specialized form. The question is equation (2.16) with $\boldsymbol{\Sigma}_a = \boldsymbol{\Sigma}_{n_{\mathrm{L}}}$ and two quantities of today's window. The first is the health index against commissioning, $\ln h_{u,36} = g_{u,36} - g_{u,0}$, with the gradient $\mathbf{g}_{\ln h_u}$ of equation (24.9), the variance of Chapter 10's index (10.17), and the band $\exp(\widehat{\ln h}_{u,36} \pm 1.6449\,\mathrm{sd})$ about the base fit's $\widehat{\ln h}_{1,36} = -0.12719$ and $\widehat{\ln h}_{2,36} = -0.18760$. The second is the hall's total conductance, with the gradient $\mathbf{g}_G$ at the base fit's $G_{1,36} + G_{2,36} = 11.7897\,\mathrm{kW/K}$, read against the capacity study's stated requirement $G^* = 10.761\,\mathrm{kW/K}$ as the margin in standard deviations, $(11.7897 - G^*)/\mathrm{sd}$. Each spread is $\sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_{n_{\mathrm{L}}}\mathbf{g}}$. The two quantities have no floor. By the budget (10.13), which holds for $\boldsymbol{\Sigma}_{n_{\mathrm{L}}}$ for the reason given in Section 24.3.1,

$$[\boldsymbol{\Sigma}_{n_{\mathrm{L}}}]_{\gamma_u\gamma_u} = \frac{\mathbf{g}_{\ln h_u}^{\mathsf{T}}\boldsymbol{\Sigma}_{n_{\mathrm{L}}}\mathbf{g}_{\ln h_u}}{(3\,\mathrm{yr})^2} + \frac{\sigma_q^2}{36\,\Delta t^2}\,: \tag{24.15}$$

the health index's variance is the only part of the rate's that logging touches, and the index itself has no term that logging cannot.

Building blocks. The preposterior (24.10); the gradients (24.9) at the base fit's MAP state; the requirement G^*; the tie (24.15) between the rate and the index.

The level at each logging rate. Table 24.6 gives the spreads at monthly, daily and hourly logging. Each row is a rate; the columns give each coil's standard deviation of $\ln h_{u,36}$ and its index's band, the standard deviation of $G_{1,36} + G_{2,36}$, and the margin to G^* in standard deviations.

Daily logging takes 72.6 % and 69.7 % off the standard deviation of the two health indices and 69.5 % off the capacity margin's. Hourly logging takes a further 15.0 %, 13.9 % and 15.3 % off daily, where the rate gained 1.10 % and 0.84 %. The tie (24.15) carries the difference: at hourly logging coil 1's index standard deviation, 0.00745, divided by three years and added to the floor in quadrature, gives back the rate's 0.01030 of Table 24.4. The index's term, its variance over $(3\,\mathrm{yr})^2$, is already small beside the floor's at daily logging, so the rate stops gaining there while the index does not. Beside the reported values, the mean posterior standard

Table 24.6: Today's health index against commissioning and the capacity margin at three logging rates, windows 0–36 logged at that rate, by the preposterior (24.10); bands about the base fit's MAP state.

Rate	sd $\ln h_{1,36}$ [–]	$h_{1,36}$ band [–]	sd $\ln h_{2,36}$ [–]	$h_{2,36}$ band [–]	sd $G_1 + G_2$ [kW/K]	margin [sd]
monthly	0.03203	0.835 to 0.928	0.02599	0.794 to 0.865	0.1193	8.6
daily	0.00876	0.868 to 0.893	0.00787	0.818 to 0.840	0.0364	28.3
hourly	0.00745	0.870 to 0.891	0.00678	0.820 to 0.838	0.0308	33.4

deviation of $\ln h_{1,36}$ over the simulated records is 0.03233 at monthly, 0.00899 at daily and 0.00765 at hourly logging.

The three relative spreads are set on one axis in Figure 24.4. The horizontal axis is the logging interval, as before; the vertical axis is each quantity's standard deviation over its value at monthly logging. Coil 1's rate (blue), its health index today (green) and the capacity margin (gold) start together at one; the rate's curve flattens at about a day and ends 29.6 % below one, and the other two keep falling to about an hour, where they stand 76.7 % and 74.2 % below one at hourly logging, and barely move from there to the minute.

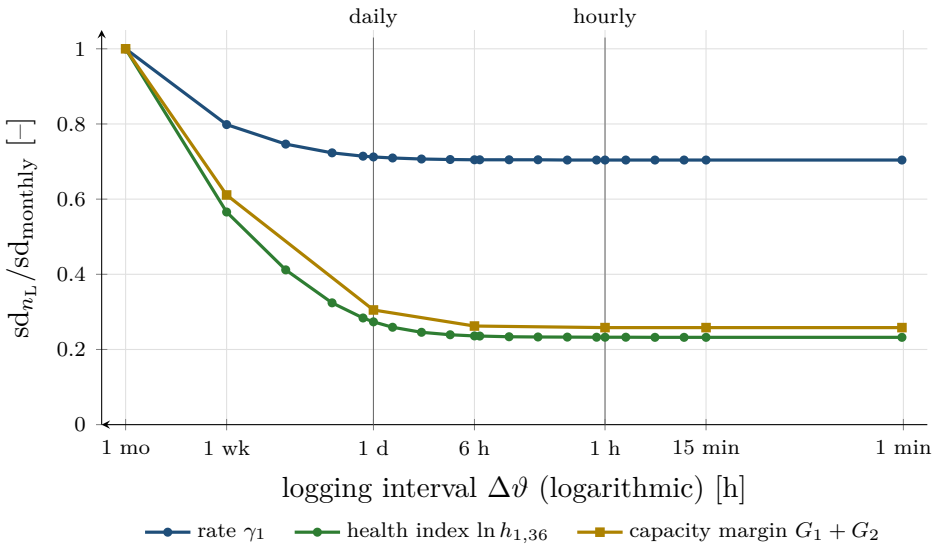


Figure 24.4: The standard deviations of coil 1's degradation rate, its health index today and the hall's total conductance today, each over its value at monthly logging, against the logging interval; windows 0–36 logged at that interval, by the preposterior (24.10). Coil 2's curves lie close to coil 1's (Tables 24.4 and 24.6).

The level has no wander floor. A path known exactly at every window knows its level exactly, so logging keeps paying for the level until the effective number saturates, at hourly logging, while it stops paying for the rate at daily. The crossing date of Chapter 14 rests on both the level and the rate, and that is why hourly logging is worth more than daily to the asset manager, though not to the rate.

24.3.5 The logged readings treated as independent

Asset manager's question. What does the band look like if the logged readings are treated as independent? It is the proposal's own arithmetic, carried through the fit.

General form. The design form, equation (2.16), with the spread as its quantity, under the precision (24.12) of a block of readings with its correlation matrix taken as the identity, against the same block with $\mathbf{R}$ full.

Specialized form. The question is equation (2.16) with $\mathbf{\Sigma}_a$ the preposterior (24.10) at $n_{\text{eff}} = n_{\text{L}}$, the independent count, set beside the same quantity at the effective number (24.5). For a quantity the readings alone set, such as the straight line's rate, the two spreads differ by nearly the factor $\sqrt{n_{\text{L}}/n_{\text{eff}}}$, the square root of Table 24.5's readings per effective reading; for a quantity the trend rows also inform, the factor is smaller, because the floor's term of equation (24.15) is the same in both. The three quantities compared are the straight line's rate, the rate in the base drawing, and coil 1's health index today.

Building blocks. The preposterior (24.10) at n_{eff} and at n_{L}; the straight line's preposterior $\mathbf{\Sigma}^{\text{line}}_{n_{\text{L}}}$; the gradients (24.9).

The two treatments. Table 24.7 sets them side by side for coil 1. Each row is a logging rate; each group of three columns is one quantity, its spread with the readings' correlation carried (the rows' own), its spread with the readings treated as independent, and their ratio.

Table 24.7: Coil 1's spreads with the logged readings' correlation carried and with the readings treated as independent ($n_{\text{eff}} = n_{\text{L}}$), windows 0–36; the ratio is the first over the second. Coil 2's are alike.

	straight line's rate [yr^{-1}]			rate with wander [yr^{-1}]			$\ln h_{1,36}$ [–]		
Rate	carried	independent	ratio [–]	carried	independent	ratio [–]	carried	independent	ratio [–]
daily	0.00192	0.00178	1.08	0.01042	0.01037	1.005	0.00876	0.00825	1.06
every 6 hours	0.00159	0.00089	1.79	0.01031	0.01012	1.019	0.00756	0.00470	1.61
hourly	0.00156	0.00036	4.30	0.01030	0.01002	1.028	0.00745	0.00211	3.54
every 15 minutes	0.00156	0.00018	8.60	0.01030	0.01001	1.030	0.00744	0.00108	6.90
every minute	0.00156	0.00005	33.31	0.01030	0.01000	1.030	0.00744	0.00028	26.49

The proposal's ±0.058 %/yr is the straight line with hourly readings treated as independent. Carried with their correlation, the straight line gives ±0.248 %/yr, and the record's drawing with its wander ±1.625 (±1.581 with the readings treated as independent). The trend rows shield the rate: with the wander in the drawing, the rate's standard deviation with the correlation carried is only 1.028 times the independent one at hourly logging. The level is not shielded. The health index's band treated as independent is 3.54 times too narrow at hourly logging (coil 2: 0.00201 against 0.00678, 3.36), and 26.49 times at one reading a minute. Beside the reported values, over the 100 simulated hourly records the 5–95 % band on $\ln h_{1,36}$ from the readings treated as independent covers the drawing's pseudo-true value in 0.24 of the records (coil 2: 0.25), against 0.97 for the band with the correlation carried; the hourly simple averages scatter 4.3625 times the standard deviation the independent treatment claims for them.

The ratios are plotted in Figure 24.5, against the interval as before, on a logarithmic vertical axis: one marks no error. The straight line's ratio (dashed) and the health index's (green) grow without bound as the interval shrinks; the rate's (blue) stays near one.

At a daily interval the readings are nearly independent, $\varrho = 0.07977$, and the independent treatment errs by 6 % on the level. At hourly logging and faster, independence is the error the proposal makes, and it falls hardest on the quantity logging helps most.

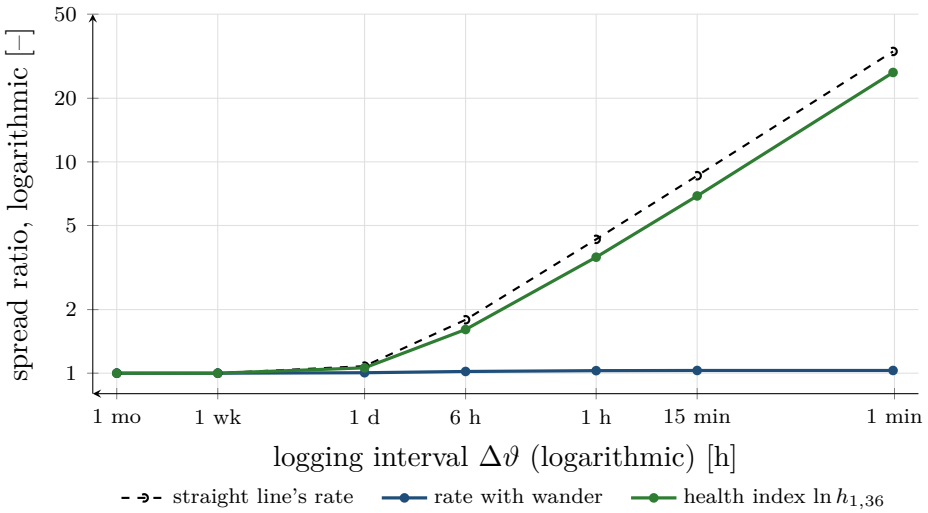


Figure 24.5: How much too narrow coil 1's spreads are when the logged readings are treated as independent: the spread with the correlation carried over the spread with $n_{\text{eff}} = n_{\text{L}}$, against the logging interval (Table 24.7; monthly and weekly readings are independent at $\varrho = 0$, ratio one). Dashed: the straight line's rate. Blue: the rate with the wander. Green: the health index today.

24.3.6 The logger's setting for year 4

Asset manager's question. What should the logger be set to for year 4, against a crew visit at the same price?

General form. The design form, equation (2.16), with the spread of a quantity after the outcome as its quantity and a candidate's cost entered as a loss of its own beside it.

Specialized form. The question is equation (2.16) on Chapter 10's extended state, equation (10.18), carried to the refit at window 45: windows 37 to 45 added, each with its load and two log-derates, 142 entries in all, tied to the record and to each other by their trend rows. The state is linearized with the new windows' path on the fitted line and their loads at the record's forecast growth, and a future load carries a spread of 1000 kW until its window is read. With $\mathbf{J}_y$ the record's reading rows, $\mathbf{J}_{y,\text{f}}$ the new windows' reading rows and $\mathbf{J}_o$ the priors and all the trend rows, a candidate's preposterior is

$$\boldsymbol{\Sigma}_a = \left(\mathbf{J}_o^{\mathsf{T}}\mathbf{J}_o + n_{\text{p}}\,\mathbf{J}_y^{\mathsf{T}}\mathbf{J}_y + n_{\text{eff}}(n_{\text{L}})\,\mathbf{J}_{y,\text{f}}^{\mathsf{T}}\mathbf{J}_{y,\text{f}} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a + \boldsymbol{\Lambda}_Q\right)^{-1}, \tag{24.16}$$

with $n_{\text{p}} = 1$ for the record as it stands and $n_{\text{p}} = n_{\text{eff}}$ when the logger's archive of windows 0–36 is read at the same rate, $\mathbf{J}_a$ the lines of a crew visit's rows if one is made, and $\boldsymbol{\Lambda}_Q$ the future loads' prior precision, $(1000\,\text{kW})^{-2}$ on each future load's diagonal entry and zero elsewhere. The quantities are the rates γ_u, the health index at the refit $\ln h_{1,45} = g_{1,45} - g_{1,0}$, and the total conductance $G_{1,45} + G_{2,45}$, each by $\sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_a\mathbf{g}}$. The two crew visits at the same price are Chapter 10's: a reference condition test now, reference sensors of 0.1 K on window 36's four temperatures, and the side-by-side check of unit 1's discharge sensor now, a row on $\dot{b}_{d,1}$ alone by equation (7.27). One choice enters here: *the horizon is the refit at window 45*, the date at which Chapter 11's entry asked for the record to be read again, not the end of year 4. This buys the answer at the date the shutdown's work order uses it. It fails if the refit slips past the shutdown: the work order is then written on the record of window 36, and the table's widths describe a refit that did not inform it; a later refit lowers every row of the table together, the floor with them.

Building blocks. The base fit's information and the extended state's rows; the effective number (24.5); the preposterior (24.16); the gradients at the extended state; the stated costs.

The candidates for year 4. With no new rows the rates' standard deviations stay at 0.01463 and 0.01323 yr^{-1}, and the total conductance at window 45 has a standard deviation of 0.2010 kW/K; the floor at 45 windows is $\sigma_q/(\sqrt{45}\,\Delta t) = 0.00894\,\mathrm{yr}^{-1}$. Table 24.8 ranks the candidates. Each row is one; the columns give each coil's rate spread with the width it takes off the spread with no new rows, the spread of coil 1's health index at the refit, that of the total conductance, and the cost.

Table 24.8: The logger's setting for year 4: windows 37–45 logged at each rate, refit at window 45 (Sep Y4), by the preposterior (24.16); the width taken off is against the extended record with no new rows. Costs are for year 4: the logging's yearly price, the chapter's stated input, and each crew visit made once in the year at R$ 1 500, the plant's cost schedule.

Candidate	sd γ_1 [yr^{-1}]	off [%]	sd γ_2 [yr^{-1}]	off [%]	sd $\ln h_{1,45}$ [–]	sd $G_1 + G_2$ [kW/K]	cost [R$/yr]
year 4 logged monthly (as now)	0.01195	18.3	0.01109	16.2	0.02970	0.1091	0
year 4 logged weekly	0.01070	26.8	0.01034	21.9	0.02204	0.0716	1 500
year 4 logged daily	0.00996	31.9	0.00991	25.1	0.01646	0.0366	1 500
year 4 logged hourly	0.00989	32.4	0.00986	25.5	0.01584	0.0309	1 500
year 4 logged every 15 minutes	0.00989	32.4	0.00986	25.5	0.01584	0.0309	1 500
year 4 daily, and the archive of windows 0–36 at daily	0.00921	37.0	0.00918	30.6	0.00821	0.0331	1 500
year 4 hourly, and the archive of windows 0–36 at hourly	0.00914	37.5	0.00912	31.1	0.00700	0.0280	1 500
year 4 monthly, a reference condition test now	0.01113	23.9	0.01077	18.6	0.02482	0.0990	1 500
year 4 monthly, the side-by-side check now	0.01106	24.4	0.01084	18.1	0.02441	0.0958	1 500
year 4 daily, readings treated as independent	0.00993	32.1	0.00989	25.3	0.01621	0.0344	–
year 4 hourly, readings treated as independent	0.00973	33.5	0.00975	26.3	0.01435	0.0081	–
the path known exactly at windows 0–45 (the floor)	0.00894	38.9	0.00894	32.4	0	0	–

At the same price, nine months of daily logging, R$ 1 125 of the year's R$ 1 500, take 31.9 % off coil 1's rate band, against 23.9 % for a reference condition test and 24.4 % for the side-by-side check. On the capacity margin daily logging takes the standard deviation at window 45 from monthly logging's 0.1091 to 0.0366 kW/K, a cut of 66.5 %. Hourly logging takes a further 0.71 % and 0.45 % off daily's rate bands, 3.7 % and 2.5 % off the two health indices and 15.5 % off the capacity margin, at no extra cost under the stated price. The readings treated as independent would report the margin's standard deviation at hourly logging as 0.0081 kW/K, against the 0.0309 the rows carry.

The candidates are ranked in Figure 24.6. The left panel gives the width each takes off coil 1's rate spread, in percent, with the floor at 45 windows dashed; the right panel the standard deviation of the total conductance at window 45. Each row is a candidate; the two crew visits at R$ 1 500 are in gold. Daily and hourly logging beat both crew visits in both panels, and on the right by the most: 0.0366 and 0.0309 kW/K against the crew visits' 0.0958 and 0.0990.

The answer to the drill question is *hourly, carried in the fit as 38.8 effective readings a window, and no faster.* Daily is where the rate saturates; hourly is where the level does; nothing faster changes either. On windows 0–36, hourly logging against monthly takes 29.6 % off coil 1's rate spread and 74.2 % off the capacity margin's.

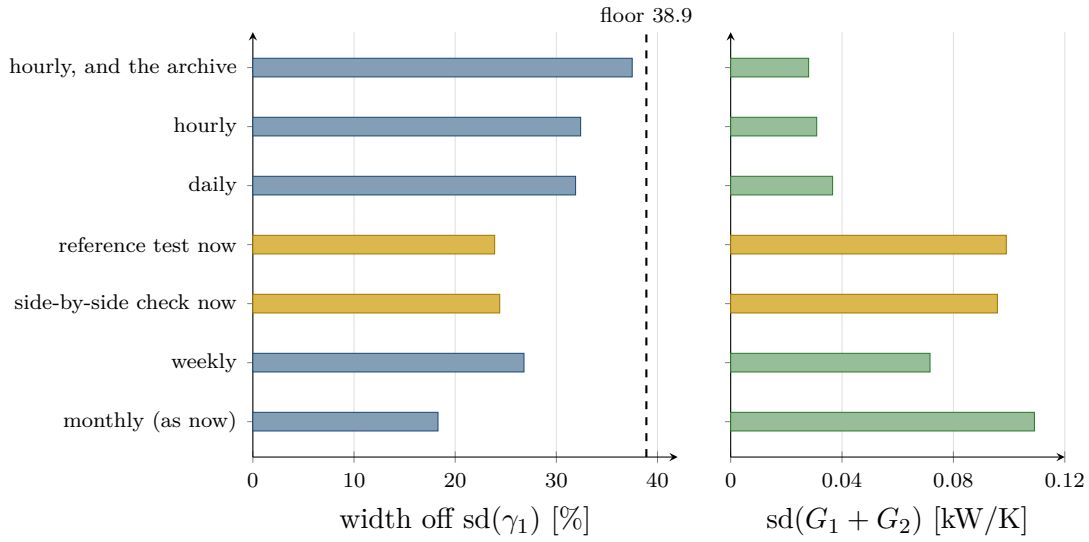


Figure 24.6: The candidates for year 4 of Table 24.8, refit at window 45. Left: the width taken off coil 1's rate spread against the record with no new rows, with the floor of 45 windows exactly known (dashed). Right: the standard deviation of the hall's total conductance at window 45. Gold: the two crew visits at R$ 1 500; the logging rates finer than monthly cost R$ 1 500 a year.

24.4 The checks (step 5)

Each answer of Section 24.3 rests on a closed-form effective number, the reduced state, the Gauss–Newton covariance at the base fit's linearization point, a stated error model and the base drawing, and this hall lets each be tested against the calculation it stands for.

24.4.1 The effective number by simulation

The 20 000 simulated windows of Section 24.2.4 reproduce the effective numbers of Table 24.5 within one Monte Carlo standard error at hourly and daily logging, for the whitened rows and for the simple average, as Section 24.3.2 set beside them. The same windows, logged contiguously, check the boundary between windows: at hourly logging the correlation of a window's whitened mean with the next window's first hour is 0.1565 simulated (standard error 0.0071) against 0.1444 in closed form, within two standard errors. In the simulated records the collapsed readings' errors, in units of $\sigma_k/\sqrt{n_{\text{eff}}}$, have standard deviations of 1.0054 at monthly, 1.0062 at daily and 1.0067 at hourly logging over 18 500 values each (standard error 0.0052), within two standard errors of one: the whitened weights of equation (24.6) give the mean the standard deviation the rows claim for it.

24.4.2 The limit of exact rows and the curvature term

The answers run on the reduced state, the limit $\sigma_p \to 0$ of the full state by equation (4.48). The check builds the chapter's own full state, the 115 entries with each window's five temperatures, 300 entries under 185 physical rows at a common tolerance, and forms the daily preposterior from its Jacobian. Table 24.9 gives the result: each row is a tolerance, the columns the rates' standard deviations at monthly and daily logging and coil 1's health index at daily logging, and the last row the reduced state.

At 0.001 the full state reproduces the reduced state in every printed digit, and the monthly columns reproduce Chapter 10's limit table (Table 10.6). The daily columns show what the logging changes in the check. At a tolerance of 0.1 the health index's spread at daily logging is

Table 24.9: The limit of exact physical rows on the 300-entry full state of the record: the rates' standard deviations at monthly and daily logging and that of coil 1's health index at daily logging, at three tolerances, against the reduced state.

Tolerance σ_p [kW or K]	sd γ_1, monthly [yr^{-1}]	sd γ_2, monthly [yr^{-1}]	sd γ_1, daily [yr^{-1}]	sd γ_2, daily [yr^{-1}]	sd $\ln h_{1,36}$, daily [–]
0.1	0.01469	0.01324	0.01048	0.01035	0.00944
0.01	0.01463	0.01323	0.01042	0.01034	0.00878
0.001	0.01463	0.01323	0.01042	0.01034	0.00876
reduced state	0.01463	0.01323	0.01042	0.01034	0.00876

0.00944 against 0.00876, a larger gap than any at monthly logging, because a tolerance of 0.1 K is no longer small beside a window read to 0.05916 K (Table 24.5): the better the readings, the tighter the tolerance at which the physical rows count as exact.

The covariance (24.10) is the Gauss–Newton form at the base fit's MAP state. On one simulated daily record (seed 24001) the full Hessian by differences gives the Newton covariance, whose standard deviations against Gauss–Newton's are in the ratios 0.9998, 1.0000 and 0.9996 for γ_1, γ_2 and $\dot{b}_{d,1}$; the largest entry of the curvature term $\mathbf{C}$ of equation (4.44) is 1.03×10^{-3} of the largest of $\mathbf{J}^{\mathsf{T}}\mathbf{J}$. The same record's own Gauss–Newton standard deviations, 0.01043 and 0.01034, match the preposterior's 0.01042 and 0.01034: refitted at its own MAP state, a logged record has the band the preposterior gave it before it was logged. That record's $2\mathcal{L}$ is 213.30 on 146 degrees of freedom, the first sign of the finding of Section 24.4.4.

24.4.3 Sensitivity to the error model and the wander

The AR(1) at $\varrho_1 = 0.9$ is a stated input, and Table 24.10 replaces it. Each row is an error model: the readings independent, the AR(1) at five correlations an hour apart, the AR(1) at 0.9 with a part f_{sh} of the variance shared by the whole window, the same shared part persisting from month to month, and the simple average at 0.9. The columns give the effective number at daily and hourly logging, each coil's rate spread at daily and hourly logging, and coil 1's health index spread at hourly logging. A shared part enters equation (24.5) as a window-wide component: with the AR(1) effective number n^{AR} of the remaining part and $\tilde{n} = n^{\mathrm{AR}}/(1 - f_{\mathrm{sh}})$, the effective number is $\tilde{n}/(1 + f_{\mathrm{sh}}\tilde{n})$, which never exceeds $1/f_{\mathrm{sh}}$ (Exercise 24.3). Errors that persist across months, the case Chapter 10 left to this chapter, are carried by letting the shared part of each window's error be AR(1) from one window to the next, at a stated correlation of 0.5 a month: each channel's window means then have the covariance, in units of σ_k^2, $(1 - f_{\mathrm{sh}})/n^{\mathrm{AR}}$ on the diagonal plus $f_{\mathrm{sh}}\, 0.5^{|t-s|}$ between windows t and s, and the reading rows enter the information through its inverse rather than one window at a time. At monthly logging every model but this one gives the base fit's 0.01463 and 0.01323, because one reading a window has no partner within the window to correlate with; the persistent model gives 0.01520 and 0.01363 even then.

The rate's verdict holds over every model. At daily or hourly logging its standard deviation stays between 0.01002 and 0.01220 in every row, against 0.01463 and 0.01323 at monthly, because the rate sits near its floor whatever the readings are worth. The level's verdict does not hold over every model. It depends on how fast the error forgets, and most on a part the window shares: a ninth of the variance shared caps the effective number at 9 and takes the hourly health index's spread from 0.00745 to 0.01418, a third shared caps it at 3 and takes it to 0.02074. Which model is right is a property of the hall's sensors and of its hours, and the first logged months would show it.

An error that persists across months reaches the rate as well. With a ninth of the variance shared by the window and carried from month to month at 0.5, the rates' spreads are 0.01190 and 0.01132 at daily logging, against 0.01113 and 0.01084 with the same part independent across

Table 24.10: The answers under other error models for the logged readings, windows 0–36: the effective number at daily and hourly logging, the rates' standard deviations, and coil 1's health index at hourly logging. The chapter's model is the fourth row.

Error model	n_{eff} daily [–]	n_{eff} hourly [–]	sd γ_1 daily [yr^{-1}]	sd γ_2 daily [yr^{-1}]	sd γ_1 hourly [yr^{-1}]	sd γ_2 hourly [yr^{-1}]	sd $\ln h_{1,36}$ hourly [–]
independent ($\varrho_1 = 0$)	30.000	720.000	0.01037	0.01030	0.01002	0.01002	0.00211
$\varrho_1 = 0.5$	30.000	240.667	0.01037	0.01030	0.01007	0.01006	0.00348
$\varrho_1 = 0.8$	29.727	80.889	0.01037	0.01031	0.01017	0.01015	0.00554
$\varrho_1 = 0.9$ (the chapter's)	25.715	38.842	0.01042	0.01034	0.01030	0.01025	0.00745
$\varrho_1 = 0.95$	16.892	19.436	0.01058	0.01045	0.01052	0.01041	0.00977
$\varrho_1 = 0.98$	7.896	8.263	0.01102	0.01076	0.01098	0.01074	0.01363
$\varrho_1 = 0.9$, a ninth shared by the window	6.864	7.463	0.01113	0.01084	0.01106	0.01079	0.01418
$\varrho_1 = 0.9$, a ninth shared, AR(1) across months at 0.5	6.864	7.463	0.01190	0.01132	0.01184	0.01128	0.01899
$\varrho_1 = 0.9$, a third shared by the window	2.784	2.853	0.01220	0.01156	0.01216	0.01154	0.02074
simple average, $\varrho_1 = 0.9$	25.705	38.400	0.01042	0.01034	0.01031	0.01025	0.00748

windows, and the hourly health index's is 0.01899 against 0.01418. It changes the saturation too. The effective number within a window is capped at 9 as before, but the cap is no longer the end: with unlimited readings a window, the rates' spreads stop at 0.01170 and 0.01118 and the health index's at 0.01822, above the wander floor of 0.01000 and above the 0.01106 and 0.01079 that hourly logging gives with the same part independent across months. The persistent part moves neighboring windows together the way the coils' path does, so the rows read part of it as the path, and no logging rate averages it away.

The floor is the chosen wander's. At daily logging and $\sigma_q = 0.0025$, half the chapter's, the rates' spreads are 0.00559 and 0.00544 (from 0.01118 and 0.00947 at monthly) against a floor of 0.00500, and the readings' share is 11.8 and 8.0 %; at $\sigma_q = 0.01$ they are 0.02028 and 0.02024 (from 0.02338 and 0.02239) against 0.02000, with shares of 0.9 and 0.6 %. The floor moves one for one with the wander, and with a smaller wander the readings keep a larger share, so a faster logging rate pays for the rate longer.

24.4.4 Fresh records at the stated truth

The simulated records of Section 24.2.4, 100 per rate, test the answers against the record's stated truth. Table 24.11 collects them. Each row is a rate, with the hourly records also fitted as independent readings; the columns give, for each coil, the scatter of the 100 rate estimates, their mean posterior standard deviation and the share of 5–95 % bands that cover the pseudo-true value at that rate; the coverage of coil 1's health index; the median $2\mathcal{L}$ on 146 degrees of freedom and the share of records rejecting the drawing at $P < 0.05$; the median of the trend rows' part of $2\mathcal{L}$, 72 rows; and the share of records reading coil 1 the faster.

A band of assumed wander. At daily logging the estimates scatter by 0.00268 against a mean posterior standard deviation of 0.01044, and every rate band covers its pseudo-true value. The truth's path carries no wander, so the floor the drawing allows for it is spent on nothing: the band is conservative at this truth, as Chapter 10 found at monthly logging, and more so, since the readings' part that would have filled it shrinks with the logging.

The drawing's rejection by logged records. This is the step's surprise. On 146 degrees of freedom the base drawing's $2\mathcal{L}$ has a median of 150.09 over the monthly records, with 9 % of

Table 24.11: One hundred simulated records per rate at the stated truth, AR(1) errors drawn at the logging interval and collapsed by the whitened weights; seeds 24100, 24200 and 24300. "Covers" is the share of 5–95 % bands containing the drawing's pseudo-true value at that rate's weights; its Monte Carlo standard error near 0.9 is 0.030.

Rate	scatter $\hat{\gamma}_1$ [yr^{-1}]	mean sd [yr^{-1}]	covers [–]	scatter $\hat{\gamma}_2$ [yr^{-1}]	mean sd [yr^{-1}]	covers [–]	$\ln h_{1,36}$ covers [–]	median $2\mathcal{L}$ [–]	share $P < 0.05$ [–]	trend rows' part [–]	coil 1 faster [–]
monthly	0.00827	0.01470	1.00	0.00776	0.01320	0.99	0.96	150.09	0.09	3.81	0.18
daily	0.00268	0.01044	1.00	0.00209	0.01034	1.00	0.94	192.47	0.84	51.86	0.54
hourly	0.00218	0.01032	1.00	0.00185	0.01025	1.00	0.97	203.66	0.92	66.18	0.73
hourly, independent	0.00298	0.01003	1.00	0.00281	0.01002	1.00	0.24	1699.89	1.00	288.09	0.86

them rejecting it at $P < 0.05$. Over the daily records the median is 192.47 and 84 % reject it; over the hourly records, 203.66 and 92 %. The excess sits in the trend rows: their part is 51.86 on 72 rows at daily logging, against 3.81 at monthly. Read well, the windows show the path's departures from a straight line with wander, the departures that Chapters 11 and 12 test. The monthly record shows them in its path, within its errors, untested (Chapter 10); read well, they exceed what the wander allows. Whether a daily record decides those tests is a question of power, Chapter 22's computation, and it is named here and not computed.

The drawing's moving answer and the coils' order. The pseudo-true value of each rate, the base drawing fitted to the truth's noise-free readings with each rate's weights, is −3.793 and −5.179 %/yr at monthly logging. At daily logging it is −4.594 and −4.472, and at hourly −4.664 and −4.417: it moves toward the truth's net change over the record, −4.877 and −4.241. With well-read windows the drawing's path follows the true path, and the rate is that path's telescoped end-to-end change, equation (10.7). Coil 1 reads the faster in 0.18 of the monthly records, 0.54 of the daily and 0.73 of the hourly. Chapter 10 found coil 2 the faster and said that the order holds under the base fit's constant rates, while the truth's net change runs the other way (Section 10.6.4); the logged records show that the order is the monthly record's answer, not the coils'.

The fresh records are drawn in Figure 24.7, one panel per coil. The horizontal axis holds the four treatments; the vertical axis is the rate in %/yr. Each dot is one record's estimate, spread sideways only to be seen; the bar to its right is the mean 5–95 % band about the records' mean estimate (the dot on the bar); the thick black dash is the pseudo-true value at that rate; the dotted line is the truth's net change. The dots gather as the logging gets faster, far inside the bars; the dashes move toward the dotted line, coil 1's down and coil 2's up.

24.4.5 The answers against the stated truth

The stated truth's net change per year over the record is −0.05000 for coil 1 and −0.04333 for coil 2, −4.877 and −4.241 %/yr, with no wander, and its $\ln h$ at window 36 is −0.15000 and −0.13000. The plant's record draws independent reading errors; the AR(1) is this chapter's stated input, and only the simulated records carry it. Table 24.12 collects the routes to the daily answers beside it. Each row is a route; the columns give a center and a spread for each coil's rate and for coil 1's health index, where the route has one.

A spread has no truth to be compared with: it is a property of the drawing and the logging. The preposterior routes agree on it to the fourth decimal, with the full state at $\sigma_p = 0.1$ the one exception. The simulated records agree on the rates, 0.01043 and 0.01044 against 0.01042 for coil 1, and give a wider health index, 0.00899 against 0.00876: each record is fitted at its own MAP state, not at the base fit's. What the truth checks is the center a logged record would report, the pseudo-true value. At daily weights it is −0.04703 and −0.04575 for the rates and −0.14109 for coil 1's index, against the truth's −0.05000, −0.04333 and −0.15000, where

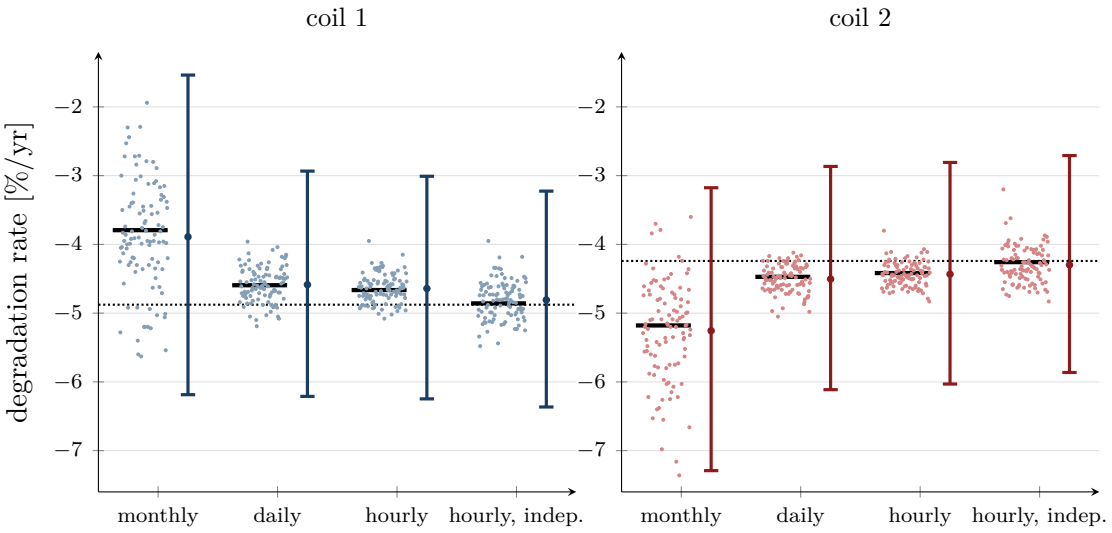


Figure 24.7: One hundred simulated records per rate at the stated truth, each coil's rate in %/yr. Dots: the records' estimates, spread sideways for legibility. Bars: the mean 5–95 % band, the records' mean estimate ±1.6449 times their mean posterior standard deviation, with the mean estimate as the dot on the bar. Thick black dashes: the drawing's pseudo-true value at that rate's weights. Dotted: the stated truth's net change over the record, −4.877 (coil 1) and −4.241 %/yr (coil 2), a check marker. "Hourly, indep.": the hourly records fitted as independent readings.

Table 24.12: The routes to the answers at daily logging, and the stated truth. A preposterior has a spread and no center of its own; the centers shown are the base fit's (monthly), the drawing's pseudo-true values at daily weights, and the truth.

Route	γ_1 [yr^{-1}]	sd γ_1 [yr^{-1}]	γ_2 [yr^{-1}]	sd γ_2 [yr^{-1}]	$\ln h_{1,36}$ [–]	sd $\ln h_{1,36}$ [–]
base fit, the record as it stands (monthly)	−0.04240	0.01463	−0.06253	0.01323	−0.12719	0.03203
preposterior, reduced state, daily (reported)	–	0.01042	–	0.01034	–	0.00876
preposterior, full state, $\sigma_p = 10^{-3}$	–	0.01042	–	0.01034	–	0.00876
preposterior, full state, $\sigma_p = 0.1$	–	0.01048	–	0.01035	–	0.00944
one simulated daily record, Gauss–Newton	–	0.01043	–	0.01034	–	–
100 simulated daily records, mean posterior sd	–	0.01044	–	0.01034	–	0.00899
100 simulated daily records, scatter of the estimates	–	0.00268	–	0.00209	–	0.00804
pseudo-true value at daily weights	−0.04703	–	−0.04575	–	−0.14109	–
stated truth	−0.05000	–	−0.04333	–	−0.15000	–

the record as it stands reads −0.04240, −0.06253 and −0.12719: the drawing's answer at daily weights lies nearer each coil's net change, and its two rates fall in the truth's order. A single logged record need not: coil 1 reads the faster in 0.54 of the daily records and 0.73 of the hourly ones.

24.5 The report (step 6)

Logging 720 times faster does not make the rate 27 times sharper. An hour's error carries into the next, so a month of hourly logging holds 39 independent readings, not 720, and a one-minute log holds the same 39. The rate's band narrows by 29.6 % for coil 1 and 22.5 % for coil 2 from monthly to hourly logging, and daily logging does nine-tenths of what any logging can; the rest of the band is the path's month-to-month wander, which no sensor removes. Today's health index and the capacity margin gain far more, 73.9 to 76.7 % of their spread at hourly logging, with hourly logging a further 13.9 to 15.3 % beyond daily. Treated as independent, hourly readings would make the health index's band 3.5 times too narrow. And the path's departures from a straight line, which the monthly record shows within its errors, untested, a daily record would show beyond what the wander allows: the straight line with its wander does not fit the coils' paths, and the order of the coils' rates is the monthly record's, not the coils'.

24.5.1 The answers and their bands

Table 24.13 gives the answers in the order of Table 24.1. Each row is a question; the second column gives the reported value from the preposterior, the third the simulated records' or simulated windows' value beside it, and the fourth the range the answer spans: across the logging rates, the error models of Table 24.10, or the candidates.

The answers rest on the base fit's rows and on two stated inputs above all. The wander σ_q sets the rate's floor and therefore where logging stops paying for the rate; the error model sets the level's gain, and a part of the error shared by the window would cut it.

24.5.2 The entry in the maintenance log

The logging rate is the next entry of the two-unit hall's maintenance log, at window 36 with the entries of Part III and of this volume's earlier chapters. It uses nothing later than window 36, and in the values of Table 24.13 it reads:

> *Window 36 (Dec Y3), the logging rate of the condition record.*
> Set the trend export to hourly for the five channels from window 37 (Jan Y4), and read it into the fit as 38.8 effective readings per window (AR(1) errors, 0.9 an hour apart). Do not log faster; do not treat the hourly readings as independent.
> Expected at the window-45 refit (Sep Y4): rate spreads 0.0099 yr^{-1} for each coil (from 0.0146 and 0.0132 today; monthly logging would give 0.0120 and 0.0111); capacity margin spread 0.031 kW/K (monthly 0.109).
> If the logger's archive of windows 0–36 holds hourly readings, read them too: at window 45 the margin's spread falls to 0.028 kW/K and coil 1's health-index spread to 0.0070 (0.0158 without the archive).
> Cost: R$ 1 500 a year (stated); it buys more than a reference condition test or the side-by-side check at the same price.
> Open: refit the drawing on the first logged months. The record may then show the departures from the straight line that the monthly record shows within its errors, beyond what the wander allows (Chapters 11 and 12), and the order of the coils' rates is to be read again.
> Next measurement: none beyond the logging for the rate; for the level, a permanent instrument (step 7).

Table 24.13: The report for the two-unit hall's logging rate at window 36 (Dec Y3). Reported values from the preposterior (24.10) on windows 0–36, or (24.16) to window 45; beside them the simulated records (Section 24.2.4).

Asset manager's question	Reported value	Simulated	Range
Hourly, daily, or monthly?	sd γ_1, γ_2: monthly 0.01463, 0.01323; daily 0.01042, 0.01034; hourly 0.01030, 0.01025 [yr^{-1}]	daily 0.01044, 0.01034; hourly 0.01032, 0.01025 [yr^{-1}]	band width 4.613 to 3.249 (coil 1), 4.089 to 3.168 (coil 2) [%/yr]; floor 0.01000 [yr^{-1}]
How many independent readings does a month of logging hold?	n_{eff}: daily 25.7153, hourly 38.8421, limit 38.9298 [–]	25.653 (0.257), 38.642 (0.386) [–]	simple average 25.7051, 38.4000 [–]
From which rate on does logging stop narrowing the rate's band, and why?	readings' share $\leq 10\,\%$ at $n_{\mathrm{L}} = 8$ (coil 1) and 5 (coil 2), $\leq 5\,\%$ at 18 and 11 [–]; daily share 3.4, 2.2 %; path-ends share daily 7.9, 6.4 %, coil 1's never $\leq 5\,\%$	–	daily share of coil 1 11.8 to 0.9 % for $\sigma_q = 0.0025$ to 0.01
Does a faster rate help today's health index and the capacity margin?	sd $\ln h_{1,36}$: 0.03203, 0.00876, 0.00745 [–]; sd $G_1 + G_2$: 0.1193, 0.0364, 0.0308 [kW/K] (monthly, daily, hourly)	sd $\ln h_{1,36}$ 0.03233, 0.00899, 0.00765 [–]	$h_{1,36}$ hourly 0.870 to 0.891 [–]; margin 8.6 to 33.4 [sd]
What does the band look like if the logged readings are treated as independent?	hourly sd $\ln h_{1,36}$ 0.00211 against 0.00745, 3.54 times too narrow; rate 1.028 [–]	coverage 0.24 against 0.97 [–]	sd $\ln h_{1,36}$ hourly 0.00211 to 0.02074 over the error models [–]
What should the logger be set to for year 4, against a crew visit at the same price?	hourly, as 38.8 effective readings: sd γ 0.00989, 0.00986 [yr^{-1}], sd $G_1 + G_2$ 0.0309 [kW/K] at window 45; R$ 1 500 a year	–	crew visits sd γ_1 0.01113 and 0.01106 [yr^{-1}]; floor 0.00894 [yr^{-1}]

The entry records what the spreads rest on as well as their values: the base fit's drawing, its wander of 0.005 a month, and the AR(1) errors stated for the logged readings, which the first logged months can test.

Common mistake

Counting logged readings as independent. Readings logged faster than their error forgets are worth an effective number of independent readings that saturates with the window's span, not their count; the fit must carry their correlation, or the whitened mean with its true standard deviation, or it reports a health index several times sharper than the record supports.

24.6 The next measurement (step 7)

The report sets the logger to hourly from window 37. What else is worth buying is asked with that logging in place: each candidate is a row now, or a row logged hourly in every window to the refit, on the extended record of Section 24.3.6.

24.6.1 The preposterior with the logging set

With hourly logging adopted, the base of every comparison is the preposterior (24.16) at $n_{\mathrm{eff}}(720)$ for windows 37–45 and $n_{\mathrm{p}} = 1$, written $\boldsymbol{\Sigma}_{\mathrm{h}}$. A candidate a adds rows with lines $\mathbf{J}_a$ in the columns of the extended state, and by equation (2.17)

$$\boldsymbol{\Sigma}_a = \left(\boldsymbol{\Sigma}_{\mathrm{h}}^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a\right)^{-1}, \qquad 1 - \frac{\sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_a\mathbf{g}}}{\sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_{\mathrm{h}}\mathbf{g}}}, \tag{24.17}$$

the second the width a candidate takes off a quantity, against hourly logging alone: 0.00989 and 0.00986 yr^{-1} on the rates and 0.0309 kW/K on the total conductance at window 45. A permanent instrument is logged hourly like the record, so its row in each window 37 to 45 carries the weight $n_{\mathrm{eff}}(720)$, with errors of the same AR(1). The candidates are the two crew visits now, the side-by-side check of unit 1's discharge sensor (ten pairs at 0.1 K) and a reference condition test (0.1 K on four temperatures); water-side heat meters on both coils ($\sigma_W = 3.0$ kW), permanent; a second discharge sensor on unit 1 (0.3 K), permanent, the permanent instrument at the place of the question, where the drift and coil 1 share a reading; and the logger's archive of windows 0–36 read at hourly. The last row is not a measurement: the path known exactly at windows 0–45, the floor.

24.6.2 The candidates

Table 24.14 ranks the candidates. Each row is one; the columns give each coil's rate spread after it by equation (24.17) with the width it takes off, the drift rate's spread, coil 1's health index at the refit, and the total conductance's spread with the width taken off it.

The widths taken off are compared in Figure 24.8, two bars per candidate: coil 1's rate (blue) and the capacity margin (green), with the floor of the rate dashed. Every blue bar is short; the green bars of the two permanent instruments are the longest.

With the logging set, the side-by-side check that Chapters 7, 10 and 11 kept open takes 0.1 % off coil 1's rate: given the onset at window 20, the logged record has already narrowed the drift, whose spread is 0.0223 K/yr with hourly logging alone, from 0.1681 today, and 0.0213 with the check. Every candidate on the rate is capped by the floor, 9.6 and 9.3 %, and only the archive comes near it, at 7.6 and 7.5 %. The permanent instruments act on the level, not on the rate: the second discharge sensor takes 19.6 % off the capacity margin's spread and 1.1 % off coil 1's rate, and the heat meters 15.3 % and 1.1 %. The second sensor's row is the first sensor's

Table 24.14: Candidate measurements with hourly logging from window 37 adopted, refit at window 45 (Sep Y4), by the preposterior (24.17); the widths taken off are against hourly logging alone.

Candidate	sd γ_1 [yr^{-1}]	off [%]	sd γ_2 [yr^{-1}]	off [%]	sd $\dot{b}_{d,1}$ [K/yr]	sd $\ln h_{1,45}$ [–]	sd $G_1 + G_2$ [kW/K]	off [%]
hourly logging alone	0.00989	0.0	0.00986	0.0	0.0223	0.01584	0.0309	0.0
the side-by-side check now (ten pairs, 0.1 K)	0.00988	0.1	0.00986	0.0	0.0213	0.01577	0.0305	1.4
a reference condition test now (0.1 K, four temperatures)	0.00989	0.0	0.00986	0.0	0.0217	0.01580	0.0306	1.0
water-side heat meters on both coils, permanent, logged hourly	0.00979	1.1	0.00980	0.6	0.0136	0.01489	0.0262	15.3
a second discharge sensor on unit 1 (0.3 K), permanent, logged hourly	0.00978	1.1	0.00983	0.3	0.0118	0.01484	0.0249	19.6
the archive of windows 0–36 at hourly	0.00914	7.6	0.00912	7.5	0.0168	0.00700	0.0280	9.4
the path known exactly at windows 0–45 (the floor)	0.00894	9.6	0.00894	9.3	0.0089	0	0	100.0

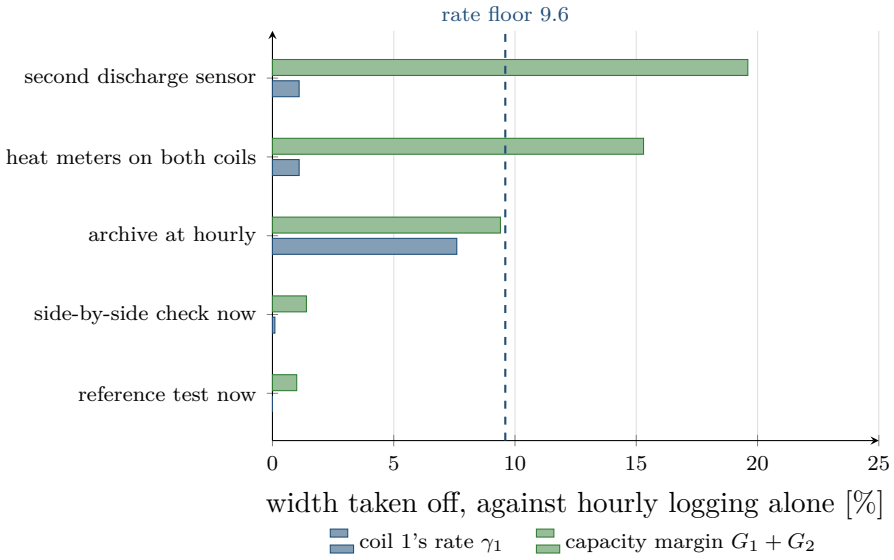


Figure 24.8: The candidates of Table 24.14, with hourly logging from window 37 adopted: the width each takes off coil 1's rate spread and off the total conductance's spread at window 45, by the preposterior (24.17). Dashed: the rate's floor, the path known exactly at windows 0–45; on the margin that limit takes off 100 %.

line without the drift's column, so it reads coil 1's discharge free of the drift and turns the ridge of Section 24.2.3: the drift's spread falls to 0.0118 K/yr. The heat meter's row reads the split of the heat between the coils, as in Chapter 10: at window 37 its slopes are +36.4112 on $g_{1,37}$, −36.4112 on $g_{2,37}$ and 0.5201 on Q_{37}, in kW per unit entry, which the return air cannot give.

24.6.3 The recommendation

Set the logger to hourly from window 37 and carry the readings as 38.8 effective readings a window; read the logger's archive of windows 0–36 at hourly as well, if it exists. The side-by-side check of unit 1's discharge sensor is no longer needed for the rate. Its remaining value is the drift's onset, and through the onset the drift's size, which logging does not place: over onsets the size runs from 0.4738 to 1.5370 K/yr (Section 10.4.4). Of the permanent instruments, a second discharge sensor on unit 1 serves the level, with the larger cut of the margin's spread; its value for the crossing date is Chapter 23's question. The logging's cost, R$ 1 500 a year, is the chapter's stated input, and the recommendation holds for any price of a finer rate that is the same at every rate.

24.6.4 What the analysis showed

The proposal said that 720 readings a month would make the rate 27 times sharper; read through the rows, a month of hourly readings is 39 effective readings, and the rate's band narrows by 29.6 %. The rows carry the readings' memory: whitened, a window's readings collapse to one row whose weight is an effective number, and that number saturates at the window's span over two correlation times. The trend rows carry the rate, and their wander sets its floor, so the rate stops gaining at daily logging; the level has no floor and gains until the effective number saturates, at hourly, which is why the logger is set to hourly and no faster. The probability adds the honest band and shows where the independent treatment fails: on the level, whose band would cover the drawing's pseudo-true value in 0.24 of the hourly records. And the logging the record needs for its band is also the logging that would test its drawing: a daily record rejects the straight line with wander in 84 % of the simulated records, and the drawing's answer moves toward the truth, its pseudo-true rates in the truth's order, while a single record reads coil 1 the faster in 0.54 of the daily records and 0.73 of the hourly ones.

24.6.5 Limits of the model

A constructed record. The condition record is generated from a stated truth, with the stated sensor errors, and step 5 compares the answers with that truth (Section 24.4.5). It is not field data. The plant draws independent reading errors; the AR(1), its correlation and the window's span are this chapter's stated inputs, and no logged record exists to estimate them.

One steady condition per window. An hour's load departure that moves all five channels together is treated as error, not state. Where such departures are large, each logged hour needs its own load entry, the form of Chapter 10's archive candidate, which this chapter does not compute.

An error that forgets within a day. A slower part of the error, a calibration wander over weeks or a bias at a fixed hour of the day, caps the effective number at the inverse of its share (Table 24.10); the level's gain from logging would shrink with it, and the rate's would not.

Readings independent across windows. Checked under the AR(1): two consecutive windows' means correlate by 0.0232 at hourly logging. An error that persists across months, a

sensor's slow drift other than unit 1's, is carried only as the sensitivity of Table 24.10, at one stated share and correlation; it would raise the rate's band above the floor at any logging rate.

Chapter 10's drawing. The wander σ_q is chosen, a straight line with wander, with the drift from window 20; the rate's floor is the chosen wander's, and a daily record would test the drawing itself. The answers are widths about the base fit's MAP state; a logged record would move the centers to its own fit.

The cost as stated. The logging's cost is a stated input, not a priced decision, and the value of a narrower band, as a crossing date or a trigger, is not computed here.

24.7 What to retain

Definition 24.1 (Effective number of readings and the preposterior of a logging rate)**.** The n_{L} logged readings of one channel in a window, with errors correlated as $\mathrm{Cov} = \sigma_k^2\mathbf{R}$, whitened by $\mathbf{R}^{-1/2}$, collapse to one row on their generalized least-squares mean, of standard deviation $\sigma_k/\sqrt{n_{\mathrm{eff}}}$, with $n_{\mathrm{eff}} = \mathbf{1}^{\mathsf{T}}\mathbf{R}^{-1}\mathbf{1}$; for AR(1) errors of correlation ϱ at the interval, $n_{\mathrm{eff}} = (n_{\mathrm{L}} - (n_{\mathrm{L}} - 2)\varrho)/(1+\varrho)$, which saturates at $1 + \frac{1}{2}\vartheta_w \ln(1/\varrho_1)$ over a span ϑ_w. A record logged at n_{L} readings a window has the preposterior $\mathbf{\Sigma}_{n_{\mathrm{L}}} = (\mathbf{\Sigma}^{-1} + (n_{\mathrm{eff}} - 1)\mathbf{J}_y^{\mathsf{T}}\mathbf{J}_y)^{-1}$, and a degradation rate's variance under it is the wander floor's plus the health index's variance over $(3\,\mathrm{yr})^2$; only the second is a reading's to narrow.

Three conclusions carry beyond the hall. A logging rate is a candidate measurement like a sensor, and its value is read from the preposterior of the rows it adds, with their correlation carried. A degradation rate is set by the path's wander once the windows are read well, so it stops gaining from faster logging long before the level does. And the independent treatment of logged readings is wrong most where logging helps most, on the level.

Table 24.15 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 24.15: The terms of the sampling frequency and where each is defined. Units for the two-unit hall.

Term	Expression	Units	Meaning	Defined in
Logged readings as rows	eq. (24.2), $\mathrm{Cov} = \sigma_k^2\mathbf{R}$, $R_{ii'} = \varrho^{\lvert i-i'\rvert}$	–	a window's readings with a full correlation	24.1.4
AR(1) errors	eq. (24.1), $\varrho = \varrho_1^{\Delta\vartheta/1\,\mathrm{h}}$	–	each error remembers the last	24.1.1
Whitened rows	eq. (24.4)	–	each row the new part of its error	24.1.4
Collapse to the mean	eq. (24.6)	–	weights 1 at the ends, $1-\varrho$ inside	24.1.4
Effective number	eq. (24.5)	–	independent readings a window is worth	24.1.4
Saturation	eq. (24.7)	–	one reading per two correlation times	24.1.4
Preposterior and answers				
Preposterior of a logging rate	eq. (24.10)	varies	the reading rows weighted by n_{eff}	24.2.1
Readings' share	eq. (24.14)	–	the straight line's variance over the rate's	24.3.3
Wander floor	eq. (10.12)	yr^{-1}	the rate's limit as logging grows	24.3.1
The level and the rate tied	eq. (24.15)	–	the level has no floor	24.3.4
Independent treatment	eq. (24.10) at $n_{\mathrm{eff}} = n_{\mathrm{L}}$	–	overconfident on the level	24.3.5
Error shared by a window	$n_{\mathrm{eff}} = \tilde{n}/(1+f_{\mathrm{sh}}\tilde{n})$	–	caps n_{eff} at $1/f_{\mathrm{sh}}$	24.4.3

24.8 Exercises

Exercise 24.1. *Formulation.* Write Chapter 10's "second logged hour" as two correlated rows on one window value, and give the effective number of two readings an hour apart, 12 h apart and a day apart, under the chapter's AR(1). (Two rows $(f-y_i)/\sigma$ with $\mathbf{R} = \left[\begin{smallmatrix} 1 & \varrho \\ \varrho & 1 \end{smallmatrix}\right]$ give $n_{\mathrm{eff}} = \mathbf{1}^{\mathsf{T}}\mathbf{R}^{-1}\mathbf{1} = 2/(1+\varrho)$: 1.0526 at $\varrho = 0.9$, 1.5595 at 0.2824, and 1.8523 at a day's 0.0798.)

Exercise 24.2. *Derivation.* Show that $\mathbf{R}^{-1}$ of an AR(1) is tridiagonal, that its row sums give the weights $(1, 1-\varrho, \ldots, 1-\varrho, 1)/(1+\varrho)$, and that $\mathbf{1}^{\mathsf{T}}\mathbf{R}^{-1}\mathbf{1} \to 1 + \frac{1}{2}\vartheta_w \ln(1/\varrho_1)$ at a fixed span. (Multiply $\mathbf{R}$ by the matrix of equation (24.3); each end row sums to $(1-\varrho)/(1-\varrho^2) = 1/(1+\varrho)$ and each interior row to $(1-\varrho)^2/(1-\varrho^2)$; with $n_{\mathrm{L}}(1-\varrho) \to \vartheta_w \ln(1/\varrho_1)$ the sum tends to the limit. At hourly logging, $73.8/1.9 = 38.8421$; the limit is 38.9298.)

Exercise 24.3. *Derivation.* A part f_{sh} of each reading's error variance is shared by the whole window and the rest is AR(1) with effective number n^{AR}. Show that $n_{\mathrm{eff}} = \tilde{n}/(1+f_{\mathrm{sh}}\tilde{n})$ with $\tilde{n} = n^{\mathrm{AR}}/(1-f_{\mathrm{sh}})$, that it never exceeds $1/f_{\mathrm{sh}}$, and evaluate it for $f_{\mathrm{sh}} = 1/9$. (By the Sherman–Morrison identity the whitened mean keeps the AR(1) part's weights, and its variance is the shared part's $f_{\mathrm{sh}}\sigma^2$ plus the rest's $(1-f_{\mathrm{sh}})\sigma^2/n^{\mathrm{AR}} = \sigma^2/\tilde{n}$; its effective number is $1/(f_{\mathrm{sh}} + 1/\tilde{n}) = \tilde{n}/(1+f_{\mathrm{sh}}\tilde{n}) < 1/f_{\mathrm{sh}}$. Daily 6.8645, hourly 7.4629, cap 9.)

Exercise 24.4. *Application.* Redo the daily row of Table 24.4 at $\varrho_1 = 0.95$, and say whether the verdict for the rate changes. ($\varrho = 0.95^{24} = 0.2920$, $n_{\mathrm{eff}} = 16.8920$; the rates' spreads are 0.01058 and 0.01045, against 0.01042 and 0.01034 at $\varrho_1 = 0.9$ and the floor's 0.01000, and the readings' shares 5.0 and 3.3 %: the verdict for the rate stands.)

Exercise 24.5. *Application.* Re-value Chapter 10's archive candidate, a second logged hour of every window, on the base fit, when the second hour is adjacent to the first, 12 h away, or

independent of it. (With $n_{\text{eff}} = 1.0526$, 1.5595 and 2 in equation (24.10): $\text{sd}(\gamma_1) = 0.01446$, 0.01335 and 0.01280, taking 1.1, 8.7 and 12.5 % off; $\text{sd}(\gamma_2) = 0.01311$, 0.01235 and 0.01197, taking 0.9, 6.7 and 9.5 % off. The independent case is Chapter 10's 12.5 %.)

Appendix A

Atlas of asset-management questions

This appendix can be read independently of the chapters. Each row is a question an asset manager asks. The first column gives the question in the asset manager's words, as the chapter's questions table gives it; the second the general form of its category, written out or cited by equation number; the third its mathematical form for the plant of that chapter, once the plant graph and its state $\mathbf{z}$ are fixed, in that chapter's symbols; the fourth its category and the fifth its computation, the five categories and the five computations of Chapter 2; and the last the section in which the question is answered, or the chapter when a whole chapter answers it. The rows grow as chapters are added.

Asset manager's question	General form	Mathematical form	Category	Computation	Answered in
Chapter 1. Physical Assets as Graphs					
What temperature is the hall running at now?	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$, $\mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y})$	$p(T \mid \mathbf{y})$, $T = z_1$	estimation	posterior	1.3.1
How far is the as-commissioned conductance from its design value?	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$, $\mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y})$, $\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$	$p(h \mid \mathbf{y})$, $h = G/G_{\mathrm{ref}} = \mathrm{e}^{z_2}$; $\Pr(h < 1 \mid \mathbf{y})$	estimation of a derived quantity	posterior	1.3.1
Is the hall within its 24°C limit now?	$\Pr(f(\mathbf{z}) < f^* \mid \mathbf{y})$	$\Pr(T < 24^\circ\mathrm{C} \mid \mathbf{y})$	prediction under a threshold	posterior	1.3.2
Can the path carry 60 kW within 27.5°C?	$\Pr(f(\mathbf{z}) < f^* \mid \mathbf{y})$	$\Pr(T_2 < 27.5^\circ\mathrm{C} \mid \mathbf{y})$, $T_2 = T_c + (Q_2/G_{\mathrm{ref}})\,\mathrm{e}^{-z_2}$	prediction under a threshold	posterior	1.3.2
Which added sensor would narrow the health index?	preposterior covariance $(\boldsymbol{\Sigma}^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a)^{-1}$	$\boldsymbol{\Sigma}_a$ for each candidate row	design	preposterior	1.6.1
Chapter 2. The Language of Asset-Management Questions					
Where does the air path stand against its design sheet?	(2.9), (2.10)	$p(h \mid \mathbf{y})$, $h = \mathrm{e}^{z_3}$: $\hat{z}_3$ and $\sqrt{\Sigma_{33}}$ through $\mathbf{g} = (0,\ 0,\ h)^{\mathsf{T}}$	estimation	posterior	2.4.1
Can the hall carry the requested 60 kW below 27.5°C?	(2.12)	$\Pr(T_2 < 27.5^\circ\mathrm{C} \mid \mathbf{y})$, the normal of the entry T_2	prediction under a threshold	posterior	2.4.2
Will the summer season pass without a corrective work order?	(2.13)	$\mathbb{E}[\mathrm{e}^{-\lambda\tau} \mid \mathbf{y}]$ over $\mathcal{N}(\hat{\omega}, \sigma_\omega^2)$; $(b'/(b'+\tau))^{a'}$ exact	prediction under a threshold	posterior	2.4.4

Asset manager's question	General form	Mathematical form	Category	Computation	Answered in
How fast is the conductance falling, and when does it reach its functional limit?	(2.21)	$\Pr(\gamma < 0 \mid \mathbf{y})$; $\Pr(t^* \le s \mid \mathbf{y}) = \Pr(\gamma < 0,\ g_s \le g_F \mid \mathbf{y})$	prediction under a threshold; estimation for the crossing date	trend	2.3.4
Will both replacements be needed this year?	(2.12), an intersection	$\Pr(F_A)\ \Pr(F_B \mid F_A)$	prediction under a threshold	posterior	2.4.7
Is the air path fouled, or is the sensor drifting?	(2.14), (2.15)	$\Pr(M_j \mid \mathbf{y}) \propto p(\mathbf{y} \mid M_j)\ \Pr(M_j)$	hypothesis comparison; diagnosis	evidence	2.5.1, 2.5.3
Change the filters before the summer peak, or wait?	(2.23)	$\mathcal{R}(a) = c_a + C_e\ \Pr(e \mid \mathbf{y}, a)$, $a^* = \arg\min_a \mathcal{R}(a)$	design	expected consequence	2.5.4, 2.5.5
Is another condition test or an inspection worth scheduling?	(2.16)	$\boldsymbol{\Sigma}_a = (\boldsymbol{\Sigma}^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a)^{-1}$; EVSI	design	preposterior; expected consequence	2.6.1, 2.6.2
Chapter 3. Uncertainty and Conditional Probability					
Will one test read the coil below 5.0 kW/K when nothing has changed?	(2.13)	$\Pr(y_T > T_c + Q/G^* \mid T)$	prediction under a threshold	posterior	3.1.1
What is known about the coil's conductance, before and after the commissioning test?	(2.9)	$p(g \mid \mathbf{y})$; $\hat{G}$, $h = \mathrm{e}^g$ and band	estimation	posterior	3.1.2
Will the next test read below 5.0 kW/K, and do more tests remove that?	(2.13)	$\int \Pr(y_{\text{new}} > c \mid T)\, p(T \mid \mathbf{y})\, \mathrm{d}T$	prediction under a threshold	posterior	3.1.3
Is a flagged coil fouled?	(2.15)	$\Pr(F \mid H)$	diagnosis	evidence	3.3.3
Which inspection grade is a flagged coil in?	(2.15)	$\Pr(S = s \mid H)$	diagnosis	evidence	3.4.1
Will the unit run the next year without a corrective failure?	(2.13)	$\int \mathrm{e}^{-\lambda\tau} p(\lambda \mid \mathbf{y})\, \mathrm{d}\lambda$	prediction under a threshold	posterior	3.4.2
A fouled coil was found; should the crew open the fan?	(2.12)	$\Pr(F_B \mid F_A)$	prediction under a threshold	posterior	3.4.6
Does the hall exceed its limit after the test?	(2.12)	$\Pr(T > T_{\max} \mid y_T)$	prediction under a threshold	posterior	3.5.5
How tight is the health index from two tests by one sensor?	(2.9)	$\sqrt{\Sigma_{gg}}$ with $\mathbf{V}_y$ full, or b_T in the state	estimation	posterior	3.5.7
How many corrective work orders next year, and does a longer log help?	(2.10)	$\mathrm{Var}(K_\tau \mid \mathbf{y})$, (3.35)	prediction under a threshold	posterior	3.6.2

Asset manager's question	General form	Mathematical form	Category	Computation	Answered in
How well is the change since commissioning known?	(2.10)	$\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$, $\mathbf{g} = (-1,\ 1)^{\mathsf{T}}$	estimation	posterior	3.6.4
Clean path or blocked filter; what is the hall's exceedance?	(2.14)	(3.39)	hypothesis comparison	evidence	3.6.6
Chapter 4. The Anatomy of Probabilistic Inference					
What is the thermal resistance of the air path at commissioning, and how well is it known?	(2.9)	$p(z_i \mid \mathbf{y})$: $\hat{z}_i$ and Σ_{ii}	estimation	posterior	4
What will the hall temperature be at full load?	(2.10)	$f(\hat{\mathbf{z}})$ and $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ by the delta method	estimation	posterior	4
Can the hall be released for 60 kW within the 28°C acceptance limit?	(2.12)	$\Pr(T_2 < T_{\max} \mid \mathbf{y}) = \Phi\big((T_{\max} - \hat{z}_3)/\sqrt{\Sigma_{33}}\big)$	prediction under a threshold	posterior	4
Is the conductance falling between two tests, and how fast?	(2.21)	$\Pr(\gamma < 0 \mid \mathbf{y})$: a normal tail on the trend entry; $\mathrm{Var}(\gamma)$ from the two windows' variances and $n_{\mathrm{w}}\sigma_q^2$	prediction under a threshold	trend	4.6.5
Chapter 5. Health Indices from Physical Parameters					
How far has this coil's conductance fallen from its commissioned value?	(2.9)–(2.11)	$p(h_{u,18} \mid \mathbf{y})$, $\ln h_{u,18} = g_{u,18} - g_{u,0}$: median and 5–95 % band	estimation	posterior	5.3.1
Has it lost any conductance at all, beyond what the readings' scatter allows?	(2.12)	$\Pr(g_{u,18} - g_{u,0} < 0 \mid \mathbf{y})$, a normal tail	prediction under a threshold	posterior	5.3.2
Which coil has fallen more, and how sure is that?	(2.12)	$\Pr(\ln h_{2,18} < \ln h_{1,18} \mid \mathbf{y})$	prediction under a threshold	posterior	5.3.2
Is the dashboard's 21 % and 29 % the coils, or the load and the season?	(2.10), (2.11)	$p(h_u^{\mathrm{raw}} \mid \mathbf{y})$, $h_u^{\mathrm{raw}} = A_{u,0}(\mathbf{z})/A_{u,18}(\mathbf{z})$	estimation	posterior	5.3.3
Do the coils still carry the hall's capacity scenario today?	(2.12)	$\Pr(G_{1,18} + G_{2,18} \geq G^* \mid \mathbf{y})$	prediction under a threshold	posterior	5.3.4
Chapter 6. Hidden Degradation Under Control					
The unit still meets its set-point; is it working harder to do so?	(2.12), (2.10), (2.11)	$\Pr(\eta_{A,u} > 0 \mid \mathbf{y})$, $\eta_{A,u} = A_{u,18}(\mathbf{z}) - \breve{A}_u(\mathbf{z})$: median and 5–95 % band	prediction under a threshold	posterior	6.3.1
How much colder must the chilled water run, because of the coils?	(2.9)–(2.11)	$p(\eta_s \mid \mathbf{y})$, $\eta_s = T_{r,18} - \breve{T}_r$; $\Pr(\eta_s > 0 \mid \mathbf{y})$	estimation	posterior	6.3.2

Asset manager's question	General form	Mathematical form	Category	Computation	Answered in
Why does the cold aisle still read as it did at handover?	(2.10), (2.11)	$T_{c,18} - T_{c,0} = s_{\text{season}} + s_{\text{load}} + \eta_s$, each term's $p(\cdot \mid \mathbf{y})$	estimation	posterior	6.3.3
Which unit carries the heat the fouled coil no longer passes?	(2.12)	$\Pr(\eta_{D,1} > 0 \mid \mathbf{y})$, $\eta_{D,1} = -\eta_{D,2}$	prediction under a threshold	posterior	6.3.4
Chapter 7. Distinguishing Degradation from Sensor Drift					
The approach temperature has risen 1.5 K; is the exchanger fouled or is the sensor drifting?	(2.14), (2.15)	$B(M_{\text{foul}} : M_{\text{drift}})$ by the Laplace evidence; $\Pr(M \mid \mathbf{y})$ over M_{neither}, M_{foul}, M_{drift}, M_{both}	hypothesis comparison	evidence	7.3.1
How much of the 1.5 K is in question at all, once the load, the season and the other coil are taken out?	(2.10)	$A_{1,25} - A_{1,0}$ split into heat, coil 2, coil 1, sensor and reading errors	estimation	posterior	7.3.2
Could it be both, and how much of each?	(2.10), (2.11)	$p(\ln h_{1,25}, b_{d,1} \mid \mathbf{y}, M_{\text{both}})$	estimation	posterior	7.3.3
What is coil 1's health index once the drift is allowed for, and has it lost conductance?	(2.11), (2.12)	$p(h_{1,25} \mid \mathbf{y}, M_{\text{both}})$; $\Pr(\ln h_{1,25} < 0 \mid \mathbf{y}, M_{\text{both}})$; the model average beside it	estimation	posterior	7.3.4
Would the readings themselves have flagged a faulty sensor?	(2.12)	$\Pr(\chi^2_{n_r - n_z} > 2\mathcal{L})$ and the largest $\lvert r_k^{\text{N}} \rvert$ under Chapter 5's drawing	prediction under a threshold	posterior	7.3.5
Chapter 8. Battery State of Health					
What is the usable capacity now, and the internal resistance?	(2.9)–(2.11)	$p(C_{\text{bat},36}, R_{\text{bat},36} \mid \mathbf{y})$: the correlation first, then each median and 5–95 % band	estimation	posterior	8.3.1
How much capacity has the string lost since commissioning, and how far has its resistance risen?	(2.10), (2.11), (2.12)	$\ln h_C = \xi_{36} - \xi_0$, ΔR; $\Pr(\Delta R > 0 \mid \mathbf{y})$	estimation; prediction under a threshold for $\Pr(\Delta R > 0)$	posterior	8.3.2
Is the dashboard's 83.2 % the capacity?	(2.10), (2.11)	$p(T_{\text{run},36}/T_{\text{run},0} \mid \mathbf{y})$: $\ln h_C$ plus the log-ratio of the loaded heads	estimation	posterior	8.3.3
Is the string still above the replacement criterion of 80 %?	(2.12)	$\Pr(h_C < 0.80 \mid \mathbf{y})$; $\Pr(C_{\text{bat},36} < 80\,\text{Ah} \mid \mathbf{y})$	prediction under a threshold	posterior	8.3.4
Chapter 9. Fleet Pooling and Hierarchical Models					
Where does this unit sit among its fifty siblings?	(2.9)–(2.11)	$p(\mathcal{K}_0 \mid \mathbf{y})$; $p(\pi_0 \mid \mathbf{y})$, from composition draws	estimation	posterior	9.3.1
Is it fouling faster than the design's typical coil, or as slowly as the chart says?	(2.12)	$\Pr(\gamma_0 < \mu_\gamma \mid \mathbf{y})$; $\Pr(\pi_0 > 0.9 \mid \mathbf{y})$; $\Pr(\gamma_0 > \hat{\gamma}_0^{\text{chart}} \mid \mathbf{y})$	prediction under a threshold	posterior	9.3.1

Asset manager's question	General form	Mathematical form	Category	Computation	Answered in
How fast is it fouling, and what is its health index today, read with the fleet?	(2.9)–(2.11)	$p(\gamma_0 \mid \mathbf{y})$ by mixture quantiles; $d_0(3)$, $h_0(3) = \mathrm{e}^{3\gamma_0}$	estimation	posterior	9.3.2
How fast does this design foul, how alike are its coils, and what does one coil tell the fleet?	(2.9), (2.13)	$p(\mu_\gamma, \tau_\gamma, \mu_\mathrm{c}, \tau_\mathrm{c} \mid \mathbf{y})$; $p(\gamma_\mathrm{new} \mid \mathbf{y})$	estimation, prediction	posterior	9.3.3
Are the coils alike enough to share one line, or different enough to need their own?	(2.14)	$\ln p(\mathbf{y} \mid M)$ of the separate, hierarchical and fully pooled drawings	hypothesis comparison	evidence	9.3.4
Chapter 10. Degradation Rates from Condition Records					
How many percent of conductance per year is it losing?	(2.21), first line; (2.9), (2.11)	$p(\gamma_u \mid \mathbf{y})$, read as $100(\mathrm{e}^{\gamma_u} - 1)\,\%$/yr: median and 5–95 % band, $u = 1, 2$	estimation	trend	10.3.1
Is it losing conductance at all, beyond three years of scatter?	(2.21), first line	$\Pr(\gamma_u < 0 \mid \mathbf{y})$	prediction under a threshold	trend	10.3.1
Which coil is losing conductance faster, and how sure is that?	(2.12)	$\Pr(\gamma_2 < \gamma_1 \mid \mathbf{y})$	prediction under a threshold	trend	10.3.2
Is the dashboard's 15 % a year the coils, or the load, the season and the sensor?	(2.10)	$s_u^{\mathrm{raw}}(\mathbf{z})$, the slope of the line through $h_{u,t}^{\mathrm{raw}}(\mathbf{z})$, split into load and season, coils, sensor and reading errors	estimation	trend	10.3.3
How far has each coil fallen since commissioning, read with the whole record?	(2.9), (2.11)	$p(h_{u,36} \mid \mathbf{y})$, $\ln h_{u,36} = g_{u,36} - g_{u,0}$	estimation	posterior	10.3.4
Does the record still need the drift of unit 1's sensor, and how large is it now?	(2.14); (2.9)	$B(\text{drift} : \text{none})$; $p(b_{d,1}(36) \mid \mathbf{y})$	hypothesis comparison	evidence	10.3.5
Chapter 11. Accelerating Degradation					
Is the loss accelerating?	(2.12) on the curvature; (2.21), first line, one order up	$\Pr(\dot{\gamma}_u < 0 \mid \mathbf{y}, M_{12})$, $u = 1, 2$	prediction under a threshold	trend	11.3.1
Does the record need a changing rate at all, or does the constant rate do?	(2.14), (2.47)	$B(M_u : M_0)$; $\Pr(M \mid \mathbf{y})$ over $\{M_0, M_1, M_2, M_{12}\}$; accelerating, constant and slowing as the family's split	hypothesis comparison	evidence	11.3.2
How fast is it losing now, against the record's average?	(2.9)–(2.11), over the family	$\gamma_u(3) = \gamma_u + \dot{\gamma}_u\,(3 - \bar{t}_\mathrm{y})$ in %/yr, averaged over the four drawings	estimation	trend	11.3.3

Asset manager's question	General form	Mathematical form	Category	Computation	Answered in
Is the dashboard's bend in year 3 the coil?	(2.10)	$\dot{s}_u^{\text{raw}}(\mathbf{z})$, the bend of the quadratic through $h_{u,t}^{\text{raw}}(\mathbf{z})$, split into load and season, coils, sensor and reading errors	estimation	trend	11.3.4
With the coil allowed to accelerate, does unit 1's sensor still need its drift?	(2.14)	$B(\text{drift} : \text{none} \mid M_1)$; the drift alone, the acceleration alone, both	hypothesis comparison	evidence	11.3.5
Chapter 12. Change-Point Detection					
Did something happen in March?	(2.15) at the named window; (2.14)	$\Pr(t_0 = 15 \mid \mathbf{y}) = \bar{B}_{15}/(1+\bar{B}_{15})$, $\bar{B}_{15} = \frac{1}{8}\sum_j B_{j,15}$	diagnosis	evidence	12.3.1
What stepped: a coil, or a sensor reading it?	(2.15)	$\Pr(M_{j,15} \mid \mathbf{y}, t_0 = 15)$ over the eight kinds j	diagnosis	evidence	12.3.2
Was it March, or the month either side?	(2.15) over windows	$\Pr(t_0 = s \mid \mathbf{y}, \text{coil 2})$, $s \in \{14, 15, 16\}$	diagnosis	evidence	12.3.3
How large was the step, and how much of coil 2's loss is it?	(2.9), (2.11), (2.10)	$p(\delta_2 \mid \mathbf{y}, M_{2,15})$, as $100(\mathrm{e}^{\delta_2} - 1)\,\%$ of conductance; $\delta_2 / \ln h_{2,36}$	estimation	posterior	12.3.4
With the step allowed for, how fast is each coil losing conductance, and which is faster?	(2.21), first line; (2.12); (7.10)	$p(\gamma_u \mid \mathbf{y})$ mixed over the named-window family; $\Pr(\gamma_2 < \gamma_1 \mid \mathbf{y})$	estimation	trend	12.3.5
Did a step happen anywhere in the three years?	(2.15) over the whole record	$\Pr(\text{a step} \mid \mathbf{y}) = \bar{B}/(1+\bar{B})$ over windows 1–36; the ranking of the 288 candidates	diagnosis	evidence	12.3.6
Chapter 13. Separating Seasonal Effects from Aging					
Is the summer drop in performance degradation or just the ambient?	(2.10), (2.11)	$y_{36} - y_{31} = s_{\text{season}} + s_{\text{load}} + s_{\text{coils}} + s_{\text{err}}$, each term's $p(\cdot \mid \mathbf{y})$; on $\ln G^{\text{app}}$, $s_{\text{coils}} = \ln(G_{\text{tot},36}/G_{\text{tot},31})$	estimation	trend	13.3.1
Is more of the drop the coils than the season?	(2.12)	$\Pr(s_{\text{coils}} < s_{\text{season}} \mid \mathbf{y})$ on $\ln G^{\text{app}}$	prediction under a threshold	trend	13.3.1
What does the trend line through the raw performance read, and how much of it is the season?	(2.10)	$s^{\text{line}} = (\mathbf{w}^{\text{line}})^{\mathsf{T}}\mathbf{y}^{\text{raw}}$ over windows 0 to t_{as}, split into season (the leak), load, coils and errors	estimation	trend	13.3.2
Does a seasonal term in the regression take the season out?	(2.10)	$s^{\text{seas}} = (\mathbf{w}^{\text{seas}})^{\mathsf{T}}\mathbf{y}^{\text{raw}}$, split the same way; its fitted harmonic against the known inputs'	estimation	trend	13.3.3
How fast is the hall's conductance falling, with the season taken out by the physics?	(2.21), first line; (2.11)	$p(\Gamma \mid \mathbf{y})$, $\Gamma = \ln(G_{\text{tot},36}/G_{\text{tot},0})$ over $36\,\Delta t$, as $100(\mathrm{e}^{\Gamma} - 1)\,\%/\text{yr}$	estimation	trend	13.3.4
Does the coils' conductance itself fall in summer?	(2.14)	$B(M_{\text{sea}} : M_{\text{base}})$; $p(g_u^c, g_u^s \mid \mathbf{y})$	hypothesis comparison	evidence	13.3.5

Asset manager's question	General form	Mathematical form	Category	Computation	Answered in
Chapter 14. Remaining Useful Life and Survival to the Next Outage					
When does the conductance cross the level at which the hall can no longer be held?	(2.21), second line; (2.11) by inversion; (7.10)	$\Pr(t^* \le s \mid \mathbf{y}) = \sum_M P_M \Pr(t^* \le s \mid \mathbf{y}, M)$: median and 5–95 % band, as windows and dates; the remaining life $t^* - 36$	estimation	trend	14.3.1
Will it hold until the planned shutdown in October?	(2.21), second line at $s = t_{\text{sd}}$	$\Pr(t^* > 46 \mid \mathbf{y})$, with $\Pr(t^* \le 45 \mid \mathbf{y})$ beside it	prediction under a threshold	trend	14.3.2
How much margin will be left at the shutdown?	(2.21), third line	$G_{\text{tot},46} - G^*$: median and band	estimation	trend	14.3.3
Which open question moves the date, the acceleration or the step?	(2.9), (7.10), (7.11)	the forecast within each factor of the family; the band's width within the drawings against the mixture's	estimation	evidence	14.3.4
Is the hall already past its limit, as the dashboard says?	(2.12)	$\Pr(G_{\text{tot},36} \le G^* \mid \mathbf{y})$	prediction under a threshold	posterior	14.3.5
Chapter 15. Failure Probability over a Mission Window					
Given its rate today, what is the probability it fails in the next 90 days?	the predictive, (2.13), in its event form (2.8)	$\bar{p}_\tau = \Pr(T_f \le \tau \mid \mathbf{y}) = \mathbb{E}[1 - \mathrm{e}^{-\lambda\tau} \mid \mathbf{y}]$, beside the plug-in $1 - \mathrm{e}^{-\hat{\lambda}\tau}$	prediction under a threshold	posterior	15.3.1
What is its failure rate today, and its MTBF?	(2.9), (2.11)	$p(\lambda \mid \mathbf{y})$, $\lambda = \lambda_{\text{ref}}\mathrm{e}^{\omega}$, and $1/\lambda$: medians and 5–95 % bands	estimation	posterior	15.3.2
How well is that probability known, and is it below 5 % with 95 % assurance?	(2.12)	$p(p_\tau(\lambda) \mid \mathbf{y})$, median and band; $\Pr(p_\tau(\lambda) < p^* \mid \mathbf{y}) \ge \kappa$	prediction under a threshold	posterior	15.3.3
Will it run to the October shutdown, and in which window does its failure probability pass 5 %?	the predictive, (2.13); its inverse, (2.9)	$1 - \bar{p}_{\tau_{\text{sd}}}$, $\tau_{\text{sd}} = 0.75\,\text{yr}$; $\tau_p(\lambda) = -\ln(1 - p^*)/\lambda$, median and inversion band, as a window	prediction under a threshold	posterior	15.3.4
It has never failed: is its record evidence that it is better than its siblings?	(2.14)	$\ln B$ of its own rate against the shared rate, on its year in the band; Pr(three clean years | the siblings)	hypothesis comparison	evidence	15.3.5
Chapter 16. Age-Dependent Failure Rates Across a Fleet					
Does the failure rate rise with age, and where is this unit on that curve?	the trend form over age, its first line, (2.21); moments, credible interval, (2.9)–(2.11)	$\Pr(\theta_a > 0 \mid \mathbf{y})$; $p(\lambda(a) \mid \mathbf{y})$ along age; $p(\lambda(3) \mid \mathbf{y})$ and $\Pr(\lambda(3) < 0.1107 \mid \mathbf{y})$	prediction under a threshold; estimation	trend	16.3.1
How fast does it rise, and in how many years does it double?	moments, credible interval; a ratio by inversion, (2.11)	$100(\mathrm{e}^{\theta_a} - 1)$ in %/yr; $t_{\text{dbl}} = \ln 2/\theta_a$, median and inversion band, with $\Pr(\theta_a \le 0 \mid \mathbf{y})$ at $t_{\text{dbl}} = \infty$	estimation	trend	16.3.2

Asset manager's question	General form	Mathematical form	Category	Computation	Answered in
What rate should I plan with for coil 0 over the next year and the next 90 days, and for the oldest coils?	moments; the predictive, (2.13)	$\int_3^4 \lambda(a)\,\mathrm{d}a$; $\mathbb{E}[1-\mathrm{e}^{-\int_3^{3+\tau}\lambda(a)\,\mathrm{d}a} \mid \mathbf{y}]$; $\lambda(46/12)$; $\lambda(11.5)$	estimation; prediction under a threshold	trend	16.3.3
Is one constant rate good enough, or does the record need the age curve?	$B_{\text{age}} = p(\mathbf{y} \mid M_{\text{age}})/p(\mathbf{y} \mid M_{\text{const}})$, (2.14)	$\ln B$(age curve : constant), $\theta_a \sim \mathcal{N}(0, s_\theta^2)$, over s_θ	hypothesis comparison	evidence	16.3.4
Do the coils in worse condition fail more, once age is counted?	moments; B_{age}, (2.14)	$p(\theta_g \mid \mathbf{y})$ with and without θ_a; $\ln B$ of the age, derate and age-plus-derate drawings	estimation; hypothesis comparison	posterior; evidence	16.3.5
Chapter 17. Repair or Replace					
Is restoring this unit's conductance worth more than its replacement, given what the remaining life band says?	(2.23)	$\mathcal{R}(a) = \min_{\mathcal{P}\in a} \mathbb{E}[L(\mathcal{P}, \mathbf{z}) \mid \mathbf{y}]$ for $a \in$ {clean now, clean at the shutdown, replace now, replace at the shutdown, wait}; $a^* = \arg\min_a \mathcal{R}(a)$	design	expected consequence	17.3.1
How firm is that choice, and what would knowing the path be worth?	(2.24), second line; perfect information	$\Pr(L(\mathcal{P}^*, \mathbf{z}) < L(\mathcal{P}, \mathbf{z}) \mid \mathbf{y})$ for each action's best plan $\mathcal{P}$; $\text{EVPI} = \mathcal{R}(a^*) - \mathbb{E}[\min_{\mathcal{P}} L(\mathcal{P}, \mathbf{z}) \mid \mathbf{y}]$	design	expected consequence	17.3.2
Should the work be done now, or can it wait for the October shutdown?	(2.21), third line; (2.23)	$\text{SF}\sum_{t=37}^{46}\text{df}(t) \times \Pr(G_{\text{tot},t} \le G^* \mid \mathbf{y})$ with no work, against OW	prediction under a threshold; design	trend; expected consequence	17.3.3
How long does each action hold the hall?	(2.21), second line, from the visit	$t_c(\mathbf{z}; \mathcal{P}_{\text{v}}) = \min\{t > t_{\text{v}} : G_{\text{tot},t}(\mathbf{z}; \mathcal{P}_{\text{v}}) \le G^*\}$: median and band	estimation	trend	17.3.4
What do the dashboard's and the calendar's schedules cost?	(2.23)	$\mathcal{R}(\mathcal{P}_{\text{dash}})$, $\mathcal{R}(\mathcal{P}_{\text{cal}})$ against $\mathcal{R}(a^*)$; $\Pr(L(\mathcal{P}^*, \mathbf{z}) < L(\mathcal{P}_{\text{cal}}, \mathbf{z}) \mid \mathbf{y})$	design	expected consequence	17.3.5
Chapter 18. Inspection Intervals					
How often should the exchanger be inspected so that the probability of an undetected crossing stays under 5 %?	(2.23) with the indicator of an undetected crossing as the loss; (2.21), second line	$P_{\text{U}}(\Delta_{\text{I}}) = \mathbb{E}\big[\prod_{s\in\mathcal{S},\, s<t^*} p_{\text{no},s} \mid \mathbf{y}\big]$, averaged over the phase s_1; $\Delta_5 = \max\{\Delta_{\text{I}} : P_{\text{U}}(\Delta') \le 0.05\ \forall\, \Delta' \le \Delta_{\text{I}}\}$	design	trend; expected consequence	18.3.2
How long is this hall's P–F interval?	(2.21), second line at two levels; (2.11)	$\Delta_{\text{PF}} = t^* - t_{\text{P}}$: median and 5–95 % band	estimation	trend	18.3.1
What does each interval cost, which is cheapest, and how firm is that choice?	(2.23); (2.24); perfect information, (17.9)	$\mathcal{R}(\Delta_{\text{I}}) = \mathbb{E}[C_V N_V + \text{SF}\, N_{\text{sf}} \mid \mathbf{y}]$; $\Delta^* = \arg\min \mathcal{R}$; $\Pr(L(\Delta^*, \mathbf{z}) < L(\Delta_5, \mathbf{z}) \mid \mathbf{y})$; EVPI	design	expected consequence	18.3.3
What do the calendar and the rule of thumb give?	the same, at the naive actions	P_{U} and $\mathcal{R}$ at $\Delta_{\text{I}} = 12$, and at $\Delta_{\text{I}} = 3$ (half the handbook's 6 months)	design	expected consequence	18.3.4

Asset manager's question	General form	Mathematical form	Category	Computation	Answered in
If the hall is carried to the October shutdown with no work done, which tests before it does it need?	the same, with the program ended at window 46	$P_{\mathrm{U}}(\mathcal{S}) = \mathbb{E}\big[\mathbf{1}(t^* \le 46)\prod_{s\in\mathcal{S},\, s<t^*} p_{\mathrm{no},s} \mid \mathbf{y}\big]$ for dated sets $\mathcal{S} \subset \{37,\dots,45\}$ of one, two and three tests; the probability of an order before the shutdown	design	trend; expected consequence	18.3.5
Chapter 19. Condition-Based Maintenance Triggers					
At what measured value should the work order fire?	(2.23), (2.24); (2.12)	at one look, $\mathcal{R}(\text{order}) = \mathrm{OW} + \mathrm{FO}\,(1-q_{45})$, $\mathcal{R}(\text{wait}) = \mathrm{SF}\,q_{45}$; order when $q_{45} \ge c_1^*$	design	expected consequence	19.3.1
With a refit every month, which level costs least to the shutdown, and what do 0.5, 0.9 and 0.95 cost?	(2.16) with the rule's loss; (2.17)	$c^* = \arg\min_c \mathcal{R}(c)$, $\mathcal{R}(c) = \mathbb{E}[L(c, \mathbf{z}, \tilde{\mathbf{y}}_{37:45}) \mid \mathbf{y}_{0:36}]$; $\mathcal{R}(c)$ at c^*, 0.5, 0.9, 0.95; $\Pr(L(c^*) < L(c'))$; EVPI	design	expected consequence; preposterior	19.3.2
When would this record fire?	(2.16), its first look	$\Pr(T_{\mathrm{ord}} \le s \mid \mathbf{y}_{0:36})$ over the predicted records; $q_{t_{\mathrm{as}}}$ on the as-of fits of the logged record	prediction under a threshold	preposterior	19.3.3
How often does the rule fire on a coil that is not degrading?	(2.16) at a truth with no loss	$\Pr(\exists\, t_{\mathrm{as}} \le 45 : q_{t_{\mathrm{as}}} \ge c^* \mid \text{no loss})$, from today and on a restarted record; the dashboard's beside it	design	preposterior	19.3.4
Chapter 20. Prioritizing Maintenance Across a Fleet					
Fifty units, one crew; which one does the crew visit next?	expected consequence over actions, (2.23); how firm a ranking is, (2.24), second line	$\mathcal{V}^* = \arg\min_{\mathcal{V}} \mathrm{FC} \sum_i P_i(t_{\mathrm{ah},v_i})$ and its week 1; $\Pr(L(\mathcal{V}^*, \mathbf{z}) < L(\mathcal{V}', \mathbf{z}) \mid \mathbf{y})$	design	expected consequence	20.3.1
What does it cost to make each coil wait, and for how long does that hold?	expected consequence of one action, (2.23); first passage, (2.21), second line	$\mathcal{W}_i(\Delta) = \mathrm{FC}\,P_i(\Delta)$ for Δ of 1, 13 and 50 weeks	prediction under a threshold; design	trend; expected consequence	20.3.2
Which coils are already past their functional limit?	$\Pr(f(\mathbf{z}) \ge f^* \mid \mathbf{y})$, (2.12)	$\Pr(g_i(a_i) \le g_F \mid \mathbf{y})$	prediction under a threshold	posterior	20.3.3
What does the year's plan cost against the dashboard's order, and what would knowing every coil's condition be worth?	expected consequence, (2.23); perfect information, (17.9)	$\mathcal{R}(\mathcal{V}^*)$, $\mathcal{R}(\mathcal{V}_{\mathrm{dash}})$, $\Pr(L(\mathcal{V}^*, \mathbf{z}) < L(\mathcal{V}_{\mathrm{dash}}, \mathbf{z}) \mid \mathbf{y})$; $\mathrm{EVPI} = \mathcal{R}(\mathcal{V}^*) - \mathbb{E}[\min_{\mathcal{V}} L(\mathcal{V}, \mathbf{z}) \mid \mathbf{y}]$	estimation; design	expected consequence	20.3.4
Chapter 21. Warranty Claims and Specification Compliance					
Has this unit degraded faster than its specification allows, with enough confidence to make the claim?	(2.14) on the two halves of one entry; (2.12)	$P_u = \Pr(\gamma_u < \gamma_{\mathrm{S}} \mid \mathbf{y})$, mixed over the reported family; $B_u^{\mathrm{S}} = \dfrac{P_u/(1-P_u)}{P_u^0/(1-P_u^0)}$	hypothesis comparison	trend	21.3.1
Should the claim be filed, at R$ 5 000 to win R$ 60 000?	(2.23), (2.24)	$\mathcal{R}(\text{file}) = \mathrm{CF} - \mathrm{CV}\,P_u$, $\mathcal{R}(\text{not}) = 0$; a^*; $\Pr(a^* \text{ better} \mid \mathbf{y})$; EVPI	design	expected consequence	21.3.2

Asset manager's question	General form	Mathematical form	Category	Computation	Answered in
What do the maker's arguments do to the claim: the March step, the drifting sensor, the summer?	(2.14), each argument a drawing; (2.12) under each	$\Pr(\gamma_u < \gamma_{\mathrm{S}} \mid \mathbf{y}, M_{\mathrm{arg}})$; $\ln B(M_{\mathrm{arg}} : M_{\mathrm{claim}})$	hypothesis comparison	evidence	21.3.3
Read as the loss so far, or over the specification's five years, is it still faster?	(2.21) (line 1), (2.12)	$\Pr(\gamma_u^{\mathrm{td}} < \gamma_{\mathrm{S}} \mid \mathbf{y})$, $\Pr(\gamma_u^{5\mathrm{y}} < \gamma_{\mathrm{S}} \mid \mathbf{y})$	prediction under a threshold	trend	21.3.4
Chapter 22. Detectability of Degradation					
Can the existing sensors resolve a 5 % per year loss in the time I have?	(2.16) with the indicator of the rule (2.12) as the utility, conditioned on the rate: (22.1)	$\mathrm{POD}_u(45) = \Pr\big(\Pr(\gamma_u < 0 \mid \tilde{\mathbf{y}}_{0:45}) \geq \kappa \mid \gamma_u = \gamma_{\mathrm{D}}\big)$, and its curve over t_{as}	design	preposterior	22.3.1
How long a record does a power of 0.9 need, and what would better sensors save?	the same, solved for the record	$T_{90,u} = \min\{t_{\mathrm{as}} : \mathrm{POD}_u(t_{\mathrm{as}}) \geq \mathrm{POD}^*\}$, over the temperature sensors' grade σ_T and over (κ, POD^*)	design	preposterior	22.3.2
What is the smallest loss the record resolves by the shutdown?	the same, solved for the stated rate	the smallest $\lvert\gamma\rvert$ with $\mathrm{POD}_u(45) \geq \mathrm{POD}^*$ at $\gamma_u = \gamma$, in %/yr	design	preposterior	22.3.3
If I refit every month and act on the first verdict, when does it come, and how often is it false?	the same, with the rule read at every look	$\Pr\big(\exists\, t_{\mathrm{as}} \leq 45 : \Pr(\gamma_u < 0 \mid \tilde{\mathbf{y}}_{0:t_{\mathrm{as}}}) \geq \kappa \mid \gamma_u\big)$ at γ_{D} and at 0; the first such t_{as}	design	preposterior	22.3.4
Does a stricter rule, a Bayes factor of 19 against no trend, change the answer?	(2.16) with $\mathbf{1}[B \geq 19]$ as the utility, B by (2.14)	$\Pr\big(B(\mathrm{trend} : \mathrm{none}) \geq 19 \mid \gamma_u = \gamma_{\mathrm{D}}\big)$ on $\tilde{\mathbf{y}}_{0:t_{\mathrm{as}}}$; $T_{90,u}^{B}$; its value at $\gamma_u = 0$	design; hypothesis comparison	preposterior; evidence	22.3.5
Chapter 23. Sensor Selection for Prognostics					
Which added measurement most tightens the remaining-life band?	(2.16) with U_{p} the band's width; (2.17); (7.10)	$\mathcal{B}_a$ of (23.1), the band of t^* from window 45 under each $\boldsymbol{\Sigma}_{M,a}$ mixed at the held P_M; $\Delta\mathcal{B}_a$; the own bands $\bar{\mathcal{B}}_a$ beside it	design	preposterior	23.3.1
Which one makes the call at the shutdown surest?	(2.16); (2.21), third line	the 5–95 % band of $G_{\mathrm{tot},46}$ under each $\boldsymbol{\Sigma}_{M,a}$	design	preposterior; trend	23.3.2
Which is worth its price?	(2.16) with the price c_a beside the width	$\Delta\bar{\mathcal{B}}_a/c_a$ per R$ 1 000; the price at which a candidate matches the best	design	preposterior	23.3.3
Why does a sensor on unit 1 narrow the date and two more return-air sensors do not?	(2.17); (2.9) with a direction known	the delta form's $\mathrm{Var}\, t^*$ under the continuation, conditioned on each direction known; the spread of $G_{\mathrm{tot},45}$ under each candidate; one window's conductance by each route	estimation	preposterior	23.3.4

Asset manager's question	General form	Mathematical form	Category	Computation	Answered in
How much of the band can any sensor reach, and what is left?	(2.17); (7.10); (2.14)	the floors (rates known; rates and path known); $\mathcal{B}_a$ against $\bar{\mathcal{B}}_a$; the share of predicted records that decide coil 1's curvature	design; hypothesis comparison	preposterior; evidence	23.3.5
Chapter 24. Sampling Frequency and Correlated Readings					
Hourly, daily, or monthly?	(2.16) with U_{p} the spread; (2.17)	$\sqrt{[\boldsymbol{\Sigma}_{n_{\mathrm{L}}}]_{\gamma_u\gamma_u}}$ for windows 0–36 logged at n_{L} readings, against the wander floor $\sigma_q/(\sqrt{36}\,\Delta t)$; the band's width in %/yr	design	preposterior; trend	24.3.1
How many independent readings does a month of logging hold?	(2.16); (3.29) with the readings' covariance full	$n_{\mathrm{eff}}(n_{\mathrm{L}}) = \mathbf{1}^{\mathsf{T}}\mathbf{R}^{-1}\mathbf{1}$ for the AR(1) correlation $\mathbf{R}$, and its limit as $n_{\mathrm{L}} \to \infty$	design	preposterior	24.3.2
From which rate on does logging stop narrowing the rate's band, and why?	(2.16); (4.74)	the readings' share $[\boldsymbol{\Sigma}^{\mathrm{line}}_{n_{\mathrm{L}}}]_{\gamma_u\gamma_u}/[\boldsymbol{\Sigma}_{n_{\mathrm{L}}}]_{\gamma_u\gamma_u}$ over n_{L}; the first n_{L} with a share of 10 % or less, and of 5 %	design	preposterior; trend	24.3.3
Does a faster rate help today's health index and the capacity margin?	(2.16)	$\sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_{n_{\mathrm{L}}}\mathbf{g}}$ for $\ln h_{u,36} = g_{u,36} - g_{u,0}$ and for $G_{1,36} + G_{2,36}$	design	preposterior	24.3.4
What does the band look like if the logged readings are treated as independent?	(2.16); (3.29), the covariance diagonal against full	the same spreads with $n_{\mathrm{eff}} = n_{\mathrm{L}}$, against the rows' own; their ratio	design	preposterior	24.3.5
What should the logger be set to for year 4, against a crew visit at the same price?	(2.16), a candidate's cost a loss of its own	$\sqrt{[\boldsymbol{\Sigma}_a]_{\gamma_u\gamma_u}}$ and $\sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_a\mathbf{g}}$ for $\ln h_{u,45}$ and $G_{1,45} + G_{2,45}$, windows 37–45 logged at each rate, beside the crew visits	design	preposterior	24.3.6