

Operations and Control by Derivation

Scenarios, interventions and the decisions of one night shift

Draft · 220 pages

Ricardo Marquez, Ph.D.

ricardo-marquez.com

Operations and Control by Derivation

Volume II

Operator actions, model validation and alarms

Ricardo Marquez

November 1, 2026

How to read this book

This book teaches one skill: to take an operator's question stated in the language of the control room, write it as a probability or an expectation over the unknowns of a physical model, derive what that expression requires, compute it, check it, and act on it. Volume I holds Parts I to III, Chapters 1 to 11: the method itself, the state of the plant read from its telemetry (virtual sensors, unit heats, meters, the operating envelope), and the detection and diagnosis of abnormal operation. This second volume holds Parts IV to VI, Chapters 12 to 20. It asks what happens if the operator acts (a set point, a unit out, a load moved, nothing done), then how far the model can be trusted tonight (goodness of fit, sensor loss, the next field measurement), and then where the alarms should sit and how far an alarm should be believed. Chapter numbers run continuously across the two volumes; a reference to a chapter numbered 11 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 and restated here in one line each: 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 telemetry attached as rows; step 3, the posterior of the state, derived and computed once; step 4, the answers, each question read from that posterior down to its number; step 5, evaluation and checks against the full calculation and the constructed truth; step 6, the report with its band; step 7, the next reading, ranked by the width it removes. The five kinds of question and the five computations beneath them are those of Volume I.

The plant of this volume is two data halls on one chilled-water loop, followed through a twelve-hour night shift in five-minute windows: six cooling units, ten racks, sixteen readings. On that plant many of the questions are the same posterior read at a set point not yet applied, and the chapters say so where it happens.

The framework the book teaches is physics-constrained probabilistic modeling: a plant written as a graph, its unknowns as a state vector, every reading as a row, and every question as a probability over that state. TERRA is the author's own implementation of this framework; the book names the framework and the computations, not the software, and every number in it is reproducible from the derivations as written. The notation chapter is reproduced here so that this volume's symbols can be read with Volume I closed, and the atlas of questions in the appendix covers Volumes I and II.

Each chapter closes with a statement of what to retain, a table of its terms, and exercises at three levels, formulation, derivation and application, each followed by its answer.

Contents

How to read this book **ii**

Notation and vocabulary **vi**

IV Predicting the Effect of Operator Actions **1**

12 Set-point changes as scenarios **2**
12.1 A raise of the chilled water in the night . . . 2
12.2 The posterior of the state (step 3) . . . 8
12.3 The answers (step 4) . . . 13
12.4 The checks (step 5) . . . 18
12.5 The report (step 6) . . . 21
12.6 The next reading (step 7) . . . 22
12.7 What to retain . . . 24
12.8 Exercises . . . 25

13 Equipment outages and maintenance windows **26**
13.1 A hall with a unit to spare . . . 26
13.2 The posterior of the state (step 3) . . . 32
13.3 The answers (step 4) . . . 34
13.4 The checks (step 5) . . . 41
13.5 The report (step 6) . . . 43
13.6 The next reading (step 7) . . . 44
13.7 What to retain . . . 46
13.8 Exercises . . . 47

14 Load transfer between halls **48**
14.1 Forty kilowatts between two halls . . . 48
14.2 The posterior of the state (step 3) . . . 51
14.3 The answers (step 4) . . . 52
14.4 The checks (step 5) . . . 59
14.5 The report (step 6) . . . 61
14.6 The next reading (step 7) . . . 62
14.7 What to retain . . . 64
14.8 Exercises . . . 65

15 Look-ahead risk without intervention **66**
15.1 The hours before the morning shift . . . 66
15.2 The posterior of the state (step 3) . . . 71
15.3 The answers (step 4) . . . 74

15.4 The checks (step 5) . . . 83
15.5 The report (step 6) . . . 86
15.6 The next reading (step 7) . . . 87
15.7 What to retain . . . 89
15.8 Exercises . . . 90

V Validating the Model in Operation 92

16 Goodness of fit in operation 93
16.1 A fit-quality light that stays green . . . 93
16.2 The posterior of the state (step 3) . . . 97
16.3 The answers (step 4) . . . 99
16.4 The checks (step 5) . . . 108
16.5 The report (step 6) . . . 109
16.6 The next reading (step 7) . . . 111
16.7 What to retain . . . 112
16.8 Exercises . . . 113

17 Sensor loss and measurement redundancy 114
17.1 A work order for the return sensor . . . 114
17.2 The posterior of the state (step 3) . . . 118
17.3 The answers (step 4) . . . 119
17.4 The checks (step 5) . . . 131
17.5 The report (step 6) . . . 132
17.6 The next reading (step 7) . . . 134
17.7 What to retain . . . 135
17.8 Exercises . . . 136

18 Selecting the next field measurement 137
18.1 One technician and one trip . . . 137
18.2 The posterior of the state (step 3) . . . 140
18.3 The answers (step 4) . . . 141
18.4 The checks (step 5) . . . 148
18.5 The report (step 6) . . . 150
18.6 The next reading (step 7) . . . 151
18.7 What to retain . . . 152
18.8 Exercises . . . 153

VI Alarm Thresholds and Alarm Credibility 154

19 Calibrating alarm thresholds 155
19.1 An alarm that has never fired . . . 155
19.2 The posterior of the state (step 3) . . . 159
19.3 The answers (step 4) . . . 162
19.4 The checks (step 5) . . . 172
19.5 The report (step 6) . . . 173
19.6 The next reading (step 7) . . . 176
19.7 What to retain . . . 177
19.8 Exercises . . . 178

20 Alarm credibility and base rates **179**
20.1 An alarm that has just fired . . . 179
20.2 The posterior of the state (step 3) . . . 183
20.3 The answers (step 4) . . . 184
20.4 The checks (step 5) . . . 193
20.5 The report (step 6) . . . 195
20.6 The next reading (step 7) . . . 195
20.7 What to retain . . . 197
20.8 Exercises . . . 198

A Atlas of operator questions **199**

Notation and vocabulary

Every symbol below has the same meaning in every chapter, except where a chapter's table says otherwise. The unknown always stands on the left of a conditioning bar and the evidence on the right. Vectors and matrices are bold. On the two-unit hall of Parts II and III the units are $u \in \{1, 2\}$ and the racks $k \in \{1, 2, 3\}$. On the shift plant of Parts IV to VI, hall subscripts are $h \in \{A, B\}$, cooling-unit subscripts $u \in \{A1, A2, A3, B1, B2, B3\}$ and rack subscripts $k \in \{A1, \ldots, A5, B1, \ldots, B5\}$; the hall of a unit or a rack is the letter of its name.

Symbol	Unit	Meaning
The plant graph and its state		
$\mathbf{z}$, z_i	varies	State vector of the plant graph: every free quantity of one telemetry window at once; its i-th entry.
$\mathbf{z}_t$, t	varies, –	The state of window t; the window index, each window five minutes long: on the shift plant $t = 1, \ldots, 144$, window 1 starting at 19:00; on the two-unit hall's run $t = 1, \ldots, 36$, window 1 starting at 23:00.
$\mathbf{y}$, y	varies	The evidence: all readings and prior information treated as known; a single reading, subscripted by what it reads (y_{T_s}, y_{Q_A}).
$r_k(\mathbf{z})$	varies	Row k of the model: a physical relation that must hold $(r_k = 0)$, a reading compared with the state, or a prior.
w_k, σ_k	varies	Weight (precision) of row k and the standard deviation it corresponds to, $w_k = 1/\sigma_k^2$. A row may equivalently be written normalized, r_k/σ_k, with unit weight.
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; its reference value.
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–III and the shift plant		
T_s	°C	Chilled-water supply temperature; the actuator, set by the chiller set point.
Q_h, $Q_{o,h}$	kW	IT heat load of hall h; its other heat (fans, lighting), pinned.

Symbol	Unit	Meaning
G_u, $\bar{G}_u$	kW/K	Effective conductance of the coil of cooling unit u; its rated value, pinned.
g_u, d_u	–	Log-derate $\ln(G_u/\bar{G}_u)$, the entry of the state that carries the unit's conductance; the derate factor $d_u = \mathrm{e}^{g_u}$, one at the rating.
C_u	kW/K	Air capacity rate of unit u (air mass flow times specific heat), pinned at its fixed fan speed.
Q_u	kW	Heat unit u rejects to the chilled-water loop, $G_u(T_{r,h} - T_s)$.
$T_{r,h}$, $T_{c,h}$	°C	Return (hot-aisle) and cold-aisle air temperature of hall h.
$T_{d,u}$	°C	Discharge air temperature of unit u.
T_k, β_k	°C, –	Inlet temperature of rack k; its recirculation fraction, $T_k = T_{c,h} + \beta_k(T_{r,h} - T_{c,h})$.
$T_{\max}$	°C	A limit on a temperature: the hall's 24°C now and 27.5°C at full load in Part I (28°C for the illustrative halls of Chapters 3 and 4); a rack inlet's 27°C, the recommended envelope, from Part II.
f^*, G^*, T^*	varies	A requirement level on a state entry or a derived quantity, as in the companion book: the event $f \geq f^*$ is *the requirement holds*. From Part IV a starred plant quantity (Q_A^*, T_k^*) is also a scenario value: the quantity at a set point, unit state or load not yet applied, as in the companion book's Part V.
b_j	varies	Bias of reading j, an entry freed when a meter is suspected.
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; $\Phi_2(u_1, u_2; \varrho)$ that of two standard normals of correlation ϱ.
χ^2_ν	–	Chi-square distribution with ν degrees of freedom.
Inference objects		
$\mathcal{L}(\mathbf{z})$	–	The objective, minus the log posterior up to a constant: $\mathcal{L} = \frac{1}{2} r(\mathbf{z})^\mathsf{T} r(\mathbf{z})$ for normalized Gaussian rows.
$\hat{\mathbf{z}}$	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, $\mathbf{J}^\mathsf{T}\mathbf{W}\mathbf{J}$ with $\mathbf{W} = \mathrm{diag}(w_k)$ for unnormalized ones.

Symbol	Unit	Meaning
$\boldsymbol{\Sigma}$, Σ_{ij}	varies	Posterior covariance of the state under the Laplace approximation, $\boldsymbol{\Lambda}^{-1}$; its entries. The posterior spread of entry i is written $\sqrt{\Sigma_{ii}}$; σ_k always names a row's standard deviation.
$f(\mathbf{z})$, $\mathbf{g}$	varies	A derived quantity of the state and its gradient $\partial f/\partial \mathbf{z}$ at $\hat{\mathbf{z}}$.
$\mathbf{A}$, $\mathbf{A}_{\text{free}}$, $\mathbf{A}_{\text{dep}}$	varies	Jacobian of the physics rows on the full state; its blocks on the free and the dependent entries (Chapter 4).
$\mathbf{N}$, $\mathbf{J}_{\text{rows}}$, $\mathbf{J}_{\text{red}}$	varies	Null-space basis of the physics rows, $[\mathbf{I};\ -\mathbf{A}_{\text{dep}}^{-1}\mathbf{A}_{\text{free}}]$; the Jacobian of the evidence rows (readings and priors); the reduced Jacobian $\mathbf{J}_{\text{rows}}\mathbf{N}$ of the reduced state.
f_L, f_U	varies	The ends of the credible interval of a derived quantity.
y_{new}, $\tilde{y}_a$	varies	A future reading (the predictive) and the outcome of a candidate reading a before it is taken (the preposterior).
e_k	–	Normalized residual of reading k: its residual divided by the standard deviation of that residual.
M, M_j, B_{AB}	–	A drawing (model) of the plant; the members of a family of candidates; the Bayes factor $p(\mathbf{y} \mid M_A)/p(\mathbf{y} \mid M_B)$.
$p(\mathbf{y} \mid M)$	–	Evidence (marginal likelihood) of model M.
$\mathbf{J}_a$, $\boldsymbol{\Lambda}_a$, $\boldsymbol{\Sigma}_a$	varies	The Jacobian rows of a candidate reading, its information $\mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a$, and the preposterior covariance $(\boldsymbol{\Sigma}^{-1} + \boldsymbol{\Lambda}_a)^{-1}$.
γ, t^*	varies/h, h	Slope of an entry's trend across windows; the time at which the trend crosses a limit (as in the companion book).
Alarms and decisions		
c	varies	Decision level (alarm threshold) on a monitored quantity.
α	–	False-alarm probability per window.
$U(a)$, $L(a, \cdot)$	varies	Preposterior value of a candidate reading a; the loss of an action a.
Scenario nodes (Part IV); a row naming a chapter is that chapter's own symbol		
T_s^*, T_k^*, $\mathbf{T}^*$	°C	The supply at a scenario node; the inlet of rack k there, a derived quantity of the state; the inlets stacked. A star marks the value at a set point, unit state or load not yet applied.
$T_{r,h}^*$, $T_{c,h}^*$, $T_{d,u}^*$, Q_u^*, Q_h^*	°C, kW	The return, cold aisle, discharge, unit heat and hall load at the node (Chapter 13 writes T_{rA}^* and T_{cA}^* for hall A).
$\mathbf{g}_{T_k}$, $\mathbf{g}_{T_k^*}$	varies	The gradient of the inlet of rack k at the present state and at the node.
$\mathbf{m}^*$, $\mathbf{D}^*$	°C, varies	The inlets at the node evaluated at $\hat{\mathbf{z}}$; their gradients stacked as rows, so that the inlets are jointly $\mathcal{N}(\mathbf{m}^*, \mathbf{D}^*\boldsymbol{\Sigma}\mathbf{D}^{*\mathsf{T}})$.
κ, T_s^κ	–, °C	The assurance a requirement is held to; the largest set point whose joint under the limit is at least κ (Chapter 12).
β_k^{s}	–	The airflow survey's value of the recirculation fraction of rack k, the center of its prior.

Symbol	Unit	Meaning
Q_{loop}	kW	The heat the loop meter reads: both hall loads and both halls' other heat.
$\mathcal{R}$, $\sum G^*$, $\sum C^*$, β_k^*	–, kW/K, kW/K, –	The units of the hall still running with one out; the sums of their conductances and capacity rates; the fraction at the node with the rack air flow held (Chapter 13).
α, γ, $\bar{t}$, d_t	kW, kW per window, –, kW	The line and wander of a load, $Q_{A,t} = \alpha + \gamma(t - \bar{t}) + d_t$: its level, its slope, the center window, and the wander (Chapter 13).
ϕ, s, ρ	–, kW, –	The wander's persistence, its stationary spread, and the part of it left after some windows, $\rho = \phi^{t-73}$ (Chapters 13 and 15).
$Q_f(t)$, $\mathbf{J}_{\text{load}}$, $\mathbf{\Sigma}_{\text{load}}$	kW, varies	The forecast load at window t; the Jacobian and covariance of the second fit, of the load's line and wander (Chapter 13).
Δ, $\Delta_{\max}(\kappa)$	kW	The load removed from a hall (Chapter 13) or moved between halls (Chapter 14), pinned; the largest transfer held at assurance κ.
$\sum G_h$, $\sum C_h$	kW/K	The sums of the conductances and of the capacity rates of the units of hall h.
$T_{k,t}$, $T_{s,t}$, $Q_{h,t}$, $T_{r,h,t}$, $T_{c,h,t}$, $T_{d,u,t}$	°C, kW	The quantities of the plant at window t, when more than one window is in the state.
a_h, b_h, $w_{h,t}$, $\varepsilon_{h,t}$	kW, kW per window, kW, kW	The trend of the load of hall h, $Q_{h,t} = a_h + b_h(t - \bar{t}) + w_{h,t}$: level, slope, the wander $w_{h,t} = \phi w_{h,t-1} + \varepsilon_{h,t}$ and its innovation (Chapter 15).
$g_{\text{B2},t}$, γ, τ, $(x)_+$	–, per window, –, –	The coil's log-derate at window t, $g_{\text{B2}} - \gamma(t - \tau)_+$; the rate of its ramp; the window of its onset; the positive part $\max(x, 0)$ (Chapter 15).
s_j, b_{Q_A}	K, kW	The jitter of each window's set point about one level; the freed bias of the hall-A meter (Chapter 15).
E_s, E, M	–, –, °C	The first-passage event, some inlet at or over the limit at some window through s; the event at the horizon, E_{144}; the largest inlet over the racks and windows (Chapter 15; M_j remains a drawing).
$h(s_1, s_2)$	h^{-1}	The conditional hazard of the first passage between two horizons, given none by the first (Chapter 15).
$\mathbf{G}$, $\widetilde{\mathbf{T}}$	varies	The stacked gradients of the inlets over the horizon; the inlets linearized at $\hat{\mathbf{z}}$, of which the first passage is the complement of an orthant (Chapter 15).
$\mathbf{z}^{(j)}$, w_j	varies, –	A draw of the state and its importance weight (Chapter 15); w_k remains the weight of a row.
m_k, δ_{k,Q_A}, $\mathbf{y}_{1:T}$	varies, –, varies	The prediction of reading k; one on the hall-A meter's row and zero elsewhere; the record through decision window T (Chapter 15).
The model in operation (Part V)		
$\mathbf{H}_t$, h_{kk}	–	The hat matrix $\mathbf{J}_t\mathbf{\Sigma}_t\mathbf{J}_t^{\mathsf{T}}$ of window t; the leverage of row k, its diagonal entry, with $1 - h_{kk}$ the redundancy number.

Symbol	Unit	Meaning
X_t^2, ν_t, r_k^{N}	–	The sum of the squared reading rows at $\hat{\mathbf{z}}_t$; the effective degrees of freedom it is referred to, $\sum_k (1 - h_{kk})$ over the reading rows; the normalized residual $r_k(\hat{\mathbf{z}}_t)/\sqrt{1 - h_{kk}}$.
$\mathbf{d}$, $\boldsymbol{\ell}$, $\boldsymbol{\ell}_{\mathrm{R}}$, λ	–	A stated change on the reading rows; what the fit leaves of it, $(\mathbf{I} - \mathbf{H}_t)$ applied to it; the part on the reading rows; the test's noncentrality, $\|\boldsymbol{\ell}_{\mathrm{R}}\|^2$ (Chapter 16).
$\mathbf{z}_{\mathrm{ref}}$, $\hat{\mathbf{z}}_{\mathrm{s}}$, $\boldsymbol{\Sigma}_{\mathrm{s}}$	varies	The state of the reference fit, shared coils and fractions with each window's supply and loads; the shared entries' MAP values and their covariance (Chapter 16).
$\mathbf{y}_t$, $\mathbf{q}_t$, $\mathbf{q}_0$	varies	The readings of window t; its supply and two loads, free per window against the reference; the center of their prior (Chapter 16).
$\mathbf{A}_{\mathrm{s}}$, $\mathbf{A}_{\mathrm{q}}$, $\mathbf{P}$, $\mathbf{L}$	varies	The gradients of the readings on the shared entries and on $\mathbf{q}_t$; the set point's prior information; $\mathbf{A}_{\mathrm{q}}^{\mathsf{T}}\mathbf{V}_t^{-1}\mathbf{A}_{\mathrm{q}} + \mathbf{P}$ (Chapter 16).
$\mathbf{V}_t$, $\mathbf{e}_t$, $e_{k,t}$	varies	The predictive covariance of the readings, sensors plus the reference's spread; observed minus predicted at the found $\hat{\mathbf{q}}_t$, the sign of the screen; its entry for reading k (Chapter 16).
Q_t, $z_{k,t}$	–	The reference test, $\mathbf{e}_t^{\mathsf{T}}\mathbf{V}_t^{-1}\mathbf{e}_t$ (a statistic, not a load; a load carries a hall, $Q_{h,t}$); the standardized $e_{k,t}$ (Chapter 16).
$\mathbf{a}_k$, $\mathbf{A}_S$	varies	The row of $\mathbf{J}$ for reading k; the rows of a set S of readings (Chapter 17).
$\mathbf{y}_{(k)}$, $\boldsymbol{\Lambda}_{(k)}$, $\boldsymbol{\Sigma}_{(k)}$; $\mathbf{y}_{(S)}$, $\mathbf{J}_{(S)}$, $\boldsymbol{\Sigma}_{(S)}$, $\hat{\mathbf{z}}_{(S)}$	varies	The evidence, information and covariance with reading k removed; the same with the set S removed, the MAP state of its refit (Chapter 17).
c_k, b	varies, –	The coupling $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{a}_k^{\mathsf{T}}$ of row k to a stated quantity, whose loss adds variance $c_k^2/(1 - h_{kk})$; the signed departure of a stuck reading in its own spreads (Chapter 17).
m_a, σ_a, r_a, $\mathbf{a}$	varies	What candidate reading a reads, as a function of the state; its spread; its scenario row $(m_a(\mathbf{z}) - \tilde{y}_a)/\sigma_a$; its gradient over its spread at $\hat{\mathbf{z}}$ (Chapter 18).
$\hat{\mathbf{z}}_a$, ρ_a, P_a	varies, –, –	The MAP state after outcome $\tilde{y}_a$; the correlation of the reading, error included, with a worry; the joint at the node after that outcome (Chapter 18).
$T_{\mathrm{in}k}$, T_{rB}, Q_{B2}	°C, °C, kW	The worries by name: the inlet of rack k, written T_k elsewhere; hall B's return, written $T_{r,\mathrm{B}}$ elsewhere; the heat of unit B2 (Chapter 18).
Alarm levels and credibility (Part VI)		
$\mathbf{K}$, $\mathbf{K}_t$	–	The residual map, $\mathbf{I} - \mathbf{A}_{\mathrm{q}}\mathbf{L}^{-1}\mathbf{A}_{\mathrm{q}}^{\mathsf{T}}\mathbf{V}_t^{-1}$, which takes the readings' departures to the residuals with this window's supply and loads refound (Chapter 19; Chapter 20 writes $\mathbf{K}_t$).
e	K	The residual of a monitored reading, its entry of $\mathbf{e}_t$, in kelvin and not in its spreads; in Chapter 20 also the value at which the alarm fired. It is not the normalized e_k of Chapter 17.

Symbol	Unit	Meaning
sd, $\mathrm{sd_n}$, $\mathrm{sd_s}$, x	K	The predictive spread of a monitored quantity; its noise part, new in each window; its shared part, the same in every window; the reference's shared error at that quantity, normal on $\mathrm{sd_s}$ (Chapter 19).
$\mathbf{z}^*$, g_u^*, β_k^*	varies, –, –	The state at the scenario node of an excursion route; the coil's log-derate and the rack's fraction pinned there (Chapters 19 and 20).
s_r, s	K	The shift route r leaves on the residual (Chapter 19); the shift a stated excursion leaves on it (Chapter 20). In Part IV and in the plain reading's spread of Chapter 19, s is the wander's spread.
M	windows	The windows of a month, 8640 (Chapter 19); M_j remains a drawing.
α, π, z_α, z_π	–	In Chapter 19: the requested false-alarm tail per window; the requested catch; their standard normal points.
L_α, L_π	K	The level on the monitored quantity that gives the rate, z_α sd, and the level that gives the catch, $s_r - z_\pi$ sd; both requests hold on the interval $[L_\alpha, L_\pi]$ when it is not empty (Chapter 19).
$\alpha(L)$, $\pi_r(L)$, $\pi_t(L)$	–	The false-alarm tail at level L; the catch of route r there; the catch of the plain reading at window t (Chapter 19).
$\bar{y}_t$, sd_t, $\mathcal{W}$	°C, K, –	The forecast of the plain reading at window t, a superscript ref when it stands on the reference; its spread; the span of windows a fixed level covers (Chapter 19).
a, a_r, σ_T, σ_j	K/kW, K/kW, K, K	The inlet's slope on hall B's load at the one fit and at the reference; the sensor's spread; the set point's jitter (Chapter 19).
π, c, α, L	–, –, –, K	In Chapter 20: the base rate of excursions per window; the catch, the tail past the level under an excursion (elsewhere c is a decision level, and in Chapter 19 π is the catch); the false-alarm tail per window; the level.
$\boldsymbol{\delta}$, $\mathbf{d}$	varies	What a hypothesis moves in the prediction; the same through the map, $\mathbf{K}_t\boldsymbol{\delta}$ (Chapter 20; in Chapter 19 $\mathbf{d}$ is the readings' change before the map).
$B(e)$, S, $w(S)$	–, K, 1/K	The density ratio at the value e; an excursion's size; its stated density (Chapter 20).
d, π_d	windows, 1/window	The windows an excursion lasts; the prior of a window drawn from all the windows that fire (Chapter 20).
γ, σ_γ, $\mathbf{h}_t$, $\mathbf{C}$	1/window, 1/window, varies, varies	The ramp's slope, signed, g_{B2} moved by $\gamma(t-\tau)_+$, so a derate has $\gamma < 0$ (Chapter 15 writes the ramp with a minus sign); its prior spread; its column at window t; the stacked covariance (Chapter 20).
S_{eh}, S_{hh}, Λ	varies	The ramp's two weighted products, $\mathbf{h}^\mathsf{T}\mathbf{C}^{-1}\mathbf{e}$ and $\mathbf{h}^\mathsf{T}\mathbf{C}^{-1}\mathbf{h}$; its information (Chapter 20).
r, α_r, c_r	K, –, –	The ramp's part in the window's residual; the two tails with it in both means (Chapter 20; r_k remains a row).

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. Built before any question is asked.

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

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

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

Telemetry window The span of readings fitted at once, five minutes here, within which the plant is treated as steady.

State estimation The fit of the state of one window to its readings and priors: the MAP state and its covariance.

Virtual (soft) sensor An unmeasured quantity read from the fitted state with its band, in place of a missing instrument.

Observability Whether the readings determine an entry of the state, or only a combination of entries; read from the Jacobian and from the posterior spread against the prior spread.

Bad-data detection Testing whether a reading is consistent with the others: the chi-square test of the whole residual vector and the largest normalized residual.

Fault detection and isolation Deciding that something changed, and where: a family of candidate faults generated from the graph, ranked by evidence.

What-if (contingency) analysis The plant re-solved at a scenario node, with an actuator or a unit state pinned to a value not yet applied, and the covariance carried through.

Look-ahead A forecast of the plant over the coming windows, with its band.

Operating envelope The region of the state inside which every limit holds.

Alarm rationalization Setting each alarm's decision level from its purpose, its false-alarm rate and its probability of detection.

False-alarm rate, probability of detection The rate at which an alarm fires without a real excursion; the probability that it fires when one of stated size occurs.

Standing alarm, alarm flood An alarm that stays active for long periods; more alarms than an operator can handle in a short time.

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 $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$.

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 reading not yet taken; the object that values a reading before it is requested.

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

Scenario node A copy of the plant's temperature nodes at values of the pinned or actuated quantities not yet applied. It adds no reading row and no entry of the state; its temperatures are derived quantities with their bands.

Pin, shift A scenario's set point fixed at a value; the same set point moved from its fitted entry by an amount. Under a pin the band at the node is that of the inlet less the supply; under a shift it is the present inlet's.

Assurance The probability a requirement is held to; a largest set point, unit load or transfer is the threshold question inverted at it.

N−1 contingency The plant with one unit out, asked of each unit in turn.

Stop rule, rack air flow held What a stopped unit does to the air: the plant's rule drops its flow; the second form holds the rack air flow.

Forecast peak The load the trend line and the wander give at the horizon's last window.

Load transfer (load migration) Load moved between halls at a scenario node: a donor hall and a receiving hall.

Margin The limit less an inlet's temperature: a capacity margin, not an availability.

First passage Some inlet at or over the limit at some window of the horizon, a union over racks and windows; its conditional hazard is the probability of a first passage in an hour, given none earlier.

Linearized joint (orthant) The state route for a first passage: an orthant of the normal joint of the forecast inlets, integrated by draws.

Load wander A persistent first-order autoregressive departure of a load about its trend line, with its persistence and stationary spread.

Epistemic and aleatory spread The spread from entries the record could fix; the spread of the future's own wander.

Decision window A window at which the answer is read again, using nothing later.

Goodness of fit, effective degrees of freedom The agreement between the readings and the fitted model, tested by the sum of the squared reading rows against the effective degrees of freedom $\sum(1 - h_{kk})$, not rows minus entries.

Fit-quality light A screen's summary of the window's own test of fit.

Noncentrality, power The squared size of a stated change left on the rows a statistic adds up; the noncentral chi-square tail past the level, which is the test's probability of detection.

Reference fit, reference test One fit of a stretch of windows with shared coils and fractions; the test of a later window's readings against its predictive.

First verdict The first window past a level fixed in advance, reported with its spread over fresh sensor errors.

Repaired model The model with an added entry, fitted and tested again.

Sensor loss (removed row) A reading's row taken out of the fit. A sensor that fails and stays connected is a different row, a stuck-at failure, whose band does not widen.

Down-date (deletion formula) The covariance with a row removed, read from $\boldsymbol{\Sigma}$ at $\hat{\mathbf{z}}$; it needs no outcome. Block deletion removes a named group of readings together.

Coupling A row's link to a stated quantity; its square over the redundancy number is the variance its loss adds.

Critical measurement A reading whose loss lowers the rank of the reading rows.

Candidate reading, scenario row An instrument, a place and a spread, not yet used; the reading row it would add, written before its value exists.

Worry A named quantity to narrow, a function of the state; a ranking of candidates is against a worry.

Width taken off The share by which a band falls after a row.

Settled share The probability over a reading's outcomes that a joint ends at or above an upper reference level or below a lower one.

Greedy order The best reading, then the best given it, tested against the best pair.

Decision level The value of a monitored quantity past which an alarm fires.

Monitored quantity, residual alarm What an alarm compares with its level; an alarm on a reading's residual against a reference, with the window's supply and loads refound, as against a fixed level on the plain reading.

Excursion route A physical way a stated rise comes about, injected as a scenario node. A level's catch is stated for a route.

Catch The probability of detection of a stated excursion at a level.

Shared part of a residual's spread The reference's own error, the same in every window, as against the per-window reading noise; it does not average over windows and makes a false-alarm rate itself uncertain.

Base rate How often the thing an alarm is meant for occurs, regardless of any alarm.

Positive predictive value (alarm credibility) The posterior probability of an excursion given the alarm: the two tails are the evidences of excursion and of no excursion, and the base rate gives the prior odds.

Tail ratio, density ratio The likelihood ratio of the observation that the residual passed the level; that of the observation that it took this value.

First alarm The first fired window of a run, as against any window that fires.

Part IV

Predicting the Effect of Operator Actions

Chapter 12

Set-Point Changes as Scenarios

12.1 A raise of the chilled water in the night

It is 01:00, window 73 of the night, the loop is steady, and the operator is considering a raise of the chilled-water set point from 14 to 16°C. Warmer chilled water is commonly asked for because the chiller then does less work, so a raise could be a saving. The screen shows the supply at 13.88°C, the returns at 29.97 and 29.15°C, and the four inlet sensors at 19.51 (rack A1), 22.03 (A3), 22.13 (B2) and 22.67°C (B4). The hottest read inlet, rack B4, is 4.33 K inside the 27°C limit. Adding 2 K to the screen gives 24.67°C, 2.33 K inside, so the dashboard's reading is that the raise is safe.

Chapter 11 closed on a forecast that assumed nobody acts, and here somebody may. That reading adds a number to four sensors, and the plant has ten racks. The plant also carries no chiller-power row, so this chapter prices the risk side of the raise only. The question it answers is

what the probability is that every inlet stays under the limit at the raised set point, including the six racks no sensor reads, and what stands behind that probability.

It is answered on a bigger plant than the two-unit hall, with the method of Chapter 4 and one new object. The quantities asked for lie at a value of the actuator that has not been applied, and no reading can touch them.

Jargon: set point, actuator, what-if (contingency) analysis, scenario node, pinned, carried, assurance, joint tail, pin and shift

The *set point* is the value a controller is told to hold; the *actuator* is what the operator moves to change the plant, here the chilled-water set point. *What-if (contingency) analysis* computes the plant at a value of an actuator or a unit's run state not yet applied (in the framework, a *scenario node*). A quantity at the node is *pinned* if the scenario fixes it, and *carried* if it keeps the posterior of the present state. The *assurance* κ is the probability the operator requires of a requirement, as in Chapter 2. A *joint tail* is the probability of an event on several quantities at once, as in Chapter 8. The *pin* sets T_s^* to a value; the *shift* adds an amount to the fit's entry for T_s.

Key idea

A scenario node adds no reading and no entry to the reduced state: its temperatures are derived quantities of the state, the posterior is unchanged, and $\boldsymbol{\Sigma}$ is carried by $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$. A raise of the supply moves every temperature by the same amount, so the physics of the raise is one line. The answer is not, because that amount lands on inlets whose fractions no reading has seen.

12.1.1 A scenario node and what it carries

Chapter 5 read a temperature of the present plant at a place with no sensor. A what-if reads a temperature of a plant that does not yet exist: the same racks and coils, at a value of the actuator the operator has not entered. A *scenario node* is a copy of the plant's temperature nodes at values of the pinned or actuated quantities not yet applied. Write $\mathbf{u}^*$ for those values. Each scenario temperature is the forward chain of the physical rows solved at $\mathbf{u}^*$ and at the free entries,

$$T_k^*(\mathbf{z}) = T_k(\mathbf{z};\ \mathbf{u}^*), \qquad T_k^* \mid \mathbf{y} \approx \mathcal{N}(T_k^*(\hat{\mathbf{z}}),\ \mathbf{g}_{T_k^*}^{\mathsf{T}} \boldsymbol{\Sigma}\, \mathbf{g}_{T_k^*}), \tag{12.1}$$

a function of the state like the virtual sensor of Chapter 5. The node adds no reading row, since nobody can read a plant that has not happened, and no entry to the reduced state, so the posterior is the one the readings already gave. On the full state it adds one copy of the temperatures and one copy of the physical rows, which the limit of exact rows removes again.

Principle 12.1. *What-if analysis needs no new reading and no new fit: pin what the operator would change, re-solve the temperatures, and carry $\boldsymbol{\Sigma}$ through their gradients. What a scenario can know is what the present posterior knows.*

12.1.2 The two-hall shift plant (step 0)

The two-unit hall of Chapter 5 had two units, three racks and a state of six entries. The shift plant has two halls A and B on one chilled-water loop at supply temperature T_s, three cooling units per hall (A1 to A3, B1 to B3) and five racks per hall (A1 to A5, B1 to B5); a unit and a rack may share a name, and the subscript of $T_{d,\mathrm{A1}}$ is a unit's while that of T_{A1} is a rack's.

The physical relations. They are the rows of Chapter 5, repeated per hall h with $G_u = \bar{G}_u \mathrm{e}^{g_u}$:

$$\begin{aligned}
r_{\mathrm{h},h} &= \Big[\sum_{u\in h} G_u\,(T_{r,h} - T_s) - Q_h - Q_{o,h}\Big]/\sigma_p && \text{hall balance,}\\
r_{\mathrm{a},u} &= \big[C_u\,(T_{r,h} - T_{d,u}) - G_u\,(T_{r,h} - T_s)\big]/\sigma_p && \text{air side of unit } u,\\
r_{\mathrm{m},h} &= \Big[\sum_{u\in h} C_u\,T_{c,h} - \sum_{u\in h} C_u T_{d,u}\Big]/\sigma_p && \text{cold-aisle mixing,}\\
r_{\mathrm{i},k} &= \big[T_k - T_{c,h} - \beta_k\,(T_{r,h} - T_{c,h})\big]/\sigma_p && \text{rack inlet } k,
\end{aligned} \tag{12.2}$$

ten rows per hall, twenty in all, with σ_p the tolerances of those rows taken to zero together. The pinned quantities are the ratings $\bar{G}_u = 8\,\mathrm{kW/K}$ in hall A and $7\,\mathrm{kW/K}$ in hall B, the air capacity rates $C_u = 10\,\mathrm{kW/K}$ and the other heat $Q_{o,\mathrm{A}} = 12\,\mathrm{kW}$, $Q_{o,\mathrm{B}} = 10\,\mathrm{kW}$. The free quantities and the twenty temperatures form a full state of 39 entries, and with one scenario copy of the temperatures, 59.

The reduced state. On the reduced state of Chapter 4 the fit runs on the 19 free entries,

$$\mathbf{z}_t = \big(T_s,\ Q_{\mathrm{A}},\ Q_{\mathrm{B}},\ g_{\mathrm{A1}},\dots,g_{\mathrm{B3}},\ \beta_{\mathrm{A1}},\dots,\beta_{\mathrm{B5}}\big)^{\mathsf{T}}, \qquad t = 73, \tag{12.3}$$

in °C, kW and sixteen dimensionless entries. The dependent block $\mathbf{A}_{\mathrm{dep}}$ of the rows (12.2) is triangular with a nonzero diagonal, so it is invertible, and the temperatures of each hall follow in order,

$$\begin{aligned}
T_{r,h} &= T_s + \frac{Q_h + Q_{o,h}}{\sum_{u\in h} G_u}, \qquad T_{d,u} = T_{r,h} - \frac{G_u\,(T_{r,h} - T_s)}{C_u},\\
T_{c,h} &= \frac{\sum_{u\in h} C_u T_{d,u}}{\sum_{u\in h} C_u}, \qquad T_k = T_{c,h} + \beta_k\,(T_{r,h} - T_{c,h}).
\end{aligned} \tag{12.4}$$

As in equation (5.5), the air balance of each hall,

$$\sum_{u \in h} C_u \left(T_{r,h} - T_{c,h}\right) = Q_h + Q_{o,h}, \tag{12.5}$$

holds at every point of the set: the rise from the cold aisle to the return is set by the heat and the airflow alone.

Modeling choices at the graph. The plant is taken in steady state within a window, which makes the rows (12.2) algebraic and lets the node say where the plant settles at T_s^*; it says nothing of how long the loop and the air take to get there or of any controller's overshoot on the way, and the record holds no set-point step from which either could be read, so a statement of the form "safe for the next hour" is not available from this node. The units' airflow is pinned and no controller acts on the coils within a window; this lets a discharge reading inform its coil, and a valve that held a discharge temperature would break it. Each coil carries one effective conductance from the return air to the chilled water; this keeps one entry per coil, and it fails if a coil's heat transfer depends on its discharge temperature in a way one conductance cannot carry.

The graph. Figure 12.1 draws the rows. The two halls stand left and right of the chilled water in the middle (green). In each hall the load (green) enters the return air (blue), the return feeds the three coils (green, each with its log-derate), the coils' discharge mixes in the cold aisle (blue), and the cold aisle feeds five racks, each carrying its own fraction. Dashed arrows leave for the sixteen readings. Rack B3 is drawn in red. The red dashed box at the bottom is the scenario node: a copy of every temperature, with T_s^* pinned, and no dashed arrow to a reading.

12.1.3 The operator's questions (step 1)

The operator asks five questions of window 73, and step 1 writes each over the reduced state (12.3). Table 12.1 lists them. The first column is the question in the operator'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 plant; 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.

12.1.4 The evidence and the node as rows (step 2)

Step 2 attaches window 73 to the graph as rows on the reduced state. The sixteen readings are those of Figure 12.1: the supply, of standard deviation 0.1 K; the two returns, six discharges and four inlets, each 0.3 K, the discharges of units A1 to B3 reading 17.87, 16.48, 17.60, 18.07, 19.99 and 18.97°C; the two hall meters, each 4 kW; and the loop meter on $Q_{\mathrm{A}} + Q_{\mathrm{B}} + Q_{o,\mathrm{A}} + Q_{o,\mathrm{B}}$, 8 kW. The priors are the plant's: $g_u \sim \mathcal{N}(0,\ 0.15^2)$ from the commissioning record, $\beta_k \sim \mathcal{N}(\beta_k^{\mathrm{s}},\ 0.08^2)$ from the airflow survey, $T_s \sim \mathcal{N}(14,\ 0.5^2)$. The loads have no prior row. That makes $16+17 = 33$ rows on 19 entries. Each reading is taken as Gaussian, unbiased and independent of the others given the state; this makes the objective a sum of squares, and it fails if a meter carries a bias the rows do not model, which Section 12.2.3 examines. The airflow survey is taken as the only knowledge of an unread rack's fraction. This gives the six empty columns of the Jacobian a proper posterior, and it fails if the survey's center or width is wrong or has gone stale, since no reading can correct it; the answer inherits whatever the survey got wrong, and Section 12.4.3 varies both.

What the node pins, leaves free and carries. Table 12.2 lists each quantity at the node of this chapter, with what it does there and why.

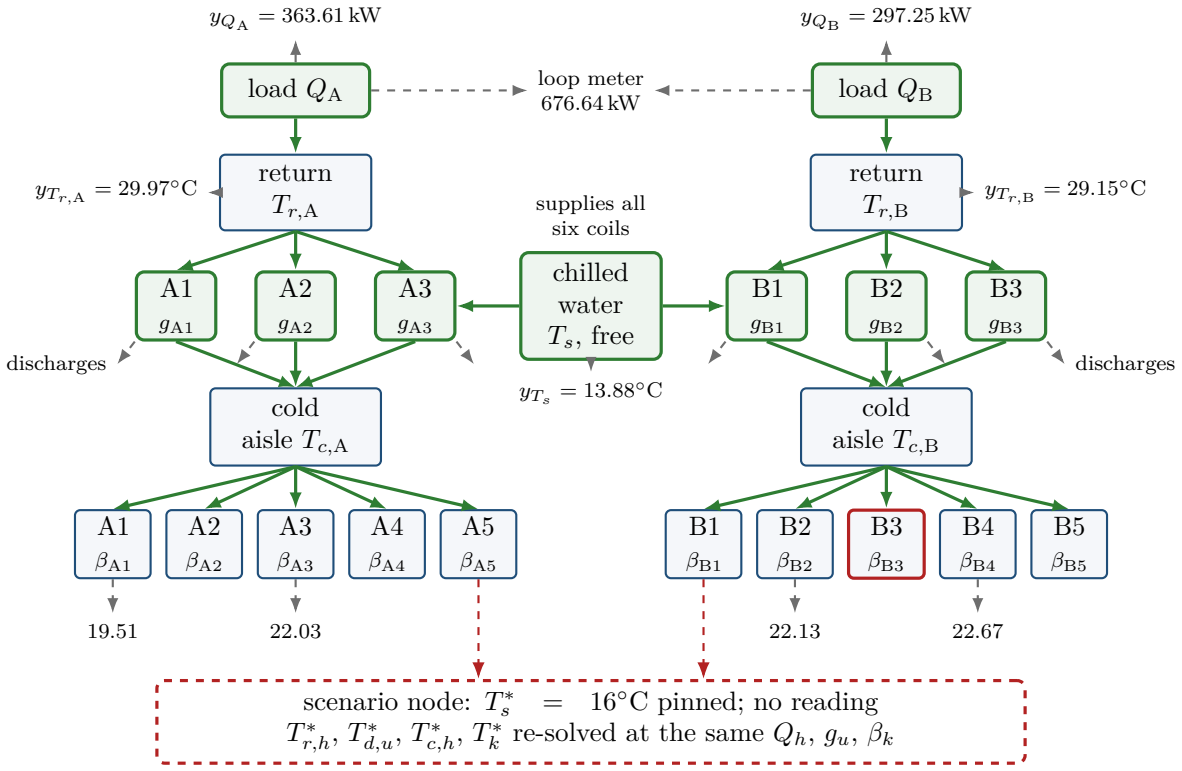


Figure 12.1: The rows of the shift plant at window 73, 01:00. Free (green): the supply T_s, the loads Q_h and the log-derates g_u; the recirculation fractions β_k are free entries carried by the rack nodes, and a rack's mix of cold aisle and return is its inlet row, not drawn from the return. Blue: temperatures fixed by the rows (12.2). Observed (dashed gray): the supply, two returns, six discharges, two hall meters, the loop meter and the inlets of racks A1, A3, B2 and B4, in °C and kW. Not drawn: the pinned ratings, capacity rates and other heat, and the priors. Red: rack B3, and the scenario node with the arrows that give it the state's entries.

Table 12.1: The questions asked of window 73. $\mathbf{y}$ is the sixteen readings and the seventeen priors; T_k^* the inlet of rack k at the node (12.1); the limit is 27°C.

Operator question	General form	Mathematical form	Category	Computation	Answered in
If I raise the supply temperature 2 K, what is the probability the inlets stay under the limit?	(2.12), a joint tail at a scenario node	$\Pr(\max_k T_k^* < 27^\circ\mathrm{C} \mid \mathbf{y},\ T_s^* = 16^\circ\mathrm{C})$	prediction under a threshold	posterior	12.3.1
Which rack carries it, and what stands behind its band?	(2.9), (2.11)	T_k^* and its band; the budget of $T_{\mathrm{B}3}^*$	estimation	posterior	12.3.2
Is the raise seen by the sensors on the screen?	(2.12) on two subsets	the measured four against the unread six; the sensors see the raise and not the tail	prediction under a threshold	posterior	12.3.3
Does the probability fall with the raise, and how large a raise holds at a stated assurance?	(2.12) read along T_s^*	$\Pr(\cdot \mid T_s^*)$; the largest T_s^* with $\Pr \geq \kappa$	prediction under a threshold	posterior	12.3.4
Is the raise a pinned value or a shift?	(2.12), two nodes	$T_s^* = 16^\circ\mathrm{C}$ against $T_s + 2\,\mathrm{K}$	prediction under a threshold	posterior	12.3.5

Table 12.2: The scenario node of the raise.

Quantity	At this node	Why
T_s^*	pinned, 16.0°C	the actuator; the set point the controller is told
Q_{A}, Q_{B}	free, window 73's entries	loads do not respond to the water
g_u	free, window 73's entries, assumed independent of T_s	a choice, below
β_k	free, window 73's entries	the air pattern does not respond to the water
C_u, $Q_{o,h}$	pinned as in the plant	fixed fans, other heat
every temperature	re-solved	$T_{r,h}^*$, $T_{d,u}^*$, $T_{c,h}^*$, T_k^*
loop and unit heats	unchanged at the node	Q_u depends on $T_{r,h} - T_s$, which does not move

The physics before any number. By the chain (12.4), every temperature is T_s plus a quantity that does not depend on T_s, so at the node

$$T_k^* = T_s^* + \frac{Q_h + Q_{o,h}}{\sum_{u \in h} G_u} - (1 - \beta_k)\, \frac{Q_h + Q_{o,h}}{\sum_{u \in h} C_u}, \tag{12.6}$$

the second term the rise from the supply to the return and the third the rise (12.5) scaled. Every row is linear in T_s with coefficient one: a raise moves every temperature of the chain by the same amount and no unit heat at all. A raised supply must therefore make the probability fall. The chapter's physics is that simple, and its content is that the amount is added to a rack that no sensor reads.

Modeling choices at the node. Four choices make this node: the coil and the actuator are varied in Section 12.4.3, the pin in Section 12.3.5; the loads are not varied (Section 12.6.4).

The conductance independent of the supply temperature. This buys the coefficient one above. The record cannot test it: its supply readings span 13.75 to 14.23°C with a standard deviation of 0.102 K against the reading's 0.10 K, so the spread cannot be told from reading noise, and the node lies 1.8 K above the largest of them. It fails if warmer water lowers a coil's effective conductance, through less latent duty or less temperature difference at the valve.

The set point pinned at 16.0°C*, not the entry plus* 2 K. A controller holds a set point, and the operator enters that number. This takes the entry T_s out of the gradient of every scenario temperature; what that does to the band is computed in Section 12.3.5, where the shift stands beside it.

The loads held at window 73's values. The node is the plant at 01:00 with a different set point. This keeps the question to the actuator, and it fails as a statement about any later window, since it holds no load change.

The actuator exact. The node's T_s^* has no spread. This keeps the answer to what the fit knows; it fails to the extent the loop misses the value it is told.

The limit and the assurance. The limit, 27°C, is the plant's. The assurance κ in equation (2.12) belongs to the operator. The chapter computes at the stated values 0.90, 0.95 and 0.99 and names none of them as the rule.

Which row acts on which entry. Figure 12.2 draws the sixteen reading rows of the Jacobian on the 19 columns: each row of the figure is a reading, each column an entry, and a filled square a nonzero; the seventeen prior rows, each a single entry on its own column, are not drawn. Each hall's temperature rows see T_s, that hall's load and its three coils; a measured inlet adds its own fraction; the meters see the loads only. The six shaded columns are empty.

The block structure reads as observability. The six fractions of the unread racks are not observable from the readings, and only their prior rows hold them; the two halls share the column of T_s and, through the loop meter, those of the loads. The question's gradient is not a row. It has the pattern

$$\mathbf{g}_{T^*_{\mathrm{B3}}}^{\mathsf{T}} = (0,\ 0,\ *,\ 0,\ 0,\ 0,\ *,\ *,\ *,\ \underbrace{0,\ \ldots,\ 0}_{\beta_{\mathrm{A1}},\ldots,\beta_{\mathrm{B2}}},\ *,\ 0,\ 0), \tag{12.7}$$

with entries on Q_{B}, the three coils of hall B and β_{B3}, one of the empty columns, and none on T_s: the pin removes it.

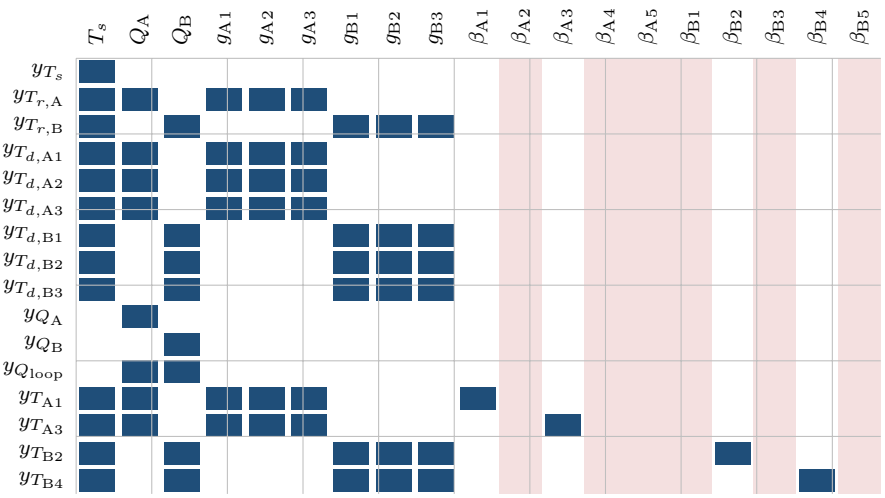


Figure 12.2: The sparsity of the reading rows of $\mathbf{J}$. Shaded: the columns of β_{A2}, β_{A4}, β_{A5}, β_{B1}, β_{B3} and β_{B5}, in which no reading has a nonzero.

12.2 The posterior of the state (step 3)

This is the fit of window 73, 01:00, computed by the plant's own rows, and it is the one fit Part IV reads. No concrete reason against this window was found: its test passes. Everything an answer below uses from the record's events is computed here too.

12.2.1 The rows and the normal equations

Stacked in the order of Figure 12.2, each divided by its standard deviation, the rows on the reduced state are

$$\mathbf{r}(\mathbf{z}) = \begin{bmatrix} (T_s - y_{T_s})/0.1 \\ (T_{r,h}(\mathbf{z}) - y_{T_{r,h}})/0.3 \quad (h = \mathrm{A}, \mathrm{B}) \\ (T_{d,u}(\mathbf{z}) - y_{T_{d,u}})/0.3 \quad (\text{six units}) \\ (Q_h - y_{Q_h})/4 \quad (h = \mathrm{A}, \mathrm{B}) \\ (Q_{\mathrm{A}} + Q_{\mathrm{B}} + Q_{o,\mathrm{A}} + Q_{o,\mathrm{B}} - y_{Q_{\mathrm{loop}}})/8 \\ (T_k(\mathbf{z}) - y_{T_k})/0.3 \quad (k = \mathrm{A1}, \mathrm{A3}, \mathrm{B2}, \mathrm{B4}) \\ (T_s - 14)/0.5 \\ g_u/0.15 \quad (\text{six units}) \\ (\beta_k - \beta_k^{\mathrm{s}})/0.08 \quad (\text{ten racks}) \end{bmatrix}, \qquad \mathcal{L}(\mathbf{z}) = \tfrac{1}{2}\,\mathbf{r}(\mathbf{z})^{\mathsf{T}}\mathbf{r}(\mathbf{z}), \tag{12.8}$$

with the temperatures from the chain (12.4) and the reduced Jacobian $\mathbf{J}$ of equation (4.47). The MAP state minimizes $\mathcal{L}$ by the normal equations of equation (4.29),

$$(\mathbf{J}^{\mathsf{T}}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^{\mathsf{T}}\mathbf{r}(\mathbf{z}), \qquad \boldsymbol{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1}\big|_{\hat{\mathbf{z}}}, \tag{12.9}$$

from the set point, the hall meters, the ratings and the survey, the Gauss–Newton form of equation (4.38); the posterior is approximately normal, $\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$. The empty columns of Figure 12.2 separate the objective: the terms in the six unread fractions are their prior terms alone, so their posterior is the survey's exactly, not only under the linearization, and they are uncorrelated with every other entry, as in equation (5.10).

The exact posterior. The density $p(\mathbf{z} \mid \mathbf{y}) \propto \exp(-\mathcal{L}(\mathbf{z}))$ has no closed form on 19 entries. It is set beside the Laplace normal by importance sampling: 400 000 draws from that normal, each weighted by the ratio of the two densities, with an effective size of 396 168 and a largest weight of 7.58×10^{-6}. Each draw is pushed through the chain with T_s replaced by the pin.

12.2.2 The fit of window 73

The iterations. Table 12.3 runs equation (12.9). Each row is an iteration; the columns give the objective and the norm of its gradient. The objective falls from 31.8994 to 1.3134 in the first step and settles at 1.2748 by the second, and the norm falls to 8.46×10^{-9} by the fifth.

Table 12.3: Gauss–Newton iterations of window 73, equation (12.9).

Iteration k [–]	Objective $\mathcal{L}(\mathbf{z}^{(k)})$ [–]	$\lVert\mathbf{J}^{\mathsf{T}}\mathbf{r}\rVert$ [–]
0	31.8994	4.37×10^{2}
1	1.3134	1.98×10^{1}
2	1.2748	1.43×10^{-2}
3	1.2748	6.52×10^{-5}
4	1.2748	3.22×10^{-7}
5	1.2748	8.46×10^{-9}

The Jacobian at the MAP state. Table 12.4 gives the reading rows of $\mathbf{J}(\hat{\mathbf{z}})$ on the thirteen observed columns: each row a reading over its standard deviation, each cell its derivative in that column. The six unread fractions' columns hold zeros in these rows. The prior rows are single entries on their own columns: 2 for T_s, 1/0.15 for each coil and 1/0.08 for each fraction. An inlet's row repeats its hall's return row in the coil columns, because by the air balance (12.5) a derate leaves the rise unchanged.

Table 12.4: The reading rows of $\mathbf{J}(\hat{\mathbf{z}})$ on the observed columns, window 73; the Q columns in kW^{-1}, the others per unit of the entry, T_s per K.

[–]	T_s	Q_{A}	Q_{B}	g_{A1}	g_{A2}	g_{A3}	g_{B1}	g_{B2}	g_{B3}	β_{A1}	β_{A3}	β_{B2}	β_{B4}
y_{T_s}	10.000	0	0	0	0	0	0	0	0	0	0	0	0
$y_{T_{r,\mathrm{A}}}$	3.333	0.142	0	−16.977	−18.907	−17.358	0	0	0	0	0	0	0
$y_{T_{r,\mathrm{B}}}$	3.333	0	0.166	0	0	0	−18.487	−15.409	−17.038	0	0	0	0
$y_{T_{d,\mathrm{A1}}}$	3.333	0.036	0	−44.140	−4.745	−4.357	0	0	0	0	0	0	0
$y_{T_{d,\mathrm{A2}}}$	3.333	0.024	0	−2.815	−47.549	−2.878	0	0	0	0	0	0	0
$y_{T_{d,\mathrm{A3}}}$	3.333	0.033	0	−3.976	−4.428	−44.839	0	0	0	0	0	0	0
$y_{T_{d,\mathrm{B1}}}$	3.333	0	0.045	0	0	0	−42.072	−4.208	−4.653	0	0	0	0
$y_{T_{d,\mathrm{B2}}}$	3.333	0	0.066	0	0	0	−7.287	−36.932	−6.716	0	0	0	0
$y_{T_{d,\mathrm{B3}}}$	3.333	0	0.055	0	0	0	−6.102	−5.086	−39.746	0	0	0	0
$y_{Q_{\mathrm{A}}}$	0	0.250	0	0	0	0	0	0	0	0	0	0	0
$y_{Q_{\mathrm{B}}}$	0	0	0.250	0	0	0	0	0	0	0	0	0	0
$y_{Q_{\mathrm{loop}}}$	0	0.125	0.125	0	0	0	0	0	0	0	0	0	0
$y_{T_{\mathrm{A1}}}$	3.333	0.050	0	−16.977	−18.907	−17.358	0	0	0	41.689	0	0	0
$y_{T_{\mathrm{A3}}}$	3.333	0.072	0	−16.977	−18.907	−17.358	0	0	0	0	41.689	0	0
$y_{T_{\mathrm{B2}}}$	3.333	0	0.090	0	0	0	−18.487	−15.409	−17.038	0	0	34.001	0
$y_{T_{\mathrm{B4}}}$	3.333	0	0.096	0	0	0	−18.487	−15.409	−17.038	0	0	0	34.001

The MAP state. Table 12.5 gives the fit: each row an entry, the columns its value, its posterior spread, its prior spread and the rows that inform it. Four fractions move and narrow under their inlet sensors; the other six stay at the survey's values with the survey's spread. At the MAP state $T_{r,\mathrm{A}} = 29.8674$, $T_{r,\mathrm{B}} = 29.1749$, $T_{c,\mathrm{A}} = 17.3607$ and $T_{c,\mathrm{B}} = 18.9745$°C; the conductances sum to 23.4901 and 20.0268 kW/K; and the rises from the supply to the return are 15.9727 and 15.2802 K.

The covariance. On the observed entries the variances, the diagonal of $\boldsymbol{\Sigma}$, are 9.5167×10^{-3} for T_s; 11.794 and 11.766 for Q_{A} and Q_{B}; 5.4524, 4.6074, 5.2627, 6.2183, 8.4368 and 7.1135×10^{-4} for the six coils in order; and 6.4333, 6.4406, 9.1888 and 9.1992×10^{-4} for the four read fractions. Table 12.6 gives their correlations, each cell the correlation of its row's entry and its column's.

Table 12.5: The MAP state of window 73, 01:00, with the posterior and prior spreads.

Entry [unit]	$\hat{z}_i$ [entry's]	$\sqrt{\Sigma_{ii}}$ [entry's]	Prior sd [entry's]	Informed by
T_s [°C]	13.8947	0.0976	0.5	supply reading
Q_{A} [kW]	363.2006	3.4343	none	hall and loop meters, hall A's temperatures
Q_{B} [kW]	296.0123	3.4302	none	hall and loop meters, hall B's temperatures
g_{A1} [–]	−0.0659	0.0234	0.15	hall A's temperatures
g_{A2} [–]	0.0418	0.0215	0.15	hall A's temperatures
g_{A3} [–]	−0.0437	0.0229	0.15	hall A's temperatures
g_{B1} [–]	0.0377	0.0249	0.15	hall B's temperatures
g_{B2} [–]	−0.1444	0.0290	0.15	hall B's temperatures
g_{B3} [–]	−0.0439	0.0267	0.15	hall B's temperatures
β_{A1} [–]	0.1698	0.0254	0.08	inlet sensor and survey
β_{A3} [–]	0.3715	0.0254	0.08	inlet sensor and survey
β_{B2} [–]	0.3082	0.0303	0.08	inlet sensor and survey
β_{B4} [–]	0.3665	0.0303	0.08	inlet sensor and survey
β_{A2}, β_{A4}, β_{A5} [–]	0.25, 0.30, 0.45	0.0800	0.08	survey only
β_{B1}, β_{B3}, β_{B5} [–]	0.20, 0.50, 0.35	0.0800	0.08	survey only

Table 12.6: Correlations of $\boldsymbol{\Sigma}$ on the observed entries, window 73. The six unread fractions have a variance of 0.08^2 and no correlation with any entry.

[–]	T_s	Q_{A}	Q_{B}	g_{A1}	g_{A2}	g_{A3}	g_{B1}	g_{B2}	g_{B3}	β_{A1}	β_{A3}	β_{B2}	β_{B4}
T_s	1.00	−0.00	−0.00	0.26	0.28	0.26	0.25	0.22	0.24	−0.00	−0.00	−0.00	−0.00
Q_{A}	−0.00	1.00	−0.18	0.18	0.15	0.17	−0.03	−0.04	−0.04	0.02	−0.04	0.00	0.01
Q_{B}	−0.00	−0.18	1.00	−0.03	−0.03	−0.03	0.19	0.23	0.21	−0.00	0.01	−0.02	−0.04
g_{A1}	0.26	0.18	−0.03	1.00	−0.10	−0.12	0.06	0.05	0.05	0.17	0.16	−0.00	0.00
g_{A2}	0.28	0.15	−0.03	−0.10	1.00	−0.10	0.07	0.05	0.06	0.18	0.17	−0.00	0.00
g_{A3}	0.26	0.17	−0.03	−0.12	−0.10	1.00	0.06	0.05	0.06	0.17	0.16	−0.00	0.00
g_{B1}	0.25	−0.03	0.19	0.06	0.07	0.06	1.00	−0.17	−0.15	−0.00	0.00	0.15	0.15
g_{B2}	0.22	−0.04	0.23	0.05	0.05	0.05	−0.17	1.00	−0.18	−0.00	0.00	0.13	0.12
g_{B3}	0.24	−0.04	0.21	0.05	0.06	0.06	−0.15	−0.18	1.00	−0.00	0.00	0.14	0.14
β_{A1}	−0.00	0.02	−0.00	0.17	0.18	0.17	−0.00	−0.00	−0.00	1.00	0.18	0.00	0.00
β_{A3}	−0.00	−0.04	0.01	0.16	0.17	0.16	0.00	0.00	0.00	0.18	1.00	−0.00	−0.00
β_{B2}	−0.00	0.00	−0.02	−0.00	−0.00	−0.00	0.15	0.13	0.14	0.00	−0.00	1.00	0.17
β_{B4}	−0.00	0.01	−0.04	0.00	0.00	0.00	0.15	0.12	0.14	0.00	−0.00	0.17	1.00

The rows fix what to read in it. The supply correlates with every coil at +0.22 to +0.28: the return is read, so a lower supply leaves a larger difference to the return, and the same heat then needs a smaller conductance. The two loads correlate at −0.18 through the loop meter, which reads their sum. Across the two halls the other correlations are at most 0.07 in size, so the halls are otherwise nearly separate. The fractions read by an inlet sensor correlate with their own hall's coils; the six unread fractions correlate with nothing.

12.2.3 What the record shows of its events

These are the methods of Parts II and III computed at 01:00, stated once here for Part IV. A one-window fit has its own loads, so a load that drifts between windows is no misspecification of it; it would matter only to a forecast. Fitted one window at a time, the entries do drift: per window, 0.28135 (0.02802) kW on Q_{A} and 0.19141 (0.02404) kW on Q_{B} over windows 1 to 73, and −0.00333 (0.00044) on g_{B2} over windows 37 to 73, by least squares whose standard errors ignore autocorrelation.

The window's own test. The chi-square test of Chapter 7, on the reading rows with the effective degrees of freedom ν of equation (7.14), gives $X^2 = 0.755$ on 3.691, an upper tail of 0.925, and a largest normalized residual, equation (7.15), of −0.718 on the discharge of unit B2. The hall-A meter's own is −0.20. The window's test sees nothing.

The coil of unit B2. The refit stacks windows 1 to 36, with the shared coils and fractions as one entry each and T_s, the two loads and g_{B2} free per window. The mean of g_{B2} over those first three hours is −0.0022. Window 73's entry, −0.1444, lies −4.90 of its spreads from that mean, a difference of −0.1423. A second refit of window 73 with g_{B2} pinned at that mean is misspecified: its reading rows give $X^2 = 28.421$ on effective degrees of freedom of 4.644, with an upper tail below 0.001 and a largest normalized residual of −5.078 on the discharge of B2, and the supply reading's is +4.12. The free entry absorbs the derate, so the plain fit needs no further term. The derate is therefore already in every scenario read from this fit, and what is in it is the coil's level at 01:00, not a slope.

The hall-A meter. The meter reconciliation of Section 7.2.5 does not act on one window; it acts on the record stacked to window T, the same stacked rows with one constant bias entry on one meter, prior $\mathcal{N}(0,\ 10^2)$ in kW, and the Laplace Bayes factor of equation (7.17) against no bias. Each window of that stack has eighteen rows: its sixteen readings, its T_s prior and its g_{B2} prior, with one prior row for each shared entry. Table 12.7 gives the results. Each row is a stack; each of the three meters has its logarithm of the Bayes factor, its bias estimate and that estimate's spread. Figure 12.3 draws the logarithms against T, with the level $\ln 20 = 2.9957$ dashed.

Only the hall-A meter passes the level. It is past it at $T = 24$, under it at 36, past it again at 48, and above 10 at 72 and 73, where it stands 8.742 above the loop meter and 12.209 above the hall-B meter. This is one night's draw. Over 100 nights of fresh sensor errors on the same truth, $\ln B$ of this meter has medians 1.512, 3.085, 4.379, 5.986 and 7.725 at $T = 24$, 36, 48, 61 and 73, and the fractions of nights past the level are 0.33, 0.52, 0.73, 0.86 and 0.86. At $T = 73$ the range from the 10th to the 90th percentile is 2.681 to 13.968: this night's 10.139 is above the median, and 0.14 of nights are not yet past the level at 01:00. So no first window is reported as a date; by 01:00 this night's record shows a high hall-A meter, and the one-window fit absorbs a part of it.

The fit carrying the bias. A refit of window 73 subtracts one bias entry from the meter's reading, its prior the stack of windows 1 to 72, 4.636 (0.918) kW. It gives $Q_{\mathrm{A}} = 359.815$ (3.499) kW

Table 12.7: The meter reconciliation on the record stacked to window T: $\ln B$ against no bias [–], bias estimate and its spread [kW].

		y_{Q_A}		y_{Q_B}		$y_{Q_{loop}}$	
T [–]	Clock [h:min]	$\ln B$ [–]	Bias (sd) [kW]	$\ln B$ [–]	Bias (sd) [kW]	$\ln B$ [–]	Bias (sd) [kW]
12	19:55	−1.198	1.74 (2.20)	−0.670	2.73 (2.20)	−1.204	−1.34 (2.67)
24	20:55	4.048	5.43 (1.58)	−1.273	1.58 (1.57)	−0.144	−3.37 (1.92)
36	21:55	2.881	4.08 (1.29)	−2.043	0.05 (1.29)	−1.184	−1.85 (1.58)
48	22:55	4.065	3.99 (1.12)	−2.142	−0.46 (1.12)	−1.271	−1.67 (1.37)
61	00:00	5.227	3.90 (1.00)	−2.294	0.09 (1.00)	−0.720	−2.05 (1.22)
72	00:55	10.197	4.64 (0.92)	−2.114	0.59 (0.92)	1.239	−2.96 (1.12)
73	01:00	10.139	4.60 (0.91)	−2.070	0.64 (0.91)	1.396	−3.01 (1.11)

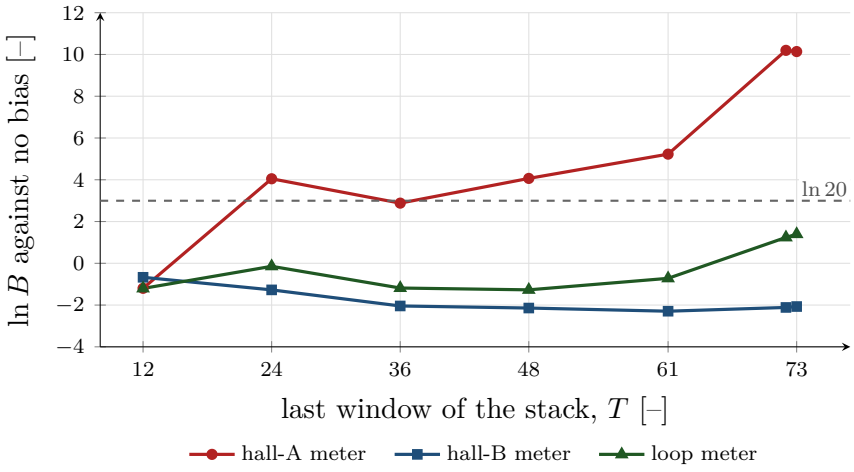


Figure 12.3: The logarithm of the Bayes factor of a bias on each meter against no bias, on the record stacked to window T (Table 12.7).

against the plain fit's 363.201 (3.434) kW, a difference of −3.385 kW, −0.99 plain spreads, and through the loop meter $Q_{\mathrm{B}} = 296.614$ against 296.012 kW.

The fit the answers read. The reported fit is the plain fit of the plant's rows, for three reasons: it passes its window's test, the one fit of a Part should be the one its plant states, and the bias is the size of the meter's noise. The fit carrying the bias is a second named route, set beside each answer.

12.3 The answers (step 4)

Every answer is read from the pair $(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of Section 12.2.2, and that value is the one reported. A scenario temperature is carried by its gradient at $\hat{\mathbf{z}}$, equation (12.1). The exact posterior's draws stand beside each reported value where they were run, which is at the 16.0°C node.

12.3.1 The inlets at the raised set point

Operator question. If I raise the supply temperature 2 K, what is the probability the inlets stay under the limit? The question is about all ten inlets at once, at the set point the operator would enter, at this window's loads and coils.

General form. A prediction under a threshold, equation (2.12), here on a maximum as in equation (8.6), and read on the requirement's side.

Specialized form. The question is equation (2.12) with $\mathbf{z}$ the reduced state (12.3), $f(\mathbf{z}) = -\max_k T_k^*(\mathbf{z})$ at $T_s^* = 16°\mathrm{C}$, $f^* = -27°\mathrm{C}$, and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal (12.9). The assurance κ is the operator's. By equation (12.6) the inlet's gradient has the entries

$$\frac{\partial T_k^*}{\partial T_s} = 0, \quad \frac{\partial T_k^*}{\partial Q_h} = \frac{1}{\sum_{u\in h} G_u} - \frac{1-\beta_k}{\sum_{u\in h} C_u}, \quad \frac{\partial T_k^*}{\partial g_u} = -\frac{G_u\,(T_{r,h} - T_s)}{\sum_{v\in h} G_v}, \quad \frac{\partial T_k^*}{\partial \beta_k} = T_{r,h} - T_{c,h}, \tag{12.10}$$

with every difference $T_{r,h} - T_s$ and $T_{r,h} - T_{c,h}$ at $\hat{\mathbf{z}}$, and they do not depend on the pin. Stacking the ten gradients as the rows of $\mathbf{D}^*$, as in equation (8.9), the inlets are jointly normal, and the requirement is the corner of that normal,

$$\begin{gathered} \mathbf{T}^* \mid \mathbf{y} \approx \mathcal{N}(\mathbf{m}^*,\ \mathbf{D}^*\boldsymbol{\Sigma}\mathbf{D}^{*\mathsf{T}}), \\ \Pr(\max_k T_k^* < 27°\mathrm{C} \mid \mathbf{y},\ T_s^* = 16°\mathrm{C}) = \Pr(T_k^* < 27°\mathrm{C} \text{ for all ten} \mid \mathbf{y}), \end{gathered} \tag{12.11}$$

with $\mathbf{m}^*$ the ten inlets at $\hat{\mathbf{z}}$. One choice enters here: the joint is the delta method's normal, which gives the corner from $\hat{\mathbf{z}}$, $\boldsymbol{\Sigma}$ and one matrix, and whose linearity the draws test. A corner of ten dimensions has no closed form; it is evaluated by seeded numerical integration of the normal, with an absolute error requested of 10^{-6}. The union bracket (8.13) applies to its complement with ten terms.

Building blocks. The MAP state and $\boldsymbol{\Sigma}$ (Section 12.2.2); the gradients (12.10); the integration of the corner.

The joint at the node. By equation (12.11),

$$\Pr(\text{all ten} < 27°\mathrm{C} \mid \mathbf{y},\ T_s^* = 16°\mathrm{C}) = 0.80983,$$

and the exact posterior's draws give 0.81026 (se 0.00062), a gap of +0.00043. The exceedance is 0.19017, between the largest marginal exceedance, 0.16300 at rack B3, and the union bound,

0.19587. No stated assurance from 0.90 up is met, on either value. The fit carrying the bias gives 0.81013.

12.3.2 The rack that carries the tail

Operator question. Which rack carries it, and what stands behind its band? The question is which inlet comes nearest the limit at the node, and which entries of the state make its spread.

General form. An estimation of a derived quantity, equation (2.9), with its central interval, equation (2.11), and its variance split by entry, as in equation (5.14).

Specialized form. The question is equations (2.9) and (2.11) with $f(\mathbf{z}) = T_k^*(\mathbf{z})$ and the Laplace normal, so the band is $T_k^*(\hat{\mathbf{z}}) \pm 1.960\,\mathrm{sd}$ by equation (12.1). The budget of T_{B3}^* has the gradient's five nonzero entries,

$$\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} = \sum_i [\mathbf{g}]_i^2\,\Sigma_{ii} + \sum_{i \neq j} [\mathbf{g}]_i [\mathbf{g}]_j \Sigma_{ij}, \tag{12.12}$$

with no cross term in β_{B3}, whose column is empty. The gradient's signs fix those of the others. On hall B's coils it is negative, and the coils' mutual correlations are negative, so each such term is a negative covariance on two same-signed gradients. On Q_{B} the gradient is positive and its correlations with the coils are positive, so those terms are a positive covariance on opposite-signed gradients. Every cross term is negative.

Building blocks. The gradients (12.10); $\boldsymbol{\Sigma}$ (Section 12.2.2); the budget (12.12).

The ten inlets. Table 12.8 gives all ten: each row a rack, the columns its inlet at the present state with its spread, its inlet at the node with its spread and band, its probability of staying under the limit, and the exact posterior's median and probability. Figure 12.4 draws them, with the horizontal axis the racks, the vertical axis the inlet temperature, a gray bar for the present band, a blue bar for the band at the node, a red line at the limit, and a dotted tick for the stated truth at the node, a check marker for Section 12.4.4. The unread racks' bars are longer than the read racks', and the bars move up together.

Table 12.8: The ten inlets at the present state and at the node $T_s^* = 16\,^{\circ}\mathrm{C}$, window 73. Racks A1, A3, B2 and B4 have sensors. Medians and probabilities beside them from the exact posterior.

Rack [–]	T_k [°C]	sd [K]	T_k^* [°C]	sd [K]	2.5 % [°C]	97.5 % [°C]	Pr(< 27) [–]	Exact median [°C]	Exact Pr [–]
A1	19.4840	0.2876	21.5893	0.3037	20.9941	22.1845	1.00000	21.5881	1.00000
A2	20.4874	1.0112	22.5927	1.0157	20.6018	24.5835	0.99999	22.5888	0.99999
A3	22.0065	0.2876	24.1118	0.3037	23.5166	24.7070	1.00000	24.1108	1.00000
A4	21.1127	1.0112	23.2180	1.0158	21.2271	25.2089	0.99990	23.2175	0.99991
A5	22.9887	1.0114	25.0940	1.0160	23.1026	27.0854	0.96967	25.0899	0.97002
B1	21.0145	0.8288	23.1198	0.8344	21.4844	24.7553	1.00000	23.1165	0.99999
B2	22.1185	0.2821	24.2238	0.2985	23.6389	24.8088	1.00000	24.2223	1.00000
B3	24.0747	0.8293	26.1800	0.8349	24.5436	27.8164	0.83700	26.1767	0.83720
B4	22.7126	0.2821	24.8179	0.2985	24.2330	25.4029	1.00000	24.8175	1.00000
B5	22.5446	0.8289	24.6499	0.8345	23.0144	26.2855	0.99757	24.6458	0.99758

Rack B3. Rack B3 is nearest the limit, and the upper end of its band in Table 12.8 lies past it; rack A5 comes next and rack B5 after it. The fit carrying the bias gives rack B3

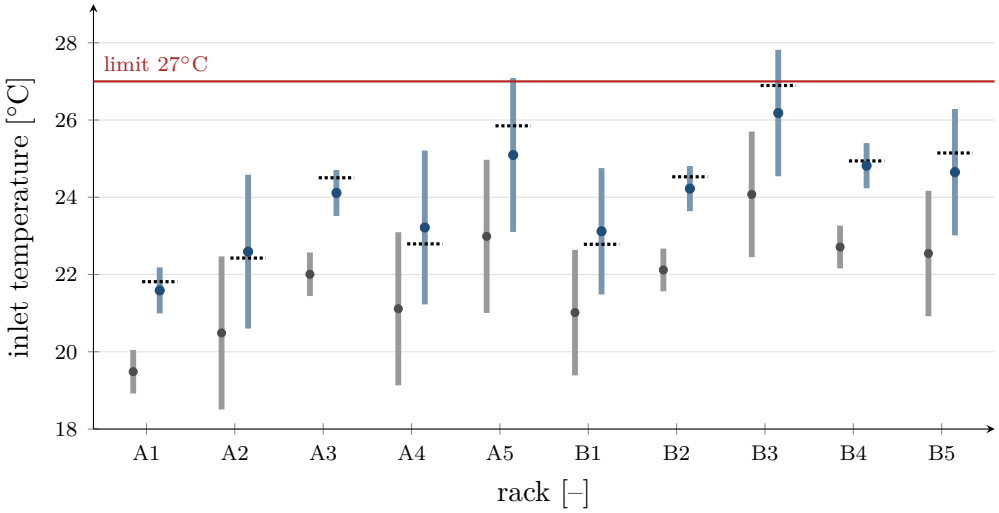


Figure 12.4: The ten inlets with their central 95 % bands. Gray: the present state. Blue: the node $T_s^* = 16°\mathrm{C}$. Dotted: the stated truth at the node, a check marker (Section 12.4.4).

26.1849 (0.8365)°C and 0.83509, and rack A5 25.0707 (1.0071)°C and 0.97229. By equation (12.10) the gradient of T^*_{B3} is 0.0333 K/kW on Q_{B}, −5.5461, −4.6226, −5.1115 K on g_{B1}, g_{B2}, g_{B3}, and 10.2004 K on β_{B3}.

The budget. Figure 12.5 draws the budget (12.12), one bar per entry and one for the cross terms, the horizontal axis in percent of the variance $0.69705\,\mathrm{K}^2$. The terms are $0.66591\,\mathrm{K}^2$ for β_{B3}, 0.01913, 0.01859 and 0.01803 for g_{B1}, g_{B3} and g_{B2}, 0.01302 for Q_{B}, and −0.03762 for the cross terms. With β_{B3} known the spread is 0.1765 K. The answer is one survey's fraction, 0.50 ± 0.08.

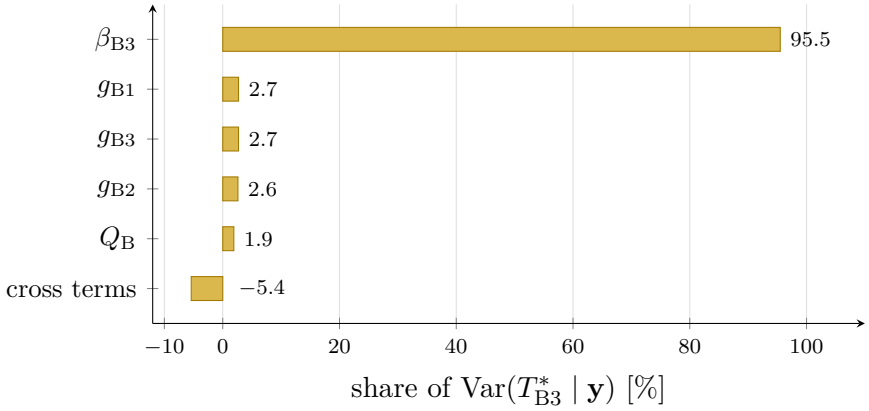


Figure 12.5: The variance budget of the inlet of rack B3 at the node, equation (12.12), in percent of $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$.

The correlations at the node. The correlation of T^*_{A5} and T^*_{B3} is 0.0108, and that of T^*_{B3} and T^*_{B4} is 0.0479. Hall A alone has the joint 0.969566 and hall B alone 0.835050, and their product, independence of the halls, is 0.809636 against the joint's 0.80983; the fit carrying the bias gives the halls 0.97221 and 0.83310. Here the joint adds little to the marginals, because the racks that carry the tail share little beyond entries the readings have fixed, while Section 8.3.4 found otherwise on a poorly read plant.

12.3.3 The sensors and the unread racks

Operator question. Is the raise seen by the sensors on the screen? The sensors see the raise and not the tail; the question is how much of the tail lies in the racks the operator can see.

General form. The threshold form, equation (2.12), on a maximum over a subset of the racks.

Specialized form. The question is equation (2.12) with $f = -\max_{k\in S} T_k^*$ and $f^* = -27°\mathrm{C}$, for S the four measured racks and S the six unread racks, each the corner of the marginal of $\mathbf{T}^*$ in equation (12.11) on those rows of $\mathbf{D}^*$.

Building blocks. The joint normal (12.11) of Section 12.3.1; the same integration.

The two subsets. The measured four alone have the joint 1.000000, to six decimals, with the hottest of them rack B4 at 24.8179°C and a spread of 0.2985 K. The unread six alone have 0.809827, with the exact posterior's 0.810259 (se 0.00062) and 1.000000 for the measured four beside them; the fit carrying the bias gives 0.810126 and 1.000000. All of the tail is in racks with no sensor. For reference, the same ten at the present state have the joint 0.99975, with rack B3 at 24.0747°C and 0.8293 K.

12.3.4 The raise as a size

Operator question. Does the probability fall with the raise, and how large a raise holds at a stated assurance? The premise of the raise was a saving, so the question is what each size costs on the risk side.

General form. The threshold form, equation (2.12), read as a function of the pinned value, and its acceptance rule read backwards.

Specialized form. The question is equation (2.12) with f and f^* as in Section 12.3.1 and the pin a free argument. By equation (12.10) neither the gradients nor their covariance depend on the pin, so only the means move, and by equation (12.6) they move together. The probability therefore falls with T_s^*, and the largest set point at an assurance κ,

$$T_s^{\kappa} = \max\{T_s^* : \Pr(\max_k T_k^* < 27°\mathrm{C} \mid \mathbf{y},\ T_s^*) \geq \kappa\}, \tag{12.13}$$

is found by bisection on the joint. The raise is measured from the set point 14°C.

Building blocks. The joint (12.11) at each T_s^*; the operator's κ.

The probability against the raise. Table 12.9 gives the joint along the pin: each row a pinned value, the columns the raise, the inlet of rack B3, its probability alone, the joint and the exact posterior's joint. Figure 12.6 draws the last two against T_s^*, with the three assurances as dotted levels and a marker where each meets the joint. The joint falls at every step, as the physics requires, and its largest step on the grid is from 16.5 to 17.0°C, 0.26693. The exact column's standard errors are at most 0.00078, and its widest departure is at 17.0°C, 0.32328 (se 0.00074) against 0.32115.

Table 12.9: The probability against the raise, pinned, window 73.

T_s^* [°C]	Raise [K]	T_{B3}^* [°C]	Pr($T_{B3}^* < 27$) [–]	Pr(all < 27) [–]	Exact [–]
14.00	0.00	24.1800	0.99963	0.99957	0.99954
14.50	0.50	24.6800	0.99727	0.99687	0.99692
15.00	1.00	25.1800	0.98537	0.98326	0.98346
15.50	1.50	25.6800	0.94307	0.93436	0.93457
16.00	2.00	26.1800	0.83700	0.80983	0.81026
16.50	2.50	26.6800	0.64926	0.58808	0.58901
17.00	3.00	27.1800	0.41466	0.32115	0.32328
17.50	3.50	27.6800	0.20770	0.11546	0.11632

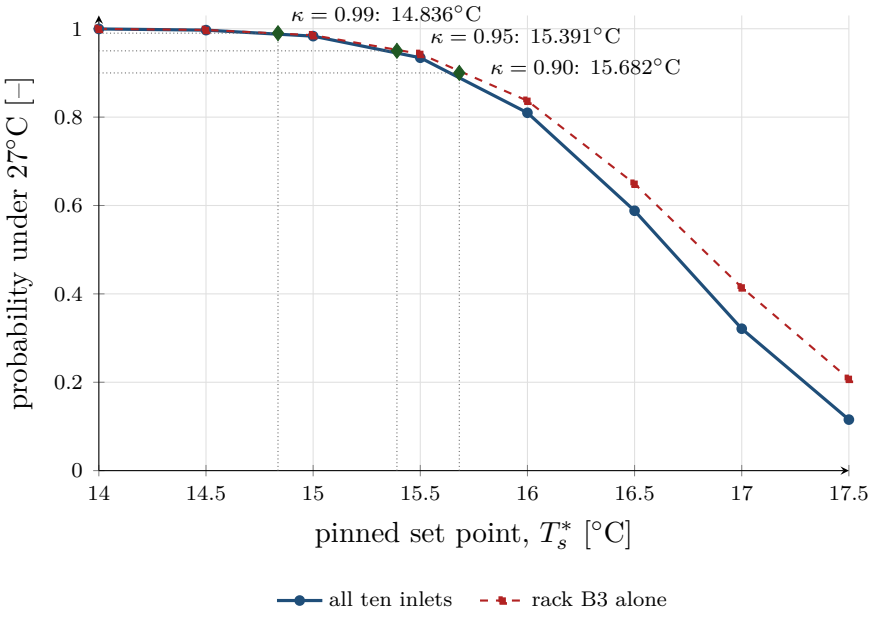


Figure 12.6: The probability of staying under the limit against the pinned set point (Table 12.9). Diamonds: the largest set point at each stated assurance, equation (12.13).

The size at a stated assurance. By equation (12.13) the largest set point is 15.682°C at $\kappa = 0.90$, a raise of 1.682 K; 15.391°C at 0.95, 1.391 K; and 14.836°C at 0.99, 0.836 K. The exact posterior gives 15.683 (se 0.002), 15.392 (se 0.003) and 14.840°C (se 0.005), and the fit carrying the bias 15.682, 15.390 and 14.834°C. This is the comparison the premise asks for. The saving at each size is not computed, since the plant has no chiller-power row, and no size is named as the choice. These sizes move with the window's draw by more than with any route; Section 12.4.3 gives the spread.

12.3.5 The pin and the shift

Operator question. Is the raise a pinned value or a shift? The operator enters 16°C; the drill question says 2 K; and the fit's supply entry is not 14.

General form. The threshold form, equation (2.12), at two nodes.

Specialized form. The question is equation (2.12) at two values of the actuator in the node (12.1): the pin of Section 12.3.1, and the shift $T_s^* = T_s + 2\,\mathrm{K}$, which keeps the entry T_s in the gradient with coefficient one. The two gradients differ by that entry alone, $\mathbf{g}_{\text{shift}} = \mathbf{g}_{\text{pin}} + \mathbf{e}_{T_s}$, and the shift's is the present temperature's. So the shift keeps the present band exactly, and the pin has

$$\mathrm{Var}(T_k^* \mid \mathbf{y},\ \text{pin}) = \mathrm{Var}(T_k - T_s \mid \mathbf{y}) = \mathrm{Var}(T_k \mid \mathbf{y}) + \Sigma_{T_sT_s} - 2\ \mathrm{Cov}(T_k, T_s \mid \mathbf{y}). \qquad (12.14)$$

A reading at rack k pins T_k and not $T_k - T_s$, so the supply's spread returns in the difference.

Building blocks. The gradients (12.10); $\Sigma_{T_sT_s}$ and the T_s row of $\mathbf{\Sigma}$ (Section 12.2.2); the joint (12.11).

The two nodes. The fit's T_s entry is 13.8947°C, so the pin moves every mean by $16 - 13.8947 = 2.1053$ K, of which 0.1053 K is the window's supply reading; the return of hall A goes from 29.8674 to 31.9727°C, and unit B2's heat is 92.5754 kW at both. The shift has $T_{\mathrm{B3}}^* = 26.0747$ (0.8293)°C and the joint 0.84615 (exact 0.84609, se 0.00057; fit carrying the bias 0.84625), against the pinned node of Section 12.3.1. At 2 K their gap crosses no κ among those stated. At another size it does: at 1.5 K the pin has 0.93436 and the shift 0.95075, either side of 0.95, and both stand. The largest raise, pin against shift, is 1.682 against 1.795 K at $\kappa = 0.90$, 1.391 against 1.506 K at 0.95, and 0.836 against 0.955 K at 0.99. A measured rack's spread is the present one under the shift, 0.2821 K at rack B4 and 0.2876 K at rack A3, and under the pin it widens, to 0.2985 and 0.3037 K. The covariance term of equation (12.14) is negligible, $0.000013\,\mathrm{K}^2$ at rack B4 and $0.000102\,\mathrm{K}^2$ at rack B3 against $\Sigma_{T_sT_s} = 0.009517\,\mathrm{K}^2$, so the widening is nearly all the supply's variance. Removing T_s from the gradient did not narrow the band; it widened it. A band at a node is not the present's band moved.

12.4 The checks (step 5)

Every route here is computed at window 73 and the node $T_s^* = 16$°C unless it says otherwise.

12.4.1 The limit of exact rows and the curvature term

The check fits the plant's own full state with the node, 59 entries: the reduced entries, the twenty present temperatures and the twenty scenario temperatures, with the physical rows (12.2) at a finite tolerance on both copies, and the reading rows on the present copy only. Table 12.10

gives the result: each row a tolerance, the columns the node's inlet of rack B3 and the load of hall A with their spreads, and the iterations. The scenario's rows add nothing to the fit of the entries: the load matches the reduced state from 0.01 on, and the spread of the inlet goes to the reduced state's as the tolerance goes to zero.

Table 12.10: The limit of exact physical rows on the 59-entry state with the node.

Tolerance σ_p [kW or K]	T^*_{B3} [°C]	sd [K]	Q_{A} [kW]	sd [kW]	Iterations [–]
0.1	26.1801	0.8409	363.2003	3.4343	3
0.01	26.1800	0.8350	363.2006	3.4343	3
0.001	26.1800	0.8349	363.2006	3.4343	2
reduced state	26.1800	0.8349	363.2006	3.4343	–

The covariance (12.9) is the Gauss–Newton form. The objective's full Hessian at $\hat{\mathbf{z}}$ gives spreads of 0.83492, 0.29846 and 1.01603 K for T^*_{B3}, T^*_{B4} and T^*_{A5} against 0.83490, 0.29846 and 1.01602 K, and the joint 0.80982 against 0.80983, a gap of −0.00001.

12.4.2 The exact posterior and the fit carrying the bias

The exact posterior's draws carry standard errors of 0.00062 on the joint and 0.0017 on the median of T^*_{B3}. The joint's gap of +0.00043 is below its standard error, so no correction to it is resolved at this count. The median's gap of −0.0033 is larger than its standard error, and every rack's exact median in Table 12.8 lies below its reported value. The band's ends differ by +0.0014 and −0.0005 K. The exact posterior's largest set points and joints along the pin stand in Section 12.3.4. The fit carrying the bias moves the joint by +0.00030 because two larger moves offset: hall A up 0.00264, hall B down 0.00195, the loop meter moving Q_{B} up 0.60 kW.

12.4.3 Sensitivity

Each choice of Section 12.1.4 is varied at the pinned node, as T^*_{B3} (spread) and the joint. Three of these variations cross a stated assurance, so both values are stated there and neither is chosen.

The coil and the warmer water. This is the most sensitive choice in the chapter. Every coil's g lower by 0.02 gives 26.4887 (0.8357)°C and 0.68166; lower by 0.05, 26.9634 (0.8370)°C and 0.43146; lower by 0.10, 27.7870 (0.8393)°C and 0.08288. No row of the record can tell these apart.

The survey's spread. The unread fractions at a spread of 0.04 give 26.1800 (0.4445)°C and 0.96730, above 0.90 and 0.95; at 0.12, 26.1800 (1.2367)°C and 0.64495. The reported 0.80983 meets neither of those assurances. The survey's center for rack B3 moved by its own spread gives 26.9960 (0.8353)°C and 0.48588 at +0.08, and 25.3639 (0.8346)°C and 0.94312 at −0.08, above 0.90.

The actuator. An error of standard deviation 0.1 K on the pin gives 26.1800 (0.8409)°C and 0.80776.

One window's draw. The plain fits of windows 70, 71, 72, 74, 75 and 76 give T^*_{B3} of 26.3611, 25.9854, 26.1759, 25.6879, 26.1398 and 26.1795°C, and the joint 0.73777, 0.87365, 0.80744, 0.92596, 0.82467 and 0.80529. Over windows 61 to 85 the joint has a minimum of 0.6221, a median of 0.8247 and a maximum of 0.9260, the trend and the window's noise together. Window 74 passes 0.90 where window 73 does not; no neighbor reaches 0.95. The largest set

point at $\kappa = 0.95$ on these fits is 15.213, 15.601, 15.373, 15.842, 15.435 and 15.384°C against window 73's 15.391°C, a range over windows 70 to 76 of 15.213 to 15.842°C; at 0.90 it runs 15.501 to 16.130°C and at 0.99, 14.664 to 15.290°C. This spread is why the fit's window is named in the log.

12.4.4 The answers against the stated truth

The record is generated from a stated truth (Section 12.6.4), and Table 12.11 collects the checks: each row a route, the columns the reported value, the value beside it, and the gap, the second minus the first. The last row is the truth, with the normalized difference in place of the gap.

Table 12.11: The checks and their gaps at the node $T_s^* = 16$°C. Exact draws: importance sampling, 400 000 draws, effective size 396 168. z is the reported value less the truth over its spread.

Check [unit]	Reported [row's]	Beside it [row's]	Gap [row's]
Joint, exact draws [–]	0.80983	0.81026 (se 0.00062)	+0.00043
T^*_{B3} median, exact [°C]	26.1800	26.1767 (se 0.0017)	−0.0033
T^*_{B3} band ends, exact [°C]	24.5436, 27.8164	24.5450, 27.8158	+0.0014, −0.0005
$\Pr(T^*_{\mathrm{B3}} < 27)$, exact [–]	0.83700	0.83720 (se 0.00059)	+0.00021
$\mathrm{sd}(T^*_{\mathrm{B3}})$, Newton against Gauss–Newton [K]	0.83490	0.83492	below the print
Joint, Newton [–]	0.80983	0.80982	−0.00001
$\sigma_p \to 0$, T^*_{B3} [°C], sd [K] at 0.1, 0.01, 0.001	26.1800, 0.8349	26.1801, 0.8409; 26.1800, 0.8350; 26.1800, 0.8349	to zero
Shift form, joint [–]	0.80983	0.84615	+0.03632
Shift at 1.5 K against the pin at 15.5°C, joint [–]	0.93436	0.95075	crosses 0.95
Largest raise, pin against shift, $\kappa = 0.90, 0.95, 0.99$ [K]	1.682, 1.391, 0.836	1.795, 1.506, 0.955	–
Largest set point, exact, $\kappa = 0.90, 0.95, 0.99$ [°C]	15.682, 15.391, 14.836	15.683, 15.392, 14.840	–
Fit carrying the bias, joint [–]	0.80983	0.81013	+0.00030
Fit carrying the bias, largest set point [°C]	15.682, 15.391, 14.836	15.682, 15.390, 14.834	–
Stated truth, T^*_{B3} [°C]	26.1800	26.8937	$z = -0.85$

Only the pin and the shift, at 1.5 K, cross a stated κ in the table; both are stated. One reported value stands at each node, with the routes beside it.

The do-it-and-see check. The scenario is run on the plant's true state at window 73, with no sensor noise, and every inlet stays under the limit: the true maximum is 26.8937°C at rack B3, 0.1063 K inside. The ten values of z run from −1.29 to +0.42, all inside the bands, as the dotted ticks of Figure 12.4 show. The true fractions explain the largest misses in kelvin, −0.7560 at rack A5 and −0.7137 at rack B3: β_{B3} is truly 0.57 against the survey's 0.50, +0.875 prior spreads, and β_{A5} is truly 0.51 against 0.45, +0.750.

Of the fit's entries, Q_{A} stands $z = +2.44$ from its true 354.837 kW, the window's high meter reading, and g_{A3} stands +1.58; the rest are at most 1.12 in size. The bias-carried fit's 359.815 kW lies nearer the true load. The meter's stated bias, 4.0 kW, stands $z = +0.69$ from the reconciliation's 4.636 (0.918) kW. A probe reading the truth's present inlet at rack B3 would read 24.857°C and leave the joint at 0.64535.

Beyond the question's window. In the truth only, the plant at 16.0°C from then on is under the limit through window 80, 01:35, and first over at window 81, 01:40, at rack B3, 27.009°C, forty minutes later; it reads 28.735°C at window 144. The node is one window's plant, not a shift's promise, and a forecast with a raise is not this chapter's.

12.5 The report (step 6)

The sensors' margin does not describe the raise. The read inlets end 2.33 K inside the limit by the screen's sum, and the chance that all ten stay under it is 0.81, because rack B3, unread, sits at 26.18 ± 0.83°C. That spread is one survey's fraction. A smaller raise buys a higher assurance at a price this plant cannot state.

12.5.1 The answers and their bands

Table 12.12 gives the answers in the order of Table 12.1. Each row is a question; the second column gives the reported answer, from the plain fit by the delta method, and the third what stands beside it.

Table 12.12: The report at window 73, 01:00, for the node $T_s^* = 16$°C at this window's loads and coils. The limit is the plant's; the assurances are the operator's.

Question	Answer, reported	Beside it
If I raise the supply temperature 2 K, what is the probability the inlets stay under the limit?	0.80983	exact 0.81026; fit carrying the bias 0.81013
Which rack carries it, and what stands behind its band?	rack B3, 26.1800°C, band 24.5436 to 27.8164°C, 0.83700 alone; 95.5 % of the variance its survey fraction	exact median 26.1767, band 24.5450 to 27.8158°C, 0.83720 alone; fit carrying the bias 26.1849°C, 0.83509 alone; sd 0.1765 K with the fraction known
Is the raise seen by the sensors on the screen?	the sensors see the raise and not the tail: measured four 1.000000, unread six 0.809827	exact 1.000000, 0.810259; fit carrying the bias 1.000000, 0.810126; hottest measured, rack B4, 24.8179°C
Does the probability fall with the raise, and how large a raise holds at a stated assurance?	falling at every step; largest set point 15.682, 15.391, 14.836°C at $\kappa = 0.90, 0.95, 0.99$	raises 1.682, 1.391, 0.836 K; exact 15.683, 15.392, 14.840°C; fit carrying the bias 15.682, 15.390, 14.834°C; at 0.95, windows 70 to 76 15.213 to 15.842°C; saving not computed
Is the raise a pinned value or a shift?	pinned, 0.80983, $T^*_{\mathrm{B3}} = 26.1800$°C	shift 0.84615 (exact 0.84609, fit carrying the bias 0.84625), 26.0747°C; no stated κ between them at 2 K; at 1.5 K pin 0.93436, shift 0.95075; largest raise, shift, 1.795, 1.506, 0.955 K

12.5.2 The entry in the shift log

The operator's entry reads: "01:00, window 73. What-if: set point 14 → 16°C at this window's loads. Read inlets stay under (hottest B4 24.82). All ten under the 27°C limit, probability 0.81 (exact 0.81); rack B3, unread, 26.18°C, band 24.54 to 27.82, set by its survey fraction. Largest

set point at assurance 0.95: 15.39°C at this window, windows 70 to 73 15.21 to 15.60; at 0.99: 14.84°C, 14.66 to 15.05. Saving not computed. Steady-state answer at 01:00 loads and coils only; the coil of B2 and the loads are moving. Action: not decided here."

Common mistake

Adding the raise to the hottest sensor and reading what is left as the plant's margin. The raise moves every inlet alike, and the inlet nearest the limit is one no sensor reads; the probability is read at all ten.

12.6 The next reading (step 7)

The report leaves a joint whose shortfall is 0.16300 of 0.19017 at rack B3, the rest mostly rack A5; both are survey fractions. A further reading is a further row, and what it does to the entries at the node is known before its value is.

12.6.1 The preposterior at the node

A candidate adds rows $\mathbf{J}_a$ to the present state. Its covariance is the preposterior of equation (4.65), $\mathbf{\Sigma}_a = (\mathbf{\Sigma}^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a)^{-1}$, by the rank-one form (4.66) for one row, and the node's joint is the corner of $\mathcal{N}(\mathbf{m}^*,\ \mathbf{D}^*\mathbf{\Sigma}_a\mathbf{D}^{*\mathsf{T}})$ at the means the reading is expected to leave. A probe at rack k of 0.3 K reads the present inlet, so its row is $\mathbf{g}_{T_k}^{\mathsf{T}}/0.3$, with T_s's entry included. A re-survey adds a prior row on β_{B3} of 0.02. A permanent inlet sensor at rack B3, the instrument at the place of the question, is that probe's row in six windows, a choice that takes the six as repeats of 01:00's state and that answers 30 minutes later than a probe; a second supply sensor adds a row on T_s; a clamp-on meter adds a row on Q_{B}. The means move with the reading, by the gain of equation (4.66),

$$\mathbf{m}_a^*(r) = \mathbf{m}^* + \frac{\mathbf{D}^*\mathbf{\Sigma}\,\mathbf{g}_{T_k}}{\mathbf{g}_{T_k}^{\mathsf{T}}\mathbf{\Sigma}\,\mathbf{g}_{T_k} + 0.3^2}\,(r - T_k(\hat{\mathbf{z}})), \tag{12.15}$$

for a probe's reading r. By the law of total probability the expected joint over a reading's outcomes does not change with the reading. What changes is the chance of ending past the operator's assurance, here 0.95, the indicator form of equation (2.16), computed over 8000 outcomes with 8000 inner draws each and a standard error of that fraction of at most 0.0055. The printed expected joints run from 0.80305 to 0.81354, around the 0.81 they should keep. Their standard errors over the outcomes are at most 0.00367, and the inner draws, shared across a candidate's outcomes, add an error these leave out, so this is the draws' scatter.

12.6.2 The candidates

Table 12.13 compares them by what each does to the entries at the node: each row a candidate, the columns the spreads of β_{B3} and of T_{B3}^* after it, the joint at the present means, and the chance that the joint ends at or above 0.95.

The candidates that touch β_{B3} take the joint at the present means past 0.96; those that do not leave a chance of 0.0000 of reaching 0.95. By the chance of reaching 0.95 the order is the permanent sensor, 0.7175, on its six repeated windows and 30 minutes later; the re-survey, 0.6479; then the two probes and the one probe at rack B3, 0.6079 and 0.6036, whose difference is within their standard errors of 0.0055. A permanent sensor also makes the node's largest term testable window by window.

Table 12.13: Candidate readings at the node $T_s^* = 16°\text{C}$, from the posterior at window 73.

Candidate	sd β_{B3} [–]	sd T^*_{B3} [K]	Pr at present means [–]	Pr(after ≥ 0.95) [–]
None	0.0800	0.8349	0.80983	0.0000
Handheld probe at rack B3, sd 0.3 K	0.0303	0.2985	0.96435	0.6036
Handheld probe at rack A5, sd 0.3 K	0.0800	0.8349	0.83496	0.0000
Probes at B3 and A5	0.0303	0.2985	0.99448	0.6079
Re-survey of β_{B3}, sd 0.02	0.0194	0.2652	0.96627	0.6479
Permanent inlet sensor at B3, six windows as repeats of 01:00's state, 30 minutes later, sd 0.3 K	0.0183	0.1555	0.96722	0.7175
Second supply sensor, sd 0.1 K	0.0800	0.8322	0.81078	0.0000
Clamp-on on hall B's load, sd 1 kW	0.0800	0.8344	0.80996	0.0000

The probe's result. A probe at rack B3 moves the answer by what it reads, by equation (12.15), always with $\text{sd}(T^*_{\text{B3}}) = 0.2985\,\text{K}$: at 23.5°C the node has $T^*_{\text{B3}} = 25.6719°\text{C}$ and the joint 0.96734; at the fit's 24.075°C, 26.1803°C and 0.96434; at 24.5°C, 26.5560°C and 0.90129; at 25.0°C, 26.9981°C and 0.48695; at 25.5°C, 27.4402°C and 0.06813. A probe need not reassure. Chapter 18 ranks readings against a named worry. The choice stays with the operator.

12.6.3 What the analysis showed

The scenario adds nothing to the fit and needs no new reading. The physics is one line: every temperature moves together, and no unit heat moves. The answer is not one line, because the amount is added to a fraction no sensor reads, and the dashboard's 2.33 K describes four racks of ten. The raise's limit is set by a survey at rack B3. The most sensitive assumption, that a coil does not care about the water's temperature, cannot be tested from a record whose supply never moved beyond its reading's noise. And the pin, which looks as if it should narrow a band by removing an entry, widens it, because the sensors read temperatures and not their difference from the supply.

12.6.4 Limits of the model

The steady state. The node says where the plant settles at T_s^*, and nothing about the settling time or a controller's overshoot, as Section 12.1.2 says. The record holds no set-point step, so neither is read from it.

The loads and coils held. The node holds them at 01:00. The coil of unit B2 enters at its level at that window, and no later window's load or coil is in the answer.

The coil and the water; the fractions. Each g_u is taken as independent of T_s, and each β_k as independent of everything. Section 12.4.3 shows how much the first carries, and the record tests neither.

No chiller-power row. The plant has no row for the chiller's work, so no saving is computed at any size, and the risk side stands alone.

One window's noise. The answer is one window's fit, and its neighbors spread as Section 12.4.3 gives.

The plain fit. The reported fit leaves a meter's bias in its load, with the fit carrying the bias beside it.

Constructed readings. The record is constructed: generated from a stated truth of the plant with the stated sensor errors, and Section 12.4.4 compares the answers with that truth. It is not field data, and the agreement is a check of the method on its own model.

12.7 What to retain

Definition 12.1 (Scenario node)**.** A scenario node is a copy of the plant's temperature nodes at values of the pinned or actuated quantities not yet applied. It adds no reading row and no entry to the reduced state. Its temperatures are derived quantities $T_k^*(\mathbf{z})$ with variance $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$, so the posterior is unchanged; on the full state it adds one copy of the temperatures and of the physical rows. A requirement over several inlets is the corner of the joint normal $\mathcal{N}(\mathbf{m}^*,\ \mathbf{D}^*\boldsymbol{\Sigma}\mathbf{D}^{*\mathsf{T}})$, and the largest actuator value at an assurance κ reads that corner along the pin.

Three conclusions carry beyond the raise. A what-if needs no new evidence, and so it knows no more than the present posterior knows: an amount added to an unread place inherits that place's spread. A pinned actuator changes the gradient and therefore the band, which is not the present's band moved. And a choice no record can test, here a coil's indifference to the water, belongs beside the answer with its effect on the answer.

Table 12.14 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 12.14: The terms of the scenario node and where each is defined. Units for the two-hall shift plant.

Term	Expression	Units	Meaning	Defined in
Scenario node	$T_k^*(\mathbf{z})$, eq. (12.1)	°C	a temperature at an actuator value not applied	12.1.1
Physical rows	eq. (12.2), twenty rows	kW, K	the plant, per hall	12.1.2
Reduced state	eq. (12.3), 19 entries	°C, kW, –	the free entries	12.1.2
Inlet at the node	eq. (12.6)	°C	the pin plus two rises	12.1.4
Posterior				
Residual vector	eq. (12.8), 33 rows	–	the readings and priors of window 73	12.2.1
Fit carrying the bias	the meter's reading less a bias entry	kW	a second named route	12.2.3
Questions and answers				
Gradient at the node	eq. (12.10)	K, K/kW	no entry on the pinned T_s	12.3.1
Joint at the node	eq. (12.11)	–	all ten inlets under the limit	12.3.1
Variance budget	eq. (12.12)	K^2	entries behind the band	12.3.2
Largest set point	T_s^κ, eq. (12.13)	°C	the size at an assurance	12.3.4
Pinned variance	eq. (12.14)	K^2	why a band at a node widens	12.3.5
Candidate's chance	$\Pr(\text{after} \geq \kappa)$	–	ending past the assurance	12.6.1
Probe's result	eq. (12.15)	°C	the means after a reading r	12.6.1

12.8 Exercises

Exercise 12.1. *Application.* From the present inlet of rack B3 and the fit's T_s entry, compute T^*_{B3} at the pin. (By equation (12.6) a pin moves every temperature by its distance from the entry: $16 - 13.8947 = 2.1053\,\mathrm{K}$, and $24.0747 + 2.1053 = 26.1800°\mathrm{C}$.)

Exercise 12.2. *Derivation.* Derive the budget's share of β_{B3}. (By equation (12.10) its entry is the rise $10.2004\,\mathrm{K}$; with $\Sigma_{\beta_{\mathrm{B3}}\beta_{\mathrm{B3}}} = 0.08^2$ and no covariance with any entry the term is $0.66591\,\mathrm{K}^2$, of the total $0.69705\,\mathrm{K}^2$: 95.5 %.)

Exercise 12.3. *Application.* Give the spread of T^*_{B3} with β_{B3} known. (The fraction's column is empty, so its term comes off whole, by equation (12.12): $0.69705 - 0.66591\,\mathrm{K}^2$, a spread of $0.1765\,\mathrm{K}$.)

Exercise 12.4. *Derivation.* Explain why a measured rack's band widens at the pin. (Its sensor reads T_k, and the supply's sensor reads T_s. At the pin the variance is that of the difference, equation (12.14), with $\Sigma_{T_sT_s}$ added and twice the covariance taken off; the shift keeps the present variance. For rack B4 the spread goes from 0.2821 to $0.2985\,\mathrm{K}$.)

Exercise 12.5. *Application.* Give the preposterior spread of T^*_{B3} after a $0.3\,\mathrm{K}$ probe at rack B3. (The row is $\mathbf{g}^{\mathsf{T}}_{T_{\mathrm{B3}}}/0.3$ with the present gradient; by equation (4.65) the covariance after it is $\mathbf{\Sigma}_a$, β_{B3} narrows from 0.0800 to 0.0303, and the node's gradient on $\mathbf{\Sigma}_a$ gives $0.2985\,\mathrm{K}$, the spread of a measured rack in hall B at the pin.)

Chapter 13

Equipment Outages and Maintenance Windows

13.1 A hall with a unit to spare

It is 01:00, window 73 of the shift, and unit maintenance is on the schedule. The operator wants to know whether a unit in hall A can come down now. The screen shows hall A's two read inlets at 19.51°C (rack A1) and 22.03°C (rack A3), the warmer of them 4.97 K inside the 27°C limit, and hall A's read inlets are the cooler pair on the screen. The return reads 29.97°C, the hall meter 363.61 kW. Three units are running, so stopping one leaves two of three. The dashboard's reading is that the hall has a unit to spare: the read inlets have about 5 K in hand.

Those 5 K were measured with three units running, at two racks of five. What the screen cannot say is where the return and the cold aisle go with a unit stopped, and what then happens at the three racks no sensor reads. The question of this chapter is

what the probability is that every inlet of the hall stays under its limit with the unit stopped, at the present load and at the load the shift is heading for, and what that probability rests on.

It is answered on the one fit of Chapter 12, read at a scenario node with a unit's run state pinned. The fit is not repeated. One thing more is needed for the load ahead: a forecast, which is a second fit, on the hall-A meter alone, by the trend method of Chapter 11.

Jargon: N−1 contingency, run state, stop rule, rack air flow, forecast peak

An *N−1 contingency* is the what-if (contingency) analysis with one unit out of the plant's N. The *run state* of a unit says whether it is in service. A *stop rule* states what a stopped unit does to the air: which fans stop, and what the racks then draw. The *rack air flow* is the air the racks' own fans draw through them, here as a capacity rate in kW/K. The *forecast peak* is the largest load the operator plans against; on this record it is named in Section 13.1.3.

Key idea

On the plant's stop rule a stopped unit moves the return far and the cold aisle little, so each inlet moves by its own recirculation fraction of a widened gap, and its band widens with the gap, because the fraction no reading sees is multiplied by it. With the rack air flow held, every inlet moves by the return's rise instead. What a stopped fan does to the racks' air is a physical claim in the model, not a detail of it: it is not an entry of the state, so no band includes it, and it must be stated beside the answer.

13.1.1 What stopping a unit asks of the state

Chapter 12 defined the scenario node: a copy of the plant's temperatures at a value of an actuated quantity not yet applied, adding no reading row and no entry of the reduced state, with $\mathbf{\Sigma}$ carried through by $\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}$. There the actuator was the set point, which enters every row with coefficient one. A run state enters differently. It removes a term from a sum in the denominator of the return, so it changes the gain of every other entry on the inlet, and it removes one entry from the answer altogether. The node is therefore not the present plant moved by a constant: both its mean and its band change.

The node also poses a question the set point did not. Removing a coil from the rows is definite. Removing a unit's fan from the hall's air is a statement about the racks, and the rows do not settle it. It is settled by a choice, which the chapter names, and whose size it reports.

> **Principle 13.1.** *A scenario pins what the operator changes and must say what that change does to every quantity it touches. Where the record cannot say, the choice is stated beside the answer, because a band carries only the entries of the state.*

13.1.2 The shift plant with a unit stopped (step 0)

The plant, its rows and its reduced state of 19 entries are those of Chapter 12, with the reduced state of Chapter 4. The contingency here is unit A1 of hall A; $\mathcal{R}$ is the set of hall A's units still running, $\sum G^*$ and $\sum C^*$ their sums. At the node the dependent block is again triangular with a nonzero diagonal, one row fewer, so it is invertible, and the forward chain is

$$\begin{gathered} T^*_{rA} = T_s + \frac{Q_A + Q_o}{\sum_{u\in\mathcal{R}} G_u}, \qquad Q^*_u = G_u\,(T^*_{rA} - T_s), \qquad T^*_{d,u} = T^*_{rA} - \frac{Q^*_u}{C_u}, \\ T^*_{cA} = \frac{\sum_{u\in\mathcal{R}} C_u T^*_{d,u}}{\sum_{u\in\mathcal{R}} C_u}, \qquad T^*_k = T^*_{cA} + \beta_k\,(T^*_{rA} - T^*_{cA}), \end{gathered} \tag{13.1}$$

with $Q_o = 12\,\mathrm{kW}$ hall A's other heat. Hall B is unmoved: the halls share only T_s, and T_s does not change with a unit's run state.

The physics at the node. Three consequences follow from equation (13.1) before any number. The hall's heat leaves through fewer coils, so the return's rise above the water, $(Q_A+Q_o)/\sum G^*$, grows by the ratio of the two sums of conductances. The running units take the stopped unit's heat, and by the chain a discharge moves by $\Delta T^*_{d,u} = \Delta T^*_{rA}\,(1 - G_u/C_u)$, far less than the return when G_u is near C_u. On the plant's rule, the air balance of equation (5.5) holds on the running units,

$$\sum_{u\in\mathcal{R}} C_u\,(T^*_{rA} - T^*_{cA}) = Q_A + Q_o, \tag{13.2}$$

which fixes the gap between the return and the cold aisle by heat and air alone, and with it the cold aisle. And each inlet stands below the return by its share of that gap,

$$T^*_k = T^*_{rA} - (1 - \beta_k)\,\frac{Q_A + Q_o}{\sum C^*}. \tag{13.3}$$

On the plant's rule an inlet's response to a unit out is set by its recirculation fraction, as its band was at the present plant. With the rack air flow held, below, the fraction's lever does not change and every inlet moves by the return's rise.

The stop rule. This is the chapter's main choice, and it is the plant's own. A stopped unit contributes neither coil nor air: the sums in (13.1) run over the running units, $\sum C^*$ falls from 30 to 20 kW/K, and the fractions β_k are unchanged. This buys one consistent plant, the stated one, on which a do-it-and-see check of the truth can run. It fails in a named way: with a fraction fixed, the air flow the racks draw, $\sum C/(1-\beta)$, falls by a third with the supply, and server fans do not do that.

Three forms that hold the rack air flow. Three other descriptions of the same stop keep the racks' air. In form (b), the coil is valved off and the fan keeps running: the stopped unit's discharge equals T^*_{rA} and $\sum C^* = 30\,\text{kW/K}$. In form (c) the fan stops and the other two are raised to 15 kW/K each, so again $\sum C^* = 30\,\text{kW/K}$. In form (d) the fan stops with $\sum C^* = 20\,\text{kW/K}$ and the racks keep drawing what they drew, so more of their air must be recirculated: $\beta^*_k = 1 - (1-\beta_k)\,20/30$. On equation (13.3) forms (b) and (c) have $\sum C^* = 30$, and form (d) has $(1-\beta^*_k)/20 = (1-\beta_k)/30$, so all three reach one closed form,

$$T^*_k = T^*_{rA} - (1-\beta_k)\,\frac{Q_A + Q_o}{30}. \tag{13.4}$$

This is not a coincidence of the forms: rack air flow and heat fix each inlet's distance below the return, and the return is set by the coils alone. The chapter therefore has two rules, not four: *the plant as modeled*, equation (13.3) with $\sum C^* = 20$, which is reported, and *rack air flow held*, equation (13.4), set beside it in every answer. Neither is an entry of the state, and the record, whose units never stop, tests neither.

Modeling choices at the graph. The plant is taken in steady state within the window, with the airflow pinned and no active control, as in Section 5.1.2. Here the first says that the node is where the hall settles with the unit stopped; it says nothing of how fast, which Section 13.6.4 takes up. The last has a particular form at this node: the two running coils keep their log-derates. The entry g_u lumps fouling, water flow and valve position, so a valve that opens further is a larger g_u. Holding them keeps one state for the present and the node, and fails on any plant whose control opens the valves; the size of that reserve is not an entry of this state, and Section 13.4.3 varies it.

The graph. Figure 13.1 draws hall A at the node. The cold aisle on the left (blue) feeds the five racks, each drawn as its mixing row with the return air in the middle (blue); the return receives the load (green) and feeds the units on the right. Unit A1 is struck through (red). Its coil has left the rows, and its discharge row now reads the return air. Units A2 and A3 (green) pass the hall's heat to the chilled water (green, since T_s is an entry of the state) and send their discharge back along the bottom. Racks A1 and A3 are read; rack A5 (red) is unread and has the largest survey fraction. Every temperature here carries a star, and the node adds no reading: the readings of window 73 attach to the present plant of Chapter 12.

13.1.3 The operator's questions (step 1)

Which unit is "unit A". The drill's name does not identify one of A1, A2 and A3, and no unit is named on the schedule in front of the operator. The lead here is unit A1, the hall's weakest coil in the fit of Chapter 12, by a stated rule: it costs least to lose, so if it cannot come out, no unit of this hall can. The rule is a choice, a maintenance schedule may name another unit, and Section 13.3.3 answers for all six.

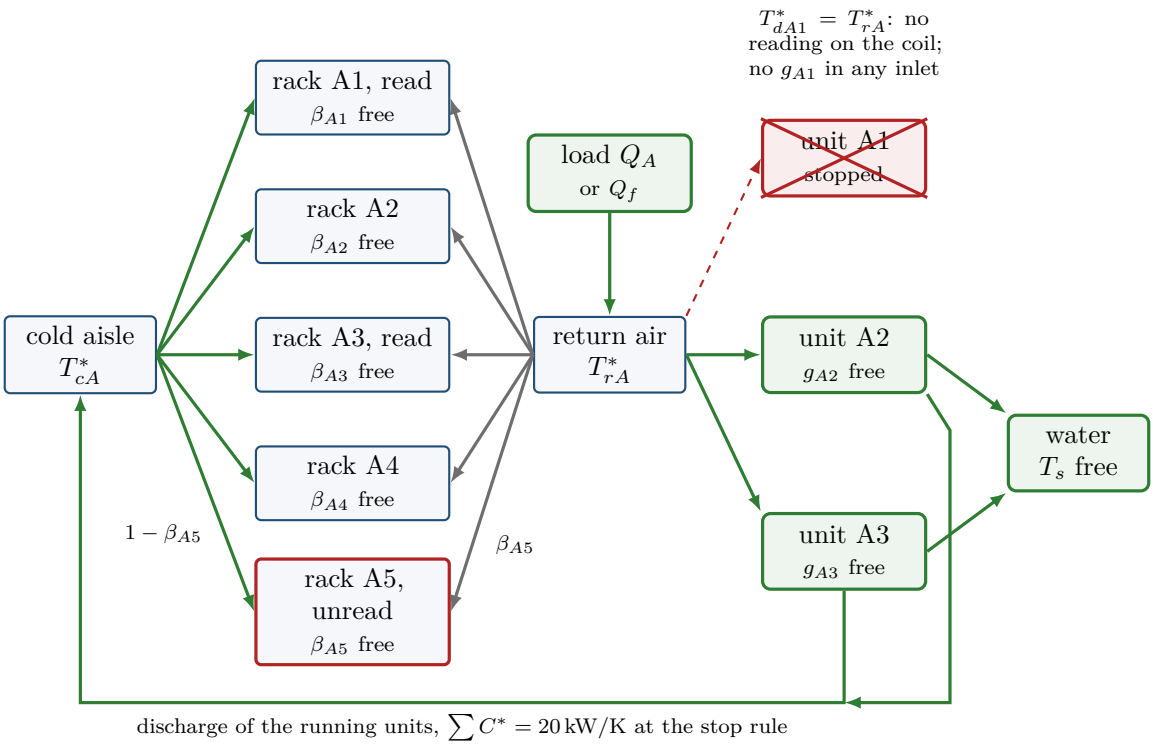


Figure 13.1: Hall A at the node with unit A1 stopped. Free (green): the load and the running coils' log-derates, with the recirculation fractions carried by the rack nodes and the supply T_s at the water. Blue: temperatures re-solved by equation (13.1). Red: the stopped unit, struck through, and the unread rack with the largest survey fraction. The dashed red arrow is the stopped unit's discharge row, which reads the return air. Rack air flow held replaces $\sum C^* = 20\,\mathrm{kW/K}$ by one of the three forms that reach equation (13.4). No reading attaches to the node.

The forecast peak. A rising line has no peak inside its span. On this record the forecast peak is therefore named as the line's value at the horizon, window 144 (06:55), the end of the shift, with windows 85, 97 and 121 beside it, because the maintenance length is not given and the table reads at each. The line says nothing about when the night's batch ramp ends.

The limit and the assurance. The 27°C limit is the plant's. The assurance κ of equation (2.12) belongs to the operator; the chapter computes at stated values, 0.90, 0.95 and 0.99, and names none as the rule.

The operator asks five questions, and step 1 writes each over the reduced state. Table 13.1 lists them, the drill question first. Its columns are the question in the operator's words, the general form in the notation of Chapter 2, the mathematical form on this plant, the category, the computation of Section 2.3.6, and the subsection that answers it. All five are predictions under a threshold, read at different nodes.

Table 13.1: The questions asked at window 73, 01:00. $\mathbf{y}$ is the window's readings and the priors; for the forecast, also the hall-A meter of windows 1 to 73. $\max_{k\in A}$ runs over hall A's five racks; $T_{\max} = 27°\mathrm{C}$.

Operator question	General form	Mathematical form	Category	Computation	Answered in
Can I take unit A down for maintenance now without a limit violation?	(2.12), a joint tail at a scenario node	$\Pr(\max_{k\in A} T_k^* < T_{\max} \mid \mathbf{y},\ \text{A1 out})$	prediction under a threshold	posterior	13.3.1
And at the forecast peak?	(2.12); (2.21), line 3	the same with $Q_A \to Q_f(t)$, $t = 144$	prediction under a threshold	trend	13.3.2
Does it matter which unit?	(2.12), per unit	the same, six nodes	prediction under a threshold	posterior	13.3.3
Where does the air go?	(2.12), two forms	the plant as modeled against rack air flow held	prediction under a threshold	posterior	13.3.4
How much load would have to leave first?	(2.12), inverted	the largest Q_A^* with $\Pr(\cdot) \geq \kappa$	prediction under a threshold	posterior	13.3.5

13.1.4 The evidence as rows (step 2)

The node adds no evidence. The evidence on the state is the 33 rows of window 73, Chapter 12, and a stopped unit's discharge reading does not exist yet. What step 2 fixes is how the node treats each quantity, and the rows of the one new fit.

What the node pins, leaves free and carries. Table 13.2 lists it: each row a quantity, the columns its treatment at the node and the reason.

Modeling choices at the node. The supply is carried at the window's entry, not at the set point, because the node reads window 73's posterior, the one fit of Chapter 12, and pins only what the operator changes. It fails to the extent the window's entry differs from what the chiller holds; the set point is set beside it in Section 13.4.3. A stopped unit's discharge row gives no evidence on its coil: by equation (13.1) it would read the return, so it would be a second return sensor.

Table 13.2: What the node of unit A1 pins, leaves free and carries.

Quantity	At this node	Why
Unit A1's run state	pinned, stopped	the action
T_s	free, window 73's entry	the window's posterior, the one fit; the run state does not touch the water
Q_A	free, the entry; replaced by Q_f at a forecast	the question's two loads
Q_B, hall B's g_u and β_k	free; hall B's temperatures unmoved	the halls share only T_s
g_{A2}, g_{A3}	free, unchanged	a choice: no valve opens further
β_k	free, unchanged at the plant's stop rule	a choice: Section 13.1.2
C_{A2}, C_{A3}	pinned, 10 kW/K	fixed fans at the plant's stop rule
Every temperature of hall A	re-solved	T^*_{rA}, T^*_d, T^*_{cA}, T^*_k

The meter alone, for the load. The forecast is fitted to the hall-A meter's readings of windows 1 to 73, $y_{Q_A,t}$ of standard deviation 4 kW, and to nothing else. This buys a slope that a constant bias of the meter cannot touch. It fails on a drifting meter, and it leaves the line's level carrying any constant bias. The window's fit and the record's reconciliation are in Chapter 12; the level's route is set beside the reported value in Section 13.3.2.

The line and the wander. The load is a line plus a persistent wander, $Q_{A,t} = \alpha + \gamma\,(t - \bar{t}) + d_t$, with $\bar{t} = 37$ and d_t first-order autoregressive with persistence ϕ and stationary spread s. The wander is the persistence Chapter 11 could not fit on twelve windows, and 73 can; without it the rows would treat the meter's windows as independent and overstate the slope's knowledge. A line fails if the ramp bends or ends, and nothing in the rows tests beyond window 73.

The priors. The slope has $\gamma \sim \mathcal{N}(0,\ 0.5^2)$ per window, one spread being 6 kW per hour; the level has a wide prior, $\alpha \sim \mathcal{N}(\bar{y},\ 100^2)$, about the mean reading. The persistence is taken by the marginal likelihood on a grid of ϕ from 0 to 0.95 and s from 0.5 to 8 kW, and then held at its best point, which ignores what is left unknown about it.

Which row acts on which entry. On the reading rows of hall A, with $*$ for an entry that can be nonzero and with the columns $T_s, Q_A, g_{A1}, g_{A2}, g_{A3}, \beta_{A1}, \ldots, \beta_{A5}$, the pattern of Chapter 12's Jacobian is

$$\mathbf{J}_A = \begin{bmatrix} * & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ * & * & * & * & * & 0 & 0 & 0 & 0 & 0 \\ * & * & * & * & * & 0 & 0 & 0 & 0 & 0 \\ * & * & * & * & * & 0 & 0 & 0 & 0 & 0 \\ * & * & * & * & * & 0 & 0 & 0 & 0 & 0 \\ 0 & * & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & * & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ * & * & * & * & * & * & 0 & 0 & 0 & 0 \\ * & * & * & * & * & 0 & 0 & * & 0 & 0 \end{bmatrix} \begin{matrix} y_{T_s} \\ y_{T_{rA}} \\ y_{T_{dA1}} \\ y_{T_{dA2}} \\ y_{T_{dA3}} \\ y_{Q_A} \\ y_{Q_{\text{loop}}} \\ y_{T_{A1}} \\ y_{T_{A3}} \end{matrix}, \tag{13.5}$$

$$\mathbf{g}^{\mathsf{T}}_{T^*_{A5}} = (*,\ *,\ 0,\ *,\ *,\ 0,\ 0,\ 0,\ 0,\ *),$$

the second line being the gradient of the answer's rack, not a row. Two facts are read here before any number. The columns of β_{A2}, β_{A4} and β_{A5} are empty in the reading rows: those

three fractions are not observable from the readings, and their prior rows are all that holds them. And the gradient has a zero in the column of g_{A1}: the column carries reading rows at the present plant, but the stopped unit's derate has left the answer, since its conductance appears in none of the sums of equation (13.1).

The load fit is a separate block with 75 entries $(\alpha, \gamma, d_1, \ldots, d_{73})$:

$$\mathbf{J}_{\text{load}} = \begin{bmatrix} \frac{1}{4}\,\mathbf{1} & \frac{1}{4}\,(\mathbf{t} - \bar{t}\,\mathbf{1}) & \frac{1}{4}\,\mathbf{I}_{73} \\ \mathbf{0} & \mathbf{0} & \mathbf{B} \\ 1/100 & 0 & \mathbf{0} \\ 0 & 1/0.5 & \mathbf{0} \end{bmatrix}, \tag{13.6}$$

with $\mathbf{B}$ the lower bidiagonal 73×73 matrix of the wander's rows. The α and γ columns meet every window, and the wander's columns each meet their own reading and their neighbors.

13.2 The posterior of the state (step 3)

The load trend is a second fit, and it shares one reading with the fit of window 73, that window's meter.

13.2.1 The fit of window 73 at the node

Chapter 12 derives the fit, and it is not refitted. Its entries for hall A, with their posterior spreads, are $T_s = 13.8947$ (0.0976), $Q_A = 363.2006$ (3.4343), $g_{A1} = -0.0659$ (0.0234), $g_{A2} = 0.0418$ (0.0215), $g_{A3} = -0.0437$ (0.0229), $\beta_{A1} = 0.1698$ (0.0254), $\beta_{A3} = 0.3715$ (0.0254), and the survey's $\beta_{A2} = 0.25$, $\beta_{A4} = 0.30$, $\beta_{A5} = 0.45$ with spread 0.08, uncorrelated with everything. The correlations that the answer's rack will use, on $(T_s, Q_A, g_{A2}, g_{A3})$, are

$$\begin{bmatrix} 1 & -0.00 & 0.28 & 0.26 \\ -0.00 & 1 & 0.15 & 0.17 \\ 0.28 & 0.15 & 1 & -0.10 \\ 0.26 & 0.17 & -0.10 & 1 \end{bmatrix}. \tag{13.7}$$

The fit is the plain fit of the plant's rows; Chapter 12 says why and states what the record shows of its events by 01:00. The fit carrying the hall-A meter's bias is the second route beside each headline.

The conductances are $G_{A1} = 7.4902$, $G_{A2} = 8.3418$ and $G_{A3} = 7.6582\,\text{kW/K}$, which is why unit A1 leads.

13.2.2 The load trend, a second fit

The state and the rows. The trend state of Section 2.3.4 and Chapter 11 is stacked here on the meter of hall A:

$$\mathbf{z}_{\text{load}} = (\alpha,\ \gamma,\ d_1,\ \ldots,\ d_{73})^{\mathsf{T}}, \qquad \mathbf{r}_{\text{load}} = \begin{bmatrix} (\alpha + \gamma\,(t - 37) + d_t - y_{Q_A,t})/4 \quad (t = 1, \ldots, 73) \\ d_1/s \\ (d_t - \phi\, d_{t-1})/(s\sqrt{1-\phi^2}) \quad (t = 2, \ldots, 73) \\ (\alpha - \bar{y})/100 \\ \gamma/0.5 \end{bmatrix}, \tag{13.8}$$

75 entries and 148 rows. The rows are linear in the state, so one solve of the normal equations of equation (4.29) is the fit, and $\mathbf{\Sigma}_{\text{load}} = (\mathbf{J}_{\text{load}}^{\mathsf{T}}\mathbf{J}_{\text{load}})^{-1}$ is the exact posterior covariance at the stated ϕ and s; no iteration table and no draws are needed.

The persistence. The marginal likelihood of the meter with the line integrated out peaks at an interior point of the grid, $\phi = 0.85$ and $s = 4.25\,\text{kW}$. Within two log units of the top lie ϕ from 0.55 to 0.95 and s from 2.50 to 8.00 kW, and those upper values are the ends of the grid: the record fixes that the wander persists, not how much. Section 13.4.3 gives the effect.

The line. The level and slope, with their spreads, are

$$\hat{\alpha} = 352.623\ (1.706)\,\text{kW}, \quad \hat{\gamma} = 0.2904\ (0.0715)\ \text{kW per window} = 3.485\ (0.858)\ \text{kW per hour}, \tag{13.9}$$

with $z = 4.064$ and $\Pr(\gamma > 0 \mid \mathbf{y}) = 0.999976$ by the first line of (2.21). With no wander the slope would read 0.2921 (0.0222) per window: the persistence widens the slope's spread by a factor of 3.22, because a wander that persists looks over a few windows like a slope. The log Bayes factor of this line against a level with its own best persistence, on the same priors, is $\ln B = 1.254$, which is not decisive. That alternative is a very persistent level, $\phi = 0.95$ and $s = 7.75\,\text{kW}$, with its ϕ held at the grid's end, so the grid bounds the comparison. The slope is read; a line beyond doubt is not (Section 13.6.4).

The forecast. The forecast load at a window t past the record is the third line of equation (2.21) with a later window's own wander added,

$$Q_f(t) = \hat{\alpha} + \hat{\gamma}\,(t - 37), \qquad \text{Var}(Q_f(t)) = \mathbf{g}(t)^{\mathsf{T}}\boldsymbol{\Sigma}_{\text{load}}\mathbf{g}(t) + s^2(1 - \phi^{2(t-73)}), \tag{13.10}$$

with $\mathbf{g}(t) = (1,\ t - 37)^{\mathsf{T}}$ on the level and slope, the form of equation (2.22). The wander's variance is conditional on window 73's wander, while the mean keeps the line alone and leaves that wander out: its fit is -1.289 (3.406) kW, of which $\rho = \phi^{t-73}$ remains at t. This is immaterial at windows 121 and 144, where ρ prints as 0.000. It is not at window 73 itself, where the same form gives the line with no wander term, so a row at that window is the line and not a predictive of the window. Table 13.3 gives each horizon: each row a window, the columns its load, its spread, the spread's two parts and ρ. At window 144 the spread's two parts are the line's 7.836 kW and the wander's 4.250 kW; the lever arm of (2.22) has taken the first past the second. The band at window 144 is 366.229 to 401.172 kW.

Table 13.3: The forecast load of hall A from the line, equation (13.10), fitted to the meter of windows 1 to 73.

Window t [–]	Clock [–]	Q_f [kW]	sd [kW]	Line sd [kW]	Wander sd [kW]	ρ [–]
85	02:00	366.565	5.690	3.832	4.207	0.142
97	03:00	370.050	6.274	4.615	4.249	0.020
121	05:00	377.021	7.551	6.242	4.250	0.000
144	06:55	383.701	8.914	7.836	4.250	0.000

Figure 13.2 draws it: the horizontal axis is the window, the vertical axis hall A's load. The solid line is the fitted line over the whole shift, with the forecast's 95 % band shaded from window 85. The filled points are the meter's hourly means inside the fit's span, and the open points are the later hours, which the fit did not use and Section 13.4.5 tests. The line runs through the filled points; the last open points lie below it and inside the band.

The extended state. An answer at a forecast uses $(\mathbf{z}_{73}, Q_f)$, 20 entries, with

$$\boldsymbol{\Sigma}_a = \begin{bmatrix} \boldsymbol{\Sigma} & \mathbf{0} \\ \mathbf{0} & \text{Var}(Q_f(t)) \end{bmatrix}. \tag{13.11}$$

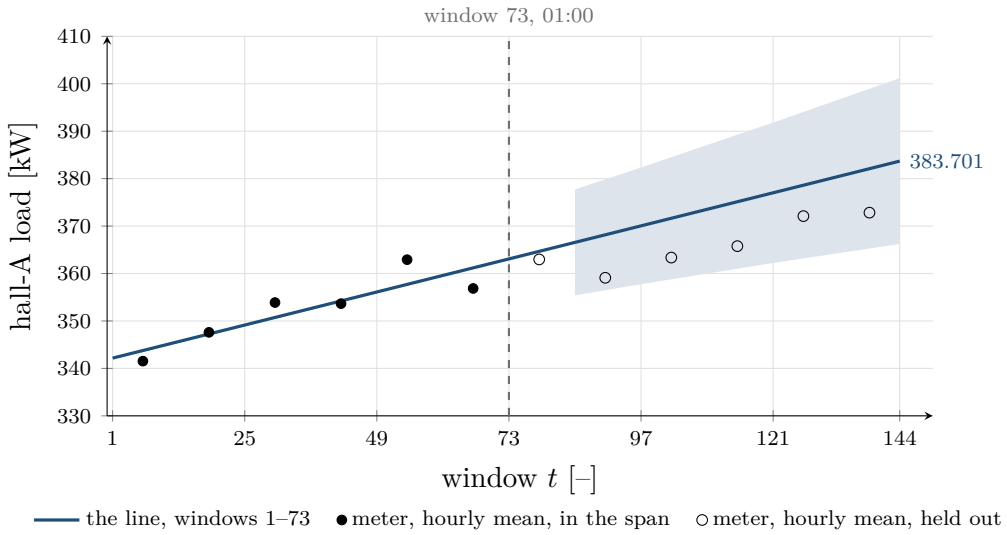


Figure 13.2: The load trend of hall A, Table 13.3. Solid: the line of equation (13.9). Shaded: the forecast's central 95 % band, drawn through the four horizons of the table. Points: the hall meter's mean over each hour; the hour starting at window 73 holds one window of the span.

Q_f is taken as independent of the window's entries. The window 73 meter reading is one of the 73 readings of the line, and a load in the future is not the entry Q_A; this is stated as a choice, and at the horizon the entry's own variance is much the smaller.

13.3 The answers (step 4)

Every answer is read from $(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of Section 13.2.1, or at a forecast from the state (13.11), and that value is the reported one, with the exact posterior beside it. The two air rules stand beside each other as two statements, because no route separates them.

13.3.1 The node at the current load

Operator question. Can I take unit A down for maintenance now without a limit violation? The question asks for the probability that all five inlets of hall A stay under the limit, once the hall has settled with unit A1 stopped.

General form. A prediction under a threshold, equation (2.12), with a joint event as in Chapter 8,

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

Specialized form. The question is equation (2.12) with $\mathbf{z}$ the reduced state of window 73, $f(\mathbf{z}) = -\max_{k \in A} T_k^*(\mathbf{z})$ of equation (13.1), $f^* = -T_{\max}$, and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal of the fit. By the delta method applied to five quantities at once, equation (8.9), the inlets at the node are jointly normal,

$$\mathbf{T}^* \mid \mathbf{y} \approx \mathcal{N}\big(\mathbf{T}^*(\hat{\mathbf{z}}),\ \mathbf{D}\boldsymbol{\Sigma}\mathbf{D}^{\mathsf{T}}\big), \qquad \Pr(\text{hall A}) = \Pr\big(\mathbf{T}^* < T_{\max}\,\mathbf{1} \mid \mathbf{y}\big), \tag{13.12}$$

with $\mathbf{D}$ the matrix of gradients and the probability from a numerical integration of the five-dimensional normal, seeded, with absolute error requested at 10^{-6}. A row of $\mathbf{D}$ differenti-

ates (13.3) and (13.1),

$$\frac{\partial T_k^*}{\partial T_s} = 1, \quad \frac{\partial T_k^*}{\partial Q_A} = \frac{1}{\sum G^*} - \frac{1-\beta_k}{\sum C^*}, \quad \frac{\partial T_k^*}{\partial g_u} = -\frac{G_u\,(T_{rA}^* - T_s)}{\sum G^*}, \quad \frac{\partial T_k^*}{\partial \beta_k} = \frac{Q_A + Q_o}{\sum C^*}, \tag{13.13}$$

for $u \in \mathcal{R}$, and zero in every other column, the column of g_{A1} among them. The last entry is the widened gap. The budget of the variance is equation (4.54) on this gradient. Its cross terms are all negative, by sign: the gradient has signs $(+,+,-,-)$ on $(T_s, Q_A, g_{A2}, g_{A3})$; T_s and Q_A correlate positively with each log-derate in (13.7), positive by negative by positive; the two log-derates correlate negatively on same-signed entries; and T_s with Q_A does not correlate to the print.

Building blocks. $\hat{\mathbf{z}}$, $\boldsymbol{\Sigma}$ and the correlations (13.7); the gradient (13.13); and the node's plant at the MAP state, as follows.

With it out, $\sum G$ falls from 23.4901 to 16.0000 kW/K. The hall's heat, 375.20 kW, then sets the return's rise above the water by equation (13.1): from 15.9727 to 23.4501 K, so the return goes from 29.8674 to 37.3447°C, up 7.48 K. Unit A2's heat goes from 133.24 to 195.62 kW and unit A3's from 122.32 to 179.59 kW, and their discharges rise only 1.24 and 1.75 K (16.5433 to 17.7832 and 17.6351 to 19.3862°C), by the factor $1 - G_u/C_u$. So the cold aisle moves from 17.3607 to 18.5847°C, and the gap $T_{rA} - T_{cA}$ of equation (13.2) widens from 12.51 to 18.76 K.

Figure 13.3 draws that: the horizontal axis has the two states, the present and the node, and the vertical axis temperature; each line is one quantity of the plant at the MAP state. The return climbs steeply, the cold aisle hardly at all, and the water not at all; the space between the top two lines is what each fraction multiplies.

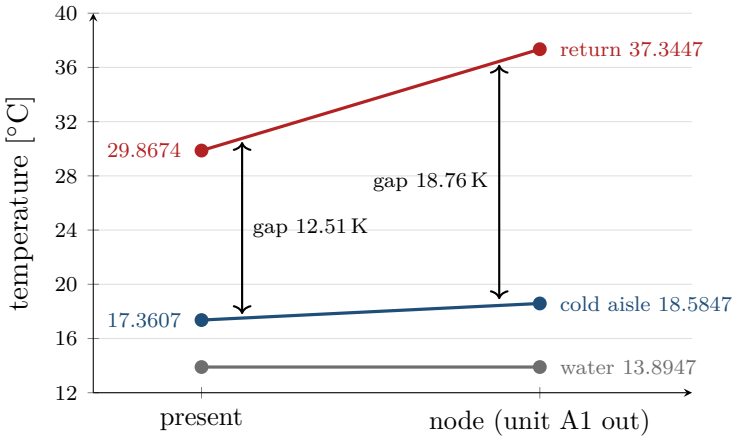


Figure 13.3: The return, the cold aisle and the chilled water at the present plant and at the node, at the MAP state of window 73, the plant's stop rule. The gap is the quantity of equation (13.2).

The gradient and the inlet. At rack A5, in K per unit of each entry,

$$\mathbf{g}_{T_{A5}^*}: \quad T_s\ 1.0000, \quad Q_A\ 0.0350\ \text{K/kW}, \quad g_{A2}\ -12.2259, \quad g_{A3}\ -11.2241, \quad \beta_{A5}\ 18.7600.$$

So $T_{A5}^* = 27.0267$°C with spread 1.5333 K and band 24.0214 to 30.0320°C, against an exact median of 27.0297°C. Its variance is, in K^2,

$$\underbrace{2.25241}_{\beta_{A5},\ 95.8\,\%} + \underbrace{0.06887}_{g_{A2}} + \underbrace{0.06630}_{g_{A3}} + \underbrace{0.01445}_{Q_A} + \underbrace{0.00952}_{T_s} \underbrace{- 0.06052}_{\text{cross}} = 2.35102.$$

The band has widened from the present's 1.0114 K to 1.5333 K, because the same fraction now multiplies 18.76 rather than 12.51 K. A band at a node is not the present's band moved.

The racks. Table 13.4 lists hall A's five inlets: each row a rack, the columns the present and the node with spread, and the probability under the limit at the node. Rack A1, with the smallest fraction, rises 2.29 K; rack A5, with the largest, rises 4.04 K.

Table 13.4: Hall A's inlets at the present and at the node, unit A1 stopped, the plant's stop rule, at window 73's load.

Rack	T_k [°C]	sd [K]	T_k^* [°C]	sd [K]	lo 2.5 % [°C]	hi 97.5 % [°C]	$\Pr(T_k^* < 27)$ [–]
A1, read	19.4840	0.2876	21.7696	0.4871	20.8149	22.7243	1.00000
A2	20.4874	1.0112	23.2747	1.5329	20.2703	26.2791	0.99246
A3, read	22.0065	0.2876	25.5534	0.4871	24.5986	26.5081	0.99851
A4	21.1127	1.0112	24.2127	1.5329	21.2082	27.2172	0.96549
A5	22.9887	1.0114	27.0267	1.5333	24.0214	30.0320	0.49305

The joint. By equation (13.12),

$$\begin{gathered}\Pr(\text{hall A} < 27°\text{C} \mid \mathbf{y},\ \text{A1 out}) = 0.47350;\\ \text{exact posterior } 0.47308 \text{ (se } 0.00079),\ \text{gap } -0.00042.\end{gathered} \tag{13.14}$$

At the present plant the same joint is 0.99996. The joint exceedance is 0.52650, against a largest marginal of 0.50695 and a union bound of 0.55050; the correlation of T_{A5}^* with T_{A3}^* is 0.0726, so the joint adds a little to rack A5's marginal, 0.49305, and rack A5 carries most of it. The other hall adds almost nothing at this load: over all ten inlets the joint is 0.47340. The fit carrying the hall-A bias gives $T_{A5}^* = 26.9905$°C (1.5202) and a joint of 0.48360. Rack air flow held gives a joint below 10^{-12} (Section 13.3.4).

A joint of 0.47 comes from a median at the limit with a wide band, and it is not a margin to weigh. At κ of 0.90, 0.95 or 0.99 the requirement of equation (2.12) is not accepted, on any route.

The screen after the action. On the plant as modeled at the MAP state, the screen would show the return at 37.345°C, unit A1's discharge at 37.345°C (up from 17.904°C), units A2 and A3 at 17.783 and 19.386°C, and the two read inlets at 21.770 and 25.553°C, both still under the limit. The unread rack that carries the answer would not be on it.

13.3.2 The node at the forecast peak

Operator question. And at the forecast peak? The question asks for the same joint at the load the shift is heading for.

General form. The threshold form (2.12), with the load replaced by the forecast of the third line of (2.21),

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

Specialized form. The question is equations (2.12) and (2.21) with $f_t = \max_{k \in A} T_k^*$ read at $(\mathbf{z}_{73}, Q_f(t))$, the entry Q_A of the chain (13.1) replaced by $Q_f(t)$ of (13.10), $f^* = T_{\max}$, and the posterior the normal of (13.11). The gradient is (13.13), with the column of Q_A moved to that of Q_f, and the joint is (13.12) with $\mathbf{\Sigma}_a$. Everything else is held at window 73.

Building blocks. The fit's entries; the forecast of Table 13.3; the gradient (13.13).

The forecast node. At the peak, $Q_f = 383.701\,\text{kW}$,

$$\begin{aligned}&\Pr(\text{hall A} < 27°\text{C} \mid \mathbf{y},\ \text{A1 out},\ Q_f(144)) = 0.27878;\\ &\text{exact posterior } 0.27794 \text{ (se } 0.00071),\ \text{gap } -0.00084.\end{aligned} \tag{13.15}$$

There rack A5 is 27.7442°C (1.6489), band 24.5123 to 30.9762°C, exact median 27.7439°C, with Pr = 0.32587; rack A3 26.1904°C (0.5941), 0.91351; rack A4 24.7765°C (1.6376), 0.91273; rack A2 23.7872°C (1.6344), 0.97533; rack A1 22.1999°C (0.5524), 1.00000. Table 13.5 reads the same node at each horizon, and Figure 13.4 draws its last two columns: the horizontal axis is the forecast's window, the vertical axis the probability, the upper line with no unit out and the lower with unit A1 out. The row at window 73 stands on the line's load, not the fit's entry, which is why its joint is 0.47451 and not the 0.47350 of (13.14).

Table 13.5: The node at the forecast load, equation (13.10), unit A1 out, hall A's joint; the last column has no unit out.

t [–]	Clock [–]	Q_f [kW]	T^*_{A5} [°C]	sd [K]	$\Pr(T^*_{A5} < 27)$ [–]	Pr(hall A) [–]	no unit out [–]
73	01:00	363.079	27.0225	1.5384	0.49417	0.47451	0.99996
85	02:00	366.565	27.1445	1.5617	0.46315	0.44033	0.99993
97	03:00	370.050	27.2665	1.5785	0.43298	0.40661	0.99989
121	05:00	377.021	27.5104	1.6136	0.37587	0.34029	0.99972
144	06:55	383.701	27.7442	1.6489	0.32587	0.27878	0.99937

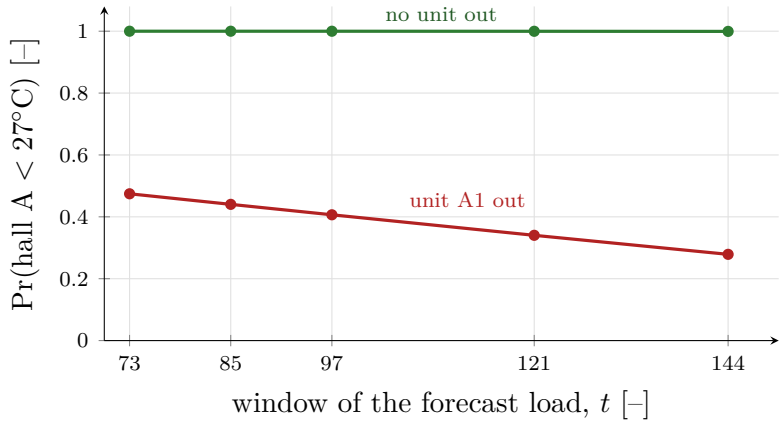


Figure 13.4: The joint of hall A against the forecast's window, Table 13.5. The stop costs far more than the load's rise.

The routes of the load. Four routes stand beside the reported load at window 144. The state of the fit carrying the bias, with the line less that bias, 4.594 (0.912) kW, gives 379.107 kW, $T^*_{A5} = 27.6699$°C and a joint of 0.29569. The fit's entry plus the slope over the span gives 383.822 kW, 27.7485°C and 0.27689. The supply pinned at 14.0°C gives 27.8495°C and 0.25048. Rack air flow held, on the meter's line, gives 31.3715°C and a joint that prints as 0.00000. None meets a stated assurance.

In kelvin the load is not the main lever. From the fit's entry to the forecast's mean is 20.50 kW, and that moves rack A5 by 0.72 K, against a band of ±3.23 K. In probability it is not small, since the median stands at the limit: the joint falls from 0.47350 to 0.27878. The forecast is for hall A's inlets only: hall B at the horizon is not asked here, since its coil and its load both move, and this chapter does not join them.

13.3.3 Each unit out

Operator question. Does it matter which unit? The rule of Section 13.1.3 chose a lead; a schedule may name another.

General form. The threshold form (2.12), once per unit.

Specialized form. The question is equation (2.12) with $f = -\max_k T_k^*$ over the racks of unit v's own hall, the chain (13.1) with $\mathcal{R}$ that hall's units other than v, and the same posterior: six nodes, each the joint (13.12) on its own five inlets.

Building blocks. The fit's entries for both halls, which Chapter 12 tables; the chain and gradient with $\mathcal{R}$ changed.

The six nodes. Table 13.6 lists them: each row a unit, the columns its conductance, what is left of the hall's, the return and cold aisle at the node, the hottest rack and its inlet, and the two joints.

Table 13.6: Each unit out at the current load, the plant's stop rule: joint over the unit's own hall and over all ten inlets.

Unit	G_u [kW/K]	$\sum G^*$ [kW/K]	T_r^* [°C]	T_c^* [°C]	Hottest	T^* [°C]	Pr(hall) [–]	Pr(all ten) [–]
A1	7.4902	16.0000	37.3447	18.5847	A5	27.0267	0.47350	0.47340
A2	8.3418	15.1484	38.6631	19.9030	A5	28.3451	0.10205	0.10203
A3	7.6582	15.8319	37.5937	18.8337	A5	27.2757	0.40285	0.40276
B1	7.2689	12.7579	37.8809	22.5802	B3	30.2306	0.00006	0.00006
B2	6.0585	13.9682	35.8024	20.5018	B3	28.1521	0.14913	0.14912
B3	6.6993	13.3274	36.8558	21.5552	B3	29.2055	0.01222	0.01222

No unit of either hall comes out at any stated assurance. With rack air flow held the joint is 0.00000 for all six. The answer is plain, and nothing has been chosen to soften it. The table holds one surprise: unit B2, the coil Chapter 12 reads as derated, is the best to lose in hall B. Losing a weak coil costs little of what the hall still has, and $\sum G^*$ is largest with it gone. That is the reason for the lead rule, seen from the other hall, and it is a statement about which stop costs least, not about which stop is acceptable.

13.3.4 The air at the node

Operator question. Where does the air go? The plant's rule loses the stopped fan's air from the racks as well; the other rule keeps it.

General form. The threshold form (2.12), on two descriptions of the same stop.

Specialized form. The question is equation (2.12) with $f = -\max_{k \in A} T_k^*$, once by equation (13.3) with $\sum C^* = 20\,\text{kW/K}$ and once by equation (13.4). The return is common to the two. On the second the gradient is (13.13) with $\sum C^* = 30$: the load's entry changes, and the fraction's entry is the smaller gap $(Q_A + Q_o)/30$. Between the two rules each inlet differs by $(1 - \beta_k)(Q_A + Q_o)/60$, the more at the rack with the less recirculation.

Building blocks. The return and the hall's heat of Section 13.3.1; the fractions.

The two rules. Table 13.7 sets the rules side by side: each row a rack, the columns the plant's rule, the three forms of Section 13.1.2, and spread and probability on the rule that holds the air. Forms (b), (c) and (d) coincide to the print at every rack, as equation (13.4) requires: at rack A1, 26.9613°C. Figure 13.5 draws the inlets: the horizontal axis lists the racks, the vertical axis is temperature; at each rack the three vertical bars are the 95 % bands at the present (gray), on the plant's rule (blue) and with rack air flow held (red), and the dotted marks are the constructed truth of Section 13.4.4 on the same two rules.

Table 13.7: The inlets of hall A at the node under the plant's stop rule and under the three forms that hold the rack air flow.

Rack	Plant's rule [°C]	(b) [°C]	(c) [°C]	(d) [°C]	sd, held [K]	$\Pr(T_k^* < 27)$, held [–]
A1	21.7696	26.9613	26.9613	26.9613	0.3772	0.54082
A2	23.2747	27.9647	27.9647	27.9647	1.0490	0.17887
A3	25.5534	29.4838	29.4838	29.4838	0.3774	0.00000
A4	24.2127	28.5901	28.5901	28.5901	1.0493	0.06483
A5	27.0267	30.4661	30.4661	30.4661	1.0502	0.00048

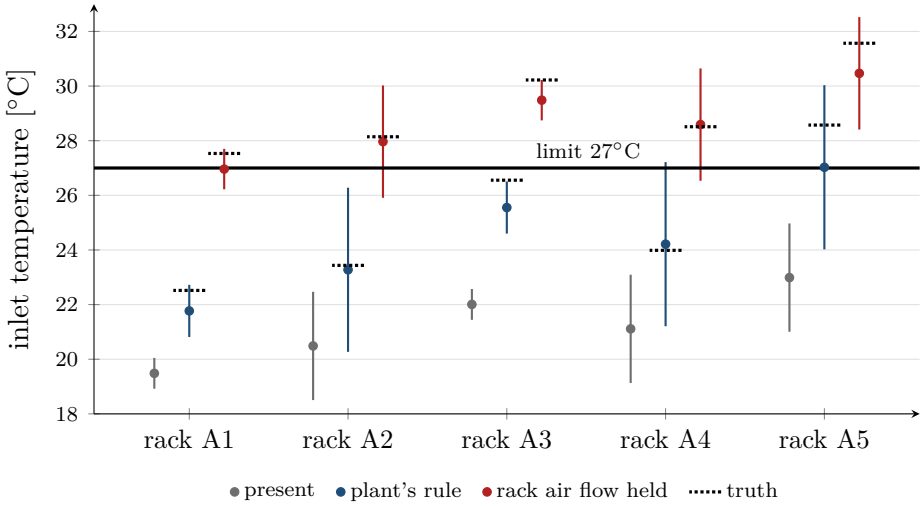


Figure 13.5: Hall A's inlets with their central 95 % bands: at the present (gray), at the node on the plant's stop rule (blue) and with rack air flow held (red). Dotted marks: the constructed truth under each stop rule, a check marker (Section 13.4.4). Both are computed on the truth's own state; the truth does not decide between the rules.

The two rules give joints of 0.47350 and below 10^{-12} (none of the 400 000 exact draws). The gap is not a route's. With the air held every inlet rises by the return's rise, 7.4774 K; the cold aisle goes from 17.3607 to 24.8381°C on forms (b) and (c), while form (d) keeps the plant's 18.5847°C and raises the fractions; and the read rack A3 is already over the limit, at 29.4838°C, with a marginal of 2.333×10^{-11}. Rack A5 differs by 3.44 K between them, because the air flow the racks are assumed to draw at the plant's rule has fallen from 43.368 to 28.912 kW/K at the hall's mean fraction, 0.3082. The threshold rule is met on neither, so the decision does not turn on the gap, and both are stated. The chapter does not choose between them.

13.3.5 The load at which the unit could come out

Operator question. How much load would have to leave first? The question inverts the joint: how far below the present must the hall's load stand before the joint reaches a stated

assurance.

General form. The threshold form (2.12), read for the level of a scenario quantity at which the rule $\Pr \geq \kappa$ first holds.

Specialized form. The question is equation (2.12) at a node with unit A1 out and $Q_A^* = Q_A - \Delta$, with Δ pinned and every other entry at window 73: the largest Q_A^* with

$$\Pr\big(\max_{k \in A} T_k^*(\mathbf{z}, \Delta) < T_{\max} \mid \mathbf{y}\big) \geq \kappa, \tag{13.16}$$

found by bisection on Δ, on each of the two air rules. The joint rises as Δ does, because the load's entry in the gradient (13.13) is positive at every rack.

Building blocks. The joint (13.12) at each Δ; the entry Q_A of the fit.

The load. Table 13.8 gives the inversion: each row an assurance, the columns the load to remove and the hall's load then, on each rule. The row at 0.50 is a reference where the joint reaches one half, not an assurance anyone is asked to hold. Figure 13.6 draws the same points: the horizontal axis is the load taken off the hall, the vertical axis the joint, one line per rule through the printed values, with the three stated assurances marked.

Table 13.8: The load at which unit A1 could come out, equation (13.16), with hall A's joint at κ.

κ [–]	Plant's rule, kW off [kW]	Plant's rule, hall load [kW]	Air held, kW off [kW]	Air held, hall load [kW]
0.50 (reference)	2.69	360.52	79.35	283.85
0.90	49.67	313.53	100.94	262.26
0.95	61.17	302.03	106.65	256.55
0.99	80.79	282.41	116.80	246.40

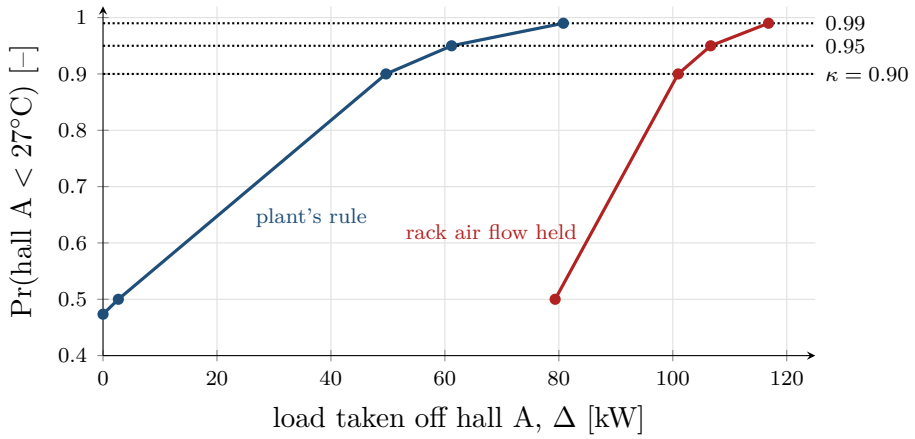


Figure 13.6: The joint of hall A against the load taken off the hall, unit A1 out, on each air rule (Table 13.8). The lines join the printed values. The joint with air held at no load removed lies below this axis.

The gap between 0.50 and 0.95 on the plant's rule, 58.48 kW, is what the band costs once the median is under the limit; it is mostly one survey's fraction. For comparison only, on the plant's rule a lower supply would restore the same joints at this load, with every temperature moving one for one: 0.094 K for 0.50 and 2.553 K for 0.95. Neither is named here as an action.

13.4 The checks (step 5)

Table 13.9 collects what can be set beside the answers: each row a check, the columns the reported value, what stands beside it and the gap. The stop rule is not among them, since the record never stops a unit.

Table 13.9: The checks of the answers. z is the reported value minus the truth, over the reported spread.

Check	Reported	Beside it	Gap
Joint, current load, exact (se 0.00079)	0.47350	0.47308	−0.00042
T^*_{A5} median, exact (se 0.0030) [°C]	27.0267	27.0297	+0.0030
T^*_{A5} band, lower end, exact (se 0.0065) [°C]	24.0214	24.0389	+0.0175
T^*_{A5} band, upper end, exact (se 0.0065) [°C]	30.0320	30.0296	−0.0024
Joint, forecast peak, exact (se 0.00071)	0.27878	0.27794	−0.00084
sd T^*_{A5}, Newton [K]	1.53331	1.53329	below the print
Joint, Newton	0.47350	0.47350	below the print
sd T^*_{A5}, 58-entry state, $\sigma_p = 0.1$ [K]	1.5333	1.5366	+0.0033
sd T^*_{A5}, 58-entry state, $\sigma_p = 0.01$, 0.001 [K]	1.5333	1.5333, 1.5333	to zero
Fit carrying the bias, joint	0.47350	0.48360	+0.01010
Plain fits of windows 70 to 76, joint, range	0.47350	0.28804 to 0.43773	one window's draw
Rack air flow held, joint	0.47350	below 10^{-12}	a choice, not a route
Truth (constructed), T^*_{A5} [°C]	27.0267	28.5720	$z = -1.01$
Truth, Q_A at window 144 [kW]	383.701	369.166	$z = 1.63$
Truth, slope [kW per window]	0.2904	0.2098	$z = 1.13$
Truth, persistence ϕ; s [kW]	0.85; 4.25	0.80; 3.00	inside the 2-log-unit range

13.4.1 The limit of exact rows and the curvature term

The full state with the node has 58 entries, one fewer than Chapter 12's, since the stopped unit has no discharge entry at the node. As the tolerances shrink through 0.1, 0.01 and 0.001, rack A5 reads 27.0266, 27.0267 and 27.0267°C, with spreads 1.5366, 1.5333 and 1.5333 K, and the return's spread settles at 0.3384 K: the reduced state's values, equation (4.48). The full Hessian against Gauss–Newton moves rack A5's spread from 1.53331 to 1.53329 K and rack A3's from 0.48711 to 0.48705 K, and the joint not at all to the print.

13.4.2 The delta method against the exact posterior

The exact posterior is drawn by importance sampling with the Laplace normal as the proposal: 400 000 draws, effective size 396 188. The gaps are in the table above, and at each rack's marginal they stay at or below 0.00067. On the joint the gap is about half a standard error, and the median's gap at rack A5 equals its standard error, so the sign of each is unresolved. The band's lower end differs by +0.0175 K with a standard error of 0.0065 K, which the draws do resolve, and its upper end by less than its standard error. At the forecast peak the gap of −0.00084 is about 1.2 standard errors, its sign unresolved too, and no gap crosses a stated assurance.

13.4.3 Sensitivity

The running coils. Raising both running coils' g by 0.05, 0.10 and 0.20, as valves opening further would, takes the joint to 0.76305, 0.92510 and 0.99724. So the whole answer could

change on a plant that holds that reserve, and the reserve's size is not an entry of this state. The record has no such event to read it from.

The survey's spread and the supply. A survey spread of 0.04 or 0.12 on the unread fractions gives 0.48634 or 0.42022. With the median at the limit, a narrower band gains little. The node carries the supply at window 73's posterior entry, 13.8947°C. Pinning it at the set point, 14.0°C, shifts every inlet by +0.1053 K, which is the cause of the change, to $T^{*}_{A5} = 27.1320$°C and a joint of 0.44357; it also raises rack A5's spread slightly, from 1.5333 to 1.5392 K, since the readings pin the inlets and not their distance from the supply.

The persistence. Table 13.10 refits the line at the far corners of the range the record allows: each row a persistence, the columns the slope, its normalized value, the forecast at window 144 and the joint at the peak. The load's band varies threefold in width, from 24.116 to 72.628 kW, and the slope's normalized value falls to 1.836 at the widest; the joint barely moves, from 0.27552 to 0.30495, because the coils and the fraction carry its band.

Table 13.10: The load trend at three persistences, unit A1 out, at window 144.

ϕ [–]	s [kW]	γ (sd) [kW per window]	z [–]	Q_f (sd) [kW]	Band [kW]	Pr(hall A) [–]
0.85	4.25	0.2904 (0.0715)	4.064	383.701 (8.914)	366.229 to 401.172	0.27878
0.95	8.00	0.2747 (0.1496)	1.836	381.497 (18.528)	345.183 to 417.811	0.30495
0.80	3.00	0.2921 (0.0491)	5.948	383.972 (6.152)	371.914 to 396.030	0.27552

One window's draw. The plain fits of the neighboring windows give, on the same node:

Window [–]	70	71	72	74	75	76
Joint [–]	0.28804	0.41625	0.35028	0.43773	0.30686	0.31547

Each of them is below the reported value, so window 73's is a generous draw among them. The verdict at every stated κ does not move.

13.4.4 The answers against the constructed truth

The record is generated from a stated truth, on which the stop can be carried out without noise. The check runs on the plant's stop rule, because the air rule belongs to the plant and the truth cannot judge itself. On it, with unit A1 out at window 73, rack A5 reads 28.5720°C, over the limit, and every other rack is under: A3 26.5544, A4 23.9865, A2 23.4363, A1 22.5192°C. The return is 37.5595°C against the fit's 37.3447°C. The errors in reported spreads are $z = -1.54$, -0.11, -2.05, $+0.15$ and -1.01 for racks A1 to A5; rack A3's 1.0010 K is 0.5234 K of fraction (fitted 0.3715), 0.2628 K of over-read load through the gap and 0.2147 K of return. With rack air flow held, on the same truth, all five racks are over: 27.5326, 28.1440, 30.2227, 28.5108 and 31.5678°C.

Every unit out, on the truth: A1 28.5720, A2 29.4403, A3 28.2355, B1 31.1964, B2 29.4531, B3 30.3286°C. On the truth it is unit A3, not unit A1, that costs least, since Chapter 12's check has the fit's g_{A3} above the truth by 1.58 spreads: the lead rule's choice of A1 is itself one window's estimate. At window 144, with A1 out and the true load 369.166 kW, rack A5 is 29.1982°C and rack A3 27.1018°C, both over.

13.4.5 The held-out load

The meter's windows 74 to 144, which the line did not see, give $\chi^2 = 71.38$ on 71 against the predictive covariance, tail 0.465. The last twelve windows' mean is 372.83 kW against the line's 382.10 kW. Against the true load the line stands 4.128 kW above it over the span, which is the bias in the line, and 13.775 kW above it over the last twelve windows. The forecast runs high by more than its bias because the slope is high, 0.2904 against the truth's $30/143 = 0.2098$ per window, and the band holds it.

13.5 The report (step 6)

The screen says 5 K in hand on the read inlets. On the plant as modeled, the probability that hall A stays under the limit with unit A1 stopped is 0.47, with rack A5, which no sensor reads, at 27.03°C. On a plant whose racks keep their air flow it is below 10^{-12}. No unit of either hall comes out at any stated assurance on either rule. At the shift-end forecast the probability falls to 0.28, on the plant as modeled. To reach assurance 0.95 at this load, 61 kW would have to leave the hall, or 107 kW if the air is held.

13.5.1 The answers and their bands

Table 13.11 gives the answers in the order of Table 13.1: each row a question, the second column the reported answer, the third what stands beside it.

Table 13.11: The report at window 73, 01:00. Limit 27°C; assurance levels the operator's; unit A1 the lead by the stated rule.

Question	Answer, reported	Beside it
Can I take unit A down for maintenance now without a limit violation?	plant as modeled: 0.47350; rack A5 27.03°C, band 24.02 to 30.03°C; not met at 0.90, 0.95 or 0.99	exact 0.47308; fit carrying the bias 0.48360; rack air flow held below 10^{-12}; both stated
And at the forecast peak?	0.27878 at 383.701 kW, band 366.229 to 401.172 kW; not met	exact 0.27794; fit carrying the bias with the line less it 0.29569; rack air flow held prints as 0.00000
Does it matter which unit?	yes; A1 0.47350, A3 0.40285, A2 0.10205, B2 0.14913, B3 0.01222, B1 0.00006; none met	rack air flow held: 0.00000 for all six
Where does the air go?	a stated choice: 3.44 K at rack A5 between the rules; three forms of holding it reach one closed form	rack A5 30.4661°C with air held, 27.0267°C at the plant's rule
How much load would have to leave first?	49.67, 61.17, 80.79 kW at κ 0.90, 0.95, 0.99	air held: 100.94, 106.65, 116.80 kW

13.5.2 The entry in the shift log

The operator's entry reads: "01:00, window 73. What-if: unit A1 out (hall A's weakest coil; other units tabled). Return 29.9 to 37.3°C; rack A5, unread, 27.0°C, band 24.0 to 30.0. Hall A under 27°C: probability 0.47 at this load (exact 0.47), 0.28 at the shift-end forecast of 384 kW (band 366 to 401; exact 0.28). If the racks keep their air flow rather than losing a third with the stopped fan: below 10^{-12} at this load, prints as 0.00000 at the forecast. Not met at assurance 0.95 on any route; would need 61 kW off the hall at this load, 107 kW if air is held. Steady state only; no credit taken for valves opening further. Action: not decided here."

Common mistake

Reading two of three units as a spare, or the read inlets' margin as the hall's. The margin was measured with three units running, at two racks of five. On the plant's stop rule the stop widens the gap each fraction multiplies, and the rack that carries the answer has no sensor. With the rack air flow held, every inlet rises by the return's 7.48 K, and the read rack A3 is already over the limit.

13.6 The next reading (step 7)

The report leaves a median near the limit, a band that is nearly one survey fraction, and a choice about the air. A reading is one more row, and what it does to the entries is known before its value.

13.6.1 The preposterior of a candidate

A candidate adds rows at the present plant, and the covariance after it is equation (4.65), carried to the node by the same $\mathbf{D}$. Probes at racks read the present inlet: at rack A5 the row acts mainly on β_{A5}. A re-survey adds a prior row on β_{A5}. The permanent sensor at rack A5 is the instrument at the place of the question, and it adds the probe's row in six windows. A clamp-on meter at the hall's load, or a water-side heat check on units A2 and A3, its spread 3 kW since the larger plant's units carry a wider stated spread than the two-unit hall's 2 kW, acts on the entries that carry the small part of the budget. The comparison across candidates is the preposterior probability of ending past an assurance, at 0.50 and 0.95, over 2000 outcomes of the reading, with the joint integrated at each outcome and a standard error of at most 0.011, on each air rule.

13.6.2 The candidates

Table 13.12 compares them: each row a candidate, the columns the spreads of β_{A5} and of T^{*}_{A5} after it, the joint at the present means with the candidate's covariance, and the chance of ending at 0.50 or at 0.95 or above, first on the plant's rule and then with rack air flow held.

Table 13.12: Candidate readings at the node: unit A1 out, current load, hall A's joint, on the plant's stop rule and with rack air flow held.

			Plant's rule			Air held		
Candidate	sd β_{A5} [–]	sd T^{*}_{A5} [K]	Pr at means [–]	Pr(≥ 0.50) [–]	Pr(≥ 0.95) [–]	sd T^{*}_{A5} [K]	Pr(≥ 0.50) [–]	Pr(≥ 0.95) [–]
None, window 73	0.0800	1.5333	0.47350	0.0000	0.0000	1.0502	0.0000	0.0000
Handheld probe at rack A5, sd 0.3 K, now	0.0254	0.4871	0.46063	0.4765	0.2080	0.3773	0.0000	0.0000
Handheld probes at A5, A4, A2, sd 0.3 K	0.0253	0.4871	0.47780	0.4795	0.2555	0.3770	0.0000	0.0000
Re-survey of β_{A5}, sd 0.02	0.0194	0.4807	0.46258	0.4775	0.2065	0.4010	0.0000	0.0000
Permanent inlet sensor at rack A5, six windows, sd 0.3 K	0.0151	0.2908	0.44822	0.4780	0.3175	0.2701	0.0000	0.0000
Clamp-on on hall A's load, sd 1 kW	0.0800	1.5329	0.47350	0.0025	0.0000	1.0482	0.0000	0.0000
Water-side heat check on A2 and A3, sd 3 kW each	0.0800	1.5310	0.47362	0.1295	0.0000	1.0479	0.0000	0.0000

On the plant's rule, narrowing the band does not move the median off the limit, and with the probe the joint at present means falls a little: at a median near 27°C a sharper band is not reassurance. The chance of ending past 0.95 is the fair measure on that rule, and it runs from 0.2065 for the re-survey to 0.3175 for the permanent sensor. The meter and the water-side heat check leave it at zero, because they do not reach the fraction. With the rack air flow held, every candidate's chance at either assurance is 0.0000: the permanent sensor's 0.3175 holds on the plant's rule only.

No listed candidate reads the choice that matters most. The air rule is not an entry of this state, and no reading at the present plant, with every unit running, reveals what the racks draw with a unit out. That is the limit of step 7 here. As a fact of the two rules, once the unit is stopped the read inlets differ between them by 5.19 K at rack A1 and 3.93 K at rack A3, against spreads of 0.38 to 0.49 K. The choice among readings, and any action, stay with the operator. Chapter 18 ranks readings on this plant against a named worry.

13.6.3 What the analysis showed

On the plant's stop rule a unit out moves the return far and the cold aisle little, so each rack answers by its fraction, and an unread rack with the largest fraction carries the hall. With the rack air flow held, every inlet rises by the return's 7.48 K, the cold aisle on forms (b) and (c) from 17.36 to 24.84°C, and the read rack A3 is already over the limit, at 29.48°C. The dashboard's 5 K measured a different plant at other racks. A stop rule in a model is a physical claim, and here it is worth 3.44 K at the rack that matters. Three ways of keeping the air give one closed form, because rack air flow and heat fix each inlet's distance below the return. The forecast peak is a small lever in kelvin beside these two, 0.72 K at rack A5, though with the median at the limit the joint falls from 0.47 to 0.28. And the rule that picked the lead unit, weakest coil first, also explains the surprise in the other hall: the derated coil is the cheapest to lose.

13.6.4 Limits of the model

Steady state. The node is where the hall settles with the unit stopped. It says nothing of how fast: not how long the return takes to climb 7.48 K, nor whether a maintenance shorter than the settling would stay under the limit. Maintenance that short would have to be argued from a dynamic model, and this plant has none. It also does not say what a unit's restart does.

No active control. Valves and fans hold their positions. A plant that opens its valves has a reserve this state does not hold, and Section 13.4.3 shows how large its effect could be.

The fractions and the air. At the plant's rule β_k is fixed. Neither air rule is read from the record, so the two answers stand as stated.

The line. The line is not decisive against a very persistent level, and it says nothing about the end of the ramp; the horizon lies 71 windows past a span of 73. The persistence is held at its best point, which is interior to the grid, though the 2-log-unit range touches the grid's ends. The forecast's wander variance is conditional on window 73's wander while its mean leaves that wander out.

Q_f independent of the window. Equation (13.11) takes the two fits as independent, though they share one reading.

The bias, in the line as read. The reported forecast carries the meter's constant bias in its level, and the carried fit's route, with the line less its bias, stands beside it.

Hall B at the horizon. It is not computed here, and the chapter joins no forecast across the halls.

One window's draw. The reported value is one window's; its neighbors are all below it.

Constructed readings. The record is constructed: generated from a stated truth, which Section 13.4.4 compares with the answers on the plant's own rule. It is not field data.

13.7 What to retain

Definition 13.1 (Unit-out scenario node)**.** A unit-out scenario node pins a unit's run state stopped. The state and $\mathbf{\Sigma}$ are unchanged and the temperatures are re-solved over the running units: the return's rise is the heat over the running conductances, the gap to the cold aisle is the heat over the running capacity rates, and an inlet stands below the return by $(1-\beta_k)$ of that gap. The stopped unit's log-derate leaves the gradient. At the plant's stop rule the gain of each fraction becomes the widened gap, so the band widens, and the air flow the racks imply falls with the supply; with rack air flow held, every inlet rises by the return's rise, and $T_k^* = T_r^* - (1-\beta_k)(Q+Q_o)/30$, however it is held. The rule is a stated choice beside the answer, not part of its band. A forecast load replaces the load's entry as one more entry of the state.

Three conclusions carry beyond the hall. A scenario can make an entry's gain larger, as the plant's stop rule does, so a survey that was tolerable at the present plant can decide the answer at the node. A choice that no entry holds must be stated beside the number, since no reading and no band can carry it. And a sharper band around a median at the limit reassures no one; on the plant's rule the fair measure of a reading is the chance of ending past the assurance, and with the rack air flow held no candidate moves that chance off zero.

Table 13.13 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 13.13: The terms of the unit-out node and where each is defined. Units for hall A of the shift plant.

Term	Expression	Units	Meaning	Defined in
Node's chain	eq. (13.1)	°C, kW	temperatures over the running units	13.1.2
Gap	eq. (13.2)	K	return to cold aisle: heat over air	13.1.2
Inlet at the node	eq. (13.3)	°C	distance below the return	13.1.2
Rack air flow held	eq. (13.4)	°C	three forms, one closed form	13.1.2
Sparsity	eq. (13.5)	–	no g_{A1} in the gradient	13.1.4
Posterior				
Load trend	eqs. (13.8), (13.9)	kW	a second fit, line and wander	13.2.2
Forecast load	eq. (13.10)	kW	line and later wander	13.2.2
Extended state	eq. (13.11)	mixed	window's entries and Q_f	13.2.2
Questions and answers				
Gradient at the node	eq. (13.13)	K, K/kW	$\mathbf{\Sigma}$ carried to the inlets	13.3.1
Joint at the node	eq. (13.12)	–	every inlet under the limit	13.3.1
Load to remove	eq. (13.16)	kW	the inverted threshold	13.3.5
Candidate's measure	Pr(after $\geq \kappa$)	–	chance of ending past the assurance	13.6.1

13.8 Exercises

Exercise 13.1. *Application.* From equation (13.1), compute the return at the node from T_s, the hall's heat and the running conductances. ($T^*_{rA} = 13.8947 + 375.20/16.0000 = 37.3447$°C.)

Exercise 13.2. *Derivation.* Derive the closed form (13.4) for form (d) from equation (13.3), and evaluate it at rack A1. (With $\sum C^* = 20$, $1-\beta^*_k = (1-\beta_k)\,20/30$, so $(1-\beta^*_k)(Q+Q_o)/20 = (1-\beta_k)(Q+Q_o)/30$. With $\beta_{A1} = 0.1698$ and the hall's heat 375.20 kW, on the unrounded entries: $T^*_{A1} = 26.9613$°C.)

Exercise 13.3. *Derivation.* Find the share of β_{A5} in the budget of T^*_{A5}. (Its gradient entry is the gap, 18.7600 K, and $\Sigma_{\beta_{A5}\beta_{A5}} = 0.08^2$, so the term is 2.25241 K^2 of the total 2.35102 K^2: 95.8 %.)

Exercise 13.4. *Application.* Assemble the variance of Q_f at window 144 from equation (13.10). (The line contributes a spread of 7.836 kW and the wander 4.250 kW, with $\rho = 0.000$ at this distance; in quadrature the spread is 8.914 kW.)

Exercise 13.5. *Derivation.* Explain why g_{A1} has left the gradient of T^*_{A5} although reading rows see it at the present plant. (Its conductance appears in none of the sums over $\mathcal{R}$ in equation (13.1), so its column in $\mathbf{D}$ is zero at the node, and its posterior spread 0.0234 contributes nothing: the printed gradient has entries for T_s, Q_A, g_{A2}, g_{A3} and β_{A5} only, and the variance they give is 2.35102 K^2.)

Chapter 14

Load Transfer Between Halls

14.1 Forty kilowatts between two halls

It is 01:00, window 73. A batch job can be rescheduled, and the operator is considering moving 40 kW of it from hall A to hall B. The screen shows hall A's load meter at 363.61 kW and hall B's at 297.25 kW, a difference of 66.37 kW, and the returns at 29.97 and 29.15°C. Moved, the meters would read 323.61 and 337.25 kW, a difference of 13.63 kW the other way, and the heat on the loop would not change. Hall A is the heavier and the warmer on the meters and the returns. Hall B's hottest read inlet, rack B4 at 22.67°C, is 4.33 K inside the limit, and hall A's, rack A3 at 22.03°C, is 4.97 K inside. The meters and returns favor the move; the four read inlets lean, by 0.64 K, the other way, and six racks are unread. The dashboard's reading is that the move levels the meters and gives margin where it is short.

The meters read what they read. What they do not say is what happens to the margin, the distance of each inlet below the limit, and six of the ten racks have no sensor. The question of this chapter is

where the margin goes when load moves from one hall to the other: how much hall A gains, how much hall B loses, at which racks, and with what probability both halls stay inside.

It is answered from the one fit of Chapter 12, read at two scenario nodes. The reported route adds no fit and no reading.

Jargon: load transfer (load migration), margin, donor and receiving hall, assurance, joint tail

A *load transfer*, or *load migration*, moves computing load, and so heat, from one hall to another on the same chilled-water loop. The *margin* of a rack is the limit less its inlet temperature, here $27°\mathrm{C} - T_k$. The *donor* hall gives up the load and the *receiving* hall takes it. The *assurance* κ is the probability an operator requires before calling a requirement met; it is the operator's choice, and this chapter computes at 0.90, 0.95 and 0.99 and names none of them the rule. A *joint tail* is the probability of an event about several quantities at once, as in Chapter 8.

Key idea

A transfer is a what-if on two scenario nodes, one per hall, and it adds nothing to the fit. Each inlet moves by its own slope in its hall's heat. That slope has a closed form, and it differs between the halls and between the racks, so a kilowatt moved is not margin moved one for one, and the numbers on the meters do not say which way the margin goes.

14.1.1 A transfer and the slope of an inlet

Chapter 12 changed a quantity that touches every temperature alike. A transfer does not. It takes heat out of one hall and puts it into the other, and each hall's temperatures answer through that hall's own conductance and each rack through its own recirculation fraction. In steady state the rise of a hall's return above the water is its heat over its conductance, and the cold aisle lies below the return by its heat over the air capacity rate. An inlet is a mix of those two, weighted by β_k. Each inlet therefore has a slope in its hall's heat, and the two halls, and the racks within them, do not share it.

Principle 14.1. *The margin a transfer moves is set by each rack's slope in its own hall's heat. The slope is larger where the hall's conductance is smaller and where the rack's recirculation fraction is larger.*

14.1.2 The plant and the two nodes (step 0)

The plant is the two-hall shift plant of Chapter 12, with its rows, its reduced state of 19 entries and its scenario node, all cited and none derived again: on the reduced state of Chapter 4, the dependent block $\mathbf{A}_{\text{dep}}$ is triangular with nonzero diagonal, and so invertible. The modeling choices of the graph are that chapter's: steady state within the window, pinned airflow and no active control in the window. They apply unchanged.

The transfer is a what-if (contingency) analysis on two scenario nodes: hall A at $Q_A^* = Q_A - \Delta$ and hall B at $Q_B^* = Q_B + \Delta$, with $\Delta = 40\,\text{kW}$. Table 14.1 lists what the nodes pin, leave free and carry: each row is a quantity or a group, the second column its status at the nodes and the third the reason. The two nodes together are one copy of the plant's 20 temperatures, ten per hall, so that the full state with the nodes has $59 = 19 + 20 + 20$ entries: the reduced entries, the present temperatures and the temperatures at the nodes.

Table 14.1: What the two nodes pin, leave free and carry.

Quantity [unit]	At these nodes	Why
transfer Δ [kW]	pinned, 40 kW	the action; a difference, so a meter's constant bias does not enter it
Q_A, Q_B [kW]	free, window 73's entries, moved by $\mp\Delta$	the sum is unchanged, so the loop's heat is
T_s [°C]	free, window 73's entry, one value on both nodes	the chiller sees the same heat
g_u [–], including g_{B2}	free, unchanged	the derate at its level at 01:00
β_k [–]	free, unchanged	a choice, below
C_u, $Q_{o,h}$ [kW/K, kW]	pinned	fixed fans, other heat
temperatures of both halls [°C]	re-solved	$T_{r,h}^*$, $T_{d,u}^*$, $T_{c,h}^*$, T_k^*

The physics at the nodes. The forward chain of Chapter 12 gives, for a rack k of hall h, with $\sum G_h = \sum_{u \in h} \bar{G}_u \mathrm{e}^{g_u}$, $\sum C_h = 30\,\text{kW/K}$ and the other heat $Q_{o,A} = 12\,\text{kW}$, $Q_{o,B} = 10\,\text{kW}$,

$$T_{r,h}^* = T_s + \frac{Q_h^* + Q_{o,h}}{\sum G_h}, \qquad T_{c,h}^* = T_{r,h}^* - \frac{Q_h^* + Q_{o,h}}{\sum C_h}, \qquad T_k^* = T_{c,h}^* + \beta_k \left(T_{r,h}^* - T_{c,h}^*\right). \quad (14.1)$$

Differentiating the three in order, an inlet moves with its hall's heat at the rate

$$\frac{\partial T_k}{\partial Q_h} = \frac{1}{\sum G_h} - \frac{1-\beta_k}{\sum C_h}. \tag{14.2}$$

The slope does not depend on the load, so each inlet is a line in Δ. It holds two asymmetries the screen does not show. Hall B has the smaller conductance, rated at 21 kW/K against hall A's 24 kW/K before any derate. And the slope grows with β_k, so the rack of the receiving hall with the most recirculation moves most.

The graph. Figure 14.1 draws the nodes on the plant's rows. The pinned transfer at the top (gray) moves the two load entries (green) in opposite directions. Each hall's return and cold aisle (blue) are re-solved, and in each hall the rack with its hall's largest survey fraction, an inlet with no sensor, is drawn in red, fed by the cold aisle and the return with its fraction's weights. The units (green) carry their free log-derates to the chilled water (green), one free entry shared by both halls. The only path between the halls in the figure is that entry and the pinned Δ; a further tie, through the fit's covariance, is not an arrow. No reading is attached to any node: the dashed arrow brings in the fit.

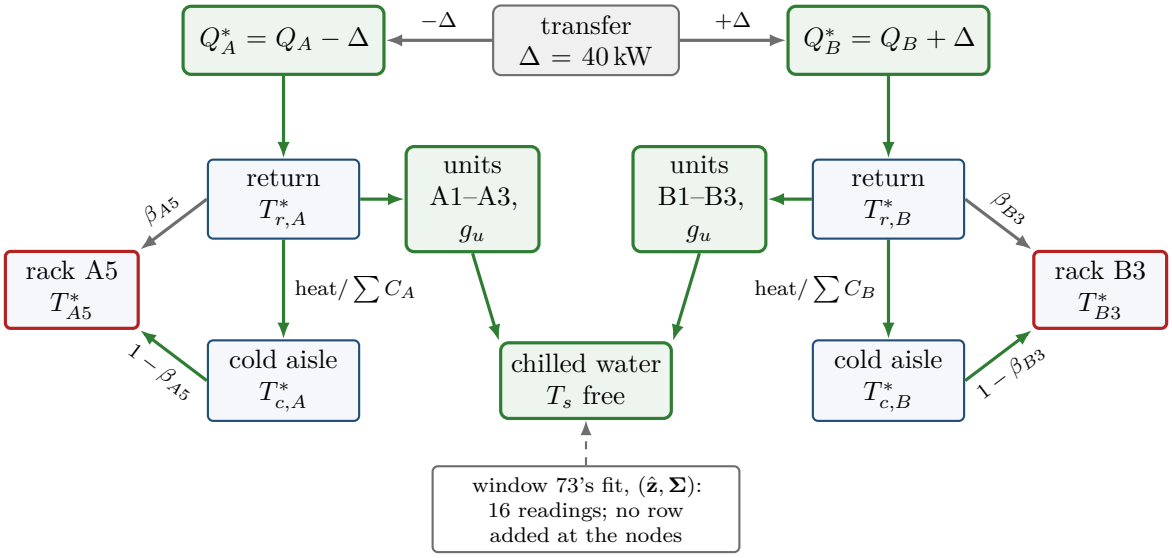


Figure 14.1: The two scenario nodes of the transfer at window 73. Gray: the pinned transfer. Green: the free entries the nodes read, the two loads, the units' log-derates and the supply. Blue: temperatures re-solved at the nodes, by equation (14.1). Red: the rack with its hall's largest survey fraction, read by no sensor; the other eight inlets are drawn the same way and left out. Dashed: the one fit, which supplies the entries and their covariance; the nodes add no reading.

Modeling choices at the nodes. Five choices enter here. The *fractions β_k do not depend on the load*: this buys the line (14.2), and it fails if a hotter or denser rack recirculates more; its size is tested in Section 14.4.2. The *transfer Δ is exact*: this buys the pin, and it fails if the job's heat is not known to the kilowatt; tested in the same section. The *supply is carried at its entry in the fit, not at the set point*: this buys the same supply as the one fit, and it fails to the extent that the window's entry is one window's reading of a controlled value; the pinned value is set beside it in Section 14.3.3. The other two have no size here and are limits (Section 14.6.4). The *hall is one heat node, so moved load is a hall's heat and not a rack's*: this buys two nodes and one number, and it fails if the 40 kW lands on a rack whose inlet is read

through its fraction, since the model has no row from a rack's load to its own inlet. The *coils and loads are held at 01:00*: the move is read on the plant at that window, and what the record shows of their motion is read in Section 14.2.3.

14.1.3 The operator's questions (step 1)

The operator asks four questions of window 73, and Table 14.2 lists them, the first three being the drill question taken in its parts. The columns give the question in the operator's words; the general form of Chapter 2; its mathematical form on this plant, with $\mathbf{y}$ the sixteen readings and the priors of the one fit; the category; the computation of Section 2.3.6; and the subsection that answers it. The limit is the plant's, $T_{\max} = 27°\mathrm{C}$.

Table 14.2: The questions asked of the transfer at window 73. T_k^* by equation (14.1); the maximum runs over the ten racks of both halls.

Operator question	General form	Mathematical form	Category	Computation	Answered in
If I move 40 kW from hall A to hall B, where does the margin go?	eq. (2.9)	$27 - T_k^*$ per rack, against the present	estimation	posterior	14.3.1
What is the band on each node?	eqs. (2.9), (2.11)	T_{A5}^* and T_{B3}^*, the rack with its hall's largest survey fraction	estimation	posterior	14.3.2
What is the probability both stay inside?	eq. (2.12), a joint tail	$\Pr(\max_k T_k^* < 27 \mid \mathbf{y},\ \Delta = 40)$	prediction under a threshold	posterior	14.3.3
What of another size, or the other direction?	eq. (2.12) along Δ	$\Pr(\cdot \mid \Delta)$; the largest Δ with $\Pr \geq \kappa$	prediction under a threshold	posterior	14.3.4

14.1.4 The evidence as rows (step 2)

The evidence is the one fit's: its sixteen reading rows and its priors, on the state $\mathbf{z}_{73}$ of window 73 (01:00), as Chapter 12 attaches them. The nodes add no row. Its sparsity is cited from that chapter, and what it says here is short. In the reading rows each hall's temperatures see T_s, that hall's load and that hall's coils; the meters see the loads; a measured inlet adds its own fraction. The columns of β_{A2}, β_{A4}, β_{A5}, β_{B1}, β_{B3} and β_{B5} hold no reading row, so those six fractions are not observable from the readings. With $*$ for an entry that can be nonzero and $\mathbf{0}$ for a block of zeros, the gradients at the rack with each hall's largest survey fraction have the pattern

$$\begin{array}{lccccccc} & T_s & Q_A & Q_B & g_{A1},\, g_{A2},\, g_{A3} & g_{B1},\, g_{B2},\, g_{B3} & \beta_{A1}, \ldots, \beta_{A5} & \beta_{B1}, \ldots, \beta_{B5} \\ \mathbf{g}_{T_{A5}^*}^{\mathsf{T}} = & * & * & 0 & * \ * \ * & \mathbf{0} & 0\ 0\ 0\ 0\ * & \mathbf{0} \\ \mathbf{g}_{T_{B3}^*}^{\mathsf{T}} = & * & 0 & * & \mathbf{0} & * \ * \ * & \mathbf{0} & 0\ 0\ *\ 0\ 0 \end{array} \tag{14.3}$$

The two gradients share one column, T_s. Each has its fraction's entry in a column with no reading row, so the rack's band will carry that prior. And Q_A and Q_B are two columns here, tied only through $\mathbf{\Sigma}$, by the loop meter's row on their sum.

14.2 The posterior of the state (step 3)

The posterior is that of the one fit, $\mathbf{z}_{73} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$, with the rows, the normal equations, the iteration table and $\mathbf{\Sigma}$ of Chapter 12. It is not refitted, and the treatment of the record's events is cited from that chapter.

14.2.1 The entries the nodes read

The nodes read these entries and spreads, in K, kW and dimensionless units: $T_s = 13.8947$ (0.0976); $Q_A = 363.2006$ (3.4343); $Q_B = 296.0123$ (3.4302); $g_{A1} = -0.0659$ (0.0234), $g_{A2} = 0.0418$ (0.0215), $g_{A3} = -0.0437$ (0.0229); $g_{B1} = 0.0377$ (0.0249), $g_{B2} = -0.1444$ (0.0290), $g_{B3} = -0.0439$ (0.0267); and $\beta_{A5} = 0.4500$, $\beta_{B3} = 0.5000$, each at its survey value with the survey's 0.0800, uncorrelated with every other entry. At those entries the conductances are

$$\sum G_A = 23.4901\,\mathrm{kW/K}, \qquad \sum G_B = 20.0268\,\mathrm{kW/K},$$

against ratings of 24 and 21, and the heat on the loop is 681.21 kW, the same at the present and at any Δ.

The block of $\boldsymbol{\Sigma}$ this chapter reads. One entry of $\boldsymbol{\Sigma}$ was not at work in Chapter 12's node: the correlation of Q_A and Q_B, -0.1782. It comes through the loop meter, which takes the sum's standard deviation from 5.27 to 4.40 kW; a sum better fixed than its parts' independent spreads, 4.85 kW, trades one against the other; here, unlike Chapter 12's node, it stands beside a shared free T_s. The other correlations the budgets below need are that chapter's: T_s with each coil of hall B, $+0.22$ to $+0.25$; Q_B with those coils, $+0.19$ to $+0.23$; and those coils with each other, -0.15 to -0.18.

14.2.2 The other fits the answers use

Two kinds of second route are read beside the reported fit, and each is defined here so that the answers can use it. The *fit carrying the hall-A bias* is Chapter 12's second named route, and the transfer, a difference, is applied to its entries unchanged. The *plain fits of the neighboring windows* have the same rows on those windows' own readings. Windows 61 to 73, the record up to the question, are read this way in the answers, each at its own entries, and none replaces the fit of window 73; later windows enter only as a check.

14.2.3 What the record reads of the motion

Chapters 12 and 13 read the events and the load's line, and the treatment is cited. Read on windows 1 to 73 only, with an ordinary least-squares line through each meter's readings and no prior on its slope, the record shows this. The hall-B meter climbs at 0.1981 kW per window, 2.377 kW per hour, with standard deviation 0.0298 per window once widened by its residuals' lag-one correlation of 0.172; the hall-A meter climbs at 0.2927 kW per window (0.0445), and a constant bias leaves a slope unchanged. The windows' own values of g_{B2} average -0.0026 over windows 1 to 36 and -0.1079 over windows 62 to 73, a slope over windows 37 to 73 of -0.00333 per window (0.00044).

14.3 The answers (step 4)

Every answer is read from $(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of Section 14.2.1, and that value is the one reported. Beside each stands the exact posterior, by importance sampling from the Laplace normal through the exact rows (400 000 draws, effective size 396 137), like with like; the fit carrying the bias stands beside each center, ratio, joint and largest transfer.

14.3.1 Where the margin goes

Operator question. If I move 40 kW from hall A to hall B, where does the margin go? The question is asked of every rack, read or not, and against the present.

General form. An estimation question about a derived quantity asks for its posterior, 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 at its center.

Specialized form. The question is equation (2.9) with $\mathbf{z} = \mathbf{z}_{73}$, $f(\mathbf{z}) = 27 - T_k^*(\mathbf{z})$ of equation (14.1) for each of the ten racks, and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal of Section 14.2.1, read at the same $\hat{\mathbf{z}}$ at $\Delta = 0$ and at $\Delta = 40\,\mathrm{kW}$. By the slope (14.2) the change of margin is $+\Delta$ times the slope in hall A and $-\Delta$ times the slope in hall B.

Building blocks. The MAP state $\hat{\mathbf{z}}$, the conductances of Section 14.2.1, the slope (14.2) and the forward chain (14.1).

The margins at 01:00. Table 14.3 gives the answer: each row is a rack; the columns give the margin at the present and at the nodes, the change, the slope, which the closed form (14.2) reproduces to the print at all ten, and the fraction. Figure 14.2 draws the same two margins as paired bars, with each one's 95 % band as a whisker and the constructed truth at the nodes as a dotted marker, a check for Section 14.4.3; a margin below the zero line is an inlet over the limit.

Table 14.3: The margin of each rack at the present and at the two nodes, $\Delta = 40\,\mathrm{kW}$ from hall A to hall B.

Rack [–]	Margin, present [K]	Margin, nodes [K]	Change [K]	$\partial T_k/\partial Q_h$ [K/kW]	β_k [–]
A1	7.5160	8.1119	+0.5959	0.01490	0.1698
A2	6.5126	7.2155	+0.7028	0.01757	0.2500
A3	4.9935	5.8583	+0.8648	0.02162	0.3715
A4	5.8873	6.6568	+0.7695	0.01924	0.3000
A5	4.0113	4.9808	+0.9695	0.02424	0.4500
B1	5.9855	5.0548	−0.9307	0.02327	0.2000
B2	4.8815	3.8065	−1.0750	0.02687	0.3082
B3	2.9253	1.5947	−1.3307	0.03327	0.5000
B4	4.2874	3.1348	−1.1526	0.02882	0.3665
B5	4.4554	3.3247	−1.1307	0.02827	0.3500

At the means, the margin goes from the rack with the least to a rack that had margin to spare. Rack A5 gains 0.9695 K, band 0.7569 to 1.1821 K (exact median 0.9692, band 0.7563 to 1.1831), and rack B3 loses 1.3307 K, band 1.1158 to 1.5455 K (exact median 1.3312, band 1.1161 to 1.5462). The ratio of the two is a ratio, so it is reported by its median 1.373 and its inversion band 1.054 to 1.829; the delta form gives standard deviation 0.190, band 1.000 to 1.745, and the exact posterior median 1.373, band 1.054 to 1.830. The probability that the ratio exceeds one is 0.9907 (exact 0.9907, se 0.0002). On the fit carrying the bias, rack A5 goes from 22.9656 to 21.9898°C and rack B3 from 24.0798 to 25.4085°C, a ratio of 1.362.

The smallest margin at the means falls from 2.9253 K to 1.5947 K, at rack B3 both times. At the means, the heavier hall was not the plant's constraint at this window, and that is a statement with a probability. At the present, rack B3 is the hottest of the ten with probability 0.733 (exact 0.735), the hottest rack is in hall B with probability 0.812 (exact 0.813), and rack A5 is hotter than rack B3 with probability 0.203 (exact 0.202). At the nodes the first two are 0.875 (exact 0.876) and 0.998 (exact 0.998).

The excess of the loss over the gain, 0.3612 K, has four terms by the slope (14.2): the rated conductances, 0.2381 K; the other coils at the fit, −0.0333 K; the derate of unit B2, 0.0897 K;

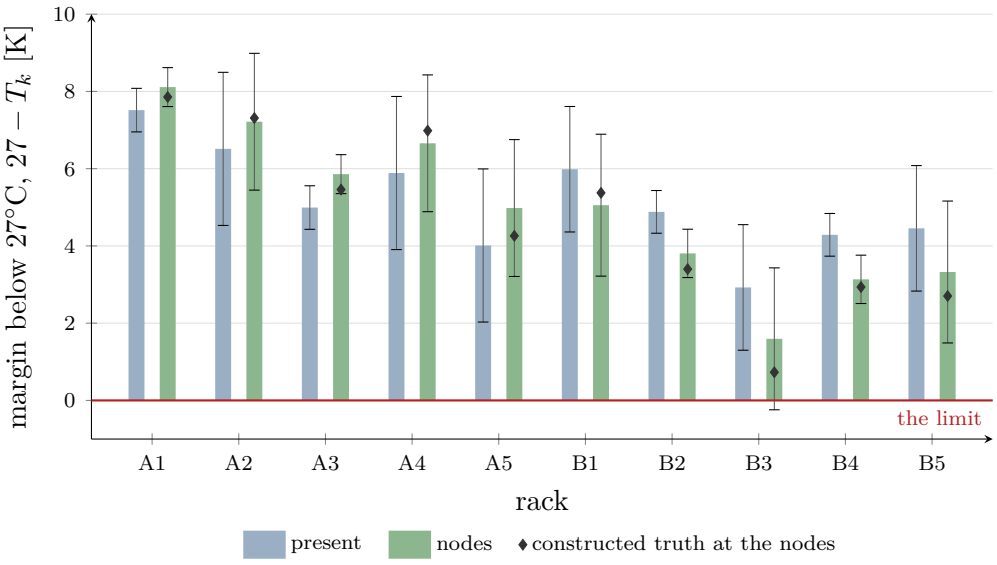


Figure 14.2: The margin of each rack at the present and at the two nodes (Table 14.3). Whiskers: the central 95 % band of each. Diamonds: the constructed truth at the nodes, a check marker (Section 14.4.3). The racks with short whiskers, A1, A3, B2 and B4, are the four with inlet sensors.

and the two fractions, 0.0667 K. With g_{B2} at its rating, rack B3 would rise 1.2410 K and the ratio would be 1.280. The derate is about a quarter of the excess, the rated conductance most of it, and the fractions less than the derate.

14.3.2 The band on each node

Operator question. What is the band on each node? The hottest inlet of each node is asked for, with what stands behind its band.

General form. The posterior of equation (2.9), summarized by its central 95 % credible interval, equation (2.11), and its variance split by entry.

Specialized form. The question is equations (2.9) and (2.11) with $f = T^*_{A5}$ and $f = T^*_{B3}$ of equation (14.1), and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal, by the delta method of equation (4.53). In the columns of the pattern (14.3), with $h = B$ and the same form for $h = A$,

$$\begin{aligned} \frac{\partial T^*_{B3}}{\partial T_s} &= 1, & \frac{\partial T^*_{B3}}{\partial Q_B} &= \frac{1}{\sum G_B} - \frac{1-\beta_{B3}}{\sum C_B}, \\ \frac{\partial T^*_{B3}}{\partial g_u} &= -\frac{G_u\,(T^*_{r,B} - T_s)}{\sum G_B}, & \frac{\partial T^*_{B3}}{\partial \beta_{B3}} &= T^*_{r,B} - T^*_{c,B}, \end{aligned} \tag{14.4}$$

and the band is $T^*_{B3}(\hat{\mathbf{z}}) \pm 1.960\sqrt{\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\,\mathbf{g}}$, equation (4.56). The budget splits $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\,\mathbf{g}$ by entry, equation (4.54). The fraction's entry is uncorrelated with the others, so its term stands alone. Every cross term with a coil is negative: T_s and Q_B have gradients of positive sign and the coils of negative sign, and their correlations with the coils are positive, which gives a negative product; two coils have gradients of the same sign and a negative correlation, which gives a negative product too.

Building blocks. $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$ of Section 14.2.1; the gradient (14.4); the temperatures of equation (14.1).

The two bands. By equation (14.1), hall A's return goes from 29.8674 to 28.1645°C and its cold aisle from 17.3607 to 16.9912°C; hall B's return from 29.1749 to 31.1722°C and its cold aisle from 18.9745 to 19.6385°C. The gradients of equation (14.4) at $\hat{\mathbf{z}}$ are, in the order T_s, the hall's load, its three coils and the rack's fraction,

$$\begin{aligned} \mathbf{g}_{T^*_{A5}} &: \ 1,\ 0.0242,\ -4.5501,\ -5.0675,\ -4.6522,\ 11.1734, \\ \mathbf{g}_{T^*_{B3}} &: \ 1,\ 0.0333,\ -6.2710,\ -5.2268,\ -5.7797,\ 11.5337, \end{aligned}$$

and so

$$\begin{aligned} T^*_{A5} \mid \mathbf{y} &\approx \mathcal{N}(22.0192^\circ\mathrm{C},\ 0.9039^2), && [20.2475,\ 23.7909]^\circ\mathrm{C}, \\ T^*_{B3} \mid \mathbf{y} &\approx \mathcal{N}(25.4053^\circ\mathrm{C},\ 0.9374^2), && [23.5681,\ 27.2426]^\circ\mathrm{C}. \end{aligned}$$

The exact medians are 22.0151 and 25.4050°C, and the exact bands [20.2435, 23.7916] and [23.5657, 27.2425]°C. The centers on the fit carrying the bias are those of Section 14.3.1. Each hall's hottest read rack is held much tighter: rack A3 at 21.1417°C (0.2573) and rack B4 at 23.8652°C (0.3194). Both hottest inlets at the nodes are racks no sensor reads. The upper end of rack B3's band lies over the limit.

The budget of the band. Figure 14.3 draws the budget of T^*_{B3}, $0.87865\,\mathrm{K}^2$ in all: one bar per entry and one for the cross terms, the horizontal axis in percent of the variance. The survey's fraction for rack B3 is $0.85137\,\mathrm{K}^2$, or 96.9 %. The three coils give 0.02445, 0.02305 and $0.02376\,\mathrm{K}^2$, the load 0.01302 and the supply $0.00952\,\mathrm{K}^2$, and the cross terms take back $0.06652\,\mathrm{K}^2$. Rack A5's band is the same kind of band: β_{A5} carries 97.8 % of it.

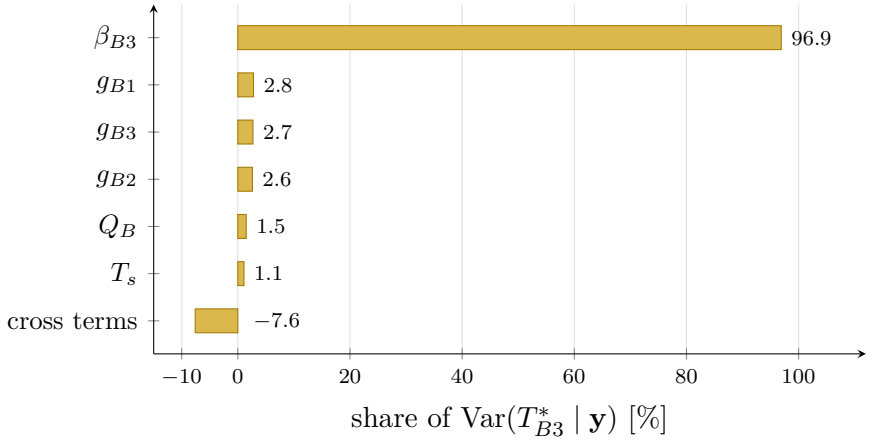


Figure 14.3: The variance budget of T^*_{B3} by equation (14.4): the diagonal terms and the cross terms, in percent of $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\,\mathbf{g} = 0.87865\,\mathrm{K}^2$.

The bands move with the transfer. The receiving hall's widens, 0.8293 to 0.9374 K, and the donor's narrows, 1.0114 to 0.9039 K, because each fraction multiplies its hall's gap $T^*_r - T^*_c$, and by equation (14.1) that gap is the hall's heat over $\sum C_h$. The move adds heat where the band is already the survey's, and it stretches that band.

14.3.3 Both halls inside

Operator question. What is the probability both stay inside? All ten inlets are asked for at once, on both nodes.

General form. A prediction under a threshold is a tail of a derived quantity, equation (2.12),

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

here a joint tail of the form of equation (8.6) in Chapter 8.

Specialized form. The question is equation (2.12) with $f = -\max_k T_k^*$ over the ten racks, $f^* = -T_{\max}$, and the two nodes read together. By the joint form of equation (8.9), the matrix $\mathbf{D}$ with rows $\mathbf{g}_{T_k^*}^{\mathsf{T}}$ carries the posterior into the ten inlets,

$$\begin{aligned} (T_{A1}^*, \ldots, T_{B5}^*)^{\mathsf{T}} \mid \mathbf{y} &\approx \mathcal{N}(\mathbf{m}^*,\ \mathbf{D}\boldsymbol{\Sigma}\mathbf{D}^{\mathsf{T}}), \\ \Pr\big(\max_k T_k^* < T_{\max} \mid \mathbf{y}, \Delta\big) &= \int_{\{\mathbf{x} < T_{\max}\}} \mathcal{N}(\mathbf{x};\ \mathbf{m}^*,\ \mathbf{D}\boldsymbol{\Sigma}\mathbf{D}^{\mathsf{T}})\, \mathrm{d}\mathbf{x}, \end{aligned} \tag{14.5}$$

integrated numerically as in Chapter 12. The off-diagonal entries between the halls come only through $\boldsymbol{\Sigma}$ and the shared column of T_s, by the pattern (14.3): the entry for a pair of racks is $\mathbf{g}_{T_{A5}^*}^{\mathsf{T}} \boldsymbol{\Sigma}\, \mathbf{g}_{T_{B3}^*}$, whose terms pair each of T_s, Q_A and hall A's coils with each of T_s, Q_B and hall B's coils.

Building blocks. The means $\mathbf{m}^*$ by equation (14.1) and the matrix $\mathbf{D}$ of equation (14.5), each row in the form of equation (14.4); $\boldsymbol{\Sigma}$; the limit.

The joint at the nodes. By equation (14.5),

$$\Pr\big(\max_k T_k^* < 27^{\circ}\mathrm{C} \mid \mathbf{y},\ \Delta = 40\big) = 0.95537,$$

with the exact posterior at 0.95548 (se 0.00033), a gap of +0.00011. By hall, Pr(hall A) = 1.00000 and Pr(hall B) = 0.95537 (exact 1.00000 and 0.95548). At the present the three are 0.99996, 0.99979 and 0.99975.

The joint reduces to one rack. The exceedance is 0.04463, the union bound 0.04464 and the largest marginal 0.04445, at rack B3. So the joint over two nodes reduces, on this plant at this size, to one rack's tail, and it does so because hall A's tail is out of reach at this size, not because of the correlation. The correlation of the two nodes' hottest inlets is −0.0003, and independence of the halls gives 0.95537, a shift of -1.32×10^{-9}. Their covariance, $-0.00023\,\mathrm{K}^2$, is a near-cancellation, in K^2: T_s with itself +0.00952, with Q_B −0.00005 and with hall B's coils −0.01059; Q_A with T_s −0.00002, with Q_B −0.00169 and with hall B's coils +0.00145; hall A's coils with T_s −0.00840, with Q_B +0.00113 and with hall B's coils +0.00843. At $\Delta = -10\,\mathrm{kW}$ the correlation is as small, −0.0002, and the reduction fails: of the exceedance 0.00017, hall A's tail is 0.00014 and hall B's 0.00002, and the largest marginal is 0.00014, at rack A5. The two nodes are what let that be found.

The routes across 0.95**.** The fit, the exact posterior and the fit carrying the bias all lie above 0.95, and that does not settle the question at that assurance. Table 14.4 gives each route with its two hottest inlets and its joint on the same rule, 40 kW and all ten racks: each row is a route, and the last column says on which side of 0.95 it falls. The centers of the exact posterior and of the fit carrying the bias are those already given.

Suppose the operator's assurance is 0.95. The reported value and the exact value stand above it, by 0.005. Pinned at the set point, the supply puts it at 0.94246, below. The two sides differ by a choice of the supply's value that the window cannot settle: the entry is the fit's

Table 14.4: The joint at $\Delta = 40\,\mathrm{kW}$ on each route. Plain fits are defined in Section 14.2.2.

Route	T^*_{A5} [°C]	T^*_{B3} [°C]	Pr(both) [–]	Against 0.95 [–]
Plain fit of window 73, reported	22.0192	25.4053	0.95537	above
Exact posterior	–	–	0.95548	above
Fit carrying the bias	–	–	0.95478	above
Supply at the set point, pinned at 14.0°C	22.1245	25.5106	0.94246	below
Plain fit of window 70 (00:45)	22.2216	25.6095	0.93207	below
Plain fit of window 71 (00:50)	22.0939	25.4932	0.94620	below
Plain fit of window 72 (00:55)	22.0906	25.4541	0.94711	below
Plain fits of windows 61–73, median	–	–	0.9501	above
Plain fits of windows 61–73, least and greatest	–	–	0.9305, 0.9782	both sides

reading of a controlled quantity and the pin is what the controller is told. Over windows 61 to 73, the record up to the question, the median is 0.9501, the range 0.9305 to 0.9782, and seven of the thirteen windows reach 0.95. At $\kappa = 0.95$ the verdict is therefore not settled by this record. Both values are stated, 0.955 at the window's entry and 0.942 at the set point, and the chapter does not choose between them. At $\kappa = 0.90$ all routes at window 73 pass, and at $\kappa = 0.99$ none does.

14.3.4 The transfer as a size and a direction

Operator question. What of another size, or the other direction? The question asks for the comparison, and the chapter answers with one; it does not name a move.

General form. The tail of equation (2.12), read along the pinned quantity Δ and inverted at a stated κ:

$$\Delta_{\max}(\kappa) = \max\{\Delta : \Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}, \Delta) \geq \kappa\}.$$

Specialized form. The question is equation (2.12) with the joint (14.5) evaluated at each Δ with the same $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$, a positive Δ moving load from hall A to hall B and a negative one the other way. The inversion is by bisection on Δ. It is a statement about the move at this window's loads and coils, and it needs no new fit.

Building blocks. The nodes at each Δ by equation (14.1); the joint (14.5); the operator's κ.

The comparison. Table 14.5 gives the answer at eleven sizes: each row is a transfer; the columns give each hottest inlet, the smallest margin over the ten racks and the three probabilities. Figure 14.4 draws the two hottest inlets against Δ with their 95 % bands, the 27°C limit and the 40 kW of the question. Figure 14.5 draws Pr(both) against Δ, with the three assurances as horizontal lines.

The direction the screen suggests is the one that spends margin. The smallest margin falls at every step from $\Delta = 0$ up. A transfer the other way enlarges it up to the point at which the means of racks A5 and B3 meet, $\Delta = -18.88\,\mathrm{kW}$, where both stand at 23.4464°C and the smallest margin is 3.5536 K, with Pr(both) = 0.99959. The largest joint on a 1 kW grid is elsewhere, at $\Delta = -7\,\mathrm{kW}$, 0.99986. The two do not coincide, because hall A's band is the wider there: equal means do not make equal tails. Table 14.6 gives the largest transfer from hall A to hall B at each assurance: each row is an assurance and each column a route. At $\kappa = 0.95$ the

Table 14.5: The transfer as a size and a direction, on the one fit. Positive: hall A to hall B. Comparison only.

Δ [kW]	T^*_{A5} [°C]	T^*_{B3} [°C]	Smallest margin [K]	Pr(hall A) [–]	Pr(hall B) [–]	Pr(both) [–]
−80	24.9277	21.4133	2.0723	0.95419	1.00000	0.95419
−60	24.4430	22.0787	2.5570	0.98534	1.00000	0.98534
−40	23.9582	22.7440	3.0418	0.99671	1.00000	0.99671
−20	23.4735	23.4093	3.5265	0.99953	1.00000	0.99953
−10	23.2311	23.7420	3.2580	0.99986	0.99998	0.99983
0	22.9887	24.0747	2.9253	0.99996	0.99979	0.99975
10	22.7463	24.4073	2.5927	0.99999	0.99877	0.99876
20	22.5039	24.7400	2.2600	1.00000	0.99474	0.99474
40	22.0192	25.4053	1.5947	1.00000	0.95537	0.95537
60	21.5344	26.0707	0.9293	1.00000	0.82354	0.82354
80	21.0497	26.7360	0.2640	1.00000	0.58940	0.58940

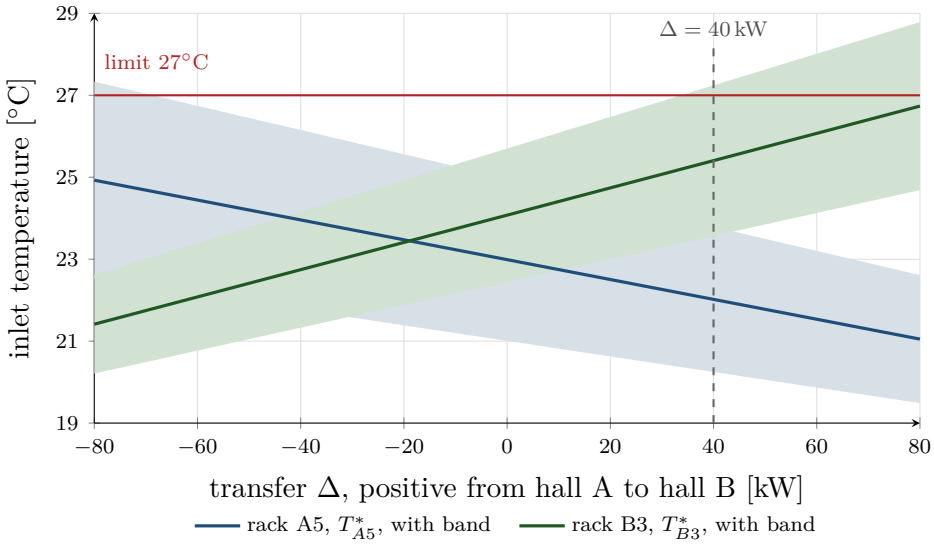


Figure 14.4: The hottest inlet of each node against the transfer, with central 95 % bands, at the sizes of Table 14.5 joined by straight lines. The band of each hall widens as that hall takes on load.

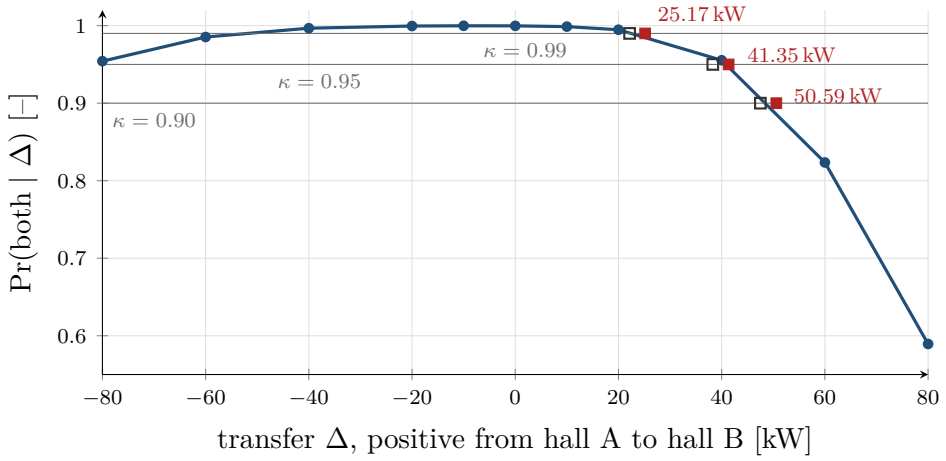


Figure 14.5: The joint of equation (14.5) against the transfer, on the one fit. Red squares: the largest transfer from hall A to hall B at each stated assurance on the window's entry. Open squares: the same with the supply at the set point (Table 14.6).

question's 40 kW lies between 38.28 kW at the set point and 41.35 kW at the window's entry, which is the unsettled verdict of Section 14.3.3 read as a size.

Table 14.6: The largest transfer from hall A to hall B at each stated assurance, by bisection on each route.

κ [–]	Window's entry, reported [kW]	Set point [kW]	Exact posterior [kW]	Fit carrying the bias [kW]
0.90	50.59	47.50	50.64	50.47
0.95	41.35	38.28	41.43	41.20
0.99	25.17	22.16	25.00	24.98

14.4 The checks (step 5)

Table 14.7 collects the checks, each row with its reported value, the value beside it and the gap; a section number stands for a value given there.

14.4.1 The limit of exact rows and the curvature term

On this chapter's own full state of 59 entries, with the physical rows of both nodes at a finite tolerance, sd T^{*}_{B3} is 0.9427, 0.9374 and 0.9374 K at $\sigma_p = 0.1$, 0.01 and 0.001, against the reduced state's 0.9374 K, and sd T^{*}_{A5} is 0.9095, 0.9040 and 0.9039 K against 0.9039 K; the fits took three, three and two iterations. The means agree to the fourth decimal at the last two. The full Hessian, with the curvature term of the nonlinear rows, gives sd $T^{*}_{B3} = 0.93738$ K against 0.93737 K, and sd $T^{*}_{A5} = 0.90392$ K on both, and a joint of 0.95536: the term moves the answer in the fifth digit.

The exact posterior's gap on the joint, +0.00011, is less than its standard error of 0.00033, so it is not resolved at this count. Neither is the gap of T^{*}_{B3}'s median, −0.0003 against a standard error of 0.0019. The gap at T^{*}_{A5}'s median, −0.0041 against 0.0018, is about two standard errors, and it moves nothing, since hall A's tail is out of reach.

Table 14.7: The checks of the answers. Exact: importance sampling through the exact rows, 400 000 draws, effective size 396 137. Truth: the constructed truth at window 73 with the transfer. z: (reported − truth) over the reported spread.

Check	Reported	Beside it	Gap
Joint, exact [–]	0.95537	0.95548 (se 0.00033)	+0.00011
T^*_{B3} median, exact, se 0.0019 [°C]	14.3.2	14.3.2	−0.0003
T^*_{B3} band, exact [°C]	14.3.2	14.3.2	−0.0024, −0.0001
T^*_{A5} median, exact, se 0.0018 [°C]	14.3.2	14.3.2	−0.0041
sd T^*_{B3}, Newton [K]	0.93737	0.93738	+0.00001
Joint, Newton [–]	0.95537	0.95536	−0.00001
sd T^*_{B3}, 59-entry state, $\sigma_p = 0.1,\ 0.01,\ 0.001$ [K]	0.9374	0.9427, 0.9374, 0.9374	to zero
Fit carrying the bias, joint [–]	14.3.3	14.3.3	−0.00059
Supply at the set point, joint [–]	0.95537	0.94246	−0.01291, across 0.95
Truth, T^*_{B3} [°C]	25.4053	26.2705	$z = -0.92$
Truth, T^*_{A5} [°C]	22.0192	22.7394	$z = -0.80$
Truth, the change at A5 and B3 [K]	−0.9695, +1.3307	−1.0741, +1.4134	–
Truth, the ratio [–]	1.373	1.316	inside its band

14.4.2 Sensitivity

Two choices of Section 14.1.2, the fractions independent of load and the exact transfer, and the survey's width are sized here, on the joint, with the reported value at 0.95537. If rack B3's fraction rises with its load, by 0.03 or 0.06 with the survey unchanged, the joint falls to 0.90840 or 0.83202. If the moved load is uncertain, with standard deviation 2 or 5 kW, it is 0.95496 or 0.95283. With the survey's spread on the unread fractions at 0.04 or 0.12, it is 0.99943 or 0.86610. A rise of 0.03 in the fraction takes the joint below 0.95 and a rise of 0.06 below 0.90, and no row of this model can test either. An uncertain Δ costs little. With unit A1 stopped as in Chapter 13, the 40 kW moves hall A's probability from 0.47350 to 0.83988 and leaves the joint at 0.80239.

14.4.3 The answers against the stated truth

The record is constructed, and the do-it-and-see check runs the transfer on the plant's true state at window 73, with no noise. Every inlet stays under; the maximum is 26.2705°C at rack B3, 0.7295 K inside. The ten z values run from −1.55 to +0.36, all inside their bands. The true change at the two hottest is a little larger than the fit's: −1.0741 and +1.4134 K against −0.9695 and +1.3307 K, in the same direction and the same order.

Beyond the question's window, the truth only. Holding the transfer through the shift, the true plant is first over 27°C at window 90 (02:25) and is over for 53 windows, with rack B3 at 28.304°C at window 144. With no transfer it is over for 0 windows, with rack B3 at 26.754°C at window 144. With the transfer from hall B to hall A it is over for 0 windows, with rack B3 at 25.204°C and rack A5 at 25.328°C. The answer at 01:00 is the plant at 01:00, and this is what that leaves unsaid. It is a check, not a forecast: the report does not read it.

14.4.4 One window's draw

The center of the band is one window's readings, so the joint at 40 kW is itself a draw. Figure 14.6 draws it over the plain fits of windows 61 to 85: the horizontal axis is the window, the vertical the joint, the solid line 0.95, the dashed line the median of those windows, and the open diamond the supply at its set point at window 73. Later windows use readings the operator does not have at 01:00, and they enter here as a check only. The plain fits of windows 74, 75 and 76

give 0.97505, 0.92984 and 0.93902, with rack B3 at 25.1536, 25.6160 and 25.5549°C. Over windows 61 to 85 the median is 0.9484, the least 0.8852 and the greatest 0.9782, and twelve of the twenty-five reach 0.95.

The spread of that draw can be read directly. With 400 fresh sets of sensor errors at window 73's constructed truth, the plain fit's joint at 40 kW has mean 0.9470, standard deviation 0.0191 and 5 to 95 % range 0.913 to 0.973, and 50.0 % of the sets reach 0.95. The record's own draw is 0.95537.

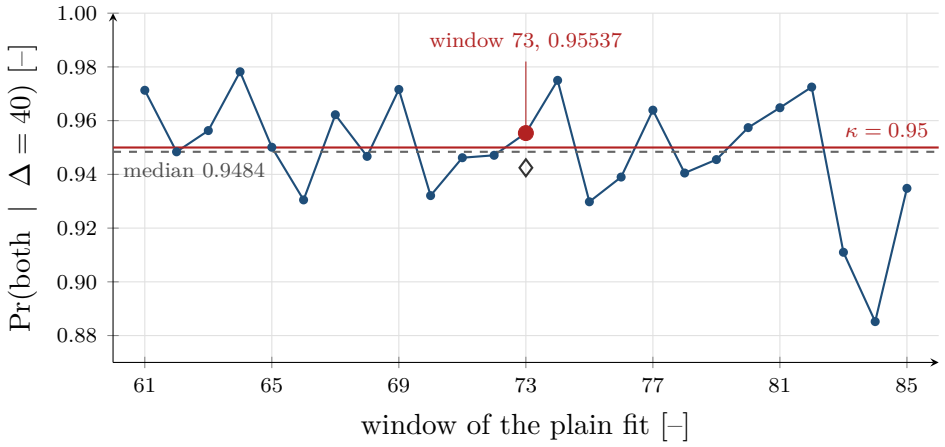


Figure 14.6: The joint of equation (14.5) at $\Delta = 40$ kW over the plain fits of windows 61 to 85. The points after window 73 use readings later than the question. Red dot: the reported fit. Open diamond: the supply at the set point, at window 73. Dashed line: the median over all twenty-five windows, a check; the median over windows 61 to 73 is the one reported.

14.4.5 The steady-state limit

Moving a batch job is not instantaneous, and the plant's response to it is not modeled in time. The nodes say where the halls settle with 40 kW across. They say nothing of a transition in which both halls carry the job, which no row here represents, or of its order.

14.5 The report (step 6)

Moving load out of the heavier, warmer hall gains rack A5 0.97 K (band 0.76 to 1.18) and costs rack B3 1.33 K (band 1.12 to 1.55), a ratio of 1.37 (band 1.05 to 1.83; exact 1.37), and at the means rack B3 already held the smallest margin; it falls from 2.93 K to 1.59 K. Rack B3 is the hottest of the ten now with probability 0.73 and at the nodes 0.88. The probability that both halls stay inside at 40 kW is 0.955 (exact 0.955), and one rack carries its tail. At an assurance of 0.95 the verdict depends on the supply's value, 0.942 at the set point, and on the window, with a median of 0.950 over windows 61 to 73. Both values stand, and this report does not choose between them.

14.5.1 The answers and their bands

Table 14.8 gives the answers in the order of Table 14.2: each row is a question, the second column the answer with the exact posterior and the fit carrying the bias beside it, and the third its band or its routes. The comparison by size and direction is a comparison, and no move is named.

Table 14.8: The report for the transfer at window 73 (01:00). Answers from the one fit by the delta method; the joint by the ten-dimensional normal.

Question	Answer	Band or routes
Where does the margin go?	A5 gains 0.97 K, B3 loses 1.33 K; smallest margin at the means 2.93 to 1.59 K, at B3	bands 0.76–1.18 and 1.12–1.55 K; ratio 1.37, band 1.05–1.83 (exact 1.37; bias carried 1.36)
The band on node A	rack A5, 22.02°C (exact 22.02; bias carried 21.99)	20.25 to 23.79°C
The band on node B	rack B3, 25.41°C (exact 25.41; bias carried 25.41)	23.57 to 27.24°C; 96.9 % of the variance the survey's β_{B3}
Both inside at 40 kW?	0.955 (exact 0.955; bias carried 0.955)	0.942 with the supply at set point; median 0.950 and range 0.93 to 0.98 over windows 61–73
Other sizes and direction	largest Δ from A to B: 50.59, 41.35, 25.17 kW at $\kappa = 0.90, 0.95, 0.99$; at the set point 47.50, 38.28, 22.16 kW	smallest margin grows the other way up to 18.88 kW, to 3.55 K

14.5.2 The entry in the shift log

01:00, window 73. What-if: 40 kW hall A → hall B at this window's loads. Rack A5 22.02°C (band 20.25–23.79), rack B3 25.41°C (band 23.57–27.24), both unread. A5 gains 0.97 K, B3 loses 1.33 K, ratio 1.37 (band 1.05–1.83); smallest margin at the means 2.93 → 1.59 K, at B3. Both halls under 27°C: 0.955 (exact 0.955; 0.942 with the supply at set point; 0.93–0.98 on windows 61–73), so not settled at assurance 0.95. Largest move A → B at assurance 0.99: 25 kW (22 at set point). The other direction enlarges the smallest margin up to 19 kW. Steady state at 01:00 loads and coils; B2's coil and the loads are moving. Action: not decided here.

Common mistake

Reading a transfer as margin moved one for one, or the heavier hall on the meters as the hall short of margin. The margin is read at each rack, by its slope and from its own hall's temperatures. And a probability that clears a stated assurance by a small amount has not cleared it until the other routes are set beside it.

14.6 The next reading (step 7)

The report leaves a joint that one rack's survey fraction carries. A further reading is a further row, and what it would do to the entries is known before its value is.

14.6.1 The preposterior of a candidate

A candidate adds rows to $\mathbf{J}$ at $\hat{\mathbf{z}}$, and the covariance after it is the preposterior of equation (4.65), with $\mathbf{\Sigma}_a$ of the rank-one form (4.66) for one row. A handheld probe of standard deviation 0.3 K adds the row of a rack's present inlet, gradient over 0.3, and so acts on that rack's fraction; rack B3 and rack A5 are the two candidates. A re-survey of β_{B3} with standard deviation 0.02 adds one prior row. A permanent inlet sensor at rack B3, the instrument at the place of the question, adds the probe's row in each of six windows. A clamp-on meter on hall B's load, of standard deviation 1 kW, adds a row on Q_B. A water-side heat check on unit B2, of standard deviation 3 kW, adds a row on that unit's heat. For each, the outcome is drawn from its predictive, the mean moves with the outcome, and the joint (14.5) is recomputed at $\mathbf{\Sigma}_a$; the

quantity reported is the chance of ending at or past the operator's assurance,

$$\Pr\Big(\Pr\big(\max_{k} T_k^{*} < T_{\max} \mid \mathbf{y},\ y_a,\ \Delta = 40\big) \geq \kappa\Big), \tag{14.6}$$

over 2000 outcomes y_a, with the joint at each outcome by integration and a standard error of the fraction of at most 0.011.

14.6.2 The candidates

Table 14.9 compares them by what each does to the entries: each row a candidate; the columns the spread of β_{B3} and of T_{B3}^{*} after it, the joint with the means unmoved and the covariance $\mathbf{\Sigma}_a$, and the quantity (14.6) at $\kappa = 0.95$ and 0.99, on the window's entry and with the supply at the set point. 0.95 is met at the means on the window's entry and not at the set point.

Table 14.9: Candidate readings at the nodes, from the posterior at window 73, on both routes of the supply.

			Window's entry			Set point		
Candidate	sd β_{B3} [–]	sd T_{B3}^{*} [K]	Pr at means [–]	≥ 0.95 [–]	≥ 0.99 [–]	Pr at means [–]	≥ 0.95 [–]	≥ 0.99 [–]
None	0.0800	0.9374	0.95537	1.0000	0.0000	0.94246	0.0000	0.0000
Handheld probe at rack B3, sd 0.3 K	0.0303	0.3194	0.99981	0.8990	0.8455	0.99967	0.8695	0.8005
Handheld probe at rack A5, sd 0.3 K	0.0800	0.9374	0.95537	1.0000	0.0000	0.94246	0.0000	0.0000
Re-survey of β_{B3}, sd 0.02	0.0194	0.2781	0.99981	0.8930	0.8485	0.99968	0.8585	0.8080
Permanent inlet sensor at rack B3, six windows	0.0183	0.1383	0.99981	0.9330	0.9180	0.99968	0.9105	0.8880
Clamp-on on hall B's load, sd 1 kW	0.0800	0.9372	0.95539	0.9995	0.0000	0.94249	0.0000	0.0000
Water-side heat check on unit B2, sd 3 kW	0.0800	0.9374	0.95537	1.0000	0.0000	0.94246	0.0000	0.0000

This is the same rack as in Chapter 12: the reading that chapter asks for serves this question as well, because the node's largest term is the same entry. A single reading at rack B3 can take the answer past 0.99 or below 0.95, on either route: on the window's entry the 0.95 column falls from 1.0000 to 0.8990 with the probe, since the band can now say no, and at the set point it rises from 0.0000 to 0.8695. The probe at rack A5 and the heat check on unit B2 change nothing at either assurance, and the clamp-on on hall B's load changes nothing the standard error can resolve (0.9995). The permanent sensor does most, and it also makes the term testable window by window. The choice among them, and any action, stay with the operator. Chapter 18 ranks readings against a named worry.

14.6.3 What the analysis showed

The transfer is not margin moved one for one. At the means the donor was not the constraint at this window, and the meters that made it look so were reading a load, not a margin. The slope has a closed form, which favors the larger conductance and penalizes the larger fraction, so that the rack that at the means already held the least margin lost the most. On this plant at this size, two nodes collapse to one rack's tail, because hall A's tail is out of reach, and the nodes were needed to find that out. At a stated assurance the record cannot always answer, because the center of the band is itself one window's draw and the supply's value is a choice the window does not settle. Hall B's derate is about a quarter of the excess, and the rated conductance most of it.

14.6.4 Limits of the model

The steady state. There is no transition (Section 14.4.5), and the limits of Chapter 5 hold in this window as well.

The hall as one heat node. The model has no row from a rack's load to its own inlet, so the answer does not depend on which rack receives the job, and this choice has no size here.

Fractions independent of load. The β_k are unchanged at the nodes; Section 14.4.2 gives what a fraction that moves would cost.

Loads and coils held at 01:00. Hall B's coil stands below its early level and the loads have risen across the record to this window (Chapters 12 and 13). Every number here is at this window's level, and the transfer held later in the shift is not asked here.

The plain fit and one window's draw. The reported route is the plain fit, with the bias carried beside it, and it is one window's draw (Table 14.4).

Constructed readings. The record is constructed: generated from a stated truth of the plant with the stated sensor errors, and Section 14.4.3 compares the answers with that truth. It is not field data, and the agreement is a check of the method on its own model.

14.7 What to retain

Definition 14.1 (Load transfer as two scenario nodes)**.** A load transfer Δ between two halls on one loop is read at two scenario nodes, $Q_A^* = Q_A - \Delta$ and $Q_B^* = Q_B + \Delta$, with the rest of the fitted state carried and $\boldsymbol{\Sigma}$ carried by $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{g}$. An inlet moves along $\partial T_k/\partial Q_h = 1/\sum G_h - (1-\beta_k)/\sum C_h$. The probability both halls stay inside is the joint tail over all inlets of both nodes, and it is compared along Δ and inverted at the operator's κ.

Three conclusions carry beyond this plant. Margin is read rack by rack, and a transfer sends it where the slopes send it, not where the meters suggest. Two nodes, whose one hall's tail is out of reach, can come down to a single rack's tail, and the computation is what shows whether they do. And when a reported probability lies close to a stated assurance, the routes and the neighboring windows belong beside it, and a verdict that depends on them is reported as not settled.

Table 14.10 lists the terms of the chapter with their expressions, units and where each is defined.

Table 14.10: The terms of the load transfer and where each is defined. Units for the shift plant.

Term	Expression	Units	Meaning	Defined in
The two nodes	$Q_A^* = Q_A - \Delta$, $Q_B^* = Q_B + \Delta$	kW	the transfer as a what-if	14.1.2
Forward chain	eq. (14.1)	°C	temperatures re-solved	14.1.2
Slope of an inlet	$1/\sum G_h - (1-\beta_k)/\sum C_h$	K/kW	the inlet's move per kW	14.1.2
Gradient pattern	eq. (14.3)	–	only T_s is a shared column	14.1.4
Margin	$27 - T_k^*$	K	distance below the limit	14.3.1
Gradient at the node	eq. (14.4)	K, K/kW	carries $\mathbf{\Sigma}$ to the inlet	14.3.2
Joint over the nodes	eq. (14.5)	–	both halls inside	14.3.3
Largest transfer	$\Delta_{\max}(\kappa)$, Table 14.6	kW	the inversion at an assurance	14.3.4
Chance of meeting κ	eq. (14.6)	–	compares candidates	14.6.1

14.8 Exercises

Exercise 14.1. *Derivation.* Derive the slope of rack B3's inlet in hall B's heat from the forward chain and evaluate it at the fit. (By equation (14.1), $T_k^* = T_{r,h}^* - (1-\beta_k)(Q_h^* + Q_{o,h})/\sum C_h$, so the slope is equation (14.2). With $\sum G_B = 20.0268\,\text{kW/K}$, $\sum C_B = 30\,\text{kW/K}$ and $\beta_{B3} = 0.5000$, it is $1/20.0268 - 0.5/30 = 0.03327\,\text{K/kW}$.)

Exercise 14.2. *Application.* From the gains and losses of Table 14.3, form the ratio of the hottest pair's movements and say what a one-for-one transfer would have made it. (One for one would give a ratio of 1. Rack B3 loses 1.3307 K and rack A5 gains 0.9695 K, so the median is 1.373; its inversion band, from the normal of the loss less a multiple of the gain, is 1.054 to 1.829.)

Exercise 14.3. *Derivation.* Using the gradient (14.4) and the survey's spread, find the share of β_{B3} in the variance of T_{B3}^*. (The entry is $T_{r,B}^* - T_{c,B}^* = 11.5337\,\text{K}$, and $\Sigma_{\beta_{B3}\beta_{B3}} = 0.08^2$, so the term is $0.85137\,\text{K}^2$. The variance is $0.87865\,\text{K}^2$, so its share is 96.9 %.)

Exercise 14.4. *Application.* Find the largest transfer from hall A to hall B at $\kappa = 0.99$, by inversion of the joint. (From Table 14.5, the joint lies between $\Delta = 20\,\text{kW}$, 0.99474, and $\Delta = 40\,\text{kW}$, 0.95537; it is above 0.99 at the first and below at the second, and bisection on equation (14.5), the nodes re-solved at each step, ends at 25.17 kW on the window's entry; pinned at the set point it ends at 22.16 kW.)

Exercise 14.5. *Formulation.* Write the operator's question "which rack will be the hottest after the move" for rack B3 as a probability over the state, in the general form of equation (2.12), and give its answer. ($\Pr(T_{B3}^* - \max_{k\neq B3} T_k^* \geq 0 \mid \mathbf{y},\ \Delta = 40)$, with $f = T_{B3}^* - \max_{k\neq B3} T_k^*$ and $f^* = 0$; read from the normal of equation (14.5) it is 0.875.)

Chapter 15

Look-Ahead Risk Without Intervention

15.1 The hours before the morning shift

It is 01:00, window 73 of the shift, and six hours remain before the morning shift arrives. The four inlet sensors read 19.51°C at rack A1, 22.03°C at A3, 22.13°C at B2 and 22.67°C at B4. The hottest read inlet, B4, is 4.33 K inside the 27°C limit, and the largest inlet reading of the whole shift so far is 23.47°C. Straight lines through the readings since 19:00 reach 23.54°C at B4 and 23.37°C at B2 by the end of the shift. The steepest line on the screen, B2's last hour, reaches 25.16°C and crosses 27°C only at window 197, more than four hours after the morning shift has arrived. From the screen alone the operator concludes that nothing reaches the limit before 07:00.

Those lines are drawn through the racks that have sensors, one reading at a time, and the screen holds no line through a rack without one. Nor does a line carry the load's wander from window to window, or say how well its own slope is known. The operator's question is the one the screen cannot answer,

if nothing is done, what is the probability that some rack, read or unread, is over 27°C *at any window before the morning shift.*

It is answered on the two-hall shift plant, and it extends the trend of Chapter 11 in one respect: the crossing there was a line against a level at one steady load, and the event here is a statement about a wandering path, over ten racks and every window to the horizon.

Jargon: first passage, do-nothing, trend path, wander, persistence, orthant, conditional hazard, decision window, fresh sensor errors, shared survey offset

A *first-passage* event is the event that a quantity reaches a level at some time up to a stated horizon, whatever it does afterward. *Do-nothing* is the premise of the forecast: the set point is held and no unit is changed. The *trend path* is the path of the plant with each window's load at its trend and the set point at its level. The *wander* is the load's departure from that trend, and its *persistence* ϕ is the share of one window's departure that remains in the next. An *orthant* of a joint normal is the region in which every one of its entries lies below its own level. The *conditional hazard* of the first passage is the probability that it occurs in the next interval, given that it has not occurred yet, per hour. The *decision window* is the window at which the question is asked; the record used stops there. *Fresh sensor errors* redraw the reading errors of the constructed record and hold its true loads and set point, to show how far one night's sensor errors move an answer. A *shared survey offset* is an error common to all the survey's recirculation fractions.

Key idea

A first passage is the union of a threshold event over racks and windows. It has no gradient, since it reads a maximum, but the state still defines it: the inlets, linearized about the MAP state, are jointly normal with the mean and covariance read from $(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$, and the event is the complement of an orthant of that normal. Such an integral has no closed form, so it is taken by draws, with their standard error. What a look-ahead probability is made of is then a separate question from its size: entries the record reads, entries the record never reads, and a future that has not happened.

15.1.1 A probability over a path

Chapter 11 carried a line forward and read a date from it. A date is the right report when the question is when, and when the quantity has one path. An operator who is deciding whether a night can be left alone asks something else: the probability of the event by the time the night is over. The event is true of a path if the path is over the limit in any window, even if it comes back under, so a reading of the last window alone can miss it.

Two facts decide how far that matters. A path that rises in every window is over at some window exactly if it is over at the last, and the trend path is such a path. The wander is not: a window can stand above the trend and the next below it. Which of those two pictures describes the night, and by how much the answer differs between them, is measured and not assumed.

> **Principle 15.1.** *A look-ahead probability is a statement about a path, computed on the joint distribution of the path. It reduces to the last window only for a path that cannot come back, and its size says nothing about whether waiting can settle it: that depends on whether the windows ahead read the entries that carry it.*

15.1.2 The stacked shift plant through the night (step 0)

The plant is the two-hall shift plant on the reduced state of Chapter 4: halls A and B on one chilled-water loop, three units per hall with pinned ratings $\bar{G}_u$ of 8.0 kW/K in hall A and 7.0 kW/K in hall B, capacity rates $C_u = 10.0$ kW/K, other heat $Q_{o,A} = 12$ kW and $Q_{o,B} = 10$ kW, five racks per hall. Windows are five minutes long, window 1 at 19:00. The dependent block of each window's physical rows is triangular with a nonzero diagonal, hence invertible, and the temperatures of hall h follow from the window's free entries in the chain

$$G_u = \bar{G}_u \, \mathrm{e}^{g_{u,t}}, \qquad T_{r,h,t} = T_{s,t} + \frac{Q_{h,t} + Q_{o,h}}{\sum_{u \in h} G_u}, \qquad T_{d,u,t} = T_{r,h,t} - \frac{G_u \, (T_{r,h,t} - T_{s,t})}{C_u},$$
$$T_{c,h,t} = \frac{\sum_{u \in h} C_u T_{d,u,t}}{\sum_{u \in h} C_u}, \qquad T_{k,t} = T_{c,h,t} + \beta_k \, (T_{r,h,t} - T_{c,h,t}). \tag{15.1}$$

Summed over the units, the air sides and the mixing give $(\sum_u C_u)(T_r - T_c) = Q_h + Q_{o,h}$: the rise is set by the heat and the airflow, not by the coils, as in equation (5.5).

Why not the one fit. Part IV's one fit is window 73 alone (Chapter 12). A forecast needs a slope, and one window has none, so this chapter stacks windows 1 to 73 in its place. Section 15.2.3 sets the two side by side where they overlap.

The stacked state. The windows through the horizon form one state,

$$\mathbf{z} = (T_s, \; a_A, \; b_A, \; a_B, \; b_B, \; g_{A1}, \ldots, g_{B3}, \; \gamma, \; \beta_{A1}, \ldots, \beta_{B5}, \; b_{Q_A}; \\ T_{s,t}, \; Q_{A,t}, \; Q_{B,t} \text{ for } t = 1, \ldots, 144)^{\mathsf{T}}, \tag{15.2}$$

23 shared entries and 3×144 per-window entries, 455 in all, 242 of them belonging to windows 1 to 73. Each hall's load is a trend and a wander about it, and unit B2's log-derate is a level with a ramp from an onset τ,

$$Q_{h,t} = a_h + b_h\,(t - \bar{t}) + w_{h,t}, \qquad w_{h,t} = \phi\, w_{h,t-1} + \varepsilon_{h,t}, \qquad g_{\mathrm{B2},t} = g_{\mathrm{B2}} - \gamma\,(t - \tau)_+, \quad (15.3)$$

with $\bar{t} = 37$ the center of windows 1 to 73 and $(x)_+ = \max(x, 0)$. The load's line is the trend form (2.19), and the coil's is its right arm, with $\gamma > 0$ worse: the ramp of Chapter 11. The windows from 74 on carry the forecast in the same state.

The event. The horizon, window 144, 06:55, ending at 07:00, is the operator's stated horizon. The chapter computes for that horizon and for each earlier one and sets no level of assurance. Over the ten racks, each window at its own state, the event is

$$E_s = \Big\{ \max_{k} \max_{73 \le t \le s} T_{k,t}(\mathbf{z}) \ge T_{\max} \Big\}, \qquad T_{\max} = 27^{\circ}\mathrm{C}, \quad (15.4)$$

with E alone meaning E_{144}. The window is the grain of the event: the plant is steady within a window, and the question is not defined more finely. The event is a statement about a path, so it is not the threshold form (2.12) at one window; it is the union of that form across windows and racks.

Modeling choices at the graph. The choices of the shift plant hold in every window: steady state within the window, one effective conductance per coil, pinned airflow and no active control, each failing as in Chapter 5. Do-nothing is the premise and not a pin: the set point is held and no unit changes, so no scenario node is drawn, because nothing is pinned to a value not yet applied.

The graph. Figure 15.1 draws the stacked state. At the top are the shared entries (green). Below them are the record's windows, each with its own set point and two loads and each sending its 16 readings out along dashed arrows. To the right are the forecast windows, carrying the same entries with no reading. The process rows (blue) tie each window's load to the one before and to the trend. The event (red) stands at the right and takes the forecast windows and the fractions; among those fractions the six of the unread racks have no dashed arrow anywhere in the figure.

15.1.3 The operator's questions (step 1)

The operator asks six questions of the record at window 73, and step 1 writes each over the stacked state. Table 15.1 lists them: the question in the operator's words, the first of them the drill question; the general form of its category in the notation of Chapter 2; its mathematical form on this plant; its category; the computation of Section 2.3.6 that answers it; and the subsection where it is answered.

15.1.4 The evidence as rows (step 2)

Step 2 attaches the record of windows 1 to 73 to the graph: 16 readings per window, 1168 reading rows. They are the supply temperature (standard deviation 0.1 K), the two returns, the six discharges and the four read inlets (0.3 K), the two hall meters (4 kW) and the loop meter (8 kW), each predicted through the chain (15.1) at its window's entries. Priors and process rows make up 453 rows more. Each choice they carry is named here.

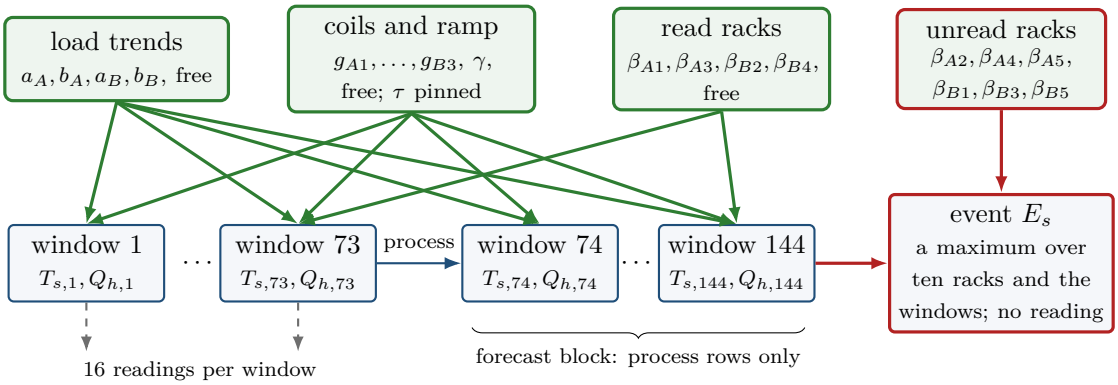


Figure 15.1: The stacked state, equation (15.2). Free (green): the shared entries. Blue: each window's own entries, tied to the trend and the previous window by process rows; only one arrow of that chain is drawn. Observed (dashed): the readings of windows 1 to 73. The priors, the set point's level and the meter bias are rows not drawn. The event (red) has no reading, and neither do the unread racks' fractions. Inside each window the plant is the chain (15.1).

Table 15.1: The questions asked of the record at window 73, 01:00. $\mathbf{y}$ is the 1168 reading rows and the prior and process rows of equation (15.6); E_s the event (15.4); M the largest inlet over the racks and windows.

Operator question	General form	Mathematical form	Category	Computation	Answered in
If I do nothing, what is the probability of a violation before the morning shift? Trend forecast of the load against the margin band, integrated over the horizon.	(2.21), line 2; (2.12) per window	$\Pr(E_{144} \mid \mathbf{y})$	prediction under a threshold	trend	15.3.1
Does the path matter, or only the last window?	(2.21), line 3	$\Pr(T_{\mathrm{B3},144} \geq T_{\max} \mid \mathbf{y})$ with and without wander, against $\Pr(E_{144} \mid \mathbf{y})$	prediction under a threshold	trend	15.3.1
How does it grow with the horizon?	(2.21), line 2, per s	$\Pr(E_s \mid \mathbf{y})$; the conditional hazard per hour	prediction under a threshold	trend	15.3.2
Which rack carries it?	(2.12)	$\Pr(\max_t T_{k,t} \geq T_{\max} \mid \mathbf{y})$ per rack; the budget; $\Pr(E \mid \beta_{\mathrm{B3}}, \mathbf{y})$	prediction under a threshold; estimation	trend	15.3.3
How much is what we do not know, how much is what has not happened yet?	(2.10)	the split of $\mathrm{Var}(T_{\mathrm{B3},144} \mid \mathbf{y})$; $\Pr(E)$ with the entries at the fit	estimation	trend	15.3.4
What would the same question have said earlier and later?	(2.21) per T	$\Pr(E \mid \mathbf{y}_{1:T})$ from each decision window; its spread over fresh sensor errors	prediction under a threshold	trend	15.3.5

The ramp with a pinned onset. The onset τ is a parameter of the drawing and not an entry, set at the mode of its evidence, because a kink has no gradient. This buys one Gaussian state. It fails by leaving the onset's spread out of $\mathbf{\Sigma}$, and Section 15.3.1 measures what that leaves out.

A straight load trend to the horizon. The night's batch ramp is taken to go on as it has gone. This buys a slope that reaches window 144 with its lever arm, and it fails if the ramp bends or ends; no row can test that at 01:00. The slopes carry a prior of 1 kW per window, one spread being 12 kW per hour.

The wander as one persistent process. Both halls share a persistence ϕ and a stationary spread s, pinned at the mode of their evidence. Unlike the one steady forecast load of Chapter 11, the wander enters the event. This buys a load that moves from window to window inside one Gaussian state. It fails if the wander's size changes through the night or has more than one time scale.

The survey. Each fraction has its airflow-survey prior, $\mathcal{N}(\beta_k^{\mathrm{s}},\ 0.08^2)$, with independent errors. This keeps the fractions of the unread racks untouched by the read ones, and it fails if the survey's errors are shared; Section 15.4.3 tests that against a shared offset.

The set point and the meter. Each window's set point has a jitter s_j about one level. The hall-A meter carries a freed bias entry b_{Q_A} with a prior of 5 kW, which the one-window fit of Chapter 12 does not free, because over seventy-three windows the air-side rise of each hall against its own meter, with the loop balance, can read a bias and name its meter, which one window cannot. This buys a load read free of its meter's offset; it fails if the jitter is not steady or the bias drifts through the night.

Which row acts on which entry. With $*$ for an entry that can be nonzero, the reading rows of one window have the pattern

$$\begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & * & 0 & 0 \\ 0 & * & 0 & 0 & 0 & 0 & 0 & * & * & 0 \\ 0 & 0 & * & * & 0 & 0 & 0 & * & 0 & * \\ 0 & 0 & 0 & 0 & 0 & 0 & * & 0 & * & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & * \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & * & * \\ 0 & * & 0 & 0 & * & 0 & 0 & * & * & 0 \\ 0 & 0 & * & * & * & 0 & 0 & * & 0 & * \end{bmatrix} \quad \begin{array}{c} \text{supply} \\ \text{return and discharges, hall A} \\ \text{return and discharges, hall B} \\ \text{meter } Q_A \\ \text{meter } Q_B \\ \text{loop meter} \\ \text{read inlets, hall A} \\ \text{read inlets, hall B} \end{array} , \tag{15.5}$$

$$\mathbf{g}_{T_{\mathrm{B3},t}}^{\mathsf{T}} = (*,\ 0,\ *,\ *,\ 0,\ *,\ 0,\ *,\ 0,\ *),$$

with the columns grouped as (a_h, b_h); the g_u of hall A; the g_u of hall B; γ; the fractions of the read racks; the fractions of the unread racks; b_{Q_A}; $T_{s,t}$; $Q_{A,t}$; $Q_{B,t}$. The set point's level T_s acts through the jitter rows alone. Over the windows the Jacobian is an arrowhead: each window's rows see the shared columns and that window's own three. The second line is the gradient of a question's quantity, and it is not a row; the load reaches it through (a_B, b_B) on the trend path and through $(T_{s,t}, Q_{B,t})$ at a window's own state, and both routes are marked; the level T_s, which has no group of its own, also has a nonzero gradient.

Three facts are read here before any number. The columns of the fractions of the six unread racks, A2, A4, A5, B1, B3 and B5, are empty in every reading row of every window, read as observability: those six are unobservable from this record at any T, and the gradient of the

question has a nonzero in one of those columns. The trend entries reach no reading directly: they act only through the process rows that tie the windows' loads. And the windows from 74 on have process rows only, so at window 73 their block is a Gaussian given the shared entries and the last window.

15.2 The posterior of the state (step 3)

The inputs are the rows of Section 15.1.4, and every fit an answer uses is derived here.

15.2.1 The rows and the normal equations

With $w_{h,t} = Q_{h,t} - a_h - b_h(t - \bar{t})$ read from the state and each row divided by its standard deviation, the rows are

$$\mathbf{r}(\mathbf{z}) = \begin{bmatrix} (m_k(\mathbf{z}_t) + \delta_{k,Q_A} b_{Q_A} - y_{k,t})/\sigma_k \quad (t = 1, \ldots, 73;\ k = 1, \ldots, 16) \\ (T_s - 14)/0.5, \quad g_u/0.15, \quad \gamma/0.02, \quad b_h/1, \quad (\beta_k - \beta_k^s)/0.08, \quad b_{Q_A}/5 \\ (T_{s,t} - T_s)/s_j \quad (t = 1, \ldots, 144) \\ w_{h,1}/s, \quad (w_{h,t} - \phi w_{h,t-1})/(s\sqrt{1-\phi^2}) \quad (t = 2, \ldots, 144;\ h = A, B) \end{bmatrix}, \tag{15.6}$$

with $\mathcal{L}(\mathbf{z}) = \frac{1}{2}\mathbf{r}^\mathsf{T}\mathbf{r}$, m_k the prediction of reading k by the chain (15.1) at window t's entries $\mathbf{z}_t$, and δ_{k,Q_A} one on the hall-A meter's row and zero elsewhere. The second line holds 21 rows, the third 144 and the last 288, the 453 of step 2. Only the reading rows are nonlinear. The MAP state minimizes $\mathcal{L}$ by the normal equations of equation (4.29),

$$(\mathbf{J}^\mathsf{T}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^\mathsf{T}\mathbf{r}(\mathbf{z}), \qquad \boldsymbol{\Sigma} = (\mathbf{J}^\mathsf{T}\mathbf{J})^{-1}\big|_{\hat{\mathbf{z}}}, \tag{15.7}$$

the Gauss–Newton form of equation (4.38), with the posterior approximately normal, $\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$. By the pattern (15.5) the rows of the six unread fractions are their prior rows alone, so the objective separates, as in equation (5.10), and their posteriors are the survey's exactly, not only under the linearization.

The exact posterior. The density $p(\mathbf{z} \mid \mathbf{y}) \propto \exp(-\mathcal{L}(\mathbf{z}))$ is set beside the Laplace normal by importance sampling, with draws $\mathbf{z}^{(j)}$ from $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ weighted by $w_j \propto \exp\big(\frac{1}{2}(\mathbf{z}^{(j)} - \hat{\mathbf{z}})^\mathsf{T}\boldsymbol{\Sigma}^{-1}(\mathbf{z}^{(j)} - \hat{\mathbf{z}}) - \mathcal{L}(\mathbf{z}^{(j)})\big)$. With 200 000 draws the effective size, $(\sum_j w_j)^2 / \sum_j w_j^2$, is 199 790, and the largest weight's share is 0.00001: the reading rows are close to linear over the posterior's width.

15.2.2 The drawing the record selects

Before the fit, six parts of the drawing are found from the record by the methods of Part III, with the Laplace evidence of equation (11.8). Each is reported as the record shows it at window 73.

The coil. A ramp on g_{B2} beats no change by 89.73 log units, with its slope at $z = 13.54$. A line over the record on B1 or B3 alone looks significant, $z = 3.60$ and 3.78, only through smearing: on top of B2's ramp every other coil's line loses, by 3.70 to 5.04 log units, with the largest $|z|$ at 1.75, so the record supports no second coil trend. A drift of B2's discharge sensor at the same onset loses to the coil by 25.76 log units, because a coil moves the return and the inlets as well, by the reasoning of Chapter 10. Every answer below is conditional on the coil.

The onset. On the grid of windows 2 to 72 the onset has mode 35, 21:50, median 35.6 and band [30.3, 41.2], 21:25 to 22:25, a posterior read directly from the grid's evidence. The drawing pins $\tau = 35$; the band stays with the answer.

The loads. With no meter bias freed, hall A's load has slope 0.2733 ± 0.0737 kW per window, $z = 3.71$, and hall B's 0.1834 ± 0.0736, $z = 2.49$, 2.20 kW per hour. Hall B flat against sloped scores -0.49 log units: the record at 01:00 does not insist by evidence on hall B's slope, which is therefore carried with its spread and its prior. The meter-only line of Chapter 13 gives hall A 0.2904 (0.0715) kW per window and this stacked fit 0.2733 (0.0737), because this one uses all readings; they are about 0.24 spreads apart, 1.8 kW over the 107 windows from the center to the horizon. Persistence does not cause the gap: under this chapter's slope prior of 1 kW per window the meter alone gives 0.2950 at Chapter 13's persistence and 0.2953 at this one, and Chapter 13's 0.2904 is lower only because its slope prior is the tighter 0.5 kW per window. A curvature on the load scores -0.03 log units.

The meter. The meter balance, loop less the halls less the other heat, has mean -3.507 kW, $z = -3.36$. A bias on the hall-A meter beats no bias by 10.56 log units, at $+4.463 \pm 0.901$ kW; a bias on the hall-B meter scores -1.37 and on the loop meter $+1.92$. The drawing frees the first, which moves hall A's slope to the 0.2727 of the fit.

The wander. The evidence over persistence and stationary spread has its mode on the grid at $\phi = 0.90$, $s = 4.0$ kW, the record's selection. Its marginal on ϕ spreads over 0.70 to 0.95 and on s over 3 to 6 kW. The meter's 4 kW noise is as large as what it measures, so the wander is read mostly through the temperatures. Chapter 13's meter-only fit selects $\phi = 0.85$, $s = 4.25$ kW on a finer grid that this one does not contain; on all the readings those values score -0.46 log units against this selection and give 0.0864 against 0.0829 on the same draws, so the two agree within what the wander's evidence can tell apart.

The set point. The excess variance of the supply reading over its sensor gives a jitter $s_j = 0.0198$ K, carried.

15.2.3 The fit at the decision window

The iterations. Table 15.2 runs equation (15.7) at window 73. Each row is an iteration; the columns give the objective and the norm of its gradient. The objective falls from 2267.0272 to 591.5974 in the first step and settles at 590.9284 by the second.

Table 15.2: Gauss–Newton iterations of the stacked state at window 73, equation (15.7).

Iteration k [–]	Objective $\mathcal{L}(\mathbf{z}^{(k)})$ [–]	$\lVert\mathbf{J}^{\mathsf{T}}\mathbf{r}\rVert$ [–]
0	2267.0272	7.39×10^{4}
1	591.5974	1.58×10^{3}
2	590.9284	1.29
3	590.9284	2.89×10^{-2}
4	590.9284	7.02×10^{-5}
5	590.9284	3.36×10^{-6}

The MAP state. Table 15.3 gives the shared entries, each row an entry, the columns its value, posterior spread and prior or survey center. Six fractions stand at the survey with the survey's spread; the four read fractions have moved and their spreads have fallen to about 0.004. The coil at window 73 is $g_{\mathrm{B2},73} = -0.11818 \pm 0.00715$, $G_{\mathrm{B2}} = 6.220$ kW/K, and the slope is 3.589 % of the conductance per hour.

Table 15.3: The MAP state at window 73 with the per-window entries omitted; the prior center is the survey's for the fractions.

Entry [unit]	$\hat{z}_i$	$\sqrt{\Sigma_{ii}}$	Prior center	Entry [unit]	$\hat{z}_i$	$\sqrt{\Sigma_{ii}}$	Prior center
T_s [°C]	14.00011	0.01193	14.000	γ [per window]	0.00305	0.00022	0
a_A [kW]	347.73304	2.00339	none	β_{A1} [–]	0.18084	0.00327	0.150
b_A [kW per window]	0.27267	0.07369	0	β_{A2} [–]	0.25000	0.08000	0.250
a_B [kW]	292.59456	1.89203	none	β_{A3} [–]	0.40101	0.00328	0.350
b_B [kW per window]	0.18363	0.07364	0	β_{A4} [–]	0.30000	0.08000	0.300
g_{A1} [–]	−0.04582	0.00293	0	β_{A5} [–]	0.45000	0.08000	0.450
g_{A2} [–]	0.02140	0.00275	0	β_{B1} [–]	0.20000	0.08000	0.200
g_{A3} [–]	−0.07619	0.00301	0	β_{B2} [–]	0.34229	0.00389	0.300
g_{B1} [–]	0.03210	0.00312	0	β_{B3} [–]	0.50000	0.08000	0.500
g_{B2} [–]	−0.00245	0.00397	0	β_{B4} [–]	0.37839	0.00389	0.400
g_{B3} [–]	−0.04630	0.00334	0	β_{B5} [–]	0.35000	0.08000	0.350
				b_{Q_A} [kW]	4.46330	0.90100	0

The covariance. On $(a_B, b_B, g_{\mathrm{B2}}, g_{\mathrm{B1}}, g_{\mathrm{B3}}, \gamma, \beta_{\mathrm{B3}}, \beta_{\mathrm{B2}}, \beta_{\mathrm{B4}})$, the entries the answers use, the spreads are those of Table 15.3 and the correlations are

$$\begin{bmatrix} 1 & 0.00088 & 0.03128 & 0.04295 & 0.04735 & -0.01335 & 0 & -0.00726 & -0.0101 \\ 0.00088 & 1 & -0.05605 & 0.00151 & 0.00161 & -0.10183 & 0 & 0.00018 & 0.00014 \\ 0.03128 & -0.05605 & 1 & -0.10446 & -0.11577 & 0.55387 & 0 & 0.13169 & 0.12948 \\ 0.04295 & 0.00151 & -0.10446 & 1 & -0.13988 & 0.01767 & 0 & 0.16473 & 0.16216 \\ 0.04735 & 0.00161 & -0.11577 & -0.13988 & 1 & 0.01725 & 0 & 0.15309 & 0.15027 \\ -0.01335 & -0.10183 & 0.55387 & 0.01767 & 0.01725 & 1 & 0 & 0.00461 & 0.00478 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ -0.00726 & 0.00018 & 0.13169 & 0.16473 & 0.15309 & 0.00461 & 0 & 1 & 0.20069 \\ -0.0101 & 0.00014 & 0.12948 & 0.16216 & 0.15027 & 0.00478 & 0 & 0.20069 & 1 \end{bmatrix}.$$

The rows fix these signs. The coil's level before the onset and its slope correlate at $+0.554$, since they share a kink: a higher level needs a steeper ramp to reach the same recent readings. The load's slope and the coil's slope correlate at -0.102, since a warmer return can be explained by either. And β_{B3} is correlated with nothing.

The reading rows. The reading rows have $X^2 = 1129.78$ on 1168 rows with effective degrees of freedom 1105.15, an upper tail of 0.297: the fit leaves them within their noise.

The overlap with the one fit. Where the stacked state and the one-window fit of Chapter 12 overlap, $T_{\mathrm{B3},73}$ is 24.0747 ± 0.8293°C on the one fit and 24.1573 ± 0.8244 here, and β_{B3} is 0.5000 ± 0.0800 on both, the survey unmoved. The stack narrows what the readings see, $g_{\mathrm{B2},73}$ going from -0.14444 ± 0.02905 to -0.11818 ± 0.00715, β_{B2} from 0.30823 ± 0.03031 to 0.34229 ± 0.00389 and $Q_{B,73}$ from 296.01 ± 3.43 to 298.12 ± 1.92 kW, and it leaves what they do not see alone.

15.2.4 The refits the answers use

Each of these repeats equation (15.7) on a changed drawing. *Onset and wander:* the state at each node of the onset grid and of the wander's grid, the nodes weighted by evidence from the fits with no bias freed, the answers at the nodes from fits with it freed. *Decision windows:* the rows at T in place of 73, with the wander and the jitter unchanged and the onset refound on windows 25 to 49 by 2, the grid stopping at $T - 4$. *Priors:* the slope priors, the jitter and the bias on another meter, varied one at a time. *Fresh errors:* the same at each T, on 50 nights of readings with the true loads and set point and the errors redrawn. *Forms:* a coil curvature, $g_{\mathrm{B2},t} = g_{\mathrm{B2}} - \gamma u_t - c\,u_t^2/1000$ with $u_t = (t - \tau)_+$ and prior $c \sim \mathcal{N}(0,\ 0.05^2)$; a load curvature

on hall B, $q\,(t-\bar{t})^2/1000$ with prior $q \sim \mathcal{N}(0,\ 1)$. *Survey:* spreads from 0.03 to 0.12; and errors sharing an offset of spread 0.04 at the same total spread of 0.08. *The whole record:* the rows at $T = 144$, used only among the checks.

15.3 The answers (step 4)

Every answer is read from the state of Section 15.2 at window 73 unless it says otherwise. The reported value is computed from $(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$; normal draws through the plant and the exact posterior stand beside it, with the gap.

15.3.1 The first passage

Operator questions. If I do nothing, what is the probability of a violation before the morning shift? Does the path matter, or only the last window? The first asks for the event (15.4) at the horizon; the second asks whether a reading of window 144 alone would have done.

General form. A forecast beyond the record is the third line of the trend form, equation (2.21); a level reached by a time is its second line, and at one window and rack the event is the threshold form, equation (2.12).

Specialized form. The two questions are equations (2.21) and (2.12) with $\mathbf{z}$ the stacked state (15.2), f_t an inlet $T_{k,t}$ of the chain (15.1), the level $f^* = T_{\max}$ and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal of equation (15.7). The line of the second form is replaced by a union, because the wander makes the path other than a straight line:

$$\Pr(E_s \mid \mathbf{y}) = \int \mathbf{1}\Big[\max_{k,\ 73\le t\le s} T_{k,t}(\mathbf{z}) \ge T_{\max}\Big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}. \tag{15.8}$$

The indicator $\gamma > 0$ of the second line falls away with the straight line. The union is a violation, an inlet at or over the limit at some window of the horizon, so a state that already stands over at window 73 is in it whichever way its path then runs, where a crossing from below would leave it out. Two routes from $(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ apply. The closed form $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ is defined for one window and one rack. Rack B3 at the horizon on the trend path, with the delta form of equation (4.53) and the tail of equation (4.57), gives

$$\Pr(T_{\mathrm{B3},144} \ge T_{\max} \mid \mathbf{y}) \approx \Phi\left(-\frac{T_{\max} - T_{\mathrm{B3},144}(\hat{\mathbf{z}})}{\sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}}}\right), \tag{15.9}$$

with the gradient, by the chain and the trends (15.3),

$$\frac{\partial T_{\mathrm{B3},s}}{\partial \beta_{\mathrm{B3}}} = T_{r,B,s} - T_{c,B,s}, \quad \frac{\partial T_{\mathrm{B3},s}}{\partial b_B} = \frac{\partial T_{\mathrm{B3},s}}{\partial Q_B}\,(s-\bar{t}), \quad \frac{\partial T_{\mathrm{B3},s}}{\partial \gamma} = -\frac{\partial T_{\mathrm{B3},s}}{\partial g_{\mathrm{B2}}}\,(s-\tau). \tag{15.10}$$

The first entry is the rise, which by the air balance is fixed by the heat and the airflow. The second and third are the lever arms of the trend form, one measured from the record's center and one from the onset. The first passage has no such gradient, since it reads a maximum over ten racks in 72 windows, 720 entries. The state route still defines it: the map from the state to those 720 inlets, linearized at $\hat{\mathbf{z}}$, makes them jointly normal,

$$\widetilde{\mathbf{T}} = \mathbf{T}(\hat{\mathbf{z}}) + \mathbf{G}\,(\mathbf{z} - \hat{\mathbf{z}}) \sim \mathcal{N}(\mathbf{T}(\hat{\mathbf{z}}),\ \mathbf{G}\boldsymbol{\Sigma}\mathbf{G}^{\mathsf{T}}), \qquad \Pr(E_{144} \mid \mathbf{y}) \approx 1 - \Pr(\widetilde{\mathbf{T}} < T_{\max}\,\mathbf{1}), \tag{15.11}$$

with $\mathbf{G}$ the stacked gradients of the inlets, and the event is the complement of an orthant. This integral has no closed form; the state route here means the orthant, integrated by draws of the

state through the linear map, with their Monte Carlo error. That is the reported value. A lower bound is the largest single marginal, an upper bound the sum of the marginals (Bonferroni).

Two choices enter here. The onset and the wander's parameters are pinned in the reported value, so the headline leaves their uncertainty out; the integrated route stands beside it to say by how much. And with the wander removed, a monotone trend path is over at some window exactly if it is over at the last, which draws confirm below.

Building blocks. $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$ (Section 15.2.3); the gradient (15.10); the map $\mathbf{G}$ of equation (15.11); the exact draws; the refits at other onsets and wander parameters (Section 15.2.4).

The path of the inlets. At the fit the inlets rise in every window: rack B3's smallest step on the trend path is +0.02038 K. Figure 15.2 sets the screen's view against it: the horizontal axis is the window, the vertical axis the temperature, the points the readings of the four read inlets through window 73, every third one, and the lines the straight lines through windows 1 to 73, with B2's last-hour line dashed, against the 27°C limit. Every line ends well below the limit. Figure 15.3 then draws the forecast. Its left panel has rack B3's trend path with its 95 % band, the limit, a bar at the horizon for the band with the wander added, and the stated truth dotted, a check marker for Section 15.4.4. Its right panel has racks B2 and B4, with their readings after window 73 as hollow points: the fit has not read them. The upper edge of B3's band passes the limit between windows 121 and 133, and at the horizon its spread is 0.910 K against 0.263 and 0.272 K at B2 and B4.

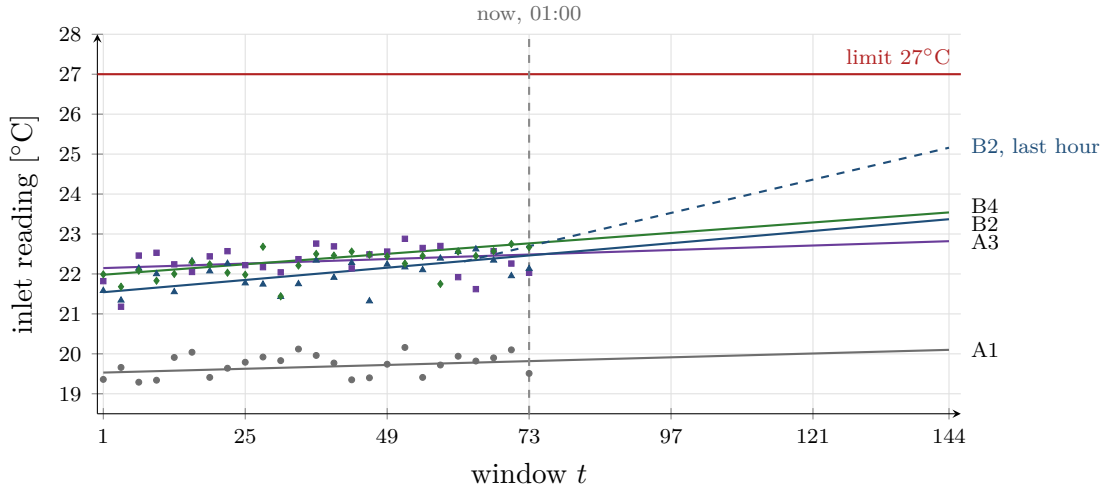


Figure 15.2: The screen's view at window 73: the four read inlets, every third reading, with straight lines through windows 1 to 73 carried to 06:55, and B2's line through its last twelve windows dashed. No rack without a sensor appears.

The end point. By equation (15.9), $T_{\mathrm{B3},144} = 25.6596$°C, $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} = 0.82859\,\mathrm{K}^2$, a spread of 0.9103 K, a margin of 1.3404 K, $z = 1.4726$ and

$$\Pr(T_{\mathrm{B3},144} \geq 27°\mathrm{C} \mid \mathbf{y}) = 0.07043.$$

The gradient on the shared entries is +10.7414 K on β_{B3}, +3.8715 K per kW per window on b_B, +491.52 K per unit slope on γ, +0.0362 K/kW on a_B, and −6.51, −4.51, −6.02 K on g_{B1}, g_{B2}, g_{B3}. Its lever arms are visible: $\partial T/\partial Q_B = 0.03618$ K/kW across 107 windows from the center, and 4.5094 K across 109 windows from the onset.

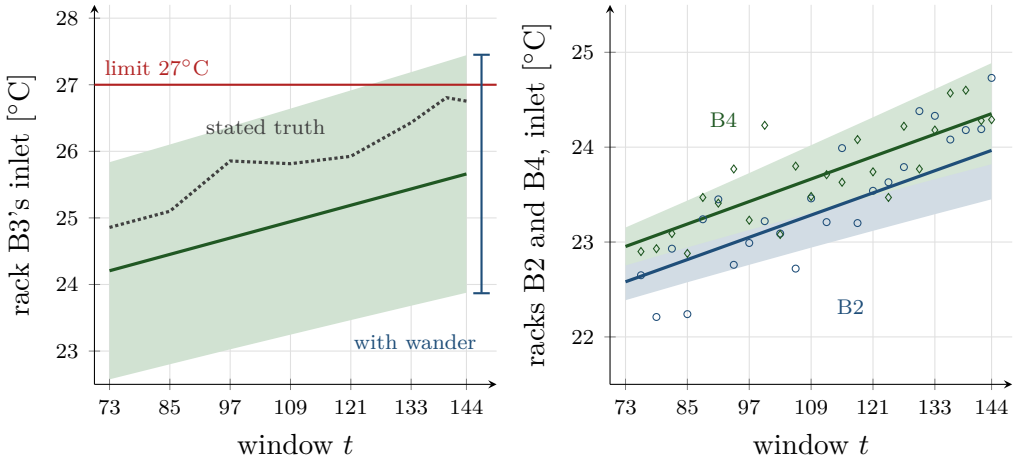


Figure 15.3: The forecast at window 73 on the trend path, centers and 95 % bands at every twelfth window. Left: rack B3, unread, against the limit; the bar at the right is the band at the horizon with the wander added; dotted, the stated truth at every twelfth window and at its largest, window 140, a check marker. Right: racks B2 and B4, with every third held-out reading as a hollow point; the readings include sensor error and wander, which the bands do not.

The orthant. With the wander removed, the linearized draws put any of the ten racks at the end point at 0.07065 (0.00057) and at any window at 0.07065: the same draws, gap 0.00000. On the trend path the first passage is its end point. The closed form above is one rack's: B3 alone on these draws is 0.06957 (0.00057), and ten racks' closed form, their marginals taken as independent, is 0.07153. With the wander the first passage is not its end point. Table 15.4 sets the routes beside each other, each row a route and the columns its value, standard error and gap from the reported value, and Figure 15.4 draws the same numbers as dots on one axis, with the part of the event that B3's end point misses marked. The reported value is

$$\Pr(E_{144} \mid \mathbf{y}) = 0.08152 \quad \text{(Monte Carlo se 0.00061)}.$$

The bounds are 0.07297, B3's marginal at the horizon, and 1.3382, the sum of the marginals, which says nothing.

Table 15.4: The routes to $\Pr(E_{144})$ at window 73. The closed forms are one rack at one window. Draws number 200 000 on the first three first-passage routes; the integrated routes use 20 000 or 10 000 per node, with independent seeds, and cover 0.9928 and 0.9168 of the weight.

Route, event	Value [–]	se [–]	− reported [–]
closed form, trend path, B3 end point, no wander	0.07043	–	−0.01109
closed form, B3 marginal at window 144 with wander, a lower bound	0.07297	–	−0.00855
first passage, linearized joint on $(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$, reported	**0.08152**	0.00061	–
first passage, normal draws, the plant solved on each	0.08421	0.00062	+0.00269
first passage, exact posterior by importance sampling	0.08486	0.00062	+0.00334
first passage, onset integrated, linearized	0.09169	0.00066	+0.01017
first passage, onset and wander integrated, normal draws	0.09566	0.00038	+0.01414

Like with like, at the end point with wander the closed form gives 0.07297, the linearized draws 0.07229 (0.00058), normal draws 0.07372 (0.00058) and the exact posterior 0.07458 (0.00059). By its parts on the linearized route, B3 is over at the end, 0.07229; or B3 is over and back under by the end, 0.00771; or B3 is never over and another rack is, 0.00153. So the

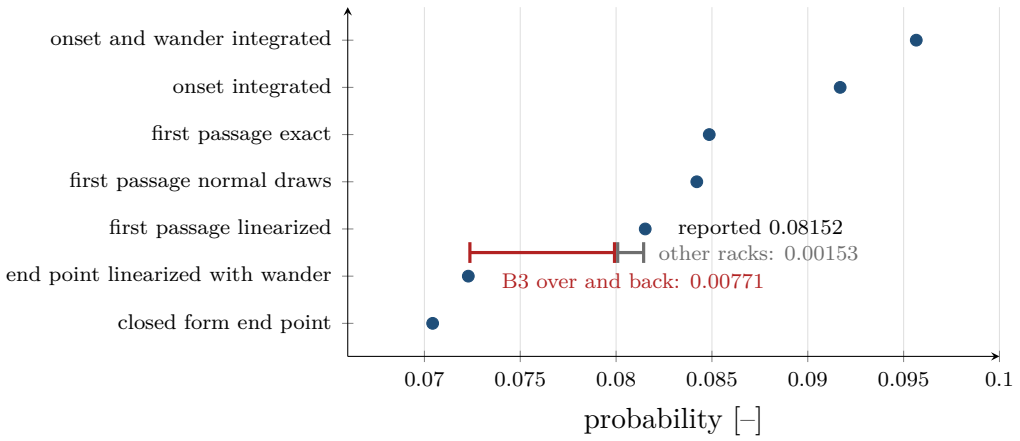


Figure 15.4: The routes of Table 15.4 and the linearized end point with wander on one axis. Red: the part of the first passage that is B3's path, over and back under by the horizon. Gray: the part on racks other than B3. The axis does not start at zero; it compares routes, not sizes.

ratio of first passage to B3's end point, 1.128, is two things: B3's path adds about 11 % (1.107 for B3 alone, 1.108 with all racks on both sides), and the other racks about 2 % (1.018). The largest inlet over the horizon has a linearized median of 25.780°C and band [24.346, 27.515], and an exact median of 25.772°C and band [24.350, 27.543]. Given a passage, its window has median 131, 05:50, and band [89, 144].

The onset and the wander left out. The reported value pins $\tau = 35$, $\phi = 0.90$ and $s = 4\,\text{kW}$, the modes of their evidence. Integrating the onset on the linearized route gives 0.09169 (0.00066), 0.01017 above the reported value (gap's se 0.00090), and integrating the wander alone gives 0.08395 (0.00072), 0.00243 above it. Integrating both on normal draws gives 0.09566 (0.00038), whose like is the pinned 0.08421 on the same route: 0.01145 more (gap's se 0.00073). That run drops 8.3 % of the weight, 1.3 % of it onset nodes above the mode, whose answers are larger, so its value is probably low. The onset, not the route, is the larger gap: pinning it leaves out about 0.010 and linearizing 0.003.

15.3.2 The horizon

Operator question. How does it grow with the horizon? The answer at 07:00 hides when the risk arrives.

General form. Each horizon is one value of the second line of the trend form, equation (2.21).

Specialized form. The question is equation (2.21), line 2, at each s, with its line replaced by the union (15.8). The conditional hazard between horizons $s_1 < s_2$, per hour, is

$$h(s_1, s_2) = \frac{\Pr(E_{s_2} \mid \mathbf{y}) - \Pr(E_{s_1} \mid \mathbf{y})}{1 - \Pr(E_{s_1} \mid \mathbf{y})} \cdot \frac{12}{s_2 - s_1}. \tag{15.12}$$

It is not a failure rate: nothing has failed, and it describes one night's look-ahead on one posterior.

Building blocks. The orthant (15.11) truncated at s, and the end-point form (15.9) at s.

The growth. Table 15.5 gives the routes side by side, each row a horizon. Figure 15.5 draws them: on the left the probability by horizon, on the right the conditional hazard by hour. The curves leave zero slowly and part only slightly. The hazard is 0.00113, 0.00283, 0.00673, 0.01296, 0.02394 and 0.03934 per hour, the last on the 55 minutes from 06:00 to 06:55, growing by a factor of 2.518 in the first step and 1.643 in the last. A slowly rising mean, 0.2448 K per hour at the end, is climbing the tail of a normal whose spread grows only from 0.832 to 0.910 K. Half the night's probability is reached by window 130, 05:45. Over now, read at the window's own state of Section 15.2.3, is 2.82×10^{-4}; the states over at window 73, which the union counts and a crossing from below would not, carry 0.00031 on the linearized joint against the horizon's 0.08152.

Table 15.5: The probability by horizon at window 73: end point on the trend path by the closed form, first passage on the linearized joint, reported, and on the exact posterior; Monte Carlo standard errors in parentheses.

s (clock)	End, no wander [–]	First passage, linear [–]	First passage, exact [–]
85 (02:00)	0.00124	0.00144 (0.00008)	0.00169 (0.00009)
97 (03:00)	0.00349	0.00426 (0.00015)	0.00491 (0.00016)
109 (04:00)	0.00878	0.01097 (0.00023)	0.01190 (0.00024)
121 (05:00)	0.01980	0.02379 (0.00034)	0.02556 (0.00035)
133 (06:00)	0.04019	0.04716 (0.00047)	0.05027 (0.00049)
144 (06:55)	0.07043	0.08152 (0.00061)	0.08486 (0.00062)

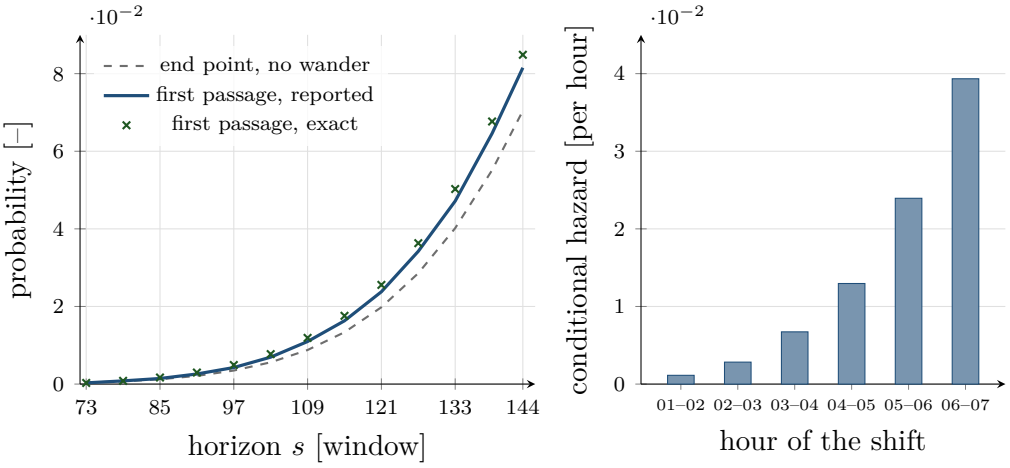


Figure 15.5: Left: the probability by horizon by the routes of Table 15.5, every sixth window. Right: the conditional hazard (15.12) per hour on the reported route; the last bar covers 06:00 to 06:55.

15.3.3 The rack that carries it

Operator question. Which rack carries it? The screen's four racks, or the six it does not show.

General form. Each rack's share is a tail of its own maximum, equation (2.12), and the budget is a variance split by entry, as in equation (2.10).

Specialized form. The question is equation (2.12) with $f(\mathbf{z}) = \max_t T_{k,t}(\mathbf{z})$ and $f^* = T_{\max}$, integrated as in the orthant (15.11). The budget is equation (5.14) on the gradient (15.10),

so a cross term $2[\mathbf{g}]_i[\mathbf{g}]_j\Sigma_{ij}$ takes the sign of the two gradients times that of the covariance. The answer as a function of the unread number fixes β_{B3} and leaves the other entries on the posterior,

$$\Pr(E \mid \beta_{\mathrm{B3}}, \mathbf{y}) = \int \mathbf{1}[\,E\,]\, p(\mathbf{z}_{\setminus\beta_{\mathrm{B3}}} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}_{\setminus\beta_{\mathrm{B3}}}, \tag{15.13}$$

which needs no new fit, because β_{B3} is uncorrelated with every other entry.

Building blocks. The gradient (15.10); $\mathbf{\Sigma}$ (Section 15.2.3); the draws of equation (15.11); the survey's density.

The racks. Table 15.6 gives each rack's fraction and probability, each row a rack. Rack B3 carries 0.0800 (exact 0.0834), A5 9.35×10^{-4} and B5 7.50×10^{-4}; every other rack is zero in the 200 000 draws. The measured B2 and B4 stand 10.4 and 8.8 spreads clear at the horizon. Given the event, B3 is over with probability 0.9813 and B3 is alone over with 0.9794.

Table 15.6: Which rack carries the first passage at window 73: fraction with its spread, and the probability that the rack is over at some window to the horizon.

Rack	Survey [–]	$\hat{\beta}_k$ [–]	sd [–]	Linear [–]	Exact [–]	Read
A1	0.150	0.1808	0.0033	0	0	yes
A2	0.250	0.2500	0.0800	0	0	no
A3	0.350	0.4010	0.0033	0	0	yes
A4	0.300	0.3000	0.0800	0	0	no
A5	0.450	0.4500	0.0800	9.35×10^{-4}	9.47×10^{-4}	no
B1	0.200	0.2000	0.0800	0	0	no
B2	0.300	0.3423	0.0039	0	0	yes
B3	0.500	0.5000	0.0800	0.0800	0.0834	no
B4	0.400	0.3784	0.0039	0	0	yes
B5	0.350	0.3500	0.0800	7.50×10^{-4}	7.40×10^{-4}	no

The budget. Figure 15.6 draws the variance of $T_{\mathrm{B3},144}$ on the trend path as bars, the horizontal axis in percent. Of $0.82859\,\mathrm{K}^2$, β_{B3} gives 0.738422, 89.12 %; b_B gives 0.081279, 9.81 %; γ gives 0.012197, 1.47 %; and a_B gives 0.004686, 0.57 %. The cross terms sum to −0.009273, mostly b_B with γ, −0.006412: their correlation is negative and both gradients positive, so a hotter load read and a worse coil read partly exclude each other. The term of g_{B2} with γ, −0.002189, has a negative gradient and a positive gradient against a positive correlation.

The answer as a function of the unread number. Figure 15.7 reads the answer off the entry that carries it: the horizontal axis is β_{B3}, the blue curve the conditional probability (15.13), and the green area behind it the survey's density scaled to a peak of one. The dotted line is the stated truth, a check marker. The probability is 0.00144 up to $\beta_{\mathrm{B3}} = 0.48$ and 0.00146 at 0.50, a floor set by the other racks, A5 and B5. It is 0.0315 at 0.56, 0.1114 at 0.58, 0.2762 at 0.60, 0.5098 at 0.62 and 0.8887 at 0.66. It stands at one half at $\beta_{\mathrm{B3}} = 0.6192$, 1.49 prior spreads above the survey, with $\Pr(\beta_{\mathrm{B3}} \geq 0.6192) = 0.0682$; with no wander the level is 0.6248. The integral of the curve against the density on its grid of 0.02, from 0.40 to 0.72, is 0.08064; the tails outside add 0.00015 and at most 0.00298, so it is at most 0.08378, which holds the reported 0.08152. The chapter's question is, nearly, whether B3's survey is 0.12 low.

What the record cannot see. That. The four read racks' survey errors are +0.39, +0.64, +0.53 and −0.27 prior spreads, an rms of 0.476. They are the only evidence about the survey.

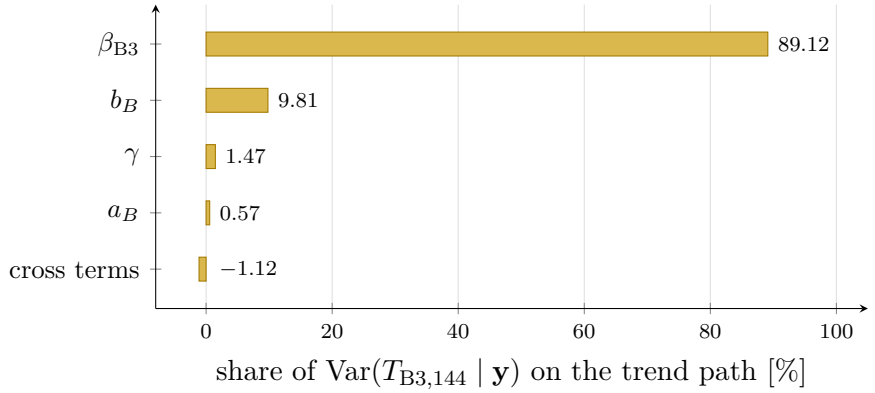


Figure 15.6: The budget of rack B3's inlet at the horizon, equation (5.14) on the gradient (15.10). The coils' own levels and the set point's level, each below 0.06 %, are not drawn. The largest term has no reading row.

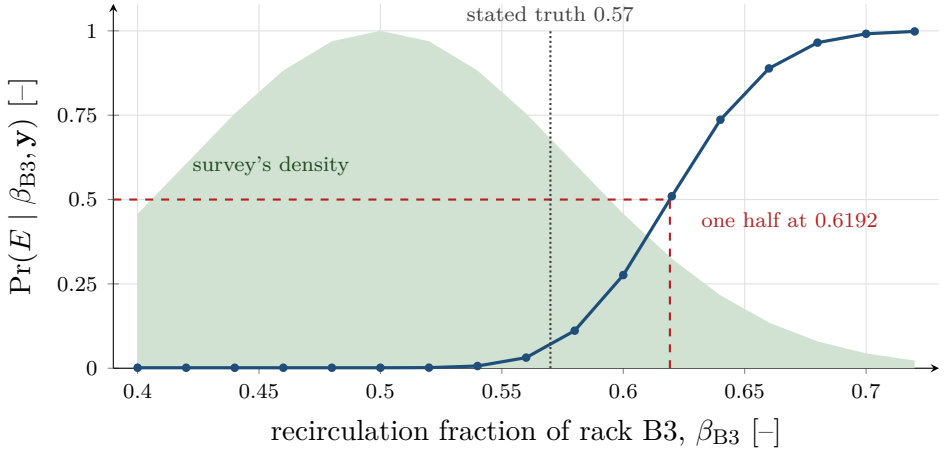


Figure 15.7: The first passage as a function of rack B3's fraction, the other entries on the posterior, equation (15.13), in front of the survey's density scaled to a peak of one. Dotted: the stated truth (Section 15.4.4).

Given the survey's spread, independent errors make them irrelevant to B3. They do bear on the spread: 0.04 scores +1.41 log units against 0.08, not decisive, and at 0.04 the answer is 0.008. The stated 0.08 is kept.

15.3.4 Epistemic and aleatory

Operator question. How much is what we do not know, how much is what has not happened yet? The first can be read now; the second cannot be read by anyone at 01:00.

General form. A variance of a derived quantity, equation (2.10).

Specialized form. The question is equation (2.10) with $f(\mathbf{z}) = T_{\mathrm{B3},144}(\mathbf{z})$, its variance split in two by the law of total variance of Chapter 3: the epistemic part, $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ on the trend path, entries of the state read from the record and the survey; and the aleatory part, the future's innovations of the wander and the window's jitter. Since each inlet is fixed by its window's entries, three events can be read off one state: the wander alone with the entries at the fit, the entries' uncertainty with no wander, and both. The effective limit is the level L with $\Pr(T^{\mathrm{trend}}_{\mathrm{B3},144} \geq L \mid \mathbf{y}) = \Pr(E \mid \mathbf{y})$.

Building blocks. The gradient (15.10), the window's own gradient on Q_B and T_s, and the draws of Section 15.3.1.

The split. The trend's part is $0.82859\,\mathrm{K}^2$, 97.49 %; the wander's is $0.02130\,\mathrm{K}^2$, 2.51 %, a spread of 0.146 K, that is $\partial T/\partial Q_B = 0.03618\,\mathrm{K/kW}$ times 4.0 kW, with 0.0198 K from the jitter. The epistemic band of the trend at the horizon runs from 23.875 to 27.444°C, and without β_{B3} its spread is 0.3003 K. The three numbers from one state are:

- wander alone, entries at the fit: none in 50 000 nights;
- entries' uncertainty alone: 0.07135, on the draws of the next item;
- both: 0.08152.

The wander cannot reach the limit from the fit, whose B3 ends at 25.66°C; all but 4.5 % of what it adds lies in draws with β_{B3} above about 0.55. On the same draws it adds 0.01017 (0.00030), equal, on those draws, to a loss of 0.0636 K of effective margin: the no-wander maximum exceeds 26.9364°C with the event's probability. The largest inlet exceeds B3's end point by a mean of 0.1519 K, and by 0.0461 K in the draws within 0.5 K of the limit. The window as grain matters less: on one set of 100 000 draws, every window gives 0.08063, and reading only every second, sixth or twelfth window counted back from the horizon gives 0.07864, 0.07529 and 0.07375.

15.3.5 From each decision window

Operator question. What would the same question have said earlier and later? An answer given once at 01:00 is one record's; the operator wants to know whether the question was settled earlier, and whether it would be settled by staying.

General form. Each decision window is the trend form at its own record, equation (2.21).

Specialized form. The question is equation (2.21) at each T, through the union (15.8) on windows T to 144, with $\mathbf{z}$ the stacked state of that record, the onset refound and $p(\mathbf{z} \mid \mathbf{y}_{1:T})$ the Laplace normal of its fit (Section 15.2.4). These answers use 30 000 draws; the row $T = 73$ thus prints 0.08077 for the reported 0.08152, inside its Monte Carlo error of 0.0016. A record's value is one night's draw of the sensor errors, and its spread is shown over 50 nights with the true

loads and set point held and the errors redrawn; each answer on those nights uses 4000 draws. That spread is not the whole night-to-night spread, since the true wander is held fixed.

Building blocks. The refits of Section 15.2.4, the orthant (15.11) at each T.

The answers. Table 15.7 gives them, each row a decision window: the onset, the slope's z, the forecast, the record's probability on the linearized and normal routes, and the fresh nights' median, spread and band. Figure 15.8 draws them, the horizontal axis the decision window, the vertical axis the probability, with the fresh nights' band shaded.

Table 15.7: The do-nothing answer from each decision window. Record: this night's sensor errors. Fresh: 50 nights with the errors redrawn and the true loads and set point held; a check, not available to the operator at the window.

T (clock)	Onset [win.]	z_γ [–]	$T_{\mathrm{B3},144}$ (sd) [°C]	Linear [–]	Normal [–]	Fresh median [–]	Fresh sd [–]	Fresh 2.5–97.5 % [–]
49 (23:00)	45	3.83	29.946 (1.333)	0.99317	0.98253	0.26687	0.30229	0.06566–0.96016
61 (00:00)	35	7.60	25.633 (0.953)	0.08787	0.09313	0.15325	0.11772	0.05779–0.48343
73 (01:00)	35	13.55	25.660 (0.910)	0.08077	0.08603	0.12550	0.05766	0.06077–0.27656
85 (02:00)	35	18.16	25.433 (0.880)	0.04560	0.04933	0.07137	0.01478	0.05123–0.10327
97 (03:00)	35	24.69	25.843 (0.897)	0.11227	0.11733	0.14637	0.02138	0.11192–0.18447
109 (04:00)	37	32.28	25.825 (0.884)	0.10397	0.10883	0.12062	0.01448	0.09742–0.14637
121 (05:00)	39	40.66	25.827 (0.875)	0.09960	0.10377	0.10075	0.01018	0.08464–0.12219
133 (06:00)	39	48.67	25.839 (0.872)	0.10250	0.10567	0.10437	0.01238	0.08756–0.12899



Figure 15.8: The answer from each decision window: this record's on the linearized route, and the median and central 95 % band over fresh sensor errors with the true wander held fixed.

Three things stand out in the table. *At 23:00 the record says* 0.99. The onset of that row, window 45, is the top of its grid, which stops at $T - 4$. The evidence has two modes, near windows 39 and 46, and the pin takes the later: $\gamma = 0.01653$, a forecast of 29.95°C. Onset 39 scores 0.79 log units below it and gives 0.43150. Integrating the onset gives 0.57218 (0.00124) on the grid and 0.66968 (0.00067) on every window from 25 to 48. Over fresh errors the same moment gives 0.07 to 0.96, and the record's 0.99317 lies outside its own fresh band, whose top is 0.96016. A forecast this early is one night's draw. *The answer does not resolve as the morning approaches.* From 03:00 to 06:00 the coil's slope firms from $z = 24.69$ to $z = 48.67$ and the fresh nights' spread falls from 0.05766 at 01:00 to 0.01238 at 06:00. It stops moving but settles near one in ten, not near 0 or 1. The racks that carry the event have no reading, so a passage would

not be seen and "not yet" cannot be observed; the readings move the answer only through the load and the coil. *The fall at 02:00 and the rise at 03:00 are the load, not the sensors.* Both the record and the fresh nights fall at window 85 and rise at 97, and they share the true loads. The fits' b_B runs 0.1836, 0.1512 and 0.2435 kW per window at windows 73, 85 and 97, and their $Q_{B,144}$ runs 312.24, 308.27 and 319.01 kW; Section 15.4.4 gives the cause. This night's 0.08077 at 01:00 lies inside the fresh nights' band, below their median of 0.12550. The half-probability window is one night's draw too: over the fresh nights at 01:00 it has median 130, spread 1.0 and band [128.2, 132.0].

15.4 The checks (step 5)

A check compares an answer with the calculation it stands for.

15.4.1 The held-out test

All 16 readings of windows 74 to 144 are predicted from window 73, the load's wander included, and scored against the joint predictive: $\chi^2 = 1148.46$ on 1136 rows, a tail of 0.392; 96.4 % lie within 1.96; the largest $|z|$ is 3.09, B2's inlet at window 81. Each hour's block has a tail between 0.180 and 0.867. Table 15.8 gives the test by reading: each row a reading, the columns its mean z, its largest $|z|$ and its mean z over the last twelve windows. At window 144 the forecasts against the readings are $T_{r,B}$ 31.030 (0.565) against 31.39°C; $T_{d,\mathrm{B2}}$ 22.498 (0.442) against 22.81°C; Q_B 312.24 (9.88) against 321.15 kW; and T_{inB4} 24.353 (0.425) against 24.29°C.

Table 15.8: The held-out test from window 73, windows 74 to 144, by reading.

Reading	Mean z [–]	Largest $\lvert z\rvert$ [–]	Last 12 [–]	Reading	Mean z [–]	Largest $\lvert z\rvert$ [–]	Last 12 [–]
T_s	+0.108	2.31	+0.109	$T_{d,\mathrm{B3}}$	−0.093	2.42	−0.100
$T_{r,A}$	−0.417	2.22	−0.568	Q_A	−0.681	2.28	−0.734
$T_{r,B}$	+0.271	1.88	+0.675	Q_B	+0.306	2.34	+0.530
$T_{d,\mathrm{A1}}$	−0.004	2.40	+0.215	Q_{loop}	−0.234	1.82	−0.305
$T_{d,\mathrm{A2}}$	−0.135	2.60	+0.218	inlet A1	−0.097	2.06	−0.234
$T_{d,\mathrm{A3}}$	−0.119	2.54	−0.150	inlet A3	−0.168	2.84	−0.024
$T_{d,\mathrm{B1}}$	+0.200	2.54	+0.240	inlet B2	+0.219	3.09	+0.748
$T_{d,\mathrm{B2}}$	+0.463	2.52	+0.583	inlet B4	+0.251	2.09	+0.711

The fit of the whole record, on the same drawing, moves the slope from 0.00305 to $\gamma = 0.00326$, +0.97 spreads of window 73's, and b_B from 0.18363 to 0.21366, +0.41. The onset's mode on a three-window grid becomes 37. On the whole record a ramp curvature scores +0.02 log units and a load curvature on hall B −0.45.

What the held-out test cannot test. No reading of windows 74 to 144 has a column on β_{B3}, β_{B5} or β_{A5}: the largest entry is 0. The test also scored the wander, so it passes on 13.12 % of $T_{\mathrm{B3},144}$'s variance with the wander, and the rest cannot be tested by any of the plant's sixteen readings, in the night or after it. A share of variance is not a share of probability: with β_{B3} at its survey the event is 0.00146. The test supports the trend and says nothing of the answer's main term.

15.4.2 The limit of exact rows and the curvature term

The fit runs on the reduced state. The check fits the full state of windows 1 to 73, with 20 temperature entries per window held by physical rows of tolerance σ_p. Table 15.9 gives the result: each row a tolerance, the columns the slope, rack B3's forecast, its spread and the

reading rows' X^2. At $\sigma_p = 0.001$ the forecast agrees with the reduced state to the fifth decimal and X^2 within 0.02. The covariance is the Gauss–Newton form; the objective's full Hessian at $\hat{\mathbf{z}}$ gives sd(γ) a gap of -1.09×10^{-7} and sd($T_{\mathrm{B3},144}$) = 0.910267 against 0.910269 K, and the end-point probability is 0.07043 on both.

Table 15.9: The limit of exact physical rows on the full state of windows 1 to 73, against the reduced state.

Tolerance σ_p [kW or K]	γ [per window]	$T_{\mathrm{B3},144}$ [°C]	sd [K]	X^2 [–]
0.1	0.003053	25.66908	0.91114	966.287
0.01	0.003046	25.65967	0.91028	1127.764
0.001	0.003046	25.65955	0.91027	1129.763
reduced state	0.003046	25.65955	0.91027	1129.783

15.4.3 Sensitivity

Table 15.10 varies the choices of Section 15.1.4, the fit repeated each time: each row a change, the columns rack B3's forecast, its spread, the end point with wander, the first passage and the log evidence against the reported drawing. Each uses 30 000 linearized draws, with a standard error of about 0.0016 at 0.08. Freeing the hall-A bias moves the first passage by less than the draws resolve. The shared offset of the survey and a curvature on the coil stand as sensitivities and neither is adopted. A jitter of 0.05 K scores +2.50 log units against the excess variance's 0.0198 K, and moves the answer within 0.008.

Table 15.10: Sensitivity at window 73, 30 000 linearized draws each, se about 0.0016 at 0.08.

Change	$T_{\mathrm{B3},144}$ [°C]	sd [K]	End point, wander [–]	Pr(E) [–]	Log evidence [–]
baseline, this seed	25.6596	0.9103	0.07297	0.08287	–
wander sd 2 kW	25.6637	0.8774	0.06456	0.06767	−8.667
wander sd 6 kW	25.6548	0.9617	0.08625	0.10650	−3.225
persistence $\phi = 0.0$	25.6599	0.8698	0.06431	0.08763	−24.314
persistence $\phi = 0.5$	25.6635	0.8774	0.06647	0.08630	−9.062
persistence $\phi = 0.8$	25.6628	0.8947	0.07009	0.08427	−0.849
persistence $\phi = 0.95$	25.6568	0.9199	0.07430	0.08083	−2.442
set-point jitter sd 0.05 K	25.6819	0.9112	0.07687	0.09007	+2.500
no bias freed	25.6548	0.9089	0.07196	0.08223	−10.563
bias on hall B's meter	25.6505	0.9076	0.07105	0.08093	−11.936
bias on the loop meter	25.6581	0.9099	0.07267	0.08287	−8.641
coil slope prior 0.005	25.6578	0.9103	0.07270	0.08263	+1.212
coil slope prior 0.2	25.6597	0.9103	0.07299	0.08290	−2.291
load slope prior 0.1 kW per window	25.4207	0.8781	0.03802	0.04463	+0.683
load slope prior 10 kW per window	25.6632	0.9108	0.07363	0.08350	−4.546
onset at 31	25.5249	0.9087	0.05444	0.06307	−1.833
onset at 37	25.7345	0.9113	0.08519	0.09567	−0.030
onset at 43	26.0062	0.9158	0.14211	0.15490	−3.085
onset at 49	26.3869	0.9232	0.25623	0.27523	−10.160
survey sd 0.03 on every rack	25.6602	0.4405	0.00194	0.00220	+1.158
survey sd 0.04 on every rack	25.6599	0.5242	0.00689	0.00827	+1.407
survey sd 0.05 on every rack	25.6597	0.6153	0.01703	0.02080	+1.168
survey sd 0.06 on every rack	25.6596	0.7110	0.03240	0.03800	+0.795
survey sd 0.12 on every rack	25.6595	1.3235	0.15702	0.19300	−1.367
survey errors sharing an offset, sd 0.04	25.8169	0.8518	0.08558	0.09500	+0.156
coil ramp with curvature, prior 0.05	26.1363	1.2180	0.24080	0.25600	−0.646
hall B load with curvature, prior 1	25.5953	0.9802	0.07818	0.08723	−0.030

What overturns the screen, and what overturns the answer. The screen's "nothing reaches 27°C before 07:00" is overturned by a rack it does not read. The largest consideration in the chapter is the survey's spread: the answer of one in twelve falls a factor of ten at 0.04, which the read racks prefer by 1.41 log units, not decisively, and rises more than a factor of two at 0.12. It is moved by a reading at B3 (Section 15.6.2); by the onset, inside its band and past it (Table 15.10); by a tight prior on the load's slope; and by a curvature of the coil, which the record neither supports nor rules out at 01:00 and which the whole record does not support. It is not overturned by the wander's size or persistence, nor by any meter's bias.

15.4.4 The answers against the stated truth

The record is generated from a stated truth (Section 15.6.4), and Table 15.11 collects the checks: each row a route or an entry, the columns the reported value, the value beside it and the gap, the second minus the first, or z against the reported spread. *No rack reaches* 27°C *in any window.* The largest true inlet is B3, 26.8058°C at window 140, 06:35, a margin of 0.1942 K. Fixing β_{B3} at its truth, the rest on the posterior, gives 0.06084; with every truth entry and fresh wander of the truth's size, 0.00110. The no-violation outcome is consistent with the answer, and one night cannot test a probability of one in twelve. The truth's wander in hall B has a sample spread of 3.663 kW and lag-one correlation 0.854, against the stated 3.0 kW and 0.8; at the stated wander the first passage is 0.07500. That wander is the cause of the fall and the rise of Section 15.3.5: hall B's load stands a mean of 5.31 kW below trend over windows 74 to 85 and 4.55 kW above it over windows 86 to 97, 10.58 kW above at window 96.

Table 15.11: The checks and their gaps at window 73. The truth is the stated truth of the constructed record, and z is against the reported spread.

Route or entry [–]	Reported	Beside it	Gap
Routes, first passage to window 144			
Normal draws through the plant	0.08152	0.08421 (0.00062)	+0.00269
Exact posterior	0.08152	0.08486 (0.00062)	+0.00334
Onset integrated, linearized	0.08152	0.09169 (0.00066)	+0.01017
Onset and wander integrated, normal draws, against the normal route	0.08421	0.09566 (0.00038)	+0.01145
Median of the largest inlet, exact [°C]	25.780	25.772	−0.008
Approximations			
Full state, $\sigma_p = 0.001$, $T_{\mathrm{B3},144}$ [°C]	25.65955	25.65955	0.00000
Newton against Gauss–Newton, sd $T_{\mathrm{B3},144}$ [K]	0.910269	0.910267	-2.21×10^{-6}
Stated truth			
Largest inlet of the night [°C]	median 25.780	26.8058, B3 at 06:35	no passage
β_{B3}	0.5000	0.57	$z = +0.87$
γ [per window]	0.003046	0.003303	$z = +1.14$
Load slope [kW per window], hall A	0.2727	0.2098	$z = -0.85$
Load slope [kW per window], hall B	0.1836	0.2098	$z = +0.36$
Hall-A meter bias [kW]	4.463	4.0	spread 0.901
Onset [window]	35, band [30.3, 41.2]	36	inside the band
$T_{\mathrm{B3},73}$ [°C]	24.1573	24.8571	inside the band
$T_{\mathrm{B3},144}$ on the truth's load trend [°C]	25.6596	26.5934	$z = +1.03$

The exact gap of 0.00334 is about four combined standard errors, most of it the plant's curvature, since normal draws through the plant already give 0.00269 of it; the remainder is not resolved. No assurance level is set, so no gap is judged against one.

15.5 The report (step 6)

The screen's "nothing reaches the limit" reads the wrong racks. The probability that some inlet reaches 27°C by 07:00 is 0.082 with the onset and the wander pinned at their modes, which leaves out about 0.010, carried almost entirely by rack B3, which has no sensor. It rests on one survey: at a survey spread of 0.04, which the read racks mildly prefer, it is 0.008. It grows through the night, and waiting will not settle it, in the sense that no outcome of the next hour's readings takes it below 0.01 or above 0.5.

15.5.1 The answers and their bands

Table 15.12 gives the answers in the order of Table 15.1. Each row is a question; the second column gives the reported answer and the third what stands beside it.

Table 15.12: The report at window 73, 01:00. Limit 27°C and horizon 07:00, both the operator's. The last row reads the record at other windows and the constructed truth's fresh nights; it is a check and no part of the entry at 01:00.

Question	Answer, reported	Beside it
Probability of a violation before the morning shift?	0.08152, se 0.00061, first passage on the linearized joint	exact 0.08486; onset and wander integrated 0.09566 on normal draws, against 0.08421; at survey spread 0.04, 0.00827
Does the path matter?	1.128 times B3's end point: B3's path about 11 % (0.00771), other racks about 2 % (0.00153)	end point closed form 0.07043 with no wander, 0.07297 with it
How does it grow?	0.01097 by 04:00; half by 05:45; 0.03934 per hour in the last hour	exact 0.01190 by 04:00
Which rack carries it?	B3: over in 0.9813 of the event; 89.12 % of the variance is β_{B3}	one half at $\beta_{\mathrm{B3}} = 0.6192$, with probability 0.0682
What we do not know, what has not happened?	97.49 % and 2.51 % of the variance at the horizon	wander alone none in 50 000; it adds 0.01017 to the entries' 0.07135, same draws
Earlier and later?	0.99 at 23:00; about one in ten from 03:00	over fresh sensor errors 0.07 to 0.96 at 23:00, 0.06 to 0.28 at 01:00

15.5.2 The entry in the shift log

The operator's entry as of window 73, 01:00, using nothing later, reads: "Do-nothing, to 07:00. Probability that some inlet reaches 27°C at some window: 0.082 with the onset and wander pinned (exact 0.085; 0.096 with them integrated, on normal draws, against 0.084). Carried by rack B3, unread: forecast 25.66°C at 06:55, 95 % band on the trend path 23.88–27.44; 98 % of the probability. To 04:00 it is 0.011; half of it arrives after 05:45. The forecast rests on B2's coil worsening 3.6 % per hour from about 21:50 (band 21:25 to 22:25), hall B's load rising 2.2 kW per hour, and B3's survey (0.008 at a survey spread of 0.04). The next hour's readings move it within 0.04 to 0.14 and will not settle it. Action: not decided here."

Common mistake

Reading a quiet hour as good news. No reading on the plant sees rack B3, so an hour with no alarm says nothing about whether its inlet has passed the limit, and the probability moves only through what the readings say of the load and the coil, not because nothing was seen. A look-ahead probability is also not a forecast of a violation: a night with no passage does not show that one in twelve was wrong.

15.6 The next reading (step 7)

The report leaves a probability that nearly all rests on one unread fraction. A further reading is a further row, and what it would do to the entries is known before its value is.

15.6.1 The preposterior of a candidate

A candidate adds rows at the present MAP state, and the covariance after it is the preposterior of equation (4.65), in the linear Gaussian update of the state. Since a probability, unlike a variance, depends on the outcome, the distribution of the probability after each candidate is taken over 201 outcomes and 20 000 joint draws. Waiting twelve windows adds those windows' readings. A handheld probe of 0.3 K at rack B3, or at B5 or A5, adds one row now, on the window's own state. A re-survey of β_{B3} of spread 0.02 adds one prior row. A permanent inlet sensor at B3, the instrument at the place of the question, of 0.3 K adds its row in windows 74 to 85, alone or with the other readings of those windows. A second discharge sensor on B2 of 0.3 K and a water-side heat meter on B2 of 4 kW, against the two-unit hall's 3 kW since the larger plant's units carry a wider stated spread, add rows on the coil in the same windows.

15.6.2 The candidates

Table 15.13 compares them by what each does to the entries: each row a candidate, the columns the spreads of γ, b_B and β_{B3} and of the forecast, the share of outcomes that would leave the probability below 0.01 or above 0.5, reference levels and not an assurance, and the central 90 % of the probability after it. The expectation of the probability after a candidate should equal the probability now; it lies from 0.0827 to 0.0887 across them, against 0.0815. The update is exact on the linear joint; the excess is Monte Carlo, 0.0835 (0.0020) from the shared draws and up to 0.017 from the 201 outcomes. The "above 0.5" shares, each with a standard error near 0.02, are not separated by this run.

Table 15.13: Candidate readings from the posterior at window 73.

Candidate	sd γ [per win.]	sd b_B [kW/win.]	sd β_{B3} [–]	sd $T_{\mathrm{B3},144}$ [K]	below 0.01 [–]	above 0.5 [–]	5–95 % of Pr [–]
None, window 73	0.00022	0.07364	0.0800	0.9103	–	–	0.08
Wait twelve windows	0.00016	0.06083	0.0800	0.8908	0.000	0.000	0.04–0.14
Handheld probe at B3, now	0.00022	0.07363	0.0280	0.4134	0.632	0.080	0.00–0.57
Re-survey of β_{B3}, sd 0.02	0.00022	0.07364	0.0194	0.3655	0.677	0.085	0.00–0.65
Permanent B3 sensor, windows 74–85, alone	0.00022	0.07167	0.0136	0.2568	0.741	0.070	0.00–0.75
Permanent B3 sensor, with the wait	0.00016	0.06080	0.0089	0.2464	0.756	0.085	0.00–0.77
Handheld probe at B5 or A5	0.00022	0.07364	0.0800	0.9103	0.000	0.000	0.08
Second discharge sensor on B2, windows 74–85	0.00018	0.07157	0.0800	0.9007	0.000	0.000	0.05–0.13
Water-side heat meter on B2, windows 74–85	0.00021	0.06777	0.0800	0.9058	0.000	0.000	0.06–0.11

Waiting and better coil instruments move the probability, here within 0.04 to 0.14, but no outcome takes it below 0.01 or above 0.5, because none reads β_{B3}. A probe answers the question at the place it lives. Its predictive is 24.157 ± 0.877°C, and Table 15.14 gives the outcome of each reading, each row a value and the columns β_{B3} after it and the probability. The permanent sensor does most to the entries, and it alone would also make B3's first passage observable. The ranking and any action stay with the operator. Chapter 18 ranks readings against a named worry.

Table 15.14: The probability after a handheld probe at rack B3 at window 73, of 0.3 K; β_{B3} has spread 0.0280 after each. The last row is the constructed truth's inlet, with no sensor error, a check.

Probe reads [°C]	β_{B3} after [–]	Pr(E) after [–]
24.000	0.4866	0.0024
24.500	0.5293	0.0153
25.000	0.5720	0.1208
25.500	0.6147	0.4588
26.000	0.6574	0.8360
24.857	0.5598	0.0707

15.6.3 What the analysis showed

The screen had nothing now within four degrees of the limit and no line reaching it before the morning shift. The analysis puts the probability of a violation at about one in twelve, at a rack the screen does not show. With no wander a monotone trend's first passage is its end point; with the wander B3's path adds about 11 % and the other racks about 2 %, almost only in draws with β_{B3} above about 0.55. The value is 0.082 with the onset and the wander pinned, and the survey's spread moves it more than anything else. The load's trend and the coil's ramp are read by the record, and the rack that carries the answer is not. The onset, pinned, leaves out more than the linearization does. The answer stops moving without resolving: from 03:00 it stands near one in ten while the coil's slope doubles its z. And the held-out test passes on 13 % of the variance, so its passing is no evidence for the rest, which holds nearly all of the probability.

15.6.4 Limits of the model

The load's straight trend. The trend is assumed to go on to the horizon, and no row at 01:00 tests it. A batch that ends at 05:00 would lie outside every band above.

The wander. One persistent process with a shared ϕ and s cannot carry a wander of changing size or two time scales.

The pinned onset and wander. The reported value leaves out their spread, by the amount of Section 15.3.1.

The survey. Its errors are taken as independent and its spread as the survey's own 0.08. The read racks bear on that spread, mildly preferring 0.04, and at the rack that matters the record can check neither.

The grain. Windows are five minutes long, steady within, so a shorter excursion lies outside the event.

The condition. Every answer is conditional on the coil (Section 15.2.2, by the method of Chapter 10).

The fresh errors. They redraw one night's sensor errors with the true wander fixed, so they do not span the whole spread from one night to another.

Repeated looks. The table of decision windows reads the same record again and again, and those looks stand uncorrected.

Constructed readings. The record is constructed: generated from a stated truth of the plant with the stated sensor errors, and Section 15.4.4 compares the answers with that truth. It is not field data, and the agreement is a check of the method on its own model.

15.7 What to retain

Definition 15.1 (First passage on the linearized joint)**.** The first passage to a horizon s is the event that some inlet is at or over the limit at some window up to s, each window at its own state. Its reported probability is that of the complement of an orthant of the normal $\mathcal{N}(\mathbf{T}(\hat{\mathbf{z}}),\ \mathbf{G}\boldsymbol{\Sigma}\mathbf{G}^{\mathsf{T}})$, integrated by draws with their Monte Carlo error; it has no closed form. With the wander removed, a monotone trend path's first passage equals its end point, and one rack's end point has the closed form $\Phi(-(T_{\max} - T(\hat{\mathbf{z}}))/\sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}})$. Parameters pinned in the drawing, here the onset and the wander's, are integrated beside the reported value to say what the pin leaves out.

Three conclusions carry beyond the plant. A look-ahead probability can be large for a reason no future reading tests, and the budget says when. The windows ahead settle a question only through the entries they read, so waiting is itself a candidate reading, with its own preposterior, and here no outcome of it takes the probability below 0.01 or above 0.5. And an early forecast's confident answer is one night's draw until its spread over fresh sensor errors has been shown.

Table 15.15 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 15.15: The terms of the look-ahead and where each is defined. Units for the shift plant.

Term	Expression	Units	Meaning	Defined in
Stacked state	eq. (15.2)	mixed	windows through the horizon in one state	15.1.2
Trends and wander	eq. (15.3)	kW, –	load trend, persistence, coil's ramp	15.1.2
Event	E_s, eq. (15.4)	–	some rack over at some window	15.1.2
Sparsity pattern	eq. (15.5)	–	six fractions in no reading row	15.1.4
Posterior				
Rows	eq. (15.6)	–	readings, priors, process rows	15.2.1
Selected drawing	τ, ϕ, s, s_j, b_{Q_A}	mixed	found from the record by evidence	15.2.2
Questions and answers				
End point	eq. (15.9)	–	one rack, one window, closed form	15.3.1
Lever arms	eq. (15.10)	K per unit	from the center and from the onset	15.3.1
First passage	eq. (15.11)	–	an orthant by draws	15.3.1
Conditional hazard	eq. (15.12)	per hour	the next interval, given none yet	15.3.2
Conditional on the fraction	eq. (15.13)	–	the answer as a function of β_{B3}	15.3.3
Split	$\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ and the wander's part	K^2	epistemic and aleatory	15.3.4
Decision window	the answer at T, fresh errors	–	how far one record moves it	15.3.5
Held-out test	χ^2 on windows 74–144	–	tests what later readings read	15.4.1
Candidate	preposterior, eq. (4.65)	mixed	entries after a reading	15.6.1

15.8 Exercises

Exercise 15.1. *Application.* From $T_{\mathrm{B3},144} = 25.6596°\mathrm{C}$ and $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} = 0.82859\,\mathrm{K}^2$ compute the end point's normalized distance and its probability. (The margin is $27 - 25.6596 = 1.3404\,\mathrm{K}$ and the spread $0.9103\,\mathrm{K}$, so by equation (15.9) the probability is 0.07043 at $z = 1.4726$.)

Exercise 15.2. *Derivation.* Derive the budget's share of β_{B3}. (By the gradient (15.10) its entry is the rise, $10.7414\,\mathrm{K}$; with $\Sigma_{\beta\beta} = 0.08^2$ and no covariance with any other entry the term is $10.7414^2 \times 0.08^2 = 0.738422\,\mathrm{K}^2$, of the total 0.82859: 89.12 %.)

Exercise 15.3. *Application.* At what β_{B3} does rack B3's trend path, with no wander, reach 27°C at the horizon? (The path is linear in the fraction with slope $10.7414\,\mathrm{K}$; with the wander the conditional probability is one half a little lower, at 0.6192, and with none the level is $0.5 + 1.3404/10.7414 = 0.6248$.)

Exercise 15.4. *Derivation.* Explain why, with no wander, the first passage equals the end point on the draws. (A rising path is over at some window if and only if it is over at the last. At the fit rack B3's trend path rises in every window, with a smallest step of $+0.02038\,\mathrm{K}$, and the wander, the only thing that could bring a path back under, is removed. The draws test the argument across the posterior: over the ten racks the two events hold on the same draws, 0.07065 and 0.07065, gap 0.00000.)

Exercise 15.5. *Application.* From the conditional hazards of Section 15.3.2 find the ratio between the last two hours, and say what sets it. (The mean rises by 0.2448 K per hour at the end against a spread that grows only from 0.832 to 0.910 K, so the hazard climbs a normal tail, and the ratio falls as the tail is climbed: $0.03934/0.02394 = 1.643$.)

Part V

Validating the Model in Operation

Chapter 16

Goodness of Fit in Operation

16.1 A fit-quality light that stays green

It is 01:00, window 73 of the shift. Each of the sixteen readings on the screen sits close to the value the fit gives it, the farthest being the loop meter at 0.572 of its noise spread. The screen's fit-quality light is green, and it has been green all night. From the screen alone the operator concludes that the model describes the plant.

The light compares tonight's readings with tonight's fit, and that fit was free to move with the readings. What it does not settle is whether the model is still the plant it was at 19:00, or whether it has followed the plant somewhere. The operator's question is the one the light cannot answer,

is the model still describing the plant, what would that test have to compare with to be able to say no, and what can no such test see.

It is answered on the two-hall shift plant, and it extends Chapter 7, which built the leverage, the normalized residual and the effective degrees of freedom, and Chapter 9, which built the baseline, both on the two-unit hall: here they are tested against a change they might not be able to see.

Jargon: goodness of fit, fit-quality light, predictive check, effective degrees of freedom, noncentrality, probability of detection, reference fit, first verdict, repaired model

Goodness of fit is the agreement between the readings and what the model, fitted, says they should be; the operations form is bad-data detection by the chi-square test. A *fit-quality light* is a screen's summary of that test. A *predictive check* compares a reading, or all of a window's readings, with the distribution a fit predicts for them. The *effective degrees of freedom* of the reading rows are $\sum_k (1 - h_{kk})$, the sum of their redundancy numbers, called below the count. The *noncentrality* of a test under a stated change is the squared size of that change left on the rows the statistic adds up, and the test's *probability of detection*, called below its power, is the probability that its statistic passes its level under that change. A *reference fit* is one fit of a stretch of windows taken as what the plant was, the reference period of Chapter 9 on this plant. A *first verdict* is the first window at which a test passes a level fixed in advance. A *repaired model* is the model with an added entry that the test's failure suggested, fitted again and tested again.

Key idea

A test of a window's readings against that window's own fit asks whether some state explains the readings. A change that lies along a column of the reading rows' Jacobian is such a state: the fit takes it, and what it costs falls on a prior row, which the statistic does not add up. A test of "still" needs a "then": the readings of a window against a fit that did not read them. That test sees only what its refound entries cannot take and what is large against the sensors' spread, and a persistent offset of one spread passes it with or without a reference. No test reaches an entry that no reading row touches.

16.1.1 A test against the window and a test against the past

Chapter 9 set the two tests side by side on a small hall. The window's own test refits the state to the readings, then adds up what is left, and what is left has lost every direction the state could move in. The reference test holds the shared entries where a stretch of earlier windows put them, with their spread, and asks how far the new readings fall from what that predicts. The first has a count of however few directions the state leaves spare, and the second keeps all the rows it does not refit.

The difference is not one of sensitivity at a chosen level. It is a difference of question: *does a state explain these readings*, against *does the state of two hours ago explain them.* An answer of yes to the first says nothing about the second.

> **Principle 16.1.** *A goodness-of-fit test can only reject what the fit could not absorb. What it says about a change is read from the columns of the Jacobian and the rows its statistic adds up, before any reading is taken; a passing test is a statement about those rows, and not about the entries of the state or the answers read from them.*

16.1.2 The two-hall shift plant and its windows (step 0)

The plant is the two-hall shift plant of Chapter 12, with its physical rows (12.2), pinned ratings, capacity rates and other heat, on the reduced state of Chapter 4: the 19 entries (12.3), indexed by window, $\mathbf{z}_t$, with windows five minutes long and window 1 at 19:00. The dependent block of the physical rows is triangular with a nonzero diagonal, so it is invertible, and the temperatures follow by the chain (12.4). The readings' prediction at a window's entries is $\mathbf{m}(\mathbf{z}_t)$, the sixteen values the chain gives.

Modeling choices at the graph. The choices of the plant hold in every window, each failing as in Chapter 12: steady state within the window, one effective conductance per coil, pinned airflow and no active control. They matter more here, because a test of fit is the one place a broken choice could show.

The graph. Figure 16.1 draws one window's rows on its entries, because the chapter's result depends on which rows are which. The sixteen reading rows stand above the entries, each group joined by dashed arrows to the entries it sees; the seventeen prior rows stand below, joined by solid purple arrows. The entries are green; the loads have no prior row; the six unread fractions, in red, have no reading row; and the log-derate g_{B2} is drawn apart, also in red, as the entry this chapter follows. The box around the reading rows marks what the window's test adds up.

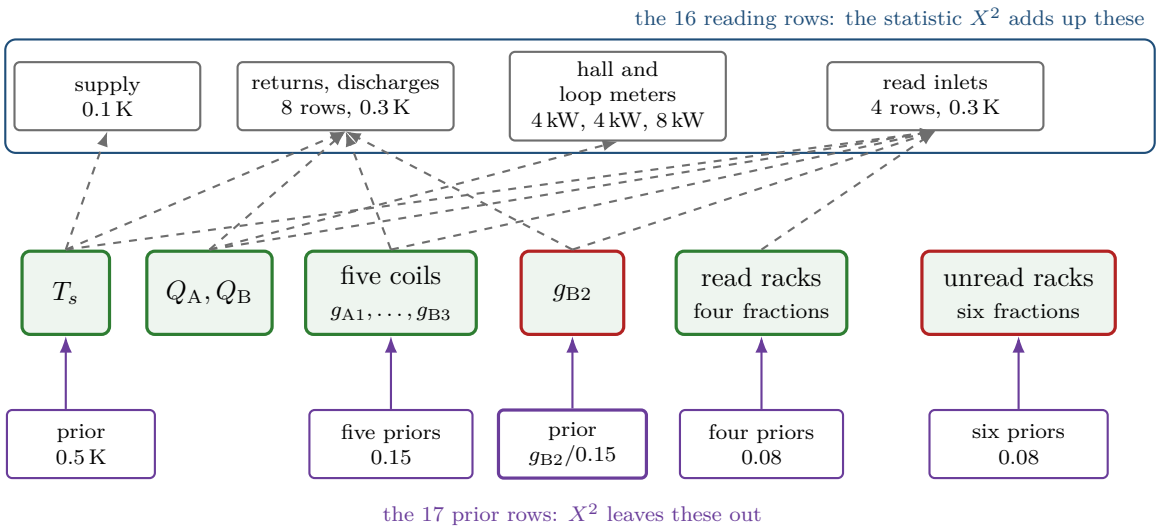


Figure 16.1: One window's rows on the reduced state. Free (green and red): the 19 entries, grouped. Observed (dashed): the sixteen reading rows, with their standard deviations. Solid purple: the seventeen prior rows, $g_u/0.15$, $(\beta_k - \beta_k^s)/0.08$, $(T_s - 14)/0.5$. Red: the unread fractions, with no reading row, and the log-derate this chapter follows. Inside the window the plant is the chain (12.4), with the ratings and capacity rates pinned.

16.1.3 The operator's questions (step 1)

The operator asks six questions at window 73, and step 1 writes each over the state. Table 16.1 lists them: the question in the operator's words, the first of them the drill question; the general form of its category in the notation of Chapter 2; its mathematical form on this plant; its category; the computation of Section 2.3.6 that answers it; and the subsection where it is answered. Answers 1 to 4 read the record through window 144 to test the test; the shift-log entry of Section 16.5.2 does not.

16.1.4 The evidence as rows (step 2)

Step 2 attaches each window's sixteen readings as rows, each Gaussian, unbiased and independent given the state, and the seventeen priors of Chapter 12. That makes 33 rows on 19 entries for the fit of one window. The reference fit of the next paragraphs stacks windows on the same rows.

Which row acts on which entry. With $*$ for a block that can be nonzero and the columns grouped as T_s; Q_A; Q_B; the coils of hall A; the coils of hall B; the read fractions of hall A; those of hall B; the unread fractions, the reading rows have the pattern

$$\mathbf{J}_R = \begin{bmatrix} * & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ * & * & 0 & * & 0 & 0 & 0 & 0 \\ * & 0 & * & 0 & * & 0 & 0 & 0 \\ 0 & * & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & * & 0 & 0 & 0 & 0 & 0 \\ 0 & * & * & 0 & 0 & 0 & 0 & 0 \\ * & * & 0 & * & 0 & * & 0 & 0 \\ * & 0 & * & 0 & * & 0 & * & 0 \end{bmatrix} \begin{matrix} \text{supply} \\ \text{return and discharges, hall A} \\ \text{return and discharges, hall B} \\ \text{meter } Q_A \\ \text{meter } Q_B \\ \text{loop meter} \\ \text{read inlets, hall A} \\ \text{read inlets, hall B} \end{matrix}, \tag{16.1}$$

$$\mathbf{g}_{T_{B3}}^{\mathsf{T}} = (*,\ 0,\ *,\ 0,\ *,\ 0,\ 0,\ *),$$

Table 16.1: The questions asked at window 73, 01:00. X_t^2 and ν_t are the statistic and count of equation (16.6); Q_t the statistic of equation (16.10); $\mathbf{y}$ the readings and prior rows of the fit each answer names.

Operator question	General form	Mathematical form	Category	Computation	Answered in
Is the model still describing the plant? Goodness of fit on the reading residuals against their noise; a posterior predictive check per sensor.	(2.13)	X_t^2 of the reading rows on $\nu_t = \sum_k (1 - h_{kk})$; r_k^{N} per sensor	prediction under a threshold	posterior	16.3.1
Why has the light stayed green all night?	(2.13)	a stated derate split between reading rows and prior rows; probability of detection at window 144	prediction under a threshold	posterior	16.3.2
Is tonight's plant the plant of 19:00?	(2.13)	$Q_t = \mathbf{e}_t^{\mathsf{T}} \mathbf{V}_t^{-1} \mathbf{e}_t$ against the reference, on 13; per sensor by the hour	prediction under a threshold	posterior	16.3.3
When did that show, and how firmly?	(2.12) per window	first t with tail below 10^{-3}; its spread over fresh sensor errors	prediction under a threshold	posterior	16.3.4
What restores the fit?	(2.14)	X^2 and evidence with g_{B2} free per window, with and without the meter's bias	hypothesis comparison	evidence	16.3.5
What can no test see?	(2.10)	the empty columns; the share of $\mathrm{Var}(T_{\mathrm{B3}} \mid \mathbf{y})$ on them	estimation	posterior	16.3.6

and each prior row is a single entry on its own column. The second line is the gradient of the unread rack's inlet, and it is not a row. Three facts are read here before any number. The last column is empty in every reading row: the six fractions β_{A2}, β_{A4}, β_{A5}, β_{B1}, β_{B3}, β_{B5} are not observable, and the gradient has a nonzero there. The reading rows' Jacobian has rank 13 of 19, so the state can move in thirteen directions the readings see: at most $16 - 13 = 3$ directions of the readings cannot be matched by it, before the priors hold anything back. And rows minus entries, $33 - 19 = 14$, is not the readings' effective degrees of freedom: the six empty columns cancel six prior rows, and $14 = 3 + 11$ counts the eleven prior rows on seen entries as spare, while $16 - 19$ is not a count at all. The count that belongs to the statistic is computed from the leverages in step 4.

The reference window. Windows 1 to 24, 19:00 to 20:55, are chosen as the reference before any verdict. A test of "still" needs a "then", and one window has none, so this chapter reads a reference fit beside the one fit of Part IV. This buys a "then". It fails for a fault already in the reference that a shared entry can fully absorb, such as a constant offset of a coil, which becomes part of the "then" and which no later comparison can show; a meter's offset is a different case (Section 16.3.5).

Coils and fractions constant in the reference. The sixteen coils and fractions are shared across the twenty-four windows, and the supply and both loads are free in each. This is the reference drawing's own claim, and it buys narrow shared entries from many windows; it fails if a coil moved in those two hours, and the reference's own X^2 tests that.

The repaired drawings. The drawing that Section 15.2.2 selects on this plant, by the method of Chapter 10, holds a change in unit B2's coil, and its reasoning is cited here and not repeated. The repaired fits free g_{B2} per window over windows 1 to 73, with its rating prior in each and no assumed shape. This buys a test of the rest of the model with that change taken out; it fails if the change is not in that coil, since a free entry can absorb what it does not cause. One of the two also frees a constant bias on the hall-A meter, with a prior of $\mathcal{N}(0,\ 10^2)\,\mathrm{kW}$. Beside them stands the unrepaired stack over the same windows, with g_{B2} shared.

16.2 The posterior of the state (step 3)

Parts IV to VI read one fit, and Section 16.2.1 cites it. The reference and repaired fits differ from it in the rows they stack, for the reasons of Section 16.1.4, and each is derived here because an answer uses it.

16.2.1 The one fit and its hat matrix

The one fit of window 73 is that of Section 12.2.2: $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$ by the normal equations (12.9) on the rows (12.8), with its iterations, its Jacobian and its covariance as shown there. For each other window the same rows give that window's own $\hat{\mathbf{z}}_t$, $\boldsymbol{\Sigma}_t$ and $\mathbf{J}_t$. The one object this chapter adds is the hat matrix of Chapter 7 on those 33 rows,

$$\mathbf{H}_t = \mathbf{J}_t \boldsymbol{\Sigma}_t \mathbf{J}_t^{\mathsf{T}}, \qquad h_{kk} = [\mathbf{H}_t]_{kk}, \qquad \operatorname{tr} \mathbf{H}_t = 19, \tag{16.2}$$

whose complement $\mathbf{I} - \mathbf{H}_t$ gives the residuals of equation (7.12) with $\boldsymbol{\Sigma}_t = (\mathbf{J}_t^{\mathsf{T}} \mathbf{J}_t)^{-1}$. With every row normalized the information matrix is $\mathbf{J}^{\mathsf{T}}\mathbf{J}$; with unnormalized rows it is $\mathbf{J}^{\mathsf{T}}\mathbf{W}\mathbf{J}$. The diagonal of $\mathbf{H}_t$ at window 73 has 12.309 on the reading rows and 6.691 on the prior rows, 6.000 of that on the six unread fractions' own rows.

16.2.2 The reference fit

The reference state stacks the shared entries with each window's own three,

$$\mathbf{z}_{\mathrm{ref}} = \left(g_{\mathrm{A1}}, \ldots, g_{\mathrm{B3}},\ \beta_{\mathrm{A1}}, \ldots, \beta_{\mathrm{B5}};\ T_{s,t},\ Q_{\mathrm{A},t},\ Q_{\mathrm{B},t} \text{ for } t = 1, \ldots, 24\right)^{\mathsf{T}}, \tag{16.3}$$

$16 + 72 = 88$ entries, with the rows

$$\mathbf{r}(\mathbf{z}_{\mathrm{ref}}) = \begin{bmatrix} (m_k(\mathbf{z}_t) - y_{k,t})/\sigma_k \quad (k = 1, \ldots, 16;\ t = 1, \ldots, 24) \\ (T_{s,t} - 14)/0.5 \quad (t = 1, \ldots, 24) \\ g_u/0.15 \quad \text{(six units)} \\ (\beta_k - \beta_k^{\mathrm{s}})/0.08 \quad \text{(ten racks)} \end{bmatrix}, \tag{16.4}$$

$384 + 24 + 6 + 10 = 424$ rows, $\mathcal{L} = \frac{1}{2}\mathbf{r}^{\mathsf{T}}\mathbf{r}$. The Jacobian is an arrowhead: each window's rows see the 16 shared columns and that window's own three. Its reading rows have the pattern (16.1) in each window, so the six empty columns stay empty however many windows are stacked. The normal equations are those of equation (4.29),

$$\left(\mathbf{J}^{\mathsf{T}}\mathbf{J}\right)\Delta\mathbf{z}_{\mathrm{ref}} = -\mathbf{J}^{\mathsf{T}}\mathbf{r}, \qquad \mathbf{\Sigma}_{\mathrm{ref}} = \left(\mathbf{J}^{\mathsf{T}}\mathbf{J}\right)^{-1}\big|_{\hat{\mathbf{z}}_{\mathrm{ref}}}, \tag{16.5}$$

the Gauss–Newton form of equation (4.38), from the ratings, the survey, the set point and the meters. Its block on the shared entries is written $\mathbf{\Sigma}_{\mathrm{s}}$, and the shared entries' MAP values $\hat{\mathbf{z}}_{\mathrm{s}}$.

The iterations. Table 16.2 runs equation (16.5). Each row is an iteration; the columns give the objective and the norm of its gradient. The objective falls from 743.9383 to 165.3726 in the first step and settles at 165.2402 by the second.

Table 16.2: Gauss–Newton iterations of the reference fit, equation (16.5).

Iteration k [–]	Objective $\mathcal{L}(\mathbf{z}^{(k)})$ [–]	$\lVert\mathbf{J}^{\mathsf{T}}\mathbf{r}\rVert$ [–]
0	743.9383	1.046×10^{4}
1	165.3726	1.373×10^{2}
2	165.2402	6.947×10^{-1}
3	165.2402	8.576×10^{-3}
4	165.2402	9.032×10^{-5}
5	165.2402	1.424×10^{-6}

The MAP state. Table 16.3 gives the shared entries, each row an entry, the columns its value and posterior spread. The coil that the rest of the chapter follows stands at $g_{\mathrm{B2}} = +0.0008 \pm 0.0057$. The six unread fractions stand at the survey, $\beta_{\mathrm{B3}} = 0.5000 \pm 0.0800$ among them, as the empty columns require. The four read fractions have moved from the survey and their spreads have fallen to 0.0058 and 0.0069.

The covariance. On $(g_{\mathrm{B1}}, g_{\mathrm{B2}}, g_{\mathrm{B3}}, \beta_{\mathrm{B2}}, \beta_{\mathrm{B4}})$, with the spreads of Table 16.3, the correlations are

$$\begin{bmatrix} 1 & -0.139 & -0.148 & 0.164 & 0.162 \\ -0.139 & 1 & -0.152 & 0.160 & 0.158 \\ -0.148 & -0.152 & 1 & 0.153 & 0.151 \\ 0.164 & 0.160 & 0.153 & 1 & 0.200 \\ 0.162 & 0.158 & 0.151 & 0.200 & 1 \end{bmatrix}.$$

The rows fix the signs. The hall's return reading fixes the three conductances' sum, so a larger coil is explained by a smaller neighbor, and the coils correlate negatively. An inlet warms when a

Table 16.3: The reference fit's shared entries, windows 1 to 24, with the per-window entries omitted.

Entry [–]	$\hat{z}_i$ [–]	$\sqrt{\Sigma_{ii}}$ [–]	Entry [–]	$\hat{z}_i$ [–]	$\sqrt{\Sigma_{ii}}$ [–]
g_{A1}	−0.0460	0.0050	β_{A3}	0.3988	0.0058
g_{A2}	+0.0174	0.0047	β_{A4}	0.3000	0.0800
g_{A3}	−0.0766	0.0051	β_{A5}	0.4500	0.0800
g_{B1}	+0.0267	0.0056	β_{B1}	0.2000	0.0800
g_{B2}	+0.0008	0.0057	β_{B2}	0.3440	0.0069
g_{B3}	−0.0477	0.0060	β_{B3}	0.5000	0.0800
β_{A1}	0.1754	0.0058	β_{B4}	0.3818	0.0069
β_{A2}	0.2500	0.0800	β_{B5}	0.3500	0.0800

coil falls and when its fraction rises, so an inlet reading ties a larger coil to a larger fraction, and those correlate positively; the two fractions share the cold aisle's error and correlate positively too. Every unread fraction is uncorrelated with everything.

The reference's own test. The reading rows have $X^2 = 328.60$ on a count of 302.52, upper tail 0.145. The drawing's claim of constant coils and fractions is not rejected on its own hours.

16.2.3 The repaired fits

The repaired state has 15 shared entries (five coils and the ten fractions) and each window's $T_{s,t}$, $Q_{\mathrm{A},t}$, $Q_{\mathrm{B},t}$ and $g_{\mathrm{B2},t}$ for $t = 1, \ldots, 73$, so 307 entries. Its rows are the sixteen readings, the set point's prior and the rating prior of $g_{\mathrm{B2},t}$ in each window, and the 15 shared priors, 1329 rows, the same normal equations as (16.5). The second drawing adds one entry $b_{Q_{\mathrm{A}}}$, added to the prediction of the hall-A meter in every window, and one prior row. It gives the bias $+4.595 \pm 0.912\,\mathrm{kW}$. The unrepaired stack shares g_{B2} across the windows, with 235 entries and 1257 rows. The Laplace evidence of equation (11.8) is $\ln p(\mathbf{y} \mid M_0) = -831.12$ for the drawing without the bias and $\ln p(\mathbf{y} \mid M_1) = -820.98$ for the one with it, up to the constant both share.

16.3 The answers (step 4)

Each answer is read from the fit it names: the window's own in the first two, the reference in the next two, the repaired fits in the fifth and the one fit in the last. Where a reported tail rests on a linear predictive, draws stand beside it.

16.3.1 The window's own test

Operator question. Is the model still describing the plant? Goodness of fit on the reading residuals against their noise; a posterior predictive check per sensor. The first question asks what the screen's light computes.

General form. A test of fit is a tail of the predictive, equation (2.13), of a statistic of replicate readings.

Specialized form. The question is equation (2.13) with y_{new} the sixteen replicate readings of a window, $\mathbf{z} = \mathbf{z}_t$ and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal of the one fit at that window. The statistic and its reference are those of equation (9.14) on sixteen rows, and per sensor the normalized

residual of equation (9.12),

$$X_t^2 = \sum_{k=1}^{16} r_k(\hat{\mathbf{z}}_t)^2, \qquad \nu_t = \sum_{k=1}^{16} (1 - h_{kk}), \qquad r_k^{\mathrm{N}} = \frac{r_k(\hat{\mathbf{z}}_t)}{\sqrt{1 - h_{kk}}}, \tag{16.6}$$

with h_{kk} from equation (16.2) and the tail $\Pr(\chi^2_{\nu_t} > X_t^2)$, on a non-integer count, as in Chapter 9. The statistic leaves out the prior rows, for the reason given there: their residuals are estimates of the entries' departures, not misfit.

Building blocks. The rows at $\hat{\mathbf{z}}_t$; the hat matrix (16.2); the chi-square tail.

Window 73. The reading rows give $X^2 = 0.755$ on a count of 3.691, upper tail 0.925; the prior rows, outside the statistic, add 1.795. Table 16.4 gives the rows, each row a reading, the columns its leverage, redundancy, residual, normalized residual and two-sided predictive tail. Figure 16.2 draws the leverages: one bar per reading, the horizontal axis the leverage from zero to one. The bars are long at the supply, the discharges and the inlets, each of which is largely its own evidence; they are short only at the two returns and the loop meter. Sixteen readings leave a count of 3.691: a little above the three directions of equation (16.1) because the priors hold the state back slightly. The largest normalized residual is -0.718 at the discharge of unit B2, with a two-sided tail of 0.4728, so no sensor stands out.

Table 16.4: The window's own test at window 73, sensor by sensor; residuals as predicted minus observed, in spreads.

Reading	h_{kk} [–]	$1 - h_{kk}$ [–]	r_k [–]	r_k^{N} [–]	Two-sided tail [–]
T_s	0.9517	0.0483	+0.105	+0.477	0.6336
$T_{r,\mathrm{A}}$	0.3204	0.6796	−0.346	−0.420	0.6747
$T_{r,\mathrm{B}}$	0.3158	0.6842	+0.095	+0.115	0.9083
$T_{d,\mathrm{A1}}$	0.8949	0.1051	+0.101	+0.312	0.7554
$T_{d,\mathrm{A2}}$	0.8979	0.1021	+0.216	+0.677	0.4986
$T_{d,\mathrm{A3}}$	0.8956	0.1044	+0.127	+0.392	0.6948
$T_{d,\mathrm{B1}}$	0.8869	0.1131	+0.004	+0.013	0.9899
$T_{d,\mathrm{B2}}$	0.8797	0.1203	−0.249	−0.718	0.4728
$T_{d,\mathrm{B3}}$	0.8840	0.1160	−0.098	−0.288	0.7732
Q_{A}	0.7371	0.2629	−0.103	−0.201	0.8405
Q_{B}	0.7354	0.2646	−0.309	−0.600	0.5483
Q_{loop}	0.3025	0.6975	+0.572	+0.685	0.4935
T_{A1}	0.9191	0.0809	−0.074	−0.261	0.7944
T_{A3}	0.9191	0.0809	−0.080	−0.283	0.7772
T_{B2}	0.8843	0.1157	−0.038	−0.111	0.9115
T_{B4}	0.8843	0.1157	+0.154	+0.453	0.6506

Through the shift. Figure 16.3 draws the same test at every window: the horizontal axis is the window, the vertical axis the tail, with a line at 0.05. The points fill the square evenly from the first window to the last, with nothing in the second half that the first half lacks. The count stays between 3.631 and 3.765; X^2 has mean 3.577 and standard deviation 2.645. Nine of the 144 tails fall below 0.05, against 7.2 expected, and one below 0.01, against 1.44, the smallest 0.0100 at window 19; a Kolmogorov–Smirnov test of uniformity gives a statistic of 0.0783 and $p = 0.323$. The largest normalized residual of the shift is 3.247 in magnitude, at $T_{r,\mathrm{A}}$ in window 16. Table 16.5 sums the test by quarter of the night, each row a quarter: the last quarter's tail, 0.2535, is above the first's, 0.2105.

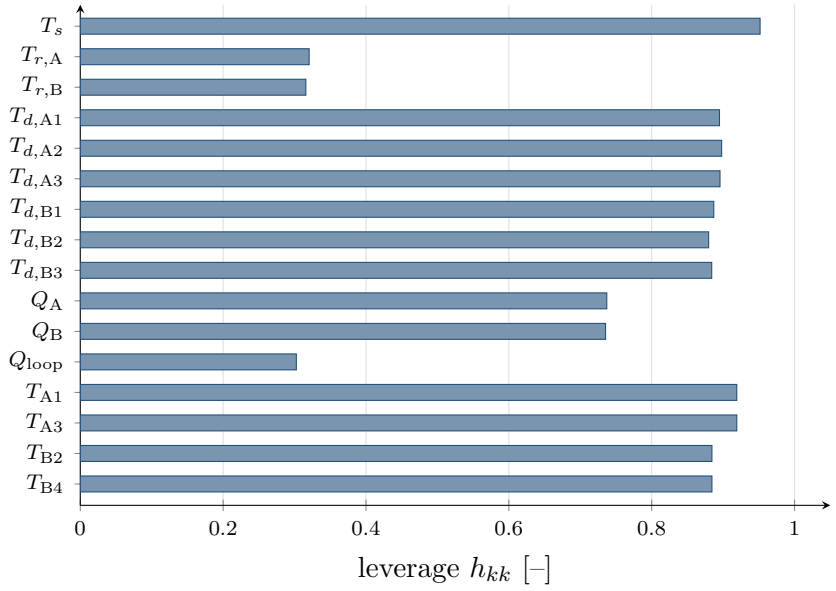


Figure 16.2: The leverages of the sixteen reading rows at window 73, from equation (16.2). A reading's redundancy is what its bar leaves to one; the sum of those is the count.

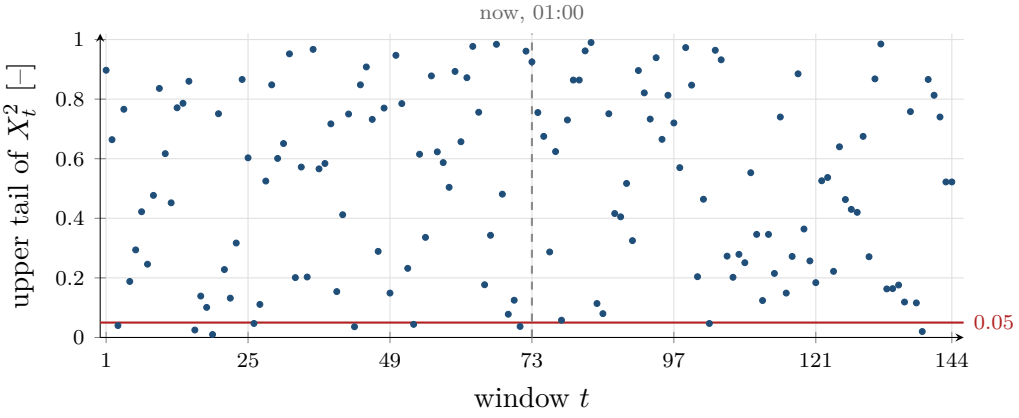


Figure 16.3: The window's own test at each of the 144 windows: the tail of equation (16.6). The whole night is drawn to test the test; only windows through 73 had happened at 01:00.

Table 16.5: The window's own test summed by quarter of the night.

Windows	Sum of X^2 [–]	Sum of counts [–]	Upper tail [–]
1–36	147.25	134.33	0.2105
37–72	121.91	133.23	0.7497
73–108	103.95	132.39	0.9678
109–144	141.95	131.59	0.2535

16.3.2 What the window's test cannot see

Operator question. Why has the light stayed green all night? A flat test may mean a flat plant, or a test that could not have responded.

General form. The power of the test under a stated change is again a tail of the predictive (2.13), with the replicates drawn from the changed plant.

Specialized form. The question is equation (2.13) with y_{new} the sixteen readings of a plant whose coil has a stated derate, and the event X_t^2 past the level of equation (16.6). By equation (9.17) what a change does to the residuals follows from the hat matrix. Let $\mathbf{d}$ be the change in the sixteen readings, in spreads, made by a derate of the stated size at the window's state. The rows' residuals then shift by

$$\boldsymbol{\ell} = (\mathbf{I} - \mathbf{H}_t) \begin{bmatrix} \mathbf{d} \\ \mathbf{0} \end{bmatrix}, \qquad \|\mathbf{d}\|^2 = \underbrace{\|\mathbf{d}\|^2 - \|\boldsymbol{\ell}\|^2}_{\text{taken by the state}} + \underbrace{\|\boldsymbol{\ell}_{\text{R}}\|^2}_{\text{reading rows}} + \underbrace{\|\boldsymbol{\ell}_{\text{P}}\|^2}_{\text{prior rows}}, \tag{16.7}$$

because $\mathbf{I} - \mathbf{H}_t$ is an orthogonal projection, so no cross term remains. The test's noncentrality is $\lambda = \|\boldsymbol{\ell}_{\text{R}}\|^2$ alone, and its power at a tail α is the noncentral chi-square tail beyond the level on the count ν_t. One choice enters here: the split is linear about the fit, and fresh nights of readings stand beside it.

Building blocks. The hat matrix (16.2); the gradient of the sixteen readings along g_{B2}; the count (16.6).

The coil's move. Each window's fit moves g_{B2} as the night goes on, -0.0010, -0.0063, -0.1444, -0.2649 and -0.3258 at windows 1, 37, 73, 109 and 144, each with a spread near 0.03. Figure 16.4 draws it: on the left, the horizontal axis is the window and the vertical axis the fitted g_{B2} with its 95 % band and the stated truth dotted, a check marker for Section 16.4.2; on the right, at the same windows, the reading rows' X^2 and the size of the prior row of g_{B2}. The fitted coil walks away from its rating while X^2 does not follow it. The prior row reads -0.963, -1.766 and -2.172 spreads at windows 73, 109 and 144: an estimate to watch, as Chapter 11 watches it, and not a fit statistic.

Where the move goes. The size $\mathbf{d}$ uses the stated truth of the constructed record, a size given here and not estimated. At window 73, the derate to its value there gives $\|\mathbf{d}\|^2 = 34.35$, of which the state takes 33.65, the prior rows hold 0.6663 and the reading rows 0.0257. As a test of the test, window 144 has the largest move of the night, a derate to $g_{\text{B2}} = -0.3567$. It changes the readings by +13.96 spreads at the discharge of unit B2, +5.83 at the return of hall B and at each read inlet of hall B, and +1.63 and +1.91 at the other two discharges, so $\|\mathbf{d}\|^2 = 303.26$. Figure 16.5 draws the split (16.7): one bar per part, the horizontal axis in squared spreads on a logarithmic scale. The state takes 297.12, the prior rows hold 5.8909, and the reading rows 0.2472. The change lies along the coil's own column, so the state takes it; what is left is the small price the prior charges. With the prior of g_{B2} removed, the reading rows would hold 0.00001. Of the prior rows' part, 5.8892 is the row of g_{B2}, with an entry of -2.4268 spreads in the linear split, against a mean of -2.291 on the fresh nights below.

The power. At window 73, with $\lambda = 0.0257$, the power at a tail of 0.05 is 0.051; at window 144, with $\lambda = 0.2472$, it is 0.063. These lean on the chi-square on ν_t, which is approximate at a fixed state: its mean there is $\text{tr}[(\mathbf{I} - \mathbf{H}_t)^2_{\text{RR}}] = 3.060$ at window 144, not the count of 3.633, because the fit's prior rows are not redrawn. So fresh nights stand beside the power. Table 16.6

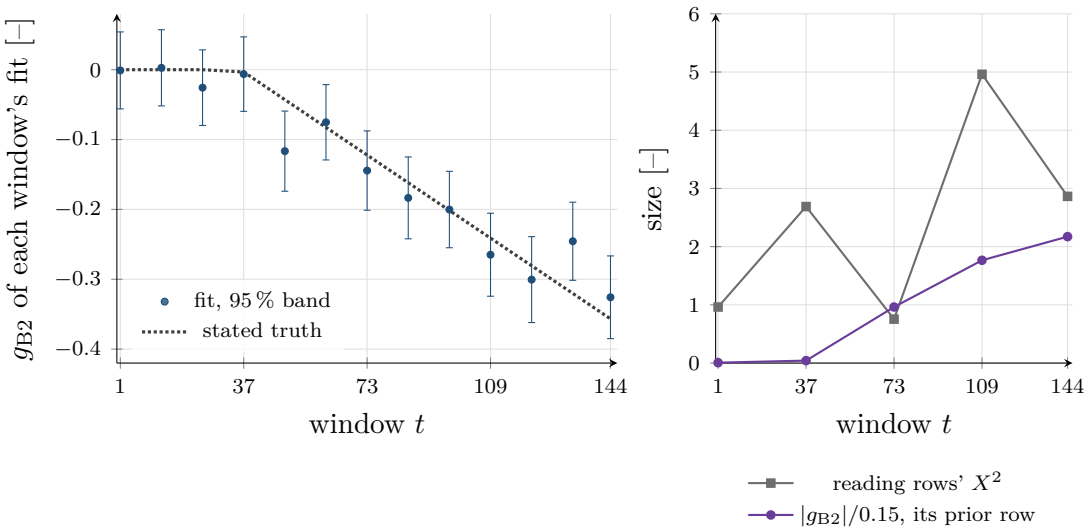


Figure 16.4: Left: the log-derate g_{B2} of each window's own fit, every twelfth window and the last, with 1.96 posterior spreads, and the stated truth. Right: at five windows, the statistic of the window's test and the size of the prior row that the statistic leaves out.

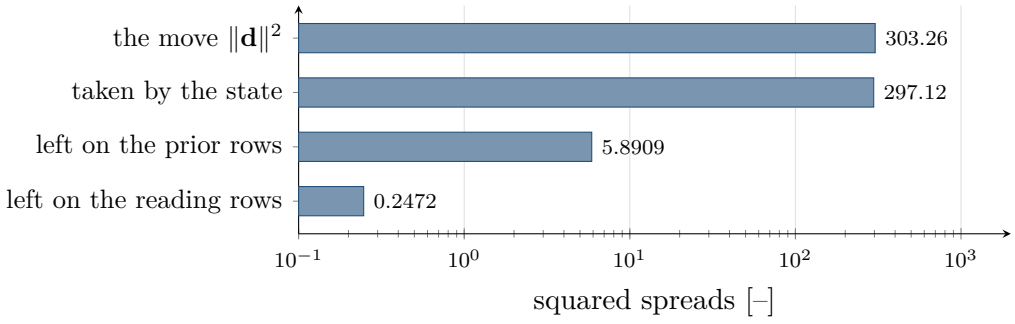


Figure 16.5: Where a stated derate of g_{B2} goes in the fit of window 144, equation (16.7). Logarithmic axis. The window's test adds up only the bottom bar.

gives 40 000 nights of sensor errors about the constructed state at each case, each row a case, with and without the hall-A meter's +4 kW offset. At window 144 with no change and no offset X^2 has mean 3.152. With the derate and the offset the tail falls below 0.05 on 0.0483 of nights, against 0.0447 for no change with the same offset, each with a standard error near 0.001: the derate adds almost nothing. The whole objective on 14, which does add the prior rows up, rejects at 0.05 on only 0.0022 of nights. The row of g_{B2} alone, watched as an estimate, does respond: $|g_{\mathrm{B2}}/0.15| > 1.96$ on 0.9573 of nights at window 144, and on none of 40 000 at window 73. The test is flat because a change along a column cannot move it.

Table 16.6: The window's own test on fresh nights, 40 000 per row; standard errors about 0.001 at 0.05 and 0.0005 at 0.01. "Offset" is +4 kW on the hall-A meter.

Case	Mean X^2 [–]	Tail < 0.05 [–]	Tail < 0.01 [–]	Whole objective on 14, < 0.05 [–]	$\lvert g_{\mathrm{B2}}/0.15\rvert > 1.96$ [–]
Window 73, no change, no offset	3.182	0.0316	0.0059	0.0001	0.0000
Window 73, no change, offset	3.446	0.0440	0.0089	0.0002	0.0000
Window 73, derate, offset	3.442	0.0430	0.0091	0.0004	0.0000
Window 144, no change, no offset	3.152	0.0333	0.0061	0.0001	0.0000
Window 144, no change, offset	3.401	0.0447	0.0092	0.0004	0.0000
Window 144, derate, no offset	3.364	0.0360	0.0070	0.0013	0.9520
Window 144, derate, offset	3.619	0.0483	0.0103	0.0022	0.9573

16.3.3 The plant against its first two hours

Operator question. Is tonight's plant the plant of 19:00? The test now needs what the window's own test lacked: a fit that did not read this window.

General form. This is the predictive of equation (2.13), with the posterior now that of an earlier stretch of the record.

Specialized form. The question is equation (2.13) with $y_{\mathrm{new}} = \mathbf{y}_t$, the readings of window t, and $p(\mathbf{z} \mid \mathbf{y})$ the reference's Laplace normal on the shared entries, $\mathcal{N}(\hat{\mathbf{z}}_{\mathrm{s}}, \boldsymbol{\Sigma}_{\mathrm{s}})$. One choice enters here. The supply and both loads are free per window in the predictive, written $\mathbf{q}_t = (T_{s,t}, Q_{\mathrm{A},t}, Q_{\mathrm{B},t})^{\mathsf{T}}$: the test asks about the plant, not about the batch load. This costs three of the sixteen rows, and it fails to see any change a load or the supply could explain. With $\mathbf{A}_{\mathrm{q}}$ and $\mathbf{A}_{\mathrm{s}}$ the Jacobian of the prediction on $\mathbf{q}_t$ and on the shared entries, the predictive covariance is

$$\mathbf{V}_t = \operatorname{diag}(\sigma_k^2) + \mathbf{A}_{\mathrm{s}} \boldsymbol{\Sigma}_{\mathrm{s}} \mathbf{A}_{\mathrm{s}}^{\mathsf{T}}. \tag{16.8}$$

The entries $\mathbf{q}_t$ are found by the weighted normal equations on that covariance, with the set point's prior row as the information $\mathbf{P}$,

$$\mathbf{L} = \mathbf{A}_{\mathrm{q}}^{\mathsf{T}} \mathbf{V}_t^{-1} \mathbf{A}_{\mathrm{q}} + \mathbf{P}, \qquad \mathbf{L}\, \Delta\mathbf{q}_t = \mathbf{A}_{\mathrm{q}}^{\mathsf{T}} \mathbf{V}_t^{-1} \mathbf{e}_t - \mathbf{P}(\mathbf{q}_t - \mathbf{q}_0), \tag{16.9}$$

and at their solution the statistic and the per-sensor form are

$$\mathbf{e}_t = \mathbf{y}_t - \mathbf{m}(\hat{\mathbf{q}}_t, \hat{\mathbf{z}}_{\mathrm{s}}), \qquad Q_t = \mathbf{e}_t^{\mathsf{T}} \mathbf{V}_t^{-1} \mathbf{e}_t \;\sim\; \chi^2_{13}, \qquad z_{k,t} = \frac{e_{k,t}}{\sqrt{\left[\mathbf{V}_t - \mathbf{A}_{\mathrm{q}} \mathbf{L}^{-1} \mathbf{A}_{\mathrm{q}}^{\mathsf{T}}\right]_{kk}}}. \tag{16.10}$$

Here $\mathbf{e}_t$ is observed minus predicted, the sign of the screen: a positive $z_{k,t}$ is a reading above what the reference predicts, the opposite of the rows' sign. An hour's block sums twelve Q_t on 156, and a sensor's hour takes the mean of its twelve $z_{k,t}$ times $\sqrt{12}$. The noncentrality of a stated change $\mathbf{d}$ is Q with $\mathbf{d}$ in place of $\mathbf{e}_t$, after removing what $\mathbf{q}_t$ can take.

Building blocks. $\hat{\mathbf{z}}_{\mathrm{s}}$ and $\mathbf{\Sigma}_{\mathrm{s}}$ of the reference fit (Section 16.2.2); the covariance (16.8); the solve (16.9); the chi-square tail.

The night. Figure 16.6 draws Q_t: the horizontal axis is the window from 25 on, the vertical axis Q_t on a logarithmic scale, with a line at 34.53, where the tail on 13 is 10^{-3}, and a line at the count of 13. Up to about window 50 the points scatter about the count. From then on they lift off and keep rising. At window 49 $Q = 33.30$, tail 1.54×10^{-3}; at window 73, 38.90, tail 2.08×10^{-4}; at window 97, 74.97, tail 9.65×10^{-11}; at window 144, 189.03, tail 2.42×10^{-33}. By the hour on 156: 21:00 has 168.1, tail 0.241; 22:00 142.8, tail 0.767; and 23:00 238.0, tail 2.57×10^{-5}. The same stated derate as above has noncentrality 194.03 in this test at window 144, against 0.2472 in the window's own, and power 1.000 at a tail of 10^{-3}; at window 73, the derate there has 22.18 and power 0.490. What this says is that the 19:00–21:00 model no longer fits. It does not say why: a drifting discharge sensor on unit B2 would fail the same way.

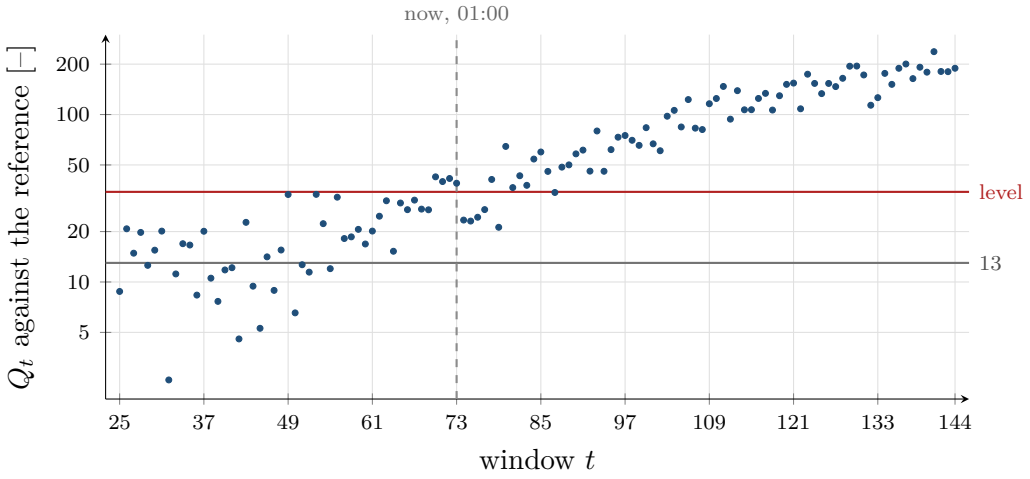


Figure 16.6: The statistic Q_t of equation (16.10) at each window from 25, against the fit of windows 1 to 24, with its count of 13 and the value 34.53 at which its tail is 10^{-3}. Logarithmic axis. Windows after 73 are drawn to test the test.

Sensor by sensor. Figure 16.7 gives the hour's z of each sensor: each row is a sensor, each column an hour, red positive and blue negative, the number in each cell. At window 73 the largest is the discharge of unit B2, $z = +5.25$, where the reading stands +1.568 K above a prediction with a residual spread of 0.299 K; the one fit, which has moved, leaves it +0.075 K on Table 16.4's row. The next is +1.96 at the hall-A meter. An hour's z is not of unit spread, because the windows share the reference: on nights with no event its measured spread is 1.00 to 1.22 across the sensors, the larger on the temperatures. By the hour, that discharge reaches 7.7 in the 23:00 hour and 37.9 in the 06:00 hour. It does not stand alone, and the direct and the smeared parts can be told apart by what a derate alone would do in the 06:00 hour. It moves the readings directly at the discharge of B2 by +13.21 spreads, with $z = +38.90$; at the inlets of B2 and B4 by +5.52, with $z = +10.22$ and +10.01; and at hall B's return by +5.52, with $z = +6.75$. Those are the coil's own signature, and the record's +10.7 and +10.1 at the two inlets match them. It moves the hall-B meter, the supply, the hall-A meter, the loop meter and hall A's discharges directly by nothing, yet produces $z = -20.80$, −17.96, +8.79, −4.92 and −3.40 to −3.63. Those are smear: the reference holds the coil fixed, so the refound $\mathbf{q}_t$ moves to explain what it cannot, and the rows it touches take the opposite. The record's −19.2 at the hall-B meter, −17.4 at the supply and +10.7 at the hall-A meter are mostly this, with

nothing wrong in those instruments under this drawing. The largest z points at a row; it does not isolate a fault, which is the work of Chapter 10.

	21:00	22:00	23:00	00:00	01:00	02:00	03:00	04:00	05:00	06:00
T_s	2.1	0.6	−2.1	−4.3	−5.6	−6.9	−11.1	−13.1	−13.8	−17.4
$T_{r,\mathrm{A}}$	−2.8	−2.4	−1.6	−2.4	−1.7	−2.1	−1.5	−1.7	−2.6	−2.8
$T_{r,\mathrm{B}}$	1.7	1.6	1.6	2.6	2.7	2.0	5.6	4.7	5.6	7.3
$T_{d,\mathrm{A1}}$	−2.7	−0.3	0.3	−0.8	−0.8	−2.0	0.0	−3.1	−4.4	−2.3
$T_{d,\mathrm{A2}}$	−1.9	−0.5	−1.8	−1.7	−0.9	−3.2	−4.0	−4.0	−5.0	−3.2
$T_{d,\mathrm{A3}}$	0.0	−0.7	−1.8	−0.6	−1.8	−1.2	−1.8	−2.5	−4.5	−3.6
$T_{d,\mathrm{B1}}$	−1.1	−0.6	−0.4	−1.8	−0.8	−1.5	−0.8	−0.3	−2.1	−2.3
$T_{d,\mathrm{B2}}$	1.1	2.2	7.7	12.0	13.6	20.6	25.2	32.1	36.5	37.9
$T_{d,\mathrm{B3}}$	0.9	−1.5	0.8	−0.9	0.0	−0.7	−0.9	−3.4	−1.7	−2.7
Q_{A}	2.6	2.8	4.2	5.4	5.9	5.8	7.9	7.1	9.2	10.7
Q_{B}	−2.1	−2.0	−3.0	−4.7	−7.7	−8.7	−11.5	−14.5	−16.7	−19.2
Q_{loop}	0.8	−0.4	−2.3	−3.4	−2.6	−2.7	−5.9	−4.3	−5.1	−7.1
T_{A1}	1.4	−1.3	−1.0	0.8	−0.3	−3.1	−1.8	−0.9	−0.1	−2.7
T_{A3}	−1.5	0.1	−0.9	−1.6	−4.1	0.0	−2.0	−1.2	−1.6	−1.6
T_{B2}	−0.6	1.1	1.2	3.5	3.5	5.0	4.6	6.7	9.0	10.7
T_{B4}	0.3	0.8	1.0	1.7	4.7	4.0	5.8	7.0	6.4	10.1

Figure 16.7: The z of each sensor's hour against the reference, from equation (16.10): the mean of twelve windows times $\sqrt{12}$. Red: a reading above the reference's prediction; blue: below. Color saturates at $|z| = 10$. Each column is named by its hour's first window; hours after 00:00 are drawn to test the test.

16.3.4 The first verdict

Operator question. When did that show, and how firmly? A time at which a verdict first holds is one night's draw.

General form. Each window's verdict is the event of the threshold form, equation (2.12), with the assurance replaced by a stated level of its tail.

Specialized form. The question is equation (2.12) with f the statistic Q_t of equation (16.10) and f^* the value at which its tail is the stated level, 34.53 on 13 for one window and 216.32 on 156 for a block of twelve; the first verdict is the first t past it. One choice enters here: a level of 10^{-3} per test. It is the chapter's stated level for a first verdict and it is not an alarm level; it buys a time to report, and it fails as a rule because 120 windows are tested in a row, with no correction. A time alone is one night's sensor errors, so the same calculation is repeated with those errors redrawn around the constructed truth, whose loads and set point are held.

Building blocks. The statistic (16.10), with the reference refitted on each fresh night's first 24 windows; the chi-square tail.

The times. On this night a single window first passes at window 70, 00:45, and an hour block first at the block ending at window 57, 23:40. Table 16.7 gives these with their spread, each row a form of the test. Over 1000 fresh nights a single window's first verdict has median 62.0 and 2.5–97.5 % range 42.0 to 72.0, and a block's has median 58.0 and range 48.0 to 65.0, 22:55 to 00:20 for the block's end. This night's window 70 is late in its spread: 0.093 of nights come at 70 or later. Both ranges end before window 73, where the window's own test gives no verdict

at all. The false-alarm side uses 1000 fresh nights with the meter's offset and no coil event, Q of mean 13.279 against 13. A single window passes the level at a rate of 0.00138 per test, and 0.153 of nights hold at least one single-window verdict and 0.088 at least one block verdict. With no offset either, over 600 nights, those shares are 0.093 and 0.053. A verdict at this level is therefore not silent on a night with no event, which is what 120 uncorrected looks at 10^{-3} cost.

Table 16.7: The first verdict at a level of 10^{-3} per test. Windows; clock times where the night's value is a single window.

Form of the test	This night [window]	Median, 1000 nights [window]	2.5–97.5 % [window]	Nights with a verdict, no coil event [–]
Single window on 13	70 (00:45)	62.0	42.0–72.0	0.153
Hour block on 156	57 (23:40)	58.0	48.0–65.0	0.088

16.3.5 The repaired model

Operator question. What restores the fit? A failed test only says that something changed; the fit with the change admitted says whether anything is left.

General form. Two drawings are compared by the ratio of their evidences, equation (2.14), and each is tested by its own reading rows, the predictive (2.13).

Specialized form. The question is equation (2.14) with M_A the repaired drawing with the bias free and M_B the repaired drawing without it, both of Section 16.2.3, and $\ln B_{AB} = \ln p(\mathbf{y} \mid M_1) - \ln p(\mathbf{y} \mid M_0)$. Each drawing's reading rows are tested by equation (16.6) summed over windows 1 to 73, on the count of its own fit.

Building blocks. The two fits and their evidences (Section 16.2.3); the hat matrix on their stacked rows; each sensor's mean residual, normalized by its own null spread through $\mathbf{I} - \mathbf{H}$ (for a sensor over n windows, the mean times $\sqrt{n}$ has null spreads such as 0.817 at $T_{r,\mathrm{A}}$, 0.511 at Q_{A}, 0.835 at Q_{loop} and 0.274 at hall A's discharges).

The fits. Table 16.8 gives each drawing a row, with the reference's alongside. The unrepaired stack fails, $X^2 = 1141.78$ on 940.44, tail 6.44×10^{-6}. With g_{B2} free per window the test passes, $X^2 = 900.55$ on 870.75, tail 0.235, but the normalized means stay off in hall A's rows: $T_{r,\mathrm{A}}$ at $+4.03$, Q_{A} at -5.06, Q_{loop} at $+2.72$ and the three discharges at -3.97 to -4.07. The signs have one mechanism, with residuals as predicted minus observed: a meter reading high pulls Q_{A} high, though less than the reading; the higher load predicts a warmer return and a higher loop heat than observed; and the coils of hall A rise to hold the temperatures, which predicts discharges cooler than observed. With the bias on the hall-A meter also free, $X^2 = 875.11$ on 869.75, tail 0.443; the bias is $+4.595 \pm 0.912$ kW; those means fall to $+1.33$, -0.46, -1.20 and -1.27 to -1.39, the meter's by construction, since its null spread there is only 0.047 with its bias entry free; and the evidence rises by $+10.14$ log units. Q_{B} at $+2.10$ remains. The method of ranking meters is that of Chapter 7 and of ranking coils that of Chapter 10, and Section 15.2.2 supplies these drawings on this plant.

What the omnibus statistics miss. The first row of Table 16.8 holds the same signs in the reference's own means, Q_{A} at -3.49 null spreads, and the reference still passes at 0.145. So a sensor's mean can read what the sum does not. The sum's failure is not caused by the reference holding the offset. On fresh nights with no coil event the mean of Q_t exceeds the

Table 16.8: The reference, the unrepaired stack and the repaired fits. Each sensor's mean residual, normalized by its null spread, predicted minus observed; $T_{d,\mathrm{A}}$ the three discharges of hall A.

Fit	X^2 [–]	Count [–]	Tail [–]	$T_{r,\mathrm{A}}$ [–]	$T_{d,\mathrm{A}}$ [–]	Q_{A} [–]	Q_{B} [–]	Q_{loop} [–]	$\ln p(\mathbf{y} \mid \cdot)$ [–]
Reference, windows 1–24	328.60	302.52	0.145	+2.87	−2.73 to −2.75	−3.49	−1.02	+1.79	–
g_{B2} shared, windows 1–73	1141.78	940.44	6.44×10^{-6}	+4.03	−4.03 to −4.12	−4.99	−0.98	+2.66	–
g_{B2} free, windows 1–73	900.55	870.75	0.235	+4.03	−3.97 to −4.07	−5.06	−0.71	+2.72	−831.12
g_{B2} and bias free, windows 1–73	875.11	869.75	0.443	+1.33	−1.27 to −1.39	−0.46	+2.10	−1.20	−820.98

no-offset mean by +0.29 with the offset from window 1, and by +0.43 with it from window 25 against a clean reference. A fresh +4 kW step against a clean reference has noncentrality 0.391, and its hour block has power 0.003 at the level. The cause is the size and the freedom: the offset is one meter spread, and Q_{A} is refound per window. The omnibus statistic cannot see a one-spread persistent offset, with or without a reference, and a per-sensor mean can. Nor is $z = +1.96$ at window 73 evidence of it: $z_{Q_{\mathrm{A}}}$ there has a mean of 0.49 with the offset alone, 0.91 with the coil alone and 1.31 with both.

16.3.6 What no test can see

Operator question. What can no test see? The tests above read the readings, and some of the state reaches none.

General form. The share is a variance split by entry, of the moments form, equation (2.10).

Specialized form. The question is equation (2.10) with $f(\mathbf{z}) = T_{\mathrm{B3}}(\mathbf{z})$ of the chain (12.4) at window 73, $p(\mathbf{z} \mid \mathbf{y})$ the one fit, and the variance $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ of equation (4.53) on the pattern's second line (16.1). The entry of $\mathbf{g}$ in the column of β_{B3} is the rise of hall B and that column's only row is its prior, so the term $[\mathbf{g}]^2_{\beta_{\mathrm{B3}}}\, 0.08^2$ has no cross term with any other entry.

Building blocks. $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$ of the one fit; the gradient; the share of one term of the budget.

The rack. The columns empty in every reading row are those of β_{A2}, β_{A4}, β_{A5}, β_{B1}, β_{B3} and β_{B5}. Rack B3's inlet at window 73 is 24.0747 ± 0.8293°C, with 96.83 % of its variance on β_{B3}, which stands at the survey's 0.5000 ± 0.0800. By the empty column X_t^2 and Q_t would take the same value at any β_{B3}, and so would every other statistic of this chapter: all of them could be perfect with this fraction wrong by any amount.

Common mistake

Reading a passing test as validation of an answer. A green light at window 73 says the reading rows are within their noise of some state; it does not say that the state is the one the plant was in at 19:00, and it does not reach the inlet of rack B3, which is the answer most of Part IV reads.

16.4 The checks (step 5)

The checks test the linear predictive against draws, the Gauss–Newton count against Newton's, the reduced state against the limit of exact rows, and the answers against the truth of the

constructed record.

16.4.1 The limit of exact rows and the curvature term

The chapter's rows run on the reduced state, which is the limit $\sigma_p \to 0$ of the full state with its physical rows (12.2). Table 16.9 repeats the fit of window 73 on the full state: each row a tolerance, the columns the return of hall B and the inlet of rack B3 with their spreads, and the iterations. At 0.001 the means match the reduced state's 29.17487 and 24.07466, and the spreads have settled by the same tolerance. The count taken on Newton's covariance, with the curvature term kept, is 3.6968 against 3.6914 on Gauss–Newton's, so the curvature does not move any statement of this chapter.

Table 16.9: The limit of exact physical rows on the full state at window 73, against the reduced state.

Tolerance σ_p [kW or K]	$T_{r,\mathrm{B}}$ [°C]	sd [K]	T_{B3} [°C]	sd [K]	Iterations [–]
0.100	29.17496	0.16868	24.07476	0.83532	3
0.010	29.17487	0.16858	24.07466	0.82936	3
0.001	29.17487	0.16858	24.07466	0.82930	2
Reduced	29.17487	–	24.07466	0.8293	–

16.4.2 The answers against draws and the stated truth

The reported tail of Q_t rests on a linear predictive and a chi-square. Beside it stand 200 000 draws at each window: the shared entries drawn from their Laplace normal, the plant solved on each at the window's own refound $\hat{\mathbf{q}}_t$, a fresh sensor error added, and the statistic formed with $\mathbf{q}_t$ removed linearly, as the statistic does. These test the linearity of the plant over $\mathbf{\Sigma}_{\mathrm{s}}$ and the chi-square, and nothing about the reference itself. Their means are 13.00 and 12.99. On 40 000 of them, with $\mathbf{q}_t$ re-solved on each, the tails are 1.500×10^{-4} (se 6.1×10^{-5}) at window 73 and 1.575×10^{-3} (se 2.0×10^{-4}) at window 49. Table 16.10 collects these gaps with the other checks: each row a quantity, the columns the reported value, what stands beside it and the gap. At window 73 the draws' tail lies within one standard error of the reported one. At window 49 every route lies above the level, the re-solved one by less than three of its standard errors, so that window's standing is close. The same table holds the truth of the constructed record, with the sensor bias removed at the true state, and the first verdict's range over fresh nights, which uses that truth and so is a check and no part of the shift log. Rack B3's fraction, which every test passed over, is 0.5000 on the fit and 0.57 in truth.

16.5 The report (step 6)

The light is green because the state can move, and it would be green whether or not the plant held: the window's test has a count of 3.691 from sixteen readings, and the derate at window 73 leaves 0.0257 of its 34.35 squared spreads on the rows that test adds up. The 19:00–21:00 model no longer fits, 38.90 on 13, tail 2.08×10^{-4} at window 73. A one-spread offset on a meter passes both tests, and a sensor's mean finds it. A passing test says nothing about the fraction of rack B3.

16.5.1 The answers and their bands

Table 16.11 gives the answers in the order of Table 16.1. Each row is a question; the second column gives the reported answer and the third what stands beside it.

Table 16.10: The checks: the reported value against draws, a second route and the stated truth of the constructed record.

Quantity	Reported [–]	Beside it [–]	Gap [–]
Tail of Q_{73}, against draws (se 3.1×10^{-5})	2.075×10^{-4}	1.900×10^{-4}	$+1.751 \times 10^{-5}$
Tail of Q_{49}, against draws (se 8.7×10^{-5})	1.538×10^{-3}	1.505×10^{-3}	$+3.287 \times 10^{-5}$
Count at window 73, against Newton	3.6914	3.6968	−0.0054
Stated truth			
g_{B2}, window 1	−0.0010	0	−0.0010
g_{B2}, window 37	−0.0063	−0.0033	−0.0030
g_{B2}, window 73	−0.1444	−0.1222	−0.0222
g_{B2}, window 109	−0.2649	−0.2411	−0.0238
g_{B2}, window 144	−0.3258	−0.3567	+0.0309
Hall-A meter's bias [kW]	$+4.595 \pm 0.912$	+4	+0.595
β_{B3}	0.5000 ± 0.0800	0.57	−0.07
T_{B3} at window 73 [°C]	24.0747 ± 0.8293	24.857	−0.782
X^2 of the readings at the true state, bias removed, against 16	16	13.260	−2.740
First verdict, block, this night against 1000 fresh nights' median [window]	57	58.0 (band 48.0–65.0)	–

Table 16.11: The report at window 73, 01:00. Later windows enter the second to fourth rows only to test the test.

Question	Answer, reported	Beside it
Is the model still describing the plant?	on the window's rows, yes: 0.755 on 3.691, tail 0.925; largest r_k^{N} −0.718 at $T_{d,\mathrm{B2}}$	9 of 144 tails below 0.05, 7.2 expected
Why has the light stayed green?	the test cannot see the change: $\lambda = 0.0257$, power 0.051 at window 73	as a test of the test, at window 144 power 0.063; on fresh nights 0.0483 against 0.0447 with no change
Is tonight's plant that of 19:00?	the 19:00–21:00 model no longer fits: $Q = 38.90$ on 13, tail 2.08×10^{-4}; largest z +5.25 at $T_{d,\mathrm{B2}}$	draws 1.90×10^{-4} (se 3.1×10^{-5})
When did that show?	hour block ending 23:40, single window 00:45	block median window 58.0, range 48.0–65.0 over 1000 fresh nights; 0.088 of nights give a block verdict with no coil event
What restores the fit?	g_{B2} free per window, with a bias of $+4.595 \pm 0.912$ kW on the hall-A meter: tail 0.443	without the bias, tail 0.235; evidence +10.14 log units for it
What can no test see?	six fractions; $T_{\mathrm{B3}} = 24.0747 \pm 0.8293$°C, 96.83 % of its variance on β_{B3}	the survey's 0.5000 ± 0.0800, untested

16.5.2 The entry in the shift log

The operator's entry as of window 73, 01:00, using nothing later, reads: "Model fit, 01:00. This window's fit passes (0.755 on 3.7), and that test cannot see a coil the state can follow. The 19:00–21:00 model no longer fits: 38.90 on 13, tail 2×10^{-4}; first verdict 23:40 by the hour, one night's draw; largest $T_{d,\mathrm{B2}}$ +5.25. With B2's coil free the fit passes best with a bias on the hall-A meter (+4.6 kW, tail 0.443; without it 0.235). Rack B3 cannot be tested from here. Action: not decided here."

16.6 The next reading (step 7)

A further reading is a further row, and what it does to the test is known before its value is: a new leverage, a new count and a new noncentrality under each stated fault.

16.6.1 The candidates

Table 16.12 gives what each candidate does at window 73: each row a candidate, the columns the count of the window's test with its row added and the power at a tail of 0.01 against three stated faults, by the noncentral chi-square on that count. A fault's noncentrality is the reading-row part of equation (16.7) with its shift placed on its own reading, and its power carries the caveat of Section 16.3.2: the chi-square on the count is approximate at a fixed state. Each candidate adds nearly one to the count except the probe at rack B3, 3.834 against 3.691, because that row's own empty column takes almost all of it. A permanent inlet sensor at B3, the instrument at the place of the question, adds the same row in this window, so its line is the probe's.

Table 16.12: Candidate readings at window 73: the count of the window's test and its power at a tail of 0.01 against three stated faults.

Candidate	Count [–]	Hall-A meter +4 kW [–]	B2 discharge +1 K [–]	Hall-B return +1 K [–]
None	3.691	0.014	0.027	0.346
Second return sensor, hall B, 0.3 K	4.681	0.014	0.033	0.360
Second discharge sensor, B2, 0.3 K	4.669	0.014	0.213	0.323
Water-side heat check, B2, 3 kW	4.675	0.014	0.121	0.311
Clamp-on, hall A load, 1 kW	4.689	0.026	0.024	0.308
Clamp-on, hall B load, 1 kW	4.687	0.014	0.026	0.355
Probe or permanent sensor at B3, 0.3 K	3.834	0.014	0.027	0.341

No single reading makes the window's test see the meter: the largest power against it is 0.026, with a clamp-on on the very load it reads, and the stack of windows in Section 16.3.5 is what reads it. A second discharge sensor is the largest single gain against a sensor fault, its noncentrality on the B2 discharge rising from 0.87 to 5.78 and its power from 0.027 to 0.213. The power against the return can fall when a candidate's row raises the count more than the noncentrality. The probe adds least to the count, and Chapter 18 ranks it for what it does to the entries, where a test of fit does not reach. Nothing in the table recommends an instrument on the strength of this test alone.

16.6.2 What the analysis showed

The screen had a green light and sixteen readings close to their fit, and the analysis finds a test that could not have turned red. Sixteen readings leave under four effective degrees of freedom

on this plant, and a change along a column goes almost entirely to the state, with only a small price on a prior row; that row, watched as an estimate, is what responds. The test that did respond asks whether the first two hours' model still fits: it rejects before 01:00 on nearly every fresh night, and at its level it also gives a verdict on some nights with no event. It says the model no longer fits, not why. It has a blind side of its own, because an offset of one spread on a meter passes it with or without a reference, and that took a sensor's mean in a repaired fit to read. And the rack that most of the night's decisions rest on lies outside every test here, so a clean statistic is no support for an answer read from that entry.

16.6.3 Limits of the model

The reference. Its hours are assumed free of any change except what the comparison cannot see, and one reference is used, not a rolling one; a slow drift across a rolling reference would meet the window test's difficulty again.

The linear predictive. The tail of Q_t uses a linear predictive and a chi-square, with the gaps of Section 16.4.2.

The freed entries. The supply and loads are refound in every window, so a change those three entries can express lies outside Q_t.

The fresh nights. They hold the true wander fixed: the first verdict's range spans sensor errors, not the whole spread from one night to another.

The level. The level of 10^{-3} is stated, per test, uncorrected across windows. It is not an alarm level.

The repaired drawings. The coil's place in them is taken from Section 15.2.2, by the method of Chapter 10, and a free entry passing a test does not by that show it is the cause.

Constructed readings. The record is constructed: generated from a stated truth of the plant with the stated sensor errors, and Section 16.4.2 compares the answers with that truth. It is not field data.

16.7 What to retain

Definition 16.1 (Two tests of fit)**.** The window's own test is $X_t^2 = \sum_k r_k(\hat{\mathbf{z}}_t)^2$ on $\nu_t = \sum_k (1 - h_{kk})$ over the reading rows, which asks whether some state explains the window. The reference test is $Q_t = \mathbf{e}_t^{\mathsf{T}} \mathbf{V}_t^{-1} \mathbf{e}_t$ on the readings less the entries refound per window, with $\mathbf{V}_t$ carrying the reference's $\mathbf{\Sigma}_{\mathrm{s}}$, which asks whether the state of a stated earlier stretch explains it. A change is seen only through the part of its shift the statistic's rows retain, equation (16.7).

Three conclusions carry beyond the plant. Effective degrees of freedom are read from the leverages, not from rows less entries, and they say how little room a window's own test has. A test's power against a stated change belongs in the report with its verdict, since a flat test and a flat plant look alike. And an omnibus test passes what is small against its spread or what its own entries can take: an offset of one spread, a fault a shared entry absorbs in its reference, an entry in no row. A sensor's mean or an entry watched as an estimate may read what the sum does not.

Table 16.13 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 16.13: The terms of the goodness-of-fit tests and where each is defined. Units for the shift plant.

Term	Expression	Units	Meaning	Defined in
Sparsity pattern	eq. (16.1)	–	rank 13 of 19; six empty columns	16.1.4
Projection	eq. (16.2)	–	leverages h_{kk} on 33 rows	16.2.1
Reference fit	eqs. (16.3), (16.4)	mixed	windows 1–24, shared coils and fractions	16.2.2
Repaired fits	g_{B2} per window; a bias	mixed	the model with the change admitted	16.2.3
Questions and answers				
Window test	X_t^2, ν_t, eq. (16.6)	–	some state explains the window	16.3.1
Split of a change	eq. (16.7)	–	state, reading rows, prior rows	16.3.2
Noncentrality, power	$\lambda = \lVert \boldsymbol{\ell}_{\mathrm{R}} \rVert^2$	–	what a test could see	16.3.2
Predictive covariance	eq. (16.8)	mixed	sensors plus the reference's spread	16.3.3
Reference test	Q_t, $z_{k,t}$, eq. (16.10)	–	the 19:00–21:00 model against the window	16.3.3
First verdict	first t past the level	–	one night's draw, with its spread	16.3.4
Evidence for the bias	$\ln p(\mathbf{y} \mid M_1) - \ln p(\mathbf{y} \mid M_0)$	–	what the repaired model prefers	16.3.5

16.8 Exercises

Exercise 16.1. *Application.* From the column $1 - h_{kk}$ of Table 16.4, compute the count of the window's test at window 73. (Adding the sixteen entries from 0.0483 to 0.1157, the count by equation (16.6) is 3.691.)

Exercise 16.2. *Derivation.* Compute the normalized residual of the discharge of unit B2 from its row and leverage in Table 16.4. (By equation (16.6), $r_k = -0.249$ and $1 - h_{kk} = 0.1203$, so $r_k^{\mathrm{N}} = -0.249/\sqrt{0.1203} = -0.718$.)

Exercise 16.3. *Application.* Why does the statistic of window 73 against the reference have 13 as its reference, and what is its tail at $Q_{73} = 38.90$? (Three of the sixteen rows go to the supply and the two loads, which are refound per window by equation (16.9), and $\Pr(\chi^2_{13} > 38.90) = 2.08 \times 10^{-4}$.)

Exercise 16.4. *Formulation.* State the two drawings compared in Section 16.3.5 and the evidence for the bias. (Both free g_{B2} per window over windows 1 to 73; the second adds a constant bias on the hall-A meter with prior $\mathcal{N}(0,\ 10^2)$ kW. The log ratio of equation (2.14) is $-820.98 - (-831.12) = +10.14$.)

Exercise 16.5. *Application.* With a second discharge sensor on unit B2 the count of the window's test is 4.669 and a +1 K fault on the first sensor has noncentrality 5.78 on the reading rows. Find the power at a tail of 0.01. (The noncentral chi-square on 4.669 with $\lambda = 5.78$, beyond the central chi-square's point at 0.01 on the same count, gives 0.213.)

Chapter 17

Sensor Loss and Measurement Redundancy

17.1 A work order for the return sensor

It is 01:00, window 73 of the shift, and a work order asks to take hall B's return sensor out for recalibration. It reads 29.15°C, and it is the one number on the screen called the hot-aisle temperature of hall B. The operator concludes that this is the sensor to protect, since without it the screen has no hot-aisle number at all, and that a discharge sensor can go first, since hall B has three of them and would still have two.

That reading counts instruments. It does not say what the hot-aisle temperature is known from, and the plant's other rows also bear on it. The operator's question is the one the screen cannot answer,

if this one sensor fails, how wide does the band on the hot-aisle temperature get, and which sensor is it that holds that band.

It is answered on the two-hall shift plant, from the one fit of Chapter 12. Nothing is added to the plant in this chapter; a row is taken away.

Jargon: hot-aisle temperature, sensor loss, down-date, importance measure, critical measurement, group loss, stuck-at failure

The *hot-aisle temperature* of a hall is its return-air temperature $T_{r,h}$, a function of the state. A *sensor loss* is the removal of one reading's row from the fit (in the framework, a row of $\mathbf{J}$ taken away). The *down-date* is the covariance with that row removed, at the fit's own state. An *importance measure* ranks the readings by what their loss does to one stated quantity; a ranking without its quantity means nothing. A *critical measurement* is a reading whose loss lowers the rank of the reading rows. A *group loss* is the removal of a named set of readings together. A *stuck-at failure* is a sensor that stays connected and goes on reporting one value.

Key idea

The information matrix is a sum with one term per row, so a reading's whole contribution to the posterior is one term, and removing it has a closed form: the covariance grows by a rank-one matrix divided by one minus the row's leverage. The band after a loss therefore depends on two numbers of that row at the fit, how strongly it is coupled to the question through $\boldsymbol{\Sigma}$, and how little the other rows could say in its place. Neither number is visible from a count of sensors.

17.1.1 A loss as a row taken away

Sensor placement asks what a row would add; this question is its reverse. Earlier chapters added rows to a posterior through the preposterior of equation (4.65), and that update needs no outcome of the reading, since a covariance of linear rows does not depend on the values read. Removal has the same property. The band after a loss can be read before the sensor is touched, from what the fit already holds.

A loss differs from a failure, and that distinction needs its own computation. A sensor that is out sends no row. A sensor that has failed and stays connected sends a row that the fit goes on trusting at its full weight. The first widens the band and says so. What the second does is a separate question, and it is asked below.

Principle 17.1. *What a sensor is worth depends on the quantity asked for and on what the other rows can say in its place. A reading that only it can supply and a reading that many rows repeat can sit side by side under the same label on the screen, and only the rows tell them apart.*

17.1.2 The shift plant with a row taken away (step 0)

The plant is the two-hall shift plant of Section 12.1.2: the physical rows (12.2), the pinned ratings, capacity rates and other heat, and the reduced state (12.3) at window 73, 01:00,

$$\mathbf{z}_{73} = (T_s,\ Q_{\mathrm{A}},\ Q_{\mathrm{B}},\ g_{\mathrm{A}1}, \dots, g_{\mathrm{B}3},\ \beta_{\mathrm{A}1}, \dots, \beta_{\mathrm{B}5})^{\mathsf{T}}, \tag{17.1}$$

on the reduced state of Chapter 4. The dependent block of the physical rows is triangular with a nonzero diagonal, hence invertible, and the temperatures follow from the state by the forward chain (12.4). The hot-aisle temperature of hall B is the first link of that chain,

$$T_{r,\mathrm{B}}(\mathbf{z}) = T_s + \frac{Q_{\mathrm{B}} + Q_{o,\mathrm{B}}}{\sum_{u \in \mathrm{B}} \bar{G}_u \, \mathrm{e}^{g_u}}, \tag{17.2}$$

so it is a derived quantity and not an entry. The return reading is one row that reads it directly. The air balance (12.5) gives a second way to it that involves no coil at all: the cold aisle is the mix of the discharges, so with the three units' equal capacity rates

$$T_{r,\mathrm{B}} = \frac{1}{3} \sum_{u \in \mathrm{B}} T_{d,u} + \frac{Q_{\mathrm{B}} + Q_{o,\mathrm{B}}}{\sum_{u \in \mathrm{B}} C_u}, \tag{17.3}$$

the heat balance of the hall. It holds at every point of the set, whatever the conductances are. It needs all four of its terms: the three discharges and the hall load.

Modeling choices at the graph. The choices of the shift plant stand as in Chapter 12: steady state within a window, pinned airflow and no active control on the coils, one effective conductance per coil, each failing as stated there. Nothing else is chosen at the graph, because a loss adds no node.

The graph. Figure 17.1 draws hall B's part of the plant with one row struck out. The load (green) feeds the return air (blue), the return and the chilled water (green) feed the three coils (green; the water's arrow is drawn to B1 alone, and to hall A), their discharge mixes in the cold aisle (blue), and the cold aisle feeds five racks. Dashed arrows are readings. The red cross marks the removed reading, the discharge of unit B1, and the red outlines mark what it reached: the log-derate $g_{\mathrm{B}1}$ and, through the balance (17.3), the return.

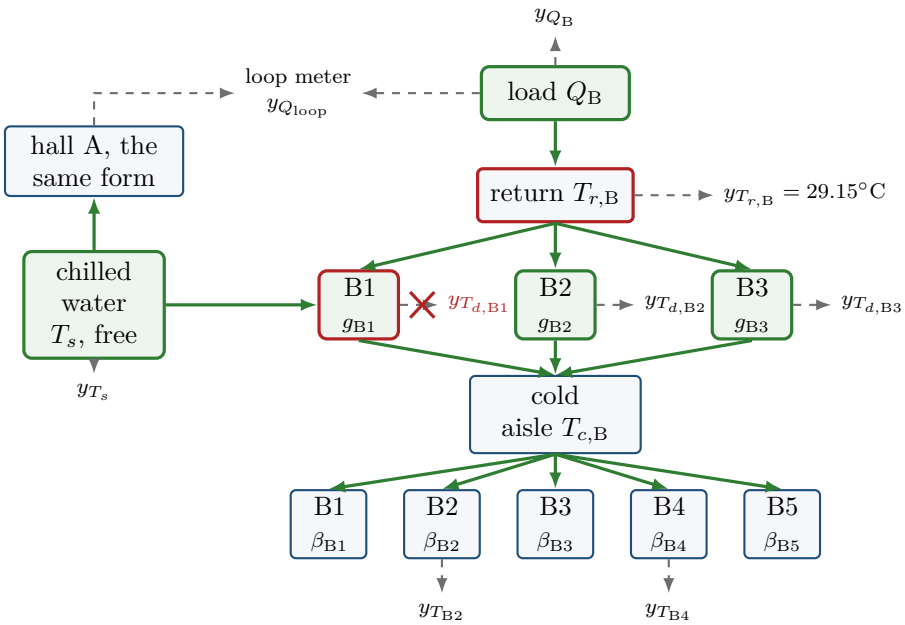


Figure 17.1: Hall B at window 73 with one reading removed. Free (green): the supply, which reaches all six coils, the load and the log-derates; the recirculation fractions are free entries carried by the rack nodes. Blue: temperatures fixed by the physical rows. Observed (dashed): the readings. Red cross: the removed row; red outlines: the entry it reads and the quantity asked for. Hall A has the same form and is not drawn. Priors are not drawn.

17.1.3 The operator's questions (step 1)

The operator asks six questions of window 73, and step 1 writes each over the state (17.1). Table 17.1 lists them: the question in the operator's words, the first of them the drill question; the general form of its category from Chapter 2; its mathematical form on this plant; its category; the computation of Section 2.3.6 that answers it; and the subsection where it is answered.

17.1.4 The evidence and the loss as rows (step 2)

The evidence is that of Section 12.1.4: sixteen readings and seventeen priors, 33 rows on 19 entries, each divided by its standard deviation. The supply reads to 0.1 K; the returns, discharges and inlets to 0.3 K; the hall meters to 4 kW and the loop meter to 8 kW. The priors are $g_u \sim \mathcal{N}(0,\ 0.15^2)$, $\beta_k \sim \mathcal{N}(\beta_k^{\mathrm{s}},\ 0.08^2)$ and $T_s \sim \mathcal{N}(14,\ 0.5^2)$; the loads have none. Three choices enter with the loss.

A lost sensor is a removed row. The reading sends nothing, so its term leaves the objective and the other rows stay as they are. This buys a closed form for the covariance after the loss. It fails for a sensor that stays on the wall reading wrong, which is a different row and not an absent one; Section 17.3.5 tests that case.

The band at the fit's own state. The reported band is the down-date at $\hat{\mathbf{z}}$. It needs no outcome and can be asked before the loss, which is when the operator asks it. It fails to the extent that removing the row would move the state and the gradients with it; the refit and the draws stand beside it to measure that.

Independent losses, one at a time, and groups stated. Each reading is removed alone, and a group is a named set the plant could lose together, such as one hall's discharges. This

Table 17.1: The questions asked of window 73, 01:00. $\mathbf{y}$ is the sixteen readings and the seventeen priors, $\mathbf{y}_{(k)}$ the same with reading k removed; $\mathbf{a}_k$ the row of $\mathbf{J}$ for reading k, h_{kk} its leverage; X^2 and e_k the window's test statistic and normalized residual.

Operator question	General form	Mathematical form	Category	Computation	Answered in
If this one sensor fails, how wide does my band on the hot-aisle temperature get? Remove its row from $\mathbf{\Lambda}$, re-read the band; the reverse of sensor placement.	(2.10)	$\mathbf{g}^\mathsf{T}(\mathbf{\Lambda} - \mathbf{a}_k^\mathsf{T}\mathbf{a}_k)^{-1}\mathbf{g}$	estimation	posterior	17.3.1
Which reading is carrying the answer?	(2.10)	the ranking over k; $\mathbf{g}^\mathsf{T}\mathbf{\Sigma}\mathbf{a}_k^\mathsf{T}$, h_{kk}	estimation	posterior	17.3.2
What if several go together?	(2.10)	the block deletion	estimation	posterior	17.3.3
Which reading has no replacement?	(2.10)	entries back at their prior; the rank; the effective degrees of freedom	estimation	posterior	17.3.4
What if it fails and stays connected?	(2.13)	X^2, e_k, $T_{r,\mathrm{B}}$ with the reading held at its last value	prediction under a threshold	posterior	17.3.5
What does a loss do to tonight's answers?	(2.12)	Chapter 12's joint, and its largest set point, with the row removed	prediction under a threshold	posterior	17.3.6

keeps the question to what a loss would do, and it needs no failure probability; none is assigned to any sensor, so nothing here says which loss is likely. It fails as a statement of risk, which it does not attempt.

Which row acts on which entry. With $*$ for an entry that can be nonzero, the sixteen reading rows of Figure 12.2 have, by group, the pattern

$$\begin{bmatrix} * & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ * & * & 0 & * & 0 & 0 & 0 & 0 \\ * & 0 & * & 0 & * & 0 & 0 & 0 \\ 0 & * & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & * & 0 & 0 & 0 & 0 & 0 \\ 0 & * & * & 0 & 0 & 0 & 0 & 0 \\ * & * & 0 & * & 0 & * & 0 & 0 \\ * & 0 & * & 0 & * & 0 & * & 0 \end{bmatrix} \begin{array}{c} \text{supply} \\ \text{return and discharges, hall A} \\ \text{return and discharges, hall B} \\ \text{meter } Q_\mathrm{A} \\ \text{meter } Q_\mathrm{B} \\ \text{loop meter} \\ \text{read inlets, hall A} \\ \text{read inlets, hall B} \end{array}, \tag{17.4}$$

$$\mathbf{g}_{T_{r,\mathrm{B}}}^\mathsf{T} = (*,\ 0,\ *,\ 0,\ *,\ 0,\ 0,\ 0),$$

with the columns grouped as T_s; Q_A; Q_B; the g_u of hall A; the g_u of hall B; the read fractions of hall A; the read fractions of hall B; the six unread fractions. The second line is the question's gradient, and it is not a row. Read as observability, the pattern says three things before any number. The question's gradient has its nonzeros in hall B's columns and the supply, so a row with nothing in those columns can reach it only through $\mathbf{\Sigma}$. Each read fraction has one reading in its column, so its inlet is the only reading that reads that fraction. And the six unread columns hold no reading to lose.

17.2 The posterior of the state (step 3)

This is the one fit of Part IV, Section 12.2.2, read once more; it is not fitted again here. What is new is the form of its information matrix, which turns a loss into one term, and every refit an answer sets beside the down-date.

17.2.1 The one fit and its rows as a sum

The MAP state $\hat{\mathbf{z}}$, the Jacobian and the covariance are those of Tables 12.5, 12.4 and 12.6, from the normal equations (12.9). Written with $\mathbf{a}_k$ for row k of $\mathbf{J}$ at $\hat{\mathbf{z}}$, a row vector, the information of the normalized rows is

$$\mathbf{\Lambda} = \mathbf{J}^{\mathsf{T}}\mathbf{J} = \sum_{k=1}^{33} \mathbf{a}_k^{\mathsf{T}}\mathbf{a}_k, \qquad \mathbf{\Sigma} = \mathbf{\Lambda}^{-1}, \tag{17.5}$$

so a reading's whole contribution is one term. With unnormalized rows the same sum is $\mathbf{J}^{\mathsf{T}}\mathbf{W}\mathbf{J}$.

The quantities the chapter reads. Table 17.2 gives them at the fit. Each row is a quantity; the columns give its value and its band, by equation (4.53). The hot-aisle gradient, from equation (17.2),

$$\frac{\partial T_{r,\mathrm{B}}}{\partial T_s} = 1, \qquad \frac{\partial T_{r,\mathrm{B}}}{\partial Q_{\mathrm{B}}} = \frac{1}{\sum_v G_v}, \qquad \frac{\partial T_{r,\mathrm{B}}}{\partial g_u} = -\frac{G_u\,(T_{r,\mathrm{B}} - T_s)}{\sum_v G_v}, \tag{17.6}$$

has the values $+1.0000$ on T_s, $+0.0499$ on Q_{B}, and -5.5461, -4.6226 and -5.1115 on g_{B1}, g_{B2} and g_{B3}, zero elsewhere, the pattern (17.4). With it $\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g} = 0.028419\,\mathrm{K}^2$, the square of the band in the first row.

Table 17.2: The quantities read at the one fit, window 73.

Quantity	Value	Band
$T_{r,\mathrm{B}}$ [°C; K]	29.1749	0.1686
$T_{r,\mathrm{A}}$ [°C; K]	29.8674	0.1698
T_{B3}, inlet, unread [°C; K]	24.0747	0.8293
T_{B4}, inlet, read [°C; K]	22.7126	0.2821
Q_{B2}, unit heat [kW]	92.5754	2.6545
Q_{B}, hall load [kW]	296.0123	3.4302

17.2.2 The deletion

Removing reading k removes one term of equation (17.5), and the inverse of the result follows from the Sherman–Morrison identity,

$$\mathbf{\Lambda}_{(k)} = \mathbf{\Lambda} - \mathbf{a}_k^{\mathsf{T}}\mathbf{a}_k, \qquad \mathbf{\Sigma}_{(k)} = \mathbf{\Sigma} + \frac{\mathbf{\Sigma}\mathbf{a}_k^{\mathsf{T}}\,\mathbf{a}_k\mathbf{\Sigma}}{1 - h_{kk}}, \qquad h_{kk} = \mathbf{a}_k\mathbf{\Sigma}\mathbf{a}_k^{\mathsf{T}}, \tag{17.7}$$

checked by multiplying out: $\mathbf{\Lambda}_{(k)}\mathbf{\Sigma}_{(k)} = \mathbf{I} - \mathbf{a}_k^{\mathsf{T}}\mathbf{a}_k\mathbf{\Sigma} + \mathbf{a}_k^{\mathsf{T}}\mathbf{a}_k\mathbf{\Sigma}\,(1 - h_{kk})/(1 - h_{kk}) = \mathbf{I}$. The number h_{kk} is the row's leverage of equation (7.12), read as in Chapter 9, and $1 - h_{kk}$ its redundancy number. The formula shows why that name is apt: as $1 - h_{kk}$ goes to zero the other rows have nothing to say in the row's place, and the covariance grows without limit along $\mathbf{\Sigma}\mathbf{a}_k^{\mathsf{T}}$. Several rows, stacked as a matrix $\mathbf{A}_S$, come out the same way,

$$\mathbf{\Sigma}_{(S)} = \mathbf{\Sigma} + \mathbf{\Sigma}\mathbf{A}_S^{\mathsf{T}}(\mathbf{I} - \mathbf{A}_S\mathbf{\Sigma}\mathbf{A}_S^{\mathsf{T}})^{-1}\mathbf{A}_S\mathbf{\Sigma}. \tag{17.8}$$

Each reading's two numbers. For a derived quantity with gradient $\mathbf{g}$, equation (17.7) gives

$$\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_{(k)}\mathbf{g} = \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} + \frac{c_k^2}{1-h_{kk}}, \qquad c_k = \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{a}_k^{\mathsf{T}}. \tag{17.9}$$

Table 17.3 gives the two numbers of each reading row for $T_{r,\mathrm{B}}$. Each row is a reading; the columns give its leverage, its redundancy number and its coupling c_k. The sign of c_k follows the correlations of Table 12.6: the rows of hall B's own quantities are positive, and hall A's return and meter are negative because the loop meter reads the sum of the loads: a reading that raises Q_{A} lowers Q_{B}. The sign has no effect on the band, since c_k enters squared, and it does say which way the reading's error would push the quantity.

Table 17.3: The rows as a sum: leverage, redundancy number and coupling to $T_{r,\mathrm{B}}$ of the sixteen readings at the one fit. The seventeen prior rows are not listed.

Reading	h_{kk} [–]	$1-h_{kk}$ [–]	c_k [K]	Reading	h_{kk} [–]	$1-h_{kk}$ [–]	c_k [K]
T_s	0.9517	0.0483	+0.00078	$T_{d,\mathrm{B3}}$	0.8840	0.1160	+0.06186
$T_{r,\mathrm{A}}$	0.3204	0.6796	−0.00443	Q_{A} meter	0.7371	0.2629	−0.01320
$T_{r,\mathrm{B}}$	0.3158	0.6842	+0.09473	Q_{B} meter	0.7354	0.2646	+0.07395
$T_{d,\mathrm{A1}}$	0.8949	0.1051	+0.00143	loop meter	0.3025	0.6975	+0.03038
$T_{d,\mathrm{A2}}$	0.8979	0.1021	+0.00145	T_{A1}	0.9191	0.0809	+0.00004
$T_{d,\mathrm{A3}}$	0.8956	0.1044	+0.00144	T_{A3}	0.9191	0.0809	−0.00006
$T_{d,\mathrm{B1}}$	0.8869	0.1131	+0.06186	T_{B2}	0.8843	0.1157	+0.00857
$T_{d,\mathrm{B2}}$	0.8797	0.1203	+0.06188	T_{B4}	0.8843	0.1157	+0.00880

17.2.3 The refits the answers use

Each of these repeats the normal equations on changed rows, from the plant's own starting state, and is computed at window 73 unless it says otherwise. A *refit* removes the rows of a set S and iterates

$$(\mathbf{J}_{(S)}^{\mathsf{T}}\mathbf{J}_{(S)})\,\Delta\mathbf{z} = -\mathbf{J}_{(S)}^{\mathsf{T}}\mathbf{r}_{(S)}(\mathbf{z}), \tag{17.10}$$

to its own $\hat{\mathbf{z}}_{(S)}$ and $\boldsymbol{\Sigma}_{(S)}$, for every single reading and for the groups of Section 17.3.3. It uses the reading values of the other rows, so unlike the down-date it can be computed only with the night's values in hand, and it is one more Laplace fit and not a full calculation. The *exact posterior with a row removed* is set beside it by importance sampling: 200 000 draws from $\mathcal{N}(\hat{\mathbf{z}}_{(k)}, \boldsymbol{\Sigma}_{(k)})$ weighted by $\exp(-\mathcal{L}_{(k)})$ over the normal density. A *stuck fit* keeps all rows and replaces one reading by its value at window 49, $y_{k,t} \to y_{k,49}$, at window t; it fits each of windows 49, 61, 73, 97, 121 and 144 on its own. What a stuck reading does develops across windows, which one window cannot show, so these fits leave the one fit, and they are repeated over 200 nights with the truth held and the sensor errors redrawn, at windows 97, 121 and 144. Two further runs test the route of equation (17.3): the coil priors widened to a spread of 10^4, and the plant refitted with its ratings or its capacity rates changed, with readings at the plant's own prediction.

17.3 The answers (step 4)

Every answer is read from the pair $(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of Section 12.2.2 through the deletion, and that value is the one reported, except the stuck sensor's, which has no deletion and reads the stuck fits. The refit and the draws stand beside it.

17.3.1 One sensor lost

Operator question. If this one sensor fails, how wide does my band on the hot-aisle temperature get? Remove its row from $\mathbf{\Lambda}$, re-read the band; the reverse of sensor placement.

General form. An estimation question, whose location and spread are the moments of equation (2.10), under the posterior given the evidence that remains.

Specialized form. The question is equation (2.10) with $f(\mathbf{z}) = T_{r,\mathrm{B}}(\mathbf{z})$ of equation (17.2), $\mathbf{y}$ replaced by $\mathbf{y}_{(k)}$, and $p(\mathbf{z} \mid \mathbf{y}_{(k)})$ approximated by $\mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma}_{(k)})$ at the fit's own state, so that by equation (17.9) the band is

$$\mathrm{band}_{(k)} = \sqrt{\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g} + \frac{c_k^2}{1 - h_{kk}}}, \tag{17.11}$$

with $\mathbf{g}$ of equation (17.6). The second term cannot be negative: at a fixed point, removing a row never narrows a band.

Building blocks. $\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}$ and the gradient (Section 17.2.1); c_k and $1 - h_{kk}$ (Table 17.3); the refit (17.10).

The worked row. For the discharge of unit B1, by equation (17.11),

$$0.028419 + \frac{0.06186^2}{0.1131} = 0.062255\ \mathrm{K}^2, \qquad \mathrm{band} = 0.2495\,\mathrm{K}.$$

All the readings. Table 17.4 gives every reading, sorted by the reported band. Each row is a removed reading; the columns give the band on $T_{r,\mathrm{B}}$ at $\hat{\mathbf{z}}$, its ratio to the band with all rows, the refit's band, how far the refit moves the value of $T_{r,\mathrm{B}}$, and the band and ratio on $T_{r,\mathrm{A}}$. Figure 17.2 draws the same results around the value: the horizontal axis is the temperature, each line is one removed reading, its bar is the fit's value ± one reported band, the small circle is the refit's value, and the dotted line is the stated truth, a check marker for Section 17.4.3.

Table 17.4: One reading removed, sorted by the reported band on $T_{r,\mathrm{B}}$. With all rows the bands are 0.1686 K on $T_{r,\mathrm{B}}$ and 0.1698 K on $T_{r,\mathrm{A}}$.

Removed	$T_{r,\mathrm{B}}$ band, at $\hat{\mathbf{z}}$ [K]	Ratio [–]	Refit band [K]	Value moves [K]	$T_{r,\mathrm{A}}$ band, at $\hat{\mathbf{z}}$ [K]	Ratio [–]
$T_{d,\mathrm{B1}}$	0.2495	1.480	0.2495	+0.0023	0.1699	1.000
$T_{d,\mathrm{B3}}$	0.2478	1.470	0.2481	−0.0521	0.1699	1.000
$T_{d,\mathrm{B2}}$	0.2454	1.456	0.2465	−0.1270	0.1699	1.000
Q_{B} meter	0.2215	1.314	0.2215	−0.0862	0.1718	1.011
$T_{r,\mathrm{B}}$, the sensor itself	0.2038	1.209	0.2038	+0.0132	0.1699	1.000
loop meter	0.1725	1.023	0.1725	+0.0249	0.1737	1.023
T_{B4}	0.1706	1.012	0.1706	+0.0117	0.1698	1.000
Q_{A} meter	0.1705	1.012	0.1705	+0.0052	0.2231	1.314
T_{B2}	0.1705	1.011	0.1704	−0.0028	0.1698	1.000
$T_{r,\mathrm{A}}$	0.1687	1.001	0.1687	+0.0023	0.2060	1.213
$T_{d,\mathrm{A2}}$	0.1686	1.000	0.1686	+0.0031	0.2606	1.534
$T_{d,\mathrm{A3}}$	0.1686	1.000	0.1686	+0.0017	0.2589	1.524
$T_{d,\mathrm{A1}}$	0.1686	1.000	0.1686	+0.0014	0.2584	1.522
T_s	0.1686	1.000	0.1686	+0.0016	0.1699	1.000
T_{A3}	0.1686	1.000	0.1686	+0.0001	0.1712	1.008
T_{A1}	0.1686	1.000	0.1686	−0.0000	0.1710	1.007

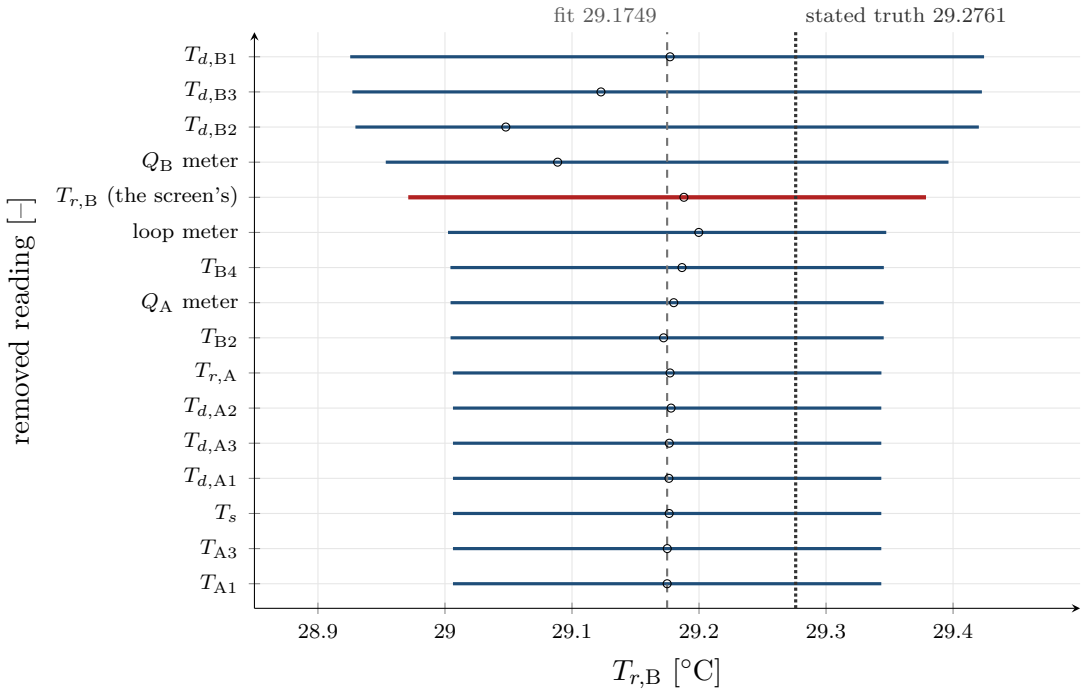


Figure 17.2: The hot-aisle temperature of hall B with each reading removed in turn. Bars: the fit's value ± one reported band, at $\hat{\mathbf{z}}$; the red bar is the sensor the screen names. Circles: the refit's value. Dashed: the fit with all rows. Dotted: the stated truth.

The answer. If the screen's sensor fails, the band goes from 0.1686 to 0.2038 K, a ratio of 1.209. If the discharge of B1 fails, it goes to 0.2495 K, a ratio of 1.480, the largest of any one loss. The refit's band matches the down-date to the fourth decimal for thirteen of the sixteen readings and differs by at most 0.0011 K, at the discharge of B2, which is also where the value moves most, by −0.1270 K.

17.3.2 The reading that carries it

Operator question. Which reading is carrying the answer?

General form. The moments of equation (2.10), compared over the readings, an importance measure for one stated quantity.

Specialized form. The question is equation (2.10) with $f = T_{r,\mathrm{B}}$ under each $\mathbf{y}_{(k)}$ in turn, ranked by the added variance of equation (17.9). The ranking belongs to this quantity; a different f has a different $\mathbf{g}$ and therefore a different c_k.

Building blocks. c_k and $1 - h_{kk}$ (Table 17.3); the bands of Table 17.4.

The ranking. The hot-aisle sensor is fifth for its own quantity, behind the three discharges of hall B and the hall-B meter. Figure 17.3 shows why. The horizontal axis is the redundancy number $1 - h_{kk}$ and the vertical axis the coupling c_k; each point is a reading, and the curve joins the points with the same added variance as the discharge of B1, 0.033837 K^2.

The three discharges have nearly the same coupling, +0.06186, +0.06188 and +0.06186, and small redundancy numbers, 0.1131, 0.1203 and 0.1160. The return's coupling is larger, +0.09473, but its redundancy number is 0.6842: the physics reads it elsewhere. Its row is the

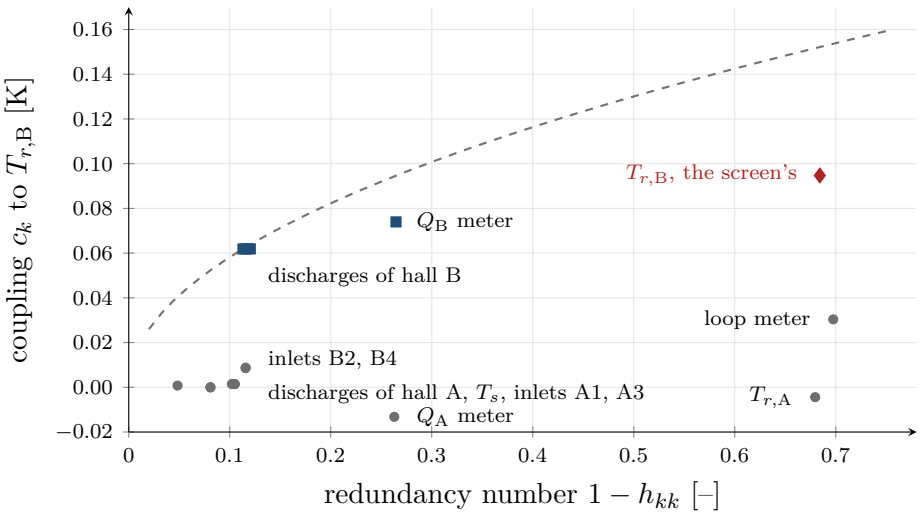


Figure 17.3: Each reading's two numbers for $T_{r,\mathrm{B}}$ (Table 17.3). Dashed: $c = \sqrt{0.033837\,(1-h)}$, the readings on it adding the same variance as the discharge of B1. The added variance grows toward the upper left.

question's own gradient over 0.3 K, so its numbers are $c_k = \mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}/0.3$ and $h_{kk} = \mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}/0.3^2$. Its added variance is $0.013115\,\mathrm{K}^2$ against 0.033837 for the discharge of B1.

The reason is the balance (17.3). Three discharges of 0.3 K and one meter of 4 kW read $T_{r,\mathrm{B}}$ through it to $\sqrt{0.3^2/3 + (4/30)^2} = 0.2186\,\mathrm{K}$, closer than the return's own 0.3 K: averaging three sensors beats one, and that is why the return is fifth, with the meter, the balance's fourth term, just ahead of it. Figure 17.4 draws the routes: from the left, the discharges and the meter through the balance; from the top, the return reading; from the right, the read inlets with the survey. Their precisions add. The total of $35.19\,\mathrm{K}^{-2}$ is 11.11 from the return, 20.96 from the balance (the route's own 0.2186 K gives 20.93), 2.56 from the inlets and 0.555 from the coils' priors.

The fifth place is not a property of this one window: at windows 1, 13, 37, 49, 61, 73, 97, 121 and 144 the return is fifth, behind the same four. Which discharge is first does depend on the window: B1 at window 73, B2 at windows 1 and 13.

The couplings show that the coils are not what ties the three discharges to the band. They stay equal, at +0.0619, with the ratings changed to 11, 7 and 3 kW/K, and they follow the capacity rates, +0.0885, +0.0594 and +0.0302 at 15, 10 and 5 kW/K. With the coil priors at a spread of 10^4 the band with all rows barely moves, 0.1689 K, and the discharges and the meter alone give 0.2186 K, the route's formula. Removing any one discharge then gives 0.2701 K for each: the balance has lost a term, and only the coil ratings could say how the heat divides among the units to replace it. At the plant's priors of 0.15 the ratings do that in part, which is the difference between that 0.2701 and the reported 0.2495 K.

Between the halls. The discharges of hall A move hall A's band and leave hall B's alone: their ratios are 1.522 to 1.534 on $T_{r,\mathrm{A}}$ and 1.000 on $T_{r,\mathrm{B}}$, with couplings of +0.00143 to +0.00145. The channel between the halls is the loop meter. With it gone, the coupling of $T_{d,\mathrm{A1}}$ falls from +0.001431 to +0.000003; with T_s held hard and the loop meter kept it stays at +0.001427. The supply passes essentially nothing, its correlation with $T_{r,\mathrm{B}}$ being 0.0048. The transfer number is the hall-A meter's ratio on $T_{r,\mathrm{B}}$, 1.012.

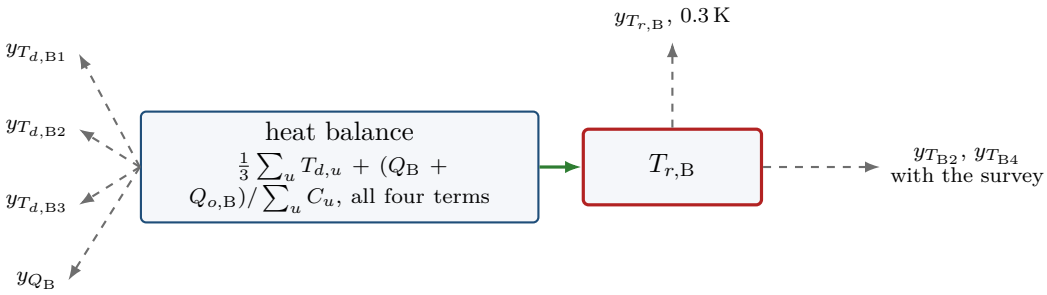


Figure 17.4: The three routes to the hot-aisle temperature, whose precisions add. Left: the three discharges and the hall meter through the heat balance (17.3), which involves no coil and needs all four. Top: the return reading. Right: the read inlets, whose fractions rest on them and the survey. Other rows, and the coils' priors, are not drawn.

Common mistake

Protecting the sensor that carries the quantity's name. A sensor's label says what it reads, and the band depends on what the others could say in its place. On this plant a discharge sensor, of which the screen shows three, is worth more to the hot-aisle band than the hot-aisle sensor. That ranking holds for $T_{r,\mathrm{B}}$ and for no other quantity by itself.

17.3.3 Groups lost together

Operator question. What if several go together?

General form. The moments of equation (2.10), under the evidence that remains after a named set is removed.

Specialized form. The question is equation (2.10) with $f = T_{r,\mathrm{B}}$, $\mathbf{y}$ replaced by $\mathbf{y}_{(S)}$ and the Laplace covariance replaced by the block deletion (17.8), at $\hat{\mathbf{z}}$. The sets are chosen for what the plant could lose together, such as a hall's temperatures on one data concentrator, and they carry no probability.

Building blocks. The rows $\mathbf{A}_S$ of each set; $\boldsymbol{\Sigma}$ and $\mathbf{g}$; the refit (17.10).

The sets. Table 17.5 gives the sets: each row a set, the columns the band on $T_{r,\mathrm{B}}$ at $\hat{\mathbf{z}}$, the refit's band, the refit's value, and the bands on T_{B3} and Q_{B2} at $\hat{\mathbf{z}}$. Figure 17.5 draws the band of each set as a bar, with the refit's band as a circle and the band with all rows dashed. The two meter groups give the same numbers because, with the loop meter gone, the hall-A meter has nothing left that reaches hall B: its coupling is -0.000002.

Series and parallel. The added variance of the three discharges of hall B, each removed alone and summed, is $0.09863\,\mathrm{K}^2$; removed together it is 0.04188, a ratio of 0.425. This is not a rule about groups. The three are in series on one route: losing any one already breaks the balance, and the next two add little, band 0.2495 with one lost and 0.2651 K with all three. Readings on parallel routes do the opposite, since each was standing in for the other: the hall-B and loop meters, both reading the load, have a ratio of 2.640 beside that 0.425; the return and the discharges 2.619; every hall-B temperature 15.70. The ratio for the three discharges also depends on what else stands. It is 0.618 with the return already gone and 1.825 with the return and the inlets gone.

Table 17.5: Groups lost together, at window 73. With every hall-B reading and the loop meter gone there is no fit: Q_{B} has no row left, reading or prior.

Set removed	$T_{r,\mathrm{B}}$ band, at $\hat{\mathbf{z}}$ [K]	Refit band [K]	Refit value [°C]	T_{B3} band [K]	Q_{B2} band [kW]
Hall B discharges	0.2651	0.2651	29.0988	0.8586	11.882
Hall B return and its discharges	0.5666	0.5647	28.9288	0.9893	11.891
Hall B inlets	0.1725	0.1725	29.1840	0.8301	2.655
Hall B return and inlets	0.2109	0.2109	29.2027	0.8362	2.685
All three heat meters	0.2940	0.2940	29.1600	0.8333	4.140
Hall B meter and loop meter	0.2940	0.2940	29.1600	0.8333	4.140
Every hall B temperature	1.3429	1.2779	28.4661	1.5655	11.917
Every hall B reading and the loop meter	no fit				

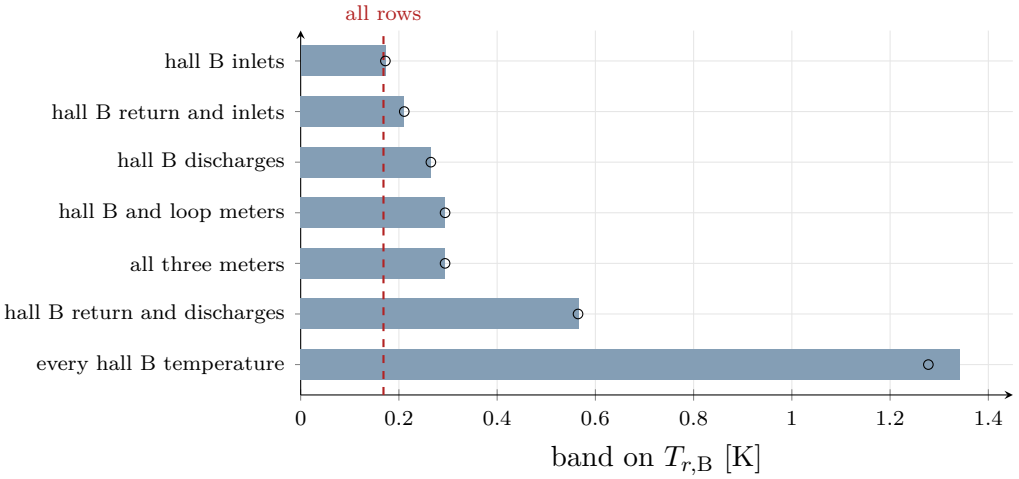


Figure 17.5: The band on $T_{r,\mathrm{B}}$ with each named group removed (Table 17.5). Bars: the block deletion at $\hat{\mathbf{z}}$; circles: the refit; dashed: all rows.

Losing the return as well. What holds the band at 0.2651 K is the return reading. With the return and the discharges gone, the band is 0.5666 K, from the inlets and the survey with the coils' priors, which are 0.178 of that precision. Those two routes combine by adding their precisions: a reading of 0.3 K with a route of 0.5666 K gives 0.2651 K, the group's band. So the return is the reserve behind the discharges, which is why their group is cheap. The band with every hall-B temperature gone is 1.3429 K, at the coils' priors; there the state moves far enough for the refit to differ visibly, 1.2779 K, and its value falls to 28.4661°C.

The work order's next question. With the return out for recalibration, a discharge is no longer a cheap loss. Taking out the discharge of B1 as well takes the band from 0.2038 to 0.4494 K, because the return's absence lowers that discharge's redundancy number from 0.1131 to 0.0509.

17.3.4 The reading with no replacement

Operator question. Which reading has no replacement?

General form. The moments of equation (2.10), here of the entries of the state and of the virtual sensors they carry.

Specialized form. The question is equation (2.10) with f an entry z_i or a derived quantity and $\mathbf{\Sigma}_{(k)}$ of equation (17.7) as the covariance. With a prior row behind an entry, a reading's leverage cannot reach one, so $1 - h_{kk}$ does not decide it. A critical measurement is one whose removal lowers the rank of the reading rows; a band that grows a great deal with the rank unchanged is called here large growth. Beside the rank stand each entry's band against its prior width and the effective degrees of freedom (7.14), called below the count, as used on this plant in Chapter 16, each with the reading removed.

Building blocks. $\mathbf{\Sigma}_{(k)}$; the reading rows of $\mathbf{J}$; $\sum_k (1 - h_{kk})$.

The leverages. The largest leverage is the supply's, 0.9517; hall A's inlets follow at 0.9191, hall A's discharges at 0.8949 to 0.8979, the discharge of B1 at 0.8869 and hall B's inlets at 0.8843. None reaches one.

The entries. Table 17.6 gives the critical measurements and two cases of large growth. Each row is a removed reading; the columns give the entry it mainly reads, that entry's band before and after, and its prior spread. Table 17.7 gives the wider picture: each row a removed reading, each column a quantity or an entry, each cell a band at $\hat{\mathbf{z}}$. Only the fractions come back to their prior. Losing the inlet at rack B4 takes β_{B4} from 0.0303 to 0.0800, the survey's spread, and the band on that rack's inlet from 0.2821 to 0.8293 K, close to the band of rack B3: a read rack becomes an unread one, and no other row can replace it. The two print alike and are not equal, 0.829272 and 0.829299: each is mostly 0.08 times the same rise, 10.2004 K. Losing the discharge of B2 is large growth and not critical: it takes g_{B2} from 0.0290 to 0.0800 (0.079962, no relation to the survey's 0.08), inside its prior of 0.15, and the unit's heat from 2.6545 to 6.4010 kW. The other discharges, the meters and the return together bound that entry: with the return gone as well it is 0.1143, with the discharge of B1 0.1197, with both the hall-B and loop meters 0.1350, and 0.1500 only with the return, the inlets and the other discharges gone. Losing the supply reading takes T_s from 0.0976 to 0.4438 K, close to its prior of 0.50, and leaves the band on $T_{r,\mathrm{B}}$ at a ratio of 1.000: the hot-aisle temperature, read against the other temperatures, does not need the absolute. The set-point answer below does. Losing the hall-B meter, large growth too, takes Q_{B} from 3.4302 to 6.6683 kW, and it has no prior to come back to.

Table 17.6: Entries after a loss, down-date at $\hat{\mathbf{z}}$. The first four and T_s are critical measurements, and the four fractions go back to the survey's spread. $T_{d,\mathrm{B2}}$ and the Q_{B} meter are large growth.

Removed	Entry	Band, all rows [entry's]	Band after [entry's]	Prior spread [entry's]
T_{B4}	β_{B4} [–]	0.0303	0.0800	0.08
T_{B2}	β_{B2} [–]	0.0303	0.0800	0.08
T_{A1}	β_{A1} [–]	0.0254	0.0800	0.08
T_{A3}	β_{A3} [–]	0.0254	0.0800	0.08
$T_{d,\mathrm{B2}}$	g_{B2} [–]	0.0290	0.0800	0.15
T_s	T_s [K]	0.0976	0.4438	0.50
Q_{B} meter	Q_{B} [kW]	3.4302	6.6683	none

Table 17.7: What else moves: bands at $\hat{\mathbf{z}}$ with each reading removed. Temperatures in K, heats in kW, g_u and β dimensionless.

Removed	$T_{r,\mathrm{B}}$ [K]	$T_{r,\mathrm{A}}$ [K]	T_{B3} [K]	T_{B4} [K]	Q_{B2} [kW]	Q_{B} [kW]	T_s [K]	g_{B1} [–]	g_{B2} [–]	g_{B3} [–]	β_{B4} [–]
none	0.1686	0.1698	0.8293	0.2821	2.6545	3.4302	0.0976	0.0249	0.0290	0.0267	0.0303
T_s	0.1686	0.1699	0.8293	0.2821	2.6547	3.4308	0.4438	0.0375	0.0402	0.0386	0.0303
$T_{r,\mathrm{A}}$	0.1687	0.2060	0.8293	0.2821	2.6554	3.4368	0.0976	0.0249	0.0290	0.0267	0.0303
$T_{r,\mathrm{B}}$	0.2038	0.1699	0.8347	0.2823	2.6840	3.6314	0.0976	0.0252	0.0292	0.0269	0.0313
$T_{d,\mathrm{A1}}$	0.1686	0.2584	0.8293	0.2821	2.6551	3.4346	0.0976	0.0249	0.0290	0.0267	0.0303
$T_{d,\mathrm{A2}}$	0.1686	0.2606	0.8293	0.2821	2.6552	3.4349	0.0976	0.0249	0.0290	0.0267	0.0303
$T_{d,\mathrm{A3}}$	0.1686	0.2589	0.8293	0.2821	2.6551	3.4347	0.0976	0.0249	0.0290	0.0267	0.0303
$T_{d,\mathrm{B1}}$	0.2495	0.1699	0.8531	0.2832	3.6749	3.5659	0.0976	0.0712	0.0329	0.0297	0.0351
$T_{d,\mathrm{B2}}$	0.2454	0.1699	0.8516	0.2831	6.4010	3.5478	0.0976	0.0271	0.0800	0.0294	0.0348
$T_{d,\mathrm{B3}}$	0.2478	0.1699	0.8525	0.2831	3.6452	3.5583	0.0976	0.0272	0.0328	0.0750	0.0350
Q_{A} meter	0.1705	0.2231	0.8293	0.2821	2.6761	3.5797	0.0976	0.0250	0.0291	0.0267	0.0303
Q_{B} meter	0.2215	0.1718	0.8307	0.2821	3.2619	6.6683	0.0976	0.0261	0.0311	0.0282	0.0304
loop meter	0.1725	0.1737	0.8294	0.2821	2.6975	3.7229	0.0976	0.0250	0.0292	0.0268	0.0303
T_{A1}	0.1686	0.1710	0.8293	0.2821	2.6545	3.4302	0.0976	0.0249	0.0290	0.0267	0.0303
T_{A3}	0.1686	0.1712	0.8293	0.2821	2.6545	3.4302	0.0976	0.0249	0.0290	0.0267	0.0303
T_{B2}	0.1705	0.1698	0.8297	0.2821	2.6545	3.4303	0.0976	0.0250	0.0291	0.0267	0.0304
T_{B4}	0.1706	0.1698	0.8297	0.8293	2.6546	3.4306	0.0976	0.0250	0.0291	0.0267	0.0800

The rank and the count. The reading rows have rank 13, the thirteen observed columns of Figure 12.2. The rank falls to 12 if the supply or any one inlet is removed, and it stays at 13 for every other reading: those five are the plant's critical measurements, each losing one direction to the priors. The count falls too. It is 3.691 with all readings; removed one at a time, a return, a discharge or a meter takes it to 2.693–3.042, and an inlet or the supply to 3.549–3.640. So a loss also leaves fewer degrees of freedom to test the remaining readings.

17.3.5 The sensor stuck on the wall

Operator question. What if it fails and stays connected?

General form. The predictive of equation (2.13), whose tail is the window's test.

Specialized form. The question is equation (2.13) with y_{new} the reading rows of a window, $p(\mathbf{z} \mid \mathbf{y})$ the fit of that window, and the reference that of Chapter 7: the statistic X^2 of the reading rows against χ^2_ν on the count (7.14), and the normalized residual of equation (7.15), written e_k here for r_k^{N}. One choice enters here: the failure mode is stuck-at, the reading held at its value of window 49, 23:00. This buys a failure the rows can show, and it fails to cover a sensor that drifts or goes noisy. Here the row stays with its full weight, so the band does not change at all. By the hat matrix in equation (7.12), for linear rows, a departure of b standard deviations in reading k moves its normalized residual and the quantity by

$$\mathbb{E}[e_k] = -b\,\sqrt{1 - h_{kk}}, \qquad \delta f = c_k\, b, \tag{17.12}$$

with b signed, positive for a reading above its truth, and c_k of equation (17.9), and the statistic X^2 gains about $b^2(1 - h_{kk})$. What a loss costs is $c_k^2/(1 - h_{kk})$; what a stuck reading shows is $b^2(1 - h_{kk})$, and that depends on the size of the departure. A small $1 - h_{kk}$ alone settles neither: the supply has the smallest and costs a ratio of 1.000, since its c_k is small.

Building blocks. The stuck fits and fresh nights of Section 17.2.3; c_k and $1 - h_{kk}$ (Table 17.3).

The six windows. Table 17.8 gives four sensors at the six windows run, and these statements cover only those windows. Each row is a stuck sensor at a window; the columns give X^2, the count, the upper tail, the sensor's own e_k, the reading with the largest $|e|$, $T_{r,\mathrm{B}}$ and its band, and the stated truth, which is a check and which the fit cannot see. At window 49 the reading is its own value, so all four sensors give the same fit, and the table prints that window for the first two only. Figure 17.6 draws two of them: the left panels have the tail against the window with 0.05 dashed, the right panels the fit's $T_{r,\mathrm{B}}$ with its band and the stated truth dotted, the upper row for the return sensor stuck and the lower for the discharge of B1.

The return stuck. The test catches it. The tail is 0.0006 at window 97 and 0.0000 at window 144, and the sensor's own e_k reaches +4.16 and +6.13 and names it. The band stays at 0.169 K while $T_{r,\mathrm{B}}$ reads 30.765°C at window 144 against a stated truth of 31.465.

The discharge of B1 stuck. On this night the test does not catch it. The tail is past 0.05 at none of the six windows, the nearest being 0.0616 at window 121, and the largest own e_k is +2.28. The band is the same, and $T_{r,\mathrm{B}}$ reads 31.149°C at window 144 against 31.465.

Table 17.8: Sensors stuck from window 49, at the six windows run. The truth column is the stated truth, a check.

Stuck sensor, value	Window [–]	X^2 [–]	ν [–]	Tail [–]	Own e_k [–]	Largest $\lvert e\rvert$	$T_{r,\mathrm{B}}$ (band) [°C]	Truth [°C]
$T_{r,\mathrm{B}}$, 29.25°C	49	6.35	3.71	0.1491	−1.75	loop meter	28.811 (0.168)	28.772
	61	2.52	3.70	0.5940	−0.56	Q_{B} meter	29.106 (0.169)	29.148
	73	0.74	3.69	0.9282	−0.16	loop meter	29.206 (0.169)	29.276
	97	18.86	3.66	0.0006	+4.16	$T_{r,\mathrm{B}}$	30.277 (0.169)	30.519
	121	12.80	3.68	0.0094	+2.75	Q_{B} meter	29.927 (0.168)	30.472
	144	39.70	3.65	0.0000	+6.13	$T_{r,\mathrm{B}}$	30.765 (0.169)	31.465
$T_{d,\mathrm{B1}}$, 17.28°C	49	6.35	3.71	0.1491	+1.88	loop meter	28.811 (0.168)	28.772
	61	2.84	3.70	0.5368	+1.21	Q_{B} meter	28.857 (0.168)	29.148
	73	1.01	3.69	0.8817	+0.89	$T_{d,\mathrm{B1}}$	29.013 (0.169)	29.276
	97	0.68	3.65	0.9339	+0.82	$T_{d,\mathrm{B1}}$	30.305 (0.169)	30.519
	121	8.46	3.67	0.0616	+2.28	loop meter	30.077 (0.169)	30.472
	144	4.09	3.63	0.3405	+1.91	$T_{d,\mathrm{B1}}$	31.149 (0.169)	31.465
$T_{d,\mathrm{B2}}$, 19.50°C	61	2.28	3.70	0.6384	−0.51	Q_{A} meter	29.103 (0.169)	29.148
	73	0.76	3.69	0.9245	−0.14	loop meter	29.073 (0.169)	29.276
	97	0.64	3.65	0.9408	+0.19	$T_{d,\mathrm{A3}}$	30.286 (0.169)	30.519
	121	13.67	3.66	0.0063	+2.40	$T_{d,\mathrm{B1}}$	29.811 (0.169)	30.472
	144	10.67	3.62	0.0231	+2.64	$T_{d,\mathrm{B1}}$	30.752 (0.169)	31.465
T_{B4}, 22.44°C	61	2.19	3.70	0.6544	+0.50	Q_{A} meter	29.032 (0.169)	29.148
	73	0.80	3.69	0.9178	+0.71	$T_{d,\mathrm{B2}}$	29.168 (0.169)	29.276
	97	2.16	3.65	0.6536	+1.84	$T_{d,\mathrm{B2}}$	30.628 (0.169)	30.519
	121	6.26	3.67	0.1511	+1.89	loop meter	30.286 (0.169)	30.472
	144	3.65	3.63	0.3983	+2.59	T_{B4}	31.400 (0.169)	31.465

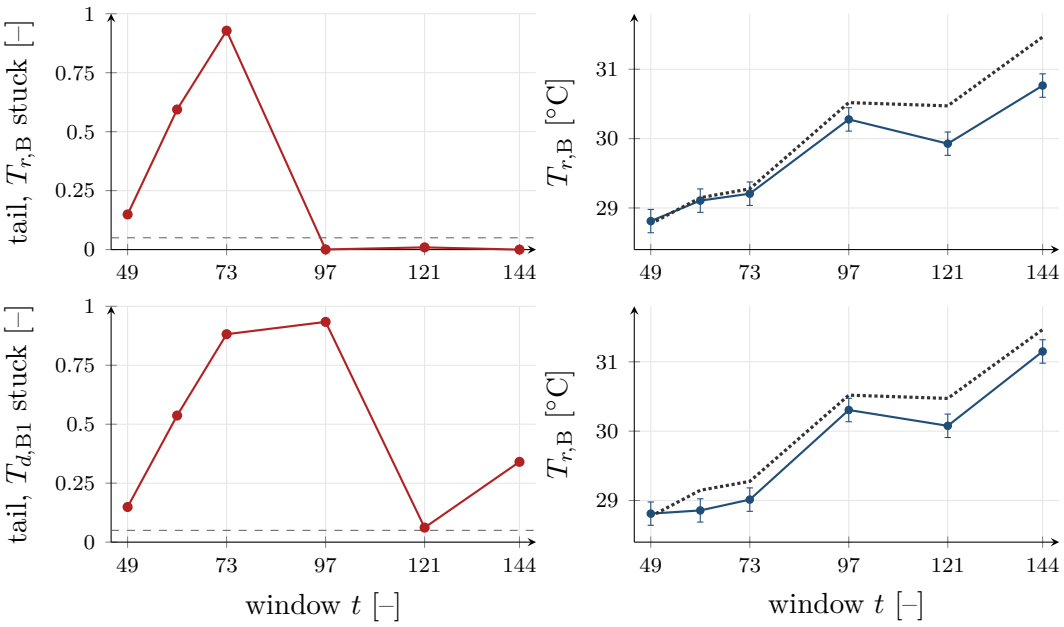


Figure 17.6: Two stuck sensors at the six windows run (Table 17.8); points are joined only to guide the eye. Left: the test's upper tail, with 0.05 dashed. Right: the fit's $T_{r,\mathrm{B}}$ with its band as bars, and the stated truth dotted. Upper row: the return sensor stuck. Lower row: the discharge of B1 stuck.

The others. The discharge of B2 stuck is past 0.05 at two of the six, 0.0063 at window 121 and 0.0231 at window 144, and the largest residual stands on the discharge of B1, the wrong sensor; $T_{r,\mathrm{B}}$ reads 30.752°C at window 144 against 31.465. The inlet at B4 stuck is past 0.05 at none of the six. At window 73 itself, none of the four has been caught: the tails are 0.8817, 0.9245, 0.9282 and 0.9178.

Why they differ. Table 17.9 evaluates equation (17.12) at window 144. Each row is a stuck sensor; the columns give the departure b from the truth, the redundancy number, what the test gains, $c_k b$, and the fit's actual miss on $T_{r,\mathrm{B}}$ where the table above has it. The discharges of B1 and B2 have nearly the same redundancy number, but the derate of unit B2 has carried its discharge further from the stuck value, so its test gain is 15.6 against 3.0. No statement holds of "a stuck discharge" as a class.

Table 17.9: A stuck reading's departure at window 144, each sensor held at its value of window 49.

Stuck sensor	b [–]	$1 - h_{kk}$ [–]	$b^2(1 - h_{kk})$ [–]	$c_k b$ [K]	Fit less truth, $T_{r,\mathrm{B}}$ [K]
$T_{d,\mathrm{B1}}$	−5.33	0.1065	3.0	−0.333	−0.316
$T_{d,\mathrm{B2}}$	−11.39	0.1202	15.6	−0.713	−0.713
$T_{r,\mathrm{B}}$	−7.40	0.6825	37.3	−0.705	−0.700
T_{B4}	−7.44	0.1020	5.6	−0.058	−0.065

Fresh nights. One night's tails are one draw. Table 17.10 repeats the test over 200 nights with the truth held and the sensor errors redrawn: each row a stuck sensor or none, the columns the share of nights past 0.05 at each of three windows and at any of them, with its standard error, and the share on which the stuck reading has the largest $|e|$. With nothing stuck the test itself goes past 0.05 on 0.130 of nights at one window or more. The discharge of B1 at 0.190 and the inlet at B4 at 0.145 stay near that rate. The discharge of B2 is caught, on 0.955 of nights, but named on only 0.12 to 0.38. The return is caught and named nearly always.

Table 17.10: Stuck sensors over 200 nights of fresh sensor errors, each held at its own reading of window 49.

Stuck sensor	Past 0.05, window 97 [–]	Window 121 [–]	Window 144 [–]	Any of three (se) [–]	Own reading largest, the three windows [–]
none	0.045	0.055	0.030	0.130 (0.024)	–
$T_{d,\mathrm{B1}}$	0.060	0.060	0.075	0.190 (0.028)	0.30, 0.31, 0.51
$T_{d,\mathrm{B2}}$	0.265	0.560	0.875	0.955 (0.015)	0.12, 0.15, 0.38
$T_{r,\mathrm{B}}$	0.995	0.985	1.000	1.000 (0.000)	0.98, 0.92, 0.99
T_{B4}	0.050	0.065	0.030	0.145 (0.025)	0.56, 0.54, 0.82

The answer. A removed sensor widens the band and says so; a stuck one keeps the band narrow. Whether the test sees it depends on $b^2(1 - h_{kk})$. At these windows the return stuck is caught and named, the discharge of B2 is caught but seldom named, and the discharge of B1 and the inlet at B4 stay near the test's own false-alarm rate. The stuck discharge of B1 still leaves $T_{r,\mathrm{B}}$ low by 0.316 K at window 144, against a band of 0.169 K.

17.3.6 The answers without it

Operator question. What does a loss do to tonight's answers?

General form. A prediction under a threshold, equation (2.12), on the joint event of Chapter 12.

Specialized form. The question is equation (2.12) with $f(\mathbf{z}) = -\max_k T_k^*(\mathbf{z})$ at the scenario node (12.1) with $T_s^* = 16°\mathrm{C}$, $f^* = -27°\mathrm{C}$, and $p(\mathbf{z} \mid \mathbf{y})$ replaced by $\mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma}_{(k)})$; it is Chapter 12's joint (12.11) carried on the down-date, and its largest set point (12.13) at the stated assurance $\kappa = 0.95$.

Building blocks. The inlets' gradients (12.10) and $\mathbf{\Sigma}_{(k)}$; the refit (17.10).

The joint and the set point. Table 17.11 gives them: each row a removed reading, the columns the joint at $\hat{\mathbf{z}}$, the joint on the refit, and the largest set point at $\kappa = 0.95$ at $\hat{\mathbf{z}}$. With all rows the joint is 0.80983 and the largest set point 15.391°C. The largest move in the reported value is the supply's, to 0.77679, and on the refit that loss goes the other way, to 0.84002, with the draws at 0.84581. The routes differ because the refit moves the state as well as the band: with its reading gone T_s goes from 13.8947 to 14.1018°C, and at a pinned T_s^* a higher present supply means a smaller rise to the return. The gap does not cross 0.90 or 0.95, and both values are stated. Every other reading stays within 0.013 of 0.80983 on either route, except the discharge of B2, whose refit gives 0.83971 against 0.80344 at $\hat{\mathbf{z}}$, and the inlet at B4, whose refit gives 0.79535 against 0.80625. The largest set point falls most without the supply reading, to 15.211°C, by 0.180 K: once the set point is pinned the absolute level matters, as it did not for the band on $T_{r,\mathrm{B}}$. The discharges of hall B come next, at 15.356 to 15.358°C, and every other reading leaves it at 15.383°C or above.

Table 17.11: Chapter 12's joint at $T_s^* = 16°\mathrm{C}$ and its largest set point at $\kappa = 0.95$, each reading removed.

Removed	Joint, at $\hat{\mathbf{z}}$ [–]	Joint, refit [–]	Largest set point, at $\hat{\mathbf{z}}$ [°C]
none	0.80983	–	15.391
T_s	0.77679	0.84002	15.211
$T_{r,\mathrm{A}}$	0.80936	0.81133	15.389
$T_{r,\mathrm{B}}$	0.80827	0.80498	15.383
$T_{d,\mathrm{A1}}$	0.80746	0.80339	15.383
$T_{d,\mathrm{A2}}$	0.80738	0.79761	15.383
$T_{d,\mathrm{A3}}$	0.80745	0.80221	15.383
$T_{d,\mathrm{B1}}$	0.80298	0.80225	15.356
$T_{d,\mathrm{B2}}$	0.80344	0.83971	15.358
$T_{d,\mathrm{B3}}$	0.80318	0.81887	15.357
Q_{A} meter	0.80973	0.80985	15.391
Q_{B} meter	0.80941	0.82007	15.389
loop meter	0.80980	0.80584	15.391
T_{A1}	0.80980	0.81010	15.391
T_{A3}	0.80802	0.80939	15.389
T_{B2}	0.80938	0.81029	15.390
T_{B4}	0.80625	0.79535	15.387

Band and answer disagree on what matters. The discharge of B1, which most widens the band on $T_{r,\mathrm{B}}$, moves this joint to 0.80298. No single loss moves it as much as the survey's unknown does: with the unread fractions at a spread of 0.04 and 0.12 the joint is 0.96730 and 0.64495, as in Section 12.4.3. This answer's band is carried by a rack with no sensor, so that part has no row of its own to lose.

The band that narrows. With the hall-B meter removed, the refit gives T_{B3} a band of 0.8217 K against 0.8293 with all rows. Equation (17.11) forbids that at a fixed point, and at $\hat{\mathbf{z}}$ the band does widen, to 0.8307 K. The refit's state has moved: Q_{B} falls from 296.012 to 292.590 kW, and the rise $T_{r,\mathrm{B}} - T_{c,\mathrm{B}}$, the gradient of T_{B3} on β_{B3} by equation (12.10), falls from 10.2004 to 10.0863 K. The draws agree with the refit, at 0.8216 K. The narrower band is the same uncertainty on β_{B3} multiplied by a smaller rise; the reading did not hurt, the point moved.

17.4 The checks (step 5)

The answers above used the band at the fit's own state and the linearization. The checks test both, and the stated truth stands last.

17.4.1 Draws with the row removed

Table 17.12 sets the three routes side by side for the four removals that were sampled: each row a removal, the columns the effective size of 200 000 draws, the band at $\hat{\mathbf{z}}$, the refit's and the draws' standard deviation with its Monte Carlo standard error. The refit and the draws agree within that error at all four; their gaps of at most 0.0003 cannot be told from zero. The down-date lies below the draws by 0.0008 K at the discharge of B2, where the refit moves the value most. On the value, the draws' mean lies below the refit's MAP by 0.0067 to 0.0143 K, a mean beside a mode, so that gap is the posterior's skew and not an error of either.

Table 17.12: The band on $T_{r,\mathrm{B}}$ by three routes, with the row removed.

Removed	Effective size [–]	At $\hat{\mathbf{z}}$ [K]	Refit [K]	Draws [K]	Draws' se [K]
$T_{d,\mathrm{B1}}$	192 347	0.2495	0.2495	0.2498	0.0004
$T_{r,\mathrm{B}}$	198 071	0.2038	0.2038	0.2040	0.0003
Q_{B} meter	196 205	0.2215	0.2215	0.2215	0.0004
$T_{d,\mathrm{B2}}$	194 021	0.2454	0.2465	0.2462	0.0004

17.4.2 The limit of exact rows and the curvature term

The fit runs on the reduced state. With the discharge of B1 removed, Table 17.13 fits the full state with the physical rows at tolerance σ_p and the removed reading's row at a standard deviation of 10^6: each row a tolerance, the columns $T_{r,\mathrm{B}}$, its band and the iterations. At $\sigma_p = 0.01$ the full state already gives the reduced refit's five decimals. The curvature term, checked numerically, is small: the objective's full Hessian at the refit gives 0.249428 K against the Gauss–Newton 0.249497 K.

Table 17.13: The limit of exact physical rows with the discharge of B1 removed.

Tolerance σ_p [kW or K]	$T_{r,\mathrm{B}}$ [°C]	Band [K]	Iterations [–]
0.100	29.17748	0.24975	4
0.010	29.17719	0.24950	3
0.001	29.17719	0.24950	2
reduced refit	29.17719	0.24950	–

17.4.3 The answers against the stated truth

The record comes from a stated truth (Section 17.6.3). Table 17.14 collects the checks: each row a route or a quantity, the columns the reported value, the value beside it and the gap, the second minus the first. The truth $T_{r,\mathrm{B}} = 29.2761\,^{\circ}\mathrm{C}$ lies within one reported band of the refit's value in every row of Tables 17.4 and 17.5, and it is the dotted line of Figure 17.2. The largest gaps of the chapter are in the joint, where the state's move counts at a pinned node, and there the draws side with the refit. A truth inside these bands supports them; it cannot show that the wider band after a loss was needed, because one night tests one value. The stuck sensors stand against the truth in Table 17.8, window by window: at window 144 the fit with the return stuck is low by 0.700 K and with the discharge of B1 stuck by 0.316 K, while the band in both is 0.169 K.

Table 17.14: The checks and their gaps at window 73. The truth is the stated truth.

Route or quantity [unit]	Reported	Beside it	Gap
Band on $T_{r,\mathrm{B}}$ at $\hat{\mathbf{z}}$, against the draws (se 0.0003 *to* 0.0004*)*			
$T_{d,\mathrm{B1}}$ removed [K]	0.2495	0.2498	+0.0003
$T_{r,\mathrm{B}}$ removed [K]	0.2038	0.2040	+0.0002
Q_{B} meter removed [K]	0.2215	0.2215	0.0000
$T_{d,\mathrm{B2}}$ removed [K]	0.2454	0.2462	+0.0008
Joint at $T_s^* = 16\,^{\circ}\mathrm{C}$, the largest route gaps; se of the draws in parentheses			
T_s removed, at $\hat{\mathbf{z}}$ against the refit [–]	0.77679	0.84002	+0.06323
T_s removed, at $\hat{\mathbf{z}}$ against the draws [–]	0.77679	0.84581 (0.00081)	+0.06902
$T_{d,\mathrm{B2}}$ removed, at $\hat{\mathbf{z}}$ against the refit [–]	0.80344	0.83971	+0.03627
$T_{d,\mathrm{B2}}$ removed, at $\hat{\mathbf{z}}$ against the draws [–]	0.80344	0.84087 (0.00084)	+0.03743
Groups, at $\hat{\mathbf{z}}$ against the refit			
Hall B return and its discharges removed [K]	0.5666	0.5647	−0.0019
Every hall B temperature removed [K]	1.3429	1.2779	−0.0650
Approximations, $T_{d,\mathrm{B1}}$ removed			
Newton against Gauss–Newton, band [K]	0.249497	0.249428	−0.000069
Full state, $\sigma_p = 0.001$, band [K]	0.24950	0.24950	0.00000
Stated truth, the fit with all rows (band)			
$T_{r,\mathrm{B}}$ [°C]	29.1749 (0.1686)	29.2761	+0.1012
$T_{r,\mathrm{A}}$ [°C]	29.8674 (0.1698)	29.8052	−0.0622
T_{B3} [°C]	24.0747 (0.8293)	24.8571	+0.7824
T_{B4} [°C]	22.7126 (0.2821)	22.9046	+0.1920
Q_{B2} [kW]	92.5754 (2.6545)	94.8592	+2.2838
Q_{B} [kW]	296.0123 (3.4302)	298.2975	+2.2852

17.5 The report (step 6)

The screen's reading is overturned twice. The sensor named hot-aisle is fifth for its own quantity: its loss takes the band from 0.1686 to 0.2038 K, and a hall-B discharge's takes it to 0.2495 K. The three discharges and the hall meter read it through the heat balance, closer than the return does. And three discharges lost together cost little more than one, at 0.2651 K, because they

are one route, with the return behind them; with the return out, one discharge costs 0.4494 K. A stuck sensor leaves the band unchanged, and over fresh nights the test sees the return and the discharge of B2 and stays near its own false-alarm rate for the discharge of B1 and the inlet at B4.

17.5.1 The answers and their bands

Table 17.15 gives the answers in the order of Table 17.1. Each row is a question; the second column gives the reported answer, at $\hat{\mathbf{z}}$, and the third what stands beside it.

Table 17.15: The report at window 73, 01:00. Bands on $T_{r,\mathrm{B}}$ in K, with all rows 0.1686. No failure probability is assigned to any sensor.

Question	Answer, reported	Beside it
How wide if this one fails?	return sensor 0.2038, ratio 1.209; discharge of B1 0.2495, ratio 1.480, the largest	refit 0.2038, 0.2495; draws 0.2040, 0.2498
Which reading carries it?	for $T_{r,\mathrm{B}}$: the three hall-B discharges, then the hall-B meter, then the return, fifth: the heat balance reads it to 0.2186	precisions 11.11 return, 20.96 balance, 2.56 inlets, 0.555 coil priors; return fifth at all nine windows read
Several together?	hall-B discharges 0.2651; with the return 0.5666; every hall-B temperature 1.3429; return then B1's discharge 0.4494	refit 0.2651, 0.5647, 1.2779; ratio to the sum of singles 0.425 for the discharges, 2.640 for the two load meters
No replacement?	critical: each inlet, for its own rack, and the supply; β_{B4} 0.0303 to 0.0800, rack B4 0.2821 to 0.8293; rank 13 to 12	the count 3.691 falls to 2.693–3.640 with any one removed
Stuck and connected?	band unchanged at 0.169; past 0.05 at one of three windows, 200 nights: return 1.000, discharge of B2 0.955, discharge of B1 0.190, inlet B4 0.145, none stuck 0.130	this night's six windows: return 0.0006 at window 97, discharge of B1 past 0.05 at none
Tonight's answers?	joint at $T_s^* = 16$°C from 0.80983 to at most 0.033 away at $\hat{\mathbf{z}}$, the supply's 0.77679; largest set point 15.211°C without the supply, 15.356 without B1's discharge, against 15.391	supply's refit 0.84002, draws 0.84581; survey's spread 0.04 and 0.12: 0.96730, 0.64495

17.5.2 The entry in the shift log

The operator's entry as of window 73, 01:00, using nothing later, reads: "Sensor loss. If $T_{r,\mathrm{B}}$'s sensor goes, the hall-B hot-aisle band goes 0.17 to 0.20 K; if any one hall-B discharge goes, to 0.25 K; all of them with the return, 0.57 K. Inlet B4 has no substitute for its own rack (0.28 to 0.83 K). With the return out, a discharge going as well takes the band to 0.45 K. The joint at $T_s^* = 16$°C moves by at most 0.033 for any single loss at the fit's state; the largest set point at 0.95 falls from 15.39 to 15.21°C without the supply reading. On these bands the return sensor is the less costly to take out, and a discharge the more costly without a replacement reading. A sensor stuck since 23:00 has not been ruled out by the test: at this window the return, the B1 and B2 discharges and inlet B4, each held at its 23:00 value, give tails of 0.88 to 0.93. Action: not decided here."

Common mistake

Reading a narrower band after a loss, or a quiet test, as good news. A refit's band can narrow because the state moved, which says nothing against the reading removed. A test with a large tail says the rows agree with one another, and a stuck reading whose $b^2(1-h_{kk})$ is small can agree with them while it is wrong.

17.6 The next reading (step 7)

The report leaves the worst single loss on the band at this window, the discharge of B1. A further reading is a further row, and what it would do to the entries is known before its value is.

17.6.1 What restores the worst loss

A candidate adds a row to $\boldsymbol{\Sigma}_{(k)}$ with k the discharge of B1, by the rank-one form of equation (4.65). Table 17.16 compares the candidates: each row a candidate, the columns the bands on $T_{r,\mathrm{B}}$, on rack B3's inlet and on the heat of unit B2. A second permanent return sensor and a handheld reading at the return put the same row into this window, so they do the same to the entries. Between them the difference is what happens in later windows and the rows do not speak to it.

Table 17.16: With the discharge of B1 lost, the candidates and the bands after each.

Candidate	$T_{r,\mathrm{B}}$ [K]	T_{B3} [K]	Q_{B2} [kW]
None	0.2495	0.8531	3.675
Repair, the same reading back	0.1686	0.8293	2.655
Second return sensor or handheld at the return, hall B, 0.3 K	0.1918	0.8390	3.376
Water-side heat check on unit B1, 3 kW	0.2066	0.8375	2.836
Clamp-on, hall-B load, 1 kW	0.2475	0.8527	3.655
Probe at rack B3, 0.3 K	0.2404	0.2830	3.626

A return reading restores most of the band, 0.1918 against the repair's 0.1686 K. The heat check on unit B1 restores most of the unit-heat column, 2.836 against 2.655 kW for Q_{B2}, and part of the band, since it says how the heat divides. The clamp-on does almost nothing to any of them, since the balance has lost a discharge and not its meter. The probe at B3 does most for that rack and little else. The ranking depends on the quantity, and the choice stays with the operator. Chapter 18 ranks readings against a named worry.

17.6.2 What the analysis showed

The screen said to protect the hot-aisle sensor and to let a discharge go first because there are three. The rows say the reverse for the band on the hot-aisle temperature: three discharges and the hall meter read that temperature through a heat balance with no coil in it, better than the sensor with its name, and they work as one route, in series, with the return in parallel behind them. So the return can go first, and after it a discharge costs far more. Which reading matters also depends on what is asked: the set-point answer barely notices the loss of most readings, is carried in its band by a rack with no sensor, and does need the supply reading. And the two failures differ: a lost discharge widens the band and shows it, while a stuck sensor leaves the band in place, and the test sees it only as far as the departure and the redundancy number together allow.

17.6.3 Limits of the model

One window. The bands are read at window 73, steady. A leverage and a coupling depend on the state; the order of the first five readings was read at nine windows, and the other answers at this one.

Named sets. Losses are named sets with no probabilities, so nothing here ranks them by how likely they are, and the common causes that would make a group lose together are not modeled.

One failure mode. Stuck-at is one failure mode, tried on four sensors, at six windows on the record's errors and at three over fresh nights with the truth held. A sensor that drifts, or a value held from another window, may behave differently, and the windows between those run were not read.

Down-date and refit. They differ where the state moves, as they do with the coils' readings removed; the down-date needs no outcome and that is why it is the one reported.

Constructed readings. The record is constructed: generated from a stated truth of the plant with the stated sensor errors, and Section 17.4.3 compares the answers with that truth. It is not field data.

17.7 What to retain

Definition 17.1 (The band after a lost reading)**.** With $\mathbf{a}_k$ the row of $\mathbf{J}$ of reading k, the covariance without it is $\mathbf{\Sigma}_{(k)} = \mathbf{\Sigma} + \mathbf{\Sigma}\mathbf{a}_k^{\mathsf{T}}\mathbf{a}_k\mathbf{\Sigma}/(1 - h_{kk})$ with $h_{kk} = \mathbf{a}_k\mathbf{\Sigma}\mathbf{a}_k^{\mathsf{T}}$, and the variance of a derived quantity with gradient $\mathbf{g}$ grows by $c_k^2/(1 - h_{kk})$ with $c_k = \mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{a}_k^{\mathsf{T}}$. It holds at the fit's own state and needs no outcome. Any ranking it gives belongs to the quantity $\mathbf{g}$ names.

Three conclusions carry beyond the plant. A count of sensors does not measure redundancy, and each row's redundancy number comes closer, read with its coupling. Rows in series on one route add less together than apart, and rows on parallel routes add more, so a group's added variance is computed and not summed. And what a loss costs, $c_k^2/(1 - h_{kk})$, and what a stuck reading shows, $b^2(1 - h_{kk})$, are different numbers, so the band and the test have to be asked separately.

Table 17.17 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 17.17: The terms of sensor loss and where each is defined. Units for the shift plant.

Term	Expression	Units	Meaning	Defined in
Hot-aisle temperature	eq. (17.2)	°C	the return, a function of the state	17.1.2
Heat balance	eq. (17.3)	°C	the coil-free route to the return	17.1.2
Rows as a sum	eq. (17.5)	mixed	one term per row	17.2.1
Down-date	eq. (17.7)	mixed	$\boldsymbol{\Sigma}$ with a row removed, at $\hat{\mathbf{z}}$	17.2.2
Block deletion	eq. (17.8)	mixed	several rows removed	17.2.2
Coupling c_k	eq. (17.9)	K	a row's link to the quantity	17.2.2
Refit	eq. (17.10)	mixed	the fit without the row	17.2.3
Band after a loss	eq. (17.11)	K	the reported answer	17.3.1
Critical measurement	the reading rows' rank falls	–	one direction left to the priors	17.3.4
Stuck reading	eq. (17.12)	–	test gains $b^2(1-h_{kk})$, answer takes $c_k b$	17.3.5

17.8 Exercises

Exercise 17.1. *Application.* From $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} = 0.028419\,\mathrm{K}^2$, $c_k = +0.06186\,\mathrm{K}$ and $1 - h_{kk} = 0.1131$, compute the band on $T_{r,\mathrm{B}}$ with the discharge of B1 removed. (By equation (17.11), $0.028419 + 0.06186^2/0.1131 = 0.062255\,\mathrm{K}^2$, and the band is $0.2495\,\mathrm{K}$.)

Exercise 17.2. *Derivation.* The return sensor's row is $\mathbf{g}^{\mathsf{T}}$ over its standard deviation of $0.3\,\mathrm{K}$. Show that its loss subtracts a precision, and find the ratio. (Then $c_k = \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}/0.3$ and $h_{kk} = \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}/0.3^2$, and equation (17.9) reduces to $1/\mathrm{band}_{(k)}^2 = 1/(\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}) - 1/0.3^2$, the precision of one reading taken away. A band of $0.2038\,\mathrm{K}$ with the reading put back gives $0.1686\,\mathrm{K}$, and the ratio is $0.2038/0.1686 = 1.209$.)

Exercise 17.3. *Formulation.* Say why removing the inlet at rack B4 returns β_{B4} to its prior spread. (By the pattern (17.4) that reading is the only reading row with a nonzero in the column of β_{B4}. With it gone the fraction is held by its prior row alone, its objective term separates from the rest, and its band is 0.0800.)

Exercise 17.4. *Application.* Explain why a discharge of hall A leaves the band on $T_{r,\mathrm{B}}$ alone. (The question's gradient has no entry in hall A's columns, so the row reaches $T_{r,\mathrm{B}}$ only through $\boldsymbol{\Sigma}$, by way of the loop meter. Its coupling c_k is $+0.00143$ to $+0.00145$, against $+0.06186$ for hall B's, and the added variance is about $0.00002\,\mathrm{K}^2$, so the ratio to the band with all rows is 1.000.)

Exercise 17.5. *Application.* From the band of $0.2651\,\mathrm{K}$ with the three hall-B discharges removed, find the added variance of the group and compare it with the sum of the three alone. (Together, the band squared less 0.028419, which is $0.04188\,\mathrm{K}^2$ on the unrounded bands; alone, the three added variances are 0.033837, 0.031826 and 0.032970, which sum to 0.09863. The group adds less because the three rows are in series on one route, the balance, which any one of them breaks, while the return stands in parallel; the added variance together is 0.04188.)

Chapter 18

Selecting the Next Field Measurement

18.1 One technician and one trip

It is 01:00, window 73 of the shift. A field technician is on site for one trip, with a handheld probe, a clamp-on meter, access to the water-side heat check, and a second return sensor and a second discharge sensor to place. The hottest read inlet on the screen is rack B4's, 22.67°C, and the unit that is changing is B2. From the screen alone the operator concludes that the technician should go to B2 and check its coil.

The screen shows where the plant is read and where it has been moving. It does not show which reading would change what the operator knows, and the operator has more than one thing to know: the hottest rack, the coil, the raise of Chapter 12. The question is the one the screen cannot answer,

of the readings available, which one narrows the quantity the operator is worried about most, and does that depend on which worry it is.

Every earlier step 7 ranked candidates for one chapter's question. Here the ranking is the question, and it is asked of the one fit of Chapter 12.

Jargon: worry, candidate reading, scenario row, width taken off, settled share, reference level, greedy order, repeat

A *worry* is a quantity the operator wants narrowed, a function of the state. A *candidate reading* is an instrument, a place and a spread, not yet used. Its *scenario row* is the reading row it would add, written before its value exists. The *width taken off* a worry is the share by which its band falls after the row. The *settled share* of a candidate is the probability, over the reading's outcomes, that a stated joint probability ends at or above an upper *reference level* or below a lower one. The *greedy order* takes the best reading, then the best given it, and so on. A *repeat* is a further reading of the same quantity with an independent error.

Key idea

A reading not yet taken is a row with a known gradient and an unknown value. The covariance after it depends on the gradient and on the spread and not on the value, so the width it takes off every worry is computed before the trip. Those widths form a matrix of candidates against worries, not a list. Beside the width stands a different question, whether the reading can settle a decision, and that one depends on the outcomes.

18.1.1 A reading as a row before its value

Each earlier chapter's posterior was computed from rows that had values. A candidate's row has the same gradient whatever the instrument later shows, because the gradient belongs to the plant and to the place of the instrument. Covariance is built from gradients, so the band after a reading can be known in advance, and only the mean waits for the reading.

That makes two different things worth asking of a candidate. How much narrower does a worry's band become? And does the narrower band, wherever the reading places it, leave a decision open? The first has one number per candidate and worry. The second has a distribution over outcomes, since the same band can land on either side of a limit.

Principle 18.1. *A candidate is ranked against a named worry. The band after it needs no outcome; a decision after it does, and is computed over the predictive of the outcome.*

18.1.2 The shift plant at window 73 (step 0)

The plant is the two-hall shift plant on the reduced state of Chapter 4, with the state (12.3), 19 entries, and the chain (12.4) from the free entries to the temperatures. The dependent block of the physical rows is triangular with a nonzero diagonal, hence invertible. A candidate adds no entry and no physical row: it reads a quantity the chain already defines. The choices of the graph, steady state, pinned airflow and one effective conductance per coil, are those of Chapter 12, and they fail as they do there.

Figure 18.1 draws the plant for this question. The middle row holds the free entries (green), with T_s left out because it reaches every temperature. The top row holds six worries, each joined to the entries its gradient touches. The bottom row holds some of the candidates, each joined by a dashed line to the entries its row touches.

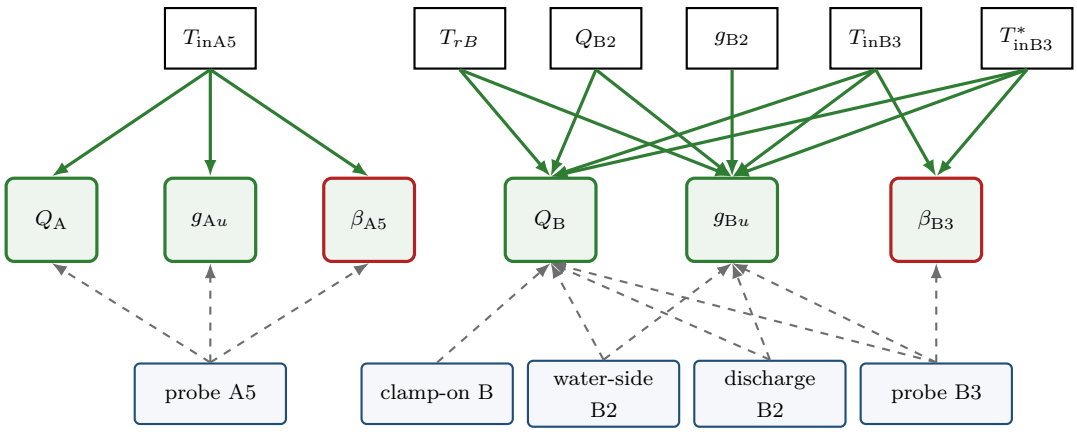


Figure 18.1: Worries (top), free entries (middle) and candidate rows not yet taken (bottom, dashed). Free entries with a red border are read by no row of the record. The set point and the priors, the read fractions, the other unread fractions and the sixteen readings of window 73 are not drawn.

18.1.3 The operator's questions (step 1)

The operator asks five questions of window 73, and step 1 writes each over the state. Table 18.1 lists them: the question in the operator's words, the first of them the drill question; the general form of its category in the notation of Chapter 2; its mathematical form on this plant; its category; the computation that answers it; and the subsection where it is answered. All five are design questions on the preposterior; **a** and **g** are defined in Section 18.1.4.

Table 18.1: The questions asked of the one fit at window 73, 01:00. $\mathbf{g}$ is a worry's gradient, $\mathbf{a}$ a candidate's gradient over its spread, $\tilde{y}_a$ its outcome, and P_a the joint at $T_s^* = 16°\text{C}$ after that outcome.

Operator question	General form	Mathematical form	Category	Computation	Answered in
Of the readings I could request from the field technician, which one narrows the quantity I am worried about most? Preposterior over the candidates, ranked by width taken off that quantity.	(2.16)	$\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g} - (\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{a})^2/(1 + \mathbf{a}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{a})$	design	preposterior	18.3.1
Is there one list?	(2.16)	the matrix of candidates against worries	design	preposterior	18.3.2
Does the reading settle anything?	(2.16), (2.12)	$\Pr(P_a(\tilde{y}_a) \geq 0.95$ or $P_a(\tilde{y}_a) < 0.05 \mid \mathbf{y})$	design	preposterior	18.3.3
What should the second and third reading be?	(2.16)	the greedy order against the best pair, by worry	design	preposterior	18.3.4
How good must the reading be?	(2.16)	width against the reading's spread; repeats; the band with β_{B3} known	design	preposterior	18.3.5

18.1.4 The worries and the candidates as rows (step 2)

The evidence is the window's 33 rows, the readings and priors of equation (12.8), unchanged. What step 2 adds is not evidence yet: one scenario row per candidate.

The worries. Six are carried: T_{inB3}, the inlet of the rack the survey puts highest in hall B; T_{inA5}, the same for hall A; T_{rB}, hall B's return; $Q_{\text{B2}} = G_{\text{B2}}(T_{rB} - T_s)$, the heat of the unit that is changing; g_{B2}, its log-derate; and T^*_{inB3}, rack B3's inlet at the scenario node $T_s^* = 16°\text{C}$ of equation (12.6). For each, $\mathbf{g} = \partial f/\partial \mathbf{z}$ at $\hat{\mathbf{z}}$.

The candidates. A candidate reads a function $m_a(\mathbf{z})$ with spread σ_a, and its scenario row and normalized gradient are

$$r_a(\mathbf{z}) = \frac{m_a(\mathbf{z}) - \tilde{y}_a}{\sigma_a}, \qquad \mathbf{a} = \frac{1}{\sigma_a}\frac{\partial m_a}{\partial \mathbf{z}}\bigg|_{\hat{\mathbf{z}}}. \tag{18.1}$$

Table 18.2 lists the fourteen, each row a kind of instrument, the columns where it may go, what it reads and its spread. A second sensor's reading is taken with an error independent of the first. Chapter 15's permanent water-side heat meter over twelve windows is stated at 4 kW; the water-side heat check here is the portable 3 kW of Chapters 13 and 14.

Four choices enter with these rows. *A reading of one window:* each row reads window 73's state, which buys one covariance and one set of worries, and it fails for an instrument left in place, whose later rows see later states. The permanent instrument is cited in Section 18.6.2. *A linear reading at $\hat{\mathbf{z}}$:* the gradient is taken at the MAP state whatever the outcome, which buys a covariance free of the outcome, and it fails if an outcome moves the state far enough to change the gradient; Section 18.2.3 defines the refit that tests it. *Width as the ranking:* the band buys a closed form that needs no outcome, and it fails when the question has a threshold, because a band can narrow and the decision stay open, or stay wide and the decision be settled;

Table 18.2: The candidate readings.

Instrument	Places	$m_a(\mathbf{z})$	σ_a
Handheld probe	racks A2, A4, A5, B1, B3, B5	$T_{\text{in}k}$	0.3 K
Water-side heat check	units B1, B2, B3, A1	Q_u	3 kW
Clamp-on	hall A's load, hall B's load	Q_h	1 kW
Second return sensor	hall B	T_{rB}	0.3 K
Second discharge sensor	unit B2	$T_{d,\text{B2}}$	0.3 K

so the settled share stands beside it. *No costs:* a technician's trip has a cost that the record does not hold, so the chapter ranks by what a reading does and the operator holds the cost; this fails as a basis for choosing wherever the cost differs enough to outweigh the width, and no row here can say when.

Which row acts on which entry. With $*$ for an entry that can be nonzero, the candidates' rows and the worries' gradients have the pattern

$$\begin{bmatrix} * & 0 & * & 0 & * & 0 & 0 & * & 0 \\ * & * & 0 & * & 0 & 0 & * & 0 & 0 \\ 0 & 0 & * & 0 & * & 0 & 0 & 0 & 0 \\ 0 & 0 & * & 0 & 0 & 0 & 0 & 0 & 0 \\ * & 0 & * & 0 & * & 0 & 0 & 0 & 0 \\ * & 0 & * & 0 & * & 0 & 0 & 0 & 0 \\ 0 & 0 & * & 0 & * & 0 & 0 & * & 0 \\ 0 & 0 & 0 & 0 & * & 0 & 0 & 0 & 0 \end{bmatrix} \quad \begin{array}{c} \text{probe at B3, and worry } T_{\text{inB3}} \\ \text{probe at A5, and worry } T_{\text{inA5}} \\ \text{water-side B}u\text{, and worry } Q_{\text{B2}} \\ \text{clamp-on, hall B} \\ \text{second return, and worry } T_{rB} \\ \text{second discharge, B2} \\ \text{worry } T^*_{\text{inB3}} \\ \text{worry } g_{\text{B2}} \end{array} \tag{18.2}$$

with the columns T_s; Q_{A}; Q_{B}; the g_u of hall A; the g_u of hall B; the read fractions; β_{A5}; β_{B3}; the other unread fractions. Read as observability, the two unread columns that matter are empty in all of the window's reading rows (Figure 12.2), and each has exactly one candidate that reaches it. The water-side row has no T_s entry, because the coil heat is the hall's heat shared in proportion to conductance. Only a probe touches an unread fraction, and every other candidate acts in the columns the record already reads.

18.2 The posterior of the state (step 3)

The posterior is the one fit of Chapter 12, window 73, with no new row. What is derived here is the worries' place in it and the posterior a scenario row would leave.

18.2.1 The one fit and the worries

The rows (12.8), the normal equations (12.9) with the information $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ of normalized rows, which with unnormalized rows is $\mathbf{J}^{\mathsf{T}}\mathbf{W}\mathbf{J}$, the MAP state (Table 12.5) and the covariance $\boldsymbol{\Sigma}$ (Table 12.6) are cited and not repeated. Table 18.3 reads the worries from them: each row a worry, the columns its value at $\hat{\mathbf{z}}$ and its band $\sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}}$. Table 18.4 gives the gradients the bands use, each row a quantity and the columns the entries; a quantity's column is blank where its entry is zero. Rows of the candidates are the same gradients: $\mathbf{a} = \mathbf{g}_{T_{\text{inB3}}}/0.3$ for the probe at B3.

Table 18.3: The worries on the one fit, window 73, 01:00.

Worry	Value at $\hat{\mathbf{z}}$	Band $\sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}}$
T_{inB3} [°C; K]	24.0747	0.8293
T_{inA5} [°C; K]	22.9887	1.0114
T_{rB} [°C; K]	29.1749	0.1686
Q_{B2} [kW]	92.5754	2.6545
g_{B2} [–]	−0.1444	0.0290
T^{*}_{inB3} at $T^{*}_{s} = 16$°C [°C; K]	26.1800	0.8349

Table 18.4: Gradients at $\hat{\mathbf{z}}$. Entries of $\mathbf{g}$ in the quantity's unit per unit of the column's entry; hall A's T_{inA5} is given in its own row, its columns being those of hall A.

Quantity	T_s [–]	Q_h [per kW]	g_{h1} [–]	g_{h2} [–]	g_{h3} [–]	β_k [–]
T_{inB3} [K]	+1.0000	+0.0333	−5.5461	−4.6226	−5.1115	+10.2004
T^{*}_{inB3} [K]		+0.0333	−5.5461	−4.6226	−5.1115	+10.2004
T_{inA5} [K]	+1.0000	+0.0242	−5.0931	−5.6722	−5.2074	+12.5067
T_{rB} [K]	+1.0000	+0.0499	−5.5461	−4.6226	−5.1115	
Q_{B2} [kW]		+0.3025	−33.6011	+64.5693	−30.9683	
$T_{d,\mathrm{B2}}$ [K]	+1.0000	+0.0197	−2.1860	−11.0795	−2.0147	

18.2.2 The posterior after a scenario row

A scenario row adds $\mathbf{a}\mathbf{a}^{\mathsf{T}}$ to the information. The preposterior covariance is equation (4.65) with $\boldsymbol{\Lambda}_a = \mathbf{a}\mathbf{a}^{\mathsf{T}}$, and by the rank-one form (4.66) with $\mathbf{h} = \sigma_a\mathbf{a}$ and $S = \sigma_a^2(1 + \mathbf{a}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{a})$,

$$\boldsymbol{\Sigma}_a = \boldsymbol{\Sigma} - \frac{\boldsymbol{\Sigma}\mathbf{a}\mathbf{a}^{\mathsf{T}}\boldsymbol{\Sigma}}{1 + \mathbf{a}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{a}}, \tag{18.3}$$

in which no outcome appears. The outcome enters the mean alone, with the predictive of the reading from the same fit,

$$\hat{\mathbf{z}}_a(\tilde{y}_a) = \hat{\mathbf{z}} + \frac{\boldsymbol{\Sigma}\mathbf{a}}{1 + \mathbf{a}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{a}}\,\frac{\tilde{y}_a - m_a(\hat{\mathbf{z}})}{\sigma_a}, \qquad \tilde{y}_a \mid \mathbf{y} \approx \mathcal{N}\big(m_a(\hat{\mathbf{z}}),\ \sigma_a^2(1 + \mathbf{a}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{a})\big). \tag{18.4}$$

For several rows at once, $\mathbf{a}$ becomes a matrix with one row each and equation (4.65) applies as it stands.

18.2.3 The refit on one outcome

Equations (18.3) and (18.4) hold for a row linear at $\hat{\mathbf{z}}$. The row itself is nonlinear, so on a given outcome the fit can be repeated: the 33 rows of equation (12.8) with r_a of equation (18.1) appended, the iteration (12.9) from $\hat{\mathbf{z}}$, and the gradient of the worry retaken at the new MAP state. This refit is another route and it is still a Laplace fit. Its outcomes are drawn from $\mathbf{z} \sim \mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ with an independent reading error. Chapter 12 set the exact posterior beside the Laplace normal on this window (Section 12.2.1), with an effective size of 396 168 of 400 000; for a candidate the same draws serve here, on $n = 400\,000$ with an effective size of 396 212, each reweighted by the reading's likelihood at a stated value, which gives the joint after that reading on the exact posterior.

18.3 The answers (step 4)

Each answer reads $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$ of Section 18.2.1 through the preposterior of Section 18.2.2, and the refit stands beside it among the checks.

18.3.1 The width a reading takes off

Operator question. Of the readings I could request from the field technician, which one narrows the quantity I am worried about most? Preposterior over the candidates, ranked by width taken off that quantity.

General form. A design question is the expected utility of a candidate over its outcomes, equation (2.16), with u the spread of a derived quantity after the outcome as one of its four utilities.

Specialized form. This is equation (2.16) with a a candidate of Table 18.2, $\tilde{y}_a$ its outcome, $\mathbf{z}$ the state (12.3) and u the variance of the worry f after the outcome. Since $\mathbf{\Sigma}_a$ in equation (18.3) is free of the outcome the integral falls away, and

$$u(a) = \mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}_a\mathbf{g} = \mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g} - \frac{(\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{a})^2}{1+\mathbf{a}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{a}} = \mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}\,(1-\rho_a^2), \qquad \rho_a = \frac{\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{a}}{\sqrt{\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}\,(1+\mathbf{a}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{a})}}. \quad (18.5)$$

Here ρ_a is the correlation of the reading, error included, with the worry. Its numerator is a sum over entries with either sign, and ρ_a is the whole of what a candidate does to a worry: two gradients that share no column with a nonzero covariance leave the band alone.

Building blocks. $\mathbf{\Sigma}$ (Section 18.2.1); $\mathbf{g}$ and $\mathbf{a}$ (Table 18.4); equation (18.5).

The probe at B3, on T_{inB3}. Equation (18.5) has $\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{a} = 2.29246\,\text{K}$ and $\mathbf{a}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{a} = 7.64152$, so

$$\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}_a\mathbf{g} = 0.68774 - \frac{2.29246^2}{1+7.64152} = 0.07959\,\text{K}^2,$$

a band of $0.2821\,\text{K}$ against $0.8293\,\text{K}$ now. The correlation is $\rho_a = 0.9404$. Rack B3's band is the survey's: the row of Table 18.4 gives β_{B3} the entry $10.2004\,\text{K}$, with $\Sigma_{\beta\beta} = 0.08^2$, a term of $0.66591\,\text{K}^2$ of the 0.68774, with no cross terms, since β_{B3} has zero covariance with every other entry. The probe's predictive by equation (18.4) is $24.0747 \pm 0.8819°\text{C}$, most of it that same term and little of it the probe's own error.

18.3.2 The matrix

Operator question. Is there one list?

General form. The form is again equation (2.16) with u a spread, read now for every candidate and every worry.

Specialized form. Equation (2.16) becomes the entries of a matrix: entry (a, f) is the root of equation (18.5) with $\mathbf{a}$ the candidate of row a and $\mathbf{g}$ the worry of column f. There is one list only if the rows keep their order in every column.

Building blocks. Equation (18.5); the six gradients of the worries; the fourteen rows.

The band after each candidate. Table 18.5 gives the band of each worry after each candidate, each row a candidate and each column a worry, with the first choice for each worry in bold. Figure 18.2 draws the share of the band taken off, $1 - \sqrt{\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}_a\mathbf{g}/\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}}$, on the same rows and columns, as a shaded grid whose number is the share in percent, darker for more.

Table 18.5: The band after each candidate, from equation (18.5). Temperatures' bands in K, Q_{B2}'s in kW, g_{B2}'s dimensionless. Bold: the first choice for that worry.

Candidate	T_{inB3} [K]	T_{inA5} [K]	T_{rB} [K]	Q_{B2} [kW]	g_{B2} [–]	T^*_{inB3} [K]
None	0.8293	1.0114	0.1686	2.6545	0.0290	0.8349
Probe at A2	0.8293	1.0112	0.1686	2.6545	0.0290	0.8349
Probe at A4	0.8293	1.0112	0.1686	2.6545	0.0290	0.8349
Probe at A5	0.8293	**0.2876**	0.1686	2.6545	0.0290	0.8349
Probe at B1	0.8290	1.0114	0.1670	2.6545	0.0290	0.8346
Probe at B3	**0.2821**	1.0114	0.1665	2.6542	0.0290	**0.2985**
Probe at B5	0.8289	1.0114	0.1667	2.6545	0.0290	0.8345
Water-side heat check,B1	0.8293	1.0114	0.1668	2.6252	0.0284	0.8349
Water-side heat check,B2	0.8293	1.0114	0.1668	**1.9880**	0.0233	0.8349
Water-side heat check,B3	0.8293	1.0114	0.1668	2.6256	0.0285	0.8349
Water-side heat check,A1	0.8293	1.0114	0.1685	2.6539	0.0290	0.8349
Clamp-on, hall A	0.8293	1.0112	0.1679	2.6474	0.0290	0.8349
Clamp-on, hall B	0.8288	1.0114	**0.1469**	2.4196	0.0283	0.8344
Second return, hall B	0.8265	1.0114	0.1470	2.6390	0.0290	0.8321
Second discharge, B2	0.8279	1.0114	0.1624	2.2079	**0.0221**	0.8335

	T_{inB3}	T_{inA5}	T_{rB}	Q_{B2}	g_{B2}	T^*_{inB3}
probe at A2	0.0	0.0	0.0	0.0	0.0	0.0
probe at A4	0.0	0.0	0.0	0.0	0.0	0.0
probe at A5	0.0	71.6	0.0	0.0	0.0	0.0
probe at B1	0.0	0.0	1.0	0.0	0.1	0.0
probe at B3	66.0	0.0	1.3	0.0	0.1	64.3
probe at B5	0.0	0.0	1.1	0.0	0.1	0.0
water-side heat checkB1	0.0	0.0	1.1	1.1	2.1	0.0
water-side heat checkB2	0.0	0.0	1.1	25.1	19.9	0.0
water-side heat checkB3	0.0	0.0	1.1	1.1	2.0	0.0
water-side heat checkA1	0.0	0.0	0.0	0.0	0.0	0.0
clamp-on hall A	0.0	0.0	0.4	0.3	0.1	0.0
clamp-on hall B	0.1	0.0	12.9	8.8	2.4	0.1
second return hall B	0.3	0.0	12.8	0.6	0.3	0.3
second discharge B2	0.2	0.0	3.7	16.8	23.9	0.2

Figure 18.2: The share of each worry's band taken off by each candidate, in percent, from the bands of Table 18.5. Shade proportional to the share.

The first choice by worry. Each worry has its own first choice, and the first choice for one worry is close to useless for most of the rest. The probe at B3 is first for two worries, the present inlet and the inlet at the raise, and takes between 0.0 and 1.3 % off the other four. The probe at A5 takes 71.6 % off hall A's hottest rack and nothing off the others. The operator's instinct, B2, takes 25.1 % off Q_{B2} and 19.9 % off g_{B2} by the water-side heat check, and the second discharge sensor takes 23.9 % off g_{B2}, the largest there; the check at B2 takes 0.0 % off rack B3. A reading's share is set by ρ_a; it is largest when an entry with no reading row carries the worry's variance and the candidate reads through that entry; elsewhere it adds precision to rows already there. The probes at B1 and B5 read other fractions and take 0.0 % off T_{inB3}. The second return is plain replication of a band already there: $(1/0.1686^2 + 1/0.3^2)^{-1/2} = 0.1470\,\mathrm{K}$. Table 18.6 sets out the first three for each worry, each row a worry and the columns the three candidates in order with the band after each.

Table 18.6: The three best candidates for each worry, with the band after each. Units as in Table 18.5.

Worry	First	Second	Third
T_{inB3}	probe at B3, 0.2821	second return, 0.8265	second discharge, 0.8279
T_{inA5}	probe at A5, 0.2876	clamp-on A, 1.0112	probe at A4, 1.0112
T_{rB}	clamp-on B, 0.1469	second return, 0.1470	second discharge, 0.1624
Q_{B2}	water-side heat check B2, 1.9880	second discharge, 2.2079	clamp-on B, 2.4196
g_{B2}	second discharge, 0.0221	water-side heat check B2, 0.0233	clamp-on B, 0.0283
T^{*}_{inB3}	probe at B3, 0.2985	second return, 0.8321	second discharge, 0.8335

18.3.3 What the reading settles

Operator question. Does the reading settle anything?

General form. The form is equation (2.16) with u the indicator of an acceptance rule of equation (2.12) applied after the outcome.

Specialized form. Equation (2.16) has here u the indicator that the requirement of equation (2.12), every inlet under 27°C at $T_s^{*} = 16$°C, ends decided one way or the other. By the means (18.4) and the covariance (18.3), the joint after an outcome is the corner of equation (12.11) with a moved mean,

$$P_a(\tilde{y}_a) = \Pr(\mathbf{T}^{*} < 27^{\circ}\mathrm{C}\,\mathbf{1}), \quad \mathbf{T}^{*} \sim \mathcal{N}\Big(\mathbf{m}^{*} + \mathbf{D}^{*}(\hat{\mathbf{z}}_a(\tilde{y}_a) - \hat{\mathbf{z}}),\ \mathbf{D}^{*}\boldsymbol{\Sigma}_a\mathbf{D}^{*\mathsf{T}}\Big), \tag{18.6}$$

and the settled share is

$$U(a) = \Pr(P_a(\tilde{y}_a) \geq 0.95 \mid \mathbf{y}) + \Pr(P_a(\tilde{y}_a) < 0.05 \mid \mathbf{y}). \tag{18.7}$$

A settled share is a probability over outcomes and not a share of a band, so the 0.640 below and the 64.3 % of Figure 18.2 are different quantities. The levels 0.95 and 0.05 are reference levels: the first is the assurance of Chapter 12, the second its mirror. They buy a statement that a reading has left the decision closed, and they fail as a description of any operator who holds a different assurance. By the law of total probability $\mathbb{E}[P_a(\tilde{y}_a) \mid \mathbf{y}]$ equals the joint now, 0.80983 (equation (12.11)). A reading does not change the expectation, only how it splits.

Building blocks. $\mathbf{m}^{*}$ and $\mathbf{D}^{*}$ (equation (12.11)); equations (18.3) and (18.4). Each share is taken over 2000 outcomes, each outcome's corner integrated numerically.

The settled share. Table 18.7 gives the result: each row a candidate, the columns the share ending below 0.05, the share ending at or above 0.95, their sum with its binomial standard error, and the expected joint with its standard error over the outcomes. The expectations stand at the 0.80983 they must keep, within those errors. A zero in 2000 outcomes has an upper 95 % bound of 0.0015.

Table 18.7: The joint at $T_s^* = 16$°C after each candidate, now 0.80983, over 2000 outcomes. Standard errors in parentheses. The pair is on separate draws from the single probes.

Candidate	Below 0.05 [–]	At or above 0.95 [–]	Settled [–]	$\mathbb{E}[P_a]$ [–]
Probe at B3	0.047	0.593	0.640 (0.0107)	0.8087 (0.0064)
Probes at B3 and at A5	0.051	0.613	0.664 (0.0106)	0.8164 (0.0068)
Probe at A5	0.005	0.000	0.005 (0.0016)	0.8118 (0.0024)
Each other candidate	0.000	0.000	0.000	0.8094 to 0.8099

Apart from the probe at B3, no candidate settles this decision on more than 1 % of outcomes; the probe at A5's 0.005 stands beside a quadrature of 0.00702. Figure 18.3 shows why: the vertical axis is the settled share, and the horizontal axis a share of a band taken off, with a filled point per candidate for the decision's own band T^*_{inB3} and a hollow point for the largest share off any of the six worries. Against the decision's own band the two agree: the probe at B3 takes 64.3 %, every other candidate at most 0.3 %, and those settle next to nothing. The hollow points show the restriction: width on a worry the decision does not read settles nothing. The probe at A5 takes the largest share of any candidate, 71.6 %, off a band the decision hardly depends on: its joint cannot exceed about 0.835, 0.8347 at −5 predictive spreads, since B3's exceedance stays and leaves 0.83700, and it falls below 0.05 only from +2.456 predictive spreads.

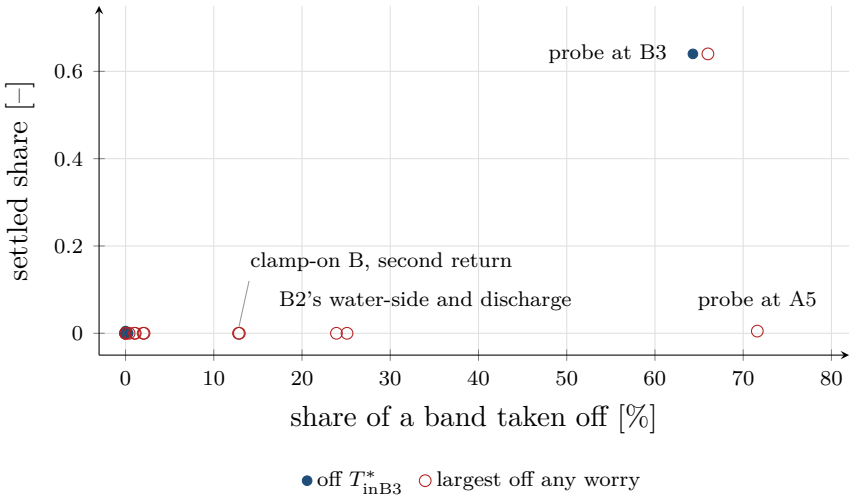


Figure 18.3: Settled share, equation (18.7), against the share of a band taken off, from Tables 18.5 and 18.7: filled, off the decision's own band; hollow, the largest off any worry. Fourteen candidates; those at the lower left are the rest.

The probe's outcomes. A one-reading candidate's joint is a curve in its outcome. Figure 18.4 draws it for the probe at B3: the horizontal axis is the reading, the solid curve is the joint (18.6), and the shaded curve behind it is the reading's predictive density from equation (18.4), on the same vertical scale in its own unit. The horizontal lines are the reference levels. The dotted line is the constructed truth's inlet, and the dashed line one reading drawn with an error about that truth. The joint reaches 0.95 at a reading of 24.293°C and 0.05 at 25.553°C, so the share

can also be read as two tails of the predictive, 0.5979 and 0.0469, sum 0.6448. That quadrature stands beside the reported 0.640 with a gap of 0.0048, within the latter's standard error of 0.0107; reweighting 400 000 draws of the exact posterior (Section 18.2.3) gives thresholds of 24.297 and 25.554°C, tails of 0.6008 and 0.0463 and a settled share of 0.6472, a gap of +0.0024 to the quadrature. As the reading falls the joint tends to about 0.968, 0.96800 at −5 predictive spreads, under the 0.96967 that rack A5 leaves, whose exceedance at the raise is 0.03033 and which this probe does not read. With both probes the same quadrature gives 0.6088 at or above 0.95 and 0.0542 below 0.05, settled 0.6630, a rise of +0.0182 from 0.6448, which is resolved. Two effects make it up: +0.0109 above, where a cool A5 lifts the ceiling, against 7.52 % of A5 outcomes on which the ceiling falls below 0.95; and +0.0074 below, where a hot A5 pushes the joint under 0.05. The decision is mostly B3's.

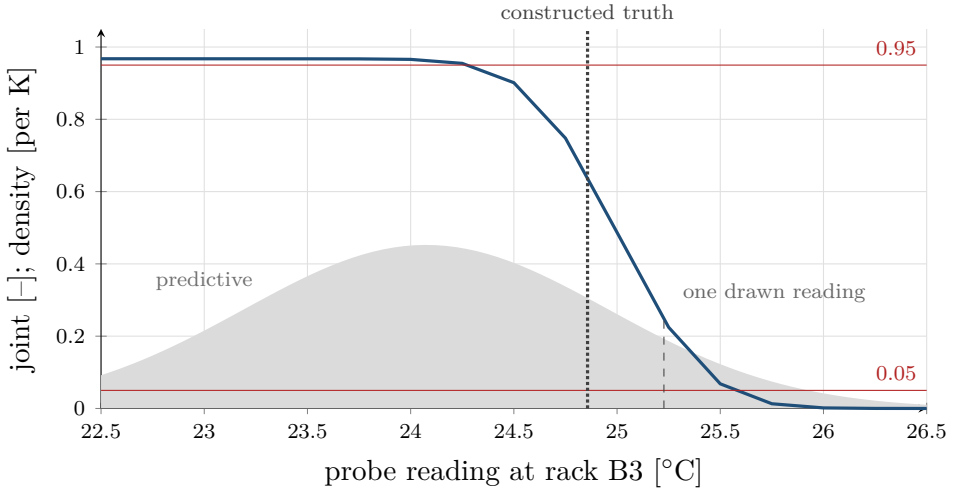


Figure 18.4: The joint at $T_s^* = 16$°C after a probe of 0.3 K at rack B3, equation (18.6), against the reading. Shaded: the normal predictive $\mathcal{N}(24.0747,\ 0.8819^2)$. Dotted: the constructed truth's inlet, 24.857°C. Dashed: one drawn reading, 25.227°C.

18.3.4 The second and third reading

Operator question. What should the second and third reading be?

General form. The form is equation (2.16) with a a set of readings and u a spread.

Specialized form. This is equation (2.16) with a a set of rows and $\mathbf{\Sigma}_a$ of equation (4.65) with their combined information; u is the band of the worry, or for the unread racks together the mean band,

$$\bar{s}(\mathbf{\Sigma}_a) = \frac{1}{6}\sum_k \sqrt{\mathbf{g}_{T_{\text{in}k}}^{\mathsf{T}}\,\mathbf{\Sigma}_a\,\mathbf{g}_{T_{\text{in}k}}}, \qquad k \in \{\text{A2}, \text{A4}, \text{A5}, \text{B1}, \text{B3}, \text{B5}\}. \tag{18.8}$$

The greedy order takes at each stage the candidate with the smallest u given those already chosen, and the best pair is searched among all pairs of the fourteen. The order buys a route for one trip that can stop at any stage, and it fails if two readings are worth more together than either first; the search for the pair tests it. The mean band is one way of pooling six racks, and it counts a band the same whichever rack holds it.

Building blocks. The rows of Table 18.2; equation (4.65); equation (18.8) with the six bands now, 1.0112, 1.0112, 1.0114, 0.8288, 0.8293 and 0.8289 K.

The orders. Table 18.8 gives the order for three worries, each row a worry, the columns the band now and after each stage. Figure 18.5 draws the same bands in three panels, the horizontal axis the number of readings. On T_{inB3} the probe gives 0.2821 and the next reading 0.2820: nothing else matters once it is read. On Q_{B2} each of the three keeps taking width. On the mean band the order goes to hall A first, A5, A4 and then A2, because hall A's unread bands are the wider (the order within hall A is a near tie, 1.0112 to 1.0114 K, and says nothing), though at the raise the rack carrying the risk is in hall B: exceedance 0.16300 at B3 against 0.03033 at A5, 0.00010 at A4 and 0.00001 at A2. The mean band and the risk rank differently. Greedy equals the best pair on each of the three.

Table 18.8: The greedy order by worry, with the band now and after each reading. The best pair has the band of the second stage on each row.

Worry	Now	First	Second	Third
T_{inB3} [K]	0.8293	probe at B3, 0.2821	second return, 0.2820	second discharge, 0.2820
Q_{B2} [kW]	2.6545	water-side heat check B2, 1.9880	second discharge, 1.7782	clamp-on B, 1.6461
Mean band, six unread racks [K]	0.9201	probe at A5, 0.7994	probe at A4, 0.6788	probe at A2, 0.5583

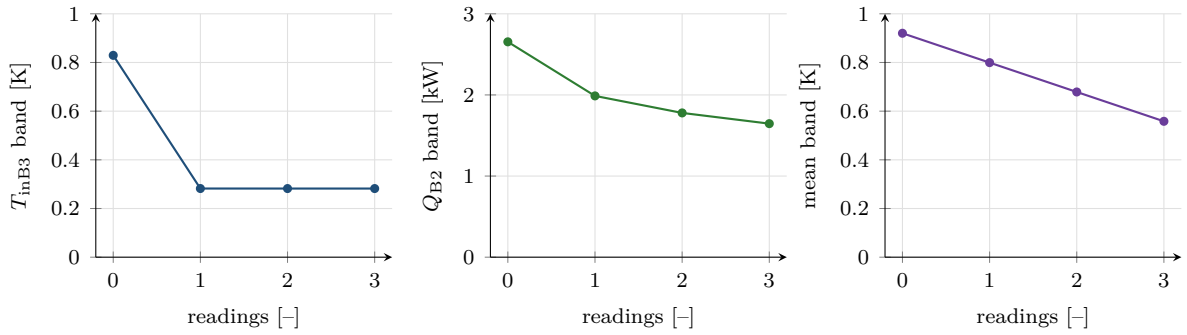


Figure 18.5: The greedy order on three worries, from Table 18.8.

18.3.5 How good a reading must be

Operator question. How good must the reading be?

General form. The form is equation (2.16) with u a spread and the instrument's spread as the design variable.

Specialized form. Equation (2.16) has here a the probe at B3 with spread σ_a or n repeats, and u of equation (18.5) with $\mathbf{a} = \mathbf{g}/\sigma_a$. Since the probe reads the worry itself, $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{a} = \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}/\sigma_a$, and the equation becomes the sum of two precisions,

$$\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_a\mathbf{g} = \left(\frac{1}{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}} + \frac{n}{\sigma_a^2}\right)^{-1}, \tag{18.9}$$

which goes to zero with σ_a. Repeats are taken with independent errors, which buys the factor n; it fails if an instrument's offset is shared across its readings, and no row of the plant can test that.

Building blocks. $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} = 0.68774\,\mathrm{K}^2$ (Section 18.3.1); the β_{B3} term $0.66591\,\mathrm{K}^2$ of the same section.

The spread and the repeats. Table 18.9 gives the band after the probe, each row a spread or a number of repeats at 0.3 K, the columns the band and the share taken off. Figure 18.6 draws the band against the spread, with the band now above and the band with β_{B3} known exactly as a dashed line. A poor reading is not worthless: a reading at 1 K still takes 23.0 % off.

Table 18.9: The band on T_{inB3} after the probe at B3, now 0.8293 K, by equation (18.9).

One reading, spread [K]	Band [K]	Share off [%]	Repeats at 0.3 K [–]	Band [K]
0.05	0.0499	94.0	1	0.2821
0.10	0.0993	88.0	2	0.2055
0.20	0.1944	76.6	3	0.1695
0.30	0.2821	66.0	12	0.0861
0.50	0.4282	48.4		
1.00	0.6383	23.0		

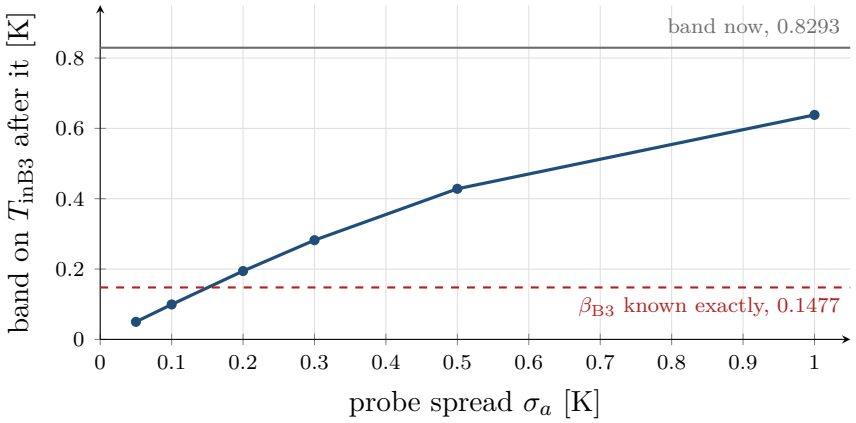


Figure 18.6: The band on T_{inB3} against the probe's spread, equation (18.9). Dashed: the band with β_{B3} known exactly, which is a perfect re-survey.

The probe and the re-survey. A perfect re-survey would fix β_{B3} and take its term off whole, leaving $0.68774 - 0.66591\,\mathrm{K}^2$, a band of 0.1477 K: the cold aisle and the return still carry their own spreads, and no re-survey, however good, goes below that. A good inlet probe does go below it, because it reads the inlet itself and not the fraction. Equation (18.9) crosses 0.1477 K at a spread of 0.1501 K: at 0.1 K the band is 0.0993 K, two readings at that spread give 0.0705 K, and twelve repeats at 0.3 K give 0.0861 K. That is what the operator wants when the worry is the inlet. A probe carries over to another load through its fraction: for the inlet at 30 kW more load the band is 0.9103 K now, 0.3100 K after the probe and 0.1606 K with the fraction known exactly.

18.4 The checks (step 5)

A check compares an answer with the calculation it stands for. The closed form is here set against the refit, and the constructed truth appears once, as one draw.

18.4.1 The width over outcomes

Over 300 outcomes, with the refit of Section 18.2.3 on each, the band stays at equation (18.5) for a candidate on its own worry. The probe at B3 on T_{inB3} has closed form 0.2821, refits' mean 0.2821 (se 0.0000), 95 % of them between 0.2821 and 0.2822; the water-side heat check at B2 on Q_{B2} has 1.9880 and 1.9880 (se 0.0000), from 1.9867 to 1.9891. The clamp-on at hall B on T_{inB3} has 0.8288 and 0.8296 (se 0.0005), from 0.8131 to 0.8453. There the outcome does change the band, by a small amount: the outcome moves the load and the load moves the rise that multiplies β_{B3}. The mean's gap of +0.0008 is mostly Monte Carlo: over 1500 refits the mean is 0.8290 (se 0.0002), a gap of +0.0002. The statement that the outcome moves the mean only belongs to the linear row, and this is its size.

The move of the mean over those outcomes should have the spread of the band taken off, $\sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} - \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_{a}\mathbf{g}}$. For the probe at B3 it is 0.7955 against 0.7798, with the standard error of a spread from 300 draws 0.0319. For the water-side heat check at B2 on Q_{B2} it is 1.6917 against 1.7590, with 0.0719.

18.4.2 Coverage over fresh errors

Over 300 sets of fresh sensor and probe errors with the truth fixed, the 95 % band on T_{inB3} covers 1.000 of the time before the probe and 0.947 after (se 0.0130), with a mean band after of 0.2823 K. Before the probe the band is carried by a survey's prior wider than the survey's actual error on this truth, so 1.000 is this truth's draw and not a calibration. After the probe the band is carried by the probe's noise, which is redrawn each time, and 0.947 lies within one standard error of the 0.95 the band states.

18.4.3 The limit of exact rows and the curvature term

The full state with its temperatures and physical rows at tolerance σ_p gives Table 18.10: each row a tolerance, the columns the band before and after the probe. Newton against Gauss–Newton after the probe at its predicted value gives 0.282109 against the refit's 0.282108 and the closed form's 0.282108, so the curvature term does not move the band at this precision.

Table 18.10: The limit of exact physical rows, the band on T_{inB3} on the full state.

Tolerance σ_p [kW or K]	Band now [K]	After the probe [K]
0.100	0.83532	0.28234
0.010	0.82936	0.28211
0.001	0.82930	0.28211

18.4.4 The answers against the stated truth

The record is generated from a stated truth (Section 18.6.4), and Table 18.11 collects the checks: each row a route or a drawn outcome, the columns the reported value, the value beside it and the gap or the truth. Each truth row draws one reading of the truth's own value with one error.

The draw for the probe at B3 reads 25.227°C, against a true inlet of 24.857°C, an error of +1.23 of the probe's spreads. The refit gives $T_{\mathrm{inB3}} = 25.093 \pm 0.282$°C, truth 24.857. The joint at $T_s^{*} = 16$°C falls from 0.81 to 0.2453 and T_{inB3}^{*} stands at 27.1986 ± 0.2985°C, while the plant at 16°C in 01:00's own state would in truth hold, with $T_{\mathrm{inB3}}^{*} = 26.8937$°C, under the limit by 0.1063 K and 1.02 of that band below the fit. The reading brought the estimate closer, 0.78 to 0.24 K off, yet the verdict crossed, since the truth's margin of 0.1063 K is under one band; of the fall from 0.81, an error-free reading accounts for the part down to 0.6452. At the raise the

estimate's error goes from -0.714 to $+0.305$ K. One reading is one draw of 0.3 K. Nothing here tests the ranking, which is a statement over outcomes.

Table 18.11: The checks and their gaps at window 73. The truth rows are the stated truth of the constructed record, one drawn reading each.

Route or outcome	Reported	Beside it	Gap or truth
Routes			
Band, probe at B3 on T_{inB3}, refits [K]	0.2821	0.2821 (se 0.0000)	0.0000
Band, water-side B2 on Q_{B2}, refits [kW]	1.9880	1.9880 (se 0.0000)	0.0000
Band, clamp-on B on T_{inB3}, refits [K]	0.8288	0.8296 (se 0.0005); 1500: 0.8290 (se 0.0002)	+0.0008; +0.0002
Settled, probe at B3, quadrature over the predictive [–]	0.640 (se 0.0107)	0.6448	+0.0048
Settled, probe at B3, exact posterior against quadrature [–]	0.6448	0.6472	+0.0024
Settled, probe at A5, quadrature [–]	0.005 (se 0.0016)	0.00702	+0.0020
Settled, both probes, quadrature [–]	0.664 (se 0.0106)	0.6630	−0.0010
Joint at the drawn reading, linear against refit [–]	0.2455	0.2453	−0.0002
Approximations			
Band after the probe, Newton [K]	0.282108	0.282109	+0.000001
Band after the probe, full state, $\sigma_p = 0.001$ [K]	0.2821	0.28211	agrees
Stated truth, one drawn outcome each			
Probe at B3 reads 25.227, true 24.857: T_{inB3} [°C]	25.093 (0.282)	joint 0.2453	24.857
Water-side B2 reads 93.122, true 94.859: Q_{B2} [kW]	92.815 (1.988)	joint 0.8094	94.859
Clamp-on B reads 296.857, true 298.297: T_{inB3} [°C]	24.081 (0.831)	joint 0.8075	24.857
Probe at A5 reads 23.999, true 23.813: T_{inA5} [°C]	23.918 (0.288)	joint 0.8345	23.813

18.5 The report (step 6)

The operator's instinct, a trip to B2, takes nothing off the rack that carries the risk of the raise. No single list exists: each worry has its own first reading. The largest width removed from a mean band over the racks points to hall A, not to the risk. And width on a worry the decision does not read settles nothing: the largest share in the matrix, the probe at A5's 71.6 %, settles under 1 % of outcomes.

18.5.1 The answers and their bands

Table 18.12 gives the answers in the order of Table 18.1. Each row is a question, and the columns give the reported answer and what stands beside it.

18.5.2 The entry in the shift log

The operator's entry as of window 73, 01:00, using nothing later, reads: "Next reading. If the worry is the raise to 16°C or the B3 forecast: handheld at rack B3 inlet, 0.3 K. The band goes 0.83 to 0.28 K, and the raise's probability is settled either way on 64 % of outcomes; the only

Table 18.12: The report at window 73, 01:00. Candidates ranked by what they do; no cost is in the record.

Question	Answer, reported	Beside it
Which reading narrows the worry?	depends on the worry: on T_{inB3} the probe at B3, 0.8293 to 0.2821 K	refits' mean 0.2821
Is there one list?	no: first choices differ across the six worries (Table 18.6)	the B3 probe takes 0.0 to 1.3 % off the other four
Does it settle anything?	probe at B3, 0.640 (se 0.0107); probe at A5, 0.005, under 1 %; the others none in 2000	quadrature 0.6448, exact 0.6472; A5 0.00702; expected joint 0.8087 (se 0.0064) against 0.80983
Second and third?	on T_{inB3} nothing further, 0.2820; on Q_{B2} 1.7782, 1.6461; on the mean band A4, A2	greedy equals the best pair on all three
How good?	0.1 K: 0.0993; 0.3 K: 0.2821; 1 K: 0.6383; three repeats at 0.3 K: 0.1695	β_{B3} known exactly: 0.1477

other reading that settles it at all is a probe at A5, on under 1 %. If the worry is the coil: water-side heat check on B2 (2.65 to 1.99 kW on its heat) or a second discharge sensor. These are not the same job. Action: technician's route not decided here."

Common mistake

Sending the technician to the place that is moving. A reading is worth what it does to a named worry's band, and that is set by the correlation ρ_a of the reading with the worry. The unit that is changing and the rack that carries the risk are different places on this plant.

18.6 The next reading (step 7)

This chapter's own next reading is the one that would improve the ranking itself. Every band a probe narrows was first set by the survey's stated spread.

18.6.1 The survey behind the ranking

The probe's share of a band and the settled share both depend on the prior $\mathcal{N}(\beta_k^{\text{s}},\ 0.08^2)$, because that prior is most of each unread rack's band. Chapter 15 tested that spread (Section 15.4.3): at 0.04 on every rack, against 0.08, on the sensitivity's draws that chapter's answer falls from 0.08287 to 0.00827, tenfold, its reported value being 0.08152, and the read racks prefer 0.04 by 1.41 log units, not decisively. A re-survey of the unread racks' fractions is one job against six entries, where a probe is one row against one rack. It is cited here and not reranked, since the ranking above holds at the stated spread.

18.6.2 The permanent instrument

A window's reading is the only kind this chapter computes. Part IV's chapters computed the instrument at the place of the question, compared by what it does to the entries. At the raise (Table 12.13) a permanent inlet sensor at B3, on six windows taken as repeats of 01:00's state, gives $\text{sd}(T^{*}_{\text{B3}}) = 0.1555$ K and a chance of 0.7175 of ending at or above 0.95, against the re-survey's 0.2652 K and 0.6479; its answer comes 30 minutes later than a probe's. Before the morning shift (Table 15.13) it alone would make B3's first passage observable. The re-survey there has a spread of 0.02. Those figures rest on fixed inner draws, where each outcome here is integrated: the same chance for the probe stands there at 0.6036 (se 0.0055) and here at 0.593,

and the quadrature of Section 18.3.3 at 0.5979. Chapter 12's second supply sensor, with T^*_{B3}'s band at 0.8322 K, is not in this chapter's list.

18.6.3 What the analysis showed

The screen pointed to B2, because B2 is what moves. The analysis finds that a trip to B2 narrows the coil and its heat and leaves the rack carrying the raise's risk where it was, and that this is not a verdict against B2: for a worry about the coil it is the first choice. What decides is the correlation of the reading with the worry. Two unread fractions hold most of two bands, so one probe removes most of either. Only one of those probes bears on the decision at the raise, so the largest width in the matrix settles almost nothing. A reading can also bring the estimate closer and still cross the verdict, as the one draw on the truth did, because the truth's margin is under one band. The route belongs to the operator, who holds the worry and the cost.

18.6.4 Limits of the model

Linear rows. The rows are linear at $\hat{\mathbf{z}}$, and the refits give the size of what that leaves out on the outcomes tried.

No costs, no time on site. The record holds neither a trip's cost nor how long a reading takes, nor the order of places on a route, so the matrix is not a schedule.

The candidates' spreads. They are stated and not measured, and the portable water-side heat check's 3 kW differs from the permanent water-side heat meter's 4 kW of Chapter 15. Section 18.3.5 shows what they do.

A reading reads what it says. A probe in the wrong place, or a reading held at an old value, breaks the row, and Chapter 17 takes up a stuck sensor. The repeats' independence is an assumption.

One window. The answer is read at 01:00's state. The truth's margin of 0.1063 K at the raise is at that state and does not last: in the truth only, the plant held at 16°C is first over the limit at window 81, 01:40 (Section 12.4.4).

Constructed readings. The record is constructed: generated from a stated truth of the plant with the stated sensor errors, and Section 18.4.4 compares the answers with that truth. It is not field data, and the agreement is a check of the method on its own model.

18.7 What to retain

Definition 18.1 (Candidate's preposterior)**.** A candidate reading is a row $(m_a(\mathbf{z}) - \tilde{y}_a)/\sigma_a$ not yet taken, with $\mathbf{a}$ its gradient over its spread. The covariance after it is $\boldsymbol{\Sigma}_a = \boldsymbol{\Sigma} - \boldsymbol{\Sigma}\mathbf{a}\mathbf{a}^{\mathsf{T}}\boldsymbol{\Sigma}/(1 + \mathbf{a}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{a})$, which holds no outcome. A worry's band after it is $\sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_a\mathbf{g}}$, or the band now times $\sqrt{1-\rho_a^2}$, with ρ_a the correlation of the reading and the worry. A decision after it is a distribution over the predictive of the outcome.

Three conclusions carry beyond the plant. A next reading has no meaning without a named worry, so a ranking is a matrix. A width and a settled share are different answers, and width on a worry the decision does not read settles nothing. And a reading's outcome is a draw: a reading can bring the estimate closer and still cross a verdict whose margin is under one band, which is no fault of the ranking.

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 next reading and where each is defined. Units for the two-hall shift plant.

Term	Expression	Units	Meaning	Defined in
Scenario row	eq. (18.1)	–	a reading row before its value	18.1.4
Row pattern	eq. (18.2)	–	which entries each row touches	18.1.4
Posterior				
Preposterior covariance	$\mathbf{\Sigma}_a$, eq. (18.3)	mixed	free of the outcome	18.2.2
Mean after an outcome	eq. (18.4)	mixed	the gain times the surprise	18.2.2
Refit	rows with r_a appended	mixed	the fit repeated on one outcome	18.2.3
Questions and answers				
Band after a reading	eq. (18.5)	per worry	the width a reading leaves	18.3.1
Correlation	ρ_a, eq. (18.5)	–	reading against worry	18.3.1
Joint after an outcome	P_a, eq. (18.6)	–	the corner with moved means	18.3.3
Settled share	$U(a)$, eq. (18.7)	–	the decision left closed	18.3.3
Mean band	eq. (18.8)	K	six racks pooled	18.3.4
Precisions add	eq. (18.9)	K^2	a probe of the worry itself	18.3.5

18.8 Exercises

Exercise 18.1. *Application.* From $\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g} = 0.68774\,\mathrm{K}^2$, $\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{a} = 2.29246\,\mathrm{K}$ and $\mathbf{a}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{a} = 7.64152$, compute the band on T_{inB3} after the probe. (By equation (18.5) the variance is $0.68774 - 2.29246^2/8.64152 = 0.07959\,\mathrm{K}^2$, a band of $0.2821\,\mathrm{K}$.)

Exercise 18.2. *Application.* Give the share of Q_{B2}'s band taken off by the water-side heat check on B2. (The band goes from 2.6545 to $1.9880\,\mathrm{kW}$, and $1 - 1.9880/2.6545$ is 25.1 %.)

Exercise 18.3. *Derivation.* Derive the band after three repeats of the probe at $0.3\,\mathrm{K}$. (Each row adds $1/0.3^2$ of precision to the worry, by equation (18.9): $(1/0.68774 + 3/0.09)^{-1}$ has the root $0.1695\,\mathrm{K}$.)

Exercise 18.4. *Derivation.* Show that a reading does not change the expected joint, and give its value here. (By total probability $\mathbb{E}[\Pr(A \mid \mathbf{y}, \tilde{y}_a) \mid \mathbf{y}] = \Pr(A \mid \mathbf{y})$, whichever row is added, so the expectation over the outcomes of the joint at $T_s^* = 16°\mathrm{C}$ stays at the present joint, 0.80983.)

Exercise 18.5. *Formulation.* Write the correlation of a reading and a worry in terms of $\mathbf{\Sigma}$, and evaluate it for the probe at B3 on T_{inB3}. (By equation (18.5) $\rho_a = \mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{a}/\sqrt{\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}\,(1 + \mathbf{a}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{a})} = 2.29246/\sqrt{0.68774 \times 8.64152}$, which is 0.9404.)

Part VI

Alarm Thresholds and Alarm Credibility

Chapter 19

Calibrating Alarm Thresholds

19.1 An alarm that has never fired

It is 01:00, window 73 of the shift. The inlet alarm on the screen sits at the 27°C limit. The four inlet sensors read 19.51, 22.03, 22.13 and 22.67°C, the largest reading of the night so far is 23.47°C, and no alarm has fired. A review has asked for an alarm setting with one false alarm a month and nine real 2 K excursions caught in ten. From the screen alone the operator concludes that the limit is safe and proposes pulling the alarm 2 K closer to the readings, to 25°C.

The screen shows a level and readings below it. It does not show how far a reading wanders with nothing wrong, and it does not show what a real excursion would do to the quantity the alarm watches. The operator's question is the one the screen cannot answer,

where should the alarm sit, and on what quantity, so that both requests hold.

The chapter answers it on the two-hall shift plant, against the reference of Chapter 16, and it does not choose a setting: it states which settings meet the review's request and on what conditions.

Jargon: decision level, false-alarm rate, probability of detection, catch, alarm rationalization, monitored quantity, plain reading, residual map, shared part, excursion route, standing alarm, alarm flood

A *decision level* is the value of a monitored quantity past which the alarm fires. The *false-alarm rate* is how often it fires with nothing wrong, per window or per month. The *probability of detection*, called here the *catch*, is the probability that it fires in a window during a stated excursion. *Alarm rationalization* is setting each alarm's level from its purpose, its false-alarm rate and its catch. The *monitored quantity* is what the alarm compares with the level: here a *plain reading* as the screen shows it, or the residual of a reading against a stated reference. The *residual map* turns a window's readings into those residuals. The *shared part* of a residual's spread is the reference's own error, the same in every window. An *excursion route* is one physical way for an inlet to rise. A *standing alarm* stays active for long periods. An *alarm flood* is more alarms in a short time than an operator can handle, by the common industry guideline (ISA-18.2, EEMUA 191) about ten new alarms in ten minutes.

Key idea

A false-alarm rate fixes a level from the spread of the monitored quantity with nothing wrong. A catch fixes a level from what the excursion does to that quantity, less its spread. Both requests hold only if the second level lies above the first, and that depends on the quantity and on what is known behind its baseline, because a quantity with a smaller spread may also be moved less by the excursion. An excursion the model can explain moves a residual little or not at all. So a catch has no meaning without its route and its duration, and a rate carried from a baseline inherits the baseline's error for as long as the baseline is used.

19.1.1 Two requests on one quantity

The review's two numbers look like one setting. They are two conditions, each giving a level on its own. The rate asks for a level far enough above the no-event distribution that its tail is the requested size. The catch asks for a level far enough below the excursion's distribution that its upper part is the requested share. Between the two levels lies every setting that meets both, and if they cross no setting does.

The width between them is a property of the quantity watched and of the information behind its baseline, not of the alarm's position. A setting on the screen cannot be moved to make the two requests compatible; a different monitored quantity, a better-known baseline or a further reading can change the distance.

> **Principle 19.1.** *State the monitored quantity and the excursion's route before the level. A level alone, a rate alone or a catch alone is not an alarm specification.*

19.1.2 The two-hall shift plant at window 73 (step 0)

The plant is the two-hall shift plant of Chapter 12, with its physical rows (12.2) and pinned ratings, on the reduced state of Chapter 4: the 19 entries (12.3), indexed by window, $\mathbf{z}_t$. The dependent block of the physical rows is triangular with a nonzero diagonal, so it is invertible, and the temperatures follow by the chain (12.4). The readings' prediction is $\mathbf{m}(\mathbf{z}_t)$, sixteen values. This chapter reads window 73, 01:00.

Modeling choices at the graph. Steady state within the window, one effective conductance per coil, pinned airflow and no active control hold as in Chapter 12, and they fail as they do there. An excursion is also taken to reach steady state within its window, and the catch is that of one window fully inside it. The request states no duration, so this is a choice: one window buys a catch that needs no assumption about how long a rise lasts, and it understates the chance of being caught in a rise that lasts longer, which Section 19.3.3 gives by duration.

The graph. Figure 19.1 draws what the question adds. The bottom row holds five scenario nodes, each pinning an entry to a value not yet applied; the second row the entries they move, the unread fraction in red; the third row the readings those entries reach; the top row the monitors that a level could be set on. The red arrows are the routes. No dashed arrow leaves β_{B3}.

19.1.3 The operator's questions (step 1)

The review's request does not state three of the things it depends on, and step 1 states them as choices.

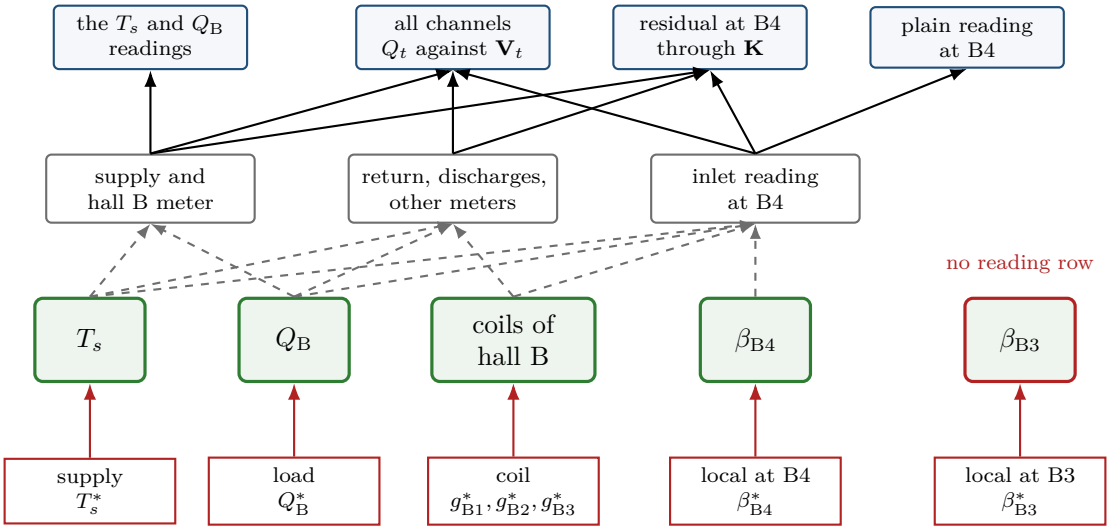


Figure 19.1: The question's graph. Free (green, with the unread fraction in red): the entries the routes move; the other entries are left out. Pinned at a scenario value (red boxes): the five routes, not yet applied. Observed (dashed): the readings, grouped. Black: the monitors, the T_s and Q_{B} readings sharing one box. The residual at B4 reads every channel through the map $\mathbf{K}$, which refinds T_s and the loads in the window.

A month. A month is the plant's own count of 8640 windows, $M = 8640$, so one false alarm a month is $\alpha = 1/M = 1.1574 \times 10^{-4}$ per window, with normal point $z_\alpha = 3.6819$. A catch of $\pi = 0.90$ has normal point $z_\pi = 1.2816$. The count treats a month as repeats of the windows modeled. It buys a per-window tail; it fails because the 8640 windows include hours and loads this night does not hold, a limit stated in Section 19.6.3.

The arithmetic of the request. On a normal spread sd and a shift of the monitored quantity of 2 K, a level exists for both requests only if $z_\alpha\,\mathrm{sd} \le 2 - z_\pi\,\mathrm{sd}$, so

$$\mathrm{sd} \le \frac{2}{z_\alpha + z_\pi} = \frac{2}{3.6819 + 1.2816} = 0.4029\,\mathrm{K}. \tag{19.1}$$

This line is arithmetic on the request, with no plant in it. On one window, any quantity whose no-event spread is above it, and which an excursion moves by at most 2 K, cannot meet the request. What counts in that spread depends on what is known behind the baseline.

Table 19.1 lists the six questions asked at window 73: the question in the operator's words, the first of them the drill question; its general form in the notation of Chapter 2; its mathematical form on this plant; its category; the computation of Section 2.3.6 that answers it; and the subsection where it is answered. The last answer also reads the record past window 73, to test the levels; the shift-log entry does not.

19.1.4 The evidence as rows (step 2)

Step 2 attaches each window's sixteen readings as rows, each Gaussian, unbiased and independent given the state, and the seventeen priors of Chapter 12; the reference stacks them over windows 1 to 24 as in Chapter 16. Three further choices enter here.

The reference, and not the one fit. An alarm's residual is taken against a stated earlier plant, the reference fit of Chapter 16, windows 1 to 24, because a window's own fit follows a change along a column of its rows and leaves little to alarm on. That buys a quantity an

Table 19.1: The questions asked at window 73, 01:00. e is a reading's residual against the reference, sd its predictive spread, s_r the shift of route r on it, $\bar{y}_t$ and sd_t the forecast of the plain reading at window t, and x the reference's shared error.

Operator question	General form	Mathematical form	Category	Computation	Answered in
Where should the alarm sit so the false-alarm rate is one per month and a real excursion of 2 K is caught nine times in ten? Calibrated decision levels: the predictive spread of the monitored entry gives the false-alarm tail, the detection power gives the miss rate.	(2.13), (2.12)	$L_\alpha = z_\alpha \,\mathrm{sd}$; $L_\pi = s_r - z_\pi \,\mathrm{sd}$; the interval $[L_\alpha, L_\pi]$	prediction under a threshold	posterior	19.3.1
Can a fixed level on the plain reading do it?	(2.21)	$\bar{y}_{144} + z_\alpha \,\mathrm{sd}_{144}$ or the mean tail, against $\bar{y}_t + 2 - z_\pi \,\mathrm{sd}_t$, on one fit and on the reference	prediction under a threshold	trend	19.3.2
Which excursions does it catch?	(2.12)	the catch by route and monitor	prediction under a threshold	posterior	19.3.3
And the rack with no sensor?	(2.12)	the shift of the virtual sensor by route	prediction under a threshold	posterior	19.3.4
Is "one a month" a number?	(2.11)	the rate over x; two windows in a row	estimation	posterior	19.3.5
What would tonight have done with it?	(2.12) per window	windows past each level; standing alarms; the first fire over fresh errors	prediction under a threshold	posterior	19.3.6

excursion can move. It fails once the 19:00–21:00 model no longer fits, and the record's count in Section 19.3.6 shows what that does.

The monitored quantity. Two are carried. The plain reading is what a level on the screen would watch. The residual is the reading against the reference with this window's supply and loads refound, $\mathbf{e}_t$ of equation (16.10). Refinding them buys a spread free of the batch load and of the shift's rise; it fails on whatever those three entries can explain.

The excursion's route. The record holds no 2 K excursion, and the plant states only how often one comes. So four physical routes are stated, each raising the inlet of the read rack B4 by 2 K at a scenario node from the reference's state at window 73: *local*, recirculation at the rack, β_{B4}; *coil*, the three coils of hall B losing conductance together; *load*, hall B's load; *supply*, the chilled-water supply. This buys a catch that can be computed; it fails for a route not on the list, and the question does not choose among them. A fifth node, local at the unread rack B3, serves the fourth question.

Independence across windows. The noise of the readings is taken as independent from window to window, as the rows state. It buys a rate per month from a tail per window. It fails for the part of the spread that does not change between windows, which Section 19.3.5 carries separately.

Which row acts on which entry. The reading rows have the pattern (16.1), with the columns grouped as there. The monitored inlet's gradient and the routes' columns are, on the same grouping,

$$\begin{gathered}\mathbf{g}_{T_{\mathrm{B4}}}^{\mathsf{T}} = (*,\ 0,\ *,\ 0,\ *,\ 0,\ *,\ 0),\\ \text{supply: column 1,}\quad \text{load: column 3,}\quad \text{coil: column 5,}\\ \text{local at B4: column 7,}\quad \text{local at B3: column 8.}\end{gathered} \tag{19.2}$$

Three facts are read here before any number. The first and third columns are the ones the residual refinds. The fifth and seventh are held at the reference. The eighth is empty in every reading row, so its entry is not observable.

19.2 The posterior of the state (step 3)

Part IV reads one fit of window 73, Section 12.2.2, and the plain reading's band is read from it here. The residual is read from the reference fit instead, for the reason of Section 19.1.4, and both the map and the load's forecast are derived below because the answers use them. With every row normalized the information matrix is $\mathbf{J}^{\mathsf{T}}\mathbf{J}$; with unnormalized rows it is $\mathbf{J}^{\mathsf{T}}\mathbf{W}\mathbf{J}$.

19.2.1 The fits cited

The one fit of window 73 gives $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$ by the normal equations (12.9) on the rows (12.8), with its iterations and covariance in Section 12.2.2. Rack B4's inlet on it stands at 22.7126°C with band 0.2821 K, read through the gradient of the chain by $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$.

The reference fit gives the shared entries $\hat{\mathbf{z}}_{\mathrm{s}}$ and their covariance $\boldsymbol{\Sigma}_{\mathrm{s}}$ by equations (16.3) to (16.5), with its iterations, entries and correlations in Section 16.2.2. Chapter 16 also found that the 19:00–21:00 model no longer fits by window 73, and that result stays in force throughout: the levels below are for a plant that still fits its reference.

19.2.2 The residual map and its two parts

The predictive covariance is $\mathbf{V}_t$ of equation (16.8), and the refound entries $\mathbf{q}_t$ solve equation (16.9). Table 19.2 runs that solve at window 73, each row an iteration, the columns the entries and the statistic Q_t; it starts from the set point and the two hall meters. The refound entries are $T_s = 13.9441 \pm 0.0710$°C, $Q_{\mathrm{A}} = 358.716 \pm 3.119$ kW and $Q_{\mathrm{B}} = 302.236 \pm 2.970$ kW.

Table 19.2: The solve (16.9) at window 73, against the reference.

Iteration [–]	T_s [°C]	Q_{A} [kW]	Q_{B} [kW]	Q_t [–]	Step norm [–]
0	14.0000	363.613	297.248	44.3897	6.99
1	13.9444	358.717	302.243	38.8990	6.72×10^{-3}
2	13.9441	358.716	302.236	38.8988	1.12×10^{-5}

The residual is $\mathbf{e}_t = \mathbf{y}_t - \mathbf{m}(\hat{\mathbf{z}}_t^{\mathrm{ref}})$, with $\hat{\mathbf{z}}_t^{\mathrm{ref}}$ the shared entries $\hat{\mathbf{z}}_{\mathrm{s}}$ and the window's $\mathbf{q}_t$ at the solution. A small change $\mathbf{d}$ in the readings, about that solution, moves the refound entries by $\mathbf{L}^{-1}\mathbf{A}_{\mathrm{q}}^{\mathsf{T}}\mathbf{V}_t^{-1}\mathbf{d}$ and the residuals by $\Delta\mathbf{e}_t = \mathbf{K}\mathbf{d}$, with the map

$$\mathbf{K} = \mathbf{I} - \mathbf{A}_{\mathrm{q}}\mathbf{L}^{-1}\mathbf{A}_{\mathrm{q}}^{\mathsf{T}}\mathbf{V}_t^{-1}. \tag{19.3}$$

Since $\mathbf{L} = \mathbf{A}_{\mathrm{q}}^{\mathsf{T}}\mathbf{V}_t^{-1}\mathbf{A}_{\mathrm{q}} + \mathbf{P}$, the map on the refound columns is

$$\mathbf{K}\mathbf{A}_{\mathrm{q}} = \mathbf{A}_{\mathrm{q}}\mathbf{L}^{-1}\mathbf{P}, \tag{19.4}$$

which is what the first and third columns of (19.2) come to. The loads have no prior row, so $\mathbf{P}$ is zero on their columns and $\mathbf{K}$ removes them entirely. The supply has its set point's row, so $\mathbf{K}$ removes nearly all of it. The trace of $\mathbf{K}$ at window 73 is 13.020, the 13 of equation (16.10) and a little from that prior row. The row of the B4 inlet has +0.9122 on its own reading and −0.3583 on the supply reading, and it has −0.0117 per kilowatt on the hall B meter, with the rest below 0.14 in size.

The readings' noise and the reference's error are independent, since the reference did not read this window, so the residuals' covariance has two parts and no cross term,

$$\mathrm{Cov}(\mathbf{e}_t) = \underbrace{\mathbf{K}\,\mathrm{diag}(\sigma_k^2)\mathbf{K}^{\mathsf{T}}}_{\text{noise, new in each window}} + \underbrace{\mathbf{K}\mathbf{A}_{\mathrm{s}}\mathbf{\Sigma}_{\mathrm{s}}\mathbf{A}_{\mathrm{s}}^{\mathsf{T}}\mathbf{K}^{\mathsf{T}}}_{\text{shared, the same in every window}}. \tag{19.5}$$

Table 19.3 gives them at window 73: each row a read inlet, the columns the residual and the square roots of the two diagonal parts and of their sum. The noise is close to 0.3 K scaled down by the map; the shared part is about a fifth of it. On (A1, A3, B2, B4) the correlations of the residuals are

$$\begin{bmatrix} 1 & -0.072 & -0.027 & -0.026 \\ -0.072 & 1 & -0.022 & -0.021 \\ -0.027 & -0.022 & 1 & -0.089 \\ -0.026 & -0.021 & -0.089 & 1 \end{bmatrix},$$

all small and negative: a reading that pulls the refound supply up leaves the other inlets' residuals lower, and the two inlets of one hall share its load as well.

19.2.3 The routes through the map

A route pins entries at a scenario node, $\mathbf{z}^*$, with the size found so that $T_{\mathrm{B4}}(\mathbf{z}^*) - T_{\mathrm{B4}}(\mathbf{z}) = 2$ K at the reference's state of window 73. The chain gives the readings' change $\mathbf{d} = \mathbf{m}(\mathbf{z}^*) - \mathbf{m}(\mathbf{z})$, and the residual's shift is $s_r = [\mathbf{K}\mathbf{d}]_{\mathrm{B4}}$. Table 19.4 gives each route a row: the entry pinned and

Table 19.3: The residuals of the read inlets at window 73 and their spreads, equation (19.5), against the reference of windows 1 to 24.

Inlet	e [K]	Noise sd [K]	Shared sd [K]	Total sd [K]
A1	−0.2325	0.2910	0.0619	0.2975
A3	−0.4696	0.2877	0.0612	0.2942
B2	+0.0515	0.2878	0.0615	0.2943
B4	+0.1949	0.2871	0.0613	0.2936

by how much, what the refound entries take from the 2 K, and the shift left. Hall B's load is 0.02731 K/kW in the inlet at this state, by the row of $\mathbf{A}_{\mathrm{q}}$.

On the local route the refound entries take 0.1755 K, the load 0.1052 and the supply 0.0703. On the coil's route, which moves six readings, they take 0.9805 K. The load's route lies along its own column and leaves nothing, in agreement with equation (19.4); the supply's leaves the small part the set point's row holds.

Table 19.4: The four routes of a 2 K rise at rack B4, from the reference's state at window 73: the scenario node, what the refound entries take and the shift left on the residual.

Route	Entry pinned	Moved by	By T_s [K]	By Q_{B} [K]	Shift s_r [K]
Local	β_{B4}	+0.1922 [–]	+0.0703	+0.1052	+1.8245
Coil	each coil of hall B, g_u	−0.1255 [–]	+0.4440	+0.5365	+1.0195
Load	Q_{B}	+73.231 [kW]	−0.0000	+2.0000	−0.0000
Supply	T_s	+2.000 [K]	+1.9597	+0.0117	+0.0287

19.2.4 The other references

The rate's uncertainty depends on how long the reference is, so the same rows (16.4) and normal equations (16.5) are solved on windows 1 to 12, 1 to 36 and 1 to 73. The shared part at B4 is 0.0874, 0.0498 and 0.0344 K, against 0.0613 K on windows 1 to 24. The last of these stacks window 73 itself, so its noise and shared parts are no longer independent, and equation (19.5) is an approximation there. It also holds g_{B2} shared across windows on which Chapter 16's stack already fails, $X^2 = 1141.78$ on 940.44; its spread is what such a fit reports, not a statement that the fit is sound.

19.2.5 The forecast of the plain reading

A plain reading moves with the load, so a fixed level must allow for the load's rise. The load's forecast is that of Chapter 13: the rows (13.8) and the forecast (13.10), here on hall B's meter through window 73, with $Q_f(t)$ and its spread. The reading's mean and variance at window t follow from the one fit with only the load moved,

$$\bar{y}_t = T_{\mathrm{B4}}(\hat{\mathbf{z}}) + a\,(Q_f(t) - \hat{Q}_{\mathrm{B}}), \qquad \mathrm{sd}_t^2 = \mathbf{g}_{\neg Q}^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g}_{\neg Q} + a^2\,\mathrm{Var}(Q_f(t)) + \sigma_T^2, \tag{19.6}$$

with $a = \partial T_{\mathrm{B4}}/\partial Q_{\mathrm{B}} = 0.02882$ K/kW at the one fit, $\mathbf{g}_{\neg Q}$ the gradient with its Q_{B} entry set to zero, and $\sigma_T = 0.3$ K. At window 73 the form is the line, with no wander term, as Chapter 13 notes of it; the later windows carry the wander.

19.2.6 The plain reading on the reference

The forecast (19.6) gives the plain reading the band of one window's fit, while the residual has 24 windows behind it. That band, 0.2984 K with the load entry left out, is the error of coils

and fractions read from one window, mostly constant over a month. Like with like, the same reading can stand on the reference,

$$\bar{y}_t^{\mathrm{ref}} = T_{\mathrm{B4}}(\bar{T}_s,\ Q_f(t),\ \hat{\mathbf{z}}_{\mathrm{s}}), \qquad \left(\mathrm{sd}_t^{\mathrm{ref}}\right)^2 = \mathbf{g}_{\mathrm{s}}^{\mathsf{T}}\boldsymbol{\Sigma}_{\mathrm{s}}\mathbf{g}_{\mathrm{s}} + \sigma_j^2 + \sigma_T^2 + a_{\mathrm{r}}^2(\mathrm{Var}(\mathrm{line}_t) + s^2), \quad (19.7)$$

with $\bar{T}_s$ the reference's mean supply, $\mathbf{g}_{\mathrm{s}}$ the gradient on the shared entries, $\sigma_j = 0.050\,\mathrm{K}$ the set point's stated jitter, $a_{\mathrm{r}} = 0.02731\,\mathrm{K/kW}$, and the load's wander of $s = 2.50\,\mathrm{kW}$ taken whole at every window. The band from the shared entries is 0.0704 K, and the spread runs 0.3225 to 0.3389 K over windows 73 to 144. Beside it, the record's own plain reading at B4 over windows 1 to 73, with a straight line taken out, has a spread of 0.283 K.

19.3 The answers (step 4)

Each answer is read from the reference's map, or for the plain reading from the one fit and from the reference, at window 73.

19.3.1 The level on the residual

Operator question. Where should the alarm sit so the false-alarm rate is one per month and a real excursion of 2 K is caught nine times in ten? Calibrated decision levels: the predictive spread of the monitored entry gives the false-alarm tail, the detection power gives the miss rate.

General form. A level on a future reading uses the predictive, equation (2.13), and an event past it the threshold form, equation (2.12).

Specialized form. The question is equation (2.13) with $y_{\mathrm{new}} = e$, the residual at B4, and $p(\mathbf{z} \mid \mathbf{y})$ the reference's normal on the shared entries; and equation (2.12) with $f = e$, $f^* = L$ and no assurance, the tail itself being read. By equation (19.5), e is normal with mean 0 with nothing wrong and mean s_r under route r, on the same spread sd, so

$$\alpha(L) = \bar{\Phi}\left(\frac{L}{\mathrm{sd}}\right), \quad \pi_r(L) = \bar{\Phi}\left(\frac{L - s_r}{\mathrm{sd}}\right); \qquad L_\alpha = z_\alpha\,\mathrm{sd}, \quad L_\pi = s_r - z_\pi\,\mathrm{sd}. \quad (19.8)$$

Every level in $[L_\alpha, L_\pi]$ meets both requests, and if $L_\pi < L_\alpha$ none does. With the shift of a smaller rise scaled along the same route, the smallest rise with both held is $2\,(z_\alpha + z_\pi)\,\mathrm{sd}/s_r$.

Building blocks. sd from Table 19.3; s_r from Table 19.4; the normal tail $\bar{\Phi}$.

At B4, the local route. The level for the rate is $L_\alpha = 0.2936 \times 3.6819 = 1.0810\,\mathrm{K}$ above the fitted value, and the level for the catch $L_\pi = 1.8245 - 0.2936 \times 1.2816 = 1.4482\,\mathrm{K}$. **The interval is** $[1.0810,\ 1.4482]\,\mathrm{K}$, 0.3673 K wide. The two requests leave room between them: at L_α the rate is one a month and the catch 0.99434; at L_π the catch is 0.90 and the rate 0.0035 a month. Figure 19.2 draws this: the horizontal axis is the residual, the two curves are the no-event density and the local route's, and the shaded band is the interval. The no-event tail past the left edge is too small to see, which is what one a month looks like, and nine tenths of the excursion curve lies right of the band's right edge.

The read inlets. Table 19.5 gives each read inlet a row with its own local route: its spread, level for the rate, shift, level for the catch, and the catch at the level for the rate. The levels for the rate run from 1.0810 to 1.0955 K, and the catches from 0.99434 to 0.99573. The four read racks need nearly the same level.

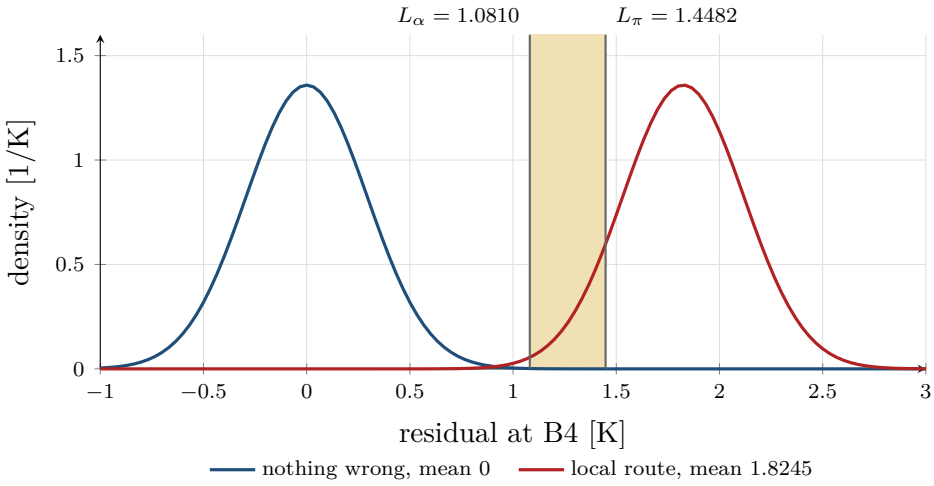


Figure 19.2: The two densities of the residual at B4 at window 73: normal with spread 0.2936 K and means 0 and 1.8245 K. Shaded: the levels that meet both requests on the local route, equation (19.8).

Table 19.5: The interval at each read inlet, local route, at window 73, equation (19.8).

Inlet	sd [K]	L_α [K]	s_r [K]	L_π [K]	$\pi_r(L_\alpha)$ [–]	false alarms at L_π [per month]
A1	0.2975	1.0955	1.8778	1.4965	0.99573	0.00212
A3	0.2942	1.0831	1.8334	1.4564	0.99462	0.00319
B2	0.2943	1.0836	1.8339	1.4567	0.99460	0.00321
B4	0.2936	1.0810	1.8245	1.4482	0.99434	0.00350

The size of a rise. Figure 19.3 draws the catch at L_α against the size of the rise, local route: the horizontal axis is the rise at the inlet, the solid curve the present sensor. The dotted curve belongs to Section 19.6.1. A 1 K rise is caught 0.2827 of the time and a 1.5 K rise 0.8362. The smallest rise with both requests held is $2 \times 4.9635 \times 0.2936/1.8245 = 1.5974$ K.

19.3.2 The level on the plain reading

Operator question. Can a fixed level on the plain reading do it? The screen's alarm watches a reading, not a residual, and the reading rises with the night.

General form. The reading's forecast is the third line of the trend form, equation (2.21).

Specialized form. The question is equation (2.21) with f_t the plain reading at B4 at window t and f^* the level, the reading's own noise added as a predictive. Three choices enter here. A straight load forecast with the coil and fractions held buys a forecast from what is known at 01:00; it fails if the plant keeps changing, which is what the record does. The rate is read either at the forecast's highest point, the peak at window 144, which holds it at every window and overspends on the others, or as the mean tail over the span, which is what "a month" averages and which lets the hot hours run above it. The catch is read either at the worst window, which holds it whenever the rise comes, or as a mean over the span, which fails on the cool windows. And the span itself decides how much of the load's ramp the level must cover. With $\mathcal{W}$ the

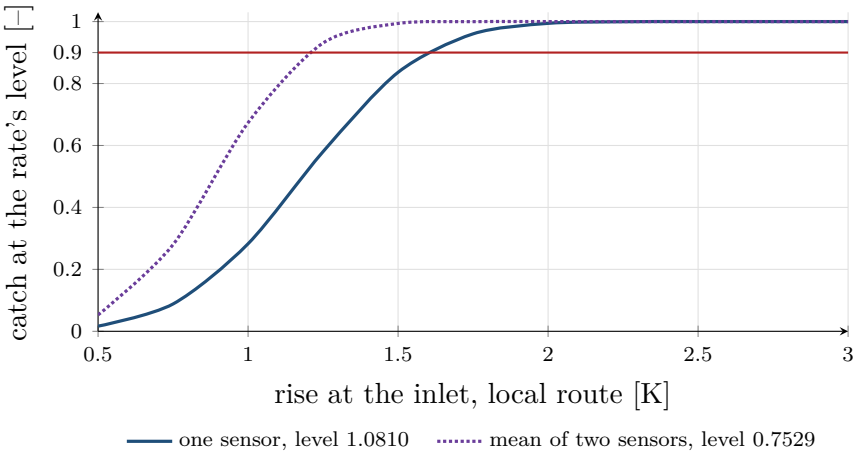


Figure 19.3: The catch at the level for one false alarm a month against the size of the rise, local route, from equation (19.8) with the shift scaled along the route. The red line is the request's 0.90.

span and the mean and spread those of (19.6) or (19.7),

$$\begin{aligned} L_\alpha^{\text{peak}} &= \bar{y}_{144} + z_\alpha \,\text{sd}_{144}, & \frac{1}{|\mathcal{W}|}\sum_{t\in\mathcal{W}} \bar{\Phi}\left(\frac{L_\alpha^{\text{mean}} - \bar{y}_t}{\text{sd}_t}\right) &= \alpha, \\ L_{\pi,t} &= \bar{y}_t + 2 - z_\pi \,\text{sd}_t, & \pi_t(L) &= \bar{\Phi}\left(\frac{L - \bar{y}_t - 2}{\text{sd}_t}\right), \end{aligned} \tag{19.9}$$

and the mean catch is the mean of π_t. A plain reading moves by the full 2 K on every route, so the catch does not depend on the route. The constant part of the spread does not average over a month, so the rate is also read over it, $M\bar{\Phi}((L - \bar{y}_t - x)/\text{sd}_\text{n})$ with x normal on the constant part.

Building blocks. The forecasts (19.6) and (19.7); the line (19.1); the normal tail.

On one window's fit. Table 19.7 gives four windows a row each: the load forecast, the reading's mean and spread, the spread's excess over the line, the state's and load's parts of the spread, the window's own two levels and the catch at the peak level. Every spread is over the line (19.1), by 0.0227 to 0.0425 K, so on one window's fit no level meets both requests even on one window: its own levels cross by 0.112 to 0.211 K. Across the night the peak level for the rate, 24.860°C, stands 0.590 K above window 73's level for the catch, 24.270°C; the mean's rise is 0.404 K of that, and the rest, 0.185 K, is the spread's. At 24.860°C the catch is 0.45849 at window 73 and 0.79052 at window 144.

The parts say where the excess comes from. The sensor alone is 0.3000 K, and sensor and load alone are 0.3035 to 0.3307 K, under the line. The state's part, 0.2984 K, is what crosses it: what one window's fit does not know of the coils and fractions. It is not a spread the reading shows from window to window, and it would not stay the same over a month's further evidence.

The constant part. That part also means the rate at a fixed level is not one number. At 24.860°C the peak's rate has mean one a month and runs from 7.46×10^{-8} to 6.132 a month with the constant part at ∓ 1.96 of its spreads.

On the reference. With the band of equation (19.7) every spread is under the line, and each window's own interval is open, 0.318 to 0.399 K wide, 0.399 K at window 73. A level fixed

over a span is another matter. Table 19.6 gives each baseline and span a row, and each pair of conventions a column; the entry is the level for the catch less the level for the rate, positive when a level exists. **The result is: empty on one window's fit; on a reference it depends on the rate and catch conventions and on how much of the load's ramp is covered.** At the mean-tail level on the reference over windows 73 to 144, 23.945°C, the rate has mean one a month, 0.138 to 3.400 over the constant part of 0.0704 K, and the catch runs 0.945, 0.976, 0.990 and 0.996 at windows 73, 97, 121 and 144. The same level on the reference is not offered as a setting: it holds the request by averaging the hot hours' tail with the cool hours'.

Table 19.6: Level for the catch less level for the rate on the plain reading at B4, for a fixed level. Positive means a level exists.

Baseline, span	Peak tail, worst catch [K]	Mean tail, worst catch [K]	Mean tail, mean catch [K]
One fit, windows 73–144	−0.590	−0.425	−0.257
Reference, windows 73–144	−0.044	+0.103	+0.262
Reference, windows 1–144	−0.433	−0.223	+0.069

Table 19.7: The plain reading at B4, forecast from window 73: equations (19.6) and (19.9), against the line 0.4029 K, with the catch of a 2 K rise at the fixed level 24.860°C.

Window	Q_f [kW]	$\bar{y}_t$ [°C]	sd_t [K]	Over line [K]	State [K]	Load [K]	$L_{\alpha,t}$ [°C]	$L_{\pi,t}$ [°C]	Catch [–]
73, 01:00	299.57	22.815	0.4256	+0.0227	0.2984	0.0459	24.382	24.270	0.45849
97, 03:00	304.32	22.952	0.4348	+0.0319	0.2984	0.1002	24.553	24.395	0.58407
121, 05:00	309.06	23.089	0.4395	+0.0366	0.2984	0.1189	24.707	24.525	0.69883
144, 06:55	313.61	23.220	0.4454	+0.0425	0.2984	0.1392	24.860	24.649	0.79052

Figure 19.4 draws the one fit's case: the horizontal axis is the window, the vertical axis the inlet temperature. From the bottom, the lines are the reading's mean, the window's level for the catch, the window's level for the rate, the mean with 2 K added, and the peak level and the screen's 27°C limit, drawn level. The level for the rate lies above the level for the catch at all four windows, and the peak level meets the rate's line at window 144.

The screen's level and the proposal. The screen's 27°C stands 3.780 K above the peak's mean and catches 3.21×10^{-5} of a 2 K rise there. The operator's 25°C is a real setting with a real result: on one window's fit, at the peak, its rate has mean 0.2770 a month and reaches 1.297 a month with the constant part at +1.96 of its spreads, and its catch is 0.33211 at window 73 and 0.68901 at window 144. Moving there meets the rate on that fit and misses the catch. On the reference the level below stands at 23.945°C, because the baseline is better known and lower, so 25°C is not close to a level that meets both.

19.3.3 Monitors against routes

Operator question. Which excursions does it catch? The interval above belongs to one monitor and one route.

General form. Each cell is the threshold form, equation (2.12), with its own monitor as f.

Specialized form. The question is equation (2.12) with f one of five monitors, each at the level giving one false alarm a month, and $\mathbf{z}$ at each route's scenario node. For the residual at B4 this is $\pi_r(L_\alpha)$ of equation (19.8). For all channels together, the statistic Q_t of equation (16.10)

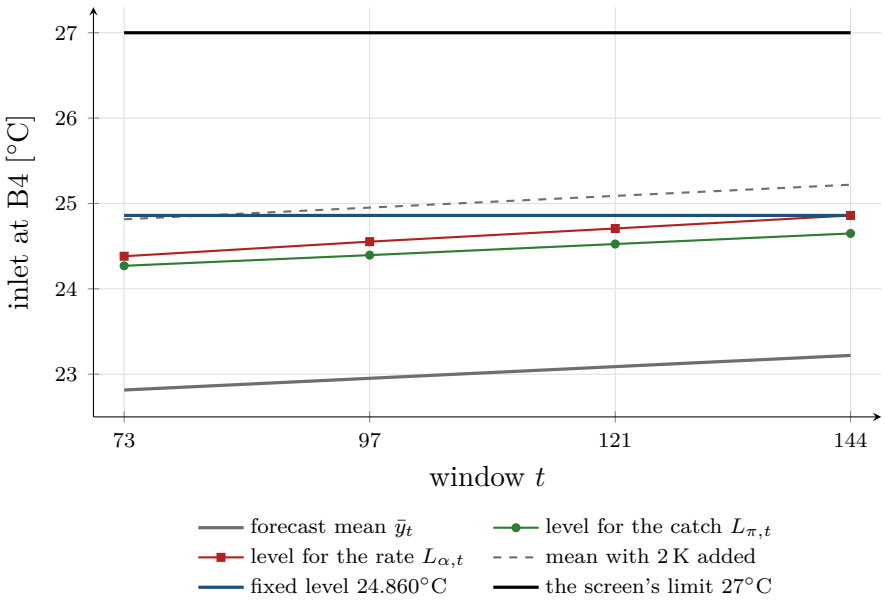


Figure 19.4: The plain reading's forecast at B4 on one window's fit, from equations (19.6) and (19.9). A level meets the rate on or above the red line, and the catch on or below the green line.

has level χ^2_{13} at the same tail, 40.478, and its catch is the noncentral chi-square tail past that with noncentrality $(\mathbf{Kd})^{\mathsf{T}}\mathbf{V}_t^{-1}(\mathbf{Kd})$. For the plain reading it is π_t of equation (19.9) at the fixed level. For the T_s reading the shift is d_{T_s} on its own spread of 0.1 K; for the Q_{B} reading $d_{Q_{\mathrm{B}}}$ on a spread that adds the meter's 4 kW to the load forecast's at the peak.

Building blocks. The map (19.3); each route's **d** of Section 19.2.3; the normal and noncentral chi-square tails.

The table. Table 19.9 gives each route a row and each monitor a column, and the largest catch of a route in bold. The residual at B4 meets the catch only on the local route. Its catch on the coil's route, 0.41709, fails it: the coil's shift of 1.0195 K gives $L_\pi = 0.6433$ K, below L_α, with an empty interval, and setting the level there would run at 122.9 false alarms a month. **On the load and supply routes the residual is blind**, with catches of 0.00012 and 0.00017 against a false-alarm rate of 0.00012 per window. The residual is free of the load because it refinds the load, and for the same reason it cannot see a load excursion: equation (19.4) gives the shift of −0.0000 K in Table 19.4, and the model explains it. Beside the linear shifts stand the residuals refound at each route's own state: coil 1.0257 K on a spread of 0.2946 K, catch 0.42553 against 0.41709; load on a spread of 0.2977 K, catch 0.00014; local 0.99432; supply 0.00017. No decision moves on those gaps.

A plain reading sees all four alike. The table's columns are at the one fit's 24.860°C, at no catch over 0.79052. At the reference's mean-tail level of Section 19.3.2 the same reading catches every route 0.945 to 0.996, on that level's conventions. So among monitors set on the worst window no one does the job on all four routes and each route has one that does, and a plain level on a reference covers all four only by way of a mean rate.

The duration. The table is for one window. Table 19.8 gives the probability of at least one window past the level in a rise lasting n windows, over the constant part, each row a monitor and route and each column a duration. The request does not state a duration, and the answer to "caught" depends on it: the coil's route, 0.417 in one window, passes 0.90 inside half an hour

at the level for one false alarm a month, and the plain reading at 24.860°C does so within the hour.

Table 19.8: The catch of a 2 K rise lasting n five-minute windows: at least one window past the level.

Monitor, route	$n = 1$ [–]	$n = 2$ [–]	$n = 3$ [–]	$n = 6$ [–]	$n = 12$ [–]
Plain at 24.860°C, window 73	0.458	0.626	0.713	0.830	0.905
Residual at B4, coil	0.417	0.654	0.790	0.949	0.996
Residual at B4, local	0.99434	0.99995	1.00000	1.00000	1.00000

Table 19.9: The catch of a 2 K rise at B4, by route and monitor, each monitor set for one false alarm a month.

Route	Residual at B4 [–]	All channels Q_t [–]	Plain, window 73 [–]	Plain, window 144 [–]	T_s reading [–]	Q_{B} reading [–]
Local	**0.99434**	0.79009	0.45849	0.79052	0.00012	0.00012
Coil	0.41709	**1.00000**	0.45849	0.79052	0.00012	0.00012
Load	0.00012	0.00012	0.45849	0.79052	0.00012	**1.00000**
Supply	0.00017	0.00016	0.45849	0.79052	**1.00000**	0.00012

Figure 19.5 draws the residual's column as a curve for every level: the horizontal axis is the level, one curve per route. The local and coil curves stand apart, and the two blind routes lie on top of each other at the bottom, with the coil crossing 0.9 left of L_α and the local crossing it right of it.

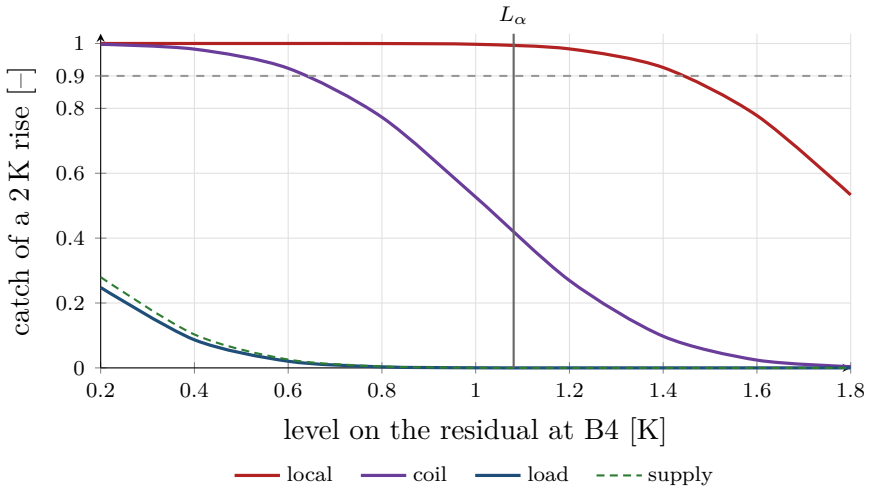


Figure 19.5: The catch $\pi_r(L)$ of equation (19.8) against the level, one curve per route, with spread 0.2936 K and the shifts of Table 19.4. The vertical line is the level for the rate; the dashed line the request's 0.90.

Sharing the rate. A set of monitors needs each route covered, and one false alarm a month then has to be divided among them. Dividing it equally is the Bonferroni level of Chapter 9. Table 19.10 gives one, four and sixteen channels a row each, with the normal point, the residual level at B4 and the local catch. At sixteen channels the level 1.2732 K is still below $L_\pi = 1.4482$ K, so on the local route a set of channels still holds the request. Those columns say nothing of the other routes.

Table 19.10: One false alarm a month shared among channels, residual at B4, local route.

Channels [–]	z [–]	Level [K]	Local catch [–]
1	3.6819	1.0810	0.99434
4	4.0213	1.1806	0.98585
16	4.3366	1.2732	0.96979

19.3.4 The unread rack

Operator question. And the rack with no sensor? Its inlet exists only as a virtual sensor, and an alarm on it would watch an estimate.

General form. The event is the threshold form, equation (2.12).

Specialized form. The question is equation (2.12) with $f = T_{\mathrm{B3}}(\hat{\mathbf{z}}_t)$, the estimate of each window's fit, the state at a scenario node of each route at B3. A route moves the estimate only through the readings, so the shift is the estimate's change when the readings carry $\mathbf{m}(\mathbf{z}^*)$ with no error. The virtual sensor's band is the one fit's $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$, split as in equation (12.12): the entry of β_{B3} against the rest.

Building blocks. $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$ of the one fit; each route at B3; the empty column of (19.2).

The rack. The virtual sensor reads 24.0747°C with a band of 0.8293 K, of which 0.8160 K is β_{B3} and 0.1477 K the rest. Over windows 1 to 24 the estimate moves with a window-to-window spread of 0.1903 K. These answer different questions. The band says how far the estimate may be from the inlet, an error that does not move from window to window; a level drawn from it would be set by an error the alarm can never see change. A local excursion at B3, with β_{B3} up 0.1961, moves no reading at all, the largest change being 0. So **its catch equals the false-alarm rate of whatever level is set**, on any monitor, 1.16×10^{-4} at L_α. The coil, load and supply routes do reach it, moving the estimate by +1.9841, +2.0020 and +1.9992 K of the 2 K. This is the rack Chapter 15 follows, and on it the one route the alarm cannot hear is the one that needs no change elsewhere in the plant.

19.3.5 How sure the rate is

Operator question. Is "one a month" a number? The rate comes from a reference, and the reference has an error that stays for as long as the level is used.

General form. The rate is a function of the state, with a central interval by equation (2.11).

Specialized form. The question is equation (2.11) with f the rate given the reference's shared error x at B4, whose posterior is the shared part of equation (19.5), $x \sim \mathcal{N}(0, \mathrm{sd}_{\mathrm{s}}^2)$, and with the noise spread sd_{n} left to the window,

$$f(x) = M\,\bar{\Phi}\!\left(\frac{L-x}{\mathrm{sd}_{\mathrm{n}}}\right), \qquad f_2(x) = M\,\bar{\Phi}\!\left(\frac{L-x}{\mathrm{sd}_{\mathrm{n}}}\right)^2. \tag{19.10}$$

The second is the alarm that fires on two windows in a row, with independent noise and the same x. Choosing a constant x over a month is the modeling choice: it buys the rate's band from the reference's posterior, and it fails if the reference is renewed, or if other slow errors exist that the reference does not hold.

Building blocks. $\mathrm{sd_n} = 0.2871$ and $\mathrm{sd_s} = 0.0613$ K from Table 19.3; 200 000 draws of x.

The rate at the level. At $L_\alpha = 1.0810$ K the rate has mean 1.0000 a month, by construction, since L_α was set on the total spread, median 0.7203, and a 95 % range of 0.1235 to 3.5365.
"One a month" is a mean over a factor of 28.6**.** Ignoring the shared part would give a level of 1.0571 K, which runs at a mean of 1.3713 a month. Figure 19.6 draws $f(x)$: the horizontal axis is the shared error in its own spreads, the vertical axis the false alarms a month on a logarithmic scale, one curve for one window and one for two in a row. The month sits at one point on the curve. The count of alarms would not say which, but the mean residual over clean windows would: its standard error over one day's 288 windows is 0.0169 K, against the shared part's 0.0613 K.

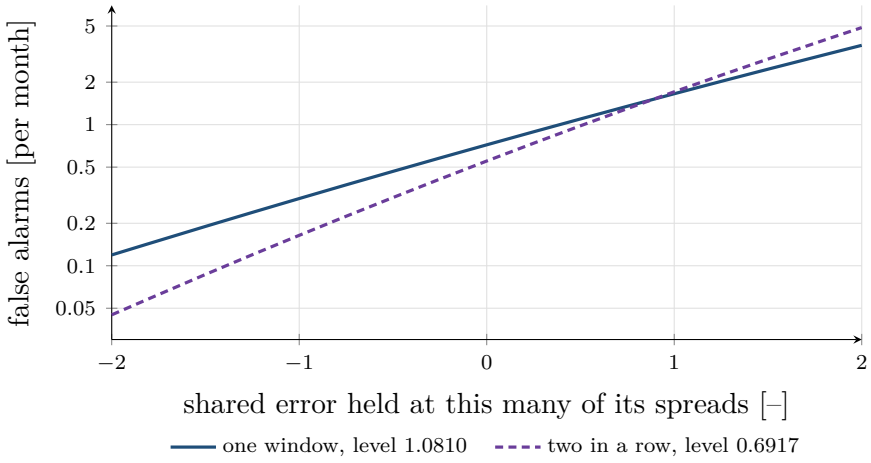


Figure 19.6: The rate of equation (19.10) against the shared error held for the month, at B4, $\mathrm{sd_n} = 0.2871$, $\mathrm{sd_s} = 0.0613$ K. Both levels have a mean of one a month over the shared error.

By reference. Table 19.11 gives each reference of Section 19.2.4 a row: spread, level, band and factor. The band narrows steadily with the length, and the level falls slowly. The narrowing is real for windows 1 to 36, which end before Chapter 16's first verdict. For windows 1 to 73 it carries the warning given there.

Table 19.11: The rate at the level for one a month, by reference. All at B4, window 73.

Reference	sd [K]	Shared sd [K]	Level [K]	Rate, 2.5–97.5 % [per month]	Factor [–]
Windows 1–12	0.3014	0.0874	1.1097	0.039–4.898	125.6
Windows 1–24	0.2936	0.0613	1.0810	0.1235–3.5365	28.6
Windows 1–36	0.2910	0.0498	1.0714	0.197–2.937	14.9
Windows 1–73	0.2882	0.0344	1.0612	0.347–2.220	6.4

Two windows in a row. Taking the noise alone, the tail per window is $\sqrt{\alpha} = 1.0758 \times 10^{-2}$, the level 0.6600 K, and the catch 0.99995 for a delay of one window. That reads as a pure gain until the shared part is carried by equation (19.10): the same level runs at a mean of 1.7315 a month, in a range of 0.0930 to 7.7977. Set for a mean of one a month over the shared error, like with like with the single window, the level is 0.6917 K, with median 0.5532 and range 0.0472 to 4.6770, a factor of 99.1. Its local catch over the shared error is 0.99989 and its coil catch 0.75541. At the same span, by the second window, the single-window alarm has local 0.99995

and coil 0.65356, and two in a row 0.99989 and 0.75534. So at the same mean rate the delay gives up a little on the local route and gains on the coil's, and it costs certainty about the rate, because the noise averages across windows and the shared error does not.

19.3.6 Tonight against the levels

Operator question. What would tonight have done with it? The levels have been read from the reference; the record through window 73 can be held against them, and the rest of the night tests them.

General form. Each window's alarm is the threshold form, equation (2.12).

Specialized form. The question is equation (2.12) at each window t, with f the residual $e_{k,t}$ of equation (19.3) on channel k and f^* the level L_α of that channel's spread, read at window 73 and kept: one-sided on the inlets, as asked, and two-sided on all sixteen channels, as a board of alarms. That board has 32 tails, each at one a month, so it is 32 false alarms a month on a plant at its reference, not the shared rate of Table 19.10. The statistic Q_t of equation (16.10) stands beside it at 40.478. A time at which an alarm first fires is one night's draw, so sensor errors are redrawn about the constructed truth, 300 nights, the reference refitted on each, and this is a check.

Building blocks. The map at each window; the channels' levels; fresh nights.

The channels. Table 19.12 gives each channel that passed its level a row: the level, the first window past it, the windows past it through window 73 and in the night, and for the inlets the largest residual. The other ten channels never pass. The board's 32 tails predict 0.444 channel-windows past a level in these 120 windows; the night holds 150, 8 of them through window 73. The statistic Q_t passes 40.478 twice through window 73, at windows 70 and 72, and stands at 38.90 at window 73. Those numbers do not say that the level is wrong. They say that the 19:00–21:00 model no longer fits, as Chapter 16 found, and a rate holds only while it does. From the one fit's g_{B2} at window 73, -0.1444 against the reference's $+0.0008$, the map leaves $+0.363$ K on the residual at B4 and the rate at 1.0810 K becomes 53.7 a month, 12.1 to 186.2 over that entry's band.

Table 19.12: The channels past their levels, windows 25 to 144, with residuals against the reference. Inlets match one-sided and two-sided here.

Channel	Level	First	Past through window 73 [–]	Past in the night [–]	Largest residual on one-sided inlets [K]
$T_{d,\mathrm{B2}}$	1.0990 K	53 (23:20)	7	72	–
T_{B2}	1.0836 K	72 (00:55)	1	6	+1.3378 (window 130)
T_{B4}	1.0810 K	138 (06:25)	0	2	+1.2351 (window 140)
T_s	0.2566 K	83 (01:50)	0	29	–
Q_{B}	9.8566 kW	100 (03:15)	0	34	–
Q_{A}	9.2134 kW	103 (03:30)	0	7	–
All channels, Q_t	40.478 [–]	70 (00:45)	2	65	–

Standing alarms and concurrent alarms. Figure 19.7 draws the record. In the upper panel the horizontal axis is the window and the vertical axis the residual in K, for the discharge of B2 and the inlet at B4, each with its level; the dotted lines are what the constructed truth's derate alone would leave on those residuals, a check marker for Section 19.4.4. The lower panel

counts the channels past their levels in each window. The discharge's alarm comes and goes at first: 1 window past in the 23:00 hour, 5 in the 00:00 hour and 6 in the 01:00 hour, windows 73 to 84, so through window 73 it is past on 7 of the 21 windows since its first, the 01:00 hour having begun. From the 02:00 hour on it is past on all 12 of each hour. As a test of the level, after its first fire it stays past on 0.783 of windows: a **standing alarm**. Channels past by the hour run 0.00, 0.00, 0.08, 0.50, 0.67, 1.00, 1.75, 2.17, 2.83, 3.50, and the largest in one window is 5, at window 123: T_s, $T_{d,\mathrm{B2}}$, Q_A, Q_B and T_B2. Those are several concurrent alarms, mostly standing, and not an alarm flood: new activations reach at most 4 in any ten minutes, 11 in any hour and 40 in the night, below the guideline's ten in ten minutes.

Of the 150, 80 channel-windows lie on the three channels a coil derate on unit B2, the change Section 15.2.2 selects by the method of Chapter 10, moves directly, and 70 on the three it does not: the supply and the two hall meters. Most of those 70 are the smear of Chapter 16, the refound entries moving to explain what they cannot, and in the largest window those channels are three of the five. The hall-A meter's bias has a share: with it taken off the same record gives Q_A 4, Q_B 39 and 152 in all. Each of those three alarms, seen alone, points at an instrument the change did not move.

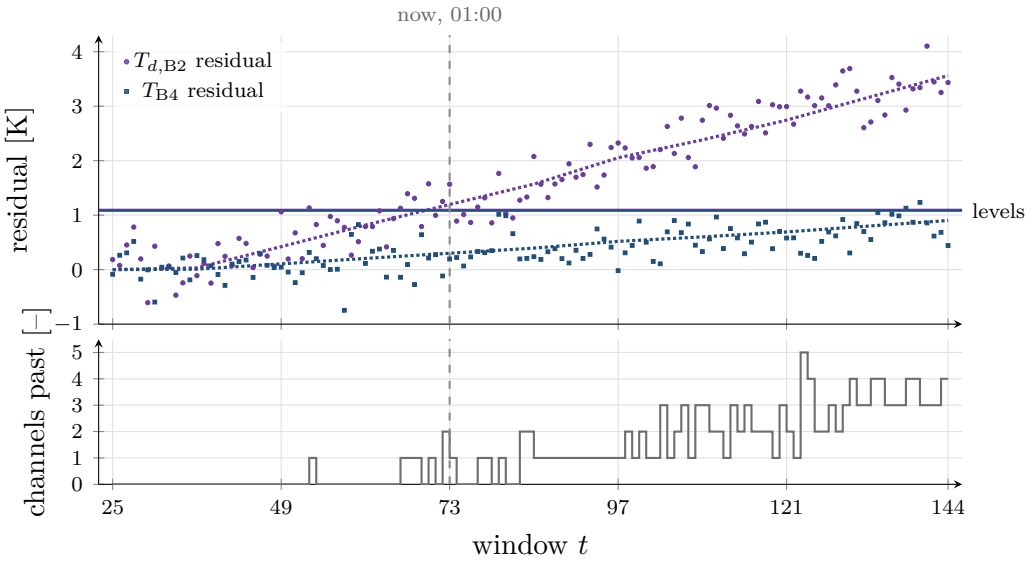


Figure 19.7: The record's channels against their levels, windows 25 to 144, read through the map (19.3). Upper: the two residuals with their levels, 1.0990 and 1.0810 K, in the same colors; dotted, the stated truth's derate alone. Lower: the number of channels past a two-sided level. Only windows through 73 had happened at 01:00.

The plain reading tonight. Against 24.860°C the record's reading at B4 is past on one window, the largest of the night, 25.15°C at window 140, and that reading is also past 25°C. Against the reference's 23.945°C it is past on none through window 73 and on 27 in the night, first at window 100. The forecasts held the coil, so these are not tests of a 2 K excursion either.

The first fire. Over 300 fresh nights the discharge's alarm first fires at a median of window 61.0, with a 95 % range of 49.0 to 69.0, and this night's window 53 is early in that range: a share of 0.113 (se 0.018) of nights come at 53 or earlier. Every night has fired by window 73. The B4 alarm first fires at a median of window 108.0, range 67.0 to 134.5, and 2 nights of the 300 have no fire, which the range's upper end falls short of; this night's window 138 is late, 0.020 of nights coming at 138 or later.

Common mistake

Treating a constant error as window-to-window noise, or the reverse. The band of a virtual sensor, and the state's part of a plain reading's spread, do not move between windows, so a rate read through them is a mean over a month's single draw and must carry that band. One quiet night's readings hold nothing about a baseline's own error or about what an excursion does.

19.4 The checks (step 5)

The checks test the reported tail against draws through the plant, the routes through the map against the plant, the fit against its limit, and the levels against their assumptions; the stated truth closes the section.

19.4.1 The limit of exact rows and the curvature term

The chapter runs on the reduced state, the limit $\sigma_p \to 0$ of the full state with its physical rows (12.2). Table 19.13 repeats the fit of window 73 on the full state: each row a tolerance, the columns rack B4's inlet with its spread and the iterations. The curvature term, checked numerically by Newton's covariance, moves the band of T_{B4} from 0.282107 to 0.282105 K, which no statement here feels.

Table 19.13: The limit of exact physical rows on the full state at window 73, against the reduced state.

Tolerance σ_p [kW or K]	T_{B4} [°C]	sd [K]	Iterations [–]
0.100	22.71202	0.28234	3
0.010	22.71261	0.28211	3
0.001	22.71261	0.28211	2
Reduced	22.71261	0.28211	–

19.4.2 The tail and the routes against the plant

The rate rests on a normal tail through a linear map. Over 200 000 draws, with the shared entries drawn from their posterior, the plant solved on each and fresh sensor errors added, the residual at B4 has spread 0.2940 against 0.2936 K, excess kurtosis +0.003, and a tail at three spreads of 1.415×10^{-3} (se 8.4×10^{-5}) against 1.350×10^{-3}. At the level itself 4 000 000 such draws give 441 past it, 1.102×10^{-4} (se 5.2×10^{-6}) against 1.157×10^{-4}, or 0.953 a month, a gap of 1.05 standard errors. The draws cannot tell the plant's curvature from the normal form at this level.

The routes are run from the constructed truth with the meter's bias through the plant and 300 fresh windows, each with its supply and loads refound. Table 19.15 carries the results: each route's share of windows past the level, against the closed form at its mean on the noise spread, the shared part being fixed in one reference. The largest gap is on the coil's route, 0.373 against 0.390, inside its standard error of 0.028. The local route's standard error is 0.005. The means differ from the shifts of Table 19.4, 1.7801 against 1.8245 and 1.0010 against 1.0195, mostly for two reasons: this reference's no-event offset at window 73, −0.0281 K, a draw of its shared error under the truth, and the Monte Carlo error of a mean over 300 windows, 0.0166 K. From the truth with no noise the shifts over that offset are +1.8244 and +1.0241 K. Given that offset, this reference's own rate at the level is 0.484 a month.

19.4.3 Sensitivity to the spreads and to the tail

Table 19.14 varies what the levels assume: each row is a case, the columns the spread or the error law, the level for the rate, the catch at it, and the rate if the level were kept at the 1.0810 K computed for the stated spreads and the stated normal errors. **The rate rests on a tail no shift has measured.** With reading spreads 1.2 times larger, the kept level runs at 8.678 a month, and at 1.5 times, at 55.976 a month and an empty interval. The heavier law has two forms, and the choice between them is stated. With a Student-t law of the same spread on each reading's error, carried through the map with the shared part normal, the kept level runs at 2.36, 7.17 and 18.32 a month on 30, 10 and 5 degrees of freedom, by 2×10^7 draws. A level raised to hold the rate catches 0.9864, 0.9282 and 0.2310. With the same law put on the total residual, the map's averaging is lost: the rate runs at 2.764, 9.023 and 21.986 and the catch is 0.9841, 0.8981 and 0.1155. So at 10 degrees the request holds on the reading-error law and misses by a little on the total-residual law, and at 5 it fails on both. The errors of the rows are on the readings, so the first form is the one reported.

Table 19.14: The alarm at B4, local route, under different spreads and error laws.

Case	Spread [K]	Level [K]	Catch [–]	Rate at 1.0810 K [per month]
Reading spreads ×0.8	0.2377	0.8753	0.99997	0.024
Reading spreads ×1.0	0.2936	1.0810	0.99434	1.000
Reading spreads ×1.2	0.3500	1.2885	0.93719	8.678
Reading spreads ×1.5	0.4350	1.6017	0.69573	55.976
Student-t on each reading, 30 degrees	0.2937	1.1678	0.9864	2.36
Student-t on each reading, 10 degrees	0.2936	1.4050	0.9282	7.17
Student-t on each reading, 5 degrees	0.2935	2.0137	0.2310	18.32
Student-t on the total, 30 degrees	0.2936	1.1860	0.9841	2.764
Student-t on the total, 10 degrees	0.2936	1.4674	0.8981	9.023
Student-t on the total, 5 degrees	0.2936	2.1345	0.1155	21.986

19.4.4 The answers against the stated truth

Table 19.15 collects the gaps with the stated truth of the constructed record, each row a quantity, the columns the reported value, what stands beside it and the gap. Rack B4's inlet at window 73 was 22.9046°C in truth, against 22.7126°C on the one fit. The coil's derate alone had raised it by +0.5888 K at window 73 and +1.7497 K at window 144, so the record's own night holds no 2 K excursion and no catch can be read from it. The dotted lines of Figure 19.7 are that derate through the map, +1.196 and +0.301 K on the discharge's and B4's residuals at window 73, and +3.561 and +0.902 K at window 144: the map's refound entries take about half of the inlet's rise, just as they do on the coil's route. Through the map at the level of 1.0810 K, those moves at B4 give rates of 28.4 a month at window 73 and 2302.9 a month at window 144. Those are the truth, and no log entry uses them.

19.5 The report (step 6)

A residual against the 19:00–21:00 reference meets both requests on [1.0810, 1.4482] K for a rise at the rack alone, for a plant still fitting the reference, which Chapter 16 finds this one no longer does; on the one fit's own g_{B2} the rate at 1.0810 K would be 53.7 a month now. That residual is blind to two of the four ways the rise could come, the load and the supply. A fixed level on the plain reading is empty on one window's fit, and pulling the alarm to 25°C catches 0.33211 now; on a reference it depends on the rate and catch conventions and on how much of

Table 19.15: The checks: the reported value against draws through the plant, second routes and the stated truth of the constructed record.

Quantity	Reported	Beside it	Gap
Residual at B4, spread, draws [K]	0.2936	0.2940	−0.0004
Tail at the level, 4 000 000 draws (se 5.2×10^{-6}) [–]	1.157×10^{-4}	1.102×10^{-4}	$+5.5 \times 10^{-6}$
Tail at three spreads, draws (se 8.4×10^{-5}) [–]	1.350×10^{-3}	1.415×10^{-3}	-6.5×10^{-5}
Coil route's catch, linear against refound at its state [–]	0.41709	0.42553	−0.00844
Local route, past the level, closed form on the noise spread against 300 windows (se 0.005) [–]	0.9926	0.993	−0.0004
Coil route, past the level, same (se 0.028) [–]	0.3903	0.373	+0.0173
Load route, past the level, same [–]	0.000	0.000	0.000
Supply route, past the level, same [–]	0.000	0.000	0.000
Band of T_{B4}, Gauss–Newton against Newton [K]	0.282107	0.282105	+0.000002
Stated truth			
T_{B4} at window 73 [°C]	22.7126	22.9046	−0.1920
Coil's rise at B4, largest of the night [K]	no excursion stated	+1.7497	–
This reference's no-event offset at B4 [K]	0	−0.0281	+0.0281
Rate at 1.0810 K given that offset [per month]	1.0000	0.484	+0.516
Rate at 1.0810 K given the coil's move at window 73 [per month]	53.7 (one fit)	28.4	+25.3

the load's ramp is covered. Rack B3's local rise cannot alarm above the false-alarm rate. "One a month" is a mean over a factor of 28.6; two in a row widen that to 99.1 at the same mean. And "caught" here means in one window, a duration the request does not state.

19.5.1 The answers and their bands

Table 19.16 gives the answers in the order of Table 19.1. Each row is a question; the second column gives the reported answer and the third what stands beside it.

19.5.2 The entry in the shift log

The operator's entry as of window 73, 01:00, using nothing later, reads: "Alarm settings, inlet B4. Plain reading: on this window's fit no fixed level gives one false a month and 2 K nine in ten in one window; 24.86°C catches 0.46 now, 25°C catches 0.33. On the 19:00–21:00 reference it depends on how the rate and catch are averaged and over which hours; 23.95°C averages one a month to 07:00 and catches 0.95 now. Residual against 19:00–21:00: 1.08 to 1.45 K over the fitted value meets both, for a rise at the rack alone (0.994 at 1.08). It will not see a load or supply rise; the T_s and Q_{B} readings do. A coil rise it catches 0.42; all channels together catch it. Rate at 1.08: one a month on average, 0.12 to 3.5 given the reference's error. Rack B3 has no sensor: a rise at the rack alone cannot alarm. So far tonight against these levels: $T_{d,\mathrm{B2}}$ first past at 23:20, 7 of 21 windows; inlet B2 once at 00:55; all channels together past 40.5 at 00:45 and 00:55, 38.9 now. The 19:00–21:00 model no longer fits, so these counts are not rates, and none of the rates above holds until it does. Setting: not decided here."

Table 19.16: The report at window 73, 01:00. Later windows enter the last row only to test the levels.

Question	Answer, reported	Beside it
Where should the alarm sit?	on the residual at B4, local route, anywhere in [1.0810, 1.4482] K; catch 0.99434 at the lower end, rate 0.0035 a month at the upper; for one window, and for a plant still fitting the reference, which this one no longer does. No level chosen here	tail at the level by draws 1.102×10^{-4} (se 5.2×10^{-6}); smallest rise with both held 1.5974 K
Can a fixed level on the plain reading do it?	on one window's fit no: spreads over 0.4029 K; at 24.860°C catch 0.45849 now, rate 7.46×10^{-8} to 6.132 a month. On the reference it depends on the conventions and span: widths −0.044, +0.103, +0.262 K over windows 73–144	23.945°C on the reference: catch 0.945 now, rate 0.138 to 3.400; the screen's 27°C catches 3.21×10^{-5}; 25°C catches 0.33211 now
Which excursions does it catch?	the residual: local 0.99434, coil 0.41709, load 0.00012, supply 0.00017	all channels catches the coil at 1.00000; the T_s and Q_{B} readings catch their own routes at 1.00000
And the rack with no sensor?	a local rise at B3 is caught at the false-alarm rate of the level set, 1.16×10^{-4} at L_α	the other routes move its estimate by +1.9841 to +2.0020 K
Is "one a month" a number?	a mean over 0.1235 to 3.5365 a month, a factor of 28.6	two in a row at the same mean: 0.0472 to 4.6770; windows 1–36 reference: 0.197 to 2.937
What would tonight have done?	through window 73: $T_{d,\mathrm{B2}}$ first at 23:20, past on 7 of 21 windows; T_{B2} once at 00:55	night's total 150 channel-windows against 0.444 expected; first fire of $T_{d,\mathrm{B2}}$ median window 61.0, range 49.0–69.0, over 300 fresh nights

19.6 The next reading (step 7)

A further reading changes the map, and what it does to the alarm is known before its value is: a new spread, a new shift, a new level and a new catch.

19.6.1 The candidates

Table 19.17 gives what each candidate does: each row a candidate, the columns the spread, the level for the rate, the local shift and the local catch at the rack it serves, with the rate's band or the smallest rise where they move. The second sensor at B4 and the permanent sensor at B3 re-solve the map (19.3) with the new row. The B3 sensor also needs a reference over it, taken here as 24 windows.

The second sensor at B4 is worth a caution. Its spread falls to 0.2045 K, and the catch at a 2 K rise reaches 1.00000, but the shift falls with it, 1.8245 to 1.6897 K, since a better inlet reading gives the refound load and supply more to follow. The smallest rise with both held comes to 1.2013 K, drawn as the dotted curve of Figure 19.3. The longer references narrow the rate's band and barely move the shift; the windows 1 to 73 row carries the warning of Section 19.2.4.

The permanent sensor at B3 is the largest change on the table. A local rise there goes from a catch of 1.16×10^{-4} to 0.99395, on an interval of $[1.0775,\ 1.4368]$ K, and β_{B3}'s band goes from 0.0800 to 0.0069. It is the instrument at the place of the question, and of the candidates only it reaches an empty column. Chapter 18 ranks the same reading against other worries. No candidate here makes one monitor catch every route; the load and supply routes stay with their own readings.

Table 19.17: What each candidate does to the alarm, for a 2 K local rise at the rack named.

Candidate	sd [K]	Level [K]	Shift [K]	Catch [–]	Beside it
None, residual at B4	0.2936	1.0810	1.8245	0.99434	smallest rise 1.5974 K; rate 0.1235–3.5365
Second inlet sensor at B4, mean of two	0.2045	0.7529	1.6897	1.00000	smallest rise 1.2013 K
Longer reference, windows 1–36	0.2910	1.0714	1.8219	0.99505	rate 0.197–2.937
Longer reference, windows 1–73	0.2882	1.0612	1.8180	0.99568	rate 0.347–2.220
None, at B3	–	1.0810	0	1.16×10^{-4}	at another level, that level's rate
Permanent inlet sensor at B3	0.2926	1.0775	1.8118	0.99395	β_{B3} band 0.0069

19.6.2 What the analysis showed

The screen had an alarm at the limit and readings well below it, and the operator's plan was to move the alarm closer. The analysis finds that where the alarm sits matters less than what it watches and what stands behind its baseline. On one window's fit a plain reading has no room, because that fit does not know the coils and fractions well; on two hours of reference it has room or not according to how a month's rate and a catch are averaged and how much of the night's ramp is covered, and those conventions are the review's to state. A residual against the reference opens an interval at the worst window, and it does so by refinding the very entries that a load or a supply excursion moves, so that alarm is deaf to two of the four ways a rise could come. A board of monitors must share the rate. At the rack with no sensor, a local rise cannot be told from nothing. A rate is not one number but a mean over a wide range, set by the baseline's own error, and a delay on the alarm trades some of the local catch for the coil's while

widening that range. “Caught” itself assumes a duration. And the night’s count says something else again: a rate is a statement about a plant that still fits its reference, and this plant stopped fitting it 1 h 20 min before 01:00 by the hour block and 15 min before by the single window.

19.6.3 Limits of the model

The routes. They are stated, not observed; the record holds no 2 K excursion. A rise by another route has another shift, and sizes are from the reference’s state at window 73.

One reference, one month. One reference is extended for a month, and the shared error stays constant over it. The month’s 8640 windows are treated as repeats of window 73’s map or of the night’s span, with no change of load, season or set point, and they include daytime hours this record does not model. Which tail and which catch a fixed level answers to is a convention, Table 19.6, that the request does not state.

The tail. Gaussian errors are assumed at a depth no record has tested; the sensitivity of Section 19.4.3 is the size of what that assumption carries.

The plain forecast. A straight load forecast, with the coil and fractions held at the stated baseline, lies under the plain level.

Window by window. The alarm is evaluated per window with no deadband or latching, so the standing and activation counts are not those of a real annunciator, and a catch is of one whole window at steady state.

Constructed readings. The record is constructed: generated from a stated truth of the plant with the stated sensor errors, and Section 19.4.4 compares the answers with that truth. It is not field data.

19.7 What to retain

Definition 19.1 (A calibrated decision level)**.** For a monitored quantity with no-event spread sd and shift s_r under route r, the level for a false-alarm tail α is $L_\alpha = z_\alpha\,\mathrm{sd}$, the level for a catch π is $L_\pi = s_r - z_\pi\,\mathrm{sd}$, and the request is met on $[L_\alpha, L_\pi]$, empty when $s_r < (z_\alpha + z_\pi)\,\mathrm{sd}$. The spread has a part that is new in each window and a part shared across them, equation (19.5); the second makes the rate uncertain, equation (19.10).

Three conclusions carry beyond the plant. A request for a rate and a catch can be tested before any fit by the arithmetic of equation (19.1), against the spread of the quantity watched given what is known behind its baseline, and a fixed level needs the request’s conventions on rate, catch, span and duration stated. A residual trades load for sight: it is free of what it refinds, and it sees almost nothing of what it refinds, which equation (19.4) says column by column. And a rate read from a baseline is a mean over that baseline’s error, conditional on the plant still fitting it, so its band and that condition belong in the report with the level.

Table 19.18 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 19.18: The terms of alarm calibration and where each is defined. Units for the shift plant.

Term	Expression	Units	Meaning	Defined in
Line of the request	eq. (19.1)	K	the largest spread a 2 K shift allows	19.1.3
Residual map	$\mathbf{K}$, eq. (19.3)	–	readings to residuals, entries refound	19.2.2
What it removes	eq. (19.4)	–	refound columns, up to their priors	19.2.2
Noise, shared part	eq. (19.5)	K	new each window; the same each window	19.2.2
Route, shift	$s_r = [\mathbf{Kd}]_k$	K	a rise's size on the residual	19.2.3
Questions and answers				
Levels, catch	eq. (19.8)	K, –	the interval that meets both	19.3.1
Plain level	eqs. (19.6)–(19.9)	°C	the fixed level, its baseline and conventions	19.3.2
Monitors against routes	Table 19.9	–	which monitor catches which route	19.3.3
Sharing the rate	Table 19.10	–	the cost of a set of channels	19.3.3
Rate over the shared error	eq. (19.10)	per month	the band of a rate; two in a row	19.3.5
Standing alarm, activations	Table 19.12	–	what a board does off its reference	19.3.6

19.8 Exercises

Exercise 19.1. *Formulation.* Show that one false alarm in 8640 windows and a catch of 0.90 on a 2 K shift limit the spread of the monitored quantity, and find the limit. (The tail 1/8640 has normal point 3.6819 and the catch 1.2816; a level for both needs $3.6819\,\text{sd} \le 2 - 1.2816\,\text{sd}$, so by equation (19.1) $\text{sd} \le 2/4.9635 = 0.4029\,\text{K}$.)

Exercise 19.2. *Application.* From the noise and shared spreads at B4 in Table 19.3, find the level for the rate. (By equation (19.5), $\text{sd} = \sqrt{0.2871^2 + 0.0613^2} = 0.2936\,\text{K}$; by equation (19.8), $L_\alpha = 0.2936 \times 3.6819 = 1.0810\,\text{K}$.)

Exercise 19.3. *Application.* Find the level for the catch of the local route at B4, and say why the shift is not 2 K. (The refound load and supply take 0.1755 K, so $s_r = 1.8245\,\text{K}$; by equation (19.8), $L_\pi = 1.8245 - 0.2936 \times 1.2816 = 1.4482\,\text{K}$.)

Exercise 19.4. *Derivation.* Scale the local shift to a 1.5 K rise and find the catch at the level for the rate. (The shift is $1.8245 \times 1.5/2 = 1.3684\,\text{K}$; by equation (19.8), $\pi = \bar{\Phi}((1.0810 - 1.3684)/0.2936) = 0.8362$.)

Exercise 19.5. *Application.* Using the rate's two ends at L_α, find the factor between them, and say why a longer reference narrows it. (A longer reference has a smaller shared spread, as Table 19.11 gives, and the band of equation (19.10) comes from that spread alone; the ends 0.1235 and 3.5365 a month have a ratio of 28.6.)

Chapter 20

Alarm Credibility and Base Rates

20.1 An alarm that has just fired

This chapter is a hindcast. That night the alarm on the screen stayed at the 27°C limit and the alarm review chose no level. Here the rate's end of the interval of Chapter 19 is applied to the record from window 25, as if the alarm had stood there, and the shift-log entries are written as what the log would have read. It is 00:55, window 72 of the shift. The inlet alarm on rack B2 has just fired on that level: the sensor reads 23.43°C, 3.57 K under the 27°C limit. It is the first inlet alarm the level gives that night. The review's own numbers are on the wall: one false alarm a month, nine real excursions in ten caught, excursions twice a month. The operator reasons that an alarm nine-in-ten reliable is probably real, and is about to send a technician to the rack.

The nine in ten is the probability of the alarm given an excursion, and the operator wants the reverse. Chapter 3 set out the difference on a small example, and Chapter 19 produced the two tails it needs. The question the wall cannot answer is

this alarm just fired; what is the probability it is real.

Jargon: base rate, probability of detection (the catch), positive predictive value, likelihood ratio, tail ratio, density ratio, Bayes factor, first alarm

The *base rate* is how often the thing an alarm is meant for occurs, regardless of any alarm. The *probability of detection* (the catch) is the probability of the alarm given an excursion. The probability of an excursion given the alarm is the *positive predictive value* (the alarm's credibility). A *likelihood ratio* is the probability of an observation under one hypothesis divided by its probability under another. Here it takes two forms: the *tail ratio*, for the observation "the residual passed the level", and the *density ratio*, for the observation "the residual took this value". A *Bayes factor* is the ratio of two evidences, equation (2.14), of which both of these are the simplest cases. A *first alarm* is the first fired window of a run, as against any window that fires.

Key idea

The credibility of an alarm is a ratio of counts: of the alarms in a long record, how many came with the thing the alarm names. Which count applies depends on what is conditioned on, on what the hypothesis says the excursion is, and on what the plant is taken to be doing when no excursion is on. A credibility without those three is not yet an answer.

20.1.1 A tail and a value

An alarm is a statement about a residual: it passed a level. Chapter 3 gives the probability of a hypothesis given that statement by Bayes' theorem (3.7), in the odds form (3.10): posterior odds are the likelihood ratio times the prior odds. The two tails of a calibrated level are the two likelihoods.

A control room holds more than that statement. The screen shows the residual itself, and an alarm a hair past its level is a different observation from one far past it. The same odds form then takes the density of that value in place of the tail. The value is the finer event, on the same model: taking this value lies inside passing the level, and the tail's answer is the value's answer averaged over the alarms past the level. A finer event also leans harder on what the model assumes. The chapter states both, and step 4 shows what each one rests on.

> **Principle 20.1.** *State what is conditioned on before stating a credibility. "Given the alarm" and "given the alarm's value" are two conditionings of one residual, and a number that does not say which is not yet an answer.*

20.1.2 The shift plant and the alarm (step 0)

The plant is the two-hall shift plant of Chapter 12 on the reduced state of Chapter 4, the 19 entries (12.3), indexed by window, $\mathbf{z}_t$. The dependent block of the physical rows is triangular with a nonzero diagonal, so it is invertible, and the temperatures follow by the chain (12.4); the sixteen readings' prediction is $\mathbf{m}(\mathbf{z}_t)$. The window read here is window 72, 00:55.

The fit read here. The chapter reads the reference fit of Section 16.2.2, windows 1 to 24, and not the one fit of Part IV at this window. An alarm's residual is against a stated earlier plant, for the reason Chapter 19 gives: a window's own fit would follow the excursion. This is the same reference the levels of that chapter rest on.

Modeling choices at the graph. The choices of the plant hold in the window as in Chapter 12: steady state, one effective conductance per coil, pinned airflow and no active control.

The graph. Figure 20.1 draws the comparison. The readings stand at the top, dashed to the entries they see. The residual map below them refinds the supply and loads per window, and its output is the residual the alarm watches. The entries are grouped along the middle, with the fraction β_{B2} in red. Below stand the hypotheses as scenario nodes: the plant with no excursion, and each route of Chapter 19 with the entry it pins moved from its value by the amount that raises the inlet of rack B2 by 2 K. The ramp and the sensor fault stand apart, because the base rate names neither.

20.1.3 The operator's questions (step 1)

Six questions follow from the alarm, and step 1 writes each over the state. "Real" means what the base rate counts, a 2 K inlet excursion, since nothing else here has a stated rate. Table 20.1 lists them: the question in the operator's words, the first of them the drill question; the general form; its mathematical form on this plant, with π the base rate per window, c and α the two tails, s the excursion's shift on the residual and sd its spread; the category; the computation; and the subsection where it is answered.

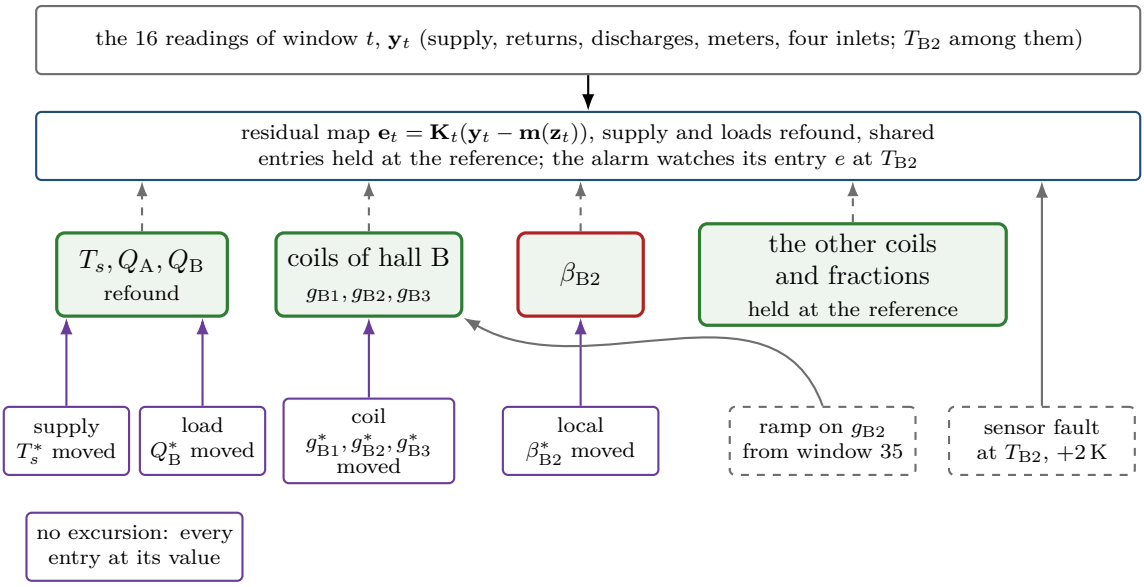


Figure 20.1: The comparison at window 72 on the reduced state. Free (green): the entries, grouped, with β_{B2} in red. Observed (dashed): the readings through the residual map. Purple: the plant with no excursion and the scenario nodes of the four routes of a 2 K rise at the inlet, each pinning the entries it moves, with the sizes of Section 20.2.2. Gray: two causes the base rate does not name. The sensor fault moves a reading and no entry.

Table 20.1: The questions asked at window 72, 00:55; the last row reads later windows at their own times.

Operator question	General form	Mathematical form	Category	Computation	Answered in
This alarm just fired; what is the probability it is real? Pr(excursion | alarm) by Bayes over the two tails of Chapter 19 and the base rate of excursions.	(2.14)	$\pi c/(\pi c + (1-\pi)\alpha)$	hypothesis comparison	evidence	20.3.1
What does twice a month mean here?	(2.14)	π by duration; first alarm against any	hypothesis comparison	evidence	20.3.2
It fired at this value; does that matter?	(2.14)	density ratio at e; $\ln = (s/\mathrm{sd}^2)(e - s/2)$;	hypothesis comparison	evidence	20.3.3
What do the other fifteen readings say?	(2.14)	the size integrated ratio on all the residuals, by route; the plain readings	hypothesis comparison	evidence	20.3.4
If not an excursion, what?	(2.15)	ramp on g_{B2}, and the two tails with it in; a sensor at this window	diagnosis	evidence	20.3.5
What does the next window say?	(2.14)	the odds updated window by window	hypothesis comparison	evidence	20.3.6

20.1.4 The evidence as rows (step 2)

Step 2 attaches the sixteen readings of window 72 as rows, each Gaussian, unbiased and independent given the state, against the reference fit's shared entries with their spread; the supply and both loads are refound on those rows. The alarm is a function of those rows.

Which row acts on which entry. With $*$ for an entry that can be nonzero and the columns T_s, Q_{B}, g_{B2}, β_{B2}, a part of the pattern (16.1) reads

$$\begin{bmatrix} * & 0 & 0 & 0 \\ * & * & * & 0 \\ 0 & * & 0 & 0 \\ * & * & * & * \\ * & * & * & 0 \end{bmatrix} \begin{matrix} \text{supply} \\ \text{return and discharges, hall B} \\ \text{meter } Q_{\mathrm{B}}\text{, loop meter} \\ \text{inlet } T_{\mathrm{B2}} \\ \text{inlet } T_{\mathrm{B4}} \end{matrix}, \qquad \mathbf{g}_{T_{\mathrm{B2}}}^{\mathsf{T}} = (*,\ *,\ *,\ *). \tag{20.1}$$

The second line is the gradient of the inlet the alarm reads, and it is not a row. The column of β_{B2} holds one nonzero, in the inlet's own row. So a rise at the rack alone touches one reading, as a fault in that sensor would, and no row can tell them apart in one window. The column of g_{B2} reaches the return, the discharges and the inlets of hall B, so a coil leaves a pattern across rows.

The base rate per window. The rate of twice a month, over a month of 8640 windows, is $\pi = 2.3148 \times 10^{-4}$ per window, with prior odds 2.3154×10^{-4}; the false-alarm rate is 1.1574×10^{-4} per window. The rate is taken as a rate of onsets. This gives a prior for a window with no excursion yet under way. It fails for a window inside an excursion, which Section 20.3.2 tests.

The rate at this rack. The stated rate is taken at this alarm. This buys a single prior. It fails if the twice a month is shared among racks, a case the sensitivities carry.

The excursion's size, fixed at 2 K**.** The request says "real 2 K excursions", which can mean exactly 2 K, about 2 K, or at least 2 K. The answers take exactly 2 K first. It buys one alternative hypothesis with one shift; it fails if excursions come in other sizes, and Section 20.3.3 integrates over size.

What is conditioned on. Two events: that the residual passed the level, and that it took this value. The drill's Bayes uses the first and the operator knows the second. Using both shows how far the finer event moves the answer, at the cost of a single number.

Two hypotheses. The base rate names one thing, an excursion, so the comparison has two sides, with no excursion meaning the plant at its reference. This buys a prior for each. It fails if something else moves the residual, and the causes without a stated rate are taken separately, in Section 20.3.5, with no base rate given to them.

The local route. The answers below use the local route of Chapter 19, a rise at the rack alone, as the excursion the alarm is calibrated for. This buys the shift the level was set against. It fails for the other three routes, which Section 20.3.4 compares.

Independence across windows. The odds carried from window to window can treat each window's residual as independent. This buys a sum of terms. It fails because the windows share the reference's error, and the shared form stands beside each sum.

20.2 The posterior of the state (step 3)

Parts IV to VI read one fit at the window a chapter states, and an alarm's residual is read against the reference fit for the reason of Section 20.1.2. That fit is cited. The residual map and the ramp's slope are what this chapter derives from it.

20.2.1 The reference fit and the residual map

The reference fit is that of Section 16.2.2, with its shared block $\mathbf{\Sigma}_{\mathrm{s}}$ and shared entries $\hat{\mathbf{z}}_{\mathrm{s}}$. With every row normalized the information matrix is $\mathbf{J}^{\mathsf{T}}\mathbf{J}$; with unnormalized rows it is $\mathbf{J}^{\mathsf{T}}\mathbf{W}\mathbf{J}$.

At window t the entries $\mathbf{q}_t$ are refound by the solve (16.9) with the covariance (16.8). The residual of equation (16.10) is then, to first order, a linear map of the difference between the readings and any prediction at those entries, the map of Chapter 19,

$$\mathbf{K}_t = \mathbf{I} - \mathbf{A}_{\mathrm{q}}\mathbf{L}^{-1}\mathbf{A}_{\mathrm{q}}^{\mathsf{T}}\mathbf{V}_t^{-1}, \qquad \mathbf{V}_{e,t} = \mathbf{K}_t\mathbf{V}_t\mathbf{K}_t^{\mathsf{T}}, \tag{20.2}$$

and sd^2 is the entry of $\mathbf{V}_{e,t}$ at the inlet. With the refound entries integrated out, the readings have the precision $\mathbf{V}_t^{-1}\mathbf{K}_t$. A hypothesis that moves the state to a scenario node $\mathbf{z}^*$ moves the prediction and the residuals by

$$\boldsymbol{\delta} = \mathbf{m}(\mathbf{z}^*) - \mathbf{m}(\mathbf{z}_t), \qquad \mathbf{d} = \mathbf{K}_t\boldsymbol{\delta}, \qquad s = [\mathbf{d}]_{T_{\mathrm{B2}}}. \tag{20.3}$$

These objects carry the posterior of the state into every answer: sd through $\mathbf{\Sigma}_{\mathrm{s}}$ in $\mathbf{V}_t$, and s through $\hat{\mathbf{z}}_{\mathrm{s}}$ at which the scenario is solved.

20.2.2 The residual at window 72

Table 20.2 gives the sixteen residuals at window 72 in their own spreads, observed minus fitted. The reading 23.4252°C stands against a fitted value of 22.3282°C. The residual at the inlet is

$$e = +1.0970\,\mathrm{K}, \qquad \mathrm{sd} = 0.2946\,\mathrm{K}\ (\text{noise } 0.2879,\ \text{shared } 0.0626), \qquad e/\mathrm{sd} = 3.724.$$

The level for one false alarm a month at this window's spread is 1.0847 K; Chapter 19's table is at window 73, 1.0836 K. The residual is 0.0123 K past it.

Table 20.2: The residuals at window 72, in their spreads. Observed minus fitted, against the reference.

Reading	z [–]	Reading	z [–]	Reading	z [–]	Reading	z [–]
T_s	−1.10	$T_{d,\mathrm{A2}}$	+0.61	$T_{d,\mathrm{B2}}$	+4.19	Q_{loop}	−0.18
$T_{r,\mathrm{A}}$	−0.29	$T_{d,\mathrm{A3}}$	−1.18	$T_{d,\mathrm{B3}}$	−2.02	T_{A1}	−0.67
$T_{r,\mathrm{B}}$	+0.22	$T_{d,\mathrm{B1}}$	+0.28	Q_{A}	+1.51	T_{A3}	−1.61
$T_{d,\mathrm{A1}}$	−0.12			Q_{B}	−2.17	T_{B2}	+3.72
						T_{B4}	−0.39

Two readings stand out, the alarm's own inlet and the discharge of unit B2, +1.2521 K against that channel's own level of 1.1000 K. Their residuals correlate at −0.062, so they are close to separate evidence. The statistic of equation (16.10) at this window is $Q = 41.507$ on 13.

The scenario nodes. Equation (20.3) at the four routes pins β_{B2} up by 0.1887 (local); g_{B1}, g_{B2}, g_{B3} down by 0.1234 together (coil); Q_{B} up by 76.7707 kW (load); or T_s up by 2 K (supply). Each raises the inlet by 2.000 K at the plant. On the residual the local route's s is +1.8342 K, or 6.226 spreads: the refound entries take the remaining 0.1658 K. Both tails and both ratios follow from s, sd and the level.

20.2.3 The ramp on the coil

Section 20.3.5 needs a fit that no tail can give: the night so far under a ramp in g_{B2}. Its reasoning is that of Chapter 15, whose fit selects an onset (Section 15.2.2) on a record through window 73. The onset here is selected on windows through 72 and stays at $\tau = 35$: the log evidence below is 76.704 there, against 76.667 at 34 and 76.626 at 36. The slope γ is a single unknown, a number. Each window from 25 to 72 contributes $\mathbf{e}_t$ and the column $\mathbf{h}_t = (\partial \mathbf{m}/\partial g_{\mathrm{B2}}) \max(t - \tau, 0)$. Stack both over the windows as $\mathbf{e}$ and $\mathbf{h}$, and let $\mathbf{C}$ be the stacked covariance, the noise of each window with its $\mathbf{q}_t$ integrated out, plus the reference's error $\mathbf{\Sigma}_{\mathrm{s}}$, which is one error shared by every window. Then with a stated prior spread of $\sigma_\gamma = 0.005$ per window,

$$\mathcal{L}(\gamma) = \tfrac{1}{2}(\mathbf{e} - \gamma\mathbf{h})^{\mathsf{T}}\mathbf{C}^{-1}(\mathbf{e} - \gamma\mathbf{h}) + \frac{\gamma^2}{2\sigma_\gamma^2}. \tag{20.4}$$

The objective is quadratic, so one Newton step from zero lands on its minimum: with $S_{eh} = \mathbf{h}^{\mathsf{T}}\mathbf{C}^{-1}\mathbf{e}$ and $S_{hh} = \mathbf{h}^{\mathsf{T}}\mathbf{C}^{-1}\mathbf{h}$,

$$\Lambda = \mathcal{L}'' = S_{hh} + \frac{1}{\sigma_\gamma^2}, \quad \sigma^2 = \frac{1}{\Lambda}, \quad \hat{\gamma} = \frac{S_{eh}}{\Lambda}, \quad \ln B = \frac{S_{eh}^2}{2\Lambda} - \tfrac{1}{2}\ln(1 + \sigma_\gamma^2 S_{hh}), \tag{20.5}$$

the last being the evidence of the ramp against no change for a Gaussian in one unknown. The record gives $S_{eh} = -52\,974.568$, $S_{hh} = 17\,555\,093.6$ and $\Lambda = 17\,595\,093.6$, so $\hat{\gamma} = -0.00301$ per window with $\sqrt{\sigma^2} = 0.00024$. At window 72 the lever is 37 windows, $g_{\mathrm{B2}} = -0.1114$.

20.3 The answers (step 4)

Each answer reads s, sd and the residuals of Section 20.2.2, or the slope of Section 20.2.3, and takes the prior odds of Section 20.1.4. No answer uses a window later than its own.

20.3.1 The two tails and the base rate

Operator question. This alarm just fired; what is the probability it is real? Pr(excursion | alarm) by Bayes over the two tails of Chapter 19 and the base rate of excursions.

General form. Two hypotheses are compared by the ratio of their evidences, equation (2.14), and the posterior odds are that ratio times the prior odds, equation (2.49).

Specialized form. The question is equation (2.14) with M_A the local excursion, M_B no change from the reference and the evidence the one event $e \geq L$. Each evidence is a tail of the normal with spread sd and, by equation (2.49),

$$\begin{gathered} c = \Pr(e \geq L \mid M_A) = \bar{\Phi}\Big(\frac{L - s}{\mathrm{sd}}\Big), \qquad \alpha = \Pr(e \geq L \mid M_B) = \bar{\Phi}\Big(\frac{L}{\mathrm{sd}}\Big), \\ \Pr(M_A \mid e \geq L) = \frac{\pi c}{\pi c + (1 - \pi)\alpha}. \end{gathered} \tag{20.6}$$

With π and α both small, the probability is close to $\pi c/(\pi c + \alpha)$: a ratio of two rates.

Building blocks. s and sd from equations (20.2) and (20.3); the base rate; the normal tail $\bar{\Phi}$.

The drill's corner. With the request's own numbers, catch 0.90 and one false a month, Pr = 0.6429. No level on this residual has both: the rate's level catches more, and the catch's level has far fewer false alarms, as the next paragraph shows.

Along the level. Table 20.3 runs equation (20.6) along the interval of Chapter 19, here $[1.0847, 1.4567]$ K: each row a level, the columns its false alarms a month, catch, tail ratio, posterior odds and probability. At the level this hindcast applies, the rate's end, the catch is 0.99453, the tail ratio 8592.7, the posterior odds 1.9895, and at no change from the reference

$$\Pr(\text{excursion} \mid e \geq L) = 0.6655.$$

It is a count. Figure 20.2 draws one month's alarms as boxes: 1.989 real and 1.000 false, so of three alarms two are real and one is false.

Table 20.3: The two tails and the base rate along the level, from equation (20.6), at window 72's spread. The first row is the rate's end of the interval, the last the catch's.

Level [K]	False a month [1/month]	Catch [–]	Tail ratio [–]	Posterior odds [–]	Pr [–]
1.0847	1.0000	0.99453	8592.7	1.9895	0.6655
1.1500	0.4092	0.98990	20 899.2	4.8389	0.8287
1.2000	0.2001	0.98433	42 496.6	9.8394	0.9077
1.2500	0.0952	0.97632	88 597.3	20.5134	0.9535
1.3000	0.0441	0.96511	189 233.8	43.8143	0.9777
1.3500	0.0198	0.94988	413 704.4	95.7871	0.9897
1.4000	0.0087	0.92975	924 763.7	214.1152	0.9954
1.4500	0.0037	0.90392	2 111 063.0	488.7851	0.9980
1.4567	0.0033	0.90000	2 360 429.7	546.5223	0.9982

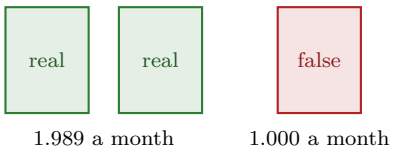


Figure 20.2: One month's alarms at the level 1.0847 K under the tails of equation (20.6). Three boxes, two real and one false, drawn to the nearest whole alarm.

The catch and the level. The count shows why, on the tails, the catch hardly matters: from 0.90 to 0.99453 the real count moves from 1.800 to 1.989, and the corner's 0.6429 becomes 0.6655. What moves the answer is the false count. Figure 20.3 draws both against the level, the horizontal axis the level across the interval and the vertical axis a probability: the catch falls from 0.99453 to 0.90000 while the probability rises from 0.6655 to 0.9982. Chapter 19 leaves the level inside its interval undecided; this curve is what that decision does to the answer here.

Over the rate's own uncertainty. Chapter 19 found that the rate at a level is not a number, because the reference's shared error stays for the month. At this inlet's noise 0.2879 K and shared spread 0.0626 K, that error's 2.5 % point, median and 97.5 % point give 0.1182, 0.7108 and 3.5977 false a month, and 0.9439, 0.7368 and 0.3561. So the reference's error alone carries the tails' answer across one half.

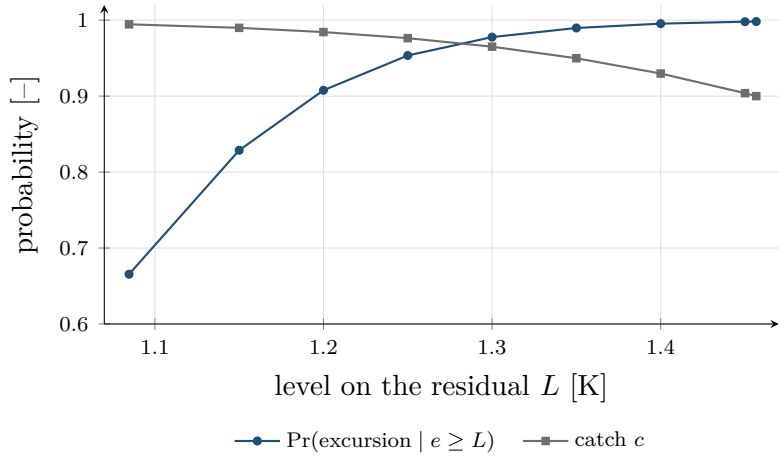


Figure 20.3: The tails' probability and the catch along the interval, the columns of Table 20.3.

20.3.2 The base rate's unit

Operator question. What does twice a month mean here? The rate counts onsets and says nothing of how long an excursion lasts.

General form. The comparison is again equation (2.14), with the odds of equation (2.49).

Specialized form. The question is equation (2.14) with an excursion of d windows as M_A, in two readings of the fired window. A window drawn from all the windows that fire has prior $\pi_d = d\,\pi$ and likelihood c. A first alarm has prior π and likelihood $1 - (1 - c)^d$, the chance that the run fires at all.

Building blocks. c and α of equation (20.6) at the set level; π.

The durations. Table 20.4 gives each duration a row. For a first alarm the probability is 0.6655, then 0.6667, and stays there: the catch has nothing left to give. For any fired window it climbs from 0.6655 through 0.9228 at half an hour to 0.9932 at six hours, because an excursion of d windows supplies d fired windows against a false alarm's one, and that count weights the late windows of a long run. Twice a month does not say which d applies. Tonight's alarm is a first alarm, so the first-alarm column reads it. That column moves only from 0.6655 to 0.6667, so at window 72 the answer does not depend on the unstated d.

Table 20.4: The base rate's unit. Two onsets a month, an excursion of d windows, at the set level.

d [windows]	Duration [h]	π_d [1/window]	Pr, first alarm [–]	Pr, any fired window [–]
1	0.08	2.3148×10^{-4}	0.6655	0.6655
3	0.25	6.9444×10^{-4}	0.6667	0.8566
6	0.50	1.3889×10^{-3}	0.6667	0.9228
12	1.00	2.7778×10^{-3}	0.6667	0.9599
24	2.00	5.5556×10^{-3}	0.6667	0.9796
72	6.00	1.6667×10^{-2}	0.6667	0.9932

20.3.3 The value it fired at

Operator question. It fired at this value; does that matter? The screen shows more than that the level was passed.

General form. The ratio of evidences is again equation (2.14), with the odds of equation (2.49), and the evidence is now the value.

Specialized form. The question is equation (2.14) with the same M_A and M_B and the evidence the event that the residual took the value e. Each evidence is a normal density, and their ratio is

$$B(e) = \frac{\phi\big((e-s)/\mathrm{sd}\big)}{\phi(e/\mathrm{sd})}, \qquad \ln B(e) = \frac{s}{\mathrm{sd}^2}\Big(e - \frac{s}{2}\Big). \tag{20.7}$$

It is linear in the value, and it crosses zero at the midpoint between the two hypotheses' means. The tails' answer is the average of this curve over the alarms: the probability at each e, weighted by how often an alarm fires there, integrates to the probability of equation (20.6) by the law of total probability. For a size unknown within a stated density $w(S)$ at the same rate, B and c are averaged over w, with the shift proportional to S.

Building blocks. e, s and sd of Section 20.2.2; the prior odds.

The alarm's value. At $e = +1.0970$, $e - s/2 = 0.1799$, so $\ln B = (1.8342/0.2946^2)(0.1799) = 3.8013$ and $B = 44.761$. The posterior odds are 0.01036 and, for an excursion of exactly 2 K at no change from the reference,

$$\Pr(\text{excursion} \mid e = +1.0970) = 0.01026.$$

Beside the tails' 0.6655 the likelihood ratio has fallen by a factor of 192.0, 8592.7 against 44.761, a factor that belongs to the point hypothesis. Figure 20.4 draws where it comes from: the horizontal axis is the residual, the vertical axis the density on a logarithmic scale, the left curve no excursion and the right curve the local route, with the level and both tails shaded and the fired value marked on each curve. The ratio of the two tails' areas is the tail ratio, though a logarithmic axis does not show areas to scale, and the ratio of the two heights at the fired value is the density ratio.

Why they differ. Past the level, a false alarm stands on average 0.0711 K over it and a caught excursion 0.7542 K over it. This one is 0.0123 K over. Of the false alarms 0.4932 fire within 0.05 K of the level, and of the caught excursions 0.00333. The tails ask about all the alarms past the level, and the second reading about the one seen.

Against the residual. Table 20.5 gives the second reading at other values, each row a residual, and Figure 20.5 draws the curve: the horizontal axis is the residual and the vertical axis the probability, with the fired value and the tails' value marked and the line starting at the level. The curve sits near zero at the level, reaches one half at 1.3132 K, inside Chapter 19's interval, and turns over quickly. Its height at the alarm is the 0.01026; the horizontal line is its average over alarms, the 0.6655.

The size. The 0.01026 is a property of the 2 K point. Table 20.6 integrates over size at the same rate, each row a stated density for the size, the columns the density ratio and both readings. A broad density around 2 K gives the value's reading 0.0459 to 0.0652 and the tails' 0.528 to 0.653: the gap is about ten times, not 64.9 times. Read as at least 2 K, a size that

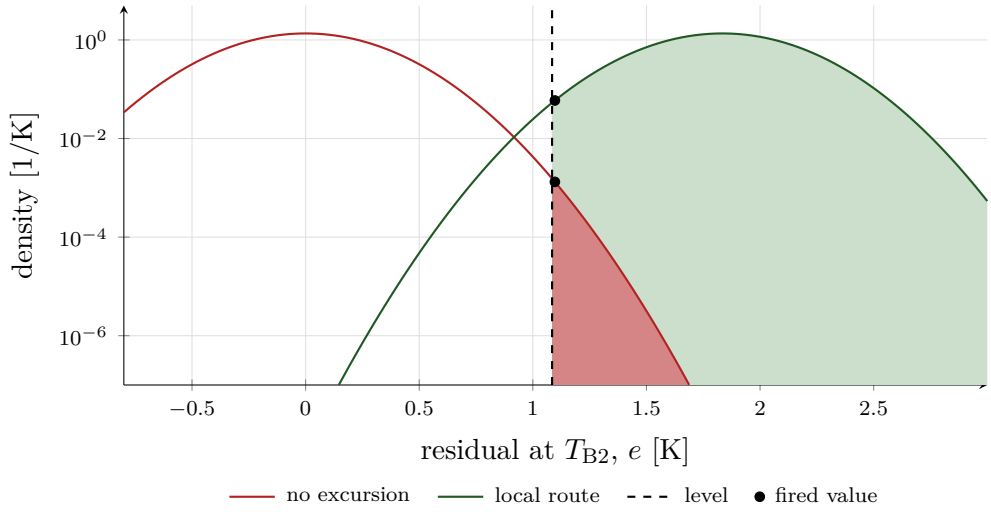


Figure 20.4: The two densities of the residual at window 72: normal, spread 0.2946 K, means 0 and 1.8342 K. Logarithmic axis. Shaded: the two tails past the level 1.0847 K. The dots on the two curves stand at the fired value, 1.0970 K.

Table 20.5: The second reading against the residual, from equation (20.7), for exactly 2 K.

Residual e [K]	Density ratio $B(e)$ [–]	Pr [–]
1.0847	34.517	0.00793
1.0970	44.761	0.01026
1.2000	395.080	0.08381
1.4000	27 068.635	0.86240
1.6000	1 854 588.589	0.99768

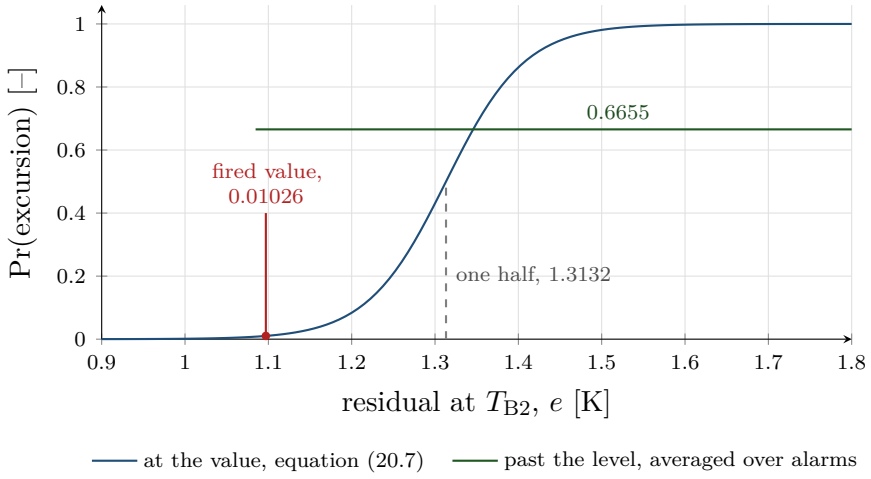


Figure 20.5: The probability of an excursion of exactly 2 K against the residual: prior odds 2.3154×10^{-4}, $s = 1.8342$ K, sd = 0.2946 K. The horizontal line starts at the level 1.0847 K.

cannot sit near this residual, the value's reading is 0.0011, 43 to 61 times lower than the broad densities give. So the words "real 2 K excursions" carry that factor on the value's reading, and the tails' reading barely feels them. The gap crosses any break-even probability between the two readings; both are stated and the chapter does not choose, because the finer event rests on the size, which no record states.

Table 20.6: The two readings with the excursion's size integrated, at twice a month.

Size [K]	Density ratio [–]	Value Pr [–]	Tails Pr [–]
Half-normal, scale 2	273.26	0.0595	0.528
Exponential, mean 2	230.02	0.0506	0.529
Lognormal, median 2, $\sigma = 0.25$	207.94	0.0459	0.653
Lognormal, median 2, $\sigma = 0.5$	288.95	0.0627	0.626
Uniform, 1 to 3	301.04	0.0652	0.638
Uniform, 0.5 to 4	232.33	0.0510	0.617
Normal, mean 2, sd 0.5	222.02	0.0489	0.647
Normal, mean 2, sd 0.2	91.14	0.0207	0.663
Exactly 2	44.761	0.01026	0.6655
At least 2: 2 plus exponential, mean 1	4.61	0.0011	0.667

Common mistake

Reporting one credibility for an alarm with no event named. 0.6655 is a statement about alarms at this level; 0.01026 is a statement about an alarm at this value, for exactly 2 K. The second is the finer event of the same model, and it is the one that moves most when the size is stated differently.

20.3.4 The other readings

Operator question. What do the other fifteen readings say? The alarm watches one row of sixteen.

General form. The ratio of evidences is again equation (2.14), on the whole residual.

Specialized form. The question is equation (2.14) with M_A an excursion by a stated route, M_B no excursion, and the evidence all sixteen readings with the entries $\mathbf{q}$ integrated out, of precision $\mathbf{V}^{-1}\mathbf{K}$ in Section 20.2.1. With $\boldsymbol{\delta}$ of equation (20.3),

$$\ln B = \boldsymbol{\delta}^{\mathsf{T}}\mathbf{V}^{-1}\mathbf{e} - \tfrac{1}{2}\,\boldsymbol{\delta}^{\mathsf{T}}\mathbf{V}^{-1}\mathbf{K}\boldsymbol{\delta}. \tag{20.8}$$

The first term is how far the residuals agree with the route, and the second is what the route costs if they do not follow it.

Building blocks. The residuals, $\mathbf{V}$ and $\mathbf{K}$ of Section 20.2.1; each route's $\boldsymbol{\delta}$; Q and its chi-square tail on 13.

The routes. With no excursion, $Q = 41.507$ has tail 0.0001 on 13, so something is off, and it need not be an excursion. Table 20.7 gives the routes, each row a route, the columns the two terms of equation (20.8), its log ratio and its shift on this inlet. The local route has +3.672. The other fifteen readings take −0.129 from B2's own 3.8013: nothing in them speaks for a rise at the rack.

Table 20.7: The routes of a 2 K rise on all sixteen readings at window 72, from equation (20.8).

Route	$\boldsymbol{\delta}^{\mathsf{T}}\mathbf{V}^{-1}\mathbf{e}$ [–]	$\boldsymbol{\delta}^{\mathsf{T}}\mathbf{V}^{-1}\mathbf{K}\boldsymbol{\delta}$ [–]	$\ln B$ [–]	Shift at T_{B2} [K]
Local	+23.048	38.751	+3.672	+1.834
Coil	+36.315	132.241	−29.806	+1.046
Load	+0.000	0.000	+0.000	−0.000
Supply	+0.372	15.677	−7.467	+0.029

The hall-wide coil route is rejected, and by the pattern across rows and not by one meter. It agrees with the residuals most, +36.315, and it costs far more than that. In spreads, with the observed value in parentheses, it asks for −7.33 (−2.17) at Q_{B}, −6.36 (−1.11) at the supply, +4.32 (+0.28) at the discharge of B1, +4.27 (+4.19) at that of B2 and +4.17 (−2.02) at that of B3. One discharge of three has done what it asks. Taking a reading out, with the entries refound and the route sized again, leaves it rejected whichever goes: −18.836 without Q_{B}, −13.340 without the discharge of B3, −16.924 without the supply reading, −22.006 without the discharge of B1 and −38.483 without that of B2.

What the residuals cannot speak for. The load's ratio of 0.000 and the supply's −7.467 are not verdicts of the readings. The refound entries take both routes: Q_{B} and T_s are refound at 307.898 kW and 14.0465°C, so the load leaves the residuals nothing, and the supply's cost is its own prior row on T_s. The plain readings do speak. T_s reads 13.969°C with a spread of 0.1 K, so a +2 K supply stands 20.0 spreads away. The hall-B meter reads 302.093 kW against the previous hour's mean of 299.499 kW, with a spread of 4 kW, so the load route's +76.7707 kW stands 19.2 spreads away. Neither route is on at this window, on evidence the residual had discarded.

On all readings. With every excursion local the Bayes factor is 39.3465 and Pr = 0.009028, for exactly 2 K. With the routes split equally it is 10.0868 and Pr = 0.002330. Both stand close to the second reading of Section 20.3.3, because this conditioning is on values and not on a tail.

20.3.5 Causes without a base rate

Operator question. If not an excursion, what? The base rate names one cause, and the record holds others.

General form. Candidates are ordered by evidence times prior, equation (2.15).

Specialized form. The question is equation (2.15) with two candidates against no change, the ramp on g_{B2} and a sensor fault at this inlet, and $\Pr(M_j)$ unstated for both. For the ramp, $p(\mathbf{y} \mid M_j)$ against no change is the $\ln B$ of equation (20.5). Its part in this window's residual is $r = [\mathbf{K}\mathbf{h}_{72}]_{T_{\mathrm{B2}}}\hat{\gamma}$. With the ramp on, the no-excursion mean is no longer zero, and equations (20.6) and (20.7) hold with r added to both means: $\alpha_r = \bar{\Phi}((L - r)/\mathrm{sd})$, $c_r = \bar{\Phi}((L - r - s)/\mathrm{sd})$ and $e - r$ in place of e. For the sensor fault, the shift is 2 K on the reading alone, through $\mathbf{K}$, which is the local route's by the single nonzero in pattern (20.1).

Building blocks. $\hat{\gamma}$ and Λ of Section 20.2.3; $\mathbf{K}\mathbf{h}_{72}$; the local route's $\ln B$ of Table 20.7.

The ramp. Equation (20.5) gives $\ln B = +76.704$, its parts being 79.747 and 3.043. Summed as if the windows were independent it would be +136.091, of which 43.6 % is the reference's one error counted as fresh in each window. Chapter 15's +89.73 at 01:00, on its own stacked fit,

stands beside it (Section 15.2.2). Between the ramp and no change alone, at any prior from 10^{-6} to 0.5 for such a ramp on a night, $\Pr(\text{ramp} \mid \text{record}) = 1.00000000$ to that many places. The evidence overwhelms any prior that sweep holds; it is evidence and not a base rate, and it compares two hypotheses only. The same section of Chapter 15 sets a discharge sensor's drift below the coil.

What it explains. The ramp predicts $+1.1532\,\mathrm{K}$ at $T_{d,\mathrm{B2}}$, against the residual $+1.2521\,\mathrm{K}$, leaving $+0.0989\,\mathrm{K}$, or 0.331 spreads. At this inlet it predicts only $r = +0.2984\,\mathrm{K}$ against $+1.0970\,\mathrm{K}$, leaving $+0.7986\,\mathrm{K}$, or 2.711 spreads, with a one-sided tail of 3.36×10^{-3}. So the ramp accounts for the discharge and for part of the alarm.

The answers with the coil in. One false alarm a month describes a plant at its reference, and the record says tonight's is not. With r in the means the false-alarm tail at this window is 3.803×10^{-3}, 32.9 a month at this window's rate, the catch 0.99981 and the tail ratio 262.9, and

$$\Pr(\text{excursion} \mid e \geq L,\ \text{coil in}) = 0.0574.$$

The value's reading, for exactly $2\,\mathrm{K}$, has $\ln B = -2.505$ and $\Pr = 1.891 \times 10^{-5}$. So the 0.6655 of Section 20.3.1 is an answer about a plant at its reference, and it is not tonight's: at 00:55 the windows through 72 already hold the coil.

The sensor. A fault of $+2\,\mathrm{K}$ on this reading at this window leaves the same residuals as the local route, so its log ratio is +3.672, and no reading of the sixteen parts them. "What else" has a clear answer for the coil and none for the sensor.

20.3.6 The next windows

Operator question. What does the next window say? An excursion that holds shows again at 01:00.

General form. The odds of equation (2.49) are updated by the ratio of evidences, equation (2.14), of each new reading.

Specialized form. The question is equation (2.14) with M_A a $2\,\mathrm{K}$ excursion that holds and M_B no excursion, the plant at its reference, and the evidence the inlet's residuals from window 72 on. With the windows independent the log odds add the terms $\ln B(e_t)$ of equation (20.7). With the reference's error shared, the ratio is that of two joint normals with the covariance of Section 20.2.3 restricted to this inlet. For the tails, the update of 0.6655 on no fire uses $1 - c$ and $1 - \alpha$.

Building blocks. Each window's e_t from equation (16.10); s and sd; c and α of equation (20.6).

The next readings. Table 20.8 gives the windows, each row a window, and Figure 20.6 draws them: above, the residual at the inlet with the level and with the mean a held excursion predicts; below, on a logarithmic axis, the probability the odds give. The dotted marker is a check, read in Section 20.4.3. The residual fires once and falls away. At 01:00 the residual is $+0.0515\,\mathrm{K}$, and a held $2\,\mathrm{K}$ excursion has $\Pr = 2.117 \times 10^{-10}$ with the reference shared, 1.176×10^{-10} with the windows independent. Independence overstates more with each window: the cumulative log ratio is -46.852 against -54.893 through window 76, -101.422 against -154.331 through window 84. From even odds, independent, window 73 alone leaves 5.080×10^{-7}.

Table 20.8: The odds carried forward from window 72 by the value at each window, with the windows independent and with the reference shared.

Window [–]	Clock	e_t [K]	$\ln B(e_t)$ [–]	Cumulative, independent [–]	Cumulative, shared [–]	Pr(2 K holds), shared [–]
72	00:55	+1.0970	+3.801	+3.801	+3.801	0.010257
73	01:00	+0.0515	−18.294	−14.493	−13.905	2.117×10^{-10}
74	01:05	+0.4306	−10.282	−24.775	−22.831	2.815×10^{-14}
76	01:15	+0.3809	−11.334	−54.893	−46.852	1.040×10^{-24}
84	01:55	+0.5474	−7.814	−154.331	−101.422	2.079×10^{-48}

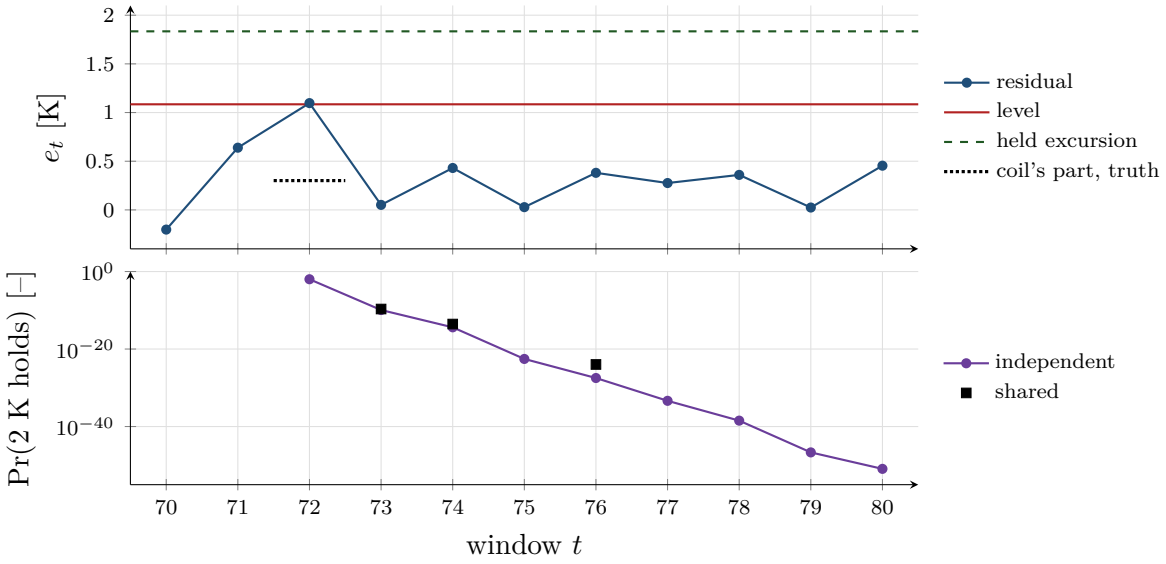


Figure 20.6: The residual at the inlet of rack B2 at windows 70 to 80 with the level of window 72, and below it the probability of a 2 K excursion that holds, from window 72 on. Windows after 72 are read at their own clock times. The dotted marker at window 72 is the constructed truth's coil through the map, a check and no part of any entry in the log.

If it holds. The update assumes an excursion lasts past one window, and at 01:00 that matters as it did not at 00:55. If a share of excursions last one window only, the reading at 01:00 cannot speak for those. At shares of 0.01, 0.1 and 0.5, the value's reading at 01:00 is 1.036×10^{-4}, 1.035×10^{-3} and 5.155×10^{-3}, and the tails' is 0.0298, 0.1727 and 0.5001. A one-window excursion at 00:55 has gone and is not tested.

By tails alone. If it holds, one fire and then silence at window 73, with the miss 0.00547 and the no-fire 0.99988, takes 0.6655 to 0.01077 (0.01068 with the reference shared). Two fires in a row, the shared part ignored, would have given 0.99994. So at 01:00 both readings stand low, still a factor of 9.16×10^{7} apart.

The two channels. Through window 73 the inlet has fired at window 72 alone. The discharge of unit B2 has been past its level since 23:20, at windows 53, 66, 67, 68, 70, 72 and 73, the channel the ramp of Section 20.3.5 accounts for.

20.4 The checks (step 5)

The answers rest on normal tails, a linear map, a fixed size, and a stated rate. The checks test the arithmetic by counts, the fit by its limit, and the answers against the constructed truth and against the rate and size they assume.

20.4.1 The limit of exact rows and the curvature term

The reference and its alarm run on the reduced state, the limit $\sigma_p \to 0$ of the full state with the physical rows (12.2). At this window, on the window's own fit, the inlet of rack B2 at tolerances of 0.100, 0.010 and 0.001 is 23.27143, 23.26952 and 23.26950°C, with spreads 0.28298, 0.28277 and 0.28276 K, in 4, 3 and 3 iterations, against the reduced state's 23.26950°C and 0.28276 K. Newton's covariance, with the curvature term kept, gives 0.282781 K against Gauss–Newton's 0.282765 K, a difference too small to move any tail here. Chapter 19 checks the same at window 73.

20.4.2 The tails by counts

Table 20.9 counts what the closed forms assert, each row a simulation. Over 20 000 simulated months false alarms come at 0.9860 a month, excursions at 2.0132 and caught ones at 2.0028, a real share of 0.6701 (se 0.0019), above 0.6655 by 2.4 standard errors, on a draw whose false alarms ran low and excursions high. Over 20 000 000 windows drawn directly, 6881 are past the level, with a real share of 0.6628 (se 0.0057). Of those within 0.05 K over the level, 1137, the real share is 0.0097 (se 0.0029), against 0.0132 for that band's closed form and 0.0134 for the density form at its middle. Averaging the value's curve over the alarms past the level returns 0.6655, as Section 20.3.3 said. These counts test the arithmetic at no change from the reference and exactly 2 K, not those two assumptions.

Table 20.9: The tails and the value by counts on the linear alarm.

Count	Size [–]	Real share [–]	Closed form [–]
Months, Poisson false and real	20 000 months	0.6701	0.6655
Windows past the level	6881 windows	0.6628	0.6655
Windows within 0.05 K over the level	1137 windows	0.0097	0.0132

The same months give the operator's sense of scale: 0.3720 of months have no false alarm, 0.1338 have no excursion and 0.0492 have no alarm at all. A given shift of 144 windows holds 0.0167 false alarms and 0.0333 excursions.

20.4.3 The answers against the stated truth

Table 20.10 collects the gaps, each row a quantity, the columns the reported value, what stands beside it and the gap. The record holds no 2 K excursion. At window 72 the true inlet is 22.5057°C, against 21.9311°C with the coil at its rating: the coil's part at the plant is +0.5746 K. The reading's own error is +0.9195 K, +3.065 spreads.

Through the residual map, the truth divides the +1.0970 K as +0.3010 K from the coil, +0.8433 K from the reading's own error, +0.0202 K from the other readings' errors and −0.0675 K from the fit against the truth. The coil's +0.3010 K is the like of the ramp's +0.2984 K, a gap of −0.0026 K. The plant's +0.5746 K is not its like, because the refound supply and load take 0.1274 and 0.1462 K of it. Without the coil's part the residual would have been +0.7960 K, under the level. So in the constructed record the alarm was the coil and an unlucky reading, and neither alone would have passed the level. No operator could have known the second at 00:55.

Over 60 fresh nights, the reference held, the first B2 alarm has median window 118.0 and 2.5–97.5 % range 82.4 to 142.0, and none fires by window 73. The chance of a fire by window 73, with the ramp in the mean and the windows taken as independent, is 0.0452 a night, and none in 60 then has probability 0.062, so tonight's 00:55 is early by a draw that the sixty nights could easily miss.

Table 20.10: The checks: the reported value against counts, a second route and the stated truth of the constructed record.

Quantity	Reported [–]	Beside it [–]	Gap [–]
Pr, tails, against months (se 0.0019)	0.6655	0.6701	−0.0046
Pr, tails, against direct draws (se 0.0057)	0.6655	0.6628	+0.0027
Pr, the 0.05 K band, against draws (se 0.0029)	0.0132	0.0097	+0.0036
Band of T_{B2}, window's fit, against Newton [K]	0.282765	0.282781	−0.000016
An assumption moved			
Pr, tails, at the reference against coil in	0.6655	0.0574	–
Pr, value, exactly 2 K, at the reference against coil in	0.01026	1.891×10^{-5}	–
Stated truth			
A 2 K excursion at window 72	0.6655; 0.01026	none	–
Coil's part on the residual, the ramp's against the truth's own [K]	+0.2984	+0.3010	−0.0026
First B2 alarm, this night against 60 fresh nights' median [window]	72	118.0 (range 82.4–142.0)	–

20.4.4 Sensitivity

Table 20.11 varies the stated quantities one at a time, each row a value, the columns both readings. The base rate moves the two together. If the twice a month is shared among ten racks, the 0.2 a month gives 0.1659 and 0.00104. The false-alarm rate is the tails' lever. The point size is the second reading's: at 1 or 1.5 K this value is more likely than at 2 K, because a smaller excursion predicts a residual closer to it, while at 3 K the value excludes it. These rows

keep the same rate for an excursion of each size, itself a choice. The size that fits this value best is 1.196 K, with $B = 1025.398$ and $\Pr = 0.19186$: at twice a month no single size lifts this value's reading past that.

Table 20.11: The two readings against the stated rate, the size and the false-alarm rate, at the reference.

Varied	Tails [–]	Value [–]	Varied	Tails [–]	Value [–]
Excursions 0.2 a month	0.1659	0.00104	Size 1.0 K	0.3629	0.16461
Excursions 0.5 a month	0.3321	0.00258	Size 1.5 K	0.6265	0.13177
Excursions 1 a month	0.4987	0.00515	Size 2.0 K	0.6655	0.01026
Excursions 2 a month	0.6655	0.01026	Size 3.0 K	0.6667	0.00000
Excursions 4 a month	0.7992	0.02031	False 0.1 a month	0.9521	–
Excursions 8 a month	0.8884	0.03983	False 1 a month	0.6655	–
Excursions 30 a month	0.9677	0.13492	False 10 a month	0.1659	–
			False 100 a month	0.0195	–

20.5 The report (step 6)

On a plant at its reference, an alarm with a catch of 0.99453 and one false alarm a month is two-thirds right, 0.6655, and on those tails its catch hardly matters while its level does. Tonight's plant is not at its reference, and with the coil in the same tails give 0.0574. The value it fired at, 0.0123 K over, gives 0.01026 for exactly 2 K and 0.0459 to 0.0652 for a size about 2 K. One more window takes a held excursion far lower.

20.5.1 The answers and their bands

Table 20.12 gives the answers in the order of Table 20.1. Each row is a question; the second column gives the reported answer and the third what stands beside it.

20.5.2 The entries in the shift log

These are the entries the log would have read with the hindcast's level in place. Each uses nothing later than its window. At 00:55: "Inlet B2 past its level, first of the night. As one alarm of a rate-set level, twice-monthly rate assumed: 0.67 that a 2 K excursion is on for a plant at its reference, 0.06 with B2's coil ramp in, which the record through this window supports. At this value, 0.012 K over: 0.01 for exactly 2 K, about 0.05 for a size near it. Those are different questions. $T_{d,\mathrm{B2}}$ past its level since 23:20; the ramp accounts for that channel and leaves 0.80 K of this alarm. Action: not decided here."

At 01:00: "Next reading +0.05. A held 2 K excursion: below 10^{-9}. An excursion over in one window not tested. The sensor at 00:55 unexplained."

20.6 The next reading (step 7)

A further row changes the odds by its own ratio, and what it can do is known before its value, by the separation between the two hypotheses in its own spread. This is the preposterior of Chapter 18 for two hypotheses.

Table 20.12: The report at window 72, 00:55; the last row is read at 01:00.

Question	Answer, reported	Beside it
This alarm just fired; what is the probability it is real?	as one alarm past a rate-set level, on a plant at its reference, 0.6655 that a 2 K excursion is on	the rate's band, 0.3561 to 0.9439; draws 0.6628 (se 0.0057); the request's corner 0.6429; coil in, 0.0574
What does twice a month mean here?	onsets; for a first alarm 0.6667 at any d of three windows or more	for any fired window 0.9228 at half an hour, 0.9932 at six hours
It fired at this value; does that matter?	yes: at the value, 0.01026 for exactly 2 K; $B = 44.761$ against a tail ratio of 8592.7; both are stated	about 2 K, 0.0459 to 0.0652; at least 2 K, 0.0011; draws in the band 0.0097 (se 0.0029)
What do the other fifteen readings say?	nothing for a rise at the rack: Pr = 0.009028 all local, 0.002330 routes equal; the hall-wide coil route at −29.806; the plain readings exclude load and supply	Q tail on 13: 0.0001, something is off
If not an excursion, what?	a ramp on g_{B2}, $\ln B = +76.704$, accounting for +0.2984 of +1.0970 K; 0.7986 K left. A sensor fault: not parted	Chapter 15's +89.73; each window alone +136.091
What does the next window say?	at 01:00, $e = +0.0515$ K: a held excursion 2.117×10^{-10}	by tails alone 0.01077; with a share 0.1 of one-window excursions, 1.035×10^{-3} and 0.1727

20.6.1 The candidates

Table 20.13 gives the candidates at 00:55, each row a candidate, the columns its separation from a held excursion, its separation from a sensor fault that stays, and what it needs. The handheld at the inlet of rack B2, with a 0.3 K spread, stands 5.99 spreads between the hypotheses on the residual's shift; on the 2 K a probe would see at the rack, and with the refound entries' variance, the same separation runs from 5.74 to 6.53. From the tails' 0.66550 it ends below 0.01 on 0.3281 of outcomes and above 0.99 on 0.6597. From the value's 0.01026 it ends below 0.01 on 0.9882 and above 0.99 on 0.0095. The mean of the posterior holds in each, 0.66608 (se 0.00074) and 0.01028 (se 0.00016). A reading at the rack's fitted value would give a ratio of 1.666×10^{-8}, and at the fitted value plus 2 K a ratio of 1.529×10^{9}.

Table 20.13: Candidate readings at 00:55: separation of a held 2 K excursion from none, and of an excursion from a sensor fault that stays.

Candidate	From none [spreads]	From the fault [spreads]	Needs
Sensor's own next window	6.23	0.00	nothing; five minutes
Handheld at the inlet, 0.3 K	5.99	5.99	one trip
Permanent second inlet sensor, mean of two	8.97	4.71 on their difference	an installation; sd 0.2044 K, level 0.7526 K

The sensor's own next window and the handheld separate an excursion from none about equally, and the plant takes the first by itself. What the next window cannot do is part an excursion from a sensor that stays wrong, since both read high again; the handheld can, and so can a second sensor at the inlet, on the difference of the two. A permanent second sensor also lowers the level on the mean of the two to 0.7526 K, with catch 1.00000. For an excursion that

has already ended, none of these reaches back.

20.6.2 What the analysis showed

The operator read nine in ten as the chance that the alarm was real. On a plant at its reference the analysis finds that the nine in ten hardly enters: as one alarm past a level, the answer is a count of two real to one false a month, and it is the false count, set by where the level sits, that moves it. But the 19:00–21:00 model no longer fit the record by then. With the coil in, the false count at this window is many times higher, and the same tails give about one in twenty. Taken at the value it fired at, a hair past the level where false alarms gather, the alarm stands near one in a hundred at exactly 2 K and near one in twenty for a size about 2 K, a spread that comes from words in the request and no reading. The log therefore carries each reading with what it assumes. The coil needed no base rate, because the record's evidence for it swamps any prior in the sweep, and it accounts for a discharge channel and part of this alarm. The rest could be an unlucky reading or a failed sensor, which that window's readings cannot part. The next window bears hard on an excursion that holds, and on none that has ended.

20.6.3 Limits of the model

The base rate. It is stated, not read from any record, and it is taken at this rack. Its duration is unstated.

One route. The two readings take the local route. Section 20.3.3 shows what size does; a mixture over routes is reported only at equal weights.

The coil in the means. The tails with the coil in use the ramp's mean part at this window alone, and a month's rate at that part is a statement about this window and not about a month.

Independent windows. The tails' update and the carried odds of the first column treat the shared part as fresh; the shared column stands beside them.

Gaussian tails. Both reported tails and the density ratio are normal; Chapter 19 shows how the rate moves under heavier tails. A density ratio this far out rests on the same untested part of the distribution.

First alarm only. The reasoning is for the first alarm of a run. An alarm that stands is a different event.

No consequence. No cost of a wrong response enters. A break-even probability would need one, and this chapter names none.

Constructed readings. The record is constructed: generated from a stated truth of the plant with the stated sensor errors, and Section 20.4.3 compares the answers with that truth. It is not field data.

20.7 What to retain

Definition 20.1 (Two conditionings of one alarm)**.** Given the residual passed its level, $\Pr = \pi c/(\pi c + (1-\pi)\alpha)$ with the tails c and α. Given the residual's value, the posterior odds are the prior odds times $B(e)$ with $\ln B(e) = (s/\mathrm{sd}^2)(e - s/2)$. The first is the average of the second over the alarms past the level, and the second is the finer event.

Four conclusions carry beyond this plant. On the tails and for a rise the alarm is calibrated to, credibility is a ratio of two rates, so the false-alarm rate and the base rate set it, and a high catch does not. Those rates describe the plant they were set on, and a change the record already holds moves them at once. A credibility states what it conditions on and what size it means. And a cause the record has its own evidence for does not need a base rate, while a cause a reading cannot part from another stays open however good the prior.

Table 20.14 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 20.14: The terms of alarm credibility and where each is defined. Units for the shift plant.

Term	Expression	Units	Meaning	Defined in
Row pattern	eq. (20.1)	–	a rack's fraction reaches one reading	20.1.4
Base rate	π	1/window	a rate of onsets	20.1.4
Residual map	$\mathbf{K}_t$, eq. (20.2)	–	readings to the residuals, entries refound	20.2.1
Shift of a hypothesis	$\boldsymbol{\delta}$, $\mathbf{d}$, s, eq. (20.3)	K, kW	a scenario node through the map	20.2.1
Ramp's slope	$\hat{\gamma}$, eq. (20.5)	1/window	one unknown, the reference shared	20.2.3
Questions and answers				
Tail ratio, credibility	c/α, eq. (20.6)	–	given past the level, at the reference	20.3.1
First alarm, any fired	$1-(1-c)^d$; $d\,\pi$	–	two readings of the rate's unit	20.3.2
Density ratio	$B(e)$, eq. (20.7)	–	given the value; integrated over size	20.3.3
Routes on all readings	eq. (20.8)	–	agreement less cost	20.3.4
Tails with the coil in	α_r, c_r	–	r added to both means	20.3.5

20.8 Exercises

Exercise 20.1. *Application.* With a catch of 0.90, one false alarm a month and two excursions a month, over 8640 windows, find the probability of equation (20.6). ($\pi = 2.3148 \times 10^{-4}$, $\alpha = 1.1574 \times 10^{-4}$, so $\Pr = 0.6429$.)

Exercise 20.2. *Application.* Repeat with the catch of the level the hindcast applies, from $s = 1.8342\,\mathrm{K}$, $\mathrm{sd} = 0.2946\,\mathrm{K}$ and $L = 1.0847\,\mathrm{K}$. ($c = \bar{\Phi}((1.0847 - 1.8342)/0.2946) = 0.99453$, the tail ratio is 8592.7, the posterior odds 1.9895, so $\Pr = 0.6655$.)

Exercise 20.3. *Derivation.* Derive the density ratio's logarithm from the two normal densities and evaluate it at $e = +1.0970\,\mathrm{K}$. (The squares in e cancel, $\ln B = [e^2 - (e-s)^2]/(2\,\mathrm{sd}^2) = (s/\mathrm{sd}^2)(e - s/2)$, and with $e - s/2 = 0.1799$ this is 3.8013.)

Exercise 20.4. *Derivation.* At what residual is the second reading one half? (Setting $\ln B(e)$ equal to minus the log prior odds gives $e = s/2 - (\mathrm{sd}^2/s)\ln(\pi/(1-\pi))$, which is $1.3132\,\mathrm{K}$.)

Exercise 20.5. *Formulation.* State why the tails' answer has a band, and find its low end. (The false-alarm rate at a fixed level depends on the reference's shared error, which stays for the month; at its 97.5 % point, 3.5977 a month, equation (20.6) gives 0.3561.)

Appendix A

Atlas of operator questions

This appendix can be read independently of the chapters. Each row is a question an operator asks. The first column gives the question in the operator'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 Section 2.3.6; and the last the section in which the question is answered, or the chapter when a whole chapter answers it. The rows of Chapter 2 are the forms of its classification, Table 2.11. Chapter 3 develops the probability rules the rows rest on and holds no questions table, so it has no rows.

Operator question	General form	Mathematical form	Category	Computation	Answered in
Chapter 1. Operating Plants as Graphs					
What is the hall air temperature now?	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$, $\mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y})$	$p(T \mid \mathbf{y})$, $T = z_1$	estimation	posterior	1.3.1
What is the effective conductance of the air path?	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$, $\mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y})$	$p(G \mid \mathbf{y})$, $G = G_{\mathrm{ref}}\,\mathrm{e}^{z_2}$	estimation of a derived quantity	posterior	1.3.1
Is the hall below 24°C 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 hall take the second 30 kW tonight?	$\Pr(f(\mathbf{z}) < f^* \mid \mathbf{y})$	$\Pr(T_2 < 27.5^\circ\mathrm{C} \mid \mathbf{y})$, $T_2 = T_c + (Q_2/G_{\mathrm{ref}})\,\mathrm{e}^{-z_2}$	prediction under a threshold	posterior	1.3.2
Would another reading narrow the answers?	preposterior covariance $(\mathbf{\Sigma}^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a)^{-1}$	$\mathbf{\Sigma}_a$ for each candidate row	design	preposterior	1.6.1
Chapter 2. The Language of Operational Questions					
What is unknown, and how well is it known?	(2.9), (2.10)	$p(z_i \mid \mathbf{y})$, $\mathbb{E}[z_i \mid \mathbf{y}]$	estimation	posterior	2.4.1
Does a requirement hold, with the stated assurance?	(2.12)	$\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}) \geq \kappa$, f^* the requirement	prediction under a threshold	posterior	2.4.2
What happens in the coming shift?	(2.13)	$\int \Pr(A_{\mathrm{future}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	prediction under a threshold	posterior	2.4.4

Operator question	General form	Mathematical form	Category	Computation	Answered in
Is it rising, and when does it cross its limit?	(2.21)	$\Pr(\gamma > 0 \mid \mathbf{y})$; $\Pr(t^* \le s \mid \mathbf{y}) = \Pr(\gamma > 0,\ f_s(\mathbf{z}) \ge f^* \mid \mathbf{y})$	prediction under a threshold for the sign and a crossing by a date; estimation for a forecast and the distribution of a crossing time	trend	2.3.4, Chapter 11
Which explanation is correct?	(2.14), (2.15)	$\Pr(M_j \mid \mathbf{y}) \propto p(\mathbf{y} \mid M_j)\ \Pr(M_j)$	hypothesis comparison; diagnosis over a family of faults	evidence	2.5.1, 2.5.2, 2.5.3
What is the expected consequence?	(2.23)	$\mathcal{R}(a) = \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}] = \int L(a, \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	estimation of a derived quantity	expected consequence	2.5.4
What should be done or measured?	(2.23), (2.16)	$\arg\min_a \mathcal{R}(a)$; EVSI	design	expected consequence; preposterior	2.5.5, 2.6.1, 2.6.2
Chapter 4. The Anatomy of Probabilistic Inference					
What is the thermal resistance of the hall's air path?	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$, moments	$p(\rho \mid \mathbf{y})$: $\hat{z}_1$, $\sqrt{\Sigma_{11}}$	estimation	posterior	4.3.6
How warm is the hall at the 30 kW test?	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$, moments	$p(T_1 \mid \mathbf{y})$: $\hat{z}_2$, $\sqrt{\Sigma_{22}}$	estimation	posterior	4.3.6
How warm will the hall be at 60 kW?	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$, moments, quantiles	$p(T_2 \mid \mathbf{y})$: $\hat{z}_3$, $\sqrt{\Sigma_{33}}$, credible interval	estimation	posterior	4.3.6
Can the hall take 60 kW and stay below 28°C?	$\Pr(f(\mathbf{z}) \ge f^* \mid \mathbf{y})$	$\Pr(T_2 < T_{\max} \mid \mathbf{y})$, with $f = -T_2$	prediction under a threshold	posterior	4.3.7
Chapter 5. Virtual Sensors for Unmeasured Temperatures					
What is the inlet temperature at the rack with no sensor?	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$	$p(T_3 \mid \mathbf{y})$, rack 3: mean and 95 % band	estimation	posterior	5.3.1
What is that band made of?	$\mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y})$, split by entry	the budget of $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$	estimation	posterior	5.3.1
Is rack 3 over 27°C right now?	$\Pr(f(\mathbf{z}) \ge f^* \mid \mathbf{y})$	$\Pr(T_3 > T_{\max} \mid \mathbf{y})$	prediction under a threshold	posterior	5.3.2
Is rack 3 hotter than rack 1, the rack I can see?	$\Pr(f(\mathbf{z}) \ge f^* \mid \mathbf{y})$	$\Pr(T_3 - T_1 > 0 \mid \mathbf{y})$	prediction under a threshold	posterior	5.3.2
If I send someone with a probe, what will it read?	$\int p(y_{\mathrm{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	$p(y_{\mathrm{p}} \mid \mathbf{y})$, a 0.3 K probe at rack 3	prediction under a threshold	posterior	5.3.3
Chapter 6. Estimating Unmeasured Heat Flows					
How much heat is this unit actually rejecting?	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$	$p(Q_1 \mid \mathbf{y})$, unit 1: mean and 95 % band	estimation	posterior	6.3.1
And the other unit?	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$	$p(Q_2 \mid \mathbf{y})$, unit 2: mean and 95 % band	estimation	posterior	6.3.1

Operator question	General form	Mathematical form	Category	Computation	Answered in
Is unit 1 carrying more than unit 2?	$\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y})$	$\Pr(Q_1 - Q_2 > 0 \mid \mathbf{y})$	prediction under a threshold	posterior	6.3.2
What share of the hall's heat goes through unit 1?	$\Pr(f(\mathbf{z}) \leq f^* \mid \mathbf{y})$	$p(s_1 \mid \mathbf{y})$, $s_1 = Q_1/(Q_1 + Q_2)$	estimation	posterior	6.3.3
Chapter 7. Reconciling Inconsistent Meters					
Two meters should sum to the third and do not; which one is off?	$\Pr(M_j \mid \mathbf{y}) \propto p(\mathbf{y} \mid M_j)\,\Pr(M_j)$	ranking of M_0, M_Q, M_{W_1}, M_{W_2}	diagnosis	evidence	7.3.1
Is anything off at all?	$\Pr\big(t(y_{\text{new}}) \geq t^* \mid M_0\big)$, a predictive tail	reading-row χ^2 against χ^2_ν; largest $\lvert r^{\text{N}}_k \rvert$	prediction under a threshold	posterior	7.3.2
By how much is it off?	$\Pr(f(\mathbf{z}) \leq f^* \mid \mathbf{y}, M_j)$	$p(b_j \mid \mathbf{y}, M_j)$ under the two leading drawings	estimation	posterior	7.3.3
Which check would settle it?	$U(a)$, the preposterior form	expected $\ln B$ of a check between the two leaders	design	preposterior	7.3.4
Chapter 8. Joint Limits and the Operating Envelope					
What is the probability the hot aisle is already above 27°C somewhere I cannot see?	$\Pr(\bigcup_k \{f_k(\mathbf{z}) > f^*\} \mid \mathbf{y})$	$\Pr(\max(T_2, T_3) > T_{\max} \mid \mathbf{y})$ from the joint of (T_2, T_3)	prediction under a threshold	posterior	8.3.1
How much do the two racks move together?	$\mathrm{Cov}(f_a(\mathbf{z}), f_b(\mathbf{z}) \mid \mathbf{y})$	$\varrho(T_2, T_3)$ from $\mathbf{D}\boldsymbol{\Sigma}\mathbf{D}^{\mathsf{T}}$	estimation	posterior	8.3.2
Could both be over at once?	$\Pr(\bigcap_k \{f_k(\mathbf{z}) > f^*\} \mid \mathbf{y})$	$\Pr(T_2 > T_{\max},\ T_3 > T_{\max} \mid \mathbf{y})$	prediction under a threshold	posterior	8.3.3
When do the marginals one by one mislead?	the same forms, under $p(\mathbf{z} \mid \mathbf{y}_{\text{II}})$	the first three without the return-air and discharge readings	prediction under a threshold	posterior	8.3.4
Chapter 9. Anomaly Detection from Residuals					
Is this residual noise or a change?	$\int p(y_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$, its tail	$r^{\text{N}}_k = r_k(\hat{\mathbf{z}}_t)/\sqrt{1 - h_{kk}}$ of unit 2's discharge, window 24; $2\Phi(-\lvert r^{\text{N}}_k \rvert)$	prediction under a threshold	posterior	9.3.1
Is the plant as a whole still behaving like the first hour?	$\Pr\big(X^2(\mathbf{y}_{\text{rep}}) \geq X^2(\mathbf{y})\big)$	$X^2_t = \sum_{k \leq 5} r_k(\hat{\mathbf{z}}_t)^2$ against $\chi^2_{\nu_t}$, $\nu_t = \sum_{k \leq 5}(1 - h_{kk})$	prediction under a threshold	posterior	9.3.2
When did it start to show?	the first t with the tail below α	first t with $\Pr(\chi^2_{\nu_t} > X^2_t) < 10^{-3}$; first t with $\max_k \lvert r^{\text{N}}_k \rvert > c$	prediction under a threshold	posterior	9.3.3
Why not refit everything each window?	$\Pr(f(\mathbf{z}_t) \leq f^* \mid \mathbf{y}_t)$, the window alone	$\hat{g}_2$ and r^{N} of unit 2's discharge under the rating priors of Chapter 5	estimation	posterior	9.3.4
Chapter 10. Fault Isolation and Derate Estimation					
If something changed, where?	(2.15)	the ranking of M_0 and the 12 candidates at window 24	diagnosis	evidence	10.3.1
A coil or a sensor?	(2.15), two drawings	$\ln B(M_{g_2} : M_{b_{T_{d,2}}})$ per window	diagnosis	evidence	10.3.2
When can the ranking be acted on?	(2.15), per window	the first window past the rule, and the first after which it never fails	diagnosis	evidence	10.3.3

Operator question	General form	Mathematical form	Category	Computation	Answered in
Can these readings say it is the coil?	(2.15), the extended family	the substitution $L_{d,2} = (G_2^{\mathrm{b}} - G_2)(T_r - T_s)$	diagnosis	evidence	10.3.4
Given it is unit B, how far is it from its rating?	(2.9)	$p(g_2 \mid \mathbf{y}_t, M_{g_2})$, $d_2 = \mathrm{e}^{g_2}$	estimation	posterior	10.3.5
Chapter 11. Trends and Limit-Crossing Forecasts					
Is it getting worse.	(2.21), line 1	$\Pr(\gamma > 0 \mid \mathbf{y})$	prediction under a threshold	trend	11.3.1
How fast?	(2.9)	γ per hour; the conductance lost per hour	estimation	trend	11.3.1
When did it begin?	(2.9)	$\tau = \bar{t} + (\alpha - g_b)/\gamma$, a ratio	estimation	trend	11.3.2
When does the plant cross its limit, and with what band?	(2.21), line 2; (2.9)	the median $t^* = \bar{t} + (\alpha - g^*)/\gamma$ and its band by inversion; $\Pr(t^* \le s \mid \mathbf{y})$	prediction under a threshold; estimation	trend	11.3.3
Could it be over already?	(2.12)	$\Pr(T_{3,24} \ge 27^\circ\mathrm{C} \mid \mathbf{y})$	prediction under a threshold	trend	11.3.3
How does the answer firm with more windows, and can it answer at all yet?	(2.21) per T	the answers above at each T; the band's upper end	prediction under a threshold	trend	11.3.4
Is a straight line the right form?	(2.14)	evidences of no change, a line over the whole record, a ramp from τ, a curved ramp, a step	hypothesis comparison	evidence	11.3.5
Chapter 12. Set-Point Changes as Scenarios					
If I raise the supply temperature 2 K, what is the probability the inlets stay under the limit?	(2.12), a joint tail at a scenario node	$\Pr(\max_k T_k^* < 27^\circ\mathrm{C} \mid \mathbf{y},\ T_s^* = 16^\circ\mathrm{C})$	prediction under a threshold	posterior	12.3.1
Which rack carries it, and what stands behind its band?	(2.9), (2.11)	T_k^* and its band; the budget of T_{B3}^*	estimation	posterior	12.3.2
Is the raise seen by the sensors on the screen?	(2.12) on two subsets	the measured four against the unread six; the sensors see the raise and not the tail	prediction under a threshold	posterior	12.3.3
Does the probability fall with the raise, and how large a raise holds at a stated assurance?	(2.12) read along T_s^*	$\Pr(\cdot \mid T_s^*)$; the largest T_s^* with $\Pr \ge \kappa$	prediction under a threshold	posterior	12.3.4
Is the raise a pinned value or a shift?	(2.12), two nodes	$T_s^* = 16^\circ\mathrm{C}$ against $T_s + 2\,\mathrm{K}$	prediction under a threshold	posterior	12.3.5
Chapter 13. Equipment Outages and Maintenance Windows					
Can I take unit A down for maintenance now without a limit violation?	(2.12), a joint tail at a scenario node	$\Pr(\max_{k \in A} T_k^* < T_{\max} \mid \mathbf{y},\ \text{A1 out})$	prediction under a threshold	posterior	13.3.1

Operator question	General form	Mathematical form	Category	Computation	Answered in
And at the forecast peak?	(2.12); (2.21), line 3	the same with $Q_A \to Q_f(t)$, $t = 144$	prediction under a threshold	trend	13.3.2
Does it matter which unit?	(2.12), per unit	the same, six nodes	prediction under a threshold	posterior	13.3.3
Where does the air go?	(2.12), two forms	the plant as modeled against rack air flow held	prediction under a threshold	posterior	13.3.4
How much load would have to leave first?	(2.12), inverted	the largest Q_A^* with $\Pr(\cdot) \geq \kappa$	prediction under a threshold	posterior	13.3.5
Chapter 14. Load Transfer Between Halls					
If I move 40 kW from hall A to hall B, where does the margin go?	(2.9)	$27 - T_k^*$ per rack, against the present	estimation	posterior	14.3.1
What is the band on each node?	(2.9), (2.11)	T_{A5}^* and T_{B3}^*, the rack with its hall's largest survey fraction	estimation	posterior	14.3.2
What is the probability both stay inside?	(2.12), a joint tail	$\Pr(\max_k T_k^* < 27 \mid \mathbf{y},\ \Delta = 40)$	prediction under a threshold	posterior	14.3.3
What of another size, or the other direction?	(2.12) along Δ	$\Pr(\cdot \mid \Delta)$; the largest Δ with $\Pr \geq \kappa$	prediction under a threshold	posterior	14.3.4
Chapter 15. Look-Ahead Risk Without Intervention					
If I do nothing, what is the probability of a violation before the morning shift? Trend forecast of the load against the margin band, integrated over the horizon.	(2.21), line 2; (2.12) per window	$\Pr(E_{144} \mid \mathbf{y})$	prediction under a threshold	trend	15.3.1
Does the path matter, or only the last window?	(2.21), line 3	$\Pr(T_{\mathrm{B3},144} \geq T_{\max} \mid \mathbf{y})$ with and without wander, against $\Pr(E_{144} \mid \mathbf{y})$	prediction under a threshold	trend	15.3.1
How does it grow with the horizon?	(2.21), line 2, per s	$\Pr(E_s \mid \mathbf{y})$; the conditional hazard per hour	prediction under a threshold	trend	15.3.2
Which rack carries it?	(2.12)	$\Pr(\max_t T_{k,t} \geq T_{\max} \mid \mathbf{y})$ per rack; the budget; $\Pr(E \mid \beta_{\mathrm{B3}}, \mathbf{y})$	prediction under a threshold; estimation	trend	15.3.3
How much is what we do not know, how much is what has not happened yet?	(2.10)	the split of $\mathrm{Var}(T_{\mathrm{B3},144} \mid \mathbf{y})$; $\Pr(E)$ with the entries at the fit	estimation	trend	15.3.4
What would the same question have said earlier and later?	(2.21) per T	$\Pr(E \mid \mathbf{y}_{1:T})$ from each decision window; its spread over fresh sensor errors	prediction under a threshold	trend	15.3.5

Operator question	General form	Mathematical form	Category	Computation	Answered in
Chapter 16. Goodness of Fit in Operation					
Is the model still describing the plant? Goodness of fit on the reading residuals against their noise; a posterior predictive check per sensor.	(2.13)	X_t^2 of the reading rows on $\nu_t = \sum_k (1 - h_{kk})$; r_k^{N} per sensor	prediction under a threshold	posterior	16.3.1
Why has the light stayed green all night?	(2.13)	a stated derate split between reading rows and prior rows; probability of detection at window 144	prediction under a threshold	posterior	16.3.2
Is tonight's plant the plant of 19:00?	(2.13)	$Q_t = \mathbf{e}_t^{\mathsf{T}} \mathbf{V}_t^{-1} \mathbf{e}_t$ against the reference, on 13; per sensor by the hour	prediction under a threshold	posterior	16.3.3
When did that show, and how firmly?	(2.12) per window	first t with tail below 10^{-3}; its spread over fresh sensor errors	prediction under a threshold	posterior	16.3.4
What restores the fit?	(2.14)	X^2 and evidence with g_{B2} free per window, with and without the meter's bias	hypothesis comparison	evidence	16.3.5
What can no test see?	(2.10)	the empty columns; the share of $\mathrm{Var}(T_{\mathrm{B3}} \mid \mathbf{y})$ on them	estimation	posterior	16.3.6
Chapter 17. Sensor Loss and Measurement Redundancy					
If this one sensor fails, how wide does my band on the hot-aisle temperature get? Remove its row from $\mathbf{\Lambda}$, re-read the band; the reverse of sensor placement.	(2.10)	$\mathbf{g}^{\mathsf{T}} (\mathbf{\Lambda} - \mathbf{a}_k^{\mathsf{T}} \mathbf{a}_k)^{-1} \mathbf{g}$	estimation	posterior	17.3.1
Which reading is carrying the answer?	(2.10)	the ranking over k; $\mathbf{g}^{\mathsf{T}} \mathbf{\Sigma} \mathbf{a}_k^{\mathsf{T}}$, h_{kk}	estimation	posterior	17.3.2
What if several go together?	(2.10)	the block deletion	estimation	posterior	17.3.3
Which reading has no replacement?	(2.10)	entries back at their prior; the rank; the effective degrees of freedom	estimation	posterior	17.3.4
What if it fails and stays connected?	(2.13)	X^2, e_k, $T_{r,\mathrm{B}}$ with the reading held at its last value	prediction under a threshold	posterior	17.3.5
What does a loss do to tonight's answers?	(2.12)	Chapter 12's joint, and its largest set point, with the row removed	prediction under a threshold	posterior	17.3.6

Operator question	General form	Mathematical form	Category	Computation	Answered in
Chapter 18. Selecting the Next Field Measurement					
Of the readings I could request from the field technician, which one narrows the quantity I am worried about most? Preposterior over the candidates, ranked by width taken off that quantity.	(2.16)	$\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} - (\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{a})^2/(1 + \mathbf{a}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{a})$	design	preposterior	18.3.1
Is there one list?	(2.16)	the matrix of candidates against worries	design	preposterior	18.3.2
Does the reading settle anything?	(2.16), (2.12)	$\Pr(P_a(\tilde{y}_a) \geq 0.95 \text{ or } P_a(\tilde{y}_a) < 0.05 \mid \mathbf{y})$	design	preposterior	18.3.3
What should the second and third reading be?	(2.16)	the greedy order against the best pair, by worry	design	preposterior	18.3.4
How good must the reading be?	(2.16)	width against the reading's spread; repeats; the band with β_{B3} known	design	preposterior	18.3.5
Chapter 19. Calibrating Alarm Thresholds					
Where should the alarm sit so the false-alarm rate is one per month and a real excursion of 2 K is caught nine times in ten? Calibrated decision levels: the predictive spread of the monitored entry gives the false-alarm tail, the detection power gives the miss rate.	(2.13), (2.12)	$L_\alpha = z_\alpha \,\mathrm{sd}$; $L_\pi = s_r - z_\pi \,\mathrm{sd}$; the interval $[L_\alpha, L_\pi]$	prediction under a threshold	posterior	19.3.1
Can a fixed level on the plain reading do it?	(2.21)	$\bar{y}_{144} + z_\alpha \,\mathrm{sd}_{144}$ or the mean tail, against $\bar{y}_t + 2 - z_\pi \,\mathrm{sd}_t$, on one fit and on the reference	prediction under a threshold	trend	19.3.2
Which excursions does it catch?	(2.12)	the catch by route and monitor	prediction under a threshold	posterior	19.3.3
And the rack with no sensor?	(2.12)	the shift of the virtual sensor by route	prediction under a threshold	posterior	19.3.4
Is "one a month" a number?	(2.11)	the rate over x; two windows in a row	estimation	posterior	19.3.5
What would tonight have done with it?	(2.12) per window	windows past each level; standing alarms; the first fire over fresh errors	prediction under a threshold	posterior	19.3.6

Operator question	General form	Mathematical form	Category	Computation	Answered in
Chapter 20. Alarm Credibility and Base Rates					
This alarm just fired; what is the probability it is real? Pr(excursion | alarm) by Bayes over the two tails of Chapter 19 and the base rate of excursions.	(2.14)	$\pi c/(\pi c + (1-\pi)\alpha)$	hypothesis comparison	evidence	20.3.1
What does twice a month mean here?	(2.14)	π by duration; first alarm against any	hypothesis comparison	evidence	20.3.2
It fired at this value; does that matter?	(2.14)	density ratio at e; $\ln = (s/\mathrm{sd}^2)(e - s/2)$; the size integrated	hypothesis comparison	evidence	20.3.3
What do the other fifteen readings say?	(2.14)	ratio on all the residuals, by route; the plain readings	hypothesis comparison	evidence	20.3.4
If not an excursion, what?	(2.15)	ramp on g_{B2}, and the two tails with it in; a sensor at this window	diagnosis	evidence	20.3.5
What does the next window say?	(2.14)	the odds updated window by window	hypothesis comparison	evidence	20.3.6