

Operations and Control by Derivation

The operator's questions, from the plant graph to inference in the control room

Draft · 337 pages

Ricardo Marquez, Ph.D.

ricardo-marquez.com

Operations and Control by Derivation

Volume I

From the operator's questions to probabilistic inference and decisions

Ricardo Marquez

November 1, 2026

How to read this book

This book is written for the operator of a running plant and for the operations engineer who supports the control room. Its questions are the ones a shift asks: where the plant is right now, whether something is wrong, what happens if the operator acts, how far the model can be trusted today, and where the alarms should sit. A dashboard answers them from point readings and fixed limits. It shows the readings that exist; it cannot show the inlet temperature of a rack with no sensor, which of three meters that should balance is wrong, the probability of a hot spot nobody measures, or whether a unit can come out for maintenance at the night's peak. Each of these is a probability statement over the unknowns of one model of the plant, and the book derives and computes each one.

The plant grows with the book. Part I works on one equipment hall cooled by a single unit, a state of two and then three entries, so that every object of the method can be seen whole. Parts II and III add a second cooling unit and the racks the hall serves, and follow that hall through a short run of telemetry windows. From Part IV the plant is two data halls on one chilled-water loop, followed through a twelve-hour night shift in five-minute windows. On each plant one fitted model serves every question put to it: many of the questions are the same posterior read at a different place of the plant or at a set point not yet applied, and the chapters say so where it happens. The book has two volumes: this one holds Parts I to III, and chapters 12–20 are in Volume II.

Every chapter that answers a question follows the same progression, and the progression is the method.

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

Step 1: the question. The operator's question in the operator's words, then as a probability or an expectation of a function of the state, with the unknown on the left of the conditioning bar and everything treated as known on the right.

Step 2: the graph extended for the question. Scenario nodes for a set point, a unit state or a load not yet applied; candidate faults when something may have changed; and the telemetry attached as rows.

Step 3: the posterior of the state. Derived and computed once, before any question is answered: the MAP state, the information, the covariance, and the exact posterior where one exists.

Step 4: the answers. Each question read from that posterior down to its number: an entry of the state or the gradient of a derived quantity; a trend across windows; or, when explanations compete, a ratio of evidences.

Step 5: evaluation and checks. The gaps between the reported values and the exact values or draws collected, the approximations tested and, because the record is constructed from a stated truth, the answers tested against that truth.

Step 6: the report. The probability, the band or the ranking the operator acts on, with the exact value beside it where one exists.

Step 7: the next reading. If the band is too wide, the readings that could be requested from the field ranked by the width each would remove, or the sensor that most deserves trust.

Five kinds of question cover what the book asks, distinguished by the mathematical object that answers them rather than by the word the control room uses: *estimation*, *prediction under a threshold*, *hypothesis comparison*, *diagnosis* and *design*. Beneath them sit five computations: the *posterior* of the state of a window, the *trend* of an entry across windows, the *evidence* of a drawing of the plant, the *preposterior* after a reading not yet taken, and the *expected consequence* of an action. Operations is mostly posterior and trend. The notation chapter that follows fixes the symbols, and the appendix collects the operator's questions in an atlas that maps each to its mathematical form, its computation and the section that answers it.

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 method and its notation are those of the companion book *Reliability Engineering by Derivation*, which asks the reliability engineer's questions of a plant of the same kind.

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 **vii**

I Thinking in Probabilities **1**

1 Operating plants as graphs **2**
1.1 One system and several questions . . . 2
1.2 The posterior of the state (step 3) . . . 15
1.3 The answers (step 4) . . . 20
1.4 The checks (step 5) . . . 25
1.5 The report (step 6) . . . 28
1.6 The next reading (step 7) . . . 29
1.7 What to retain . . . 33
1.8 Exercises . . . 34

2 The language of operational questions **37**
2.1 Operator questions and their mathematical form . . . 37
2.2 The grammar of probability statements . . . 40
2.3 The general forms and the five computations . . . 43
2.4 Estimation, prediction and joint events . . . 54
2.5 Model comparison, diagnosis and decisions . . . 66
2.6 Measurements, classification and workflow . . . 69
2.7 What to retain . . . 75
2.8 Exercises . . . 78

3 Uncertainty and conditional probability **81**
3.1 Uncertainty and evidence in operating questions . . . 81
3.2 Conditional probability and Bayes' theorem . . . 86
3.3 Bayesian diagnosis of a cooling fault . . . 89
3.4 Marginalization and dependence . . . 91
3.5 Conditioning in a physical model . . . 96
3.6 Predictive variance and the posterior covariance . . . 104
3.7 What to retain . . . 111
3.8 Exercises . . . 112

4 The anatomy of probabilistic inference **115**
4.1 From a plant graph to an operating probability . . . 115
4.2 The posterior of the state (step 3) . . . 123
4.3 The answers (step 4) . . . 135

4.4 The checks (step 5) . . . 141
4.5 The report (step 6) . . . 147
4.6 The next reading (step 7) . . . 148
4.7 What to retain . . . 152
4.8 Exercises . . . 155

II Estimating the State of the Plant 157

5 Virtual sensors for unmeasured temperatures 158
5.1 Rack inlets the hall does not measure . . . 158
5.2 The posterior of the state (step 3) . . . 164
5.3 The answers (step 4) . . . 167
5.4 The checks (step 5) . . . 172
5.5 The report (step 6) . . . 174
5.6 The next reading (step 7) . . . 175
5.7 What to retain . . . 177
5.8 Exercises . . . 178

6 Estimating unmeasured heat flows 180
6.1 Heat the units do not meter . . . 180
6.2 The posterior of the state (step 3) . . . 184
6.3 The answers (step 4) . . . 184
6.4 The checks (step 5) . . . 191
6.5 The report (step 6) . . . 193
6.6 The next reading (step 7) . . . 195
6.7 What to retain . . . 197
6.8 Exercises . . . 198

7 Reconciling inconsistent meters 200
7.1 Three meters that do not balance . . . 200
7.2 The posterior of the state (step 3) . . . 205
7.3 The answers (step 4) . . . 210
7.4 The checks (step 5) . . . 217
7.5 The report (step 6) . . . 219
7.6 The next reading (step 7) . . . 222
7.7 What to retain . . . 224
7.8 Exercises . . . 225

8 Joint limits and the operating envelope 227
8.1 The operating envelope over racks without sensors . . . 227
8.2 The posterior of the state (step 3) . . . 231
8.3 The answers (step 4) . . . 233
8.4 The checks (step 5) . . . 241
8.5 The report (step 6) . . . 242
8.6 The next reading (step 7) . . . 244
8.7 What to retain . . . 247
8.8 Exercises . . . 248

III Detecting and Diagnosing Abnormal Operation 249

9 Anomaly detection from residuals 250
9.1 A discharge reading that climbs . . . 250
9.2 The posterior of the state (step 3) . . . 255
9.3 The answers (step 4) . . . 258
9.4 The checks (step 5) . . . 267
9.5 The report (step 6) . . . 270
9.6 The next reading (step 7) . . . 271
9.7 What to retain . . . 273
9.8 Exercises . . . 274

10 Fault isolation and derate estimation 276
10.1 One discharge of two . . . 276
10.2 The posterior of the state (step 3) . . . 281
10.3 The answers (step 4) . . . 283
10.4 The checks (step 5) . . . 290
10.5 The report (step 6) . . . 293
10.6 The next reading (step 7) . . . 294
10.7 What to retain . . . 297
10.8 Exercises . . . 298

11 Trends and limit-crossing forecasts 299
11.1 A limit no sensor reads . . . 299
11.2 The posterior of the state (step 3) . . . 303
11.3 The answers (step 4) . . . 306
11.4 The checks (step 5) . . . 316
11.5 The report (step 6) . . . 319
11.6 The next reading (step 7) . . . 320
11.7 What to retain . . . 321
11.8 Exercises . . . 322

A Atlas of operator questions 324

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.

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.

Part I

Thinking in Probabilities

Chapter 1

Operating Plants as Graphs

1.1 One system and several questions

At 20:00 the night operator of a data center equipment hall has a work order to bring 30 kW of servers online in the hall tonight, doubling its load from 30 to 60 kW, and the hall air must then stay below 27.5°C. The record is thin: one hall sensor, with an error of about half a kelvin, reads 23.2°C at the present load, and the design handbook gives the conductance of the hall's air path to the supply air, held at 18°C by the cooling unit's discharge controller, as 7.5 kW/K, right to within a factor of about 1.35. The dashboard's nameplate tile works from that value. Heat leaves the hall in proportion to its excess over the supply air, so at 7.5 kW/K the hall sits at 22°C under 30 kW, comfortably below its present limit of 24°C, and at 26°C under 60 kW, well below 27.5°C: the nameplate tile shows the full load passing with room to spare.

The reading tile beside it works from the reading and tells another story. Taken as the truth, 23.2°C at 30 kW means a conductance of 5.77 kW/K, and at that conductance the full load puts the hall at 28.4°C, over the limit. One number on the screen says the work order can go ahead and the other that it cannot; each treats its own input as exact, neither is weighed against the other, and neither says how sure the operator can be. The decision needs a probability, that the hall stays below 27.5°C at a load at which it has never run, and that probability can only come from holding the handbook value, the reading and the energy balance that ties them together in one description of the hall, each with its own uncertainty. The question of this chapter is therefore

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

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

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

Key idea

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

1.1.1 Several questions about one equipment hall

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

1. What is the hall air temperature right now?
2. What is the effective conductance of the air path?
3. Can the hall take the second 30 kW tonight without going over its limit?
4. Is the hall sensor reading true, or is it biased?
5. Would another reading narrow the answer?
6. What happens if the cooling unit loses half its airflow?

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

concern the same physical system: the same hall, the same air path, the same sensor.

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

1.1.2 The system, the question and the evidence

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

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

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

1.1.3 Larger plant graphs and factor graphs

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

A cooling network. Real plants may include hundreds or thousands of components, each with its own equations and uncertainties. Consider a cooling network with two chillers, two pumps, a common electrical supply and several temperature sensors. Its plant graph may contain temperatures, pressures, flows and electrical states at its nodes; physical components represented by paths or by component relations; heat, mass and power conservation equations; telemetry and historical records; the run and trip states of equipment; and shared parameters and shared dependencies.

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

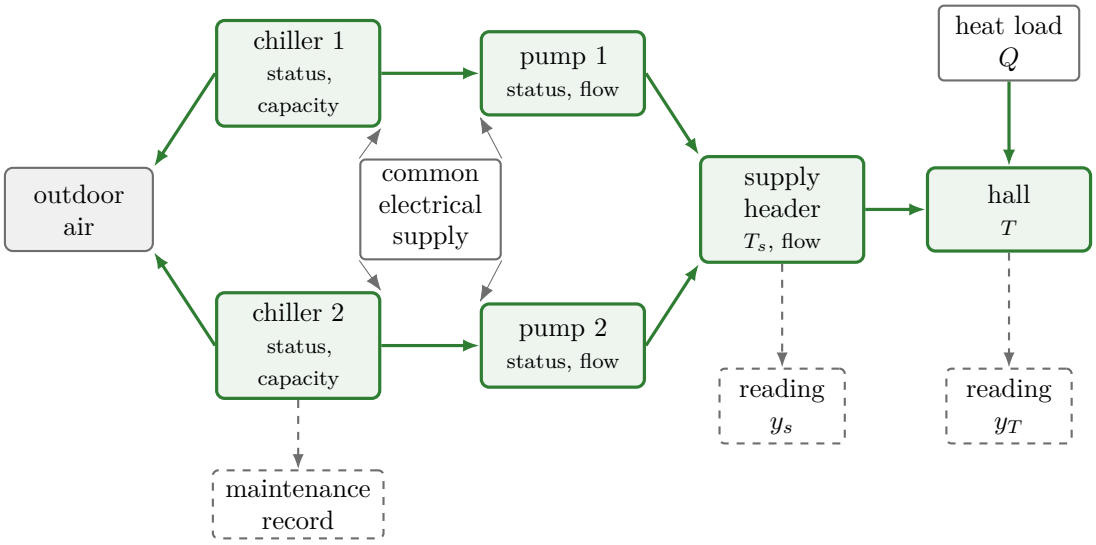


Figure 1.1: Plant graph of a reduced cooling network (illustrative; return piping omitted). Free: the status and capacity of each chiller, the status and flow of each pump, the header water temperature T_s and flow, and the hall temperature T. Pinned: the outdoor air, to which the chillers reject heat,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

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

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

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

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

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

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

1.1.4 The plant graph (step 0)

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

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

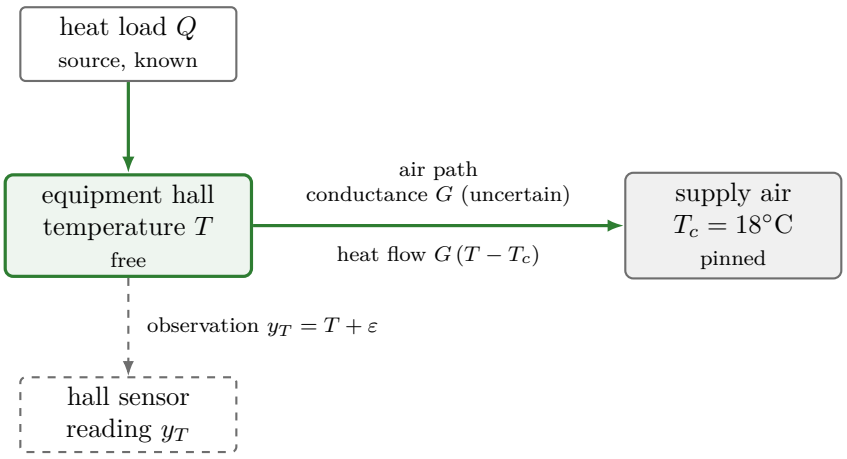


Figure 1.2: Plant graph of the equipment hall in steady state. Free: the hall air temperature T and the conductance G of the air path. Pinned: the supply-air temperature $T_c = 18°\text{C}$, held by the cooling unit's discharge controller, and the heat load Q. Observed: T, through the reading y_T with error ε. Solid arrows are physical relations; the dashed arrow is an observation.

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

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

Element	Role	Symbol	Unit	In the example
Equipment hall	free node	T	°C	air temperature unknown
Supply air	pinned node	T_c	°C	held at 18°C by the discharge controller
Heat load	source	Q	kW	thermal power supplied by the servers, known
Air path	path (edge)	G	kW/K	effective conductance, uncertain
Hall sensor	observation	y_T	°C	a reading related to the hall temperature

For now the supply-air temperature and the heat load are assumed known, and the effective conductance is uncertain. The discharge controller holds T_c at its set point, which is what makes the supply air a pinned boundary; the choice fails if the unit saturates and can no longer hold it (Section 1.6.5). The conductance lumps every mechanism by which the hall's heat reaches the supply air, the airflow through the cooling unit and the mixing of the hall air alike, into one number; how well that lumping represents the hall is a modeling assumption, examined in Section 1.6.5. This is the starting physical model.

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

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

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

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

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

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

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

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

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

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

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

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

the steady-state temperature of the hall at its design conductance. The calculation used no probability, estimated no uncertainty and used no observational evidence. It is a deterministic *forward calculation*: from pinned inputs and a specified conductance to the temperature they imply. Every later chapter needs it, because every inference compares temperatures computed this way with what the sensor reports. Doubling the load doubles the excess $T - T_c$ at every conductance: at $G = 6\,\text{kW/K}$ the hall at 60 kW would be at 28°C.

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

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

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

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

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

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

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

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

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

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

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

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

or, moving everything to one side,

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

The function $r_{\mathrm{phys}}(\mathbf{z})$ is a *residual*, in kW. A residual measures how far a candidate state is from satisfying a relation: it is zero when the relation holds exactly and grows as the candidate departs from it. With $G_{\mathrm{ref}} = 6\,\mathrm{kW/K}$, $T_c = 18°\mathrm{C}$ and $Q = 30\,\mathrm{kW}$, the candidate $\mathbf{z} = (23°\mathrm{C},\ 0)$ gives $r_{\mathrm{phys}} = 6 \times 1 \times 5 - 30 = 0$: it is the forward calculation. The candidate $\mathbf{z} = (23.2°\mathrm{C},\ 0)$ gives $6 \times 5.2 - 30 = 1.2\,\mathrm{kW}$: at that temperature and conductance the path would remove 1.2 kW more than the servers supply, so the hall could not be in steady state. The candidate $(23.2°\mathrm{C},\ \ln(30/31.2)) = (23.2°\mathrm{C},\ -0.0392)$ again gives zero: the temperature 23.2°C is consistent with the physics if the conductance is $30/5.2 = 5.769\,\mathrm{kW/K}$. The zero set

of r_{phys} is the curve $T = T_c + Q/(G_{\text{ref}}\,\mathrm{e}^{g})$ in the plane of (g, T); the physics alone allows every point of that curve and no point off it.

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

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

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

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

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

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

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

1.1.5 The record and the handbook value

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

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

The sensor is not perfectly accurate. Its measurement model is

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

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

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

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

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

1.1.6 The questions (step 1)

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

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

Operator question	General form	Mathematical form	Category	Answered in
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	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	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	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	1.3.2
Would another reading narrow the answers?	preposterior covariance $(\boldsymbol{\Sigma}^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a)^{-1}$	$\boldsymbol{\Sigma}_a$ for each candidate row	design	1.6.1

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

are not in the table. Whether the hall sensor is biased compares two complete models of the same reading, the plant with an unbiased sensor and the plant with a bias entry on the reading row, and its answer is a ratio of their evidences rather than a property of either posterior; when every sensor of a larger plant is a candidate, the same comparison run over the candidates is a diagnosis. Chapter 7 answers it, and the residual check of Section 1.4.3 takes a first, coarse look at it. What happens if the cooling unit loses half its airflow is a question about a transient that the steady-state graph cannot represent, and Section 1.6.5 states what the graph would need for it.

1.1.7 The evidence as rows (step 2)

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

The modeling choices. One choice makes the physics row, and it buys something and can fail. The row carries the balance of Section 1.1.4 in the limit $\sigma_p \to 0$: this is what confines the posterior to the physics curve, so that a reading of the temperature informs the conductance, and it fails if heat leaves the hall by a mechanism the single conductance does not lump or if the hall is not in steady state (Section 1.6.5).

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

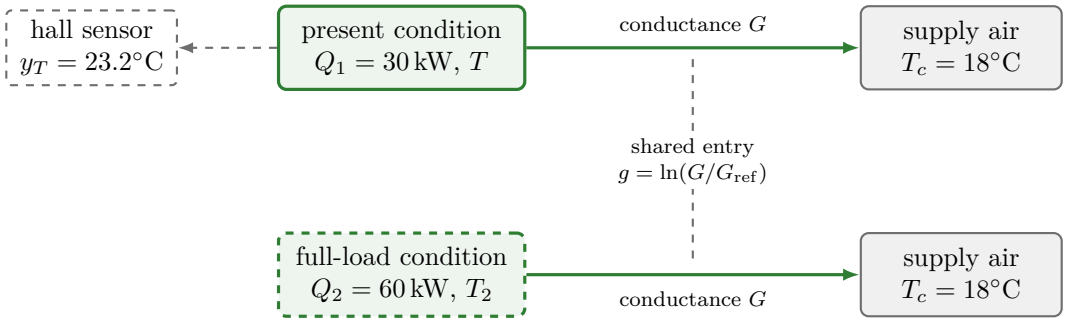


Figure 1.3: Two operating conditions of the equipment hall linked through the same physical conductance. Free: the present-condition temperature T and the shared log-conductance g, the two entries of the state; the full-load temperature T_2 (dashed node) is a function of g. Pinned: the supply-air temperature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

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

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

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

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

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

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

with T in °C and g dimensionless. Three rows act on it, the physics row (1.5), the reading row (1.6) and the prior row, each divided by its standard deviation. Stacked into one residual vector, the three rows are

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

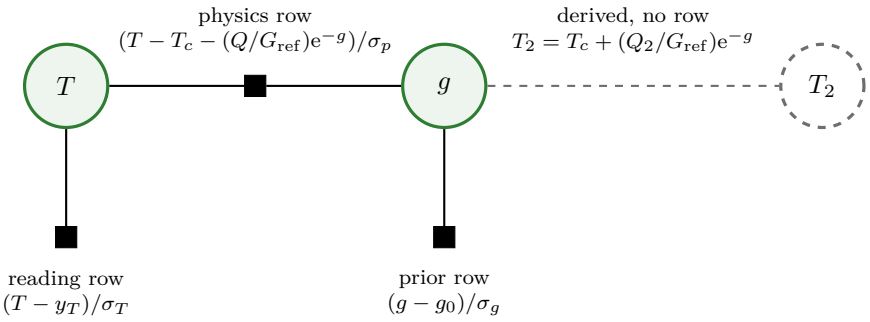


Figure 1.4: Factor graph of the equipment hall with its scenario node. Solid circles: the entries of the state, T and g. Squares: the rows of the model, each linked to the entries it involves, the pattern (1.9) of the Jacobian. Dashed: the full-load temperature T_2, a function of g with no row of its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

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

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

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

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

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

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

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

in kW/K and kW; on the zero set the second equals the load Q by equation (1.1). A one-kelvin error in the temperature unbalances the hall by G kilowatts, and a unit change in g, a factor e in the conductance, by about Q kilowatts to first order. Divided by G, to reach the temperature form, and by σ_p, the pair becomes $(1/\sigma_p,\ (Q/G)/\sigma_p)$, the first row of equation (1.12) on the zero set.

1.2 The posterior of the state (step 3)

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

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

1.2.1 The posterior from the rows

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

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

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

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

The normal equations and the Gauss–Newton step. Near a state $\mathbf{z}$, each row is replaced by its tangent, $\mathbf{r}(\mathbf{z} + \Delta\mathbf{z}) \approx \mathbf{r}(\mathbf{z}) + \mathbf{J}\,\Delta\mathbf{z}$, and the objective by $\frac{1}{2}\,\|\mathbf{r} + \mathbf{J}\,\Delta\mathbf{z}\|^2$, a quadratic in the step $\Delta\mathbf{z}$. Its gradient with respect to the step is $\mathbf{J}^{\mathsf{T}}(\mathbf{r} + \mathbf{J}\,\Delta\mathbf{z})$, and setting it to zero gives the *normal equations*

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

a linear system of two equations in the two entries of the step. The Gauss–Newton method solves it, moves to $\mathbf{z} + \Delta\mathbf{z}$, evaluates $\mathbf{r}$ and $\mathbf{J}$ there, and repeats until the step no longer changes

the state at the precision wanted. The vector $\mathbf{J}^\mathsf{T}\mathbf{r}$ is the gradient of the objective (1.11), so the iteration stops where that gradient vanishes, at the MAP state. The iteration starts here from the design state $\mathbf{z}^0 = (23°\mathrm{C},\ 0)^\mathsf{T}$ of the forward calculation, at which the physics row is zero.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

$$\lambda_{\|} = \mathbf{u}^{\mathsf{T}}(\mathbf{J}_2^{\mathsf{T}}\mathbf{J}_2 + \mathbf{J}_3^{\mathsf{T}}\mathbf{J}_3 + \mathbf{C})\mathbf{u} = \frac{n_1^2/\sigma_T^2 + 1/\sigma_g^2 + n_1 r_2/\sigma_T}{n_1^2 + 1}, \tag{1.18}$$

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

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

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

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

One-dimensional quadrature of this density over g gives every probability and moment of a quantity derived from the state exactly, for T through $T(g)$ and for G through $G_{\mathrm{ref}}\,\mathrm{e}^{g}$.

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

1.2.2 The posterior of the hall

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

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

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

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

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

and the first step is their solution,

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

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

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

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

The first step does almost all the work, because over a step this short the rows are nearly linear. It leaves the physics row slightly violated, $r_1 = -0.042$, through the curvature of e^{-g} that the tangent ignores; the second step restores it, and the third changes nothing at the printed precision. The MAP state is

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

There the physics holds to $r_1\,\sigma_p = 5\times10^{-5}\,\mathrm{K}$, and the reading and the prior share the disagreement: the temperature lies 0.26 standard deviations below the reading and g lies 0.79 standard deviations below the handbook value.

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

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

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

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

The physics term has rank one: all its information lies along its own row (100, 507.0), across the physics curve, and none along the curve; only the reading's 4 and the prior's 11.11 inform along it. The eigenvalues of $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ are 267067 across the curve and 4.27 along it, with the eigenvector $(0.9811,\ -0.1935)^{\mathsf{T}}$ of the small one pointing along the curve; the first two terms of equation (1.18) give it directly, $(5.070^2\times4+11.11)/(5.070^2+1) = 4.27$. The Gauss–Newton covariance is the inverse,

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

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

The curvature term and the Laplace covariance. The second derivative of the physics row at the MAP state is $\partial^2 r_1/\partial g^2 = -507.0$; multiplied by $r_1 = 0.0052$, it adds -2.63 to the g–g entry, $257067 \to 257064$, the same value as $(r_2/\sigma_T)(Q/G_{\mathrm{ref}})\,\mathrm{e}^{-\hat{g}} = (-0.2598/0.5)\times5.0701$. Projected on the along-curve eigenvector $(0.9811,\ -0.1935)^{\mathsf{T}}$, the term is $-2.63\times0.1935^2 = -0.099$ against the along-curve information 4.27, a loss of 2.3 % in the one direction in which the information is small; it raises the variances by 2.4 % and the standard deviations by 1.2 %. The Laplace covariance (1.16) is

$$\mathbf{\Sigma} = \begin{bmatrix} 0.23096 & -0.045552 \\ -0.045552 & 0.0089881 \end{bmatrix}, \qquad \sqrt{\Sigma_{11}} = 0.4806\,\mathrm{K}, \quad \sqrt{\Sigma_{22}} = 0.09481. \tag{1.21}$$

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

$$\frac{\Sigma_{12}}{\sqrt{\Sigma_{11}\,\Sigma_{22}}} = \frac{-0.045552}{0.48058\times0.094806} = -0.9998.$$

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

$$\lambda_{\parallel} = \frac{5.070^2\times4 + 11.11 + 5.070\times(-0.260)/0.5}{5.070^2+1} = 4.168,$$

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

the square roots of the two diagonal entries of equation (1.21).

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

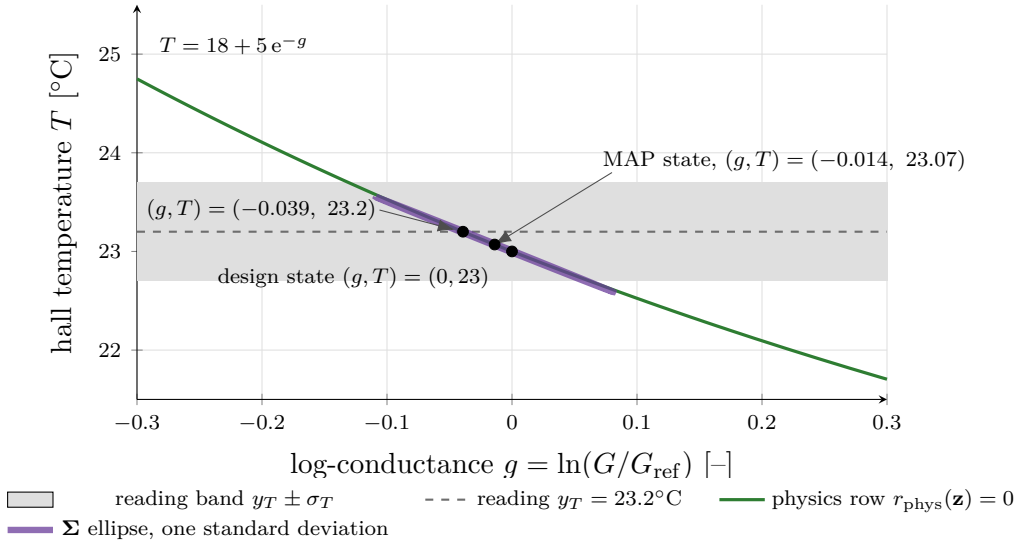


Figure 1.5: The state plane of the equipment hall with $T_c = 18$°C, $Q = 30\,\mathrm{kW}$ and $G_{\mathrm{ref}} = 6\,\mathrm{kW/K}$. Green: the zero set of the physics row, $T = 18 + 5\,\mathrm{e}^{-g}$ (equation (1.4)). Gray: the reading $y_T = 23.2$°C with its band of one standard deviation, $\sigma_T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

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

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

1.3 The answers (step 4)

The first four questions of Table 1.2 take two general forms, written in the book's notation: $p(\mathbf{z} \mid \mathbf{y})$ is the posterior density of the state, $f(\mathbf{z})$ a quantity derived from the state and f^* a threshold on it. The next-reading question, a design question about a row not yet attached, is derived in Section 1.6.1. Each is answered from the posterior of Section 1.2: its operator

question, the general form of its category, the form it takes for the hall, the quantities of the posterior it uses, and the hall's numbers. The questions name the limit of 24°C at the present load and the limit of 27.5°C at $Q_2 = 60\,\mathrm{kW}$.

1.3.1 The hall temperature and the conductance

Operator questions. What is the hall air temperature now? What is the effective conductance of the air path? The first asks for the temperature of the hall at the present load, given its one reading and the handbook value; the second for the conductance of its air path, the quantity the operations engineer compares with the design value.

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

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

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

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

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

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

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

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

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

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

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

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

The value beside each comes from the exact posterior: the moments (1.23) under the density (1.19), evaluated by one-dimensional quadrature over g, with T through $T(g) = T_c + (Q/G_{\mathrm{ref}})\, \mathrm{e}^{-g}$ and G through $G_{\mathrm{ref}}\, \mathrm{e}^{g}$, and the median of each, the value f^* at which equation (1.22) equals one half. The Laplace location, the entry of $\hat{\mathbf{z}}$, is the median of the Laplace normal, and it is set beside the exact median, like with like; the exact mode of g is $\hat{g}$ itself, and the exact mean is given as well, labeled as a mean. They differ by the skew of the posterior along the curve.

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

The hall temperature now. The hall temperature is the first entry of the state: $\hat{T} = 23.070\,^{\circ}\mathrm{C}$ with $\sqrt{\Sigma_{11}} = 0.481\,\mathrm{K}$ (Gauss–Newton alone, 0.475 K). The value beside it comes from the exact posterior: the median of T along the curve is 23.028 °C, 0.042 K below $\hat{T}$ (the mean is 23.030 °C), with the standard deviation 0.4813 K, 0.0007 K wider.

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

$$\hat{G} = 5.917\,\mathrm{kW/K}, \qquad \sigma_G = \sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{g}} = 5.917\,\sqrt{0.0089881} = 0.561\,\mathrm{kW/K},$$

a band of 9.5 % of its value, against a factor $\mathrm{e}^{0.3} = 1.35$ in the handbook before the reading. The values beside them come from the exact posterior: the conductance entry has the median -0.0056, 0.0083 above $\hat{g}$ (the mean is -0.0013), and the standard deviation 0.0967, 0.0019 wider; the conductance has the median 5.967 kW/K, 0.049 kW/K above $\hat{G}$ (the mean is 6.020 kW/K), and the standard deviation 0.591 kW/K, 0.030 kW/K or 5 % wider. The exact posterior extends further toward large conductances than the normal at $\hat{\mathbf{z}}$ does, the skew that the full-load question below meets in its tail.

1.3.2 The limits at the present load and at 60 kilowatts

Operator questions. Is the hall below 24 °C now? Can the hall take the second 30 kW tonight? The first asks about the hall as it runs now. The second is the work order's question, which for the operator reads

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

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

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

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

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

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

Specialized form. The two questions are equation (1.26) with the state $\mathbf{z} = (T, g)^{\mathsf{T}}$ of the hall, $p(\mathbf{z} \mid \mathbf{y})$ the posterior of step 3, the Laplace normal $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of equations (1.20) and (1.21), and the derived quantity $f(\mathbf{z})$ and the threshold f^{*} given for each below: $f(\mathbf{z}) = T$ with $f^{*} = 24\,^{\circ}\mathrm{C}$, and $f(\mathbf{z}) = T_2$ with $f^{*} = T_{\max} = 27.5\,^{\circ}\mathrm{C}$. Two choices enter here. The full-load condition is taken to share the conductance of the present condition, with the load and the supply air pinned at both: this is what makes a condition the hall has never run predictable from a reading at 30 kW, and it fails if the conductance changes with the load, for instance

through an airflow that rises with it (Section 1.1.4), or if the unit saturates and can no longer hold T_c at 60 kW (Section 1.6.5), and no check of this chapter can test either because nothing has been observed at 60 kW. The full-load answer is the steady temperature at 60 kW: this is what lets the event be written on the conductance alone, and it says nothing about the period just after the load rises, while the hall's stored heat still matters (Section 1.6.5).

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

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

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

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

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

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

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

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

with the thresholds

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

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

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

Figure 1.6 draws the equivalence (1.29) on the forward relation (1.2). The horizontal axis is the conductance G in kW/K and the vertical axis the hall temperature in °C. The solid blue curve is the hall at the load of 30 kW, the dashed green curve the same hall at 60 kW; the dot on the blue curve is the forward calculation of Section 1.1.4, $T = 23°\mathrm{C}$ at $G = 6\,\mathrm{kW/K}$, and the dot above it the same conductance at 60 kW, 28°C. The red line is the limit of 27.5°C, and the

dotted vertical line marks the threshold G^* of equation (1.30), where the dashed curve crosses the limit. The dashed green curve lies below the red line exactly to the right of G^*: because the curve is strictly decreasing, the event "below the limit" on the vertical axis is the event "above G^*" on the horizontal axis.

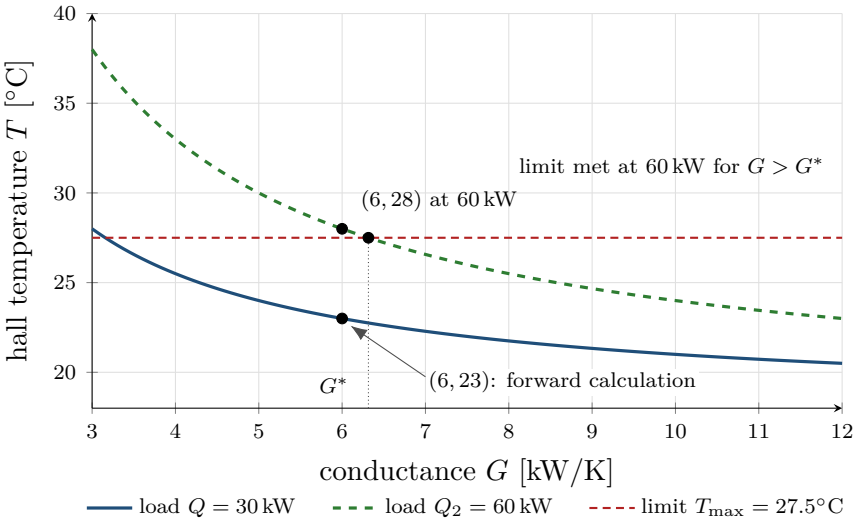


Figure 1.6: The forward relation $T = T_c + Q/G$ (equation (1.2)) with $T_c = 18$°C, for the loads $Q = 30$ kW and $Q_2 = 60$ kW. Dots: the forward calculation $T = 23$°C at $G = 6$ kW/K; the same conductance at 60 kW, 28°C; and the threshold $G^* = Q_2/(T_{\max} - T_c)$ of equation (1.30) at which the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

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

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

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

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

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

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

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

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

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

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

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

the temperature at 60 kW stays below 27.5°C if and only if the temperature at 30 kW is below 22.75°C; on the curve of Figure 1.5, the states with $g > g^*$ are those below 22.75°C. By equation (1.32),

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

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

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

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

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

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

1.4 The checks (step 5)

The fit rests on two approximations and one assumption: a physics row with a finite tolerance in place of an exact law, a normal distribution in place of the posterior along a curved ridge,

and a reading and a handbook value that are compatible with each other. Each is set against a full calculation, and the weight of the handbook value is measured by removing it.

1.4.1 The exact-physics limit

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

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

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

1.4.2 The Laplace posterior against the exact posterior

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

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

Quantity	Unit	Laplace, at $\hat{\mathbf{z}}$	Exact, along the curve
Median of T	°C	23.070	23.028
Mean of T	°C	–	23.030
Standard deviation of T	K	0.4806	0.4813
$\Pr(T < 24\text{°C} \mid \mathbf{y})$	–	0.973	0.977
Median of g	–	−0.0139	−0.0056
Mean of g	–	–	−0.0013
Standard deviation of g	–	0.0948	0.0967
Median of G	kW/K	5.917	5.967
Mean of G	kW/K	–	6.020
Standard deviation of G	kW/K	0.561	0.591
$\Pr(T_2 < 27.5\text{°C} \mid \mathbf{y})$ at 60 kW	–	0.246	0.281

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

its mode, median and mean differ: the mode of g is $\hat{g}$, the median lies 0.0083 and the mean 0.013 above it, the median conductance lies 0.049 kW/K and the mean 0.10 kW/K above $\hat{G}$, and the exact standard deviation of G is 0.591 kW/K, 5 % wider than the Laplace 0.561 kW/K. The skew matters most for the full-load question, whose event lies in the tail of g toward large conductances, the side on which the exact posterior extends further: the Laplace tail 0.246 is low by 0.036 against the exact 0.281. Further into the tail the understatement grows: in Table 1.5 the Laplace tail on g is smaller than the exact one by a factor of 2.5 to 4 in the two smallest rows, 0.0062 against 0.0155 and 0.0012 against 0.0049. Written on the first entry instead, as $T < T^* = 22.75°\mathrm{C}$, the full-load event has the Laplace tail 0.253, because the normal follows the straight tangent rather than the curve; the exact value is 0.281 on either entry. The exact values rest on the base of equation (1.19), flat in g along the curve; with a base flat in T along the curve instead, the full-load value is 0.249, a reminder that an exact tail depends on the base measure as well as on the rows. The conductance band and the full-load probability are therefore slightly narrow and slightly low, and the report says so.

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

1.4.3 The residual check

Three rows act on two entries, and with the physics exact the reading and the handbook both inform the one direction along the curve: one row is redundant, and the rows at the MAP state measure how well the two agree. They are $r_1 = 0.005$, $r_2 = -0.260$ and $r_3 = -0.790$ (Table 1.4), and their squares sum to $2\mathcal{L}(\hat{\mathbf{z}}) = 0.692$. When the model and the standard deviations of the rows are right, that sum follows a chi-square distribution with as many degrees of freedom as there are rows less entries, here one; Chapter 9 derives the result, the chi-square test of bad-data detection. A sum of 0.692 or more then has probability 0.41, while a sum above 3.84 would have probability below 0.05 and would say that the reading or the handbook value is wrong. The two agree, which is a first, coarse look at the question of a biased sensor.

1.4.4 The weight of the handbook prior

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

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

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

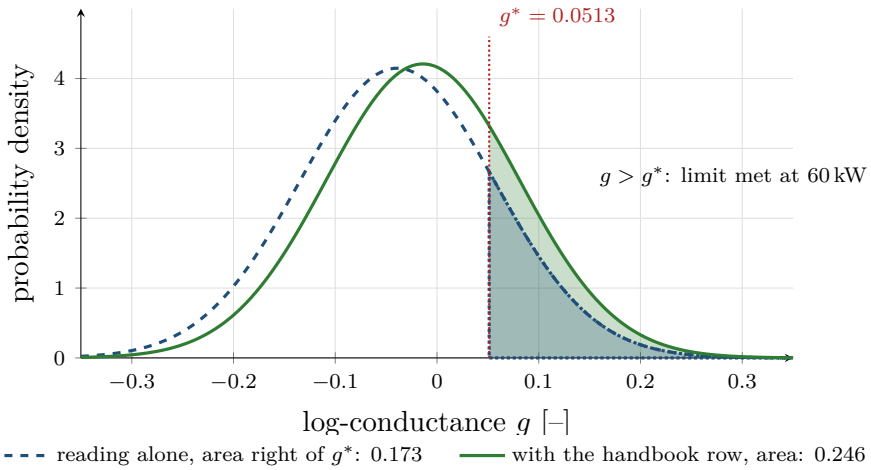


Figure 1.7: The weight of the handbook prior in the full-load answer: Laplace marginals of the conductance entry, $\mathcal{N}(-0.0392,\ 0.0962^2)$ under the reading alone (dashed) and $\mathcal{N}(-0.0139,\ 0.0948^2)$ with the handbook row (solid), densities $\varphi((g-\hat{g})/s)/s$ with φ the standard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

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

1.5 The report (step 6)

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

1.5.1 The answers and their bands

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

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

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

Question	Answer (state route)	Exact	Band
What is the hall air temperature now?	$\hat{T} = 23.07°\text{C}$	median 23.03°C	±0.48 K (0.481 K exact)
What is the effective conductance of the air path?	$\hat{G} = 5.92\,\text{kW/K}$	median 5.97 kW/K (mean 6.02 kW/K)	±0.56 kW/K (0.59 kW/K exact)
Is the hall below 24°C now?	$\Pr(T < 24°\text{C} \mid \mathbf{y}) =$ 0.973, about 0.97	0.977; gap 0.004	–
Can the hall take the second 30 kW tonight?	$\Pr(T_2 < 27.5°\text{C} \mid \mathbf{y}) =$ 0.246, about 0.25	0.281; gap 0.036	without the handbook value 0.173 (0.212 exact)

1.5.2 The full-load decision

The work-order question has the answer of the last row: at 60 kW the hall stays below 27.5°C with probability of about one in four, 0.246 by the Laplace route and 0.281 exactly: on the present cooling the record does not support taking the second 30 kW tonight. The answer assumes that the unit still holds the supply air at 18°C at 60 kW; if it cannot, T_c rises with the load and the hall is warmer than this answer says (Section 1.6.5). The reason is read from the equivalent threshold: the limit is met only if the conductance exceeds $G^* = 6.316\,\text{kW/K}$, and the posterior places the conductance at $5.92 \pm 0.56\,\text{kW/K}$; in the conductance entry the threshold $g^* = 0.0513$ lies 0.69 standard deviations above $\hat{g} = -0.0139$. The answer leans on the handbook value: without it the probability would be 0.173 by the Laplace route and 0.212 exactly. A reading cannot raise the conductance; it can only show where it lies, and with the estimate below the threshold, further evidence is more likely to confirm the shortfall than to remove it. Running the hall at 60 kW under this limit calls for more cooling: more airflow (a larger conductance) or a colder supply air, if the unit can deliver it; not more readings.

1.6 The next reading (step 7)

Whether the bands of Table 1.8 are narrow enough is the operator's decision. If they are not, the next-reading question of Table 1.2 asks which reading to request from the field technician to narrow them most, and a reading not yet taken is a row not yet attached.

1.6.1 The preposterior covariance

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

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

with $\boldsymbol{\Sigma}$ of equation (1.21), curvature term included. The value the reading will return does not appear in it. Chapter 2 derives $\boldsymbol{\Sigma}_a$ as the preposterior covariance, and Chapter 18 ranks candidate readings by it for quantities derived from the state. The band on the temperature after the reading is $\sqrt{\Sigma_{a,11}}$, and the band on the conductance, the quantity on which the full-load answer rests, is $\sigma_{G,a} = \sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_a\mathbf{g}}$ with the gradient of equation (1.24).

1.6.2 The candidate readings

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

- a second hall sensor of the same accuracy, with an error independent of the first: the row $(T - y'_T)/\sigma_T$, with the Jacobian row $\mathbf{J}_a = (2,\ 0)$, the same as the present reading row;
- a reference sensor of 0.2 K in the hall: $\mathbf{J}_a = (5,\ 0)$;
- a reading of the supply-air temperature, or a load meter on the servers: T_c and Q are pinned, not entries of the state, so the row has no entry in either column, $\mathbf{J}_a = (0,\ 0)$;
- a coil heat-removal check (the airflow through the unit times the temperature drop across its coil), $\sigma_q = 1$ kW on the heat removed $G\,(T - T_c)$: $\mathbf{J}_a = \big(\hat{G},\ \hat{G}\,(\hat{T} - T_c)\big)/\sigma_q = (5.917,\ 30.00)$.

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

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

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

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

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

1.6.3 The recommendation

The recommendation, and the answer to the next-reading question, is a second sensor in the hall: it removes 0.134 K of the 0.481 K on the temperature and 0.156 kW/K of the 0.561 kW/K

on the conductance, 28 % of each, and a 0.2 K reference sensor, if one is available, removes 62 %. A supply-air reading, a load meter and a coil heat-removal check are not worth requesting for these answers. Two conditions apply. The second sensor's error must be independent of the first: a second reading from the same sensor at the same condition would share any bias of the first, and a shared error is a discrepancy that repeats (Section 1.1.4), which a repeated reading does not reduce. And the gain is computed at $\hat{\mathbf{z}}$ under the Laplace approximation, which the checks found adequate for T and slightly narrow for G. Once the reading is taken, its row joins the three of equation (1.10), and the method returns to step 3 with four rows. For the full-load decision the narrower conductance band sharpens the answer but does not move the threshold: the decision of Section 1.5.2 turns on more cooling, more airflow or a colder supply air if the unit can deliver it, not on more readings.

1.6.4 What the analysis showed

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

1.6.5 Limits of the model

The posterior is exact only under the model, the steady-state graph of one hall with one path to a supply air that the discharge controller holds, and that model can be wrong before any inference is run. The supply air is pinned only while the controller holds it: a cooling unit driven to full output can no longer hold T_c, which then rises with the load and would have to become a free quantity of the graph, including at the 60 kW of the scenario node, which the full-load answer assumes the unit can still hold at 18°C. The record itself is illustrative: the 23.2°C reading and the handbook value are constructed for the example, not field data, and no comparison with a true state is claimed.

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

The first two mistakes confuse the roles of Table 1.3. A sensor reporting 23.2°C does not establish that the true temperature is exactly 23.2°C: the observation has an uncertainty, and the physics row of Section 1.1.4 showed that 23.2°C and 23°C imply conductances of 5.77 and 6 kW/K, a difference of 4 %. The conductance is a physical property of the path; the reading is evidence about the temperature, which is connected to the conductance through the model. The fourth mistake mixes two kinds of uncertainty that require different treatments. Uncertainty about the numerical value of G within the model of Figure 1.2 is represented by the posterior

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

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

of g. Uncertainty about whether the hall has one or two effective air paths is uncertainty about the model M itself, and it is represented by comparing the evidence of the two drawings. The fifth mistake is the converse of the first. A graph provides relations; the posterior also requires evidence, and some entries may remain poorly identified even when the model is perfectly specified. In the two-path drawing, for example, a reading of the hall temperature determines only the total conductance of the two paths and leaves their split to the prior: the Jacobian of the physics and reading rows has a null vector along the split, and only a prior row adds information in that direction.

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

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

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

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

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

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

The steady-state relation of equation (1.2) is the limit $t \to \infty$. Suppose (illustratively) that at $t = 0$ the cooling unit loses half its airflow (one of its two fans trips), so that the conductance falls from 6 to 3 kW/K while the hall is at 23°C under 30 kW. The new steady temperature is $18 + 30/3 = 28$°C, and equation (1.34) gives $T(t) = 28 - 5\,\mathrm{e}^{-t/\tau}$: after one time constant the hall is at $28 - 5\,\mathrm{e}^{-1} = 26.16$°C. A steady-state model would place it at 28°C at once. It is therefore inadequate for predicting temperatures immediately after a fan trip, where the time constant

decides whether the operators have minutes or seconds; the graph for the airflow question, the partial loss of cooling, needs the storage term and C_{th} as a further entry or a pinned input.

The chapter has built the representation on which every question will be asked. The next chapter, Chapter 2, turns to the questions themselves: how an operator's question stated in the control room's language is recognized as estimation, prediction, comparison, diagnosis or design, and written as a precise probability or expectation over the state of such a graph. From Part II the hall grows: the supply air becomes the cold aisle that the cooling coils produce from chilled water.

1.7 What to retain

Definition 1.1 (Plant graph)**.** A plant graph represents a physical system as free and pinned nodes, paths and sources, with the physical relations that must hold among them written as residual rows; its free quantities, physical states and parameters alike, form the state vector $\mathbf{z}$; readings, records and prior knowledge are attached as further rows; and every operator's question is then a probability or an expectation over $\mathbf{z}$, conditional on the evidence $\mathbf{y}$, computed from the posterior $p(\mathbf{z} \mid \mathbf{y})$ that the rows define.

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

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

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

1.8 Exercises

Exercise 1.1. (a) Which statement best describes the role of the hall sensor: it defines the physical law; it reports the exact true state; or it supplies evidence about an uncertain state? (The third: the sensor provides noisy evidence, while the physical model relates the state to the inputs

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

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

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

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

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

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

its first row gives $\Sigma_{11} = 1/c^2 = \sigma_T^2 = 0.25\,\mathrm{K}^2$, and its second $\Sigma_{22} = (\sigma_p^2 + \sigma_T^2)/(5\,\mathrm{e}^{-\hat{g}})^2$, so $\sqrt{\Sigma_{22}} = \sqrt{0.2501}/5.2 = 0.0962$.) (c) Explain why no curvature term is needed here, and how the result gives the reading-alone probability 0.173 of Section 1.4.4. (Both rows are zero at $\hat{\mathbf{z}}$, so $\sum_k r_k\, \partial^2 r_k / \partial \mathbf{z}\, \partial \mathbf{z}^{\mathsf{T}}$ vanishes; the entry g then has the Laplace marginal $\mathcal{N}(-0.0392,\ 0.0962^2)$, and its tail above $g^* = 0.0513$ is $\Phi((-0.0392 - 0.0513)/0.0962) = \Phi(-0.941) = 0.173$.)

Chapter 2

The Language of Operational Questions

2.1 Operator questions and their mathematical form

Chapter 1 drew the equipment hall as a plant graph: one hall node whose air temperature balances a heat load against one heat-removal path to the supply air, which the cooling unit's discharge controller holds at a pinned 18°C, with the conductance of that path unknown and a temperature sensor attached to the hall. The graph was built before any question was asked of it. At 19:00, the start of the night shift, the hall tile on the screen reads 23.2°C, 0.8 K inside the 24°C limit, and a work order asks for the load to go from 30 to 60 kW tonight; on the dashboard alone the hall is green and can take the load. With the hall's alarm log on the next screen, the operator asks several questions. They are not the same kind of question even though they concern the same plant: what the conductance of the air path is, whether the hall will exceed its limit at the new load, whether the hall sensor is reading high, whether a second sensor is worth installing, whether the alarm will stay quiet through the night, and whether to start the standby unit before the night peak. A dashboard gives each of them a number, and an inference algorithm can return a mathematically correct result to a question nobody asked; the step that prevents both is the translation of the operator's question into the mathematical quantity that answers it. The question of this chapter is

how an operator's question about a running plant is translated into a mathematical expression that probabilistic inference can evaluate, and how that expression determines what must be computed.

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

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

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

Key idea

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

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

2.1.1 The plant graph of the equipment hall

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

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

where Q is the heat load in kW, G the conductance in kW/K, T the hall air temperature and T_c the supply-air temperature, both in °C. The supply air is pinned because the unit's

controller holds it; the pin fails when the unit saturates, at a load its coil cannot carry, which the questions of this chapter do not reach. Solved for the hall temperature, it gives the relation used throughout the chapter,

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

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

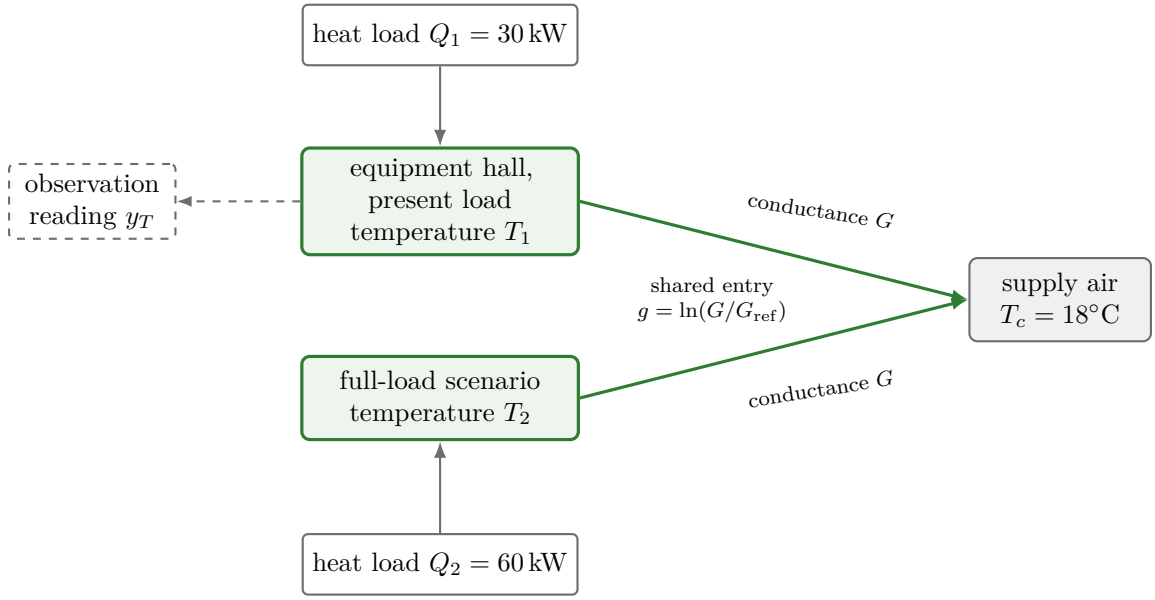


Figure 2.1: The plant graph of the equipment hall with the full-load scenario node. Free: the present-load temperature T_1 and the full-load temperature T_2 (green nodes) and the log-conductance g shared by the two heat-removal paths (solid green edges). Pinned: the supply-air temperature $T_c = 18°\mathrm{C}$,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

2.1.2 Four questions about one plant

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

Question A estimates an unknown physical quantity. Question B asks for the probability of an event at a load not yet applied. Question C distinguishes between two explanations of the same readings, one in which the sensor reads the hall temperature with zero-mean error and one in which it carries a bias. Question D predicts how a reading that does not yet exist would change an operating decision. All four concern the same graph, and all four are computed from the same physical model; none of them is answered by the same calculation as another.

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

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

Operator question	Calculation	Answering quantity
A. What is the effective conductance of the hall's air path?	estimation of an unknown physical quantity	$p(G \mid \mathbf{y})$
B. Can the hall take the second 30 kW tonight and stay below 27.5°C?	probability of an event	$\Pr(T_2 < 27.5^\circ\mathrm{C} \mid \mathbf{y})$
C. Is the hall sensor reading high?	comparison of competing explanations	$\Pr(M_{\text{bias}} \mid \mathbf{y})$ against $\Pr(M_{\text{unbiased}} \mid \mathbf{y})$
D. Is a second hall sensor worth installing?	prediction of what a future reading would change	the preposterior, and the expected loss after the reading

2.2 The grammar of probability statements

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

2.2.1 Events

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

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

where N_a is the number of activations of the hall's high-temperature alarm during a shift. A probability assigns to an event a number between zero and one. The statement

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

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

2.2.2 Conditional probability and the evidence

The probability of event A given the information B is written

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

and read "the probability of A, given B". In operations the information on the right of the bar is the evidence: the telemetry of the present window (temperatures, loads, meter readings), the alarm log and the event log, readings taken by the field technician, the set points and run states of the equipment, nameplate and handbook values, and what commissioning established. The book collects all of it in the bold symbol $\mathbf{y}$, and keeps a plain y with a subscript, such as y_T, for a single reading. The statement

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

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

2.2.3 The role of the model

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

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

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

2.2.4 Unknown quantities that are not events

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

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

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

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

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

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

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

2.2.5 A reusable notation

The questions of this chapter, and of the book, are written with the symbols of Table 2.2. The first column is the symbol, the second its meaning. The state $\mathbf{z}$ holds every free quantity of the plant graph at once, the physical quantities such as the log-conductance as well as the temperatures; there is no separate parameter vector. The action a stands for an operating decision such as starting a standby unit, changing a set point, taking a unit out for maintenance or requesting a field reading; the symbol $\tilde{y}_a$ stands for the outcome of a candidate reading a that has not yet been taken, and y_{new} for a future outcome of the plant itself, such as the time of the next alarm.

With these symbols, questions as different as the conductance of an air path, a quiet night shift and the choice of a sensor are written in one mathematical language, and the differences between them become visible as differences in the expression, not in the words.

Table 2.2: The reusable notation of probabilistic questions.

Symbol	Meaning
$\mathbf{z}$	state vector of the plant graph: every unknown quantity, physical parameters included
$\mathbf{y}$	observed evidence: readings, counts, records, priors
M, M_j	assumed model or system structure; one of several competing models
$p(\mathbf{y} \mid M)$	evidence of model M: the probability of the evidence with the model's state integrated out
B_{AB}	Bayes factor of model M_A against M_B, $p(\mathbf{y} \mid M_A)/p(\mathbf{y} \mid M_B)$; B_{10} is M_1 against M_0
A	an event
$f(\mathbf{z})$, f^*	a quantity derived from the state; a requirement or limit on it
κ	assurance level: the posterior probability a requirement must reach
y_{new}	a future observation that stays random when the state is known
a	an operating action or decision, or a candidate reading
$\tilde{y}_a$	the outcome of candidate reading a, not yet observed
u, $U(a)$	a quantity of the posterior after $\tilde{y}_a$; its expectation over the outcomes, which ranks the candidates
$\mathbf{\Sigma}_a$	covariance of the state after candidate measurement a (the preposterior)
$L(a, \mathbf{z})$	loss or cost of action a when the state is $\mathbf{z}$: its consequence
$\mathcal{R}(a)$	expected loss of action a over the posterior, $\mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]$
c_a, C_e	the fixed cost of action a (a cost, not the decision level c of an alarm); the consequence of event e
$\mathbf{g}_a$, $\mathbf{H}_a$	gradient and curvature of $L(a, \mathbf{z})$ in the state at $\hat{\mathbf{z}}$
$f_t(\mathbf{z})$, $\mathbf{g}(t)$	a quantity at time t carried by a state indexed by time; its gradient
α, γ	the level and the trend entries of such a state: the quantity at the record's center $\bar{t}$ and its change per unit time (from Chapter 19 on, α is a false-alarm probability)
t^*	crossing time: the time at which a rising f_t reaches a level f^*

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

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

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

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

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

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

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

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

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

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

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

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

2.3.2 Future events and the posterior predictive

Some events remain random even when the state is known exactly. An alarm whose rate is known still fires at a random time; knowing the rate fixes only the probability that a shift passes without it. For such a future event A_{future}, the model supplies the conditional probability $\Pr(A_{\text{future}} \mid \mathbf{z}, M)$, and the question asks for its probability given the evidence. The derivation uses three rules. By the law of total probability, the probability of the event is the integral of its joint probability with the state:

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

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

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

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

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

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

2.3.3 The general forms of the five categories

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

2.3.4 The general form of a trend

Some questions concern a quantity that moves in time: whether an alarm rate or a unit's derate is rising, when it will cross the limit of the operating envelope, what it will be at the end of the shift. By its integrand each is an estimation or a prediction under a threshold; what sets it apart is that the quantity is indexed by time, so that the state must carry how the quantity changes as well as where it stands. Write the quantity, on the scale on which the state carries it (a log-rate, a log-derate, a temperature), as $f_t(\mathbf{z})$, its value at time t. In the simplest drawing it is a straight line in time about the center $\bar{t}$ of the record, with a level entry α, the quantity at $\bar{t}$, and a *trend entry* γ, its change per unit time:

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

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

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

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

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

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

The first line is the sign of the trend, the probability that the quantity is rising: the threshold form on the trend entry. The second is the distribution function of the crossing time $t^*(\mathbf{z})$, the time at which the quantity reaches the level f^* from below. A rising path, $\gamma > 0$, has crossed by a time s if and only if it is at or above f^* at s, whether s lies before or after $\bar{t}$, and a path with $\gamma \le 0$ never crosses, a mass $\Pr(\gamma \le 0 \mid \mathbf{y})$ at $t^* = \infty$: the joint indicator is the crossing, the first passage of a rising path. The threshold probability $\Pr(f_s(\mathbf{z}) \ge f^* \mid \mathbf{y})$ alone is not the crossing, because it also counts the states with $\gamma \le 0$ that stand at or above the level at s, paths that started above it and have never crossed it; the two agree only when the posterior puts no such state there. The third is the forecast, the estimation form at a time beyond the last window of the record. Nothing of a particular plant appears in the three lines; a plant enters through the scale of f_t, the rows of its windows and the level f^*.

Under the Laplace posterior $\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ the quantity at any time is linear in the state, so the sign and the forecast are exact normal tails in two numbers, the center $\mathbf{g}(t)^{\mathsf{T}}\hat{\mathbf{z}}$ and the variance $\mathbf{g}(t)^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{g}(t)$, and the crossing is an exact bivariate normal probability of the pair (γ, f_s), whose correlation $\varrho(s)$ is read from the same $\boldsymbol{\Sigma}$:

$$\begin{gathered}
\Pr(\gamma > 0 \mid \mathbf{y}) \approx \Phi\!\left(\frac{\hat{\gamma}}{\sqrt{\Sigma_{\gamma\gamma}}}\right), \qquad \Pr(t^* \le s \mid \mathbf{y}) \approx \Phi_2\!\left(\frac{\hat{\gamma}}{\sqrt{\Sigma_{\gamma\gamma}}},\ \frac{\mathbf{g}(s)^{\mathsf{T}}\hat{\mathbf{z}} - f^*}{\sqrt{\mathbf{g}(s)^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{g}(s)}};\ \varrho(s)\right), \\
\mathbf{g}(t)^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{g}(t) = \Sigma_{\alpha\alpha} + 2\,(t-\bar{t})\,\Sigma_{\alpha\gamma} + (t-\bar{t})^2\,\Sigma_{\gamma\gamma}, \qquad \varrho(s) = \frac{\Sigma_{\alpha\gamma} + (s-\bar{t})\,\Sigma_{\gamma\gamma}}{\sqrt{\Sigma_{\gamma\gamma}\,\mathbf{g}(s)^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{g}(s)}},
\end{gathered} \tag{2.22}$$

where $\Phi_2(u_1, u_2; \varrho)$ is the distribution function of two standard normals of correlation ϱ, the probability that the first lies below u_1 and the second below u_2. It is one integral, $\Phi_2(u_1, u_2; \varrho) = \int_{-\infty}^{u_1} \varphi(x)\, \Phi\big((u_2 - \varrho x)/\sqrt{1-\varrho^2}\big)\, \mathrm{d}x$, and it lies below the single tail $\Phi(u_2)$, the *threshold line*, by the mass of the states that stand beyond the level without a rising trend. The sign is the single tail with the gradient $(0,\ 1)^{\mathsf{T}}$, which reads the trend entry alone, so that its variance is $\Sigma_{\gamma\gamma}$; the forecast is the same with the time asked. The variance of the quantity grows with the square of the distance $t - \bar{t}$, the lever arm: a forecast pays for the uncertainty of the trend with that square. The crossing time itself, $t^* = \bar{t} + (f^* - \alpha)/\gamma$ for $\gamma > 0$, is a ratio with the trend entry in its denominator. When the posterior of γ has a positive density at zero, as a Gaussian posterior always has, t^* has no posterior mean and no variance, so it is reported by its median, the date at which the second line is one half, which lies at the plug-in date $t^*(\hat{\mathbf{z}})$ when no part of the posterior stands beyond the level without a rising trend, by the band whose ends solve the second line at 0.025 and 0.975, and by the mass that never crosses; the delta band $t^*(\hat{\mathbf{z}}) \pm 1.96\sqrt{\mathbf{g}_*^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{g}_*}$, with $\mathbf{g}_* = \partial t^*/\partial \mathbf{z}$ at $\hat{\mathbf{z}}$, is a finite number standing in for an infinite one. For the alarm log of a sister hall's older unit, its discharge-temperature alarm, counted in six periods of forty days and rising from 5 to 14 activations a period (Exercise 2.5), the band of the crossing of the alarm rate's target of 0.5 activations a day runs from 5.90 to 28.3 periods by inversion, with period j centered at $t = j$ and the log ending at $t = 6.5$, against a delta band from 4.44 to 11.45 periods that is wrong at both ends. A flat or falling line on that log stands at or above the target at those dates with a mass of at most 5.1×10^{-11}, so the first passage

and the threshold line give the same median and band to every digit printed, and the band is read from the line. Chapter 11 reads the same form on the derate of a coil across telemetry windows, where the mass is not always so small.

2.3.5 The general form of an expected consequence

Other questions ask what to do: whether to start the standby unit before the night peak, whether to raise a set point, which unit to take out for maintenance. Each compares actions a from a set of candidates by what they are expected to cost. Let $L(a, \mathbf{z})$ be the loss of taking action a when the state is $\mathbf{z}$, its consequence in whatever unit the decision is taken in: reais, minutes above the limit, energy. It is a derived quantity of the state, one for each action. The general form of an expected consequence is its posterior expectation, the expected loss $\mathcal{R}(a)$ of the action, and the action that minimizes it. When the action costs c_a for certain and leaves events e that occur with probability $\Pr(e \mid \mathbf{z}, a)$ and have the consequence C_e each, the loss given the state is the cost of the action plus the consequence of each event weighted by its probability, and by the predictive (2.8) the posterior expectation turns each conditional probability into its predictive one:

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

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

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

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

The loss at the most probable state, $L(a,\hat{\mathbf{z}})$, misses the curvature term. The curvature is taken in the coordinates of the state, and its sign there, not in a physical quantity, sets the direction of the second-order term. The alarm log of Section 2.4.4 shows the difference over the next 100 h of operation, about eight night shifts. The probability of no alarm, $\mathrm{e}^{-\lambda\tau}$, is convex in the rate λ, so by Jensen's inequality the predictive 0.7725 lies above the plug-in 0.7659; but the state carries the log-rate ω, in which the probability is concave, with curvature -0.150 at $\hat{\omega}$, and the second-order reading with $\sigma_\omega^2 = 0.25$ moves the plug-in down to 0.747, close to the average 0.748 over the Laplace normal of the log-rate and away from the exact 0.7725. The term is right about the Laplace posterior and wrong about the alarm, because the error lies in the Laplace posterior itself; where the second-order term and an exact route move the plug-in in opposite directions, the value read from the state is stated with the exact one beside it, and the gap is what the Laplace posterior costs. The second line of equation (2.24) asks how firm a ranking is: the probability that action a costs less than a' at the true state. The variance of the difference is the sum of the two variances minus $2\,\mathbf{g}_a^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{g}_{a'}$, so the band of the comparison is narrower than the two bands taken apart exactly when $\mathbf{g}_a^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{g}_{a'} > 0$, when the two losses move

together across the posterior. Section 2.5.4 applies the form to the choice between starting the standby unit before the night peak and waiting, and that section is where this book applies it.

2.3.6 The five computations beneath the forms

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

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

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

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

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

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

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

It is the normalizing constant of the posterior within model M_j, and it is the only place where the comparison between models enters: the posterior inside either model, however sharp, does not say which model is true. Under the Laplace approximation it is read from the model's own fit, as the volume of its posterior times the height of the integrand at its MAP state, $p(\mathbf{y} \mid M_j) \approx (2\pi)^{d_j/2} \det(\boldsymbol{\Sigma}_j)^{1/2}\, p(\mathbf{y} \mid \hat{\mathbf{z}}_j, M_j)\, p(\hat{\mathbf{z}}_j \mid M_j)$ with d_j the number of entries of its state, exact when the objective is quadratic (Chapter 4). The comparison and diagnosis forms are a ratio and a ranking of it.

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

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

and its curvature, the information, is the curvature of the stacked Jacobian $[\mathbf{J};\ \mathbf{J}_a]$, the sum $\mathbf{J}^\mathsf{T}\mathbf{J} + \mathbf{J}_a^\mathsf{T}\mathbf{J}_a$: the information of independent rows adds. With $\mathbf{J}^\mathsf{T}\mathbf{J} = \boldsymbol{\Sigma}^{-1}$, the covariance after the measurement under the linear-Gaussian approximation is the inverse of that sum, which is equation (2.17). When the current rows include exact constraints, $\boldsymbol{\Sigma}$ is singular in the limit $\sigma_p \to 0$ and $\boldsymbol{\Sigma}^{-1}$ does not exist there. The limit is then taken through the equivalent update form, which uses $\boldsymbol{\Sigma}$ alone,

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

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

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

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

In this book the posterior is used in every chapter, and operations is mostly posterior and trend; the trend in Chapter 11, and in Chapter 15 of Volume II; the evidence in Chapters 4, 7 and 10, and in Chapter 20 of Volume II; the preposterior at step 7 of Chapters 1 and 4 to 11, and in Chapter 18 of Volume II; and the expected consequence in Section 2.5.4, where this book applies it.

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

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

2.3.7 The finer question types inside the five categories

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

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

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

Question type	Example	Mathematical target	Category	Computation
Estimation	What is the hall's alarm rate?	$p(\lambda \mid \mathbf{y})$	estimation	posterior
Threshold probability	Will the hall stay within its limit?	$\Pr(T \le T_{\max} \mid \mathbf{y})$	prediction under a threshold	posterior
Joint and conditional events	Will both halls stay quiet through the night?	$\Pr(\bar{A}_A \cap \bar{A}_B \mid \mathbf{y})$	prediction under a threshold	posterior
Prediction	Will the night shift pass without an alarm?	$\Pr(T > 12\,\mathrm{h} \mid \mathbf{y})$	prediction under a threshold	posterior
Trend and crossing	Is the coil's derate growing, and when will it pass its limit?	$\Pr(\gamma > 0 \mid \mathbf{y})$, $\Pr(t^* \le s \mid \mathbf{y})$	prediction under a threshold	trend
Derived quantities	How much heat is the unit rejecting?	$p(Q_u(\mathbf{z}) \mid \mathbf{y})$	estimation	posterior
Model comparison	Is the load meter off, or the load different?	$\Pr(M_j \mid \mathbf{y})$	hypothesis comparison	evidence
Diagnosis	Which unit is derated?	$\Pr(F_j \mid \mathbf{y})$	diagnosis	evidence
Risk and expectation	What does an excursion tonight cost on average?	$\mathbb{E}[L \mid \mathbf{y}]$	estimation	expected consequence
Importance	Which entry moves the hall temperature most?	$\partial T/\partial z_i$, an entry of the gradient $\mathbf{g}$	estimation	posterior
Information and experiment design	Which reading should the field technician take?	expected information gain or value of information	design	preposterior
Decision optimization	Start the standby unit now or wait?	$\arg\min_a \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]$	design	expected consequence, with the preposterior when a reading comes first

2.3.8 The seventeen questions of the shift

The book is organized around seventeen questions that come up in an operator's shift, grouped by the hour they come up in: where the plant is right now, whether something is wrong, what happens if the operator acts, how far the model can be trusted today, and where the alarms should sit. Table 2.5 lists them in the operator's words. Its first column is the number of the question; the second the question as the operator puts it; the third the general form it takes, by equation number; the fourth its category; the fifth the computation beneath it; and the last the chapter that leads with it. Chapters 5 to 11 put the questions to the hall of this chapter grown to two cooling units and three racks. Chapters 12 to 20 put them to two halls on one chilled-water loop, followed through a twelve-hour night shift.

Table 2.5: The seventeen operator questions of the book, with the general form, the category and the computation of each and the chapter that leads with it.

No.	Operator question	General form	Category	Computation	Chapter
Where is the plant right now					
1	**Virtual sensor.** What is the inlet temperature at the rack with no sensor?	(2.9), (2.10)	estimation	posterior	5
2	**Hidden flow.** How much heat is this unit actually rejecting?	(2.9)	estimation	posterior	6
3	**Disagreement between meters.** Two meters should sum to the third and do not; which one is off?	(2.15)	diagnosis	evidence	7
4	**Am I inside the envelope.** What is the probability the hot aisle is already above 27°C somewhere I cannot see?	(2.12), a joint tail	prediction under a threshold	posterior	8
Is something wrong					
5	**Anomaly.** Is this residual noise or a change?	(2.13), (2.12)	prediction under a threshold	posterior	9
6	**Which component.** If something changed, where?	(2.15)	diagnosis	evidence	10
7	**How big.** Given it is unit B, how far is it from its rating?	(2.9)	estimation	posterior	10
8	**Is it getting worse.**	(2.21)	prediction under a threshold	trend	11
What happens if I act					
9	**Set-point change.** If I raise the supply temperature 2 K, what is the probability the inlets stay under the limit?	(2.12), at a scenario node	prediction under a threshold	posterior	12
10	**Taking a unit out.** Can I take unit A down for maintenance now without a limit violation?	(2.12), at a scenario node	prediction under a threshold	posterior	13
11	**Load migration.** If I move 40 kW from hall A to hall B, where does the margin go?	(2.9) at two scenario nodes, then a joint tail	estimation	posterior	14
12	**Do-nothing.** If I do nothing, what is the probability of a violation before the morning shift?	(2.21)	prediction under a threshold	trend	15
How much do I trust the model today					
13	**Fit quality.** Is the model still describing the plant?	(2.13)	prediction under a threshold	posterior	16
14	**Which reading is carrying the answer.** If this one sensor fails, how wide does my band on the hot-aisle temperature get?	(2.10), with a row removed	estimation	posterior	17

No.	Operator question	General form	Category	Computation	Chapter
15	**What to look at next.** Of the readings I could request from the field technician, which one narrows the quantity I am worried about most?	(2.16)	design	preposterior	18
Alarms and thresholds					
16	**Alarm setting.** 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?	(2.13), (2.12)	prediction under a threshold	posterior	19
17	**Alarm credibility.** This alarm just fired; what is the probability it is real?	(2.14), over two hypotheses	hypothesis comparison	evidence	20

Three readings of the table carry through the book. Questions 1 to 4 and 9 to 12 rest on the same fit of a window on each plant, read at different places of the plant or at scenario nodes (Question 3 compares it with fits that free one meter's bias; Question 12 stacks windows). Questions 5 to 7 are diagnosis in the sense of fault detection and isolation, and the candidates come from the plant graph, not from the operator's memory. Questions 16 and 17 are where operations meets reliability: an alarm threshold is a decision level, set by calibrating it against the predictive spread of the monitored quantity and the probability of detection for an excursion of stated size, and Question 17 is Bayes' rule over the two tails that calibration produces, the evidences of "excursion" and "no excursion", with the base rate of excursions as their prior odds.

2.4 Worked examples of estimation, prediction and joint events

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

2.4.1 Estimation of the conductance of the hall

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

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

The conductance enters the state as the log-conductance $g = \ln(G/G_{\mathrm{ref}})$ with $G_{\mathrm{ref}} = 1\,\mathrm{kW/K}$ in this chapter's examples (Chapter 1 takes 6 kW/K; the reference only shifts g by a constant). For this example, suppose that a longer record of the hall has been summarized as one row on the log-conductance, the value $\hat{g}$ with the standard deviation σ_g,

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

These values are illustrative. The row stands for that longer record and is not the fit of Chapter 1, whose one reading of 0.5 K cannot give a spread of 5 %: there the MAP conductance is 5.917 kW/K and the probability of staying below 27.5°C at full load is 0.246, the same question under different evidence. How readings are reduced to a row of this kind is the subject

of later chapters. The row is the only one that acts on g alone, and the fit below returns the posterior of the third entry with center $\hat{g}$ and spread $\sqrt{\Sigma_{33}} = \sigma_g$.

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

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

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

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

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

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

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

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

iteration [–]	T_1 [°C]	T_2 [°C]	g [–]	$\mathcal{L}$ [–]
0	23.0000	28.0000	1.7918	1.281
1	22.5998	27.1996	1.8718	6.083
2	22.6154	27.2308	1.8718	6×10^{-18}
3	22.6154	27.2308	1.8718	2×10^{-26}

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

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

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

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

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

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

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

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

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

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

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

so that the reported answer is $G = 6.5 \pm 0.325\,\mathrm{kW/K}$, with the central 95 % band $6.5 \pm 1.960 \times 0.325$, from 5.863 to 7.137 kW/K. The exact lognormal summaries follow from the normal of the third entry and stand beside it: the median is $G_{\mathrm{ref}}\, \mathrm{e}^{\hat{z}_3} = 6.5\,\mathrm{kW/K}$, the location above; the mean, by the moment-generating function of the normal, is $G_{\mathrm{ref}}\, \mathrm{e}^{\hat{z}_3 + \Sigma_{33}/2} = 6.5\, \mathrm{e}^{0.00125} = 6.508\,\mathrm{kW/K}$; the standard deviation is $0.3256\,\mathrm{kW/K}$, a gap of 0.0006 from the band; and the central 95 % credible interval is $6.5\, \mathrm{e}^{\pm 1.960 \times 0.05}$, from 5.893 to 7.169 kW/K, whose ends lie 0.030 and $0.032\,\mathrm{kW/K}$ above those of the band. Each of these is a summary of a distribution; the answer to question A is the distribution itself, equation (2.9) with $f(\mathbf{z}) = G$, and any single number reported from it must say which summary it is.

2.4.2 The threshold probability of the hall at full load

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

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

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

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

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

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

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

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

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

The normal tail of the entry and the exact value differ by 0.003, because T_2 is not linear in g along the ridge. The probability that the hall stays below the limit at full load is 72.0 %, with 71.7 % exact beside it, and the probability that it exceeds the limit is 28.0 %, with 28.3 % exact, conditional on the row of equation (2.29), the operating conditions and the steady-state thermal model. Whether a 28 % probability of exceeding the limit is acceptable is not a probability question; it depends on the consequences of exceedance and the operating criteria that apply. The assurance rule of equation (2.12) states such a criterion as a number, applied to the probability that the requirement holds: at an assurance level $\kappa = 0.95$ the requirement $T_2 < 27.5^\circ\text{C}$ is not accepted by either value, since 0.720 and 0.717 both lie below 0.95.

Figure 2.2 shows the two questions side by side. In panel (a) the horizontal axis is the conductance G and the vertical axis its posterior density; the shaded area to the right of the dashed line at G^* is the exact 0.717. In panel (b) the horizontal axis is the full-load temperature T_2 and the vertical axis its posterior density; the solid curve is the exact density implied by the lognormal G, the dashed curve the normal of the entry of equation (2.34), whose area to the left of the limit is the reported 0.720, marked on the panel, and the shaded area under the solid curve to the left of the limit is the same exact 0.717. The two shaded areas are one probability, seen once on the conductance and once on the temperature: the event $T_2 < 27.5^\circ\text{C}$ and the event $G > G^*$ are the same set of states.

The check against the exact value and against draws. Equation (2.37) is exact under the normal posterior of g, because the map from g to T_2 is monotone along the ridge; equation (2.34) is the route that every derived quantity has. Had T_2 not been carried in the state, the same band would come from the gradient of $f(\mathbf{z}) = T_c + (Q_2/G_{\text{ref}})\,\mathrm{e}^{-g}$, $\mathbf{g} = (0,\ 0,\ -9.231)^{\mathsf{T}}$ K, and it agrees with Σ_{22} because in the limit of equation (2.33) $\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g} = \sigma_g^2(\mathbf{g}^{\mathsf{T}}\mathbf{v})^2$ and $\mathbf{g}$ has the same projection on $\mathbf{v}$, -9.231, as the unit vector on T_2. Table 2.6 compares the reported and the exact probabilities for the limit of the example and three others. The first column is the limit,

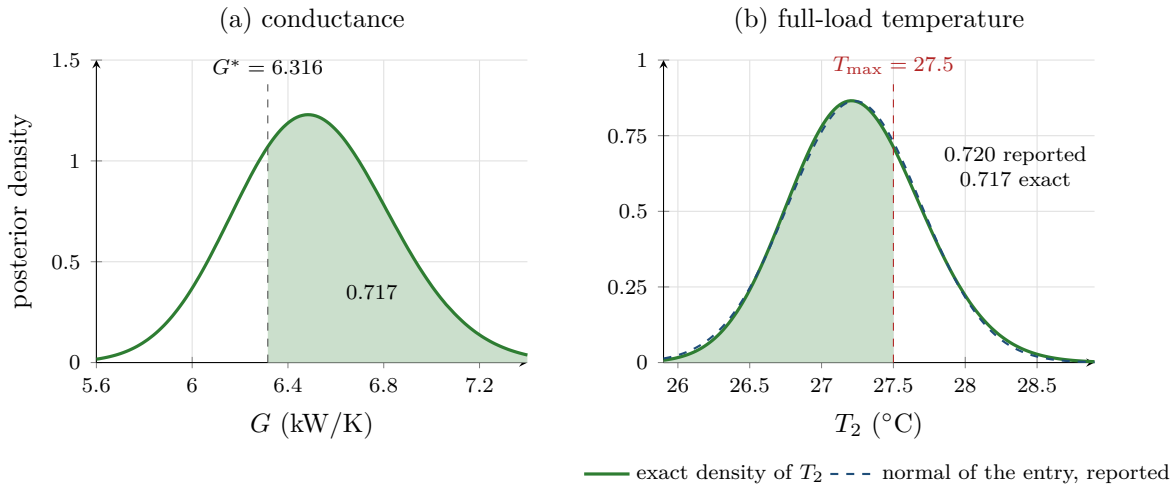


Figure 2.2: One posterior, two questions. (a) Lognormal posterior of the conductance, density $(G\,s\sqrt{2\pi})^{-1}\exp[-(\ln G - \ln 6.5)^2/(2s^2)]$ with $s = \sqrt{\Sigma_{33}} = 0.05$; shaded, $\Pr(G > 6.316\,\text{kW/K} \mid \mathbf{y}) = 0.717$. (b) Posterior of the full-load temperature $T_2 = 18 + 60/G$, lognormal in $T_2 - 18$ with median 9.231 K and log-spread 0.05 (solid), and the normal of the entry, $\mathcal{N}(\hat{z}_2, \Sigma_{22}) = \mathcal{N}(27.231, 0.4615^2)$ (dashed), whose tail is the reported 0.720; shaded, the exact $\Pr(T_2 < 27.5°\text{C} \mid \mathbf{y}) = 0.717$.

the second the equivalent conductance threshold of equation (2.35), the third the standardized distance in g, and the last two the reported normal of the entry, with mean $\hat{z}_2$ and variance Σ_{22}, and the exact probability. The two differ by at most 0.007 over these limits and by 0.003 at 27.5°C, because the curvature of $60/G$ across a posterior of 5 % spread is small; panel (b) of Figure 2.2 shows the slight right skew the normal misses. Small as it is, the gap can cross a decision level: at $T_{\max} = 28°\text{C}$ and $\kappa = 0.95$ the normal gives 0.9522, which accepts the requirement, and the exact value 0.9453, which does not. Both values are then stated, and the operator sees that the requirement sits on the level. A second check draws 200 000 values of g from its posterior, computes T_2 for each, and counts the fraction below the limit: 0.7175, against the exact 0.7174, within the sampling error of 0.001.

Table 2.6: The threshold probability $\Pr(T_2 < T_{\max} \mid \mathbf{y})$ at $Q_2 = 60\,\text{kW}$, $T_c = 18°\text{C}$, $g \mid \mathbf{y} \sim \mathcal{N}(\ln 6.5, 0.05^2)$: reported, by the normal $\mathcal{N}(\hat{z}_2, \Sigma_{22})$ of the entry T_2 (equation (2.34)), and exact (equation (2.36)).

$T_{\max}$ [°C]	$G^* = Q_2/(T_{\max} - T_c)$ [kW/K]	$(\hat{z}_3 - \ln G^*)/\sqrt{\Sigma_{33}}$ [–]	reported, normal of T_2 [–]	exact [–]
27.0	6.6667	−0.5064	0.3085	0.3063
27.5	6.3158	0.5750	0.7202	0.7174
28.0	6.0000	1.6009	0.9522	0.9453
28.5	5.7143	2.5767	0.9970	0.9950

What the two questions share. Both questions used the same posterior. Estimation, $p(G \mid \mathbf{y})$, describes the uncertainty about the conductance; the threshold probability, $\Pr(T_2 < T_{\max} \mid \mathbf{y})$, integrates that uncertainty over the states that satisfy the temperature limit. Neither required a new posterior; each asked a different question of the same one. That reuse is one of

the principal advantages of a probabilistic model of the plant: the expensive step, the posterior, is paid once, and each new question is an integral over it.

2.4.3 The posterior of the hall's alarm rate

A different record of the same hall: the log of its high-temperature alarm. Over the last 1000 h of logged operation, about six weeks, the alarm became active twice, and the night operator asks for the probability that the coming twelve-hour night shift passes without it. This is a prediction question, not only a question about the alarm rate. Let λ be the unknown alarm rate, in activations per hour, $k = 2$ the observed number of activations, $T_{\text{tot}} = 1000\,\text{h}$ the logged operation behind the count, and $\tau = 12\,\text{h}$ the length of the shift. The activations are taken to follow a homogeneous Poisson process, so that the count in the logged time is Poisson with mean λT_{tot},

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

This is a modeling choice. It is what lets the log be summarized by k and T_{tot} alone, and it fails when one excursion produces several activations (an alarm that chatters about its threshold) or when the activations crowd into the hot hours of the day; the count here is of activations after the alarm's deadband and on-delay, which merge chatter into one. As a function of λ, the likelihood is proportional to $\lambda^k \, \mathrm{e}^{-\lambda T_{\text{tot}}}$; the factor $T_{\text{tot}}^k/k!$ does not depend on the rate. The prior is a gamma distribution in the shape–rate parameterization, $\lambda \sim \text{Gamma}(a, b)$, with density proportional to $\lambda^{a-1} \, \mathrm{e}^{-b\lambda}$ and, here,

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

the alarm records of the site's sister halls, whose prior mean $a/b = 0.004\,\text{h}^{-1}$ is one activation every 250 h. Bayes' rule gives the posterior up to a constant as likelihood times prior, $p(\lambda \mid k, T_{\text{tot}}) \propto \Pr(k \mid \lambda, T_{\text{tot}})\, p(\lambda)$.

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

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

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

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

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

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

The Newton step $\Delta\omega = -\mathcal{L}'/\mathcal{L}''$ moves to the minimum of the parabola with the current slope and curvature. The table below lists the iterations from the maximum likelihood rate

$k/T_{\text{tot}} = 0.002\,\text{h}^{-1}$, that is $\omega = \ln 2$, where the count row's slope is zero: the columns are the iteration, the log-rate, the rate, the slopes of the two rows, the slope and curvature of the objective, and the step.

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

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

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

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

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

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

Set like with like, the reported posterior and the exact one differ, because the Laplace normal of the log-rate makes the rate lognormal: its median is $0.002667\,\text{h}^{-1}$ against the gamma's $0.002448\,\text{h}^{-1}$, a gap of $0.000219\,\text{h}^{-1}$; its mean $\hat{\lambda}\,\text{e}^{\sigma_\omega^2/2} = 0.003022\,\text{h}^{-1}$ against the gamma's $0.002667\,\text{h}^{-1}$, a gap of $0.000355\,\text{h}^{-1}$; and its central 95 % interval runs from 0.001001 to $0.007105\,\text{h}^{-1}$ against the gamma's 0.000727 to $0.005845\,\text{h}^{-1}$. The pair the Newton fit returned coincides with the gamma's mean and standard deviation: the mode of the log-rate maps onto the gamma mean and the linearized band onto the gamma standard deviation, for every gamma posterior, but that compares different summaries, a median of one posterior with a mean of the other. The gamma mean is one activation every 375 h. It lies between the prior mean $a/b = 0.004\,\text{h}^{-1}$, one every 250 h, and the maximum likelihood estimate $k/T_{\text{tot}} = 0.002\,\text{h}^{-1}$, one every 500 h, as the weighted form $a'/b' = \frac{b}{b'}\frac{a}{b} + \frac{T_{\text{tot}}}{b'}\frac{k}{T_{\text{tot}}} = \frac{1}{3}(0.004) + \frac{2}{3}(0.002)$ shows. An estimation question has been answered. The operator's question, a quiet night shift, has not.

2.4.4 The probability of a quiet night shift

Let T be the time to the next activation of the alarm, counted from the start of the shift. (In this example T, in hours, is a time; in the hall example, T_1 and T_2 are hall temperatures in °C. The units distinguish them.) Under the Poisson process the time to the next activation does not depend on how long ago the last one was, so the question is

$$\Pr(T > \tau \mid \mathbf{y}), \qquad \tau = 12\,\text{h}. \tag{2.40}$$

For a known constant rate the probability of no activation is exponential, $\Pr(T > \tau \mid \lambda) = \text{e}^{-\lambda\tau}$. The rate is not known, and the event is a future one that stays random even when the rate is known, so the question is a prediction of the general form (2.13), in its event form (2.8): y_{new}

is the event $T > \tau$, and $p(y_{\text{new}} \mid \mathbf{z}) = \mathrm{e}^{-\lambda\tau}$ with $\lambda = \lambda_{\text{ref}}\,\mathrm{e}^{\omega}$. The integral is a function of the one number ω averaged over its posterior, and under the Laplace approximation that posterior is $\mathcal{N}(\hat{\omega}, \sigma_\omega^2)$, so the value read from the state is a normal expectation in $\hat{\omega}$ and σ_ω^2:

$$\Pr(T > \tau \mid \mathbf{y}) \approx \mathbb{E}_{\omega\sim\mathcal{N}(\hat{\omega},\sigma_\omega^2)}\big[\mathrm{e}^{-\lambda_{\text{ref}}\mathrm{e}^{\omega}\tau}\big] \approx \sum_{j=1}^{5} w_j\,\mathrm{e}^{-\lambda_{\text{ref}}\mathrm{e}^{\omega_j}\tau}, \qquad \omega_j = \hat{\omega} + \sqrt{2}\,\sigma_\omega\,x_j. \tag{2.41}$$

The sum is the five-point Gauss–Hermite rule, which replaces a normal expectation by a weighted sum at fixed nodes x_j with weights w_j and is exact for a polynomial of degree nine. The table below evaluates it for the night shift, with $\hat{\omega} = 0.98083$ and $\sigma_\omega = 0.5$: the columns are the node, its weight, the log-rate and rate there, the probability of no activation at that rate, and the weighted term.

x_j [–]	w_j [–]	ω_j [–]	λ_j [h^{-1}]	$\mathrm{e}^{-12\lambda_j}$ [–]	term [–]
−2.02018	0.01126	−0.44766	0.000639	0.99236	0.01117
−0.95857	0.22208	0.30302	0.001354	0.98388	0.21850
0	0.53333	0.98083	0.002667	0.96851	0.51654
0.95857	0.22208	1.65864	0.005252	0.93892	0.20851
2.02018	0.01126	2.40931	0.011126	0.87501	0.00985

The terms sum to 0.9646, and a direct numerical integration agrees to six decimals. The predicted probability of a night shift without a high-temperature alarm is 96.46 %, and the probability of at least one activation is 0.0354; these are the values reported.

The exact value stands beside them, because the objective identified the posterior as a gamma and the same integral then has a closed form. The conditional probability of a quiet shift is averaged over the posterior of the rate,

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

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

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

and combine the two exponentials:

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

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

$$\Pr(T > \tau \mid \mathbf{y}) = \left(\frac{b'}{b'+\tau}\right)^{a'}. \tag{2.42}$$

For the night shift,

$$\Pr(T > 12\,\mathrm{h} \mid \mathbf{y}) = \left(\frac{1500}{1512}\right)^{4} = 0.9686,$$

and the probability of at least one activation is 0.0314. The calculation needed no Gaussian or Laplace approximation: with a gamma posterior and an exponential probability of no activation, the predictive probability has an exact closed form. The gap between the two is 0.00406 for the shift, and over the next 100 h, about eight night shifts, the state gives 0.748 against the exact 0.7725, a gap of 0.024. The state route errs low because the lognormal it implies for λ has the larger mean $0.00302\,\mathrm{h}^{-1}$ and a heavier upper tail than the gamma. No decision of this chapter sits between the two values; an assurance level of 0.965 on the shift would, and both would then be stated, the one read from the state failing it and the exact one meeting it.

2.4.5 The error of the plug-in alarm rate

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

$$\Pr_{\text{plug}}(T > 12\,\text{h}) = \mathrm{e}^{-\mathbb{E}[\lambda|\mathbf{y}]\times 12\,\text{h}} = \mathrm{e}^{-0.0320} = 0.9685,$$

against the exact predictive 0.9686; over the next 100 h it gives $\mathrm{e}^{-0.2667} = 0.7659$ against 0.7725. The two differ because averaging and exponentiation do not commute,

$$\mathbb{E}\big[\mathrm{e}^{-\lambda\tau} \mid \mathbf{y}\big] \neq \mathrm{e}^{-\mathbb{E}[\lambda|\mathbf{y}]\,\tau} \quad \text{in general,} \tag{2.43}$$

since the probability of no activation is a nonlinear function of the rate. The direction of the difference follows from the shape of that function. For $\tau > 0$ the map $\lambda \mapsto \mathrm{e}^{-\lambda\tau}$ has second derivative $\tau^2\mathrm{e}^{-\lambda\tau} > 0$, so it is convex, and Jensen's inequality, the standard result that the expectation of a convex function is at least the function of the expectation, gives $\mathbb{E}[\mathrm{e}^{-\lambda\tau} \mid \mathbf{y}] \geq \mathrm{e}^{-\mathbb{E}[\lambda|\mathbf{y}]\tau}$: the plug-in value understates the exact probability of a quiet shift. The low rates the posterior admits raise that probability by more than the high rates lower it. In the log-rate the same function is concave while $\lambda\tau < 1$, so the reported average (2.41) starts below the plug-in and crosses it at 356 h. For one shift the plug-in sits 0.0001 from the exact value, closer than the reported 0.9646. The plug-in value is wrong not because it is far from the predictive value, but because it answers a different question: the probability of a quiet shift for an alarm whose rate is known to be $0.002667\,\text{h}^{-1}$, which is not the alarm in question, and it carries no band.

Table 2.7 shows how the values change with the horizon. The first column is the horizon; the second the reported average of equation (2.41); the third the exact predictive probability of no activation, equation (2.42); the fourth their gap; the fifth the plug-in value at the posterior mean; the sixth the exact value minus the plug-in; and the last the plug-in value at the maximum likelihood estimate $0.002\,\text{h}^{-1}$, which ignores the prior as well as the uncertainty. For one shift or one day the exact and plug-in values agree to three decimals. Their ratio grows with the horizon, from 1.0005 at 24 h to 1.87 at 1000 h and 7.0 at 2000 h; the absolute difference peaks near 800 h, at 0.0625, and then falls, to 0.0289 at 2000 h, because both values tend to zero; the gap of the reported average peaks in the hundreds of hours, at 0.0387 among the rows of the table. Figure 2.3 draws the ingredients. In panel (a) the horizontal axis is the alarm rate in activations per 1000 h and the vertical axis the probability density per unit of that axis; the blue curve is the prior, the green curve the posterior, and the dotted lines mark the maximum likelihood estimate, the posterior mean and the prior mean. In panel (b) the horizontal axis is the horizon and the vertical axis the probability of no activation; the green curve is the exact predictive, the blue dash-dotted curve the reported average over the normal of the log-rate and the dashed red curve the plug-in, and the vertical dotted line marks the 12 h shift. The exact curve lies above the plug-in curve at every horizon, and their ratio grows as the horizon lengthens, because the longer the horizon, the more the probability of no activation bends across the spread of the posterior; the blue curve falls below both at first and crosses the plug-in at 356 h.

Principle 2.2. *Predicting a quiet shift requires averaging the probability of no alarm over the uncertainty in the alarm rate. Substituting a point estimate of the rate answers the question for a different alarm, one whose rate is known.*

2.4.6 Two halls quiet through the night as a joint event

A site has two other halls, A and B, each with its own temperature alarm, and one night operator for both. Figure 2.4 draws the arrangement: the white box on the left is the chiller

Table 2.7: Probability of no activation of the hall's alarm over a horizon τ: the reported average over the normal of the log-rate (equation (2.41), $\hat{\omega} = \ln(8/3)$, $\sigma_\omega = 0.5$); exact, from the posterior Gamma$(4, 1500\,\text{h})$, $(1500/(1500+\tau))^4$; plug-in $\text{e}^{-4\tau/1500}$; plug-in at the maximum likelihood estimate $\text{e}^{-0.002\tau}$. Gaps are computed before rounding.

τ [h]	reported, from the state [–]	exact [–]	gap [–]	plug-in, posterior mean [–]	exact minus plug-in [–]	plug-in, maximum likelihood [–]
12	0.9646	0.9686	0.0041	0.9685	0.0001	0.9763
24	0.9307	0.9385	0.0078	0.9380	0.0005	0.9531
100	0.7481	0.7725	0.0244	0.7659	0.0065	0.8187
500	0.2777	0.3164	0.0387	0.2636	0.0528	0.3679
1000	0.1015	0.1296	0.0281	0.0695	0.0601	0.1353
2000	0.0205	0.0337	0.0132	0.0048	0.0289	0.0183

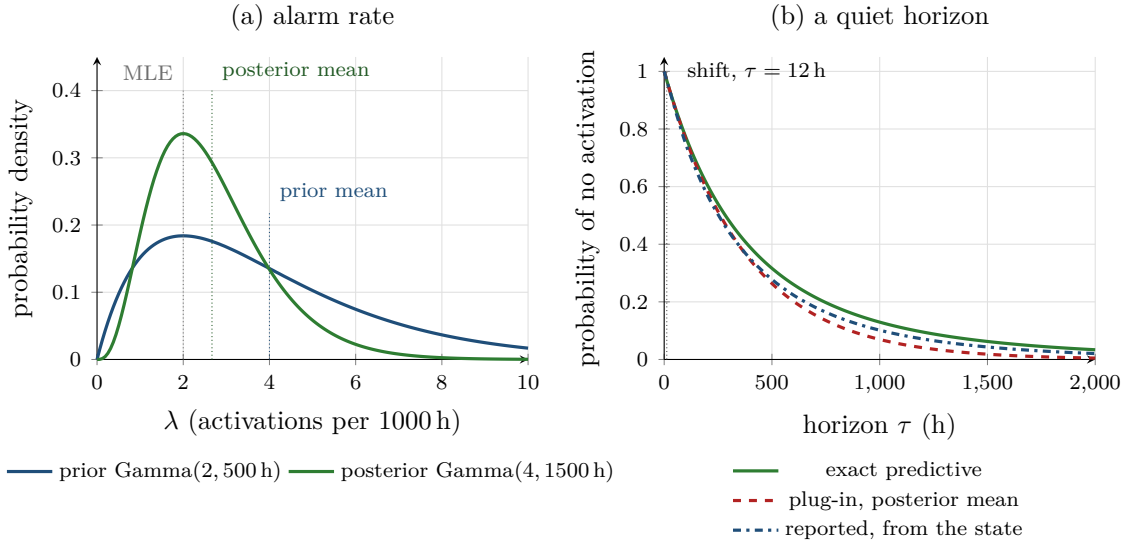


Figure 2.3: The hall's alarm rate and the probability of a quiet horizon. (a) Prior $0.25\,x\,\text{e}^{-0.5x}$ and posterior $(1.5^4/3!)\,x^3\text{e}^{-1.5x}$ in $x = 1000\lambda$; dotted lines at the maximum likelihood estimate 2.0, the posterior mean 2.67 and the prior mean 4.0 activations per 1000 h. (b) Exact predictive probability of no activation $(1500/(1500+\tau))^4$ of equation (2.42), plug-in $\text{e}^{-4\tau/1500}$ and the reported average of equation (2.41); at the 12 h shift they are 0.9686, 0.9685 and 0.9646, at 100 h 0.7725, 0.7659 and 0.7481.

that supplies the cooling units of both halls, and so the supply air of both; the two green boxes are the halls, each with its own causes of an alarm (a unit fault, a load step, a blocked floor tile); the white box on the right is the control room that receives both alarms; the solid arrows are the paths along which chilled water and alarms pass. Define the events

$$A_A = \{\text{hall A's alarm fires in the night}\}, \qquad A_B = \{\text{hall B's alarm fires in the night}\},$$

and write $\bar{A}_A$ and $\bar{A}_B$ for their complements, a quiet night in each hall. The operator wants a quiet night in both, the intersection $\bar{A}_A \cap \bar{A}_B$, and fears the opposite corner, both halls alarming in the same night, $A_A \cap A_B$, two responses one operator cannot give at once. The two are tied by inclusion and exclusion, which holds without any independence assumption:

$$\Pr(\bar{A}_A \cap \bar{A}_B) = 1 - \Pr(A_A) - \Pr(A_B) + \Pr(A_A \cap A_B). \tag{2.44}$$

Each is the general form (2.12) with the indicator of an intersection in place of the indicator of one bound, here with the marginal probabilities given directly rather than computed from a posterior, so there is no state route to report and no gap to state; every probability of this example is conditional on the same evidence, and the bar is left out below to keep the products readable. Suppose $\Pr(A_A) = 0.05$ and $\Pr(A_B) = 0.03$, illustrative values of the size the alarm log of Section 2.4.4 gives one hall, 0.0354 for a night and 0.0314 exact. If the two alarms are independent, the probability that both fire is the product of the marginal probabilities,

$$\Pr(A_A \cap A_B) = \Pr(A_A)\,\Pr(A_B) = 0.05 \times 0.03 = 0.0015,$$

and the probability of a quiet night in both halls is, by equation (2.44) or as the product of the two complements,

$$\Pr(\bar{A}_A \cap \bar{A}_B) = 0.95 \times 0.97 = 0.9215, \tag{2.45}$$

or 92.15 %.

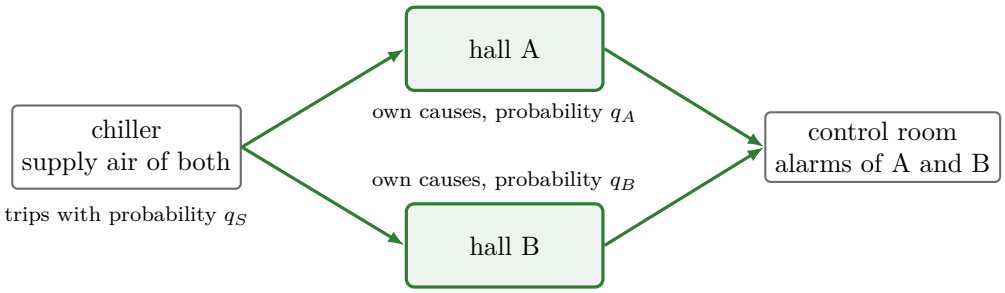


Figure 2.4: Two halls on one chiller. Each hall's alarm reaches the control room; a quiet night in both needs neither alarm to fire. A chiller trip (probability q_S in the night) raises the supply air of both halls and fires both alarms at once; the halls' own causes (probabilities q_A, q_B) are independent given a running chiller. With $q_S = 0$ the two alarms are independent.

2.4.7 One chiller feeding both halls

Both halls draw their supply air from cooling units on the same chiller. Then independence may not hold, and the probability of the intersection in the form (2.12) must be written with the product rule, which holds without any independence assumption:

$$\Pr(A_A \cap A_B) = \Pr(A_A)\,\Pr(A_B \mid A_A). \tag{2.46}$$

The conditional probability $\Pr(A_B \mid A_A)$ may be much larger than the marginal $\Pr(A_B)$: an alarm in hall A raises the probability that the chiller has tripped, and a tripped chiller raises

the supply air of hall B as well. A shared supply air is a common cause that can make the two alarms strongly dependent.

An illustrative model shows by how much. Let the chiller trip during the night with probability q_S, which fires both alarms, and let each hall's alarm otherwise fire from its own causes, independently, with probabilities q_A and q_B. The law of total probability over the two states of the chiller, the integral of the form over a state with two values, gives the marginal and the joint probabilities, exactly, with no second route beside them:

$$\Pr(A_A) = q_S + (1 - q_S)\, q_A, \qquad \Pr(A_B) = q_S + (1 - q_S)\, q_B,$$

$$\Pr(A_A \cap A_B) = q_S \cdot 1 + (1 - q_S)\, q_A q_B, \qquad \Pr(\bar{A}_A \cap \bar{A}_B) = (1 - q_S)(1 - q_A)(1 - q_B). \quad (2.47)$$

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

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

and equation (2.46) gives $\Pr(A_B \mid A_A) = 0.010808/0.05 = 0.216$, seven times the marginal 0.03. The probability of a quiet night in both halls is $0.99 \times 0.959596 \times 0.979798 = 0.930808$, the same as $1 - 0.08 + 0.010808$ by equation (2.44). Table 2.8 repeats the calculation for five chiller-trip probabilities. The first column is q_S; the next two the own-cause probabilities that keep the marginals fixed; then the probability that both alarms fire, the conditional probability of the second alarm given the first, the ratio of the joint probability to the independent value 0.0015, and the probability of a quiet night in both halls. With the marginals unchanged, a trip probability of 0.01 multiplies the probability of a night with both halls alarming by 7.2, and one of 0.02 by 13.5: the chiller, not the halls' own conditions, then governs it. The last two rows, a trip in one night of a hundred, 3.65 trips a year, and in one of fifty, are high for a maintained chiller and are there to show the trend.

Table 2.8: Two halls on one chiller through a night, with the marginals held at $\Pr(A_A) = 0.05$ and $\Pr(A_B) = 0.03$ (illustrative model of equation (2.47)).

q_S [–]	q_A [–]	q_B [–]	$\Pr(A_A \cap A_B)$ [–]	$\Pr(A_B \mid A_A)$ [–]	ratio to 0.0015 [–]	$\Pr(\bar{A}_A \cap \bar{A}_B)$ [–]
0	0.05000	0.03000	0.001500	0.0300	1.00	0.92150
0.001	0.04905	0.02903	0.002422	0.0484	1.61	0.92242
0.005	0.04523	0.02513	0.006131	0.1226	4.09	0.92613
0.01	0.04040	0.02020	0.010808	0.2162	7.21	0.93081
0.02	0.03061	0.01020	0.020306	0.4061	13.54	0.94031

The shared chiller moves the two corners in opposite directions. The quiet night in both halls becomes slightly more probable, 0.9308 against 0.9215, because when the alarms come they tend to come together; the night with both halls alarming becomes 7.2 times as probable. A product of marginals gets the first nearly right and the second badly wrong, and the second is the one that leaves one operator facing two halls at once. The wording "both halls" names an intersection of events, and an intersection directs the formulation to joint probability and to the modeling of dependence, not to a product of marginals. In a larger plant the event of concern is often a union of many overlapping events, some rack somewhere above its limit, and Chapter 8 derives its probability from the joint posterior of the racks; several alarms arriving together, an alarm flood, return in Volume II.

2.5 Model comparison, diagnosis and decisions

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

2.5.1 Posterior probabilities of competing models

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

$$M_0 : \text{the air path operates normally}, \qquad M_1 : \text{the air path is partially obstructed},$$

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

$$\Pr(M_j \mid \mathbf{y}) = \frac{p(\mathbf{y} \mid M_j)\, \Pr(M_j)}{\sum_k p(\mathbf{y} \mid M_k)\, \Pr(M_k)}, \tag{2.48}$$

where $\Pr(M_j)$ is the prior probability of model M_j, $\Pr(M_j \mid \mathbf{y})$ its posterior probability, and $p(\mathbf{y} \mid M_j)$ its evidence, equation (2.26). Dividing equation (2.48) for M_1 by the same equation for M_0 cancels the denominator and gives the odds form,

$$\frac{\Pr(M_1 \mid \mathbf{y})}{\Pr(M_0 \mid \mathbf{y})} = B_{10}\, \frac{\Pr(M_1)}{\Pr(M_0)}, \qquad B_{10} = \frac{p(\mathbf{y} \mid M_1)}{p(\mathbf{y} \mid M_0)} : \tag{2.49}$$

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

2.5.2 A numerical comparison of two explanations

Suppose the evidences are given, as illustrative values, not read from fits, so there is no Laplace value and exact value to set side by side: $p(\mathbf{y} \mid M_0) = 0.01$ and $p(\mathbf{y} \mid M_1) = 0.04$, and the prior probabilities are $\Pr(M_0) = 0.8$ and $\Pr(M_1) = 0.2$. Equation (2.48) gives

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

The Bayes factor is

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

and the odds form checks the result: posterior odds $4 \times (0.2/0.8) = 1$, a posterior probability of $1/(1+1) = 0.50$. The evidence is four times more probable under the obstruction model, yet its posterior probability is only 50 %, because the prior probability assigned to obstruction was

low. The Bayes factor expresses the relative support that the evidence provides; it is not by itself a posterior model probability, and turning it into one requires the prior probabilities of the models.

Figure 2.5 makes the distinction visible. In panel (a) the two pairs of bars are the prior and the posterior probabilities of the two models: the evidence moves probability from M_0 to M_1 until they are equal. In panel (b) the horizontal axis is the prior probability of the obstruction model and the vertical axis its posterior probability; the green curve is $\Pr(M_1 \mid \mathbf{y}) = 4\,\Pr(M_1)/(3\,\Pr(M_1)+1)$, which equation (2.49) gives for $B_{10} = 4$, and the dashed diagonal is the case $B_{10} = 1$, evidence that does not discriminate. The same Bayes factor of 4 turns a prior of 0.05 into 0.174, a prior of 0.2 into 0.50 (the marked point) and a prior of 0.5 into 0.80: the factor is fixed by the evidence, the conclusion is not.

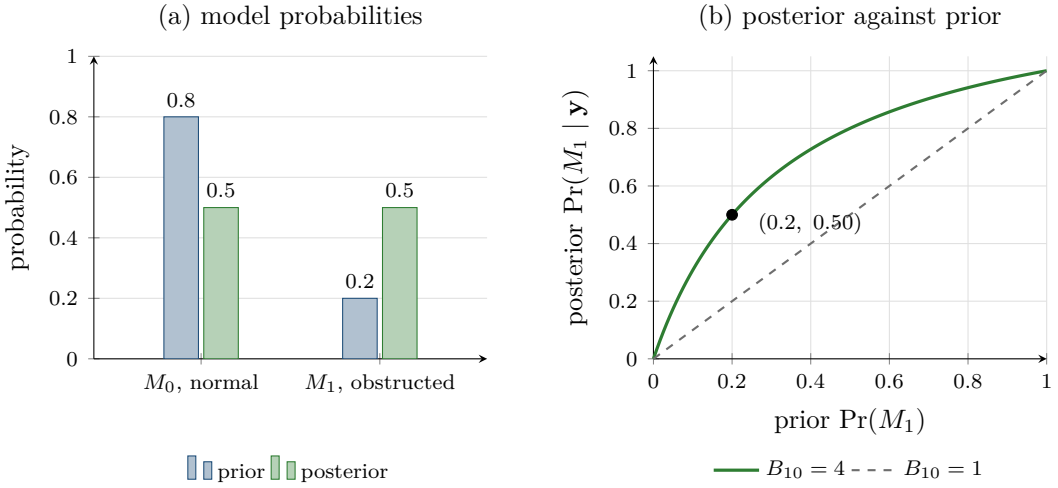


Figure 2.5: Model comparison for the air path of the hall. (a) Prior and posterior probabilities of the normal (M_0) and obstructed (M_1) models, with evidences 0.01 and 0.04. (b) Posterior probability of M_1 as a function of its prior, $4x/(3x+1)$ for a Bayes factor of 4 (equation (2.49)); the dashed diagonal is a Bayes factor of 1; the point is the example.

2.5.3 Diagnosis over a family of candidate faults

Fault detection and isolation usually involves many competing explanations, not two. For a hall that runs warmer than its model expects they might be a fouled coil, a fan running slow, a biased hall sensor, a load meter reading low or a floor tile left out of place. The mathematical target is the general form (2.15), normalized as in equation (2.48): the posterior probability of each candidate,

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

where F_j is a candidate fault condition, which is the drawing M_j of equation (2.15): the plant graph with that fault written into its rows, so that $p(\mathbf{y} \mid F_j) = p(\mathbf{y} \mid M_j)$. The output is the ranking of the candidates by it, not a single verdict. A fault condition can be represented in two ways: as a separate model, whose evidence enters equation (2.48), or as a latent quantity within one common model, such as a bias entry of the state that is zero when the sensor is sound. The second representation means that diagnosis does not always require fitting a separate model for every possible fault. Diagnosis earns its own category for two reasons: the candidates are generated systematically from the plant graph, one for each component or reading that could be at fault, rather than written by hand; and the answer is a ranking. The critical step in either

representation is to define the candidate fault events or hypotheses clearly enough that their probabilities can be calculated and compared.

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

2.5.4 Expected loss of starting the standby unit before the night peak

Questions about a running plant also lead to decisions. Before the night peak the operator can start the hall's standby cooling unit, a unit outside the one-path model of Chapter 1, or wait, and the probability of an excursion is not enough to decide: the consequence of an excursion and the cost of the action enter as well. Let e be the event that the hall air stays above its limit at the night peak long enough for the IT equipment's protection to act, an excursion, with consequence $C_e = \text{R\$}\,100\,000$ in service credits and recovery, and suppose that if the operator waits its probability is $\Pr(e \mid \mathbf{y}, \text{wait}) = 0.0015$, an illustrative value. An excursion in this sense, the hall air held above the limit long enough for the IT protection to act, is far rarer than the hall crossing the 27.5°C envelope of Question B, whose probability at 60 kW, under the row of equation (2.29), is 0.280, with 0.283 exact. The excursion is a future event that stays random when the state is known, so the loss of waiting given the state is itself an expectation, the derived quantity

$$L(\text{wait}, \mathbf{z}) = C_e \Pr(e \mid \mathbf{z}, \text{wait}) + 0 \cdot \Pr(\bar{e} \mid \mathbf{z}, \text{wait}) = C_e \Pr(e \mid \mathbf{z}, \text{wait}),$$

with no loss when no excursion occurs, and the expected loss is its location (2.10), an estimation question about a consequence, answered by the expected consequence (2.23) of waiting. Averaging over the posterior with equation (2.8) turns the conditional probability into its predictive one, and the expected loss is the probability-weighted average of the two outcomes:

$$\mathcal{R}(\text{wait}) = \mathbb{E}[L(\text{wait}, \mathbf{z}) \mid \mathbf{y}] = \Pr(e \mid \mathbf{y}, \text{wait})\, C_e + \Pr(\bar{e} \mid \mathbf{y}, \text{wait}) \cdot 0, \tag{2.50}$$

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

$$\mathcal{R}(\text{wait}) = 0.0015 \times \text{R\$}\,100\,000 + 0.9985 \times 0 = \text{R\$}\,150$$

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

Now suppose starting the standby unit costs $c_{\text{start}} = \text{R\$}\,80$ for the night, in fan energy and run hours (a cost, the c_a of equation (2.23)), and lowers the probability of an excursion from

0.0015 to 0.0005. Its loss given the state is $L(\text{start}, \mathbf{z}) = c_{\text{start}} + C_e \Pr(e \mid \mathbf{z}, \text{start})$, and its expected loss is the certain cost of the action plus the expected loss that remains,

$$\mathcal{R}(\text{start}) = \text{R\$}\,80 + 0.0005 \times \text{R\$}\,100\,000 = \text{R\$}\,130,$$

against $\mathcal{R}(\text{wait}) = \text{R\$}\,150$. The expected net benefit of starting the unit is $\text{R\$}\,150 - \text{R\$}\,130 = \text{R\$}\,20$ per night, and under these assumptions starting it is the preferred decision. The comparison rests on three numbers, two probabilities and one consequence, and its conclusion changes if any of them moves: the break-even cost of the action is $(0.0015 - 0.0005) \times \text{R\$}\,100\,000 = \text{R\$}\,100$, and, read the other way, acting and waiting cost the same when the probability of an excursion if the operator waits is $0.0005 + 80/100\,000 = 0.0013$, the break-even probability. A night whose excursion probability the posterior puts below 0.0013 is a night to wait. The probabilities here stand for the predictives a posterior would supply; they are given, so the expected losses carry no gap of their own, and a gap in the excursion probability that straddled 0.0013 would be stated with both values and both actions.

2.5.5 The general decision formulation

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

$$a^* = \arg\min_a \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]. \tag{2.51}$$

The standby-unit example of Section 2.5.4 is this rule with one event, the excursion, and two actions. The change is an important one. The questions of the earlier sections asked what will happen; equation (2.51) asks what should be done, given what might happen. The transition from inference to decisions is what the control room is for.

2.6 Measurements, the classification of questions and the derivation workflow

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

2.6.1 Uncertainty reduction by a measurement

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

For the hall, consider n readings of the hall temperature at the present load $Q_1 = 30\,\text{kW}$, each with an error of standard deviation $\sigma_T = 0.5\,\text{K}$, that of the hall sensor of Chapter 1 (independent readings, from sensors at different places of the hall). Each reading is a row on the first entry of the state of equation (2.30),

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

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

$$\mathbf{\Sigma}_a = \frac{\mathbf{\Sigma}}{1 + 0.21302\,n}, \qquad \mathbf{\Sigma}_a = \frac{\mathbf{\Sigma}}{1.21302} = \begin{bmatrix} 0.04390 & 0.08780 & -0.00951 \\ 0.08780 & 0.17561 & -0.01902 \\ -0.00951 & -0.01902 & 0.00206 \end{bmatrix} \quad (n = 1). \tag{2.52}$$

The reading sees only T_1, yet every entry of $\mathbf{\Sigma}$ shrinks by the same factor, because the posterior has only the one direction $\mathbf{v}$ to shrink along: one reading removes $1 - 1.21302^{-1/2} = 9.2\,\%$ of the spread of T_1, T_2 and g alike, so that $\sqrt{\Sigma_{a,33}}$ falls from 0.05 to 0.0454 and $\sqrt{\Sigma_{a,22}}$ from 0.462 to 0.419 K. Table 2.9 gives the result for 0 to 64 readings. The columns are the number of readings, the precision $1/\Sigma_{a,33}$ of g, its posterior standard deviation $\sqrt{\Sigma_{a,33}}$, the standard deviation $\sqrt{\Sigma_{a,22}}$ of T_2, the fraction of that spread removed, and the expected information gain. The precision of g is $400\,(1 + 0.21302\,n) = 400 + 85.21\,n$ by equation (2.52). The expected information gain is equation (2.18); in the limit both determinants are zero, and the determinant identity $\det\mathbf{\Sigma}/\det\mathbf{\Sigma}_a = \det(\mathbf{I} + \mathbf{J}_a\mathbf{\Sigma}\mathbf{J}_a^{\mathsf{T}})$, with every entry of the $n \times n$ matrix $\mathbf{J}_a\mathbf{\Sigma}\mathbf{J}_a^{\mathsf{T}}$ equal to 0.21302, gives $\mathrm{EIG} = \frac{1}{2}\ln(1 + 0.21302\,n)$. Equation (2.17) evaluated directly at $\sigma_p = 10^{-4}\,\mathrm{K}$, where $\mathbf{\Sigma}^{-1}$ exists, gives the same entries, the precision of 64 readings to within 0.02. The reduction is steep at first and then slows, because the precision grows linearly in n and the spread only as its inverse square root: halving the spread requires a precision of 1600, about fourteen readings, and halving it again a precision of 6400, about seventy.

Table 2.9: Preposterior spread of the hall after n readings at $Q_1 = 30\,\mathrm{kW}$ with $\sigma_T = 0.5\,\mathrm{K}$ (illustrative), from the posterior $\mathbf{\Sigma}$ of equation (2.33) and equation (2.52).

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

2.6.2 The value of a field reading before deciding

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

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

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

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

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

$$\mathrm{EVSI} = \bar{L}_{\mathrm{now}} - \bar{L}_{\mathrm{after}}, \tag{2.53}$$

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

An illustrative case shows the calculation, with the state here the coil's condition, two values, so that every integral is a finite sum, exact, and no second route stands beside it. Suppose the excursion probability 0.0015 of Section 2.5.4 is an average over two conditions of the running unit's coil: fouled, with probability 0.25 and excursion probability 0.0045 if the operator waits, and clean, with probability 0.75 and excursion probability 0.0005; indeed $0.25 \times 0.0045 + 0.75 \times 0.0005 = 0.0015$. Starting the standby unit costs R$ 80 and leaves either condition at 0.0005. A field reading of the coil, the technician's measurement of its air-side pressure drop, would reveal the condition exactly. Table 2.10 lists, for each condition and for the present state of knowledge, the expected loss of each action and the better one. Its columns are the condition, its probability, the excursion probability if the operator waits, the expected loss of waiting, that of starting the standby unit, and the better action. Without the reading the best action is to start the unit, at R$ 130. With it, a fouled coil leads to starting the unit at R$ 130 and a clean one to waiting at R$ 50, so

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

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

Table 2.10: Expected loss per night of the two actions for a coil of unknown condition (illustrative): standby unit R$ 80 for the night, excursion consequence R$ 100 000.

Condition	probability [–]	Pr(e), wait [–]	wait [R$]	start standby [R$]	better action
fouled	0.25	0.0045	450	130	start standby
clean	0.75	0.0005	50	130	wait
unknown (present)	1	0.0015	150	130	start standby

Three quantities are now on the table, and they are related but not interchangeable. The reduction of the posterior covariance, equation (2.17), measures how much uncertainty disappears. The expected information gain, equation (2.18), measures how much is learned. The value of information, equation (2.53), measures how much better the resulting decision is expected to be. A reading can score well on the first two and be worth nothing on the third, which is why the choice of the next reading and the choice of an action both need the preposterior but do not ask the same question of it. The value of information is where two computations meet: the preposterior supplies the outcomes and their weights, and the expected

consequence, minimized over the actions, is what it averages. Chapter 18 ranks the readings the field technician could take by the width each removes from the quantity the operator is worried about.

2.6.3 Identifying the mathematical form of a question

The examples supply a systematic procedure: before computing anything, identify the mathematical objective of the question. Table 2.11 is the reference for that step. Each row is a form of question in the operator's words, the general form it takes, by equation number, its mathematical form, what it derives to, and its category. These forms recur in the atlas of Appendix A, which collects the questions of each chapter that has a questions table, with the section that answers each.

Table 2.11: Classification of questions by their mathematical form, with the general form of each category from Section 2.3.3 and the forms of a trend and of an expected consequence from Sections 2.3.4 and 2.3.5.

Operator question	General form	Mathematical form	Derives to	Category
What is unknown, and how well is it known?	(2.9), (2.10)	$p(z_i \mid \mathbf{y})$, $\mathbb{E}[z_i \mid \mathbf{y}]$	a distribution, a mean, a credible interval: $\hat{z}_i$ and Σ_{ii}	estimation
Does a requirement hold, with the stated assurance?	(2.12)	$\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}) \geq \kappa$, f^* the requirement	the posterior integral of the indicator of the requirement; a normal tail under linearization	prediction under a threshold
What happens in the coming shift?	(2.13)	$\int \Pr(A_{\text{future}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	the average over the state, equation (2.41), with the closed form (2.42) beside it where the posterior is a gamma	prediction under a threshold
Is it rising, and when does it cross its limit?	(2.21)	$\Pr(\gamma > 0 \mid \mathbf{y})$; $\Pr(t^* \leq s \mid \mathbf{y}) = \Pr(\gamma > 0,\ f_s(\mathbf{z}) \geq f^* \mid \mathbf{y})$	the posterior read along time, $\mathbf{g}(t)^{\mathsf{T}}\hat{\mathbf{z}}$ and $\mathbf{g}(t)^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{g}(t)$; a crossing band by inversion	prediction under a threshold for the sign and a crossing by a date; estimation for a forecast and the distribution of a crossing time
Which explanation is correct?	(2.14), (2.15)	$\Pr(M_j \mid \mathbf{y}) \propto p(\mathbf{y} \mid M_j)\, \Pr(M_j)$	the evidences and the prior model probabilities; the Bayes factor	hypothesis comparison; diagnosis over a family of faults
What is the expected consequence?	(2.23)	$\mathcal{R}(a) = \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}] = \int L(a, \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	probability times consequence; $L(a, \hat{\mathbf{z}}) + \frac{1}{2}\operatorname{tr}(\mathbf{H}_a\boldsymbol{\Sigma})$	estimation of a derived quantity
What should be done or measured?	(2.23), (2.16)	$\arg\min_a \mathcal{R}(a)$; EVSI	the expected consequence of each action; the preposterior $\boldsymbol{\Sigma}_a$	design

Many variations of these forms exist. Alarm rationalization, look-ahead over the coming windows, maintenance windows and load transfer between halls each add mathematical structure of their own: joint tails over many entries, integrals of an excursion probability over a horizon, consequences that accrue over a shift. The rows of the table, over the general forms of the

categories, of a trend and of an expected consequence, already cover a substantial range of inference problems, and a new question is usually classified by recognizing which of them it is, before any computation.

2.6.4 The general derivation workflow

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

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

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

The graph extended for the question (step 2). Scenario nodes for a set point, a unit state or a load not yet applied, candidate faults when something may have changed, and the evidence attached as rows: readings, alarm counts, handbook values as priors, each with the modeling assumption it rests on, what that assumption buys and how it can fail.

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

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

Evaluation and checks (step 5). The gaps between the reported values and the exact values or draws collected in a table, the approximations tested, the answers compared with the truth where the record is constructed.

The report (step 6). The probability, the Bayes factor or the expected consequence with its band, or a crossing time with its band by inversion, with the exact value beside it where one exists.

The next reading (step 7). If the band is too wide, the readings that could be requested from the field ranked by the width each would remove, or the sensor that most deserves trust, and a return to step 3 with the new row.

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

exact-value and sampling checks of the hall example did. Step 6 communicates the answer, its uncertainty, its limitations and what it means for the shift, and step 7 is the design question of Section 2.6.1.

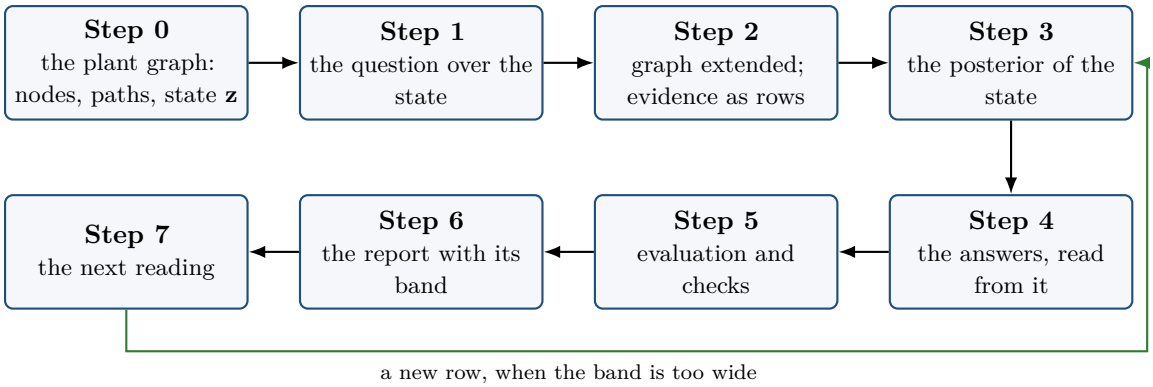


Figure 2.6: The derivation workflow of the book. Step 0 is built once per plant; steps 1 to 7 are one pass; step 3 is shared by every question put to the same plant and evidence; the green return arrow adds the reading chosen at step 7 as a new row and rebuilds the posterior from step 3.

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

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

2.6.5 Common mistakes in formulating operational questions

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

Common mistake

Confusing uncertainty about an alarm rate with the probability of a quiet shift. For an alarm with an unknown rate λ, the probability of no activation over a horizon τ at a given rate is $e^{-\lambda\tau}$. The question $\mathbb{E}[e^{-\lambda\tau} \mid \mathbf{y}]$ is the posterior predictive probability of a quiet horizon: how likely the next shift is to pass without the alarm. The question $\Pr(e^{-\lambda\tau} > 0.99 \mid \mathbf{y})$ is the posterior probability that the alarm's underlying rate meets a target written as a 99 % chance of a quiet horizon: how much support there is for the alarm meeting its rationalization target. They are different numbers and answer different questions.

Common mistake

Confusing prediction with diagnosis. When the hall runs warm, the prediction $\Pr(T > T_{\max} \mid \mathbf{y})$ concerns a temperature event, and the diagnosis $\Pr(M_{\text{fouled}} \mid \mathbf{y})$ concerns an explanation, a fouled coil against a sound one. They can use the same readings but require different calculations: a tail of a posterior in one model against a comparison of evidences across models.

Common mistake

Ignoring the time horizon. $\Pr(T > 12\,\mathrm{h} \mid \mathbf{y})$ and $\Pr(T > 8760\,\mathrm{h} \mid \mathbf{y})$ are not interchangeable statements about the alarm; for the alarm log of Section 2.4.4 the first, a quiet night shift, is 0.9646 (0.9686 exact) and the second, a quiet year, is 0.00004 (exact, $(1500/10\,260)^4 = 0.0005$). A probability of a quiet horizon is incomplete without its horizon and its operating conditions.

Common mistake

Forgetting dependence. In general $\Pr(A \cap B) \neq \Pr(A)\,\Pr(B)$; the product is justified only under independence. A shared chiller, a shared supply-air set point, the weather and a shift's maintenance work create dependence between halls and units, and Table 2.8 shows that a small common cause can multiply the probability of two halls alarming in one night several times over.

Common mistake

Treating the model assumptions as certain. A calculated probability may be numerically precise while depending strongly on uncertain modeling assumptions. A temperature-exceedance probability computed under the steady-state model of equation (2.2) may be unsuitable for the minutes after the cooling unit trips, when the hall air rises through a transient that the steady-state model does not describe and the supply air is no longer held at T_c. Every probability is conditional on its model, $\Pr(A \mid \mathbf{y}, M)$, and must be interpreted in that context.

The questions of this chapter were all written with conditional probabilities and the posterior was taken as given; the next chapter, Chapter 3, derives conditional probability and Bayes' rule, separates aleatory variability from epistemic uncertainty, and shows how readings change the uncertainty about the state of the hall through marginalization, dependence and conditional independence on the plant graph.

2.7 What to retain

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

Table 2.12 lists the chapter's terms with their expressions, units and the subsection that defines each, grouped by their role.

Table 2.12: The terms of operational questions and where each is defined.

Term	Expression	Units	Meaning	Defined in
Event	A, e.g. $T > 28\,^{\circ}\mathrm{C}$	–	a statement true or false	2.2.1

Term	Expression	Units	Meaning	Defined in
Conditional probability	$\Pr(A \mid \mathbf{y}, M)$	–	probability given evidence and model	2.2.3
Event probability from a density	$\int_5^7 p(G \mid \mathbf{y})\,\mathrm{d}G$	–	probability of the range $5 < G < 7$	2.2.4
Probability over the state	$\int \mathbf{1}_A(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\,\mathrm{d}\mathbf{z}$	–	posterior expectation of the indicator	2.3.1
Posterior predictive	$\int \Pr(A_{\text{future}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\,\mathrm{d}\mathbf{z}$	–	future event averaged over the state	2.3.2
General forms				
Estimation	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$; $\mathbb{E}[f \mid \mathbf{y}]$, $\mathrm{Var}(f \mid \mathbf{y})$	–	posterior of a derived quantity	2.3.3
Credible interval	$\Pr(f \le f_L \mid \mathbf{y}) = 0.025$, $\Pr(f \le f_U \mid \mathbf{y}) = 0.975$	–	central 95 % band; quantiles map through monotone f	2.3.3
Threshold and assurance	$\Pr(f(\mathbf{z}) \ge f^* \mid \mathbf{y}) \ge \kappa$	–	probability that the requirement $f \ge f^*$ holds, against κ	2.3.3
Predictive	$\int p(y_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\,\mathrm{d}\mathbf{z}$	–	future observation averaged over the state	2.3.3
Comparison	$B_{AB} = p(\mathbf{y} \mid M_A)/p(\mathbf{y} \mid M_B)$	–	evidence ratio of two drawings	2.3.3
Diagnosis	$\Pr(M_j \mid \mathbf{y}) \propto p(\mathbf{y} \mid M_j)\ \Pr(M_j)$, ranked	–	comparison over a family of faults	2.3.3
Design	$U(a) = \mathbb{E}_{\tilde{y}_a \mid \mathbf{y}}[u(p(\mathbf{z} \mid \mathbf{y}, \tilde{y}_a))]$	varies	preposterior expectation	2.3.3
Preposterior covariance	$\boldsymbol{\Sigma}_a = (\boldsymbol{\Sigma}^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a)^{-1}$	varies	spread after a future measurement	2.3.3
Expected information gain	$\frac{1}{2}\ln(\det \boldsymbol{\Sigma} / \det \boldsymbol{\Sigma}_a)$	nat	what a measurement teaches	2.3.3
Trend	$\Pr(\gamma > 0 \mid \mathbf{y})$; $\Pr(t^* \le s \mid \mathbf{y}) = \Pr(\gamma > 0,\ f_s(\mathbf{z}) \ge f^* \mid \mathbf{y})$; $p(f_t(\mathbf{z}) \mid \mathbf{y})$; read with $\mathbf{g}(t)^{\mathsf{T}}\hat{\mathbf{z}}$, $\mathbf{g}(t)^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{g}(t)$	–	sign, crossing and forecast: the posterior read along time	2.3.4
Rows of a record in windows	$\mathbf{J} = \mathbf{D}\mathbf{X}$, $[\boldsymbol{\Lambda}]_{\gamma\gamma} = \sum_k d_k^2 (t_k - \bar{t})^2$	varies	a window's distance enters squared, windows add linearly	2.3.4
Expected consequence	$\mathcal{R}(a) = \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]$, $a^* = \arg\min_a \mathcal{R}(a)$; spread $\sqrt{\mathbf{g}_a^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{g}_a}$	varies	probability times consequence, minimized over actions	2.3.5
Curvature term	$L(a, \hat{\mathbf{z}}) + \frac{1}{2}\,\mathrm{tr}(\mathbf{H}_a\boldsymbol{\Sigma})$	varies	what the plug-in loss misses, in the state's coordinates	2.3.5
Computations				
Posterior	$p(\mathbf{z} \mid \mathbf{y}) \propto \mathrm{e}^{-\mathcal{L}(\mathbf{z})}$; $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$	varies	what estimation and prediction integrate	2.3.6
Evidence	$p(\mathbf{y} \mid M) = \int p(\mathbf{y} \mid \mathbf{z}, M)\, p(\mathbf{z} \mid M)\,\mathrm{d}\mathbf{z}$	–	normalizing constant of a model	2.3.6
Preposterior	$p(\mathbf{z} \mid \mathbf{y}, \tilde{y}_a)$ averaged over $p(\tilde{y}_a \mid \mathbf{y})$	varies	the expected posterior under a measurement not yet taken	2.3.6
Information adds	$\mathbf{J}^{\mathsf{T}}\mathbf{J} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a$	varies	stacked rows of a candidate	2.3.6
Update form	$\boldsymbol{\Sigma} - \boldsymbol{\Sigma}\mathbf{J}_a^{\mathsf{T}}(\mathbf{I} + \mathbf{J}_a\boldsymbol{\Sigma}\mathbf{J}_a^{\mathsf{T}})^{-1}\mathbf{J}_a\boldsymbol{\Sigma}$	varies	$\boldsymbol{\Sigma}_a$ when $\boldsymbol{\Sigma}$ is singular	2.3.6

Term	Expression	Units	Meaning	Defined in
Building blocks of the fit				
State	$\mathbf{z} = (T_1, T_2, g)^{\mathsf{T}}$	varies	every free quantity of the graph at once	2.4.1
Rows and objective	$r_k(\mathbf{z})/\sigma_k$; $\mathcal{L} = \frac{1}{2} r^{\mathsf{T}} r$	–	physics, reading and prior rows in standard deviations	2.4.1
Jacobian	$\mathbf{J} = \partial r / \partial \mathbf{z}$	varies	sensitivity of each row to each entry	2.4.1
Normal equations	$(\mathbf{J}^{\mathsf{T}}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^{\mathsf{T}} r$	varies	Gauss–Newton step to $\hat{\mathbf{z}}$	2.4.1
Laplace covariance	$\boldsymbol{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1}$	varies	band of every entry of the state	2.4.1
Ridge of exact rows	$\boldsymbol{\Sigma} \to \sigma_g^2 \mathbf{v}\mathbf{v}^{\mathsf{T}}$ as $\sigma_p \to 0$	varies	entries tied by the physics move as one	2.4.1
Band of a derived quantity	$\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$, $\mathbf{g} = \partial f / \partial \mathbf{z}$	varies	delta method	2.4.1
Specialized forms and answers				
Requirement probability of the hall	$\Phi\big((T_{\max} - \hat{z}_2)/\sqrt{\Sigma_{22}}\big) = 0.720$	–	reported, the normal tail of the entry	2.4.2
Its exact value	$\Phi\big((\hat{z}_3 - \ln(G^*/G_{\mathrm{ref}}))/\sqrt{\Sigma_{33}}\big) = 0.717$	–	the bound on T_2 mapped along the ridge to the third entry; gap 0.003	2.4.2
Alarm rate	$\hat{\lambda} \pm \hat{\lambda}\,\sigma_\omega = 0.002667 \pm 0.001333$	h^{-1}	reported, from the log-rate	2.4.3
Exact posterior of the rate	$\mathrm{Gamma}(a + k,\ b + T_{\mathrm{tot}})$	–, h	the same alarm log, conjugate	2.4.3
Quiet shift	$\sum_j w_j\, \mathrm{e}^{-\lambda_{\mathrm{ref}} \mathrm{e}^{\omega_j} \tau} = 0.9646$	–	reported, no activation averaged over the normal of the log-rate	2.4.4
Its exact value	$(b'/(b' + \tau))^{a'} = 0.9686$	–	the same average over the gamma posterior; gap 0.00406	2.4.4
Plug-in quiet shift	$\mathrm{e}^{-\mathbb{E}[\lambda \mid \mathbf{y}]\tau}$	–	the same at a point estimate; below the exact value	2.4.5
Both halls quiet	$1 - \Pr(A_A) - \Pr(A_B) + \Pr(A_A \cap A_B)$	–	inclusion and exclusion	2.4.6
Both halls alarming	$\Pr(A_A \mid \mathbf{y})\ \Pr(A_B \mid A_A, \mathbf{y})$	–	intersection; product rule; a shared chiller	2.4.7
Model posterior	$\Pr(M_j \mid \mathbf{y}) \propto p(\mathbf{y} \mid M_j)\ \Pr(M_j)$	–	probability of a model	2.5.1
Bayes factor of the air path	$B_{AB} = p(\mathbf{y} \mid M_1)/p(\mathbf{y} \mid M_0)$, $M_A = M_1$, $M_B = M_0$	–	relative support from the evidence	2.5.1
Expected loss of an action	$\mathcal{R}(a) = c_a + C_e\ \Pr(e \mid \mathbf{y}, a)$	R$	starting the standby unit against waiting	2.5.4
Break-even probability	$\Pr(e \mid \mathbf{y}, \mathrm{start}) + c_{\mathrm{start}}/C_e$	–	the probability at which acting and waiting cost the same	2.5.4
Optimal action	$a^* = \arg\min_a \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]$	–	decision by expected loss	2.5.5
Readings on the hall	$\boldsymbol{\Sigma}_a = \boldsymbol{\Sigma}/(1 + 0.21302\, n)$	varies	every entry shrinks along the ridge	2.6.1

Term	Expression	Units	Meaning	Defined in
Value of information	$\mathrm{EVSI} = \bar{L}_{\text{now}} - \bar{L}_{\text{after}}$	R$	what a field reading improves	2.6.2
The book's questions				
Drill questions	seventeen operator questions	–	general form, category, computation and chapter of each	2.3.8

2.8 Exercises

Exercise 2.1 (Formulation)**.** For each statement, name the category of the question, the expression that answers it and the computation beneath it. (a) What is the probability that the hall's alarm stays quiet through the next 12 h? ($\Pr(T > 12\,\mathrm{h} \mid \mathbf{y})$, a posterior predictive probability over the stated horizon; prediction under a threshold; the posterior.) (b) What is the alarm rate of the hall, given its alarm log? ($p(\lambda \mid \mathbf{y})$ or a summary of it; estimation; the posterior.) (c) Two halls alarmed in the same night; is a common cause, the chiller, more likely than two independent causes? ($\Pr(M_{\text{common}} \mid \mathbf{y})$ against $\Pr(M_{\text{independent}} \mid \mathbf{y})$, or their Bayes factor; hypothesis comparison; the evidence.) (d) Which hour of the evening should the standby unit be started at to minimize the expected cost of the night? ($\arg\min_a \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]$ over the candidate hours; a decision, in the design category; the expected consequence.) (e) Is the activation rate of an alarm rising, and in which period will it pass its rationalization target? ($\Pr(\gamma > 0 \mid \mathbf{y})$ and $\Pr(t^* \le s \mid \mathbf{y})$ over a state that carries the log-rate and its trend; prediction under a threshold; the trend.) (f) Of the questions A to D of Table 2.1, which two are answered from the same posterior without a new calculation of it? (A and B: an estimation and a threshold probability read from one posterior.)

Exercise 2.2 (Formulation)**.** An alarm has a constant activation rate λ with a known posterior, and the site's alarm-rationalization target asks for at least a 99 % probability of a quiet 24 h day, two shifts. (a) Formulate the posterior predictive probability of a quiet day and the probability that the alarm's underlying rate meets the target. ($\mathbb{E}[\mathrm{e}^{-24\lambda} \mid \mathbf{y}] = \int_0^\infty \mathrm{e}^{-24\lambda} p(\lambda \mid \mathbf{y})\,\mathrm{d}\lambda$; $\Pr(\mathrm{e}^{-24\lambda} \ge 0.99 \mid \mathbf{y}) = \Pr(\lambda \le -\ln 0.99/24 \mid \mathbf{y}) = \Pr(\lambda \le 4.188 \times 10^{-4}\,\mathrm{h}^{-1} \mid \mathbf{y})$.) (b) Evaluate both from the posterior of the state, $\hat{\omega} = 0.98083$ and $\sigma_\omega = 0.5$, and set beside each the exact value from the posterior $\mathrm{Gamma}(4, 1500\,\mathrm{h})$ of the hall's alarm, using for the second the integer-shape identity $\Pr(\lambda \le x \mid \mathbf{y}) = 1 - \sum_{j=0}^{a'-1} \mathrm{e}^{-b'x}(b'x)^j/j!$. (From the state, equation (2.41) with $\tau = 24\,\mathrm{h}$ gives 0.9307, and $\Phi\big((\ln 0.41876 - 0.98083)/0.5\big) = \Phi(-3.703) = 0.0001$. Exactly, $(1500/1524)^4 = 0.9385$, a gap of 0.0078; $b'x = 0.6281$ and the four Poisson terms 0.5336, 0.3352, 0.1053 and 0.0220 leave 0.0039, a gap of 0.0038.) (c) Explain why they differ. (The first averages the probability of no activation over the uncertainty in the rate and asks how likely tomorrow is to be quiet; the second asks how much posterior probability supports the target on the rate itself. Here a quiet day is likely, 0.93 and 0.94 exact, while the target is almost certainly not met by either route, because the posterior puts almost no mass below $4.2 \times 10^{-4}\,\mathrm{h}^{-1}$.)

Exercise 2.3 (Derivation)**.** The hall obeys $T = T_c + Q/G$ with $T_c = 18°\mathrm{C}$ and $Q = 60\,\mathrm{kW}$. (a) With the state, rows and $\mathbf{\Sigma}$ of equations (2.30) to (2.33), find the probability that the full-load temperature remains below 28°C from the normal of the entry T_2, then the exact value from the decreasing map between G and T_2 along the ridge. ($\Pr(T_2 < 28°\mathrm{C} \mid \mathbf{y}) \approx \Phi\big((28 - \hat{z}_2)/\sqrt{\Sigma_{22}}\big) = \Phi\big((28 - 27.231)/0.4615\big) = \Phi(1.667) = 0.952$. Exactly, along $\mathbf{v}$ the bound on T_2 is $G > Q_2/(28 - T_c) = 6\,\mathrm{kW/K}$, that is $g > \ln 6$ on the third entry, so $\Pr(T_2 < 28°\mathrm{C} \mid \mathbf{y}) = \int_6^\infty p(G \mid \mathbf{y})\,\mathrm{d}G = \Phi\big((\hat{z}_3 - \ln 6)/\sqrt{\Sigma_{33}}\big) = \Phi(1.601) = 0.945$; both are rows of Table 2.6. At $\kappa = 0.95$ the first accepts and the second does not, so both are stated.) (b) Now let the supply air be uncertain as well, as it is when the unit's controller is near saturation, so that

G and T_c have a joint posterior $p(G, T_c \mid \mathbf{y})$, and derive an integral for $\Pr(T > T_{\max} \mid \mathbf{y})$ without assuming that G and T_c are independent. (The event is $T_c + Q/G > T_{\max}$, so $\Pr(T > T_{\max} \mid \mathbf{y}) = \iint \mathbf{1}\big[T_c + Q/G > T_{\max}\big]\, p(G, T_c \mid \mathbf{y})\, \mathrm{d}G\, \mathrm{d}T_c$; the supply air becomes an entry of the state, and the joint density is required because the readings can leave it and the conductance correlated, a nonzero off-diagonal entry of $\mathbf{\Sigma}$.)

Exercise 2.4 (Derivation)**.** Two halls on one site have probabilities $\Pr(A_A) = 0.08$ and $\Pr(A_B) = 0.04$ that their temperature alarms fire during the night. (a) Assuming independence, derive the probability that both fire and the probability of a quiet night in both. ($\Pr(A_A \cap A_B) = 0.08 \times 0.04 = 0.0032$; $\Pr(\bar{A}_A \cap \bar{A}_B) = 0.92 \times 0.96 = 0.8832$.) (b) Now let both halls share a chiller that trips with probability $q_S = 0.01$ in the night, with the marginals unchanged. Derive the joint probability from equation (2.47), the conditional probability $\Pr(A_B \mid A_A)$ and the probability of a quiet night in both. ($q_A = 0.07/0.99 = 0.07071$, $q_B = 0.03/0.99 = 0.03030$; $\Pr(A_A \cap A_B) = 0.01 + 0.99 \times 0.07071 \times 0.03030 = 0.01212$; $\Pr(A_B \mid A_A) = 0.01212/0.08 = 0.152$; $\Pr(\bar{A}_A \cap \bar{A}_B) = 1 - 0.12 + 0.01212 = 0.89212$; the probability of both alarming is 3.8 times the independent value.)

Exercise 2.5 (Derivation)**.** The activations of the discharge-temperature alarm of a sister hall's older unit are counted in six periods of forty days, period j centered at $t = j$ so that the log ends at $t = 6.5$, the first period with 36 logged days, and number 5, 6, 7, 9, 11 and 14. The log activation rate, in activations per logged day, is carried as a straight line in time, $f_t(\mathbf{z}) = \alpha + \gamma(t - \bar{t})$ with t in periods and $\bar{t} = 3.5$, and the fit of the six count rows with flat priors gives $\hat{\mathbf{z}} = (-1.5768,\ 0.1989)^{\mathsf{T}}$ and $\mathbf{\Sigma}$ with entries $\Sigma_{\alpha\alpha} = 0.021811$, $\Sigma_{\alpha\gamma} = -0.004328$ and $\Sigma_{\gamma\gamma} = 0.007259$. The rationalization target is 0.5 activations a day, $f^* = \ln 0.5$. (a) From the trend form (2.21) under the Laplace posterior, find the probability that the rate is rising. ($\Phi(0.1989/\sqrt{0.007259}) = \Phi(2.335) = 0.990$.) (b) Find the crossing time at the MAP state and say why it is the median of the crossing time here. ($t^*(\hat{\mathbf{z}}) = 3.5 + (\ln 0.5 + 1.5768)/0.1989 = 7.94$ periods, $(7.94 - 6.5) \times 40 = 58$ days after the log ends; the second argument of Φ_2 in equation (2.22) is zero at $s = t^*(\hat{\mathbf{z}})$, so the threshold line is one half there, and the first passage subtracts from it the states that stand at or above the target with $\gamma \le 0$, a mass of 2.8×10^{-11}.) (c) Find the probability that the rate has crossed the target by $t = 8$. ($\mathbf{g}(8) = (1,\ 4.5)^{\mathsf{T}}$, $\mathbf{g}(8)^{\mathsf{T}}\hat{\mathbf{z}} = -0.6816$, $\mathbf{g}(8)^{\mathsf{T}}\mathbf{\Sigma}\,\mathbf{g}(8) = 0.021811 - 9 \times 0.004328 + 20.25 \times 0.007259 = 0.1299$, and the covariance of γ and f_8 is $\Sigma_{\alpha\gamma} + 4.5\,\Sigma_{\gamma\gamma} = 0.02834$, a correlation of 0.923. The threshold line is $\Phi\big((-0.6816 + 0.6931)/0.3604\big) = \Phi(0.032) = 0.513$, and the first passage is $\Pr(\gamma > 0,\ f_8 \ge f^* \mid \mathbf{y}) = \Phi_2(2.335,\ 0.032;\ 0.923) = 0.513$, less by 2.8×10^{-11}.) (d) Explain why the crossing time has no posterior mean. ($t^* = \bar{t} + (f^* - \alpha)/\gamma$ divides by the trend entry, whose Gaussian posterior has a positive density at $\gamma = 0$, and the integral of $1/\gamma$ against such a density diverges.)

Exercise 2.6 (Application)**.** The probability of an excursion at tonight's peak is 0.004 if the operator waits, and an excursion costs R$ 250 000. Raising the fan speed of the hall's units for the night costs R$ 400 and lowers the probability to 0.001. (a) Derive the expected cost of each action, and state which action minimizes it. (The expected consequence (2.23) of each action, probability times consequence plus the certain cost. Waiting: $0.004 \times$ R$ 250 000 = R$ 1000; acting: R$ 400 + $0.001 \times$ R$ 250 000 = R$ 650; acting minimizes the expected cost, with an expected saving of R$ 350.) (b) Find the cost of the action at which the decision would reverse. ($(0.004 - 0.001) \times$ R$ 250 000 = R$ 750.)

Exercise 2.7 (Application)**.** Two models are competing explanations of a hall sensor that reads warmer than the model of the hall expects, M_0 a sound sensor and M_1 a biased one, with $\Pr(M_0) = \Pr(M_1) = 0.5$ and evidences $p(\mathbf{y} \mid M_0) = 0.02$ and $p(\mathbf{y} \mid M_1) = 0.08$. (a) Derive the Bayes factor B_{10} and the posterior probability $\Pr(M_1 \mid \mathbf{y})$. ($B_{10} = 0.08/0.02 = 4$, $\ln B_{10} = 1.386$ nats; $\Pr(M_1 \mid \mathbf{y}) = 0.08 \times 0.5/(0.08 \times 0.5 + 0.02 \times 0.5) = 0.80$, assuming that the two models

exhaust the possibilities under consideration.) (b) Repeat with $\Pr(M_1) = 0.2$, and say what the comparison of the two answers shows. $(0.08 \times 0.2/(0.08 \times 0.2 + 0.02 \times 0.8) = 0.50$; the Bayes factor is the same 4 in both cases, because it depends only on the evidence, while the posterior model probability also depends on the prior.)

Chapter 3

Uncertainty and Conditional Probability

3.1 Uncertainty and evidence in operating questions

At 00:00 in the night shift, the coil-approach alarm of the hall's cooling unit, the alarm meant to catch a degrading unit, sounds. It watches the gap between the discharge air and the chilled water, a quantity the screen does not show. The hall carries its 30 kW, the hall sensor reads 23.2°C, as it has all evening, and the supply air holds at its 18°C set point. The alarm's data sheet says that it catches nine degradations of the unit in ten. Read from the screen, the conclusion is quick: the alarm is right nine times in ten, so the unit is degraded with probability 0.9, and the technician should be called out now. The operator also knows that a second hall, hall B, hangs on the same chiller, and wonders whether to watch hall B more closely too.

The data sheet does not settle either point. It states how often the alarm fires when the unit is degraded, which is not how often the unit is degraded when the alarm fires; it says nothing about how often degradations occur, nor about the alarms that fire with a healthy unit; and it says nothing about what the two halls share. Four situations of this kind recur through the chapter. The hall sensor scatters from window to window, and the probability that the next reading crosses the 24°C limit is wanted. The alarm has sounded, and the probability that the unit is degraded is wanted. Two halls draw chilled water from one chiller, and the probability that both exceed their limits in the same shift is wanted. The hall has one reading, and the probability that the hall air exceeds its limit is wanted. Chapter 1 represented the hall as a plant graph whose free quantities form the state vector $\mathbf{z}$, and Chapter 2 wrote the operator's questions as probabilities over that state, with the unknown on the left of the conditioning bar and the evidence on the right. The question of this chapter is the mechanism between the two:

how new evidence changes what is known about a running plant, and how dependencies between units and between physical quantities affect a probability calculation.

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

Aleatory uncertainty is the variability of outcomes that the model treats as random even when its parameters are known, such as the reading a sensor will show in the next window when the hall air is known. *Epistemic* uncertainty is incomplete knowledge of quantities that have one true value, such as the conductance of the hall's air path or the calibration offset of a meter; more evidence can reduce it. A *conditional probability* $\Pr(A \mid B)$ is the probability of event A among the situations in which event B occurs. In Bayes' theorem the *prior* is the probability of a hypothesis before an observation, the *likelihood* is the probability of the observation under the hypothesis, the *evidence* (also *marginal likelihood*) is the probability of the observation averaged over the hypotheses, and the *posterior* is the probability of the hypothesis after the observation. *Marginalization* removes an unknown from a joint distribution by summing or integrating over it; an unknown removed this way because the question does not ask about it is a *nuisance quantity*. Two quantities are *conditionally independent* given a third when, once the third is known, learning one says nothing more about the other. A *common cause* is a shared condition, such as warm chilled water from a chiller that feeds two halls, that raises the probability of an excursion in several places at once.

Key idea

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

A control room carries two kinds of not knowing: what the next window will show, and what the plant itself is. Keeping them apart changes the probability that a reading crosses a limit, and each kind has its own distribution.

3.1.1 Aleatory variability from window to window

Suppose the hall air were known exactly: take the estimate of Chapter 1, $T = 23.07\,^{\circ}\mathrm{C}$, as if it were the truth. Even then the reading of the next five-minute window cannot be predicted. Let y_T be the reading of one window. The sensor adds an error that changes from window to window, $y_T \mid T \sim \mathcal{N}(T, \sigma_T^2)$ with $\sigma_T = 0.5\,\mathrm{K}$, so the probability that the reading of a window shows above a level c is the integral of the normal density from c to infinity,

$$\Pr(y_T > c \mid T) = \int_c^{\infty} \frac{1}{\sigma_T}\,\varphi\!\left(\frac{s - T}{\sigma_T}\right) \mathrm{d}s = 1 - \Phi\!\left(\frac{c - T}{\sigma_T}\right), \tag{3.1}$$

where s is the variable of integration and φ and Φ are the standard normal density and cumulative distribution function. For the 24°C limit,

$$\Pr(y_T > 24\,^{\circ}\mathrm{C} \mid T = 23.07\,^{\circ}\mathrm{C}) = 1 - \Phi\!\left(\frac{24 - 23.0701}{0.5}\right) = 1 - \Phi(1.8598) = 0.0315.$$

With the hall air perfectly known there is still a probability 0.0315, one window in 32, that the reading shows the hall above its limit; over a twelve-hour shift of 144 windows a dashboard that colors a reading red above 24°C would color about 4.5 windows red with the hall inside its limit throughout. This uncertainty arises from the variability that the reading process itself represents, and it is called *aleatory*. The IT load is a second source of it: server utilization

moves the load from window to window, and a scatter of 1.5 kW about 30 kW (5 %, illustrative) moves the hall air itself by $1.5/5.917 = 0.2535$ K at the estimated conductance 5.917 kW/K. The reading then scatters by $\sqrt{0.5^2 + 0.2535^2} = 0.561$ K and crosses 24°C in 0.0486 of the windows. The rest of the chapter holds the load fixed within a window and keeps the sensor's scatter alone.

Figure 3.1(a) shows it. The horizontal axis numbers twelve windows of the same hall with the same known air temperature, and the vertical axis is the reading of each. The dots are the twelve quantiles of the normal distribution $\mathcal{N}(23.07, 0.5^2)$, placed in an arbitrary order; the dashed line marks the true hall air and the red line the limit. Identical windows of an identical, exactly known hall read from 22.20 to 23.94°C: knowing the hall fixes the distribution of the reading, not the reading.

3.1.2 Epistemic uncertainty about the plant

Now suppose the hall air is not known, because what fixes it at a given load, the effective conductance G of the air path to the supply air, is not. The handbook gives 7.5 kW/K, right to a factor of about 1.35. What is known about the conductance is then a probability distribution given the evidence, $p(G \mid \mathbf{y})$, or equivalently $p(g \mid \mathbf{y})$ for the log-conductance $g = \ln(G/G_{\text{ref}})$ with $G_{\text{ref}} = 6$ kW/K. The conductance might plausibly be $7.5/1.35 = 5.556$, 7.5 or $7.5 \times 1.35 = 10.125$ kW/K, with different probabilities, or a density, assigned to those values. This is *epistemic* uncertainty: uncertainty caused by incomplete knowledge of the characteristics of the plant, which additional observations can reduce. The calibration of a meter is the other common case. A load meter that reads high by a fixed, unknown offset reads high in every window; the offset has one true value, and the evidence that reduces what is known about it is a comparison with an independent measurement, not more windows of the same meter.

Figure 3.1(b) draws such a distribution. The horizontal axis is the conductance in kW/K and the vertical axis the prior density. The curve is the handbook prior of Chapter 1, $g \sim \mathcal{N}(\ln(7.5/6), 0.3^2)$, drawn as the lognormal density of G; the dashed line is its median, 7.5 kW/K, and the three dots are the three plausible values named above. The conductance is a single number; the spread of the curve is the spread of knowledge about it.

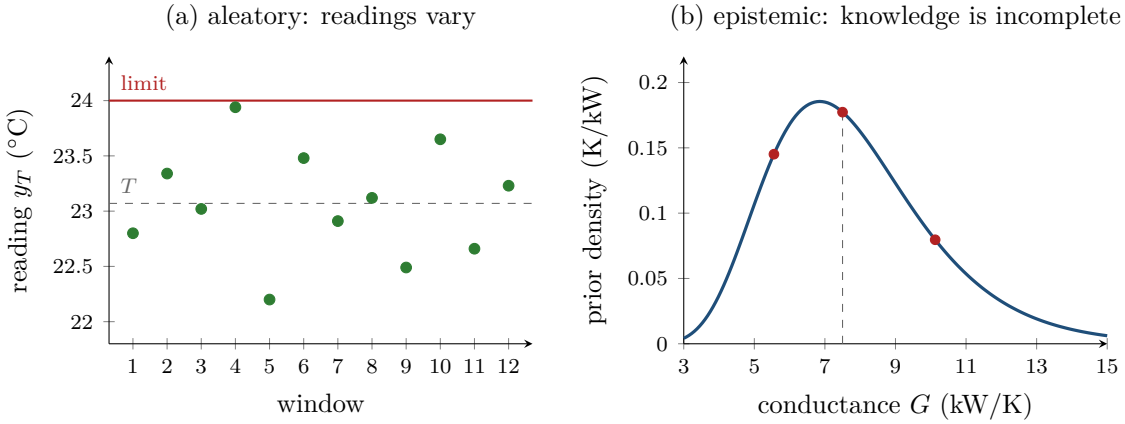


Figure 3.1: The two kinds of uncertainty. (a) Readings of twelve windows of a hall whose air is known, $T = 23.07$°C (dashed), with a sensor of standard deviation 0.5 K: the normal quantiles $T + 0.5\,\Phi^{-1}((i - 0.5)/12)$, $i = 1, \ldots, 12$, in an arbitrary order; red: the 24°C limit. (b) The handbook prior of the conductance, $\ln(G/6) \sim \mathcal{N}(\ln 1.25, 0.3^2)$, density $\exp[-\ln^2(G/7.5)/0.18]/(0.3\sqrt{2\pi}\,G)$; dashed: its median 7.5 kW/K; dots: the plausible values 5.556, 7.5 and 10.125 kW/K.

3.1.3 The distinction on the hall

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

If the hall air is known, the probability that the next reading shows above the limit is equation (3.1). If the hall air is uncertain, the probability that the next reading y_{new} crosses c is the conditional probability averaged over what is known about the hall, which is the law of total probability derived in Section 3.4.2:

$$\Pr(y_{\text{new}} > c \mid \mathbf{y}) = \int \Pr(y_{\text{new}} > c \mid T)\, p(T \mid \mathbf{y})\, \mathrm{d}T = \int \left[1 - \Phi\!\left(\frac{c - T}{\sigma_T}\right)\right] p(T \mid \mathbf{y})\, \mathrm{d}T. \tag{3.2}$$

For three plausible values the integral is a sum. Before any reading, the three conductances of Section 3.1.2 put the hall air at $T = 18 + 30/G = 23.40$, 22.00 and 20.96°C. With the illustrative probabilities 0.25, 0.50 and 0.25 on them, the probability that the next reading crosses 24°C is

$$\Pr(y_{\text{new}} > 24^\circ\text{C}) = 0.25 \times 0.1151 + 0.50 \times (3.2 \times 10^{-5}) + 0.25 \times (6 \times 10^{-10}) = 0.0288,$$

against 1.5×10^{-5} at the mean conductance 7.670 kW/K, where $T = 21.91$°C: the average is 1953 times the value at the average, and almost all of it comes from the low conductance. For a Gaussian posterior $T \mid \mathbf{y} \sim \mathcal{N}(\mu, \Sigma_{TT})$ the integral has a closed form. The next reading is $y_{\text{new}} = T + \epsilon$ with the error $\epsilon \sim \mathcal{N}(0, \sigma_T^2)$ independent of T, and the sum of two independent normal quantities is normal with the sum of their variances, a standard result, so that

$$y_{\text{new}} \mid \mathbf{y} \sim \mathcal{N}(\mu,\ \Sigma_{TT} + \sigma_T^2), \qquad \Pr(y_{\text{new}} > c \mid \mathbf{y}) = 1 - \Phi\left(\frac{c - \mu}{\sqrt{\Sigma_{TT} + \sigma_T^2}}\right). \tag{3.3}$$

After the one reading of Chapter 1, $\mu = 23.07$°C and $\sqrt{\Sigma_{TT}} = 0.481$ K, the predictive standard deviation is $\sqrt{0.4806^2 + 0.5^2} = 0.6935$ K and the probability is $1 - \Phi(1.341) = 0.0900$, against 0.0315 for the plug-in that takes 23.07°C as the truth. In the three-value sum the average of the tail exceeds the tail at the average because $1 - \Phi((c - T)/\sigma_T)$ is convex in T below the limit, where all three values lie. In the Gaussian case the reason is that the limit lies above the mean: widening the spread from 0.5 to 0.6935 K moves probability past it.

The scatter of the next reading remains even when the uncertainty about the hall is removed. Table 3.1 repeats the fit of Chapter 1 with n windows at 30 kW whose readings average 23.2°C, so that the state $(T, g)^{\mathsf{T}}$ carries one physics row, n reading rows and the handbook prior row. The first column is the number of windows, the second and third the posterior mean and standard deviation of the hall air, the fourth the predictive standard deviation of equation (3.3), and the last two the probability that the hall air exceeds 24°C and the probability that the next reading does. The epistemic spread halves each time the windows quadruple, and the probability that the hall is above its limit falls from 0.0265 to below 10^{-10}; the predictive spread falls only to the sensor's 0.5 K, and the probability that the next reading crosses the limit tends to $1 - \Phi(1.6) = 0.0548$, not to zero. The sensor's error is taken as independent from window to window, with no fixed offset; an offset would set a floor under $\sqrt{\Sigma_{TT}}$ that more windows cannot remove.

The table reports the Laplace values of the state's posterior. The exact posterior of the same rows, with T integrated out and a quadrature along g, sets the same two events beside them. The probability that the hall air exceeds 24°C is 0.0265 against the exact 0.0225 at $n = 1$, 3.76×10^{-4} against 3.27×10^{-4} at $n = 4$, and 4.43×10^{-11} against 3.88×10^{-11} at $n = 16$: the absolute gap closes while the relative gap barely moves, 14.9, 13.2 and 12.4 %. The probability that the next reading crosses is 0.0900 against 0.0813, 0.0674 against 0.0648, and 0.0583 against

0.0576. The source is the curve of the physics row: the Laplace normal is centered on the mode, 23.0701°C at $n = 1$, and the exact posterior has the mean 23.0300°C, so the exact tail above the limit is the thinner one.

Table 3.1: The hall at 30 kW after n windows whose readings average 23.2°C (the fit of Chapter 1 with n reading rows): the epistemic spread of the hall air shrinks, the scatter of the next reading does not.

n [–]	$\hat{T}$ [°C]	$\sqrt{\Sigma_{TT}}$ [K]	$\sqrt{\Sigma_{TT} + \sigma_T^2}$ [K]	$\Pr(T > 24°\text{C})$ [–]	$\Pr(y_{\text{new}} > 24°\text{C})$ [–]
1	23.070	0.4806	0.6935	0.0265	0.0900
4	23.166	0.2476	0.5580	3.8×10^{-4}	0.0674
16	23.191	0.1247	0.5153	4×10^{-11}	0.0583
64	23.198	0.0625	0.5039	$< 10^{-30}$	0.0557
∞	23.200	0	0.5000	0	0.0548

Principle 3.1. *Collecting more windows can reduce the uncertainty about the hall without eliminating the scatter of the next reading: evidence removes epistemic uncertainty, and the aleatory part remains as long as the reading is modeled as scattering from window to window.*

3.1.4 Distributions of what is uncertain

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

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

A continuous unknown. Suppose the effectiveness η of a unit's coil, the fraction of the largest possible heat transfer that it achieves, is unknown, and represent it by a density $p(\eta)$. A probability over a range is the integral of the density over that range,

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

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

A joint distribution. Suppose an illustrative hall with a weak air path has two unknown quantities, the conductance G of its air path and the supply-air temperature T_c, as it does when the discharge controller's own sensor is not trusted. Their *joint* distribution $p(G, T_c)$ describes the uncertainty about both together, including any dependence between them. The

joint distribution is essential when an operating quantity depends on several uncertain inputs, as the hall air does,

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

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

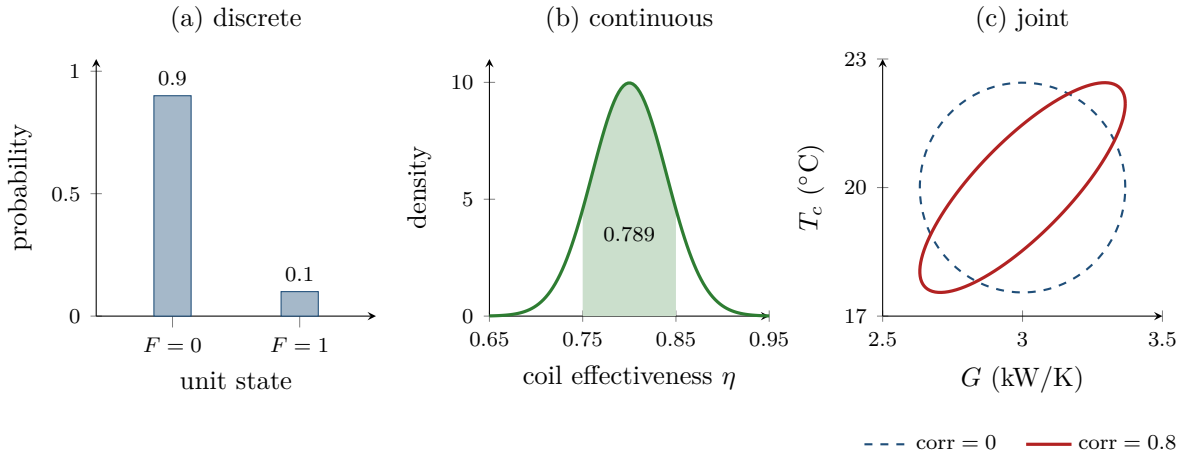


Figure 3.2: Distributions of what is uncertain (illustrative). (a) A discrete unknown: $\Pr(F = 1) = 0.10$ for a degraded cooling unit. (b) A continuous unknown: the density $\mathcal{N}(0.80, 0.04^2)$ of a coil effectiveness; shaded, $\Pr(0.75 < \eta < 0.85) = 0.789$. (c) Two joint distributions of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

3.2 Conditional probability and Bayes' theorem

Conditional probability is the foundation of every update in this book. An alarm log, read shift by shift, shows where it comes from, why the direction of the conditioning matters, and how Bayes' theorem reverses it.

3.2.1 Conditional probability from an alarm log

Let H be the event that a hall's high-temperature alarm fires during a shift and E the event that the hall air really exceeded its limit during the same shift, established afterwards from a review of independent readings. The question *given that the alarm fired, what is the probability that the excursion was real* is written $\Pr(E \mid H)$. The phrase "given that" restricts attention to the situations in which H occurred.

Imagine 1000 comparable hall-shifts of a site's alarm log, gathered over its halls and its nights and recorded in Table 3.2. The rows say whether the alarm fired, the columns whether a real excursion occurred, and the margins are the totals: alarm and excursion together in 20 shifts, the alarm alone in 60, the excursion alone in 20, and neither in 900. The shifts are

pooled over the site's halls, not drawn from this hall's log of Chapter 2: an alarm fires in 0.08 of these shifts, against a probability of 0.0354 of an alarm in a night for this hall, Chapter 2's value (0.0314 exact), 2.26 times as often (2.55 against the exact value).

Table 3.2: Outcomes of 1000 comparable hall-shifts of an alarm log (illustrative), in numbers of shifts.

	real excursion [–]	no excursion [–]	total [–]
alarm fired	20	60	80
no alarm	20	900	920
total	40	960	1000

From these counts,

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

Among the 80 shifts in which the alarm fired, the excursion was real in 20. Therefore

$$\Pr(E \mid H) = \frac{20}{80} = 0.25.$$

The conditional probability is 25 %, against an unconditional probability of 4 % that a shift has a real excursion: in these observations the alarm and the excursion are strongly associated, and still three alarms in four are false alarms.

3.2.2 The definition

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

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

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

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

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

Table 3.3 works the definition for four joint probabilities with the marginals of the log held fixed, $\Pr(H) = 0.08$ and $\Pr(E) = 0.04$. The first column is the probability that a shift has both an alarm and a real excursion, the next two are the conditional probabilities in the two directions, and the last says what kind of alarm the row describes. The third row is the log of Table 3.2: $\Pr(E \mid H) = 0.02/0.08 = 0.25$ and $\Pr(H \mid E) = 0.02/0.04 = 0.50$. The second row is the independent case, in which $\Pr(E \mid H) = 0.0032/0.08 = 0.04 = \Pr(E)$: an alarm that fires as often in quiet shifts as in hot ones carries no information. The joint probability cannot exceed the smaller marginal, 0.04; at that bound every real excursion raises the alarm.

Table 3.3: Conditional probabilities from equation (3.4) for four joint probabilities of an alarm and a real excursion, with $\Pr(H) = 0.08$ and $\Pr(E) = 0.04$ held fixed.

$\Pr(H \cap E)$ [–]	$\Pr(E \mid H)$ [–]	$\Pr(H \mid E)$ [–]	Kind of alarm
0	0	0	never fires with a real excursion
0.0032	0.04	0.08	independent of the excursion: no information
0.02	0.25	0.50	the logged alarm: positively associated
0.04	0.50	1.00	every real excursion raises the alarm

3.2.3 The multiplication rule

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

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

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

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

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

Reversing the conditioning. The two conditional probabilities of the log differ: $\Pr(E \mid H) = 0.25$, but $\Pr(H \mid E) = 20/40 = 0.50$. The first asks about the excursion among the shifts in which the alarm fired, the fraction of alarms that are real; the second asks about the alarm among the shifts with a real excursion, the alarm's probability of detection per shift. Half the real excursions are missed because they are brief or occur at racks the alarm's sensor does not see, and the alarm's on-delay merges the brief ones away. Equation (3.6) shows that they are equal only when $\Pr(H) = \Pr(E)$.

Common mistake

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

3.2.4 Bayes' theorem

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

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

For the log, the alarm's probability of detection and the base rates give the fraction of alarms that are real, $\Pr(E \mid H) = 0.50 \times 0.04/0.08 = 0.25$, the value read from the counts. In words,

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

and the four terms have distinct roles when A is a hypothesis about the plant and B an observation. The *prior* $\Pr(A)$ is what was believed about the hypothesis before the observation. The *likelihood* $\Pr(B \mid A)$ is how probable the observation is when the hypothesis holds. The *evidence* $\Pr(B)$ is the overall probability of the observation, averaged over the hypotheses; it is the normalizing constant of the posterior and, for a whole model M, the quantity the notation writes $p(\mathbf{y} \mid M)$. The *posterior* $\Pr(A \mid B)$ is what is believed about the hypothesis after the observation is accounted for. The operator's question "this alarm just fired; what is the probability it is real" is this computation, and Chapter 20 answers it on the shift plant.

3.3 Bayesian diagnosis of a cooling fault

Bayes' theorem turns the known behavior of an alarm into the probability of the fault the alarm is meant to detect, and the answer of the opening scene is far below the alarm's probability of detection; a silent alarm is evidence too.

3.3.1 The diagnosis problem

The hall's cooling unit carries a coil-approach alarm, the alarm of the opening scene, designed to detect degradation of the unit. Define the events

$$F = \{\text{the cooling unit is degraded}\}, \qquad H = \{\text{the coil-approach alarm occurs}\}.$$

The maintenance log suggests that 5 % of shifts have a degraded unit, $\Pr(F) = 0.05$. The alarm activates in 90 % of degraded conditions, $\Pr(H \mid F) = 0.90$, its probability of detection; it also produces false alarms, in 10 % of healthy shifts, $\Pr(H \mid \neg F) = 0.10$, its false-alarm probability; an alarm this noisy is what alarm rationalization, Chapter 19, targets. An alarm is observed, and the question is the probability that the unit is in fact degraded, $\Pr(F \mid H)$.

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

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

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

3.3.2 The total probability of an alarm

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

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

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

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

Alarms occur in 14 % of shifts; two thirds of them, 0.095 of the 0.14, come from shifts with a healthy unit.

3.3.3 The posterior fault probability

The third step divides,

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

The call-out of the opening scene therefore does not rest on a probability of 0.9: at 0.32 a healthy unit is the likelier explanation, and whether a technician is worth sending depends on what a visit costs against what a missed degradation costs. Figure 3.3 compares the probability of degradation before and after the observation. Each row is one state of knowledge, the horizontal scale is a probability from 0 to 1, and the green part of each bar is the probability of degradation. The alarm raises it from 5 % to 32.1 %; it increases suspicion substantially, but it does not establish that a fault is present. The third row is the case without an alarm, derived in Section 3.3.4.

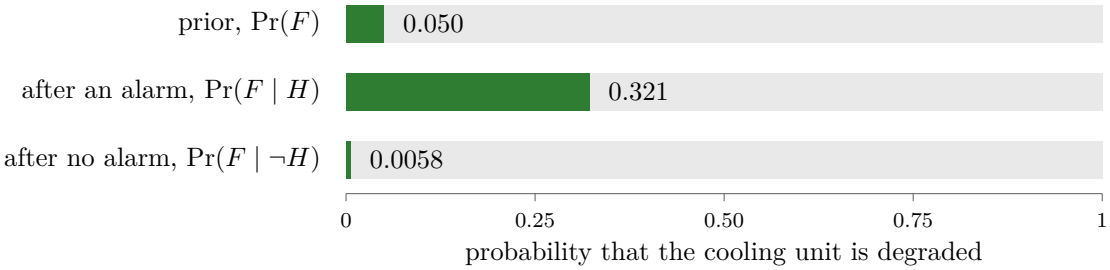


Figure 3.3: What an alarm tells about degradation. Prior $\Pr(F) = 0.05$ per shift; alarm with probability of detection $\Pr(H \mid F) = 0.90$ and false-alarm probability $\Pr(H \mid \neg F) = 0.10$. After an alarm, $0.045/0.14 = 0.321$ (equation (3.9)); after no alarm, $0.005/0.86 = 0.0058$ (equation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

The base rate. The answer is not 90 %. The alarm detects 90 % of degraded conditions, but degraded conditions are uncommon, and many alarms occur while the unit is healthy. Table 3.4 makes this concrete for 1000 shifts. The rows are the actual condition of the unit, the columns whether the alarm sounded, and each entry is the expected number of shifts: of the 50 degraded shifts, 90 %, or 45, raise the alarm; of the 950 healthy shifts, 10 %, or 95, raise it falsely. Among the 140 shifts with an alarm only 45 have a degraded unit, so $\Pr(F \mid H) = 45/140 = 0.3214$. The table and Bayes' theorem give the same result, because the table is equation (3.8) multiplied by 1000.

Table 3.4: Expected outcomes of 1000 shifts with $\Pr(F) = 0.05$, $\Pr(H \mid F) = 0.90$ and $\Pr(H \mid \neg F) = 0.10$, in numbers of shifts.

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

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

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

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

3.3.4 The case without an alarm

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

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

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

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

Therefore

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

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

Interpretation. Bayesian diagnosis does not only ask whether a reading crossed a threshold. It evaluates how strongly the observation supports competing explanations, weighted by their prior plausibility. The distinction is central in a control room because many abnormal readings are not uniquely associated with one physical fault: a high hall temperature can be caused by a reduced cooling capacity, an increased load, a warmer supply air or a biased sensor. Each of these is a hypothesis with its own prior and its own likelihood for the same reading, and the inference becomes a comparison among them, the diagnosis category of Chapter 2, which Chapter 10 develops as fault detection and isolation.

3.4 Marginalization and dependence

The probability of an alarm in Section 3.3.2 was computed without knowing whether the unit was degraded. The same operation, for discrete and continuous unknowns and for nuisance quantities, is what lets a shared chiller make two halls dependent, and what an independence assumption costs.

3.4.1 Discrete marginalization

Suppose a cooling unit (illustrative) can be in one of three operating conditions, healthy, degraded or failed, and let S denote the condition present. Table 3.5 gives, for each state, its probability and the probability that the alarm H occurs in it; the last two columns are computed below.

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

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

Table 3.5: Three operating conditions of a cooling unit: probabilities, alarm likelihoods, the joint probabilities of state and alarm, and the probabilities of the states given an alarm.

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

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

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

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

3.4.2 Continuous marginalization

When the unknown is continuous, the sum becomes an integral. Let the rate λ at which the hall's high-temperature alarm fires be uncertain and let $\mathbf{y}$ be the observed evidence; their joint distribution is $p(\lambda, \mathbf{y}) = p(\mathbf{y} \mid \lambda)\, p(\lambda)$. The probability density of the observations is obtained by integrating over every possible value of the rate,

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

the *marginal likelihood* of the observations. Bayes' theorem for a continuous unknown then reads

$$p(\lambda \mid \mathbf{y}) = \frac{p(\mathbf{y} \mid \lambda)\, p(\lambda)}{\int p(\mathbf{y} \mid \lambda')\, p(\lambda')\, \mathrm{d}\lambda'}, \tag{3.14}$$

where λ' is a dummy variable of integration. The numerator evaluates the relative plausibility of one particular rate; the denominator ensures that the posterior integrates to one.

For the hall's alarm log of Chapter 2, $k = 2$ alarms in $T_{\mathrm{tot}} = 1000\,\mathrm{h}$ of logged operation under the Poisson model $\Pr(k \mid \lambda) = (\lambda T_{\mathrm{tot}})^k \mathrm{e}^{-\lambda T_{\mathrm{tot}}}/k!$ and the prior $\mathrm{Gamma}(a, b) = \mathrm{Gamma}(2, 500\,\mathrm{h})$, with density $b^a \lambda^{a-1} \mathrm{e}^{-b\lambda}/\Gamma(a)$, the standard gamma integral $\int_0^\infty \lambda^{a-1} \mathrm{e}^{-s\lambda}\, \mathrm{d}\lambda = \Gamma(a)/s^a$ with $s = b + T_{\mathrm{tot}}$ evaluates equation (3.13) in closed form:

$$\Pr(k \mid T_{\mathrm{tot}}) = \frac{T_{\mathrm{tot}}^k}{k!} \frac{b^a}{\Gamma(a)} \frac{\Gamma(a+k)}{(b+T_{\mathrm{tot}})^{a+k}} = \frac{1000^2}{2} \cdot \frac{500^2}{1} \cdot \frac{3!}{1500^4} = 0.1481.$$

Before the log, the model gave a probability of 0.148 to exactly two alarms in 1000 h; dividing the numerator of equation (3.14) by it gives the posterior $\mathrm{Gamma}(4, 1500\,\mathrm{h})$. The probability of a quiet night shift, no alarm in the next $\tau = 12\,\mathrm{h}$, is the same operation applied to the event of no alarm: its conditional probability $\mathrm{e}^{-\lambda\tau}$ averaged over the posterior of the rate. Combining the two exponentials and using the same gamma integral with $s = b + \tau$, for the posterior $\mathrm{Gamma}(a, b)$,

$$\Pr(\text{no alarm in } \tau \mid \mathbf{y}) = \frac{b^a}{\Gamma(a)} \int_0^\infty \lambda^{a-1}\, \mathrm{e}^{-(b+\tau)\lambda}\, \mathrm{d}\lambda = \frac{b^a}{\Gamma(a)} \frac{\Gamma(a)}{(b+\tau)^a} = \left(\frac{b}{b+\tau}\right)^a, \tag{3.15}$$

which for Gamma(4, 1500 h) is $(1500/1512)^4 = 0.96863$. Chapter 2 reports this probability from the posterior of the state, the same average over the normal of the log-rate by the five-node sum of equation (2.41): 0.96457, with the gamma integral 0.96863 as its exact value beside it, a gap of 0.0041. An alarm in the night is then 0.0354 (0.0314 exact). The plug-in at the posterior mean rate, $\mathrm{e}^{-12\times4/1500} = 0.9685$, sits 0.0001 from the exact value and answers a different question, an alarm whose rate is known. The predictive probability of equation (3.2) is the same operation applied to the next reading: the conditional probability averaged over the posterior of the hall air.

3.4.3 Nuisance quantities

Suppose the log-conductance g of the hall is being estimated from the hall reading and the load meter, but the model also contains the uncertain calibration offset b_Q of the load meter, in kW. The full posterior is $p(g, b_Q \mid \mathbf{y})$, and the operator wants the distribution of g alone. It is obtained by integrating b_Q out,

$$p(g \mid \mathbf{y}) = \int p(g, b_Q \mid \mathbf{y}) \, \mathrm{d}b_Q. \tag{3.16}$$

The offset does not have to be treated as known; its uncertainty enters the answer through the integration. This matters most in physics-constrained inference, where many unknown state entries influence the quantity studied: the marginal of one entry z_i is the posterior of the whole state $\mathbf{z}$ integrated over all its other entries.

Marginalization against optimization. These two operations are different and must not be confused. Marginalization, equation (3.16), accumulates probability across all values of b_Q. Optimization over the nuisance quantity, $\max_{b_Q} p(g, b_Q \mid \mathbf{y})$, selects the best-fitting offset for each conductance. They generally give different results. A discrete illustration makes the difference plain. Let the conductance be either low or high and the offset take one of three values, with joint posterior probabilities 0.30, 0.05, 0.05 for a low conductance and 0.20, 0.20, 0.20 for a high one (illustrative). Marginalization gives Pr(low) = 0.40 and Pr(high) = 0.60; maximization over b_Q gives 0.30 against 0.20 and prefers the low conductance. The low conductance owes its single peak to one value of the offset; the high one is supported by every value. The distinction returns when later chapters compare a maximum a posteriori estimate with the marginals of the posterior.

3.4.4 Two halls on one chiller

Many capacity and alarm calculations assume that the halls of a site behave independently. Let F_A and F_B be the events that hall A and hall B exceed their limits in the same shift. If each has probability 0.05, independence gives $\Pr(F_A \cap F_B) = 0.05^2 = 0.0025$ for the probability that both are hot at once. The conclusion is invalid if the halls share an important cause. Consider two halls whose cooling units draw chilled water from one chiller. Their supply airs are taken as equal, and warm chilled water lifts them only when the discharge controllers saturate. Let C be the event that the chilled water runs warm during a shift, through a chiller trip or a staging delay, with $\Pr(C) = 0.05$. The event C counts every shift in which the chilled water runs warm enough to lift the supply air, staging delays and partial trips as well as trips, so it is more frequent than the chiller trips q_S of the two-halls example of Chapter 2; and it raises each hall's excursion probability, to 0.60 and 0.50 below, rather than forcing both. Figure 3.4 draws the dependency: the box at the top is the shared condition, the two boxes below are the excursion events of the halls, and each arrow says that the probability of the event at its head depends on the condition at its tail. There is no arrow between the halls.

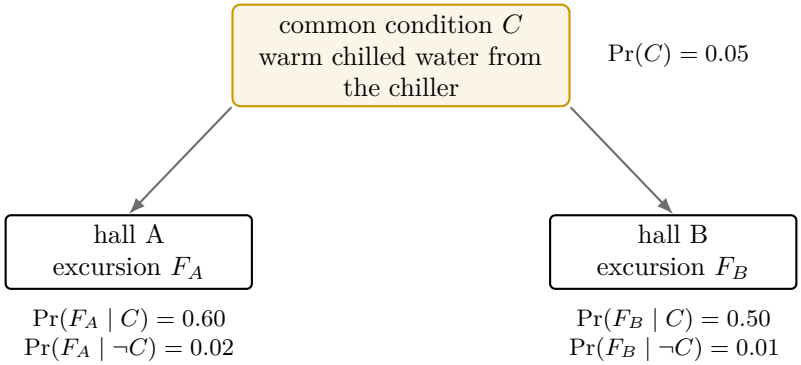


Figure 3.4: A probabilistic dependency graph for two halls on one chiller. The shared condition C, warm chilled water and with it a warm supply air in both halls, influences both excursion events; the absence of an arrow between the halls states that the excursions are conditionally independent given C. The probabilities are those of Table 3.6.

The excursion probabilities of the halls depend on whether the chilled water runs warm. Table 3.6 lists them: the rows are the two conditions and the columns the probability that each hall exceeds its limit in a shift under that condition.

Table 3.6: Probabilities that each hall exceeds its limit in a shift, with and without warm chilled water from the shared chiller (illustrative); $\Pr(C) = 0.05$.

Condition	$\Pr(F_A \mid \cdot)$ [–]	$\Pr(F_B \mid \cdot)$ [–]
warm chilled water, C	0.60	0.50
normal chilled water, $\neg C$	0.02	0.01

Assume that the two excursions are *conditionally independent given* C:

$$\Pr(F_A \cap F_B \mid C) = \Pr(F_A \mid C)\,\Pr(F_B \mid C), \qquad \Pr(F_A \cap F_B \mid \neg C) = \Pr(F_A \mid \neg C)\,\Pr(F_B \mid \neg C).$$

Once the condition of the chilled water is known, an excursion in one hall says nothing more about the other. This does not mean that the excursions are unconditionally independent, as the next subsection shows.

3.4.5 The joint excursion probability and the independence error

The law of total probability, equation (3.12), applied to the event $F_A \cap F_B$ with the two conditions C and $\neg C$ as the states, gives

$$\Pr(F_A \cap F_B) = \Pr(F_A \cap F_B \mid C)\,\Pr(C) + \Pr(F_A \cap F_B \mid \neg C)\,\Pr(\neg C).$$

Under conditional independence each conditional joint probability is a product,

$$\Pr(F_A \cap F_B) = 0.60 \times 0.50 \times 0.05 + 0.02 \times 0.01 \times 0.95 = 0.015 + 0.00019,$$

so that

$$\Pr(F_A \cap F_B) = 0.01519. \tag{3.17}$$

The probability that both halls exceed their limits in the same shift is about 1.52 %, and almost all of it, 0.015 of 0.01519, comes from the chiller.

The marginal excursion probabilities follow from the same law,

$$\Pr(F_A) = 0.60 \times 0.05 + 0.02 \times 0.95 = 0.049, \qquad \Pr(F_B) = 0.50 \times 0.05 + 0.01 \times 0.95 = 0.0345.$$

If the excursions are incorrectly treated as independent, the joint probability is the product of the marginals,

$$\Pr(F_A \cap F_B) \overset{\text{incorrect}}{=} 0.049 \times 0.0345 = 0.00169,$$

or 0.169 %. The correct value, 1.519 %, is 8.99 times larger: ignoring the shared chiller underestimates the probability that both halls are hot at once by a factor of about nine in this illustrative site, and the independent product is only 11 % of the correct value. With marginals of the same size, 0.049 and 0.0345 here against 0.05 and 0.03 there, this factor of nine sits beside the factor of 7.2 that Chapter 2 finds for two halls both alarming at a trip probability $q_S = 0.01$.

Common mistake

Multiplying the excursion probabilities of two halls to obtain the probability that both are hot at once. The product is the joint probability only when the excursions are independent; a shared chiller, a shared outdoor condition or a shared control sequence makes them dependent, and the night on which both halls need the operator at the same time is much more likely than the product suggests.

Principle 3.2. *Understanding what the units and halls of a plant share can matter more than estimating the excursion probability of each one with greater precision.*

3.4.6 Conditional independence given the supply air

Two random variables X_1 and X_2 are conditionally independent given C when, for every value c of C,

$$p(x_1, x_2 \mid c) = p(x_1 \mid c)\, p(x_2 \mid c), \tag{3.18}$$

written $X_1 \perp X_2 \mid C$. Dividing both sides by $p(x_1 \mid c)$ and using the definition of conditional probability on the left gives the equivalent statement $p(x_2 \mid x_1, c) = p(x_2 \mid c)$: after conditioning on C, learning X_1 gives no additional information about X_2 under the assumed model. The two halls in physical form are the case: with the two supply airs taken as equal, at T_c, the hall air temperatures $T_A = T_c + Q_A/G_A$ and $T_B = T_c + Q_B/G_B$ depend on separate loads and conductances, so that given T_c they are independent; with T_c unknown, a reading in hall A informs T_c and through it hall B.

When C is unknown, observing X_1 informs C, and that changes what is believed about X_2. If hall A exceeds its limit, the chiller becomes suspect: by Bayes' theorem,

$$\Pr(C \mid F_A) = \frac{\Pr(F_A \mid C)\ \Pr(C)}{\Pr(F_A)} = \frac{0.60 \times 0.05}{0.049} = 0.612,$$

up from 0.05. The probability that hall B also exceeds its limit is then the law of total probability conditioned on F_A, with conditional independence used to drop F_A from $\Pr(F_B \mid C, F_A)$:

$$\begin{aligned}\Pr(F_B \mid F_A) &= \Pr(F_B \mid C)\ \Pr(C \mid F_A) + \Pr(F_B \mid \neg C)\ \Pr(\neg C \mid F_A) \\ &= 0.50 \times 0.612 + 0.01 \times 0.388 = 0.310,\end{aligned}$$

which equals $\Pr(F_A \cap F_B)/\Pr(F_A) = 0.01519/0.049$, against the unconditional $\Pr(F_B) = 0.0345$. An excursion in one hall raises the excursion probability of the other ninefold, through the chiller they share: the operator who sees an excursion in hall A has reason to look at hall B. The alarm of the opening scene watches one unit's coil and does not raise C, so it gives no such reason unless the chiller itself is suspect. Conditional independence is therefore not unconditional independence: the variables are independent only once the common cause is known.

3.5 Conditional independence and a continuous update in a physical model

The same rules apply when the unknowns are physical quantities tied by a conservation relation. On the equipment hall they give the posterior of a state vector under Gaussian rows, the probability that the hall exceeds its limit after one reading, and, in the two-entry hall of Chapter 1, a reading that reaches the conductance through the physics row.

3.5.1 The hall as a probability model

The hall of the earlier chapters satisfies the steady heat balance

$$T = T_c + \frac{Q}{G}, \tag{3.19}$$

with the hall air T and the supply air T_c in °C, the heat load Q in kW and the conductance G of the air path in kW/K. A sensor reads the hall air with an additive error,

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

where ϵ is the measurement error and σ_T its standard deviation in K. Figure 3.5 draws the system twice. Panel (a) is the plant graph: the source supplies the load Q to the hall node, which carries the free entry T; the green path to the supply-air node carries the conductance, held in the state as the log-conductance $g = \ln(G/G_{\text{ref}})$; the supply-air node is pinned at T_c by the unit's discharge controller; and the dashed arrow is the reading y_T. Panel (b) is the dependency graph of the probability model of this section, in which Q, G and T_c are all treated as uncertain: the load, the conductance and the supply air determine the modeled hall air, and the reading depends on that temperature alone.

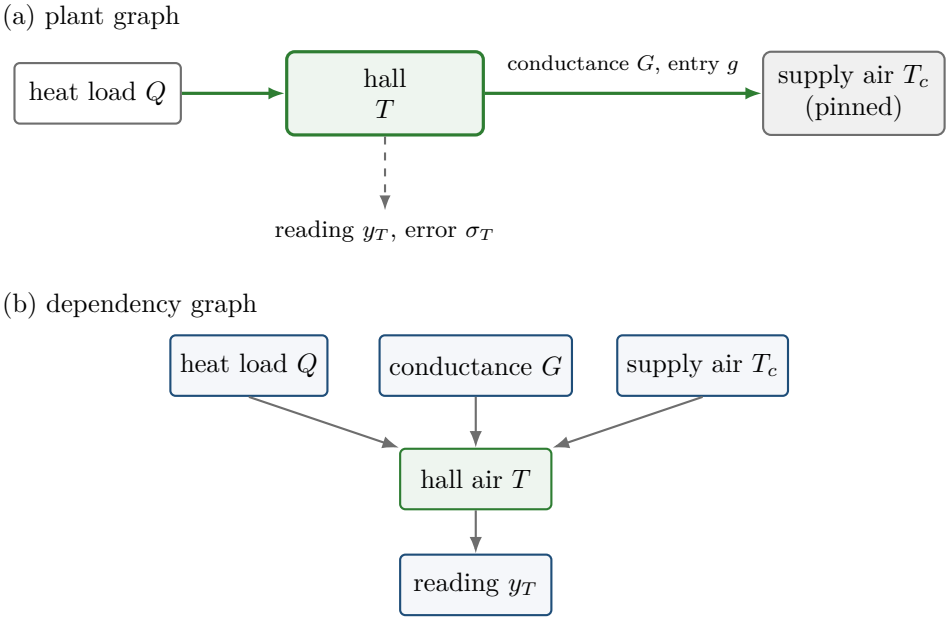


Figure 3.5: The hall in two drawings. (a) The plant graph: free entry T at the hall node and $g = \ln(G/G_{\text{ref}})$ on the air path; supply air T_c pinned by the discharge controller; source Q; reading y_T observed (dashed). (b) The dependency graph of the factorized model of equation (3.21), in which Q, G and T_c carry priors: arrows point from a quantity to what depends on it, and the reading depends on the inputs only through T.

If the measurement error is independent of the physical inputs, the reading is conditionally independent of those inputs once the true temperature is known:

$$p(y_T \mid T, G, Q, T_c) = p(y_T \mid T). \tag{3.20}$$

This does not make the reading uninformative about the conductance. The reading informs T, and the physical relation (3.19) connects T to G: evidence propagates from the sensor reading to the temperature and from the temperature to the conductance. An observation of one quantity updates the beliefs about every quantity connected to it through shared physical constraints, and this is one of the main reasons why inference over a graph is useful.

The factorized probability model. Suppose G, Q and T_c are uncertain inputs with a joint prior $p(G, Q, T_c)$, which relaxes the pins of the plant graph and carries Q and T_c as free entries. The chain rule, the multiplication rule applied twice, and equation (3.20) give the joint model

$$p(G, Q, T_c, T, y_T) = p(G, Q, T_c)\, p(T \mid G, Q, T_c)\, p(y_T \mid T). \tag{3.21}$$

The physics factor is written as a Gaussian around the heat balance, $p(T \mid G, Q, T_c) = \mathcal{N}(T;\ T_c + Q/G,\ \sigma_p^2)$, with a model tolerance σ_p in K: a physical discrepancy is a finite σ_p, and the exact deterministic balance of equation (3.19) is the limit $\sigma_p \to 0$, in which the factor becomes the Dirac delta $\delta(T - T_c - Q/G)$. The joint distribution is the complete probability model, and the posterior wanted is $p(G, Q, T_c, T \mid y_T)$, over the whole state with T kept in it, which answers questions about the conductance, the temperature, the operating margins and every other quantity of the model. Each of its three factors is a row over that state: the prior, the physics row with tolerance σ_p and the reading row; the derivations of this section keep all of them as rows and never remove an entry by substitution. In the limit $\sigma_p \to 0$, integrating T out against the delta gives the same posterior of the inputs as $p(G, Q, T_c \mid y_T) \propto p(G, Q, T_c)\, p\big(y_T \mid T = T_c + Q/G\big)$: every combination of inputs is weighted by how well the temperature it implies explains the reading.

3.5.2 Prior and reading for the hall air

The simplest continuous case has one unknown, the true hall air temperature T, so that the state is this one number,

$$T \in \mathbb{R}.$$

The general forms of the book write the state as a vector $\mathbf{z}$; here it has the single entry T, so every quantity of the fit below is a number: each row's residual and its slope in T, the objective's slope and curvature, the information and the posterior variance. On an illustrative hall running warmer than the series' hall, a commissioning model gives an uncertain prior estimate,

$$T \sim \mathcal{N}(\mu_0, \sigma_0^2) = \mathcal{N}(26, 2^2), \qquad \mu_0 = 26^\circ\mathrm{C}, \quad \sigma_0 = 2\,\mathrm{K}.$$

A sensor reports $y_T = 27^\circ$C, and its error is Gaussian with $\sigma_T = 0.5\,\mathrm{K}$, so that $y_T \mid T \sim \mathcal{N}(T, 0.5^2)$. The question is the posterior distribution of the true temperature after the reading. Bayes' theorem for a continuous unknown, with the evidence dropped because it does not depend on the state, gives

$$p(T \mid y_T) \propto p(y_T \mid T)\, p(T),$$

with the Gaussian prior and the Gaussian likelihood

$$p(T) \propto \exp\left[-\frac{(T - \mu_0)^2}{2\sigma_0^2}\right], \qquad p(y_T \mid T) \propto \exp\left[-\frac{(y_T - T)^2}{2\sigma_T^2}\right].$$

Multiplying,

$$p(T \mid y_T) \propto \exp\left[-\frac{1}{2}\left(\frac{(T-\mu_0)^2}{\sigma_0^2} + \frac{(y_T - T)^2}{\sigma_T^2}\right)\right]. \tag{3.22}$$

Each term of the exponent is the square of a difference divided by its standard deviation. The two differences are the residuals of the problem, one prior row and one reading row, each already divided by its standard deviation,

$$r_1(T) = \frac{T-\mu_0}{\sigma_0}, \qquad r_2(T) = \frac{T-y_T}{\sigma_T}, \tag{3.23}$$

and equation (3.22) reads $p(T \mid y_T) \propto \exp[-\frac{1}{2}(r_1^2 + r_2^2)] = \mathrm{e}^{-\mathcal{L}(T)}$, with the objective $\mathcal{L}(T) = \frac{1}{2}(r_1^2 + r_2^2)$ of the notation chapter. Both rows are linear in T, each with a constant slope $j_k = \mathrm{d}r_k/\mathrm{d}T$: $j_1 = 1/\sigma_0$ and $j_2 = 1/\sigma_T$.

Gaussian rows over any state. The same step holds for a state of any size and any number of rows. The joint model factorizes by the chain rule, as equation (3.21) does for the hall: a prior factor over the inputs, a physics factor for each physical relation, and a reading factor for each reading, which by conditional independence, equation (3.20), depends on the state only through the quantity read. Row k is one such factor written as a residual: the difference between the quantity computed from the state and what was observed, or what the prior or the physics expects, divided by its standard deviation σ_k. A Gaussian factor is $\exp[-\frac{1}{2}r_k(\mathbf{z})^2]$ times a normalizer, $1/(\sqrt{2\pi}\,\sigma_k)$, that is free of the state. Multiplying the factors and dropping the normalizers and the evidence, Bayes' theorem gives

$$p(\mathbf{z} \mid \mathbf{y}) \propto \prod_k \exp\left[-\tfrac{1}{2}r_k(\mathbf{z})^2\right] = \exp\left[-\tfrac{1}{2}r(\mathbf{z})^{\mathsf{T}} r(\mathbf{z})\right] = \mathrm{e}^{-\mathcal{L}(\mathbf{z})}, \tag{3.24}$$

where $r(\mathbf{z})$ stacks the reading rows, the prior rows and the physics rows with their tolerances σ_p. For the one-entry hall the product has the two factors of equation (3.22) and the sum of squares is the number $r_1^2 + r_2^2$; the vector is needed when the state has several entries, as in Section 3.5.6 and in the hall of Chapter 4. A Gaussian prior over several entries, $\mathcal{N}(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$ with correlated entries, enters as the rows $\mathbf{L}_0(\mathbf{z} - \boldsymbol{\mu}_0)$ with $\mathbf{L}_0^{\mathsf{T}}\mathbf{L}_0 = \boldsymbol{\Sigma}_0^{-1}$, since then $(\mathbf{z} - \boldsymbol{\mu}_0)^{\mathsf{T}}\boldsymbol{\Sigma}_0^{-1}(\mathbf{z} - \boldsymbol{\mu}_0)$ is the sum of their squares.

3.5.3 Completion of the square in the state

Suppose the rows of a state with several entries are linear in it, as both rows of equation (3.23) are in T. Then, for any point $\mathbf{z}_0$,

$$r(\mathbf{z}) = r(\mathbf{z}_0) + \mathbf{J}\,(\mathbf{z} - \mathbf{z}_0)$$

holds exactly, with the Jacobian $\mathbf{J} = \partial r/\partial \mathbf{z}$, a matrix of m rows and n columns for m rows and n entries. Write $\boldsymbol{\Delta} = \mathbf{z} - \mathbf{z}_0$ and $r_0 = r(\mathbf{z}_0)$. The objective expands as

$$\mathcal{L}(\mathbf{z}) = \tfrac{1}{2}(r_0 + \mathbf{J}\boldsymbol{\Delta})^{\mathsf{T}}(r_0 + \mathbf{J}\boldsymbol{\Delta}) = \tfrac{1}{2}r_0^{\mathsf{T}}r_0 + \boldsymbol{\Delta}^{\mathsf{T}}\mathbf{J}^{\mathsf{T}}r_0 + \tfrac{1}{2}\boldsymbol{\Delta}^{\mathsf{T}}\mathbf{J}^{\mathsf{T}}\mathbf{J}\,\boldsymbol{\Delta}.$$

Its gradient, $\mathbf{J}^{\mathsf{T}}r_0 + \mathbf{J}^{\mathsf{T}}\mathbf{J}\,\boldsymbol{\Delta}$, vanishes at the step $\hat{\boldsymbol{\Delta}}$ that solves the *normal equations*,

$$(\mathbf{J}^{\mathsf{T}}\mathbf{J})\,\hat{\boldsymbol{\Delta}} = -\mathbf{J}^{\mathsf{T}}r(\mathbf{z}_0), \qquad \hat{\mathbf{z}} = \mathbf{z}_0 + \hat{\boldsymbol{\Delta}}. \tag{3.25}$$

Writing $\mathbf{J}^{\mathsf{T}}r_0 = -\mathbf{J}^{\mathsf{T}}\mathbf{J}\,\hat{\boldsymbol{\Delta}}$ in the expansion completes the square,

$$\mathcal{L}(\mathbf{z}) = \tfrac{1}{2}(\boldsymbol{\Delta} - \hat{\boldsymbol{\Delta}})^{\mathsf{T}}\mathbf{J}^{\mathsf{T}}\mathbf{J}\,(\boldsymbol{\Delta} - \hat{\boldsymbol{\Delta}}) + \text{constant} = \tfrac{1}{2}(\mathbf{z} - \hat{\mathbf{z}})^{\mathsf{T}}\mathbf{J}^{\mathsf{T}}\mathbf{J}\,(\mathbf{z} - \hat{\mathbf{z}}) + \text{constant},$$

where the constant collects the terms free of $\mathbf{z}$. By equation (3.24) the posterior is the exponential of minus this quadratic form, the density of a multivariate Gaussian with mean $\hat{\mathbf{z}}$ and covariance

$$\boldsymbol{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1}. \tag{3.26}$$

The matrix $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ is the information matrix $\boldsymbol{\Lambda}$ of the normalized rows. When the rows are written unnormalized, each in its own unit with the weight $w_k = 1/\sigma_k^2$ and the weights collected in $\mathbf{W} = \operatorname{diag}(w_k)$, the same matrix is $\boldsymbol{\Lambda} = \mathbf{J}^{\mathsf{T}}\mathbf{W}\mathbf{J}$, with $\mathbf{J}$ there the Jacobian of the unnormalized rows. For rows linear in the state the normal equations reach $\hat{\mathbf{z}}$ in one step from any $\mathbf{z}_0$, and equations (3.25) and (3.26) are the exact posterior mean and covariance. For nonlinear rows the same step is the Gauss–Newton step, repeated from the new point until it vanishes, and the result is a local Gaussian approximation, the Laplace approximation, whose derivation and curvature term Chapter 4 gives, with the three-entry hall as its example; Section 3.5.6 reads the two-entry hall of Chapter 1 in the terms of this section.

Precisions add. Split the rows into prior rows $\mathbf{L}_0(\mathbf{z} - \boldsymbol{\mu}_0)$, whose Jacobian $\mathbf{L}_0$ satisfies $\mathbf{L}_0^{\mathsf{T}}\mathbf{L}_0 = \boldsymbol{\Sigma}_0^{-1}$, and reading rows $\mathbf{R}^{-1/2}(\mathbf{A}\mathbf{z} - \mathbf{y})$, where the matrix $\mathbf{A}$ maps the state to the quantities read and the diagonal matrix $\mathbf{R}$ holds the variances of the readings. The Jacobian of the reading rows is $\mathbf{J}_y = \mathbf{R}^{-1/2}\mathbf{A}$; stacking the two blocks of rows stacks their Jacobians, and $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ is the sum of the two blocks' products,

$$\boldsymbol{\Sigma}^{-1} = \boldsymbol{\Sigma}_0^{-1} + \mathbf{J}_y^{\mathsf{T}}\mathbf{J}_y = \boldsymbol{\Sigma}_0^{-1} + \mathbf{A}^{\mathsf{T}}\mathbf{R}^{-1}\mathbf{A}. \tag{3.27}$$

Starting the normal equations at $\mathbf{z}_0 = \boldsymbol{\mu}_0$, where the prior rows vanish, gives $\mathbf{J}^{\mathsf{T}}r(\boldsymbol{\mu}_0) = \mathbf{A}^{\mathsf{T}}\mathbf{R}^{-1}(\mathbf{A}\boldsymbol{\mu}_0 - \mathbf{y})$, and equation (3.25) gives

$$\hat{\mathbf{z}} = \boldsymbol{\mu}_0 + \boldsymbol{\Sigma}\,\mathbf{A}^{\mathsf{T}}\mathbf{R}^{-1}(\mathbf{y} - \mathbf{A}\boldsymbol{\mu}_0). \tag{3.28}$$

The posterior information is the prior information plus the information of the readings, and the posterior mean is moved from the prior mean toward the readings in proportion to their information.

The conditional of a joint Gaussian. The same posterior follows from the joint distribution of the state and the readings. Before the readings, $\mathbf{z} \sim \mathcal{N}(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$ and $\mathbf{y} = \mathbf{A}\mathbf{z} + \boldsymbol{\epsilon}$ with $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{R})$ independent of the state, so the stacked vector of $\mathbf{z}$ and $\mathbf{y}$ is Gaussian with means $\boldsymbol{\mu}_0$ and $\mathbf{A}\boldsymbol{\mu}_0$ and covariance blocks

$$\boldsymbol{\Sigma}_{zz} = \boldsymbol{\Sigma}_0, \qquad \boldsymbol{\Sigma}_{zy} = \boldsymbol{\Sigma}_0\mathbf{A}^{\mathsf{T}} = \boldsymbol{\Sigma}_{yz}^{\mathsf{T}}, \qquad \boldsymbol{\Sigma}_{yy} = \mathbf{A}\boldsymbol{\Sigma}_0\mathbf{A}^{\mathsf{T}} + \mathbf{R}.$$

Let $\mathbf{P}$ be the inverse of the joint covariance and $\mathbf{S} = \boldsymbol{\Sigma}_{zz} - \boldsymbol{\Sigma}_{zy}\boldsymbol{\Sigma}_{yy}^{-1}\boldsymbol{\Sigma}_{yz}$ the *Schur complement* of $\boldsymbol{\Sigma}_{yy}$. The block-inversion formula, checked by multiplying back, gives $\mathbf{P}_{zz} = \mathbf{S}^{-1}$ and $\mathbf{P}_{zy} = -\mathbf{S}^{-1}\boldsymbol{\Sigma}_{zy}\boldsymbol{\Sigma}_{yy}^{-1}$. With the readings fixed, the exponent of the joint density, times -1, is a quadratic in the state,

$$\tfrac{1}{2}(\mathbf{z}-\boldsymbol{\mu}_0)^{\mathsf{T}}\mathbf{P}_{zz}(\mathbf{z}-\boldsymbol{\mu}_0) + (\mathbf{z}-\boldsymbol{\mu}_0)^{\mathsf{T}}\mathbf{P}_{zy}(\mathbf{y}-\mathbf{A}\boldsymbol{\mu}_0) + \text{constant} = \tfrac{1}{2}(\mathbf{z}-\mathbf{m})^{\mathsf{T}}\mathbf{S}^{-1}(\mathbf{z}-\mathbf{m}) + \text{constant},$$

with $\mathbf{m} = \boldsymbol{\mu}_0 + \boldsymbol{\Sigma}_{zy}\boldsymbol{\Sigma}_{yy}^{-1}(\mathbf{y} - \mathbf{A}\boldsymbol{\mu}_0)$. The conditional density is the joint density divided by $p(\mathbf{y})$, which is free of the state, so

$$\mathbf{z} \mid \mathbf{y} \sim \mathcal{N}(\boldsymbol{\mu}_0 + \boldsymbol{\Upsilon}(\mathbf{y} - \mathbf{A}\boldsymbol{\mu}_0),\ \boldsymbol{\Sigma}_0 - \boldsymbol{\Upsilon}\mathbf{A}\boldsymbol{\Sigma}_0), \qquad \boldsymbol{\Upsilon} = \boldsymbol{\Sigma}_0\mathbf{A}^{\mathsf{T}}(\mathbf{A}\boldsymbol{\Sigma}_0\mathbf{A}^{\mathsf{T}} + \mathbf{R})^{-1}, \tag{3.29}$$

since $\boldsymbol{\Sigma}_{zy}\boldsymbol{\Sigma}_{yy}^{-1}\boldsymbol{\Sigma}_{yz} = \boldsymbol{\Upsilon}\mathbf{A}\boldsymbol{\Sigma}_0$. This is the correction form: the prior mean is moved toward the readings by the gain $\boldsymbol{\Upsilon}$ times the innovation $\mathbf{y} - \mathbf{A}\boldsymbol{\mu}_0$, and the covariance is the Schur complement. The Kalman literature writes the gain $\mathbf{K}$; here it is $\boldsymbol{\Upsilon}$, clear of the hat matrix

$\mathbf{H}$ and its entries h_{kk} of Chapters 7 and 9. The correction form gives the same posterior as equations (3.27) and (3.28), by the matrix inversion identity $(\boldsymbol{\Sigma}_0^{-1}+\mathbf{A}^\mathsf{T}\mathbf{R}^{-1}\mathbf{A})^{-1} = \boldsymbol{\Sigma}_0 - \boldsymbol{\Upsilon}\mathbf{A}\boldsymbol{\Sigma}_0$; the information form inverts a matrix of the size of the state, the correction form one of the size of the readings.

Remark 3.1. For one entry and one reading, $\mathbf{A} = 1$, $\boldsymbol{\Sigma}_0 = \sigma_0^2$ and $\mathbf{R} = \sigma_T^2$, and the three equations reduce to the scalar rules $1/\sigma_{\text{post}}^2 = 1/\sigma_0^2 + 1/\sigma_T^2$, $\mu_{\text{post}} = \sigma_{\text{post}}^2(\mu_0/\sigma_0^2 + y_T/\sigma_T^2)$ and $\mu_{\text{post}} = \mu_0 + \upsilon(y_T - \mu_0)$ with $\upsilon = \sigma_0^2/(\sigma_0^2 + \sigma_T^2)$: precisions add, and the posterior mean is the precision-weighted average of the prior mean and the reading.

3.5.4 The numbers and what changed

For the hall the state is the one number T, and the two rows of equation (3.23) have the slopes

$$j_1 = \frac{1}{\sigma_0} = 0.5\,\text{K}^{-1}, \qquad j_2 = \frac{1}{\sigma_T} = 2\,\text{K}^{-1}.$$

The row terms are $\mathcal{L}_k = \frac{1}{2}r_k^2$, with the slopes $\mathcal{L}_k' = j_k r_k$ and the curvatures $\mathcal{L}_k'' = j_k^2$ in T, the primes denoting $\text{d}/\text{d}T$; slopes and curvatures add over the rows,

$$\mathcal{L}' = j_1 r_1 + j_2 r_2, \qquad \mathcal{L}'' = j_1^2 + j_2^2.$$

Because the rows are linear, the curvature has no further term and is the same at every T; it is the information

$$\Lambda = j_1^2 + j_2^2 = 0.5^2 + 2^2 = 0.25 + 4 = 4.25\,\text{K}^{-2},$$

the prior precision $1/\sigma_0^2$ plus the reading precision $1/\sigma_T^2$, as equation (3.27) states in the scalar form of the remark above, with the prior variance $\sigma_0^2 = 4\,\text{K}^2$ and the reading variance $\sigma_T^2 = 0.25\,\text{K}^2$. It is the one-entry case of $\mathbf{J}^\mathsf{T}\mathbf{J}$, and equation (3.26) gives the posterior variance as its reciprocal,

$$\sigma_{\text{post}}^2 = \frac{1}{\Lambda} = \frac{1}{4.25} = 0.2353\,\text{K}^2, \qquad \sigma_{\text{post}} = 0.4851\,\text{K}.$$

With one entry the normal equations (3.25) are the one equation $\Lambda\,\Delta T = -\mathcal{L}'$, the Newton step $\Delta T = -\mathcal{L}'/\mathcal{L}''$. They start at the prior mean, $T_0 = 26°\text{C}$, where the prior row is 0 and the reading row $(26 - 27)/0.5 = -2$, so that $\mathcal{L}' = 0.5 \times 0 + 2 \times (-2) = -4\,\text{K}^{-1}$, and the step is $\Delta T = 4/4.25 = 0.9412\,\text{K}$. The iterations are

k [–]	T [°C]	$\mathcal{L}(T)$ [–]	$\mathcal{L}_1'$ [K^{-1}]	$\mathcal{L}_2'$ [K^{-1}]	$\mathcal{L}'$ [K^{-1}]	$\mathcal{L}''$ [K^{-2}]	ΔT [K]
0	26	2.0000	0	−4	−4	4.25	0.9412
1	26.9412	0.1176	0.2353	−0.2353	0	4.25	0

At the second point the residuals are $r_1 = 0.4706$ and $r_2 = -0.1176$, the slope is $\mathcal{L}' = 0.5 \times 0.4706 + 2 \times (-0.1176) = 0$, and the step vanishes: the rows are linear, and one step reaches the minimum. The correction form (3.29) confirms both moments: the variance of the reading before it arrives is $\sigma_0^2 + \sigma_T^2 = 4 + 0.25 = 4.25\,\text{K}^2$, $\upsilon = 4/4.25 = 0.9412$, the mean $26 + 0.9412 \times (27 - 26) = 26.9412°\text{C}$ and the Schur complement $4 - 0.9412 \times 4 = 0.2353\,\text{K}^2$. The answer is the minimum $\hat{T} = 26.9412°\text{C}$, written μ_{post} below, with the standard deviation σ_{post}:

$$T \mid y_T \sim \mathcal{N}(26.9412, 0.4851^2). \tag{3.30}$$

Figure 3.6(a) draws the update. The horizontal axis is the temperature and the vertical axis the probability density. The gray curve is the prior, the dashed blue curve the likelihood of the reading as a function of T (normalized to integrate to one), and the green curve the posterior. The posterior mean moved from 26 to 26.94°C, and the standard deviation fell from 2 to 0.49 K.

The reading was more precise than the prior, so it received the greater weight, $\upsilon = 0.94$; but the posterior mean is not exactly 27°C, because the prior also carries information.

Table 3.7 shows the trend for sensors of four precisions, all reading 27°C against the same prior. The first column is the sensor's standard deviation, the second the weight υ of the reading, and the remaining columns the posterior mean, the posterior standard deviation and the exceedance probability derived in Section 3.5.5. Each row is the same computation with the slopes $j_1 = 1/\sigma_0$ and $j_2 = 1/\sigma_T$, so that $\sigma^2_{\text{post}} = 1/\Lambda = 1/(0.25 + 1/\sigma_T^2)$. The second row is the case worked above. A sensor as uncertain as the prior, $\sigma_T = 2\,\text{K}$, receives half the weight and leaves the posterior standard deviation at $2/\sqrt{2} = 1.41\,\text{K}$.

Table 3.7: The hall-air update for four sensor precisions: prior $\mathcal{N}(26, 2^2)$, reading $y_T = 27$°C, equations (3.27) and (3.29), and the exceedance probability of $T_{\max} = 28$°C derived in Section 3.5.5.

σ_T [K]	weight υ [–]	μ_{post} [°C]	σ_{post} [K]	$\Pr(T > T_{\max} \mid y_T)$ [–]
0.25	0.9846	26.9846	0.2481	2×10^{-5}
0.5	0.9412	26.9412	0.4851	0.0145
1.0	0.8000	26.8000	0.8944	0.0899
2.0	0.5000	26.5000	1.4142	0.1444

3.5.5 The threshold question

Suppose the hall-air limit of this illustration is $T_{\max} = 28$°C, and the question is the probability that the hall exceeds it, $\Pr(T > T_{\max} \mid y_T)$. Since the posterior is Gaussian, standardizing with the mean $\mu_{\text{post}} = \hat{T}$ and the standard deviation $\sigma_{\text{post}} = 1/\sqrt{\Lambda}$ and using the standard normal cumulative distribution function Φ,

$$\begin{aligned} \Pr(T > T_{\max} \mid y_T) &= 1 - \Phi\left(\frac{T_{\max} - \mu_{\text{post}}}{\sigma_{\text{post}}}\right) \\ &= 1 - \Phi\left(\frac{28 - 26.9412}{0.4851}\right) = 1 - \Phi(2.183) = 0.0145, \end{aligned} \tag{3.31}$$

or about 1.45 %. Before the reading the same calculation with the prior gives $\Pr(T > 28\text{°C}) = 1 - \Phi\big((28 - 26)/2\big) = 1 - \Phi(1) = 0.1587$. The estimated exceedance probability fell from 15.9 % to 1.45 %, for the same fixed operating condition and the assumed Gaussian measurement model. The rows are linear, so this normal tail is the exact posterior's: a direct quadrature of the product (3.22) also gives 0.0145, and there is no gap to set beside it, as there is on the curved physics row of the next subsection. Figure 3.6(b) shades the two tails: the axes are those of panel (a), the dashed line is the limit, the gray area to its right under the prior is 0.159, and the green area under the posterior is 0.0145.

The physical hall did not necessarily change when the reading arrived; the information about it changed, and that is the mechanism of Bayesian updating. The calculation also walks the method of *How to read this book* on its smallest example. The plant graph of Figure 3.5(a) is step 0. The question $\Pr(T > T_{\max} \mid y_T)$, unknown on the left and reading on the right, is step 1. Attaching the prior and the reading as the two rows of equation (3.23) is step 2. The posterior, derived through the normal equations and the variance as the reciprocal of the information, equations (3.25) and (3.26) with one entry, and computed by the one Newton step $-\mathcal{L}'/\mathcal{L}''$ with $\Lambda = 4.25\,\text{K}^{-2}$, is step 3. Reading the normal tail from it is step 4, and checking that tail by the correction form and its Schur complement is step 5, and the sentence "the hall exceeds 28°C with probability 1.45 %, posterior $26.94 \pm 0.49\,\text{K}$" is step 6. Step 7, the next reading, reads

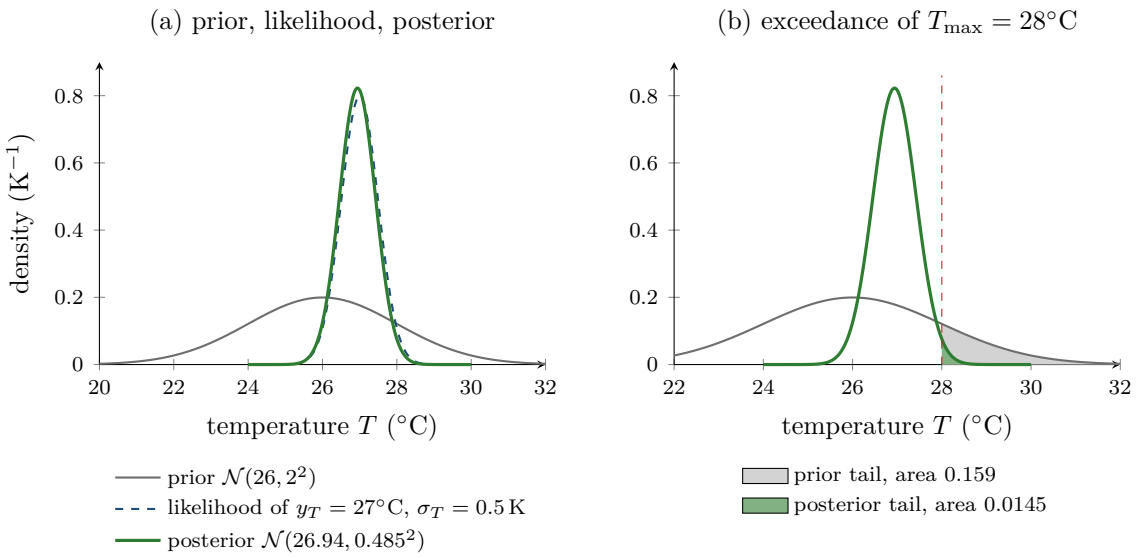


Figure 3.6: The hall-air update of equations (3.27)–(3.30). (a) Prior $\mathcal{N}(26, 2^2)$, likelihood of the reading $y_T = 27°\mathrm{C}$ with $\sigma_T = 0.5\,\mathrm{K}$ drawn as the density $\mathcal{N}(27, 0.5^2)$ in T, and posterior $\mathcal{N}(26.9412, 0.4851^2)$. (b) The probability that the hall exceeds $T_{\max} = 28°\mathrm{C}$ (dashed): $1-\Phi(1) = 0.159$ under the prior and $1-\Phi(2.183) = 0.0145$ under the posterior (equation (3.31)).

Table 3.7: if 1.45 % were too high a residual doubt, a sensor with $\sigma_T = 0.25\,\mathrm{K}$ would, for a reading near the same value, reduce it to the order of 10^{-5}; Chapter 18 ranks such candidate readings by the width each would remove. The posterior is inferred first, and the operator's question is then answered from it:

$$\underbrace{p(T \mid y_T)}_{\text{inference}} \quad \longrightarrow \quad \underbrace{\Pr(T > T_{\max} \mid y_T)}_{\text{threshold question}}.$$

3.5.6 The two-entry hall and its physics row

The reading of the one-entry hall informed only the entry it measured; Section 3.5.1 claimed that a reading also reaches the conductance through the physics row. The hall of Chapter 1 shows it. Its state is $\mathbf{z} = (T,\ g)^{\mathsf{T}}$, and its three rows, equation (1.10), are the physics row with tolerance $\sigma_p = 0.01\,\mathrm{K}$, the reading row with $\sigma_T = 0.5\,\mathrm{K}$ and the handbook prior row on g with standard deviation 0.3; Chapter 1 solves their normal equations by Gauss–Newton from (23°C, 0) (Table 1.4) and reaches $\hat{\mathbf{z}} = (23.070°\mathrm{C},\ -0.0139)^{\mathsf{T}}$. Equation (3.27) holds row by row: with $\mathbf{J}_k$ the k-th row of $\mathbf{J}$, $\mathbf{J}^{\mathsf{T}}\mathbf{J} = \sum_k \mathbf{J}_k^{\mathsf{T}}\mathbf{J}_k$, each row contributing the outer product of its own sensitivities. At $\hat{\mathbf{z}}$ the rows of $\mathbf{J}$ are (100, 507.0), (2, 0) and (0, 3.333), and

$$\mathbf{J}^{\mathsf{T}}\mathbf{J} = \underbrace{\begin{bmatrix} 10\,000 & 50\,701 \\ 50\,701 & 257\,056 \end{bmatrix}}_{\text{physics}} + \underbrace{\begin{bmatrix} 4 & 0 \\ 0 & 0 \end{bmatrix}}_{\text{reading}} + \underbrace{\begin{bmatrix} 0 & 0 \\ 0 & 11.11 \end{bmatrix}}_{\text{prior}} = \begin{bmatrix} 10\,004 & 50\,701 \\ 50\,701 & 257\,067 \end{bmatrix}.$$

The reading's and the prior's terms are diagonal, because each of those rows carries one entry. The only off-diagonal entry, 50 701, comes from the physics row, the one row with sensitivities in both columns: this is equation (3.20) in vector form, the reading informing T and the physics row carrying that information to g. Without the reading row, the physics and prior terms alone invert to $\sqrt{\Sigma_{22}} = 0.3$, the prior spread; with it, the covariance of Chapter 1, equation (1.21), has $\sqrt{\Sigma_{11}} = 0.481\,\mathrm{K}$, $\sqrt{\Sigma_{22}} = 0.0948$ and correlation -0.9998. The reading has cut the conductance spread from 0.3 to 0.0948 although its own row has a zero in the column of g, and the posterior

lies along the physics curve, on the short arc of the state plane of Figure 1.5 inside the reading band. Because the physics row is nonlinear, these spreads include the curvature term that Chapter 4 derives; $(\mathbf{J}^\mathsf{T}\mathbf{J})^{-1}$ alone gives 0.475 K and 0.094. The exact posterior of these rows has the standard deviation 0.4813 K against the Laplace 0.4806 K, a gap of 0.0007 K, and the mean 23.0300°C against the mode 23.0701°C.

3.5.7 Dependence inside and outside the physics graph

The common cause of Section 3.4.4 cost a factor of nine because it was left out of the product. A plant graph built on conservation relations does not leave out a shared cause that has a physical path: the cause is an entry of the state, and every row that depends on it carries it. The two-entry hall above shows this. The reading and the conductance share no row, yet the physics row couples them, and the posterior correlation of T and g, -0.9998, comes out of $\mathbf{J}^\mathsf{T}\mathbf{J}$ without being assumed. Had the chilled water been a node of the plant graph, with a row for each hall's excursion given its state, the joint excursion probability (3.17) would have followed from the rows, and the independent product of Section 3.4.5 could not have been formed. Conditional independence given the state is then the statement of Section 3.4.6 with the common cause known: it holds because the cause is in the state.

What the graph does not supply is the independence of the errors. Equation (3.20) assumes that a reading's error is independent of the inputs and, when there are several readings, of the other readings' errors, so that the posterior (3.24) is a product over rows. No conservation relation states that, and none checks it directly. Suppose the hall of Section 3.5.4 carries two sensors from one calibration batch, each with $\sigma_T = 0.5\,\mathrm{K}$, both reading 27°C, with errors correlated through the offset they share, $\mathrm{corr}(\epsilon_1, \epsilon_2) = 0.5$. Treated as independent rows, the pair adds the precision $2/\sigma_T^2 = 8\,\mathrm{K}^{-2}$ to the prior's 0.25. With the correlation, equation (3.27) holds with $\mathbf{R}$ a full matrix, its off-diagonal entry $0.5\,\sigma_T^2$ and its rows whitened by $\mathbf{R}^{-1/2}$, and the pair adds only

$$\mathbf{A}^\mathsf{T}\mathbf{R}^{-1}\mathbf{A} = \frac{2}{\sigma_T^2\,(1 + 0.5)} = 5.33\,\mathrm{K}^{-2}, \qquad \mathbf{A} = \begin{bmatrix} 1 \\ 1 \end{bmatrix},$$

the precision of the pair's average, whose error has the variance $\sigma_T^2\,(1 + 0.5)/2$. Table 3.8 compares the two treatments with the single sensor of Table 3.7. Its columns are the precision the readings add, the posterior mean, the posterior standard deviation and the probability of exceeding $T_{\max} = 28$°C.

Table 3.8: Two hall sensors of $\sigma_T = 0.5\,\mathrm{K}$, both reading 27°C, against the prior $\mathcal{N}(26, 2^2)$ and the limit $T_{\max} = 28$°C: the pair treated as independent, the pair with $\mathrm{corr}(\epsilon_1, \epsilon_2) = 0.5$, and one sensor, which is also the limit of a correlation of one.

Readings	precision added [K^{-2}]	mean [°C]	σ_{post} [K]	$\Pr(T > T_{\max} \mid \mathbf{y})$ [–]
two, treated as independent	8.00	26.9697	0.3482	0.00154
two, correlation 0.5	5.33	26.9552	0.4232	0.00678
one sensor	4.00	26.9412	0.4851	0.0145

The independent treatment reports an exceedance probability of 0.00154 against the 0.00678 the correlated pair supports, 4.4 times too small, and a posterior standard deviation of 0.348 K against 0.423 K. The error has the direction of Section 3.4.5's: a posterior narrower than the evidence supports, and a limit that looks safer than it is. At a correlation of one the second sensor adds nothing, and the posterior is that of one sensor.

The shared offset is a common cause, and the remedy of Section 3.4.6 applies: put it in the state. A bias entry $b \sim \mathcal{N}(0,\ 0.5\,\sigma_T^2)$ carried by both reading rows, with independent remainders of variance $0.5\,\sigma_T^2$, makes the two rows conditionally independent given (T, b), and the state (T, b) gives the posterior of T of the correlated row of Table 3.8 to every printed digit. The posterior of b, however, has the standard deviation 0.348 K against its prior's 0.354 K, with $\mathrm{corr}(T, b \mid \mathbf{y}) = -0.81$. The readings barely inform it: a shared offset moves both readings together, as a change of T does, and only the prior on b bounds it.

The same assumption hides in four other places a physics graph does not reach: a prior shared by several entries, such as two conductances taken from one handbook; a physics tolerance σ_p treated as independent from row to row where the model's error is systematic; repeated readings averaged as if their errors were independent when they are correlated in time; and a path of the conserved quantity that the graph omits. Each makes the posterior narrower than the evidence supports. What the graph adds is redundancy. When more rows bear on the state than it has entries, a misfit of the rows taken together, a reading far from what the other rows predict, or a comparison of drawings with and without the suspected cause can reveal the dependence: Chapter 7 derives the misfit test, the bias entry and the comparison of drawings on inconsistent meters, Chapter 9 the reading far from its prediction, and Chapter 12 carries the shared chilled water as an entry of the state. No test reveals a dependence that moves the readings along a direction the physics also allows, as the shared offset moves them along T.

Principle 3.3. *A physics graph carries dependence through shared physical quantities as entries of the state, so the independent product cannot be formed for them. The independence it still assumes is that of the errors: it must be stated, tested where the rows are redundant, and, where it fails, removed by putting the shared cause in the state.*

3.6 Predictive variance and the covariance of a joint posterior

Two variances complete the chapter: the law of total variance, which separates the randomness of future alarms from the uncertainty of their rate, and the covariance of a posterior over a joint state, which decides how uncertain a derived quantity is and why the shape of the posterior must be read rather than its marginals.

3.6.1 The law of total variance

Suppose the number K_τ of the hall's high-temperature alarms during a future period of length τ is Poisson given the rate, $K_\tau \mid \lambda \sim \mathrm{Poisson}(\lambda\tau)$, and the rate itself is uncertain, with posterior $p(\lambda \mid \mathbf{y})$. The predictive variance of K_τ is wanted. Two facts give it. The law of total expectation, which is the law of total probability applied inside an expectation, states that $\mathbb{E}[K_\tau \mid \mathbf{y}] = \mathbb{E}\big[\mathbb{E}[K_\tau \mid \lambda] \mid \mathbf{y}\big]$, the conditional mean averaged over the rate; and the variance of any quantity is the mean of its square minus the square of its mean. Then

$$\mathrm{Var}(K_\tau \mid \mathbf{y}) = \mathbb{E}\big[\mathbb{E}[K_\tau^2 \mid \lambda] \mid \mathbf{y}\big] - \big(\mathbb{E}\big[\mathbb{E}[K_\tau \mid \lambda] \mid \mathbf{y}\big]\big)^2,$$

and writing the conditional second moment as $\mathbb{E}[K_\tau^2 \mid \lambda] = \mathrm{Var}(K_\tau \mid \lambda) + \mathbb{E}[K_\tau \mid \lambda]^2$,

$$\mathrm{Var}(K_\tau \mid \mathbf{y}) = \mathbb{E}\big[\mathrm{Var}(K_\tau \mid \lambda) \mid \mathbf{y}\big] + \Big\{\mathbb{E}\big[\mathbb{E}[K_\tau \mid \lambda]^2 \mid \mathbf{y}\big] - \big(\mathbb{E}\big[\mathbb{E}[K_\tau \mid \lambda] \mid \mathbf{y}\big]\big)^2\Big\}.$$

The brace is the variance of the conditional mean, which gives the *law of total variance*,

$$\mathrm{Var}(K_\tau \mid \mathbf{y}) = \mathbb{E}\big[\mathrm{Var}(K_\tau \mid \lambda) \mid \mathbf{y}\big] + \mathrm{Var}\big(\mathbb{E}[K_\tau \mid \lambda] \mid \mathbf{y}\big). \tag{3.32}$$

The first term is the expected variability of the alarm count when the rate is specified, the aleatory part. The second is the additional predictive variability because the rate is uncertain,

the epistemic part. For a Poisson count the mean and the variance are equal, $\mathbb{E}[K_\tau \mid \lambda] = \mathrm{Var}(K_\tau \mid \lambda) = \lambda\tau$, so that $\mathrm{Var}(K_\tau \mid \mathbf{y}) = \mathbb{E}[\lambda\tau \mid \mathbf{y}] + \mathrm{Var}(\lambda\tau \mid \mathbf{y})$, and since τ is fixed,

$$\mathrm{Var}(K_\tau \mid \mathbf{y}) = \tau\,\mathbb{E}[\lambda \mid \mathbf{y}] + \tau^2\,\mathrm{Var}(\lambda \mid \mathbf{y}). \tag{3.33}$$

The predictive variance has been separated into a process term and a rate term.

3.6.2 Alarm counts over the next shift

Take the posterior of the hall's alarm rate from the state, the normal of the log-rate of equation (2.41), with $\hat{\omega} = \ln(8/3)$ and $\sigma_\omega = 0.5$, so that λ is lognormal about $\hat{\lambda} = 4/1500\,\mathrm{h}$. The standard lognormal moments give

$$\mathbb{E}[\lambda \mid \mathbf{y}] = \hat{\lambda}\,\mathrm{e}^{\sigma_\omega^2/2} = 3.022 \times 10^{-3}\,\mathrm{h}^{-1}, \qquad \mathrm{Var}(\lambda \mid \mathbf{y}) = \hat{\lambda}^2\,\mathrm{e}^{\sigma_\omega^2}(\mathrm{e}^{\sigma_\omega^2} - 1) = 2.593 \times 10^{-6}\,\mathrm{h}^{-2}.$$

Over the next night shift, $\tau = 12\,\mathrm{h}$, the expected alarm count is $\mathbb{E}[K_\tau \mid \mathbf{y}] = 12 \times 3.022 \times 10^{-3} = 0.0363$, and by equation (3.33) the predictive variance is

$$\mathrm{Var}(K_\tau \mid \mathbf{y}) = 0.0363 + 12^2 \times 2.593 \times 10^{-6} = 0.0363 + 0.0004 = 0.0366.$$

The process term is 99.0 % of the predictive variance and the rate term 1.02 %. The count 0.0363 is a mean, not the probability of an alarm: a night with two alarms counts twice in the mean, so the probability of at least one, 0.0354 from the same normal, lies below it.

The exact posterior $\mathrm{Gamma}(a, b) = \mathrm{Gamma}(4, 1500\,\mathrm{h})$ stands beside these, with the standard gamma moments $\mathbb{E}[\lambda \mid \mathbf{y}] = a/b = 2.667 \times 10^{-3}\,\mathrm{h}^{-1}$, one alarm in about 375 h, and $\mathrm{Var}(\lambda \mid \mathbf{y}) = a/b^2 = 1.778 \times 10^{-6}\,\mathrm{h}^{-2}$. The same two terms are

$$\mathrm{Var}(K_\tau \mid \mathbf{y}) = 12\left(\frac{4}{1500}\right) + 12^2\left(\frac{4}{1500^2}\right) = 0.032000 + 0.000256 = 0.032256,$$

the mean 0.0320 against the reported 0.0363, a gap of 0.0043, and the variance 0.0323 against 0.0366; the probability of an alarm, 0.0314, again lies below its mean. The share of the rate term under the gamma is

$$\frac{\tau^2 a/b^2}{\tau a/b + \tau^2 a/b^2} = \frac{\tau}{b + \tau},$$

which for $\tau = 12\,\mathrm{h}$ and $b = 1500\,\mathrm{h}$ is $12/1512 = 0.0079$, against the reported 0.0102. The gamma's result can be checked against its full predictive distribution. A Poisson count whose mean is averaged over a gamma rate has, by a standard result, the negative binomial distribution with shape a and success probability $\pi = b/(b + \tau)$, whose mean is $a(1 - \pi)/\pi = a\tau/b$ and whose variance is $a(1 - \pi)/\pi^2 = a\tau/b + a\tau^2/b^2$: exactly the two terms of equation (3.33). For $\pi = 1500/1512 = 0.99206$ they are 0.032000 and 0.032256.

Table 3.9 extends the calculation to longer periods. The first column is the period; the next four are the reported values from the state, the expected count, which is also the process term, the rate term, the total and the share of the rate term; and the last three the gamma's mean, variance and share. Over a shift almost all of the predictive variance comes from the random alarm process, and a longer alarm log, which shrinks only the rate term, would barely change the forecast of the next shift's alarms. The two terms are equal at about 1165 h (at 1500 h under the gamma, as long as the logged hours behind the posterior), and over 3000 h the rate term is 0.720 of the total (0.667 under the gamma). The gaps grow with the period, because the lognormal's mean rate lies above the gamma's and its upper tail is heavier. Whether a longer log will meaningfully improve an alarm forecast depends on this split: more evidence shrinks the second term, and the first remains as long as the future count is modeled as a Poisson process.

Table 3.9: Predictive mean and variance of the hall's alarm count over a future period τ, equation (3.33): reported from the normal of the log-rate, with the posterior Gamma$(4, 1500\,\mathrm{h})$ beside it; the first row is one night shift.

τ [h]	mean = process term [–]	rate term [–]	variance [–]	rate share [–]	gamma mean [–]	gamma variance [–]	gamma share [–]
12	0.0363	0.0004	0.0366	0.0102	0.0320	0.0323	0.0079
100	0.3022	0.0259	0.3281	0.0790	0.2667	0.2844	0.0625
500	1.5109	0.6483	2.1592	0.3003	1.3333	1.7778	0.2500
1500	4.5326	5.8351	10.3677	0.5628	4.0000	8.0000	0.5000
3000	9.0652	23.3405	32.4057	0.7203	8.0000	24.0000	0.6667

3.6.3 One posterior over a joint state

The examples so far had one unknown at a time. A plant graph carries many at once. Suppose the state vector holds the air temperatures of two halls, the log-conductance of an air path and the alarm rates of the two halls' high-temperature alarms,

$$\mathbf{z} = \begin{bmatrix} T_A & T_B & g & \lambda_A & \lambda_B \end{bmatrix}^{\mathsf{T}},$$

and the evidence $\mathbf{y}$ gathers temperature readings, the alarm logs and inspection results. Everything known after the evidence is the joint posterior $p(\mathbf{z} \mid \mathbf{y}, M)$, where M names the model: the graph, its rows and its priors. Different questions are asked of this same distribution:

State estimation, $p(T_A \mid \mathbf{y}, M)$: the marginal of one entry.

Threshold probability, $\Pr(T_B > T_{\max} \mid \mathbf{y}, M)$: a tail of another marginal.

Parameter comparison, $\Pr(\lambda_A > \lambda_B \mid \mathbf{y}, M)$: a tail of the derived quantity $\lambda_A - \lambda_B$, which depends on the joint distribution of the two rates.

Expected derived quantity, $\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}, M]$, such as the expected number of alarms in the two halls over the next shift: the average of a function of the state.

Prediction, $\Pr(\text{no alarm in } \tau \mid \mathbf{y}, M)$: a posterior predictive probability, equation (2.41) with the rate taken from the state.

The methods of computation differ, but every one of them depends on a consistent probability model of the physical plant and the evidence, and none needs a separate inference.

3.6.4 The posterior covariance and the variance budget

When the posterior can be approximated locally by a multivariate Gaussian, $\mathbf{z} \mid \mathbf{y}, M \approx \mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$, the covariance matrix $\boldsymbol{\Sigma}$ describes the uncertainty of each entry and how the entries vary together. The approximation takes the most probable state as the mean and the inverse curvature of the objective there as the covariance. For rows linear in the state it is exact, $\boldsymbol{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1}$ by Section 3.5.3; for nonlinear rows it is the Laplace approximation, which Chapter 4 derives with its curvature term. For two entries,

$$\boldsymbol{\Sigma} = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{12} & \Sigma_{22} \end{bmatrix}, \qquad \operatorname{corr}(z_1, z_2 \mid \mathbf{y}) = \frac{\Sigma_{12}}{\sqrt{\Sigma_{11}\Sigma_{22}}}.$$

The diagonal entries are the marginal variances, $\Sigma_{11} = \operatorname{Var}(z_1 \mid \mathbf{y})$; the off-diagonal entry is the covariance, and $\operatorname{corr}(z_1, z_2 \mid \mathbf{y})$ the correlation. The off-diagonal entries matter because physical

relations and shared observations introduce posterior dependence between unknowns; they are not zero merely because the priors of the entries were independent. The hall of Section 3.5.6 is one case. Another carries the supply air as an entry: the state $\mathbf{z} = (T,\ T_c,\ g)^\mathsf{T}$ for the illustrative hall with a weak air path of Figure 3.2(c), with $Q = 30\,\mathrm{kW}$, $G_{\mathrm{ref}} = 3\,\mathrm{kW/K}$, independent priors $T_c \sim \mathcal{N}(20, 1^2)$ and $g \sim \mathcal{N}(0, 0.05^2)$ (the $0.15\,\mathrm{kW/K}$ of the figure), a reading with $\sigma_T = 0.5\,\mathrm{K}$ equal to the prior prediction 30°C, and the physics row $\big(T - T_c - 10\,\mathrm{K}\ \mathrm{e}^{-g}\big)/\sigma_p$ with $\sigma_p = 0.01\,\mathrm{K}$. At the prior means every residual is zero and the rows of $\mathbf{J}$ are $(1,\ -1,\ 10)/\sigma_p$ for the physics, $(1,\ 0,\ 0)/\sigma_T$ for the reading, and $(0,\ 1,\ 0)$ and $(0,\ 0,\ 1)/0.05$ for the priors; equation (3.26) gives

$$\boldsymbol{\Sigma} = (\mathbf{J}^\mathsf{T}\mathbf{J})^{-1} = \begin{bmatrix} 0.2083 & 0.1667 & -0.0042 \\ 0.1667 & 0.3334 & 0.0167 \\ -0.0042 & 0.0167 & 0.0021 \end{bmatrix},$$

and the correlation of T_c and g is $0.0167/\sqrt{0.3334 \times 0.0021} = 0.632$ (unrounded entries; 0.6325 at $\sigma_p = 0.001\,\mathrm{K}$). Without the reading row the same computation gives zero. A single reading leaves T_c and G positively correlated, since a warmer supply air and a larger conductance explain the same reading, and the correlation enters through the physics row, the only row with entries in the columns of both.

Suppose a derived quantity $f(\mathbf{z})$ has the gradient $\mathbf{g} = \partial f/\partial \mathbf{z}$ at $\hat{\mathbf{z}}$, a column with one entry per entry of the state. With $\boldsymbol{\delta} = \mathbf{z} - \hat{\mathbf{z}}$ the deviation of the state from its mean, the first-order change is $f - f(\hat{\mathbf{z}}) \approx \mathbf{g}^\mathsf{T}\boldsymbol{\delta}$, a scalar equal to its own transpose, and the covariance is $\boldsymbol{\Sigma} = \mathbb{E}[\boldsymbol{\delta}\boldsymbol{\delta}^\mathsf{T}]$, so that

$$\mathrm{Var}(f) \approx \mathbb{E}\big[\mathbf{g}^\mathsf{T}\boldsymbol{\delta}\,\boldsymbol{\delta}^\mathsf{T}\mathbf{g}\big] = \mathbf{g}^\mathsf{T}\,\mathbb{E}\big[\boldsymbol{\delta}\boldsymbol{\delta}^\mathsf{T}\big]\,\mathbf{g} = \mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}. \tag{3.34}$$

This is the *variance budget* of the output. Written out entry by entry, $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$ has one term per entry of the state, the squared sensitivity of the output to the entry times the variance of the entry, plus a cross term $2\,\nabla_k \nabla_l\,\Sigma_{kl}$, with $\nabla_k = \partial f/\partial z_k$, for every pair of entries that move together. For the sum $f = z_1 + z_2$ the gradient is $\mathbf{g} = (1,\ 1)^\mathsf{T}$, and the budget is exact because f is linear,

$$\mathrm{Var}(z_1 + z_2) = \begin{bmatrix} 1 & 1 \end{bmatrix} \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{12} & \Sigma_{22} \end{bmatrix} \begin{bmatrix} 1 \\ 1 \end{bmatrix} = \Sigma_{11} + \Sigma_{22} + 2\Sigma_{12}. \tag{3.35}$$

With $\Sigma_{11} = \Sigma_{22} = 1\,\mathrm{K}^2$ (illustrative), the variance of the sum is $0.4\,\mathrm{K}^2$ for $\Sigma_{12} = -0.8\,\mathrm{K}^2$, $2\,\mathrm{K}^2$ for independent entries and $3.6\,\mathrm{K}^2$ for $\Sigma_{12} = 0.8\,\mathrm{K}^2$. A negative covariance makes the output more precisely known than the individual entries suggest; a positive one makes it less.

Table 3.10 works the budget for the hall air of equation (3.19), with the load known, $Q = 30\,\mathrm{kW}$, and the two unknowns of Figure 3.2(c) as the state $\mathbf{z} = (T_c,\ G)^\mathsf{T}$: T_c with standard deviation 1 K and G with mean $3\,\mathrm{kW/K}$ and standard deviation $0.15\,\mathrm{kW/K}$. The gradient of T and the covariance are

$$\mathbf{g} = \begin{bmatrix} \partial T/\partial T_c \\ \partial T/\partial G \end{bmatrix} = \begin{bmatrix} 1 \\ -Q/G^2 \end{bmatrix} = \begin{bmatrix} 1 \\ -3.333 \end{bmatrix}, \qquad \boldsymbol{\Sigma} = \begin{bmatrix} 1 & 0.15\,c \\ 0.15\,c & 0.0225 \end{bmatrix}, \qquad c = \mathrm{corr}(T_c, G),$$

with $\mathbf{g}$ in K per unit of each entry and $\boldsymbol{\Sigma}$ in the products of the entries' units, so that $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g} = (1.00 + 0.25 - 1.00\,c)\,\mathrm{K}^2$: the supply-air term $1^2 \times 1$, the conductance term $3.333^2 \times 0.0225$ and the cross term $2 \times 1 \times (-3.333) \times 0.15\,c$. The columns of the table are the correlation, the three terms, the variance and the standard deviation of T. The two marginals are the same in every row, and the standard deviation of the hall air ranges from 0.67 to 1.43 K with the correlation alone. The budget is first order, since T is curved in G; a quadrature over the same joint normal gives 0.6701, 1.1203 and 1.4357 K against the budget's 0.6708, 1.1180 and 1.4318 K, gaps of at most 0.0039 K.

Two consequences follow from equation (3.34). An entry can be very uncertain and still not matter: if the output is insensitive to it, its whole term is near zero, and narrowing its variance

Table 3.10: Variance budget of the hall air $T = T_c + Q/G$, equation (3.34), with $Q = 30\,\mathrm{kW}$, G of mean $3\,\mathrm{kW/K}$ and standard deviation $0.15\,\mathrm{kW/K}$, T_c of standard deviation $1\,\mathrm{K}$, and correlation $c = \mathrm{corr}(T_c, G)$ (illustrative).

c [–]	supply-air term [K^2]	conductance term [K^2]	cross term [K^2]	Var(T) [K^2]	sd [K]
-0.8	1.00	0.25	$+0.80$	2.05	1.43
0	1.00	0.25	0	1.25	1.12
0.8	1.00	0.25	-0.80	0.45	0.67

buys nothing. And an entry can be poorly known while the output is well known: when the sensitivities and the covariance have opposite signs, the cross term subtracts from the diagonal terms. A reading is worth requesting when it shrinks a variance Σ_{kk} whose sensitivity is large, and worthless when the sensitivity is near zero.

3.6.5 Reading the shape of the posterior

The second consequence deserves a worked case, because it is where marginal spreads mislead. Carry the hall air, the load and the conductance as entries, the last two through the logarithms $\ell_Q = \ln(Q/Q_{\mathrm{ref}})$ and $g = \ln(G/G_{\mathrm{ref}})$, each with an independent prior of mean 0 and standard deviation 0.10, that is about 10 %, and with the supply air pinned. The references are chosen so that $Q_{\mathrm{ref}}/G_{\mathrm{ref}} = 10\,\mathrm{K}$, a temperature rise of 10 K at the prior means, and a sensor with $\sigma_T = 0.2\,\mathrm{K}$ reads that temperature, $y_T = T_c + 10\,\mathrm{K}$. The state, its four rows and their 4×3 Jacobian are

$$\mathbf{z} = \begin{bmatrix} T \\ \ell_Q \\ g \end{bmatrix}, \qquad r(\mathbf{z}) = \begin{bmatrix} (T - y_T)/\sigma_T \\ \ell_Q/0.10 \\ g/0.10 \\ (T - T_c - 10\,\mathrm{K}\,\mathrm{e}^{\ell_Q - g})/\sigma_p \end{bmatrix},$$

$$\mathbf{J} = \begin{bmatrix} 1/\sigma_T & 0 & 0 \\ 0 & 10 & 0 \\ 0 & 0 & 10 \\ 1/\sigma_p & -10\,\mathrm{e}^{\ell_Q - g}/\sigma_p & 10\,\mathrm{e}^{\ell_Q - g}/\sigma_p \end{bmatrix},$$

with the physics row of equation (3.19) carrying the tolerance $\sigma_p = 0.001\,\mathrm{K}$. At the prior means, $\mathbf{z}_0 = (T_c + 10\,\mathrm{K},\ 0,\ 0)^{\mathsf{T}}$, every residual is zero, the normal equations (3.25) give a zero step, and $\hat{\mathbf{z}} = \mathbf{z}_0$; the curvature term that Chapter 4 adds for nonlinear rows is a sum of residuals times second derivatives and vanishes with the residuals. There the Jacobian is

$$\mathbf{J} = \begin{bmatrix} 5 & 0 & 0 \\ 0 & 10 & 0 \\ 0 & 0 & 10 \\ 1000 & -10\,000 & 10\,000 \end{bmatrix},$$

whose last two columns carry the relation between load and conductance: both are zero in the reading row and equal and opposite in the physics row, so the direction $(0,\ 1,\ 1)^{\mathsf{T}}$, load and conductance raised together, gives $\mathbf{J}\,(0,\ 1,\ 1)^{\mathsf{T}} = (0,\ 10,\ 10,\ 0)^{\mathsf{T}}$ and is seen by the prior rows alone. Multiplying by $\mathbf{J}^{\mathsf{T}}$,

$$\mathbf{J}^{\mathsf{T}}\mathbf{J}\,(0,\ 1,\ 1)^{\mathsf{T}} = 100\,(0,\ 1,\ 1)^{\mathsf{T}},$$

so $(0,\ 1,\ 1)^{\mathsf{T}}$ is an eigenvector of $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ with the eigenvalue $100 = 1/0.10^2$, the prior precision alone, and of $\boldsymbol{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1}$ with the eigenvalue 0.0100, the prior variance: the reading has left

that direction untouched. Inverting $\mathbf{J}^\mathsf{T}\mathbf{J}$ numerically, equation (3.26) gives

$$\boldsymbol{\Sigma} = \begin{bmatrix} 0.039216 & 0.001961 & -0.001961 \\ 0.001961 & 0.005098 & 0.004902 \\ -0.001961 & 0.004902 & 0.005098 \end{bmatrix}.$$

For a Gaussian, the integral of equation (3.16) over T leaves the block of $\boldsymbol{\Sigma}$ in the rows and columns of ℓ_Q and g, the lower right 2×2 block. Its precision is not the corresponding block of $\mathbf{J}^\mathsf{T}\mathbf{J}$ but, by the block-inversion formula of Section 3.5.3 applied to $\boldsymbol{\Lambda} = \mathbf{J}^\mathsf{T}\mathbf{J}$, the Schur complement of the T block, $\boldsymbol{\Lambda}_{qq} - \boldsymbol{\Lambda}_{qT}\boldsymbol{\Lambda}_{TT}^{-1}\boldsymbol{\Lambda}_{Tq}$ with q the indices of ℓ_Q and g:

$$\boldsymbol{\Lambda}_{qq} - \boldsymbol{\Lambda}_{qT}\boldsymbol{\Lambda}_{TT}^{-1}\boldsymbol{\Lambda}_{Tq} = \begin{bmatrix} 2599.94 & -2499.94 \\ -2499.94 & 2599.94 \end{bmatrix} \longrightarrow \begin{bmatrix} 2600 & -2500 \\ -2500 & 2600 \end{bmatrix} \quad (\sigma_p \to 0),$$

against the block $\boldsymbol{\Lambda}_{qq}$ of $\mathbf{J}^\mathsf{T}\mathbf{J}$ itself, whose entries are $10^8 + 100$ on the diagonal and -10^8 off it. Each marginal standard deviation is $\sqrt{0.005098} = 0.0714$, down only from 0.100, and the correlation is $0.004902/0.005098 = 0.962$. The eigenvectors of the block are $(1,\ 1)^\mathsf{T}/\sqrt{2}$ and $(1,\ -1)^\mathsf{T}/\sqrt{2}$, with the eigenvalues 0.005098 ± 0.004902, that is 0.010000 and 0.000196: the posterior is as wide as the prior along the first and 7.1 times narrower along the second. The budget (3.34) with the gradients $\mathbf{g} = (0,\ 1,\ -1)^\mathsf{T}$ and $(0,\ 1,\ 1)^\mathsf{T}$ gives the variance of the difference, $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g} = 0.000392$, a standard deviation of 0.0198, and that of the sum, 0.0200, a standard deviation of 0.1414, exactly its prior value $\sqrt{2} \times 0.10$. Repeated at $\sigma_p = 0.01\,\mathrm{K}$, the block entries print the same, 0.005098 and 0.004902, and the correlation prints 0.961 against 0.962 at 0.001 K; the limit is reached at 0.001 K.

Remark 3.2. Eliminating T through the physics row turns the reading into the linearized row $(\ell_Q - g)/0.02$, since $\sigma_T/(T - T_c) = 0.02$, and the same block follows from a 3×2 Jacobian; along the difference alone, the prior precision 50 plus the reading's 2500 gives $1/2550 = 0.000392$.

Figure 3.7 draws the two distributions in the plane of the two entries. The horizontal axis is ℓ_Q and the vertical axis g; the dashed gray circle encloses 95 % of the prior, the green ellipse 95 % of the posterior, and the thin line is $\ell_Q - g = 0$, the value of the difference that the reading fixed. The posterior is a narrow ellipse lying along that line, its axes the two eigenvectors of the block: it is long in the direction the reading cannot see, $(1,\ 1)^\mathsf{T}$, where it reaches the prior circle because the reading added nothing there, and short in the direction it measured, $(1,\ -1)^\mathsf{T}$.

Read by its marginals, this posterior says that neither the load nor the conductance is well known, about 7 % each. Read by its shape, it says that their ratio, and with it the hall air at the present load, is known to 2 %. A prediction at a different, specified load Q', $T' = T_c + Q'/G$, has a gradient with respect to (T, ℓ_Q, g) proportional to $(0,\ 0,\ 1)^\mathsf{T}$, so $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$ picks out Σ_{gg} and the prediction inherits its 7 % spread; a prediction that scales the present load, such as the hall at twice the unknown Q, has a gradient proportional to $(0,\ 1,\ -1)^\mathsf{T}$, the short axis of the ellipse, and its temperature rise is known to the same 2 % relative as the present one; the absolute spread doubles. Which statement holds depends on the gradient of the question, and the covariance, not the list of marginal variances, carries the answer.

Principle 3.4. *Read the posterior's shape, not its marginals. Wide marginal spreads with a strong correlation mean that the evidence determines a combination of the entries; a quantity that depends on that combination is well known, and the entries are not unobservable merely because each alone is uncertain.*

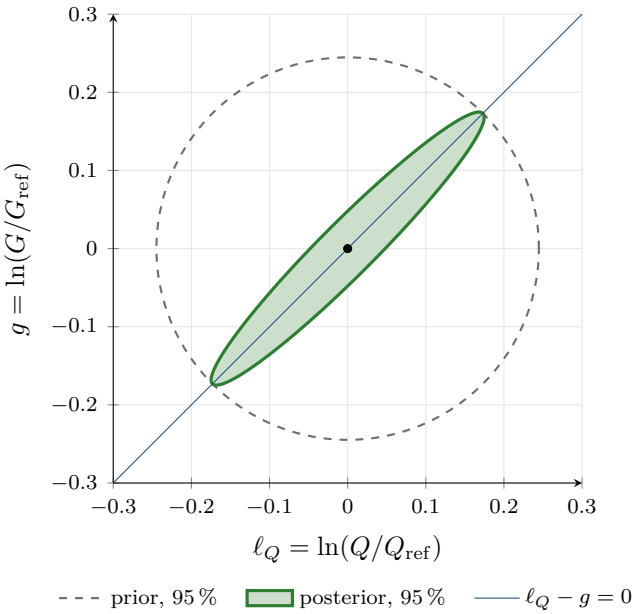


Figure 3.7: Prior and posterior of the log-load ℓ_Q and the log-conductance g after one hall-air reading that fixes $\ell_Q - g$ within 0.02 (illustrative): the (ℓ_Q, g) block of the posterior of $\mathbf{z} = (T, \ell_Q, g)^{\mathsf{T}}$. Prior: independent, standard deviations 0.10. Posterior: standard deviations 0.0714, correlation 0.962; standard deviation of $\ell_Q - g$, 0.0198; of $\ell_Q + g$, 0.1414. Contours enclose 95 % ($\chi^2_{2,0.95} = 5.991$).

3.6.6 Mixtures over structure

A Gaussian approximation is most useful for continuous uncertainty near a well-identified posterior mode. Conditional probability is more general. If the plant may operate under either of two structures, $M \in \{M_1, M_2\}$ (two drawings of the plant), the law of total probability, applied over the models with $\mathbf{z}$ conditioned on the evidence throughout, gives

$$p(\mathbf{z} \mid \mathbf{y}) = \sum_j p(\mathbf{z}, M_j \mid \mathbf{y}) = \sum_j p(\mathbf{z} \mid \mathbf{y}, M_j)\ \Pr(M_j \mid \mathbf{y}), \tag{3.36}$$

where the second step is the multiplication rule. The weights follow from Bayes' theorem over the models, $\Pr(M_j \mid \mathbf{y}) = p(\mathbf{y} \mid M_j)\ \Pr(M_j) / \sum_i p(\mathbf{y} \mid M_i)\ \Pr(M_i)$, so that the evidence of each model, the normalizing constant of its own posterior, sets its share. The expression assumes that the relevant state $\mathbf{z}$ has a common interpretation across the models. Such a posterior need not be Gaussian, even when each model's posterior is.

A two-structure illustration on the hall air shows what is lost. Let the drawing M_1, the air path as designed, give $T \mid \mathbf{y}, M_1 \sim \mathcal{N}(26.9, 0.5^2)$, and a drawing M_2 with a degraded path give $\mathcal{N}(28.5, 0.5^2)$, with posterior model probabilities 0.7 and 0.3 (illustrative). By the law of total expectation the mixture mean is $0.7 \times 26.9 + 0.3 \times 28.5 = 27.38\,^{\circ}\mathrm{C}$; by the law of total variance, equation (3.32) with the model in place of the rate, its variance is the mean of the within-model variances plus the variance of the model means,

$$\mathrm{Var}(T \mid \mathbf{y}) = 0.5^2 + 0.7 \times 0.3 \times (28.5 - 26.9)^2 = 0.25 + 0.5376 = 0.7876\,\mathrm{K}^2,$$

a standard deviation of 0.887 K. The probability of exceeding 28 °C is the weighted sum of the two tails,

$$\Pr(T > 28\,^{\circ}\mathrm{C} \mid \mathbf{y}) = 0.7\,(1 - \Phi(2.2)) + 0.3\,(1 - \Phi(-1)) = 0.7 \times 0.0139 + 0.3 \times 0.8413 = 0.262.$$

The most probable drawing alone would report 0.014, and a single Gaussian with the mixture's mean and variance would report $1 - \Phi(0.699) = 0.242$ and hide the two separate modes, 3.2 standard deviations apart. The law of total probability combines the uncertainty over the structure with the uncertainty over the numerical state, and this combination becomes central in fault detection and isolation and in deciding which drawing of the plant the telemetry supports. The next chapter assembles these pieces into the complete workflow of probabilistic inference, from the plant graph and its likelihood to the posterior, its marginals, its expectations and the probabilities of functions of the state.

3.7 What to retain

Definition 3.1 (Conditioning, marginalization and dependence)**.** A conditional probability restricts attention to the situations consistent with the evidence and renormalizes, $\Pr(A \mid B) = \Pr(A \cap B)/\Pr(B)$; Bayes' theorem reverses the conditioning through the prior and the evidence; marginalization sums or integrates over every unknown the question does not name; and the posterior covariance of a joint state, not its marginal variances, decides how uncertain a derived quantity is.

Four ideas carry forward. Operating questions contain two kinds of uncertainty: aleatory uncertainty is the scatter of the next reading, load or alarm count, and epistemic uncertainty is incomplete knowledge of the plant, its conductance or a meter's calibration; more windows reduce the second and leave the first, and the law of total variance separates them in a prediction. Conditional probability is how information is incorporated: Bayes' theorem updates a probability when an observation arrives, the base rate decides how far an alarm moves it, and the fraction of alarms that are real is not the alarm's probability of detection. Plant quantities are often dependent: a shared chiller, a conservation relation or a common supply air creates relations among halls, units and unknown states, and conditional independence simplifies a calculation only where the model justifies it. A physics graph puts the shared physical causes in the state; the independence it still assumes is that of the errors, and a shared error, such as a calibration offset, is removed by carrying it as an entry too. Marginalization answers a focused question without pretending that the other unknowns are known, whether they are nuisance quantities, competing conditions of a unit or competing drawings of the plant.

Together these establish the mechanism of inference: the prior $p(\mathbf{z})$ and the observation model $p(\mathbf{y} \mid \mathbf{z})$ combine into the posterior $p(\mathbf{z} \mid \mathbf{y})$, the update multiplying the prior by the likelihood and normalizing the product. For Gaussian rows the product is $\exp(-\frac{1}{2} r^{\mathsf{T}} r)$ over the whole state, physics rows included; the normal equations give its most probable state, $\boldsymbol{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1}$ its covariance, exact when the rows are linear, and the conditional of the joint Gaussian of state and readings, through the Schur complement, gives the same answer. Relations between quantities are read in that space: a reading informs an entry it does not measure through a row of $\mathbf{J}$ that carries both, and a combination of entries is well known when it lies along a short eigenvector of $\boldsymbol{\Sigma}$, which $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ measures. Once the posterior is available, the operator's question decides the next calculation: a marginal probability, an expectation, a prediction, a model comparison or a decision. Table 3.11 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 3.11: The terms of uncertainty and conditional probability and where each is defined.

Term	Expression	Units	Meaning	Defined in
Aleatory uncertainty	$\Pr(y_T > c \mid T) = 1 - \Phi((c - T)/\sigma_T)$	–	readings vary with the hall known	3.1.1
Epistemic uncertainty	$p(g \mid \mathbf{y})$	–	knowledge of a fixed quantity	3.1.2
Predictive exceedance	$\int \Pr(y_{\text{new}} > c \mid T)\, p(T \mid \mathbf{y})\, \mathrm{d}T$	–	next reading, averaged over the hall	3.1.3
Conditional probability	$\Pr(A \cap B)/\Pr(B)$	–	probability of A among cases of B	3.2.2
Multiplication rule	$\Pr(A \cap B) = \Pr(B \mid A)\Pr(A)$	–	joint from conditional	3.2.3
Bayes' theorem	$\Pr(B \mid A)\Pr(A)/\Pr(B)$	–	reverses the conditioning	3.2.4
Posterior odds	likelihood ratio × prior odds	–	Bayes without the evidence	3.3.3
Law of total probability	$\sum_s \Pr(H \mid s)\Pr(s)$	–	marginal of an observation	3.4.1
Marginal likelihood	$\int p(\mathbf{y} \mid \lambda) p(\lambda)\, \mathrm{d}\lambda$	–	evidence; normalizer	3.4.2
Quiet-shift probability	$\sum_j w_j\, \mathrm{e}^{-\lambda_{\text{ref}} \mathrm{e}^{\omega_j \tau}} = 0.96457$	–	reported, no alarm over the log-rate, equation (2.41)	3.4.2
Its exact value	$(b/(b+\tau))^a = 0.96863$	–	same average over the gamma; gap 0.0041	3.4.2
Nuisance marginal	$\int p(g, b_Q \mid \mathbf{y})\, \mathrm{d}b_Q$	varies	integrate out, do not maximize	3.4.3
Conditional independence	$X_1 \perp X_2 \mid C$	–	joint given c factorizes	3.4.6
Common-cause joint	$\sum_c \Pr(F_A \mid c)\Pr(F_B \mid c)\Pr(c)$	–	both halls hot with a shared cause	3.4.5
Correlated reading errors	$\mathbf{A}^{\mathsf{T}}\mathbf{R}^{-1}\mathbf{A}$, $\mathbf{R}$ full; or a shared bias entry	K^{-2}	the precision the readings truly add	3.5.7
Gaussian rows	$p(\mathbf{z} \mid \mathbf{y}) \propto \exp(-\frac{1}{2} r^{\mathsf{T}} r)$	–	posterior of the whole state	3.5.2
Normal equations	$(\mathbf{J}^{\mathsf{T}}\mathbf{J})\,\boldsymbol{\Delta} = -\mathbf{J}^{\mathsf{T}} r$, $\boldsymbol{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1}$	varies	most probable state, covariance	3.5.3
Precisions add	$\boldsymbol{\Sigma}^{-1} = \boldsymbol{\Sigma}_0^{-1} + \mathbf{J}_y^{\mathsf{T}}\mathbf{J}_y$	varies	prior plus reading information	3.5.3
Gaussian conditional	$\boldsymbol{\Sigma}_{zz} - \boldsymbol{\Sigma}_{zy}\boldsymbol{\Sigma}_{yy}^{-1}\boldsymbol{\Sigma}_{yz}$	varies	Schur complement	3.5.3
Exceedance probability	$1 - \Phi\big((T_{\max} - \mu_{\text{post}})/\sigma_{\text{post}}\big)$	–	tail of the posterior	3.5.5
Law of total variance	$\mathbb{E}[\mathrm{Var}(K_\tau \mid \lambda)] + \mathrm{Var}(\mathbb{E}[K_\tau \mid \lambda])$	–	process plus rate terms	3.6.1
Posterior covariance	$\boldsymbol{\Sigma}$, $\mathrm{corr}(z_1, z_2) = \Sigma_{12}/\sqrt{\Sigma_{11}\Sigma_{22}}$	varies	joint spread of the state	3.6.4
Variance budget	$\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$	varies	sensitivity^2 × variance, plus cross terms	3.6.4
Mixture posterior	$\sum_j p(\mathbf{z} \mid \mathbf{y}, M_j)\Pr(M_j \mid \mathbf{y})$	varies	uncertainty over structure	3.6.6

3.8 Exercises

Exercise 3.1 (Formulation)**.** (a) In 2000 comparable hall-shifts of an alarm log the alarm fired in 100, a real excursion occurred in 60, and both in 25. Compute $\Pr(H)$, $\Pr(E)$, $\Pr(H \cap E)$,

$\Pr(E \mid H)$ and $\Pr(H \mid E)$, name what the last two mean for the alarm, and state whether alarm and excursion are independent. (0.05, 0.03, 0.0125, $25/100 = 0.25$, the fraction of alarms that are real, and $25/60 = 0.417$, the probability of detection; not independent, since $0.0125 \neq 0.05 \times 0.03 = 0.0015$.) (b) Two halls on one chiller have $\Pr(F_A) = 0.05$, $\Pr(F_B) = 0.04$ for an excursion in a shift and $\Pr(F_A \cap F_B) = 0.015$. Compute the joint probability under an incorrect independence assumption, the probability that at least one hall stays within its limit, and $\Pr(F_B \mid F_A)$, and state what it means for the night. ($0.05 \times 0.04 = 0.002$; $1 - 0.015 = 0.985$, against 0.998 under independence; $0.015/0.05 = 0.30$; both halls are hot together 7.5 times as often as independence predicts, so a night with both halls needing the operator is far more likely than the product suggests.)

Exercise 3.2 (Formulation)**.** A hall satisfies $T = T_c + Q/G$ with T_c known, $G > 0$, and both Q and G uncertain. Starting from $\Pr(T > T_{\max} \mid \mathbf{y})$, write an equivalent event in Q and G and express the probability as an integral over their joint posterior, without assuming independence. (Since $G > 0$, $T_c + Q/G > T_{\max} \iff Q > G(T_{\max} - T_c)$, so $\Pr(T > T_{\max} \mid \mathbf{y}) = \iint \mathbf{1}[Q > G(T_{\max} - T_c)]\, p(Q, G \mid \mathbf{y})\, \mathrm{d}Q\, \mathrm{d}G$; the joint posterior keeps the correlation between Q and G that a reading of T creates.)

Exercise 3.3 (Derivation)**.** (a) A fan-vibration alarm monitors a cooling unit with $\Pr(F) = 0.02$ per shift, $\Pr(H \mid F) = 0.95$ and $\Pr(H \mid \neg F) = 0.05$. Derive the probability that the fan is faulty when the alarm activates, and explain why it is not 95 %. ($\Pr(H) = 0.95 \times 0.02 + 0.05 \times 0.98 = 0.068$; $\Pr(F \mid H) = 0.019/0.068 = 0.279$; 0.95 is $\Pr(H \mid F)$, not $\Pr(F \mid H)$, and the low base rate makes most alarms false.) (b) With prior $\Pr(F) = 0.10$ that the unit is degraded, the coil-approach alarm A has $\Pr(H_A \mid F) = 0.9$, $\Pr(H_A \mid \neg F) = 0.1$, and the hall-temperature alarm B has $\Pr(H_B \mid F) = 0.8$, $\Pr(H_B \mid \neg F) = 0.2$; the alarms are conditionally independent given the unit's state. Derive $\Pr(F \mid H_A, H_B)$ and explain why the assumption is needed. ($\Pr(H_A, H_B \mid F) = 0.9 \times 0.8 = 0.72$ and $\Pr(H_A, H_B \mid \neg F) = 0.1 \times 0.2 = 0.02$, so $\Pr(F \mid H_A, H_B) = 0.072/(0.072 + 0.018) = 0.80$; sequentially, alarm A gives 0.5 and alarm B then 0.8; the likelihoods multiply only under conditional independence given each state of the unit, and a control fault that raises both alarms whatever the unit's state would invalidate it.)

Exercise 3.4 (Derivation)**.** Future alarms satisfy $K_\tau \mid \lambda \sim \text{Poisson}(\lambda\tau)$ with $\lambda \mid \mathbf{y} \sim \text{Gamma}(3, 1000\,\text{h})$ (shape and rate) and $\tau = 200\,\text{h}$, about seventeen shifts. Derive the expected count, the predictive variance and its two contributions, and state what happens to each when a much longer alarm log supporting a similar rate arrives. ($\mathbb{E}[K_\tau \mid \mathbf{y}] = 200 \times 0.003 = 0.6$; by equation (3.33), $\text{Var}(K_\tau \mid \mathbf{y}) = 0.6 + 200^2 \times 3/1000^2 = 0.6 + 0.12 = 0.72$; process 0.6, rate 0.12, a share $200/1200 = 0.167$. With a and b growing in proportion, the rate term $\tau^2(a/b)/b$ tends to zero while the process term stays near 0.6.)

Exercise 3.5 (Application)**.** A hall operates in mild weather with probability 0.90 per shift and an excursion probability of 0.01, and in hot weather, with the chiller at its limit, with probability 0.10 and an excursion probability of 0.15. Derive the unconditional excursion probability of a shift, and the probability that the weather was hot given that an excursion occurred. ($\Pr(E) = 0.90 \times 0.01 + 0.10 \times 0.15 = 0.024$; $\Pr(\text{hot} \mid E) = 0.015/0.024 = 0.625$.)

Exercise 3.6 (Application)**.** (a) A hall air temperature has prior $\mathcal{N}(25, 1.5^2)$, and a sensor with $\sigma_T = 0.5\,\text{K}$ reads $y_T = 26°\text{C}$. For the state, the one number T, write the two residuals and their slopes j_k, the information Λ, take the Newton step from the prior mean, and compute σ^2_{post} and $\Pr(T > 27°\text{C} \mid y_T)$. ($r_1 = (T - 25)/1.5$ and $r_2 = (T - 26)/0.5$, $j_1 = 0.6667$ and $j_2 = 2\,\text{K}^{-1}$, $\Lambda = 0.4444 + 4 = 4.4444\,\text{K}^{-2}$; at $T = 25$, $r_1 = 0$ and $r_2 = -2$, so $\mathcal{L}' = j_1 r_1 + j_2 r_2 = -4$, $\Delta T = 4/4.4444 = 0.9$ and $\hat{T} = 25.90°\text{C}$; $\sigma^2_{\text{post}} = 1/4.4444 = 0.225\,\text{K}^2$, $\sigma_{\text{post}} = 0.4743\,\text{K}$; $1 - \Phi(1.10/0.4743) = 1 - \Phi(2.319) = 0.0102$.) (b) For the two-entry hall of equation (1.10), whose Jacobian at the forward state (23°C, 0) has the rows (100, 500), (2, 0) and (0, 3.333),

compute the three outer products $\mathbf{J}_k^{\mathsf{T}}\mathbf{J}_k$ and their sum, and state which row creates the off-diagonal entry and what the rank of its term is. (Physics $\begin{bmatrix} 10\,000 & 50\,000 \\ 50\,000 & 250\,000 \end{bmatrix}$, reading $\mathrm{diag}(4,\ 0)$, prior $\mathrm{diag}(0,\ 11.11)$; sum $\begin{bmatrix} 10\,004 & 50\,000 \\ 50\,000 & 250\,011 \end{bmatrix}$; only the physics row has entries in both columns, and its term has rank one, determinant $10^4 \times 2.5 \times 10^5 - (5 \times 10^4)^2 = 0$.)

Chapter 4

The Anatomy of Probabilistic Inference

4.1 From a plant graph to an operating probability

At the commissioning review of the equipment hall, the commissioning engineer has the record of its acceptance test on the screen: one cooling unit holding the supply air at 18°C, a test load of 30 kW, and one hall sensor reading 23.2°C. The hall is to be released for 60 kW, and the commissioning test writes its acceptance limit at full load as 28°C. The engineer has that one reading, the design sheet's value for the thermal resistance of the hall's air path and nothing at all from full load. The quick answer of the record uses the design sheet's handbook value: a resistance of 0.160 K/kW puts the hall at $18 + 60 \times 0.160 = 27.6$°C at full load, under the limit, so the hall can be released. The reading gives a second quick answer: its rise of 5.2 K above the supply air, doubled with the load, puts the hall at 28.4°C, just above the limit. These are the commissioning record's own entries, not the nameplate tile's value or the operating limit of Chapter 1: this is the commissioning engineer's version of the night operator's question.

The two quick answers disagree, and neither says how sure the engineer can be: the handbook value ignores the reading, and the scaled reading ignores the handbook and the error of the sensor. What neither settles is what one reading at the test condition, the handbook value and the physics of the hall together say about a condition that has never been run. The question of this chapter is

how a posterior distribution over the state of a plant graph is obtained from physical relations, readings and prior information, and how that posterior is turned into the probability that answers an operator's question.

Chapter 1 built the plant graph, Chapter 2 classified the questions and Chapter 3 derived the rules one at a time; here they are used together for the first time, on a state of three entries of which the reading touches only one.

Jargon: state estimation, maximum a posteriori (MAP) estimate, Hessian, Laplace approximation, delta method, change of variables, credible interval, prior and posterior predictive distributions, marginal likelihood, Bayes factor, rank-one update, reduced state

State estimation is the fit of the state of a plant to its readings and priors; its results are the MAP state and the posterior covariance defined next. The *maximum a posteriori (MAP) estimate* is the value of the unknowns at which the posterior density is highest; it is found by minimizing the negative logarithm of the posterior, the *objective*. The *Hessian* is the matrix of second derivatives of the objective; at the MAP estimate it measures the curvature of the objective, and for rows linear in the state and divided by their standard deviations it is $\mathbf{J}^\mathsf{T}\mathbf{J}$, the information matrix. The *Laplace approximation* replaces the posterior by the Gaussian whose mean is the MAP estimate and whose covariance is the inverse of the Hessian there. The *delta method* approximates the variance of a function of the unknowns by linearizing the function at the estimate. A *change of variables* carries a density from one coordinate to another, multiplying it by the derivative of the old coordinate with respect to the new so that each interval keeps its probability. A *credible interval* contains the unknown with a stated posterior probability. The *prior predictive* and *posterior predictive* distributions are the distributions of a reading not yet seen, averaged over the prior and over the posterior respectively. The *marginal likelihood*, or *evidence*, of a model is the probability density of the observed readings averaged over the prior; the *Bayes factor* is the ratio of the evidences of two models. A *rank-one update* adds one new reading to an existing Gaussian posterior by a closed-form correction of its mean and covariance, without solving the problem again. The *reduced state* is the vector of free entries that remain when the physical rows are taken as exact; every other entry of the state is then a function of it.

Key idea

An operating probability is computed in two operations. Inference, the state estimation of the plant, turns the plant graph, the readings and the priors into a posterior over the state: its mean is the MAP state and its covariance is the inverse curvature of the objective. Query evaluation turns that posterior into the operator's answer: an entry of the state and its variance, or a derived quantity $f(\mathbf{z})$ with variance $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$, and then a normal tail at the limit.

4.1.1 The complete inference workflow

The method fixed in Chapter 2 (Figure 2.6) is a progression: the plant graph (step 0), then the questions written over its state (step 1), the graph extended for them with the evidence attached (step 2), the posterior of the state that the rows give, derived and computed once (step 3), the answers read from that posterior (step 4), their checks (step 5), the report with its bands (step 6), and the next reading when a band is too wide (step 7). Figure 4.1 draws the progression. The wide box at the top is the plant graph, built once; the seven boxes beneath are one pass that carries every question, read left to right along the top row and right to left along the bottom row; the second line of each box names what the step produces. The dashed arrow from step 7 back to step 3 is the loop taken when a new reading is requested: its row is added and the posterior is rebuilt.

The distinction between steps 3 and 4 carries the method, and it is the distinction between the two operations of every inference. Step 3 is *inference*: the rows of step 2 are combined into one posterior of the state, its MAP state, information and covariance, derived and then computed for the plant; it does not depend on which question is asked. Step 4 is *query evaluation*: each question is reduced, from the operator's statement through its general and specialized forms, to quantities read from that posterior, in this chapter the posterior mean and variance

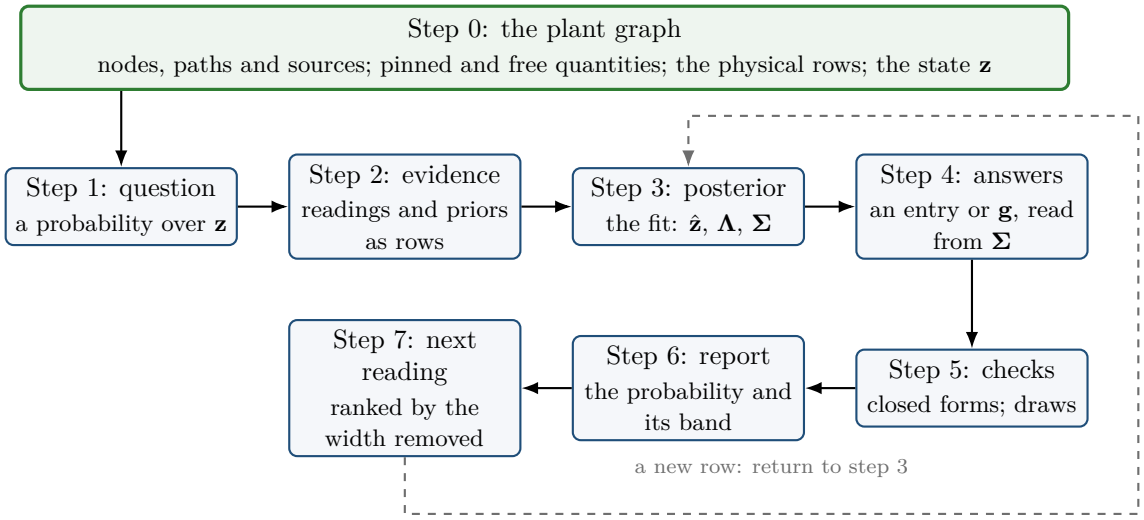


Figure 4.1: The inference workflow. Step 0 is built once; steps 1 to 7 are one pass, carrying every question of the plant. The second line of each box names what the step produces.

of one state entry, the full-load temperature, and the standard normal distribution function evaluated at the limit. Keeping them apart is what lets one posterior serve every question, and a new reading change step 3 alone.

4.1.2 The posterior and the questions put to it

The calculation divides into two operations of different kinds:

$$\underbrace{(\text{graph},\ \mathbf{y}) \longrightarrow p(\mathbf{z} \mid \mathbf{y})}_{\text{inference}}, \qquad \underbrace{p(\mathbf{z} \mid \mathbf{y}) \longrightarrow \text{operating probability or decision}}_{\text{query evaluation}}, \tag{4.1}$$

the second applied to the result of the first. Inference is done once for a given graph and a given body of evidence. Under a Gaussian posterior, exact or approximate, its result is the triple of the MAP state $\hat{\mathbf{z}}$, the posterior covariance $\boldsymbol{\Sigma}$, and the Jacobian of the model at $\hat{\mathbf{z}}$, from which the gradient $\mathbf{g}$ of any derived quantity is formed. Query evaluation is then a formula on the triple. An output that is itself an entry z_i of the state, such as a temperature at a node, has mean $\hat{z}_i$ and variance Σ_{ii}, read straight off the posterior. An output that is a function $f(\mathbf{z})$ of several entries needs its gradient and the quadratic form $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$. A probability that the output stays below a limit is a normal tail on either pair.

Two kinds of further work follow from this division and must be kept apart. A new question asked of the same evidence needs no new inference: a new function of the state is pushed through the existing $\boldsymbol{\Sigma}$. New evidence changes the posterior itself: either the inference is repeated with the new row, or the existing Gaussian is conditioned on the new row by a rank-one update, which Section 4.6.4 derives. The posterior represents what is known about the plant under the model; the queries represent what the operator wants to determine.

4.1.3 The inference in the terms of graphical models

The operations of this book are those of inference in a probabilistic graphical model, under the names a control room uses. The plant graph with its rows is a factor graph (Section 1.1.3): each row is a factor over the entries it touches, and the posterior $\mathrm{e}^{-\mathcal{L}(\mathbf{z})}$ is the normalized product of the factors. Table 4.1 sets each operation beside its name in the literature on graphical models. Its first column is the operation as this book writes it, the second the standard name, and the third the chapters that derive it.

Table 4.1: The operations of the book and their names in the literature on probabilistic graphical models.

Operation in this book	Name in graphical models	Derived in
a row: prior, physics or reading	a factor of a factor graph	Chapters 1, 3
the posterior $\mathrm{e}^{-\mathcal{L}(\mathbf{z})}$	the normalized product of the factors	Chapter 3
the normal equations, $\mathbf{\Lambda} = \mathbf{J}^{\mathsf{T}}\mathbf{J}$, $\mathbf{\Sigma} = \mathbf{\Lambda}^{-1}$	exact inference in a linear-Gaussian factor graph, in information form	Chapter 3; this chapter
the gamma updates of an alarm rate	exact inference with conjugate factors	Chapters 2, 3
Gauss–Newton and the Laplace covariance	the Gaussian approximation at the mode of a nonlinear factor graph (factor-graph optimization)	this chapter; Chapter 5
the Schur complement; the marginal of a joint Gaussian	variable elimination of the entries not asked about	Chapter 3
the ranking of fault hypotheses; total probability	exact inference over discrete hypotheses by summation	Chapter 10
the evidence $p(\mathbf{y} \mid M)$	the normalizing constant (partition function) of the product	Chapter 2
the Laplace evidence against draws	importance sampling, with the Laplace normal as the proposal	Chapter 7

The general algorithms of that literature are not needed for the plants of this book, and the table shows why: every posterior here is solved jointly, by factoring the information matrix of the whole state, which is exact for linear-Gaussian rows and costs little at a few dozen entries. Message passing, belief propagation (the sum-product algorithm) and the junction tree compute the same marginals by local operations on the graph, and they earn their place on graphs too large to factor jointly or with discrete entries of many levels; on a Gaussian graph without loops, belief propagation returns exactly the Schur complements of the joint factorization. Factoring a large sparse information matrix is itself variable elimination, and its cost is set by the order in which the entries are eliminated: eliminated in time order, a state indexed by time, such as the windows of a shift record, gives the Kalman filter and smoother. Markov chain Monte Carlo (Gibbs sampling and the Metropolis algorithm) draws from a posterior that the Laplace normal cannot approximate and reweighting cannot correct: one with several modes, or the funnel of a hierarchical model, in which the spread of a population of rates is itself unknown; a later volume of *Reliability Engineering by Derivation* meets it in a fleet of pumps. Variational inference is not used in this book.

These algorithms differ in what they give up. Variable elimination, the junction tree, belief propagation on a graph without loops and the Kalman smoother on linear-Gaussian rows are exact: they compute the posterior the joint factorization computes, in another order. The Laplace approximation, loopy belief propagation, the linearized Kalman filters and variational inference are approximate, and their error is a bias that more computation does not remove; loopy belief propagation and variational inference in particular tend to report variances too small, which for a tail probability against a limit is the error that matters most. Importance sampling and Markov chain Monte Carlo are exact in the limit of many draws, and their error is random and can be estimated, although a chain that never visits a second mode can look converged. The two kinds of error check each other, which is why the draws serve as the other route against the Laplace answer in step 5. The route also decides which objects come cheaply. The joint factorization gives $\hat{\mathbf{z}}$, $\mathbf{\Sigma}$, the gradient and the preposterior covariance $\mathbf{\Sigma}_a$ of a rank-one update in one pass, and the evidence in closed form under the Laplace approximation; draws give the posterior as a sample, from which a variance is an estimate, a preposterior needs the

candidate's outcomes simulated and the posterior refitted or reweighted, and the evidence needs a further computation, since the draws do not carry the normalizing constant. The method is unchanged by the choice: the plant graph, the questions, the rows and the objective $\mathcal{L}(\mathbf{z})$ are the same for every algorithm, the general forms of Chapter 2 are expectations over the posterior whatever computes it, and only the computations of steps 3 and 4, and the checks of step 5 that test them, depend on the route.

4.1.4 The plant graph under two operating conditions (step 0)

The plant is an equipment hall with a heat-generating load and a single effective heat-removal path to the supply air, which the cooling unit's discharge controller holds at $T_c = 18°\mathrm{C}$. Two operating conditions are considered: a test condition with a heat load of $Q_1 = 30\,\mathrm{kW}$, and the requested full load $Q_2 = 60\,\mathrm{kW}$. At steady state the heat supplied to the hall leaves it through the path, and energy conservation at the hall node reads

$$Q = G\,(T - T_c), \tag{4.2}$$

where T is the hall air temperature in °C, Q the heat load in kW, and G the effective thermal conductance of the air path in kW/K. The supply air is pinned: the controller holds it at its set point, which is what makes T_c a boundary of the graph rather than an unknown, and it fails if the unit runs out of capacity and the discharge drifts up, which Section 4.6.6 takes up. The balance is taken at steady state through one effective path: this reduces the cooling of the hall to a single unknown, and it fails in the minutes after a load change, while heat stored in the hall and its equipment is still moving, which Section 4.6.6 and Exercise 4.6 take up. It is convenient to work with the thermal resistance of the path,

$$\rho = \frac{1}{G}, \tag{4.3}$$

in K/kW, with which equation (4.2) becomes

$$T = T_c + Q\rho. \tag{4.4}$$

Carrying the resistance in the state, rather than the conductance, is a choice: the temperature is then a linear function of the state, which makes every row of the example linear, the posterior exactly Gaussian and every result below exact rather than approximate, and it fails only where a Gaussian on the resistance puts real probability below zero, which Section 4.1.7 measures for the handbook prior and Section 4.6.6 discusses. Chapter 1 carries the conductance in the state as the log-conductance $g = \ln(G/G_{\mathrm{ref}})$, which keeps it positive but makes the temperature nonlinear in the state; the coils of Chapter 5 carry their conductances the same way, as log-derates.

At the two operating conditions equation (4.4) gives

$$T_1 = 18 + 30\rho, \tag{4.5}$$

$$T_2 = 18 + 60\rho, \tag{4.6}$$

with T_1 the hall temperature during the test and T_2 the temperature at full load, both in °C with ρ in K/kW. Figure 4.2 draws the graph. Each row is one operating condition: the box on the left is the heat source, the green box in the middle the hall node carrying a free temperature, and the gray box on the right the supply-air node, pinned at 18°C. The solid arrows are the physical paths, the heat entering the hall and the heat leaving it through the resistance. The green node between the rows is the resistance ρ, free and shared by both paths, as the dashed links show. The dashed arrows leave the graph for the evidence: the reading of the test temperature and the handbook prior on ρ, attached in Section 4.1.7. The full-load node carries no reading; it is the scenario node of a what-if (contingency) analysis that the question

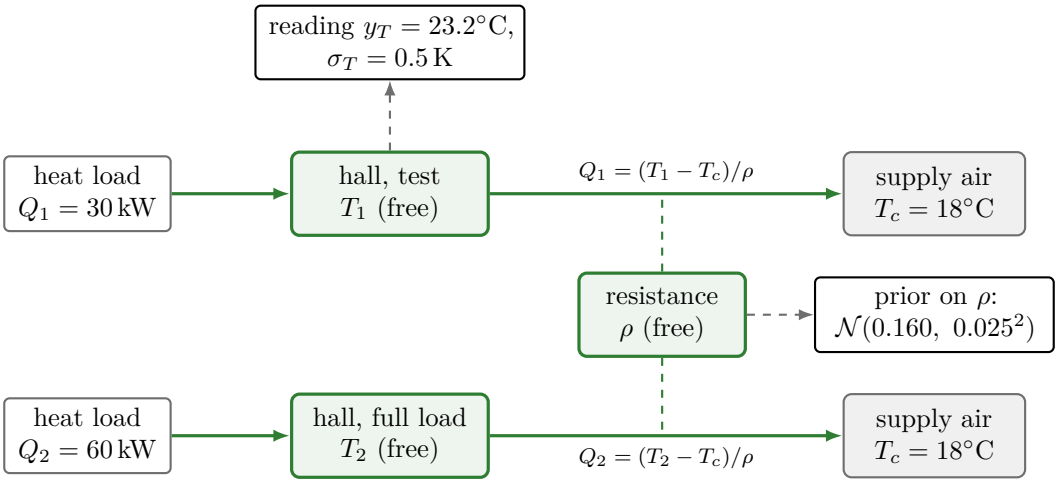


Figure 4.2: The plant graph of the equipment hall under its two operating conditions. Free: the resistance ρ and the temperatures T_1 and T_2 (green). Pinned: the supply-air temperature T_c, held by the cooling unit's discharge controller, and the loads Q_1, Q_2 (gray and white). Observed: T_1, by one reading (dashed arrow, top); ρ carries the handbook prior (dashed arrow, right). The full-load node T_2 is the scenario node of the question and has no reading; it is linked to the test only through the shared resistance.

of Section 4.1.6 requires, drawn here with the rest of the graph so that the figure shows the graph as extended for that question.

The resistance is taken to be the same at 60 kW as at 30 kW, one free entry shared by both rows: this is what lets a reading taken at the test condition carry information about the full-load condition, although the hall has never been operated at full load, and it fails if the path changes with load, for example when the unit's fans or its coil change state, which one test reading cannot detect (Section 4.6.6) and a reading at full load would test (Section 4.6.3). The link between the two rows of Figure 4.2 is not a device of the calculation; it is the physical statement that both conditions are the same hall with the same heat-removal path, and the prediction at full load rests on it.

4.1.5 The state vector and the physical rows

The free quantities of the graph form the state vector,

$$\mathbf{z} = \begin{bmatrix} \rho \\ T_1 \\ T_2 \end{bmatrix}, \tag{4.7}$$

with entries in K/kW, °C and °C. The physical relations are written as rows that vanish when they hold:

$$r_1(\mathbf{z}) = T_1 - 18 - 30\rho = 0, \tag{4.8}$$
$$r_2(\mathbf{z}) = T_2 - 18 - 60\rho = 0. \tag{4.9}$$

These are the conservation rows of the graph, one for each hall node. The balance is taken to hold at both conditions up to a small model tolerance σ_p, a standard deviation that represents model discrepancy, and the answers are reported in the limit $\sigma_p \to 0$ of exact balances: the finite tolerance keeps every matrix of the calculation invertible while the reported numbers are those of an exact balance, and the limit understates the spread if the real discrepancy is finite,

which Section 4.4.5 measures over three tolerances and Section 4.6.6 discusses. Section 4.2.3 puts the rows, with their tolerance, into one residual vector with the reading and the prior.

The three entries of the state are therefore not three independent unknowns. The two physical rows leave one direction of the state free, a change of ρ with both temperatures following it; in the limit of exact rows the state has three entries but one degree of freedom, and this fact governs the shape of the posterior covariance derived in Section 4.2.5 and the reduced state of Section 4.2.8.

4.1.6 The operator's questions (step 1)

The release raises the heat load from 30 kW to 60 kW. The maximum acceptable hall temperature, the acceptance limit of the commissioning test, is $T_{\max} = 28°\mathrm{C}$. In the operator's words the question is: given the hall sensor's reading and what is known of the air path, what is the probability that the hall stays below 28°C at full load? As a probability over the state, with the unknown on the left of the bar and everything treated as known on the right,

$$\Pr(T_2 < T_{\max} \mid \mathbf{y}, M), \tag{4.10}$$

where T_2 is the full-load entry of the state, $\mathbf{y}$ the evidence (the test reading and the handbook prior), and M the assumed steady-state thermal model, equation (4.4) with the rows (4.8)–(4.9). The model is suppressed from the notation below when it is not in question. This is a question of the category *prediction under a threshold*: a tail probability of one marginal of the posterior.

Three estimation questions are asked of the same evidence, one for each entry of the state: the thermal resistance of the air path, the hall temperature at the test condition, and the hall temperature at full load. Table 4.2 lists the four questions: the first column is the question in the operator's words; the second the general form that every question of its category takes, written for a quantity $f(\mathbf{z})$ derived from the state and a threshold f^* on it; the third its mathematical form for the hall; the fourth its category; and the last the subsection where it is answered. Which further reading would narrow the answers most asks nothing new of the present evidence: it is the next reading for these four questions, answered in Section 4.6.

Table 4.2: The questions asked of the hall's test reading, the general form of their category, their mathematical form for the hall, their category and where each is answered. $\mathbf{y}$ is the test reading and the handbook prior; $T_{\max} = 28°\mathrm{C}$.

Operator question	General form	Mathematical form	Category	Answered in
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	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	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	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	4.3.7

Because the full-load temperature is carried as an entry of the state, the threshold question needs exactly one marginal of the posterior: the mean $\hat{z}_3$ and the variance Σ_{33} of the third entry: the question is written over the state, and the position of T_2 in the state shows which part of the posterior it needs. Each estimation question decomposes in the same way, into one entry of $\hat{\mathbf{z}}$ and one diagonal entry of $\boldsymbol{\Sigma}$.

4.1.7 The evidence as rows (step 2)

The temperature reading. During the 30 kW test the hall sensor reports $y_T = 23.2°\text{C}$, with a measurement error of standard deviation $\sigma_T = 0.5\,\text{K}$. The observation model writes the random reading Y as the true test temperature, the second entry of the state, plus an error,

$$Y = T_1 + \epsilon = \mathbf{h}^\mathsf{T}\mathbf{z} + \epsilon, \qquad \mathbf{h} = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \qquad \epsilon \sim \mathcal{N}(0, \sigma_T^2), \tag{4.11}$$

where $\mathbf{h}$ selects the entry the sensor measures. The error is taken as Gaussian, unbiased, of known spread and independent of the state: this makes the reading one linear row of fixed weight $1/\sigma_T^2$, and it fails if the sensor is biased, which would shift every inferred quantity; the prior predictive check of Section 4.4.2 would flag a large bias, and the comparison of models of Section 4.2.2 is the test built for it. The reading lands on T_1 only; it is the physical row of the test hall, not a substitution, that will carry it to ρ.

Prior information. The design sheet gives a preliminary value of the resistance, the handbook figure for an air path of this kind, represented by the Gaussian prior

$$\rho \sim \mathcal{N}(\mu_0, \sigma_0^2), \qquad \mu_0 = 0.160\,\text{K/kW}, \quad \sigma_0 = 0.025\,\text{K/kW}. \tag{4.12}$$

The prior mean corresponds to a conductance $G_0 = 1/0.160 = 6.25\,\text{kW/K}$. Taking the handbook value as the prior for this hall is a choice: it is what places the full-load temperature at 27.6°C before any reading and what keeps the estimate from following the one reading alone, and it misleads if the handbook describes air paths unlike this one, a conflict that the prior predictive check of Section 4.4.2 would show as a reading far from 22.8°C. A Gaussian prior on a resistance assigns some probability to negative, nonphysical values; here it is $\Phi(-0.160/0.025) = \Phi(-6.4) \approx 8 \times 10^{-11}$, negligible, and the Gaussian is kept for its analytical simplicity. A production model enforces positivity through a logarithmic parameterization, such as the log-conductance g of Chapter 1, or a prior with positive support.

The rows and the entries they act on. The physics and the evidence put four rows on the state: the physical rows of the test hall and of the full-load hall, equations (4.8)–(4.9), the reading row on T_1, and the prior row on ρ. Which row acts on which entry is a property of the graph, and it fixes where the Jacobian of the rows can be nonzero, marked $*$:

$$\mathbf{J} = \begin{bmatrix} * & * & 0 \\ * & 0 & * \\ 0 & * & 0 \\ * & 0 & 0 \end{bmatrix} \begin{matrix} \text{physics, test hall} \\ \text{physics, full-load hall} \\ \text{reading} \\ \text{prior} \end{matrix}, \qquad \text{columns } \rho, T_1, T_2. \tag{4.13}$$

A line of the pattern lists the entries one row sees, and a column lists the rows that see one entry. The third column has a single mark: the full-load temperature is seen by its own physical row and by nothing else. The second column shows that the reading and the test hall's balance are the two rows that see T_1, and the first that the two balances and the prior see ρ. No row joins the reading to T_2; the two physical rows, through the column ρ they share, are the only path from the one to the other.

Knowns and unknowns. Table 4.3 lists every quantity of the problem. The first column is the symbol, the second its value, and the third its role in the graph: pinned boundary condition, evidence, or free entry of the state. The graph and the probabilistic model are now specified completely enough for the inference to be carried out.

Table 4.3: Knowns and unknowns of the equipment-hall example.

Quantity	Value	Role
T_c [°C]	18	pinned: supply air, held by the discharge controller
Q_1 [kW]	30	pinned: test load
Q_2 [kW]	60	pinned: requested load
y_T [°C]	23.2	evidence: the test reading
σ_T [K]	0.5	evidence: standard deviation of the reading error
μ_0 [K/kW]	0.160	evidence: prior mean of the resistance
σ_0 [K/kW]	0.025	evidence: prior standard deviation of the resistance
σ_p [K]	0.01, then $\to 0$	model: tolerance of the two physical rows
ρ [K/kW]	unknown	free: thermal resistance, the first entry of the state
T_1, T_2 [°C]	unknown	free: test and full-load temperatures, tied to ρ by the physical rows

Remark 4.1. Four spreads appear in what follows and must not be confused. The reading error has standard deviation σ_T (no index), a property of the sensor. The prior on the resistance has standard deviation σ_0. The physical rows have the model tolerance σ_p, a property of the steady-state model. The posterior standard deviation of the i-th entry of the state is $\sqrt{\Sigma_{ii}}$, the square root of a diagonal entry of the posterior covariance, written $\mathrm{sd}(T_2 \mid \mathbf{y})$ when the entry is named.

4.2 The posterior of the state (step 3)

Every question of Table 4.2 is a property of one density, the posterior of the state given the reading and the handbook prior. Each row of Section 4.1.7 contributes one factor to it, and the MAP state, the information matrix and the covariance are properties of that product. This section derives the posterior and computes it for the hall; Section 4.3 reads every answer from it.

4.2.1 The likelihood of the reading and the physical rows

The likelihood states how probable the observed reading would be under a proposed state. By equation (4.11), for a fixed state the reading is the test temperature T_1 plus a Gaussian error, so $Y \mid \mathbf{z} \sim \mathcal{N}(T_1,\ \sigma_T^2)$, and its density at the observed value is

$$p(y_T \mid \mathbf{z}) = \frac{1}{\sqrt{2\pi\sigma_T^2}} \exp\left[-\frac{(T_1 - y_T)^2}{2\sigma_T^2}\right]. \tag{4.14}$$

For the reading $y_T = 23.2°\mathrm{C}$ and $\sigma_T = 0.5\,\mathrm{K}$, so that $2\sigma_T^2 = 0.5\,\mathrm{K}^2$,

$$p(y_T \mid \mathbf{z}) \propto \exp\left[-\frac{(T_1 - 23.2)^2}{0.5}\right]. \tag{4.15}$$

The likelihood is a function of the state with the reading held fixed; it is not a density in $\mathbf{z}$ and need not integrate to one over it. It depends on the first temperature only and is flat in ρ and T_2: the sensor measures a temperature and says nothing directly about the resistance or about the full-load hall.

What connects the three entries is the pair of physical rows. A steady-state balance that holds up to a model tolerance σ_p (a standard deviation in kelvin, the amount by which the balance of equation (4.4) is allowed to fail) is the density

$$p(T_1 \mid \rho) = \mathcal{N}\big(T_1;\ 18 + 30\rho,\ \sigma_p^2\big), \qquad p(T_2 \mid \rho) = \mathcal{N}\big(T_2;\ 18 + 60\rho,\ \sigma_p^2\big), \tag{4.16}$$

where $\mathcal{N}(x; m, v)$ denotes the Gaussian density with mean m and variance v evaluated at x. The exact rows (4.8)–(4.9) are the limit $\sigma_p \to 0$, in which each density concentrates on its row. The derivation below keeps σ_p finite, carries all three entries, and takes the limit at the end, where it is checked numerically; no row is used to eliminate an entry.

4.2.2 The prior density and Bayes' theorem

By equation (4.12) the prior density of the resistance is

$$p(\rho) = \frac{1}{\sqrt{2\pi\sigma_0^2}} \exp\left[-\frac{(\rho - \mu_0)^2}{2\sigma_0^2}\right] \propto \exp\left[-\frac{(\rho - 0.160)^2}{2(0.025)^2}\right]. \tag{4.17}$$

The prior of the state is the product of this density with the two physical densities (4.16), $p(\mathbf{z}) = p(T_1 \mid \rho)\, p(T_2 \mid \rho)\, p(\rho)$, the factorization of Chapter 3 read along the graph: each hall temperature depends on the resistance, and on nothing else. Bayes' theorem, derived in that chapter, gives the posterior of the state as

$$p(\mathbf{z} \mid y_T) = \frac{p(y_T \mid \mathbf{z})\, p(\mathbf{z})}{p(y_T)}, \qquad p(y_T) = \int p(y_T \mid \mathbf{u})\, p(\mathbf{u})\, \mathrm{d}\mathbf{u},$$

with $\mathbf{u}$ a variable of integration over the state. The denominator does not depend on $\mathbf{z}$, so

$$p(\mathbf{z} \mid y_T) \propto p(y_T \mid \mathbf{z})\, p(\mathbf{z}).$$

The evidence of the model. The denominator has a name of its own. Writing the model M explicitly, it is the evidence, or marginal likelihood, of the model,

$$p(y_T \mid M) = \int p(y_T \mid \mathbf{z}, M)\, p(\mathbf{z} \mid M)\, \mathrm{d}\mathbf{z}\,: \tag{4.18}$$

the density of the observed reading averaged over the prior of the state. For the linear-Gaussian model it is the prior predictive density of the reading evaluated at the observed value (Section 4.4.2). It also follows from the objective and its curvature, which is how it is computed when no closed form exists. The integrand is $\exp[-\mathcal{L}(\mathbf{z})]$, with the objective $\mathcal{L}$ of equation (4.24), times the normalizing constants $(2\pi\sigma_k^2)^{-1/2}$ of the four rows; when the objective is quadratic about its minimum $\hat{\mathbf{z}}$ with matrix $\mathbf{J}^\mathsf{T}\mathbf{J}$, whose inverse is the covariance $\boldsymbol{\Sigma}$ of Section 4.2.5, the Gaussian integral over the three entries gives

$$p(y_T \mid M) = \frac{(2\pi)^{3/2}\ \det(\boldsymbol{\Sigma})^{1/2}}{(2\pi)^2\, \sigma_p^2\, \sigma_T\, \sigma_0}\, \mathrm{e}^{-\mathcal{L}(\hat{\mathbf{z}})}. \tag{4.19}$$

The two directions of the state that the physical rows fix have posterior widths proportional to σ_p, and they cancel the two physical rows' normalizing constants: $\det(\boldsymbol{\Sigma})^{1/2}/\sigma_p^2$ tends to the width of the posterior along the free direction. The result is the volume of the posterior, $\det(\boldsymbol{\Sigma})^{1/2}$, times the height of the integrand at the MAP state. The formula is exact when the objective is quadratic, as it is here; in general it is the Laplace approximation to the evidence.

The evidence normalizes the posterior, and its second use is to compare models. Alternatives for the hall include M_1, a single thermal resistance; M_2, two distinct heat-removal paths; and

M_3, a model with a biased hall sensor. Each gives its own evidence for the same reading, and their ratio is a Bayes factor,

$$B_{21} = \frac{p(y_T \mid M_2)}{p(y_T \mid M_1)}, \tag{4.20}$$

which multiplies the prior odds of the two models to give their posterior odds, $\Pr(M_2 \mid y_T)/\Pr(M_1 \mid y_T) = B_{21} \cdot \Pr(M_2)/\Pr(M_1)$. That comparison is the category *hypothesis comparison*, and Chapters 7 and 10 develop it.

The posterior up to its constant. The data $\mathbf{y}$, the evidence in the sense of what is known about the hall, consists of the reading and the handbook prior; the likelihood involves only the reading, and the prior enters through $p(\mathbf{z})$, so $p(\mathbf{z} \mid \mathbf{y})$ and $p(\mathbf{z} \mid y_T)$ name the same density. Substituting equations (4.14), (4.16) and (4.17) and collecting the four exponents,

$$p(\mathbf{z} \mid \mathbf{y}) \propto \exp\left[-\frac{1}{2}\left(\frac{(T_1 - 18 - 30\rho)^2}{\sigma_p^2} + \frac{(T_2 - 18 - 60\rho)^2}{\sigma_p^2} + \frac{(T_1 - y_T)^2}{\sigma_T^2} + \frac{(\rho - \mu_0)^2}{\sigma_0^2}\right)\right]. \tag{4.21}$$

This is the posterior of the state up to its normalizing constant. It combines the physical relations, the reading, the reading's uncertainty and the prior knowledge in one probability model, one term per row. Every later result of the chapter is a property of this expression.

4.2.3 The residual vector, the Jacobian and the information matrix

The maximum a posteriori estimate is the state at which the posterior density is greatest,

$$\hat{\mathbf{z}} = \arg\max_{\mathbf{z}} \; p(\mathbf{z} \mid \mathbf{y}). \tag{4.22}$$

The logarithm is increasing, so maximizing the posterior is equivalent to minimizing its negative logarithm. Dropping the terms that do not depend on $\mathbf{z}$ (the normalizing constants of the four densities and $\ln p(y_T)$) leaves half the sum of the squared exponents of (4.21). Each exponent is a row divided by its standard deviation, and the four rows stack into one residual vector,

$$\mathbf{r}(\mathbf{z}) = \begin{bmatrix} (T_1 - 18 - 30\rho)/\sigma_p \\ (T_2 - 18 - 60\rho)/\sigma_p \\ (T_1 - y_T)/\sigma_T \\ (\rho - \mu_0)/\sigma_0 \end{bmatrix} \begin{matrix} \text{physics, test hall} \\ \text{physics, full-load hall} \\ \text{reading} \\ \text{prior} \end{matrix} \tag{4.23}$$

whose entries are dimensionless, so that the objective is

$$\mathcal{L}(\mathbf{z}) = \tfrac{1}{2}\,\mathbf{r}(\mathbf{z})^{\mathsf{T}}\mathbf{r}(\mathbf{z}), \qquad p(\mathbf{z} \mid \mathbf{y}) \propto \exp[-\mathcal{L}(\mathbf{z})]. \tag{4.24}$$

In the notation of the book it is $\mathcal{L} = \frac{1}{2}\sum_k w_k r_k^2$ over the unnormalized rows r_1, r_2 of equations (4.8)–(4.9), $T_1 - y_T$ and $\rho - \mu_0$, with weights $w_k = 1/\sigma_k^2$; the normalized rows carry their weights inside them. The first two terms penalize a state that unbalances a hall, in units of the model tolerance; the third, a state that disagrees with the reading, in units of the sensor spread; the fourth, a resistance far from its handbook value, in units of the prior spread. The weight of a physical row has a statistical meaning: $1/\sigma_p^2$ is the inverse variance of a model discrepancy, the amount by which the steady-state balance may fail.

The Jacobian. The derivatives of the four rows with respect to the three entries (ρ, T_1, T_2) form the 4×3 Jacobian $\mathbf{J} = \partial\mathbf{r}/\partial\mathbf{z}$, one line per row and one column per entry:

$$\mathbf{J} = \begin{bmatrix} -30/\sigma_p & 1/\sigma_p & 0 \\ -60/\sigma_p & 0 & 1/\sigma_p \\ 0 & 1/\sigma_T & 0 \\ 1/\sigma_0 & 0 & 0 \end{bmatrix} = \begin{bmatrix} -3000 & 100 & 0 \\ -6000 & 0 & 100 \\ 0 & 2 & 0 \\ 40 & 0 & 0 \end{bmatrix} \quad (\sigma_p = 0.01\,\mathrm{K}), \tag{4.25}$$

with the first column in kW/K and the other two in K^{-1}. The entry J_{ki} is the change in row k, counted in units of that row's standard deviation, per unit change of entry i. Its nonzero entries sit where the pattern (4.13) of the graph places them. The entries $2 = 1/0.5$ and $40 = 1/0.025$ are the normalizing factors of the reading and the prior. Every row is linear in the state, so $\mathbf{J}$ is the same at every linearization point. The tolerance $\sigma_p = 0.01\,\mathrm{K}$, one-fiftieth of the sensor's spread, is tight enough that the answers no longer move (the limit is checked in Section 4.4.5) and loose enough, within two decades of σ_T, to keep the normal equations conditioned well enough for double precision.

The information matrix. For rows normalized by their standard deviations the information matrix is $\mathbf{\Lambda} = \mathbf{J}^\mathsf{T}\mathbf{J}$, equation (3.26). Written with the unnormalized rows, whose Jacobian is $\tilde{\mathbf{J}}$, and the diagonal weight matrix $\mathbf{W} = \mathrm{diag}(1/\sigma_p^2,\ 1/\sigma_p^2,\ 1/\sigma_T^2,\ 1/\sigma_0^2)$, it is the same matrix in the form used throughout the book,

$$\mathbf{\Lambda} = \mathbf{J}^\mathsf{T}\mathbf{J} = \tilde{\mathbf{J}}^\mathsf{T}\mathbf{W}\tilde{\mathbf{J}} = \mathbf{J}_\mathrm{d}^\mathsf{T}\mathbf{W}_\mathrm{d}\mathbf{J}_\mathrm{d} + \mathbf{\Lambda}_\mathrm{prior}, \tag{4.26}$$

where the last form keeps the prior rows apart: $\mathbf{J}_\mathrm{d}$ and $\mathbf{W}_\mathrm{d}$ hold the physical and reading rows, and $\mathbf{\Lambda}_\mathrm{prior}$, here $\mathrm{diag}(1/\sigma_0^2, 0, 0)$, is the prior precision. For rows that are nonlinear in the state, equation (4.26) is the Gauss–Newton approximation to the Hessian of $\mathcal{L}$: it drops the curvature term $\sum_k r_k\,\partial^2 r_k/\partial\mathbf{z}\,\partial\mathbf{z}^\mathsf{T}$, the second derivatives of the rows weighted by their residuals, which vanishes for linear rows. Writing $\mathbf{j}_k$ for the k-th line of $\mathbf{J}$, the product is a sum of one outer product per row, so the information of independent rows adds, as equation (3.27) states for priors and readings; with the numbers of equation (4.25),

$$\mathbf{J}^\mathsf{T}\mathbf{J} = \sum_{k=1}^{4} \mathbf{j}_k^\mathsf{T}\mathbf{j}_k = \begin{bmatrix} 45\,001\,600 & -300\,000 & -600\,000 \\ -300\,000 & 10\,004 & 0 \\ -600\,000 & 0 & 10\,000 \end{bmatrix}. \tag{4.27}$$

The entry in the corner is $3000^2 + 6000^2 + 40^2$: the two physical rows and the prior. The reading contributes only the $4 = 1/\sigma_T^2$ of the (T_1, T_1) entry. The reading and prior rows alone would give $\mathrm{diag}(1600, 4, 0)$: no information at all would reach T_2, and none would pass from the reading to ρ. It is the off-diagonal entries of the physical rows that carry the reading at T_1 to ρ and from ρ to T_2. The matrix does not contain the reading's value y_T, only its standard deviation: the information is fixed by the design of the measurement and by the prior, not by the value observed. Its eigenvalues are 1.155, 1.0003×10^4 and 4.5012×10^7, all positive, with a condition number of 3.9×10^7 that grows as $1/\sigma_p^2$.

4.2.4 The normal equations and the MAP state

The stationary condition. By the chain rule the gradient of the objective (4.24) is

$$\nabla\mathcal{L}(\mathbf{z}) = \mathbf{J}^\mathsf{T}\mathbf{r}(\mathbf{z}), \tag{4.28}$$

and the MAP state is the point where it vanishes, $\mathbf{J}^\mathsf{T}\mathbf{r}(\hat{\mathbf{z}}) = \mathbf{0}$: three equations, one per entry.

The Gauss–Newton step. The step that sets the gradient of the linearized rows to zero solves the normal equations of equation (3.25),

$$(\mathbf{J}^\mathsf{T}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^\mathsf{T}\mathbf{r}(\mathbf{z}), \qquad \mathbf{z}^{(k+1)} = \mathbf{z}^{(k)} + \Delta\mathbf{z}, \tag{4.29}$$

the Gauss–Newton iteration. The rows of this example are linear, so the linearization is exact and one step lands on $\hat{\mathbf{z}}$ from any starting state.

The free direction. The two physical lines of $\mathbf{J}$ annihilate one direction of the state,

$$\mathbf{J}_{\text{phys}}\,\mathbf{v} = \mathbf{0}, \qquad \mathbf{J}_{\text{phys}} = \begin{bmatrix} -3000 & 100 & 0 \\ -6000 & 0 & 100 \end{bmatrix}, \quad \mathbf{v} = \begin{bmatrix} 1 \\ 30 \\ 60 \end{bmatrix}, \tag{4.30}$$

the free direction: a change of ρ with T_1 and T_2 following by $Q_1 = 30\,\text{kW}$ and $Q_2 = 60\,\text{kW}$ times that change, which unbalances neither hall. A pull of the reading on T_1 is answered by a move along $\mathbf{v}$, because any move off it would cost the physical rows far more than it gains on the reading; from a state that balances both halls, the step of the normal equations therefore lies, to within the tolerance, along $\mathbf{v}$.

4.2.5 The posterior covariance from the information matrix

The point estimate does not answer the operating question, which depends on how uncertain the state remains.

The covariance. The rows are linear and Gaussian, so by equation (3.26) the posterior is exactly $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ with

$$\boldsymbol{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1},$$

the inverse of the matrix (4.27), with rows and columns ordered (ρ, T_1, T_2), so that its entries carry $(\text{K/kW})^2$, K^2, and K·K/kW for the products. A larger information means a narrower posterior: the covariance is the inverse of the steepness of the objective.

The marginals. The estimation questions ask for entries. The marginals of a Gaussian are read off its mean and covariance,

$$\rho \mid \mathbf{y} \sim \mathcal{N}(\hat{z}_1,\ \Sigma_{11}), \qquad T_1 \mid \mathbf{y} \sim \mathcal{N}(\hat{z}_2,\ \Sigma_{22}), \qquad T_2 \mid \mathbf{y} \sim \mathcal{N}(\hat{z}_3,\ \Sigma_{33}), \tag{4.31}$$

so the resistance is $\hat{\rho} = \hat{z}_1$ with standard deviation $\sqrt{\Sigma_{11}}$, the test temperature $\hat{z}_2$ with $\sqrt{\Sigma_{22}}$, and the full-load temperature, the entry the threshold question asks about, $\hat{T}_2 = \hat{z}_3$ with $\sqrt{\Sigma_{33}}$. How the entries move together is read off the correlations $\Sigma_{ij}/\sqrt{\Sigma_{ii}\Sigma_{jj}}$, which the physical rows drive toward one.

Coordinates along and across the physical rows. The shape of $\boldsymbol{\Sigma}$, and its limit as the tolerance shrinks, are read by measuring the state along the free direction and across the two physical rows. Write

$$\mathbf{z} = \mathbf{c} + \mathbf{B}\mathbf{u}, \qquad \mathbf{u} = \begin{bmatrix} \rho \\ d_1 \\ d_2 \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} 1 & 0 & 0 \\ 30 & 1 & 0 \\ 60 & 0 & 1 \end{bmatrix}, \quad \mathbf{c} = \begin{bmatrix} 0 \\ 18 \\ 18 \end{bmatrix}, \tag{4.32}$$

where $d_1 = T_1 - 18 - 30\rho$ and $d_2 = T_2 - 18 - 60\rho$ are the two hall imbalances in kelvin and the first column of $\mathbf{B}$ is $\mathbf{v}$. The rows are linear in $\mathbf{u}$ with Jacobian $\mathbf{JB}$, whose columns are $\mathbf{J}\mathbf{v}$, $\mathbf{J}\mathbf{e}_2$ and $\mathbf{J}\mathbf{e}_3$,

$$\mathbf{JB} = \begin{bmatrix} 0 & 1/\sigma_p & 0 \\ 0 & 0 & 1/\sigma_p \\ 30/\sigma_T & 1/\sigma_T & 0 \\ 1/\sigma_0 & 0 & 0 \end{bmatrix} = \begin{bmatrix} 0 & 100 & 0 \\ 0 & 0 & 100 \\ 60 & 2 & 0 \\ 40 & 0 & 0 \end{bmatrix}.$$

The first column, $\mathbf{J}\mathbf{v} = (0,\ 0,\ 60,\ 40)^{\mathsf{T}}$, is seen only by the reading (with slope $2 \times 30 = 60$) and the prior (slope 40); a move along $\mathbf{v}$ costs the information $\|\mathbf{J}\mathbf{v}\|^2 = 60^2 + 40^2 = 5200\,(\text{kW/K})^2$. The other two columns carry the entries $1/\sigma_p$, so every direction that unbalances a hall costs

about $1/\sigma_p^2 = 10^4\,\mathrm{K}^{-2}$: a tiny tolerance leaves the state free to move cheaply only along $\mathbf{v}$, a narrow ridge in the objective. The third column is orthogonal to the other two, so the information matrix in these coordinates is block diagonal,

$$(\mathbf{JB})^{\mathsf{T}}(\mathbf{JB}) = \begin{bmatrix} 1/\sigma_0^2 + 30^2/\sigma_T^2 & 30/\sigma_T^2 & 0 \\ 30/\sigma_T^2 & 1/\sigma_p^2 + 1/\sigma_T^2 & 0 \\ 0 & 0 & 1/\sigma_p^2 \end{bmatrix} = \begin{bmatrix} 5200 & 120 & 0 \\ 120 & 10\,004 & 0 \\ 0 & 0 & 10\,000 \end{bmatrix},$$

and its inverse, the covariance $\mathbf{\Sigma}_u$ of $\mathbf{u}$, is

$$\mathbf{\Sigma}_u = \begin{bmatrix} 1.9236 \times 10^{-4} & -2.3074 \times 10^{-6} & 0 \\ -2.3074 \times 10^{-6} & 9.9988 \times 10^{-5} & 0 \\ 0 & 0 & 1.0000 \times 10^{-4} \end{bmatrix} \quad (\sigma_p = 0.01\,\mathrm{K}).$$

Three results follow. First, the full-load imbalance d_2 is seen by its own row only: $(\mathbf{\Sigma}_u)_{33} = \sigma_p^2$, the variance that row allows, and it is uncorrelated with ρ and d_1. Second, by the block-inversion formula of Section 3.5.3, the first diagonal entry of $\mathbf{\Sigma}_u$ is the inverse of the Schur complement of the (d_1, d_1) entry in the upper 2×2 block,

$$\Sigma_{11} = (\mathbf{\Sigma}_u)_{11} = \left[\frac{1}{\sigma_0^2} + \frac{30^2}{\sigma_T^2} - \frac{(30/\sigma_T^2)^2}{1/\sigma_p^2 + 1/\sigma_T^2} \right]^{-1} = \left[\frac{1}{\sigma_0^2} + \frac{30^2}{\sigma_T^2 + \sigma_p^2} \right]^{-1}, \tag{4.33}$$

the first row of $\mathbf{B}$ being $\mathbf{e}_1^{\mathsf{T}}$; a tolerance acts on the reading as extra sensor variance and widens the spread of ρ. Third, the covariance of the state is $\mathbf{\Sigma} = \mathbf{B}\mathbf{\Sigma}_u\mathbf{B}^{\mathsf{T}}$; its third diagonal entry, with $T_2 = 18 + 60\rho + d_2$ and d_2 uncorrelated with ρ, is

$$\mathrm{Var}(T_2 \mid \mathbf{y}) = \Sigma_{33} = \mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g} + \sigma_p^2 = 60^2\,\Sigma_{11} + \sigma_p^2, \qquad \mathbf{g} = (60,\ 0,\ 0)^{\mathsf{T}}, \tag{4.34}$$

with $\mathbf{g}$ the gradient of the prediction $18 + 60\rho$. A model discrepancy therefore enters the answer twice: it widens Σ_{11}, and it adds its own variance at the node that has no reading.

The limit of exact rows. As $\sigma_p \to 0$, $(\mathbf{\Sigma}_u)_{22}$ and $(\mathbf{\Sigma}_u)_{33}$, both at most σ_p^2, vanish and $(\mathbf{\Sigma}_u)_{11}$ tends to $1/5200$ by equation (4.33), so $\mathbf{\Sigma}_u \to \mathrm{diag}(1/5200,\ 0,\ 0)$ and

$$\mathbf{\Sigma} = \mathbf{B}\mathbf{\Sigma}_u\mathbf{B}^{\mathsf{T}} \ \longrightarrow\ \frac{\mathbf{B}\mathbf{e}_1\mathbf{e}_1^{\mathsf{T}}\mathbf{B}^{\mathsf{T}}}{5200} = \frac{\mathbf{v}\mathbf{v}^{\mathsf{T}}}{5200} = \frac{1}{5200}\begin{bmatrix} 1 & 30 & 60 \\ 30 & 900 & 1800 \\ 60 & 1800 & 3600 \end{bmatrix}, \tag{4.35}$$

of rank one. The matrix has two zero eigenvalues and no ordinary inverse, and the posterior correlation of T_1 and T_2 is exactly one. This is not a failure of the inference; it is the consequence of exact physical rows, which leave the state one degree of freedom. The state-vector form needs no special treatment of it: the normal equations are solved at a finite tolerance, where $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ is invertible, and the limit is derived in the coordinates of equation (4.32) and read from the sequence of tolerances. The distinction between exact constraints and uncertain physical relations matters whenever a Hessian and a covariance are constructed over the full state, and the tolerance is the device that keeps both defined.

The objective along the free direction. On the line $\mathbf{z} = \hat{\mathbf{z}} + t\,\mathbf{v}$ through the MAP state, in the limit of exact rows, the physical rows keep their values at $\hat{\mathbf{z}}$ (zero in the limit) and only the reading and prior rows change, as $r_3(\hat{\mathbf{z}}) + 60t$ and $r_4(\hat{\mathbf{z}}) + 40t$, the slopes being the entries of $\mathbf{J}\mathbf{v} = (0,\ 0,\ 60,\ 40)^{\mathsf{T}}$. The objective on the line is the reading term $\frac{1}{2}(r_3(\hat{\mathbf{z}}) + 60t)^2$ plus the prior term $\frac{1}{2}(r_4(\hat{\mathbf{z}}) + 40t)^2$. At the minimum the slopes of the two terms cancel,

$\mathbf{v}^{\mathsf{T}}\mathbf{J}^{\mathsf{T}}\mathbf{r}(\hat{\mathbf{z}}) = 60\,r_3 + 40\,r_4 = 0$: the reading pulls with weight 60 and the prior with weight 40. The gradient vanishes at $\hat{\mathbf{z}}$, so the objective on the line is

$$\mathcal{L}(\hat{\mathbf{z}}+t\mathbf{v}) = \mathcal{L}(\hat{\mathbf{z}})+\tfrac{1}{2}\,t^2\,\|\mathbf{J}\mathbf{v}\|^2, \qquad \frac{\mathrm{d}^2}{\mathrm{d}t^2}\mathcal{L}(\hat{\mathbf{z}}+t\mathbf{v}) = \mathbf{v}^{\mathsf{T}}\mathbf{J}^{\mathsf{T}}\mathbf{J}\,\mathbf{v} = 60^2+40^2 = 5200\,(\mathrm{kW/K})^2, \tag{4.36}$$

a parabola whose curvature is the information along $\mathbf{v}$, equal to $1/\Sigma_{11}$ in the limit; the rise of $\frac{1}{2}$ is reached at $t = \pm 1/\sqrt{5200} = \pm\sqrt{\Sigma_{11}}$.

Remark 4.2. Substituting the exact rows $T_1 = 18 + 30\rho$ and $T_2 = 18 + 60\rho$ into the objective gives the same curve as a function of the resistance alone,

$$\mathcal{L}(\rho) = \frac{1}{2}\left[\frac{(\rho-\mu_0)^2}{\sigma_0^2} + \frac{(y_T - 18 - 30\rho)^2}{\sigma_T^2}\right], \tag{4.37}$$

whose second derivative is the 5200 of equation (4.36) and whose minimum is $\hat{\rho} = 880/5200$.

4.2.6 The Gaussian form of the posterior

In this example the prior is Gaussian, the reading error and the model tolerance are Gaussian, and every row is linear in the state. The negative log posterior is therefore exactly quadratic in $\mathbf{z}$, the curvature term $\sum_k r_k\,\partial^2 r_k/\partial\mathbf{z}\,\partial\mathbf{z}^{\mathsf{T}}$ is zero, the Gauss–Newton matrix $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ is the exact Hessian, and its inverse is the exact posterior covariance under the stated model. Nothing in equations (4.29)–(4.31) is an approximation; the only limit taken, $\sigma_p \to 0$, is derived in equation (4.35).

For a nonlinear physical model, for instance one written in the log-conductance g, the negative log posterior is not quadratic. The same construction is then made locally: the objective is expanded to second order about the MAP state,

$$\mathcal{L}(\mathbf{z}) \approx \mathcal{L}(\hat{\mathbf{z}}) + \tfrac{1}{2}(\mathbf{z}-\hat{\mathbf{z}})^{\mathsf{T}}\mathbf{H}(\mathbf{z}-\hat{\mathbf{z}}), \qquad \mathbf{H} = \left.\frac{\partial^2\mathcal{L}}{\partial\mathbf{z}\,\partial\mathbf{z}^{\mathsf{T}}}\right|_{\hat{\mathbf{z}}} = \mathbf{J}^{\mathsf{T}}\mathbf{J} + \sum_k r_k \left.\frac{\partial^2 r_k}{\partial\mathbf{z}\,\partial\mathbf{z}^{\mathsf{T}}}\right|_{\hat{\mathbf{z}}},$$

with no first-order term because the gradient vanishes at the minimum, and the posterior is replaced by the Gaussian that this expansion defines, the Laplace approximation:

$$\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}},\ \mathbf{H}^{-1}). \tag{4.38}$$

It holds when the posterior is close to Gaussian over its own width and when $\mathbf{H}$ is positive definite in the chosen free coordinates. The distinction matters throughout the book: the calculation of this chapter is exact, while a Laplace approximation in general is not, and its adequacy is one of the checks of step 5.

4.2.7 Rows that are not Gaussian

Every row of the hall is Gaussian, and its term in the objective is half its squared residual. Chapter 2's alarm log already has a row of another kind, a count of alarms in a number of windows. The construction above carries over to it unchanged once the objective is defined by the posterior rather than by a sum of squares.

The objective. The objective is the negative logarithm of the posterior up to a constant. Each row, a likelihood factor or a prior factor, contributes one term, its negative log-probability:

$$\mathcal{L}(\mathbf{z}) = \sum_k \mathcal{L}_k(\mathbf{z}), \qquad \mathcal{L}_k(\mathbf{z}) = -\ln p_k(\cdot \mid \mathbf{z}), \qquad p(\mathbf{z} \mid \mathbf{y}) \propto \exp[-\mathcal{L}(\mathbf{z})]. \tag{4.39}$$

A Gaussian row is the case $\mathcal{L}_k = \frac{1}{2}r_k^2$, and the objective (4.24) of the hall is this sum for its four rows.

What the fit uses. The MAP state, the Newton step and the Laplace covariance use the objective only through its gradient and its Hessian, and both are sums over the rows:

$$\nabla\mathcal{L} = \sum_k \nabla\mathcal{L}_k, \qquad \mathbf{H} = \sum_k \nabla^2\mathcal{L}_k, \qquad \mathbf{H}\,\Delta\mathbf{z} = -\nabla\mathcal{L}, \qquad \mathbf{\Sigma} = \mathbf{H}(\hat{\mathbf{z}})^{-1}. \tag{4.40}$$

The value of $\mathcal{L}$ enters only the evidence, equation (4.19). When the state has one entry ξ, the gradient is the slope $\mathcal{L}'(\xi)$, the Hessian the curvature $\mathcal{L}''(\xi)$, the step $\Delta\xi = -\mathcal{L}'/\mathcal{L}''$ and the variance $1/\mathcal{L}''(\hat{\xi})$; nothing else changes.

The Gaussian row. For $\mathcal{L}_k = \frac{1}{2}r_k^2$ with Jacobian line $\mathbf{j}_k = \partial r_k/\partial\mathbf{z}^\mathsf{T}$, the chain rule gives

$$\nabla\mathcal{L}_k = \mathbf{j}_k^\mathsf{T} r_k, \qquad \nabla^2\mathcal{L}_k = \mathbf{j}_k^\mathsf{T}\mathbf{j}_k + r_k\,\frac{\partial^2 r_k}{\partial\mathbf{z}\,\partial\mathbf{z}^\mathsf{T}}. \tag{4.41}$$

Summed over the rows these are $\mathbf{J}^\mathsf{T}\mathbf{r}$ and $\mathbf{J}^\mathsf{T}\mathbf{J}$ plus the curvature term of Section 4.2.6.

The count row. A record of n_k events (the count k of Chapter 2) whose expected number is $\mu_k(\mathbf{z})$, Poisson in windows, has the term $\mathcal{L}_k = \mu_k - n_k \ln\mu_k$. Its gradient and Hessian are

$$\nabla\mathcal{L}_k = (\mu_k - n_k)\,\nabla\ln\mu_k, \qquad \nabla^2\mathcal{L}_k = \mu_k\,\nabla\ln\mu_k\,\nabla\ln\mu_k^\mathsf{T} + (\mu_k - n_k)\,\nabla^2\ln\mu_k. \tag{4.42}$$

The first term of the Hessian is the product of one line with itself, so it is $\mathbf{j}_k^\mathsf{T}\mathbf{j}_k$ for the line and the residual

$$\mathbf{j}_k = \sqrt{\mu_k}\,\nabla\ln\mu_k^\mathsf{T}, \qquad r_k = \frac{\mu_k - n_k}{\sqrt{\mu_k}}, \tag{4.43}$$

which also give $\mathbf{j}_k^\mathsf{T} r_k = \nabla\mathcal{L}_k$. This pair, the *working residual* and the *working Jacobian line*, is how a count row joins the Jacobian beside Gaussian rows: one more line of $\mathbf{J}$, one more entry of $\mathbf{r}$, and the same products (4.40). The remaining term $(\mu_k - n_k)\,\nabla^2\ln\mu_k$ plays the part of $r_k\,\partial^2 r_k$ in a Gaussian row. It vanishes when $\ln\mu_k$ is linear in the state, which is the case when the state carries the log-rate, the canonical parameter of the Poisson distribution.

The working pair reproduces the row's gradient and the Gauss–Newton part of its Hessian, and nothing more. Its half square is not the row's term, $\frac{1}{2}r_k^2 \neq \mathcal{L}_k$, and $\mathbf{j}_k$ is not the derivative of r_k except where $r_k = 0$; the fit never uses either, because it uses only the products (4.40). For a Gaussian row that is linear in the state, the working pair is the row's own residual and Jacobian line, which is why the two forms are the same in the hall.

The information and its curvature term. With every row in this form, the Hessian is

$$\mathbf{H} = \mathbf{J}^\mathsf{T}\mathbf{J} + \mathbf{C}, \qquad \mathbf{\Lambda} = \mathbf{J}^\mathsf{T}\mathbf{J}, \tag{4.44}$$

with $\mathbf{C}$ the sum of the rows' remaining terms. The matrix $\mathbf{\Lambda}$ is the information; it is the average of $\mathbf{H}$ over the records the model could have produced, the *expected information*. The two are equal when $\mathbf{C} = \mathbf{0}$: for linear Gaussian rows, and for count rows carried on their canonical parameter. Where $\mathbf{C}$ is not zero, the chapters name it, keep it in the Newton step and report the covariance from $\mathbf{H}$. Once $\hat{\mathbf{z}}$ and $\mathbf{\Sigma}$ are in hand, the kind of row no longer matters: every answer of step 4 is read from them in the same way.

4.2.8 The reduced state

The full state carries every free quantity of the graph, and in the limit of exact rows the physical rows make some of its entries functions of the others; in the hall, two of the three entries follow from the third. The same posterior can then be written over the entries that remain free, the reduced state, as a projection of the rows onto the null space of the physical rows; the two-unit hall of Chapter 5 and every plant after it are fitted in this form, so the derivation below is written for any plant, with the hall as its instance.

The null space of the physical rows. Let the full state $\mathbf{z}$ have n entries and the graph m physical rows, with Jacobian $\mathbf{A} = \mathbf{J}_{\text{phys}}$, an $m \times n$ matrix. Split the full state into $n_{\text{f}} = n - m$ free entries $\mathbf{z}_{\text{f}}$ and m dependent entries $\mathbf{z}_{\text{d}}$, and $\mathbf{A}$ into the matching columns, $\mathbf{A} = [\,\mathbf{A}_{\text{free}}\ \ \mathbf{A}_{\text{dep}}\,]$, an $m \times n_{\text{f}}$ and an $m \times m$ block (the entries ordered free first). The split is a choice: any split serves for which $\mathbf{A}_{\text{dep}}$ is square and invertible, that is, for which the physical rows can be solved for the dependent entries at given free ones. A move $\Delta\mathbf{z}$ of the full state keeps the linearized balances when $\mathbf{A}\,\Delta\mathbf{z} = \mathbf{0}$, that is $\mathbf{A}_{\text{free}}\,\Delta\mathbf{z}_{\text{f}} + \mathbf{A}_{\text{dep}}\,\Delta\mathbf{z}_{\text{d}} = \mathbf{0}$, so

$$\Delta\mathbf{z}_{\text{d}} = \mathbf{N}_{\text{b}}\,\Delta\mathbf{z}_{\text{f}}, \qquad \mathbf{N}_{\text{b}} = -\mathbf{A}_{\text{dep}}^{-1}\mathbf{A}_{\text{free}}, \qquad \mathbf{N} = \begin{bmatrix} \mathbf{I}_{n_{\text{f}}} \\ \mathbf{N}_{\text{b}} \end{bmatrix}, \qquad \mathbf{AN} = \mathbf{0}. \tag{4.45}$$

Each physical row's tolerance divides both of its blocks and cancels in $\mathbf{N}_{\text{b}}$. The n_{f} columns of the $n \times n_{\text{f}}$ matrix $\mathbf{N}$ are independent, through its identity block, and $\mathbf{A}$ has rank m, so they span the null space of the physical rows, the tangent of the set on which the balances hold; each column moves one free entry by one unit and the dependent entries as the linearized balances require. In the hall $n = 3$ and $m = 2$; with $\mathbf{z}_{\text{f}} = \rho$ and $\mathbf{z}_{\text{d}} = (T_1, T_2)^{\mathsf{T}}$,

$$\mathbf{A}_{\text{free}} = \begin{bmatrix} -Q_1/\sigma_p \\ -Q_2/\sigma_p \end{bmatrix}, \qquad \mathbf{A}_{\text{dep}} = \begin{bmatrix} 1/\sigma_p & 0 \\ 0 & 1/\sigma_p \end{bmatrix}, \qquad \mathbf{N}_{\text{b}} = \begin{bmatrix} Q_1 \\ Q_2 \end{bmatrix}, \qquad \mathbf{N} = \begin{bmatrix} 1 \\ 30 \\ 60 \end{bmatrix},$$

the entries of $\mathbf{N}_{\text{b}}$ the loads in kW: a unit change of resistance raises each hall temperature by its load. The single column of $\mathbf{N}$ is the free direction $\mathbf{v}$ of equation (4.30). Taking T_1 as the free entry would serve as well, since the block of (ρ, T_2) is invertible too; the resistance is chosen because the prior acts on it and each hall temperature follows from it by its row of $\mathbf{N}_{\text{b}}$.

The rows of the reduced state. With exact balances the full state lies on the set where the physical rows vanish, and the reduced state $\mathbf{z}_{\text{f}}$ names a point of that set; the point itself is $\mathbf{z}_{\text{full}}(\mathbf{z}_{\text{f}}) = (\mathbf{z}_{\text{f}},\ \mathbf{z}_{\text{d}}(\mathbf{z}_{\text{f}}))$, with $\mathbf{z}_{\text{d}}(\mathbf{z}_{\text{f}})$ the solution of the physical rows at the given free entries, and its derivative with respect to $\mathbf{z}_{\text{f}}$ is $\mathbf{N}$. Write $\mathbf{r}_{\text{rows}}$ for the evidence rows, the readings and the priors, as functions of the full state. The rows of the reduced state are the evidence rows at the point of the set,

$$\mathbf{r}(\mathbf{z}_{\text{f}}) = \mathbf{r}_{\text{rows}}\big(\mathbf{z}_{\text{full}}(\mathbf{z}_{\text{f}})\big), \tag{4.46}$$

and the objective on the reduced state is the full objective (4.24) on the set, where its physical rows contribute nothing. In the hall, with $\mathbf{z}_{\text{full}}(\rho) = (\rho,\ T_c + Q_1\rho,\ T_c + Q_2\rho)^{\mathsf{T}}$, the rows are

$$\mathbf{r}(\rho) = \begin{bmatrix} (T_c + Q_1\rho - y_T)/\sigma_T \\ (\rho - \mu_0)/\sigma_0 \end{bmatrix} \begin{matrix} \text{reading} \\ \text{prior} \end{matrix},$$

and the objective is the curve (4.37), reached here as a projection and a limit rather than by a substitution made by hand.

The reduced Jacobian. A move $\Delta\mathbf{z}_{\text{f}}$ of the reduced state moves the point of the set by $\mathbf{N}\,\Delta\mathbf{z}_{\text{f}}$, so the chain rule gives the Jacobian of the reduced rows as the Jacobian of the evidence rows projected onto the null space,

$$\mathbf{J}_{\text{red}} = \mathbf{J}_{\text{rows}}\mathbf{N}, \tag{4.47}$$

where $\mathbf{J}_{\text{rows}} = \partial\mathbf{r}_{\text{rows}}/\partial\mathbf{z}$ is the Jacobian of the evidence rows (readings and priors) with respect to the full state, the lines of $\mathbf{J}$ that are not physical rows. A reading whose row touches only dependent entries sees the free entries through $\mathbf{N}_{\text{b}}$. In the hall $\mathbf{J}_{\text{rows}}$ is the reading and prior lines of equation (4.25),

$$\mathbf{J}_{\text{rows}} = \begin{bmatrix} 0 & 1/\sigma_T & 0 \\ 1/\sigma_0 & 0 & 0 \end{bmatrix}, \qquad \mathbf{J}_{\text{red}} = \begin{bmatrix} Q_1/\sigma_T \\ 1/\sigma_0 \end{bmatrix}.$$

The reading, whose row in the full state touches only T_1, sees the resistance on the reduced state with the slope Q_1/σ_T: its entry is nonzero, so the reading alone determines the one free entry (the hall is observable from its one sensor). The information of the reduced state is $\mathbf{J}_{\text{red}}^{\mathsf{T}}\mathbf{J}_{\text{red}} = Q_1^2/\sigma_T^2 + 1/\sigma_0^2$, the reading's term and the prior's, equal to $\|\mathbf{J}\mathbf{v}\|^2$ of equation (4.36).

Nonlinear physical rows. When the physical rows are nonlinear in the state, as they are when a conductance is carried as a log-derate or a recirculation fraction multiplies a temperature difference, the same objects are evaluated point by point. $\mathbf{z}_{\text{full}}(\mathbf{z}_{\text{f}})$ is the solution of the physical rows at the given free entries, found by Newton's method on those rows alone; $\mathbf{N}(\mathbf{z}_{\text{f}})$, from the blocks of $\mathbf{A}$ at that solution, and $\mathbf{J}_{\text{red}}(\mathbf{z}_{\text{f}}) = \mathbf{J}_{\text{rows}}\,\mathbf{N}(\mathbf{z}_{\text{f}})$ are evaluated there and change with the free entries. The Gauss–Newton iteration (4.29) runs on the reduced rows, $(\mathbf{J}_{\text{red}}^{\mathsf{T}}\mathbf{J}_{\text{red}})\,\Delta\mathbf{z}_{\text{f}} = -\mathbf{J}_{\text{red}}^{\mathsf{T}}\mathbf{r}(\mathbf{z}_{\text{f}})$, each iterate solving the physical rows again. At the MAP state, $\mathbf{\Sigma}_{\text{red}} = (\mathbf{J}_{\text{red}}^{\mathsf{T}}\mathbf{J}_{\text{red}})^{-1}$ is the Gauss–Newton form of the Laplace covariance (4.38); the term it drops holds the second derivatives of the reduced rows, which now include the curvature of the set, and each chapter checks it numerically. The hall's rows are linear, $\mathbf{N}$ is constant, one step is exact and $\mathbf{\Sigma}_{\text{red}}$ is the exact covariance.

The reduced covariance as the limit. The covariance of the reduced state is $\mathbf{\Sigma}_{\text{red}} = (\mathbf{J}_{\text{red}}^{\mathsf{T}}\mathbf{J}_{\text{red}})^{-1}$, and it is the free-entry block $\mathbf{\Sigma}_{\text{ff}}$ of the full $\mathbf{\Sigma}$ in the limit of exact rows. Write a move of the full state in the coordinates $\Delta\mathbf{z} = \mathbf{B}\,\Delta\mathbf{u}$ with $\mathbf{B} = [\,\mathbf{N}\;\;\mathbf{B}_{\text{d}}\,]$, where $\mathbf{B}_{\text{d}} = [\,\mathbf{0}\;\;\mathbf{I}_m\,]^{\mathsf{T}}$ is the $n \times m$ matrix that moves only the dependent entries. The square matrix $\mathbf{B}$ is invertible and its first block row is $[\,\mathbf{I}_{n_{\text{f}}}\;\;\mathbf{0}\,]$, so the first n_{f} coordinates are the move of the free entries; in the hall $\mathbf{B}$ is the matrix of equation (4.32). Since $\mathbf{A}\mathbf{N} = \mathbf{0}$ and $\mathbf{A}\mathbf{B}_{\text{d}} = \mathbf{A}_{\text{dep}}$, the Jacobian in these coordinates, physical rows first, and its information matrix are

$$\mathbf{J}\,[\,\mathbf{N}\;\;\mathbf{B}_{\text{d}}\,] = \begin{bmatrix} \mathbf{0} & \mathbf{A}_{\text{dep}} \\ \mathbf{J}_{\text{red}} & \mathbf{J}_{\text{rows}}\mathbf{B}_{\text{d}} \end{bmatrix}, \qquad \begin{bmatrix} \mathbf{J}_{\text{red}}^{\mathsf{T}}\mathbf{J}_{\text{red}} & \mathbf{E} \\ \mathbf{E}^{\mathsf{T}} & \mathbf{F} + \mathbf{A}_{\text{dep}}^{\mathsf{T}}\mathbf{A}_{\text{dep}} \end{bmatrix},$$

with $\mathbf{E} = \mathbf{J}_{\text{red}}^{\mathsf{T}}\mathbf{J}_{\text{rows}}\mathbf{B}_{\text{d}}$ and $\mathbf{F} = \mathbf{B}_{\text{d}}^{\mathsf{T}}\mathbf{J}_{\text{rows}}^{\mathsf{T}}\mathbf{J}_{\text{rows}}\mathbf{B}_{\text{d}}$. By the block-inversion formula of Section 3.5.3 the free-entry block of the covariance in these coordinates, which is $\mathbf{\Sigma}_{\text{ff}}$ because the first block row of $\mathbf{B}$ is $[\,\mathbf{I}_{n_{\text{f}}}\;\;\mathbf{0}\,]$, is the inverse of the Schur complement of the lower block, and that block grows without bound as the tolerances shrink: each line of $\mathbf{A}_{\text{dep}}$ carries its row's $1/\sigma_p$, so with a common tolerance $\mathbf{A}_{\text{dep}}^{\mathsf{T}}\mathbf{A}_{\text{dep}} = \sigma_p^{-2}\,\tilde{\mathbf{A}}_{\text{dep}}^{\mathsf{T}}\tilde{\mathbf{A}}_{\text{dep}}$, with $\tilde{\mathbf{A}}_{\text{dep}}$ the invertible block of the unnormalized rows:

$$\mathbf{\Sigma}_{\text{ff}} = \left[\mathbf{J}_{\text{red}}^{\mathsf{T}}\mathbf{J}_{\text{red}} - \mathbf{E}\big(\mathbf{F} + \mathbf{A}_{\text{dep}}^{\mathsf{T}}\mathbf{A}_{\text{dep}}\big)^{-1}\mathbf{E}^{\mathsf{T}}\right]^{-1} \;\longrightarrow\; \big(\mathbf{J}_{\text{red}}^{\mathsf{T}}\mathbf{J}_{\text{red}}\big)^{-1} = \mathbf{\Sigma}_{\text{red}} \qquad (\sigma_p \to 0). \quad (4.48)$$

The correction the bracket subtracts is of order σ_p^2; for nonlinear rows the equality holds with every block evaluated at the point of linearization, the MAP state. In the hall $\mathbf{\Sigma}_{\text{ff}}$ is the scalar Σ_{11}, $\mathbf{E} = (Q_1/\sigma_T^2,\;0)$, $\mathbf{F} = \text{diag}(1/\sigma_T^2,\;0)$ and $\mathbf{A}_{\text{dep}}^{\mathsf{T}}\mathbf{A}_{\text{dep}} = \sigma_p^{-2}\,\mathbf{I}_2$, and the bracket is the Schur complement of equation (4.33). The same coordinates give the covariance of the dependent entries in the limit, $\mathbf{N}_{\text{b}}\mathbf{\Sigma}_{\text{red}}\mathbf{N}_{\text{b}}^{\mathsf{T}}$, in the hall the lower block of the limit (4.35): the dependent entries are known through the free ones, each by its row of $\mathbf{N}_{\text{b}}$. The exact balance is this limit, never a substitution made by hand.

4.2.9 The fit of the hall state

The fit is run at the tolerance $\sigma_p = 0.01\,\text{K}$ of Table 4.3, and the limit of exact rows is read off equation (4.35).

The step from the handbook state. The iteration starts at the forward calculation with the handbook resistance, $\mathbf{z}^{(0)} = (0.160,\ 22.8,\ 27.6)^{\mathsf{T}}$, where both physical rows and the prior row vanish and the reading row is $(22.8 - 23.2)/0.5 = -0.8$, so that $\mathcal{L}(\mathbf{z}^{(0)}) = 0.32$ and $\mathbf{J}^{\mathsf{T}}\mathbf{r}(\mathbf{z}^{(0)}) = (0,\ -1.6,\ 0)^{\mathsf{T}}$: only the reading is out of balance, and it pulls on T_1 alone. The normal equations (4.29) are

$$\begin{bmatrix} 45\,001\,600 & -300\,000 & -600\,000 \\ -300\,000 & 10\,004 & 0 \\ -600\,000 & 0 & 10\,000 \end{bmatrix} \Delta\mathbf{z} = \begin{bmatrix} 0 \\ 1.6 \\ 0 \end{bmatrix}, \qquad \Delta\mathbf{z} = \begin{bmatrix} 0.0092296 \\ 0.27694 \\ 0.55378 \end{bmatrix}.$$

Divided by its first entry, the step is $\Delta\mathbf{z} = 0.0092296 \times (1,\ 30.005,\ 60.000)^{\mathsf{T}}$: it lies, to within the tolerance, along the free direction $\mathbf{v}$ of equation (4.30). Table 4.4 lists the iteration. Each line is one iterate: its three entries, the objective, and the largest entry of the gradient (4.28). After the first step the gradient is at round-off and the second step moves no entry by more than 10^{-13}.

Table 4.4: Gauss–Newton iterations for the hall state from the handbook state (0.160, 22.8, 27.6), with $\sigma_p = 0.01\,\mathrm{K}$.

Iterate k	ρ [K/kW]	T_1 [°C]	T_2 [°C]	$\mathcal{L}$ [–]	$\max_i \lvert(\mathbf{J}^{\mathsf{T}}\mathbf{r})_i\rvert$ [kW/K or K^{-1}]
0	0.160000	22.8000	27.6000	0.32000	1.6
1	0.169230	23.0769	28.1538	0.09845	$< 10^{-8}$
2	0.169230	23.0769	28.1538	0.09845	$< 10^{-8}$

The MAP state. The fit returns

$$\hat{\mathbf{z}} = \begin{bmatrix} \hat{\rho} \\ \hat{T}_1 \\ \hat{T}_2 \end{bmatrix} = \begin{bmatrix} 0.16923 \\ 23.0769 \\ 28.1538 \end{bmatrix}: \tag{4.49}$$

a resistance of 0.16923 K/kW, a test temperature of 23.077°C and a full-load temperature of 28.154°C, each read straight off its entry, with the effective conductance $\hat{G} = 1/\hat{\rho} = 5.909\,\mathrm{kW/K}$ evaluated at the estimated resistance.

4.2.10 The covariance and the resistance

The covariance. Inverting the matrix (4.27),

$$\boldsymbol{\Sigma} = \left(\mathbf{J}^{\mathsf{T}}\mathbf{J}\right)^{-1} = \begin{bmatrix} 1.9236 \times 10^{-4} & 5.7685 \times 10^{-3} & 1.1542 \times 10^{-2} \\ 5.7685 \times 10^{-3} & 0.17309 & 0.34611 \\ 1.1542 \times 10^{-2} & 0.34611 & 0.69260 \end{bmatrix} \quad (\sigma_p = 0.01\,\mathrm{K}). \tag{4.50}$$

Its eigenvalues are 0.866 along (nearly) $\mathbf{v}$, and 1.0×10^{-4} and 2.2×10^{-8} in the two directions the physics fixes, both shrinking as σ_p^2.

The entries read off. The resistance has

$$\hat{\rho} = \hat{z}_1 = 0.16923\,\mathrm{K/kW}, \qquad \sqrt{\Sigma_{11}} = \sqrt{1.9236 \times 10^{-4}} = 0.01387\,\mathrm{K/kW},$$

the test temperature $\sqrt{\Sigma_{22}} = 0.416\,\mathrm{K}$, and the full-load temperature $\hat{T}_2 = \hat{z}_3 = 28.154$°C with $\sqrt{\Sigma_{33}} = \sqrt{0.69260} = 0.8322\,\mathrm{K}$. The correlations $\Sigma_{ij}/\sqrt{\Sigma_{ii}\Sigma_{jj}}$ are 0.99971 between ρ and T_1, 0.99993 between ρ and T_2, and 0.99964 between T_1 and T_2.

The limit of exact rows. The limit covariance (4.35) gives $\Sigma_{11} = 1/5200$, so that

$$\sqrt{\Sigma_{11}} = 0.013868\,\mathrm{K/kW}, \qquad \hat{\rho} = 880/5200 = 0.169231\,\mathrm{K/kW}, \qquad \hat{T}_2 = 18{+}60\hat{\rho} = 28.1538^{\circ}\mathrm{C},$$

the resistance being the minimum of equation (4.37); the rest of the chapter carries these limit values.

The objective along the free direction. Figure 4.3 draws the objective (4.36) on the line $\mathbf{z} = \hat{\mathbf{z}} + t\,\mathbf{v}$ in the limit of exact rows, with t read on the horizontal axis as the resistance $\rho = \hat{z}_1 + t$ in K/kW and the dimensionless objective on the vertical axis: the solid curve is the objective, the dashed curve the prior term and the dotted curve the reading term; the dot marks the minimum and the double arrow joins the two points where the objective has risen by $\frac{1}{2}$. At the minimum $r_3(\hat{\mathbf{z}}) = -0.2462$ and $r_4(\hat{\mathbf{z}}) = 0.3692$, and the slopes cancel, $60 \times (-0.2462) + 40 \times 0.3692 = 0$.

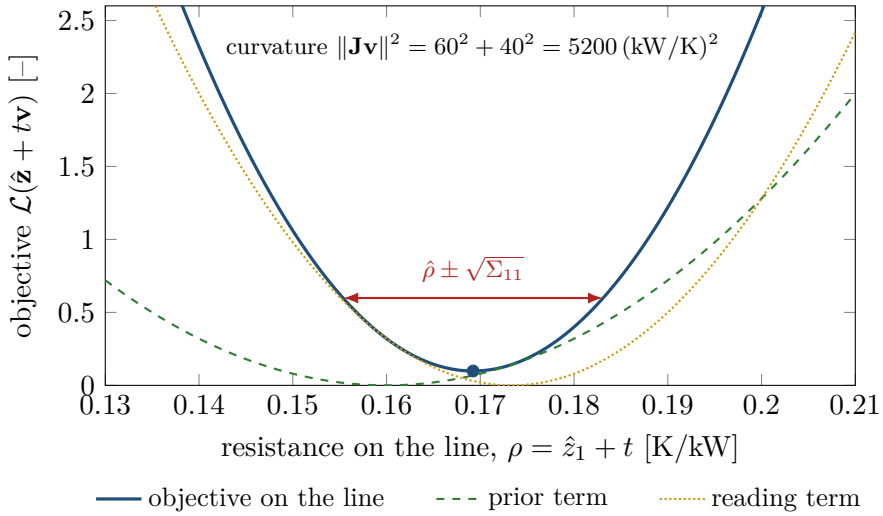


Figure 4.3: The objective on the line $\mathbf{z} = \hat{\mathbf{z}} + t\,\mathbf{v}$ along the free direction $\mathbf{v} = (1, 30, 60)^{\mathsf{T}}$, in the limit of exact physical rows, with its prior term $\frac{1}{2}(r_4(\hat{\mathbf{z}}) + 40t)^2$ and reading term $\frac{1}{2}(r_3(\hat{\mathbf{z}}) + 60t)^2$; the horizontal axis is the resistance $\hat{z}_1 + t$. Dot: the minimum at $\hat{\rho} = 0.16923\,\mathrm{K/kW}$, $\mathcal{L} = 0.0985$. Double arrow: the points $\hat{\rho} \pm \sqrt{\Sigma_{11}}$, $\sqrt{\Sigma_{11}} = 0.01387\,\mathrm{K/kW}$, where the objective has risen by $\frac{1}{2}$.

The posterior of the resistance. By the first marginal of equation (4.31),

$$\rho \mid \mathbf{y} \sim \mathcal{N}(\hat{z}_1,\ \Sigma_{11}) = \mathcal{N}(0.169231,\ 0.013868^2), \tag{4.51}$$

a standard deviation of $0.013868/0.169231 = 0.082$ of the estimate. Figure 4.4 compares it with the prior. The horizontal axis is the resistance and the vertical axis the probability density in kW/K; the dashed curve is the prior of equation (4.12) and the solid curve the posterior of equation (4.51). The reading has shifted the resistance upward, from 0.160 to 0.169, and narrowed it, from a standard deviation of 0.025 to 0.0139.

4.2.11 The reduced state of the hall

At the tolerance $\sigma_p = 0.01\,\mathrm{K}$ of equation (4.25), $\mathbf{A}_{\mathrm{free}} = (-3000,\ -6000)^{\mathsf{T}}$ and $\mathbf{A}_{\mathrm{dep}} = \mathrm{diag}(100,\ 100)$. From the handbook resistance 0.160 the reduced rows (4.46) are $(-0.8,\ 0)$, $\mathbf{J}_{\mathrm{red}}^{\mathsf{T}}\mathbf{r} = -48$, and one Gauss–Newton step $48/5200 = 0.0092308$ lands on $\hat{\rho} = 0.169231\,\mathrm{K/kW}$, which a second step does not move. The carried entries have $\mathbf{N}_{\mathrm{b}}\mathbf{\Sigma}_{\mathrm{red}}\mathbf{N}_{\mathrm{b}}^{\mathsf{T}}$ with diagonal 0.17308 and $0.69231\,\mathrm{K}^2$ and off-diagonal $0.34615\,\mathrm{K}^2$.

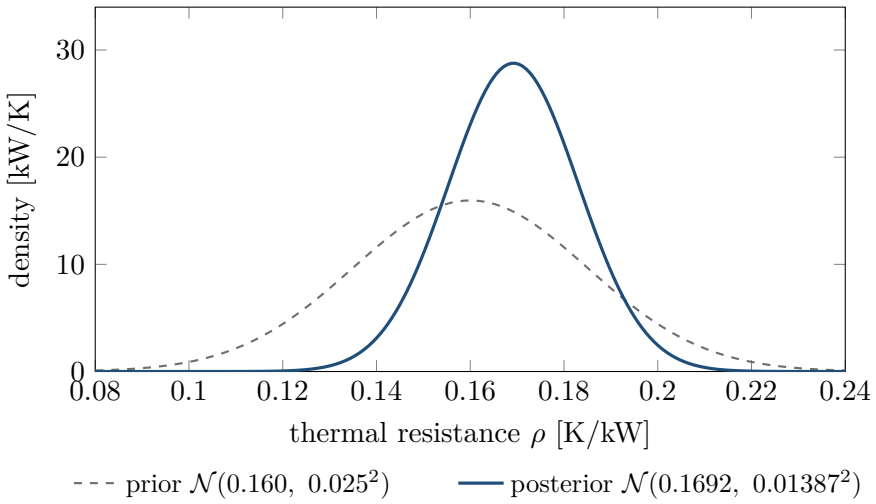


Figure 4.4: Prior and posterior densities of the thermal resistance, the first entry of the state: $\mathcal{N}(0.160, 0.025^2)$ and $\mathcal{N}(0.169231, 0.013868^2)$. The reading moves the resistance up and reduces its spread.

Table 4.5 sets the reduced state against the full one. Its first two rows give $\mathbf{N}$ and $\mathbf{J}_{\text{red}}$, which do not depend on the tolerance; the remaining rows give, for the full state at three tolerances and for the reduced state in the last column, the (ρ, ρ) entry of the covariance, its square root, the MAP resistance, and the relative difference of the full entry from $\mathbf{\Sigma}_{\text{red}}$. The difference shrinks a hundredfold per decade of σ_p, as the order σ_p^2 of the correction in equation (4.48) requires, and at $\sigma_p = 0.001\,\text{K}$ the two agree to every printed digit.

Table 4.5: The reduced state of the hall against the free-entry marginal of the full state. $\mathbf{N}$ is in [–, kW, kW]; the last column is the reduced state, exact in the limit of exact rows.

Quantity	$\sigma_p = 0.1\,\text{K}$	$0.01\,\text{K}$	$0.001\,\text{K}$	reduced state
null-space basis $\mathbf{N}$	$(1,\ 30,\ 60)^\mathsf{T}$ at every tolerance			
$\mathbf{J}_{\text{red}} = \mathbf{J}_{\text{rows}}\mathbf{N}$ [kW/K]	$(60,\ 40)^\mathsf{T}$, information $5200\,(\text{kW/K})^2$			
Σ_{11}, $\mathbf{\Sigma}_{\text{red}}$ $[(\text{K/kW})^2]$	1.97568×10^{-4}	1.92361×10^{-4}	1.92308×10^{-4}	1.92308×10^{-4}
standard deviation [K/kW]	0.014056	0.013869	0.013868	0.013868
$\hat{\rho}$ [K/kW]	0.169119	0.169230	0.169231	0.169231
relative difference from $\mathbf{\Sigma}_{\text{red}}$ [–]	2.7×10^{-2}	2.8×10^{-4}	2.8×10^{-6}	–

4.3 The answers (step 4)

Each question of Table 4.2 is read from the posterior of Section 4.2.

4.3.1 Uncertainty propagation and the delta method

The rows of the hall are linear, so the entries of $\mathbf{\Sigma}$ are exact variances. In most plant models an output is a nonlinear function of the state, $f(\mathbf{z})$, and the posterior is Gaussian at best approximately, $\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$.

Linearization of the output. The first-order Taylor expansion of f about $\hat{\mathbf{z}}$ is

$$f(\mathbf{z}) \approx f(\hat{\mathbf{z}}) + \nabla f(\hat{\mathbf{z}})^\mathsf{T}(\mathbf{z} - \hat{\mathbf{z}}).$$

Writing $\mathbf{g} = \nabla f(\hat{\mathbf{z}})$ for the column gradient,

$$f(\mathbf{z}) \approx f(\hat{\mathbf{z}}) + \mathbf{g}^{\mathsf{T}}(\mathbf{z} - \hat{\mathbf{z}}). \tag{4.52}$$

Taking the posterior expectation, with $\mathbb{E}[\mathbf{z} - \hat{\mathbf{z}} \mid \mathbf{y}] = \mathbf{0}$ because $\hat{\mathbf{z}}$ is the mean of the Gaussian,

$$\mu_f = \mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}] \approx f(\hat{\mathbf{z}}).$$

Taking the variance, the constant $f(\hat{\mathbf{z}})$ drops out, and the variance of the linear combination $\mathbf{g}^{\mathsf{T}}(\mathbf{z} - \hat{\mathbf{z}}) = \sum_i g_i(z_i - \hat{z}_i)$ is the double sum of its covariances,

$$\operatorname{Var}(f(\mathbf{z}) \mid \mathbf{y}) \approx \operatorname{Var}\Big(\sum_i g_i z_i \;\Big|\; \mathbf{y}\Big) = \sum_i \sum_j g_i\, g_j\, \operatorname{Cov}(z_i, z_j \mid \mathbf{y}),$$

which in matrix form is

$$\operatorname{Var}(f(\mathbf{z}) \mid \mathbf{y}) \approx \mathbf{g}^{\mathsf{T}} \boldsymbol{\Sigma}\, \mathbf{g}. \tag{4.53}$$

This is the first-order delta method, one of the most reused formulas in probabilistic plant modeling. Its only approximation is the linearization (4.52); when f is linear and the posterior Gaussian, it is exact.

Application to the thermal system. The full-load temperature is the third entry of the state, so its gradient is the third unit vector and the quadratic form picks one diagonal entry,

$$\mathbf{g} = \frac{\partial T_2}{\partial \mathbf{z}} = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}, \qquad \mathbf{g}^{\mathsf{T}} \boldsymbol{\Sigma}\, \mathbf{g} = \Sigma_{33}.$$

The same temperature can also be predicted as a derived quantity of the resistance alone, $f(\mathbf{z}) = T_c + Q_2\rho = 18 + 60\rho$, with

$$\mathbf{g} = \frac{\partial f}{\partial \mathbf{z}} = \begin{bmatrix} 60 \\ 0 \\ 0 \end{bmatrix}, \qquad \mathbf{g}^{\mathsf{T}} \boldsymbol{\Sigma}\, \mathbf{g} = 60^2\, \Sigma_{11}.$$

In the limit the two are equal, since the limit covariance (4.35) has $\Sigma_{33} = 3600/5200 = 60^2\, \Sigma_{11}$; at a finite tolerance they differ by σ_p^2, equation (4.34), because the entry carries the discrepancy of the full-load balance and the prediction from ρ does not. The two calculations, reading the variance off the covariance and propagating the resistance's uncertainty through a gradient, agree in the limit because they are the same transformation of the same uncertainty: an output that is an entry of the state needs no propagation, and an output that is not needs its gradient.

4.3.2 The gradient and the variance budget

Physical meaning of the gradient. The first entry of the gradient $\mathbf{g} = (60,\ 0,\ 0)^{\mathsf{T}}$ of the prediction $18 + 60\rho$, 60 kW, means that a change of 0.01 K/kW in the resistance changes the full-load temperature by $\Delta T_2 = 60 \times 0.01 = 0.6$ K. The gradient is the link between physical uncertainty and operating consequence. A state entry can have a large posterior variance and still contribute little to an output that is insensitive to it, with $g_i \approx 0$; conversely, a modest uncertainty in a highly influential entry can dominate the uncertainty of a prediction.

Several uncertain entries. Suppose the supply-air temperature is uncertain as well, the controller holding it only near its set point, so that the full-load temperature $T = T_c + Q\rho$ depends on the sub-vector $\boldsymbol{\theta} = (T_c, \rho)^\mathsf{T}$ of the state, with posterior covariance $\boldsymbol{\Sigma}_\theta$. The gradient is $\mathbf{g} = (\partial T/\partial T_c,\ \partial T/\partial \rho)^\mathsf{T} = (1,\ Q)^\mathsf{T}$, and equation (4.53) gives

$$\mathrm{Var}(T) = \mathbf{g}^\mathsf{T} \boldsymbol{\Sigma}_\theta\, \mathbf{g} = \begin{bmatrix} 1 & Q \end{bmatrix} \begin{bmatrix} \mathrm{Var}(T_c) & \mathrm{Cov}(T_c, \rho) \\ \mathrm{Cov}(T_c, \rho) & \mathrm{Var}(\rho) \end{bmatrix} \begin{bmatrix} 1 \\ Q \end{bmatrix}.$$

Multiplying out,

$$\mathrm{Var}(T) = \mathrm{Var}(T_c) + Q^2\ \mathrm{Var}(\rho) + 2Q\ \mathrm{Cov}(T_c, \rho). \tag{4.54}$$

This is the variance budget of the output: one term per entry, each its sensitivity squared times its variance, plus a cross term from their covariance. The first term is the spread the supply air alone puts on the temperature; the second, the spread the resistance alone puts on it, $Q^2\ \mathrm{Var}(\rho)$; the third corrects for the two moving together in the posterior. If the readings cannot separate the two entries, the posterior leaves them correlated, and the cross term either adds to the diagonal terms or cancels part of them. With a negative correlation and same-signed sensitivities it is negative, and the output can be well known although each entry is poorly known: the readings then fix the combination $T_c + Q\rho$ that the output depends on, not the entries separately. Adding the two individual variances without the cross term is wrong in either direction, which is why the joint posterior, not its marginals, must be kept; Exercise 4.2 works a case in which ignoring the cross term overstates the spread by more than half.

4.3.3 The density of a transformed quantity

The delta method gives a derived quantity a mean and a variance. Some questions need more: the whole density of the quantity, or a tail probability far from the center, where the shape matters and a Gaussian with the right variance can still be wrong. The tool for these is the *change of variables*, and it is used throughout the book whenever a quantity is computed in one coordinate and reported in another.

Probability is conserved; density is not. Let x have the density $p_x(x)$, and let $u = h(x)$ be a monotone function with inverse $x = h^{-1}(u)$. A short interval $[x, x + \mathrm{d}x]$ is mapped onto a short interval $[u, u + \mathrm{d}u]$, and the two carry the same probability, because they hold the same outcomes:

$$p_u(u)\,|\mathrm{d}u| = p_x(x)\,|\mathrm{d}x|.$$

Dividing by $|\mathrm{d}u|$ gives the density of the transformed quantity,

$$p_u(u) = p_x(h^{-1}(u))\ \left|\frac{\mathrm{d}x}{\mathrm{d}u}\right|. \tag{4.55}$$

The first factor evaluates the old density at the point that maps to u. The second, the *change-of-variable factor*, records how much the map stretches or squeezes the axis there: where a small change of u corresponds to a large change of x, the probability of the long x-interval is packed into a short u-interval, and the density of u must be larger to hold it. Without the factor, $p_x(h^{-1}(u))$ is not a density of u at all: it does not in general integrate to one. A monotone map also carries quantiles over unchanged: if h is increasing, $\Pr(u \le h(x^*)) = \Pr(x \le x^*)$, so the median of u is h of the median of x, while the mode, which depends on the factor, moves.

The conductance. For the conductance $G = 1/\rho$ of the hall, the inverse map is $\rho = 1/G$ and $|\mathrm{d}\rho/\mathrm{d}G| = 1/G^2$, so equation (4.55) gives

$$p_G(G) = \frac{1}{G^2}\ p(\rho = 1/G \mid \mathbf{y}).$$

It is skewed to the right of $\hat{G}$ because the factor $1/G^2$ is larger at small G, which squeezes the left of the curve and stretches its right.

Two routes to a derived quantity. The delta method and the change of variables answer the same question with different approximations. For a scalar entry with posterior $x \mid \mathbf{y} \approx \mathcal{N}(\hat{x}, \sigma^2)$ and a monotone output $u = h(x)$:

- the *delta method* linearizes h at $\hat{x}$, with gradient $g = \mathrm{d}u/\mathrm{d}x$ evaluated there, and reports $u \mid \mathbf{y} \approx \mathcal{N}(h(\hat{x}),\ g^2\sigma^2)$, a Gaussian in u;
- the *change of variables* keeps h exact and reports $p_u(u) = p_x(h^{-1}(u))\,|\mathrm{d}x/\mathrm{d}u|$, with the Gaussian of x as its only approximation.

The two factors are reciprocals: the delta method's gradient $\mathrm{d}u/\mathrm{d}x$ is the inverse of the change-of-variable factor $\mathrm{d}x/\mathrm{d}u$, but the delta method evaluates it once, at $\hat{x}$, while the change of variables evaluates it at every u. The delta method is the change of variables with the map replaced by its tangent, which is why the two agree on the spread near the center and differ in the tails. When h is linear the factor is constant and the two coincide. When h bends across the posterior, the delta method puts a symmetric Gaussian on a quantity whose density is skewed, and can put probability where the quantity cannot go: below zero for a conductance or a rate, outside $(0, 1)$ for a probability such as a probability of detection.

> **Principle 4.1.** *The gradient carries a spread; the change-of-variable factor carries a density. Report $\sqrt{\mathbf{g}^\top \boldsymbol{\Sigma} \mathbf{g}}$ when a standard deviation is the answer, and transform the density when the answer is a shape, a tail probability or a bound near the edge of what the quantity can be.*

4.3.4 Credible intervals

The control room often needs an interval for an uncertain quantity as well as its mean. For the Gaussian marginal $T_2 \mid \mathbf{y} \sim \mathcal{N}(\hat{z}_3, \Sigma_{33})$ of equation (4.31) the central 95 % credible interval is the interval that leaves 2.5 % of the posterior in each tail. Since $\Phi(1.960) = 0.975$,

$$\left[\hat{z}_3 - 1.960\,\sqrt{\Sigma_{33}},\ \hat{z}_3 + 1.960\,\sqrt{\Sigma_{33}}\right]. \tag{4.56}$$

The same form serves any entry, with $\hat{z}_i$ and $\sqrt{\Sigma_{ii}}$, and a derived quantity, with $f(\hat{\mathbf{z}})$ and $\sqrt{\mathbf{g}^\top \boldsymbol{\Sigma} \mathbf{g}}$ from the delta method (4.53).

The interval and a threshold probability carry different information. The interval describes the range that holds most of the posterior probability; a threshold probability answers directly whether the plant meets a stated limit. For an operating decision the threshold probability is usually the more actionable of the two. Neither is meaningful apart from the model assumptions, the quality of the measurement and the sources of uncertainty that the model includes; Section 4.6.6 lists those it leaves out.

4.3.5 The temperature-threshold probability

The question of equation (4.10) needs one marginal of the posterior. The marginal of the full-load temperature is read off the third entry of the state, equation (4.31): $T_2 \mid \mathbf{y} \sim \mathcal{N}(\hat{z}_3,\ \Sigma_{33})$.

The standardized temperature. For a Gaussian variable the standardized variable $(T_2 - \hat{z}_3)/\sqrt{\Sigma_{33}}$ has the standard normal distribution, and subtracting a constant and dividing by a positive constant preserve the inequality, so

$$\Pr(T_2 < T_{\max} \mid \mathbf{y}) = \Pr\Big(\frac{T_2 - \hat{z}_3}{\sqrt{\Sigma_{33}}} < \frac{T_{\max} - \hat{z}_3}{\sqrt{\Sigma_{33}}} \;\Big|\; \mathbf{y}\Big) = \Phi\Big(\frac{T_{\max} - \hat{z}_3}{\sqrt{\Sigma_{33}}}\Big). \tag{4.57}$$

The probability of exceeding the limit is its complement, $1 - \Pr(T_2 < T_{\max} \mid \mathbf{y})$.

The prior marginal. Before the test reading the state had the prior formed by the physical and prior rows alone. Its Jacobian $\mathbf{J}_0$ is $\mathbf{J}$ of equation (4.25) without the reading's line, its mean is the handbook state $\mathbf{z}^{(0)} = (0.160,\ 22.8,\ 27.6)^{\mathsf{T}}$, and its covariance $\mathbf{\Sigma}_0 = (\mathbf{J}_0^{\mathsf{T}}\mathbf{J}_0)^{-1}$ tends, by the argument of equation (4.35) with $\mathbf{J}_0\mathbf{v} = (0,\ 0,\ 40)^{\mathsf{T}}$ in place of $\mathbf{J}\mathbf{v}$, to $\mathbf{v}\mathbf{v}^{\mathsf{T}}/40^2 = \sigma_0^2\,\mathbf{v}\mathbf{v}^{\mathsf{T}}$. Its third diagonal entry is $(\mathbf{\Sigma}_0)_{33} = 60^2 \times 0.025^2 = 2.25\,\mathrm{K}^2$, so the full-load temperature had the prior marginal $\mathcal{N}(27.6,\ 1.5^2)$, the third entry of $\mathbf{z}^{(0)}$ with $\sqrt{(\mathbf{\Sigma}_0)_{33}} = 1.5\,\mathrm{K}$.

The question through the resistance. In the limit of exact rows the posterior lies on the line $\mathbf{z} = \hat{\mathbf{z}} + t\,\mathbf{v}$ through the MAP state along the free direction, $\mathbf{\Sigma} = \mathbf{v}\mathbf{v}^{\mathsf{T}}/5200$ by equation (4.35). On that line the third entry is $\hat{z}_3 + 60t$ and the first $\hat{z}_1 + t$, so the event $\{T_2 < T_{\max}\}$ is the event $\{t < t^*\}$ with $t^* = (T_{\max} - \hat{z}_3)/60$, which is the event that the resistance stays below

$$\rho^* = \hat{z}_1 + t^* = \hat{z}_1 + \frac{T_{\max} - \hat{z}_3}{60}. \tag{4.58}$$

With the marginal (4.31) of the first entry,

$$\Pr(\rho < \rho^* \mid \mathbf{y}) = \Phi\Big(\frac{\rho^* - \hat{z}_1}{\sqrt{\Sigma_{11}}}\Big) = \Phi\Big(\frac{t^*}{\sqrt{\Sigma_{11}}}\Big),$$

which is the probability (4.57): both standardized distances are t^* divided by the spread of t, since the third entry moves 60 times as far as the first along $\mathbf{v}$ and $\sqrt{\Sigma_{33}} = 60\sqrt{\Sigma_{11}}$ in the limit. Reading the third entry and propagating the first one along the free direction are the same calculation; the full-load question is a question about the resistance.

The dependence on the value read. The information matrix (4.27), and with it $\mathbf{\Sigma}$, does not contain the value observed, so a different reading leaves every standard deviation where it is and moves only the mean, along one vector. The reading enters the normal equations only through the part $-y_T/\sigma_T$ of the third entry of the residual vector that depends on y_T; that row's line of the Jacobian is $\mathbf{h}^{\mathsf{T}}/\sigma_T$, so differentiating $\hat{\mathbf{z}} = \mathbf{z} - \mathbf{\Sigma}\mathbf{J}^{\mathsf{T}}\mathbf{r}(\mathbf{z})$ with respect to y_T gives the gain vector

$$\frac{\partial \hat{\mathbf{z}}}{\partial y_T} = \frac{\mathbf{\Sigma}\mathbf{h}}{\sigma_T^2}, \tag{4.59}$$

the second column of $\mathbf{\Sigma}$ divided by σ_T^2. The threshold probability (4.57) is computed with a spread fixed before the reading and a mean set by its value, so the reading can move it in either direction.

4.3.6 The posterior state and the full-load temperature

General form. The three estimation questions, the resistance of the air path and the hall temperature at the test and at full load, share one form: the posterior of a quantity $f(\mathbf{z})$, equation (2.9), summarized by its moments (2.10) and its central interval (2.11).

Specialized form. Equations (2.10) and (2.11) with $f(\mathbf{z}) = z_i$, the entry asked about, and $p(\mathbf{z} \mid \mathbf{y})$ the Gaussian posterior of step 3: the mean is $\hat{z}_i$, the variance Σ_{ii}, by the marginals (4.31), and the interval is equation (4.56). The rows are linear and Gaussian, so these are the exact moments, with no approximation to set beside them.

Building blocks. The MAP state and the limit covariance (4.35), with its limit values of Section 4.2.10; the prior marginal $\mathcal{N}(27.6,\ 1.5^2)$ of Section 4.3.5, for the prior interval.

The posterior state. Table 4.6 collects it: each row is one entry of $\hat{\mathbf{z}}$, with its posterior mean and its standard deviation $\sqrt{\Sigma_{ii}}$, in the limit $\sigma_p \to 0$; the full-load spread is twice the test spread, the ratio 60/30 of the entries of $\mathbf{v}$.

Table 4.6: The posterior state of the equipment hall after the test reading, in the limit of exact physical rows.

State entry	Posterior mean	Standard deviation
ρ [K/kW]	0.16923	0.01387
T_1 [°C; sd in K]	23.077	0.416
T_2 [°C; sd in K]	28.154	0.832

The diagonal of the limit covariance. The limit covariance (4.35) gives $\Sigma_{11} = 1/5200 = 1.9231 \times 10^{-4}\,(\mathrm{K/kW})^2$, $\Sigma_{22} = 900/5200 = 0.17308\,\mathrm{K}^2$ ($\sqrt{\Sigma_{22}} = 0.4160\,\mathrm{K}$) and

$$\Sigma_{33} = \mathrm{Var}(T_2 \mid \mathbf{y}) = \frac{3600}{5200} = 0.69231\,\mathrm{K}^2, \qquad \sqrt{\Sigma_{33}} = 0.8321\,\mathrm{K}. \tag{4.60}$$

The credible interval. By equation (4.56), $28.1538 \pm 1.960 \times 0.8321 = 28.1538 \pm 1.6308$, so that

$$T_2 \in [26.52,\ 29.78]^\circ\mathrm{C}$$

with 95 % posterior probability under the stated model; the prior interval was $[24.66,\ 30.54]^\circ$C, 1.8 times as wide.

The conductance beside the resistance. The conductance is a derived quantity, $G(\mathbf{z}) = 1/z_1$, so the delta method (4.53) reads it from the same posterior, with

$$\mathbf{g} = \frac{\partial G}{\partial \mathbf{z}} = \begin{bmatrix} -1/\hat{z}_1^2 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} -34.92 \\ 0 \\ 0 \end{bmatrix} (\mathrm{kW/K})^2, \qquad \mathrm{Var}(G \mid \mathbf{y}) \approx \mathbf{g}^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g} = \frac{\Sigma_{11}}{\hat{z}_1^4},$$

which gives $\hat{G} = 1/\hat{\rho} = 5.909\,\mathrm{kW/K}$ with $\mathrm{sd}(G \mid \mathbf{y}) \approx 0.484\,\mathrm{kW/K}$ and the 5–95 % band $[5.113,\ 6.706]\,\mathrm{kW/K}$. The exact band, by the change of variables (4.55), is $[5.207,\ 6.830]\,\mathrm{kW/K}$, higher by 0.095 and $0.124\,\mathrm{kW/K}$; the median is $5.909\,\mathrm{kW/K}$ on both routes.

4.3.7 The threshold probability at full load

General form. A requirement that a quantity stay below a limit is the posterior probability of the event, equation (2.12), compared with an assurance level κ.

Specialized form. Equation (2.12) with $f(\mathbf{z}) = -T_2 = -z_3$, $f^* = -T_{\max}$ and $p(\mathbf{z} \mid \mathbf{y})$ the Gaussian posterior of step 3 becomes the normal tail (4.57) of the third entry. The commissioning test states no assurance level κ; the probability is reported without an accept verdict. The posterior is exactly Gaussian, so the tail is the exact probability under the model.

Building blocks. The third entry $\hat{z}_3$ and the variance Σ_{33} of equation (4.60), and the standard normal distribution function Φ.

The probability. With $\hat{z}_3 = 28.1538\,^\circ\mathrm{C}$ and $\sqrt{\Sigma_{33}} = 0.8321\,\mathrm{K}$, the standardized distance in equation (4.57) is $(28 - 28.1538)/0.8321 = -0.1849$, and

$$\Pr(T_2 < 28\,^\circ\mathrm{C} \mid \mathbf{y}) = \Phi(-0.1849) = 0.4267. \tag{4.61}$$

The probability of exceeding the limit is $1 - 0.4267 = 0.5733$. The fit at $\sigma_p = 0.01\,\mathrm{K}$ gives 0.4267 as well, with $\sqrt{\Sigma_{33}} = 0.83223\,\mathrm{K}$ against the limit's $0.83205\,\mathrm{K}$.

Against the prior. Under the prior marginal $\mathcal{N}(27.6,\ 1.5^2)$ of Section 4.3.5, before the test reading, the probability that the full-load temperature stays below the limit was

$$\Pr(T_2 < 28) = \Phi\left(\frac{28 - 27.6}{1.5}\right) = \Phi(0.2667) = 0.6051.$$

Figure 4.5 draws the answer. The horizontal axis is the full-load temperature and the vertical axis its density in K^{-1}. The solid curve is the posterior marginal, the dashed curve the prior marginal, and the vertical line the limit $T_{\max}$. The shaded area under the posterior to the left of the line is the probability (4.61); the unshaded area under it to the right is the probability of exceeding the limit. The posterior is narrower than the prior and centered to the right of it, across the limit.

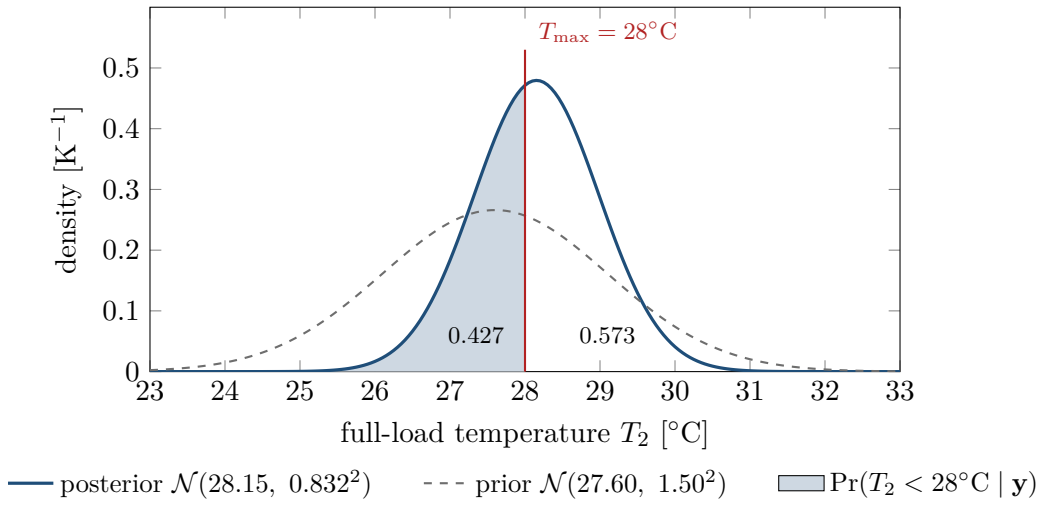


Figure 4.5: The threshold probability of equation (4.57). Posterior marginal of the full-load temperature, $\mathcal{N}(28.1538, 0.8321^2)$, with the area below $T_{\max} = 28\,^\circ\mathrm{C}$ shaded ($\Phi(-0.1849) = 0.4267$), and the prior marginal $\mathcal{N}(27.6, 1.5^2)$, under which the same probability is $\Phi(0.2667) = 0.6051$.

4.4 The checks (step 5)

Obtaining a posterior is not the same as establishing that it is credible.

4.4.1 Algebraic and physical verification

Algebraic verification. At the fitted state of equation (4.49), with $\sigma_p = 0.01\,\mathrm{K}$, the residual vector (4.23) is $\mathbf{r}(\hat{\mathbf{z}}) = (0.0049225,\ 0,\ -0.24612,\ 0.36919)^{\mathsf{T}}$, and the gradient (4.28), column by column of $\mathbf{J}$,

$$\mathbf{J}^{\mathsf{T}}\mathbf{r}(\hat{\mathbf{z}}) = \begin{bmatrix} -3000 \times 0.0049225 + 40 \times 0.36919 \\ 100 \times 0.0049225 + 2 \times (-0.24612) \\ 100 \times 0 \end{bmatrix} = \begin{bmatrix} -14.767 + 14.767 \\ 0.4922 - 0.4922 \\ 0 \end{bmatrix} = \mathbf{0},$$

to numerical precision: the state satisfies the first-order optimality condition. Each line reads as a balance of rows on one entry: the prior's pull on ρ is held by the first physical row, which in turn is held at T_1 by the reading row; nothing pulls on T_2 but its own physical row. The information matrix (4.27) has positive eigenvalues, so the objective is strictly convex and the stationary point is its unique minimum.

Physical consistency. The physical rows at the fitted state, in kelvin, are $\sigma_p r_1 = 4.9 \times 10^{-5}\,\mathrm{K}$ and $\sigma_p r_2 = 0$: both halls balance to well within the tolerance, and the first imbalance shrinks with σ_p. The difference between the observed and the fitted test temperature is $23.2 - 23.0769 = 0.1231\,\mathrm{K}$, a quarter of the assumed sensor standard deviation of $0.5\,\mathrm{K}$. The model is consistent with this reading. Agreement with one reading does not validate the physical model: with one reading and one free direction, the fit has nearly enough freedom to reproduce any value.

4.4.2 Prior predictive check

Before the reading was used, the model and the prior predicted a distribution for it. By equation (4.11), $Y = \mathbf{h}^{\mathsf{T}}\mathbf{z} + \epsilon$, with the state under its prior $\mathcal{N}(\mathbf{z}^{(0)}, \boldsymbol{\Sigma}_0)$ of Section 4.3.5 and ϵ independent of it, so Y is Gaussian with

$$\mathbb{E}[Y] = \mathbf{h}^{\mathsf{T}}\mathbf{z}^{(0)} = 22.8^{\circ}\mathrm{C}, \qquad \mathrm{Var}(Y) = \mathbf{h}^{\mathsf{T}}\boldsymbol{\Sigma}_0\mathbf{h} + \sigma_T^2 = (\boldsymbol{\Sigma}_0)_{22} + \sigma_T^2 = 0.5625 + 0.25 = 0.8125\,\mathrm{K}^2,$$

the variances adding because the two terms are independent, and $(\boldsymbol{\Sigma}_0)_{22} = 30^2 \times 0.025^2$ in the limit of exact rows. Thus

$$Y \sim \mathcal{N}(22.8,\ 0.9014^2). \tag{4.62}$$

The observed reading lies $(23.2 - 22.8)/0.9014 = 0.444$ standard deviations from the predicted mean, an unremarkable distance: the reading is compatible with the prior information and the model. This check is distinct from computing the posterior. It asks whether the evidence was plausible under the assumptions before those assumptions were used to interpret it; a reading several predictive standard deviations away would have signaled a wrong prior, a wrong model or a faulty sensor before any posterior was formed.

4.4.3 The evidence by two routes

The evidence (4.18) of the model is the prior predictive density (4.62) evaluated at the reading,

$$p(y_T \mid M) = \mathcal{N}(y_T;\ \mathbf{h}^{\mathsf{T}}\mathbf{z}^{(0)},\ \mathbf{h}^{\mathsf{T}}\boldsymbol{\Sigma}_0\mathbf{h} + \sigma_T^2),$$

with $\mathcal{N}(y; m, v)$ the Gaussian density of Section 4.2.1. With $m = 22.8^{\circ}\mathrm{C}$ and $v = 0.8125\,\mathrm{K}^2$,

$$p(y_T = 23.2 \mid M) = \frac{1}{\sqrt{2\pi \times 0.8125}} \exp\left[-\frac{(23.2 - 22.8)^2}{2 \times 0.8125}\right] = 0.44259 \times 0.90623 = 0.4011\,\mathrm{K}^{-1}. \tag{4.63}$$

The Laplace form (4.19) gives the same number from the fit. At $\sigma_p = 0.01\,\mathrm{K}$, $\det(\boldsymbol{\Sigma})^{1/2} = 1.3867 \times 10^{-6}$ and $\mathcal{L}(\hat{\mathbf{z}}) = 0.09845$ (Table 4.4), so

$$p(y_T \mid M) = 0.39894 \times \frac{\det(\boldsymbol{\Sigma})^{1/2}}{\sigma_p^2} \times \frac{\mathrm{e}^{-\mathcal{L}(\hat{\mathbf{z}})}}{\sigma_T\,\sigma_0} = 0.39894 \times 0.013867 \times 0.90624 \times 80 = 0.4011\,\mathrm{K}^{-1},$$

with $\det(\boldsymbol{\Sigma})^{1/2}/\sigma_p^2 = 0.013867$ against its limit $1/\sqrt{5200} = 0.013868$, the width of the posterior along the free direction.

Common mistake

Reading an evidence value as a probability that the model is correct. The density $0.4011\,\mathrm{K}^{-1}$ is not a 40 % confidence in the thermal model: it is a density, in inverse units of the reading, whose value changes by a factor of 1000 if the reading is expressed in millikelvin. A posterior probability of a model requires the evidence of the competing models and their prior probabilities. Likewise, the posterior probability 0.4267 is a confidence in the claim about T_2 under the model, while the evidence ratio is a confidence in the model the claim was computed under; the two answer different questions.

4.4.4 Posterior predictive validation with a new observation

Suppose a further, independent reading is to be taken at the same 30 kW condition. Before its value is seen, its distribution given the evidence already in hand is the posterior predictive. The new reading is $Y_{\mathrm{new}} = \mathbf{h}^{\mathsf{T}}\mathbf{z} + \epsilon_{\mathrm{new}}$, with $\mathbf{z} \mid \mathbf{y} \sim \mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ and $\epsilon_{\mathrm{new}} \sim \mathcal{N}(0, \sigma_T^2)$ independent of it, so

$$Y_{\mathrm{new}} \mid \mathbf{y} \sim \mathcal{N}(\mathbf{h}^{\mathsf{T}}\hat{\mathbf{z}},\ \mathbf{h}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{h} + \sigma_T^2) = \mathcal{N}(23.0769,\ 0.17308 + 0.25) = \mathcal{N}(23.0769,\ 0.6504^2), \quad (4.64)$$

with $\mathbf{h}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{h} = \Sigma_{22}$ from equation (4.35). The predictive variance has two parts: the remaining uncertainty about the test temperature, $0.173\,\mathrm{K}^2$, and the error of the new reading itself, $0.25\,\mathrm{K}^2$. A further reading of the same sensor can therefore never be predicted more precisely than σ_T, however well the resistance becomes known.

An independent new reading compared with this distribution assesses the predictive performance of the model: a reading a fraction of a predictive standard deviation from 23.0769°C passes, and one several standard deviations away does not. Repeated large discrepancies, or a systematic sign in them, would indicate that the model, the prior, the measurement-error assumptions or the representation of the operating condition need revision.

4.4.5 The limit of exact rows

Table 4.7 repeats the fit at three tolerances and in the limit $\sigma_p \to 0$. Each column is one tolerance; the rows give the MAP resistance and full-load temperature, their standard deviations, the correlation of the two hall temperatures, the threshold probability of Section 4.3.5, and the condition number of the information matrix. By $\sigma_p = 0.01\,\mathrm{K}$ the full-load temperature and the threshold probability have stopped moving at the printed precision; the resistance and its spread still move in their sixth decimal, the full-load spread in its fourth; at $\sigma_p = 0.001\,\mathrm{K}$ all have stopped. The fit at $\sigma_p = 0.001\,\mathrm{K}$ reproduces the limit values of Section 4.2.10 to every printed digit.

The $\sqrt{\Sigma_{11}}$ row widens with the tolerance as equation (4.33) requires, the tolerance acting on the reading as extra sensor variance, and the limit column is the matrix of equation (4.35); its entry for $\sqrt{\Sigma_{11}}$ is the square root of $\boldsymbol{\Sigma}_{\mathrm{red}}$, the covariance of the reduced state (Section 4.2.11). The fitted covariances satisfy equation (4.34) at every tolerance of the table ($0.69260 = 0.69250 + 0.0001$ at $\sigma_p = 0.01\,\mathrm{K}$, $0.72125 = 0.71125 + 0.01$ at $0.1\,\mathrm{K}$). A tighter tolerance than

Table 4.7: The fit of the hall state as the model tolerance of the physical rows shrinks. The last column is the limit of exact rows.

Quantity	$\sigma_p = 0.1\,\mathrm{K}$	$0.01\,\mathrm{K}$	$0.001\,\mathrm{K}$	$\sigma_p \to 0$
$\hat{\rho}$ [K/kW]	0.169119	0.169230	0.169231	0.169231
$\sqrt{\Sigma_{11}}$ [K/kW]	0.014056	0.013869	0.013868	0.013868
$\hat{T}_2$ [°C]	28.1471	28.1538	28.1538	28.1538
$\sqrt{\Sigma_{33}}$ [K]	0.84926	0.83223	0.83205	0.83205
correlation of T_1, T_2 [–]	0.96522	0.99964	0.999996	1
$\Pr(T_2 < 28°\mathrm{C})$ [–]	0.4312	0.4267	0.4267	0.4267
condition number of $\mathbf{J}^\mathsf{T}\mathbf{J}$ [–]	4.0×10^5	3.9×10^7	3.9×10^9	–

0.001 K buys nothing and costs conditioning: the condition number grows a hundredfold per decade, and at $\sigma_p = 10^{-6}\,\mathrm{K}$ it reaches 4×10^{15} and round-off already moves the fourth digit of $\sqrt{\Sigma_{33}}$.

4.4.6 Second routes to the full-load answers

The propagated full-load temperature. As the third entry, with $\mathbf{g} = (0,\ 0,\ 1)^\mathsf{T}$, the full-load temperature has $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g} = \Sigma_{33} = 0.69231\,\mathrm{K}^2$; as the prediction $18{+}60\rho$, with $\mathbf{g} = (60,\ 0,\ 0)^\mathsf{T}$, it has $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g} = 60^2\,\Sigma_{11} = 3600/5200 = 0.69231\,\mathrm{K}^2$, the same number, as Section 4.3.1 requires in the limit.

The threshold through the resistance. By equation (4.58), $t^* = (28 - 28.1538)/60 = -0.0025641\,\mathrm{K/kW}$ and

$$\rho^* = \hat{z}_1 + t^* = 0.169231 - 0.002564 = 0.166667\,\mathrm{K/kW},$$
$$\Pr(\rho < \rho^* \mid \mathbf{y}) = \Phi\!\left(\frac{-0.002564}{0.013868}\right) = \Phi(-0.1849) = 0.4267,$$

the number of equation (4.61).

4.4.7 The propagation check against draws

For a linear output the delta method is exact, and the check is the algebraic agreement of the entry and the gradient already shown in Section 4.4.6. For a nonlinear output the linearized result must be compared with the full nonlinear evaluation at draws from the posterior. The procedure is the following.

1. Factor the covariance as $\boldsymbol{\Sigma} = \mathbf{L}\mathbf{L}^\mathsf{T}$ with $\mathbf{L}$ lower triangular (the Cholesky factor, which exists at any finite tolerance), and draw N independent states $\mathbf{z}^{(j)} = \hat{\mathbf{z}} + \mathbf{L}\boldsymbol{\varepsilon}^{(j)}$, $j = 1, \ldots, N$, each $\boldsymbol{\varepsilon}^{(j)}$ a vector of three independent standard normal numbers, so that $\mathbf{z}^{(j)} \sim \mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$. At $\sigma_p = 0.01\,\mathrm{K}$,
$$\mathbf{L} = \begin{bmatrix} 0.013869 & 0 & 0 \\ 0.41592 & 0.0099980 & 0 \\ 0.83217 & 0 & 0.010000 \end{bmatrix},$$
whose first column is $\sqrt{\Sigma_{11}}$ times (nearly) $\mathbf{v}$ and whose other two columns have the width σ_p of the physical rows. (In the limit of exact rows the factor has the single column $\mathbf{v}/\sqrt{5200}$, and a draw is $\hat{\mathbf{z}} + \mathbf{v}\,\varepsilon^{(j)}/\sqrt{5200}$, along the free direction.)
2. Evaluate the output at each drawn state through the full model, for example $T_2^{(j)} = z_3^{(j)}$ or $G^{(j)} = 1/z_1^{(j)}$.

3. Estimate a threshold probability as the fraction of draws satisfying the event, $\hat{P} = N^{-1}\sum_j \mathbf{1}[f(\mathbf{z}^{(j)}) < f^*]$, whose standard error is $\sqrt{\hat{P}(1-\hat{P})/N}$; estimate quantiles as the empirical quantiles of the $f(\mathbf{z}^{(j)})$.
4. Compare with the closed forms of the delta method. Agreement within the standard error means the closed form can be trusted.

A few hundred draws resolve a probability near one half to about ± 0.03; $N = 10^5$ resolves it to about ± 0.002.

Table 4.8 applies the procedure to the full-load temperature, which is linear in ρ, and to the conductance $G = 1/\rho$, which is not. The first column names the quantity; the second gives the delta-method value, from a Gaussian with mean $f(\hat{\mathbf{z}})$ and standard deviation $\sqrt{\mathbf{g}^\top \boldsymbol{\Sigma} \mathbf{g}}$; the third gives the exact value of the push-forward of the posterior, the limit to which the draw estimates converge, computed here in closed form because both outputs are monotone in ρ (the q-quantile of G is $1/(\hat{\rho} - z_q\sqrt{\Sigma_{11}})$ with z_q the standard normal quantile); the fourth gives one run of $N = 10^5$ drawn states at $\sigma_p = 0.01\,\mathrm{K}$, generated with a seeded pseudo-random generator so that the run can be repeated exactly. The conductance carries the gradient of Section 4.3.6.

Table 4.8: The propagation check: delta method against the exact push-forward of the posterior, whose first entry is $\mathcal{N}(0.169231, 0.013868^2)$, and one run of 10^5 states drawn as $\hat{\mathbf{z}} + \mathbf{L}\boldsymbol{\varepsilon}$ (standard error of a probability about 0.0015). The threshold $G = 6.25\,\mathrm{kW/K}$ is the handbook conductance.

Quantity	Delta method	Exact	10^5 draws
T_2, median [°C]	28.154	28.154	28.152
T_2, 5 % and 95 % [°C]	26.785, 29.522	26.785, 29.522	26.786, 29.516
$\Pr(T_2 < 28°\mathrm{C})$ [–]	0.4267	0.4267	0.4269
G, median [kW/K]	5.909	5.909	5.910
G, 5 % and 95 % [kW/K]	5.113, 6.706	5.207, 6.830	5.210, 6.828
$\Pr(G > 6.25\,\mathrm{kW/K})$ [–]	0.2407	0.2528	0.2527

For the temperature the three columns agree to within the sampling error, and the closed form needs no correction. For the conductance the median is right in all three columns, since $1/\rho$ is monotone and maps the median of ρ to the median of G, but the delta-method band is shifted: the exact band $[5.207, 6.830]$ is skewed to the right of $\hat{G}$, while the linearized band $[5.113, 6.706]$ is symmetric about it, and the delta method understates the probability that the conductance exceeds its handbook value by about 0.012. The draws reproduce the exact column to within their standard error. Figure 4.6 shows the cause. The horizontal axis is the conductance and the vertical axis its density; the solid curve is the exact density of $G = 1/\rho$, obtained by the change of variables $p_G(G) = p(\rho = 1/G \mid \mathbf{y})/G^2$ of Section 4.3.3, and the dashed curve the delta-method Gaussian. The exact density has its mode to the left of 5.909 and a longer right tail, and the Gaussian cannot represent the asymmetry.

The conductance also shows why "the posterior mean of G" must be handled with care. Under a Gaussian posterior for ρ, the mean of $1/\rho$ does not exist in the strict sense, because the Gaussian has positive density at $\rho = 0$, where $1/\rho$ is not integrable; the mass involved here is about 10^{-34} and has no practical effect, but it is a reminder that moments of a transformed quantity are properties of the parameterization. Quantiles and threshold probabilities, which the push-forward preserves exactly for a monotone output, are the robust summaries, and they are what the check compares.

4.4.8 Sensitivity to the value read

Table 4.9 repeats the assessment for other values of the single test reading. The first numeric column is the prior; each further column is the posterior after a different value of the reading,

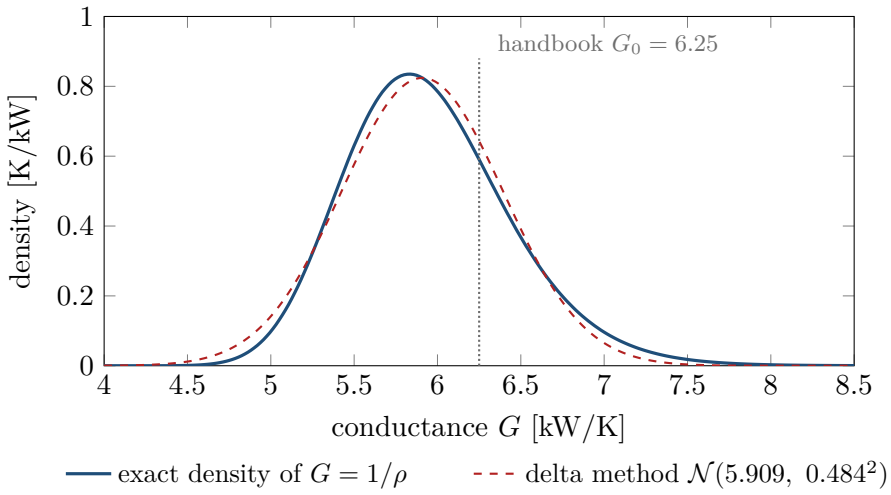


Figure 4.6: Posterior density of the conductance $G = 1/\rho$ under $\rho \mid \mathbf{y} \sim \mathcal{N}(0.169231, 0.013868^2)$: exact, by change of variables, $p_G(G) = \varphi((1/G - 0.169231)/0.013868)/(0.013868\,G^2)$ with φ the standard normal density; and the delta-method Gaussian $\mathcal{N}(1/\hat{\rho}, \Sigma_{11}/\hat{\rho}^4)$. Dotted line: the handbook conductance.

the actual one, 23.2°C, among them. The rows give the MAP resistance from equation (4.29), the mean and standard deviation of the full-load temperature, and the probability (4.57). The standard deviation is the same for every reading, and the mean moves along the gain vector (4.59), which with the limit covariance (4.35) is

$$\frac{\partial \hat{\mathbf{z}}}{\partial y_T} = \frac{\boldsymbol{\Sigma}\mathbf{h}}{\sigma_T^2} = 4 \begin{bmatrix} 0.0057692 \\ 0.17308 \\ 0.34615 \end{bmatrix} = \begin{bmatrix} 0.023077 \\ 0.69231 \\ 1.38462 \end{bmatrix},$$

four times the second column of the limit covariance, in K/kW per K and K per K. One column of the table, for $y_T = 22.4$°C, is the actual state moved by -0.8 K along it: $\hat{\rho} = 0.169231 - 0.8 \times 0.023077 = 0.150769$, $\mathbb{E}[T_2 \mid \mathbf{y}] = 28.1538 - 0.8 \times 1.38462 = 27.046$°C, and $\Phi((28 - 27.046)/0.8321) = \Phi(1.1464) = 0.8742$.

Table 4.9: The full-load assessment before the test and after a single test reading of various values ($\sigma_T = 0.5$ K, prior $\mathcal{N}(0.160, 0.025^2)$). The column 23.2°C is the observed reading.

Quantity	Prior	$y_T = 22.4$	22.8	23.2	23.6	24.0
$\hat{\rho}$ [K/kW]	0.16000	0.15077	0.16000	0.16923	0.17846	0.18769
$\mathbb{E}[T_2]$ [°C]	27.60	27.05	27.60	28.15	28.71	29.26
sd(T_2) [K]	1.500	0.832	0.832	0.832	0.832	0.832
$\Pr(T_2 < 28°\text{C})$ [–]	0.6051	0.8742	0.6846	0.4267	0.1975	0.0647

A reading of 22.8°C, equal to what the prior predicted, leaves the mean where it was and only narrows the distribution; because the mean then lies below the limit, narrowing alone raises the probability from 0.605 to 0.685. The observed 23.2°C moves the mean across the limit and lowers the probability to 0.427. The posterior probability equals the prior one for a reading of 22.93°C and falls below one half for any reading above 23.09°C.

4.5 The report (step 6)

One reading at the test condition overturns the handbook's answer. The handbook resistance put the hall at 27.6°C at full load, under the limit, with probability 0.6051 of staying below it; the reading of 23.2°C moves the resistance up, from 0.160 to 0.169 K/kW, and leaves the probability at 0.4267, near even odds, although it narrows the full-load spread from 1.5 to 0.832 K. The answers to the questions of Table 4.2 are given in the units in which the questions were asked, conditional on the steady-state model, the test reading and the handbook prior; the operating decision follows from them.

4.5.1 The answers and their bands

Table 4.10 gives the answers. Each row is a question; the second column gives the answer in the limit of exact physical rows, and the third its band: the posterior standard deviation of an estimate, with the central 95 % credible interval of the full-load temperature, and for the probability the range over the model tolerances of Table 4.7.

Table 4.10: The report for the equipment hall at 60 kW, after the test reading of 23.2°C. Answers in the limit of exact physical rows; bands are posterior standard deviations, a credible interval, or the range over model tolerances σ_p from 0 to 0.1 K.

Operator question	Answer	Band
What is the thermal resistance of the hall's air path?	0.16923 [K/kW]; conductance 5.909 [kW/K]	standard deviation 0.01387 [K/kW]; conductance 5–95 % band 5.113 to 6.706 [kW/K], exact 5.207 to 6.830 [kW/K]
How warm is the hall at the 30 kW test?	23.077 [°C]	standard deviation 0.416 [K]
How warm will the hall be at 60 kW?	28.15 [°C]	standard deviation 0.83 [K]; 26.52 to 29.78 [°C], central 95 % interval
Can the hall take 60 kW and stay below 28°C?	probability 0.427 [–]; of exceeding, 0.573 [–]	0.4267 to 0.4312 [–] over the model tolerances

The observed temperature was slightly warmer than the prior model predicted (22.8°C, the second entry of $\mathbf{z}^{(0)}$), so the inferred resistance rose above its handbook value and the implied conductance fell below 6.25 kW/K. The inferred test temperature is not the observed 23.2°C: the reading is noisy, and the prior, carried to T_1 by the physical row, also contributes information. The conductance in the first row is $1/\hat{\rho}$, the conductance evaluated at the estimated resistance; it is not the posterior mean of the conductance (Section 4.4.7). With correlations above 0.9996 the three entries are known almost only as one combination, the free direction $\mathbf{v}$, and the full-load interval contains temperatures both below and above the limit.

4.5.2 The operating decision

Conditional on the steady-state model, the reading and the prior, the estimated full-load temperature lies slightly above the limit, with substantial uncertainty on both sides of it. The reading reduced the uncertainty and, at the same time, lowered the probability of staying within the operating limit, from 0.605 before it to 0.427. The evidence does not support a high-confidence claim that the requested load keeps the hall within the limit: the posterior probability of doing so is about 0.43. A decision to raise the load to full would require an assurance level κ, stated explicitly, and a consideration of the consequences of exceeding the limit. Further readings or more cooling capacity may be appropriate, but their value should be

computed, not assumed. That is what an inference report gives the control room: the physical interpretation of a mathematical result, its uncertainty, and its implications for a decision.

Table 4.9 shows that the direction of the change was set by the value read, not by the act of reading. New evidence does not necessarily make the plant look safer: it can make it look less safe, or move the probability in either direction, depending on the value observed. The purpose of inference is to improve the knowledge of the plant, not to produce reassuring probabilities.

Principle 4.2. *A reading narrows the posterior by an amount fixed before it is taken, but the direction in which it moves a threshold probability is set by the value it returns. Evidence improves knowledge; it does not promise reassurance.*

4.6 The next reading (step 7)

The report leaves the operating question near even odds, with a band that a single reading at the test condition has set. A further reading is one more row of the residual vector, and the information it adds is known before its value is.

4.6.1 The preposterior covariance

A further reading at the test condition is the row $(T_1 - y)/\sigma_T$, whose line of the Jacobian is $\mathbf{j} = (0,\ 2,\ 0)$, the line of the first reading; a reading at full load, by a sensor of the same spread, is the row $(T_2 - y)/\sigma_T$, with line $\mathbf{j}_\mathrm{f} = (0,\ 0,\ 2)$. The information that a reading adds, $\mathbf{j}^\mathsf{T}\mathbf{j}$, is known before the reading is taken, because its Jacobian depends on the sensor and on where it is placed, not on the value. For a candidate reading a with information $\mathbf{\Lambda}_a$ the preposterior covariance is

$$\mathbf{\Sigma}_a = \left(\mathbf{\Sigma}^{-1} + \mathbf{\Lambda}_a\right)^{-1}, \tag{4.65}$$

formed at a finite tolerance, where $\mathbf{\Sigma}^{-1} = \mathbf{J}^\mathsf{T}\mathbf{J}$ exists, the object on which candidate readings are ranked. In the limit of exact rows the posterior stays on the free direction, and a candidate adds to the information $\|\mathbf{J}\mathbf{v}\|^2 = 5200$ along it the square of its own line times $\mathbf{v}$: $(\mathbf{j}\mathbf{v})^2 = 60^2 = 3600$ for a reading at 30 kW and $(\mathbf{j}_\mathrm{f}\mathbf{v})^2 = 120^2 = 14\,400$ for a reading at 60 kW. The full-load standard deviation after the candidate is 60 divided by the square root of the total. For n further readings at 30 kW, $\mathbf{\Lambda}_a = n\,\mathbf{j}^\mathsf{T}\mathbf{j}$, which adds $4n$ to the (T_1, T_1) entry, and the third diagonal entry of $\mathbf{\Sigma}_a$ gives the standard deviation of the full-load temperature: 0.640 K for $n = 1$, 0.539 K for $n = 2$, 0.429 K for $n = 4$ and 0.325 K for $n = 8$. More readings resolve the question; they do not decide its answer. If the full-load temperature truly lies above 28°C, the probability tends to zero as the readings accumulate, and the decision then rests on the air path, not on the readings.

4.6.2 The candidate readings

Table 4.11 ranks three candidates. Each row is a reading the operator can request; the columns give the information along $\mathbf{v}$ after it, the standard deviations of the resistance and of the full-load temperature that the preposterior covariance (4.65) predicts, and the width it removes from the current full-load standard deviation of 0.8321 K. One reading at full load carries four times the information of one at the test condition, because the full-load temperature moves twice as far as the test temperature along $\mathbf{v}$, and it removes exactly the width of four test readings.

Table 4.11: Candidate readings for the hall, in the limit of exact physical rows, from the posterior after the test reading. Each reading has $\sigma_T = 0.5\,\mathrm{K}$; the width removed is taken from the current $\sqrt{\Sigma_{33}} = 0.8321\,\mathrm{K}$.

Candidate	Information along $\mathbf{v}$ [(kW/K)^2]	sd(ρ) after [K/kW]	sd(T_2) after [K]	Width removed [K]
One further reading at 30 kW	8800	0.01066	0.640	0.192
Four further readings at 30 kW	19 600	0.00714	0.429	0.403
One reading at 60 kW	19 600	0.00714	0.429	0.403

4.6.3 The recommendation

Four further readings at the test condition and one reading at full load narrow the answer equally, and the choice between them rests on what else each tests. Readings at 30 kW need no operation at the requested load, but repeated readings of one sensor share its bias, and no number of them can show whether the resistance at 60 kW is the resistance at 30 kW. A reading at full load tests exactly that assumption, on which the prediction rests (Section 4.6.6), at the cost of running the hall at the load in question. The next reading is therefore one reading at full load, taken in a supervised run, the load raised in steps and held at 60 kW until the hall settles, which the operator can stop if the hall approaches its limit; where such a run is not acceptable, it is four further readings at the test condition, by more than one sensor where possible. Chapter 18 ranks the readings an operator can request on the shift plant in the same way.

4.6.4 A second observation as an update

Before any run at full load, the first candidate of Table 4.11 arrives: a second, independent reading at the test condition, $y_{T,2} = 23.0^\circ\mathrm{C}$, with the same standard deviation $\sigma_T = 0.5\,\mathrm{K}$; the first reading is now written $y_{T,1} = 23.2^\circ\mathrm{C}$. Against the posterior predictive (4.64) it lies $(23.0 - 23.0769)/0.6504 = -0.12$ standard deviations from the predicted mean and passes that check. The posterior after the first reading serves as the prior for the second.

The precision form. Information from independent rows adds. The new reading is one more row of the residual vector, $(T_1 - y_{T,2})/\sigma_T$, whose line of the Jacobian is $\mathbf{j} = (0,\ 2,\ 0)$, the same as the first reading's. Its contribution to the information matrix is the outer product $\mathbf{j}^{\mathsf{T}}\mathbf{j}$, which changes the (T_1, T_1) entry alone,

$$\boldsymbol{\Lambda}_2 = \mathbf{J}^{\mathsf{T}}\mathbf{J} + \mathbf{j}^{\mathsf{T}}\mathbf{j} = \mathbf{J}^{\mathsf{T}}\mathbf{J} + \begin{bmatrix} 0 & 0 & 0 \\ 0 & 4 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \qquad \boldsymbol{\Sigma}_2 = \boldsymbol{\Lambda}_2^{-1}.$$

The old rows are balanced at $\hat{\mathbf{z}}$, $\mathbf{J}^{\mathsf{T}}\mathbf{r}(\hat{\mathbf{z}}) = \mathbf{0}$, so the gradient of the enlarged objective at $\hat{\mathbf{z}}$ comes from the new row alone, $\mathbf{j}^{\mathsf{T}}(23.0769 - 23.0)/0.5 = (0,\ 0.30769,\ 0)^{\mathsf{T}}$, and one Gauss–Newton step (4.29) from $\hat{\mathbf{z}}$ gives $\hat{\mathbf{z}}_2 = \hat{\mathbf{z}} - 0.30769\,(\boldsymbol{\Sigma}_2)_{\cdot 2}$, with $(\boldsymbol{\Sigma}_2)_{\cdot 2}$ the second column of $\boldsymbol{\Sigma}_2$. In the limit of exact rows this column is $(3.4091 \times 10^{-3},\ 0.10227,\ 0.20455)^{\mathsf{T}}$, and

$$\hat{\mathbf{z}}_2 = \begin{bmatrix} 0.169231 \\ 23.0769 \\ 28.1538 \end{bmatrix} - 0.30769 \begin{bmatrix} 0.0034091 \\ 0.10227 \\ 0.20455 \end{bmatrix} = \begin{bmatrix} 0.168182 \\ 23.0455 \\ 28.0909 \end{bmatrix}.$$

The first entry of $\boldsymbol{\Sigma}_2$ is $1.1364 \times 10^{-4}\,(\mathrm{K/kW})^2$, so that $\sqrt{(\boldsymbol{\Sigma}_2)_{11}} = 0.010660\,\mathrm{K/kW}$.

The rank-one form. The same result follows from the first posterior alone, without forming and inverting $\mathbf{\Lambda}_2$. The second reading measures $\mathbf{h}^\mathsf{T}\mathbf{z} = T_1$ with $\mathbf{h} = (0,\ 1,\ 0)^\mathsf{T}$; its predicted value under the first posterior is $\mathbf{h}^\mathsf{T}\hat{\mathbf{z}} = 23.0769\,^\circ\mathrm{C}$ and its predictive variance, from equation (4.64), is $S = \mathbf{h}^\mathsf{T}\mathbf{\Sigma}\mathbf{h} + \sigma_T^2 = \Sigma_{22} + \sigma_T^2 = 0.42308\,\mathrm{K}^2$. Conditioning the Gaussian on the new row gives the update

$$\hat{\mathbf{z}}_2 = \hat{\mathbf{z}} + \mathbf{K}\,(y_{T,2} - \mathbf{h}^\mathsf{T}\hat{\mathbf{z}}), \qquad \mathbf{\Sigma}_2 = \mathbf{\Sigma} - S\,\mathbf{K}\mathbf{K}^\mathsf{T}, \qquad \mathbf{K} = \frac{\mathbf{\Sigma}\mathbf{h}}{S}, \tag{4.66}$$

where $\mathbf{K}$ is the gain vector, one entry per state entry, and $y_{T,2} - \mathbf{h}^\mathsf{T}\hat{\mathbf{z}}$ the innovation, the part of the reading the first posterior did not predict. Because $\mathbf{\Sigma}_2\mathbf{h} = \mathbf{\Sigma}\mathbf{h} - \mathbf{K}\,\mathbf{h}^\mathsf{T}\mathbf{\Sigma}\mathbf{h} = (\sigma_T^2/S)\,\mathbf{\Sigma}\mathbf{h}$, the gain is also $\mathbf{K} = \mathbf{\Sigma}_2\mathbf{h}/\sigma_T^2$, the same vector as the gain (4.59) of the posterior that includes the new reading. The gain follows from writing the posterior mean as $\hat{\mathbf{z}} + \mathrm{Cov}(\mathbf{z}, Y_\mathrm{new})\,(y_{T,2} - \mathbb{E}[Y_\mathrm{new}])/S$ for the jointly Gaussian state and new reading, with $\mathrm{Cov}(\mathbf{z}, Y_\mathrm{new}) = \mathbf{\Sigma}\mathbf{h}$, the second column of $\mathbf{\Sigma}$. With the limit covariance (4.35),

$$\mathbf{K} = \frac{1}{0.42308}\begin{bmatrix} 0.0057692 \\ 0.17308 \\ 0.34615 \end{bmatrix} = \begin{bmatrix} 0.013636 \\ 0.40909 \\ 0.81818 \end{bmatrix},$$

in kW^{-1} for the first entry and dimensionless for the other two; the innovation is $23.0 - 23.0769 = -0.0769\,\mathrm{K}$, and $\hat{\mathbf{z}}_2 = \hat{\mathbf{z}} - 0.0769\,\mathbf{K} = (0.168182,\ 23.0455,\ 28.0909)^\mathsf{T}$, the state of the precision form. The covariance update subtracts the outer product $S\mathbf{K}\mathbf{K}^\mathsf{T}$; on its first and last diagonal entries,

$$(\mathbf{\Sigma}_2)_{11} = 1.9231 \times 10^{-4} - 0.42308 \times 0.013636^2 = 1.1364 \times 10^{-4}\,(\mathrm{K/kW})^2,$$
$$(\mathbf{\Sigma}_2)_{33} = 0.69231 - 0.42308 \times 0.81818^2 = 0.40909\,\mathrm{K}^2,$$

the values of the precision form. Because $S\,\mathbf{K}^\mathsf{T} = \mathbf{h}^\mathsf{T}\mathbf{\Sigma}$, the covariance update is also

$$\mathbf{\Sigma}_2 = (\mathbf{I} - \mathbf{K}\mathbf{h}^\mathsf{T})\,\mathbf{\Sigma},$$

with $\mathbf{I}$ the identity matrix: each column of $\mathbf{\Sigma}$ loses the gain vector times that column's second entry. The correction is the outer product of one vector, which is the reason for the name rank-one update. For a linear-Gaussian model it is exact, and one more independent reading costs one product of $\mathbf{\Sigma}$ with a vector rather than a new inference. For a strongly nonlinear model, or after a shift of the estimate large enough to make the Jacobian at the old estimate stale, the model is linearized again or the full fit is repeated.

The updated assessment. Table 4.12 follows the assessment through the two readings. Each column is a state of knowledge, before any reading, after the first, after both; the rows are the information about the resistance, the reciprocal of the first diagonal entry of the covariance, the MAP resistance, the mean and standard deviation of the full-load temperature, the threshold probability, and the 95 % credible interval.

After the second reading the full-load temperature is the third entry of $\hat{\mathbf{z}}_2$, 28.091°C, with standard deviation $\sqrt{(\mathbf{\Sigma}_2)_{33}} = \sqrt{0.40909} = 0.640\,\mathrm{K}$, and the probability of staying below the limit becomes $\Phi(-0.1421) = 0.4435$. The interval narrowed by a quarter, yet the probability moved only from 0.427 to 0.443, because two effects acted in opposite directions. The narrowing alone, with the mean held at 28.154°C, would have lowered the probability to 0.405, since the mean lies above the limit; the shift of the mean alone, with the spread held at 0.832 K, would have raised it to 0.456. The effect of actual new evidence cannot be described by its reduction of the covariance alone: the value observed matters.

Table 4.12: The full-load assessment before the test, after the first reading (23.2°C) and after the second (23.0°C).

Quantity	Prior	One reading	Two readings
$1/\Sigma_{11}$ [(kW/K)^2]	1600	5200	8800
$\hat{\rho}$ [K/kW]	0.160000	0.169231	0.168182
$\mathbb{E}[T_2]$ [°C]	27.600	28.154	28.091
sd(T_2) [K]	1.500	0.832	0.640
Pr($T_2 < 28$°C) [–]	0.6051	0.4267	0.4435
95 % interval [°C]	24.66, 30.54	26.52, 29.78	26.84, 29.34

4.6.5 What the analysis showed

The handbook said the hall passes at 27.6°C, and the reading scaled to full load said it fails at 28.4°C; the analysis says neither. With the reading, the handbook and the physics combined, the probability that the hall stays below 28°C at 60 kW is 0.4267, less than even, and the full-load temperature is not known well enough to call it either way. One reading at the test condition fixes the thermal resistance of the air path to about ±8 % and the full-load temperature to about ±0.8 K, one standard deviation each, although the hall has never run at 60 kW: the physics carries the reading along the free direction $\mathbf{v}$, from T_1 to ρ and from ρ to T_2. The posterior is a ridge along that direction, the three entries known almost only as one combination, and the reduced state is that ridge named by its one free entry: every question asked of the hall is a question about where it sits on the ridge. The 60 kW question in particular is a resistance question: the hall stays below 28°C exactly when the resistance stays below ρ^*, and the reading at 30 kW places the resistance close to ρ^*, which leaves the question near even odds. Further readings at the test condition narrow the ridge and resolve the question without deciding its answer. A reading at full load would act on the third column of $\mathbf{J}$ directly, and its worth is what its preposterior covariance $\mathbf{\Sigma}_a$ says before it is taken: as much as four readings at the test condition, with the shared resistance itself put to the test.

4.6.6 Limits of the model

The steady-state balance (4.2) neglects the heat stored in the hall and its equipment, and it describes the temperature long after a load change, not minutes after it; Exercise 4.6 formulates the transient model that a question about the first minutes of full load would need. The single effective path assumes that the resistance is the same at 60 kW as at 30 kW; a path whose resistance changes with load, for example because the unit's fans or its coil change state, would break the link drawn in Figure 4.2, and one test reading cannot detect it. The supply air is pinned at 18°C by the discharge controller, which holds it only while the unit has capacity to spare; a unit at its limit at full load lets the supply air rise, which shifts the full-load temperature up, not only widens it; an uncertain supply air adds its variance through the budget (4.54).

The probabilistic assumptions have limits of their own. The Gaussian prior on a positive resistance is harmless here, with a negative mass of 8×10^{-11}, but a wide prior on a small quantity would place real probability below zero; a logarithmic parameterization avoids this at the price of a nonlinear model and an approximate posterior. The sensor error is taken as Gaussian, unbiased and of known spread; a bias would shift every inferred quantity, and Section 4.2.2 names the comparison of models that would detect it. The physical rows are taken in the limit of zero tolerance; a real model discrepancy is a finite σ_p, which widens Σ_{11} and adds σ_p^2 to the full-load variance by equation (4.34), and the reduced state of Section 4.2.8 is exact only in that limit. The reading of 23.2°C is a stated example value, not field data,

and the chapter has no truth to check its answers against; from Chapter 5 on the readings are constructed from a stated truth and the checks compare the answers with it. Each of these is a statement about the model, and the probability 0.4267 is conditional on all of them.

Chapter 5 grows the hall: two coils fed with chilled water produce the supply air, now the cold aisle, and three racks draw from it; the method is unchanged, and it runs on the reduced state of this chapter.

4.7 What to retain

Definition 4.1 (The anatomy of an inference)**.** An operating probability over a plant graph is obtained in two operations. Inference sums one term per row into the objective $\mathcal{L} = -\ln p(\mathbf{z} \mid \mathbf{y})$ up to a constant, and uses only its gradient and Hessian. A Gaussian row's term is half its squared residual; a count row's is its negative log-probability, which joins the Jacobian through its working residual and line, reproducing the row's gradient and information but not its value. The physical rows, the readings and the priors, each divided by its standard deviation, stack into one residual vector $\mathbf{r}(\mathbf{z})$ with Jacobian $\mathbf{J}$; the normal equations $(\mathbf{J}^\mathsf{T}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^\mathsf{T}\mathbf{r}$ lead to the MAP state $\hat{\mathbf{z}}$, and the inverse of the information matrix, $\mathbf{\Sigma} = (\mathbf{J}^\mathsf{T}\mathbf{J})^{-1}$, is the posterior covariance, exactly for a linear-Gaussian model and under the Laplace approximation, in its Gauss–Newton form, otherwise. An exact physical relation is the limit of a row's tolerance, never a substitution; in that limit the state reduces to its free entries, with Jacobian $\mathbf{J}_{\text{rows}}\mathbf{N}$ on the null space $\mathbf{N}$ of the physical rows. Query evaluation then reads the answer off the posterior: an entry of the state with its variance Σ_{ii}, or a derived quantity $f(\mathbf{z})$ with variance $\mathbf{g}^\mathsf{T}\mathbf{\Sigma}\mathbf{g}$, and a threshold probability as the normal tail $\Phi\big((f^* - f(\hat{\mathbf{z}}))/\sqrt{\mathbf{g}^\mathsf{T}\mathbf{\Sigma}\mathbf{g}}\big)$, which for an entry is $\Phi\big((f^* - \hat{z}_i)/\sqrt{\Sigma_{ii}}\big)$.

Three conclusions carry beyond the example. Physical relations connect the quantities that are measured to those that are not, which is what allows an unobserved quantity, here a full-load temperature never reached, to be inferred from indirect evidence. Uncertainty is not an error bar attached to a point estimate but a probability distribution, and it must be propagated through the physical model to answer an operator's question. And the posterior is reusable: once the joint posterior is in hand, many questions are answered without repeating the inference, and new evidence updates it rather than replacing it. A derived quantity is propagated in one of two ways: its spread by the gradient, $\mathbf{g}^\mathsf{T}\mathbf{\Sigma}\mathbf{g}$, and its density by the change of variables, which multiplies the density by the change-of-variable factor of the map and so keeps the skew and the bounds that a linearization loses.

The inference of the hall returns the posterior state $\hat{\mathbf{z}}$ of equation (4.49), its covariance $\mathbf{\Sigma}$ of equations (4.50) and (4.35), and the Jacobian $\mathbf{J}$ of equation (4.25), from which the gradients of derived quantities are formed. Table 4.13 lists what the chapter has computed from them without repeating the inference: the four answers of Table 4.2, and three further quantities that later chapters reuse. Each row is a quantity, the calculation on the posterior that gives it, and the place where it is carried out.

Table 4.13: The four answers of Table 4.2 and three further quantities reused later, all computed from one posterior of the equipment hall.

Quantity	Calculation on the posterior	Section
The answers		
Thermal resistance of the air path	the entry $\hat{z}_1$ and $\sqrt{\Sigma_{11}}$	4.3.6
Hall temperature at the test condition	the entry $\hat{z}_2$ and $\sqrt{\Sigma_{22}}$	4.3.6
Hall temperature at full load	the entry $\hat{z}_3$, $\sqrt{\Sigma_{33}}$ and the credible interval	4.3.6
Probability of staying below 28°C at 60 kW	the Gaussian tail $\Phi\big((T_{\max} - \hat{z}_3)/\sqrt{\Sigma_{33}}\big)$	4.3.7
Further quantities reused later		
Spread of the conductance	$\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ with $\mathbf{g} = (-1/\hat{z}_1^2, 0, 0)^{\mathsf{T}}$, and the exact push-forward	4.3.6
The entry that drives the temperature	the gradient and the variance budget	4.3.2
Effect of another reading	the information matrix plus the reading's $\mathbf{j}^{\mathsf{T}}\mathbf{j}$	4.6.1

The table is one inference followed by a column of cheap queries on its result, the division of Section 4.1.2. Two kinds of question fall outside it: a question that compares two models, which needs the evidence of each (Section 4.2.2), and the arrival of new evidence, which changes the posterior itself (Section 4.6.4).

Table 4.14 lists the five mathematical objects of the chapter that recur throughout the book, with the form of each and its purpose. The first three are the working objects when the posterior is well represented by a local Gaussian. The posterior distribution itself is the more fundamental object: it may be Gaussian, multimodal, discrete, a mixture or otherwise non-Gaussian, and its structure decides which computational method is appropriate.

Table 4.14: The five mathematical objects of an inference.

Object	Form	Purpose
State estimate	$\hat{\mathbf{z}}$	the most probable (or otherwise selected) central state
Posterior covariance	$\boldsymbol{\Sigma}$	local uncertainty of the state and the dependence among its entries
Output gradient	$\mathbf{g} = \nabla f(\hat{\mathbf{z}})$	connects state uncertainty to a derived quantity
Posterior distribution	$p(\mathbf{z} \mid \mathbf{y}, M)$	the uncertainty about the state given the evidence
Model evidence	$p(\mathbf{y} \mid M)$	normalizes the posterior and compares models

Table 4.15 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 4.15: The terms of the anatomy of an inference and where each is defined. Units for the equipment hall.

Term	Expression	Units	Meaning	Defined in
Thermal resistance	$\rho = 1/G$, $T = T_c + Q\rho$	K/kW	makes every row linear in the state	4.1.4
Physical rows	$r_1 = T_1 - 18 - 30\rho$, r_2	K	conservation at each hall node	4.1.5
The question as one marginal	$\Pr(T_2 < T_{\max})$ from $\hat{z}_3$, Σ_{33}	–	one entry's mean and variance	4.1.6
Model tolerance	σ_p; exact rows as $\sigma_p \to 0$	K	model discrepancy of a physical row	4.1.5
Likelihood	$p(y_T \mid \mathbf{z})$, eq. (4.14)	K^{-1}	the reading as a function of the state	4.2.1
Residual vector	$\mathbf{r}(\mathbf{z})$, eq. (4.23)	–	four rows, each over its standard deviation	4.2.3
Objective	$\mathcal{L} = \frac{1}{2}\mathbf{r}^{\mathsf{T}}\mathbf{r}$	–	minus log posterior	4.2.3
Jacobian	$\mathbf{J} = \partial\mathbf{r}/\partial\mathbf{z}$, 4×3	mixed	one line per row, one column per entry	4.2.3
Information matrix	$\mathbf{J}^{\mathsf{T}}\mathbf{J} = \sum_k \mathbf{j}_k^{\mathsf{T}}\mathbf{j}_k = \tilde{\mathbf{J}}^{\mathsf{T}}\mathbf{W}\tilde{\mathbf{J}}$	mixed	the Hessian when $\mathbf{C} = \mathbf{0}$	4.2.3
General objective	$\mathcal{L} = \sum_k \mathcal{L}_k$, $\mathcal{L}_k = -\ln p_k$	–	one term per row, any distribution	4.2.7
Count row, working form	$r_k = (\mu_k - n_k)/\sqrt{\mu_k}$, $\mathbf{j}_k = \sqrt{\mu_k}\,\nabla \ln \mu_k^{\mathsf{T}}$	–	gradient and information, not the term	4.2.7
Hessian	$\mathbf{H} = \mathbf{J}^{\mathsf{T}}\mathbf{J} + \mathbf{C}$	mixed	the information plus a named curvature term	4.2.7
Normal equations	$(\mathbf{J}^{\mathsf{T}}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^{\mathsf{T}}\mathbf{r}$	mixed	Gauss–Newton step to $\hat{\mathbf{z}}$	4.2.4
Posterior covariance	$\mathbf{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1}$	mixed	spread is inverse information	4.2.5
Laplace approximation	$\mathcal{N}(\hat{\mathbf{z}}, \mathbf{H}^{-1})$	–	local Gaussian posterior	4.2.6
Free direction	$\mathbf{v}$, $\mathbf{J}_{\mathrm{phys}}\mathbf{v} = \mathbf{0}$	mixed	$\mathbf{\Sigma} \to \mathbf{v}\mathbf{v}^{\mathsf{T}}/5200$, rank one	4.2.4
Delta method	$\mathrm{Var}(f) \approx \mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}$	varies	linearized propagation	4.3.1
Change of variables	$p_u(u) = p_x(h^{-1}(u))\,\lvert \mathrm{d}x/\mathrm{d}u \rvert$, eq. (4.55)	varies	density of a transformed quantity	4.3.3
Variance budget	eq. (4.54)	K^2	one term per entry plus cross terms	4.3.2
Threshold probability	$\Phi\big((T_{\max} - \hat{z}_3)/\sqrt{\Sigma_{33}}\big)$	–	normal tail at the limit	4.3.5
Credible interval	$\hat{z}_3 \pm 1.960\sqrt{\Sigma_{33}}$	°C	central 95 % posterior range	4.3.4
Reduced state	$\mathbf{N}$, $\mathbf{J}_{\mathrm{red}} = \mathbf{J}_{\mathrm{rows}}\mathbf{N}$, $\mathbf{\Sigma}_{\mathrm{red}}$	mixed	the free entries with exact physical rows	4.2.8
Prior predictive	$\mathcal{N}(22.8, 0.9014^2)$	°C	reading predicted before use	4.4.2
Posterior predictive	$\mathcal{N}(23.077, 0.6504^2)$	°C	next reading predicted	4.4.4
Evidence	$p(y_T \mid M)$, eq. (4.18)	K^{-1}	normalizer; compares models	4.2.2
Bayes factor	$B_{21} = p(y_T \mid M_2)/p(y_T \mid M_1)$	–	evidence ratio of two models	4.2.2
Rank-one update	eq. (4.66)	–	adds one reading to a posterior	4.6.4
Preposterior covariance	$\mathbf{\Sigma}_a = (\mathbf{\Sigma}^{-1} + \mathbf{\Lambda}_a)^{-1}$, eq. (4.65)	mixed	ranks candidate readings	4.6.1

4.8 Exercises

Exercise 4.1 (Application: a complete inference)**.** A different equipment hall obeys $T = T_c + Q\rho$ with its supply air held at $T_c = 20°\mathrm{C}$ and a test load $Q = 20\,\mathrm{kW}$. The prior on the resistance is $\rho \sim \mathcal{N}(0.20, 0.04^2)$ in K/kW, and a sensor with $\sigma_T = 0.4\,\mathrm{K}$ reads $y_T = 24.5°\mathrm{C}$. The question is the temperature at a requested load of 40 kW. (a) Write the state $\mathbf{z} = (\rho, T_1, T_2)^\mathsf{T}$, with T_1 the test temperature and T_2 the temperature at 40 kW, the four normalized rows, and the 4×3 Jacobian at $\sigma_p = 0.01\,\mathrm{K}$. ($\mathbf{r} = ((T_1 - 20 - 20\rho)/\sigma_p,\ (T_2 - 20 - 40\rho)/\sigma_p,\ (T_1 - 24.5)/0.4,\ (\rho - 0.20)/0.04)^\mathsf{T}$; $\mathbf{J}$ has the lines $(-2000, 100, 0)$, $(-4000, 0, 100)$, $(0, 2.5, 0)$ and $(25, 0, 0)$.) (b) Take one Gauss–Newton step from the handbook state and compute $\boldsymbol{\Sigma}$ in the limit of exact rows. (At $\mathbf{z}^{(0)} = (0.20,\ 24.0,\ 28.0)^\mathsf{T}$ only the reading row is nonzero, -1.25, and $\mathbf{J}^\mathsf{T}\mathbf{r} = (0,\ -3.125,\ 0)^\mathsf{T}$; the normal equations give $\hat{\mathbf{z}} = (0.2200,\ 24.400,\ 28.800)^\mathsf{T}$, and a second step does not move it. The free direction is $\mathbf{v} = (1,\ 20,\ 40)^\mathsf{T}$, with $\mathbf{J}\mathbf{v} = (0,\ 0,\ 50,\ 25)^\mathsf{T}$ and information $50^2 + 25^2 = 3125$, so $\boldsymbol{\Sigma} \to \mathbf{v}\mathbf{v}^\mathsf{T}/3125$: $\Sigma_{11} = 3.2 \times 10^{-4}$, $\sqrt{\Sigma_{11}} = 0.01789\,\mathrm{K/kW}$; at $\sigma_p = 0.01\,\mathrm{K}$ the fit gives $\Sigma_{11} = 3.2016 \times 10^{-4}$ and a correlation of 0.9996 between ρ and T_1.) (c) Read off the posterior mean and standard deviation of the temperature at 40 kW, and the probability that it stays below 30°C. ($\hat{z}_3 = 28.8°\mathrm{C}$ and $\sqrt{\Sigma_{33}} = \sqrt{1600/3125} = 0.7155\,\mathrm{K}$; $\Pr(T_2 < 30 \mid \mathbf{y}) = \Phi(1.2/0.7155) = \Phi(1.677) = 0.9532$, against $\Phi(2/1.6) = 0.8944$ under the prior alone.)

Exercise 4.2 (Derivation: a variance budget with correlated entries)**.** For $T = T_c + Q\rho$ at the known load $Q = 50\,\mathrm{kW}$, the posterior covariance of (T_c, ρ) is $\boldsymbol{\Sigma} = \begin{bmatrix} 0.25 & -0.003 \\ -0.003 & 0.0001 \end{bmatrix}$, in K^2, K·K/kW and $(\mathrm{K/kW})^2$. (a) Derive the temperature variance from $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$ and compute the standard deviation. ($\mathbf{g} = (1, 50)^\mathsf{T}$; $\mathrm{Var}(T) = 0.25 + 50^2 \times 0.0001 + 2 \times 50 \times (-0.003) = 0.25 + 0.25 - 0.30 = 0.20\,\mathrm{K}^2$, $\mathrm{sd}(T) = 0.4472\,\mathrm{K}$.) (b) State the result if the covariance term were ignored, and explain the difference. ($\sqrt{0.50} = 0.7071\,\mathrm{K}$, 58 % too large; the posterior correlation is $-0.003/\sqrt{0.25 \times 0.0001} = -0.6$, and with sensitivities of the same sign the negative cross term cancels part of the two diagonal terms: the evidence pins the combination $T_c + 50\rho$ better than either entry.)

Exercise 4.3 (Derivation: a nonlinear derived quantity)**.** With $\rho \mid \mathbf{y} \sim \mathcal{N}(0.169231, 0.013868^2)$ from the chapter's example, derive the first-order approximation to the variance of the conductance $G = 1/\rho$, compute its standard deviation, and explain why the result is only an approximation. (With $G(\mathbf{z}) = 1/z_1$, $\mathbf{g} = (-1/\hat{z}_1^2,\ 0,\ 0)^\mathsf{T}$, so $\mathrm{Var}(G \mid \mathbf{y}) \approx \mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g} = \Sigma_{11}/\hat{z}_1^4$; $\hat{G} = 1/0.169231 = 5.909\,\mathrm{kW/K}$ and $\mathrm{sd}(G \mid \mathbf{y}) \approx 0.013868/0.169231^2 = 0.484\,\mathrm{kW/K}$. The map $1/\rho$ is nonlinear, so G is not Gaussian: its exact 5–95 % band is $[5.207, 6.830]$, skewed, against the linearized $[5.113, 6.706]$; strictly, the moments of $1/\rho$ do not exist under a Gaussian ρ, which has positive density at zero, so positivity and the transformed distribution must be handled explicitly when exact conductance statements are needed.)

Exercise 4.4 (Formulation: observability and the plant graph)**.** In $T = T_c + Q\rho$ suppose that both T_c and ρ are unknown, the discharge controller's set point not being trusted, and that there is a single temperature reading at one known load. (a) Write the observation Jacobian with respect to (T_c, ρ) and state whether the reading can determine both unknowns without further information. ($y = T_c + Q\rho + \epsilon$, so $\mathbf{J} = [1 \quad Q]$, of rank one for two unknowns: the reading determines only the combination $T_c + Q\rho$, and every pair with the same combination explains it equally well.) (b) Name a reading or operating condition that resolves the ambiguity. (A second reading at a different known load $Q_2 \neq Q_1$ gives $\mathbf{J} = \begin{bmatrix} 1 & Q_1 \\ 1 & Q_2 \end{bmatrix}$, with determinant $Q_2 - Q_1 \neq 0$ and full rank; alternatively, a direct reading of the supply air, by a discharge sensor, a new node observed in the graph, removes the ambiguity.)

Exercise 4.5 (Application: Bayesian evidence)**.** A model predicts $Y \sim \mathcal{N}(24, 1^2)$ in °C, and the reading $y = 25$°C arrives. (a) Compute the marginal likelihood $p(y \mid M)$. ($\frac{1}{\sqrt{2\pi}} \exp[-(25-24)^2/2] = 0.39894 \times 0.60653 = 0.2420\,\mathrm{K}^{-1}$.) (b) Explain why this value is not the probability that M is correct. (It is a probability density in inverse temperature units, 2.420×10^{-4} per millikelvin for the same reading, not $\Pr(M \mid y)$; a posterior model probability requires prior probabilities and the evidence of the alternative models under comparison.)

Exercise 4.6 (Formulation: model validity and thermal storage)**.** The chapter assumed the steady-state balance $Q = G(T - T_c)$. The load now steps from 30 kW to 60 kW, and the operator asks for the temperature five minutes later. (a) Explain why the steady-state equation may be insufficient. (It neglects the energy stored in the air, structure and equipment of the hall, so it gives only the temperature reached long after the change; five minutes may be short against the thermal time constant.) (b) Starting from an energy balance, formulate a model with thermal storage and name the new parameter and the condition it needs. (With an effective thermal capacitance C_{th} in kJ/K, $C_{\mathrm{th}}\,\mathrm{d}T/\mathrm{d}t = Q - G(T - T_c)$; the new parameter is C_{th} and an initial temperature $T(0)$ is required. For constant Q, G and T_c after the step, $T(t) = T_c + Q/G + [T(0) - T_c - Q/G]\,\mathrm{e}^{-Gt/C_{\mathrm{th}}}$, with time constant C_{th}/G; the probability for the temperature at five minutes propagates the posterior of G, C_{th} and $T(0)$ through this time-dependent model rather than through the steady-state equation.)

Part II

Estimating the State of the Plant

Chapter 5

Virtual Sensors for Unmeasured Temperatures

5.1 Rack inlets the hall does not measure

It is 23:00, and the night operator is reading the hall's tile on the dashboard. The hall now has two cooling units and three racks, and only rack 1 has an inlet sensor; it reads 22.18°C, 4.82 K inside the 27°C limit of the recommended envelope. The tile labeled "rack inlets" shows that one number, in green. Rack 3 has no sensor, and the last airflow survey put the most recirculation at it. The design sheet, worked forward at rated units, surveyed airflow and 200 kW, puts rack 3 at 23.66°C. Both numbers say the same thing to the operator: the inlets are several degrees inside the limit, so the racks are fine.

Neither number is a statement about rack 3 at 23:00. Rack 1's sensor reads rack 1, and the design sheet reads a hall that runs at its ratings. What neither settles is what rack 3's inlet is tonight, and how well the five readings the hall does have, its return air, the two discharge temperatures, the hall load meter and rack 1's inlet, know it. The question of this chapter is

how the temperature at a place with no sensor is read from the fitted state of the plant, with its band, what that band is made of, and which reading would narrow it.

It is answered with the method of Chapter 4, now on a state of six entries and the reduced state derived there. What changes is the quantity asked for: it is not an entry of the state and no reading row touches it.

Jargon: state estimation, virtual (soft) sensor, cold aisle, return air, chilled water, discharge air temperature, air capacity rate, log-derate, recirculation fraction, airflow survey, recommended envelope, observability, variance budget, probe reading

State estimation is the fit of the state of one telemetry window to its readings and priors. A *virtual sensor*, or *soft sensor*, is a quantity of the plant read from that fit, with its band, in place of an instrument that is missing. The *cold aisle* is the space in front of the racks that the cooling units fill with cooled air; the *return air* is the warm air leaving the racks on its way back to the units, here the hall air of Part I. *Chilled water* reaches each unit's coil from the plant's chiller at the supply temperature T_s. The *discharge air temperature* of a unit is the temperature of the air leaving its coil. The *air capacity rate* C_u of a unit, in kW/K, is its air mass flow times the specific heat of air. A unit's *log-derate* $g_u = \ln(G_u/\bar{G}_u)$ compares its coil's effective conductance with the rating; it is zero at the rating. The *recirculation fraction* β_k of rack k is the share of the air at its inlet that is return air drawn back over or around the rack instead of cold-aisle air. An *airflow survey* is a one-time measurement campaign that estimates those fractions. The *recommended envelope* is the range of inlet temperatures in which the racks are meant to run; its upper limit here is 27°C. *Observability* (in the framework, the rank of the Jacobian's columns and the posterior spread of an entry against its prior spread) is whether the readings determine an entry of the state. The *variance budget* splits the variance of a derived quantity into one term per entry and the cross terms. A *probe reading* is a spot reading taken by a technician with a handheld instrument.

Key idea

A virtual sensor is a derived quantity of the fitted state. Its value is the quantity at the MAP state and its variance is $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$, with $\mathbf{g}$ its gradient; the budget of that variance says which entries, and so which rows, stand behind the band. When the entry that carries most of the budget is seen by no reading, the band is the prior's, pushed through the physics, and a sensor anywhere else in the plant cannot narrow it.

5.1.1 A virtual sensor and what stands behind its band

Chapter 4 read a full-load temperature that the hall had never reached as an entry of its state, and the band of that entry came from one reading carried along the free direction by the physics. The inlet of rack 3 is a temperature of the present plant, not of a scenario, and it could be measured tonight by a probe. It is read the same way: on the reduced state the rack inlets are functions of the six free entries, and the inlet of rack 3 is the function $T_3(\mathbf{z})$, with its mean at the MAP state and its variance through its gradient by the delta method of Chapter 4. The answer is called a virtual sensor because it stands where an instrument would stand, and because it can be compared, on one scale, with the instrument that could be installed there.

Its band is computed from the same covariance $\boldsymbol{\Sigma}$ as every other output of the fit, and it is only as good as the rows that inform the entries T_3 depends on. Where those rows are several independent readings tied together by the physics, the band can be narrower than the error of the instrument it replaces. Where the only row behind an entry is a prior, the band is the prior's spread pushed through the physics, and the virtual sensor reports the prior rather than the plant. Which of the two holds for rack 3 is read off the variance budget, not guessed.

Principle 5.1. *A virtual sensor is the posterior of a derived quantity of the fitted state, with the variance* $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$ *and the checks of any other derived quantity. Its band measures what the rows behind its gradient say about the quantity, which may be much, or nothing beyond a prior.*

5.1.2 The two-unit hall (step 0)

The hall of Part I grows here. Its supply air T_c becomes the cold aisle, produced by two cooling coils from chilled water at $T_s = 14°\mathrm{C}$; its hall air T becomes the return air T_r; the units' airflow C_u is pinned and the coils' log-derates g_u are the unknowns; and three racks draw their inlet air from the cold aisle with recirculation fractions β_k, rack 1's inlet measured. The paths change with the plant. Part I's unknown path ran from the hall air to the supply air, which the discharge controller held at 18°C. Here the path from the cold aisle to the return is the air loop itself, and its rise is fixed by the pinned capacity rates $C_1 + C_2$; the unknown conductances are a new path, from the return air to the chilled water; and the pinned sink has moved from $T_c = 18°\mathrm{C}$ to $T_s = 14°\mathrm{C}$, so that the cold aisle now floats with the coils.

The physical relations. The racks put the IT load Q into the hall and the lighting and fans add the other heat Q_o. Unit u's coil passes heat to the chilled water at the rate

$$Q_u = G_u\,(T_r - T_s), \qquad G_u = \bar{G}_u\,\mathrm{e}^{g_u},$$

the coil heat, with $\bar{G}_u$ the rated conductance. In steady state the hall stores no heat, so the heat the racks and the other sources put in leaves through the two coils, the hall balance; the air crossing unit u gives up to its coil what the coil passes to the water, the air side of the unit, which sets its discharge temperature $T_{d,u}$; the two discharge streams mix in the cold aisle in proportion to their capacity rates; and the inlet of rack k is cold-aisle air with a share β_k of return air mixed in. Written as rows that vanish when they hold, each with its own tolerance, $\sigma_{p,\mathrm{h}}$, $\sigma_{p,\mathrm{a}}$ and $\sigma_{p,\mathrm{m}}$ in kW and $\sigma_{p,\mathrm{i}}$ in K, they are

$$\begin{aligned}
r_{\mathrm{h}} &= \big[(G_1 + G_2)(T_r - T_s) - Q - Q_o\big]/\sigma_{p,\mathrm{h}} && && \text{hall balance,}\\
r_{\mathrm{a},u} &= \big[C_u\,(T_r - T_{d,u}) - G_u\,(T_r - T_s)\big]/\sigma_{p,\mathrm{a}}, && u = 1,2 && \text{air side of unit } u,\\
r_{\mathrm{m}} &= \big[(C_1 + C_2)\,T_c - C_1 T_{d,1} - C_2 T_{d,2}\big]/\sigma_{p,\mathrm{m}} && && \text{cold-aisle mixing,}\\
r_{\mathrm{i},k} &= \big[T_k - T_c - \beta_k\,(T_r - T_c)\big]/\sigma_{p,\mathrm{i}}, && k = 1,2,3 && \text{rack inlet } k,
\end{aligned} \tag{5.1}$$

seven physical rows; the limit of exact rows takes the four tolerances to zero together, written $\sigma_p \to 0$ below. The pinned quantities are the chilled-water supply $T_s = 14.0°\mathrm{C}$, read by a precise supply sensor, the ratings $\bar{G}_1 = \bar{G}_2 = 7.0\,\mathrm{kW/K}$, the air capacity rates $C_1 = C_2 = 10.0\,\mathrm{kW/K}$ at the units' fixed fan speed, and the other heat $Q_o = 8\,\mathrm{kW}$. The free quantities of the graph form the full state of thirteen entries,

$$\mathbf{z}_{\mathrm{full}} = \big(Q,\ g_1,\ g_2,\ \beta_1,\ \beta_2,\ \beta_3 \,\big|\, T_r,\ T_{d,1},\ T_{d,2},\ T_c,\ T_1,\ T_2,\ T_3\big)^{\mathsf{T}}, \tag{5.2}$$

six free entries and seven temperatures that the seven rows fix.

Modeling choices at the graph. The hall is taken in steady state within the window: this is what makes every relation of equation (5.1) algebraic, so that one window's readings fix one state, and it fails in the minutes after a load step or a unit start, while the hall's air and equipment still store or release heat (Section 5.6.5). The heat path through each unit is one effective conductance between the return air and the chilled water, the hall of Part I's single path split in two: this keeps one entry per coil, and it fails if a coil's heat transfer depends on its own discharge temperature in a way one conductance cannot carry. The units' airflow is pinned and nothing controls the discharge temperature within the window: the chilled-water valves and the fans hold their positions, so the cold aisle T_c floats with the coils instead of being held at a set point. This is what lets the two discharge readings inform the log-derates, since a coil that passes less heat leaves its discharge warmer; an active controller that pinned $T_{d,u}$ by moving its valve would break it, because the discharge would then report the set point and the valve, not the coil (Section 5.6.5).

The graph. Figure 5.1 draws the hall's rows, not its airflow. The cold aisle on the left (blue) feeds the three racks; the return air in the middle (blue) receives the IT load Q from above, a free node (green), the racks' exhaust carried in the model by the hall balance rather than drawn as arrows; the return feeds the two units on the right (green, each carrying its free log-derate), which pass heat to the pinned chilled water at the far right (gray) and send their discharge air back to the cold aisle along the top and bottom of the figure. Each rack's inlet is drawn as the mixing row it is, a node fed by the cold aisle, with weight $1 - \beta_k$, and by the return air, with weight β_k; rack 3 is drawn in red, the place where the question is asked. The dashed arrows leave the graph for the five readings. Blue nodes are temperatures fixed by the rows (5.1); they carry no entry of the reduced state.

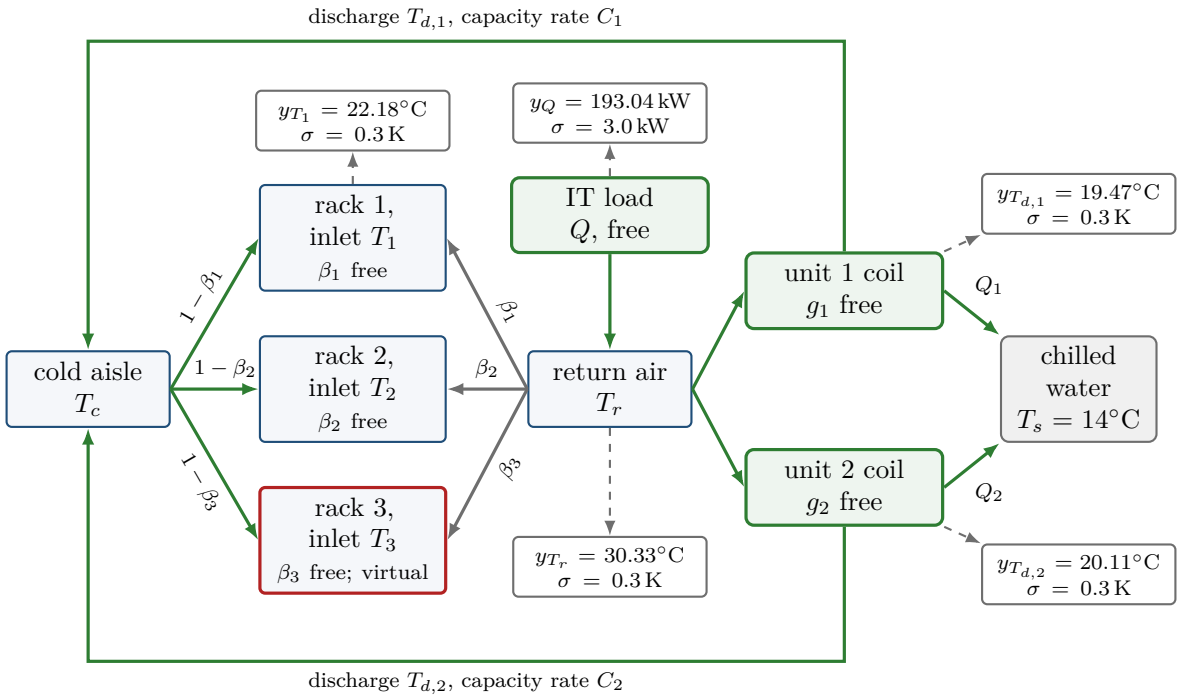


Figure 5.1: The rows of the two-unit hall in its 23:00 window. Free (green): the IT load Q and the coils' log-derates g_1, g_2; the recirculation fractions β_1, β_2, β_3 are free entries carried by the rack nodes. Blue: temperatures fixed by the physical rows (5.1), the cold aisle T_c, the return air T_r and the rack inlets T_k; rack 3 (red) is the virtual sensor. Pinned (gray): the chilled-water supply $T_s = 14°\mathrm{C}$; also pinned, not drawn as nodes, the ratings $\bar{G}_u = 7.0\,\mathrm{kW/K}$, the capacity rates $C_u = 10.0\,\mathrm{kW/K}$ and the other heat $Q_o = 8\,\mathrm{kW}$. Observed (dashed): the return air, the two discharge temperatures, the hall load meter on Q and rack 1's inlet. The priors on g_1, g_2 and the β_k are rows not drawn. The arrows into the racks are the mixing rows; the racks' exhaust to the return is carried by the hall balance, through Q.

The reduced state. The seven rows fix the seven temperatures at given free entries, and in the limit $\sigma_p \to 0$ of exact rows the fit runs on the reduced state of Chapter 4, the six free entries,

$$\mathbf{z} = (Q,\ g_1,\ g_2,\ \beta_1,\ \beta_2,\ \beta_3)^{\mathsf{T}}, \tag{5.3}$$

in kW and five dimensionless entries. Split as in equation (4.45), the dependent block $\mathbf{A}_{\mathrm{dep}}$ of the physical rows is triangular in the order of equation (5.2), with nonzero diagonal $(G_1 + G_2,\ C_1,\ C_2,\ C_1 + C_2,\ 1,\ 1,\ 1)$ over the tolerances: the hall balance fixes T_r, the air sides then fix $T_{d,1}$ and $T_{d,2}$, the mixing fixes T_c and the inlets fix the T_k. It is therefore invertible, the split

is valid, and the dependent temperatures follow in that order,

$$\begin{aligned} T_r &= T_s + \frac{Q + Q_o}{G_1 + G_2}, \qquad T_{d,u} = T_r - \frac{G_u\,(T_r - T_s)}{C_u}, \\ T_c &= \frac{C_1 T_{d,1} + C_2 T_{d,2}}{C_1 + C_2}, \qquad T_k = T_c + \beta_k\,(T_r - T_c), \end{aligned} \tag{5.4}$$

the point $\mathbf{z}_{\text{full}}(\mathbf{z})$ of the set on which the balances hold. The rows of $\mathbf{N}_{\text{b}} = -\mathbf{A}_{\text{dep}}^{-1}\mathbf{A}_{\text{free}}$ are the gradients of these seven temperatures with respect to the reduced state.

One combination of the rows is worth having in view. The two air-side rows summed, with the mixing row and the hall balance, give

$$(C_1 + C_2)\,(T_r - T_c) = Q + Q_o, \tag{5.5}$$

the air balance of the units taken together: the rise from the cold aisle to the return is set by the heat and the airflow alone, and not by the coils. It holds at every point of the set and is not a substitution; it is a property of the null space that the gradients below make visible.

5.1.3 The operator's questions (step 1)

The operator asks five questions of the 23:00 window of the two-unit hall, and step 1 writes each as a probability or an expectation over the reduced state, with the unknown on the left of the bar and the five readings and the priors on the right. Table 5.1 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 in the notation of Chapter 2, for a quantity $f(\mathbf{z})$ derived from the state, a level f^* of it, the posterior $p(\mathbf{z} \mid \mathbf{y})$ and a future reading y_{new}; the third its mathematical form for the hall; the fourth its category; the fifth the computation of Section 2.3.6 that answers it; and the last the subsection where it is answered. The first two are estimation, the inlet of rack 3 with its band and the split of the band's variance. The third and fourth are predictions under a threshold, tails of a derived quantity: rack 3 over the limit, and rack 3 warmer than the rack the operator can see. The fifth predicts a reading not yet taken, which Chapter 2 also files under prediction under a threshold. Which reading would narrow the band most asks nothing new of the present evidence; it is the next reading for these five questions, Section 5.6.

Table 5.1: The questions asked of the two-unit hall at 23:00. $\mathbf{y}$ is the five readings and the five priors; $T_k(\mathbf{z})$ the inlet of rack k by equation (5.4); $T_{\max} = 27°\text{C}$; y_{p} a probe reading at rack 3.

Operator question	General form	Mathematical form	Category	Computation	Answered in
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?	$\int p(y_{\text{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	$p(y_{\text{p}} \mid \mathbf{y})$, a 0.3 K probe at rack 3	prediction under a threshold	posterior	5.3.3

5.1.4 The evidence as rows (step 2)

Step 2 attaches the evidence of the window at 23:00 to the graph as rows on the reduced state. Five readings are shown in Figure 5.1: the return air $y_{T_r} = 30.33°\mathrm{C}$, the two discharge temperatures $y_{T_{d,1}} = 19.47°\mathrm{C}$ and $y_{T_{d,2}} = 20.11°\mathrm{C}$, each by a sensor of standard deviation 0.3 K; the hall load meter $y_Q = 193.04\,\mathrm{kW}$, of standard deviation 3.0 kW, which reads Q directly; and rack 1's inlet sensor $y_{T_1} = 22.18°\mathrm{C}$, of standard deviation 0.3 K. The readings are constructed, generated from a stated truth of the hall with the stated sensor errors. Five priors carry what was known before the window. The commissioning record puts each coil at its rating to within about 15 %, $g_u \sim \mathcal{N}(0,\ 0.15^2)$; the airflow survey gives the recirculation fractions $\beta_1 = 0.20$, $\beta_2 = 0.35$ and $\beta_3 = 0.50$, each with standard deviation 0.08, $\beta_k \sim \mathcal{N}(\beta_k^{\mathrm{s}},\ 0.08^2)$. The load Q has no prior row: the meter reads it.

Which row acts on which entry. On the reduced state a reading row sees the free entries through the row of $\mathbf{N}_{\mathrm{b}}$ of the temperature it reads, equation (4.47). By equation (5.4) the return air and the two discharge temperatures depend on (Q, g_1, g_2) only; rack 1's inlet on those three and β_1; the meter on Q alone. With $*$ for an entry that can be nonzero, the 10×6 Jacobian has the pattern

$$\mathbf{J} = \begin{bmatrix} * & * & * & 0 & 0 & 0 \\ * & * & * & 0 & 0 & 0 \\ * & * & * & 0 & 0 & 0 \\ * & 0 & 0 & 0 & 0 & 0 \\ * & * & * & * & 0 & 0 \\ 0 & * & 0 & 0 & 0 & 0 \\ 0 & 0 & * & 0 & 0 & 0 \\ 0 & 0 & 0 & * & 0 & 0 \\ 0 & 0 & 0 & 0 & * & 0 \\ 0 & 0 & 0 & 0 & 0 & * \end{bmatrix} \begin{array}{l} \text{return air } y_{T_r} \\ \text{discharge } y_{T_{d,1}} \\ \text{discharge } y_{T_{d,2}} \\ \text{load meter } y_Q \\ \text{rack-1 inlet } y_{T_1} \\ \text{prior on } g_1 \\ \text{prior on } g_2 \\ \text{prior on } \beta_1 \\ \text{prior on } \beta_2 \\ \text{prior on } \beta_3 \end{array}, \qquad (5.6)$$

$$\mathbf{g}_{T_3}^{\mathsf{T}} = (*,\ *,\ *,\ 0,\ 0,\ *),$$

with columns $Q, g_1, g_2, \beta_1, \beta_2, \beta_3$. The gradient of the question's quantity, on the second line, is not a row; it has the pattern of the inlet of rack 3 in equation (5.4). Two facts of the graph are read here before any number. No reading row has a nonzero in the column of β_2 or of β_3: each holds only its prior row, so the five readings have rank at most four on six entries, and neither fraction is observable from the readings. The gradient of T_3 has its last entry in exactly one of those two columns.

Modeling choices at the rows. Three choices make these rows. Each reading is taken as Gaussian, unbiased and independent of the others given the state, the form of the reading rows of Chapter 4: this makes each reading one normalized row and the objective a sum of squares, and it fails if a sensor carries a bias that the rows do not model, which is what the bad-data detection of Chapter 9 looks for. Each rack carries one recirculation fraction, constant over the window and over the height of its inlet, with an independent prior: this is what makes T_3 a mix of two temperatures with one entry of its own, and it fails if the inlet is stratified, cooler at the bottom than at the top, so that no single fraction describes it, or if the racks share a cause of recirculation that would tie their fractions together (Section 5.6.5). The airflow survey is taken as the only knowledge of rack 3's airflow, $\beta_3 \sim \mathcal{N}(0.50,\ 0.08^2)$: by the pattern (5.6) no reading sees β_3, so this prior row is the whole of what the fit knows of it, and the answer inherits whatever the survey got wrong; the survey's width and its center are varied in Section 5.4.3.

5.2 The posterior of the state (step 3)

All five questions of Table 5.1 are read from one posterior: that of the reduced state under the ten rows of step 2. Its inputs are the five readings and five priors of Section 5.1.4 and the pinned quantities of Section 5.1.2.

5.2.1 The rows on the reduced state

The rows of the reduced state are the evidence rows at the point $\mathbf{z}_{\text{full}}(\mathbf{z})$ of the set, equation (4.46), with the temperatures of equation (5.4); stacked in the order of the pattern (5.6), each divided by its standard deviation, they form the residual vector

$$\mathbf{r}(\mathbf{z}) = \begin{bmatrix} (T_r(\mathbf{z}) - y_{T_r})/0.3 \\ (T_{d,1}(\mathbf{z}) - y_{T_{d,1}})/0.3 \\ (T_{d,2}(\mathbf{z}) - y_{T_{d,2}})/0.3 \\ (Q - y_Q)/3.0 \\ (T_1(\mathbf{z}) - y_{T_1})/0.3 \\ g_1/0.15 \\ g_2/0.15 \\ (\beta_1 - 0.20)/0.08 \\ (\beta_2 - 0.35)/0.08 \\ (\beta_3 - 0.50)/0.08 \end{bmatrix}, \qquad \mathcal{L}(\mathbf{z}) = \tfrac{1}{2}\,\mathbf{r}(\mathbf{z})^{\mathsf{T}}\mathbf{r}(\mathbf{z}), \tag{5.7}$$

temperatures in °C over standard deviations in K, the load in kW over kW, every entry dimensionless. Ten rows act on six entries. The reduced Jacobian is the Jacobian of these rows with respect to the reduced state, $\mathbf{J} = \mathbf{J}_{\text{rows}}\mathbf{N}$ of equation (4.47): each reading row is its temperature's row of $\mathbf{N}_{\text{b}}$ divided by the reading's standard deviation, the meter row is $(1/3.0)$ in the column of Q, and each prior row is the reciprocal of its standard deviation in its own column. The rows are nonlinear in the state through e^{g_u} and through the product $\beta_k(T_r - T_c)$, so $\mathbf{J}$ changes with the state and is evaluated at each iterate.

5.2.2 The normal equations and the covariance

The MAP state $\hat{\mathbf{z}}$ minimizes the objective (5.7). Each iteration solves the normal equations of equation (4.29) on the reduced rows,

$$(\mathbf{J}^{\mathsf{T}}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^{\mathsf{T}}\mathbf{r}(\mathbf{z}), \qquad \mathbf{z}^{(k+1)} = \mathbf{z}^{(k)} + \Delta\mathbf{z}, \tag{5.8}$$

with $\mathbf{J}$ and $\mathbf{r}$ at $\mathbf{z}^{(k)}$ and the temperatures re-solved by equation (5.4) at every iterate. The iteration starts from the state the record suggests, the meter's load, the coils at their ratings and the survey's fractions, $\mathbf{z}^{(0)} = (y_Q,\ 0,\ 0,\ 0.20,\ 0.35,\ 0.50)^{\mathsf{T}}$. With every row normalized the information matrix is $\mathbf{\Lambda} = \mathbf{J}^{\mathsf{T}}\mathbf{J}$. At the minimum the Laplace covariance is

$$\mathbf{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1}\big|_{\hat{\mathbf{z}}}, \tag{5.9}$$

the Gauss–Newton form of equation (4.38), which drops the curvature term $\sum_k r_k\,\partial^2 r_k/\partial\mathbf{z}\,\partial\mathbf{z}^{\mathsf{T}}$ of the nonlinear rows; Section 5.4.1 sets the full Hessian beside it. The posterior is then approximately normal, $\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$. By equation (4.48) this $\mathbf{\Sigma}$ is the free-entry block of the full state's covariance in the limit of exact physical rows.

The block structure. The pattern (5.6) fixes the shape of $\mathbf{\Lambda}$ before any number. The columns of β_2 and β_3 each hold one prior row, so their columns of $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ have a single nonzero entry, $1/0.08^2$, on the diagonal, and

$$\mathbf{\Lambda} = \begin{bmatrix} \mathbf{\Lambda}_{\text{o}} & \mathbf{0} \\ \mathbf{0} & 0.08^{-2}\,\mathbf{I}_2 \end{bmatrix}, \qquad \mathbf{\Sigma} = \begin{bmatrix} \mathbf{\Lambda}_{\text{o}}^{-1} & \mathbf{0} \\ \mathbf{0} & 0.08^{2}\,\mathbf{I}_2 \end{bmatrix}, \tag{5.10}$$

with $\mathbf{\Lambda}_{\mathrm{o}}$ the 4×4 block of (Q, g_1, g_2, β_1), the entries the readings see. The objective itself separates: the terms in β_2 and β_3 are their prior terms alone, so their MAP values are the survey's and their posterior is the prior exactly, not only under the linearization. This is the observability rule of Chapter 4 read in matrix form: an entry whose posterior spread equals its prior spread is one the readings do not see.

5.2.3 The fit of the window

The iterations. Table 5.2 runs the normal equations (5.8) from $\mathbf{z}^{(0)}$. Each row is an iteration; the columns give the objective and the norm of its gradient $\mathbf{J}^{\mathsf{T}}\mathbf{r}$. The objective falls from 66.355 to 1.5125 in the first step and to 1.31849 in the second, and the norm of the gradient falls from 796 to 5.1×10^{-9} by the fifth.

Table 5.2: Gauss–Newton iterations of the two-unit hall, equation (5.8), on the reduced state, from $\mathbf{z}^{(0)}$ with the meter's $y_Q = 193.04\,\mathrm{kW}$.

Iteration k [–]	Objective $\mathcal{L}(\mathbf{z}^{(k)})$ [–]	$\lVert\mathbf{J}^{\mathsf{T}}\mathbf{r}\rVert$ [–]
0	66.355014	7.960×10^{2}
1	1.512462	4.876×10^{1}
2	1.318486	1.701×10^{-1}
3	1.318483	1.596×10^{-4}
4	1.318483	1.795×10^{-7}
5	1.318483	5.122×10^{-9}

The Jacobian at the MAP state. With the columns $Q, g_1, g_2, \beta_1, \beta_2, \beta_3$ and the rows in the order of equation (5.7),

$$\mathbf{J}(\hat{\mathbf{z}}) = \begin{bmatrix} 0.2640 & -27.5596 & -25.9396 & 0 & 0 & 0 \\ 0.0923 & -44.4298 & -9.0695 & 0 & 0 & 0 \\ 0.1024 & -10.6895 & -42.8098 & 0 & 0 & 0 \\ 0.3333 & 0 & 0 & 0 & 0 & 0 \\ 0.1341 & -27.5596 & -25.9396 & 33.7713 & 0 & 0 \\ 0 & 6.6667 & 0 & 0 & 0 & 0 \\ 0 & 0 & 6.6667 & 0 & 0 & 0 \\ 0 & 0 & 0 & 12.5 & 0 & 0 \\ 0 & 0 & 0 & 0 & 12.5 & 0 \\ 0 & 0 & 0 & 0 & 0 & 12.5 \end{bmatrix}, \tag{5.11}$$

the first column in kW^{-1} and the others dimensionless. The zeros sit where the pattern (5.6) places them. The rack-1 row repeats the return-air row in the log-derate columns, because by the air balance (5.5) a derate leaves the rise unchanged and moves the inlet as it moves the return air, and its entry 33.7713 in the column of β_1 is the rise divided by the sensor's 0.3 K. The five reading rows have rank 4 of 6.

The information and the covariance. The information matrix $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ has the block form of equation (5.10). Its observed block, in the order (Q, g_1, g_2, β_1), and the two diagonal entries of β_2 and β_3, each $156.25 = 0.08^{-2}$, are

$$\mathbf{\Lambda}_{\mathrm{o}} = \begin{bmatrix} 0.2178 & -16.1684 & -15.5486 & 4.5287 \\ -16.1684 & 3651.7832 & 2290.3444 & -930.7236 \\ -15.5486 & 2290.3444 & 3305.1062 & -876.0139 \\ 4.5287 & -930.7236 & -876.0139 & 1296.7476 \end{bmatrix}, \qquad \mathbf{\Lambda}_{66} = \mathbf{\Lambda}_{55} = 156.25,$$

with zeros everywhere else in the rows and columns of β_2 and β_3. Its inverse, the Laplace covariance (5.9), is

$$\boldsymbol{\Sigma} = \begin{bmatrix} 7.66947874 & 0.02055535 & 0.02271429 & 0.00331342 & 0 & 0 \\ 0.02055535 & 0.00056368 & -0.00025057 & 0.00016351 & 0 & 0 \\ 0.02271429 & -0.00025057 & 0.00062655 & 0.00016409 & 0 & 0 \\ 0.00331342 & 0.00016351 & 0.00016409 & 0.0009878 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0.0064 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.0064 \end{bmatrix},$$

in kW^2, kW and dimensionless units, and the correlations of the observed block (Q, g_1, g_2, β_1) are

$$\begin{bmatrix} 1 & 0.3126 & 0.3277 & 0.0381 \\ 0.3126 & 1 & -0.4216 & 0.2191 \\ 0.3277 & -0.4216 & 1 & 0.2086 \\ 0.0381 & 0.2191 & 0.2086 & 1 \end{bmatrix}.$$

The load is correlated positively with both log-derates and the two log-derates negatively with each other, and the rows fix those signs. The return-air reading fixes $T_r - T_s = (Q+Q_o)/(G_1+G_2)$, the load over the two conductances together: a larger load is explained equally well by larger conductances; and the sum $G_1 + G_2$ being fixed, a larger g_1 is explained by a smaller g_2. The lower right block of $\boldsymbol{\Sigma}$ is the survey's, untouched.

The MAP state. Table 5.3 gives the MAP state with its posterior spreads $\sqrt{\Sigma_{ii}}$ and the prior spread of each entry. Each row is an entry; the last column says which rows inform it. The load is read by the meter and the coil rows together. Rack 1's fraction moves from the survey's 0.20 to 0.220 and its spread falls from 0.08 to 0.031, the effect of its inlet sensor. Rack 2's and rack 3's fractions stay at the survey's values with the survey's spread, as equation (5.10) requires.

Table 5.3: The MAP state of the two-unit hall at 23:00 with the posterior and prior spreads of each entry.

Entry [unit]	$\hat{z}_i$ [entry's]	$\sqrt{\Sigma_{ii}}$ [entry's]	Prior sd [entry's]	Informed by
Q [kW]	194.62752	2.76938	none	meter and the temperature readings
g_1 [–]	−0.07355	0.02374	0.15	discharge, return and rack-1 readings; the meter, through Q
g_2 [–]	−0.13413	0.02503	0.15	discharge, return and rack-1 readings; the meter, through Q
β_1 [–]	0.22043	0.03143	0.08	rack-1 inlet sensor and survey
β_2 [–]	0.35000	0.08000	0.08	survey only
β_3 [–]	0.50000	0.08000	0.08	survey only

The temperatures at the MAP state. By equation (5.4), $T_r = 30.0498°\mathrm{C}$, $T_{d,1} = 19.6116°\mathrm{C}$, $T_{d,2} = 20.2252°\mathrm{C}$ and $T_c = 19.9184°\mathrm{C}$, a rise $T_r - T_c = 10.1314\,\mathrm{K}$, which is $(Q + Q_o)/(C_1 + C_2)$ of equation (5.5).

5.3 The answers (step 4)

Every answer is read from the pair $(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of Section 5.2.3, and that value is the one reported. A derived quantity f is carried by its gradient $\mathbf{g} = \partial f / \partial \mathbf{z}$ at $\hat{\mathbf{z}}$, with variance $\mathbf{g}^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g}$. Draws of the state pushed through the exact physics (5.4) stand beside each reported value, with the gap.

5.3.1 The inlet of rack 3 and the budget of its band

Operator questions. What is the inlet temperature at the rack with no sensor? What is that band made of? The first asks for the inlet of rack 3 at 23:00 and how well it is known; the second asks which entries of the state stand behind that band.

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

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

summarized by its mean and variance, equation (2.10), and by its central 95 % credible interval, equation (2.11); the second question asks for that variance split by entry.

Specialized form. The two questions are equations (2.9), (2.10) and (2.11) with $\mathbf{z}$ the reduced state (5.3), $f(\mathbf{z}) = T_3(\mathbf{z})$ of equation (5.4) and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of Section 5.2.2. One choice enters here: the answer is reported by the delta method on that normal, which gives it in closed form from $\hat{\mathbf{z}}$, $\boldsymbol{\Sigma}$ and one gradient, and which is accurate only if T_3 is nearly linear in the state over the posterior's width; the draws set beside it test that.

The gradient of T_3 is the row of $\mathbf{N}_\mathrm{b}$ that belongs to it, formed from the rows in their triangular order. The hall balance gives the gradient of the return air,

$$\frac{\partial T_r}{\partial \mathbf{z}} = \Big(\frac{1}{G_1 + G_2},\ -\frac{G_1\,(T_r - T_s)}{G_1 + G_2},\ -\frac{G_2\,(T_r - T_s)}{G_1 + G_2},\ 0,\ 0,\ 0\Big)^{\mathsf{T}},$$

since $\partial G_u / \partial g_u = G_u$. The air balance (5.5) gives the gradient of the rise, $\partial (T_r - T_c)/\partial \mathbf{z} = (1/(C_1{+}C_2),\ 0,\ 0,\ 0,\ 0,\ 0)^{\mathsf{T}}$, a function of Q alone. Writing the inlet as $T_3 = T_r - (1{-}\beta_3)(T_r{-}T_c)$ and differentiating,

$$\mathbf{g}_{T_3} = \frac{\partial T_3}{\partial \mathbf{z}} = \begin{bmatrix} \dfrac{1}{G_1 + G_2} - \dfrac{1 - \beta_3}{C_1 + C_2} \\ -\dfrac{G_1\,(T_r - T_s)}{G_1 + G_2} \\ -\dfrac{G_2\,(T_r - T_s)}{G_1 + G_2} \\ 0 \\ 0 \\ T_r - T_c \end{bmatrix} \begin{matrix} [\mathrm{K/kW}] \\ [\mathrm{K}] \\ [\mathrm{K}] \\ \\ \\ [\mathrm{K}] \end{matrix}, \tag{5.12}$$

evaluated at $\hat{\mathbf{z}}$. The last entry, $\partial T_3 / \partial \beta_3 = T_r - T_c$, is the rise from the cold aisle to the return: a change of 0.01 in the fraction moves the inlet by one hundredth of the rise, and by equation (5.5) the rise is fixed by the load and the airflow. The two log-derate entries equal those of $\partial T_r / \partial \mathbf{z}$: a derate warms the return and the cold aisle by the same amount, because it leaves the rise unchanged, so it moves every rack inlet as it moves the return air. The load enters twice, through the return air and through the rise, with opposite signs. The gradient of any other rack k is the same vector with β_k in the first entry and $T_r - T_c$ moved to the column of β_k.

By the delta method, equation (4.53), the virtual reading and its band are

$$T_3 \mid \mathbf{y} \approx \mathcal{N}\big(T_3(\hat{\mathbf{z}}),\ \mathbf{g}_{T_3}^{\mathsf{T}} \boldsymbol{\Sigma}\, \mathbf{g}_{T_3}\big), \qquad \big[T_3(\hat{\mathbf{z}}) \mp 1.960 \sqrt{\mathbf{g}_{T_3}^{\mathsf{T}} \boldsymbol{\Sigma}\, \mathbf{g}_{T_3}}\,\big], \tag{5.13}$$

the central 95 % interval of equation (4.56). The budget splits the variance by entry, as equation (4.54) does for two entries: one diagonal term per entry and the cross terms,

$$\mathbf{g}_{T_3}^{\mathsf{T}}\,\boldsymbol{\Sigma}\,\mathbf{g}_{T_3} = \sum_i [\mathbf{g}_{T_3}]_i^2\,\Sigma_{ii} + \sum_{i\neq j} [\mathbf{g}_{T_3}]_i\,[\mathbf{g}_{T_3}]_j\,\Sigma_{ij}. \tag{5.14}$$

By the block structure (5.10) $\boldsymbol{\Sigma}$ has no covariance between β_3 and any other entry and Σ_{66} is the survey's 0.08^2. The β_3 term is then $(T_r - T_c)^2 \times 0.08^2$, the survey's spread times the rise, and every cross term lies among Q, g_1 and g_2.

The sign of those cross terms is fixed by the correlations of Section 5.2.3. The gradient (5.12) has the signs $(+,-,-)$ on (Q, g_1, g_2). A cross term $2[\mathbf{g}]_i[\mathbf{g}]_j\Sigma_{ij}$ in Q and g_u is then a positive times a negative times a positive covariance, and the term in g_1 and g_2 is a negative times a negative times a negative covariance: every cross term is negative, and the readings know T_3's share in (Q, g_1, g_2) better than the three diagonal terms suggest. The variance that would remain if β_3 were known exactly is the budget without that term, the quadratic form on the other five entries with the row and column of β_3 removed.

Building blocks. The MAP state $\hat{\mathbf{z}}$ and the covariance $\boldsymbol{\Sigma}$ (Section 5.2.3); the temperatures $T_r(\hat{\mathbf{z}})$ and $T_c(\hat{\mathbf{z}})$ of the same section; the gradient (5.12); the quadratic form and its terms (5.14).

The virtual reading at 23:00. The gradient (5.12) at $\hat{\mathbf{z}}$ is

$$\mathbf{g}_{T_3} = (0.05421\ \text{K/kW},\ -8.26789\ \text{K},\ -7.78189\ \text{K},\ 0,\ 0,\ 10.13138\ \text{K})^{\mathsf{T}}.$$

Figure 5.2 draws it: one bar per entry of the state, the horizontal axis the change of the inlet of rack 3 per unit of the entry, in K for the five dimensionless entries and in K per kW for the load. The fraction β_3 has the longest bar, the rise; the two log-derates have bars of nearly the same length and a sign opposite to β_3's; the load's bar is short because it is per kilowatt; the fractions of racks 1 and 2 have none.

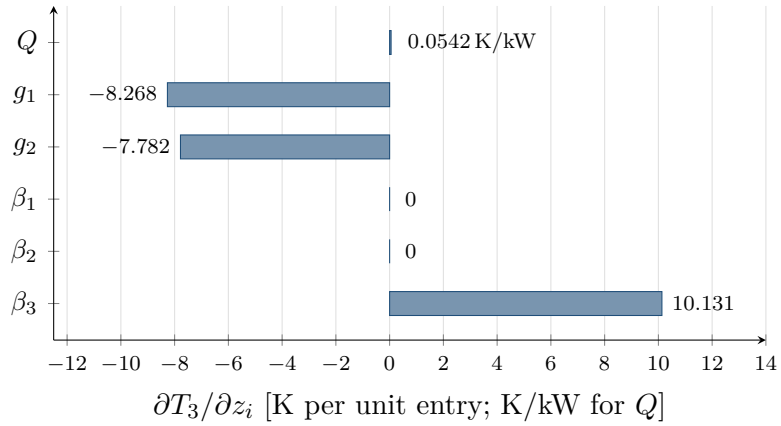


Figure 5.2: The gradient of the inlet of rack 3 at the MAP state, equation (5.12): the change of T_3 per unit change of each entry of the reduced state. The fraction's entry is the rise $T_r - T_c$; the log-derates' entries are those of the return air; the fractions of the other racks do not enter.

By equation (5.13), with $T_3(\hat{\mathbf{z}}) = T_c + \beta_3(T_r - T_c) = 19.9184 + 0.50 \times 10.1314 = 24.9841$°C and the quadratic form of this gradient and $\boldsymbol{\Sigma}$ equal to $0.686105\,\text{K}^2$,

$$T_3 \mid \mathbf{y} \approx \mathcal{N}(24.9841°\text{C},\ 0.8283^2), \qquad T_3 \in [23.361,\ 26.608]°\text{C with 95 \% probability.} \tag{5.15}$$

Figure 5.3 draws it: the horizontal axis is the inlet temperature of rack 3, the vertical axis its posterior density; the green area is the 95 % band and the red area the probability beyond the

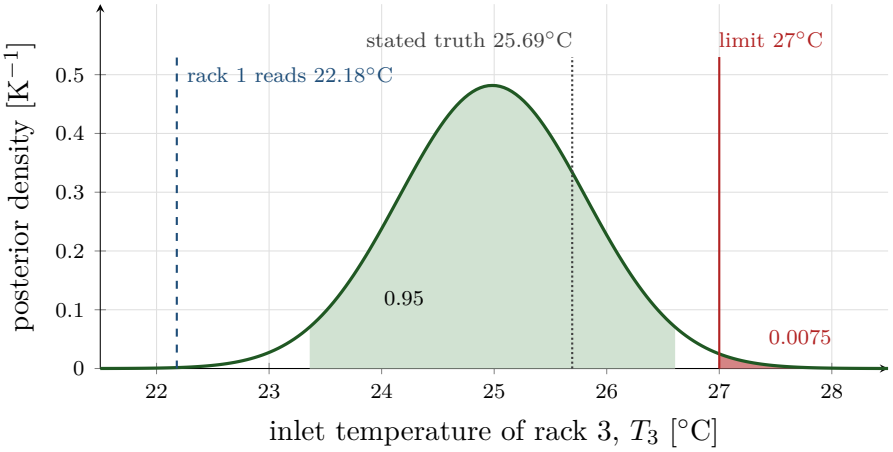


Figure 5.3: The virtual sensor at rack 3: the normal of equation (5.15). Green: the central 95 % band. Red: the probability beyond 27°C. Dashed: rack 1's reading. Dotted: the stated truth of the record, a check marker (Section 5.4.4).

27°C limit; the dashed blue line is rack 1's reading, the dashboard's number, and the dotted line the stated truth, a check marker for Section 5.4.4. Rack 1's reading lies below the whole band.

Beside the reported value stand 100 000 draws of the state from $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$, each pushed through the exact equations (5.4). They give a mean of 24.9815°C, a standard deviation of 0.8252 K and a 95 % band [23.366, 26.602]°C: the mean 0.0026 K lower, the standard deviation 0.0031 K lower, and the band's ends 0.0055 K higher and 0.0060 K lower than the delta method's.

The budget of the band. By equation (5.14), the variance 0.686105 K^2 splits as Table 5.4 shows: each row is a term, the middle column its value and the right column its share of the total. Figure 5.4 draws the shares as bars, one per entry and one for the cross terms, the horizontal axis in percent of the variance.

Table 5.4: The variance budget of the inlet of rack 3, equation (5.14), at the MAP state.

Term	Value [K^2]	Share [%]
Q, $[\mathbf{g}]_1^2\,\Sigma_{11}$	0.022537	3.3
g_1, $[\mathbf{g}]_2^2\,\Sigma_{22}$	0.038532	5.6
g_2, $[\mathbf{g}]_3^2\,\Sigma_{33}$	0.037943	5.5
β_1, β_2	0	0.0
β_3, $(T_r - T_c)^2 \times 0.08^2$	0.656927	95.7
cross terms among Q, g_1, g_2	−0.069833	−10.2
total $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$	0.686105	100

The survey's fraction for rack 3 carries 95.7 % of the variance. The load and the two log-derates carry the other diagonal terms, 3.3, 5.6 and 5.5 %, and the cross terms among them take back 10.2 %, all three negative as the signs of the gradient and of the correlations require: $-0.0184\,\mathrm{K}^2$ for Q with g_1, $-0.0192\,\mathrm{K}^2$ for Q with g_2 and $-0.0322\,\mathrm{K}^2$ for g_1 with g_2. With β_3 known exactly, its row and column removed, the standard deviation of the inlet would be 0.1708 K, against 0.8283 K with it. The same quadratic form on the gradient of the cold aisle, $\mathbf{g}_{T_c} = (0.02921\ \mathrm{K/kW},\ -8.26789\ \mathrm{K},\ -7.78189\ \mathrm{K},\ 0,\ 0,\ 0)^\mathsf{T}$, gives 0.1747 K, and on the gradient of the rise 0.1385 K: what the readings fix, they fix tightly.

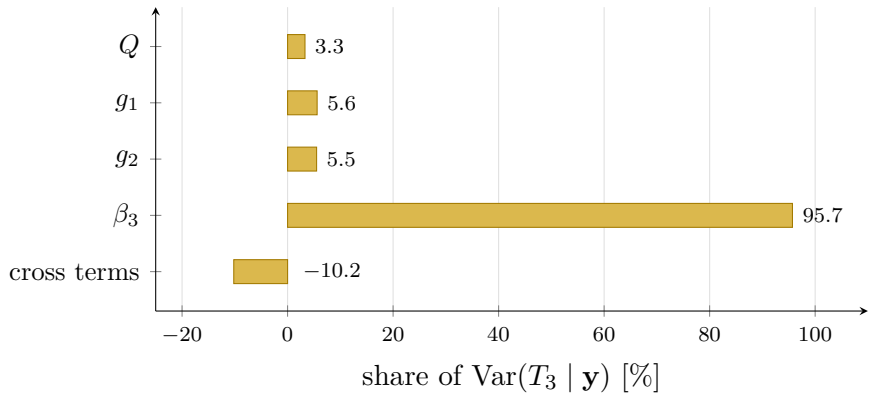


Figure 5.4: The variance budget of the inlet of rack 3 (Table 5.4): the diagonal terms $[\mathbf{g}]_i^2\Sigma_{ii}$ and the cross terms, in percent of $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$. The fractions of racks 1 and 2 contribute nothing.

5.3.2 The limit and the rack that can be seen

Operator questions. Is rack 3 over 27°C right now? Is rack 3 hotter than rack 1, the rack I can see? The first asks for the probability that the inlet of rack 3 already exceeds the limit at 23:00; the second, whether the one measured rack is a safe stand-in for the hottest.

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

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

Specialized form. The two questions are equation (2.12) with $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal of Section 5.2.2; for the limit, $f(\mathbf{z}) = T_3(\mathbf{z})$ and $f^* = T_{\max} = 27°\mathrm{C}$, and for the comparison, $f(\mathbf{z}) = T_3(\mathbf{z}) - T_1(\mathbf{z})$ and $f^* = 0$. The assurance level κ of that equation is not set here: no acceptance rule is in question at 23:00, and the probability itself is reported. The tail is that of the delta method's normal, and a skew of the true distribution would move a tail more than it moves the mean.

For the limit, the event is the exceedance, the complement of the requirement $T_3 \leq T_{\max}$ of the envelope. With the band (5.13), by equation (4.57),

$$\Pr(T_3 > T_{\max} \mid \mathbf{y}) = 1 - \Phi\left(\frac{T_{\max} - T_3(\hat{\mathbf{z}})}{\sqrt{\mathbf{g}_{T_3}^\mathsf{T}\boldsymbol{\Sigma}\,\mathbf{g}_{T_3}}}\right) = \Phi\left(-\frac{T_{\max} - T_3(\hat{\mathbf{z}})}{\sqrt{\mathbf{g}_{T_3}^\mathsf{T}\boldsymbol{\Sigma}\,\mathbf{g}_{T_3}}}\right). \tag{5.16}$$

For the comparison, the difference is a derived quantity of its own, with the gradient $\mathbf{g}_\Delta = \mathbf{g}_{T_3} - \mathbf{g}_{T_1}$; by equation (5.12) and the gradient of rack 1 read off it,

$$\mathbf{g}_\Delta = \Big(\frac{\beta_3 - \beta_1}{C_1 + C_2},\ 0,\ 0,\ -(T_r - T_c),\ 0,\ T_r - T_c\Big)^\mathsf{T}, \qquad \Pr(T_3 > T_1 \mid \mathbf{y}) = \Phi\left(\frac{T_3(\hat{\mathbf{z}}) - T_1(\hat{\mathbf{z}})}{\sqrt{\mathbf{g}_\Delta^\mathsf{T}\boldsymbol{\Sigma}\,\mathbf{g}_\Delta}}\right). \tag{5.17}$$

The log-derates drop out of the difference: they move both inlets alike. What remains are the two fractions, through the rise, and a small term in the load. How far rack 1 stands in for rack 3 is the correlation of the two inlets, $\mathrm{Cov}(T_3, T_1)/\sqrt{\mathrm{Var}(T_3)\,\mathrm{Var}(T_1)}$, with $\mathrm{Cov}(T_3, T_1) = \mathbf{g}_{T_3}^\mathsf{T}\boldsymbol{\Sigma}\,\mathbf{g}_{T_1}$. The entry of $\mathbf{g}_{T_3}$ in β_3 meets nothing in $\boldsymbol{\Sigma}$ by the block structure (5.10), so only the part of $\mathbf{g}_{T_3}$ in Q, g_1, g_2 enters the covariance, and by the Cauchy–Schwarz inequality the correlation is at most $\mathrm{sd}(T_3)$ with β_3 known over $\mathrm{sd}(T_3)$. That bound is why the correlation must be small when the fraction dominates the variance of T_3.

Building blocks. The virtual reading and its variance (5.15); $T_1(\hat{\mathbf{z}})$ by equation (5.4); the gradient $\mathbf{g}_\Delta$ of equation (5.17); $\mathbf{g}_{T_1}$, which is 0.3 times the rack-1 row of $\mathbf{J}(\hat{\mathbf{z}})$ (5.11); $\boldsymbol{\Sigma}$ of Section 5.2.3; and the normal distribution function Φ.

The probability over the limit. By equation (5.16), the limit lies $(27 - 24.9841)/0.8283 = 2.4337$ standard deviations above the virtual reading, and

$$\Pr(T_3 > 27^\circ\mathrm{C} \mid \mathbf{y}) = \Phi(-2.4337) = 7.472 \times 10^{-3},$$

about 0.75 %, the red area of Figure 5.3. The draws of Section 5.3.1 put a fraction 0.00711 of the states above 27°C, 0.00036 lower.

Rack 3 against rack 1. The fitted inlet of rack 1 is $T_1(\hat{\mathbf{z}}) = 22.1516^\circ$C with standard deviation 0.2821 K, close to its reading. The gradient of the difference, equation (5.17), is $\mathbf{g}_\Delta = (0.01398\ \mathrm{K/kW},\ 0,\ 0,\ -10.13138\ \mathrm{K},\ 0,\ 10.13138\ \mathrm{K})^\mathsf{T}$, and its quadratic form with $\boldsymbol{\Sigma}$ has the terms in Q, β_1 and β_3 and the cross term of Q with β_1,

$$\mathbf{g}_\Delta^\mathsf{T} \boldsymbol{\Sigma}\, \mathbf{g}_\Delta = 0.00150 + 0.10139 + 0.65693 - 0.00094 = 0.75888\,\mathrm{K}^2 = 0.8711^2,$$

so that

$$T_3 - T_1 \mid \mathbf{y} \approx \mathcal{N}(2.8324\,\mathrm{K},\ 0.8711^2), \qquad [1.125,\ 4.540]\,\mathrm{K} \text{ at } 95\,\%,$$

$$\Pr(T_3 > T_1 \mid \mathbf{y}) = \Phi\left(\frac{2.8324}{0.8711}\right) = 0.99943.$$

The draws give a mean of 2.8345 K, a standard deviation of 0.8678 K and a probability of 0.99941: gaps of +0.0021, −0.0033 K and −0.00002.

The correlation of the two inlets is 0.0145, a near-cancellation: of the covariance $0.00340\,\mathrm{K}^2$, the shared entries give $+0.02821\,\mathrm{K}^2$ and β_1, which $\boldsymbol{\Sigma}$ connects to Q, g_1 and g_2, gives $-0.02481\,\mathrm{K}^2$. The shared entries alone would give 0.1207, and the bound is $0.1708/0.8283 = 0.2062$. The rack-1 sensor is what connects them: the β_1 term of $\mathrm{Var}(T_1)$, $0.10139\,\mathrm{K}^2$, exceeds the whole variance, $0.07957\,\mathrm{K}^2$.

5.3.3 The reading of a probe

Operator question. If I send someone with a probe, what will it read? The technician would hold a handheld probe of standard deviation 0.3 K at rack 3's inlet; the question is what that reading will be, before it is taken.

General form. A reading not yet taken is predicted by its predictive distribution, equation (2.13),

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

Specialized form. The question is equation (2.13) with y_{new} the probe's reading y_p, with $p(y_{\mathrm{new}} \mid \mathbf{z}) = \mathcal{N}(T_3(\mathbf{z}),\ 0.3^2)$ and with $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal of Section 5.2.2. One choice enters here: the probe reads the inlet temperature that T_3 stands for, unbiased with a Gaussian error of 0.3 K, independent of the hall's sensors. This is what makes the predictive a convolution of two normals, and it fails if the probe is held where the inlet differs from the one fraction of the model.

The probe reading is $y_\mathrm{p} = T_3(\mathbf{z}) + \epsilon_\mathrm{p}$, with $\epsilon_\mathrm{p} \sim \mathcal{N}(0,\ 0.3^2)$ independent of the state. Under the band (5.13) the sum of two independent normals is normal, with the variances added, as in equation (4.64):

$$y_\mathrm{p} \mid \mathbf{y} \approx \mathcal{N}(T_3(\hat{\mathbf{z}}),\ \mathbf{g}_{T_3}^\mathsf{T} \boldsymbol{\Sigma}\, \mathbf{g}_{T_3} + 0.3^2). \tag{5.18}$$

The predictive is wider than the virtual sensor's band by the probe's own error, and it is never narrower than 0.3 K, however well the state is known.

Building blocks. The virtual reading and its variance (5.15), and the probe's 0.3 K.

The probe predictive. By equation (5.18), $\sqrt{0.686105 + 0.09} = 0.8810$ K, and a probe at rack 3 will read from

$$y_{\mathrm{p}} \mid \mathbf{y} \approx \mathcal{N}\big(24.9841^{\circ}\mathrm{C},\ 0.8810^2\big), \qquad y_{\mathrm{p}} \in [23.257,\ 26.711]^{\circ}\mathrm{C} \text{ with } 95\,\% \text{ probability.}$$

The draws, each with its own probe error, give a mean of 24.9804°C, a standard deviation of 0.8798 K and a 95 % range [23.256, 26.703]°C: gaps of −0.0037 and −0.0012 K, and of −0.001 and −0.008 K at the ends.

5.4 The checks (step 5)

The answers of Section 5.3 rest on approximations and inputs that can be tested on this hall, against the calculation each stands for.

5.4.1 The limit of exact rows and the curvature term

The fit runs on the reduced state, which equation (4.48) makes the limit $\sigma_p \to 0$ of the full state. The check fits this hall's own full state of thirteen entries, equation (5.2), with the seven physical rows (5.1) at a finite tolerance beside the readings and priors, and lets the tolerance shrink. Table 5.5 gives the result: each row is a tolerance, the columns the inlet of rack 3 and the load with their spreads, and the last row the reduced state. At 0.1 the spread of T_3 is 0.8344 K, wider by the discrepancy the rows are allowed; the full state reproduces the reduced state to the third decimal at 0.01 and to the fourth at 0.001.

Table 5.5: The limit of exact physical rows on the full state of the two-unit hall, equation (5.2): the inlet of rack 3 and the load at three tolerances, against the reduced state. Each fit took five iterations.

Tolerance σ_p [kW or K]	T_3 [°C]	sd(T_3) [K]	Q [kW]	sd(Q) [kW]
0.1	24.9840	0.8344	194.6270	2.7695
0.01	24.9841	0.8284	194.6275	2.7694
0.001	24.9841	0.8283	194.6275	2.7694
reduced state	24.9841	0.8283	194.6275	2.7694

The covariance (5.9) is the Gauss–Newton form, without the curvature term of the nonlinear rows. The full Hessian of the objective at $\hat{\mathbf{z}}$, curvature term included, gives a Newton covariance with sd(T_3) = 0.8283 K, unchanged, and sd(Q) = 2.7675 kW against 2.7694 kW: the term moves the load's spread in the fourth digit and the answer not at all.

5.4.2 The delta method against draws through the physics

The delta method linearizes T_3 about $\hat{\mathbf{z}}$; the draws of Section 5.3.1 keep the nonlinearity of T_3 in the log-derates and the product $\beta_3(T_r - T_c)$. The second-order correction to the mean, $\frac{1}{2}\,\mathrm{tr}(\nabla^2 T_3\,\boldsymbol{\Sigma})$, is −0.00372 K: the inlet is linear in β_3, and its cross derivative with the load meets no covariance in $\boldsymbol{\Sigma}$, so the correction comes from the curvature in Q, g_1 and g_2 alone, which the meter and the temperature readings hold tight. The draws carry Monte Carlo standard

errors of 0.0026 K on the mean, 0.0018 K on the standard deviation and 0.00027 on the tail. The correction is about 1.4 standard errors of the mean, so not resolved at this count, and the signs of the draws' gaps carry no information at this count. The check tests the delta method, not the Laplace posterior: both routes draw from the same normal.

5.4.3 Sensitivity to the airflow survey

The survey is the only knowledge of β_3, so the answer is tested against it. Figure 5.5 draws the virtual sensor under three survey widths, the fit repeated with each: the horizontal axis is the inlet of rack 3, the vertical axis its density, each curve one width of the β priors, and the shaded areas the probability beyond the limit, which grows from invisible to visible as the survey widens. The center barely moves, from 25.0001 to 24.9841 to 24.9793°C; the standard deviation follows the survey's width almost in proportion, 0.4376, 0.8283 and 1.2280 K for 0.04, 0.08 and 0.12; and the tail moves by four orders of magnitude, 2.4×10^{-6}, 0.0075 and 0.050. Moving the survey's center for rack 3 by its own standard deviation, to 0.42 or to 0.58, moves the virtual reading to 24.17 or to 25.79°C, symmetrically about 24.98°C and by the rise times 0.08 (Exercise 5.2).

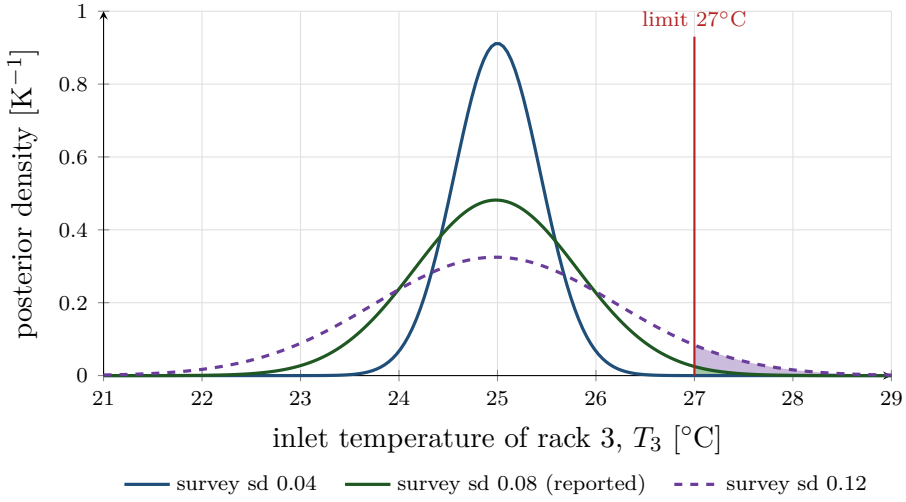


Figure 5.5: The virtual sensor at rack 3 under three widths of the airflow survey's priors, each the normal of the fit repeated with that width. Shaded: the probability beyond 27°C for the reported width (green) and the widest (purple); for the narrowest it is too small to see.

5.4.4 The answers against the stated truth

The record of this hall is generated from a stated truth (Section 5.1.4), and the last check compares the answers with it. Table 5.6 collects the checks: each row is a route or a variation, the columns the value, its standard deviation and the probability, each with its gap from the reported value, the route minus the delta method. The true inlet of rack 3 is 25.6933°C. The virtual reading is 0.7092 K low, −0.86 of its standard deviation, inside its band: the survey's 0.50 for rack 3 is below the true 0.57, and the fit, which no reading of β_3 corrects, carries that error into the answer by the lever of Figure 5.2 times the survey's error, $0.07 \times 10.1314 = 0.709$ K; the fit's other errors cancel here, its cold aisle +0.1531 K from the truth and 0.57 times its rise error −0.1531 K. The dashboard's answer, rack 1's reading, is 3.5133 K below the truth and $22.18 - 24.9841 = -2.8041$ K from the fit, below the band; the design sheet's 23.6571°C, 1.3270 K below the fit, lies inside it. Rack 2, also unmeasured, is read from the same fit at 23.4644°C with standard deviation 0.8280 K; its true inlet, 22.9893°C, lies inside that band too.

Table 5.6: The checks of the answers and their gaps from the reported values. The draws use 100 000 states from the Laplace posterior through the exact physics (5.4); the survey variations repeat the fit. The objective and the residual in the last row are tested against their noise in Chapter 9.

Route or variation	Value [°C or K]	Gap [K]	sd [K]	Gap [K]	Probability [–]	Gap [–]
Inlet of rack 3, T_3; probability $\Pr(T_3 > 27°\mathrm{C})$						
Delta method, reported	24.9841	–	0.8283	–	0.00747	–
Draws through the physics	24.9815	−0.0026	0.8252	−0.0031	0.00711	−0.00036
Full state, $\sigma_p = 0.001$	24.9841	0.0000	0.8283	0.0000	–	–
Newton covariance	–	–	0.8283	0.0000	–	–
Survey sd 0.04	25.0001	0.0160	0.4376	−0.3908	2.4×10^{-6}	−0.00747
Survey sd 0.12	24.9793	−0.0048	1.2280	0.3997	0.04993	0.04246
Survey center for rack 3 at 0.42	24.1736	−0.8105	–	–	–	–
Survey center for rack 3 at 0.58	25.7946	0.8105	–	–	–	–
Stated truth	25.6933	0.7092	–	–	–	–
Design sheet	23.6571	−1.3270	–	–	–	–
Dashboard (rack 1's reading)	22.18	−2.8041	–	–	–	–
Rack 3 against rack 1, $T_3 - T_1$; probability $\Pr(T_3 > T_1)$						
Delta method, reported	2.8324	–	0.8711	–	0.99943	–
Draws through the physics	2.8345	0.0021	0.8678	−0.0033	0.99941	−0.00002
Probe at rack 3, y_p						
Delta method, reported	24.9841	–	0.8810	–	–	–
Draws through the physics	24.9804	−0.0037	0.8798	−0.0012	–	–
Fit of the rows	$2\mathcal{L}(\hat{\mathbf{z}}) = 2.6370$ on $10 - 6 = 4$; largest normalized residual −0.9341, return air					

5.5 The report (step 6)

The dashboard's number does not describe rack 3. Rack 1 reads 22.2°C; rack 3 is at 25.0°C, with a 95 % band of 23.4 to 26.6°C and a 0.75 % chance that it is already over 27°C, and it is warmer than rack 1 with probability 0.9994. The band is the airflow survey's, not the sensors': 95.7 % of its variance is the survey's uncertainty about rack 3's recirculation.

5.5.1 The answers and their bands

Table 5.7 gives the answers in the order of Table 5.1. Each row is a question; the second column gives the answer with the draws beside it, and the third its band: the central 95 % interval for a temperature or a difference, for a probability its range over the three survey widths of Section 5.4.3, and for the budget the standard deviation that would remain.

5.5.2 The entry in the shift log

The operator's entry for 23:00 reads the virtual sensor, not the dashboard tile: rack 3's inlet is about 25.0°C, between 23.4 and 26.6°C, two degrees inside the limit rather than five, and not over it now except with a small probability that depends on how far the airflow survey can be trusted. The entry also records what the band rests on. The temperature readings and the meter fix the cold aisle and the rise to under two tenths of a kelvin (0.175 and 0.139 K); what they cannot fix is how much return air rack 3 draws, and that number comes from a survey taken at another time. A wider survey uncertainty of 0.12 would put the chance of an exceedance at 5 %; no decision level is set here, so the entry records the probability and that range. Rack 1's sensor is a good sensor of rack 1, and says almost nothing about rack 3 that the rest of the hall does not already say.

Table 5.7: The report for the two-unit hall at 23:00. Answers from the Laplace posterior by the delta method; bands are 95 % intervals or the range over survey widths 0.04 to 0.12.

Question	Answer	Band
What is the inlet temperature at the rack with no sensor?	24.98°C, sd 0.83 K (draws: 24.98°C, sd 0.83 K)	23.36 to 26.61°C, central 95 %
What is that band made of?	95.7 % the survey's fraction for rack 3	sd 0.17 K with that fraction known
Is rack 3 over 27°C right now?	probability 0.0075 (draws: 0.0071)	2.4×10^{-6} to 0.050 over the survey widths
Is rack 3 hotter than rack 1?	probability 0.9994; by 2.83 K, sd 0.87 K (draws: 0.9994; 2.83 K, sd 0.87 K)	1.125 to 4.540 K, central 95 %
What will a probe at rack 3 read?	24.98°C, sd 0.88 K (draws: 24.98°C, sd 0.88 K)	23.26 to 26.71°C, central 95 %

Common mistake

Reading the one measured rack as the hall's inlet temperature. The hottest rack must be read from the fit.

5.6 The next reading (step 7)

The report leaves a band almost all of whose variance is one prior. A further reading is one more row of the residual vector, and what it would remove is known before its value is.

5.6.1 The preposterior of a candidate reading

A candidate reading a is a row with Jacobian line $\mathbf{J}_a$, a 1×6 row in the columns of the state. Its preposterior covariance is equation (4.65), and for one row it takes the rank-one form of equation (4.66),

$$\boldsymbol{\Sigma}_a = \left(\boldsymbol{\Sigma}^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a\right)^{-1} = \boldsymbol{\Sigma} - \frac{\boldsymbol{\Sigma}\mathbf{J}_a^{\mathsf{T}}\,\mathbf{J}_a\boldsymbol{\Sigma}}{1 + \mathbf{J}_a\boldsymbol{\Sigma}\mathbf{J}_a^{\mathsf{T}}}, \qquad \mathrm{sd}_a(T_3) = \sqrt{\mathbf{g}_{T_3}^{\mathsf{T}}\boldsymbol{\Sigma}_a\,\mathbf{g}_{T_3}}, \tag{5.19}$$

and the width it removes is $1 - \mathrm{sd}_a(T_3)/\mathrm{sd}(T_3)$. Five kinds of candidate are open to the shift. A probe at rack 3 of standard deviation σ_a reads T_3 itself, $\mathbf{J}_a = \mathbf{g}_{T_3}^{\mathsf{T}}/\sigma_a$; sandwiched between $\mathbf{g}_{T_3}^{\mathsf{T}}$ and $\mathbf{g}_{T_3}$, the rank-one form gives $1/\mathrm{sd}_a^2 = 1/\mathrm{sd}^2 + 1/\sigma_a^2$, the virtual sensor and the probe adding their precisions like two readings of one quantity. Since (Q, g_1, g_2) are already fixed by the other readings, the probe's row acts mainly on β_3: it is a one-point re-survey of rack 3. A permanent inlet sensor at rack 3 is the same row, added in every window. An airflow re-survey of rack 3 with standard deviation 0.03 is a new prior row on β_3, $\mathbf{J}_a = \mathbf{e}_6^{\mathsf{T}}/0.03$, which narrows Σ_{66} from 0.08^2 to $(0.08^{-2} + 0.03^{-2})^{-1}$ and leaves the other entries alone by the block structure (5.10). A second sensor at rack 1 reads T_1, $\mathbf{J}_a = \mathbf{g}_{T_1}^{\mathsf{T}}/0.3$, and removes from $\mathrm{Var}(T_3)$ only $(\mathbf{g}_{T_3}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_{T_1})^2/(0.3^2 + \mathbf{g}_{T_1}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_{T_1})$, a quantity proportional to the square of the correlation of the two inlets. A second return-air sensor repeats the return-air row of $\mathbf{J}$ and acts on (Q, g_1, g_2) only, the entries that carry the smallest part of the budget.

5.6.2 The candidates

Table 5.8 ranks the six candidates. Each row is a reading the operator can request; the columns give its standard deviation, the standard deviation of the inlet of rack 3 after it by

equation (5.19), the share of the current 0.8283 K that it removes, and, for the rows that touch it, the standard deviation of β_3 after it. Figure 5.6 draws the width removed as bars, ranked, the horizontal axis in percent.

Table 5.8: Candidate readings for the inlet of rack 3, from the posterior at 23:00; the width removed is taken from the reported $\mathrm{sd}(T_3 \mid \mathbf{y})$.

Candidate	Its sd [K or –]	$\mathrm{sd}_a(T_3)$ [K]	Width removed [%]	$\mathrm{sd}_a(\beta_3)$ [–]
A probe at rack 3	0.1	0.0993	88.0	0.0190
A probe at rack 3	0.3	0.2821	65.9	0.0313
A permanent inlet sensor at rack 3, every window	0.3	0.2821	65.9	0.0313
An airflow re-survey of rack 3's fraction	0.03	0.3319	59.9	0.0281
A second return-air sensor	0.3	0.8238	0.5	–
A second sensor at rack 1	0.3	0.8283	0.0	–

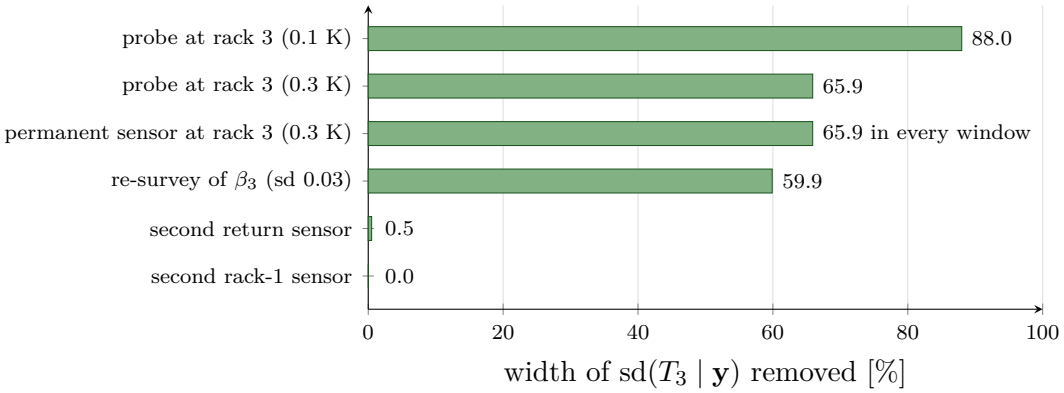


Figure 5.6: The candidate readings of Table 5.8 ranked by the width they would remove from the band of the virtual sensor at rack 3, by the preposterior (5.19).

The ranking follows the budget. Every candidate that touches β_3, directly or through T_3, removes more than half the width; every candidate that does not removes almost nothing, however good the sensor. A 0.3 K probe leaves rack 3 known to 0.2821 K, the spread rack 1's inlet already has with its own sensor: with a reading of its own, rack 3 is known as rack 1 is. Read on the entries, the probe and the re-survey do nearly the same thing: each acts mainly on β_3. The probe narrows it from 0.08 to 0.0313, the re-survey to 0.0281, and the two remove nearly the same width from T_3.

5.6.3 The recommendation

The next reading is a probe at rack 3 tonight: a 0.3 K probe removes 65.9 % of the band in one visit, and a 0.1 K probe 88.0 %. The request beyond tonight is a permanent inlet sensor at rack 3, which removes the same 65.9 % in every window instead of one, so that the hottest rack is read rather than inferred. A re-survey of the hall's airflow is the alternative where a sensor cannot be fitted, and it is worth having for reasons the rows do not show: it is taken under controlled conditions rather than at one point in one window, it covers rack 2, whose fraction is as unobserved as rack 3's, and it can find the stratification that a single fraction hides (Section 5.6.5). A second sensor at rack 1 or in the return air is not worth requesting for this question. Chapter 18 ranks the readings an operator can request on the shift plant in the same way.

5.6.4 What the analysis showed

The dashboard said the racks were five degrees inside the limit, because it showed rack 1; the analysis puts rack 3 two degrees inside, at 24.98°C, and the stated truth is warmer still. The one measured rack is the wrong proxy for the hall: rack 1's sensor fixes rack 1's recirculation and almost nothing else, and the rack that matters shares with it only the cold aisle and the rise, which the hall's other readings already fix. What is uncertain about rack 3 is its airflow, not the plant's temperatures. The return air, the two discharge temperatures and the meter pin the cold aisle and the rise to under two tenths of a kelvin, and the band of rack 3 is the survey's spread times that rise. It follows that a sensor anywhere else in the hall, a second return-air sensor or a second sensor at rack 1, cannot narrow it, and that the survey's quality is the hall's thermal margin at rack 3: the answer moves with the survey's center and widens with its width. The hall's instrumentation was placed to watch the cooling units, and it watches them well; it was not placed to watch the racks.

5.6.5 Limits of the model

One fraction per rack. On a stratified inlet the one fraction of Section 5.1.4 describes an average, and a probe at the top of rack 3 would read warmer than T_3. Racks that share a cause of recirculation, such as missing blanking panels along one row, would have correlated fractions, and a reading of one would then inform the other.

Steady state. The balances (5.1) hold within a window only when the hall has settled. After a load step or a unit start the hall's air and equipment store or release heat, the rise (5.5) no longer equals $(Q + Q_o)/(C_1 + C_2)$, and the fit misplaces every temperature until the transient has passed.

The survey's age. The prior on β_3 is a survey taken at another time, under another load and perhaps another arrangement of the racks. The model has no row that would detect a survey that has gone stale.

The pinned airflow and the absent control. The capacity rates C_u are pinned at the units' fixed fan speed. An error in them scales the rise and every fraction's lever with it, and it lies outside every band above. Under the active control that Section 5.1.2 excludes, the log-derate would be informed by the valve position, a row this model does not have.

Constructed readings. The five readings are the constructed ones of Section 5.1.4, not field data, and the agreement of the answers with the truth (Section 5.4.4) is a check of the method on its own model, not of the model against a real hall.

The next chapter. Chapter 6 asks a second question of the same fit: how much heat each unit is actually rejecting.

5.7 What to retain

Definition 5.1 (Virtual sensor)**.** A virtual sensor reads an unmeasured quantity $f(\mathbf{z})$ of the plant from the fitted state: on the reduced state of Chapter 4, f is a dependent temperature of the physical rows and its gradient $\mathbf{g}$ is its row of $\mathbf{N}_\mathrm{b}$; its value is $f(\hat{\mathbf{z}})$, its variance $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$, its band $f(\hat{\mathbf{z}}) \pm 1.960\sqrt{\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}}$, its tail at a limit a normal tail, and the reading of a probe at its place the predictive $\mathcal{N}(f(\hat{\mathbf{z}}),\ \mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g} + \sigma_a^2)$. The variance budget $\sum_i [\mathbf{g}]_i^2 \Sigma_{ii}$ plus the cross terms names the entries behind the band, and the columns of $\mathbf{J}$ name the rows behind those entries.

An entry that only its prior informs, read off a column of $\mathbf{J}$ with no reading row, carries its prior spread into every virtual sensor that depends on it, and only a reading of that entry or of the quantity itself narrows the band.

Three conclusions carry beyond the hall. A quantity with no sensor is known as well as the rows behind its gradient make it, and the budget says which they are. A reading elsewhere informs it only through the entries the two share, so the instrument the dashboard displays may be the wrong proxy for the place that matters. And the value of a further reading follows from the budget before the reading is taken: it is the preposterior of the row that would be added.

Table 5.9 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 5.9: The terms of the virtual sensor and where each is defined. Units for the two-unit hall.

Term	Expression	Units	Meaning	Defined in
Coil heat	$Q_u = \bar{G}_u \mathrm{e}^{g_u}(T_r - T_s)$	kW	heat to the chilled water	5.1.2
Physical rows	eq. (5.1), seven rows	kW, K	balance, air sides, mixing, inlets	5.1.2
Rack inlet	$T_k = T_c + \beta_k(T_r - T_c)$	°C	mix of cold aisle and return	5.1.2
Air balance	$(C_1 + C_2)(T_r - T_c) = Q + Q_o$	kW	the rise set by heat and airflow	5.1.2
Reduced state	$(Q, g_1, g_2, \beta_1, \beta_2, \beta_3)$	kW, –	the free entries	5.1.2
Sparsity pattern	eq. (5.6)	–	which row sees which entry	5.1.4
Rows and posterior				
Residual vector	eq. (5.7), ten rows	–	readings and priors on the reduced state	5.2.1
Reduced Jacobian	$\mathbf{J} = \mathbf{J}_{\text{rows}}\mathbf{N}$, 10×6	mixed	rows seen through the physics	5.2.1
Normal equations	$(\mathbf{J}^{\mathsf{T}}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^{\mathsf{T}}\mathbf{r}$	mixed	Gauss–Newton step	5.2.2
Block structure	eq. (5.10)	–	unobserved fractions keep their prior	5.2.2
Questions and answers				
Virtual sensor	$T_3(\hat{\mathbf{z}})$, $\mathbf{g}_{T_3}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_{T_3}$	°C, K^2	unmeasured inlet with its band	5.3.1
Gradient of an inlet	eq. (5.12)	K, K/kW	a row of $\mathbf{N}_{\mathrm{b}}$	5.3.1
Variance budget	$\sum_i [\mathbf{g}]_i^2 \Sigma_{ii}$ + cross terms	K^2	entries behind the band	5.3.1
Exceedance	$\Phi\big(-(T_{\max} - T_3(\hat{\mathbf{z}}))/\mathrm{sd}\big)$	–	inlet over the limit	5.3.2
Comparison of racks	$\mathbf{g}_\Delta = \mathbf{g}_{T_3} - \mathbf{g}_{T_1}$	K, K/kW	difference free of the derates	5.3.2
Probe predictive	$\mathcal{N}(T_3(\hat{\mathbf{z}}),\ \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} + \sigma_a^2)$	°C	a reading not yet taken	5.3.3
Preposterior	$\boldsymbol{\Sigma}_a = (\boldsymbol{\Sigma}^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a)^{-1}$	mixed	band after a candidate reading	5.6.1
Width removed	$1 - \mathrm{sd}_a/\mathrm{sd}$	–	ranks the candidates	5.6.2

5.8 Exercises

Exercise 5.1. *Formulation.* Rack 2 has no sensor either. Write the operator's question for rack 2 as a probability over the reduced state, give the gradient of its inlet, and state its answer from the fit of this chapter. ($\Pr(T_2 \le t \mid \mathbf{y})$ with $T_2(\mathbf{z}) = T_c + \beta_2(T_r - T_c)$; by equation (5.12) with the index moved, $\mathbf{g}_{T_2} = (1/(G_1+G_2) - (1-\beta_2)/(C_1+C_2),\ -G_1(T_r - T_s)/(G_1+G_2),\ -G_2(T_r - T_s)/(G_1+G_2),\ 0,\ T_r - T_c,\ 0)^{\mathsf{T}}$; $T_2 = 23.46$°C with standard deviation 0.83 K, and the stated

truth 22.99°C lies inside its band. Its band too is the survey's: of the total $0.6855\,\mathrm{K}^2$, 2.4 % is Q, 5.6 % g_1, 5.5 % g_2, −9.4 % the cross terms and 95.8 % β_2.)

Exercise 5.2. *Derivation.* Derive $\partial T_k/\partial\beta_k$ from the inlet row of equation (5.1) and the air balance (5.5), and use it to check the β_3 term of the budget. ($T_k = T_c + \beta_k(T_r - T_c)$ and neither T_c nor T_r depends on β_k, so $\partial T_k/\partial\beta_k = T_r - T_c = (Q + Q_o)/(C_1 + C_2) = 10.13\,\mathrm{K}$; with $\Sigma_{66} = 0.08^2$ the β_3 term has standard deviation $0.08 \times 10.13 = 0.81\,\mathrm{K}$, which is also the shift of the answer when the survey's center moves by 0.08, and its square is the $0.657\,\mathrm{K}^2$ of Table 5.4.)

Exercise 5.3. *Derivation.* Show from the gradients that the correlation of the inlets of racks 3 and 1 is small. ($\mathrm{Cov}(T_3, T_1) = \mathbf{g}_{T_3}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_{T_1}$. The two gradients have their last nonzero entries in the columns of β_3 and β_1, and by the block structure (5.10) β_3 is uncorrelated with every other entry, so the β_3 entry of $\mathbf{g}_{T_3}$ contributes nothing to the covariance; what remains is the part of $\mathbf{g}_{T_3}$ in Q, g_1, g_2, whose own variance is that of T_3 with β_3 known, $0.171^2\,\mathrm{K}^2$. By the Cauchy–Schwarz inequality the covariance is at most $0.171 \times 0.282\,\mathrm{K}^2$ in magnitude, so the correlation is at most the ratio of 0.171 to the full 0.828 K, and in the fit it is 0.0145.)

Exercise 5.4. *Application.* The technician's 0.3 K probe at rack 3 reads 25.9°C. Update the virtual sensor. (The probe is the row $(T_3(\mathbf{z}) - 25.9)/0.3$; by the rank-one update of equation (4.66) applied to T_3, the gain is $K = V/(V + 0.3^2) = 0.8840$ with $V = 0.686105\,\mathrm{K}^2$, the new mean is $24.9841 + K\,(25.9 - 24.9841)$°C, between the virtual reading and the probe and closer to the probe, and the new standard deviation is 0.2821 K, the value of Table 5.8, since the spread does not depend on the value read. The reading lies inside the probe predictive of Section 5.3.3. The gain falls on β_3, 0.0835 per K of innovation, the probe acting as the one-point re-survey of Section 5.6.1: β_3 moves from 0.50 to 0.577, and T_3 from 24.98 to 25.79°C, with standard deviation 0.282 K.)

Exercise 5.5. *Application.* The facilities team proposes a second inlet sensor at rack 1 to improve the hall's inlet monitoring. Explain from the preposterior why it removes nothing from the band of rack 3. (Its row is $\mathbf{J}_a = \mathbf{g}_{T_1}^{\mathsf{T}}/0.3$, and by equation (5.19) it removes $(\mathbf{g}_{T_3}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_{T_1})^2/(0.3^2 + \mathbf{g}_{T_1}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_{T_1})$ from $\mathrm{Var}(T_3)$, which is $\mathrm{Var}(T_3)$ times the square of the correlation 0.0145 times a factor below one; the sensor narrows β_1, which T_3 does not depend on, and the entries the two inlets share are already fixed by the other readings, so the width removed is 0.0 %. A probe at rack 3 removes 65.9 %.)

Chapter 6

Estimating Unmeasured Heat Flows

6.1 Heat the units do not meter

It is 23:00 in the two-unit hall of Chapter 5, and the shift log asks for the heat each cooling unit is rejecting. The hall load meter reads 193.04 kW; with the 8 kW of lighting and fans the two coils together carry 201.04 kW. Nothing meters either coil. The two units are the same model with the same rating, so the operator splits the hall's heat between them by rating and writes 100.52 kW for each.

The units' own tiles do not agree with that entry. The hall's one return-air sensor reads 30.33°C, and each unit's tile sets it against the unit's discharge temperature, 19.47°C for unit 1 and 20.11°C for unit 2; the air-side arithmetic, the temperature drop times the unit's air capacity rate, gives 108.60 kW for unit 1 and 102.20 kW for unit 2: unequal, and together more than the meter says the coils carry. The even split ignores the temperatures; the air-side arithmetic ignores the meter and every sensor's error. Neither settles how much heat unit 1 is actually rejecting tonight, how well the hall's readings know it, and which reading would know it better. The question of this chapter is

how a heat flow that no meter reads is estimated from the fitted state of the plant, with its band, and how the hall balance ties the flows of the two units together.

It is answered from the fit of Chapter 5, not from a new one: the same window, the same rows and the same posterior, read at the two coils instead of at rack 3. What changes is the quantity asked for, a flow along an edge of the graph rather than a temperature at a node, and the role of the physics: the balance rows, which fixed the cold aisle and the rise there, now split the hall's heat between the units.

Jargon: heat rejection, water-side heat meter, even split, load share, hall balance, derived flow, anti-correlation

The *heat rejection* of a cooling unit is the heat its coil passes to the chilled water, Q_u in kW. A *water-side heat meter* reads it directly, from the water flow through the coil and the temperature rise of the water. The *even split* is the dashboard's reading of identical units, the hall's coil heat divided by the number of units. The *load share* of unit 1 is the fraction $Q_1/(Q_1 + Q_2)$ of the coil heat that passes through it. The *hall balance* is the steady-state statement that the heat put into the hall leaves through the coils. A *derived flow* is a flow of the plant computed from the state rather than carried as an entry of it. Two quantities are *anti-correlated* when their posterior correlation is negative: a larger value of one goes with a smaller value of the other.

Key idea

The heat a unit rejects is a derived quantity of the fitted state, read with its band by $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ like any virtual sensor. The hall balance makes the gradients of the two unit heats add up to the gradient of the load alone, so the readings know their sum exactly as well as they know the load. Whether the two heats are correlated positively or negatively depends on which the readings know better, the total or the split; where the split is the less certain, they are anti-correlated, the sum is narrower and the difference wider than independence would give, and the question of which unit carries more has a band of its own.

6.1.1 A heat flow read from the fitted state

Chapter 5 read a temperature that no sensor measures, the inlet of rack 3, as a function of the six free entries of the reduced state. A heat flow is read the same way. On the reduced state the unit heats are not entries: they are functions $Q_1(\mathbf{z})$ and $Q_2(\mathbf{z})$ fixed by the coil rows once the return air is fixed, and their values and bands come from the MAP state, the covariance $\boldsymbol{\Sigma}$ and their gradients, by the delta method of Chapter 4. The one posterior of the 23:00 window serves both chapters' questions; nothing is re-fitted, and every answer of this chapter is a function of that posterior.

What a flow brings that a temperature did not is a conservation law across the quantities asked for. The two unit heats are tied by the hall balance: whatever the readings say about their sum, they say about the load. A band on each unit is then only part of the answer; the covariance of the two is the other part, and it decides how sure the operator can be that the units differ.

> **Principle 6.1.** *A flow along an edge of the plant graph is a derived quantity of the fitted state, with the variance* $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ *of any other. Where a balance row ties several flows to one entry, their gradients sum to that entry's gradient and the variance of their sum is that entry's variance; the sign of their covariance is set by how well the readings know the split against the total.*

6.1.2 The two-unit hall (step 0)

The plant is the two-unit hall of Section 5.1.2, unchanged: the IT load Q, the two coils with their log-derates g_1 and g_2 against the ratings $\bar{G}_1 = \bar{G}_2 = 7.0\,\mathrm{kW/K}$, the pinned chilled-water supply $T_s = 14^{\circ}\mathrm{C}$, the pinned air capacity rates $C_1 = C_2 = 10.0\,\mathrm{kW/K}$, the other heat $Q_o = 8\,\mathrm{kW}$, three racks with their recirculation fractions, the seven physical rows (5.1) and the full state (5.2). The question is asked at the two coil edges, from the return air through each unit to the chilled water, which carry the coil row

$$Q_u = G_u\,(T_r - T_s), \qquad G_u = \bar{G}_u\,\mathrm{e}^{g_u}, \qquad u = 1, 2. \tag{6.1}$$

Modeling choices at the graph. The choices of Section 5.1.2 carry over: steady state within the window, one effective conductance per coil, the airflow pinned and no control acting within the window. One choice is new, because the question now reads the coil edge as a heat flow: the coil row is taken to be the heat the unit rejects to the chilled water, with no losses on the water side between the coil and the loop. This is what lets Q_u answer the log's question, since the heat the coil takes from the air is then the heat the unit puts into the loop; it fails if the unit's piping gains or loses heat to the hall (Section 6.6.5).

The graph. Figure 6.1 redraws the graph of Figure 5.1 with the question moved. The racks are drawn as one node, the cold aisle on the left and the return air in the middle as before, the free IT load above the return, the two units on the right with their free log-derates and the pinned chilled water at the far right; the dashed arrows are the five readings. The two coil edges, from each unit to the chilled water, are drawn in red: they are the flows asked for, and no reading touches them.

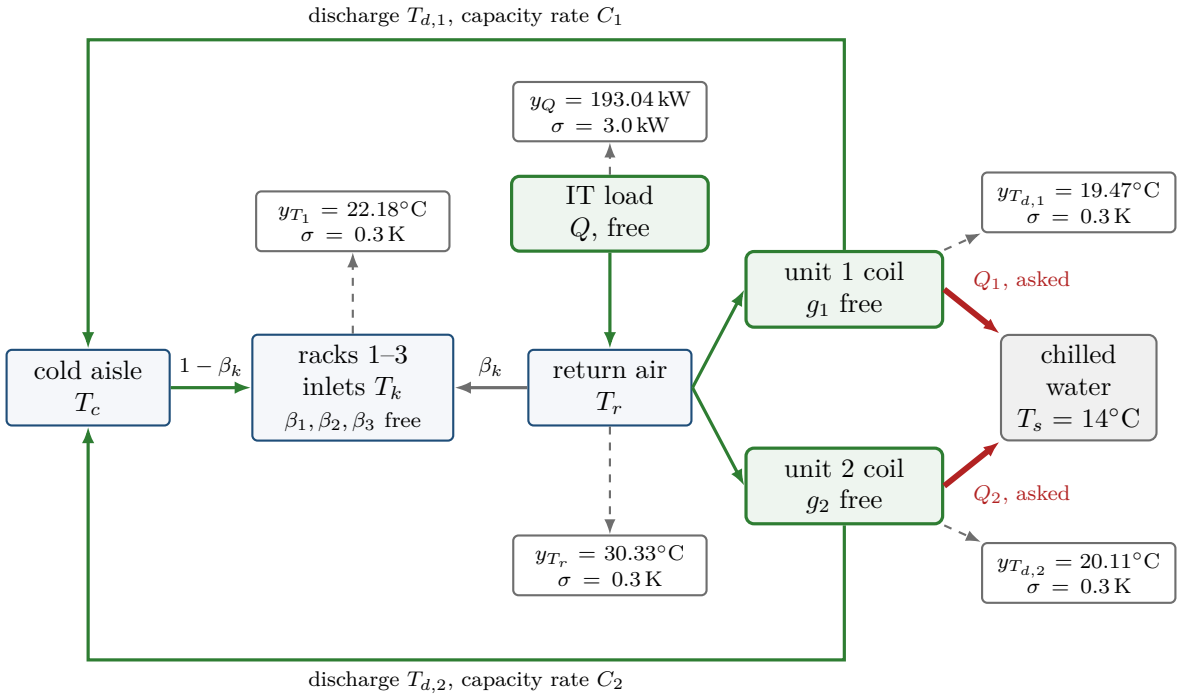


Figure 6.1: The rows of the two-unit hall at 23:00 with the question at the coils. Free (green): the IT load Q and the log-derates g_1, g_2; the recirculation fractions are free entries carried by the rack node, drawn as one. Blue: temperatures fixed by the physical rows (5.1). Pinned (gray): the chilled-water supply; also pinned, not drawn, the ratings $\bar{G}_u = 7.0\,\text{kW/K}$, the capacity rates $C_u = 10.0\,\text{kW/K}$ and the other heat $Q_o = 8\,\text{kW}$. Red: the coil edges, the unit heats Q_1 and Q_2 of equation (6.1), which no reading touches. Observed (dashed): the hall load meter, rack 1's inlet, the return air and the two discharge temperatures.

The reduced state. The fit runs on the reduced state of Chapter 4, the six free entries Q, g_1, g_2, β_1, β_2 and β_3 of equation (5.3). The dependent block $\mathbf{A}_{\text{dep}}$ of the physical rows is triangular with a nonzero diagonal, so it is invertible and the temperatures follow from the free entries in the order of equation (5.4): the return air from the hall balance, then the discharges, the cold aisle and the inlets. The coil row (6.1) at the return air of that chain gives the unit heats,

$$Q_u(\mathbf{z}) = G_u\,(T_r - T_s) = \frac{G_u}{G_1 + G_2}\,(Q + Q_o), \qquad u = 1, 2, \tag{6.2}$$

since $T_r - T_s = (Q + Q_o)/(G_1 + G_2)$. Two identities of the null space follow and are worth having in view. Summed over the units,

$$Q_1(\mathbf{z}) + Q_2(\mathbf{z}) = Q + Q_o, \tag{6.3}$$

the hall balance read on the coil edges, at every point of the set and not as a substitution. And the fraction of the coil heat through unit 1 depends on the two log-derates alone,

$$s_1(\mathbf{z}) = \frac{Q_1}{Q_1 + Q_2} = \frac{G_1}{G_1 + G_2} = \frac{1}{1 + (\bar{G}_2/\bar{G}_1)\,\mathrm{e}^{g_2 - g_1}}, \qquad s_2 = 1 - s_1, \tag{6.4}$$

which is one half at the ratings, $g_1 = g_2 = 0$, for these two equal units, and moves only with their difference $g_1 - g_2$.

6.1.3 The operator's questions (step 1)

The operator asks four questions of the 23:00 window, each written as a probability over the reduced state with the five readings and the priors on the right of the bar. Table 6.1 lists them, with the columns of Table 5.1: the question in the operator's words, the drill question first; the general form of its category in the notation of Chapter 2, for a derived quantity $f(\mathbf{z})$ and a level f^* of it; its mathematical form for the hall; its category; the computation of Section 2.3.6 that answers it; and the subsection where it is answered. Which reading would narrow the bands most is the next reading for these questions, Section 6.6.

Table 6.1: The questions asked of the two-unit hall at 23:00 about its unit heats. $\mathbf{y}$ is the five readings and the five priors of Section 5.1.4; $Q_u(\mathbf{z})$ the heat of unit u by equation (6.2); s_1 the share of unit 1 by equation (6.4).

Operator question	General form	Mathematical form	Category	Computation	Answered in
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
Is unit 1 carrying more than unit 2?	$\Pr(f(\mathbf{z}) \ge f^* \mid \mathbf{y})$	$\Pr(Q_1 - Q_2 > 0 \mid \mathbf{y})$	prediction under a threshold	posterior	6.3.2
What share of the hall's heat goes through unit 1?	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$	$p(s_1 \mid \mathbf{y})$, $s_1 = Q_1/(Q_1+Q_2)$	estimation	posterior	6.3.3

6.1.4 The evidence as rows (step 2)

The evidence is that of Section 5.1.4, attached to the graph as the same ten rows: the five readings of Figure 6.1, the commissioning priors on the two log-derates and the airflow survey's priors on the three fractions. No row reads a unit heat: the hall has no water-side meter, and the discharge readings see the coils only through the temperatures they set.

Which row acts on which entry. The Jacobian has the pattern (5.6): the return air and the two discharge rows act on (Q, g_1, g_2), rack 1's inlet on those three and β_1, the meter on Q alone, and each prior on its own entry. By equation (6.2) the unit heats depend on (Q, g_1, g_2) only, and by equation (6.4) the share on (g_1, g_2) only; with $*$ for an entry that can be nonzero and the columns $Q, g_1, g_2, \beta_1, \beta_2, \beta_3$,

$$\mathbf{g}_{Q_1}^{\mathsf{T}} = (*,\ *,\ *,\ 0,\ 0,\ 0), \qquad \mathbf{g}_{Q_2}^{\mathsf{T}} = (*,\ *,\ *,\ 0,\ 0,\ 0), \qquad \mathbf{g}_{s_1}^{\mathsf{T}} = (0,\ *,\ *,\ 0,\ 0,\ 0). \quad (6.5)$$

These gradients are not rows; they are the patterns of the questions' quantities. Read against the columns of $\mathbf{J}$, they settle the observability of the answer before any number. Rack 3's inlet in Chapter 5 had an entry in the column of β_3, which only its prior row informs. Every nonzero of the gradients above lies in the columns of Q, g_1 and g_2, which four temperature readings (the return air, both discharges and rack 1's inlet) and the meter see; so the bands of the unit heats are made of readings, not of a prior pushed through the physics.

6.2 The posterior of the state (step 3)

Nothing is fitted here. The rows of step 2 add nothing to those of Chapter 5, so the posterior of the reduced state is the fit of Section 5.2.3: the normal with the MAP state of Table 5.3 and the Laplace covariance (5.9). Its load is $\hat{Q} = 194.6275\,\text{kW}$, with spread $2.7694\,\text{kW}$, above the meter's $193.04\,\text{kW}$ because the temperature readings pull it up, and its log-derates are $\hat{g}_1 = -0.07355$ and $\hat{g}_2 = -0.13413$. The block of $\boldsymbol{\Sigma}$ in (Q, g_1, g_2) is

$$\boldsymbol{\Sigma}_{(Q,g_1,g_2)} = \begin{bmatrix} 7.66947874 & 0.02055535 & 0.02271429 \\ 0.02055535 & 0.00056368 & -0.00025057 \\ 0.02271429 & -0.00025057 & 0.00062655 \end{bmatrix}, \tag{6.6}$$

in kW^2, kW and dimensionless units, with the correlations 0.3126 for Q with g_1, 0.3277 for Q with g_2 and -0.4216 for g_1 with g_2, whose signs the rows fix (Section 5.2.3).

This block is all that the unit heats' gradients consume. By the patterns (6.5) the survey's block of β_2 and β_3 never enters; β_1 has nonzero covariances with Q, g_1 and g_2, but every gradient has a zero in its column, so they are multiplied by zero. Every answer of this chapter is a quadratic form on the 3×3 block (6.6).

6.3 The answers (step 4)

Every answer is read from the pair $(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of Section 6.2, and that value is the one reported. A derived quantity f is carried by its gradient $\mathbf{g} = \partial f / \partial \mathbf{z}$ at $\hat{\mathbf{z}}$, with variance $\mathbf{g}^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g}$. Draws of the state pushed through the exact equation (6.2) stand beside each reported value, with the gap.

6.3.1 The heat of each unit

Operator questions. How much heat is this unit actually rejecting? And the other unit? The first asks for the heat of unit 1 at 23:00 with its band; the second for the heat of unit 2, and with it how the two answers are related.

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

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

summarized by its mean and variance, equation (2.10), and its central 95 % credible interval, equation (2.11).

Specialized form. The two questions are equations (2.9), (2.10) and (2.11) with $\mathbf{z}$ the reduced state (5.3), $f(\mathbf{z}) = Q_u(\mathbf{z})$ of equation (6.2) for $u = 1, 2$, f^* a heat in kW, $[f_L, f_U]$ the band of unit u and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal of Section 5.2.2. One choice enters here: the answer is reported by the delta method on that normal, which gives it in closed form from $\hat{\mathbf{z}}$, $\boldsymbol{\Sigma}$ and one gradient, and which is accurate only if Q_u is nearly linear in the state over the posterior's width; the unit heats are nonlinear in the log-derates through the ratio of equation (6.2), and the draws set beside the answer test the linearization.

The gradient of Q_u is formed from the coil row (6.1) and the gradient of the return air of Section 5.3.1, $\partial T_r / \partial \mathbf{z} = (1,\ -G_1(T_r - T_s),\ -G_2(T_r - T_s),\ 0,\ 0,\ 0)^{\mathsf{T}} / (G_1 + G_2)$:

$$\frac{\partial Q_u}{\partial \mathbf{z}} = (T_r - T_s)\, \frac{\partial G_u}{\partial \mathbf{z}} + G_u\, \frac{\partial T_r}{\partial \mathbf{z}}, \qquad \frac{\partial G_u}{\partial g_u} = G_u.$$

With the shares $s_u = G_u/(G_1 + G_2)$ of equation (6.4) and $Q_u = s_u(Q + Q_o)$, the entries collect into

$$\mathbf{g}_{Q_1} = \begin{bmatrix} s_1 \\ \dfrac{Q_1 Q_2}{Q + Q_o} \\ -\dfrac{Q_1 Q_2}{Q + Q_o} \\ 0 \\ 0 \\ 0 \end{bmatrix}, \qquad \mathbf{g}_{Q_2} = \begin{bmatrix} s_2 \\ -\dfrac{Q_1 Q_2}{Q + Q_o} \\ \dfrac{Q_1 Q_2}{Q + Q_o} \\ 0 \\ 0 \\ 0 \end{bmatrix} \begin{matrix} [-] \\ [\mathrm{kW}] \\ [\mathrm{kW}], \end{matrix} \tag{6.7}$$

evaluated at $\hat{\mathbf{z}}$. Three readings of it follow. A kilowatt more load goes to the units in proportion to their shares. A larger g_1 takes heat from unit 2 and gives it to unit 1 by the same amount, $Q_1 Q_2/(Q + Q_o)$ per unit of log-derate, because the hall balance holds the total. And the two gradients add up to

$$\mathbf{g}_{Q_1} + \mathbf{g}_{Q_2} = (1,\ 0,\ 0,\ 0,\ 0,\ 0)^{\mathsf{T}} = \mathbf{g}_Q, \tag{6.8}$$

the gradient of the load, which is equation (6.3) differentiated: the balance read off the gradients.

By the delta method, equation (4.53), the unit heats and their bands are

$$Q_u \mid \mathbf{y} \approx \mathcal{N}(Q_u(\hat{\mathbf{z}}),\ \mathbf{g}_{Q_u}^{\mathsf{T}} \boldsymbol{\Sigma}\, \mathbf{g}_{Q_u}), \qquad [Q_u(\hat{\mathbf{z}}) \mp 1.960 \sqrt{\mathbf{g}_{Q_u}^{\mathsf{T}} \boldsymbol{\Sigma}\, \mathbf{g}_{Q_u}}\,], \tag{6.9}$$

and the two together are approximately bivariate normal with the covariance $\mathrm{Cov}(Q_1, Q_2 \mid \mathbf{y}) = \mathbf{g}_{Q_1}^{\mathsf{T}} \boldsymbol{\Sigma}\, \mathbf{g}_{Q_2}$. The identity (6.8) ties that covariance to the load: the quadratic form of $\mathbf{g}_Q$ is Σ_{11}, the posterior variance of the load, so

$$\mathrm{Var}(Q_1) + \mathrm{Var}(Q_2) + 2\ \mathrm{Cov}(Q_1, Q_2) = \mathrm{Var}(Q_1 + Q_2) = \Sigma_{11}. \tag{6.10}$$

The sum is known exactly as well as the load. The balance does not fix the sign of the covariance; the total is not fixed, it is uncertain with the load. The sign is read from the two factors of $Q_u = s_u(Q + Q_o)$. To first order $\delta Q_1 = s_1\, \delta Q + (Q + Q_o)\, \delta s_1$ and $\delta Q_2 = s_2\, \delta Q - (Q + Q_o)\, \delta s_1$, since $s_2 = 1 - s_1$; so

$$\mathrm{Cov}(Q_1, Q_2) = s_1 s_2\ \mathrm{Var}(Q) - (Q + Q_o)^2\ \mathrm{Var}(s_1) + (Q + Q_o)(s_2 - s_1)\ \mathrm{Cov}(Q, s_1). \tag{6.11}$$

The first term is positive: the shared load moves both heats the same way. The second is negative: the uncertain split moves them in opposite ways. The third is small while the shares are near one half. The heats are anti-correlated when the readings know the split less well than the total, and positively correlated when they know the split better, as discharge sensors of 0.03 K (correlation +0.95) or a load meter of 30 kW (correlation +0.48) would make them.

The budget of the variance of Q_2 splits it by entry, as equation (5.14) does for the inlet,

$$\mathbf{g}_{Q_2}^{\mathsf{T}} \boldsymbol{\Sigma}\, \mathbf{g}_{Q_2} = \sum_i [\mathbf{g}_{Q_2}]_i^2\, \Sigma_{ii} + \sum_{i \neq j} [\mathbf{g}_{Q_2}]_i\, [\mathbf{g}_{Q_2}]_j\, \Sigma_{ij}, \tag{6.12}$$

with every term among Q, g_1 and g_2 by the pattern (6.5). The signs of the cross terms follow from the correlations of Section 6.2. The gradient of Q_2 has the signs $(+, -, +)$ on (Q, g_1, g_2). The cross term in Q and g_1 is a positive times a negative times a positive covariance, negative; the one in Q and g_2 is positive; and the one in g_1 and g_2 is a negative times a positive times a negative covariance, positive. The last is the term of the split: a gradient with opposite signs on the two log-derates meets their negative correlation, and the readings' uncertainty about how the conductance divides between the coils adds to the band instead of cancelling in it. The budget of Q_1 is the mirror, with the signs $(+, +, -)$ and the same conclusion.

Building blocks. The MAP state $\hat{\mathbf{z}}$ and the block (6.6) of $\boldsymbol{\Sigma}$ (Section 6.2); the heats $Q_u(\hat{\mathbf{z}})$ and shares $s_u(\hat{\mathbf{z}})$ by equations (6.2) and (6.4); the gradients (6.7); the quadratic forms (6.9), the covariance of the pair (6.11) and the budget (6.12).

The heats at 23:00. At the MAP state the conductances are $G_1 = 6.5036$ and $G_2 = 6.1213\,\mathrm{kW/K}$, of sum $12.6249\,\mathrm{kW/K}$, the shares $s_1 = 0.51514$ and $s_2 = 0.48486$, and the coils carry $Q + Q_o = 202.6275\,\mathrm{kW}$. Table 6.2 gives the gradients (6.7) at $\hat{\mathbf{z}}$: each row is a quantity, each column an entry of the reduced state, the load's column dimensionless and the log-derates' in kW per unit of log-derate. The two rows mirror each other; their sum, the last row, is the gradient of the load, as equation (6.8) requires.

Table 6.2: The gradients of the unit heats at the MAP state, equation (6.7), and of their sum.

Quantity	$\partial/\partial Q$ [–]	$\partial/\partial g_1$ [kW]	$\partial/\partial g_2$ [kW]	$\partial/\partial \beta_1$ [kW]	$\partial/\partial \beta_2$ [kW]	$\partial/\partial \beta_3$ [kW]
Q_1	0.5151	50.6104	−50.6104	0	0	0
Q_2	0.4849	−50.6104	50.6104	0	0	0
$Q_1 + Q_2$	1.0000	0.0000	0.0000	0	0	0

By equation (6.9),

$$\begin{aligned} Q_1 \mid \mathbf{y} &\approx \mathcal{N}(104.3816\,\mathrm{kW},\ 2.5010^2), \quad Q_1 \in [99.48,\ 109.28]\,\mathrm{kW}, \\ Q_2 \mid \mathbf{y} &\approx \mathcal{N}(98.2459\,\mathrm{kW},\ 2.4983^2), \quad Q_2 \in [93.35,\ 103.14]\,\mathrm{kW}, \end{aligned} \tag{6.13}$$

the second column the central 95 % bands. The covariance of the pair is $\mathrm{Cov}(Q_1, Q_2 \mid \mathbf{y}) = -2.4134\,\mathrm{kW}^2$, a correlation of -0.3863. The sum is $Q_1 + Q_2 = 202.6275\,\mathrm{kW}$, which is $\hat{Q} + Q_o = 194.6275 + 8.0\,\mathrm{kW}$, with standard deviation $2.7694\,\mathrm{kW}$, the spread of the load: equation (6.10) holds on the numbers. The sum is not known better than either unit, whose standard deviations are 2.5010 and 2.4983 kW; it is known better than it would be if the two heats were independent, 3.5350 kW for the sum and for the difference alike, and the difference is known worse, 4.1621 kW. By equation (6.11) the covariance is

$$\mathrm{Cov}(Q_1, Q_2 \mid \mathbf{y}) = 1.9156 - 4.3323 + 0.0033 = -2.4134\,\mathrm{kW}^2 :$$

the shared load contributes $1.9156\,\mathrm{kW}^2$, the uncertain split $-4.3323\,\mathrm{kW}^2$, and the third term almost nothing. The heats are anti-correlated here because the readings know the split less well than the total.

Figure 6.2 draws the two answers: the horizontal axis is the heat of a unit, the vertical axis its posterior density, unit 1 in red and unit 2 in blue, each with its 95 % band shaded; the dashed line is the even split of the dashboard, and the triangles above the curves mark the stated truth of each unit, used in Section 6.4.4. The even split lies inside both bands, near the lower end of unit 1's.

Figure 6.3 draws the pair together. The horizontal axis is the heat of unit 1, the vertical axis that of unit 2; the closed curves are the contours of the bivariate normal of the pair at one and two standard deviations; the solid line is the hall balance $Q_1 + Q_2 = \hat{Q} + Q_o$ at the fitted load, the dashed line the same balance at the meter's reading, on which the even split (the square) sits; the thin dotted diagonal is $Q_1 = Q_2$, below which unit 1 carries more; and the diamond is the stated truth. The ellipse is long along the balance line: the readings know where the pair sits across that line, the sum, better than where it sits along it, the split.

Beside the reported values stand 100 000 draws of the state from $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$, each pushed through the exact equation (6.2). For unit 1 they give a mean of 104.3789 kW, a standard deviation of 2.4968 kW and a 95 % band $[99.4992,\ 109.3193]$ kW: gaps of -0.0027 and -0.0042 kW,

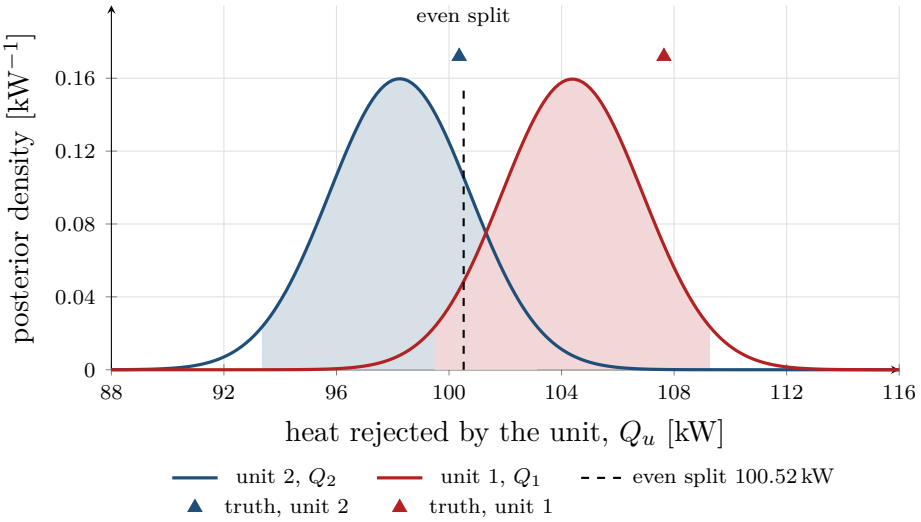


Figure 6.2: The heats of the two units at 23:00, the normals $\mathcal{N}(104.3816,\ 2.5010^2)$ for unit 1 (red) and $\mathcal{N}(98.2459,\ 2.4983^2)$ for unit 2 (blue) of equation (6.13), density $\varphi((Q_u-\mu_u)/\sigma_u)/\sigma_u$. Shaded: the central 95 % bands, 99.48 to 109.28 kW and 93.35 to 103.14 kW. Dashed: the dashboard's even split, 100.52 kW. Triangles: the stated truth of the record, 107.64 kW and 100.36 kW, check markers (Section 6.4.4); unit 2's truth and the even split coincide within 0.16 kW.

and of +0.0196 and +0.0357 kW at the band's ends. For unit 2 they give 98.2566 kW, 2.4913 kW and [93.4138, 103.1738] kW: gaps of +0.0107 and −0.0070 kW, and of +0.0645 and +0.0313 kW at the ends. The correlation of the pair is −0.3823, a gap of +0.0039.

The same normal also gives exact values without draws: given $g_1 - g_2$ the share is fixed and the load is normal, so one-dimensional quadrature over $g_1 - g_2$ gives the distribution of $Q_u = s_u(Q + Q_o)$. For unit 1 the mean is 104.3798 kW, the standard deviation 2.5004 kW and the band [99.5137, 109.3146] kW: gaps of −0.0018 and −0.0006 kW, and of +0.0341 and +0.0310 kW at the ends. For unit 2 they are 98.2477 kW, 2.4977 kW and [93.3931, 103.1830] kW: gaps of +0.0018 and −0.0006 kW, and of +0.0438 and +0.0406 kW. All four ends sit to the right of the delta method's band, the right skew of a product of two uncertain factors.

The budget of the band. By equation (6.12), the variances of the two heats split as Table 6.3 shows: each row is a term, the middle columns its value and its share of the variance for unit 1 and for unit 2. The cross terms are listed one by one, their shares together.

Table 6.3: The variance budgets of the unit heats, equation (6.12), at the MAP state. The fractions β_k contribute nothing.

Term	Q_1 [kW^2]	Share [%]	Q_2 [kW^2]	Share [%]
Q, $[\mathbf{g}]_1^2\,\Sigma_{11}$	2.0352	32.5	1.8030	28.9
g_1, $[\mathbf{g}]_2^2\,\Sigma_{22}$	1.4438	23.1	1.4438	23.1
g_2, $[\mathbf{g}]_3^2\,\Sigma_{33}$	1.6049	25.7	1.6049	25.7
cross term Q–g_1	1.0718		−1.0088	
cross term Q–g_2	−1.1844		1.1148	
cross term g_1–g_2	1.2836		1.2836	
cross terms together		18.7		22.3
total $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$	6.2550	100	6.2413	100

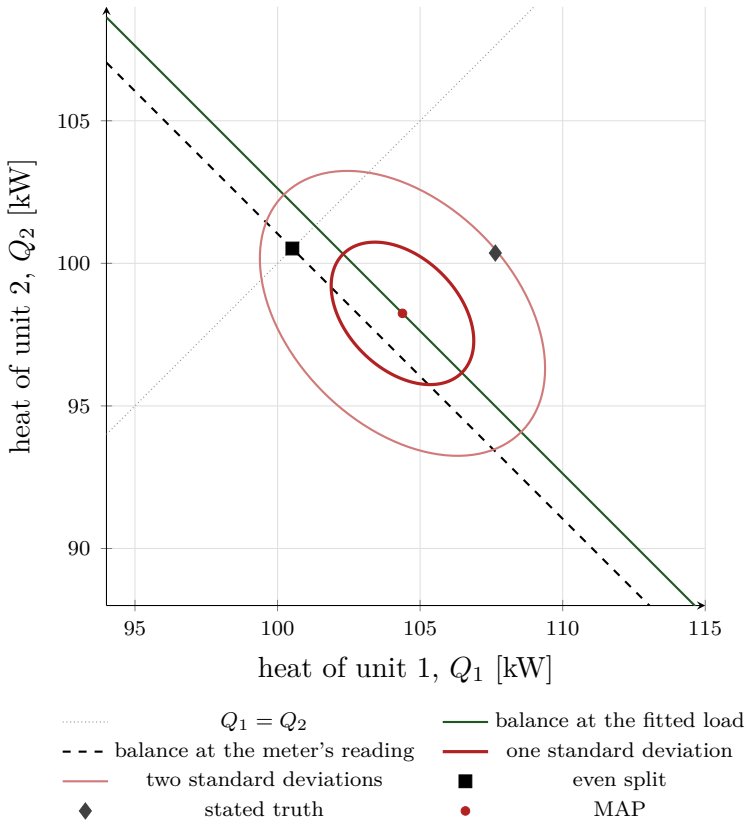


Figure 6.3: The joint posterior of the two unit heats: contours at one and two standard deviations of the bivariate normal with means 104.3816 and 98.2459 kW, standard deviations 2.5010 and 2.4983 kW and correlation −0.3863, drawn as $Q_1 = \mu_1 + k\sigma_1 \cos t$, $Q_2 = \mu_2 + k\sigma_2(\rho \cos t + \sqrt{1-\rho^2} \sin t)$ for $k = 1, 2$. Solid line: $Q_1 + Q_2 = 202.6275$ kW, the balance at the fitted load; dashed: $Q_1 + Q_2 = 201.04$ kW, the balance at the meter's reading, with the even split. Diamond: the stated truth (Section 6.4.4).

No entry dominates the band of Q_2, unlike the band of rack 3, of which the survey's fraction carried 95.7 % (Table 5.4). The load and the two log-derates carry comparable diagonal terms, and the cross terms add to the variance rather than taking from it, with the signs of the specialized form above: for Q_2 the term in Q and g_1 is $-1.0088\,\mathrm{kW}^2$, the term in Q and g_2 is $+1.1148\,\mathrm{kW}^2$, and the term in g_1 and g_2, the term of the split, is $+1.2836\,\mathrm{kW}^2$; together they add 22.3 %. The two terms in Q nearly cancel, and what remains is the split. For Q_1 the gradient's signs are $(+,+,-)$ and the two terms in Q change sign, $+1.0718$ and $-1.1844\,\mathrm{kW}^2$, while the split's term is the same $+1.2836\,\mathrm{kW}^2$.

6.3.2 The comparison of the two units

Operator question. Is unit 1 carrying more than unit 2? The units are identical by rating; the question is whether the hall's readings say that tonight they are not sharing the heat equally.

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

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

Specialized form. The question is equation (2.12) with $\mathbf{z}$ the reduced state (5.3), $f(\mathbf{z}) = Q_1(\mathbf{z}) - Q_2(\mathbf{z})$, $f^* = 0$ and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal of Section 5.2.2. The assurance level κ of that equation is not set here: the probability itself is reported. The tail is read from the normal of the delta method. The gradient of the difference is the difference of the two gradients (6.7),

$$\begin{aligned} \mathbf{g}_\Delta = \mathbf{g}_{Q_1} - \mathbf{g}_{Q_2} &= \Big(s_1 - s_2,\ \frac{2\,Q_1 Q_2}{Q + Q_o},\ -\frac{2\,Q_1 Q_2}{Q + Q_o},\ 0,\ 0,\ 0\Big)^{\mathsf{T}}, \\ \Pr(Q_1 > Q_2 \mid \mathbf{y}) &= \Phi\left(\frac{Q_1(\hat{\mathbf{z}}) - Q_2(\hat{\mathbf{z}})}{\sqrt{\mathbf{g}_\Delta^{\mathsf{T}} \boldsymbol{\Sigma}\, \mathbf{g}_\Delta}}\right). \end{aligned} \tag{6.14}$$

The load nearly drops out, since $s_1 - s_2$ is small when the shares are near one half, and the log-derates enter doubled with opposite signs: the difference is a question about the split. Its variance is also $\mathrm{Var}(Q_1) + \mathrm{Var}(Q_2) - 2\,\mathrm{Cov}(Q_1, Q_2)$. Against two independent heats, $\sqrt{\mathrm{Var}(Q_1) + \mathrm{Var}(Q_2)}$ for both the sum and the difference, a negative covariance narrows the sum and widens the difference by the same term.

The event has an exact form under the same normal. By equation (6.1) unit 1 carries more when $G_1 > G_2$, which for these equal ratings is $g_1 - g_2 > 0$, a linear function of the state, so

$$\Pr(Q_1 > Q_2 \mid \mathbf{y}) = \Phi\Big(\frac{\hat{g}_1 - \hat{g}_2}{\sqrt{\Sigma_{22} + \Sigma_{33} - 2\,\Sigma_{23}}}\Big). \tag{6.15}$$

Building blocks. The heats $Q_1(\hat{\mathbf{z}})$ and $Q_2(\hat{\mathbf{z}})$ and their variances and covariance from Section 6.3.1; the log-derates and the block (6.6); the normal distribution function Φ.

Unit 1 against unit 2. By equation (6.14), with the variance $6.2550 + 6.2413 + 2 \times 2.4134\,\mathrm{kW}^2$,

$$Q_1 - Q_2 \mid \mathbf{y} \approx \mathcal{N}\big(6.1357\,\mathrm{kW},\ 4.1621^2\big), \qquad \Pr(Q_1 > Q_2 \mid \mathbf{y}) = 0.9298,$$

unit 1 carrying more than unit 2 with probability 0.93. The standard deviation of the difference stands against the independent case of Section 6.3.1; in Figure 6.3 the difference is the long axis of the ellipse.

Beside the reported 0.92979, the exact form (6.15), with $\hat{g}_1 - \hat{g}_2 = 0.06058$ of standard deviation 0.04113, gives $\Phi(1.4730) = 0.92963$, a gap of -0.00016. The draws of Section 6.3.1 give a mean of 6.1223 kW, a standard deviation of 4.1469 kW and a probability of 0.92910: gaps of -0.0134 and -0.0152 kW and of -0.00069. The quadrature of the same section gives a mean of 6.1321 kW, a standard deviation of 4.1607 kW and a band $[-2.0297,\ 14.2795]$ kW: gaps of -0.0037 and -0.0014 kW from the delta method, and of -0.0077 and -0.0140 kW from its band $[-2.0220,\ 14.2935]$ kW.

6.3.3 The share of unit 1

Operator question. What share of the hall's heat goes through unit 1? The even split assumes one half; the question is what the window says it is, and how well.

General form. The estimation form again: the posterior distribution of a derived quantity, equation (2.9), summarized by its mean and variance, equation (2.10), and its central 95 % credible interval, equation (2.11).

Specialized form. The question is equations (2.9), (2.10) and (2.11) with $\mathbf{z}$ the reduced state (5.3), $f(\mathbf{z}) = s_1(\mathbf{z})$ of equation (6.4), f^* a share, $[f_L, f_U]$ the share's band and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal of Section 5.2.2. No new choice enters. The share is a logistic function of $g_1 - g_2$, so its gradient has two entries of opposite sign and equal size,

$$\mathbf{g}_{s_1} = s_1 s_2 \,(0,\ 1,\ -1,\ 0,\ 0,\ 0)^{\mathsf{T}}, \qquad \mathbf{g}_{s_1}^{\mathsf{T}} \boldsymbol{\Sigma}\, \mathbf{g}_{s_1} = s_1^2 s_2^2 \,(\Sigma_{22} + \Sigma_{33} - 2\,\Sigma_{23}), \tag{6.16}$$

with no entry in the load's column: the meter's error reaches the share only through $\mathrm{Cov}(Q, s_1)$. Its budget has two diagonal terms, $s_1^2 s_2^2 \Sigma_{22}$ and $s_1^2 s_2^2 \Sigma_{33}$, and one cross term, $-2 s_1^2 s_2^2 \Sigma_{23}$, positive because the two log-derates are negatively correlated: the gradient's opposite signs meet a negative covariance, the same term that widened Q_2 and the difference. The share's variance is the variance of $g_1 - g_2$, scaled.

The share also has an exact form under the same normal. It is an increasing function of $g_1 - g_2$ alone, so its median is $s_1(\hat{\mathbf{z}})$ and its central 95 % interval is the logistic function of $\hat{g}_1 - \hat{g}_2 \mp 1.960\sqrt{\Sigma_{22} + \Sigma_{33} - 2\,\Sigma_{23}}$, the same standard deviation as in equation (6.15).

Building blocks. The share $s_1(\hat{\mathbf{z}})$ of Section 6.3.1; the gradient and quadratic form (6.16) on the block (6.6); the budget of Q_2 of Table 6.3 for the comparison.

The share at 23:00. By equation (6.16), the share of unit 1 is $s_1(\hat{\mathbf{z}}) = 0.51514$ with standard deviation 0.01027, against 0.5 at the rating, and its central 95 % band is [0.4950, 0.5353]. Unit 1 most likely passes 51.5 % of the coil heat, but the band contains one half: the window does not exclude an even split of the heat at the 95 % level. The share at the MAP state equals the first entry of $\mathbf{g}_{Q_1}$ in Table 6.2, 0.5151, as equation (6.7) says it must. The load's correlation with the share is -0.0190, and the fit repeated with the meter at the true load moves the share by -0.04 of its standard deviation. The exact form gives the median 0.51514, a gap below the fifth decimal, and the band $[0.494993,\ 0.535239]$ against the reported $[0.495007,\ 0.535274]$, gaps of -0.000014 and -0.000035. The draws give a mean of 0.515109, a standard deviation of 0.010233 and a band $[0.494934,\ 0.535134]$: gaps from the reported 0.515140 and 0.010272 of -0.000031 and -0.000039, and of -0.000073 and -0.000140 at the ends.

Figure 6.4 sets the share's budget under that of Q_2. Its upper panel draws the budget of Q_2, one bar per term in kW^2, the three cross terms separately; its lower panel draws the budget of the share, of total 1.05517×10^{-4}, in percent of that total: 33.3 % for g_1, 37.0 % for g_2 and 29.6 % for its one cross term.

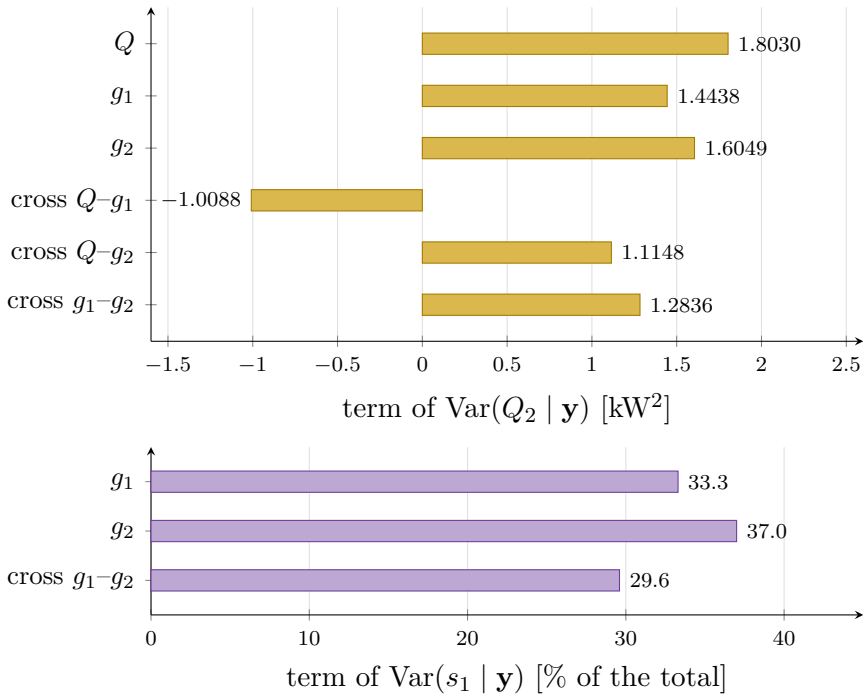


Figure 6.4: The variance budgets. Upper panel (gold): the heat of unit 2, equation (6.12), terms in kW^2 of the total $6.2413\,\mathrm{kW}^2$ (Table 6.3). Lower panel (purple): the share s_1 of unit 1, equation (6.16), in percent of the total 1.05517×10^{-4}.

The share's budget has no load term, and its single cross term, positive by the signs above, adds 29.6 %, against the 22.3 % the cross terms add to Q_2: the readings' uncertainty about the difference of the two log-derates is the whole of the share's band.

6.4 The checks (step 5)

The answers of Section 6.3 rest on approximations and inputs that can be tested on this hall, against the calculation each stands for.

6.4.1 The limit of exact rows and the curvature term

On the reduced state the unit heats are derived quantities. The check fits a full state that carries them as entries, so that the coil rows are rows of the fit rather than functions read after it: the twelve entries

$$\mathbf{z}_{\text{full}} = (Q,\ g_1,\ g_2,\ \beta_1,\ \beta_2,\ \beta_3 \mid T_r,\ Q_1,\ Q_2,\ T_{d,1},\ T_{d,2},\ T_c)^{\mathsf{T}},$$

with six physical rows at a common tolerance σ_p: the hall balance on the coil edges $Q_1 + Q_2 - Q - Q_o$, the two coil rows $Q_u - G_u(T_r - T_s)$, the two air sides $C_u(T_r - T_{d,u}) - Q_u$ and the cold-aisle mixing row, beside the five readings and five priors. Rack 1's inlet enters its reading row as $T_c + \beta_1(T_r - T_c)$; the inlets of racks 2 and 3, which no row reads, are not carried. As the tolerance shrinks, the fitted heat of unit 2 goes from 98.2466 kW with standard deviation 2.4996 kW at $\sigma_p = 0.1$ to 98.2459 kW and 2.4983 kW at 0.01 and at 0.001, the values of the reduced state to the fourth decimal. The difference and the share, read from the same full fit, behave alike (Table 6.5). The entry of the full state and the derived quantity of the reduced state are the same answer in the limit.

The covariance is the Gauss–Newton form, without the curvature term of the nonlinear rows. The full Hessian of the objective at $\hat{\mathbf{z}}$, curvature term included, gives a Newton covariance under which the heat of unit 1 has standard deviation 2.4981 kW, against 2.5010 kW, the heat of unit 2 2.4933 kW, against 2.4983 kW, and the difference 4.1539 kW, against 4.1621 kW: the term moves the standard deviations in the third decimal, and narrows them, so the reported Gauss–Newton bands are the slightly wider of the two. It moves the probability that unit 1 carries more by +0.00039.

6.4.2 The delta method against draws through the physics

The delta method linearizes the heats about $\hat{\mathbf{z}}$, while equation (6.2) is a ratio in the log-derates; the draws of Section 6.3.1 keep the ratio. Their Monte Carlo standard errors are 0.0079 kW on the mean of a heat, 0.0056 kW on its standard deviation and 0.021 kW on an end of its band; 0.0131 and 0.0093 kW on the mean and standard deviation of the difference; 0.00081 on the probability; and 0.00003, 0.00002 and 0.00009 on the mean, the standard deviation and a band end of the share. The second-order correction to the mean, $\frac{1}{2}\operatorname{tr}(\nabla^2 f\,\boldsymbol{\Sigma})$, is −0.00184 kW for Q_1, +0.00184 kW for Q_2 and −0.00367 kW for the difference, at most 0.28 of the standard error of the mean it would move, so whatever the ratio's curvature does to the mean is below what 100 000 draws resolve. The gaps on the means are −0.34 standard errors for Q_1, +1.36 for Q_2 and −1.02 for the difference, and those on the standard deviations −0.75, −1.25 and −1.64, and −1.72 for the share: several exceed one standard error, and none is resolved on its own. The largest gap is at the lower end of Q_2's band, +0.0645 kW or 3.06 standard errors. It is not all noise and not all signal: the quadrature of Section 6.3.1 puts the skew's shift there at +0.0438 kW, and the rest, +0.0207 kW, is about one standard error of this seed. Over 200 seeds of 100 000 draws the same gap averages 1.95 standard errors, with a spread of 0.97, and 12.0 % of the seeds reach 3.06 or more. The quadrature's mean gaps, −0.00183 and +0.00183 kW, are the second-order correction. The draws fall −0.00053 from the exact tail (6.15), −0.65 of a standard error, and that exact form resolves the reported tail's gap, −0.00016, which the draws cannot. As in Chapter 5, the check tests the delta method, not the Laplace posterior: all the routes use the same normal.

6.4.3 Sensitivity to the pinned airflow

The capacity rates set how the discharge temperatures inform the log-derates, so the answers are tested against them. Table 6.4 repeats the fit with one unit's rate moved by 5 %: each row is one such fit, the columns its heats, its share with the move in the reported standard deviation of the share, and its probability that unit 1 carries more. The split moves by more than one standard deviation in every row, and the probability ranges from 0.60 to 0.997; no band of this chapter holds that error.

Table 6.4: The answers with one capacity rate moved by 5 %, the fit repeated, against the reported $s_1 = 0.51514$ and $\Pr(Q_1 > Q_2) = 0.9298$ at $C_1 = C_2 = 10.0$ kW/K.

C_1 [kW/K]	C_2 [kW/K]	Q_1 [kW]	Q_2 [kW]	s_1 [–]	Move [sd]	$\Pr(Q_1 > Q_2)$ [–]
10.5	10.0	107.1102	96.1840	0.52687	+1.14	0.9948
9.5	10.0	101.4691	100.4282	0.50258	−1.22	0.6012
10.0	10.5	102.4364	100.8103	0.50400	−1.08	0.6486
10.0	9.5	106.4395	95.5094	0.52706	+1.16	0.9965

6.4.4 The answers against the stated truth

The record of this hall is generated from a stated truth (Section 5.1.4), and the last check compares the answers with it. Table 6.5 collects the checks: each row is a route, a tolerance or a reference, the columns give the value, its standard deviation and the probability, each with its gap from the reported value, the route minus the delta method. The true heats are 107.6385 kW for unit 1 and 100.3615 kW for unit 2. The MAP heats are 3.2569 kW and 2.1156 kW low, −1.30 and −0.85 of their standard deviations, both inside their bands. Taken jointly, with the covariance of the pair, the truth lies at a Mahalanobis distance of 1.959 from the fitted pair: just inside the two-standard-deviation contour of Figure 6.3, because both errors have the same sign, the direction of the sum, which the fit knows best.

The two errors have the same sign, and the balance says why. By equation (6.3) the fitted heats add up to the fitted load plus the other heat, so their errors add up to the error of the fitted load; both being negative, the fitted load lies below the true one. The total is 202.6275 kW against a true 208.0000 kW, an error of −5.3725 kW, −1.94 of its standard deviation; the share is 0.51514 against a true 0.51749, an error of −0.00235, −0.23 of its standard deviation. The fitted load rests mostly on the hall meter, which read 193.04 kW against a true load of 200.0 kW, 6.96 kW low, 2.32 of its standard deviation; why it read low is taken up in Chapter 7. The split, which the temperatures carry, is close to the truth; the total, which the meter carries, is not. The dashboard's even split is 7.12 kW low for unit 1 and 0.16 kW high for unit 2: right for unit 2 by chance, because the meter's low reading and the even split's wrong share happen to offset there, and wrong for unit 1 by 2.19 times the fit's error. The units' tiles stand about 4 kW above the fit for each unit and, this night, nearer the truth, +0.96 and +1.84 kW against the fit's −3.26 and −2.12 kW, +2.80 kW on the total against −5.37 kW, because the hall meter read low and the fit leans on it while the tiles ignore the meter. They are the less certain reading all the same, with standard deviations of 4.24 kW on a unit and 7.35 kW on their sum, against the fit's 2.50 and 2.77 kW. Their share, 0.51518, lies 0.00004 from the fit's.

6.5 The report (step 6)

The even split is not the best reading of the units. Unit 1 most likely rejects 104.4 kW, with a 95 % band of 99.5 to 109.3 kW, and unit 2 98.2 kW; the probability that unit 1 carries more than unit 2 is 0.93. The even split's 100.5 kW is still inside unit 1's band, and the share's band, 0.4950 to 0.5353, contains one half: the window favors unequal units without excluding equal ones.

6.5.1 The answers and their bands

Table 6.6 gives the answers in the order of Table 6.1. Each row is a question; the second column gives the answer with the draws and any exact form beside it, and the third its band: the central 95 % interval for a heat and for the share, and for the comparison the difference with its standard deviation.

6.5.2 The entry in the shift log

The operator's entry for 23:00 replaces the even split with the fitted heats: unit 1 most likely 104.4 kW, between 99.5 and 109.3 kW; unit 2 most likely 98.2 kW, between 93.3 and 103.1 kW; unit 1 most likely carrying 51.5 % of the coil heat, between 49.5 and 53.5 %, and more than unit 2 with probability 0.93. The entry also records what the numbers rest on. The two heats add up to the fitted load and the other heat, so they share the load meter's error, and the sum is known exactly as well as the load is, 2.77 kW. The split between them comes from the return and discharge temperatures, and the difference has a standard deviation of 4.16 kW, wider than

Table 6.5: The checks of the answers and their gaps from the reported values. The full state carries Q_1 and Q_2 as entries with six physical rows at the tolerance σ_p; the draws use 100 000 states from the Laplace posterior through the exact physics (6.2); the exact forms hold under the same normal; the tiles' row is their air-side arithmetic. Gaps are taken before rounding.

Route or reference	Value [kW or –]	Gap [kW or –]	sd [kW or –]	Gap [kW or –]	Probability [–]	Gap [–]
Heat of unit 1, Q_1						
Delta method, reported	104.3816	–	2.5010	–	–	–
Exact under the normal	104.3798	−0.0018	2.5004	−0.0006	–	–
Draws through the physics	104.3789	−0.0027	2.4968	−0.0042	–	–
Full state, $\sigma_p = 0.1$	104.3821	+0.0004	2.5023	+0.0013	–	–
Full state, $\sigma_p = 0.001$	104.3816	0.0000	2.5010	0.0000	–	–
Newton covariance	–	–	2.4981	−0.0029	–	–
Stated truth	107.6385	+3.2569	–	–	–	–
Dashboard (even split)	100.52	−3.8616	–	–	–	–
Units' tiles	108.60	+4.2184	–	–	–	–
Heat of unit 2, Q_2						
Delta method, reported	98.2459	–	2.4983	–	–	–
Exact under the normal	98.2477	+0.0018	2.4977	−0.0006	–	–
Draws through the physics	98.2566	+0.0107	2.4913	−0.0070	–	–
Full state, $\sigma_p = 0.1$	98.2466	+0.0007	2.4996	+0.0013	–	–
Full state, $\sigma_p = 0.001$	98.2459	0.0000	2.4983	0.0000	–	–
Newton covariance	–	–	2.4933	−0.0050	–	–
Stated truth	100.3615	+2.1156	–	–	–	–
Dashboard (even split)	100.52	+2.2741	–	–	–	–
Units' tiles	102.20	+3.9541	–	–	–	–
Unit 1 against unit 2, $Q_1 - Q_2$; probability $\Pr(Q_1 > Q_2)$						
Delta method, reported	6.1357	–	4.1621	–	0.92979	–
Exact under the normal	6.1321	−0.0037	4.1607	−0.0014	0.92963	−0.00016
Draws through the physics	6.1223	−0.0134	4.1469	−0.0152	0.92910	−0.00069
Full state, $\sigma_p = 0.1$	6.1355	−0.0003	4.1643	+0.0022	0.92967	−0.00011
Full state, $\sigma_p = 0.001$	6.1357	0.0000	4.1621	0.0000	0.92979	0.00000
Newton covariance	–	–	4.1539	−0.0082	0.93018	+0.00039
Stated truth	7.2770	+1.1413	–	–	–	–
Dashboard (even split)	0.00	−6.1357	–	–	–	–
Units' tiles	6.40	+0.2643	–	–	–	–
Share of unit 1, s_1						
Delta method, reported	0.51514	–	0.01027	–	–	–
Exact median under the normal	0.51514	0.00000	–	–	–	–
Draws through the physics	0.51511	−0.00003	0.01023	−0.00004	–	–
Full state, $\sigma_p = 0.1$	0.51514	0.00000	0.01028	+0.00001	–	–
Full state, $\sigma_p = 0.001$	0.51514	0.00000	0.01027	0.00000	–	–
Newton covariance	–	–	0.01025	−0.00002	–	–
Stated truth	0.51749	+0.00235	–	–	–	–
Dashboard (even split)	0.50	−0.01514	–	–	–	–
Units' tiles	0.51518	+0.00004	–	–	–	–
Central 95 % bands: reported; then each route with the gaps at the two ends						
Q_1 [kW]	[99.4797, 109.2836]; draws [99.4992, 109.3193], +0.0196, +0.0357					
	exact [99.5137, 109.3146], +0.0341, +0.0310					
Q_2 [kW]	[93.3493, 103.1425]; draws [93.4138, 103.1738], +0.0645, +0.0313					
	exact [93.3931, 103.1830], +0.0438, +0.0406					
$Q_1 - Q_2$ [kW]	[−2.0220, 14.2935]; draws [−2.0514, 14.2346], −0.0294, −0.0589					
	exact [−2.0297, 14.2795], −0.0077, −0.0140					
s_1 [–]	[0.495007, 0.535274]; draws [0.494934, 0.535134], −0.000073, −0.000140					
	exact [0.494993, 0.535239], −0.000014, −0.000035					

Table 6.6: The report on the unit heats of the two-unit hall at 23:00. Answers from the Laplace posterior of Chapter 5's fit by the delta method; bands are 95 % intervals.

Question	Answer	Band
How much heat is this unit actually rejecting?	unit 1: 104.38 kW, sd 2.50 kW (exact: 104.38 kW, sd 2.50 kW; draws: 104.38 kW, sd 2.50 kW)	99.48 to 109.28 kW, central 95 % (exact: 99.51 to 109.31 kW)
And the other unit?	unit 2: 98.25 kW, sd 2.50 kW (exact: 98.25 kW, sd 2.50 kW; draws: 98.26 kW, sd 2.49 kW)	93.35 to 103.14 kW, central 95 % (exact: 93.39 to 103.18 kW); correlation with unit 1 −0.386
Is unit 1 carrying more than unit 2?	probability 0.930 (exact: 0.930; draws: 0.929)	difference 6.14 kW, sd 4.16 kW, 95 % band −2.02 to 14.29 kW (draws: 6.12 kW, sd 4.15 kW, −2.05 to 14.23 kW)
What share of the hall's heat goes through unit 1?	0.5151, sd 0.0103 (exact median: 0.5151; draws: 0.5151, sd 0.0102)	0.4950 to 0.5353, central 95 % (exact: 0.4950 to 0.5352); contains 0.5, the rating's

it would be for two independent heats. The entry notes that the units' tiles stand about 4 kW above the fit on each unit, and that the fit leans on the one hall meter, so a low meter would pull both heats down. The tiles' air-side arithmetic is not entered as the reading: it has a standard deviation of 4.24 kW on a unit and 7.35 kW on the sum, against 2.50 and 2.77 kW.

Common mistake

Splitting the hall's heat between identical units by rating. The rating says what the units would share if their coils matched their nameplates; the window's temperatures say how the conductance actually divides between them tonight, and the fitted split is read from them. The even split here is 7.12 kW low for unit 1.

6.6 The next reading (step 7)

The report leaves unit 1's heat, the lead question's, with a band of 99.48 to 109.28 kW, made of the load, the two log-derates and the split between them. A further reading is one more row of the residual vector, and what it would remove is known before its value is.

6.6.1 The preposterior of a candidate reading

A candidate reading a is a row with Jacobian line $\mathbf{J}_a$, and its preposterior covariance is the rank-one form (5.19), with the width removed $1 - \mathrm{sd}_a(Q_1)/\mathrm{sd}(Q_1)$, where $\mathrm{sd}_a(Q_1) = \sqrt{\mathbf{g}_{Q_1}^{\mathsf{T}} \boldsymbol{\Sigma}_a \, \mathbf{g}_{Q_1}}$. Four candidates are open, and their rows say what each does to the entries. A water-side heat meter on unit 1, of standard deviation $\sigma_a = 3$ kW, reads Q_1 itself, $\mathbf{J}_a = \mathbf{g}_{Q_1}^{\mathsf{T}}/\sigma_a$, and the rank-one form gives $1/\mathrm{sd}_a^2 = 1/\mathrm{sd}^2 + 1/\sigma_a^2$: the virtual sensor and the meter add their precisions. The same meter on unit 2, $\mathbf{J}_a = \mathbf{g}_{Q_2}^{\mathsf{T}}/\sigma_a$, removes from $\mathrm{Var}(Q_1)$ only $\mathrm{Cov}(Q_1, Q_2)^2/(\sigma_a^2 + \mathrm{Var}(Q_2))$, a quantity proportional to the square of the correlation of the two heats: its row reaches Q_1 through the load and through the split, both by way of that correlation, and the two partly cancel in it (equation (6.11)). A second, more precise hall load meter of 1 kW, added beside the present one, is the row $\mathbf{e}_1^{\mathsf{T}}/1$ in the column of Q. A second discharge sensor on unit 1, of 0.3 K, repeats the unit-1 discharge row of $\mathbf{J}$ (5.11) and acts on (Q, g_1, g_2), most on g_1.

6.6.2 The candidates

Table 6.7 ranks the four. Each row is a reading the operator can request; the columns give its standard deviation, the standard deviation of Q_1 after it by the preposterior, and the share of the present 2.5010 kW it removes. Figure 6.5 draws the width removed as bars, ranked, the horizontal axis in percent.

Table 6.7: Candidate readings for the heat of unit 1, from the posterior at 23:00; the width removed is taken from $\mathrm{sd}(Q_1 \mid \mathbf{y}) = 2.5010\,\mathrm{kW}$.

Candidate	Its sd [kW or K]	$\mathrm{sd}_a(Q_1)$ [kW]	Width removed [%]
A water-side heat meter on unit 1	3	1.9210	23.2
A second hall load meter, added beside the present one	1	2.1337	14.7
A second discharge sensor on unit 1	0.3	2.1816	12.8
A water-side heat meter on unit 2	3	2.4234	3.1

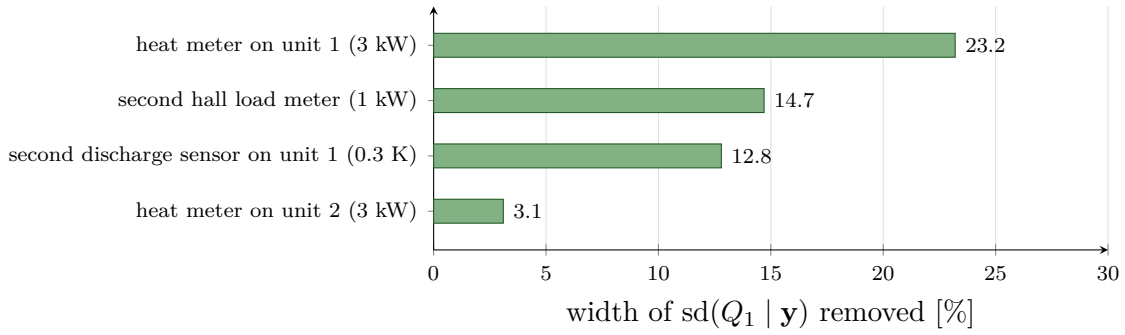


Figure 6.5: The candidate readings of Table 6.7 ranked by the width they would remove from the band of the heat of unit 1, by the preposterior (5.19).

The ranking follows the rows. The meter on unit 1 reads the quantity asked for, so it acts on the load, both log-derates and the split at once, and it removes the most. The added hall meter's row lies in the load's column and takes the load term from 2.0352 to 0.2348 kW^2; through $\boldsymbol{\Sigma}$ the other terms move too, g_1 from 1.4438 to 1.3190, g_2 from 1.6049 to 1.4524, the cross terms in Q from 1.0718 to 0.1236 and from -1.1844 to -0.1366, and the split's from 1.2836 up to 1.5595 kW^2, and the total falls by 1.7023 of 6.2550 kW^2. The discharge sensor acts mostly on g_1, and removes part of the split. The meter on unit 2 removes almost nothing from unit 1: its row reaches Q_1 only through the correlation of -0.386. For unit 2 the ranking is the mirror: its own meter removes 23.2 %, the added hall meter 14.6 %, a second discharge sensor on unit 2 12.7 % and the meter on unit 1 3.1 %.

6.6.3 The recommendation

The next reading is a water-side heat meter on unit 1, first in Table 6.7: it is the permanent instrument at the place of the question, reading the unit's heat rather than inferring it. A second hall load meter improves the heat of both units together, through the sum. A heat meter on unit 2 is not worth requesting for the heat of unit 1, though it is the corresponding instrument for unit 2's own question. Chapter 18 ranks the readings an operator can request on the shift plant in the same way.

6.6.4 What the analysis showed

The dashboard said the two units were sharing the heat equally, because they have the same rating; the analysis puts unit 1 at 104.4 kW and unit 2 at 98.2 kW as the most likely values, with probability 0.93 that unit 1 carries more, without excluding an even split; the stated truth is further apart still. The units' own tiles saw the inequality and, this night, stood nearer the truth than the fit, because the hall meter read low and the fit leans on it; they remain the less certain reading. The fit uses every reading and the balance: the meter and the temperatures fix the total heat, the temperatures fix how it divides, and because the readings know the split less well than the total, the two heats come out anti-correlated, the sum narrower and the difference wider than for two independent heats. It follows that a sensor that reads the hall's total, however good, cannot settle the split, and that a meter on one unit tells little about the other. The hall was instrumented to know its load and its temperatures; it does not know its units' heats except through them. Why the two coils divide the heat as they do is not a question of this chapter.

6.6.5 Limits of the model

The coil row without losses. The heat a unit's coil takes from the air is taken to be the heat it puts into the chilled-water loop (Section 6.1.2). Where piping breaks that, a water-side heat meter reads a different quantity from the one estimated here, and a comparison of the two would need a row for the difference.

The fan. A fan downstream of the coil warms the discharge air after the coil. The coil heat and the water-side heat still agree; what fails is the air-side row of equation (5.1), which then carries $C_u(T_r - T_{d,u}) = Q_u - Q_{\mathrm{fan},u}$, and the fitted log-derates absorb the difference. A water-side meter would read the coil heat, and it is the estimate that is biased, through the discharge rows.

The pinned airflow. The capacity rates C_u are pinned at the units' fixed fan speed. They set how the discharge temperatures inform the log-derates, and an error in them is in no band above: 5 % on one unit's rate moves the share by 1.08 to 1.22 of its standard deviation and the probability that unit 1 carries more from 0.93 to between 0.60 and 0.997 (Table 6.4). The even split's assumption of equal units is replaced here by the assumption of known airflow.

Steady state. The balance (6.3) holds only when the hall has settled. After a load step the hall's air and equipment store or release heat, the coils do not carry the load and the other heat alone, and the fitted heats misplace the difference.

Constructed readings. The five readings are the constructed ones of Section 5.1.4, not field data, and the agreement of the answers with the truth (Section 6.4.4) is a check of the method on its own model, not of the model against a real hall.

The next chapter. Chapter 7 installs a water-side heat meter on each unit's coil and asks what to do when the three meters do not balance.

6.7 What to retain

Definition 6.1 (Unmeasured heat flow)**.** The heat along an edge of the plant graph that no meter reads is a derived quantity $Q_u(\mathbf{z})$ of the fitted state: on the reduced state of Chapter 4 it is the coil row at the dependent temperatures, $Q_u = G_u(T_r - T_s)$; its value is $Q_u(\hat{\mathbf{z}})$, its variance

$\mathbf{g}_{Q_u}^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{g}_{Q_u}$ and its band $Q_u(\hat{\mathbf{z}}) \pm 1.960\sqrt{\mathbf{g}_{Q_u}^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{g}_{Q_u}}$. Where a balance row ties several flows to an entry, their gradients sum to that entry's gradient and the variance of their sum is that entry's variance. Written as shares of the total, $Q_u = s_u(Q + Q_o)$, their covariance is a positive term from the shared total and a negative term from the uncertain split: the flows are anti-correlated when the split is less certain than the total, their sum then narrower and their difference wider than for independent flows, and positively correlated when the split is the better known.

Three conclusions carry beyond the hall. A flow that no meter reads is known through the entries its gradient touches, and when those are entries the readings see, its band is made of readings rather than priors. A balance turns several such flows into one question: the band of each is not the whole answer, and the covariance among them decides comparisons between them. And a reading of the total, or of one of the flows, informs another flow only through their posterior covariance, $\mathbf{J}_a\boldsymbol{\Sigma}\mathbf{g}$, which the preposterior gives before the reading is taken.

Table 6.8 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 6.8: The terms of the unmeasured heat flow and where each is defined. Units for the two-unit hall.

Term	Expression	Units	Meaning	Defined in
Coil row	$Q_u = \bar{G}_u \mathrm{e}^{g_u}(T_r - T_s)$	kW	heat to the chilled water	6.1.2
Unit heat	$Q_u = s_u(Q + Q_o)$, eq. (6.2)	kW	a derived flow	6.1.2
Hall balance on the coils	$Q_1 + Q_2 = Q + Q_o$	kW	the identity of the null space	6.1.2
Load share	$s_1 = G_1/(G_1 + G_2)$	–	fraction through unit 1	6.1.2
Gradient pattern	eq. (6.5)	–	entries behind each answer	6.1.4
Block of $\boldsymbol{\Sigma}$	eq. (6.6)	mixed	all the answers consume	6.2
Questions and answers				
Gradient of a unit heat	eq. (6.7)	–, kW	shares and $Q_1Q_2/(Q + Q_o)$	6.3.1
Gradients of the sum	$\mathbf{g}_{Q_1} + \mathbf{g}_{Q_2} = \mathbf{g}_Q$	–	the balance on the gradients	6.3.1
Covariance of the heats	$\mathbf{g}_{Q_1}^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{g}_{Q_2}$, eq. (6.11)	kW^2	shared total against uncertain split	6.3.1
Comparison of the units	$\Phi\big((Q_1 - Q_2)/\mathrm{sd}\big)$, $\mathbf{g}_\Delta$	–	which unit carries more	6.3.2
Exact tail	eq. (6.15)	–	the same event on $g_1 - g_2$	6.3.2
Share variance	$s_1^2 s_2^2(\Sigma_{22} + \Sigma_{33} - 2\Sigma_{23})$	–	the split's band	6.3.3
Checks and the next reading				
Full state with the heats	twelve entries, six physics rows	kW, K	the limit $\sigma_p \to 0$	6.4.1
Preposterior	$\boldsymbol{\Sigma}_a = (\boldsymbol{\Sigma}^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a)^{-1}$	mixed	band after a candidate reading	6.6.1
Width removed	$1 - \mathrm{sd}_a/\mathrm{sd}$	–	ranks the candidates	6.6.2

6.8 Exercises

Exercise 6.1. *Formulation.* The chilled-water plant asks for the total heat the two units reject. Write the question as a probability over the reduced state, give the gradient of the quantity, and state its answer from the fit. ($\Pr(Q_1 + Q_2 \le q \mid \mathbf{y})$; by equation (6.8) the gradient is $(1,\ 0,\ 0,\ 0,\ 0,\ 0)^{\mathsf{T}}$, so the answer is the fitted load plus the other heat with the load's spread: $Q_1 + Q_2 = 194.6275 + 8.0 = 202.6275\,\mathrm{kW}$, standard deviation $2.7694\,\mathrm{kW}$.)

Exercise 6.2. *Derivation.* Derive $\partial Q_1/\partial g_1$ from the coil row (6.1) and the hall balance on the reduced state, and check it against Table 6.2. ($Q_1 = (Q + Q_o)\,G_1/(G_1 + G_2)$ with $\partial G_1/\partial g_1 = G_1$, so $\partial Q_1/\partial g_1 = (Q + Q_o)\,G_1G_2/(G_1 + G_2)^2 = Q_1 s_2 = Q_1Q_2/(Q + Q_o)$, which is $104.3816 \times 98.2459/202.6275$ at the MAP state, the 50.6104 kW of the table.)

Exercise 6.3. *Derivation.* Show from the gradients that the covariance of the two unit heats is negative, and give it. (By equation (6.8), $\mathrm{Var}(Q_1) + \mathrm{Var}(Q_2) + 2\,\mathrm{Cov}(Q_1, Q_2) = \mathbf{g}_Q^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_Q = \Sigma_{11}$, so $\mathrm{Cov}(Q_1, Q_2) = \frac{1}{2}(\Sigma_{11} - \mathrm{Var}(Q_1) - \mathrm{Var}(Q_2))$, negative whenever the two variances together exceed the load's, which the balance alone does not guarantee; by equation (6.11) it is the shared load's 1.9156 less the split's 4.3323 plus $0.0033\,\mathrm{kW}^2$, negative because the split is the less certain: $-2.4134\,\mathrm{kW}^2$, a correlation of -0.3863.)

Exercise 6.4. *Application.* The facilities team proposes a water-side heat meter of 3 kW on unit 1, arguing that with the hall balance it will also give unit 2's heat. Explain from the preposterior how much it narrows the band of unit 2. (Its row is $\mathbf{J}_a = \mathbf{g}_{Q_1}^{\mathsf{T}}/3$, and by equation (5.19) it removes from $\mathrm{Var}(Q_2)$ only $\mathrm{Cov}(Q_1, Q_2)^2/(3^2 + \mathrm{Var}(Q_1))$, proportional to the square of the correlation -0.3863: the balance gives $Q_2 = Q + Q_o - Q_1$, and the load itself only narrows from 2.7694 to 2.5888 kW, so $\mathrm{sd}(Q_2)$ goes from 2.4983 to 2.4206 kW, 3.1 % of the width.)

Exercise 6.5. *Application.* The operator reports the share of unit 1 instead of its heat. Explain how far the hall meter's error reaches the share, and give its band. (By equation (6.16) the gradient of s_1 is $s_1 s_2(0,\ 1,\ -1,\ 0,\ 0,\ 0)^{\mathsf{T}}$, with a zero in the column of Q: the meter's error reaches the share only through $\mathrm{Cov}(Q, s_1)$, a correlation of -0.0190, and setting the meter at the true load moves the share by -0.04 of its standard deviation. The share depends on $g_1 - g_2$ alone, and its variance is the variance of that difference times $s_1^2 s_2^2$, of which the cross term is 29.6 %; the share is 0.51514 with standard deviation 0.01027, and its band, which contains one half, is $[0.4950,\ 0.5353]$.)

Chapter 7

Reconciling Inconsistent Meters

7.1 Three meters that do not balance

It is 23:00 in the two-unit hall of Chapter 5, and each cooling unit now carries a water-side heat meter on its coil. The two unit meters read 106.45 and 110.15 kW, 216.60 kW together. The hall load meter reads 193.04 kW, and with the 8.0 kW of other heat the coils should be carrying 201.04 kW. The balance is off by 15.56 kW, and its spread from the three meters' own errors alone is 5.20 kW, so the mismatch stands 2.99 of its standard deviations from zero. The dashboard turns its meter-mismatch flag red. The flag's work-order template asks which meter to send for recalibration, and the reading the flag invites is that one of the three meters is off by the whole 15.56 kW.

The flag cannot say which. A unit-1 meter that reads 15.56 kW high, a unit-2 meter that reads 15.56 kW high and a hall meter that reads 15.56 kW low each clear the balance on their own, and the balance is one equation in three suspects. What the balance does not use is the rest of the hall: the return air, the two discharge temperatures and rack 1's inlet, which Chapter 5 fitted, also know how much heat each coil carries and how much the racks put in. The question of this chapter is

which of the three meters is off, how sure that is, by how much it is off, and which check would settle it.

It is answered on the reduced state of Chapter 5, with two reading rows added. What changes is the computation: the answer is not a band on the state of one drawing of the hall but a ranking of four drawings by their evidence, one with every meter sound and one with a bias freed on each meter in turn.

Jargon: water-side heat meter, meter balance, bias, bad-data detection, chi-square test, leverage, redundancy number, normalized residual, fault detection and isolation, candidate fault, Bayes factor, Occam factor, clamp-on meter, UPS

A *water-side heat meter* reads the heat a coil passes to the chilled water from the water's flow and its temperature rise across the coil. The *meter balance* sets meters that should sum to one another against each other; here the two unit meters against the hall meter plus the other heat. A meter's *bias* b_j is a constant offset of its reading from the quantity it reads. *Bad-data detection* tests whether the readings of one window are consistent with one another under the model, by the *chi-square test* of the whole residual vector at the fit and by the *largest normalized residual*. The *leverage* h_{kk} of row k is the share of its own reading in its fitted value, and $1 - h_{kk}$, its *redundancy number*, is the variance of its residual at the fit; the *normalized residual* is the residual divided by the root of the redundancy number. *Fault detection and isolation* decides that something is wrong and where; in the framework it is the diagnosis computation, a family of *candidate faults* generated from the plant graph and ranked by their evidence. The *Bayes factor* of two drawings is the ratio of their evidences, and the *Occam factor* is the part of it that charges a drawing for an entry it adds. A *clamp-on meter* reads a current or a power without breaking the circuit; the *UPS*, the uninterruptible power supply, feeds the racks, so its output is the IT load.

Key idea

A meter balance that does not close names a set of suspects, not a culprit. Each suspect becomes a drawing of the plant with a bias freed on that one meter, and the drawings are ranked by their evidence, the fit each buys against the price of the entry it adds. The rest of the plant's readings decide the ranking, and when two drawings explain the window almost equally well, the answer is the two of them with their probabilities and the check that would separate them.

7.1.1 A meter balance as a set of suspects

The balance of the three meters is a derived quantity of the readings, $y_{W_1} + y_{W_2} - y_Q - Q_o$. If every meter reads its quantity with an independent error of standard deviation 3 kW, the balance has mean zero and the standard deviation $\sqrt{3 \times 3^2} = 5.20$ kW that the dashboard divides by. A large balance says that the three readings are not all right. It says nothing about which is wrong, because a bias on any one of them enters the balance with the same weight, and the balance is a single number.

The hall has more to say than the balance. Each meter reads a quantity that the fitted state of Chapter 5 also carries: the hall meter reads the load Q, an entry of the state, and each unit meter reads the heat Q_u of its coil, a derived quantity of the state that the return air and the discharge temperatures inform through the coil rows. A bias on one meter leaves the other rows where they are and pulls the fit toward that meter's reading only as far as the other rows allow, so each suspect leaves its own pattern in the residuals of the fit. The diagnosis compares those patterns as drawings of the plant. Each drawing is the hall of Chapter 5 with one more entry, the bias of one meter, and its evidence is the probability of the window's readings with that entry integrated out over its prior. The evidence pays a drawing for the misfit its bias removes and charges it for the spread of the bias it had to admit, so a drawing that frees a bias where none is needed loses to the drawing without one.

> **Principle 7.1.** *A meter is isolated as faulty by ranking the drawings of the plant that free a bias on one meter at a time, by their evidence, against the drawing with every meter sound. The answer is a probability per meter, not a verdict, and the readings that are not meters decide it.*

7.1.2 The two-unit hall with its heat meters (step 0)

The plant is the two-unit hall of Chapter 5, Figure 5.1, unchanged: the seven physical rows (5.1), the pinned chilled-water supply $T_s = 14°\mathrm{C}$, the ratings $\bar{G}_u = 7.0\,\mathrm{kW/K}$, the capacity rates $C_u = 10.0\,\mathrm{kW/K}$ and the other heat $Q_o = 8\,\mathrm{kW}$, and the reduced state (5.3),

$$\mathbf{z} = (Q,\ g_1,\ g_2,\ \beta_1,\ \beta_2,\ \beta_3)^{\mathsf{T}}, \tag{7.1}$$

on which the dependent temperatures follow by the forward chain (5.4). Its dependent block $\mathbf{A}_{\mathrm{dep}}$ is triangular with nonzero diagonal in the order of that chain, so it is invertible and the reduced state of Chapter 4 is valid here as there. What is new is two quantities the meters read. The heat that coil u passes to the chilled water is the coil row of Chapter 5,

$$Q_u(\mathbf{z}) = G_u\,(T_r - T_s) = \frac{G_u}{G_1 + G_2}\,(Q + Q_o), \qquad G_u = \bar{G}_u\,\mathrm{e}^{g_u}, \quad u = 1, 2, \tag{7.2}$$

the second form by the hall balance, which fixes $T_r - T_s = (Q + Q_o)/(G_1 + G_2)$. Two structural facts of the null space are worth having in view. The two unit heats sum to the load and the other heat,

$$Q_1(\mathbf{z}) + Q_2(\mathbf{z}) = Q + Q_o, \tag{7.3}$$

at every point of the set on which the balances hold: this is the balance the meters should close, and on the reduced state it is a property of the rows, not a row of its own. And the split between the units, $Q_u/(Q + Q_o) = G_u/(G_1 + G_2)$, depends on the log-derates alone, so a unit meter informs the split and the load together while the hall meter informs the load alone.

Modeling choices at the graph. The hall is taken in steady state within the window, its airflow pinned and no controller acting, as in Chapter 5: these are what make the forward chain (5.4) hold within the window, and they fail after a load step or a unit start (Section 7.6.4). A water-side heat meter is taken to read its coil's heat Q_u, the heat the air gives to the water; this is what ties each unit meter to the coil row, and it fails if the meter's flow or temperature sensors sit where the chilled water of the two units mixes, so that one meter sees part of the other unit's heat.

The graph. Figure 7.1 draws the part of the hall's graph that the meters see. The IT load Q at the upper left (green, free) and the pinned other heat Q_o at the lower left (gray) enter the return air (blue), whose temperature the hall balance fixes; the return feeds the two coils (green, each carrying its free log-derate), which pass the heats Q_1 and Q_2 to the pinned chilled water on the right. The three meters leave the graph as dashed arrows: the hall meter from the load, and each unit meter from its coil's heat path. The red box states the balance that the three should close. The other four readings, rack 1's inlet among them, and the three racks are those of Figure 5.1, drawn there and listed here in the note.

7.1.3 The operator's questions (step 1)

The operator asks four questions of the 23:00 window, and step 1 writes each as a probability over the state of a drawing of the hall, with the unknown on the left of the bar and the seven readings and five priors on the right. Table 7.1 lists them, with the columns of Chapter 5's table: 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, with M_j a drawing of the family, $f(\mathbf{z})$ a derived quantity, f^* a level of it and $U(a)$ the value of a candidate check a; the mathematical form for this hall; the category; the computation of Section 2.3.6 that answers it; and the subsection where it is answered. The first is a diagnosis over four drawings: M_0, every meter sound, and

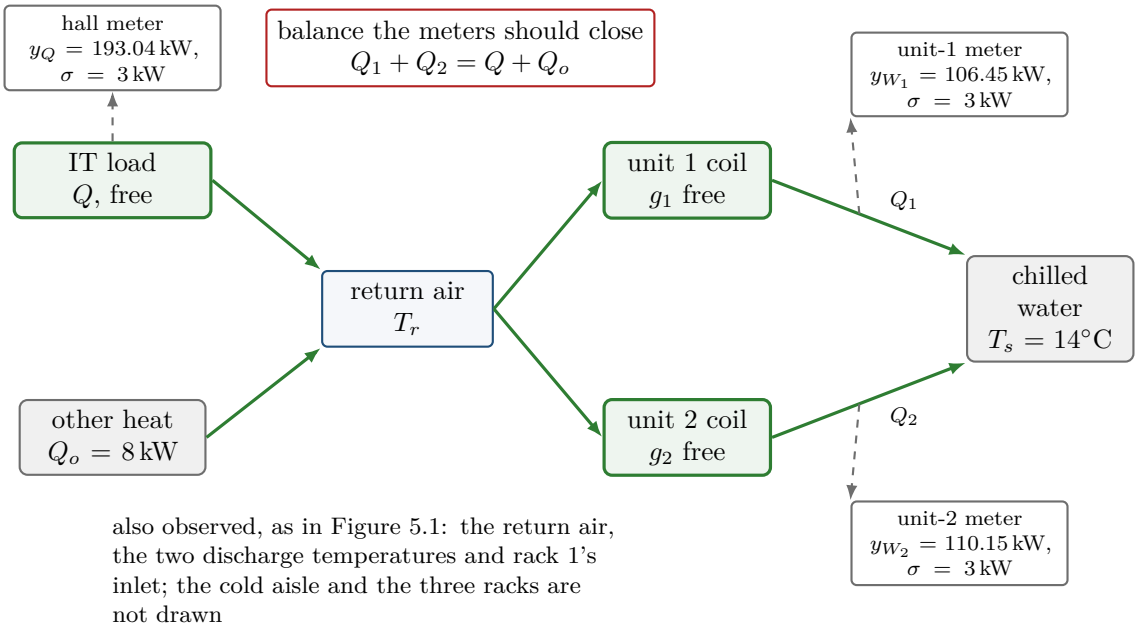


Figure 7.1: The part of the two-unit hall's graph that the three meters see, at 23:00. Free (green): the IT load Q and the log-derates g_1, g_2 of the coils. Blue: the return air, fixed by the hall balance. Pinned (gray): the other heat $Q_o = 8\,\text{kW}$ and the chilled-water supply $T_s = 14°\text{C}$. Observed (dashed): the hall meter on Q and the water-side meters on the coil heats Q_1, Q_2 of equation (7.2), each of standard deviation 3 kW. Red: the balance (7.3), which holds on the reduced state. The other readings, the racks and their fractions are as in Figure 5.1.

M_Q, M_{W_1}, M_{W_2}, a bias on the hall meter, the unit-1 meter or the unit-2 meter. The second asks whether the window is consistent with M_0 at all, a tail of the predictive of a statistic of the residuals, which Chapter 2 files under prediction under a threshold. The third estimates the bias within each of the leading drawings. The fourth is a design question, the value of a check before it is made; Section 7.6 then turns its answer into the request.

7.1.4 The evidence as rows and the candidate faults (step 2)

Step 2 attaches the evidence of the 23:00 window to the graph. The five readings and five priors of Chapter 5 are unchanged: the return air, the two discharge temperatures and rack 1's inlet, each of standard deviation 0.3 K; the hall meter $y_Q = 193.04\,\text{kW}$ of standard deviation 3 kW; the coil priors $g_u \sim \mathcal{N}(0,\ 0.15^2)$; and the survey's priors on the three fractions, of standard deviation 0.08. Two reading rows are added, one per unit meter,

$$r_{W_u}(\mathbf{z}) = \frac{Q_u(\mathbf{z}) - y_{W_u}}{3}, \qquad u = 1, 2, \qquad y_{W_1} = 106.45\,\text{kW}, \quad y_{W_2} = 110.15\,\text{kW}, \tag{7.4}$$

each Gaussian with standard deviation 3 kW and independent of the other readings given the state, the form of every reading row of Chapter 4. The load still has no prior row.

The candidate faults. The family is generated from the graph: one candidate per meter, since a meter is a dashed edge that a bias can sit on. Candidate M_j, for $j \in \{Q, W_1, W_2\}$, is the hall with the bias b_j of meter j freed as a seventh entry of the state, added to that meter's predicted reading, with the prior row b_j/σ_b, $b_j \sim \mathcal{N}(0,\ \sigma_b^2)$, $\sigma_b = 10\,\text{kW}$; M_0 has every bias at zero. The prior probability of a meter fault in the window is taken as one half, shared equally among the three meters,

$$\Pr(M_0) = 0.5, \qquad \Pr(M_Q) = \Pr(M_{W_1}) = \Pr(M_{W_2}) = 0.5/3. \tag{7.5}$$

Table 7.1: The questions asked of the two-unit hall at 23:00. $\mathbf{y}$ is the seven readings and the five priors; M_0 the drawing with every meter sound; M_Q, M_{W_1}, M_{W_2} a bias freed on one meter; b_j that bias; ν the reading rows' share of the degrees of freedom; a a candidate check.

Operator question	General form	Mathematical form	Category	Computation	Answered in
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 $\lvert r_k^{\text{N}}\rvert$	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

Three modeling choices enter the family itself. One meter is at fault at a time: this keeps the family to four drawings of at most one added entry each, and it cannot explain two meters biased together (Section 7.6.4). The bias is constant within the window: one window gives one value per reading, so a constant is all the window can see, and a bias that drifts across windows would need a run of windows. The width σ_b of the bias prior is a choice of what size of meter error is plausible, and the evidence depends on it; Section 7.4.4 varies it from 5 to 20 kW.

Which row acts on which entry. By equation (7.2) each unit meter sees (Q, g_1, g_2), the entries of the return air and the discharge rows. With $*$ for an entry that can be nonzero, the 12×6 Jacobian of M_0 is the pattern (5.6) of Chapter 5 with two rows inserted after rack 1's inlet:

$$\mathbf{J} = \begin{bmatrix} * & * & * & 0 & 0 & 0 \\ * & * & * & 0 & 0 & 0 \\ * & * & * & 0 & 0 & 0 \\ * & 0 & 0 & 0 & 0 & 0 \\ * & * & * & * & 0 & 0 \\ * & * & * & 0 & 0 & 0 \\ * & * & * & 0 & 0 & 0 \\ 0 & * & 0 & 0 & 0 & 0 \\ 0 & 0 & * & 0 & 0 & 0 \\ 0 & 0 & 0 & * & 0 & 0 \\ 0 & 0 & 0 & 0 & * & 0 \\ 0 & 0 & 0 & 0 & 0 & * \end{bmatrix} \begin{matrix} \text{return air } y_{T_r} \\ \text{discharge } y_{T_{d,1}} \\ \text{discharge } y_{T_{d,2}} \\ \text{hall meter } y_Q \\ \text{rack-1 inlet } y_{T_1} \\ \text{unit-1 meter } y_{W_1} \\ \text{unit-2 meter } y_{W_2} \\ \text{prior on } g_1 \\ \text{prior on } g_2 \\ \text{prior on } \beta_1 \\ \text{prior on } \beta_2 \\ \text{prior on } \beta_3 \end{matrix}, \tag{7.6}$$

$$\frac{\partial \mathbf{r}}{\partial b_Q} = \frac{\mathbf{e}_4}{3}, \qquad \frac{\partial \mathbf{r}}{\partial b_{W_1}} = \frac{\mathbf{e}_6}{3}, \qquad \frac{\partial \mathbf{r}}{\partial b_{W_2}} = \frac{\mathbf{e}_7}{3},$$

with columns $Q, g_1, g_2, \beta_1, \beta_2, \beta_3$ and $\mathbf{e}_k$ the k-th unit vector of the twelve rows. The second line is the column each candidate adds, one nonzero in its meter's row, together with its prior row. Two facts are read here before any number. The three meter rows and the four temperature rows share the columns of (Q, g_1, g_2), so the three meters are checked against one another and against the temperatures through those three entries, and not through the balance alone. And a bias column $\mathbf{e}_k/3$ is observable from the readings only to the extent that it does not lie in the span of the columns of $\mathbf{J}$: a reading that no other row can predict would carry a bias that

no fit could see. The share of $\mathbf{e}_k$ outside that span is the redundancy number of row k, which Section 7.2.4 defines.

7.2 The posterior of the state (step 3)

All four questions of Table 7.1 are read from the posteriors of four drawings: the reduced state under the twelve rows of step 2, and the same state with the bias of one meter freed, once per meter. The inputs are the seven readings and five priors of Section 7.1.4, the pinned quantities of Section 7.1.2 and the bias prior $\sigma_b = 10\,\mathrm{kW}$.

7.2.1 The rows of the four drawings

The rows of M_0 are the ten rows of Chapter 5 with the two meter rows (7.4) inserted after rack 1's inlet, in the order of the pattern (7.6), each divided by its standard deviation,

$$\mathbf{r}(\mathbf{z}) = \begin{bmatrix} (T_r(\mathbf{z}) - y_{T_r})/0.3 \\ (T_{d,1}(\mathbf{z}) - y_{T_{d,1}})/0.3 \\ (T_{d,2}(\mathbf{z}) - y_{T_{d,2}})/0.3 \\ (Q - y_Q)/3 \\ (T_1(\mathbf{z}) - y_{T_1})/0.3 \\ (Q_1(\mathbf{z}) - y_{W_1})/3 \\ (Q_2(\mathbf{z}) - y_{W_2})/3 \\ g_1/0.15 \\ g_2/0.15 \\ (\beta_1 - 0.20)/0.08 \\ (\beta_2 - 0.35)/0.08 \\ (\beta_3 - 0.50)/0.08 \end{bmatrix}, \qquad \mathcal{L}_0(\mathbf{z}) = \tfrac{1}{2}\,\mathbf{r}(\mathbf{z})^{\mathsf{T}}\mathbf{r}(\mathbf{z}). \tag{7.7}$$

Drawing M_j adds the bias to its meter's row and appends its prior row,

$$\mathbf{r}_j(\mathbf{z}, b_j) = \begin{bmatrix} \mathbf{r}(\mathbf{z}) + b_j\,\mathbf{e}_{k(j)}/3 \\ b_j/\sigma_b \end{bmatrix}, \qquad \mathcal{L}_j = \tfrac{1}{2}\,\mathbf{r}_j^{\mathsf{T}}\mathbf{r}_j, \tag{7.8}$$

with $k(j)$ the row of meter j: the fourth for the hall meter, the sixth and seventh for the unit meters. Twelve rows act on six entries in M_0, thirteen on seven in each M_j. The reduced Jacobian is $\mathbf{J} = \mathbf{J}_{\text{rows}}\mathbf{N}$ of equation (4.47); a meter row's line is the gradient of Q_u divided by 3, which by equation (7.2) is

$$\frac{\partial Q_1}{\partial \mathbf{z}} = \Big(\frac{G_1}{G_1+G_2},\ \frac{G_1G_2\,(T_r - T_s)}{G_1+G_2},\ -\frac{G_1G_2\,(T_r - T_s)}{G_1+G_2},\ 0,\ 0,\ 0\Big)^{\mathsf{T}}, \tag{7.9}$$

and the same for Q_2 with the signs of the two log-derate entries exchanged. The two meters' lines are equal and opposite in the log-derates and add to $(1, 0, \ldots, 0)$ in the load, the balance (7.3) read off the Jacobian. A bias column adds $\mathbf{e}_{k(j)}/3$ and, in the prior row, $1/\sigma_b$. The meter rows are nonlinear in the log-derates through $G_u/(G_1+G_2)$, so $\mathbf{J}$ is evaluated at each iterate.

The other heat freed. The checks and the report also use one fit outside the family. Drawing M_{Q_o} holds every meter sound and frees the pinned other heat as a seventh entry: its rows are those of equation (7.7) with Q_o an entry of the forward chain (5.4), and one more prior row, $(Q_o - 8)/10$. The gradient of Q_u gains the entry $G_u/(G_1+G_2)$ in the column of Q_o.

The check rows. A check a of standard deviation σ_a on a quantity $f_a(\mathbf{z})$, once read as y_a, is one more row of a drawing,

$$r_a(\mathbf{z}) = \frac{f_a(\mathbf{z}) - y_a}{\sigma_a}, \tag{7.10}$$

with $f_a = Q$ for a clamp-on reading of the UPS output and $f_a = Q_2$ for a portable water-side heat check of unit 2, whose Jacobian line is the gradient of Q_2 of equation (7.9) divided by σ_a.

7.2.2 The normal equations and the covariance

Each drawing's MAP state minimizes its objective by the normal equations of Chapter 5, equation (5.8), on its own rows,

$$(\mathbf{J}_j^{\mathsf{T}}\mathbf{J}_j)\,\Delta\mathbf{z} = -\mathbf{J}_j^{\mathsf{T}}\mathbf{r}_j, \qquad \mathbf{\Sigma}_j = (\mathbf{J}_j^{\mathsf{T}}\mathbf{J}_j)^{-1}\big|_{\hat{\mathbf{z}}_j}, \tag{7.11}$$

started for M_0 from the state of Chapter 5, the meter's load, the coils at their ratings and the survey's fractions, and for each M_j from $\hat{\mathbf{z}}_0$ with $b_j = 0$. With every row normalized the information is $\mathbf{\Lambda}_j = \mathbf{J}_j^{\mathsf{T}}\mathbf{J}_j$. The Gauss–Newton covariance drops the curvature term of the nonlinear rows; Section 7.4.1 sets the full Hessian beside it, and for the evidence, which integrates the whole objective, Section 7.4.2 measures what that costs.

The block structure. The pattern (7.6) keeps the shape of Chapter 5's information, equation (5.10). No meter row touches a fraction, so the columns of β_2 and β_3 still hold their prior rows alone, their posterior is the survey's in every drawing, and their entries in $\mathbf{\Sigma}_j$ are the survey's 0.08^2 with no covariance; a bias column couples only to (Q, g_1, g_2), through its meter's row, and to β_1 only through those entries. The diagnosis is therefore decided in the four entries the readings see, (Q, g_1, g_2, β_1), and the bias: the fractions of racks 2 and 3, which carried the band of Chapter 5, play no part in it.

7.2.3 The fits of the four drawings

The fit of M_0 ends with $2\mathcal{L}_0(\hat{\mathbf{z}}_0) = 12.9934$ on 12 rows and 6 entries, at the load $\hat{Q} = 198.796\,\text{kW}$ and the log-derates $\hat{g}_1 = -0.0854$ and $\hat{g}_2 = -0.0982$. The two unit meters pull the load above the hall meter's 193.04 kW: with them in the fit the hall carries more heat than its own meter says. The three bias drawings end at $2\mathcal{L}_j(\hat{\mathbf{z}}_j) = 5.5196$ for the hall meter, 12.0808 for the unit-1 meter and 4.2201 for the unit-2 meter, each on 13 rows and 7 entries. Freeing the bias on either the hall meter or the unit-2 meter removes most of the misfit; freeing it on the unit-1 meter removes little.

The iterations. Table 7.2 runs the normal equations (7.11) for M_0 from the starting state of Chapter 5 and for M_{W_2} from $\hat{\mathbf{z}}_0$ with $b_{W_2} = 0$. Each row is an iteration; the columns give each fit's objective and the norm of its gradient $\mathbf{J}^{\mathsf{T}}\mathbf{r}$. The fit of M_0 falls from 73.4607 to 6.4967 in two steps; the fit of M_{W_2} starts where M_0 ended and falls to 2.1100, half of the $2\mathcal{L} = 4.2201$ above. Both gradients reach 10^{-8} by the fifth iteration.

The MAP states. Table 7.3 gives the MAP state of each drawing: each row is an entry, each column a drawing, with the spread $\sqrt{(\mathbf{\Sigma}_j)_{ii}}$ in parentheses. The fractions of racks 2 and 3 sit at the survey's values and spreads in every drawing, as the block structure of Section 7.2.2 requires. The drawings part most on the load: 205.778 kW when the hall meter carries a bias, 195.620 kW when the unit-2 meter does.

Table 7.2: Gauss–Newton iterations of the drawing with every meter sound and of the unit-2 drawing, equation (7.11). M_0 starts from $(193.04,\ 0,\ 0,\ 0.20,\ 0.35,\ 0.50)^{\mathsf{T}}$; M_{W_2} from $\hat{\mathbf{z}}_0$ with $b_{W_2} = 0$.

Iteration k [–]	$\mathcal{L}_0(\mathbf{z}^{(k)})$ [–]	$\lVert\mathbf{J}^{\mathsf{T}}\mathbf{r}\rVert$, M_0 [–]	$\mathcal{L}_{W_2}(\mathbf{z}^{(k)})$ [–]	$\lVert\mathbf{J}^{\mathsf{T}}\mathbf{r}\rVert$, M_{W_2} [–]
0	73.4607	7.9×10^{2}	6.4967	8.2×10^{-1}
1	6.7873	6.0×10^{1}	2.1113	3.6×10^{0}
2	6.4967	2.4×10^{-1}	2.1100	3.1×10^{-3}
3	6.4967	1.8×10^{-4}	2.1100	7.8×10^{-6}
4	6.4967	2.0×10^{-7}	2.1100	1.6×10^{-8}
5	6.4967	2.5×10^{-8}	2.1100	7.6×10^{-9}

Table 7.3: The MAP state of each drawing at 23:00, with its spread in parentheses.

Entry [unit]	M_0	M_Q	M_{W_1}	M_{W_2}
Q [kW]	198.7964 (2.3193)	205.7782 (3.4525)	197.7703 (2.5564)	195.6199 (2.5556)
g_1 [–]	−0.0854 (0.0188)	−0.0670 (0.0194)	−0.0964 (0.0222)	−0.0690 (0.0195)
g_2 [–]	−0.0982 (0.0190)	−0.0794 (0.0196)	−0.0929 (0.0197)	−0.1332 (0.0231)
β_1 [–]	0.2223 (0.0309)	0.2252 (0.0300)	0.2218 (0.0310)	0.2209 (0.0313)
β_2 [–]	0.3500 (0.0800)	0.3500 (0.0800)	0.3500 (0.0800)	0.3500 (0.0800)
β_3 [–]	0.5000 (0.0800)	0.5000 (0.0800)	0.5000 (0.0800)	0.5000 (0.0800)
b_j [kW]	–	−11.6864 (4.2766)	3.4359 (3.5980)	10.6494 (3.5969)

The posterior under the unit-2 drawing. The covariance of M_{W_2} has zeros in the rows and columns of β_2 and β_3 beside their 0.0064 on the diagonal. The rest of it, in the order $(Q, g_1, g_2, \beta_1, b_{W_2})$, is

$$\boldsymbol{\Sigma}_{W_2} = \begin{bmatrix} 6.531142 & 0.008286 & 0.028911 & 0.002798 & -3.862044 \\ 0.008286 & 0.000379 & -0.000128 & 0.000157 & 0.019975 \\ 0.028911 & -0.000128 & 0.000532 & 0.000166 & -0.043604 \\ 0.002798 & 0.000157 & 0.000166 & 0.000979 & -0.001643 \\ -3.862044 & 0.019975 & -0.043604 & -0.001643 & 12.937813 \end{bmatrix},$$

in kW^2, kW and dimensionless units. The correlations of the entries that the diagnosis turns on, in the order (Q, g_1, g_2, b_{W_2}), are

$$\begin{bmatrix} 1 & 0.1665 & 0.4907 & -0.4201 \\ 0.1665 & 1 & -0.2850 & 0.2852 \\ 0.4907 & -0.2850 & 1 & -0.5258 \\ -0.4201 & 0.2852 & -0.5258 & 1 \end{bmatrix}.$$

They are read by sign. The unit-2 meter row reads $Q_2(\mathbf{z}) + b_{W_2}$, so a larger bias is explained equally well by a smaller unit-2 heat, which a smaller g_2 or a smaller load gives: the bias trades against unit 2's log-derate, −0.53, and against the load, −0.42, and with g_1 it is positively correlated, 0.29, because a larger g_1 takes heat from unit 2. The load and the two log-derates keep the signs of Chapter 5, which the return-air reading's hold on $(Q + Q_o)/(G_1 + G_2)$ gives them.

The meter rows' Jacobian. At the MAP state of M_{W_2} the two meter rows of $\mathbf{J}$, in the columns $(Q, g_1, g_2, \beta_1, \beta_2, \beta_3, b_{W_2})$, are

$$\begin{bmatrix} 0.1720 & 16.9508 & -16.9508 & 0 & 0 & 0 & 0 \\ 0.1613 & -16.9508 & 16.9508 & 0 & 0 & 0 & 0.3333 \end{bmatrix} \begin{matrix} \text{unit-1 meter} \\ \text{unit-2 meter'} \end{matrix}$$

the gradient (7.9) divided by the meters' 3 kW. The log-derate entries are equal and opposite, $16.9508 \times 3 = 50.85$ kW for $G_1G_2(T_r - T_s)/(G_1 + G_2)$; the load entries 0.1720 and 0.1613 add to the 1/3 of the balance (7.3); and the bias column is the 1/3 of $\mathbf{e}_7/3$, in the unit-2 row alone.

The fit of the other heat. M_{Q_o} ends at $2\mathcal{L} = 5.5196$ with $Q = 194.0918$ (2.8876) kW and $Q_o = 19.6864$ (4.2766) kW, its log-derates and fractions those of M_Q. The two drawings are one fit: the substitution $b_Q = 8 - Q_o$, $Q' = Q + b_Q$ carries one into the other with the priors of equal width, so no window of these rows can separate them; only a row on Q or on Q_o alone can. Their log evidences are equal, -12.2294.

The refits with a check row. Table 7.4 gives the fits with a check row (7.10) added, $\sigma_a = 1$ kW for the clamp-on reading and 2 kW for the portable check, read at the values of Section 7.3.4. Each row is a first reading; the columns give what each refit then predicts for the other check, with its own σ_a in the spread.

Table 7.4: The refits with one check row: the predicted reading of the other check, mean (sd) in kW, under each drawing.

First reading [kW]	Other check	M_{W_2} [kW]	M_Q [kW]	M_{Q_o} [kW]
Portable 98.54	clamp-on	195.62 (2.38)	200.80 (3.01)	–
Portable 100.36	clamp-on	196.42 (2.38)	201.98 (3.01)	–
Portable 102.40	clamp-on	197.32 (2.38)	203.30 (3.01)	–
Portable 106.23	clamp-on	199.00 (2.38)	205.78 (3.01)	194.09 (3.05)
Clamp-on 195.62	portable	98.54 (2.68)	101.54 (2.54)	106.31 (3.03)
Clamp-on 200.0	portable	100.99 (2.69)	103.56 (2.54)	106.54 (3.03)
Clamp-on 200.7	portable	101.38 (2.69)	103.88 (2.54)	106.57 (3.03)
Clamp-on 205.78	portable	104.22 (2.69)	106.23 (2.54)	106.84 (3.03)

7.2.4 The residuals at the fit

The residual tests read the fit of M_0. One choice enters here: the rows are taken as linear over the spread of the fit, which makes the residuals at the fit Gaussian with a covariance the Jacobian gives; it holds for the meter and temperature rows to the extent that the fit moves little when the errors are redrawn, which the agreement of the Laplace and sampled evidences of Section 7.4.2 bears on.

For linear rows the residual vector at the MAP state is a projection of the rows' normalized errors $\boldsymbol{\zeta} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$,

$$\mathbf{r}(\hat{\mathbf{z}}_0) = -(\mathbf{I} - \mathbf{H})\,\boldsymbol{\zeta}, \qquad \mathbf{H} = \mathbf{J}(\mathbf{J}^\mathsf{T}\mathbf{J})^{-1}\mathbf{J}^\mathsf{T} = \mathbf{J}\boldsymbol{\Sigma}\mathbf{J}^\mathsf{T}, \tag{7.12}$$

since the fit absorbs the part of the errors that the columns of $\mathbf{J}$ can explain. $\mathbf{H}$ is symmetric and idempotent, so $\mathbf{I} - \mathbf{H}$ is too, and the residual vector is Gaussian with mean zero and covariance $(\mathbf{I} - \mathbf{H})(\mathbf{I} - \mathbf{H})^\mathsf{T} = \mathbf{I} - \mathbf{H}$. Row k's residual has variance $1 - h_{kk}$, its redundancy number. Summed over all rows, by the cyclic property of the trace,

$$\sum_{k=1}^{n} (1 - h_{kk}) = n - \operatorname{tr}\!\left(\mathbf{J}\boldsymbol{\Sigma}\mathbf{J}^\mathsf{T}\right) = n - \operatorname{tr}\!\left(\boldsymbol{\Sigma}\mathbf{J}^\mathsf{T}\mathbf{J}\right) = n - \operatorname{tr}\mathbf{I}_d = n - d, \tag{7.13}$$

the degrees of freedom of the fit, with $n = 12$ rows and $d = 6$ entries for M_0: the leverages are the share of the entries each row takes, and they add to the number of entries. The expected misfit of any subset of rows is the sum of their redundancy numbers. For the seven reading

rows, the subset that holds the meters, the statistic is their part of the misfit and its reference is their share of the degrees of freedom,

$$t = \sum_{k \in \text{readings}} r_k(\hat{\mathbf{z}}_0)^2, \qquad \mathbb{E}[t \mid M_0] = \nu = \sum_{k \in \text{readings}} (1 - h_{kk}). \tag{7.14}$$

This part is a sum of correlated squares, not an exact chi-square. With $\mathbf{P}$ the selector of the reading rows, $t = \boldsymbol{\zeta}^{\mathsf{T}}(\mathbf{I}-\mathbf{H})\mathbf{P}^{\mathsf{T}}\mathbf{P}(\mathbf{I}-\mathbf{H})\boldsymbol{\zeta}$, and t is distributed as $\sum_i \lambda_i \, \chi^2_{1,i}$, a weighted sum of independent one-degree chi-squares whose weights λ_i are the eigenvalues of the reading block of $\mathbf{I}-\mathbf{H}$; their sum is ν. When those eigenvalues are near zero or one, as a reading block with a few redundant directions makes them, t is close to χ^2_ν on the non-integer ν, the reference used here. The bad-data detection for one reading is its normalized residual,

$$r_k^{\mathrm{N}} = \frac{r_k(\hat{\mathbf{z}}_0)}{\sqrt{1 - h_{kk}}} \sim \mathcal{N}(0, 1), \tag{7.15}$$

standard normal under M_0, and the largest normalized residual names the reading that departs most. For linear rows r_k^{N} is also the standardized departure of the reading from what the other rows predict for it, a virtual sensor of the reading's own quantity: the other rows' prediction minus the reading, over the spread of that difference. Its sign carries the direction: with every residual written predicted minus observed, a negative r_k^{N} is a reading above what the rest of the hall predicts.

7.2.5 The evidence of a drawing

One choice enters here: the evidence of each drawing is its Laplace evidence, the Gaussian integral of the objective about the drawing's own MAP state with its own curvature. This gives every evidence from the drawing's fit alone, and it is accurate only as far as each objective is quadratic over its posterior's width, which the meter rows, nonlinear in the log-derates, make approximate.

The evidence of a drawing is the density of the window's readings under it, averaged over its prior, equation (4.18). For M_0, with the twelve normalized Gaussian rows of standard deviations σ_k and a flat prior of constant density c_Q on the load, the Laplace form of equation (4.19) on six entries is

$$\ln p(\mathbf{y} \mid M_0) \approx \ln c_Q - \sum_{k=1}^{12} \ln(\sqrt{2\pi}\,\sigma_k) - \mathcal{L}_0(\hat{\mathbf{z}}_0) + \frac{6}{2}\ln 2\pi - \frac{1}{2}\ln\det \boldsymbol{\Lambda}_0. \tag{7.16}$$

The drawing M_j has the same twelve rows with b_j added to the predicted reading of meter j, one more row b_j/σ_b with the constant $-\ln(\sqrt{2\pi}\,\sigma_b)$, and one more entry, which adds $\frac{1}{2}\ln 2\pi$ to the Gaussian integral; its curvature $\boldsymbol{\Lambda}_j$ is 7×7. The flat-prior constant and the twelve row constants are common to all four drawings and cancel in a ratio, and the $\sqrt{2\pi}$ of the bias row cancels against the added dimension:

$$\ln B_{j0} = \left[\mathcal{L}_0(\hat{\mathbf{z}}_0) - \mathcal{L}_j(\hat{\mathbf{z}}_j)\right] - \frac{1}{2}\ln\frac{\det \boldsymbol{\Lambda}_j}{\det \boldsymbol{\Lambda}_0} - \ln\sigma_b. \tag{7.17}$$

The first bracket is the gain in fit that the bias buys, half the drop in $2\mathcal{L}$; the rest subtracts the Occam factor, $\frac{1}{2}\ln(\sigma_b^2 \det \boldsymbol{\Lambda}_j / \det \boldsymbol{\Lambda}_0)$, for linear rows, and to within 0.08 here, the logarithm of the ratio of the bias's prior width to its posterior width: $\frac{1}{2}\ln(1+\psi_k)$ of the closed form below is 1.0222, 0.8496 and 1.0221 against the factors 0.9867, 0.9247 and 1.0110 for M_{W_2}, M_Q and M_{W_1}. Its kilowatts cancel: σ_b^2 is in kW^2, and the determinant ratio is in kW^{-2} because the extra column of $\mathbf{J}_j$ is in inverse kW. Two drawings of the family compare by the difference of their logarithms, $\ln B_{jl} = \ln B_{j0} - \ln B_{l0}$.

The evidence and the residual test. For linear rows the Bayes factor (7.17) of a bias on reading k has a closed form in the statistic of the residual test, derived in *Reliability Engineering by Derivation*, Chapter 22: with

$$\psi_k = \frac{\sigma_b^2\,(1 - h_{kk})}{\sigma_k^2}, \qquad \ln B_{k0} = \frac{1}{2}\left[\frac{\psi_k}{1+\psi_k}\,(r_k^{\mathrm{N}})^2 - \ln(1+\psi_k)\right], \tag{7.18}$$

the ratio of the bias's prior variance to the variance of the reading's departure from the other rows. The evidence depends on the window only through the normalized residual, and on the meter's place in the graph and the bias prior only through ψ_k. A drawing therefore gains evidence when its normalized residual is large and its redundancy number is large; a meter with $1 - h_{kk}$ near zero, which the other rows cannot predict, gets $\ln B_{k0}$ near zero whatever it reads, and a bias drawing whose $(r_k^{\mathrm{N}})^2$ falls below $(1+\psi_k)\ln(1+\psi_k)/\psi_k$ loses to M_0. Where two meters have the largest normalized residuals, of opposite sign and nearly equal size, the two tests agree that one of them is at fault and cannot say firmly which; the evidence then splits between them. The rows here are not linear, and the chapter reports the Laplace evidence; the closed form is the reading of it.

For linear rows the same derivation reads the bias within a drawing,

$$\hat{b}_k = \frac{\sigma_b^2\,(y_k - \hat{y}_k)}{\sigma_k^2 + \sigma_b^2\,(1 - h_{kk})}, \qquad \mathrm{sd}(b_k \mid \mathbf{y}, M_k) = \frac{\sigma_b}{\sqrt{1+\psi_k}}, \tag{7.19}$$

with $\hat{y}_k$ the fitted value of reading k under M_0. The bias is the reading's departure from the fit, scaled up because the fit itself had leaned toward the biased reading, and shrunk by the bias prior.

7.3 The answers (step 4)

Every answer is read from the fits of Section 7.2.3, and that value is the one reported. Importance sampling, the closed forms and the exact tail stand beside each reported value, with the gap.

7.3.1 The ranking of the candidate faults

Operator question. Two meters should sum to the third and do not; which one is off? The answer is a probability for each meter and for none.

General form. Diagnosis, equation (2.15), the ranking of a family of candidate faults by their posterior probabilities, built from the evidences (2.26) and their ratios, the Bayes factors of equation (2.14),

$$\Pr(M_j \mid \mathbf{y}) = \frac{p(\mathbf{y} \mid M_j)\,\Pr(M_j)}{\sum_i p(\mathbf{y} \mid M_i)\,\Pr(M_i)}, \qquad B_{j0} = \frac{p(\mathbf{y} \mid M_j)}{p(\mathbf{y} \mid M_0)}.$$

Specialized form. The question is equation (2.15) with the family the four drawings M_0, M_Q, M_{W_1}, M_{W_2} of Section 7.1.4, $\Pr(M_j)$ the prior (7.5) and $p(\mathbf{y} \mid M_j)$ the Laplace evidence (7.16); each Bayes factor is equation (2.14) with $M_A = M_j$ and $M_B = M_0$, which is equation (7.17). Divided through by $p(\mathbf{y} \mid M_0)$,

$$\Pr(M_j \mid \mathbf{y}) = \frac{B_{j0}\,\Pr(M_j)}{\Pr(M_0) + \sum_l B_{l0}\,\Pr(M_l)}, \qquad l \in \{Q, W_1, W_2\}. \tag{7.20}$$

Building blocks. The minimized objectives $\mathcal{L}_j(\hat{\mathbf{z}}_j)$ and the information matrices of Section 7.2.3; the bias prior's σ_b; the prior probabilities (7.5); beside them, the normalized residuals and redundancy numbers of M_0 (Section 7.2.4) in the closed form (7.18).

The ranking at 23:00. Table 7.5 gives the four drawings. Each row is a drawing; the columns give twice its minimized objective, the logarithm of its determinant, the gain in fit and the Occam factor of equation (7.17), its log Bayes factor against M_0, its log evidence by equation (7.16), and its posterior probability by equation (7.20). The unit-2 meter's gain of 4.3867 less its Occam factor of 0.9867 gives its 3.4000.

Table 7.5: The four drawings of the two-unit hall at 23:00, ranked by their Laplace evidence, with the constant $\ln c_Q$ of the flat prior, common to all four, left out. Bias prior $\mathcal{N}(0,\ 10^2)$ kW; prior probabilities 0.5 for M_0 and 0.5/3 for each meter.

Drawing	$2\mathcal{L}_j(\hat{\mathbf{z}}_j)$ [–]	$\ln\det\mathbf{\Lambda}_j$ [–]	Gain [–]	Occam [–]	$\ln B_{j0}$ [–]	$\ln p(\mathbf{y} \mid M_j)$ [–]	$\Pr(M_j \mid \mathbf{y})$ [–]
Unit-2 meter, M_{W_2}	4.2201	29.2138	4.3867	0.9867	3.4000	−11.6416	0.5971
Hall meter, M_Q	5.5196	29.0897	3.7369	0.9247	2.8122	−12.2294	0.3317
Every meter sound, M_0	12.9934	31.8456	0	0	0	−15.0416	0.0598
Unit-1 meter, M_{W_1}	12.0808	29.2625	0.4563	1.0110	−0.5548	−15.5964	0.0114

Figure 7.2 draws the prior and the posterior probabilities side by side: one pair of bars per drawing, the vertical axis the probability, the light bar the prior (7.5) and the dark bar the posterior. The no-bias drawing falls from 0.5 to 0.060; the unit-2 meter rises to 0.597 and the hall meter to 0.332; the unit-1 meter falls to 0.011.

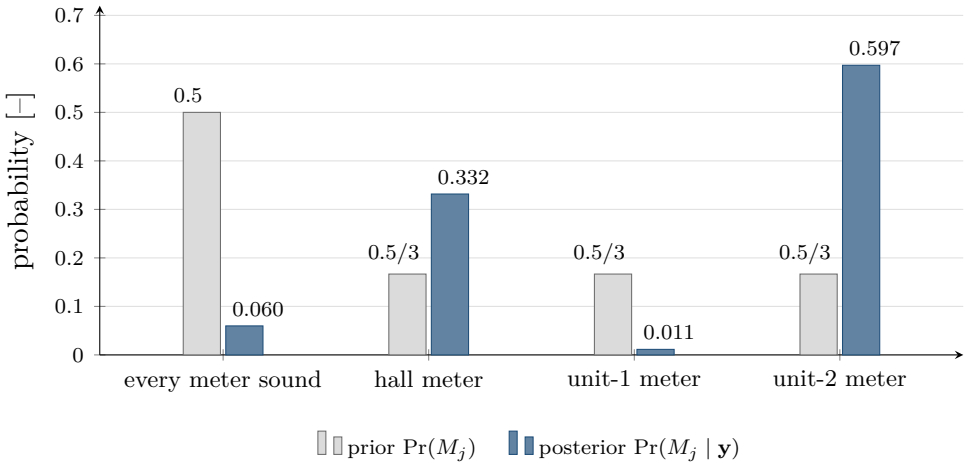


Figure 7.2: The prior and posterior probabilities of the four drawings at 23:00 (Table 7.5), by equation (7.20) with the Laplace evidences. Light: the prior; dark: the posterior.

Against each other the leaders stand at

$$\ln B_{W_2Q} = 0.5877, \qquad \ln B_{W_2W_1} = 3.9547, \qquad \ln B_{QW_1} = 3.3670.$$

The unit-2 meter leads the hall meter by 0.59 in the logarithm, far short of a settled ranking; both lead the unit-1 meter by more than three.

The values beside them. Importance sampling of the exact objective (Section 7.4.2) gives the log evidences −11.6422, −12.2289, −15.0416 and −15.5969 for M_{W_2}, M_Q, M_0 and M_{W_1}, gaps of −0.0006, +0.0004, 0.0000 and −0.0005, and the probabilities 0.5968, 0.3319, 0.0598 and 0.0114, gaps of at most 0.0002; $\ln B_{W_2Q}$ is 0.5867.

The closed form (7.18) reads the same Bayes factors from the residual test of M_0 as if the rows were linear. Table 7.6 evaluates it: each row is a meter; the columns give ψ_k, the squared

normalized residual, its break-even $(1+\psi_k)\ln(1+\psi_k)/\psi_k$, and the closed-form and Laplace $\ln B_{k0}$ with their gap. The unit-2 meter's 10.089 and the hall meter's 9.152 lie far above their break-evens, 2.348 and 2.079; the unit-1 meter's 1.049 lies below its 2.348, and its drawing loses. The two unit meters have nearly the same ψ_k, 6.725 and 6.723, because their leverages are nearly equal, 0.3948 and 0.3949, so they differ only in their residuals. The closed form narrows the lead of the unit-2 meter over the hall meter, $\ln B_{W_2Q} = 0.4796$ against the Laplace 0.5877, a gap of -0.1081; the larger part of it is the hall meter's, and importance sampling sides with the Laplace value. Its probabilities are 0.5741, 0.3554, 0.0593 and 0.0112, gaps of -0.0229, $+0.0237$, -0.0005 and -0.0002: the order of the drawings does not change, and no threshold is crossed.

Table 7.6: The closed form (7.18) for each meter at 23:00, from the residual test of the drawing with every meter sound, against the Laplace evidence of Table 7.5. Bias prior $\sigma_b = 10\,\mathrm{kW}$.

Meter	ψ_k [–]	$(r_k^{\mathrm{N}})^2$ [–]	Break-even [–]	$\ln B_{k0}$, closed [–]	$\ln B_{k0}$, Laplace [–]	Gap [–]
Unit-2 meter	6.725	10.089	2.348	3.3694	3.4000	−0.0306
Hall meter	4.470	9.152	2.079	2.8898	2.8122	+0.0776
Unit-1 meter	6.723	1.049	2.348	−0.5655	−0.5548	−0.0107

7.3.2 The misfit and the largest normalized residual

Operator question. Is anything off at all? The question asks whether the window's readings are consistent with the hall with every meter sound, before any meter is blamed, and which reading departs most.

General form. A tail of the predictive, equations (2.13) and (2.12): the probability, under M_0, that a replicate window y_{new} of the same hall would show a statistic t of its residuals at least as large as the observed t^*,

$$\Pr\big(t(y_{\mathrm{new}}) \geq t^* \mid M_0\big) = \int \mathbf{1}\big[t(y_{\mathrm{new}}) \geq t^*\big]\, p(y_{\mathrm{new}} \mid M_0)\, \mathrm{d}y_{\mathrm{new}}.$$

Specialized form. The question is equations (2.13) and (2.12) with y_{new} a replicate window of the seven readings, f the statistic t of equation (7.14) read on the replicate's residuals and f^* its observed value t^*. The state $p(\mathbf{z} \mid \mathbf{y})$ of those equations drops out: by equation (7.12), for linear rows the law of t under M_0 does not depend on the state, so the replicate may be drawn at any state. That law is the weighted sum of one-degree chi-squares of Section 7.2.4; its tail is reported from χ^2_ν, with the tail of the weighted sum beside it. The reading-level test is the normalized residual (7.15) of each reading against the standard normal. The assurance level of equation (2.12) is not set here: the tail itself is reported.

Building blocks. The MAP state $\hat{\mathbf{z}}_0$ and its covariance (Section 7.2.3); the residuals, leverages and sums (7.13) and (7.14) at that fit; the eigenvalues of the reading block of $\mathbf{I} - \mathbf{H}$.

The misfit at 23:00. Table 7.7 gives the seven reading rows of M_0 at its fit. Each row is a reading; the columns give its residual $r_k(\hat{\mathbf{z}}_0)$ of equation (7.12), its leverage h_{kk} and its normalized residual r_k^{N} of equation (7.15).

Table 7.7: The reading rows of the fit with every meter sound, at 23:00: the residual at the fit (predicted minus observed, over the reading's standard deviation), the leverage and the normalized residual.

Reading	$r_k(\hat{\mathbf{z}}_0)$ [–]	h_{kk} [–]	r_k^{N} [–]
Return air y_{T_r}	−0.4616	0.3868	−0.5895
Discharge of unit 1 $y_{T_{d,1}}$	1.0518	0.5786	1.6203
Discharge of unit 2 $y_{T_{d,2}}$	−0.6405	0.5785	−0.9866
Hall meter y_Q	1.9188	0.5977	3.0252
Rack-1 inlet y_{T_1}	−0.1009	0.8881	−0.3015
Unit-1 meter y_{W_1}	−0.7968	0.3949	−1.0243
Unit-2 meter y_{W_2}	−2.4711	0.3948	−3.1764

The chi-square test. By equation (7.13) the redundancy numbers of all twelve rows add to 6.0000, the $12 - 6$ degrees of freedom. The seven reading rows take $\nu = 3.1807$ of them, and their part of the misfit is

$$t = \sum_{k \in \text{readings}} r_k(\hat{\mathbf{z}}_0)^2 = 12.1628 \quad \text{against} \quad \mathbb{E}[t \mid M_0] = 3.1807,$$

$$\Pr(\chi^2_{3.1807} > 12.1628) = 0.0081.$$

The weighted sum has as weights the eigenvalues of the reading block,

$$\{0,\ 0.011,\ 0.018,\ 0.152,\ 1,\ 1,\ 1\}, \qquad \text{sum } 3.1807:$$

three directions of the readings fully redundant, three with little redundancy, and one with none. Drawn 2 000 000 times, it exceeds 12.1628 with probability 0.0075, a gap of −0.0007 from the reported 0.0081. The readings are inconsistent with every meter sound: a window of this hall with sound meters shows a reading-row misfit this large less than once in a hundred windows, by either tail.

The largest normalized residual. Figure 7.3 draws the normalized residuals of the last column of Table 7.7 as bars, one per reading, the horizontal axis in standard deviations; negative bars are readings above what the rest of the hall predicts. Two bars stand out, of opposite sign and nearly equal length: the unit-2 meter at −3.18, reading high, and the hall meter at +3.03, reading low. The next is unit 1's discharge at +1.62; the others are under 1.1 in magnitude. The largest normalized residual names the unit-2 meter, with the hall meter close behind. The unit-2 meter also has the smaller leverage of the two, 0.3948 against 0.5977, so it has the larger redundancy number, and by the closed form (7.18) both its larger residual and its larger redundancy favor it in the evidence (Table 7.6).

7.3.3 The size of the bias

Operator question. By how much is it off? The answer is the bias of the suspected meter with its band, given under each drawing that remains in contention.

General form. Estimation, equation (2.9), of an entry within a drawing,

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

summarized by its mean, its spread and its central 95 % interval, equation (2.11).

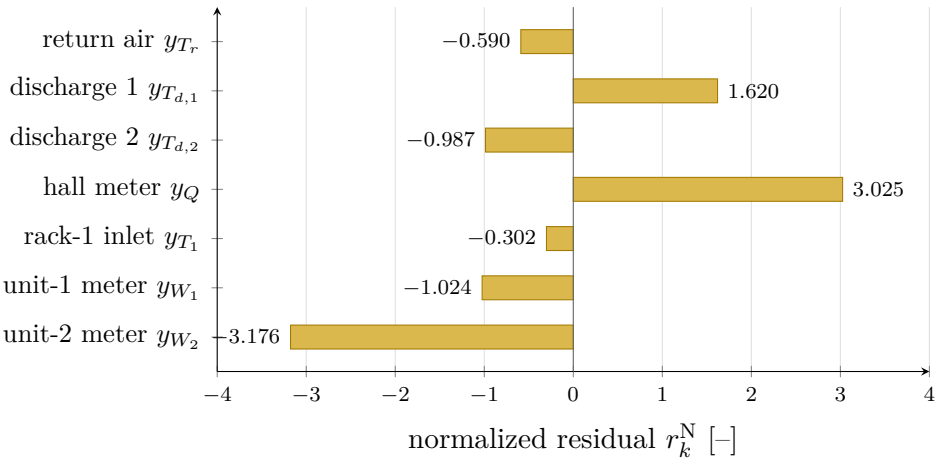


Figure 7.3: The normalized residuals of the seven readings under the drawing with every meter sound, equation (7.15) (Table 7.7). Each is standard normal under that drawing; residuals are predicted minus observed, so a negative bar is a reading above what the rest of the hall predicts.

Specialized form. The question is equations (2.9) and (2.11) with $f(\mathbf{z}) = b_j$, the seventh entry of the state of M_j, f^* any level of it, and $p(\mathbf{z} \mid \mathbf{y}, M_j)$ the Laplace normal $\mathcal{N}(\hat{\mathbf{z}}_j, \boldsymbol{\Sigma}_j)$ of equation (7.11). One choice enters here: the bias is read within each leading drawing, not averaged over the family. A mixture over the drawings would put most of its weight on two lumps of opposite sign, one per leader, and its mean would describe neither meter; the report gives the bias under each drawing with that drawing's probability.

The estimate is an entry of the MAP state and its spread a diagonal entry of the covariance,

$$b_j \mid \mathbf{y}, M_j \approx \mathcal{N}(\hat{b}_j,\ (\boldsymbol{\Sigma}_j)_{77}), \qquad [\hat{b}_j \mp 1.960\sqrt{(\boldsymbol{\Sigma}_j)_{77}}\,]. \tag{7.21}$$

Building blocks. The MAP state and the covariance of each leading drawing (Table 7.3; $\boldsymbol{\Sigma}_{W_2}$ of Section 7.2.3); beside them, the closed form (7.19).

The biases at 23:00. By equation (7.21), under the unit-2 drawing

$$b_{W_2} \mid \mathbf{y}, M_{W_2} \approx \mathcal{N}(10.6494\,\mathrm{kW},\ 3.5969^2), \qquad b_{W_2} \in [3.60,\ 17.70]\,\mathrm{kW},$$

with the load at $\hat{Q} = 195.620\,\mathrm{kW}$; and under the hall-meter drawing

$$b_Q \mid \mathbf{y}, M_Q \approx \mathcal{N}(-11.6864\,\mathrm{kW},\ 4.2766^2), \qquad b_Q \in [-20.07,\ -3.30]\,\mathrm{kW},$$

with the load at $\hat{Q} = 205.778\,\mathrm{kW}$. The closed form (7.19) stands beside them at 10.6632 kW (sd 3.5980), gaps of +0.0138 and +0.0011, and at −11.6930 kW (sd 4.2757), gaps of −0.0066 and −0.0009. Figure 7.4 draws the two: the horizontal axis is the bias in kW, the vertical axis its posterior density within each drawing; the solid blue curve is the unit-2 meter's, with its 95 % band shaded, the solid red curve the hall meter's, the dashed gray curve the bias prior and the dotted line the stated truth used in Section 7.4.6. Neither band contains zero, and neither is centered on the 15.56 kW of the balance: the unit-2 meter is read about 10.6 kW high and the hall meter about 11.7 kW low. The two drawings also disagree on the load, 195.620 against 205.778 kW, which is where a check can separate them.

7.3.4 The value of a check

Operator question. Which check would settle it? A technician can make one spot reading at a time; the question asks which would move the ranking of the two leading drawings most, before it is made.

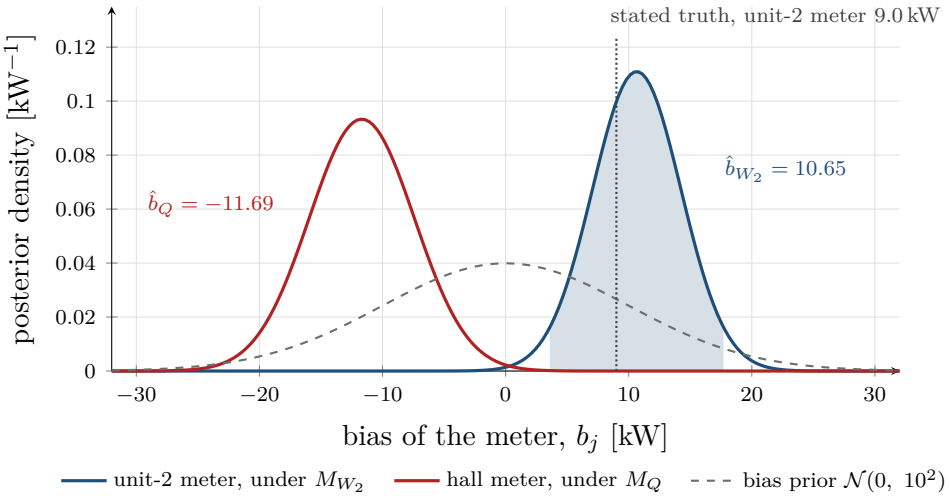


Figure 7.4: The bias of the meter under each leading drawing, equation (7.21): $\mathcal{N}(10.6494,\ 3.5969^2)$ for the unit-2 meter under M_{W_2}, its 95 % band $[3.60,\ 17.70]$ kW shaded, and $\mathcal{N}(-11.6864,\ 4.2766^2)$ for the hall meter under M_Q; dashed, the bias prior. Dotted: the stated truth of the record, a check marker (Section 7.4.6).

General form. Design, equation (2.16), the preposterior value of a candidate check a whose outcome $\tilde{y}_a$ is not yet known,

$$U(a) = \int u(p(\cdot \mid \mathbf{y}, \tilde{y}_a))\, p(\tilde{y}_a \mid \mathbf{y})\, \mathrm{d}\tilde{y}_a.$$

Specialized form. The question is equation (2.16) with a one of the three checks below, $\tilde{y}_a$ its reading, $p(\tilde{y}_a \mid \mathbf{y})$ its predictive under the drawing taken as true, and u the change of the log Bayes factor of the two leading drawings. Three choices enter here. Each check is valued between the two leading drawings only: this is what makes its value one change of one Bayes factor, and it says nothing of an explanation outside the family, which can predict the same reading as one of the leaders. Each check is given the precision of the instrument the shift would use, 1 kW for a clamp-on reading of the UPS output and 2 kW for a portable water-side heat check: the value of a check falls as its standard deviation grows, and the ranking of the checks can turn on it, which Table 7.9 sweeps for the clamp-on reading. And each check is valued before any other is made: a check made after another is valued against the posterior that the first leaves, which can be a different number.

By the product rule $p(\mathbf{y}, \tilde{y}_a \mid M) = p(\mathbf{y} \mid M)\, p(\tilde{y}_a \mid \mathbf{y}, M)$, so for the leaders M_l and M_m the check changes the log Bayes factor by the log ratio of their predictive densities at the value read,

$$\ln B_{lm}(\mathbf{y}, \tilde{y}_a) = \ln B_{lm}(\mathbf{y}) + \ln \frac{p(\tilde{y}_a \mid \mathbf{y}, M_l)}{p(\tilde{y}_a \mid \mathbf{y}, M_m)}.$$

Within each drawing the check is predicted as a probe was in Chapter 5, equation (5.18): a check of standard deviation σ_a on a quantity $f_a(\mathbf{z})$ of gradient $\mathbf{g}_a$ reads from

$$\tilde{y}_a \mid \mathbf{y}, M_l \approx \mathcal{N}(\mu_l,\ s_l^2), \qquad \mu_l = f_a(\hat{\mathbf{z}}_l), \quad s_l^2 = \mathbf{g}_a^{\mathsf{T}} \boldsymbol{\Sigma}_l\, \mathbf{g}_a + \sigma_a^2. \tag{7.22}$$

If M_l is true, the expected change is the Kullback–Leibler divergence of the M_m predictive from the M_l predictive, which is never negative: on average a check moves the evidence toward the true drawing, though a particular reading can move it the other way. For two normals it is

$$U_l(a) = \mathbb{E}[\Delta \ln B_{lm} \mid M_l] = \frac{1}{2}\left[\ln \frac{s_m^2}{s_l^2} + \frac{s_l^2 + (\mu_l - \mu_m)^2}{s_m^2} - 1\right], \tag{7.23}$$

and the same with the indices exchanged, as the expected change in favor of M_m, if M_m is true. The two directions differ when the spreads differ: the drawing with the wider predictive is the easier one to confirm, because its readings fall in the far tail of the narrower predictive more often than the reverse. For a common spread s both directions equal $(\mu_l - \mu_m)^2/(2s^2)$. With s^2 the mean of the two variances, the separation of the check and its equal-spread summary are

$$\delta_a = \frac{|\mu_l - \mu_m|}{\sqrt{\frac{1}{2}\,(s_l^2 + s_m^2)}}, \qquad U(a) \approx \tfrac{1}{2}\,\delta_a^2, \tag{7.24}$$

which ranks the checks by one number and is reported beside the two directions of equation (7.23). Three checks are open to the shift. A clamp-on reading of the UPS output reads the load, $f_a = Q$, $\mathbf{g}_a = \mathbf{e}_1$, with $\sigma_a = 1\,\text{kW}$. A portable water-side heat check of unit 2 reads $f_a = Q_2(\mathbf{z})$ of equation (7.2), with $\sigma_a = 2\,\text{kW}$, and a portable check of unit 1 reads $Q_1(\mathbf{z})$, also with $\sigma_a = 2\,\text{kW}$. The two drawings differ in their fitted states, and a check separates them by as much as its quantity differs between those states, in units of its spread.

Building blocks. The MAP states and covariances of the two leading drawings (Section 7.2.3); the quantity f_a of each check and its gradient, equation (7.9) for a unit heat; the check's σ_a.

The separation of the checks. Table 7.8 evaluates the three checks by equations (7.22), (7.23) and (7.24). Each row is a check; the columns give its standard deviation, the predicted reading with its spread under each leader, the separation δ_a, its equal-spread summary $\frac{1}{2}\delta_a^2$, and the expected change of the log Bayes factor in favor of the true drawing if the unit-2 meter is biased and if the hall meter is.

Table 7.8: The three checks open to the shift, valued between the two leading drawings at 23:00. Predicted readings $\mu \pm s$ in kW, the spread s including the check's own σ_a; the last two columns are equation (7.23) with M_{W_2} and with M_Q true.

Check	σ_a [kW]	Under M_{W_2} [kW]	Under M_Q [kW]	δ_a [–]	$\frac{1}{2}\delta_a^2$ [–]	If M_{W_2} [–]	If M_Q [–]
Clamp-on reading of the UPS output, Q	1	195.620 ± 2.744	205.778 ± 3.594	3.18	5.05	4.05	6.94
Portable check of unit 2, Q_2	2	98.542 ± 3.092	106.227 ± 3.032	2.51	3.15	3.21	3.09
Portable check of unit 1, Q_1	2	105.078 ± 2.770	107.552 ± 3.032	0.85	0.36	0.34	0.41

The clamp-on reading separates the leaders by 3.18 of their spreads. Its two predictive spreads differ, 2.744 against 3.594 kW, so its value depends on which leader is true: an expected 4.05 in the log Bayes factor if the unit-2 meter is biased and 6.94 if the hall meter is, around the equal-spread 5.05. The portable check of unit 2 separates them by 2.51, an expected 3.21 or 3.09, nearly the same in both directions because its two spreads are nearly equal; the portable check of unit 1 by 0.85, an expected 0.34 or 0.41.

The precision of the clamp-on reading. Table 7.9 repeats the clamp-on reading's separation, equal-spread value and two directions for standard deviations from 1 to 4 kW, against the portable check of unit 2's 3.15, 3.21 and 3.09 at its 2 kW; the last row is its separation of a hall-meter bias from the other heat of M_{Q_o}. The clamp-on reading's lead survives to a standard deviation between 2.6 and 2.7 kW, 3.23 and 3.12 against 3.15; a clamp-on instrument coarser

than that ranks second. By direction the crossing lies between 2 and 2.5 kW if the unit-2 meter is biased and between 3 and 4 kW if the hall meter is. The portable check's own value falls too, 4.63, 3.15, 2.05 and 1.38 at 1, 2, 3 and 4 kW; at equal precision the clamp-on reading leads throughout, 5.05, 3.90, 2.83 and 2.05.

Table 7.9: The clamp-on reading of the UPS output against its precision: the separation of the two leaders, the equal-spread expected change of the log Bayes factor, equation (7.24), and the two directions of equation (7.23); last, its separation of M_Q from M_{Q_o}. The portable check of unit 2 is worth 3.15 at its 2 kW.

σ_a of the clamp-on reading [kW]	1	2	2.5	2.6	2.7	2.8	3	4
Separation δ_a [–]	3.18	2.79	2.58	2.54	2.50	2.46	2.38	2.02
$\frac{1}{2}\delta_a^2$ [–]	5.05	3.90	3.33	3.23	3.12	3.02	2.83	2.05
If M_{W_2} [–]	4.05	3.28	2.87	2.79	2.71	2.63	2.49	1.86
If M_Q [–]	6.94	4.95	4.07	3.91	3.76	3.62	3.35	2.30
M_Q from M_{Q_o} [–]	3.50	3.11	2.89	2.84	2.80	2.76	2.67	2.29

One check after the other. From the refits of Table 7.4, take first a clamp-on reading at 200.7 kW, the midpoint of its two predictions, or at 200.0 kW beside it. The check of unit 2 then separates the leaders by 0.96 and 0.99, worth 0.46 and 0.49 against its 3.15 made first; in the two directions, 0.49 and 0.44 at 200.7 kW, 0.52 and 0.46 at 200.0 kW. Over the range of the clamp-on predictions the value runs from 0.66 at 195.62 kW down to 0.30 at 205.78 kW. Against the other heat of M_{Q_o} it keeps more: separations of 2.71, 1.94 and 1.82, worth 3.68, 1.88 and 1.65, at 195.62, 200.0 and 200.7 kW. Taken the other way, the clamp-on reading after a portable reading of 98.54, 100.36, 102.40 and 106.23 kW separates the leaders by 1.91, 2.05, 2.20 and 2.50, worth 1.82, 2.10, 2.43 and 3.12. Read together, equation (7.24) with the mean of the two joint covariances, the two checks are worth 5.55, and equation (7.23) in two dimensions gives 4.76 if the unit-2 meter is biased and 7.23 if the hall meter is. The pair is worth 5.55 in either order, so whichever check comes second is worth what the first left: the clamp-on reading after a portable reading midway, 102.40 kW, keeps 2.43.

7.4 The checks (step 5)

The answers of Section 7.3 rest on approximations, inputs and a pinned quantity that can be tested on this hall, against the calculation each stands for or the alternative it sets aside. Gaps are taken before rounding.

7.4.1 The limit of exact rows and the curvature term

The fit runs on the reduced state, the limit $\sigma_p \to 0$ of the full state. The check fits the full state of the leading drawing, the thirteen entries of equation (5.2) and the bias, with the seven physical rows (5.1) at a finite tolerance beside the seven readings, the five priors and the bias prior, and lets the tolerance shrink. Each unit meter reads $G_u(T_r - T_s)$ of entries of that state. Table 7.10 gives the result: each row is a tolerance, the columns the load and the bias with their spreads, and the last row the reduced state. At 0.1 the bias is 0.0011 kW low and its spread 0.0006 kW wider; at 0.01 and 0.001 the full state reproduces the reduced state in all four digits shown.

The Gauss–Newton covariance drops the curvature term of the nonlinear meter and temperature rows. The full Hessian of the objective of M_{W_2} at its MAP state, curvature term included, gives a Newton spread of the bias of 3.5948 kW against the Gauss–Newton 3.5969 kW: the term moves the answer in its fourth digit.

Table 7.10: The limit of exact physical rows on the full state of the unit-2 drawing: the load and the bias at three tolerances, against the reduced state.

Tolerance σ_p [kW or K]	Q [kW]	sd(Q) [kW]	b_{W_2} [kW]	sd(b_{W_2}) [kW]
0.1	195.6189	2.5559	10.6484	3.5975
0.01	195.6199	2.5556	10.6494	3.5969
0.001	195.6199	2.5556	10.6494	3.5969
reduced state	195.6199	2.5556	10.6494	3.5969

7.4.2 The Laplace evidence against importance sampling

The Laplace evidence (7.16) replaces each objective by its quadratic expansion. The check integrates the exact objective instead: it draws 200 000 states from each drawing's Laplace normal, weights each by the exact density of the rows over the density it was drawn from, and averages, which is importance sampling of the integral $\int \mathrm{e}^{-\mathcal{L}_j}\,\mathrm{d}\mathbf{z}$ with the nonlinear rows and their curvature kept. Its standard errors, from the spread of the weights, are 0.00009 to 0.00013; repeating it in 20 batches of 20 000 draws spreads the estimates by 0.00026 for M_0 and 0.00032 for M_{W_2}, which at 200 000 draws is $0.00026/\sqrt{10} = 0.00008$ for M_0, and agrees. Three further seeds at 4 000 000 draws per drawing, standard errors 0.00002 to 0.00003, give gaps of -0.00059 to -0.00065 for M_{W_2}, $+0.00024$ to $+0.00033$ for M_Q and -0.00032 to -0.00036 for M_{W_1}, so the gaps of 2 to 7×10^{-4} for the bias drawings are real Laplace errors, the price of the nonlinear rows. They are of either sign and of the same size, and they move no probability by more than 0.0002: the Laplace evidence is adequate for the ranking.

7.4.3 The chi-square reference

The tail of the weighted sum in Section 7.3.2 carries a Monte Carlo standard error of 0.00006, so the gap of -0.0007 is resolved: the reference χ^2_ν is slightly conservative here, and the conclusion that the readings are inconsistent with every meter sound does not move.

7.4.4 Sensitivity to the bias prior

The width σ_b of the bias prior enters the Occam factor (7.17) directly, so the ranking is repeated for $\sigma_b = 5$, 10 and 20 kW. Table 7.11 gives the result: each row is a width, the columns the four posterior probabilities. A narrow prior penalizes a bias of 10 to 12 kW as implausible and returns probability to M_0, 0.1156 at 5 kW; a wide prior charges every bias drawing more for its entry, returning a little to M_0, 0.0671 at 20 kW, and it costs least to the drawing whose bias is largest, the hall meter, which rises to 0.3587. Across a factor of four in the width the unit-2 meter stays between 0.567 and 0.597 and the hall meter between 0.282 and 0.359: the ranking and the split between the two leaders hold.

Table 7.11: The posterior probabilities of the four drawings at three widths of the bias prior. The middle row is the reported ranking.

Bias prior sd σ_b [kW]	Pr(M_0) [–]	Pr(M_Q) [–]	Pr(M_{W_1}) [–]	Pr(M_{W_2}) [–]
5	0.1156	0.2819	0.0330	0.5695
10	0.0598	0.3317	0.0114	0.5971
20	0.0671	0.3587	0.0071	0.5671

7.4.5 The pinned other heat

The other heat $Q_o = 8\,\text{kW}$ is pinned, and an error in it is not a member of the family. The fit M_{Q_o} of Section 7.2.3 frees it and is the hall-meter drawing again, so a hall-side explanation is a pair. This exposes a choice the prior (7.5) made: the half is shared among the three meters alone. With the other heat as a fifth drawing, at 0.5/4 each, the probabilities are unit-2 meter 0.4417, hall meter 0.2454, other heat 0.2454, none 0.0590 and unit-1 meter 0.0085, and the hall-side pair, 0.4908, then leads the unit-2 meter. The gap crosses a decision, which explanation leads, and the evidence does not settle it: both sets of values are stated, and the report does not choose between them.

The checks part ways on the pair. Under M_{Q_o} a clamp-on reading of the UPS output is predicted at 194.09 kW (sd 3.06), next to the unit-2 drawing's 195.620 kW, a separation of 0.53, and far from the hall-meter drawing's 205.778 kW, a separation of 3.50, or 3.86 after a portable reading of 106.23 kW. The portable check of unit 2 is predicted at 106.227 kW (sd 3.032), the hall-meter drawing's own prediction: it separates the unit-2 explanation from both by 2.51, and those two from each other not at all.

7.4.6 The answers against the stated truth

The record of this hall is generated from a stated truth (Section 7.6.4), and the last check compares the answers with it. In the truth the unit-2 meter carries a bias of 9.0 kW, and its error in this window, bias and noise together, is 9.79 kW; the unit-1 meter's error is $-1.19\,\text{kW}$; the true unit heats are 107.64 and 100.36 kW. The leading drawing is the true one, and the true bias lies inside its band $[3.60,\ 17.70]\,\text{kW}$, below its center (Figure 7.4). The hall meter carries no bias, but its reading of 193.04 kW against the true load of 200.0 kW is 2.32 of its standard deviations low, by noise alone. That draw is why the evidence is split: a sound hall meter that happens to read low and a unit-2 meter that reads high both open the balance in the same direction, and the hall-meter drawing explains the window nearly as well as the true one. It is also why the unit heats of Chapter 6, fitted from the hall meter without the unit meters, sat below their true values.

Against the truth, the clamp-on reading of the load would fall between the two predictions of Table 7.8, 200.0 kW against 195.620 and 205.778 kW, and the portable check of unit 2 near the unit-2 drawing's prediction, 100.36 kW against 98.542 and 106.227 kW. The expected change of $\ln B_{W_2Q}$ when each check reads its true quantity with its own error is $+1.69$ for the portable check of unit 2, $+0.26$ for the clamp-on reading and -0.38 for the portable check of unit 1. The values of Table 7.8 rank the checks before any is made, over every state the drawings allow; at the true state of this window the clamp-on reading lands near the middle and moves the ranking little, while the check of unit 2 moves it toward the true drawing.

Table 7.12 collects the checks and their gaps: each row is a route or a variation, grouped by the answer it tests, and each gap is the route minus the reported value.

7.5 The report (step 6)

The flag's implied single culprit does not survive the analysis. Two explanations remain: the unit-2 meter reading about 10.6 kW high, with probability 0.60, and the hall meter reading about 11.7 kW low, with probability 0.33. Every meter sound keeps 0.06 and the unit-1 meter 0.01. Neither meter is off by the 15.56 kW the balance shows, and the evidence between the two leaders, 0.59 in the logarithm, is not enough to send either meter for recalibration. The second explanation is a hall-side one, a hall meter or an other heat larger than pinned (Section 7.4.5): with the other heat given a prior share of its own, the unit-2 meter stands at 0.44 and the hall-side pair at 0.49. Both statements are reported, and the request below does not depend on which is taken.

Table 7.12: The checks of the answers and their gaps from the reported values. Importance sampling uses 200 000 draws per drawing and the exact tail 2 000 000; the prior variations repeat the four fits; the truth is the stated truth of the constructed record.

Route or variation	Value	Gap	Value	Gap
The ranking: $\Pr(M_{W_2} \mid \mathbf{y})$ *and* $\Pr(M_Q \mid \mathbf{y})$ *[–]*				
Laplace evidence, $\sigma_b = 10\,\text{kW}$, reported	0.5971	–	0.3317	–
Importance sampling	0.5968	−0.0002	0.3319	+0.0002
Closed form for linear rows	0.5741	−0.0229	0.3554	+0.0237
Bias prior $\sigma_b = 5\,\text{kW}$	0.5695	−0.0275	0.2819	−0.0498
Bias prior $\sigma_b = 20\,\text{kW}$	0.5671	−0.0300	0.3587	+0.0270
Other heat freed instead of a meter	–	–	same evidence as M_Q	
Stated truth	unit-2 meter biased		sound, read 2.32 sd low	
The leaders: $\ln B_{W_2Q}$ *[–]*				
Laplace evidence, reported	0.5877	–		
Importance sampling	0.5867	−0.0010		
Closed form for linear rows	0.4796	−0.1081		
The misfit: tail of the reading-row statistic [–]				
χ^2_ν on $\nu = 3.1807$, reported	0.0081	–		
Weighted sum, exact	0.0075	−0.0007		
The bias under M_{W_2}*:* $\hat{b}_{W_2}$ *and its sd [kW]*				
Fit, reported	10.6494	–	3.5969	–
Closed form for linear rows	10.6632	+0.0138	3.5980	+0.0011
Newton covariance	–	–	3.5948	−0.0021
Full state, $\sigma_p = 0.001$	10.6494	0.0000	3.5969	0.0000
Stated truth	9.0	−1.6494	–	–
Dashboard (the balance)	15.56	+4.9106	–	–
The bias under M_Q*:* $\hat{b}_Q$ *and its sd [kW]*				
Fit, reported	−11.6864	–	4.2766	–
Closed form for linear rows	−11.6930	−0.0066	4.2757	−0.0009

7.5.1 The answers and their bands

Table 7.13 gives the answers in the order of Table 7.1. Each row is a question; the second column gives the answer with the exact routes beside it and the third its band: the range over the three bias priors of Section 7.4.4 for a probability, the central 95 % interval for a bias, and for the check the equal-spread value of the other candidate over its precision.

Table 7.13: The report for the two-unit hall's meters at 23:00. Answers from the Laplace evidences and the fits of the drawings, with the exact routes beside them; bands are ranges over bias priors of 5 to 20 kW, or 95 % intervals, and for the check the other candidate's value over its precision.

Question	Answer	Band
Which one is off?	unit-2 meter 0.597; hall meter 0.332; none 0.060; unit-1 meter 0.011 (importance sampling: the same to three decimals; closed form: 0.574, 0.355, 0.059, 0.011); with the other heat a fifth drawing: 0.442, 0.245, other heat 0.245, 0.059, 0.009, the hall-side pair 0.491	unit-2 meter 0.567 to 0.597; hall meter 0.282 to 0.359
Is anything off at all?	yes: reading-row misfit 12.16 against 3.18, tail 0.0081 (exact: 0.0075); largest normalized residuals −3.18 (unit-2 meter) and +3.03 (hall meter)	no-bias drawing 0.060 to 0.116
By how much is it off?	unit-2 meter +10.65 kW, sd 3.60 kW, if it is the one; hall meter −11.69 kW, sd 4.28 kW, if it is (closed form: +10.66, sd 3.60; −11.69, sd 4.28)	3.60 to 17.70 kW; −20.07 to −3.30 kW
Which check would settle it?	a portable water-side heat check of unit 2: separation 2.51 sd, expected $\ln B$ 3.21 or 3.09	clamp-on reading of the UPS output, larger between the leaders: 5.05 at 1 kW, 2.05 at 4 kW; not first because it cannot tell unit 2 from other heat (0.53)

7.5.2 The entry in the shift log

The operator's entry for 23:00 records a confirmed inconsistency and an open diagnosis, not a work order. The meters do not balance by 15.56 kW, and the hall's other readings confirm the inconsistency: with every meter sound and the other heat as pinned, a misfit this large has a tail probability of 0.0081. The leading explanations are the unit-2 meter, reading about 10.6 kW high, with probability 0.60, and the hall meter, reading about 11.7 kW low, with probability 0.33; none has 0.060 and the unit-1 meter 0.011. With the other heat as a fifth drawing the unit-2 meter has 0.44 and the hall-side pair 0.49. Requested tonight: a portable water-side heat check of unit 2. Until it is in, the hall's heat is reported with both explanations, and no meter is recalibrated; one check that lands between the predictions can leave the question open into the next window.

Common mistake

Recalibrating the meter with the largest normalized residual. The largest normalized residual names the unit-2 meter at −3.18, but the hall meter stands at +3.03 beside it, and the evidence gives the hall meter a probability of 0.33. A recalibration on the residual alone carries that probability of adjusting a sound meter and leaving the biased one in service.

7.6 The next reading (step 7)

The report leaves the two leaders at 0.59 in the log Bayes factor. A check is one more row of the residual vector of both drawings, and its value to the diagnosis is known before its reading is.

7.6.1 The candidate checks

Read on the entries, the three checks of Table 7.8 differ in what they see. The clamp-on reading is a row on the load alone, $\mathbf{J}_a = \mathbf{e}_1^{\mathsf{T}}/1$, the entry in which the two drawings differ most: under the unit-2 drawing the hall meter is right and the load sits near its reading, at 195.620 kW, and under the hall-meter drawing the load is what the unit meters and the temperatures say, 205.778 kW. The portable check of unit 2 is a row on (Q, g_1, g_2) through the gradient of Q_2, equation (7.9): it reads the quantity whose meter is suspected, without that meter's bias. The portable check of unit 1 reads Q_1, which the two drawings put at 105.078 and 107.552 kW, close together because both hold the unit-1 meter sound; it can move the ranking little. Read against the alternative outside the family, the two best checks part ways (Section 7.4.5). The permanent instrument at the place of the question is a water-side heat meter on unit 2, of the same 2 kW: it is the same row as the portable check, added in every window, so it separates the leaders by the same 2.51 tonight and gives the unit-2 meter a partner in every later window. A permanent meter on the UPS output of 1 kW is likewise the clamp-on row in every window; it would check the hall meter directly, but it would leave a unit-meter bias and an other-heat error to be told apart through the rows.

The predicted readings. Figure 7.5 draws the predictive densities (7.22) of the two best checks. In each panel the horizontal axis is the reading the check would show, in kW, and the vertical axis its predictive density; blue is the prediction under the unit-2 drawing and red under the hall-meter drawing, each a normal with the mean and spread of Table 7.8. The clamp-on curves overlap less than the portable ones, which is what their separations of 3.18 and 2.51 say, and their spreads differ, the red curve wider, which is why the clamp-on reading is worth more if the hall meter is the one. A portable reading near 99 kW would point to the unit-2 meter, one near 106 kW to the hall side; a reading midway, where the two curves cross, would leave the question open.

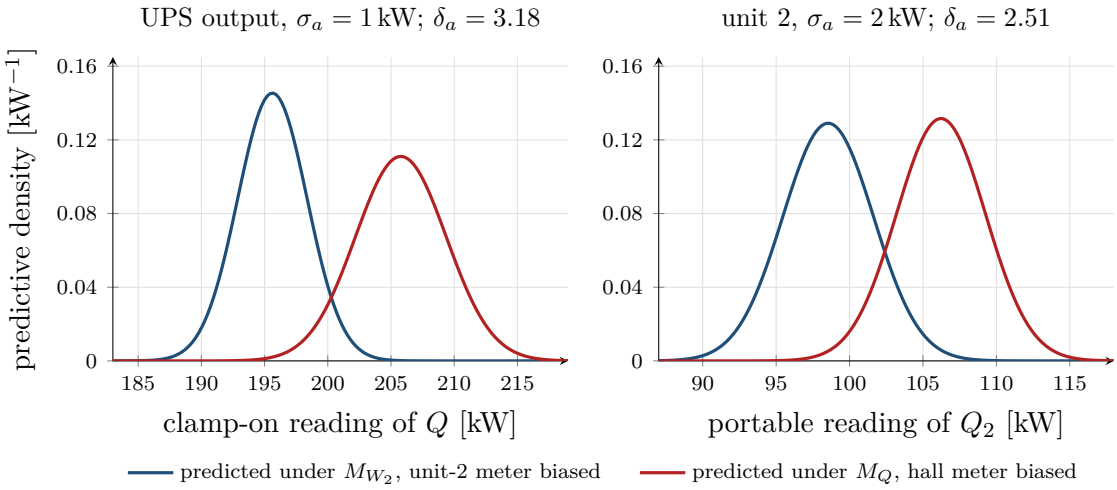


Figure 7.5: The predicted readings of the two best checks under the two leading drawings, equation (7.22). Left, a clamp-on reading of the UPS output: $\mathcal{N}(195.620,\ 2.744^2)$ under M_{W_2} and $\mathcal{N}(205.778,\ 3.594^2)$ under M_Q. Right, a portable water-side heat check of unit 2: $\mathcal{N}(98.542,\ 3.092^2)$ and $\mathcal{N}(106.227,\ 3.032^2)$. The spreads include each check's own σ_a.

7.6.2 The recommendation

The next reading is the portable water-side heat check of unit 2 tonight. It is expected to add 3.21 to the log Bayes factor in favor of the true drawing if the unit-2 meter is biased and 3.09 if the hall meter is, against the present 0.59; the check of unit 2 separates the leading explanation from both alternatives, the hall-meter bias and an other heat larger than pinned. The clamp-on reading of the UPS output is worth more between the leaders, 4.05 or 6.94, and keeps that lead to about 2.7 kW (Table 7.9), but it cannot tell the unit-2 explanation from an other-heat error (0.53). That is the reason for the order. It is the second request: if the check of unit 2 is inconclusive, or if it points to the hall side, where, of the three checks, only the clamp-on reading tells a hall-meter bias from an other-heat error. Made after a clamp-on reading of 200.0 kW, the check of unit 2 adds 0.49 between the leaders, against its 3.15 made first, but 1.88 against the other-heat explanation; made first, it leaves the clamp-on reading 2.43 at a midway reading (Section 7.3.4). The portable check of unit 1 is not worth the visit for this question. One check whose reading lands near the middle of Figure 7.5 can leave the question open tonight, and the entry then says so. No meter is recalibrated before the checks report; if they convict a meter, that meter is recalibrated and the window refitted without its bias, and if they leave the question open, the hall's heat is reported with both explanations into the next window. The request beyond tonight is a permanent water-side heat meter on unit 2. Chapter 18 ranks the readings an operator can request on the shift plant in the same way.

7.6.3 What the analysis showed

The dashboard showed a mismatch of 15.56 kW and invited the recalibration of one meter by that amount; the analysis found two meters under suspicion, at 0.60 and 0.33, neither off by the full mismatch. A meter balance is one equation, and it cannot isolate a meter on its own. What isolates a meter is the rest of the hall: the temperatures, through the coil rows, know each unit's heat and the hall's load independently of the meters, and they put the unit-1 meter in the clear and keep the other two in contention. The split between the two survivors comes from the hall's design, not from the method. The hall meter and the unit-2 meter are checked against each other mainly through the balance and the temperature rows, and a low hall reading and a high unit-2 reading open the balance in the same direction, so a window in which the hall meter happens to read low by noise looks much like a window in which the unit-2 meter is biased. The hall's metering was placed to bill the load and to watch each coil, and it watches each coil well; it was not placed to check its own meters. The check to make first is therefore a second view of unit 2's heat, which no explanation in play can imitate on both sides; if it points to the hall side, only a view of the load, or of the other heat, finishes the question.

7.6.4 Limits of the model

One meter at a time. The family holds one biased meter per drawing. Two biased meters at once, the unit-2 meter high and the hall meter low, are not a member, and their evidence is not computed; a window in which both are off would be ranked as one of the two, with the bias of each overstated.

A constant bias. Within one window a bias is a constant, because one window gives one reading per meter. A meter whose error grows with its flow, or drifts across the night, is a bias that varies in time, and only a run of windows could see it. The chapter does not treat it.

The width of the bias prior. The evidence charges each drawing for its bias through σ_b, and the probabilities move with it (Table 7.11). The ranking holds from 5 to 20 kW, but the

width is a judgment of how large a meter error is plausible, not a measured quantity, and a calibration history of the meters would replace it.

The pinned other heat. The lighting, fans and other heat are pinned at 8 kW, and the family holds no drawing with that figure wrong; whether one is given a prior share moves the headline (Section 7.4.5). A reading of the other heat itself would replace the pinned figure and test the freed fit's 19.69 kW, sd 4.2766 kW, against it.

Steady state and the pinned airflow. The coil rows that tie the unit meters to the temperatures hold only once the hall has settled, with the units' airflow at its pinned value and no controller acting, as in Chapter 5. After a load step or with a valve moving, the unit meters and the temperatures disagree for reasons that are not a meter's, and the family would blame a meter for a transient.

Constructed readings. The seven readings are constructed: they are generated from a stated truth of the hall with the stated sensor errors and a stated bias on the unit-2 meter, and Section 7.4.6 compares the answers with that truth. They are not field data, and the agreement of the ranking with the truth is a check of the method on its own model, not of the model against a real hall.

Chapter 8 returns to the racks and asks whether the hall is inside its operating envelope where no sensor reads it.

7.7 What to retain

Definition 7.1 (Meter reconciliation by evidence)**.** For a set of meters that should balance, each meter j defines a drawing M_j of the plant with a bias $b_j \sim \mathcal{N}(0,\ \sigma_b^2)$ freed on its reading, against M_0 with every bias at zero. Each drawing is fitted on its own state, and its Laplace evidence gives the log Bayes factor $\ln B_{j0} = [\mathcal{L}_0(\hat{\mathbf{z}}_0) - \mathcal{L}_j(\hat{\mathbf{z}}_j)] - \frac{1}{2}\ln(\det \mathbf{\Lambda}_j / \det \mathbf{\Lambda}_0) - \ln \sigma_b$, a gain in fit less an Occam factor. The posterior probabilities $\Pr(M_j \mid \mathbf{y}) \propto B_{j0}\ \Pr(M_j)$ rank the meters, the bias within a drawing is the entry $\hat{b}_j$ with $\sqrt{(\mathbf{\Sigma}_j)_{77}}$, and the check that separates two leaders is the one, among those no outside explanation can imitate, whose predicted readings under them differ most in units of their spreads.

Three conclusions carry beyond the hall. A balance names suspects and the rest of the plant ranks them: the bad-data tests, the chi-square of the residuals and the largest normalized residual, say that something is off and point at a reading, and the evidence turns the pointing into a probability. The leverage of a reading decides how far it can be checked at all: the redundancy numbers add to the degrees of freedom, and a reading with none left has no bias that any fit can see. And when two drawings explain a window almost equally well, the next reading is one their fitted states disagree on, valued before it is taken by the evidence it is expected to add in each direction, and preferably one that a confounded explanation outside the family, such as a pinned quantity in error, cannot imitate.

Table 7.14 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 7.14: The terms of meter reconciliation and where each is defined. Units for the two-unit hall.

Term	Expression	Units	Meaning	Defined in
Unit heat	$Q_u = G_u(Q + Q_o)/(G_1 + G_2)$	kW	what a unit meter reads	7.1.2
Meter balance	$Q_1 + Q_2 = Q + Q_o$	kW	holds on the reduced state	7.1.2
Meter row	$(Q_u(\mathbf{z}) - y_{W_u})/3$	–	a unit meter on the state	7.1.4
Candidate fault	M_j, bias $b_j \sim \mathcal{N}(0,\ \sigma_b^2)$ on meter j	kW	one meter at fault	7.1.4
Rows and posterior				
Rows of a drawing	eq. (7.8)	–	bias added to one meter's row	7.2.1
Gradient of a unit heat	eq. (7.9)	kW	a meter's Jacobian line	7.2.1
Leverage, redundancy	h_{kk}, $1 - h_{kk}$ of $\mathbf{H} = \mathbf{J}\mathbf{\Sigma}\mathbf{J}^{\mathsf{T}}$	–	share of a row in its own fit	7.2.4
Degrees of freedom	$\sum_k (1 - h_{kk}) = n - d$	–	expected misfit	7.2.4
Normalized residual	$r_k/\sqrt{1 - h_{kk}}$	–	bad-data test of one reading	7.2.4
Check row	eq. (7.10)	–	a check once read, in a refit	7.2.1
Laplace evidence	eq. (7.16)	–	volume times height of the posterior	7.2.5
Log Bayes factor	eq. (7.17)	–	gain in fit less Occam factor	7.2.5
Closed form	eq. (7.18)	–	evidence from the residual test	7.2.5
Questions and answers				
Posterior of a drawing	eq. (7.20)	–	the ranking of the meters	7.3.1
Bias of a drawing	$\hat{b}_j$, $\sqrt{(\mathbf{\Sigma}_j)_{77}}$	kW	size of the fault	7.3.3
Value of a check	eq. (7.23), each direction	–	expected change of $\ln B$	7.3.4
Separation of a check	δ_a, $U(a) \approx \delta_a^2/2$	–	one-number summary	7.3.4

7.8 Exercises

Exercise 7.1. *Formulation.* Write the drawing M_{W_1} of a bias on the unit-1 meter: its state, the row that changes, the row it adds, and its standing in the ranking. (State $(Q, g_1, g_2, \beta_1, \beta_2, \beta_3, b_{W_1})^{\mathsf{T}}$; the unit-1 meter row becomes $(Q_1(\mathbf{z}) + b_{W_1} - 106.45)/3$ and the prior row $b_{W_1}/10$ is appended. Its fit removes little misfit, $2\mathcal{L}$ from 12.9934 to 12.0808, too little to pay for the entry: $\ln B_{W_1 0} = -0.5548$, a fitted bias of 3.436 kW, and a posterior probability of 0.0114.)

Exercise 7.2. *Derivation.* Show that the redundancy numbers of all rows of a fit add to the number of rows less the number of entries, and evaluate the sum for M_0 at 23:00. ($\sum_k h_{kk} = \mathrm{tr}(\mathbf{J}\mathbf{\Sigma}\mathbf{J}^{\mathsf{T}}) = \mathrm{tr}(\mathbf{\Sigma}\mathbf{J}^{\mathsf{T}}\mathbf{J}) = \mathrm{tr}\,\mathbf{I}_d = d$ by the cyclic property of the trace, so $\sum_k (1 - h_{kk}) = n - d = 12 - 6$; the fit gives 6.0000, of which the seven reading rows take 3.1807.)

Exercise 7.3. *Derivation.* From the closed form (7.18), explain why the unit-1 meter's drawing loses to M_0 while the other two win. (The gain term $\frac{1}{2}[\psi_k/(1 + \psi_k)](r_k^{\mathrm{N}})^2$ is at most $\frac{1}{2}(r_k^{\mathrm{N}})^2$, and the Occam term is $\frac{1}{2}\ln(1 + \psi_k)$ with $\psi_k = 10^2(1 - h_{kk})/3^2$. The two unit meters have nearly equal leverages, 0.3949 and 0.3948, so they pay nearly the same Occam term, $\psi = 6.723$ and 6.725, and differ only in their normalized residuals. The drawing wins when $(r_k^{\mathrm{N}})^2$ exceeds the break-even $(1 + \psi_k)\ln(1 + \psi_k)/\psi_k$, 2.348 for both: the unit-1 meter's 1.049 lies below it

and the unit-2 meter's 10.089 far above, and the hall meter's 9.152 lies far above its 2.079, at $\psi = 4.470$. The closed form gives $\ln B_{W10} = -0.565$, and the Laplace evidence -0.5548.)

Exercise 7.4. *Application.* The maintenance history of the hall meter shows no drift in five years, and the engineer wants to narrow the bias prior to 5 kW for every meter. What happens to the ranking? (A narrow prior makes a bias of 10 to 12 kW less plausible and returns probability to the drawing without one; the ranking holds, with $\Pr(M_{W_2}) = 0.5695$, $\Pr(M_Q) = 0.2819$, $\Pr(M_{W_1}) = 0.0330$ and $\Pr(M_0) = 0.1156$.)

Exercise 7.5. *Application.* From the predicted readings of Table 7.8, compute the separation and the expected change of evidence of the portable check of unit 2. ($\delta_a = |98.542 - 106.227|/\sqrt{\tfrac{1}{2}(3.092^2 + 3.032^2)}$, the difference of the two predictions over the root mean square of their spreads, which is 2.51; by equation (7.24) the equal-spread expected change of the log Bayes factor is $\tfrac{1}{2} \times 2.51^2 = 3.15$, and by equation (7.23), with spreads this close, the two directions are 3.21 if the unit-2 meter is biased and 3.09 if the hall meter is.)

Chapter 8

Joint Limits and the Operating Envelope

8.1 The operating envelope over racks without sensors

It is 23:00 again, in the window of the two-unit hall that Chapter 5 read. Rack 1's inlet sensor reads 22.18°C, and the tile labeled "rack inlets" is green. Racks 2 and 3 have no sensor. Chapter 5 has already put rack 3, the rack with the most recirculation, at a 0.75 % chance of being over the 27°C limit, and the night operator has read that number. The worry now is wider than one rack: two inlets go unwatched, and either of them could be the one that is over. The dashboard's answer is the rack it measures, 4.82 K inside the limit, in green.

That number is about rack 1, and Chapter 5's number is about rack 3 alone. Neither is the probability that *some* unmeasured inlet is over the limit, which is a statement about two temperatures at once; and the two temperatures are read from the same fit, through the same cold aisle, return air and coils, so they need not be independent. Rack 1 is left out of that event: it reads 4.82 K inside the limit with a sensor of standard deviation 0.3 K, and its own exceedance is negligible. The question of this chapter is

what the probability is that any unmeasured inlet is over its limit, how far the joint of the inlets moves it from the inlets taken one by one, and when the difference matters.

It is answered from the posterior of Chapter 5: the same hall, the same five readings and five priors, the same fit. One posterior serves Chapters 5, 6 and 8. What changes is the event. It is no longer a tail of one derived quantity but a region of the plane of two, and its probability needs their joint distribution, not their marginals.

Jargon: operating envelope, recommended envelope, inlet limit, hot aisle, joint event, union, intersection, joint exceedance, correlation, union bound, bivariate normal

The *recommended envelope* is the range of rack inlet temperatures in which the racks are meant to run, the ASHRAE recommended range; its upper end, the *inlet limit*, is 27°C here, and it applies to the air entering a rack. The *operating envelope* is the region of the state inside which every limit of the plant holds at once; for the inlet limit of three racks it is the set of states with T_1, T_2 and T_3 all at or below 27°C. The *hot aisle* is the space behind the racks into which they exhaust, the return air of Chapter 5; it is above the inlet limit by design, and the limit is not read against it. A *joint event* is an event about several quantities at once. A *union* of events holds when at least one of them holds ("some rack is over"); an *intersection* holds when all of them hold ("both racks are over"). A *joint exceedance* is the union of the exceedances of the quantities under one limit, the event that the plant has left its operating envelope. The *correlation* of two quantities is their covariance divided by the product of their standard deviations, a number between -1 and 1. The *union bound* is the sum of the probabilities of the events, an upper bound on the probability of their union. The *bivariate normal* is the joint normal distribution of two quantities, fixed by their two means, their two standard deviations and their correlation.

Key idea

An exceedance anywhere among several unmeasured quantities is the union of their exceedances, and its probability is one minus the probability of the corner of the operating envelope, computed from the joint posterior of the quantities. On the fitted state that joint is the normal with covariance $\mathbf{D\Sigma D}^{\mathsf{T}}$, $\mathbf{D}$ the matrix whose rows are the quantities' gradients. The union always lies between the largest single exceedance and their sum; when one quantity sits much nearer its limit than the others, the bracket is narrow and the union is that quantity's own tail, whatever the correlation. The quantities are correlated only through the entries their gradients share, and the correlation matters for the intersection: "both over" is the product of the marginals only when the gradients are orthogonal under $\mathbf{\Sigma}$, which nothing here arranges, and it is farthest from that product when the shared entries are poorly read. There the linearization behind the joint normal errs as well, and draws through the physics, set beside the reported value, show by how much.

8.1.1 The operating envelope as a joint event

Chapter 2 wrote every question as the posterior probability of a region of the state, equation (2.7), and noted that a joint event changes the region, not the category. A limit on one rack inlet cuts the state space into two parts. A limit on two cuts the plane of their inlets into four: both inside, the first over, the second over, both over. The operating envelope is the first of the four, the lower-left corner below 27°C in each, and the operator's question is about the other three together, the complement of the corner. *Reliability Engineering by Derivation* meets the same structure twice, in its Chapter 2's two pumps as a joint event and in its Chapter 13's series system, which fails when any component fails: a hall is inside its envelope only if every rack is.

Two facts follow from the probability rules alone and need no model. The probability of the union lies between the largest probability of one event and the sum of them all, so the inlets taken one by one already bracket the answer. And the probability of the intersection is fixed by the union and the marginals through inclusion and exclusion, so the two joint questions, "some rack over" and "both racks over", are one computation. Where in the bracket the answer falls, and how far the intersection is from the product of the marginals, depends on how the inlets move together. That is read from the joint posterior, and on the fitted state it is read from the

gradients and $\boldsymbol{\Sigma}$.

Principle 8.1. *The probability that some quantity leaves its limit is the posterior probability of the complement of the operating envelope, a union of exceedances; it lies between the largest single exceedance and their sum, and the joint posterior places it within that bracket.*

8.1.2 The two-unit hall (step 0)

The plant is the two-unit hall of Chapter 5, unchanged: two cooling coils between the return air T_r and chilled water pinned at $T_s = 14°\text{C}$, the cold aisle T_c mixed from the two discharge streams, and three racks whose inlets mix cold-aisle and return air with the recirculation fractions β_k. Its seven physical rows are equation (5.1), and the fit runs on the reduced state of Chapter 4,

$$\mathbf{z} = (Q,\ g_1,\ g_2,\ \beta_1,\ \beta_2,\ \beta_3)^{\mathsf{T}}, \tag{8.1}$$

in kW and five dimensionless entries. The dependent block $\mathbf{A}_{\text{dep}}$ of the physical rows is triangular with nonzero diagonal, so it is invertible and the seven temperatures follow from the free entries in the forward chain of equation (5.4): the hall balance fixes T_r, the air sides the discharges, the mixing T_c, and the inlet rows $T_k = T_c + \beta_k(T_r - T_c)$. The two inlets of this chapter are T_2 and T_3. Both are functions of the same T_c and T_r, and they differ only in their fractions. The structural identity that matters here is the air balance (5.5): the rise from the cold aisle to the return is common to every rack and is set by the heat and the airflow alone.

Modeling choices at the graph. The choices that no question owns are Chapter 5's and hold unchanged (Section 5.1.2), and no new row enters the plant.

The graph. Figure 8.1 draws the hall's rows as Chapter 5 drew them, with the two unmeasured racks in red. The cold aisle on the left (blue) feeds the three racks; the return air in the middle (blue), the hot aisle, receives the IT load Q, a free node (green); the two units on the right (green) pass heat to the pinned chilled water (gray) and return their discharge air to the cold aisle along the top and bottom. Each rack's inlet is a node fed by the cold aisle with weight $1 - \beta_k$ and by the return air with weight β_k. The dashed arrows are the five readings: in gray the two that read the load and rack 1, in purple the three that read the return air and the two discharges, which are withheld in the second reading set of Section 8.1.4. The return air reads 30.33°C, above 27°C as a hot aisle is; the limit is read at the racks' inlets.

8.1.3 The operator's questions (step 1)

The operator asks four questions of the 23:00 window, and step 1 writes each as a probability or a moment over the reduced state, with the unknown on the left of the bar and the readings and priors on the right. Table 8.1 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 in the notation of Chapter 2, for derived quantities $f_k(\mathbf{z})$ and a level f^*; the third its mathematical form for the hall; the fourth its category; the fifth the computation of Section 2.3.6 that answers it; and the last the subsection where it is answered. The first and third are predictions under a threshold whose regions are a union and an intersection. The second is estimation, of a moment shared by two derived quantities. The fourth asks the first three again under a second reading set, in which the entries the two inlets share are poorly read; it is the same category on a second posterior. Which reading would lower the probability of the first question most asks nothing new of the present evidence; it is the next reading, Section 8.6.

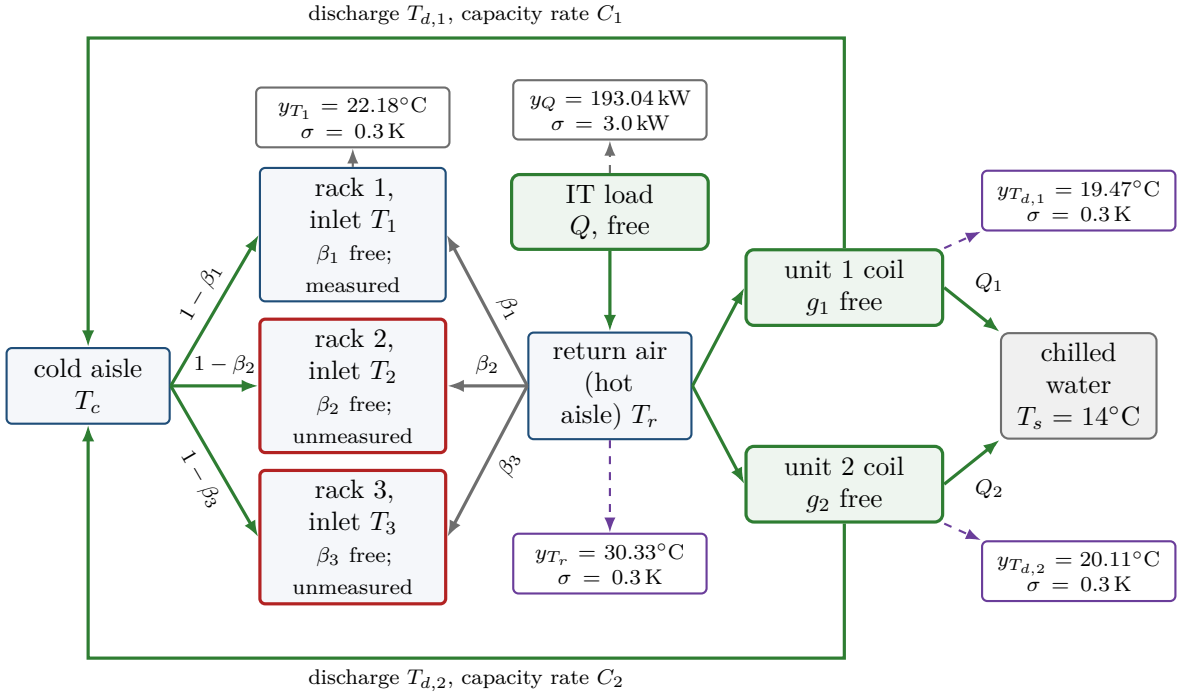


Figure 8.1: The rows of the two-unit hall in its 23:00 window, as in Figure 5.1. Free (green): the IT load Q and the coils' log-derates g_1, g_2; the fractions β_1, β_2, β_3 are free entries carried by the rack nodes. Blue: temperatures fixed by the physical rows, the cold aisle, the return air (the hot aisle) and the rack inlets; racks 2 and 3 (red) are the unmeasured inlets of the joint question. Pinned (gray): the chilled water $T_s = 14°\mathrm{C}$, and, not drawn, the ratings $\bar{G}_u = 7.0\,\mathrm{kW/K}$, the capacity rates $C_u = 10.0\,\mathrm{kW/K}$ and the other heat $Q_o = 8\,\mathrm{kW}$. Observed (dashed): gray, the load meter and rack 1's inlet, kept in both reading sets; purple, the return air and the two discharges, withheld in the second. The priors on g_1, g_2 and the β_k are rows not drawn.

Table 8.1: The questions asked of the two-unit hall at 23:00 about its two unmeasured inlets. $\mathbf{y}$ is the five readings and the five priors; $\mathbf{y}_{\mathrm{II}}$ the meter, rack 1's inlet and the five priors; $T_k(\mathbf{z})$ the inlet of rack k by equation (5.4); $T_{\max} = 27°\mathrm{C}$; $\mathbf{D}$ the 2×6 matrix of the two inlets' gradients.

Operator question	General form	Mathematical form	Category	Computation	Answered in
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 $p(\mathbf{z} \mid \mathbf{y}_{\mathrm{II}})$	the first three without the return-air and discharge readings	prediction under a threshold	posterior	8.3.4

8.1.4 The evidence as rows (step 2)

Step 2 attaches the evidence of the 23:00 window as rows on the reduced state, and the first reading set is Chapter 5's (Section 5.1.4): the return air $y_{T_r} = 30.33$°C, the two discharges $y_{T_{d,1}} = 19.47$°C and $y_{T_{d,2}} = 20.11$°C, each of standard deviation 0.3 K; the load meter $y_Q = 193.04$ kW, of standard deviation 3.0 kW; rack 1's inlet $y_{T_1} = 22.18$°C, of standard deviation 0.3 K; and the five priors $g_u \sim \mathcal{N}(0,\ 0.15^2)$ and $\beta_k \sim \mathcal{N}(\beta_k^{\mathrm{s}},\ 0.08^2)$ with the survey's 0.20, 0.35 and 0.50. The second reading set withholds the return-air and the two discharge readings and keeps the meter, rack 1's inlet and the five priors: it is the same window with the three readings that see the coils taken out.

Modeling choice at the rows. The second set is the first with three reading rows removed and nothing else changed. It stands for a window in which the return-air and discharge sensors are out of service: this is what isolates the effect of poorly read shared entries on the joint questions, and it fails if a real outage comes with other changes, of the load, the coils or the other sensors, that the variation does not carry. It is a variation on the 23:00 window, not a record of another one.

Which row acts on which entry. With $*$ for an entry that can be nonzero and the columns $Q, g_1, g_2, \beta_1, \beta_2, \beta_3$, the Jacobian of the first set is the 10×6 pattern of equation (5.6). The second set drops its first three rows and leaves the 7×6 pattern

$$\mathbf{J}_{\mathrm{II}} = \begin{bmatrix} * & 0 & 0 & 0 & 0 & 0 \\ * & * & * & * & 0 & 0 \\ 0 & * & 0 & 0 & 0 & 0 \\ 0 & 0 & * & 0 & 0 & 0 \\ 0 & 0 & 0 & * & 0 & 0 \\ 0 & 0 & 0 & 0 & * & 0 \\ 0 & 0 & 0 & 0 & 0 & * \end{bmatrix} \begin{matrix} \text{load meter } y_Q \\ \text{rack-1 inlet } y_{T_1} \\ \text{prior on } g_1 \\ \text{prior on } g_2 \\ \text{prior on } \beta_1 \\ \text{prior on } \beta_2 \\ \text{prior on } \beta_3 \end{matrix}, \tag{8.2}$$

$$\mathbf{g}_{T_2}^{\mathsf{T}} = (*,\ *,\ *,\ 0,\ *,\ 0), \qquad \mathbf{g}_{T_3}^{\mathsf{T}} = (*,\ *,\ *,\ 0,\ 0,\ *).$$

The two gradients on the second line are not rows; they have the pattern of the inlets in equation (5.4), and they are the same in both reading sets. Three facts of the graph are read here before any number. The two gradients share the columns of Q, g_1 and g_2 and nothing else: each has its own fraction, in a column the other does not touch. In both sets the columns of β_2 and β_3 hold their prior rows alone, so neither fraction is observable, as in Chapter 5. And the shared columns are read very differently by the two sets. In the first, three temperature readings act on (Q, g_1, g_2) together with the meter. In the second, the log-derates are reached by one reading row only, rack 1's inlet, which also carries β_1: the difference of the log-derates is held by its priors alone and their total by that one row.

8.2 The posterior of the state (step 3)

Each of the two reading sets gives one posterior of the reduced state, and the four questions of Table 8.1 are read from them. They differ in the block of (Q, g_1, g_2, β_1) alone.

8.2.1 The fit with every reading

The first set is Chapter 5's, and so is its posterior (Section 5.2.3): the residual vector (5.7) of ten rows, the normal equations (5.8), five Gauss–Newton iterations (Table 5.2), the MAP state $\hat{\mathbf{z}}$ of Table 5.3 and the covariance $\boldsymbol{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1}$ displayed there with its correlations. The posterior is $\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$.

8.2.2 The refit with the meter and rack 1 only

The rows. The residual vector of the second set is that of the first with its first three rows removed,

$$\mathbf{r}_{\mathrm{II}}(\mathbf{z}) = \begin{bmatrix} (Q - y_Q)/3.0 \\ (T_1(\mathbf{z}) - y_{T_1})/0.3 \\ g_1/0.15 \\ g_2/0.15 \\ (\beta_1 - 0.20)/0.08 \\ (\beta_2 - 0.35)/0.08 \\ (\beta_3 - 0.50)/0.08 \end{bmatrix}, \qquad \mathcal{L}_{\mathrm{II}}(\mathbf{z}) = \tfrac{1}{2}\,\mathbf{r}_{\mathrm{II}}(\mathbf{z})^{\mathsf{T}}\mathbf{r}_{\mathrm{II}}(\mathbf{z}), \tag{8.3}$$

seven rows on six entries, every entry dimensionless. It is minimized by the normal equations (5.8) with $\mathbf{J}_{\mathrm{II}}$ in place of $\mathbf{J}$, from the same start $\mathbf{z}^{(0)} = (y_Q,\ 0,\ 0,\ 0.20,\ 0.35,\ 0.50)^{\mathsf{T}}$, and its Laplace covariance is the Gauss–Newton form

$$\boldsymbol{\Sigma}_{\mathrm{II}} = \left(\mathbf{J}_{\mathrm{II}}^{\mathsf{T}}\mathbf{J}_{\mathrm{II}}\right)^{-1}\Big|_{\hat{\mathbf{z}}_{\mathrm{II}}}, \qquad \mathbf{z} \mid \mathbf{y}_{\mathrm{II}} \approx \mathcal{N}(\hat{\mathbf{z}}_{\mathrm{II}},\ \boldsymbol{\Sigma}_{\mathrm{II}}). \tag{8.4}$$

The columns of β_2 and β_3 hold one prior row each, so this covariance has the block form (5.10): the fractions of the two unmeasured racks keep the survey's center and spread.

The iterations. Table 8.2 runs the normal equations from $\mathbf{z}^{(0)}$. Each row is an iteration; the columns give the objective and the norm of its gradient. The objective falls from 19.253 to 0.5505 in the first step, and the norm of the gradient falls from 295.5 to 7.5×10^{-8} by the fifth.

Table 8.2: Gauss–Newton iterations of the second reading set, the residual vector (8.3), from $\mathbf{z}^{(0)}$.

Iteration k [–]	Objective $\mathcal{L}_{\mathrm{II}}(\mathbf{z}^{(k)})$ [–]	$\lVert\mathbf{J}_{\mathrm{II}}^{\mathsf{T}}\mathbf{r}_{\mathrm{II}}\rVert$ [–]
0	19.253081	2.955×10^{2}
1	0.550544	1.175×10^{1}
2	0.523998	1.043×10^{-2}
3	0.523998	2.267×10^{-4}
4	0.523998	3.606×10^{-6}
5	0.523998	7.487×10^{-8}

The Jacobian at the MAP state. With the columns $Q, g_1, g_2, \beta_1, \beta_2, \beta_3$ and the rows in the order of equation (8.3),

$$\mathbf{J}_{\mathrm{II}}(\hat{\mathbf{z}}_{\mathrm{II}}) = \begin{bmatrix} 0.3333 & 0 & 0 & 0 & 0 & 0 \\ 0.1347 & -26.3829 & -26.3829 & 33.5395 & 0 & 0 \\ 0 & 6.6667 & 0 & 0 & 0 & 0 \\ 0 & 0 & 6.6667 & 0 & 0 & 0 \\ 0 & 0 & 0 & 12.5 & 0 & 0 \\ 0 & 0 & 0 & 0 & 12.5 & 0 \\ 0 & 0 & 0 & 0 & 0 & 12.5 \end{bmatrix}, \tag{8.5}$$

the first column in kW^{-1} and the others dimensionless. The zeros sit where the pattern (8.2) places them. The two reading rows have rank 2 of 6, and the rack-1 row has equal entries in the two log-derate columns: it sees the coils only through their total conductance.

The information and the covariance. The observed block of the information matrix, in the order (Q, g_1, g_2, β_1), is

$$\mathbf{\Lambda}_{\mathrm{o,II}} = \begin{bmatrix} 0.1293 & -3.5535 & -3.5535 & 4.5174 \\ -3.5535 & 740.5028 & 696.0583 & -884.8698 \\ -3.5535 & 696.0583 & 740.5028 & -884.8698 \\ 4.5174 & -884.8698 & -884.8698 & 1281.1479 \end{bmatrix},$$

with $\mathbf{\Lambda}_{55} = \mathbf{\Lambda}_{66} = 156.25$ and zeros elsewhere in the rows and columns of β_2 and β_3. The matching block of the covariance (8.4) and its correlations are

$$\begin{bmatrix} 8.96297357 & 0.01813206 & 0.01813206 & -0.00655659 \\ 0.01813206 & 0.01362064 & -0.00887936 & 0.00321080 \\ 0.01813206 & -0.00887936 & 0.01362064 & 0.00321080 \\ -0.00655659 & 0.00321080 & 0.00321080 & 0.00523897 \end{bmatrix},$$

$$\begin{bmatrix} 1 & 0.0519 & 0.0519 & -0.0303 \\ 0.0519 & 1 & -0.6519 & 0.3801 \\ 0.0519 & -0.6519 & 1 & 0.3801 \\ -0.0303 & 0.3801 & 0.3801 & 1 \end{bmatrix},$$

in kW^2, kW and dimensionless units, the fractions of racks 2 and 3 with the survey's variance 0.0064 and no covariance with any entry. The log-derates' variance, 0.01362, is now of the order of their priors' 0.15^2, against 0.00056 and 0.00063 with every reading. Their correlation is strongly negative, -0.6519. The meter fixes the load, and rack 1's inlet then fixes its temperature through the return air, which depends on the coils only through their total conductance $G_1 + G_2$; the readings fix that total better than either coil, so a larger g_1 is explained by a smaller g_2. The inlets' gradients will be equal in g_1 and g_2, so $g_1 + g_2$ is what matters: its posterior spread is 0.0974 against the priors' 0.2121 and 0.0263 with every reading, while $g_1 - g_2$ stays at its priors' 0.2121. The load is read by the meter almost alone, and its correlation with each log-derate falls to 0.0519.

The MAP state. Table 8.3 gives the MAP state with its posterior spreads and the prior spread of each entry; each row is an entry. The two log-derates come out equal, since no row tells the coils apart and their priors are the same. By equation (5.4), $T_r = 29.8298°\mathrm{C}$ and $T_c = 19.7679°\mathrm{C}$, a rise of 10.0618 K, which is $(Q + Q_o)/(C_1 + C_2)$ of equation (5.5).

Table 8.3: The MAP state of the second reading set with the posterior and prior spreads of each entry.

Entry [unit]	$\hat{z}_i$ [entry's]	$\sqrt{\Sigma_{\mathrm{II},ii}}$ [entry's]	Prior sd [entry's]
Q [kW]	193.23699	2.99382	none
g_1 [–]	−0.09647	0.11671	0.15
g_2 [–]	−0.09647	0.11671	0.15
β_1 [–]	0.23488	0.07238	0.08
β_2 [–]	0.35000	0.08000	0.08
β_3 [–]	0.50000	0.08000	0.08

8.3 The answers (step 4)

Every answer is read from the posterior of the state, $\mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$ of Section 8.2.1 or $\mathcal{N}(\hat{\mathbf{z}}_{\mathrm{II}}, \mathbf{\Sigma}_{\mathrm{II}})$ of Section 8.2.2, through the joint normal of the two inlets, and that value is the one reported.

Beside it stand draws: states drawn from the same normal, each pushed through the exact equations (5.4), and the fraction $\hat{p}$ of n draws in the event, with Monte Carlo standard error $\sqrt{\hat{p}(1-\hat{p})/n}$. Each set has one draw of 200 000, and the second set a further draw of 1 000 000, because at 200 000 the gap of its intersection is not resolved; every statistic of a draw comes from that draw, and where the second set's draws are quoted without a count they are the 1 000 000.

8.3.1 The hottest unmeasured inlet

Operator question. What is the probability the hot aisle is already above 27°C somewhere I cannot see? The operator's "hot aisle" here means the air the racks draw in; the return air behind them is above 27°C by design (Section 8.1.2). The question asks for the probability that at least one of the two unmeasured inlets is over the inlet limit at 23:00.

General form. A joint tail, the tail of equation (2.12) with $f = \max_k f_k$, whose region is a union of exceedances, the complement of the corner of the operating envelope in which every $f_k \le f^*$:

$$\Pr\Big(\bigcup_k \{f_k(\mathbf{z}) > f^*\} \mid \mathbf{y}\Big) = \int \mathbf{1}\big[\max_k f_k(\mathbf{z}) > f^*\big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}. \tag{8.6}$$

Specialized form. The question is equation (8.6) with $\mathbf{z}$ the reduced state (8.1), $f_2 = T_2(\mathbf{z})$ and $f_3 = T_3(\mathbf{z})$ of equation (5.4), $f^* = T_{\max} = 27$°C and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal of Section 8.2.1. The assurance level κ of equation (2.12) is not set here: no acceptance rule is in question at 23:00, and such a rule applies to the requirement, never to the exceedance, so the probability itself is reported. One choice enters here: the joint of the two inlets is read as the bivariate normal of the delta method. This gives the union in closed form from $\hat{\mathbf{z}}$, $\boldsymbol{\Sigma}$ and one matrix of gradients, and it is exact only if both inlets are linear in the state over the posterior's width; a curvature of the inlets in the log-derates would bend the joint's tails, which the draws beside each answer measure at the corner.

The event $\max(T_2, T_3) > T_{\max}$ is the union $\{T_2 > T_{\max}\} \cup \{T_3 > T_{\max}\}$. Its complement is the corner $\{T_2 \le T_{\max},\ T_3 \le T_{\max}\}$, so

$$\Pr(\max(T_2, T_3) > T_{\max} \mid \mathbf{y}) = 1 - \Pr(T_2 \le T_{\max},\ T_3 \le T_{\max} \mid \mathbf{y}). \tag{8.7}$$

The matrix whose rows are the gradients of the two inlets, equation (5.12) with each rack's own index, carries the posterior of the state into the plane of (T_2, T_3):

$$\mathbf{D} = \begin{bmatrix} \mathbf{g}_{T_2}^{\mathsf{T}} \\ \mathbf{g}_{T_3}^{\mathsf{T}} \end{bmatrix} = \begin{bmatrix} \dfrac{1}{G_1+G_2} - \dfrac{1-\beta_2}{C_1+C_2} & -\dfrac{G_1(T_r-T_s)}{G_1+G_2} & -\dfrac{G_2(T_r-T_s)}{G_1+G_2} & 0 & T_r - T_c & 0 \\ \dfrac{1}{G_1+G_2} - \dfrac{1-\beta_3}{C_1+C_2} & -\dfrac{G_1(T_r-T_s)}{G_1+G_2} & -\dfrac{G_2(T_r-T_s)}{G_1+G_2} & 0 & 0 & T_r - T_c \end{bmatrix}, \tag{8.8}$$

2×6, the first column in K/kW and the others in K. The two rows agree in the log-derate columns, where a derate warms the return and the cold aisle alike and so every rack alike; they differ in the load column only by $(\beta_3 - \beta_2)/(C_1 + C_2)$; and each has the rise $T_r - T_c$ in the column of its own fraction and zero in the other's.

By the delta method of equation (4.53), applied to the two quantities at once, the vector of inlets is linear in the state near $\hat{\mathbf{z}}$, $(T_2, T_3)^{\mathsf{T}} \approx (T_2(\hat{\mathbf{z}}), T_3(\hat{\mathbf{z}}))^{\mathsf{T}} + \mathbf{D}(\mathbf{z} - \hat{\mathbf{z}})$, and a linear map of a normal vector is normal, with the covariance carried through the map:

$$\begin{bmatrix} T_2 \\ T_3 \end{bmatrix} \Big|\ \mathbf{y} \approx \mathcal{N}\left(\begin{bmatrix} T_2(\hat{\mathbf{z}}) \\ T_3(\hat{\mathbf{z}}) \end{bmatrix},\ \mathbf{D}\boldsymbol{\Sigma}\mathbf{D}^{\mathsf{T}} \right), \qquad \mathbf{D}\boldsymbol{\Sigma}\mathbf{D}^{\mathsf{T}} = \begin{bmatrix} \mathbf{g}_{T_2}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_{T_2} & \mathbf{g}_{T_2}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_{T_3} \\ \mathbf{g}_{T_3}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_{T_2} & \mathbf{g}_{T_3}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_{T_3} \end{bmatrix}. \tag{8.9}$$

The diagonal holds the two variances of Chapter 5's band (5.13), one per rack; the off-diagonal entry is the covariance, and with $s_k = \sqrt{\mathbf{g}_{T_k}^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g}_{T_k}}$ the correlation is $\varrho = \mathbf{g}_{T_2}^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g}_{T_3} / (s_2 s_3)$. Standardizing each inlet as in equation (4.57), with the distances to the limit

$$h_k = \frac{T_{\max} - T_k(\hat{\mathbf{z}})}{s_k}, \qquad k = 2, 3, \tag{8.10}$$

the corner becomes the corner $\{u_2 \le h_2,\ u_3 \le h_3\}$ of two standard normals u_2, u_3 of correlation ϱ, and

$$\Pr(\max(T_2, T_3) > T_{\max} \mid \mathbf{y}) = 1 - \Phi_2(h_2, h_3;\ \varrho), \tag{8.11}$$

with Φ_2 the standard bivariate normal distribution function,

$$\begin{aligned} \Phi_2(a, b;\ \varrho) &= \int_{-\infty}^{a} \int_{-\infty}^{b} \varphi_2(u, v;\ \varrho)\, \mathrm{d}v\, \mathrm{d}u, \\ \varphi_2(u, v;\ \varrho) &= \frac{1}{2\pi\sqrt{1-\varrho^2}} \exp\left(-\frac{u^2 - 2\varrho u v + v^2}{2(1-\varrho^2)} \right), \end{aligned} \tag{8.12}$$

the second line the bivariate normal density. The marginal exceedances are the normal tails of equation (5.16), $p_k = \Pr(T_k > T_{\max} \mid \mathbf{y}) = \Phi(-h_k)$.

Without any model, the union of two events lies between the larger of them and their sum,

$$\max(p_2, p_3) \ \le\ \Pr(\max(T_2, T_3) > T_{\max} \mid \mathbf{y}) \ \le\ p_2 + p_3, \tag{8.13}$$

the right-hand side the union bound. Inside the bracket the correlation places the answer. A standard identity of the bivariate normal gives the derivative of Φ_2 with respect to the correlation as the density at the corner, $\partial \Phi_2(a, b;\ \varrho) / \partial \varrho = \varphi_2(a, b;\ \varrho)$; integrating from $\varrho = 0$, where the two standard normals are independent and $\Phi_2 = \Phi(a)\Phi(b)$,

$$\Phi_2(h_2, h_3;\ \varrho) = \Phi(h_2)\, \Phi(h_3) + \delta(\varrho), \qquad \delta(\varrho) = \int_0^{\varrho} \varphi_2(h_2, h_3;\ r)\, \mathrm{d}r. \tag{8.14}$$

With $\Phi(h_k) = 1 - p_k$, the union of equation (8.11) becomes

$$\Pr(\max(T_2, T_3) > T_{\max} \mid \mathbf{y}) = \underbrace{1 - (1 - p_2)(1 - p_3)}_{\text{independence}} - \delta(\varrho). \tag{8.15}$$

For a positive correlation $\delta(\varrho)$ is positive, because the density is: inlets that move together tend to be over together, so the union is smaller than under independence. Its size is bounded by the smaller marginal, by equation (8.13); when one rack sits much nearer the limit than the other, the bracket is narrow, and the union is the nearer rack's own tail whatever the correlation.

Building blocks. The MAP state $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$ (Section 8.2.1); the inlets $T_2(\hat{\mathbf{z}})$, $T_3(\hat{\mathbf{z}})$ by equation (5.4); the matrix $\mathbf{D}$ of equation (8.8); the 2×2 covariance $\mathbf{D}\boldsymbol{\Sigma}\mathbf{D}^{\mathsf{T}}$; the distances h_k; and the bivariate normal distribution function Φ_2 of equation (8.12), evaluated by numerical integration of its density.

The union with every reading. At $\hat{\mathbf{z}}$, equation (8.8) is, with the columns $Q, g_1, g_2, \beta_1, \beta_2, \beta_3$,

$$\mathbf{D} = \begin{bmatrix} 0.0467 & -8.2679 & -7.7819 & 0 & 10.1314 & 0 \\ 0.0542 & -8.2679 & -7.7819 & 0 & 0 & 10.1314 \end{bmatrix},$$

the first column in K/kW and the others in K. By equation (8.9) the joint is

$$\begin{bmatrix} T_2 \\ T_3 \end{bmatrix} \Big|\ \mathbf{y} \approx \mathcal{N}\left(\begin{bmatrix} 23.4644 \\ 24.9841 \end{bmatrix},\ \begin{bmatrix} 0.68550 & 0.02866 \\ 0.02866 & 0.68611 \end{bmatrix} \right), \qquad s_2 = 0.8280, \quad s_3 = 0.8283, \tag{8.16}$$

in °C and K^2, with s_k in K, the diagonal the two bands of Chapter 5. The marginal exceedances are $p_2 = 9.7596 \times 10^{-6}$ and $p_3 = 7.4718 \times 10^{-3}$: rack 3 is Chapter 5's 0.75 %, and rack 2, cooler with the same spread, is smaller by nearly three orders of magnitude. By equation (8.11), with the distances $h_2 = 4.2703$ and $h_3 = 2.4337$ and the correlation of the covariance in equation (8.16),

$$\Pr(\max(T_2, T_3) > 27°\mathrm{C} \mid \mathbf{y}) = 1 - \Phi_2(4.2703,\ 2.4337;\ 0.0418) = 0.007481,$$

against 0.007481 under independence, 0.007482 for the union bound and 0.007472 for the larger marginal: the whole bracket (8.13) is narrower than the fourth decimal. The bracket is narrow because rack 2 is far inside the limit: its exceedance, the width of the bracket, is too small to matter, and the union is rack 3's own tail whatever the correlation. The draws give 0.00755 (se 0.00019), a gap of +0.00007, 0.35 of a standard error, which they do not resolve.

8.3.2 How the two inlets move together

Operator question. How much do the two racks move together? The question asks whether a warm rack 2 tells anything about rack 3, and so whether the two unwatched inlets can be treated one at a time.

General form. A moment shared by two derived quantities under the posterior, the two-quantity extension of equation (2.10), the covariance

$$\mathrm{Cov}(f_a(\mathbf{z}), f_b(\mathbf{z}) \mid \mathbf{y}) = \int \big(f_a(\mathbf{z}) - \mathbb{E}[f_a]\big)\big(f_b(\mathbf{z}) - \mathbb{E}[f_b]\big)\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \tag{8.17}$$

reported with the correlation, the covariance over the product of the two standard deviations.

Specialized form. The question is equation (8.17) with $f_a = T_2(\mathbf{z})$, $f_b = T_3(\mathbf{z})$ and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal, which by the linearization of the joint (8.9) makes the covariance its off-diagonal entry $\mathbf{g}_{T_2}^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g}_{T_3}$. One choice enters here: the survey's priors on the three fractions are independent, so that by the pattern (8.2) the fractions of racks 2 and 3 are tied to nothing but their own prior rows. This is what keeps the fractions out of the correlation, and Section 8.6.5 says what breaks it.

Write $\mathbf{D}_\mathrm{s}$ for the 2×3 block of $\mathbf{D}$ in the shared columns (Q, g_1, g_2) and $\boldsymbol{\Sigma}_{\mathrm{ss}}$ for the matching 3×3 block of $\boldsymbol{\Sigma}$. By the block structure (5.10), $\boldsymbol{\Sigma}$ has no covariance between β_2 or β_3 and any other entry, and their variances are the survey's 0.08^2; the column of β_1 is zero in both rows. The covariance of the inlets therefore splits into a shared part and a survey part,

$$\mathbf{D}\boldsymbol{\Sigma}\mathbf{D}^{\mathsf{T}} = \mathbf{D}_\mathrm{s}\boldsymbol{\Sigma}_{\mathrm{ss}}\mathbf{D}_\mathrm{s}^{\mathsf{T}} + (T_r - T_c)^2\, 0.08^2\, \mathbf{I}_2, \tag{8.18}$$

an identity on the fitted state, not a further approximation. The off-diagonal entry, the covariance, is the shared part's alone. Because the two rows of $\mathbf{D}_\mathrm{s}$ are nearly the same vector and $\boldsymbol{\Sigma}_{\mathrm{ss}}$ is positive definite, the covariance is positive and close to the shared variance v_s of either rack, so that

$$\varrho \approx \frac{v_\mathrm{s}}{v_\mathrm{s} + (T_r - T_c)^2\, 0.08^2}. \tag{8.19}$$

The correlation is the share of each inlet's variance that comes from the shared entries. It is small when the readings fix the cold aisle and the rise and the survey dominates each band; it is one half when the shared variance equals the survey's, and it tends to one as the shared variance grows beyond it.

Building blocks. $\boldsymbol{\Sigma}$ with its shared block $\boldsymbol{\Sigma}_{\mathrm{ss}}$ (Section 5.2.3), and the matrix $\mathbf{D}$ at $\hat{\mathbf{z}}$ (Section 8.3.1).

The correlation with every reading. The covariance in equation (8.16) is $0.02866\,\mathrm{K}^2$, the shared part of the split (8.18); over 0.8280×0.8283 it is the correlation already used in Section 8.3.1. The covariance is close to the variance $0.1708^2 = 0.0292\,\mathrm{K}^2$ that rack 3's inlet keeps with its fraction known (Section 5.3.1), as equation (8.19) requires, and it is a small part of each variance, so the two inlets are nearly uncorrelated. The draws give a covariance of $0.02971\,\mathrm{K}^2$ and a correlation of 0.0433 (se 0.0022), a gap of +0.0015.

8.3.3 Both inlets over the limit

Operator question. Could both be over at once? The question asks for the probability that racks 2 and 3 are both over 27°C at 23:00, the corner opposite the operating envelope.

General form. The probability of a region, as in Section 8.3.1, with the region an intersection of exceedances:

$$\Pr\Big(\bigcap_k \{f_k(\mathbf{z}) > f^*\} \mid \mathbf{y}\Big) = \int \mathbf{1}\big[\min_k f_k(\mathbf{z}) > f^*\big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}. \tag{8.20}$$

Specialized form. The question is equation (8.20) with the same f_2, f_3, f^* and Laplace normal as in Section 8.3.1. No new choice enters: the region is another corner of the same bivariate normal (8.9). By inclusion and exclusion, which holds for any joint distribution,

$$\begin{aligned} \Pr(T_2 > T_{\max},\ T_3 > T_{\max} \mid \mathbf{y}) &= p_2 + p_3 - \Pr(\max(T_2, T_3) > T_{\max} \mid \mathbf{y}) \\ &= 1 - \Phi(h_2) - \Phi(h_3) + \Phi_2(h_2, h_3;\ \varrho), \end{aligned} \tag{8.21}$$

and by the identity (8.14),

$$\Pr(T_2 > T_{\max},\ T_3 > T_{\max} \mid \mathbf{y}) = p_2\, p_3 + \delta(\varrho). \tag{8.22}$$

The same $\delta(\varrho)$ that equation (8.15) subtracts from the union under independence is added here to the product of the marginals. Against the union, which is at least the larger marginal, $\delta(\varrho)$ is a small correction; against the product $p_2 p_3$, which is the product of two small numbers, it can be many times the whole. A correlation the union barely registers can multiply the intersection, and the product of marginals is then the wrong answer to "both", as Chapter 2 found for two halls on one chiller (Section 2.4.6).

Building blocks. p_2, p_3, ϱ and Φ_2 of Section 8.3.1.

The intersection with every reading. By equation (8.21),

$$\Pr(T_2 > 27°\mathrm{C},\ T_3 > 27°\mathrm{C} \mid \mathbf{y}) = 1.1985 \times 10^{-7},$$

against 7.2922×10^{-8} for the product of the marginals, 1.64 times it. The correlation of 0.0418 that leaves the union unchanged in its fourth decimal already moves the intersection by more than half, as equation (8.22) says it can; both numbers are below one in a million. The draws cannot be set beside it: 200 000 draws are expected to put 0.024 draws in the corner, and none fell there.

8.3.4 The shared entries poorly read

Operator question. When do the marginals one by one mislead? The question asks under what evidence the inlets read one at a time give a wrong answer to the joint questions, and it is answered by asking the first three again of the second reading set.

General form. The probability of a region A of the state and the covariance, under the posterior of the second reading set,

$$\Pr(\mathbf{z} \in A \mid \mathbf{y}_{\mathrm{II}}) = \int \mathbf{1}_A(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y}_{\mathrm{II}})\, \mathrm{d}\mathbf{z}, \tag{8.23}$$

the forms (8.6) and (8.20) and the moment (8.17) with $\mathbf{y}_{\mathrm{II}}$ on the right of the bar.

Specialized form. The question is equation (8.23) with A the union and then the intersection of the two exceedances at $T_{\max}$, with $f_a = T_2$ and $f_b = T_3$ in the moment, and with $p(\mathbf{z} \mid \mathbf{y}_{\mathrm{II}})$ the Laplace normal (8.4). No new choice enters, but the delta method of Section 8.3.1 is now used on a posterior whose spread is wider, of the log-derates 0.117 against the priors' 0.15, so the linearity it needs is asked of a wider region.

The derivation is unchanged: equations (8.9) to (8.22) with $\hat{\mathbf{z}}_{\mathrm{II}}$, $\mathbf{\Sigma}_{\mathrm{II}}$ and $\mathbf{D}_{\mathrm{II}}$, the matrix (8.8) at $\hat{\mathbf{z}}_{\mathrm{II}}$. By the pattern (8.2) the fractions of racks 2 and 3 still hold their prior rows alone, so the split (8.18) still holds with the survey's 0.08^2; what changes is the shared block, which the coil readings no longer narrow. Each inlet's band widens with it, so each marginal grows; this, not the correlation, is what moves the union. By equation (8.19) the correlation also grows, and by equations (8.15) and (8.22) it moves the union below independence and the intersection above it by the same $\delta(\varrho)$.

Building blocks. $\hat{\mathbf{z}}_{\mathrm{II}}$ and $\mathbf{\Sigma}_{\mathrm{II}}$ (Section 8.2.2); $\mathbf{D}_{\mathrm{II}}$ by equation (8.8); the covariance $\mathbf{D}_{\mathrm{II}}\mathbf{\Sigma}_{\mathrm{II}}\mathbf{D}_{\mathrm{II}}^{\mathsf{T}}$; Φ_2 of equation (8.12).

The joint without the coil readings. The gradient matrix is

$$\mathbf{D}_{\mathrm{II}} = \begin{bmatrix} 0.0462 & -7.9149 & -7.9149 & 0 & 10.0618 & 0 \\ 0.0537 & -7.9149 & -7.9149 & 0 & 0 & 10.0618 \end{bmatrix},$$

with equal log-derate entries, since the two coils are read alike, and the joint is

$$\begin{bmatrix} T_2 \\ T_3 \end{bmatrix} \,\Big|\, \mathbf{y}_{\mathrm{II}} \approx \mathcal{N}\!\left(\begin{bmatrix} 23.2895 \\ 24.7988 \end{bmatrix},\ \begin{bmatrix} 1.23458 & 0.58759 \\ 0.58759 & 1.23698 \end{bmatrix} \right), \qquad \varrho = 0.4755.$$

Table 8.4 sets it beside the joint with every reading: each row a quantity, one column per reading set. Both standard deviations grow, to 1.1111 and 1.1122 K; the survey's part of each variance is the new rise times 0.08, squared; and the covariance, the shared part, grows from 0.02866 to 0.58759 K^2, near half of each variance, so that the correlation is near one half, as equation (8.19) predicts. The draws give a correlation of 0.4796 (se 0.0008), a gap of +0.0041.

Table 8.4: The joint posterior of the two unmeasured inlets under the two reading sets: the means, the gradient matrix $\mathbf{D}$, the 2×2 covariance $\mathbf{D\Sigma D}^{\mathsf{T}}$ and the correlation, equations (8.9) and (8.8). Rows of $\mathbf{D}$ in the columns $(Q, g_1, g_2, \beta_1, \beta_2, \beta_3)$, K/kW for Q and K for the others.

Quantity [unit]	Every reading	Meter and rack 1 only
$T_2(\hat{\mathbf{z}})$, $T_3(\hat{\mathbf{z}})$ [°C]	23.4644, 24.9841	23.2895, 24.7988
s_2, s_3 [K]	0.8280, 0.8283	1.1111, 1.1122
$\mathbf{g}_{T_2}^{\mathsf{T}}$ [K/kW, K]	$(0.0467, -8.2679, -7.7819, 0, 10.1314, 0)$	$(0.0462, -7.9149, -7.9149, 0, 10.0618, 0)$
$\mathbf{g}_{T_3}^{\mathsf{T}}$ [K/kW, K]	$(0.0542, -8.2679, -7.7819, 0, 0, 10.1314)$	$(0.0537, -7.9149, -7.9149, 0, 0, 10.0618)$
$\mathbf{D\Sigma D}^{\mathsf{T}}$ [K^2]	$\begin{bmatrix} 0.68550 & 0.02866 \\ 0.02866 & 0.68611 \end{bmatrix}$	$\begin{bmatrix} 1.23458 & 0.58759 \\ 0.58759 & 1.23698 \end{bmatrix}$
Correlation ϱ [–]	0.0418	0.4755

Figure 8.2 draws the two joints: in each panel the horizontal axis is the inlet of rack 2 and the vertical axis the inlet of rack 3; the ellipses are curves of constant joint density, the shaded L-shaped region is the complement of the operating envelope's corner at 27°C, the dot is the pair of virtual readings and the open diamond the stated truth of Section 8.4.3. In the left panel, with every reading, the ellipses are nearly circles and only the outermost reaches the shaded band, and only above the corner, across the line $T_3 = 27$°C; the band to the right, rack 2 over, lies beyond it. In the right panel the ellipses are larger and tilted along the diagonal, and the outer ones reach into the shaded band above the corner and lean toward the top right square.

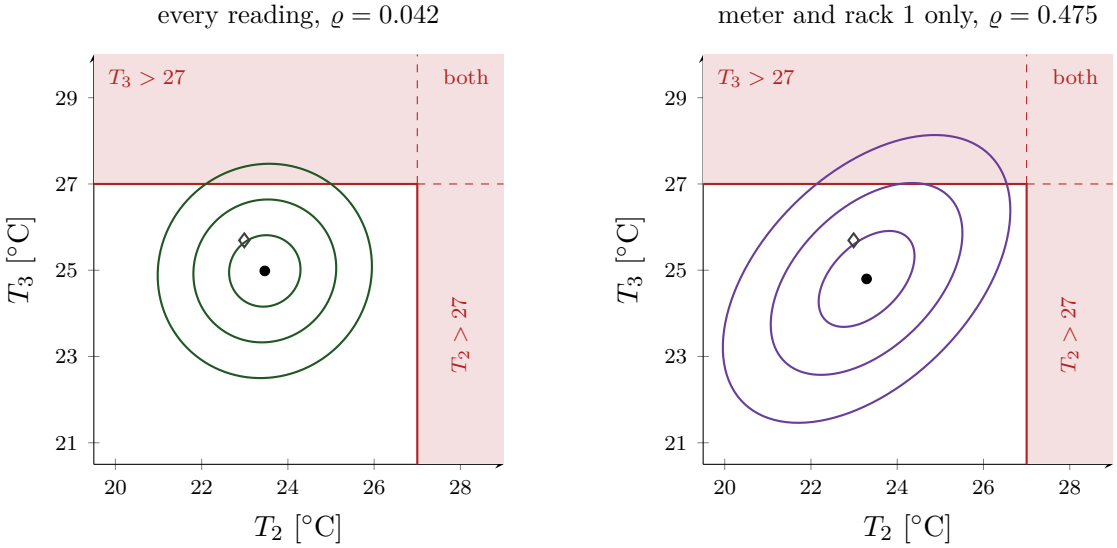


Figure 8.2: The joint posterior of the two unmeasured inlets, the bivariate normal of equation (8.9), under the two reading sets of Table 8.4. Ellipses: constant joint density at standardized (Mahalanobis) distance 1, 2 and 3 from the mean, drawn as $T_2 = m_2 + d\,s_2 \cos t$, $T_3 = m_3 + d\,s_3(\varrho \cos t + \sqrt{1-\varrho^2}\sin t)$ for $d = 1, 2, 3$. Shaded: the complement of the operating envelope's corner, $\max(T_2, T_3) > 27$°C; the top right square is "both". Dot: the virtual readings $(T_2(\hat{\mathbf{z}}), T_3(\hat{\mathbf{z}}))$. Open diamond: the stated truth of the constructed record, (22.99, 25.69)°C, a check marker (Section 8.4.3).

The union and the intersection without the coil readings. The distances are $h_2 = 3.3394$ and $h_3 = 1.9791$, the marginals $p_2 = 4.1981 \times 10^{-4}$ and $p_3 = 2.3901 \times 10^{-2}$, and by equations (8.11) and (8.21)

$$\Pr(\max(T_2, T_3) > 27°\text{C} \mid \mathbf{y}_{\text{II}}) = 0.024161, \qquad \Pr(T_2 > 27°\text{C},\ T_3 > 27°\text{C} \mid \mathbf{y}_{\text{II}}) = 1.6012 \times 10^{-4}.$$

The union is 0.024311 under independence, 0.024321 for the union bound and 0.023901 for the larger marginal. It has grown by a factor of 3.23 from 0.007481 because rack 3's own band has widened and its marginal with it, from 7.4718×10^{-3} to 2.3901×10^{-2}, even though its mean falls 0.185 K: at the old mean the wider band would give 0.0350; rack 2 is still far enough inside the limit that the bracket stays narrow, and the correlation only pulls the union below independence, by $\delta = 1.501 \times 10^{-4}$, as equation (8.15) requires. The intersection is 1.0034×10^{-5} under independence, and the same δ puts it at 15.96 times that.

Beside these stand 1 000 000 draws: the union 0.02285 (se 0.00015), a gap of −0.00131, 8.77 standard errors, 5.4 % of the reported value; and the intersection 2.160×10^{-4} (se 1.5×10^{-5}), a gap of $+5.59 \times 10^{-5}$, the joint normal 3.80 standard errors low, 21.53 times independence against its 15.96. The draws show the linearization errs here. The draws' union lies below the joint normal's larger marginal 0.023901, which is no breach of the bracket (8.13): the bracket

holds within each route, and the draws' own marginals are 0.02256 (se 0.00015) and 5.07×10^{-4}, so their bracket is 0.02256 to 0.02307, and their union lies inside it. No decision level is set at 23:00, so both values stand: by either route the union has about tripled, 3.03 times by the draws, and "both" is many times the product of the marginals.

Table 8.5 collects the probabilities; each row is a probability, the columns the two reading sets, with the draws beside the joint normal. Figure 8.3 sets them on one logarithmic axis: each row is a probability, the horizontal axis its value from 10^{-8} to 10^{-1}, the filled circles the fit with every reading, the open squares the fit with the meter and rack 1 only, and the open triangles that fit's draws. The estimates of the union sit on top of one another in each set; the estimates of the intersection are far apart in the second.

Table 8.5: The probabilities of the joint questions under the two reading sets: the joint normal of the delta method, the reported value, equations (8.11), (8.13) and (8.21), and beside it the draws through the physics, 200 000 with every reading and 1 000 000 with the meter and rack 1 only.

Probability [–]	Every reading	Meter and rack 1 only
$p_2 = \Pr(T_2 > 27)$	9.7596×10^{-6}	4.1981×10^{-4}
$p_3 = \Pr(T_3 > 27)$	7.4718×10^{-3}	2.3901×10^{-2}
$\Pr(\max(T_2, T_3) > 27)$, joint normal	0.007481	0.024161
under independence	0.007481	0.024311
union bound $p_2 + p_3$	0.007482	0.024321
larger marginal	0.007472	0.023901
draws through the physics	0.00755	0.02285
the draws' own bracket	0.00754 to 0.00755	0.02256 to 0.02307
$\Pr(T_2 > 27,\ T_3 > 27)$, joint normal	1.1985×10^{-7}	1.6012×10^{-4}
under independence $p_2 p_3$	7.2922×10^{-8}	1.0034×10^{-5}
draws through the physics	unresolvable	2.160×10^{-4}

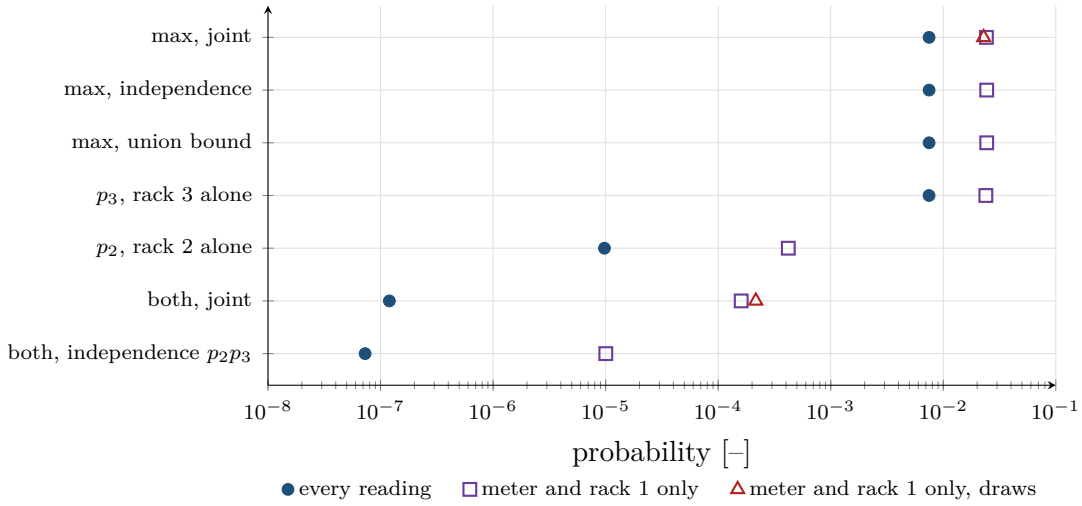


Figure 8.3: The probabilities of Table 8.5 on a logarithmic axis, for the fit with every reading (filled circles) and with the meter and rack 1 only (open squares), from the joint normal of the delta method, the reported values, and for the second fit by 1 000 000 draws through the physics (open triangles), beside them.

8.4 The checks (step 5)

The answers of Section 8.3 rest on the delta method and on two fits, and each is tested here against the calculation it stands for.

8.4.1 The limit of exact rows and the curvature term

The fit with every reading is Chapter 5's, and so are its checks (Table 5.5 and Section 5.4.1); with every reading the Newton covariance of the two inlets gives the same standard deviations, 0.8280 and 0.8283 K, and the same correlation, 0.0418.

The refit is checked on its own full state of thirteen entries, equation (5.2), with the seven physical rows at a finite tolerance beside the two readings and five priors of the second set. Table 8.6 gives the result: each row is a tolerance, the columns the inlet of rack 3 and the load with their spreads, the last row the reduced state. The full state reproduces the reduced state at $\sigma_p = 0.01$ except for the last digit of the spread, and to the fourth decimal at 0.001.

Table 8.6: The limit of exact physical rows on the full state of the second reading set: the inlet of rack 3 and the load at three tolerances, against the reduced state. Each fit took six iterations.

Tolerance σ_p [kW or K]	T_3 [°C]	sd(T_3) [K]	Q [kW]	sd(Q) [kW]
0.1	24.7944	1.1195	193.2363	2.9938
0.01	24.7988	1.1123	193.2370	2.9938
0.001	24.7988	1.1122	193.2370	2.9938
reduced state	24.7988	1.1122	193.2370	2.9938

The curvature term depends on the readings, and it is checked on the second fit directly: the Newton covariance gives the inlets standard deviations of 1.1164 and 1.1175 K and a correlation of 0.4804, against 1.1111, 1.1122 K and 0.4755 for Gauss–Newton, and its joint normal puts the union at 0.02470, a gap of +0.00054 from the reported value. The draws come from the Gauss–Newton normal, so the curvature term plays no part in their gap; it is a separate gap of its own.

8.4.2 The joint normal against draws through the physics

The delta method replaces the two inlets by their linearization about the MAP state; the draws of Section 8.3 keep the exact physics. Table 8.7 collects the gaps: each row is a route, the columns the probability of the union and of the intersection, each with its Monte Carlo standard error in parentheses and its gap from the reported joint normal, the route minus the joint normal; the last block holds the moments and the truth.

With every reading the routes agree within what the draws resolve, and the same draw puts rack 3 alone at 0.00754 (se 0.00019), its union differing by the one draw in which only rack 2 is over.

With the meter and rack 1 only, the gaps are resolved and stated in Section 8.3.4; the draw of 200 000 already shows the union's, 0.02246 (se 0.00033), 5.13 standard errors low, and it puts the intersection at 1.900×10^{-4} (se 3.1×10^{-5}), a gap of $+2.99 \times 10^{-5}$, 0.97 of a standard error, which it does not resolve. Figure 8.4 draws the union by every route: in each panel the horizontal axis is the probability of the union, each dot with its bar a draws' estimate with two standard errors on each side, the solid red line the Gauss–Newton joint normal and, in the right panel, the dashed red line the Newton joint normal. In the left panel the line passes through the bar; in the right panel both lines lie outside both bars, the Newton line farther.

The draws' means show what the linearization drops. With the meter and rack 1 only they lie lower than the MAP values by 0.069 K (se 0.0011), the same for both racks by construction:

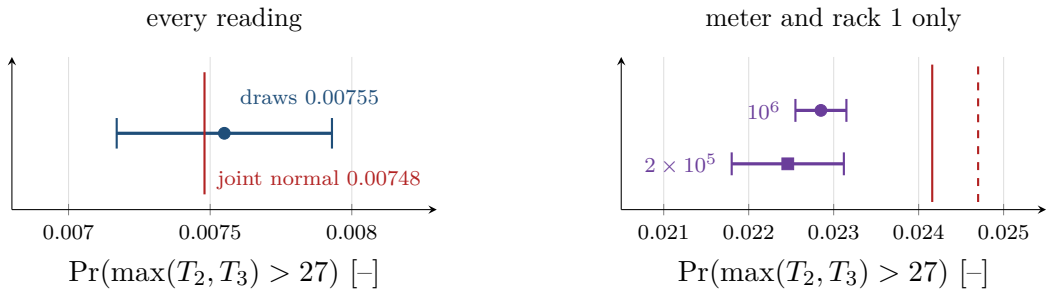


Figure 8.4: The probability that the hotter unmeasured inlet exceeds 27°C by draws from the Laplace posterior through the exact physics (markers, with a bar of two Monte Carlo standard errors on each side) and by the joint normal of the delta method (solid red line, Gauss–Newton, the reported value). Left: every reading, 200 000 draws. Right: the meter and rack 1 only, 200 000 draws (lower bar) and 1 000 000 draws (upper bar), with the Newton joint normal dashed.

$T_3 - T_2 = (\beta_3 - \beta_2)(Q + Q_o)/(C_1 + C_2)$ has independent factors, so the difference keeps its mean. With every reading they lie 0.0015 and 0.0053 K lower (se 0.0019). The shift is the second-order term $\frac{1}{2}\,\mathrm{tr}(\nabla^2 T_k\,\boldsymbol{\Sigma})$ of each inlet, which is −0.0717 K on the second fit, and it comes from the return air. Each inlet is $T_k = T_r - (1 - \beta_k)(Q + Q_o)/(C_1 + C_2)$, whose second term is linear in the load and meets β_k only in a product whose entries $\boldsymbol{\Sigma}$ does not correlate; the first is $T_r = T_s + (Q + Q_o)/(G_1 + G_2)$ with $G_u = \bar{G}_u \mathrm{e}^{g_u}$. With equal ratings and $g_1 = g_2$ at the MAP state of the second fit, the curvature of $1/(G_1 + G_2)$ in each log-derate alone vanishes, $\partial^2/\partial g_1^2 \propto \mathrm{e}^{g_1}(\mathrm{e}^{g_1} - \mathrm{e}^{g_2}) = 0$, and the cross-curvature is positive, $\partial^2 T_k/\partial g_1 \partial g_2 = 7.915$ K. Acting on the negative covariance of the two log-derates, −0.00888 (correlation −0.6519), it gives −0.0703 K of the term, and the cross terms with the load are small, its correlation with the log-derates being 0.0519. With every reading the same term acts on a covariance of −0.00025 and the second-order term is −0.0037 K. The draws' joint is therefore not the linearized normal: its center lies lower, which lowers the union, and its upper corner holds more weight than the normal's, which raises the intersection. The check tests the linearization, not the Laplace posterior: both routes draw from the same normal of the state.

8.4.3 The answers against the stated truth

The record of this hall is generated from a stated truth (Section 8.6.5), and the last check compares the answers with it. The true inlets are $T_2 = 22.9893$°C and $T_3 = 25.6933$°C, the last row of Table 8.7 and the open diamond of Figure 8.2. Neither is over 27°C: the true state lies inside the operating envelope's corner, inside the second ellipse of either panel. Rack 3's true inlet is warmer than its virtual reading because the survey's 0.50 is below the true fraction, the error Chapter 5 traced (Section 5.4.4); rack 2's is cooler than its virtual reading for the opposite reason. That the event did not occur does not test a probability of 0.75 %; one window cannot, and the check is that the truth lies where the joint puts most of its weight.

8.5 The report (step 6)

The worry about the rack no one sees adds nothing to what rack 3 already said. With every reading in, the probability that some unmeasured inlet is over 27°C is 0.75 %, and all of it is rack 3's, because rack 2 sits too far inside the limit to add to it. When the coil readings drop out the same probability rises to 2.4 %, and again because of rack 3: its own band widens, while the correlation of the two inlets lowers the union slightly. The joint matters for one question only, "both over at once", which is 1.64 times the product of the marginals with every reading and

Table 8.7: The checks of the joint questions and their gaps from the reported joint normal, the route minus the joint normal. The draws are states from the Laplace posterior of each reading set through the exact physics (5.4), with their Monte Carlo standard errors in parentheses; the Newton joint normal keeps the curvature term; the truth is the stated truth of the constructed record.

Route	Pr(max > 27) [–]	Gap [–]	Pr(both > 27) [–]	Gap [–]
Every reading				
Joint normal, reported	0.007481	–	1.1985×10^{-7}	–
200 000 draws	0.00755 (0.00019)	+0.00007	unresolvable	–
Meter and rack 1 only				
Joint normal, reported	0.024161	–	1.6012×10^{-4}	–
Newton joint normal	0.02470	+0.00054	–	–
200 000 draws	0.02246 (0.00033)	−0.00170	1.900×10^{-4} (3.1×10^{-5})	$+2.99 \times 10^{-5}$
1 000 000 draws	0.02285 (0.00015)	−0.00131	2.160×10^{-4} (1.5×10^{-5})	$+5.59 \times 10^{-5}$
The joint's moments, draws against the joint normal				
Correlation, every reading, 200 000	0.0433 (0.0022)	+0.0015	joint normal 0.0418	
Correlation, meter and rack 1 only, 10^6	0.4796 (0.0008)	+0.0041	joint normal 0.4755	
Means, meter and rack 1 only, 10^6 [K]	−0.069 (0.0011)		second-order term −0.0717	
Stated truth: $T_2 = 22.9893$°C, $T_3 = 25.6933$°C	neither over	–	neither over	–

15.96 times it with the coil readings out. There the linearization errs, and the draws through the physics say by how much: with the coil readings out they put the union at 2.3 % and "both" at 21.53 times the product. No decision level is set, so both values stand.

8.5.1 The answers and their bands

Table 8.8 gives the answers in the order of Table 8.1. Each row is a question; the second column gives the answer from the joint normal with the draws beside it, and the third what qualifies it: the approximations of the bracket (8.13) for the union and the independence value for the intersection.

8.5.2 The entry in the shift log

The operator's entry for 23:00 answers the worry about the racks without sensors in one line: the chance that any unmeasured inlet is over 27°C is 0.75 %, the same as rack 3's alone, and rack 2, read at 23.5°C, adds nothing to it. Rack 3 is the rack to watch; a second unwatched rack is not a second risk of the same size, because it sits well inside the limit. The entry also records what the coil readings do. If the return-air or discharge sensors go out of service, rack 3's own band widens, more than its lower mean offsets, and the chance of an exceedance somewhere rises from 0.75 % to 2.4 %, or 2.3 % by the draws; the two inlets also begin to move together, so that "both racks over" becomes 16 times the product of two small numbers, or 22 times by the draws. The coil readings narrow every rack's band and keep the racks' estimates nearly independent of each other, and their loss is worth a line in the log of its own.

Table 8.8: The report for the two unmeasured inlets of the two-unit hall at 23:00. Answers from the joint normal of the posterior of the state, draws through the physics beside them, with Monte Carlo standard errors.

Question	Answer	Qualification
Is the hot aisle already above 27°C somewhere I cannot see?	probability 0.75 % (0.007481), carried by rack 3 (draws: 0.00755, se 0.00019)	larger marginal 0.007472, union bound 0.007482
How much do the two racks move together?	correlation 0.0418 (draws: 0.0433, se 0.0022)	0.4755 with the meter and rack 1 only (draws: 0.4796, se 0.0008)
Could both be over at once?	probability 1.2×10^{-7} (draws: too rare to resolve)	1.64 times the product of the marginals
When do the marginals mislead?	for "both" only. With the coil readings out, union 2.4 % (0.024161; draws: 0.02285, se 0.00015) and both 1.6012×10^{-4} (draws: 2.160×10^{-4}, se 1.5×10^{-5})	larger marginal 0.023901 (draws' own: 0.02256); both 15.96 times the product of the marginals (draws: 21.53 times)

Common mistake

Multiplying the marginals to get "both racks over", or adding them to get "some rack over", without asking how the racks move together. The sum is an upper bound and is close to the answer only when one rack carries the union; the product is the answer only when the racks are independent, and two inlets that share a cold aisle and a rise are not. A second mistake is to read the hot aisle against the inlet limit: the return air behind the racks is at 30.33°C by design, and the 27°C limit applies to the air the racks draw in.

8.6 The next reading (step 7)

The report leaves a 0.75 % probability of an exceedance somewhere, nearly all of it in one inlet. A further reading is one more row of the residual vector, and what it would do to the union is known before its value is.

8.6.1 The preposterior of a candidate reading

A candidate reading a is a row with Jacobian line $\mathbf{J}_a$ and the preposterior covariance of equation (5.19),

$$\boldsymbol{\Sigma}_a = \boldsymbol{\Sigma} - \frac{\boldsymbol{\Sigma}\mathbf{J}_a^{\mathsf{T}}\,\mathbf{J}_a\boldsymbol{\Sigma}}{1 + \mathbf{J}_a\boldsymbol{\Sigma}\mathbf{J}_a^{\mathsf{T}}}, \qquad \mathbf{D}\boldsymbol{\Sigma}_a\mathbf{D}^{\mathsf{T}}, \tag{8.24}$$

applied once per row for a candidate of several rows. The joint after the reading has the covariance $\mathbf{D}\boldsymbol{\Sigma}_a\mathbf{D}^{\mathsf{T}}$, and the union is equation (8.11) with it. One choice enters here: the union after the reading is evaluated at the present means. This is the union the hall would show if the reading confirmed the present estimate; the reading's value will move the means, and the union with them, so the number ranks the candidates by what they do to the spread and is not a forecast of the union tomorrow. A probe at rack 3 of standard deviation 0.3 K reads T_3, $\mathbf{J}_a = \mathbf{g}_{T_3}^{\mathsf{T}}/0.3$; a probe at rack 2 reads T_2, $\mathbf{J}_a = \mathbf{g}_{T_2}^{\mathsf{T}}/0.3$; a permanent inlet sensor at rack 3 is the rack-3 probe's row added in every window; and a re-survey of the two unmeasured fractions with standard deviation 0.03 each is two new prior rows, $\mathbf{e}_5^{\mathsf{T}}/0.03$ and $\mathbf{e}_6^{\mathsf{T}}/0.03$, which by the block structure narrow β_2 and β_3 and nothing else.

8.6.2 The candidates

Table 8.9 gives the candidates. Each row is a reading the operator can request; the columns give the standard deviations of the two inlets after it, by equation (8.24), and the union at the present means. Figure 8.5 draws the two standard deviations as paired bars, one pair per candidate, the horizontal axis in kelvin and the present fit at the top.

Table 8.9: Candidate readings for the joint question, from the posterior with every reading at 23:00; the union is evaluated at the present means.

Candidate	$\mathrm{sd}_a(T_2)$ [K]	$\mathrm{sd}_a(T_3)$ [K]	$\Pr_a(\max > 27)$ [–]
Present fit, no further reading	0.8280	0.8283	7.48×10^{-3}
A probe at rack 3, sd 0.3 K	0.8273	0.2821	9.62×10^{-6}
A permanent inlet sensor at rack 3, every window	0.8273	0.2821	9.62×10^{-6}
A probe at rack 2, sd 0.3 K	0.2821	0.8277	7.43×10^{-3}
A re-survey of β_2 and β_3, sd 0.03 each	0.3310	0.3319	6.26×10^{-10}

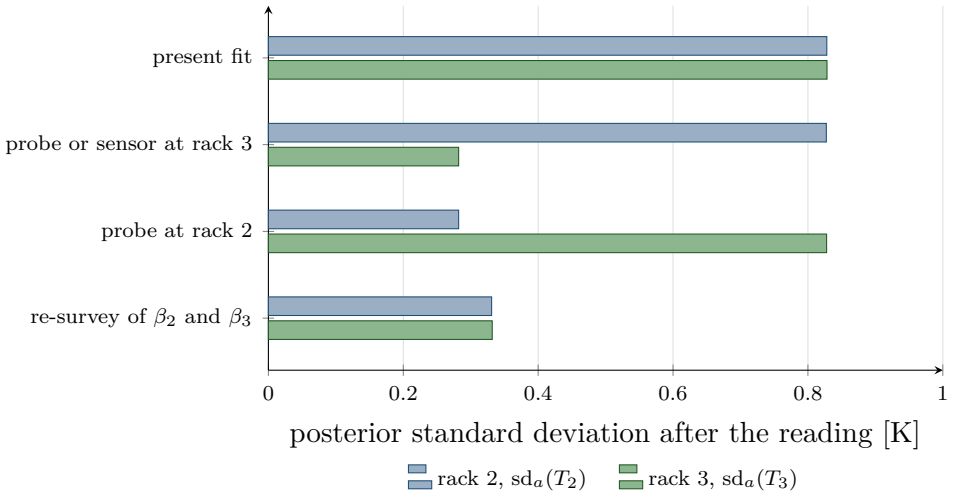


Figure 8.5: The standard deviations of the two unmeasured inlets after each candidate of Table 8.9, by the preposterior (8.24). A probe or a permanent sensor at rack 3 narrows rack 3 and leaves rack 2; a probe at rack 2 does the reverse; the re-survey narrows both.

Read on the entries, each candidate acts on one fraction or on both. The probe at rack 3 is, as in Chapter 5, a one-point re-survey of β_3: it takes rack 3 from 0.8283 to 0.2821 K and leaves rack 2 at 0.8273 K, and the union falls from 0.007481 to 9.62×10^{-6}, close to rack 2's own marginal: with rack 3 read, the union is carried by the rack that is left. The probe at rack 2 acts mostly on β_2 and slightly on the shared entries, which its row also reaches; rack 2 carried nothing of the union, so the union barely moves, to 7.43×10^{-3}, the small fall coming from the shared entries, which narrow rack 3 from 0.8283 to 0.8277 K. The re-survey narrows both inlets, to 0.3310 and 0.3319 K, and takes the union to 6.26×10^{-10}.

8.6.3 The recommendation

The re-survey of both unmeasured fractions lowers the union most, the one candidate that narrows both racks. The next reading is nevertheless the one Chapter 5 recommended, on that chapter's reasons: a probe at rack 3 is one visit tonight, and it already takes the probability of an exceedance anywhere from 0.75 % to 0.001 % at the present means, because the rack nearest

the limit carries the union. A probe at rack 2 is not worth a visit for this question. The request beyond tonight is the permanent inlet sensor at rack 3 where one can be fitted, the same row in every window, with the re-survey as the alternative. A request that the analysis adds is to keep the return-air and discharge sensors in service: they are the rows that hold the shared entries; without them rack 3's band widens, the union rises from 0.75 to 2.4 %, and the inlets become correlated. Chapter 18 ranks the readings an operator can request on the shift plant in the same way.

8.6.4 What the analysis showed

The operator feared that two unwatched racks were two chances of being over the limit unseen; the analysis found one. With every reading in, the probability of an exceedance anywhere is rack 3's own 0.75 %, because rack 2 sits too far inside the limit to contribute, and the inlets taken one by one already give the answer: the bracket between the larger marginal and the sum is narrower than the fourth decimal, whatever the correlation. With the coil readings out of service the union triples, to 2.4 %, and that too is rack 3's doing, its band widened more than its mean fell, not the correlation's. The joint matters for "both racks over", which is the product of the marginals only when the gradients are orthogonal under $\mathbf{\Sigma}$, which nothing here arranges: 1.64 times it with every reading and 15.96 times with the coil readings out, when the two inlets move together. In that case the linearization also errs, high for the union and low for "both", and the draws, 2.3 % and 21.53 times, show how far; the reading of the hall does not change. The hall's coil sensors were placed to watch the units; they also narrow every rack's band and keep the racks' estimates apart, a reason to keep them working that no single rack's reading would show.

8.6.5 Limits of the model

Independent fractions. Racks that share a cause of recirculation, such as missing blanking panels along one row or one hot spot in the return path, would have correlated fractions; the block of β_2 and β_3 in $\mathbf{\Sigma}$ would no longer be diagonal, the survey part of the split (8.18) would add to the covariance, and the joint would matter even with every reading in.

Two racks. The hall has two unmeasured inlets, one of them near the limit. With many unmeasured inlets at similar distances from their limits, no one marginal carries the union, the union bound is loose, and the answer must be read from the joint of all of them; the bivariate normal function becomes a multivariate one, computed by numerical integration or by draws.

The delta method at the corner. The joint normal errs when the shared entries are poorly read, by the gaps of Section 8.3.4. With every reading the intersection is too rare for the draws to resolve, so it stands on the joint normal alone. The Laplace posterior the draws come from is itself least reliable where the posterior is widest, of the log-derates 0.117 against the priors' 0.15; the Newton covariance, which keeps the curvature term, puts the union at +0.00054 against the reported value, a separate gap.

One fraction per rack and the steady state. Each inlet is one mix of two air streams, as in Chapter 5; a stratified inlet would put the top of a rack over the limit before the model's T_k, and the event would then belong to a height, not to a rack. The hall is taken in steady state within the window, and after a load step the rise no longer equals $(Q + Q_o)/(C_1 + C_2)$.

Constructed readings. The five readings are constructed: they are generated from a stated truth of the hall with the stated sensor errors, and Section 8.4.3 compares the answers with that truth. They are not field data, and the agreement is a check of the method on its own model.

The next chapter. Part III follows the same hall through the following night, window by window, and asks first, in Chapter 9, whether a residual is noise or a change.

8.7 What to retain

Definition 8.1 (Joint exceedance over unmeasured quantities)**.** For derived quantities $f_k(\mathbf{z})$ with gradients $\mathbf{g}_k$ stacked as the rows of $\mathbf{D}$, the joint posterior on the fitted state is $\mathcal{N}(f(\hat{\mathbf{z}}),\ \mathbf{D}\boldsymbol{\Sigma}\mathbf{D}^{\mathsf{T}})$ by the delta method. The probability that any of them exceeds its limit is one minus the probability of the operating envelope's corner; for two it is $1-\Phi_2(h_2, h_3;\ \varrho)$, with the standardized distances h_k and the correlation ϱ. It lies between the larger marginal and the union bound. Positive correlation lowers it below independence and raises the intersection above the product of the marginals by the same amount $\delta(\varrho)$. The correlation comes from the columns the gradients share, and it is the shared entries' share of each variance.

Three conclusions carry beyond the hall. An exceedance somewhere is a union, and the marginals bracket it before any joint is computed; when one quantity sits much nearer its limit than the others, its own tail is the answer, whatever the correlation. An intersection is the product of the marginals only when the gradients are orthogonal under $\boldsymbol{\Sigma}$, which nothing here arranges; the shared columns of the gradients set how far it is from that product. And the linearization that serves one band can err at the corner of two when the shared entries are poorly read; the value from the posterior of the state stays the reported one, the draws through the physics stand beside it, and the gap says how far it errs and whether a decision depends on it.

Table 8.10 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 8.10: The terms of the joint exceedance and where each is defined. Units for the two-unit hall.

Term	Expression	Units	Meaning	Defined in
Operating envelope	every $T_k \le T_{\max}$	°C	the region where every limit holds	8.1.1
Joint exceedance	$\max(T_2, T_3) > T_{\max}$, a union	–	some inlet over its limit	8.1.1
Sparsity pattern	eq. (8.2)	–	shared and own columns	8.1.4
Rows and posterior				
Second reading set	eq. (8.3), seven rows	–	shared entries poorly read	8.2.2
Its covariance	$\boldsymbol{\Sigma}_{\mathrm{II}}$, eq. (8.4)	mixed	the refit's posterior	8.2.2
Questions and answers				
Draws through the physics	fraction of draws over, se $\sqrt{\hat{p}(1-\hat{p})/n}$	–	set beside the reported value	8.3
Gradient matrix	$\mathbf{D}$, rows $\mathbf{g}_{T_2}^{\mathsf{T}}$, $\mathbf{g}_{T_3}^{\mathsf{T}}$	K, K/kW	state to the plane of the inlets	8.3.1
Joint of the inlets	$\mathcal{N}(\cdot,\ \mathbf{D}\boldsymbol{\Sigma}\mathbf{D}^{\mathsf{T}})$	°C, K^2	the bivariate normal	8.3.1
Union	$1-\Phi_2(h_2, h_3;\ \varrho)$	–	exceedance anywhere	8.3.1
Bracket	$\max_k p_k \le \cdot \le \sum_k p_k$	–	the marginals' bounds	8.3.1
Correlation shift	$\delta(\varrho) = \int_0^{\varrho} \varphi_2 \, \mathrm{d}r$	–	union down, intersection up	8.3.1
Covariance split	eq. (8.18)	K^2	shared part plus survey part	8.3.2
Intersection	$p_2 p_3 + \delta(\varrho)$	–	both inlets over	8.3.3
Preposterior joint	$\mathbf{D}\boldsymbol{\Sigma}_a\mathbf{D}^{\mathsf{T}}$	K^2	joint after a candidate reading	8.6.1

8.8 Exercises

Exercise 8.1. *Formulation.* Write the operator's "is any unmeasured rack over 27°C" as a region of the state, give its complement, and state the answer with every reading. (The region is $\{T_2(\mathbf{z}) > 27\} \cup \{T_3(\mathbf{z}) > 27\}$, the event $\max(T_2, T_3) > 27$; its complement is the corner $\{T_2 \le 27,\ T_3 \le 27\}$ of the operating envelope, so the probability is $1 - \Pr(T_2 \le 27,\ T_3 \le 27 \mid \mathbf{y}) = 1 - \Phi_2(h_2, h_3;\ \varrho) = 0.007481$.)

Exercise 8.2. *Derivation.* Show from the block structure of $\mathbf{\Sigma}$ that the fractions of racks 2 and 3 contribute nothing to the covariance of the two inlets, and read the covariance off the fit with every reading. ($\mathrm{Cov}(T_2, T_3) = \mathbf{g}_{T_2}^{\mathsf{T}} \mathbf{\Sigma} \mathbf{g}_{T_3}$. The column of β_2 is nonzero only in $\mathbf{g}_{T_2}$, that of β_3 only in $\mathbf{g}_{T_3}$, and $\mathbf{\Sigma}$ has no covariance between either fraction and any other entry, so every term with a fraction pairs it with a zero entry or a zero covariance; what remains is $\mathbf{D}_{\mathrm{s}} \mathbf{\Sigma}_{\mathrm{ss}} \mathbf{D}_{\mathrm{s}}^{\mathsf{T}}$ in the shared columns, and the covariance is $0.02866\,\mathrm{K}^2$.)

Exercise 8.3. *Derivation.* Derive the bracket (8.13) from the probability rules and check it on the fit with the meter and rack 1 only. ($\Pr(A \cup B) = \Pr(A) + \Pr(B) - \Pr(A \cap B) \le \Pr(A) + \Pr(B)$, and $A \subseteq A \cup B$ gives $\Pr(A \cup B) \ge \Pr(A)$, and likewise for B; with $p_3 = 0.023901$ and $p_2 + p_3 = 0.024321$ the joint value lies between them, at 0.024161.)

Exercise 8.4. *Application.* An engineer computes "both racks over" for the window with the coil sensors out as the product of the two marginals. Explain the error and give the joint answer. (The product assumes $\varrho = 0$, but the two inlets share the columns of Q, g_1 and g_2, the remaining rows read the log-derates poorly, and their correlation is 0.4755; by equation (8.22) the joint normal adds $\delta(\varrho)$ to the product 1.0034×10^{-5}. The draws through the physics, beside it, give 2.160×10^{-4}, 21.53 times the product, and show the linearization low by 3.80 of their standard errors; the reported answer is the joint normal's 1.6012×10^{-4}, 15.96 times the product.)

Exercise 8.5. *Application.* The technician is sent to rack 2 instead of rack 3 with the 0.3 K probe. What does the reading do to the probability that some unmeasured inlet is over the limit? (The row $\mathbf{g}_{T_2}^{\mathsf{T}}/0.3$ narrows β_2 and, slightly, the shared entries: rack 2 falls to 0.2821 K and rack 3 only to 0.8277 K. Rack 2 carried nothing of the union, so at the present means the union falls only to 0.007433.)

Part III

Detecting and Diagnosing Abnormal Operation

Chapter 9

Anomaly Detection from Residuals

9.1 A discharge reading that climbs

It is 00:55, window 24 of the night, and the operator is reading unit 2's tile on the dashboard. The unit's discharge air reads 21.78°C. Its fixed high-discharge alarm sits at 23.0°C, so the tile is green; the tile shows the reading and its alarm, and no band. The trend behind it shows the first hour, 23:00 to 23:55, averaging 19.98°C, and since midnight the reading has read 19.88°C at window 13, 00:00, 20.88°C at window 19, 00:30 and 21.44°C at window 22, 00:45. The operator reads the climb as noise on a warm night: the reading is below its alarm, so unit 2 is fine.

The alarm limit does not settle that. It says how warm the discharge may get before the cold aisle suffers, not whether 21.78°C is what this hall, at tonight's load, with the other four readings of the window, should show. A residual is the gap between a reading and what the plant predicts for it; what the dashboard cannot say is whether the gap of window 24 is the scatter of a sensor or the plant moving away from how it ran. The question of this chapter is

whether a residual is noise or a change, how the plant as a whole is tested against how it ran, and when a change first shows.

It is answered on the hall and the rows of Chapter 5, now over a record of windows, with the residual tests of Chapter 7. What changes is the evidence: a reference period is fitted first, its posterior becomes the prior of every later window, and the state carries the window index, $\mathbf{z}_t$.

Jargon: bad-data detection, residual, reference period (baseline), normalized residual, leverage, redundancy number, hat matrix, chi-square of the reading rows, false-alarm level, Bonferroni level, smearing, fixed alarm

Bad-data detection (in the framework, the predictive check of the rows at the fit) tests whether the readings of a window agree with one another and with what the plant is known to be: the chi-square test of the residual vector and the largest normalized residual. A *residual* is a row of the fit at the MAP state, predicted minus observed over the row's standard deviation; a negative residual is a reading above its prediction. The *reference period*, or *baseline*, is a stretch of operation taken as normal and fitted first; here the first hour. The *hat matrix* $\mathbf{H} = \mathbf{J}\boldsymbol{\Sigma}\mathbf{J}^{\mathsf{T}}$ of Chapter 7 maps the rows' errors onto the part the fit absorbs; its diagonal entry h_{kk} is the *leverage* of row k, and $1 - h_{kk}$ its *redundancy number*, the share of the row's error that the fit cannot absorb. The *normalized residual* r_k^{N} of a reading is its residual divided by the root of its redundancy number; under the model it is a standard normal variable. The *chi-square of the reading rows* is the sum of their squared residuals, referred to a chi-square distribution. The *false-alarm level* α is the probability per window that a test fires on a normal plant. A *Bonferroni level* sets the decision level of the largest of several statistics so that their combined false-alarm probability stays below α. *Smearing* is the spread of one reading's error, or of a change, into the residuals of other readings. A *fixed alarm* is a limit on one reading, set once and compared with that reading alone.

Key idea

A residual is noise or a change only against a reference. Its normalized residual is the reading's distance from its predictive, the prediction of the reading from the baseline and the window's other readings, in units of the predictive's spread; the chi-square of the reading rows, counted by their redundancy numbers, tests the window as a whole. A fit free to move the entries a change acts on absorbs the change into the state, and its residuals stay quiet; a change that a free entry can explain, a load step, is absorbed on purpose.

9.1.1 Noise, change and the reference

Chapter 5 fitted one window and recorded its residuals: the largest normalized residual of its fit was -0.9341 (Table 5.6), and testing such numbers against their noise was left for later. Chapter 7 tested a window's readings against one another. A residual of one window says how well the rows of that window agree with one another and with the priors. On its own it cannot say that the plant has changed, because the fit chooses its state to make the rows agree: a change of a coil is, to a fit with a free log-derate, only another value of the log-derate, and the readings fit as well as they did before.

A change is a departure from something. Here it is the first hour of the night, windows 1 to 12. Fitted as one state, those windows fix the entries that do not move from window to window, the coils' log-derates and the recirculation fractions, much more tightly than any one window can, and that posterior becomes the prior of each later window. A later window's readings are then tested against what the plant of the first hour predicts for them, and a change in an entry the baseline holds cannot be absorbed: it surfaces in the residuals of the readings it acts on.

> **Principle 9.1.** *A test of residuals detects a change only along the entries that a reference holds. Against the rating priors a fit absorbs a coil's change into its log-derate; against a baseline fitted on normal operation the same change surfaces in the residuals of the readings the coil acts on.*

9.1.2 The two-unit hall through the night (step 0)

The plant is the two-unit hall of Chapter 5, Figure 5.1: the IT load Q, two coils with log-derates g_1, g_2 against the ratings $\bar{G}_u = 7.0\,\mathrm{kW/K}$, the chilled-water supply pinned at $T_s = 14.0°\mathrm{C}$, the capacity rates pinned at $C_u = 10.0\,\mathrm{kW/K}$, the other heat $Q_o = 8\,\mathrm{kW}$, and three racks with recirculation fractions β_1, β_2, β_3. Its seven physical rows (5.1) are unchanged. On the reduced state of Chapter 4 their dependent block $\mathbf{A}_{\mathrm{dep}}$ is triangular with a nonzero diagonal, so it is invertible, and the temperatures follow from the free entries in the chain of equation (5.4): the hall balance fixes the return air T_r, the air sides the discharge temperatures $T_{d,u}$, the mixing the cold aisle T_c, and the inlets the T_k. The air balance (5.5), $(C_1 + C_2)(T_r - T_c) = Q + Q_o$, keeps the rise from the cold aisle to the return independent of the coils, so a change of a coil moves the return and the cold aisle together.

The record. The night is a record of 36 windows of five minutes, window t starting at 23:00 plus $5(t-1)$ minutes: window 1 at 23:00, window 13 at 00:00, window 24 at 00:55, window 36 at 01:55. This is the record of Part III, not the shift record of the two-hall plant. Each window carries the five readings of Chapter 5, and the state of window t is

$$\mathbf{z}_t = \big(Q(t),\ g_1,\ g_2,\ \beta_1,\ \beta_2,\ \beta_3\big)^{\mathsf{T}} = \big(Q(t),\ \boldsymbol{\theta}^{\mathsf{T}}\big)^{\mathsf{T}}, \qquad \boldsymbol{\theta} = \big(g_1,\ g_2,\ \beta_1,\ \beta_2,\ \beta_3\big)^{\mathsf{T}}, \tag{9.1}$$

the reduced state (5.3) with the load written $Q(t)$, a value of its own in each window. The sub-vector $\boldsymbol{\theta}$ holds the entries that the first hour shares across windows.

Modeling choices at the graph. The choices of Chapter 5 at the graph carry over to every window: the hall is in steady state within a window, one effective conductance carries each coil, the airflow is pinned and nothing controls the discharge temperature within the window (Section 5.1.2). Each is what makes one window's five readings a row set on one state, and each fails as it did there, after a load step or a unit start, or when a controller holds a discharge at a set point. The choices that are new with the record enter at the rows, Section 9.1.4.

The graph of the record. Figure 9.1 draws the record, not the hall. On the left the twelve windows of the first hour, each with its own load $Q(t)$, hang from one free node (green) carrying $\boldsymbol{\theta}$, the entries they share; together they are one stacked state. Their posterior on $\boldsymbol{\theta}$, the baseline (blue), feeds each later window on the right as a prior row (blue arrows). Every window sends its five readings out along the dashed arrows; window 24, the window of the opening, is drawn in red.

9.1.3 The operator's questions (step 1)

The operator asks four questions of the record, and step 1 writes each as a probability over the state of a window, with the unknown on the left of the bar and the baseline and the window's readings on the right. Table 9.1 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 in the notation of Chapter 2, for a future reading y_{new}, a statistic X^2 of a record, the posterior $p(\mathbf{z} \mid \mathbf{y})$ and the false-alarm level α; the third its mathematical form for the hall; the fourth its category; the fifth the computation of Section 2.3.6 that answers it; and the last the subsection where it is answered. The first three are predictions under a threshold: a reading against its predictive, the window's reading rows against theirs, and the first window at which either falls past its level. Each is the tail of a predictive, which is computed from the posterior of the window. The fourth is an estimation, the posterior of a window fitted without the baseline, asked to show what the reference contributes.

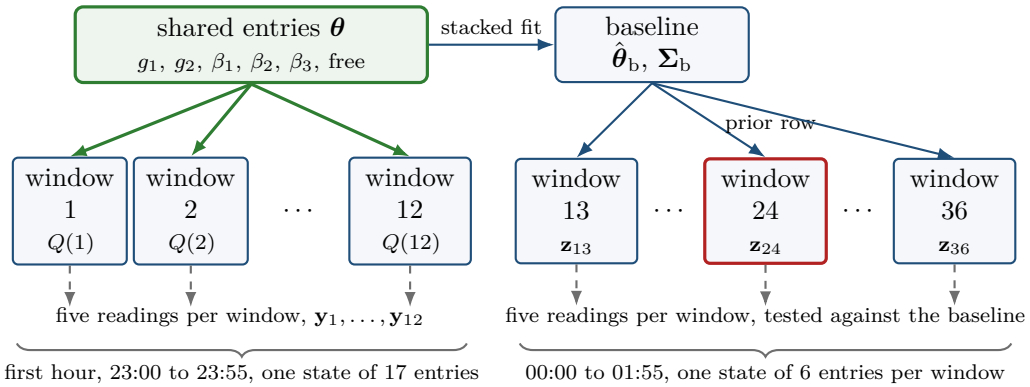


Figure 9.1: The record of the night. Free (green): the shared entries $\boldsymbol{\theta}$ of the first hour; each of its windows adds its own load $Q(t)$. Blue: the baseline, the posterior of the stacked first hour on $\boldsymbol{\theta}$, which enters every later window as a prior row. Each later window is fitted on its own state $\mathbf{z}_t$ of equation (9.1); window 24 (red) is the window of the opening. Observed (dashed): the five readings of each window, those of Chapter 5. The hall inside each window is Figure 5.1.

Table 9.1: The questions asked of the night's record. $\mathbf{y}$ is the baseline and the readings of window t; $\hat{\mathbf{z}}_t$ the MAP state of window t; r_k the residual of reading k and h_{kk} its leverage; c the decision level of the largest of five normalized residuals at $\alpha = 10^{-3}$.

Operator question	General form	Mathematical form	Category	Computation	Answered in
Is this residual noise or a change?	$\int p(y_{\mathrm{new}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$, its tail	$r_k^{\mathrm{N}} = r_k(\hat{\mathbf{z}}_t)/\sqrt{1-h_{kk}}$ of unit 2's discharge, window 24; $2\Phi(-\lvert r_k^{\mathrm{N}}\rvert)$	prediction under a threshold	posterior	9.3.1
Is the plant as a whole still behaving like the first hour?	$\Pr\big(X^2(\mathbf{y}_{\mathrm{rep}}) \geq X^2(\mathbf{y})\big)$	$X_t^2 = \sum_{k\leq 5} r_k(\hat{\mathbf{z}}_t)^2$ 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 $\Pr(\chi^2_{\nu_t} > X_t^2) < 10^{-3}$; first t with $\max_k \lvert r_k^{\mathrm{N}}\rvert > c$	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

9.1.4 The evidence as rows (step 2)

Step 2 attaches the record to the graph as rows, in two fits. The first is the baseline: the twelve windows of the first hour stacked into one state. The second is each later window on its own state $\mathbf{z}_t$, with the baseline as its prior.

The readings. Each window t carries the five readings of Chapter 5, together $\mathbf{y}_t$: the return air y_{T_r}, the two discharge temperatures $y_{T_{d,1}}$ and $y_{T_{d,2}}$, the load meter y_Q and rack 1's inlet y_{T_1}. The four temperatures are read with standard deviation 0.3 K, the load with 3.0 kW. The predicted value of reading k is $m_k(\mathbf{z}_t)$, its temperature by equation (5.4), or $Q(t)$ for the meter. In window 24, the window of the opening, the readings are $y_{T_r} = 30.95°\text{C}$, $y_{T_{d,1}} = 19.90°\text{C}$, $y_{T_{d,2}} = 21.78°\text{C}$, $y_Q = 198.45\,\text{kW}$ and $y_{T_1} = 22.95°\text{C}$.

The baseline as one stacked state. The first hour is fitted as one state of 17 entries,

$$\mathbf{z}_\text{b} = (\boldsymbol{\theta}^\mathsf{T},\ Q(1),\ Q(2),\ \ldots,\ Q(12))^\mathsf{T}, \tag{9.2}$$

the five shared entries and one load per window. Its rows are the five reading rows of each of the twelve windows, sixty in all, each of the form of equation (5.7) with $\mathbf{z}_t = (Q(t), \boldsymbol{\theta}^\mathsf{T})^\mathsf{T}$, and the five prior rows of Chapter 5 on $\boldsymbol{\theta}$, $g_u/0.15$ and $(\beta_k - \beta_k^\text{s})/0.08$ with the survey's $\beta_k^\text{s} = 0.20$, 0.35, 0.50: 65 rows on 17 entries. With $*$ for an entry that can be nonzero, the 65×17 Jacobian has the block pattern

$$\mathbf{J}_\text{b} = \begin{bmatrix} \mathbf{S}_1 & \mathbf{q}_1 & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{S}_2 & \mathbf{0} & \mathbf{q}_2 & \cdots & \mathbf{0} \\ \vdots & & & \ddots & \\ \mathbf{S}_{12} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{q}_{12} \\ \mathbf{P} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} \end{bmatrix},$$

$$\mathbf{S}_t = \begin{bmatrix} * & * & 0 & 0 & 0 \\ * & * & 0 & 0 & 0 \\ * & * & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ * & * & * & 0 & 0 \end{bmatrix}, \quad \mathbf{q}_t = \begin{bmatrix} * \\ * \\ * \\ * \\ * \end{bmatrix}, \quad \mathbf{P} = \operatorname{diag}\big(0.15^{-1}, 0.15^{-1}, 0.08^{-1}, 0.08^{-1}, 0.08^{-1}\big), \tag{9.3}$$

with the columns of $\boldsymbol{\theta}$ first, then $Q(1)$ to $Q(12)$; the rows of $\mathbf{S}_t$ and $\mathbf{q}_t$ are the return air, the two discharge temperatures, the meter and rack 1's inlet of window t. Three facts of the graph are read here. The columns of g_1, g_2 and β_1 hold rows from all twelve windows, so their information adds across the hour. The column of each load holds the five rows of its own window and nothing else: a load is an entry of one window. The columns of β_2 and β_3 hold only their prior rows, so the first hour leaves them where the survey put them.

A later window under the baseline. The baseline enters a later window as one prior row on $\boldsymbol{\theta}$. With $\hat{\boldsymbol{\theta}}_\text{b}$ the baseline mean and $\boldsymbol{\Sigma}_\text{b}$ its covariance, the 5×5 block of the stacked covariance on the shared entries (the loads of the first hour integrated out), and $\mathbf{L}_\text{b}$ the upper-triangular factor with $\mathbf{L}_\text{b}^\mathsf{T}\mathbf{L}_\text{b} = \boldsymbol{\Sigma}_\text{b}^{-1}$, the prior term is $\frac{1}{2}(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_\text{b})^\mathsf{T}\boldsymbol{\Sigma}_\text{b}^{-1}(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_\text{b})$, written as the five normalized rows $\mathbf{L}_\text{b}(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_\text{b})$. The load has no prior row: the meter and the temperatures read it. The

10×6 Jacobian of a later window has the pattern

$$\mathbf{J}_t = \begin{bmatrix} * & * & * & 0 & 0 & 0 \\ * & * & * & 0 & 0 & 0 \\ * & * & * & 0 & 0 & 0 \\ * & 0 & 0 & 0 & 0 & 0 \\ * & * & * & * & 0 & 0 \\ 0 & * & * & * & 0 & 0 \\ 0 & 0 & * & * & 0 & 0 \\ 0 & 0 & 0 & * & 0 & 0 \\ 0 & 0 & 0 & 0 & * & 0 \\ 0 & 0 & 0 & 0 & 0 & * \end{bmatrix} \begin{matrix} \text{return air } y_{T_r} \\ \text{discharge } y_{T_{d,1}} \\ \text{discharge } y_{T_{d,2}} \\ \text{load meter } y_Q \\ \text{rack-1 inlet } y_{T_1} \\ \text{baseline row 1} \\ \text{baseline row 2} \\ \text{baseline row 3} \\ \text{baseline row 4} \\ \text{baseline row 5} \end{matrix}, \qquad \mathbf{g}_{T_{d,2}}^{\mathsf{T}} = (*,\ *,\ *,\ 0,\ 0,\ 0), \tag{9.4}$$

with columns $Q(t), g_1, g_2, \beta_1, \beta_2, \beta_3$. The reading rows are those of Chapter 5, pattern (5.6); the baseline rows replace its five diagonal prior rows by an upper triangle in the observed entries (g_1, g_2, β_1), because the first hour has correlated them, and keep β_2 and β_3 diagonal, because it has not touched them. The second line is the gradient of the predicted discharge of unit 2, the reading the opening asks about; it is not a row. Two facts follow before any number. The column of $Q(t)$ has nonzero entries in reading rows only, so the load is the one entry of the window that no reference holds; and the discharge of unit 2 acts on $(Q(t), g_1, g_2)$, the entries the baseline holds except the load.

Modeling choices at the rows. Three choices make these rows. Each reading is taken as Gaussian, unbiased and independent of the others given the state, the form of the reading rows of Chapter 4: this makes each reading one normalized row, and a sensor's bias is what the tests of this chapter look for. The first hour is taken as normal operation, the reference: this is what makes its posterior on $\boldsymbol{\theta}$ the prior of every later window, and it fails if a change was already under way before 23:55, in which case the baseline holds the changed plant and the test measures departures from it (Section 9.6.5); the first hour is tested against its own baseline in Section 9.4.1. The readings' errors are taken as independent between windows: this lets the twelve windows' rows add their information in pattern (9.3) and makes a later window's errors independent of the baseline, and it fails if a sensor's error drifts slowly, so that neighboring windows share it (Section 9.6.5).

9.2 The posterior of the state (step 3)

All four questions of Table 9.1 are read from the posteriors of one stacked fit and of the windows fitted under it. The inputs are the readings and the rows of Section 9.1.4 and the pinned quantities of Section 9.1.2.

9.2.1 The rows of the baseline and of a window

The baseline minimizes the objective of the 65 rows of Section 9.1.4 on the stacked state (9.2),

$$\mathcal{L}_{\mathrm{b}}(\mathbf{z}_{\mathrm{b}}) = \frac{1}{2}\sum_{t=1}^{12}\sum_{k=1}^{5}\left(\frac{m_k(Q(t),\boldsymbol{\theta}) - y_{k,t}}{\sigma_k}\right)^2 + \frac{1}{2}\sum_{u=1}^{2}\left(\frac{g_u}{0.15}\right)^2 + \frac{1}{2}\sum_{k=1}^{3}\left(\frac{\beta_k - \beta_k^{\mathrm{s}}}{0.08}\right)^2, \tag{9.5}$$

with $y_{k,t}$ reading k of window t. A later window minimizes the objective of its ten rows on $\mathbf{z}_t$,

$$\mathbf{r}(\mathbf{z}_t) = \begin{bmatrix} (T_r(\mathbf{z}_t) - y_{T_r})/0.3 \\ (T_{d,1}(\mathbf{z}_t) - y_{T_{d,1}})/0.3 \\ (T_{d,2}(\mathbf{z}_t) - y_{T_{d,2}})/0.3 \\ (Q(t) - y_Q)/3.0 \\ (T_1(\mathbf{z}_t) - y_{T_1})/0.3 \\ \mathbf{L}_\mathrm{b}(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_\mathrm{b}) \end{bmatrix}, \qquad \mathcal{L}(\mathbf{z}_t) = \tfrac{1}{2}\,\mathbf{r}(\mathbf{z}_t)^\mathsf{T}\mathbf{r}(\mathbf{z}_t), \tag{9.6}$$

the readings of window t, every row dimensionless. The rows are nonlinear in the state through e^{g_u} and $\beta_1(T_r - T_c)$, and the Jacobians are evaluated at each iterate.

9.2.2 The normal equations and the covariance

The stacked fit runs the normal equations (4.29) on $\mathbf{J}_\mathrm{b}$, from the coils at their ratings, the survey's fractions and each window's meter reading for its load. Its covariance $(\mathbf{J}_\mathrm{b}^\mathsf{T}\mathbf{J}_\mathrm{b})^{-1}$ is 17×17; its block on $\boldsymbol{\theta}$ is $\boldsymbol{\Sigma}_\mathrm{b}$ and its mean there $\hat{\boldsymbol{\theta}}_\mathrm{b}$. By the pattern (9.3) the information is block-diagonal between (g_1, g_2, β_1) with the loads and (β_2, β_3), so $\boldsymbol{\Sigma}_\mathrm{b}$ leaves β_2 and β_3 uncorrelated with every other entry and at the survey's spread.

A later window runs

$$(\mathbf{J}_t^\mathsf{T}\mathbf{J}_t)\,\Delta\mathbf{z} = -\mathbf{J}_t^\mathsf{T}\mathbf{r}(\mathbf{z}_t), \qquad \boldsymbol{\Sigma} = (\mathbf{J}_t^\mathsf{T}\mathbf{J}_t)^{-1}\big|_{\hat{\mathbf{z}}_t}, \tag{9.7}$$

from the window's meter reading and the baseline mean, $\mathbf{z}_t^{(0)} = (y_Q,\ \hat{\boldsymbol{\theta}}_\mathrm{b}^\mathsf{T})^\mathsf{T}$. The MAP state $\hat{\mathbf{z}}_t$ is its limit. The covariance $\boldsymbol{\Sigma}$ is the Gauss–Newton form of equation (4.38), and the posterior is approximately normal, $\mathbf{z}_t \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}_t, \boldsymbol{\Sigma})$; Section 9.4.2 sets the full Hessian beside it.

The residuals at the fit. The rows at $\hat{\mathbf{z}}_t$ carry the tests. Chapter 7 derived the residual vector at the fit of linear rows, $\mathbf{r}(\hat{\mathbf{z}}) = -(\mathbf{I} - \mathbf{H})\boldsymbol{\zeta}$ with covariance $\mathbf{I} - \mathbf{H}$ and the hat matrix $\mathbf{H} = \mathbf{J}_t\boldsymbol{\Sigma}\mathbf{J}_t^\mathsf{T}$, here 10×10, equation (7.12); the effective degrees of freedom, $\sum_k (1 - h_{kk}) = n - d$ for n rows and d entries, equation (7.13); and the normalized residual (7.15), standard normal under the model. The column k of $\mathbf{I} - \mathbf{H}$ is the response of every residual to a unit error in row k; its off-diagonal entries are the smearing.

9.2.3 The baseline of the first hour

The stacked fit of the first hour, 17 entries on 65 rows, converges in 7 Gauss–Newton iterations. Table 9.2 gives the baseline on the shared entries: each row is an entry, the columns its baseline mean $\hat{\boldsymbol{\theta}}_\mathrm{b}$, its posterior spread $\sqrt{(\boldsymbol{\Sigma}_\mathrm{b})_{ii}}$, the spread of the prior of Chapter 5, and the rows that inform it. The coils are known to 0.0069 and 0.0073 in log-derate, against 0.024 and 0.025 from the one window of Chapter 5 (Table 5.3), because the twelve windows' temperature rows all act on the same two columns. The correlations of the observed block are -0.4139 for g_1 with g_2, 0.2391 for g_1 with β_1 and 0.2281 for g_2 with β_1. The first is negative for the reason of Section 5.2.3: the return air fixes the sum of the two conductances, and a larger g_1 is explained by a smaller g_2. The fractions of racks 2 and 3 stay at the survey, as the pattern (9.3) requires.

9.2.4 Window 24 under the baseline

Every window of the record is fitted this way, from its own meter reading and the baseline mean; window 24 is shown in full.

Table 9.2: The baseline: the posterior of the stacked first hour, windows 1 to 12, on the shared entries $\boldsymbol{\theta}$, with the spread of each entry's prior.

Entry [–]	$\hat{\boldsymbol{\theta}}_{\mathrm{b}}$ [–]	$\sqrt{(\boldsymbol{\Sigma}_{\mathrm{b}})_{ii}}$ [–]	Prior sd [–]	Informed by
g_1	−0.05611	0.00692	0.15	return, discharge and rack-1 readings of twelve windows
g_2	−0.11763	0.00730	0.15	return, discharge and rack-1 readings of twelve windows
β_1	0.21313	0.00971	0.08	rack-1 inlet readings of twelve windows and the survey
β_2	0.35000	0.08000	0.08	survey only
β_3	0.50000	0.08000	0.08	survey only

The iterations. Table 9.3 runs the normal equations (9.7) for window 24, 00:55, from $\mathbf{z}_{24}^{(0)} = (198.45,\ \hat{\boldsymbol{\theta}}_{\mathrm{b}}^{\mathsf{T}})^{\mathsf{T}}$. Each row is an iteration; the columns give the objective and the norm of its gradient. The objective falls from 27.1318 to 17.2878 in the first step and settles at 17.2863 by the second; the gradient's norm falls from 5.3×10^{2} to the level of the finite differences, near 10^{-7}.

Table 9.3: Gauss–Newton iterations of window 24 under the baseline, equation (9.7), from the meter's load and the baseline mean.

Iteration k [–]	Objective $\mathcal{L}(\mathbf{z}_{24}^{(k)})$ [–]	$\lVert\mathbf{J}_t^{\mathsf{T}}\mathbf{r}\rVert$ [–]
0	27.1318	5.3×10^{2}
1	17.2878	1.2×10^{1}
2	17.2863	6.6×10^{-2}
3	17.2863	2.9×10^{-3}
4	17.2863	2.0×10^{-5}
5	17.2863	1.0×10^{-6}
6 to 9	17.2863	1.4 to 2.7×10^{-7}
10	17.2863	5.6×10^{-8}

The MAP state. Window 24 ends at $Q(24) = 206.553\,\mathrm{kW}$, $g_1 = -0.05669$, $g_2 = -0.12780$, $\beta_1 = 0.21294$, with β_2 and β_3 at the survey's values, the coils close to the baseline's -0.05611 and -0.11763: the baseline holds them. The load is the entry the baseline does not hold, and the fit has raised it above the meter's 198.45 kW.

The Jacobian and the covariance. The Jacobian at the fit, columns $Q(t), g_1, g_2, \beta_1, \beta_2, \beta_3$, rows in the order of equation (9.6), is

$$\mathbf{J}_{24}(\hat{\mathbf{z}}_{24}) = \begin{bmatrix} 0.2609 & -28.9874 & -26.9976 & 0 & 0 & 0 \\ 0.0883 & -46.8442 & -9.1408 & 0 & 0 & 0 \\ 0.1002 & -11.1305 & -44.8544 & 0 & 0 & 0 \\ 0.3333 & 0 & 0 & 0 & 0 & 0 \\ 0.1298 & -28.9874 & -26.9976 & 35.7588 & 0 & 0 \\ 0 & 171.4067 & 80.243 & -42.9804 & 0 & 0 \\ 0 & 0 & 140.6607 & -24.1334 & 0 & 0 \\ 0 & 0 & 0 & 103.0359 & 0 & 0 \\ 0 & 0 & 0 & 0 & 12.5 & 0 \\ 0 & 0 & 0 & 0 & 0 & 12.5 \end{bmatrix}, \tag{9.8}$$

the first column in kW^{-1} and the others dimensionless. The zeros sit where the pattern (9.4) places them, and the baseline rows dwarf the reading rows in the columns of g_1, g_2 and β_1. The covariance on the observed block (Q, g_1, g_2, β_1) and its correlations are

$$\boldsymbol{\Sigma}_{\mathrm{o}} = \begin{bmatrix} 4.93058636 & 0.0016569 & 0.00181607 & 0.00036491 \\ 0.0016569 & 0.00004382 & -0.00001916 & 0.00001472 \\ 0.00181607 & -0.00001916 & 0.00004889 & 0.00001482 \\ 0.00036491 & 0.00001472 & 0.00001482 & 0.00008641 \end{bmatrix},$$

$$\begin{bmatrix} 1 & 0.1127 & 0.117 & 0.0177 \\ 0.1127 & 1 & -0.414 & 0.2393 \\ 0.117 & -0.414 & 1 & 0.228 \\ 0.0177 & 0.2393 & 0.228 & 1 \end{bmatrix},$$

in kW^2, kW and dimensionless units; the fractions of racks 2 and 3 keep the survey's 0.08, uncorrelated. The spreads are 2.220 kW for the load, 0.00662 for g_1, 0.00699 for g_2 and 0.00930 for β_1: the coils keep the baseline's spread, the load is known to about the meter's. The signs of the correlations are those of Chapter 5, Section 5.2.3.

9.2.5 The refits an answer uses

Window 24 without the baseline. The same window is fitted on its five readings and the rating and survey priors of Chapter 5, the rows of equation (5.7) and the normal equations (5.8). It converges in 5 iterations at $Q(24) = 198.165\,\mathrm{kW}$ (spread 2.769 kW), $g_1 = -0.06358$ (0.02217), $g_2 = -0.23876$ (0.02589) and $\beta_1 = 0.21138$ (0.03098).

Window t without one reading. Removing the row of reading k from the residual vector (9.6) leaves nine rows on six entries. Their normal equations, started from $\hat{\mathbf{z}}_t$, give the MAP state $\hat{\mathbf{z}}_{t,-k}$ and the covariance $\boldsymbol{\Sigma}_{-k}$. These refits are run for each of the five readings of every window from 13 to 36, and for unit 2's discharge in the window-alone fit.

One reading under the baseline. A test of unit 2's discharge alone uses no fit of the window: the posterior of the state is the load as the meter reads it, $Q(t) \sim \mathcal{N}(y_Q,\ 3.0^2)$, independent of the baseline, $\boldsymbol{\theta} \sim \mathcal{N}(\hat{\boldsymbol{\theta}}_{\mathrm{b}}, \boldsymbol{\Sigma}_{\mathrm{b}})$. With $\mathbf{g}$ the gradient of $T_{d,2}$ at $(y_Q, \hat{\boldsymbol{\theta}}_{\mathrm{b}})$, the delta method carries it to the predictive variance of the reading,

$$\mathrm{sd}_t^2 = \sigma^2 + \mathbf{g}^{\mathsf{T}}\,\mathrm{diag}(3.0^2,\ \boldsymbol{\Sigma}_{\mathrm{b}})\,\mathbf{g}, \qquad \sigma = 0.3\,\mathrm{K}. \tag{9.9}$$

9.3 The answers (step 4)

Every answer is read from a posterior of Section 9.2, through the residuals of its fit, and that value is the one reported. An exact refit, an exact tail or a set of draws stands beside each reported value, with the gap.

9.3.1 The residual of one reading

Operator question. Is this residual noise or a change? The operator asks it of unit 2's discharge in window 24: whether 21.78°C is a reading the plant of the first hour, at this window's load and with the window's other readings, could produce by sensor scatter alone.

General form. A reading is tested against its predictive, equation (2.13),

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

and the test is a tail of it, equation (2.12): the probability that a replicate of the reading lies at least as far from the predictive's center as the reading does.

Specialized form. The question is equations (2.13) and (2.12) with y_{new} reading k of window t, $y_{\text{new}} = m_k(\mathbf{z}_t) + \epsilon_k$ and $\epsilon_k \sim \mathcal{N}(0, \sigma_k^2)$; with $p(\mathbf{z} \mid \mathbf{y})$ the posterior of every row of the window but its own, $\mathcal{N}(\hat{\mathbf{z}}_{t,-k}, \boldsymbol{\Sigma}_{-k})$ of Section 9.2.5; and with the threshold's f the distance of a replicate from the predictive's center, f^* the distance of the reading and its level κ the false-alarm level α, the test firing when the tail falls below it. Call $\hat{m}_{-k} = m_k(\hat{\mathbf{z}}_{t,-k})$ the predicted reading and $V_{-k} = \sigma_k^2\, \mathbf{j}_k^{\mathsf{T}} \boldsymbol{\Sigma}_{-k} \mathbf{j}_k$ its variance, with $\mathbf{j}_k^{\mathsf{T}}$ the row of $\mathbf{J}_t$ that belongs to reading k. One choice enters here: the rows are taken as linear over the spread of the fit, which gives the predictive in closed form from one fit, and the exact refits set beside it test that. The predictive is then the normal of equation (4.64),

$$y_{\text{new}} \mid \mathbf{y} \approx \mathcal{N}\big(\hat{m}_{-k},\ \sigma_k^2 + V_{-k}\big), \qquad e_k = \frac{\hat{m}_{-k} - y_k}{\sqrt{\sigma_k^2 + V_{-k}}}, \tag{9.10}$$

and e_k, the standardized residual of the notation, is the reading's distance from the predictive's center in units of its spread, with the sign of every residual of the book, predicted minus observed (only the one-reading departure δ_t, not a residual of a fit, is written the other way). Under the model e_k is standard normal, and the tail at the reading is

$$\Pr\big(|y_{\text{new}} - \hat{m}_{-k}| \geq |y_k - \hat{m}_{-k}| \mid \mathbf{y}\big) = 2\Phi\big(-|e_k|\big). \tag{9.11}$$

The prediction from the other rows need not be computed by five separate fits. Add the reading back by the rank-one update (4.66). In the units of the reading, with the predictive variance $S = \sigma_k^2 + V_{-k}$, the update moves the predicted reading toward y_k by the share V_{-k}/S of the gap, and it leaves the predicted reading with the variance $V_{-k} - V_{-k}^2/S = V_{-k}\sigma_k^2/S$, so that

$$r_k(\hat{\mathbf{z}}_t) = \frac{1}{\sigma_k}\Big(\hat{m}_{-k} - y_k\Big)\frac{\sigma_k^2}{S}, \qquad h_{kk} = \frac{1}{\sigma_k^2}\,\frac{V_{-k}\,\sigma_k^2}{S} = \frac{V_{-k}}{S}, \qquad 1 - h_{kk} = \frac{\sigma_k^2}{S}.$$

Dividing the first by the root of the third,

$$e_k = \frac{\hat{m}_{-k} - y_k}{\sqrt{S}} = \frac{\hat{m}_{-k} - y_k}{\sigma_k}\,\frac{\sigma_k^2}{S}\,\frac{\sqrt{S}}{\sigma_k} = \frac{r_k(\hat{\mathbf{z}}_t)}{\sqrt{1 - h_{kk}}} = r_k^{\mathrm{N}}, \tag{9.12}$$

the normalized residual (7.15) of the fit of all ten rows: the normalized residual of bad-data detection is the leave-one-out residual, and one fit gives all five. From here on it is written r_k^{N}. The redundancy number of a reading is the share of the sensor's own variance in the variance of the reading as the rest of the window predicts it: near one when the baseline and the other readings predict it much better than the sensor measures it, near zero when they predict it much worse. A reading with a leverage near one is one the fit follows: its residual is small whatever its error, and the division restores its scale. A negative r_k^{N} is a reading above its prediction.

A second choice sets the level. The false-alarm level is set per window, $\alpha = 10^{-3}$, one false alarm in a thousand windows of a normal plant: this makes each window's test stand alone, and it leaves the chance of at least one false alarm over a night of windows larger than α, and no level here is set for the record as a whole. The five readings are tested at once, by the largest. Each r_k^{N} crosses c with probability $2\Phi(-c)$; the probability that any of the five crosses

is at most the sum of the five, whatever their correlation (the Bonferroni bound, *Reliability Engineering by Derivation*, Chapter 22),

$$\Pr\Big(\max_{k\le 5}|r_k^{\mathrm{N}}|>c\Big)\le 10\,\Phi(-c)=\alpha,\qquad c=\Phi^{-1}\Big(1-\frac{\alpha}{10}\Big)=\Phi^{-1}(1-10^{-4})=3.7190. \quad (9.13)$$

The bound is conservative, because the normalized residuals are correlated through the off-diagonal entries of $\mathbf{H}$. The reading with the largest $|r_k^{\mathrm{N}}|$ is the leading suspect, the reading whose removal would lower the misfit most.

Building blocks. The MAP state $\hat{\mathbf{z}}_{24}$, the Jacobian and the covariance (Section 9.2.4); the hat matrix and the residuals (Section 9.2.2); the normalized residuals (9.12); the normal distribution function Φ; the level (9.13); the exact refits (Section 9.2.5).

The residuals of window 24. At $\hat{\mathbf{z}}_{24}$ unit 2's discharge is predicted at 20.449°C against its reading of 21.78°C. Table 9.4 gives the five reading rows at the fit: each row is a reading, the columns its value, its residual $r_k(\hat{\mathbf{z}}_{24})$, its leverage h_{kk}, its normalized residual (9.12) and the two-sided tail (9.11). Figure 9.2 draws the normalized residuals as bars, one per reading, the vertical axis in units of the predictive's spread, with the decision level $\pm c$ of equation (9.13) as dashed lines.

Table 9.4: The reading rows of window 24, 00:55, at the fit under the baseline: residual (predicted minus observed over the sensor's standard deviation), leverage, normalized residual and its two-sided tail.

Reading	Value [°C or kW]	r_k [–]	h_{kk} [–]	r_k^{N} [–]	$2\Phi(-\lvert r_k^{\mathrm{N}}\rvert)$ [–]
return air y_{T_r}	30.95	−0.5150	0.3275	−0.6281	5.30×10^{-1}
discharge $y_{T_{d,1}}$	19.90	−0.7115	0.1057	−0.7523	4.52×10^{-1}
discharge $y_{T_{d,2}}$	21.78	−4.4363	0.1141	−4.7135	2.44×10^{-6}
load meter y_Q	198.45	2.7010	0.5478	4.0168	5.90×10^{-5}
rack-1 inlet y_{T_1}	22.95	−1.9929	0.1550	−2.1680	3.02×10^{-2}

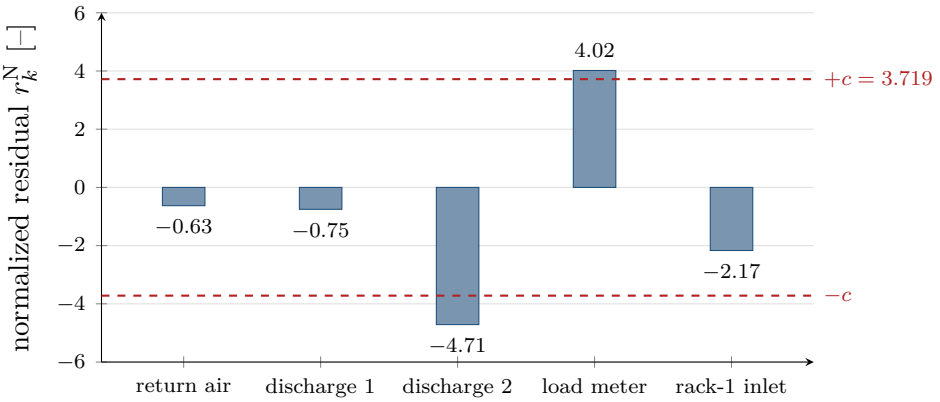


Figure 9.2: The normalized residuals of window 24, 00:55, under the baseline (Table 9.4). Dashed: the decision level $\pm c$ of the largest of five normalized residuals at a false-alarm level of 10^{-3} per window, equation (9.13). A negative bar is a reading above its prediction.

Unit 2's discharge stands at −4.7135 predictive spreads, with a two-sided tail of 2.44×10^{-6}; its leverage is 0.1141, so the fit follows it little and its residual already carries most of its scale. The meter, at +4.0168, is the second reading past the level. Its leverage, 0.5478, is the largest

of the five, because the load is the window's own entry and the meter its most direct row: the fit raises the load and moves the meter's residual with it, which is consistent with a change smeared through the load, the one entry of the window with no reference. The meter's residual does not vanish with the discharge reading: removing that reading from the window leaves the meter at 3.108 and rack 1's inlet at -2.478, with $X^2 = 11.495$ on 2.823 degrees of freedom, a tail of 7.88×10^{-3}. Rack 1's inlet, at -2.1680 here, is below the level; the return air and unit 1's discharge sit inside one spread.

The exact refits beside them. Table 9.5 sets each reading's departure by its exact refit of Section 9.2.5, through the predictive (9.10), beside the linear normalized residual: each row is a reading, the columns the two values, the gap and the two-sided tail of the refit. The widest gap is the meter's, the row with the largest leverage, and no reading changes side of c.

Table 9.5: Window 24: the standardized departure of each reading from its predictive by an exact refit without it, against the linear normalized residual (9.12).

Reading	Exact refit e_k [–]	Linear r_k^{N} [–]	Gap [–]	Tail of the refit [–]
return air y_{T_r}	-0.6289	-0.6281	-0.0009	5.294×10^{-1}
discharge $y_{T_{d,1}}$	-0.7520	-0.7523	0.0003	4.520×10^{-1}
discharge $y_{T_{d,2}}$	-4.7155	-4.7135	-0.0021	2.411×10^{-6}
load meter y_Q	3.9955	4.0168	-0.0213	6.455×10^{-5}
rack-1 inlet y_{T_1}	-2.1691	-2.1680	-0.0011	3.007×10^{-2}

9.3.2 The window as a whole

Operator question. Is the plant as a whole still behaving like the first hour? The question is asked of all five readings of a window together, without naming one.

General form. The statistic of a record is compared with the statistics that replicate records of the same plant would show, a tail of the predictive (2.13) of the statistic,

$$\Pr(X^2(\mathbf{y}_{\mathrm{rep}}) \geq X^2(\mathbf{y})) = \int \mathbf{1}\left[X^2(\mathbf{y}_{\mathrm{rep}}) \geq X^2(\mathbf{y})\right] p(\mathbf{y}_{\mathrm{rep}})\, \mathrm{d}\mathbf{y}_{\mathrm{rep}},$$

with $p(\mathbf{y}_{\mathrm{rep}})$ the distribution of replicate records under the model, a small value saying that the record misfits more than the model's own records do.

Specialized form. The question is the tail of equation (2.13) with y_{new} the replicate $\mathbf{y}_{\mathrm{rep}}$, the five readings of a window redrawn from the baseline's predictive, and $p(\mathbf{z} \mid \mathbf{y})$ the baseline: $\boldsymbol{\theta}$ from the baseline and the load at any value, since its direction carries no residual; X^2 is the misfit of the reading rows, refitted on each replicate. Two modeling choices enter here. The statistic is the sum of the squared residuals of the five reading rows, not of all ten: this asks whether the window's readings agree with the predictive of the baseline, and it leaves the residuals of the five baseline rows, the departures of the window's entries from the reference, to be read as estimates of a change in Chapters 10 and 11. The statistic is referred to a chi-square distribution on the sum of the reading rows' redundancy numbers, kept as a non-integer: this matches the statistic's mean, equation (7.14), and it is an approximation in its shape, which the eigenvalues below bound and the exact tail tests.

For window t,

$$X_t^2 = \sum_{k=1}^{5} r_k(\hat{\mathbf{z}}_t)^2, \qquad \nu_t = \sum_{k=1}^{5} (1 - h_{kk}), \qquad \Pr(X^2(\mathbf{y}_{\mathrm{rep}}) \geq X_t^2) \approx \Pr(\chi^2_{\nu_t} > X_t^2), \tag{9.14}$$

with the sums over the five reading rows. By equation (7.12) the reading residuals $\mathbf{r}_\mathrm{R}$ are normal with covariance $\mathbf{B}_\mathrm{R}$, the 5×5 reading block of $\mathbf{I}-\mathbf{H}$, and $\mathbb{E}[X_t^2] = \operatorname{tr}\mathbf{B}_\mathrm{R} = \nu_t$. By equation (7.13) the redundancy numbers of all ten rows sum to ten rows less six entries, four, shared between the reading rows and the baseline rows. For a non-integer ν_t the tail of the reference is the regularized upper incomplete gamma function, $\Pr(\chi^2_\nu > x) = \Gamma(\nu/2,\ x/2)/\Gamma(\nu/2)$.

The shape of the reference. The block $\mathbf{B}_\mathrm{R}$ is a principal block of a projection, so it is symmetric with eigenvalues λ_j between zero and one, and it is not itself a projection. Rotating the reading residuals onto its eigenvectors,

$$\begin{aligned} X_t^2 &= \sum_{j=1}^{5} \lambda_j\, u_j^2, \qquad u_j \sim \mathcal{N}(0,1) \text{ independent,} \\ \mathbb{E}[X_t^2] &= \sum_j \lambda_j = \nu_t, \qquad \operatorname{Var}(X_t^2) = 2\sum_j \lambda_j^2 \le 2\sum_j \lambda_j = 2\nu_t, \end{aligned} \tag{9.15}$$

since $\lambda_j^2 \le \lambda_j$ on $[0,1]$. The statistic has the mean of $\chi^2_{\nu_t}$ and a smaller variance: the reference is the wider of the two, and its tail beyond a high level errs toward fewer alarms. The weighted sum (9.15) also has an exact tail, its expansion as a mixture of chi-square distributions, which is computed beside the reference. One eigenvalue is zero, and the graph says which. The load has no prior row, so its column of $\mathbf{J}_t$ is $(\mathbf{v}^\mathsf{T},\ \mathbf{0}^\mathsf{T})^\mathsf{T}$ with $\mathbf{v}$ its part in the reading rows; since $(\mathbf{I}-\mathbf{H})\mathbf{J}_t = \mathbf{0}$, the reading block of that product is $\mathbf{B}_\mathrm{R}\mathbf{v} = \mathbf{0}$. The direction in which a change of the load moves the five readings carries no residual at all.

Building blocks. The residuals and the leverages of the five reading rows (Table 9.4); X_t^2, ν_t and the incomplete gamma function of equation (9.14); the eigenvalues of $\mathbf{B}_\mathrm{R}$, equation (9.15), for the exact tail and the rate.

Window 24 as a whole. By equation (9.14), the reading rows give

$$X_{24}^2 = 31.7195, \qquad \nu_{24} = \sum_{k=1}^{5}(1-h_{kk}) = 3.7498, \qquad \Pr(\chi^2_{3.7498} > 31.7195) = 1.608\times10^{-6}.$$

The five reading rows hold 3.7498 of the four degrees of freedom of equation (7.13), and the five baseline rows the rest: the baseline is tight, so its rows are hard to move and check little. The readings of window 24 are not what the plant of the first hour produces, on the evidence of all five at once. The block $\mathbf{B}_\mathrm{R}$ has the eigenvalues 0, 0.9154, 0.9171, 0.9173 and 1, with sum 3.7498 and sum of squares 3.5205: by equation (9.15) the statistic has variance 7.0410 against the 7.4996 of $\chi^2_{3.7498}$. The exact tail at 31.7195 is 8.540×10^{-7}, a gap of -7.537×10^{-7} from the reported value, and both are far below α.

The false-alarm rate beside the level. The 10^{-3} level of the reference at window 24 is $X^2 = 17.933$. The exact probability that the weighted sum (9.15) exceeds it is 7.519×10^{-4}, a gap of -2.481×10^{-4} from the level, and at the 10^{-2} level 8.506×10^{-3}, a gap of -1.494×10^{-3}: the reference is conservative, the test firing a little less often than its nominal level. These rates are at window 24's eigenvalues; at the 10^{-3} level window 1's give 7.707×10^{-4}, window 13's 7.709×10^{-4} and window 36's 7.242×10^{-4}. Two million draws of the same sum give 7.730×10^{-4} and 8.494×10^{-3}, each within about one Monte Carlo standard error (2.0×10^{-5} and 6.5×10^{-5}) of the exact rate.

Twenty thousand windows are drawn from the model, with the shared entries drawn from the baseline, a load of 200 kW and the stated sensor errors, and each is fitted and tested as the night is. A share 0.00065 (standard error 0.00018) has a chi-square tail below 10^{-3}, within one

standard error of the exact rate. The median tail is 0.4946 (standard error 0.0035), near the 0.5 of a calibrated test.

9.3.3 The first window past the level

Operator question. When did it start to show? The question asks for the first window of the record in which a test of the night says the first-hour model no longer fits.

General form. A detection time is the first window at which a tail of the predictive falls below the false-alarm level,

$$t_{\mathrm{d}} = \min\{t : \Pr(X^2(\mathbf{y}_{\mathrm{rep}}) \geq X^2(\mathbf{y}_t)) < \alpha\},$$

the threshold of equation (2.12) applied window by window.

Specialized form. The question is equation (2.12) applied to each window from 13 to 36 in turn. Two tests of the five readings are run. For the chi-square test f is the statistic of a replicate, f^* the window's observed X_t^2 and κ the level α, and it fires at the first t with $\Pr(\chi^2_{\nu_t} > X_t^2) < \alpha$ by equation (9.14). For the largest normalized residual $f = \max_k |r_k^{\mathrm{N}}|$ with the level c through the bound (9.13), and it fires at the first t with $\max_k |r_k^{\mathrm{N}}| > c$.

Two tests of one reading are set beside them. The dashboard's fixed alarm fires at the first t with $y_{T_{d,2}} > 23.0$°C, a limit with no prediction. A one-reading band tests unit 2's discharge alone against its predictive from the baseline, with the load taken from the meter: the predictive (4.64) with mean $T_{d,2}(y_Q, \hat{\boldsymbol{\theta}}_{\mathrm{b}})$ and the variance (9.9), and its standardized departure

$$\delta_t = \frac{y_{T_{d,2}} - T_{d,2}(y_Q,\ \hat{\boldsymbol{\theta}}_{\mathrm{b}})}{\mathrm{sd}_t}, \tag{9.16}$$

positive for a reading above its prediction. It uses the baseline, as the five-reading tests do, but not the window's other readings. A comparison in time holds only at a common level. One statistic at $\alpha = 10^{-3}$ has the two-sided level $\Phi^{-1}(1 - \alpha/2) = 3.2905$; the level $c = 3.7190$ of the largest of five has, on one statistic, the nominal rate $2\Phi(-c) = 2.000 \times 10^{-4}$, at which the largest of five is held against $\Phi^{-1}(1 - 2.000 \times 10^{-5}) = 4.1075$. All the tests are read at both.

Building blocks. The tail of equation (9.14) and the largest normalized residual (9.12) of each window from 13 to 36; the level of equation (9.13); the mean and the variance (9.9) of the one-reading predictive and δ_t; the exact tail and the exact refits of each window.

The reading against its alarm and its band. Figure 9.3 draws unit 2's discharge reading in every window, the horizontal axis the window number with its clock time, the vertical axis the reading in degrees. The red line is the dashboard's fixed alarm at 23.0°C. The dashed line and the shaded band are the one-reading band of equation (9.16): the predictive mean of the reading from the baseline at the meter's load, and that mean ± 3.719 predictive standard deviations, each window its own, the standard deviation 0.324 to 0.327 K over the night. In window 24 the predictive is 20.072°C with standard deviation 0.325 K. The dotted line marks the stated truth of the record, a check marker for Section 9.4.3. The reading first crosses the alarm in window 34, 01:45; its largest value of the night is 23.61°C in window 35, 01:50. It first leaves the drawn band in window 21, 00:40, at $\delta_{21} = 3.851$, and it passes the one-reading level 3.2905 in window 20, 00:35, at $\delta_{20} = 3.381$.

At window 24 the gradient of equation (9.9) on (Q, g_1, g_2) is

$$(0.0294\ \text{K/kW},\ -3.1294\ \text{K},\ -12.9477\ \text{K}),$$

and the variance $0.10550\,\mathrm{K}^2$ splits into the sensor's 0.09, the load's 0.00779, g_1's 0.00047, g_2's 0.00894 and the cross term -0.00169, negative because g_1 and g_2 correlate negatively and their gradients share a sign. Four million draws of the predictive through the exact physics give $\delta_{18} = 3.167$ against the linear 3.167, $\delta_{20} = 3.380$ against 3.381 and $\delta_{21} = 3.852$ against 3.851.

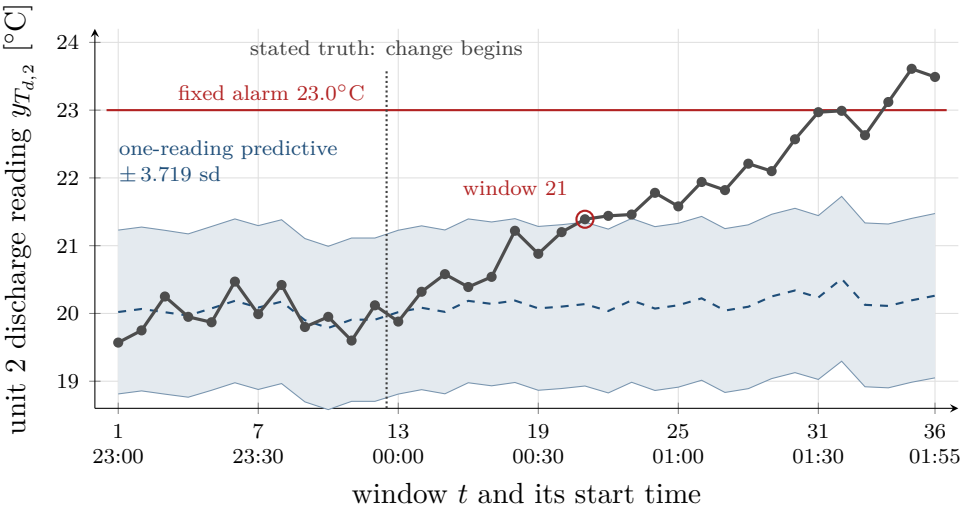


Figure 9.3: Unit 2's discharge reading through the night, every window, against the dashboard's fixed alarm at 23.0°C (red). Dashed and shaded: the one-reading band, equation (9.16), the predictive mean of the reading from the baseline at the meter's load and ±3.719 of its standard deviation, a nominal 2.000×10^{-4}, window by window. Circled: window 21, the first reading outside the band. Dotted: the start of the stated truth's change, a check marker (Section 9.4.3).

Table 9.6 gives the tests window by window from window 13: each row is a window, the columns its clock time, the one-reading departure δ_t, the normalized residual of unit 2's discharge, the largest normalized residual of the five readings and the reading that carries it, X_t^2 and its chi-square tail.

The detection times. At $\alpha = 10^{-3}$ the chi-square test fires in window 20, 00:35, with a tail of 6.13×10^{-4}, and its tail stays below 10^{-3} in every window after; the one-reading band, at 3.2905, fires in the same window. The largest normalized residual first passes 3.719 in window 22, 00:45, at 3.964, and the reading that carries it is unit 2's discharge. At the nominal 2.000×10^{-4} the band, at 3.719, fires in window 21, 00:40; the chi-square test in window 22, 00:45, by its reference (window 21's tail is 2.35×10^{-4}) and in window 21, 00:40, by the exact tail (1.63×10^{-4}), and both are stated; the largest normalized residual, at 4.1075, in window 24, 00:55. The dashboard's fixed alarm fires in window 34, 01:45. Figure 9.4 draws the chi-square tail of every window on a logarithmic vertical axis, the horizontal axis the window; the horizontal red line is the level 10^{-3} and the three vertical lines the detection times at $\alpha = 10^{-3}$ of the five-reading tests and the fixed alarm. The tail falls from 0.801 in window 13 to 8.38×10^{-25} in window 36, and at $\alpha = 10^{-3}$ the chi-square test crosses its level 14 windows, 70 minutes, before the fixed alarm.

The exact values beside the detections. At $\alpha = 10^{-3}$ neither detection moves. The exact tail of the weighted sum, each window on its own eigenvalues, first falls below 10^{-3} in window 20, at 4.50×10^{-4} against the reported 6.13×10^{-4}; in window 19 it is 5.66×10^{-2}. The largest departure by exact refits first passes 3.719 in window 22, at 3.965 against the reported 3.964. The widest gap of the largest statistic between the two is the meter's in window 34, 8.010 against 8.077. Of the quiet windows of Section 9.3.2 a share 0.00075 passes 3.719, standard error

Table 9.6: The night under the baseline, windows 13 to 36: the one-reading departure (9.16), the normalized residual of unit 2's discharge, the largest normalized residual of the five readings and its reading, the reading-row chi-square (9.14) and its tail. Bold: the first window past $|\delta_t| > 3.719$, $\max_k |r_k^{\mathrm{N}}| > 3.719$ and tail $< 10^{-3}$, levels that differ in their nominal rates.

t [–]	Start [h:min]	δ_t [–]	$r^{\mathrm{N}}_{T_{d,2}}$ [–]	$\max_k \lvert r_k^{\mathrm{N}}\rvert$ [–]	At [–]	X_t^2 [–]	$\Pr(\chi^2_{\nu_t} > X_t^2)$ [–]
13	00:00	−0.427	0.435	1.120	$y_{T_{d,1}}$	1.490	8.01×10^{-1}
14	00:05	0.725	−0.871	1.594	y_{T_r}	4.266	3.38×10^{-1}
15	00:10	1.715	−1.537	1.788	y_{T_1}	5.729	1.96×10^{-1}
16	00:15	0.628	−0.449	1.351	y_{T_1}	2.891	5.37×10^{-1}
17	00:20	1.227	−0.604	3.154	y_Q	11.238	1.98×10^{-2}
18	00:25	3.167	−3.186	3.186	$y_{T_{d,2}}$	11.536	1.75×10^{-2}
19	00:30	2.477	−2.172	2.172	$y_{T_{d,2}}$	8.635	6.05×10^{-2}
20	00:35	3.381	−2.896	3.040	y_Q	19.007	$\mathbf{6.13 \times 10^{-4}}$
21	00:40	**3.851**	−3.357	3.357	$y_{T_{d,2}}$	21.085	2.35×10^{-4}
22	00:45	4.326	−3.964	**3.964**	$y_{T_{d,2}}$	29.395	4.87×10^{-6}
23	00:50	3.894	−3.471	3.471	$y_{T_{d,2}}$	21.367	2.06×10^{-4}
24	00:55	5.258	−4.713	4.713	$y_{T_{d,2}}$	31.719	1.61×10^{-6}
25	01:00	4.489	−3.507	5.655	y_Q	39.116	4.67×10^{-8}
26	01:05	5.276	−4.521	4.868	y_Q	38.498	6.24×10^{-8}
27	01:10	5.474	−4.949	4.949	$y_{T_{d,2}}$	36.796	1.44×10^{-7}
28	01:15	6.496	−5.843	5.843	$y_{T_{d,2}}$	54.408	2.96×10^{-11}
29	01:20	5.679	−5.111	5.111	$y_{T_{d,2}}$	39.537	3.80×10^{-8}
30	01:25	6.846	−6.056	6.056	$y_{T_{d,2}}$	60.948	1.20×10^{-12}
31	01:30	8.404	−7.559	7.559	$y_{T_{d,2}}$	83.451	2.04×10^{-17}
32	01:35	7.590	−7.098	7.098	$y_{T_{d,2}}$	62.644	5.19×10^{-13}
33	01:40	7.700	−6.721	6.721	$y_{T_{d,2}}$	72.624	4.12×10^{-15}
34	01:45	9.259	−7.996	8.077	y_Q	99.583	7.46×10^{-21}
35	01:50	10.501	−9.625	9.625	$y_{T_{d,2}}$	110.978	2.74×10^{-23}
36	01:55	9.916	−8.701	8.701	$y_{T_{d,2}}$	118.014	8.38×10^{-25}

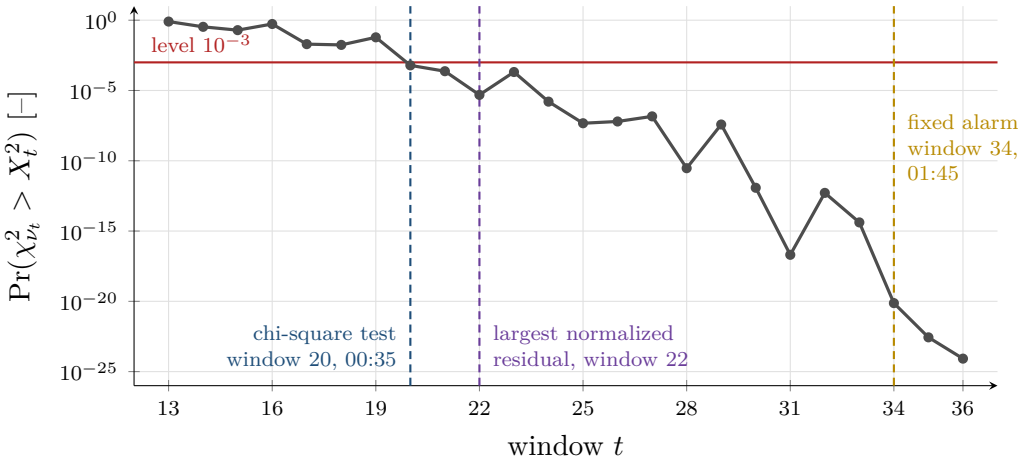


Figure 9.4: The chi-square tail of the reading rows, equation (9.14), in every window from 00:00 (Table 9.6), on a logarithmic axis. Red: the level 10^{-3}. Dashed: the first window past the level for the chi-square test (blue), for the largest normalized residual against $c = 3.719$ (purple) and for the dashboard's fixed alarm at 23.0°C (gold).

0.00019. Beside the 10^{-3} of the bound (9.13), the multivariate normal of the five normalized residuals at window 24 gives exactly $\Pr(\max_k |r_k^{\mathrm{N}}| > c) = 9.751 \times 10^{-4}$, standard error 7×10^{-6}, so the bound's slack is small.

What the five readings add to one. In time they add nothing on this record: at a common level the band and the chi-square test fire within a window of each other, and the largest normalized residual later. The band's false-alarm rate can be stated as well: in the quiet draws it passes 3.2905 in a share 0.00105 (standard error 0.00023) and 3.719 in 0.00005 (standard error 0.00005). Nor does a load step read correctly by the meter separate them: the discharge of unit 2 moves with the load by 0.0301 K/kW at window 24, so a 30 kW step moves the reading by 0.90 K; the band's center follows the meter, and the five-reading fit absorbs the step by the zero eigenvalue of Section 9.3.2. A band drawn without a fit, from the reading's own scatter in the first hour, would take that move for a change. In window 24 the band's departure, $\delta_{24} = 5.258$, is larger than unit 2's normalized residual, 4.713 in magnitude, because the five-reading fit lets the load move while the band takes the meter's load as it reads.

The five readings add two things. The test does not rest on one meter: the band's center moves with any error of the meter, by 0.0294 K/kW at window 24, and it cannot see that error, while the fit reads the load from all its rows and reports the meter's own residual. And they locate the disagreement. From window 22 the largest normalized residual is unit 2's discharge in 12 of the 15 windows; in windows 25, 26 and 34 it is the load meter's, 5.655, 4.868 and 8.077, against -3.507, -4.521 and -7.996 for unit 2's discharge, which is below the level in window 25 only. Before window 22 the meter carries the largest in windows 17 and 20, the second the window of the first chi-square detection, 3.040 against -2.896 for unit 2's discharge.

9.3.4 The window without its reference

Operator question. Why not refit everything each window? The question is why the baseline is needed at all: Chapter 5 fitted one window on its own, with the coils on their rating priors, and read its residuals.

General form. The window alone is an estimation question, a derived quantity under the posterior of its state given its own readings and the rating priors, equation (2.9),

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

Specialized form. The question is equation (2.9) with $f(\mathbf{z}_t) = g_2$, f^* any level of it, $\mathbf{y}_t$ the window's five readings and $p(\mathbf{z}_t \mid \mathbf{y}_t)$ the Laplace normal of the window-alone fit of Section 9.2.5. This is the fit that a plant without a reference period runs, and its residuals are tested by the same equations (9.12) and (9.14).

What a change does to the residuals of a fit follows from the residual at the fit. With the rows linear, the choice of Section 9.3.1, $\mathbf{r}(\mathbf{z}) = \mathbf{J}\mathbf{z} - \tilde{\mathbf{y}}$, with $\tilde{\mathbf{y}}$ the readings and prior means over their standard deviations, split into the reading rows (R) and the prior rows (P). Without a change, $\tilde{\mathbf{y}} = \mathbf{J}\mathbf{z}^{\mathrm{true}} + \boldsymbol{\zeta}$. A change moves the plant's entry j by Δ: the readings follow it, the prior rows, which hold what was known before, do not, and the shift of $\tilde{\mathbf{y}}$ is

$$\mathbf{s} = \begin{bmatrix} \mathbf{J}_{\mathrm{R}}\, \mathbf{e}_j \\ \mathbf{0} \end{bmatrix} \Delta = \mathbf{J}\, \mathbf{e}_j\, \Delta - \begin{bmatrix} \mathbf{0} \\ \mathbf{J}_{\mathrm{P}}\, \mathbf{e}_j \end{bmatrix} \Delta.$$

The residual at the fit is $-(\mathbf{I}-\mathbf{H})$ times the errors and the shift, equation (7.12), and $(\mathbf{I}-\mathbf{H})\mathbf{J} = \mathbf{0}$, so

$$\mathbf{r}(\hat{\mathbf{z}}) = -(\mathbf{I} - \mathbf{H})\, \boldsymbol{\zeta} + (\mathbf{I} - \mathbf{H}) \begin{bmatrix} \mathbf{0} \\ \mathbf{J}_{\mathrm{P}}\, \mathbf{e}_j \end{bmatrix} \Delta. \tag{9.17}$$

The part of the change that lies along a column of the Jacobian is absorbed: it moves the state, not the residuals. What reaches the residuals is the change seen through the prior rows' column of entry j, $\mathbf{J}_{\mathrm{P}}\mathbf{e}_j$. For a diagonal prior that column has one entry, the reciprocal of the prior's spread of z_j. For the baseline's whitened rows it is $\mathbf{L}_{\mathrm{b}}\mathbf{e}_j$, a vector of length $\sqrt{(\mathbf{\Sigma}_{\mathrm{b}}^{-1})_{jj}}$, the root of the baseline's precision on z_j, which is at least the reciprocal $1/\sqrt{(\mathbf{\Sigma}_{\mathrm{b}})_{jj}}$ of its spread. The reading part of a coil's change in the residuals therefore increases with how tightly the reference holds the coil, up to the full reading shift $\mathbf{J}_{\mathrm{R}}\mathbf{e}_j\Delta$. On the load, which has no prior row, $\mathbf{J}_{\mathrm{P}}\mathbf{e}_Q = \mathbf{0}$ and a load step is absorbed whole, the zero eigenvalue of Section 9.3.2.

Building blocks. The window-alone fit of window 24 and its exact refit (Section 9.2.5); its normalized residual of unit 2's discharge (9.12) and its statistic (9.14); the two columns of the prior rows on g_2, from $\mathbf{\Sigma}_{\mathrm{b}}$ (Section 9.2.3).

Window 24 alone. The fit carries the warm discharge in $g_2 = -0.23876$, and the readings fit. Unit 2's discharge has the normalized residual -0.9721, and the reading rows give $X^2 = 0.3487$ on $\nu = 1.2016$ degrees of freedom, with a tail of 0.633. The window alone sees nothing. The reading rows hold fewer degrees of freedom than under the baseline, 1.2016 against 3.7498, because loose priors let the entries follow the readings. Beside these stand the exact refit, -0.9574, a gap of 0.0147, and the exact tail on the eigenvalues 0, 0.0117, 0.0375, 0.1524 and 1, 0.7053, a gap of 0.0728: the reference is less close here, where the eigenvalues are far from one, and both values lie far inside the level.

The columns of the prior rows on g_2 in equation (9.17) set the difference. The rating prior's column is the single entry $1/0.15 = 6.67$. The baseline's whitened column has length $\sqrt{(\mathbf{\Sigma}_{\mathrm{b}}^{-1})_{g_2g_2}} = 161.9$, above the reciprocal of the baseline's spread, $1/0.00730 = 137.0$, because the baseline also correlates g_2 with g_1 and β_1. Per unit of g_2, the shift of the five reading residuals in equation (9.17) has the norm 44.716 under the baseline and 1.128 under the rating prior, against a full reading shift of 59.61 under the baseline. Under the baseline that shift is (7.020, 2.353, 34.326, -22.639, 15.934) on the return air, the two discharges, the meter and rack 1's inlet: a coil's change would put a residual in the meter of the sign the meter shows. The meter's residual fits the coil explanation; ranking the explanations is Chapter 10.

Figure 9.5 sets the two fits side by side: the horizontal axis is the window, the vertical axis the normalized residual of unit 2's discharge; the line is the residual under the baseline in every window from 00:00, the dashed lines the decision level $\pm c$, and the open diamond the residual of window 24 fitted alone. The line first leaves the band of $\pm c$ in window 22 and stays outside it from window 26; the diamond sits inside one spread.

9.4 The checks (step 5)

The answers of Section 9.3 rest on a normal first hour, on the reduced state and on approximations that can be tested on this hall, against the calculation each stands for.

9.4.1 The first hour against its own baseline

A test that fires on normal operation is no test. The twelve windows of the first hour, each refitted on the rows of equation (9.6) with the baseline as its prior, give chi-square tails from 0.151 upward, with a median of 0.666, and their largest normalized residual over the hour is 1.784; the largest one-reading $|\delta_t|$ is 1.388. No window comes near the level 10^{-3} or the decision level 3.719. The check is generous to the tests, since the first hour is both the reference and the record tested against it, and its residuals are smaller than a fresh normal hour's would be; the quiet draws of Section 9.3.2 are fresh windows.

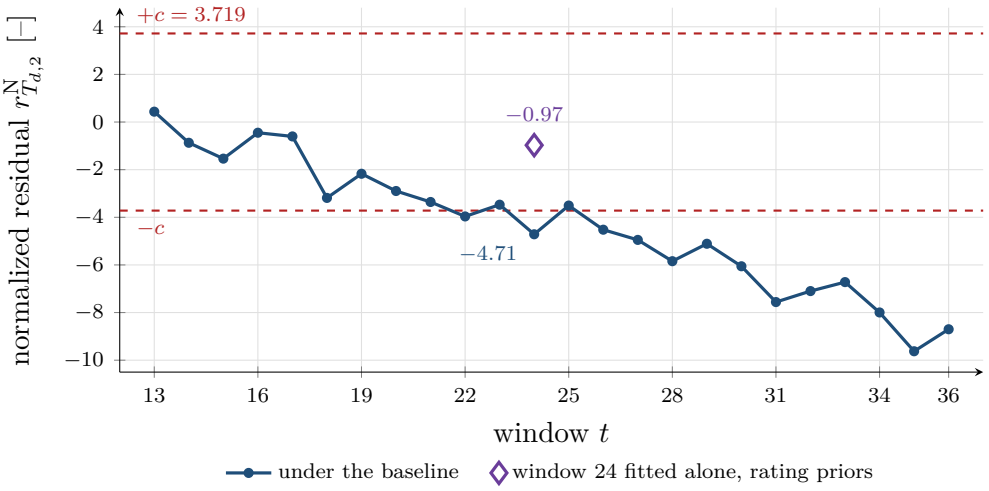


Figure 9.5: The normalized residual of unit 2's discharge through the night under the baseline (Table 9.6), and that of window 24 fitted alone on the rating priors of Chapter 5, equation (5.7). Dashed: the decision level $\pm c$, equation (9.13).

9.4.2 The limit of exact rows and the curvature term

The fit runs on the reduced state. The check fits window 24's own full state of thirteen entries, equation (5.2), with the seven physical rows (5.1) at a finite tolerance beside the readings and the baseline row, and lets the tolerance shrink. Table 9.7 gives the result: each row is a tolerance, the columns the load and g_2 with their spreads and the statistic, and the last row the reduced state. At 0.1 the load stands at 206.4317 kW and the statistic at 30.9409; at 0.001 the full state reproduces the reduced state in every column but the last digit of the statistic.

Table 9.7: The limit of exact physical rows on window 24's full state, under the baseline row, against the reduced state.

Tolerance σ_p [kW or K]	$Q(24)$ [kW]	sd(Q) [kW]	g_2 [–]	sd(g_2) [–]	X^2_{24} [–]
0.1	206.4317	2.2282	−0.12778	0.00699	30.9409
0.01	206.5518	2.2206	−0.12780	0.00699	31.7107
0.001	206.5531	2.2205	−0.12780	0.00699	31.7194
reduced state	206.5531	2.2205	−0.12780	0.00699	31.7195

The covariance (9.7) is the Gauss–Newton form, without the curvature term of the nonlinear rows. The full Hessian of the objective at $\hat{\mathbf{z}}_{24}$ gives the load the spread 2.2136 kW against 2.2205 kW and g_2 the spread 0.00695 against 0.00699.

9.4.3 The answers against the stated truth

The record of the night is generated from a stated truth (Section 9.6.5). In it unit 2's coil changes from window 13, 00:00, and its log-derate at window 24 is −0.23834 against the −0.12000 of the first hour. Table 9.8 collects the checks: each row is a route, the columns the reported value, the value beside it and the gap, the second minus the first. One gap crosses a level: at 2.000×10^{-4} the chi-square test is window 22 by its reference and window 21 by the exact tail. The tests raise no alarm in the first hour, before the change. The first windows after the start carry residuals within a sensor's scatter (Table 9.6).

Table 9.8: The checks of the answers and their gaps from the reported values. Refits and exact tails at window 24 unless stated; the draws are of the weighted sum and of quiet windows from the baseline; the truth is the stated truth of the constructed record.

Route [–]	Reported	Beside it	Gap
Under the baseline			
$r^{\mathrm{N}}_{T_{d,2}}$, exact refit	-4.7135	-4.7155	-0.0021
$r^{\mathrm{N}}_{y_Q}$, exact refit	4.0168	3.9955	-0.0213
Chi-square tail, exact	1.608×10^{-6}	8.540×10^{-7}	-7.537×10^{-7}
Rate at the 10^{-3} level, exact	10^{-3}	7.519×10^{-4}	-2.481×10^{-4}
Rate at the 10^{-2} level, exact	10^{-2}	8.506×10^{-3}	-1.494×10^{-3}
Exact rate at 10^{-3}, two million draws	7.519×10^{-4}	7.730×10^{-4} (se 2.0×10^{-5})	2.1×10^{-5}
Exact rate at 10^{-2}, two million draws	8.506×10^{-3}	8.494×10^{-3} (se 6.5×10^{-5})	-1.2×10^{-5}
Largest residual past 3.719, multivariate normal	10^{-3} (bound)	9.751×10^{-4} (se 7×10^{-6})	–
Largest residual past 3.719, 20000 quiet windows	10^{-3} (bound)	0.00075 (se 0.00019)	-0.00025
Chi-square tail below 10^{-3}, 20000 quiet windows	10^{-3}	0.00065 (se 0.00018)	–
Band past 3.2905, 20000 quiet windows	10^{-3}	0.00105 (se 0.00023)	–
Band past 3.719, 20000 quiet windows	2.000×10^{-4}	0.00005 (se 0.00005)	–
One-reading δ_{20}, draws of the predictive	3.381	3.380	–
Newton against Gauss–Newton, sd Q [kW]	2.2205	2.2136	-0.0069
Newton against Gauss–Newton, sd g_2	0.00699	0.00695	-0.00004
Full state, $\sigma_p = 0.001$, X^2_{24}	31.7195	31.7194	-0.0001
Window 24 alone, rating priors			
$r^{\mathrm{N}}_{T_{d,2}}$, exact refit	-0.9721	-0.9574	0.0147
Chi-square tail, exact	0.6325	0.7053	0.0728
The detections			
Chi-square at 10^{-3}, exact tail	window 20	window 20	0
Chi-square at 2.000×10^{-4}, exact tail	window 22	window 21	-1
Largest residual at 10^{-3}, exact refits	window 22	window 22	0
Largest residual at 2.000×10^{-4}, exact refits	window 24	window 24	0
One-reading band, at 10^{-3}; at 2.000×10^{-4}	window 20; window 21	–	–
Fixed alarm	window 34	–	–
First hour against its baseline	no detection	tails from 0.151, median 0.666; largest $\lvert r^{\mathrm{N}}_k \rvert$ 1.784; largest $\lvert \delta_t \rvert$ 1.388	–
Stated truth, when the change began	–	window 13, 00:00	–
Stated truth, g_2 at window 24	–	-0.23834, against -0.12000	–

9.5 The report (step 6)

The climb is not noise. Unit 2's discharge in window 24 stands 4.71 predictive spreads above what the plant of the first hour predicts for it, a tail of 2.4×10^{-6}, while its fixed alarm is still green; at a false-alarm level of 10^{-3} per window the chi-square test of the hall's readings fired at window 20, 00:35, 70 minutes before that alarm. A band on the one reading, from the same baseline and at the same level, fires in the same window. What the five readings add is the location of the disagreement, unit 2's discharge in most windows from window 22, 00:45, and a test that does not rest on one meter.

9.5.1 The answers and their levels

Table 9.9 gives the answers in the order of Table 9.1. Each row is a question; the second column gives the answer with the exact value beside it and the third the level it is held against.

Table 9.9: The report for the night. Tests under the baseline of the first hour at a false-alarm level of 10^{-3} per window.

Question	Answer	Against the level
Is this residual noise or a change?	not noise: unit 2's discharge at window 24, $r^{\mathrm{N}} = -4.71$, tail 2.4×10^{-6} (exact refit: -4.72, tail 2.4×10^{-6})	beyond $c = 3.719$; the meter's 4.02 is also past it, consistent with a change smeared through the load
Is the plant as a whole still behaving like the first hour?	the first-hour model no longer fits: $X^2 = 31.72$ on 3.75 effective degrees of freedom, tail 1.6×10^{-6} (exact: 8.5×10^{-7})	below 10^{-3}; exact rate of the reference at window 24 7.5×10^{-4}
When did it start to show?	at 10^{-3}: window 20, 00:35, chi-square test and one-reading band; window 22, 00:45, largest normalized residual. At 2.000×10^{-4}: window 21, 00:40, band; chi-square window 22, 00:45 (exact: window 21, 00:40); window 24, 00:55, largest normalized residual	fixed alarm window 34, 01:45
Why not refit everything each window?	alone, the fit absorbs the change into g_2: $r^{\mathrm{N}} = -0.97$, tail 0.63 (exact: -0.96, tail 0.71)	inside both levels

9.5.2 The entry in the shift log

The operator's entry for window 24, 00:55, reads the residual test, not the tile: "a departure carried by unit 2's discharge reading, flagged at window 20, 00:35; onset not dated here (Chapter 11); cause not yet isolated (Chapter 10)". The entry rests on three things the record shows. The first-hour model stopped fitting the hall's five readings together at window 20, 00:35, and has not fit them since. Unit 2's discharge reading is the one that carries the departure in most windows. And the meter's large residual fits a change smeared through a larger load, but it does not vanish without the discharge reading, so the entry does not clear the meter. What the entry does not say is when the departure began, which needs the record's windows read together, or what stands behind it: a coil and a discharge sensor both put a residual in that reading, and the test of residuals does not rank them.

Common mistake

Reading the largest residual as the faulty reading. In window 24 two readings are past the level, unit 2's discharge and the load meter; which entry or sensor changed is a question of diagnosis, because a residual says only where the rows disagree.

9.6 The next reading (step 7)

The report leaves a departure at unit 2's discharge reading and two explanations to test, each of which would put a residual there. A further reading is one more row of the residual vector, and the one worth requesting is the one whose row separates the two.

9.6.1 A coil or a sensor

The two explanations are two columns of the Jacobian. A change of the coil is the column of g_2 in the pattern (9.4): nonzero in the rows of the return air, both discharge temperatures and rack 1's inlet, because a coil that passes less heat warms the return and through it the cold aisle. A fault of the discharge sensor is a bias b on that reading, a new column with a single nonzero, $1/0.3$ in the row of $y_{T_{d,2}}$. With the present readings the two columns differ only through the other readings' entries in the column of g_2, and those are what Chapter 10 weighs when it ranks the two hypotheses by their evidence.

9.6.2 The candidates

Three readings would add a row that the two columns treat differently. A technician's probe held in unit 2's discharge air, of standard deviation 0.3 K, reads $T_{d,2}(\mathbf{z}_t)$: its row repeats the discharge row in the column of g_2 and has a zero in the column of b. Under a coil change the probe and the sensor move together; under a sensor fault only the sensor moves. It acts in one window. A permanent second discharge sensor on unit 2, reading the same air, is the same row added to every window: the instrument at the place of the question, which makes every later window's test separate the two columns instead of leaving the separation to the evidence ranking. A second return-air sensor repeats the return-air row, which already acts on g_2 and not on b; it adds precision to a row the hall has, not a row the separation lacks, and it sees g_1 and g_2 through the sum of the two conductances only.

The width of b is the wrong measure of these candidates: at window 24, under the rows with b free, a probe moves $\mathrm{sd}(b)$ only from 0.3184 to 0.3164 K. What separates them is the reading each explanation predicts, a test and not a finding. Fit window 24 twice: once with b free on the discharge row, the baseline row unchanged, and once with b absent and g_2 released from the baseline, the baseline row then taken on $(g_1, \beta_1, \beta_2, \beta_3)$ alone. The first predicts the air at 20.279°C (spread 0.107 K), the second at 21.739°C (spread 0.243 K): a probe's readings differ by 1.461 K, 3.79 probe spreads. A second return-air sensor's predictions differ by 0.468 K only.

9.6.3 The recommendation

The next reading is a probe in unit 2's discharge air tonight, which compares a second measurement of that air with the sensor's. The request beyond tonight is a second permanent discharge sensor on unit 2, so that the same comparison is made in every window by the rows rather than by a visit. Chapter 10 ranks the coil and the sensor on the present readings, and Chapter 18 ranks the readings an operator can request on the shift plant.

9.6.4 What the analysis showed

The dashboard said unit 2 was fine, because its discharge was below its alarm; the analysis found the hall's readings departing from the first hour's plant well before that alarm. The alarm compares one reading with a limit set for the cold aisle; the residual test compares the readings with what the plant of the first hour predicts at tonight's load. The gain in time is over the fixed limit, and it comes from the baseline, not from the number of readings: a predictive band on the one reading, from the same baseline and at the same level, fired in the same window or within one. What the five readings add is the location of the disagreement and a test that does not rest on one meter. The test needs a reference: the same window, fitted as Chapter 5 fitted its window, gives $X^2 = 0.3487$, a tail of 0.633, and sees nothing. And a large residual is not always the source: the load, which no reference holds, moves with whatever the fit must explain, so its meter's residual fits a change elsewhere and does not clear the meter either. The hall's readings were enough to see the departure early; they are not enough to say what changed, and the reading that would test it is a second measurement of the air the discharge sensor reads.

9.6.5 Limits of the model

The first hour as normal. The baseline is the plant of 23:00 to 23:55, taken as normal. A change already under way in that hour is folded into the reference, and the tests then measure departures from a plant that had already changed. The first hour is also the record its own check is run on (Section 9.4.1); the quiet draws test the rate, not the assumption that the hour was normal.

Independence between windows. The readings' errors are independent from window to window. A sensor whose error wanders slowly shares it between neighboring windows; it then enters the baseline's twelve windows as one error rather than twelve, the baseline is narrower than it should be, and the tests fire more often than α.

The chi-square reference and the linearization. The reading-row statistic is referred to $\chi^2_{\nu_t}$, which matches its mean and is wider than its true distribution; its exact rate and the quiet draws place the test below its nominal level at the plant's normal state. The exact refits agree with the linear normalized residuals to the gaps of Table 9.8, and the widest reading-wise gap of the night is the meter's in window 36, 01:55, 8.136 against 8.228, a gap of -0.092. The test is run on the two-unit hall only.

One window at a time. Each later window is tested alone against the baseline, and the record's evidence is not accumulated: windows 17 to 19 each fall short of the level, and pooled, $X^2 = 31.409$ on 11.261 degrees of freedom, their tail of 1.11×10^{-3} is still above it. When a change began and how it moves across windows are the questions of Chapter 11, and the false-alarm rate of a test repeated over a night is not set here.

The steady state and the rows. The limits of Chapter 5 hold in every window: steady state within the window, one fraction per rack, the pinned airflow and no active control of the discharge (Section 5.6.5). A transient after a load step would put residuals in every reading that the tests would read as a change.

Constructed readings. The record of the night is constructed: it is generated from a stated truth of the hall with the stated sensor errors, and Section 9.4.3 compares the detections with that truth. It is not field data, and the agreement is a check of the method on its own model, not of the model against a real hall.

The next chapter. Chapter 10 takes the change the tests have found and asks which component carries it.

9.7 What to retain

Definition 9.1 (Detection from residuals)**.** A window is tested against a reference: a baseline posterior of normal operation, carried into the window as a prior row on the entries it holds. At the fit of the window the residual vector is $-(\mathbf{I}-\mathbf{H})\boldsymbol{\zeta}$ with the hat matrix $\mathbf{H}=\mathbf{J}\boldsymbol{\Sigma}\mathbf{J}^{\mathsf{T}}$; the normalized residual $r_k^{\mathrm{N}}=r_k(\hat{\mathbf{z}}_t)/\sqrt{1-h_{kk}}$ of a reading is its distance from its predictive, the prediction from every other row, and is standard normal under the model; the reading rows' chi-square X_t^2 is referred to $\chi^2_{\nu_t}$ with ν_t the sum of their redundancy numbers, a reference with the statistic's mean and a larger variance; the largest of n normalized residuals is held against $c=\Phi^{-1}(1-\alpha/2n)$. A change along an entry the fit is free to move is absorbed into the state and reaches the residuals only through the prior rows of that entry; an entry with no prior row, such as the load, absorbs its changes whole.

Three conclusions carry beyond the hall. Detection is relative: a residual is a change only against a reference that holds the entries the change acts on, and the tighter the reference, the earlier the change shows. A test against a reference sees a change before a fixed limit on one reading does; against a predictive band on one reading from the same reference, at the same level, a test of many readings gains nothing in time and gains location and independence from any one meter. And a residual locates a disagreement, not its cause: the largest normalized residual is the leading suspect, smearing can lift a reading past the level whether or not its sensor is sound, and telling a component from its sensor needs a row that treats them differently.

Table 9.10 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 9.10: The terms of detection from residuals and where each is defined. Units for the two-unit hall.

Term	Expression	Units	Meaning	Defined in
State of window t	$\mathbf{z}_t = (Q(t), \boldsymbol{\theta}^{\mathsf{T}})^{\mathsf{T}}$	kW, –	the reduced state, indexed	9.1.2
Stacked baseline	eq. (9.2), 17 entries, 65 rows	kW, –	the first hour as one state	9.1.4
Baseline row	$\mathbf{L}_{\mathrm{b}}(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_{\mathrm{b}})$	–	the reference as a prior	9.1.4
Posterior and residuals at the fit (Chapter 7)				
Hat matrix	$\mathbf{H} = \mathbf{J}\boldsymbol{\Sigma}\mathbf{J}^{\mathsf{T}}$	–	what the fit absorbs	9.2.2
Leverage, redundancy number	h_{kk}, $1 - h_{kk}$	–	absorbed and checked share of a row	9.2.2
Smearing	off-diagonal of $\mathbf{I} - \mathbf{H}$	–	one error in another's residual	9.2.2
Fit without a reading	$\hat{\mathbf{z}}_{t,-k}$, $\boldsymbol{\Sigma}_{-k}$	kW, –	the exact refit	9.2.5
Questions and answers				
Predictive of a reading	$\mathcal{N}(\hat{m}_{-k},\ \sigma_k^2 + V_{-k})$	°C, kW	the reading from the other rows	9.3.1
Normalized residual	$r_k^{\mathrm{N}} = r_k(\hat{\mathbf{z}}_t)/\sqrt{1 - h_{kk}} = e_k$	–	distance from the predictive	9.3.1
Bonferroni level	$c = \Phi^{-1}(1 - \alpha/2n)$	–	decision level of the largest	9.3.1
One-reading variance	eq. (9.9)	K^2	the band's predictive	9.2.5
Reading-row chi-square	$X_t^2 = \sum_{k \le 5} r_k^2$ against $\chi^2_{\nu_t}$	–	the window as a whole	9.3.2
Shape of the reference	$X_t^2 = \sum_j \lambda_j u_j^2$, $\lambda_j \in [0, 1]$	–	variance at most $2\nu_t$	9.3.2
One-reading band	δ_t, eq. (9.16)	–	one reading against the baseline	9.3.3
Detection time	first t with the tail below α	window	when the change shows	9.3.3
Absorption	eq. (9.17)	–	a change seen only through a prior	9.3.4

9.8 Exercises

Exercise 9.1. *Formulation.* Write the test of the hall load meter in window 24 as a tail of its predictive and give its value. (The meter's future reading is $y_{\mathrm{new}} = Q(24) + \epsilon$, $\epsilon \sim \mathcal{N}(0,\ 3.0^2)$, predicted from the baseline and the window's four temperature readings, equation (9.10); its normalized residual is $r/\sqrt{1 - h_{kk}}$ with $r = 2.7010$ and $h_{kk} = 0.5478$, so $r^{\mathrm{N}} = 4.0168$, positive because the fit's load is above the meter's reading, and the two-sided tail $2\Phi(-4.0168)$ is 5.90×10^{-5}.)

Exercise 9.2. *Derivation.* Show that the variance of the reading-row chi-square is at most that of its reference, and give both for window 24. (By equation (9.15) the statistic is $\sum_j \lambda_j u_j^2$ with λ_j the eigenvalues of the reading block of $\mathbf{I} - \mathbf{H}$, a principal block of a projection, so $0 \le \lambda_j \le 1$ and $\lambda_j^2 \le \lambda_j$; hence $2\sum_j \lambda_j^2 \le 2\sum_j \lambda_j = 2\nu_t$. In window 24 the eigenvalues sum to 3.7498 and their squares to 3.5205, so the variance is 7.0410 against 7.4996.)

Exercise 9.3. *Derivation.* Show that a change of g_2 by Δ reaches the residuals of a fit only through the prior rows of g_2, and compare the rating prior with the baseline. (The readings shift by $\mathbf{J}_{\mathrm{R}}\mathbf{e}_{g_2}\Delta$ and the prior rows do not; since $(\mathbf{I}-\mathbf{H})\mathbf{J} = \mathbf{0}$ the residual shift is $(\mathbf{I}-\mathbf{H})[\mathbf{0};\ \mathbf{J}_{\mathrm{P}}\mathbf{e}_{g_2}]\Delta$, equation (9.17). The rating prior's column is the single entry $1/0.15 = 6.67$; the baseline's whitened column has length $\sqrt{(\boldsymbol{\Sigma}_{\mathrm{b}}^{-1})_{g_2 g_2}} = 161.9$, above $1/0.00730 = 137.0$. Alone, the fit moves g_2 to -0.23876 and unit 2's normalized residual is -0.9721.)

Exercise 9.4. *Application.* Derive the decision level of the largest of the five normalized residuals at one false alarm in a thousand windows, and find the first window of the night that crosses it. (By the Bonferroni bound (9.13), $10\,\Phi(-c) = 10^{-3}$, so $c = \Phi^{-1}(1 - 10^{-4}) = 3.7190$; Table 9.6 first shows $\max_k |r_k^{\mathrm{N}}| > 3.719$ at 3.964, unit 2's discharge, in window 22.)

Exercise 9.5. *Application.* A 30 kW load step arrives in a quiet window, read correctly by the meter. Say what it does to the one-reading band, to a band on the reading drawn from the first hour's scatter alone, and to the five-reading tests. (Unit 2's discharge moves by 0.0301 K/kW, 0.90 K for the step. A band from the reading's own scatter has no load in it and reads the move as a change. The one-reading band moves its center with the meter's load and follows the step as well as the meter reads it. In the five-reading fit the step lies along the load column, which has no prior row; by $\mathbf{B}_{\mathrm{R}}\mathbf{v} = \mathbf{0}$ it leaves the residuals untouched, the zero eigenvalue of the reading block, against the other eigenvalues 0.9154, 0.9171, 0.9173 and 1.)

Chapter 10

Fault Isolation and Derate Estimation

10.1 One discharge of two

It is 00:55, window 24 of the night. The residual test of Chapter 9 has been past its level for twenty minutes, and its largest normalized residual sits on one reading, unit 2's discharge. The screen shows that sensor at 21.78°C, 1.80 K above its first-hour mean and 1.22 K inside the 23.0°C alarm; unit 1's discharge, at 19.90°C, has moved 0.55 K. One discharge of two has moved, so the dashboard's reading is a sensor drifting, and the entry it would write is "check the unit 2 discharge sensor at the morning round".

A residual says where the rows disagree, not what changed. The sensor could have moved, or the coil behind it, or the airflow across that coil, or the chilled water, or heat from somewhere else could be reaching the hall. The question of this chapter is

which change the record of the window supports, how firmly, what the readings cannot tell apart, and how large the change is under the explanation that leads.

It is answered on the hall, the record and the baseline of Chapter 9, with the ranking of Chapter 7. What changes is the family: it is no longer a list of three meters but every change the plant's graph holds that a reading can see.

Jargon: fault detection and isolation, fault identification, candidate family, fault signature, non-isolable faults, decisive window, time to isolation, correct and false isolation, derate factor, entropy of a ranking

Fault detection and isolation decides that something changed and then where; Chapter 9 did the first and this chapter does the second. *Fault identification* then estimates the size of the change. The *candidate family* is the list of changes that are ranked, generated from the plant graph before the record is read, each one a drawing of the plant with one added entry. A candidate's *fault signature* is the pattern its change leaves in the residuals of the fit with no change, its Jacobian column after the fit has absorbed what it can. Two faults are *non-isolable* on a set of readings when a substitution between them leaves every reading the same, so that no value of those readings can tell them apart. A *decisive* window is one whose leader passes a rule fixed in advance. The *time to isolation* is the first such window, set against the time to detection. *Correct isolation* is a decisive ranking that names the change that was made; *false isolation* is a decisive ranking that names another. The *derate factor* of unit u is $d_u = \mathrm{e}^{g_u} = G_u/\bar{G}_u$, the share of its rated conductance that its coil has. The *entropy* of a ranking, $-\sum_j p_j \ln p_j$ in nats, is zero when one drawing holds all the probability and largest when all hold the same.

Key idea

Detection needs the misfit to be large; isolation needs its pattern to beat near neighbors, so isolation comes later and can come and go from window to window. The family is written down from the graph first, and the rule for acting on the ranking with it. And a ranking can only separate what the readings separate: a family built to hold every change the graph holds that a reading can see also shows which pairs of changes no value of these readings can ever tell apart.

10.1.1 From a residual to a family of changes

Chapter 7 ranked three biases because three meters were the suspects its balance named. Here the residual names nothing so short. A change is any entry the fit holds that could have moved, and the plant's graph is the list of those: an edge whose conductance or fraction the baseline fixes, a reading that could carry a bias, a quantity that was pinned and need not have stayed so. Writing the list from the graph, before the record is ranked, has two effects. It protects the ranking from the explanation that comes to mind first, here the sensor, and it makes the ranking's limit visible, because a change the graph has no place for cannot lead.

Each candidate is a drawing with one more entry and the rows of Chapter 9 otherwise unchanged, and its evidence pays for the misfit it removes and charges for the width it had to admit, the bookkeeping of Chapter 7. The new thing is the comparison of signatures. Two candidates whose signatures are nearly parallel need many readings, or a large change, to come apart; two whose signatures are parallel never do.

> **Principle 10.1.** *Fault isolation ranks a family generated from the plant graph by evidence, against no change, under a rule fixed before the ranking. Its answer is a probability for each listed change, and it is limited by the angle between the signatures: what the readings cannot separate, the ranking reports as a choice of prior, not as a finding.*

10.1.2 The hall and the candidates (step 0)

The plant is the two-unit hall of Chapter 5, Figure 5.1, on the state $\mathbf{z}_t$ of equation (9.1); its dependent block $\mathbf{A}_{\text{dep}}$ is triangular with nonzero diagonal, so the reduced state of Chapter 4 holds and the temperatures follow from the free entries by the forward chain (5.4). The record is the night of Chapter 9: 36 windows of five minutes from 23:00, the first twelve fitted as the baseline $(\hat{\boldsymbol{\theta}}_{\text{b}}, \boldsymbol{\Sigma}_{\text{b}})$ of Table 9.2, cited here and not refitted. The choices of that chapter at the graph carry over: steady state within a window, one effective conductance per coil, the airflow pinned, no controller acting (Section 9.1.2).

The family. Each candidate M_j adds one scalar entry θ_j to $\mathbf{z}_t$, with the prior $\theta_j \sim \mathcal{N}(0,\ w_j^2)$; the baseline row is unchanged, so every drawing has the bookkeeping of the bias in Chapter 7. Table 10.1 generates the family from the graph by rule: each row is a rule, the columns where in the graph it looks, the candidates it generates and their prior widths. Rule 1 releases an entry the baseline holds, θ_j being a change since the first hour. Rule 2 puts a bias on a reading. Rule 3 frees a pinned quantity; the ratings $\bar{G}_u$ are the units in which g_u is written and generate nothing. Chapter 5 pinned T_s because a precise supply sensor reads it and C_u at the units' fixed fan speed, so freeing them means that sensor reads wrong, or that a fan or a damper has changed: this buys a place in the ranking for causes the pins had ruled out, and it fails to say which of those causes it is, since $\ln C_u$ is one entry for fan and damper alike.

Rules 1 to 3 generate fourteen candidates. β_2 and β_3 are excluded by named choice: their columns in every reading row are zero, so their $\ln B$ against no change is 0 in every window,

Table 10.1: The candidate family, generated from the graph before the record is ranked. Each candidate adds one entry $\theta_j \sim \mathcal{N}(0,\ w_j^2)$; M_0 adds none. Rule 4 changes two rows of the graph and stands beside the family, not in it.

Rule [–]	From [–]	Generated [–]	Width w_j [entry's]
1. Release an entry the baseline holds	edges	g_1, g_2, β_1; β_2, β_3 excluded	g_u: 0.5; β_1: 0.2
2. A bias on a reading	readings	b_{T_r}, $b_{T_{d,1}}$, $b_{T_{d,2}}$, b_Q, b_{T_1}	temperatures: 2 K; b_Q: 10 kW
3. Free a pinned quantity	nodes, edges	Q_o (a leak at the return node), T_s, $\ln C_1$, $\ln C_2$	Q_o: 10 kW; T_s: 1 K; $\ln C_u$: 0.5
4. Add a heat term to a balance row that has none, and to the hall's row	nodes	$L_{d,1}$, $L_{d,2}$, L_c; the extended family only	10 kW

as a fit of each at width 0.2 prints at window 24 (0.000000), and no record could isolate them. The family is therefore $J = 12$ candidates beside no change, M_0.

Which leaks this graph supports. The full state (5.2) has four heat balance rows: the hall, coil 1, coil 2 and the mixing of the cold aisle; the rack inlets are fractions, not balances. Only the hall's row carries a heat term, Q_o, so as drawn the plant has one place for a leak, at the return node, which is in the family. The other three rows have no such term. Rule 4 adds one, $L_{d,1}$ or $L_{d,2}$ in a unit's discharge air or L_c in the cold aisle, and, since that heat must leave through the coils, to the hall's row as well; that changes the graph. These three are excluded from the family by named choice, which keeps every candidate a change of the graph as it stands and which fails if heat does enter one of those places, and they are reported beside it as the extended family, $J = 15$, never instead of it. With them the chain (5.4) becomes

$$T_r = T_s + \frac{Q + Q_o + L_{d,1} + L_{d,2} + L_c}{G_1 + G_2}, \qquad T_{d,u} = T_r - \frac{G_u\,(T_r - T_s) - L_{d,u}}{C_u}, \qquad T_c = \frac{C_1 T_{d,1} + C_2 T_{d,2} + L_c}{C_1 + C_2}. \tag{10.1}$$

The graph. Figure 10.1 is Figure 5.1 with the candidates marked where the rules find them. The cold aisle at the left feeds the three racks, the return air in the middle takes the load and the other heat, and the units at the right pass heat to the chilled water and send their discharge air back along the top and bottom; the dashed arrows leave for the five readings. Red marks a candidate of the family: on the coils and rack 1, on each reading, on the other heat, on the chilled water and on each unit's discharge path, which carries its airflow. Purple marks the three added heat terms of the extended family. Gray marks the two excluded fractions.

Modeling choices at the family. Six choices make the family a ranking, and each is named here once. One change at a time: this keeps every drawing one entry from M_0, and it cannot explain two things moving together. The change is constant within a window: one window gives one value per reading, so a constant is all that window can see. Each window is ranked alone, with no evidence accumulated across windows: this keeps the ranking free of any statement about how the change moves, and it gives up the strength a run of windows would add. The priors on the changes are symmetric: a coil could also recover, and the sign is left to the readings, at the price of an Occam charge for the half-line that is never used. The widths of Table 10.1 state what size of change is plausible: the evidence depends on them, and they are held fixed in

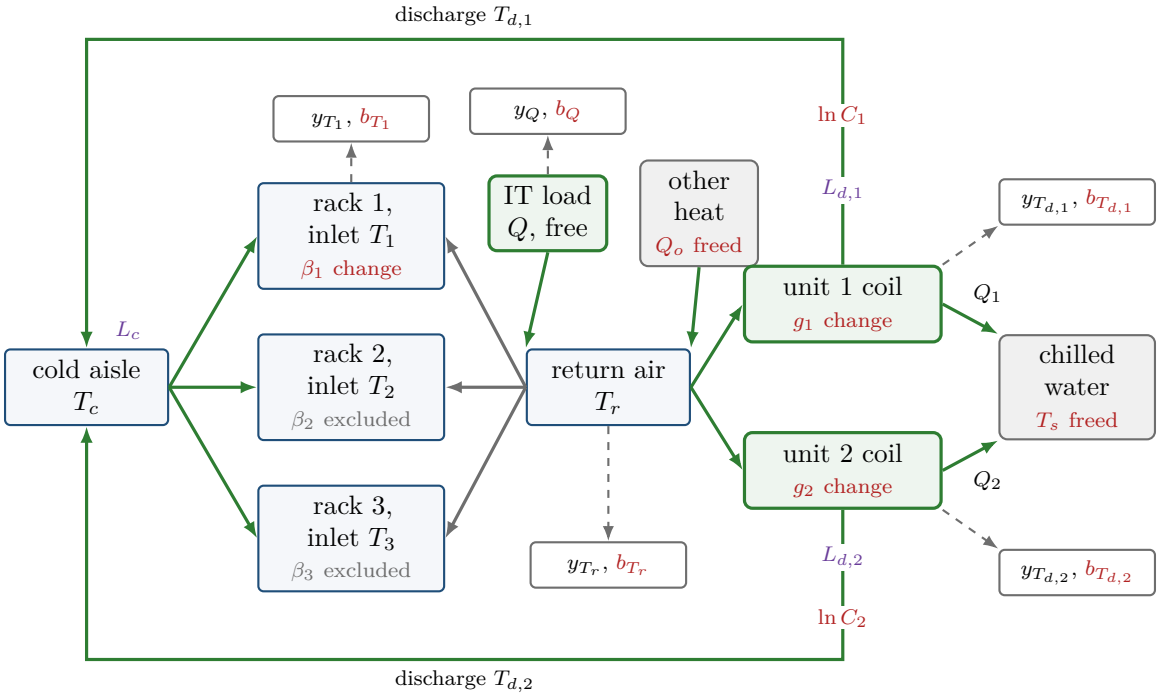


Figure 10.1: The hall's graph with the candidate family marked. Red: the twelve candidates of Table 10.1, a released edge entry on a coil or rack 1 (rule 1), a bias on each dashed reading (rule 2), and the freed other heat, supply temperature and airflows (rule 3). Purple: the heat terms of rule 4, which add to the graph and form the extended family. Gray: the two fractions excluded because no reading row has a column on them. Green free, blue fixed by the rows and gray pinned, as in Figure 5.1; the baseline row on the shared entries is not drawn.

the ranking. The prior probabilities are those of Chapter 7, equation (7.5), now on this family,

$$\Pr(M_0) = 0.5, \qquad \Pr(M_j) = 0.5/J, \quad j = 1, \ldots, J, \tag{10.2}$$

which gives no candidate an advantage and leaves the probabilities to be read as shares among the listed drawings only.

10.1.3 The operator's questions (step 1)

The operator asks five questions of the record, and step 1 writes each as a probability over the drawings, or over the state within one drawing. Table 10.2 lists them. The columns are those of Chapter 5's table, with M_j a drawing, $\mathbf{y}_t$ the baseline and the readings of window t, and $f(\mathbf{z})$ a derived quantity. The first and the fifth are the drill questions; the second is the pair Chapter 9 handed over.

What counts as decisive. The rule is fixed before the ranking and is not moved after it. A window's leader is decisive when (a) $\ln B$ against every other candidate but its structurally non-isolable partner in the family is at least $\ln 20 = 2.9957$, (b) $\ln B$ against M_0 is at least $\ln 20$, and (c) its probability, with that partner's added, is at least 0.95. The partners in the family are b_{T_1} with β_1, each of which acts on the one row y_{T_1}, and b_Q with Q_o, by the substitution of Section 7.2.3. A factor of 20 is a choice: it makes the rule demand more than a leader that is merely ahead, and it would delay an action that a smaller factor would allow.

Table 10.2: The questions asked of the night's record. $\ln B(A : B)$ is the log Bayes factor of drawing A against B; $L_{d,2}$ a heat term in unit 2's discharge air; $d_2 = \mathrm{e}^{g_2}$ the derate factor of unit 2.

Operator question	General form	Mathematical form	Category	Computation	Answered in
If something changed, where?	$\Pr(M_j \mid \mathbf{y}) \propto p(\mathbf{y} \mid M_j) \Pr(M_j)$	the ranking of M_0 and the 12 candidates at window 24	diagnosis	evidence	10.3.1
A coil or a sensor?	the same, 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?	the same, per window	the first window past the rule; from the finished record, the first after which it never fails	diagnosis	evidence	10.3.3
Can these readings say it is the coil?	the same, the extended family	whether a heat term $L_{d,2}$ gives the same readings as M_{g_2}	diagnosis	evidence	10.3.4
Given it is unit B, how far is it from its rating?	$\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$	$p(g_2 \mid \mathbf{y}_t, M_{g_2})$, $d_2 = \mathrm{e}^{g_2}$	estimation	posterior	10.3.5

10.1.4 The evidence as rows (step 2)

The evidence of a window is what it was in Chapter 9: the five reading rows of window t and the five baseline rows $\mathbf{L}_\mathrm{b}(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_\mathrm{b})$, here written on the bold vector of shared entries, which the scalar θ_j is not. Window 24, 00:55, is worked in full; windows 20, 00:35, and 22, 00:45, where the chi-square and largest-residual tests of Chapter 9 first fire, are read in the ranking; the night enters as windows 13 to 36. In window 24 the readings are those of the opening with $y_{T_r} = 30.95\,^\circ\mathrm{C}$, $y_Q = 198.45\,\mathrm{kW}$ and $y_{T_1} = 22.95\,^\circ\mathrm{C}$.

A drawing's rows. Drawing M_j keeps those ten rows, with the reading predictions $m_k^j(\mathbf{z}_t, \theta_j)$ of the chain under its change, and adds the prior row of its entry,

$$\mathbf{r}_j(\mathbf{z}_t, \theta_j) = \begin{bmatrix} (m_k^j(\mathbf{z}_t, \theta_j) - y_{k,t})/\sigma_k, \quad k = 1, \dots, 5 \\ \mathbf{L}_\mathrm{b}(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_\mathrm{b}) \\ \theta_j / w_j \end{bmatrix}, \qquad \mathcal{L}_j = \tfrac{1}{2}\, \mathbf{r}_j^\mathsf{T} \mathbf{r}_j. \tag{10.3}$$

The prediction m_k^j replaces g_u by $g_u + \theta_j$ or β_1 by $\beta_1 + \theta_j$ under rule 1; adds θ_j to one prediction under rule 2; and under rule 3 replaces Q_o by $Q_o + \theta_j$, T_s by $T_s + \theta_j$, or C_u by $C_u\, \mathrm{e}^{\theta_j}$. The drawing of the coil gets its own names, since it carries the answer of the last question: the baseline's entry is g_2^b, the change is $\delta_2 = \theta_{g_2}$ (not the one-reading departure of Chapter 9), and

$$g_2 = g_2^\mathrm{b} + \delta_2 \tag{10.4}$$

is the log-derate of the window, a derived quantity.

Which row acts on which entry. The drawing M_{g_2} has 11 rows on 7 entries. With $*$ for an entry that can be nonzero, its Jacobian is the pattern (9.4) with one column and one row added,

$$\mathbf{J}_{g_2} = \begin{bmatrix} * & * & * & 0 & 0 & 0 & * \\ * & * & * & 0 & 0 & 0 & * \\ * & * & * & 0 & 0 & 0 & * \\ * & 0 & 0 & 0 & 0 & 0 & 0 \\ * & * & * & * & 0 & 0 & * \\ 0 & * & * & * & 0 & 0 & 0 \\ 0 & 0 & * & * & 0 & 0 & 0 \\ 0 & 0 & 0 & * & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & * & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & * & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & * \end{bmatrix} \begin{matrix} \text{return air } y_{T_r} \\ \text{discharge } y_{T_{d,1}} \\ \text{discharge } y_{T_{d,2}} \\ \text{load meter } y_Q \\ \text{rack-1 inlet } y_{T_1} \\ \text{baseline row 1}, \\ \text{baseline row 2} \\ \text{baseline row 3} \\ \text{baseline row 4} \\ \text{baseline row 5} \\ \text{prior on } \delta_2 \end{matrix} \tag{10.5}$$

$$\mathbf{g}_{g_2}^\mathsf{T} = (0,\ 0,\ 1,\ 0,\ 0,\ 0,\ 1),$$

with columns $Q(t), g_1, g_2^\mathrm{b}, \beta_1, \beta_2, \beta_3, \delta_2$; the second line, the gradient of the quantity of equation (10.4), is not a row. Every other drawing swaps the last column. On the five reading rows, in the order of the pattern, the added columns are: g_u, Q_o, T_s, $L_{d,u}$ and L_c, $(*, *, *, 0, *)$; $\ln C_1$, $(0, *, 0, 0, *)$; $\ln C_2$, $(0, 0, *, 0, *)$; β_1, $(0, 0, 0, 0, *)$; and a bias, a single nonzero in its own row. Three facts are read here before any number. On the reading rows the column of δ_2 repeats that of g_2^b, so only the baseline row tells a change from the first hour's value. No raw column but a bias's has a nonzero in the meter's row: a change reaches the meter only through the load, which no reference holds. And the columns of b_{T_1} and β_1 have the same single nonzero, so those two are non-isolable from the pattern alone.

10.2 The posterior of the state (step 3)

Each of the questions of Table 10.2 reads the fits of the drawings of one window, and each drawing is one more fit by the normal equations of Chapter 9. Window 24 is shown; every window from 13 to 36 is fitted the same way.

10.2.1 The fit and the evidence of a drawing

The MAP state $\hat{\mathbf{z}}_j$ of M_j minimizes $\mathcal{L}_j$ of equation (10.3) by

$$(\mathbf{J}_j^\mathsf{T} \mathbf{J}_j)\, \Delta\mathbf{z} = -\mathbf{J}_j^\mathsf{T} \mathbf{r}_j, \qquad \boldsymbol{\Sigma}_j = (\mathbf{J}_j^\mathsf{T} \mathbf{J}_j)^{-1}\big|_{\hat{\mathbf{z}}_j} = \boldsymbol{\Lambda}_j^{-1}, \tag{10.6}$$

equation (9.7) on its own rows, started from the fit of M_0 with $\theta_j = 0$. Its evidence is the Laplace evidence of Chapter 7, equation (7.16), the choice named there and carried here,

$$\ln p(\mathbf{y} \mid M_j) \approx -\mathcal{L}_j(\hat{\mathbf{z}}_j) + \text{const} - \tfrac{1}{2} \ln \det \boldsymbol{\Lambda}_j - \ln w_j, \tag{10.7}$$

the constant the same for all the drawings, the row constants, the baseline's normalization and the $\sqrt{2\pi}$ of the added row against the added dimension. Against M_0 the constant cancels, and equation (7.17) with w_j for σ_b gives the gain in fit less the Occam factor,

$$\ln B_{j0} = \underbrace{\mathcal{L}_0(\hat{\mathbf{z}}_0) - \mathcal{L}_j(\hat{\mathbf{z}}_j)}_{\text{gain}} - \underbrace{\left[\tfrac{1}{2} \ln \frac{\det \boldsymbol{\Lambda}_j}{\det \boldsymbol{\Lambda}_0} + \ln w_j \right]}_{\text{Occam}}. \tag{10.8}$$

The signature of a candidate. The fit of M_0 carries the signatures. With $\mathbf{H}_0 = \mathbf{J}_0 \boldsymbol{\Sigma}_0 \mathbf{J}_0^{\mathsf{T}}$ the hat matrix of Section 9.2.2 and $\mathbf{c}_j$ the column the candidate would add, in the ten rows of M_0 and at $\theta_j = 0$, the part of that column the fit cannot absorb is

$$\mathbf{v}_j = (\mathbf{I} - \mathbf{H}_0)\, \mathbf{c}_j, \qquad s_j = \mathbf{v}_j^{\mathsf{T}} \mathbf{v}_j, \qquad z_j = \frac{\mathbf{v}_j^{\mathsf{T}} \mathbf{r}(\hat{\mathbf{z}}_0)}{\sqrt{s_j}}, \qquad \psi_j = w_j^2\, s_j, \tag{10.9}$$

the projected column. On linear rows the evidence reads off it as in equation (7.18), $\ln B_{j0} \approx \tfrac{1}{2}[\psi_j z_j^2/(1+\psi_j) - \ln(1+\psi_j)]$; z_j is the score of the residual on the signature, and the cosine of two signatures, $\mathbf{v}_j^{\mathsf{T}} \mathbf{v}_l / \sqrt{s_j s_l}$, is the angle the ranking has to work with.

10.2.2 The fits of window 24

No change. M_0 is Chapter 9's fit, Section 9.2.4: $2\mathcal{L}_0 = 34.5727$, at $Q(24) = 206.553\,\text{kW}$ and $g_2 = -0.12780$.

The iterations of the coil's drawing. Table 10.3 runs the normal equations (10.6) for M_{g_2} from the fit of M_0 with $\delta_2 = 0$: each row an iteration, the columns the objective and the norm of its gradient. The objective starts at Chapter 9's 17.2863 and falls to 0.2505 in one step.

Table 10.3: Gauss–Newton iterations of the drawing M_{g_2} at window 24, equation (10.6).

Iteration k [–]	Objective $\mathcal{L}_{g_2}$ [–]	$\lVert \mathbf{J}^{\mathsf{T}} \mathbf{r} \rVert$ [–]
0	17.2863	2.7×10^{2}
1	0.2505	2.7×10^{1}
2	0.2082	3.0×10^{-2}
3	0.2082	9.4×10^{-6}
4	0.2082	1.6×10^{-8}
5	0.2082	5.3×10^{-9}

The Jacobian at the fit. The reading rows of $\mathbf{J}_{g_2}$, columns as in the pattern (10.5), are

$$\begin{bmatrix} 0.2759 & -31.1378 & -25.7251 & 0 & 0 & 0 & -25.7251 \\ 0.0934 & -48.1542 & -8.7087 & 0 & 0 & 0 & -8.7087 \\ 0.1251 & -14.1215 & -42.7414 & 0 & 0 & 0 & -42.7414 \\ 0.3333 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.1448 & -31.1378 & -25.7251 & 34.344 & 0 & 0 & -25.7251 \end{bmatrix},$$

the first column in kW^{-1} and the others dimensionless, the last column repeating the third as the pattern requires.

The MAP state and the covariance. M_{g_2} ends at $2\mathcal{L} = 0.4163$ with, the spreads in parentheses,

$$Q(24) = 198.064\ (2.654)\,\text{kW}, \quad g_1 = -0.05661\ (0.00661), \quad g_2^{\text{b}} = -0.11743\ (0.00721),$$

$$\beta_1 = 0.21295\ (0.00933), \quad \delta_2 = -0.13014\ (0.02326),$$

with β_2 and β_3 at the survey's values and spreads. The covariance of the observed entries, in the order $(Q, g_1, g_2^{\text{b}}, \beta_1, \delta_2)$, and its correlations are

$$\boldsymbol{\Sigma}_{g_2} = \begin{bmatrix} 7.0423034 & 0.0016601 & -0.0007668 & 0.0003766 & 0.0365505 \\ 0.0016601 & 0.0000437 & -0.0000191 & 0.0000147 & -0.0000037 \\ -0.0007668 & -0.0000191 & 0.000052 & 0.000015 & -0.0000364 \\ 0.0003766 & 0.0000147 & 0.000015 & 0.000087 & -0.0000012 \\ 0.0365505 & -0.0000037 & -0.0000364 & -0.0000012 & 0.0005408 \end{bmatrix},$$

$$\begin{bmatrix} 1 & 0.0946 & -0.0401 & 0.0152 & 0.5923 \\ 0.0946 & 1 & -0.4 & 0.2392 & -0.0241 \\ -0.0401 & -0.4 & 1 & 0.2224 & -0.2171 \\ 0.0152 & 0.2392 & 0.2224 & 1 & -0.0053 \\ 0.5923 & -0.0241 & -0.2171 & -0.0053 & 1 \end{bmatrix},$$

in kW^2, kW and dimensionless units. Two signs carry the drawing. The correlation of the load with the change is positive, 0.5923: the temperatures hold $(Q + Q_o)/(G_1 + G_2)$, so a larger load is met by a larger conductance, a smaller fall. The correlation of the baseline's entry with the change is negative, −0.2171: the temperatures hold their sum, so a higher baseline value is met by a larger fall.

The load comes back. Under no change the fit raised the load to 206.553 kW to explain the warm hall, and the meter took the smeared residual of Table 9.4, whitened +2.701 and normalized +4.017. With the coil's entry free the load returns to the meter's 198.45 kW, and the five whitened reading residuals are 0.3629, −0.4169, −0.1452, −0.1287 and −0.0007; the meter's is −0.1287, normalized −0.276. The smeared residual is gone.

The other drawings. Table 10.4 gives each drawing's fit: each row is a drawing, the columns twice its minimized objective, its Laplace evidence (10.7), the value of its added entry and that entry's spread. The rows are in the order of the evidence, within the family and then within rule 4. $M_{L_{d,2}}$ leaves g_2 at −0.11792 (0.00720), the baseline's, since the heat term has taken the change. The drawings M_{b_Q} and M_{Q_o} are one fit, by the substitution of Section 7.2.3: their evidences are equal, −10.798041, and their entries opposite.

10.3 The answers (step 4)

Every answer is read from the fits of Section 10.2.2 and of the same fits at the other windows, and that value is the one reported. Importance sampling of the same rows and widths stands beside each evidence, as in Section 7.4.2, 200 000 draws per drawing per window, with the gap taken like with like.

10.3.1 The ranking at window 24

Operator question. If something changed, where? The operator asks it at 00:55, of the family and of no change.

Table 10.4: The fits of the drawings at window 24 with their Laplace evidence, equation (10.7). The added entry is θ_j, in the units of Table 10.1.

Drawing [–]	$2\mathcal{L}_j$ [–]	$\ln p(\mathbf{y} \mid M_j)$ [–]	$\hat{\theta}_j$ [entry's]	sd [entry's]
M_{g_2}	0.4163	−2.6276	−0.1301	0.0233
$M_{\ln C_2}$	6.4626	−5.2216	0.1684	0.0345
M_{T_s}	9.2711	−5.5711	0.9882 K	0.1958 K
$M_{b_{T_{d,2}}}$	12.9076	−7.6177	1.4643 K	0.3145 K
M_{b_Q}	21.1560	−10.7980	−14.9269 kW	4.0885 kW
M_{Q_o}	21.1560	−10.7980	14.9269 kW	4.0885 kW
M_{g_1}	24.8669	−14.9475	−0.0636	0.0208
$M_{b_{T_1}}$	29.9941	−16.1400	0.6889 K	0.3218 K
M_{β_1}	29.9798	−16.1957	0.0648	0.0304
M_0	34.5727	−16.6043	–	–
$M_{b_{T_r}}$	34.1905	−18.1273	0.2226 K	0.3598 K
$M_{b_{T_{d,1}}}$	34.0208	−18.1816	0.2327 K	0.3133 K
$M_{\ln C_1}$	32.3407	−18.3548	0.0419	0.0284
$M_{L_{d,2}}$	1.9548	−1.8200	12.3723 kW	2.1700 kW
M_{L_c}	19.7628	−10.3024	12.7546 kW	3.3192 kW
$M_{L_{d,1}}$	25.3117	−13.4957	6.6260 kW	2.1794 kW

General form. Diagnosis, equation (2.15),

$$\Pr(M_j \mid \mathbf{y}) \propto p(\mathbf{y} \mid M_j)\ \Pr(M_j).$$

Specialized form. The question is equation (2.15) with the candidates M_j the twelve of Table 10.1 and M_0, $\mathbf{y}$ the evidence $\mathbf{y}_{24}$ of the rows (10.3), $p(\mathbf{y} \mid M_j)$ the Laplace evidence (10.7) and $\Pr(M_j)$ the prior (10.2). Divided through by $p(\mathbf{y} \mid M_0)$,

$$\Pr(M_j \mid \mathbf{y}_t) = \frac{B_{j0}\,(0.5/J)}{0.5 + \sum_l B_{l0}\,(0.5/J)}, \tag{10.10}$$

with the sum over the family.

Building blocks. The evidences and the terms of equation (10.8) from the fits of Section 10.2.2.

The ranking. Table 10.5 gives the leading drawings: each row is a drawing, the columns the gain and the Occam factor of equation (10.8), the log Bayes factor against M_0, the probability (10.10), and the importance-sampled evidence with its gap from the Laplace value in Table 10.4. For the coil, the gain of 17.0782 less the Occam factor of 3.1015 gives 13.9766.

The leader is the drawing of unit 2's coil, and it is not decisive. Against its runner-up, the airflow of the same unit, $\ln B = 2.594$, short of 2.9957, and its probability is 0.881, short of 0.95. The coil pays more for its entry than any close rival, 3.1015 against the airflow's 2.6724, and buys it back with a larger gain. Importance sampling gives $\ln B = 2.593$ and the probability 0.88122; the widest gap of a log evidence in this window is −0.0074, on M_0. Neither moves the verdict.

10.3.2 The coil against the sensor

Operator question. A coil or a sensor? This is the pair Chapter 9 handed over, Section 9.6.1, and the one behind the dashboard's entry.

Table 10.5: The ranking at window 24, 00:55, in the family of 12. The other six candidates each hold a probability below 0.00001.

Drawing [–]	Gain [–]	Occam [–]	$\ln B_{j0}$ [–]	$\Pr(M_j \mid \mathbf{y}_{24})$ [–]	IS $\ln p(\mathbf{y} \mid M_j)$ [–]	Gap [–]
M_{g_2}	17.0782	3.1015	13.977	0.88122	−2.6275	+0.0001
$M_{\ln C_2}$	14.0550	2.6724	11.383	0.06585	−5.2209	+0.0008
M_{T_s}	12.6508	1.6176	11.033	0.04643	−5.5714	−0.0003
$M_{b_{T_{d,2}}}$	10.8325	1.8460	8.987	0.00600	−7.6218	−0.0041
M_{b_Q}, M_{Q_o}	6.7083	0.9021	5.806	0.00025 each	−10.8032	−0.0052
M_0	–	–	0	0.00001	−16.6117	−0.0074

General form. The comparison of two drawings within the diagnosis (2.15) is the Bayes factor, equation (2.14),

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

Specialized form. The question is equation (2.14) with $M_A = M_{g_2}$, $M_B = M_{b_{T_{d,2}}}$ and $\mathbf{y}$ the evidence $\mathbf{y}_t$ of each window, each evidence by equation (10.7), so that $\ln B = \ln p(\mathbf{y} \mid M_{g_2}) - \ln p(\mathbf{y} \mid M_{b_{T_{d,2}}})$. Its reason is read from the signatures (10.9), which differ where the rows differ.

Building blocks. The two evidences per window; the signatures of the two candidates at the fit of M_0.

The signatures. Figure 10.2 draws the projected columns of the four leading candidates, one group of bars per reading: the vertical axis the entry of $\mathbf{v}_j$ in that reading's row, each column scaled to unit length over its ten rows so that the four share an axis. The black bar in each group is the residual vector of M_0, scaled the same way, which the signatures are trying to match. The sign of a whole column is arbitrary, since θ_j can take either sign, and what matters is the pattern. Before scaling, the coil's column is $(-7.020,\ -2.353,\ -34.326,\ 22.639,\ -15.934)$ and the sensor's $(-0.453,\ -0.087,\ 2.953,\ -0.438,\ -0.186)$; the scores (10.9) are $z_j = 5.850$ for the coil, −5.300 for the airflow, −5.130 for the supply and −4.713 for the sensor, the last being Chapter 9's normalized residual of the discharge.

The answer. $\ln B(g_2 : b_{T_{d,2}}) = 4.990$ at window 24, and importance sampling gives 4.9943, a gap of +0.0042. It first passes $\ln 20$ at window 21, 00:40, at 4.156, falls under it again at window 23 at 2.467, and is past it in every window from window 24. The dashboard's reading is overturned here, for the reason the figure shows. The sensor's column has a single entry before the projection, so a biased discharge sensor would have left the return air and rack 1's inlet where the meter's load puts them, and they rose, by 0.92 and 1.06 K over their first-hour means, while the meter, up 2.18 kW, stayed within its first-hour scatter of 4.05 kW. The meter's +2.70 in Table 9.4 is a residual because M_0 raised the load, not a movement of the meter. Turned to match the residual on the discharge row, the sensor's scaled signature carries 0.14 on the meter's row and 0.06 on rack 1's, where the residual has 0.46 and −0.34 and the coil's signature 0.48 and −0.34: it is mainly the size of those entries that separates the two. The two are not orthogonal, a cosine of −0.7809, and that is why a window can still fall short, as window 23 does. This comparison is one line of the ranking and does not settle the ranking: clearing the sensor leaves the airflow and the supply, at cosines of −0.8827 and −0.8557 with the coil.

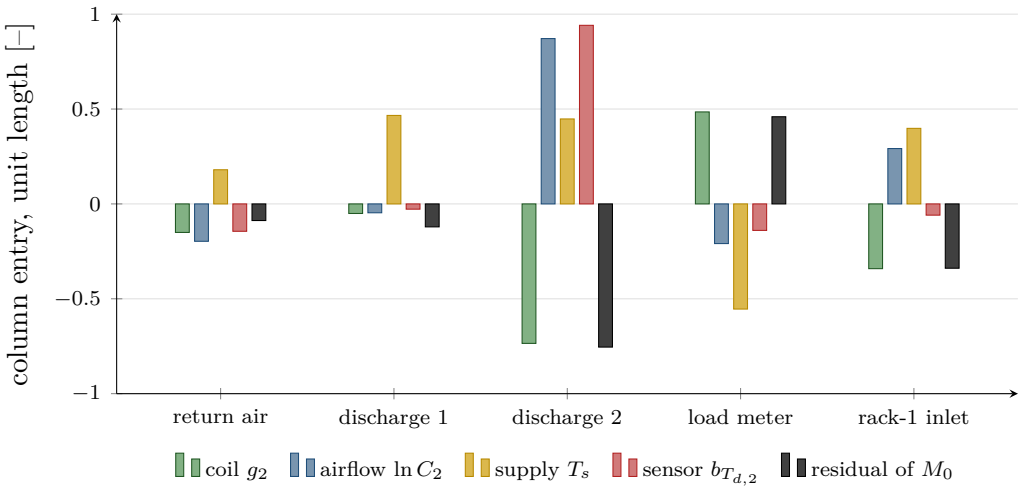


Figure 10.2: The signatures of the four leading candidates at window 24: the reading rows of the projected columns (10.9) at the fit of M_0, each scaled to unit length over its ten rows, beside the residual vector of M_0 scaled the same way (black). The coil, the airflow and the supply each move several rows; the sensor's column has nearly all its length in one.

10.3.3 The ranking across the night

Operator question. When can the ranking be acted on? The operator wants the first window at which the rule of Section 10.1.3 is met, and the first after which it is never failed.

General form. Diagnosis, equation (2.15), window by window, with the rule as the event asked of it.

Specialized form. The question is equation (2.15) with $\mathbf{y}$ the evidence $\mathbf{y}_t$ of window t, the candidates and priors of Section 10.3.1 and $p(\mathbf{y} \mid M_j)$ the evidence (10.7) in that window, for $t = 13, \ldots, 36$. The leader of window t is the most probable candidate; the runner-up is the strongest rival but the leader's partner; and the answer is the two windows

$$t_{\text{first}} = \min\{t : \text{rule met at } t\}, \qquad t_{\text{held}} = \min\{t : \text{rule met at every } t' \geq t\}. \tag{10.11}$$

of which the first is known when it happens and the second only from the finished record: at any window the next may fail, and at the record's end it is censored.

Building blocks. The evidences of all the drawings in each window; the rule; importance sampling per window.

The night. Table 10.6 gives the ranking: each row is a window, the columns its clock time, the leader and its runner-up, the log Bayes factor between them and its importance-sampling gap, the leader's probability by each route, the probability of no change, and the verdict. A leader's probability includes its partner's. In windows 13 to 16 no change holds more than any candidate.

Figure 10.3 draws the probabilities of the leading drawings through the night: the horizontal axis the window, the vertical axis the probability, one line per drawing, with the windows of Chapter 9's first detections marked by dashed lines and the decisive windows by triangles along the top. Figure 10.4 draws the log Bayes factor of the leader against the runner-up per window, with the level $\ln 20$; filled points are decisive windows.

Table 10.6: The ranking of the family through the night, one window at a time. Bold: the windows of Chapter 9's first detections, the worked window and the first decisive window. IS gap: importance sampling minus Laplace for the same pair of drawings.

t [–]	Start [h:min]	Leader [–]	Runner-up [–]	$\ln B$ [–]	IS gap [–]	Pr leader [–]	IS [–]	$\Pr(M_0)$ [–]	Decisive [–]
13	00:00	b_Q	$b_{T_{d,1}}$	0.361	+0.0009	0.05755	0.05756	0.83657	no
14	00:05	b_{T_r}	Q_o	0.368	+0.0006	0.04056	0.04054	0.79356	no
15	00:10	T_s	Q_o	0.061	+0.0016	0.04964	0.04973	0.65834	no
16	00:15	Q_o	T_s	0.030	−0.0016	0.08252	0.08254	0.76413	no
17	00:20	b_Q	T_s	0.162	−0.0019	0.54758	0.54751	0.12848	no
18	00:25	$b_{T_{d,2}}$	$\ln C_2$	1.399	+0.0048	0.50952	0.51021	0.27476	no
19	00:30	T_s	b_Q	1.448	+0.0030	0.32819	0.32884	0.32183	no
20	00:35	g_2	b_Q	1.733	+0.0027	0.55623	0.55676	0.06170	no
21	00:40	g_2	$\ln C_2$	1.757	−0.0065	0.74705	0.74659	0.00366	no
22	00:45	$\ln C_2$	g_2	0.333	+0.0117	0.56094	0.56386	0.00015	no
23	00:50	g_2	$\ln C_2$	1.080	−0.0024	0.66231	0.66221	0.01200	no
24	00:55	g_2	$\ln C_2$	2.594	−0.0007	0.88122	0.88122	0.00001	no
25	01:00	T_s	g_2	0.438	+0.0005	0.59374	0.59392	0.00000	no
26	01:05	g_2	T_s	4.311	+0.0005	0.97816	0.97812	0.00000	**yes**
27	01:10	g_2	T_s	2.050	−0.0001	0.80067	0.80046	0.00000	no
28	01:15	g_2	$\ln C_2$	2.409	−0.0079	0.91742	0.91681	0.00000	no
29	01:20	g_2	$\ln C_2$	2.453	−0.0036	0.88197	0.88174	0.00000	no
30	01:25	g_2	$\ln C_2$	4.481	−0.0067	0.98879	0.98872	0.00000	yes
31	01:30	g_2	$\ln C_2$	5.327	−0.0044	0.99515	0.99513	0.00000	yes
32	01:35	g_2	$\ln C_2$	0.032	+0.0045	0.50151	0.50264	0.00000	no
33	01:40	g_2	$\ln C_2$	8.553	−0.0032	0.99974	0.99974	0.00000	yes
34	01:45	g_2	$\ln C_2$	15.469	+0.0080	1.00000	1.00000	0.00000	yes
35	01:50	g_2	$\ln C_2$	6.554	+0.0121	0.99857	0.99859	0.00000	yes
36	01:55	g_2	T_s	10.283	−0.0002	0.99996	0.99996	0.00000	yes

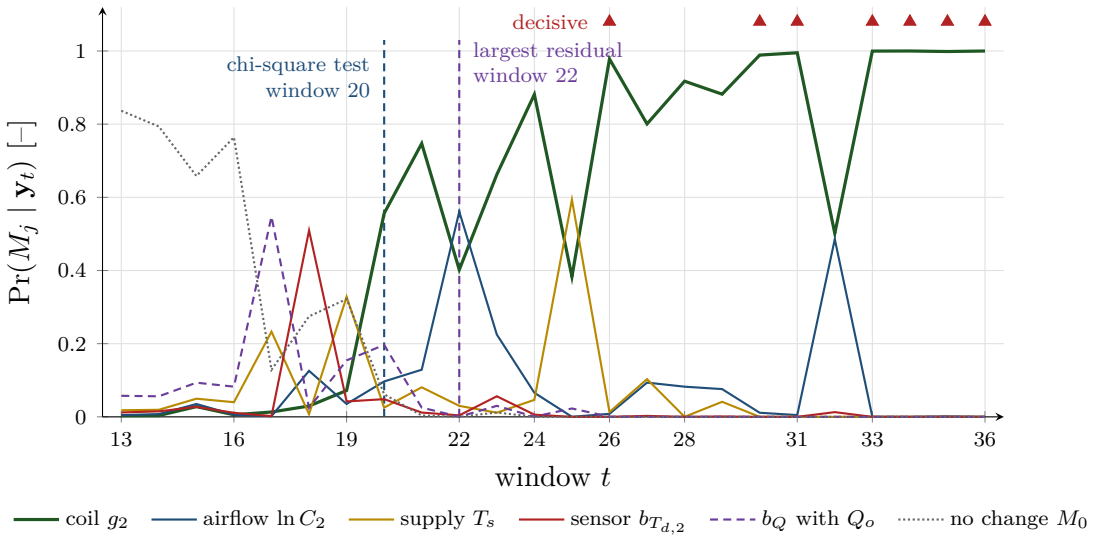


Figure 10.3: The probabilities of the leading drawings per window, each window ranked alone. Dashed: the first detections of Chapter 9 at $\alpha = 10^{-3}$. Triangles: the windows past the rule of Section 10.1.3.

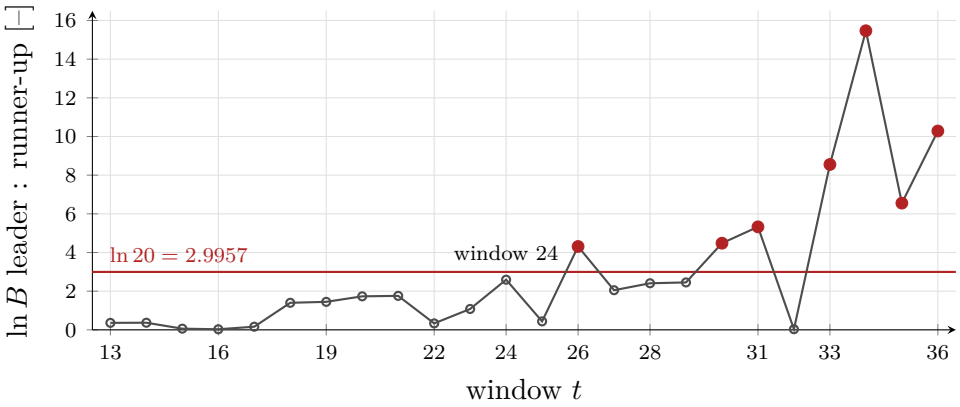


Figure 10.4: The log Bayes factor of each window's leader against its runner-up (Table 10.6), against the level of the rule. Filled: decisive windows, which also meet the rule's other two parts.

The answer. By equation (10.11), $t_{\text{first}} = 26$, 01:05, thirty minutes after the first detection of Chapter 9, and $t_{\text{held}} = 33$, 01:40. Importance sampling gives the same verdict in all 24 windows: no gap crosses the level, the widest gap of the leader-to-runner-up $\ln B$ being $+0.0121$, at window 35, and the widest of a log evidence -0.0155, on M_0 at window 36. These two windows are one night's draw of the sensor errors: over 200 nights of the same truth and baseline with fresh errors, t_{first} has a median of 26 (bootstrap standard error 0.19) and a 5 to 95 % range of 22 to 29, and t_{held} a median of 32 (0.49) and a range of 27 to 36, with 3 % of nights not held at the record's end.

The verdict is not monotone, because each window is ranked alone and each carries its own sensor errors. The detection windows do not isolate anything. At window 20 the coil leads at 0.556 with no change still at 0.062. At window 22 the leader is the airflow, 0.561, the coil at 0.402. At window 25 the leader is the supply, 0.594, the coil at 0.383. And after the first decisive window the rule fails again in windows 27 to 29 and at window 32, where the coil holds 0.502 with $\ln B = 0.032$. No change falls below 0.00001 by window 25, before the candidates come apart: detection and isolation are different events.

10.3.4 What the readings cannot separate

Operator question. Can these readings say it is the coil? The ranking's leader is a drawing of the family; the question is whether an excluded drawing would explain the same readings.

General form. Diagnosis, equation (2.15), over the extended family, and the Bayes factor of a pair, equation (2.14).

Specialized form. The question is equation (2.15) with the candidates M_j the fifteen of the extended family, $\mathbf{y}$ the evidence $\mathbf{y}_t$, $p(\mathbf{y} \mid M_j)$ the evidence (10.7) and the prior (10.2) at $J = 15$, and equation (2.14) with $M_A = M_{g_2}$ and $M_B = M_{L_{d,2}}$. The pair turns on a substitution in the chain (10.1). Write $G_2^{\text{b}} = \bar{G}_2 \mathrm{e}^{g_2^{\text{b}}}$ for the conductance before the change, $G_2 = \bar{G}_2 \mathrm{e}^{g_2^{\text{b}} + \delta_2}$ after it, and $D = T_r - T_s$ under M_{g_2}. Take the same Q and all other entries and set

$$L_{d,2} = (G_2^{\text{b}} - G_2)\, D. \tag{10.12}$$

Then $Q + Q_o + L_{d,2} = (G_1 + G_2)D + (G_2^{\text{b}} - G_2)D = (G_1 + G_2^{\text{b}})D$, so the hall balance under $M_{L_{d,2}}$ returns the same D and the same T_r; and $T_{d,2} = T_r - (G_2^{\text{b}} D - L_{d,2})/C_2 = T_r - G_2 D/C_2$, the same discharge, with $T_{d,1}$ and through them T_c and the rack inlets the same. Every reading

is the same in every window, whatever the noise. The two drawings are non-isolable on these readings: a coil passing less heat and the same heat entering unit 2's discharge air are the same five numbers. With the fit of the readings shared, their Bayes factor can only compare the prior densities at the two values the fit implies, which are the widths and the size of the implied change, and no reading separates them.

Building blocks. The fit of M_{g_2} for D, G_2^{b} and G_2; the evidences of the family and of the three drawings of rule 4; the cosines of the signatures.

The numbers. At the fit of window 24, $D = 17.0589\,\mathrm{K}$, $G_2^{\mathrm{b}} = 6.2244$ and $G_2 = 5.4649\,\mathrm{kW/K}$, so equation (10.12) gives $L_{d,2} = 12.957\,\mathrm{kW}$. The largest difference between the two reading vectors prints as 0, and the projected columns have the cosine -1.000000. The substitution with unit 1's entries does the same for g_1 and $L_{d,1}$. Together with the partners of the rule these are the non-isolable pairs of the hall: b_Q with Q_o, with equal evidences; b_{T_1} with β_1, cosine $+1.000000$; g_2 with $L_{d,2}$; and g_1 with $L_{d,1}$.

In the extended family $L_{d,2}$ leads at window 24, with probability 0.664 against the coil's 0.296, $\ln B(g_2 : L_{d,2}) = -0.808$, and importance sampling gives -0.8070, a gap of $+0.0007$. That order is the widths and the size of the implied change, not a reading: $2\mathcal{L}$ is 0.4163 against 1.9548 because of the two prior rows, the width of $L_{d,2}$ that matches the width of δ_2 is 49.8 kW, and at five times the 10 kW of Table 10.1 the difference is $+0.009$. Over the night it runs from -1.576 to $+1.305$ as the implied heat grows against a fixed width.

What the pair holds together depends on that same width. With $L_{d,2}$ at 10 kW the pair holds 0.9600 at window 24 and $\ln B(L_{d,2} : \ln C_2) = 3.402$, past $\ln 20$; at 49.8 kW it holds 0.9365 and its best member's $\ln B$ is 2.594, short of it. Both are stated, and neither is chosen. The next window the pair holds 0.625 at 10 kW, so it is no more steady than the leader. The rule is fixed for a leader of the family and is not extended to the pair.

What the record supports is a coil or heat in its discharge air, which the readings separate only from the other candidates, firmly only from window 33 (01:40). The name "coil" rests on the choice to exclude rule 4. The one quantity the substitution does not preserve is the coil's own heat, $Q_2 = G_2 D = 93.22\,\mathrm{kW}$ under M_{g_2} against $G_2^{\mathrm{b}} D = 106.18\,\mathrm{kW}$ under $M_{L_{d,2}}$ at the substitution: it is the only thing a further reading can use, and no reading of this record has a row on it.

10.3.5 The derate factor under the leader

Operator question. Given it is unit B, how far is it from its rating? In this hall unit B is unit 2. This is fault identification, and the answer is conditional on the drawing M_{g_2}, whose standing is the four subsections above.

General form. Estimation of a derived quantity, equation (2.9),

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

summarized by its central 95 % interval, equation (2.11).

Specialized form. The question is equation (2.9) with $\mathbf{z}$ the state of M_{g_2}, $f(\mathbf{z}) = d_2 = \mathrm{e}^{g_2}$ by equation (10.4), f^* a level of d_2, the ends of the interval of equation (2.11) its 2.5 and 97.5 % points, $\mathbf{y}$ the evidence $\mathbf{y}_t$ with the drawing M_{g_2} conditioned on, and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal $\mathcal{N}(\hat{\mathbf{z}}_{g_2}, \boldsymbol{\Sigma}_{g_2})$. With the gradient of line two of the pattern (10.5), the delta method (4.53) gives

$$\mathrm{Var}(g_2) = \Sigma_{g_2^{\mathrm{b}} g_2^{\mathrm{b}}} + \Sigma_{\delta_2 \delta_2} + 2\,\Sigma_{g_2^{\mathrm{b}} \delta_2}, \qquad \mathrm{sd}(d_2) = d_2\,\mathrm{sd}(g_2), \tag{10.13}$$

and the cross term is negative, the correlation of Section 10.2.2: a higher baseline value is met by a larger change, so the sum is known better than its parts added. The change alone is the factor e^{δ_2} against the first hour.

Building blocks. $\hat{\mathbf{z}}_{g_2}$ and $\mathbf{\Sigma}_{g_2}$ of Section 10.2.2; the gradient $\mathbf{g}_{g_2}$; the budget (10.13).

The derate of window 24. $g_2 = -0.11743 - 0.13014 = -0.24757$, with variance $5.2033 \times 10^{-5} + 5.4083 \times 10^{-4} + 2(-3.6412 \times 10^{-5}) = 5.2004 \times 10^{-4}$, standard deviation 0.02280 and band $[-0.29227,\ -0.20288]$. Then $d_2 = 0.7807$, with delta standard deviation 0.01780 and delta band $[0.7458,\ 0.8156]$, and $G_2 = 5.465 \pm 0.125$ kW/K, 1.54 kW/K short of the 7.0 rating. Importance sampling beside it, 2 000 000 draws of the exact posterior of the drawing, 1 997 700 effective: median 0.78049, a gap of -0.00020; standard deviation 0.01780, the gap below the print; band $[0.74574,\ 0.81553]$, gaps of -0.00006 and -0.00005 from the delta band.

The 21.9 % below the rating is not all of tonight. The baseline was already 11.1 % below, $d_2 = 0.88902$, and the change since the first hour is $e^{\delta_2} = 0.878$, band $[0.839,\ 0.919]$. At the other windows, with $\exp(g_2 \pm 1.960\,\mathrm{sd})$ as the band, 0.7466 to 0.8164 at window 24, where the exact band's gaps from it are -0.00082 and -0.00085: at window 20, $d_2 = 0.8136$, band $[0.7789,\ 0.8499]$; window 22, 0.7915, $[0.7567,\ 0.8279]$; window 26, 0.7739, $[0.7408,\ 0.8084]$; window 36, 0.6968, $[0.6657,\ 0.7294]$. These are window values under a drawing that was not the leader in every one of those windows, and the condition travels with them. The alternative at window 24: a drawing that leaves the coil alone leaves d_2 near the baseline, 0.8888 (sd 0.0064) under $L_{d,2}$, and 0.8884, 0.8845 and 0.8878 under $\ln C_2$, T_s and $b_{T_{d,2}}$.

10.4 The checks (step 5)

A ranking and a size are computed under approximations and choices, and each is set here against the calculation it stands for.

10.4.1 The limit of exact rows and the curvature term

The fit runs on the reduced state. The check fits the drawing's own full state of 14 entries, the thirteen of equation (5.2) and δ_2, with the seven physical rows (5.1) at a finite tolerance. The spread of g_2 is 0.02293, 0.02281 and 0.02280 at $\sigma_p = 0.1$, 0.01 and 0.001, against the reduced state's 0.02280; Q and g_2 hold at 198.0638 kW and -0.24757, and the spread of Q goes from 2.6540 to 2.6537 kW, the reduced value. The full Hessian of $\mathcal{L}_{g_2}$ at the fit gives $\mathrm{sd}(g_2) = 0.02280$, the gap below the print, $\mathrm{sd}(Q) = 2.6534$ against 2.6537 kW, and with the Newton determinant $\ln p(\mathbf{y} \mid M_{g_2}) = -2.6279$ against -2.6276, a gap of -0.0003.

10.4.2 The Laplace evidence and the linear form

Table 10.7 collects the gaps. The importance-sampling gaps are about ten times larger than those of Section 7.4.2, -0.0074 on M_0 at window 24 and -0.0145 within the family over the night, with relative standard errors of 0.00002 to 0.00016 at window 24, and they cross no verdict. The linear form of equation (10.9) is a reading of the evidence, not a replacement for it: its widest gap, -0.2254 for $\ln C_2$, moves the runner-up by more than the importance-sampling gap and in the leader's direction.

The mean of g_2 under the exact posterior sits 0.00049 from the MAP value with a standard error of 0.000016, a resolved difference and a small one: the delta method's median and the exact median of d_2 differ in the fourth decimal.

Table 10.7: The checks of the answers and their gaps from the reported values. IS: importance sampling, 200 000 draws per drawing for the evidences and 2 000 000 for the derate. The truth is the stated truth of the constructed record.

Check [–]	Reported [–]	Beside it [–]	Gap [–]
$\ln p(\mathbf{y} \mid M_{g_2})$, window 24, IS	−2.6276	−2.6275	+0.0001
$\ln p(\mathbf{y} \mid M_0)$, window 24, IS	−16.6043	−16.6117	−0.0074
$\ln B$ leader : runner-up, window 24, IS	2.594	2.593	−0.0007
$\ln B(g_2 : b_{T_{d,2}})$, window 24, IS	4.9901	4.9943	+0.0042
$\ln B(g_2 : L_{d,2})$, window 24, IS	−0.8077	−0.8070	+0.0007
The same, widest of the night (window 35), IS	6.554	6.566	+0.0121
$\ln B(g_2 : M_0)$, window 24, linear form	13.9766	13.9293	−0.0474
$\ln B(\ln C_2 : M_0)$, the widest linear gap	11.3826	11.1572	−0.2254
d_2 median, M_{g_2}, IS	0.78069	0.78049	−0.00020
g_2 mean, IS; Monte Carlo se 0.000016	−0.24757	−0.24806	−0.00049
sd g_2, Newton against Gauss–Newton	0.02280	0.02280	below the print
$\ln p(\mathbf{y} \mid M_{g_2})$, Newton determinant	−2.6276	−2.6279	−0.0003
sd g_2, full state, $\sigma_p = 0.001$	0.02280	0.02280	below the print
Truth (constructed), g_2 at window 24	−0.24757	−0.23834	−0.40 sd
Truth, d_2 at window 24; inside the band	0.7807	0.7879	−0.0072

10.4.3 Sensitivity to the prior probabilities and the widths

Table 10.8 repeats the ranking under other prior probabilities: each row is a setting, the columns its leader and that leader's probability at windows 20, 22 and 24, and the first decisive window. The evidences do not change, only what they multiply.

Table 10.8: The ranking under other prior probabilities, with the evidences unchanged.

Setting [–]	Window 20 [–]	Window 22 [–]	Window 24 [–]	First decisive [–]
Stated, equation (10.2)	g_2, 0.556	$\ln C_2$, 0.561	g_2, 0.881	26
Equal over the 13	g_2, 0.590	$\ln C_2$, 0.561	g_2, 0.881	26
$\Pr(M_0) = 0.9$	g_2, 0.372; M_0, 0.372	$\ln C_2$, 0.560	g_2, 0.881	26
Sensors ten times as likely as a plant change	b_Q, 0.370	$\ln C_2$, 0.535	g_2, 0.834	26
Unit 2's coil one tenth as likely as any other	b_Q, 0.199	$\ln C_2$, 0.879	g_2, 0.426	**31**

Window 20's leader is a prior's choice. The first decisive window moves only when the coil itself is marked down. The prior odds against the coil that would bring it level with its best rival are 13.4 at window 24, 74.5 at window 26 and 5.18×10^3 at window 33: an operator who holds a reason to doubt the coil can read off how strong that reason must be.

The widths move less. All widths at half or double, or the width of g_2 alone at half or double, leave the first decisive window at 26 and the held window at 33. Two of these touch window 24 and both are stated. The width of g_2 at half takes $\ln B(g_2 : \ln C_2)$ past $\ln 20$, to 3.18, and the window stays short on probability, at 0.9304; at double it gives 1.93. Halving the sensor's width alone leaves the leader's comparison where it was.

10.4.4 The decision rule on draws

The rule's behavior can be measured on the plant. Table 10.9 ranks 300 fresh windows per scenario, drawn from the baseline at a load of 200 kW with the stated sensor errors and the given change: each row is a scenario, the columns the shares of windows with a decisive and correct isolation, a decisive false isolation and no decisive verdict, and the share in which the

top drawing is the right one, in the first row $\Pr(M_0) \geq 0.5$. The standard error of each share is at most 0.029.

Table 10.9: The decision rule on 300 fresh windows per scenario; standard errors at most 0.029.

Scenario [–]	Correct isolation [–]	False isolation [–]	Not decisive [–]	Top right [–]
No change	0.000	0.000	1.000	0.943
g_2 change -0.06	0.000	0.003	0.997	0.153
g_2 change -0.12	0.160	0.013	0.827	0.697
g_2 change -0.24	0.967	0.000	0.033	0.987
Bias on $T_{d,2}$, +1.6 K	0.313	0.003	0.683	0.893
Other heat +15 kW	0.023	0.010	0.967	0.757

At a change of window 24's size, -0.12, the rule rarely fires, and when it does it is wrong in 4 of 52 windows, 7.7 % against the 5 % its 0.95 would allow; the top drawing is right about two times in three. For other heat of 15 kW it is wrong in 3 of the 10 windows where it fires. Window 24 falling short is the typical outcome, not bad luck. The first hour, a lenient test since it made the baseline, stays quiet: the largest $\ln B$ against M_0 is 0.260, the smallest $\Pr(M_0)$ is 0.7218, and no window is decisive.

The night repeated. The record is one draw of the sensor errors. Over 200 nights of the same truth and baseline with fresh errors, window 24 is decisive in 15 % (standard error 2.5 %), and g_2 tops it in 73 % (3.1 %). There $\ln B(g_2 : b_{T_{d,2}})$ has a median of 4.23, a 5 to 95 % range of -1.45 to 10.14, and is past $\ln 20$ in 65 % (3.4 %). The first decisive leader is not the coil in 9 of the 200 nights (1.5 %), and 23 nights (2.3 %) hold some decisive window naming another candidate. The record's own t_{first} and clearing of the sensor at window 24 are ordinary draws, not properties of the plant.

10.4.5 Adequacy of the leader

Evidence is relative: the leader is the best of those listed, which does not make it adequate. The test is the chi-square of the reading rows (9.14) on the fit of M_{g_2}, with the redundancy numbers of that fit. At window 24, $X^2 = 0.343$ on 2.838, a tail of 0.941, against Chapter 9's tail of 1.6×10^{-6}; no window of 13 to 36 has a tail under 0.01. The lowest is window 32, 0.026, with a normalized residual of 2.775 on unit 1's discharge, the window whose ranking also failed the rule. The departure of Chapter 9 is explained under this drawing; the airflow's also passes, a tail of 0.100, and the supply's stands at 0.049, while the sensor's, at 0.00805, does not reach 0.01. The test catches a cause that leaves misfit, not one confounded with a listed drawing: $M_{L_{d,2}}$ leaves nearly the same residuals, $X^2 = 0.412$ against 0.343, a tail of 0.928, so this test cannot speak to that pair.

10.4.6 The answers against the stated truth

The record is generated from a stated truth, which is a change in unit 2's coil. At window 24 the fit is within its band, the last two rows of Table 10.7. Over windows 20 to 36, $z = (\hat{g}_2 - g_2^{\text{truth}})/\text{sd}$ has a mean of +0.048 (standard error 0.243), with no lag. Its standard deviation over windows 13 to 36 is 0.649, narrow on this record; on the draws of Section 10.4.4 it is 1.019 (standard error 0.041), so the band is not overstated in general.

The one-window quantity. The stated truth moves across the night, so a fit could lag it. The one-window g_2 does not: under a fixed baseline it holds no memory of earlier windows, and it estimates that window's value, which the record holds constant within a window and which

on a plant would be the window's mean. The stacked fits are the ones that lag. One change held over windows 22 to 24 gives $g_2 = -0.23241$, $+0.44$ of its standard deviation from window 24's truth; over 20 to 24, -0.22568, $+1.20$; over 13 to 24, -0.18370, $+7.91$, which is close to the mean of those windows' truth, -0.18410. The chapter therefore reports the one-window value.

10.5 The report (step 6)

The record does not support the sensor: $\ln B = 4.99$ at window 24 (00:55), probability 0.006. Among the twelve changes the graph holds that a reading can see, unit 2's coil leads at 0.88 at 00:55, and it is not decisive, with the airflow at 0.066 and the supply temperature at 0.046 behind it; it is first decisive at 01:05, one night's draw, 22 to 29 in window number over the sensor errors, and from the finished record decisive in every window from 01:40. These readings name a coil or heat in its discharge air, and they cannot tell the two apart. Under the coil's drawing, unit 2 is at 0.78 of its rating; under the heat term, 0.889.

10.5.1 The answers and their bands

Table 10.10 gives the answers in the order of Table 10.2: each row a question, the answer with its importance-sampled value beside it, and what holds it back. Figure 10.5 draws the derate factor per window under the coil's drawing: the horizontal axis the window, the vertical axis d_2, the solid line with its shaded band $\exp(g_2 \pm 1.960\,\mathrm{sd})$ of Section 10.3.5, the dashed line the baseline's value and the dotted line the stated truth, a check marker. The band stays with the truth, and it is one window's width at every window: each window is its own fit.

Table 10.10: The report for the night. Laplace evidence and the delta method; importance sampling (IS) in parentheses.

Question	Answer	What holds it back
If something changed, where?	at 00:55, unit 2's coil, probability 0.881 among the 12 (IS 0.881)	not decisive; $\ln B = 2.594$ against the airflow; the coil's name rests on excluding rule 4
A coil or a sensor?	the record does not support the sensor: $\ln B = 4.990$ at window 24 (IS 4.994)	in fresh sensor errors, past $\ln 20$ at window 24 in 65 % of nights
When can the ranking be acted on?	first at window 26, 01:05 (IS the same); from the finished record, held from window 33, 01:40	one night's draw; 22 to 29 over the sensor errors; moves to window 31 only if the coil is marked down tenfold
Can these readings say it is the coil?	no: a coil or heat in its discharge air; $L_{d,2}$ and g_2 are non-isolable, pair probability 0.960 at window 24 (0.937 at the matched width), $\ln B(g_2 : L_{d,2}) = -0.808$ (IS -0.807)	the order within the pair, and the pair's own figures, depend on the width
How far from its rating?	$d_2 = 0.7807$, sd 0.0178, band 0.7458 to 0.8156 (IS median 0.7805, band 0.7457 to 0.8155); 0.878 of the first hour; under the heat term, 0.889	conditional on the coil's drawing

10.5.2 The entries in the shift log

The first entry keeps the structural limit beside the name, and the second records that the rule is met again.

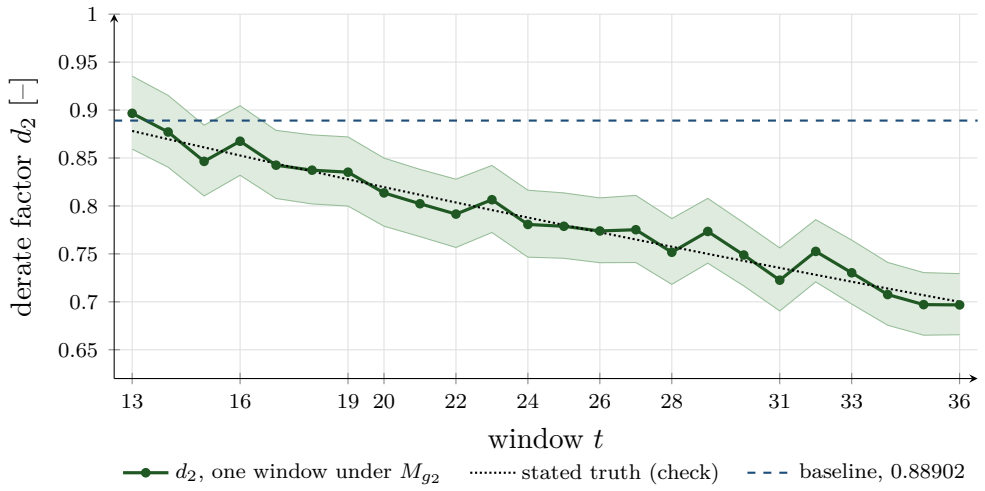


Figure 10.5: The derate factor of unit 2 per window, conditional on the drawing of the coil, with its band. Dotted: the stated truth of the constructed record, a check marker. Dashed: the value at the baseline of the first hour. Each window is fitted alone.

00:55. Unit 2 discharge rise: the record does not support the sensor ($\ln B$ 4.99, probability 0.006). Leading explanation: coil 2 derated, probability 0.88 among the 12 graph candidates, not yet decisive (airflow 0.07, chilled-water supply 0.05). These readings cannot tell the coil from heat entering unit 2's discharge air. If it is the coil, 0.78 of rating (0.75 to 0.82), 0.88 of the first hour. Requested: water-side heat check on unit 2 tonight.

01:40. Coil 2 decisive in the ranking this window, as at 01:05, 01:25 and 01:30; it failed at 01:35. The water-side heat check remains needed to rule out heat entering the discharge air, which these readings cannot.

Common mistake

Reading the probability of the leader as the probability that the coil is at fault. It is the share among the listed drawings. A change the graph has no row for holds none, and adding its row can change the order without the readings changing at all.

10.6 The next reading (step 7)

A further reading is one more row of each drawing.

10.6.1 The expected reduction in entropy

The value of a candidate reading a is the design form, equation (2.16), with $\mathbf{y}$ the evidence $\mathbf{y}_{24}$, the posterior of that form the ranking over the drawings rather than a density on one state, and its quantity the reduction of the ranking's entropy. With $p_j = \Pr(M_j \mid \mathbf{y}_{24})$ the probabilities of the family and μ_j, τ_j the mean and spread of the reading's predictive under drawing j, equation (7.22),

$$U(a) = H(\mathbf{p}) - \mathbb{E}_{\tilde{y}_a}\left[H(\mathbf{p}(\tilde{y}_a))\right], \qquad p_j(\tilde{y}_a) \propto p_j\, \mathcal{N}(\tilde{y}_a;\ \mu_j,\ \tau_j^2), \tag{10.14}$$

with H the entropy in nats and $\tilde{y}_a$ drawn from the mixture of the predictives. It is computed from 400 000 draws per candidate, with standard errors of at most 0.0006. One choice enters: each reading's value is taken before any other is made. The entropy of the ranking now is 0.4681

nats over the family of 12 and 0.7969 over the extended family of 15, and the gains are reported on both, since the family leaves out the rival that matters most to the name.

10.6.2 The candidates

Table 10.11 ranks the candidates: each row is a reading, the columns the quantity it reads, the gain (10.14) over each family, and the probability of the top drawing expected after it. Figure 10.6 draws the predictive of the portable water-side heat check under the five leading drawings: the horizontal axis the reading, the vertical axis its density, one curve per drawing, each including the check's own 2 kW.

Table 10.11: Candidate readings at window 24, by the expected reduction in the entropy of the ranking; standard errors at most 0.0006.

Candidate [–]	Reads [–]	Gain, 12 [nats]	Gain, 15 [nats]	Top after, 12 / 15 [–]
Portable water-side heat check of unit 2, sd 2 kW, one visit	Q_2	**0.2692**	**0.6258**	0.929 / 0.953
Permanent water-side heat meter on unit 2, sd 3 kW	Q_2	0.1997	0.5457	0.911 / 0.932
Airflow traverse of unit 2, sd 0.5 kW/K, one visit	C_2	0.1897	0.0802	0.930 / 0.679
Check of the chilled-water supply, sd 0.1 K, one visit	T_s	0.1871	0.0801	0.927 / 0.680
Clamp-on on the UPS output, sd 1 kW, one visit	Q	0.1025	0.0375	0.891 / 0.665
Probe in unit 2's discharge air, sd 0.3 K, one visit	$T_{d,2}$	0.0865	0.0329	0.886 / 0.665
Permanent second discharge sensor on unit 2, sd 0.3 K	$T_{d,2}$	0.0866	0.0331	0.887 / 0.665
Second return-air sensor, sd 0.3 K	T_r	0.0789	0.0292	0.889 / 0.665

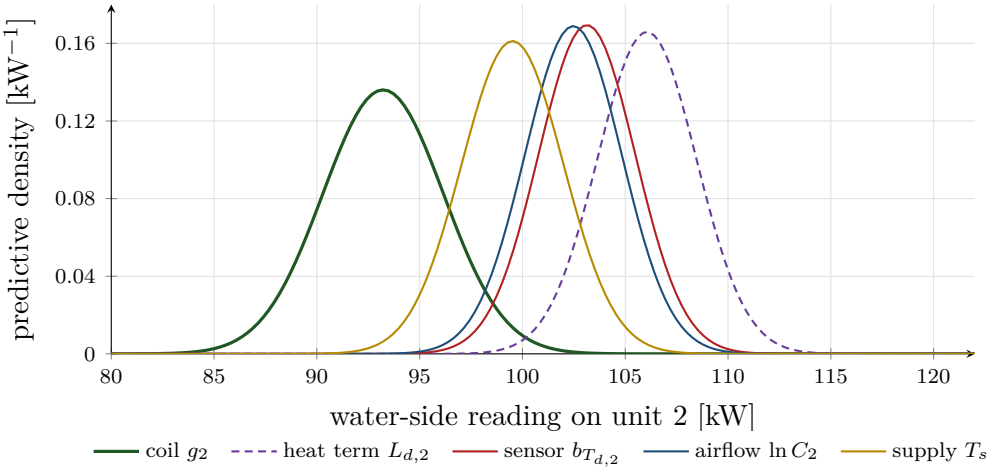


Figure 10.6: The predictive of a portable water-side heat check of unit 2, standard deviation 2 kW, under the five leading drawings at window 24, each normalized alone and not weighted by its probability.

Read on the entries, the candidates do different things. The water-side reading is the one row on Q_2: predicted 93.22 with a spread of 2.93 kW under the coil's drawing, against 106.05 under $L_{d,2}$, 103.14 (spread 2.355 kW) under the sensor, 102.47 under the airflow and 99.52 under the supply, with expected $\ln B$, if the coil's drawing holds, of 14.27, 8.91, 7.72 and 3.27. It is the

only candidate that acts on every rival, and the only one on the quantity the substitution (10.12) leaves free. The traverse and the supply check each split one rival only: in the family their gains are 0.19 against the portable check's 0.27, and in the extended family they fall to 0.0802 and 0.0801. The probe sees $L_{d,2}$ at 21.68 against 21.74°C, an expected $\ln B$ of 0.01, because its row repeats the discharge row; the same holds for the permanent second discharge sensor, whose gain matches the probe's within the standard errors.

10.6.3 The recommendation

The next reading is the portable water-side heat check of unit 2 tonight. The permanent instrument at the place of the question is the water-side heat meter on unit 2, which makes the heat row present in every window, at a gain of 0.5457 nats over the extended family against the permanent second discharge sensor's 0.0331. Chapter 7's hall carries that meter and showed it can read wrong, so its bias enters the family when it is added; this record does not hold it.

This replaces the recommendation of Chapter 9, Section 9.6.3, made at window 24. For the two-way question asked there, the coil or the sensor, the probe does come first as an ordering: if the coil's drawing holds, its expected $\ln B$ is 10.47 against the water-side heat check's 8.91. But the existing readings already favored the coil at 4.99 in that window, and that question was nearly closed when it was asked. Over the extended family the water-side heat check leads even at window 20, 0.4119 against the probe's 0.2116, and over the family of 12 there, 0.566 against 0.357. Chapter 18 ranks readings on the shift plant.

10.6.4 What the analysis showed

The dashboard pointed at a sensor, and the graph's other rows do not support it: a biased sensor would have left the return air and rack 1's inlet where the meter's load puts them, and they rose while the meter stayed within its noise. Isolation comes well after detection, from 00:35 to 01:05 on this record, because detection needs the misfit to be large and isolation needs its pattern to beat near neighbors whose signatures lie at cosines of up to 0.88 in size with the coil's. And the readings stop short of the physical cause: the same five numbers are what a coil passing less heat and a leak into its discharge air would give, a limit no amount of the record could lift. The missing row is one row, on the coil's own heat.

10.6.5 Limits of the model

One change at a time. Two changes together are outside the family: a coil and a sensor at once would be fitted as whichever single change comes closest, and the adequacy test would be the only sign.

The width choices. The priors are symmetric and the widths stated; the order within a non-isolable pair is theirs, and the Occam charge on the leader moves with them.

Each window alone. Evidence is not accumulated, and the night's non-monotone verdict is partly the price of that. How the change moves, and when it began, are Chapter 11's questions, on a stacked state with a trend.

An effective conductance. g_2 is one effective conductance, and it lumps fouling, water flow and valve position; the ranking does not separate them, and neither does the water-side heat check.

The exclusion of rule 4. The name "coil" depends on it, as stated. Other changes that would need new rows, such as a controller moving a valve, are not listed at all, and the probabilities compare listed drawings only.

Steady state and the first hour. The limits of Chapter 9 hold: steady state within a window, and the first hour taken as normal (Section 9.6.5). A baseline that had already moved would put part of the change into g_2^{b}.

Constructed record. The record is constructed: it is generated from a stated truth of the hall with the stated sensor errors, and step 5 compares the answers with that truth. It is not field data, and the agreement is a check of the method on its own model, not of the model against a real hall.

The next chapter. Chapter 11 reads the windows together.

10.7 What to retain

Definition 10.1 (Fault isolation by ranking)**.** A candidate family is generated from the plant graph: release an entry a baseline holds, put a bias on a reading, free a pinned quantity. Each candidate M_j adds one entry $\theta_j \sim \mathcal{N}(0,\ w_j^2)$ to the state of a window, and its fit gives its Laplace evidence $\ln p(\mathbf{y} \mid M_j) \approx -\mathcal{L}_j + \text{const} - \frac{1}{2}\ln\det\mathbf{\Lambda}_j - \ln w_j$. The ranking is $\Pr(M_j \mid \mathbf{y}_t) \propto p(\mathbf{y} \mid M_j)\ \Pr(M_j)$, read against a rule fixed in advance. A candidate's signature is its projected column $\mathbf{v}_j = (\mathbf{I} - \mathbf{H}_0)\mathbf{c}_j$; two candidates related by a substitution that leaves every reading the same are non-isolable, and their Bayes factor is a statement about their widths and the size of the implied change. Under the leading drawing a derate factor is read from the state as any derived quantity is, conditional on that drawing.

Three conclusions carry beyond the hall. Isolation is a harder test than detection, and comes later: it is limited by the angle between signatures as well as the size of the change. A ranking is only as wide as its family, so the family and the rule belong before the record, and a change the graph cannot hold is a limit to state beside the answer. And when a pair is non-isolable, the useful question is which quantity the substitution between them fails to preserve, because that quantity is where the next reading goes.

Table 10.12 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 10.12: The terms of fault isolation and where each is defined. Units for the two-unit hall.

Term	Expression	Units	Meaning	Defined in
Candidate family	Table 10.1, $J = 12$	–	changes the graph holds that a reading can see	10.1.2
Extended family	rule 4, eq. (10.1), $J = 15$	kW	added heat terms, beside the family	10.1.2
Decisive window	rule (a), (b), (c) at $\ln 20$ and 0.95	–	when the ranking can be acted on	10.1.3
Drawing's rows	eq. (10.3)	–	a window's rows plus one entry	10.1.4
Posterior and evidence				
Laplace evidence	eq. (10.7)	–	a drawing's fit less its price	10.2.1
Gain and Occam	eq. (10.8)	–	the bookkeeping of $\ln B_{j0}$	10.2.1
Fault signature	$\mathbf{v}_j = (\mathbf{I} - \mathbf{H}_0)\mathbf{c}_j$	–	what a change leaves in the residuals	10.2.1
Questions and answers				
Ranking	eq. (10.10)	–	probability among the listed	10.3.1
Time to isolation	eq. (10.11)	window	first and held windows	10.3.3
Non-isolable pair	eq. (10.12)	kW	same readings under a substitution	10.3.4
Derate factor	$d_2 = \mathrm{e}^{g_2^{\mathrm{b}} + \delta_2}$, eq. (10.13)	–	share of the rating, under the drawing	10.3.5
Expected reduction in entropy	eq. (10.14)	nats	value of a reading over the drawings	10.6.1

10.8 Exercises

Exercise 10.1. *Formulation.* From the bookkeeping of equation (10.8), find $\ln B$ of the airflow's drawing against no change at window 24. (The gain is $\mathcal{L}_0 - \mathcal{L}_{\ln C_2} = 14.0550$ and the Occam factor is 2.6724, so $\ln B = 14.0550 - 2.6724 = 11.3826$.)

Exercise 10.2. *Derivation.* Show from the chain (10.1) that the substitution (10.12) leaves the five readings unchanged, and evaluate it at window 24. (The hall balance gives $(G_1 + G_2^{\mathrm{b}})D = Q + Q_o + L_{d,2}$, satisfied by the D of M_{g_2}; $T_{d,2} = T_r - (G_2^{\mathrm{b}}D - L_{d,2})/C_2 = T_r - G_2 D/C_2$; the rest follows. With $G_2^{\mathrm{b}} = 6.2244$, $G_2 = 5.4649\,\mathrm{kW/K}$ and $D = 17.0589\,\mathrm{K}$, the heat term is $L_{d,2} = 12.957\,\mathrm{kW}$.)

Exercise 10.3. *Application.* An operator with a reason to doubt unit 2's coil marks it at one tenth the prior probability of any other candidate. What does the ranking say at window 24? (The prior of M_{g_2} becomes 0.1 of that of each of the other eleven, the probabilities of equation (10.10) are recomputed with the same evidences, and the first decisive window moves from 26 to 31; the coil still leads window 24, with probability 0.4259.)

Exercise 10.4. *Derivation.* Write the variance budget of g_2 at window 24 and say why its cross term is negative. (By equation (10.13), with the entries of $\mathbf{\Sigma}_{g_2}$; the readings hold the sum $g_2^{\mathrm{b}} + \delta_2$, so a higher baseline value is met by a larger fall, and the covariance is negative: $5.2033 \times 10^{-5} + 5.4083 \times 10^{-4} + 2(-3.6412 \times 10^{-5}) = 5.2004 \times 10^{-4}$.)

Exercise 10.5. *Application.* Find the expected $\ln B$ of the portable water-side heat check in favor of the coil against the sensor, if the coil's drawing holds. (Equation (7.23) with the predictives of the two drawings, means 93.224 and 103.135 kW and spreads 2.933 and 2.355 kW. The expected $\ln B$ is 8.91.)

Chapter 11

Trends and Limit-Crossing Forecasts

11.1 A limit no sensor reads

It is 00:55, window 24 of the night, and the operator has the hall's tiles and their trend lines on the screen. Unit 2's discharge reads 21.78°C, 1.22 K under its fixed 23°C alarm. Rack 1's inlet, the only inlet with a sensor, reads 22.95°C, 4.05 K inside the 27°C limit, and a straight line drawn through its last hour, windows 13 to 24, reaches 27°C at window 66.2, 04:26, 211 minutes away. The same line through the discharge reaches its alarm at window 32.1, 01:35. The operator reads the two lines together: an alarm may sound in forty minutes, and the racks have more than three hours.

Those lines are drawn through readings, and the limit is not at either reading. Chapter 9 found the departure and left it undated; Chapter 10 has unit 2's coil as its leading explanation, not yet decisive at this window. This chapter takes the coil as a condition: every answer below is conditional on the coil, and under a fault of the discharge sensor the rows of this chapter do not apply. The operator's next question is the one the screen cannot answer,

whether the coil is still getting worse, since when, and how long until the hall reaches its limit at the rack no sensor reads.

It is answered on the hall and the rows of Chapter 5 and the record of Chapter 9, with the trend form of Chapter 2. What changes is the state: the windows from midnight are stacked into one state whose entries include a level and a slope of the coil's log-derate.

Jargon: trend, slope, onset, crossing time, plant limit, reading level, forecast load, horizon, look-ahead, inversion, Fieller's construction, delta band, held-out test

A *trend* is a quantity carried by the state as a level and a *slope*, its change per window. The *onset* is the time at which the fitted line leaves the level of the reference. The *crossing time* t^* is the time at which a line that worsens reaches a stated level. A *plant limit* is a level the plant must stay within, here the 27°C rack inlet of the recommended envelope. A *reading level* is a level set on an instrument, here the discharge's fixed alarm; it says nothing of the plant by itself. The *forecast load* Q_f is the one steady load at which the line is carried forward. The *horizon* is how far ahead the operator asks, and a forecast to that horizon is a *look-ahead*. *Inversion* reads the band of a crossing time from the probability that the crossing has happened by each time, solved for the time at which that probability is 0.025 or 0.975; *Fieller's construction* is the standard closed form of that solution for a ratio of two normal quantities, when the probability is a single normal tail. The *delta band* is the band ± 1.96 spreads about a derived quantity, with the spread from its gradient. A *held-out test* predicts readings that the fit did not use.

Key idea

A crossing time is a ratio with the slope in its denominator: it has a median and a band, and no mean or variance. Its band is read by inverting the probability of its first passage, which is a normal tail at each time once the slope's band stands clear of zero. How well a forecast firms depends on which entry its band is made of: more windows read the slope and the coil's level better and better, and they read nothing about an entry no reading row sees.

11.1.1 A trend, a limit and a date

Chapter 9 tested each later window on its own against the first hour, and that test says that something has moved. It cannot say at what rate, because each window's fit has one state and the coil holds no memory between fits. A rate is a property of the windows read together. Section 2.3.4 wrote the general form: the state carries a level and a trend entry, each window's rows see the state through the quantity at their own time, and the information on the trend entry grows with the square of each window's distance from the center.

The date comes from the level the line is held against, and the two levels of the night are different things. The 23°C alarm is a level on a reading. The plant's limit is the rack inlet's 27°C of Chapter 5, and the chapter holds it at the rack the survey puts warmest, which has no sensor. Carried to the coil's scale the limit is itself a function of the state, through that rack's recirculation fraction, the load and the other coil. The crossing therefore inherits two sets of uncertainties, those of the line and those of the level, and the second has rows of a different kind.

> **Principle 11.1.** *A forecast's band is the sum of what is not known about the line and about the level. Before saying that a forecast firms, name which of the two the next window reads.*

11.1.2 The stacked hall through the night (step 0)

The plant is the two-unit hall of Chapter 5, Figure 5.1, through the record of Chapter 9: windows of five minutes, window 13 at 00:00, window 24 at 00:55, window 36 at 01:55. Its seven physical rows (5.1) are unchanged. On the reduced state of Chapter 4 their dependent block is triangular with a nonzero diagonal, so it is invertible and the temperatures of each window follow from its free entries in the chain (5.4): the return air, the discharge temperatures, the cold aisle, the inlets. The air balance (5.5) keeps the rise from the cold aisle to the return independent of the coils.

The stacked state. The windows of the span from the first window after the reference to the last window read, $t = 13, \dots, T$, form one state,

$$\mathbf{z} = (g_1,\ g_b,\ \beta_1,\ \beta_2,\ \beta_3,\ \alpha,\ \gamma,\ Q_f,\ Q_{13},\ \dots,\ Q_T)^{\mathsf{T}}, \qquad g_{2,t} = \alpha - \gamma\,(t - \bar{t}), \tag{11.1}$$

with $\bar{t}$ the mean of the span's windows. Here g_b is the coil's log-derate in the reference, α its level at $\bar{t}$, γ the slope per window, Q_t the load of window t and Q_f the forecast load. The line is the trend form (2.19) on the quantity $f_t = -g_{2,t}$, with level entry $-\alpha$ and trend entry γ, so that $\gamma > 0$ is a coil getting worse. At the decision window $T = 24$ the state has 20 entries and $\bar{t} = 18.5$.

The limit, named precisely. The plant limit is the rack inlet's $T_{\max} = 27°\mathrm{C}$. It is applied at rack 3, the rack nearest it by the survey's fractions. Holding it at one rack is a modeling choice: it buys one level and one date, and it fails if another rack is the warmer, which the survey leaves open for rack 2 with $\Pr(\beta_2 > \beta_3) = 0.092$. By the chain (5.4) and the air balance, rack 3's inlet at the forecast load is $T_3 = T_s + A/(G_1 + G_2) - (1 - \beta_3)A/(C_1 + C_2)$ with $A = Q_f + Q_o$. Setting it to $T_{\max}$ and solving for the second coil gives the limit on the coil's scale,

$$g^*(Q_f, g_1, \beta_3) = \ln \frac{1}{\bar{G}_2} \left[\frac{A}{T_{\max} - T_s + (1 - \beta_3)\, A/(C_1 + C_2)} - \bar{G}_1\, \mathrm{e}^{g_1} \right], \tag{11.2}$$

a level that the state carries, since three of its entries set it. The discharge's 23°C is a reading level, an alarm on a reading and not a limit of the plant. It is forecast beside the limit because it is the level the record can test: unit 2's discharge reaches a level $d + T_s$ at the log-derate

$$g_d(Q_f, g_1) = \ln \frac{1}{\bar{G}_2} \frac{A - d\,\bar{G}_1 \mathrm{e}^{g_1}}{d + A/C_2}, \qquad d = 23 - T_s = 9\,\mathrm{K}. \tag{11.3}$$

Where the limit is set, what assurance is required and the horizon of one hour used below, the look-ahead asked for, belong to the operator; the chapter computes for a stated limit and horizon.

Modeling choices at the graph. The choices of Chapter 5 at the graph hold in every window (Section 5.1.2): steady state within the window, one effective conductance per coil, pinned airflow and no active control. Each fails as it did there. The choices new with the trend enter at the rows, Section 11.1.4.

The graph. Figure 11.1 draws the stacked state. At the top are the shared entries (green) and the baseline of the first hour (blue), which reaches them as a prior row. The coil's line (green) is a node with its two entries, tied to the reference's g_b by a prior row. The windows below each take the shared entries and the line, and each adds its own load and sends five readings out along dashed arrows. The forecast (red) stands to the right, beyond the record: the line carried to a time s, at the forecast load, tied to the span's loads, with rack 3's fraction β_3 entering only there. No dashed arrow leaves the red node.

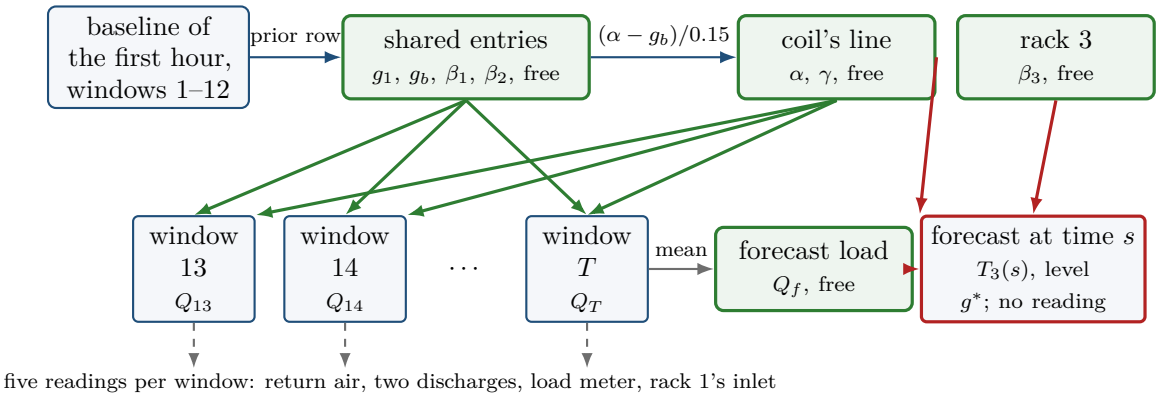


Figure 11.1: The stacked state of the trend, equation (11.1). Free (green): the shared entries, the coil's line, the forecast load and rack 3's fraction. Blue: the baseline, a prior row on $(g_1, g_b, \beta_1, \beta_2)$ and on β_3, whose arrow is not drawn, and each window's rows. Observed (dashed): the five readings of each window. The forecast (red) has no reading. Inside each window the hall is Figure 5.1.

11.1.3 The operator's questions (step 1)

The operator asks seven questions of the night's record, and step 1 writes each over the stacked state. Table 11.1 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 in the notation of Chapter 2; the third its mathematical form for the hall; the fourth its category; the fifth the computation of Section 2.3.6 that answers it; and the last the subsection where it is answered. Two are estimations of a rate and of a date, three are threshold probabilities along time, one asks for all of these as the windows arrive, and the last compares forms of the trend by their evidence.

Table 11.1: The questions asked of the night's record. $\mathbf{y}$ is the baseline, the readings of windows 13 to T and the trend's prior rows; τ the onset; g^* the level (11.2); $T_{3,t}$ rack 3's inlet at window t and that window's load.

Operator question	General form	Mathematical form	Category	Computation	Answered in
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

11.1.4 The evidence as rows (step 2)

Step 2 attaches the record of windows 13 to T to the graph. Each window carries the five readings of Chapter 5, the four temperatures with standard deviation 0.3 K and the load with 3.0 kW, each predicted through the chain (5.4) at the window's load and at the line's $g_{2,t}$. In window 24 they are the readings of Chapter 9, Section 9.1.4. Beside the $5(T - 12)$ reading rows stand eight prior rows, and each of the choices they carry is named here.

The line. The log-derate is a straight line over the span from the first window after the reference. This buys a level and a slope in two entries; it fails if the span reaches before the onset or if the coil curves. Both are tested in Section 11.3.5.

The trend prior. The slope has $\gamma \sim \mathcal{N}(0,\ 0.02^2)$ per window, one spread being a coil that loses 21.3 % of its conductance in an hour; the level has $\alpha \sim \mathcal{N}(g_b,\ 0.15^2)$, the rating's spread

about the reference's coil. These buy a proper state in the first windows. They fail by pulling the slope down if tight (Section 11.4.3).

The baseline as a prior. The baseline of Chapter 9 (Table 9.2) enters through its Gaussian summary, the five rows $\mathbf{L}_{\mathrm{b}}(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_{\mathrm{b}})$ with $\boldsymbol{\theta} = (g_1, g_b, \beta_1, \beta_2, \beta_3)^{\mathsf{T}}$, standing for the first hour's rows. This keeps the stacked state small. It excludes a change inside the first hour from the reference's own account; the ramp over the whole record tests that.

The forecast load. One steady load carries the line forward, $Q_f \sim \mathcal{N}(\bar{Q},\ 3.454^2)$, with $\bar{Q}$ the mean of the span's loads and 3.454 kW the sample spread of the first hour's fitted loads, the row $(Q_f - \bar{Q})/3.454$. This ignores the load's persistence between windows and the wander of any one window: the line crosses at a steady load, and the true inlet crosses on a wandering one. Twelve windows are too short to fit that persistence.

Independence. Loads and reading errors are taken as independent between windows, as in Chapter 9; it fails if a sensor's error drifts, and the time-ordered residuals of Section 11.2.2 look at it.

Which row acts on which entry. With $*$ for an entry that can be nonzero, the Jacobian is an arrowhead: each window's reading rows see the shared columns and their own load, and nothing of another window's load,

$$\mathbf{J} = \begin{bmatrix} \mathbf{S}_{13} & \mathbf{q}_{13} & & \\ \vdots & & \ddots & \\ \mathbf{S}_T & & & \mathbf{q}_T \\ \mathbf{P} & & \mathbf{P}_Q & \end{bmatrix},$$

$$\mathbf{S}_t = \begin{bmatrix} * & 0 & 0 & 0 & 0 & * & * & 0 \\ * & 0 & 0 & 0 & 0 & * & * & 0 \\ * & 0 & 0 & 0 & 0 & * & * & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ * & 0 & * & 0 & 0 & * & * & 0 \end{bmatrix}, \quad \mathbf{q}_t = \begin{bmatrix} * \\ * \\ * \\ * \\ * \end{bmatrix}, \tag{11.4}$$

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

with the columns $g_1, g_b, \beta_1, \beta_2, \beta_3, \alpha, \gamma, Q_f$ first, then the loads. The rows of $\mathbf{S}_t$ and $\mathbf{q}_t$ are the return air, the two discharge temperatures, the meter and rack 1's inlet of window t; $\mathbf{P}$ and $\mathbf{P}_Q$ hold the eight prior rows. The second line is the gradient of the crossing time; it is not a row.

The pattern holds the chapter's main fact. The columns of β_2 and β_3 are empty in the reading rows, as they were in Chapter 5, and no number of windows fills them: no reading of the record sees rack 3. Read as observability, β_3 is unobservable from this record at any T, and the gradient of the crossing has a nonzero in its column. The column of γ holds $-(t - \bar{t})$ times the column of α, which is the structure (2.20). The column of g_b has no reading row; the reference and the level's prior hold it.

11.2 The posterior of the state (step 3)

The inputs are the rows of Section 11.1.4, and every fit an answer uses is derived here.

11.2.1 The rows and the normal equations

Stacked with each row divided by its standard deviation, the rows are

$$\mathbf{r}(\mathbf{z}) = \begin{bmatrix} (m_k(Q_t, g_1, g_{2,t}, \beta_1) - y_{k,t})/\sigma_k \quad (t = 13, \dots, T;\ k = 1, \dots, 5) \\ \mathbf{L}_\mathrm{b}(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_\mathrm{b}) \\ (\alpha - g_b)/0.15 \\ \gamma/0.02 \\ (Q_f - \bar{Q})/3.454 \end{bmatrix}, \quad \mathcal{L}(\mathbf{z}) = \tfrac{1}{2}\,\mathbf{r}(\mathbf{z})^\mathsf{T}\mathbf{r}(\mathbf{z}), \tag{11.5}$$

with m_k the prediction of reading k as in equation (9.6): at $T = 24$, 68 rows on 20 entries, 60 reading rows and eight prior rows. The MAP state minimizes $\mathcal{L}$ by the normal equations of equation (4.29),

$$(\mathbf{J}^\mathsf{T}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^\mathsf{T}\mathbf{r}(\mathbf{z}), \qquad \boldsymbol{\Sigma} = (\mathbf{J}^\mathsf{T}\mathbf{J})^{-1}\big|_{\hat{\mathbf{z}}}, \tag{11.6}$$

from the baseline mean, $\alpha = \hat{g}_b$, $\gamma = 0$ and each window's meter reading for its load. The posterior is approximately normal, $\mathbf{z} \mid \mathbf{y} \approx \mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$, the Gauss–Newton form of equation (4.38). By the pattern (11.4) the rows of β_2 and β_3 are their prior rows alone, so their posteriors are the survey's, and the objective separates exactly 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 20 entries. It is set beside the Laplace normal by importance sampling: draws $\mathbf{z}^{(j)}$ from a Student-t proposal q with 12 degrees of freedom at $\hat{\mathbf{z}}$ with scale $\boldsymbol{\Sigma}$, each weighted by $w_j \propto \exp(-\mathcal{L}(\mathbf{z}^{(j)}))/q(\mathbf{z}^{(j)})$. With 400 000 draws at window 24 the effective sample size, $(\sum_j w_j)^2/\sum_j w_j^2$, is 308 201. Every exact value below is a weighted quantile or a weighted mean of these draws, its Monte Carlo standard error from twenty batches.

11.2.2 The fit at the decision window

The iterations. Table 11.2 runs equation (11.6) at window 24, 00:55. Each row is an iteration; the columns give the objective and the norm of its gradient. The objective falls from 135.5013 to 19.8939 in the first step and settles at 19.8852 by the second.

Table 11.2: Gauss–Newton iterations of the stacked state at window 24, equation (11.6).

Iteration k [–]	Objective $\mathcal{L}(\mathbf{z}^{(k)})$ [–]	$\lVert\mathbf{J}^\mathsf{T}\mathbf{r}\rVert$ [–]
0	135.5013	6.1×10^{3}
1	19.8939	5.9×10^{1}
2	19.8852	7.8×10^{-2}
3	19.8852	1.6×10^{-3}
4	19.8852	4.4×10^{-6}

The Jacobian of a window. The rows of window 24, in the eight shared columns of the pattern (11.4) and the column of Q_{24}, are

$$\begin{bmatrix} -31.0978 & 0 & 0 & 0 & 0 & -25.6521 & 141.0867 & 0 & 0.2749 \\ -48.1408 & 0 & 0 & 0 & 0 & -8.6091 & 47.3498 & 0 & 0.0923 \\ -14.0547 & 0 & 0 & 0 & 0 & -42.6952 & 234.8235 & 0 & 0.1243 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.3333 \\ -31.0978 & 0 & 34.4029 & 0 & 0 & -25.6521 & 141.0867 & 0 & 0.1454 \end{bmatrix}.$$

The column of γ is $-(24 - 18.5) = -5.5$ times the column of α. The zeros sit where the pattern (11.4) places them. Unit 2's discharge carries the largest entries in the coil's columns.

The covariance. On the entries the answers use, $(g_1, g_b, \beta_3, \alpha, \gamma, Q_f)$, with Q_f in kW,

$$\boldsymbol{\Sigma} = \begin{bmatrix} 2.287\times10^{-5} & -9.950\times10^{-6} & 0 & -1.093\times10^{-5} & 1.656\times10^{-7} & 8.460\times10^{-4} \\ -9.950\times10^{-6} & 4.540\times10^{-5} & 0 & 7.827\times10^{-6} & -7.320\times10^{-8} & -3.929\times10^{-4} \\ 0 & 0 & 6.400\times10^{-3} & 0 & 0 & 0 \\ -1.093\times10^{-5} & 7.827\times10^{-6} & 0 & 4.764\times10^{-5} & -4.045\times10^{-7} & 2.418\times10^{-3} \\ 1.656\times10^{-7} & -7.320\times10^{-8} & 0 & -4.045\times10^{-7} & 3.326\times10^{-6} & -2.582\times10^{-5} \\ 8.460\times10^{-4} & -3.929\times10^{-4} & 0 & 2.418\times10^{-3} & -2.582\times10^{-5} & 12.55 \end{bmatrix},$$

and the correlations of the same entries are

$$\begin{bmatrix} 1 & -0.3088 & 0 & -0.3312 & 0.0190 & 0.0499 \\ -0.3088 & 1 & 0 & 0.1683 & -0.0060 & -0.0165 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ -0.3312 & 0.1683 & 0 & 1 & -0.0321 & 0.0989 \\ 0.0190 & -0.0060 & 0 & -0.0321 & 1 & -0.0040 \\ 0.0499 & -0.0165 & 0 & 0.0989 & -0.0040 & 1 \end{bmatrix}.$$

The rows fix these signs. The level and the slope are nearly orthogonal, -0.0321, because the center $\bar{t}$ makes the two columns nearly so. The two coils are negatively correlated, as in Section 5.2.3: the return air fixes the sum of the conductances. The reference's g_b is positively correlated with α, 0.1683, through both of them against g_1. And β_3 is correlated with nothing.

The MAP state. Table 11.3 gives the fit: each row an entry, the columns its value, its posterior spread and the rows that inform it. The twelve loads run 196.01, 197.40, 197.98, 202.03, 204.17, 200.04, 198.39, 200.19, 201.40, 196.66, 200.63, 198.42 kW, each with a spread between 2.17 and 2.32 kW.

Table 11.3: The MAP state of the stacked trend at window 24, 00:55, with the loads omitted.

Entry [unit]	$\hat{z}_i$ [entry's]	$\sqrt{\Sigma_{ii}}$ [entry's]	Informed by
g_1 [–]	−0.05221	0.00478	baseline; temperature readings of the span
g_b [–]	−0.11718	0.00674	baseline; through α and g_1, the span
β_1 [–]	0.22313	0.00684	baseline; rack 1's inlet readings
β_2 [–]	0.35000	0.08000	survey only
β_3 [–]	0.50000	0.08000	survey only
α [–]	−0.18298	0.00690	temperature readings of the span
γ [per window]	0.01122	0.00182	temperature readings of the span, by distance from $\bar{t}$
Q_f [kW]	199.44355	3.54212	the span's loads and the first hour's wander

The reading rows. The reading rows have $X^2 = 38.095$ on 60 rows, with effective degrees of freedom $\nu = 44.985$: the fit leaves them within their noise, and the test of the whole window of Chapter 9 has nothing to fire on under a free slope. The time-ordered residuals of unit 2's discharge have a lag-one correlation of -0.331 and a Durbin–Watson statistic of 2.560: a drifting sensor error would make the correlation positive, and these residuals show no drift. The fitted loads' lag-one correlation is 0.197.

The square law. The slope's information is (2.20). Over the span $\sum_t (t - \bar{t})^2 = 143$. The one-window estimates of the coil, each from one window's five readings with the other entries on the baseline, have a spread that averages 0.02209 over the span, and $0.02209/\sqrt{143} = 0.00185$ against the fit's 0.00182: the slope's spread is the one-window spread over the root of that sum. At window 24 the one-window estimate is -0.2478 ± 0.0228.

11.2.3 The refits the answers use

The fit at each window. The rows (11.5) are refitted at each T from 14 to 36, on the span from window 13, with $\bar{t}$ the span's mean.

The forms on the whole record. To test the line, windows 1 to T are fitted under five forms of $g_{2,t}$, each on the same rows: all readings of those windows, the rating priors of Chapter 5 on g_1 and the survey on the fractions, with no baseline. With entries a, b, c of the form,

$$\begin{aligned} \text{no change:} \quad & g_{2,t} = a, \\ \text{line over the record:} \quad & g_{2,t} = a - b\,(t - \bar{t}_{1:T}), \\ \text{ramp from } \tau: \quad & g_{2,t} = a - b\,\max(t-\tau, 0), \\ \text{curved ramp:} \quad & g_{2,t} = a - b\,u_t - c\,u_t^2, \quad u_t = \max(t-\tau,0), \\ \text{step:} \quad & g_{2,t} = a - b\,\mathbf{1}[t > \tau], \end{aligned} \tag{11.7}$$

with $a \sim \mathcal{N}(0,\ 0.15^2)$, a slope $b \sim \mathcal{N}(0,\ 0.02^2)$, a step size $b \sim \mathcal{N}(0,\ 0.15^2)$ and a curvature $c \sim \mathcal{N}(0,\ (0.02/12)^2)$. The ramps run on a grid of τ with step 0.5 window over 1 to $T-1$, and the step over the half-integer times; τ has a uniform prior over the grid. Each fit has the Laplace evidence of its own rows,

$$\ln p(\mathbf{y} \mid M) \approx -\mathcal{L}(\hat{\mathbf{z}}_M) - \tfrac{1}{2}\ln\det(\mathbf{J}_M^{\mathsf{T}}\mathbf{J}_M) - \sum_i \ln s_i + (\text{common terms}), \tag{11.8}$$

with s_i the prior spreads; each prior-bearing entry's $\frac{1}{2}\ln 2\pi$ cancels its prior's. The loads, which have no prior, enter every form alike, and their constant cancels in the ratios. The ramps' evidence averages over the grid; the onset's posterior is that average's normalized density.

The crossing under the ramp. A mixture over the grid of the ramp's fits, weighted by the onset's posterior, gives 400 000 draws of the crossing, with the forecast load of Section 11.1.4.

11.3 The answers (step 4)

Every answer is read from the state of Section 11.2 at window 24 unless it says otherwise, and that value is the reported one. A quantity linear in the state is carried by its gradient. A ratio with the slope in its denominator is reported by its median and the band by inversion of its distribution function, as Chapter 2 and *Reliability Engineering by Derivation*, Chapter 12, report a ratio, with the delta form and the exact posterior beside it.

11.3.1 The sign and the rate

Operator questions. Is it getting worse. How fast? The first asks whether the slope is positive at all on the record to window 24; the second asks for the rate on a scale the operator can use.

General form. The sign is the first line of the trend form, equation (2.21), and under the Laplace posterior the normal tail of equation (2.22). The rate is an estimation, equation (2.9).

Specialized form. The two questions are equations (2.21) and (2.9) with $\mathbf{z}$ the stacked state (11.1), the general form's trend entry the slope γ, its level entry $-\alpha$, its quantity $f_t = -g_{2,t}$,

and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal of equation (11.6); for the rate, $f(\mathbf{z}) = 12\gamma$ per hour and the fraction of the conductance lost per hour. So

$$\Pr(\gamma > 0 \mid \mathbf{y}) = \Phi\left(\frac{\hat{\gamma}}{\sqrt{\Sigma_{\gamma\gamma}}}\right), \qquad 1 - \frac{G_{2,t+12}}{G_{2,t}} = 1 - \mathrm{e}^{-12\gamma}, \tag{11.9}$$

the second a monotone function of γ, whose band is the image of the slope's band. The level now is $g_{2,24} = \alpha - \gamma\,(24 - \bar{t})$, linear in the state.

Building blocks. $\hat{\alpha}$, $\hat{\gamma}$ and their block of $\mathbf{\Sigma}$ (Section 11.2.2); the normal distribution function Φ.

The sign and the rate at window 24. The normalized slope is $0.01122/0.00182 = 6.155$, and

$$\Pr(\gamma > 0 \mid \mathbf{y}) = \Phi(6.155) = 1 - 3.75 \times 10^{-10}.$$

The exact draws put all their weight on $\gamma > 0$ to the eighth decimal; the exact spread of the slope is 0.00183. The rate is 0.13469 ± 0.02188 per hour, band [0.09180, 0.17758]: the coil loses 12.60 % of its conductance per hour, band 8.77 to 16.27 %. The level now is $g_{2,24} = -0.24472 \pm 0.01236$, a conductance of $G_2 = 5.480\,\mathrm{kW/K}$ against the rating's $7.0\,\mathrm{kW/K}$, inside the one-window estimate of Section 11.2.2.

11.3.2 The onset

Operator question. When did it begin? This is the question Chapter 9 left open: its first flag said when the change showed, not when it started.

General form. An estimation question asks for the posterior distribution of a derived quantity, equation (2.9), summarized by a named location and the band (2.11).

Specialized form. The question is equations (2.9) and (2.11) with $\mathbf{z}$ the stacked state (11.1), $f(\mathbf{z})$ the time at which the line leaves the reference's level,

$$\tau(\mathbf{z}) = \bar{t} + \frac{\alpha - g_b}{\gamma}, \tag{11.10}$$

and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal. It is a ratio with γ in its denominator, a quantity with no posterior variance, for the reason given at equation (2.22). It is reported as the crossing of a level is: for $\gamma > 0$ the event $\tau \le s$ is the event $(\alpha - g_b) - \gamma(s - \bar{t}) \le 0$, linear in (α, g_b, γ), so

$$\Pr(\tau \le s \mid \mathbf{y}) \approx \Phi\left(\frac{\hat{\gamma}\,(s - \bar{t}) - (\hat{\alpha} - \hat{g}_b)}{\sqrt{V_\tau(s)}}\right), \quad V_\tau(s) = \mathrm{Var}\big((\alpha - g_b) - \gamma\,(s - \bar{t}) \mid \mathbf{y}\big). \tag{11.11}$$

This is the single tail of the joint with $\gamma > 0$, from which it differs by a mass below $\Pr(\gamma \le 0 \mid \mathbf{y}) = 3.75 \times 10^{-10}$ at window 24. Its median is $\tau(\hat{\mathbf{z}})$, where the argument is zero, and its band solves this at 0.025 and 0.975. Write n for the numerator, here $\alpha - g_b$, with the variance V_{nn} and its covariance with the slope $V_{n\gamma}$, and $u = s - \bar{t}$. Squared, the two ends are the roots of one quadratic,

$$(\hat{\gamma}^2 - c^2\Sigma_{\gamma\gamma})\,u^2 - 2(\hat{n}\hat{\gamma} - c^2 V_{n\gamma})\,u + (\hat{n}^2 - c^2 V_{nn}) = 0, \qquad c = 1.960, \tag{11.12}$$

Fieller's construction. Two ends exist only if the first coefficient is positive, that is only if the slope's normalized value exceeds c. The delta form carries τ by its tangent, with gradient $\partial\tau/\partial(\alpha, g_b, \gamma) = (1/\gamma,\ -1/\gamma,\ -(\alpha - g_b)/\gamma^2)$.

Building blocks. $\hat{\alpha}$, $\hat{g}_b$, $\hat{\gamma}$ and their block of $\mathbf{\Sigma}$ (Section 11.2.2); the ramp's fits on the grid (Section 11.2.3); the exact draws.

The onset at window 24. By equation (11.10),

$$\tau(\hat{\mathbf{z}}) = 18.5 + \frac{-0.18298 + 0.11718}{0.01122} = 12.638,$$

window 12.638, 23:58, and the band by inversion is [9.380, 14.594], 23:42 to 00:08. The delta form gives a spread of 1.221 windows and [10.245, 15.030], 23:46 to 00:10. The exact posterior has the median 12.649 and the band [9.375, 14.599]: gaps of +0.012, −0.005 and +0.005 from the reported values, against −0.869 and −0.431 from the delta band's ends. A skew toward early onsets is what a ratio with a small denominator has, and a symmetric band misses it.

The delta spread's budget in windows^2 names the entries behind it: α 0.378, g_b 0.360, γ 0.907, the cross term of α and γ −0.038, the rest −0.117. The last is mostly the cross term of α and g_b, −0.1242: their correlation is positive and the onset reads their difference, so the term is negative. The delta spread stays between 1.11 and 1.29 from window 22 to window 36, while the spread of the line against the reference at the onset falls from 0.01370 at window 24 to 0.01102 at window 36: the slope that divides it falls as well.

The ramp over the whole record, which sets no reference, dates the same change. On its grid the onset has the median 12.222, 23:56, the band [8.707, 14.753], 23:39 to 00:09, and the mode 13.0, with $\Pr(\tau < 12) = 0.442$ and $\Pr(\tau < 16) = 0.998$. Figure 11.2 draws that density: the horizontal axis is the onset in windows, the vertical axis its posterior density on the grid. Above it the reported band and the ramp's band are drawn as bars, and the dotted line is the stated truth, a check marker for Section 11.4.4. The density has a broad top from window 11 to window 13.5 and a long lower tail. The reported band narrows slowly, 26 minutes at window 24 and 22 at window 36, and the ramp's from 30.2 to 23.4. Its floor is the reference's level, $\mathrm{sd}(g_b)/\hat{\gamma} = 0.66$ window at window 36, which later windows barely read. The record cannot say whether the change began in the last quarter hour of the reference.

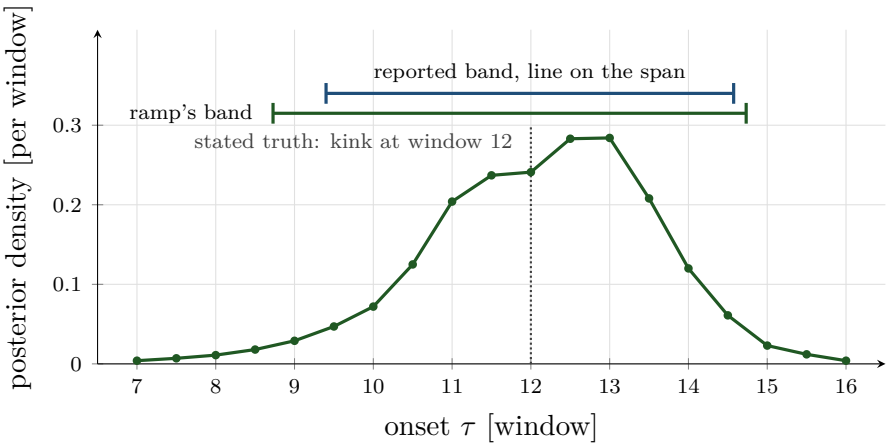


Figure 11.2: The onset's posterior under the ramp over the whole record, windows 1 to 24, on its grid of 0.5 window. Bars: the reported band on the span, [9.380, 14.594], and the ramp's, [8.707, 14.753]. Dotted: the stated truth (Section 11.4.4).

Figure 11.3 draws the coil through the night: the horizontal axis is the window, the vertical axis the log-derate g_2. The points are the one-window estimates of Section 11.2.2, filled to window 24 and hollow for the twelve windows the fit has not read. The blue line is the fit of window 24 with its 95 % band, carried to window 37; the short gray line is the reference's level; the shaded strip is the onset's band; the dotted line is the stated truth. The band pinches near $\bar{t}$ and opens along the lever arm of equation (2.22).

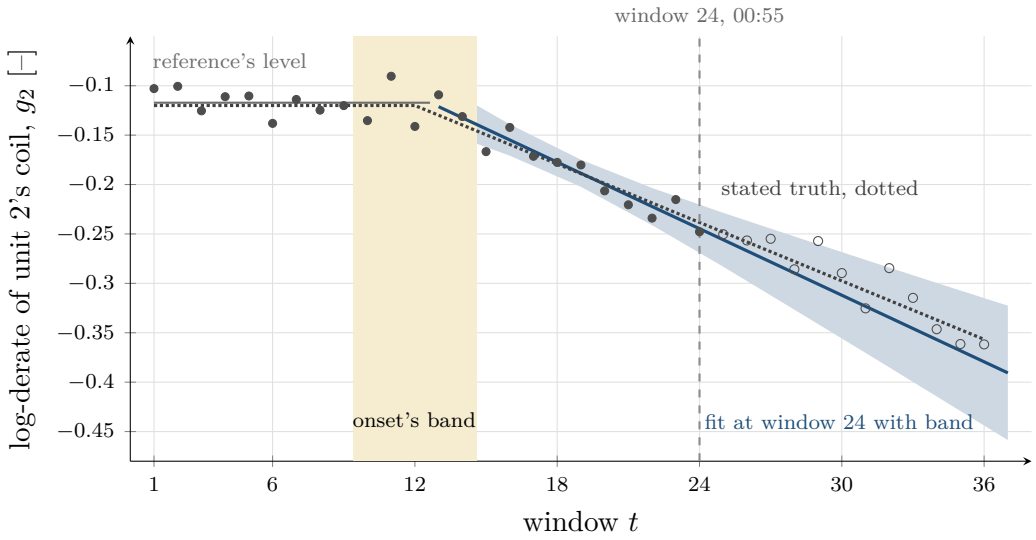


Figure 11.3: Unit 2's log-derate through the night. Points: the one-window estimates, spread about 0.022 each (not drawn). Blue: the line of the fit at window 24 and its 95 % band. Gray: the reference's level, to the onset. Shaded: the onset's band. Dotted: the stated truth of the record, a check marker (Section 11.4.4). The hollow points are held out.

11.3.3 The crossing time

Operator questions. When does the plant cross its limit, and with what band? Could it be over already? The first asks for a date at rack 3; the second for the same event at the present window, where the band already reaches.

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

$$\Pr(t^* \le s \mid \mathbf{y}) = \int \mathbf{1}[\gamma > 0]\, \mathbf{1}\big[f_s(\mathbf{z}) \ge f^*\big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z},$$

with a band by equation (2.11), and whether it is over now is a threshold, equation (2.12).

Specialized form. The two questions are equation (2.21), line 2, and equation (2.12) with $\mathbf{z}$ the stacked state (11.1), $f_s = -g_{2,s}$, the level $f^* = -g^*$ of equation (11.2), itself a function of the state, $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal, and for the threshold $f(\mathbf{z}) = T_{3,24}(\mathbf{z})$, at the window's load Q_{24}, with $f^* = T_{\max}$. No assurance level κ is set; the probabilities are reported. One choice enters here. A state with $\gamma \le 0$ is taken never to cross: this keeps every state in the distribution, with an atom at infinity, and it would mislead if a flat coil could turn worse later, which the line excludes. The crossing time is

$$t^*(\mathbf{z}) = \bar{t} + \frac{\alpha - g^*(Q_f, g_1, \beta_3)}{\gamma}. \tag{11.13}$$

The margin $m(s) = g^* - \alpha + \gamma(s - \bar{t})$ is positive once the line stands past the limit at time s. For $\gamma > 0$ the event $t^* \le s$ is $m(s) \ge 0$ at any s, and it is the joint event $\gamma > 0$, $m(s) \ge 0$ that the general form asks. The line stands below the level at $\bar{t}$ in 0.035 of the exact posterior, states

whose crossing lies before $\bar{t}$. Carrying g^* by its gradient, the joint of equation (2.22) becomes

$$\begin{aligned}
&\Pr(t^* \le s \mid \mathbf{y}) \approx \Phi_2\!\left(\frac{\hat{\gamma}}{\sqrt{\Sigma_{\gamma\gamma}}},\ u_m(s);\ \varrho_m(s)\right) = \Phi(u_m(s)) - \Pr(\gamma \le 0,\ m(s) \ge 0 \mid \mathbf{y}),\\
&u_m(s) = \frac{m(s)\big|_{\hat{\mathbf{z}}}}{\sqrt{\mathbf{g}_m(s)^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{g}_m(s)}}, \qquad \mathbf{g}_m(s) = \frac{\partial m(s)}{\partial \mathbf{z}}, \qquad \varrho_m(s) = \frac{\mathbf{g}_m(s)^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{e}_\gamma}{\sqrt{\Sigma_{\gamma\gamma}\,\mathbf{g}_m(s)^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{g}_m(s)}},
\end{aligned} \tag{11.14}$$

with $\mathbf{e}_\gamma$ the unit vector of the slope's entry. The single tail $\Phi(u_m)$ is the threshold line. It also counts the states with $\gamma \le 0$ that stand past the limit, a flat coil with a rack whose fraction already puts it over, and the crossing subtracts them. That mass is below $\Pr(\gamma \le 0 \mid \mathbf{y}) = 3.75 \times 10^{-10}$ at window 24, 1.4×10^{-11} at the median and at the ends of the band, so there the threshold line is the crossing's distribution function to every digit printed: its median is $t^*(\hat{\mathbf{z}})$, where u_m is zero, and the mass at never crossing is $\Phi(-\hat{\gamma}/\sqrt{\Sigma_{\gamma\gamma}})$. Here the numerator is $n = \alpha - g^*$, with V_{nn} and $V_{n\gamma}$ by the gradient of n. The ends of the band solve equation (11.14) at 0.025 and 0.975. On the threshold line they are the roots of the quadratic (11.12) of Section 11.3.2 with this numerator, and with $\alpha - g_d$ of equation (11.3) the ends for the reading level; where the mass shows at the digits printed, on the early records of Section 11.3.4, the ends are solved on the joint. The joint rises only to $\Phi(\hat{\gamma}/\sqrt{\Sigma_{\gamma\gamma}})$, so past the first coefficient's zero the band has a lower end and no upper end on both.

The delta form is beside it. By equation (11.13) the gradient of the crossing, with the pattern of equation (11.4), is

$$\frac{\partial t^*}{\partial(\alpha, \gamma, g_1, \beta_3, Q_f)} = \left(\frac{1}{\gamma},\ -\frac{t^* - \bar{t}}{\gamma},\ -\frac{1}{\gamma}\frac{\partial g^*}{\partial g_1},\ -\frac{1}{\gamma}\frac{\partial g^*}{\partial \beta_3},\ -\frac{1}{\gamma}\frac{\partial g^*}{\partial Q_f}\right), \tag{11.15}$$

with the delta band $t^*(\hat{\mathbf{z}}) \pm 1.96$ spreads and the budget (5.14). The form is wrong in three ways, named here. It is symmetric about a ratio whose distribution is skewed, so that both ends sit early; it carries no mass on never crossing; and its spread is finite where the crossing time has no variance at all, the tangent's spread standing in for an infinite one. It stays useful as a measure on one scale: its budget splits by entry, and the table of each window carries its spread as a width. No band of a ratio is reported from it.

The over-now question is a present-tense event, and it is read at the window's own load from the state: the tail of $T_{3,24}$ at Q_{24} by the delta method, which is reported. The threshold line of equation (11.14) at $s = 24$ asks a different thing, whether the line stands past the level at a steady forecast load. It is set beside the reported value, named as that.

Building blocks. $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$ (Section 11.2.2); the level (11.2) and its gradient; the joint (11.14), its threshold line and the line's quadratic (11.12); the gradient (11.15); the exact draws.

The limit on the coil's scale. At $\hat{Q}_f = 199.44\,\text{kW}$, $\hat{g}_1 = -0.05221$ and $\hat{\beta}_3 = 0.500$, equation (11.2) gives $g^* = -0.38507$, a conductance of $4.763\,\text{kW/K}$, with the gradient -1.3950 in g_1, 1.3659 in β_3 and 0.00825 per kW in Q_f. On the same line rack 1 reaches 27°C at window 71.73 and rack 2 at 55.14. The choice of rack moves the answers little: in the exact posterior rack 2 or rack 3 is over by window 24 with 0.1054 against rack 3's 0.1044 and by window 36 with 0.4921 against 0.4775, on the same side of one half, and the upper end for the first of the two is 58.375 against 60.913.

The crossing. By equation (11.13),

$$\hat{t}^* = 18.5 + \frac{-0.18298 + 0.38507}{0.01122} = 36.504,$$

01:58, 62.5 minutes ahead, and the reported band by the quadratic (11.12) is [16.74, 60.33], 00:19 to 03:57. Figure 11.4 reads it off on the scale of the limit. Its left panel has the window on the horizontal axis and rack 3's inlet at the forecast load on the vertical axis, with the fit's center and 95 % band carried forward from the span, the 27°C limit in red and the stated truth dotted. Its right panel is the same for unit 2's discharge against its 23°C alarm, with the readings as points, hollow where the fit has not read them. The left band meets the limit near window 17 at its upper edge and near window 60 at its lower edge, and the center meets it at window 36.5. The short red bar at window 24 is the inlet at that window's own load.

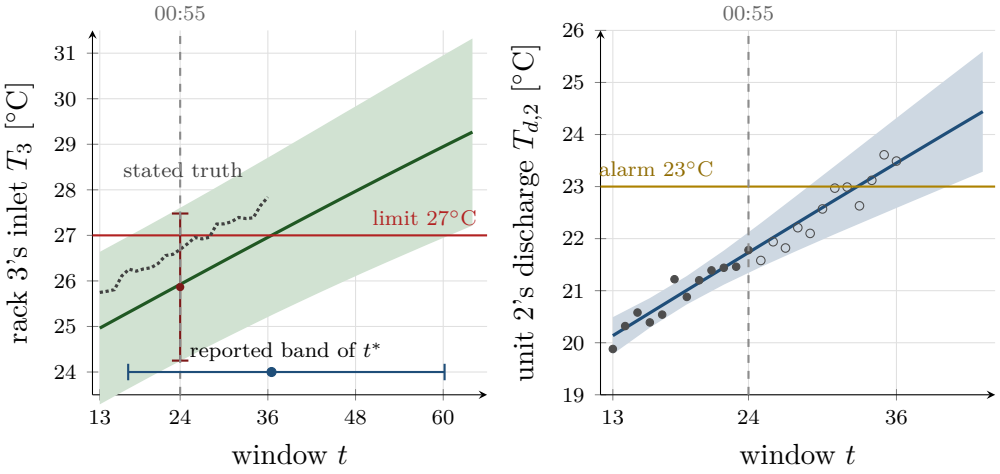


Figure 11.4: The forecast at window 24. Left: rack 3's inlet at the forecast load with its 95 % band, against the plant limit, the stated truth dotted (a check marker), and the reported band of the crossing time as a bar with its median; red bar, $T_{3,24}$ at Q_{24}. Right: unit 2's discharge, center and band, against its reading level, with the readings as points; hollow points were not read by the fit.

Table 11.4 sets the routes beside each other at window 24: each row is a quantity, the columns the reported value, the delta form, draws of the Laplace normal, the exact posterior and the gap from unrounded values, each with its Monte Carlo standard error. Figure 11.5 draws the distribution functions: the horizontal axis is the time s, the vertical axis the probability that the crossing has happened by then. The three curves nearly coincide near the median and part in the upper tail, where the delta form runs above the other two.

Table 11.4: The crossing at window 24 by route, with 400 000 draws each; standard errors in parentheses. The gap is exact minus reported.

Quantity [window or –]	Inversion, reported	Delta	Normal draws	Exact	Exact − reported
median t^*	36.504	36.504	36.550 (0.016)	36.592 (0.030)	+0.088
2.5 % point	16.743	15.937	17.044 (0.043)	17.095 (0.051)	+0.352
97.5 % point	60.327	57.072	60.780 (0.064)	60.913 (0.087)	+0.587
$\Pr(t^* \le 24)$, steady load	0.1083	0.1167	0.1062 (0.0005)	0.1044 (0.0006)	−0.0039
$\Pr(t^* \le 36)$	0.4808	0.4808	0.4790 (0.0008)	0.4775 (0.0009)	−0.0033
Pr(never)	3.75×10^{-10}	no mass	below resolution	below resolution	–

The delta form's lower end stands 0.806 window before the reported one and its upper end 3.254 windows, 16.3 minutes, before it; against the exact upper end the gap is 3.841 windows. It has no mass at never crossing, where the line has 3.75×10^{-10}. The exact values stand beside the reported ones, a small step later, and every gap above is larger than its standard error. No

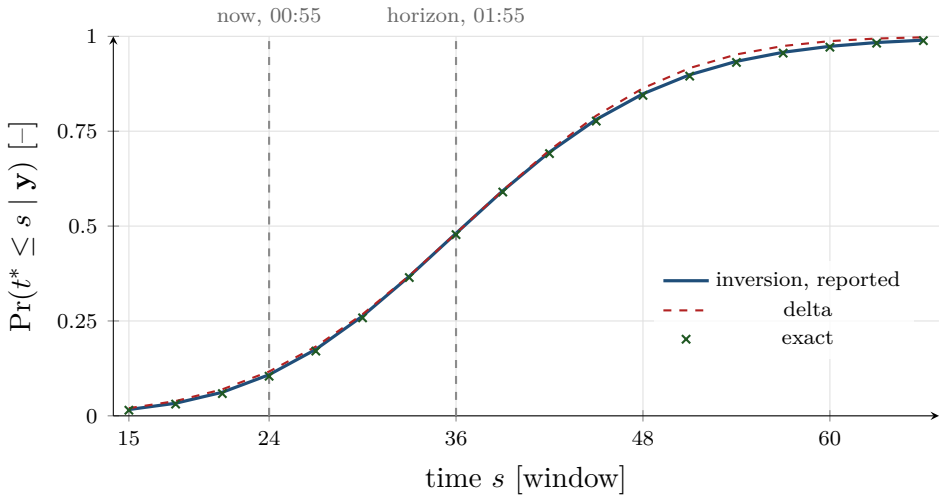


Figure 11.5: The crossing's distribution function at window 24 by the threshold line of equation (11.14), by the delta form's normal and by the exact posterior. The two dashed verticals are the present window and the operator's one-hour horizon.

gap crosses a decision level at window 24: at the horizon the reported 0.4808 and the exact 0.4775 stand on the same side of one half, and the band's lower end lies before the present on every route. One reported value stands, with the exact one beside it. Most of the exact gap at the ends is g^{*}'s curvature in β_3: normal draws with g^{*} linear give [36.501, 16.709, 60.288], and the normal draws stand +0.049, +0.334 and +0.492 from those. At the operator's one-hour horizon the crossing is a coin toss.

The budget. Figure 11.6 draws the delta spread's terms from the gradient (11.15), whose entries at $\hat{\mathbf{z}}$ are 89.09 in α, −1604.0 in γ, 124.28 in g_1, −121.69 in β_3 and −0.7353 per kW in Q_f: one bar per term, the horizontal axis in windows^2. The total of 110.123 gives a spread of 10.494 windows, 52.5 minutes. The survey's β_3 alone is 94.778, 86.1 %. The slope's term is the only large term the record reduces, 8.556 at window 24 and 0.757 at window 36. The β_3 term rises to 120.741, since the slope's fall lengthens it in windows. The cross term of α and γ is positive, 0.116, because their gradient entries have opposite signs and their covariance is negative. The other cross terms are all negative, together −0.841: α with g_1 −0.242 by the negative covariance of the coils, α with Q_f −0.317 and g_1 with Q_f −0.155 because a larger load is read as larger conductances and these gradient entries oppose, and the two with γ, −0.066 and −0.061, small.

Over already. At the window's own load the state gives $T_{3,24} = 25.865 \pm 0.836$°C, with the gradient −9.329 in g_1, −7.696 in α, 42.326 in γ, 10.321 in β_3 and 0.0575 per kW in Q_{24}, so that $\Pr(T_{3,24} \geq 27\text{°C} \mid \mathbf{y}) = 0.0871$; the exact value is 0.0866 (se 0.0005). Beside it stands the threshold line at a steady forecast load, 0.1083, with the exact 0.1044 and the tail 0.1048 of the same inlet at Q_f, 25.924 ± 0.858°C. The band's lower end, which is on that line, lies before the present because rack 3 may already be over; it is not a date. Crossing before the onset, over the limit before the coil moved, has exact probability 0.00785.

The reading level. By equation (11.3) the discharge reaches 23°C at $g_d = -0.34372$, and the same line reaches it at the median 32.820, 01:39, with the band by inversion [28.81, 40.08], 01:19 to 02:15. The delta form gives a spread of 2.611 and [27.70, 37.94]; the exact posterior the median 32.841 and [28.814, 40.096]. It is the same line without β_3, whose gradient entry

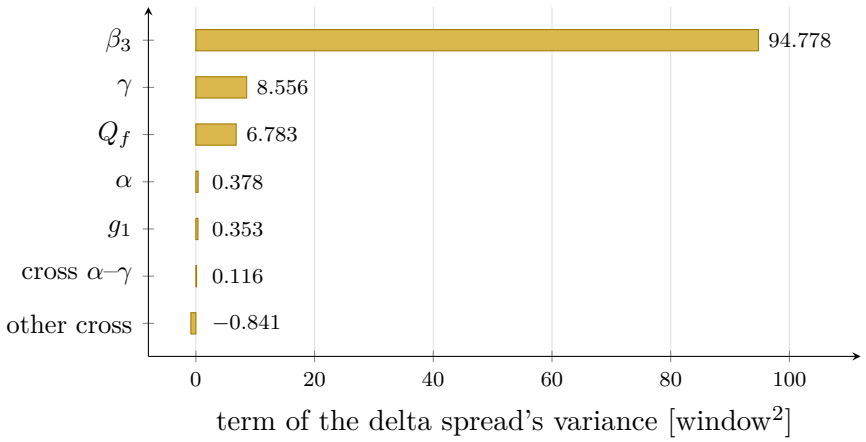


Figure 11.6: The budget of the crossing's delta spread at window 24, total 110.123 window^2. The survey's fraction for rack 3 carries 86.1 % of it.

here is zero, so that its band is the line's and the forecast load's, Q_f holding 1.159 of the delta variance 6.816, 17.0 %, and g_1 0.030. It is the level the record can test.

11.3.4 The forecast as of each window

Operator question. How does the answer firm with more windows, and can it answer at all yet? The operator at 00:15 and the operator at 01:55 ask the same thing with different records.

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

Specialized form. The question is equation (2.21) at each T from 14 to 36, with $\mathbf{z}$ the stacked state of that T, $\bar{t}$ its span's center, $f_s = -g_{2,s}$, the level $f^* = -g^*$ and $p(\mathbf{z} \mid \mathbf{y})$ the Laplace normal of its fit (Section 11.2.3). The crossing's median and band are the joint's, equation (11.14), at every T. The split of the delta spread asks what would remain with β_3 known: the conditional covariance $\boldsymbol{\Sigma} - \boldsymbol{\Sigma}\mathbf{e}_{\beta_3}\mathbf{e}_{\beta_3}^{\mathsf{T}}\boldsymbol{\Sigma}/\Sigma_{\beta_3\beta_3}$ in the same quadratic form.

Building blocks. The fit at each window (Section 11.2.3); the quadratic (11.12); the gradient (11.15).

The table. Table 11.5 gives every window: each row is a decision window, with its slope, the slope's normalized value, the onset with its delta spread, the crossing's median and reported band, its delta spread, that spread with β_3 known, the threshold line at that window at a steady load, and the reading level's median and delta spread. Figure 11.7 draws the crossing's median and band from window 20 on with the band by inversion at β_3 known beside it, the vertical axis in windows; below it, on a second panel, the slope and its band. Both horizontal axes are the decision window; the dotted lines are the stated truth.

The slope firms and the crossing does not. From window 20 to window 36 the slope's normalized value goes from 3.43 to 15.18 and the reading level's delta spread from 4.86 to 1.40, while the crossing's spread stays between 9.76 and 11.99, and from window 21 the reported band's upper end stays between 58.49 and 63.12. With β_3 known the same spread would fall from 6.27 to 2.99, and the band by inversion from [27.84, 63.26] to [32.03, 43.84]. The record's readings do not contain β_3, and waiting does not help. At window 36, with the whole record, the threshold line at a steady load is still 0.4403.

Table 11.5: The forecast as of each window. The median and reported band of t^* are by inversion of the joint (11.14); a dash marks no upper end. Spreads (sd) are delta spreads, kept as a width on one scale; the line at T is the threshold line, a single tail.

T (clock)	$\hat{\gamma}$ [per window]	z [–]	τ (sd) [window]	t^* [window]	band [window]	sd [window]	sd, β_3 known [window]	line at T [–]	t^*_{23} (sd) [window]
14 (00:05)	0.00715	0.43	13.13 (2.43)	52.63	[17.27, –]	92.58	91.31	0.0078	45.69 (75.14)
15 (00:10)	0.01847	1.52	13.03 (0.98)	28.50	[16.15, –]	11.34	9.67	0.0147	25.71 (7.74)
16 (00:15)	0.01129	1.31	12.74 (1.74)	36.92	[17.07, –]	19.81	17.30	0.0191	33.10 (14.25)
17 (00:20)	0.01271	1.99	12.87 (1.40)	33.59	[16.28, 1147.43]	12.84	9.61	0.0311	30.75 (7.97)
18 (00:25)	0.01165	2.34	12.74 (1.50)	35.38	[16.42, 147.60]	12.87	8.86	0.0372	32.12 (7.19)
19 (00:30)	0.01067	2.67	12.57 (1.60)	37.26	[16.60, 107.70]	13.21	8.43	0.0426	33.76 (6.78)
20 (00:35)	0.01131	3.43	12.71 (1.40)	36.08	[16.54, 74.57]	11.49	6.27	0.0543	32.65 (4.86)
21 (00:40)	0.01191	4.30	12.84 (1.25)	35.02	[16.44, 63.06]	10.36	4.86	0.0705	31.77 (3.62)
22 (00:45)	0.01232	5.16	12.93 (1.15)	34.71	[16.73, 58.49]	9.76	4.10	0.0839	31.39 (2.91)
23 (00:50)	0.01107	5.35	12.59 (1.29)	36.74	[16.69, 62.59]	10.81	4.39	0.0905	33.01 (3.06)
24 (00:55)	0.01122	6.16	12.64 (1.22)	36.50	[16.74, 60.33]	10.49	3.92	0.1083	32.82 (2.61)
25 (01:00)	0.01109	6.87	12.59 (1.20)	36.50	[16.55, 59.54]	10.48	3.65	0.1300	33.00 (2.36)
26 (01:05)	0.01083	7.51	12.49 (1.21)	36.75	[16.33, 59.69]	10.66	3.52	0.1518	33.30 (2.19)
27 (01:10)	0.01038	7.94	12.29 (1.25)	37.72	[16.43, 61.30]	11.08	3.57	0.1625	34.08 (2.17)
28 (01:15)	0.01048	8.79	12.34 (1.20)	37.69	[16.57, 60.61]	10.94	3.39	0.1847	33.97 (1.97)
29 (01:20)	0.00982	9.03	11.96 (1.29)	38.80	[16.27, 63.09]	11.65	3.54	0.1973	34.98 (2.03)
30 (01:25)	0.00971	9.75	11.89 (1.28)	38.80	[16.01, 63.05]	11.75	3.45	0.2249	35.02 (1.91)
31 (01:30)	0.01006	10.93	12.12 (1.19)	38.00	[16.00, 61.06]	11.30	3.20	0.2667	34.42 (1.69)
32 (01:35)	0.00954	11.21	11.74 (1.26)	38.73	[15.48, 63.00]	11.93	3.32	0.2855	35.00 (1.72)
33 (01:40)	0.00948	11.94	11.69 (1.24)	38.85	[15.46, 63.12]	11.99	3.29	0.3121	35.12 (1.66)
34 (01:45)	0.00967	13.01	11.85 (1.19)	38.42	[15.48, 62.05]	11.74	3.16	0.3529	34.77 (1.55)
35 (01:50)	0.00985	14.11	11.99 (1.14)	38.10	[15.56, 61.19]	11.53	3.06	0.3938	34.42 (1.46)
36 (01:55)	0.00994	15.18	12.07 (1.11)	37.71	[15.42, 60.44]	11.39	2.99	0.4403	34.25 (1.40)

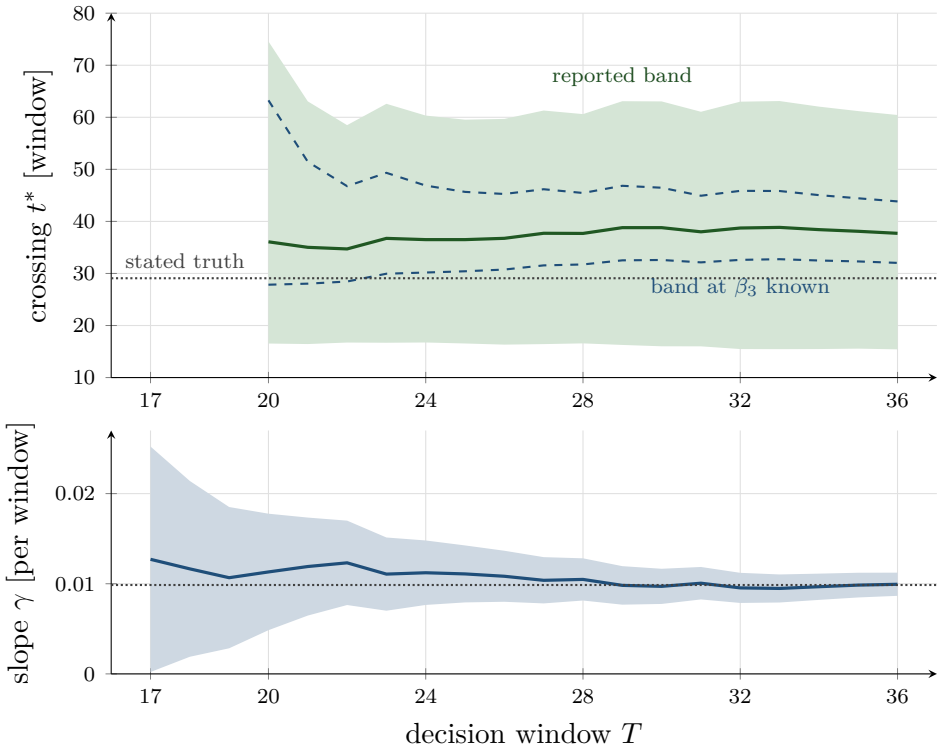


Figure 11.7: How the forecast firms. Top: the crossing's median and reported band from each decision window, and the band by inversion at β_3 known (dashed); the dotted line is the line's stated crossing, 29.08. Bottom: the slope and its 95 % band, the stated slope dotted. The band's upper end stands above the top axis before window 20 and is absent through window 16 (Table 11.5).

The span where it cannot yet answer. Through window 16 the band has no upper end: the normalized slope is 0.43, 1.52 and 1.31 at windows 14 to 16, and the first coefficient of equation (11.12) is negative. At windows 15 and 16 the draws of the normal put 0.0640 and 0.0947 on never crossing. The delta form still prints [6.3, 50.7] and [−1.9, 75.7] there. Window 17 is the first with an upper end, 1147.43 on the joint, 1144.24 on the threshold line and 1165.6 in draws of the normal, and the first from which z stays beyond 1.96; the upper end falls to 147.60 at 18, 107.70 at 19 and 74.57 at 20. The sample spread of crossing draws at window 16 is 133.5, 347.4, 4464.9 and 3535.5 at 10^3 to 10^6 draws: it does not settle, which is what no variance looks like. Here the flat coil standing past the limit has a mass that shows. The threshold line and the joint part by 1.08×10^{-3} at the lower end at window 14, moving that end from 17.17 to 17.27, by 9.79×10^{-4} at window 16, from 16.94 to 17.07, and by 3.89×10^{-4} at window 17, from 16.22 to 16.28; the medians at windows 14 to 16 move from 52.60, 28.49 and 36.90 to 52.63, 28.50 and 36.92. Normal draws with g^* linear give the joint's values within their error, 52.62, 28.50 and 36.88 at the median and 17.26, 16.14 and 17.02 at the lower end. The mass falls below 10^{-5} by window 20 and the two agree to every digit printed from window 21. The column of the line at T stays a single tail: it is the event of standing past the limit at a steady load, not of having crossed, and at window 14 it reads 0.0078 where the joint reads 0.0058. The sign passes 1.96 spreads at window 17, 00:20, before Chapter 9's flag at window 20. The two are not the same test: the trend reads the coil's entry, and its repeated looks are not corrected.

The median's wander. Over windows 20 to 36 the median moves between 34.71 and 38.85, a range of 4.14 windows, 20.7 minutes, and $\Pr(t^* \le 36)$ rises above one half at window 21 and falls below it for good at window 23: 0.4971, 0.5374, 0.5521 and 0.4728 at windows 20 to 23. That wander is the median's alone, apart from the band.

11.3.5 The form of the trend

Operator question. Is a straight line the right form? The answers above lean on it, and the operator can ask it of the record itself.

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

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

Specialized form. The question is equation (2.14) with M_A each form of equation (11.7) and M_B no change, the evidences those of equation (11.8) on the same rows. Its log ratio is reported at each record.

Building blocks. The forms' fits and evidences (Section 11.2.3).

The evidences. Table 11.6 gives the log evidence of each form against no change: each row a record, the columns the four forms and the ramp's onset. Through window 16 no form beats no change. A straight line over the whole record is wrong: through windows 1 to 24 it has the slope 0.00565 and a reading-row X^2 of 88.94 against the ramp's 74.97, on 120 reading rows, and it loses to the ramp by 5.83 log units. The line of the reported state is on the span from window 13, and it is the ramp's right arm if the onset is at or before 13; $\Pr(\tau < 13) = 0.7145$ and $\Pr(\tau < 16) = 0.9981$. Starting the span at window 16 instead moves the median by 0.84 window. Curvature is not supported: the curved ramp stands 0.99 log units below the ramp at

Table 11.6: Log evidence against no change on the whole record, and the onset under the ramp.

T (clock)	Line over the record [–]	Ramp [–]	Curved ramp [–]	Step [–]	τ median and band [window]
16 (00:15)	−0.93	−0.05	−0.30	−0.95	11.48 [1.83, 14.71]
18 (00:25)	2.85	4.36	4.09	2.83	11.88 [2.65, 15.34]
24 (00:55)	36.93	42.75	41.76	34.76	12.22 [8.71, 14.75]
36 (01:55)	229.20	241.02	238.32	195.27	11.73 [9.24, 13.92]

window 24. A step, which is a coil that dropped once, is well below the ramp from window 24 on.

Under the ramp itself at window 24, the mixture of Section 11.2.3 gives the slope 0.01082 ± 0.00190, the crossing's median 37.304, band [17.02, 62.76], and $\Pr(t^* \le 36) = 0.4526$: the median 0.800 window after the line's, on the same side of one half. A span reaching back into the reference carries two errors together, a line over windows where the coil is flat and windows 7 to 12 counted again against the baseline: from window 7 the slope is 0.00797, the onset 9.86, the crossing 44.66 and $\Pr(t^* \le 36) = 0.2742$.

11.4 The checks (step 5)

The answers of Section 11.3 rest on a line, a forecast load, the reduced state and approximations that can be tested on this hall, and the first test is the one the rest rests on.

11.4.1 The held-out test

From window 24 the four temperature readings of windows 25 to 36 are predicted at each window's metered load, with the state's covariance and the meter's error carried, and scored against a joint normal predictive. Table 11.7 gives unit 2's discharge window by window: each row a window, the columns the prediction, its spread, the reading and the standardized distance, reading minus prediction over the spread. The joint chi-square is $\chi^2 = 46.77$ on 48 rows, a tail of 0.523; 95.8 % of the 48 lie within 1.96; the largest $|z|$ is 3.16, unit 1's discharge at window 32. Unit 2's discharge has a mean z of −0.556 and a largest $|z|$ of 1.34. The line beats no change by 208.68 log units of predictive density and a line with curvature by 1.72. From every record in the printed set from $T = 15$ to $T = 33$ the tail lies between 0.187 and 0.746. The discharge first reads above 23°C at window 34, inside the reported band of the reading level; the slope at window 36 is 0.00994, −0.70 spreads from window 24's.

What the held-out test cannot test. The entry β_3, and the steady load, which the test replaces by the meter: the mean meter reading over windows 25 to 36 is 203.13 kW against $Q_f = 199.44$ kW ($z = 1.04$), with a sample spread of 4.32 kW against the wander's 3.454 kW. No later reading sees rack 3, so the readings validate the coil's line and leave the crossing's main term untested. Only the constructed truth can test it here, and in operation only a reading at rack 3 could.

11.4.2 The limit of exact rows and the curvature term

The fit runs on the reduced state. The check fits the stacked full state, in which each of the twelve windows adds T_r, $T_{d,1}$, $T_{d,2}$, T_c and T_1 with five physical rows, and lets the tolerance shrink through 0.1, 0.01 and 0.001. Table 11.8 carries the result. The spread of γ goes from 0.001836 to 0.001824, t^* from 36.5143 to 36.5044 and the reading rows' X^2 from 36.5010 to 38.0951, against the reduced state's 0.001824, 36.5044 and 38.0953. The covariance (11.6) is the

Table 11.7: The held-out test from window 24: unit 2's discharge predicted at each later window's metered load.

Window (clock)	Predicted [°C]	Spread [K]	Read [°C]	z [–]
25 (01:00)	21.903	0.366	21.58	−0.88
26 (01:05)	22.182	0.379	21.94	−0.64
27 (01:10)	22.087	0.388	21.82	−0.69
28 (01:15)	22.307	0.402	22.21	−0.24
29 (01:20)	22.661	0.420	22.10	−1.34
30 (01:25)	22.932	0.438	22.57	−0.83
31 (01:30)	22.931	0.451	22.97	0.09
32 (01:35)	23.480	0.478	22.99	−1.03
33 (01:40)	23.064	0.479	22.63	−0.91
34 (01:45)	23.183	0.495	23.12	−0.13
35 (01:50)	23.452	0.517	23.61	0.31
36 (01:55)	23.701	0.538	23.49	−0.39

Gauss–Newton form; the objective's full Hessian at $\hat{\mathbf{z}}$ gives $\mathrm{sd}(\gamma) = 0.001825$ against 0.001824, $\mathrm{sd}(Q_{24}) = 2.2597$ against 2.2602 kW and a delta spread of the crossing 10.4946 against 10.4940.

11.4.3 Sensitivity

Each choice of Section 11.1.4 is varied at window 24, the fit repeated. The prior spread of the slope at 0.005, 0.01, 0.05 and 0.5 gives the slope 0.00998 at the tightest and the median 38.75, 36.95, 36.38, 36.35. The prior spread of the level at 0.05 and 0.5 gives 36.55 and 36.50. The forecast load's wander at 0 and 6 kW gives delta spreads of 10.18 and 11.10 and the threshold line at window 24 0.1008 and 0.1222. A span from window 16, 19 or 21 gives the median of the joint 35.67, 38.46 or 55.48, against the threshold line's 35.67, 38.42 and 54.76, with delta spreads of 10.62, 14.25 or 62.94; from window 21 the onset is lost, −0.17 with a spread of 41.18.

The strict line is tested against departures, $g_{2,t} = \alpha - \gamma(t - \bar{t}) + d_t$ with one $d_t \sim \mathcal{N}(0, s^2)$ per window. The log evidence against the strict line is −0.03 at $s = 0.002$, −0.66 at 0.01 and −6.07 at 0.04; the strict line stands, and the median moves to 37.04 at 0.04, 0.53 window.

11.4.4 The answers against the stated truth

The record is generated from a stated truth (Section 11.6.4): a ramp with its kink at window 12, first changed window 13, 00:00, with $\gamma = 0.009861$ per window, and $\beta_3 = 0.57$ against the survey's 0.50, +0.875 prior spreads. The truth's load wanders, which the state does not model. The line of the truth reaches the limit at window 29.08 at 200 kW, and the true inlet is first over at window 29, 01:20, at 27.248°C, and stays over. Table 11.8 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 slope stands −0.75 spreads from the estimate, the level at window 24 0.52, and the onset's kink inside its reported band. Rack 3's inlet at window 24 is truly 26.691°C, 0.99 spreads above the estimate at Q_{24}. The crossing lies inside the reported band at every window from 14 on, with the threshold line placing 0.2338 of its probability before it at window 24. The median stands 7.43 windows, 37 minutes, late: the band is honest, and its center leans on a survey that is low.

The baseline as a prior stands in for the first hour's rows at the level of Table 11.8. The gaps of the exact posterior on the crossing exceed their standard errors by three to seven times, most of them g^*'s curvature (Section 11.3.3), and none crosses a level. On the onset and the reading level the medians' gaps are resolved and the ends' gaps are not: at this count the inversion's ends are not distinguished from the exact ones there.

Table 11.8: The checks and their gaps at window 24. Monte Carlo standard errors in parentheses: exact draws, importance sampling with 308 201 effective draws; the truth is the stated truth of the constructed record.

Route [–]	Reported	Beside it	Gap
Exact posterior against the reported value			
t^* median	36.504	36.592 (0.030)	+0.088
t^* lower end	16.743	17.095 (0.051)	+0.352
t^* upper end	60.327	60.913 (0.087)	+0.587
Over now at Q_{24}	0.0871	0.0866 (0.0005)	−0.0005
$\Pr(t^* \le 24)$, steady load	0.1083	0.1044 (0.0006)	−0.0039
$\Pr(t^* \le 36)$	0.4808	0.4775 (0.0009)	−0.0033
Onset median	12.638	12.649 (0.003)	+0.012
Onset lower end	9.380	9.375 (0.007)	−0.005
Onset upper end	14.594	14.599 (0.005)	+0.005
Reading level median	32.820	32.841 (0.004)	+0.022
Reading level lower end	28.807	28.814 (0.007)	+0.007
Reading level upper end	40.081	40.096 (0.028)	+0.014
sd(γ)	0.00182	0.00183	–
Approximations			
Full state, $\sigma_p = 0.001$, t^*	36.5044	36.5044	0.0000
Full state, $\sigma_p = 0.001$, reading rows' X^2	38.0953	38.0951	−0.0002
Newton against Gauss–Newton, sd(γ)	0.001824	0.001825	1.5×10^{-6}
Newton against Gauss–Newton, delta spread of t^*	10.4940	10.4946	0.0007
Baseline as prior against first hour as rows, γ	0.01122	0.01122	-4.6×10^{-8}
Baseline as prior against first hour as rows, t^*	36.504	36.504	below 0.001
Other forms			
Ramp, t^* median	36.504	37.304	+0.800
Ramp, $\Pr(t^* \le 36)$	0.4808	0.4526	−0.0282
Ramp, onset median	12.638	12.222	−0.416
Stated truth			
Slope γ	0.01122	0.009861	$z = -0.75$
$g_{2,24}$	−0.24472	−0.23834	$z = 0.52$
Onset (kink)	12.638	12	inside the band
$T_{3,24}$ at Q_{24} [°C]	25.865	26.691	$z = 0.99$
Crossing of the line [window]	36.504	29.08	inside the band
First window truly over	–	window 29, 01:20	–

11.5 The report (step 6)

The dashboard's date is the wrong one. The rack 1 line said 04:26. The plant's limit is at rack 3, unread, with a median crossing of 01:58, a band from 00:19 to 03:57, and a probability of 0.087 that it is over now; the one-hour crossing is a coin toss, 0.48, and the band will not tighten with waiting, because most of its variance is one survey.

11.5.1 The answers and their bands

Table 11.9 gives the answers in the order of Table 11.1. Each row is a question; the second column gives the reported answer and the third the exact value or the band beside it.

Table 11.9: The report at window 24, 00:55. Limit 27°C at rack 3 and horizon of one hour, both the operator's.

Question	Answer, reported	Beside it
Is it getting worse.	yes, $\Pr(\gamma > 0) = 1 - 3.75 \times 10^{-10}$, $z = 6.155$	exact: 1.00000000
How fast?	12.60 % of conductance per hour	band 8.77 to 16.27 %
When did it begin?	window 12.638, 23:58	band [9.380, 14.594], 23:42 to 00:08; exact median 12.649, [9.375, 14.599]
When does the plant cross its limit?	window 36.504, 01:58	band [16.74, 60.33], 00:19 to 03:57; exact median 36.592, [17.095, 60.913]; by the horizon 0.4808 (exact 0.4775)
Could it be over already?	0.0871 at the window's load	exact 0.0866; threshold line at a steady load 0.1083 (exact 0.1044)
How does it firm?	the slope firms; the crossing does not	delta spread 9.76 to 11.99 from window 20; no upper end through window 16
Is a line the right form?	on the span, if the onset is at or before window 13 (0.71); over the record, no	the span's start moves the median by 0.84; ramp over the line over the record 5.83 log units; curved ramp 0.99 below the ramp

11.5.2 The entry in the shift log

The operator's entry as of window 24, 00:55, reads: "Condition: coil of unit 2 (Chapter 10). Coil worsening, probability of a positive slope above 0.999 999 99, 12.6 % of conductance per hour (8.8 to 16.3). Began about 23:58, between 23:42 and 00:08. Unit 2's discharge reaches its 23°C alarm about 01:39 (band 01:19 to 02:15). Rack 3's inlet reaches 27°C about 01:58, band 00:19 to 03:57 (exact end window 60.9); rack 3 unread, probability 0.087 it is over now; one-hour crossing a coin toss (0.48); the band will not tighten with waiting. Action: not decided here."

Common mistake

Reading a band's lower end as a past date, or the crossing as a failure of equipment. The lower end says that rack 3 may already be over; the crossing is an inlet reaching a limit of the envelope, and nothing has stopped working.

11.6 The next reading (step 7)

The report leaves a crossing whose band is mostly one entry of the survey. A further reading is a further row, and what it would do to the entries is known before its value is.

11.6.1 The preposterior of a candidate

A candidate adds rows to $\mathbf{J}$ at the present MAP state, and the covariance after it is the preposterior of equation (4.65), on the state extended to the windows it waits for. Waiting six windows on the present readings adds those windows' rows. A handheld probe of standard deviation 0.3 K at rack 3 adds one row, the gradient of $T_{3,24}$ at the window's load over 0.3; it reads the inlet and so acts on β_3. A re-survey of β_3 with standard deviation 0.02 adds one prior row. A permanent inlet sensor at rack 3, the instrument at the place of the question, adds that row in windows 25 to 30. A second discharge sensor on unit 2 and a water-side heat meter on unit 2, of standard deviation 3 kW, add rows on the coil in the same windows.

11.6.2 The candidates

Table 11.10 compares them by what each does to the entries: each row a candidate, the columns the spreads of the slope, of the coil's level at window 24 and of β_3 after it, the crossing's delta spread and rack 3's inlet spread, and the width of the crossing's band by inversion in minutes.

Table 11.10: Candidate readings from the posterior at window 24.

Candidate	sd γ [per window]	sd $g_{2,24}$ [–]	sd β_3 [–]	sd t^* [window]	sd $T_{3,24}$ [K]	band [min]
No new reading, window 24	0.00182	0.01236	0.0800	10.49	0.858	217.9
Wait six windows	0.00101	0.00647	0.0800	10.16	0.854	202.4
Handheld probe at rack 3, now	0.00182	0.01234	0.0294	5.05	0.363	105.8
Re-survey of β_3, sd 0.02	0.00182	0.01236	0.0194	4.57	0.297	96.7
Permanent rack 3 inlet sensor, windows 25–30	0.00101	0.00647	0.0130	3.15	0.236	62.8
Second discharge sensor on unit 2, windows 25–30	0.00091	0.00555	0.0800	10.14	0.854	201.3
Water-side heat meter on unit 2, windows 25–30	0.00093	0.00574	0.0800	10.15	0.854	201.7

The instruments that read the coil better leave the crossing where it was: they halve the slope's spread, and the band moves from 217.9 to about 201 minutes, about a minute more than waiting's 202.4. A probe at rack 3 halves the band at once. The permanent sensor does most, and it is the only candidate that gives a held-out test of the crossing's main term in the windows that follow. The probe's result, in refits of the state with its row, moves the median by what it reads: at 26.0°C the median is 35.10; at 26.5°C it is 30.00; at the constructed truth's 26.691°C, with no sensor error, it is 28.06, band [19.66, 38.13], with $\Pr(t^* \le 36) = 0.9453$. The choice among them, and any action, stay with the operator. Chapter 18 ranks readings against a named worry.

11.6.3 What the analysis showed

The dashboard's line through rack 1 said three and a half hours; the plant's limit stands at a rack with no sensor, about an hour away at the median and possibly already passed. The slope is known long before the date is: it passes 1.96 spreads at window 17, while the crossing's band keeps its width to the end of the record. The onset's band narrows slowly toward a floor set by the reference's level. The crossing's band is a rack's survey, not the trend. And the held-out test passes without reaching that survey.

Parts I to III have asked what the plant is doing, how well that is known, what has changed, and where it is heading, each read from the posterior of the state on the two smaller plants. What is left open is deliberate. Repeated looks stand uncorrected. Every forecast of this chapter assumes that nobody acts: what happens if the operator does would need a scenario node, and none is computed here.

11.6.4 Limits of the model

The line across the span. A line assumes a steady rate from window 13 on. The forms of Section 11.3.5 support it on this record to window 36, and they do not say what a coil does beyond it. The crossing at window 60 lies far past the record.

One steady load. The forecast carries one load forward, and the inlet crosses on a wandering one. The stated truth's inlet is over at window 29, with its line at 29.08, on a load whose persistence the state does not model.

The survey unmoved. Nothing in the record moves β_3, so an error of the survey passes into the median in full, as it does in the truth here, and the fit has no row that would detect it.

The condition. Under a fault of the discharge sensor the rows of this chapter do not apply, as Section 11.1 says, and the slope would then describe a sensor.

Repeated looks. The table of each window reads the same record again and again. The sign at each window is a separate look, uncorrected, so no false-alarm rate is stated for the record as a whole.

The steady state. The limits of Chapter 5 hold in every window (Section 5.6.5).

Constructed readings. The record is constructed: generated from a stated truth of the hall with the stated sensor errors, and Section 11.4.4 compares the answers with that truth. It is not field data, and the agreement is a check of the method on its own model.

11.7 What to retain

Definition 11.1 (Trend and crossing time)**.** Windows stacked into one state carry a level α and a slope γ of the log-derate; the sign is a normal tail of γ, whose spread falls as the inverse root of $\sum_t (t - \bar{t})^2$. A level set on the plant, carried to the entry's scale, may depend on other entries of the state. A crossing $t^* = \bar{t} + (\alpha - g^*)/\gamma$ is a ratio with the slope in its denominator: it has no posterior variance. Its distribution function is the first passage of a worsening line, $\Pr(t^* \le s \mid \mathbf{y}) = \Pr(\gamma > 0,\ m(s) \ge 0 \mid \mathbf{y})$, the threshold line $\Phi(m(s)/\sqrt{\mathbf{g}_m^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g}_m})$ less the mass of $\gamma \le 0$ past the level. Its median and band solve that joint; once the mass is negligible the median is $t^*(\hat{\mathbf{z}})$ and Fieller's quadratic gives the ends. The band has an upper end only if the

slope's normalized value exceeds 1.96. The onset is the same ratio with the reference's level for g^*. Whether the limit is over now is a present-tense event, read at the window's own load.

Three conclusions carry beyond the hall. A rate is known sooner than a date, because a date also needs the level, and the level may rest on an entry no reading sees. A forecast firms only in the entries the next windows read, and the budget says which those are. And a held-out test of a forecast tests the rows the later readings act on, not the whole of the forecast: it passes without touching an entry that no reading sees.

Table 11.11 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 11.11: The terms of the trend and the crossing and where each is defined. Units for the two-unit hall.

Term	Expression	Units	Meaning	Defined in
Stacked state	eq. (11.1)	kW, –	one state over the span	11.1.2
Limit on the coil's scale	g^*, eq. (11.2)	–	the plant limit as a level of g_2	11.1.2
Reading level	g_d, eq. (11.3)	–	the alarm as a level of g_2	11.1.2
Trend prior	$\gamma/0.02$, $(\alpha - g_b)/0.15$	–	a proper state early	11.1.4
Forecast load	$(Q_f - \bar{Q})/3.454$	–	one steady load ahead	11.1.4
Posterior				
Square law	$\mathrm{sd}(\gamma) \approx$ one-window spread over $\sqrt{\sum_t (t-\bar{t})^2}$	per window	the slope's spread	11.2.2
Forms on the record	eq. (11.7)	–	no change, line, ramp, curved, step	11.2.3
Questions and answers				
Sign	$\Phi(\hat{\gamma}/\sqrt{\Sigma_{\gamma\gamma}})$	–	is it getting worse	11.3.1
Onset	τ, eq. (11.10)	window	when the line left the reference	11.3.2
Crossing time	t^*, eq. (11.13)	window	when the line reaches the level	11.3.3
First passage	eq. (11.14)	–	the crossing's distribution function	11.3.3
Threshold line	$\Phi(u_m(s))$, eq. (11.14)	–	its single tail, the same once $\gamma \le 0$ is negligible	11.3.3
Fieller's quadratic	eq. (11.12)	–	the ends of the band	11.3.2
Delta spread	eq. (11.15)	window	a width on one scale, split by entry	11.3.3
Spread with an entry known	$\boldsymbol{\Sigma} - \boldsymbol{\Sigma}\mathbf{e}\mathbf{e}^{\mathsf{T}}\boldsymbol{\Sigma}/\Sigma_{ee}$	window	what would remain	11.3.4
Log evidence of a form	eq. (11.8)	–	the form's support	11.2.3
Held-out test	χ^2 on later readings	–	the rows the later windows read	11.4.1
Candidate's band	preposterior, eq. (4.65)	min	the band after a reading	11.6.1

11.8 Exercises

Exercise 11.1. *Derivation.* Use the square law to predict the slope's spread on the span from window 13 to window 36, and the slope's normalized value at the slope of window 24. (Over that span $\sum_t (t-\bar{t})^2 = 1150$; with the one-window spread 0.02209 the spread is $0.02209/\sqrt{1150} = 0.00065$, as the fit's; at $\gamma = 0.01122$ the normalized value is 17.23, and at the fit's own slope 0.00994 it is 15.18.)

Exercise 11.2. *Application.* From $\hat{\alpha} = -0.18298$, $\hat{g}_b = -0.11718$, $\hat{\gamma} = 0.01122$, $\bar{t} = 18.5$ and $g^* = -0.38507$ compute the onset and the crossing in windows and in clock time. (By equations (11.10) and (11.13), $\tau = 18.5 + (-0.18298 + 0.11718)/0.01122 = 12.638$, 23:58, and $t^* = 18.5 + (-0.18298 + 0.38507)/0.01122 = 36.504$, 01:58.)

Exercise 11.3. *Derivation.* Derive the budget's share of β_3. (By the gradient (11.15) its entry is $-(1/\gamma)\,\partial g^*/\partial\beta_3 = -121.69$; with $\Sigma_{\beta_3\beta_3} = 0.08^2$ and no covariance with any other entry the term is 94.778, of the total 110.123: 86.1 %.)

Exercise 11.4. *Application.* A 0.3 K probe at rack 3 takes β_3 from a spread of 0.08 to 0.0294. Estimate what that does to the crossing's delta spread from the budget alone, and compare with the preposterior. (The term becomes $121.69^2 \times 0.0294^2 = 12.800$, the total 28.146, the spread 5.31. The preposterior gives 5.05, since the probe's row also acts on the other entries; the band by inversion falls from 217.9 to 105.8 minutes.)

Exercise 11.5. *Derivation.* Explain why the band has no upper end at window 16, and what the delta form prints there. (The first coefficient of equation (11.12) is $\hat{\gamma}^2 - c^2\Sigma_{\gamma\gamma}$, negative when the normalized slope 1.31 is below $c = 1.960$: at the 97.5 % level the record cannot exclude that the line never reaches the limit, with Pr(never) = 0.0947 in draws. The delta form still prints $[-1.9,\ 75.7]$ around the plug-in date 36.90.)

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)	$\Pr(T_{3,24} \ge 27°\mathrm{C} \mid \mathbf{y})$	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