

# Capacity Planning and Design by Derivation

Bottlenecks, margins and the cost of uncertainty in the design inputs

Draft · 272 pages



**Ricardo Marquez, Ph.D.**

ricardo-marquez.com

# Capacity Planning and Design by Derivation

Volume II

Bottlenecks, input uncertainty, failure and risk

Ricardo Marquez

October 31, 2026 (draft for review)

# How to read this book

This book teaches one skill: to take a planner's question stated in the language of the design review, write it as a probability or an expectation over the unknowns of a physical model, derive what that expression requires, compute it, check it, and decide on it. Volume I holds Parts I to III, Chapters 1 to 12: the method itself, the load a data hall can carry (headroom, growth, the design day, diversity), and how much to build (redundancy, unit size, the order date, oversizing). This second volume holds Parts IV to VII, Chapters 13 to 24. It asks where the site's bottleneck is and how much margin the plant really has, what the uncertainty in the design inputs costs and which measurement before committing is worth making, how the design holds when a unit fails, and what the design's expected cost and risk are. Chapter numbers run continuously across the two volumes; a reference to a chapter numbered 12 or below points to Volume I, whose results this volume builds on and does not repeat.

Every chapter follows the progression fixed in Volume I's preface and restated here in one line each: step 0, the plant graph; step 1, the question as a probability or expectation over the state; step 2, the graph extended for the question and the record attached as rows; step 3, the posterior of the state, derived and computed once; step 4, the answers, each question read from that posterior down to its number; step 5, evaluation and checks against the exact posterior and the constructed truth; step 6, the report with its band; step 7, the next measurement, ranked by the width it removes. The five kinds of question and the five computations beneath them are those of Volume I.

The plant of this volume is the site of Volume I's Part III: two data halls on one chilled-water plant with three chillers, one pump, a shared supply pipe, a power feed and a drawn economizer, calibrated on a year of operating records, capacity tests and an outage record. The notation is Volume I's.

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.

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**

**IV Bottlenecks and Margins** **1**

**13 Binding constraints** **2**
13.1 Which element of the site binds first . . . 2
13.2 The posterior of the state (step 3) . . . 8
13.3 The answers (step 4) . . . 9
13.4 The checks (step 5) . . . 16
13.5 The report (step 6) . . . 17
13.6 The next measurement (step 7) . . . 18
13.7 What to retain . . . 20
13.8 Exercises . . . 21

**14 Capacity importance measures** **22**
14.1 Which upgrade buys the most headroom . . . 22
14.2 The posterior of the state (step 3) . . . 26
14.3 The answers (step 4) . . . 27
14.4 The checks (step 5) . . . 34
14.5 The report (step 6) . . . 35
14.6 The next measurement (step 7) . . . 36
14.7 What to retain . . . 37
14.8 Exercises . . . 38

**15 Shared paths and common-mode limits** **40**
15.1 One pipe under two halls . . . 40
15.2 The posterior of the state (step 3) . . . 46
15.3 The answers (step 4) . . . 49
15.4 The checks (step 5) . . . 56
15.5 The report (step 6) . . . 58
15.6 The next measurement (step 7) . . . 59
15.7 What to retain . . . 61
15.8 Exercises . . . 62

**16 As-built margin** **64**
16.1 Twenty percent on the drawing . . . 64
16.2 The posterior of the state (step 3) . . . 68
16.3 The answers (step 4) . . . 69
16.4 The checks (step 5) . . . 75
16.5 The report (step 6) . . . 77

16.6 The next measurement (step 7) . . . 78
16.7 What to retain . . . 81
16.8 Exercises . . . 82

**V The Cost of Input Uncertainty 83**

**17 Variance budgets for design parameters 84**
17.1 How much of the band one handbook guess carries . . . 84
17.2 The posterior of the state (step 3) . . . 89
17.3 The answers (step 4) . . . 90
17.4 The checks (step 5) . . . 98
17.5 The report (step 6) . . . 99
17.6 The next measurement (step 7) . . . 100
17.7 What to retain . . . 102
17.8 Exercises . . . 103

**18 Measurements before commitment 105**
18.1 Which test narrows the build decision . . . 105
18.2 The posterior of the state and the deciding quantity (step 3) . . . 110
18.3 The answers (step 4) . . . 111
18.4 The checks (step 5) . . . 117
18.5 The report (step 6) . . . 118
18.6 The next measurement (step 7) . . . 120
18.7 What to retain . . . 122
18.8 Exercises . . . 123

**19 The value of a pilot 125**
19.1 A pilot before the full build . . . 125
19.2 The posterior of the state (step 3) . . . 130
19.3 The answers (step 4) . . . 131
19.4 The checks (step 5) . . . 139
19.5 The report (step 6) . . . 140
19.6 The next measurement (step 7) . . . 141
19.7 What to retain . . . 143
19.8 Exercises . . . 144

**VI Design Under Failure 146**

**20 Headroom with the largest unit out 147**
20.1 Hall B with its largest unit out . . . 147
20.2 The posterior of the state (step 3) . . . 152
20.3 The answers (step 4) . . . 153
20.4 The checks (step 5) . . . 162
20.5 The report (step 6) . . . 163
20.6 The next measurement (step 7) . . . 165
20.7 What to retain . . . 167
20.8 Exercises . . . 168

**21 Cascading overloads** **169**
21.1 A chiller trip on the design day . . . 169
21.2 The posterior of the state (step 3) . . . 175
21.3 The answers (step 4) . . . 176
21.4 The checks (step 5) . . . 185
21.5 The report (step 6) . . . 186
21.6 The next measurement (step 7) . . . 187
21.7 What to retain . . . 189
21.8 Exercises . . . 190

**22 Concurrent maintainability** **192**
22.1 A maintenance calendar for the site . . . 192
22.2 The posterior of the state (step 3) . . . 198
22.3 The answers (step 4) . . . 199
22.4 The checks (step 5) . . . 206
22.5 The report (step 6) . . . 207
22.6 The next measurement (step 7) . . . 209
22.7 What to retain . . . 211
22.8 Exercises . . . 212

**VII Expected Cost and Risk** **214**

**23 Expected unserved energy** **215**
23.1 The site's shed energy over a year . . . 215
23.2 The posterior of the state (step 3) . . . 220
23.3 The answers (step 4) . . . 222
23.4 The checks (step 5) . . . 229
23.5 The report (step 6) . . . 230
23.6 The next measurement (step 7) . . . 231
23.7 What to retain . . . 233
23.8 Exercises . . . 234

**24 Risk–cost tradeoffs in design** **236**
24.1 The design review's list of additions . . . 236
24.2 The posterior of the state (step 3) . . . 242
24.3 The answers (step 4) . . . 243
24.4 The checks (step 5) . . . 250
24.5 The report (step 6) . . . 251
24.6 The next measurement (step 7) . . . 253
24.7 What to retain . . . 256
24.8 Exercises . . . 257

**A Atlas of planning questions** **259**

# Part IV

# Bottlenecks and Margins

# Chapter 13

# Binding Constraints

## 13.1 Which element of the site binds first

The site's chilled-water plant has four elements that can stop the load from growing. The three chillers, with the economizer the drawing adds, must carry the heat on the loop; the shared supply pipe must pass the loop's flow within its allowable pressure drop of 150 kPa; the pump must drive that flow within its 12 kW motor; and the power feed must carry the IT load, the chillers and the pump within its 1 600 kVA at power factor 0.95. The design sheet works out, for each, the site IT load at which it reaches its limit on the design day, from nameplate ratings and handbook values. The pipe binds at 1 001.9 kW, the pump at 1 155.2 kW, the feed at 1 202.6 kW and the chillers at 1 228.0 kW. Its order is pipe, pump, feed, chillers, and its recommendation follows from the first entry: the pipe binds first, so upgrade the pipe.

The four numbers come from four different handbooks, a flow coefficient, a pump efficiency, a chiller efficiency and the economizer's conductance, and none of them carries a band. Two of them sit 25 kW apart, the feed and the chillers; the sheet's first and second are 153 kW apart, which looks safe, but only if the pipe's handbook coefficient is right. The question of this chapter is the drill question, verbatim,

*Which element binds first as the load grows: the chiller, the pump, the pipe, or the power feed?*

It is answered on the site's one fit of Chapter 9, with the binding load of Section 5.1.1 solved for each element in turn. What is new is that the elements are unlike: each binding load comes from its own physics and its own entries of the state, and the answer is not one binding load with a band but an order of four, which is itself a derived quantity of the state, with a probability for each order it can take.

**Jargon: binding constraint, binding element, ranking of the elements, order flip, pairwise gap, debottlenecking, allowable pressure drop, flow coefficient, motor rating, apparent power, power factor, auxiliary load, economizer duty, handbook value, Monte Carlo standard error, parameter risk (variance budget)**

The *binding constraint* of a plant at a growing load is the limit that is met first, and the element it belongs to is the *binding element*. The *ranking of the elements* is the order of their binding loads, the smallest first; an *order flip* is a state of the plant at which that order differs from the order at the MAP state. The *pairwise gap* of two elements is the difference of their binding loads; its sign says which binds first. *Debottlenecking* is the raising of the binding element's limit so that another element binds; what it buys is the gap to the next element. The *allowable pressure drop* of a pipe is the largest drop its design permits, here 150 kPa across the shared supply pipe. The pipe's *flow coefficient* $k_p$ is the flow, in L/s, that drops one kilopascal; the drop rises with the square of the flow. The *motor rating* of a pump is the shaft power its motor can deliver, here 12 kW. The *apparent power* of a feed, in kVA, is the product of its voltage and current; the real power it can deliver is that times the *power factor*, here $1\,600 \times 0.95 = 1\,520$ kW. The *auxiliary load* is the power drawn by UPS losses and lighting in proportion to the IT load. The *economizer duty* $Q_e$ is the heat the economizer takes from the loop without the chillers' compressors. A *handbook value* is a rating taken from a catalog or a design guide rather than from the plant. The *Monte Carlo standard error* of a probability estimated as a share of $N$ draws is the standard deviation of that share between independent sets of draws. The *parameter risk* of a derived quantity, its *variance budget* in the framework, splits its variance into one share per entry of the state.

**Key idea**

A binding constraint is a ranking of binding loads, and a ranking is a derived quantity of the state. Each element's binding load is a scenario solved for the load at which its limit is met, with its own gradient and its own band. The element that binds first, and the probability that the order changes, are read jointly from draws of the posterior of the state pushed through every element's binding load. When two elements bind within each other's bands, the order is decided by whatever entry carries the band of their gap, and that entry need not belong to either element's own record.

### 13.1.1 Binding loads of unlike elements

Chapter 5 solved the scenario node for the load at which a rack inlet meets its limit, and took the hall's binding load as the smallest over its racks. The racks were alike: every inlet rose with the load by the same physics, and they differed only in one fraction. The site's elements are unlike. For each element $e$ write $h_e(\mathbf{z}, Q^*)$ for the quantity its limit constrains at the scenario node, as a function of the free entries $\mathbf{z}$ and of the pinned site IT load $Q^*$, and $h_e^{\max}$ for the limit. The binding load $Q_e^\dagger$ of element $e$ is the load at which the two meet,

$$h_e(\mathbf{z},\ Q_e^\dagger(\mathbf{z})) = h_e^{\max}, \qquad \mathbf{g}_e = \nabla_{\mathbf{z}} Q_e^\dagger = -\left.\frac{\nabla_{\mathbf{z}} h_e}{\partial h_e / \partial Q^*}\right|_{(\hat{\mathbf{z}},\ Q_e^\dagger(\hat{\mathbf{z}}))}, \tag{13.1}$$

the implicit-function form of equation (5.1), with $h_e$ in the element's own unit: kilowatts of heat for the chillers, kilopascals for the pipe, kilowatts of shaft power for the pump and of real power for the feed. The division by the sensitivity to the load puts every gradient in kilowatts of site IT load, so that the four binding loads stand on one axis and can be compared.

On that axis the plant's binding load is the smallest of the four, and the binding element is the one that attains it. Writing $\pi(\mathbf{z})$ for the ranking, the order of the four elements by their

binding loads at the state $\mathbf{z}$, the smallest first,

$$Q^\dagger(\mathbf{z}) = \min_e Q_e^\dagger(\mathbf{z}), \qquad e^\dagger(\mathbf{z}) = \arg\min_e Q_e^\dagger(\mathbf{z}), \qquad \pi(\mathbf{z}) = \operatorname{argsort}_e Q_e^\dagger(\mathbf{z}), \tag{13.2}$$

three derived quantities of the state: a load, an element and an order. The first has a gradient wherever one element binds alone; the second and the third are constant over regions of the state and jump where two binding loads tie, so they have no gradient at all, and their probabilities are volumes of those regions under the posterior. The order at the MAP state, $\pi(\hat{\mathbf{z}})$, is the one a calibration reports; an order flip is the event $\pi(\mathbf{z}) \neq \pi(\hat{\mathbf{z}})$. For two elements $a$ and $b$ the pairwise gap

$$D_{ab}(\mathbf{z}) = Q_a^\dagger(\mathbf{z}) - Q_b^\dagger(\mathbf{z}), \qquad \mathbf{g}_D = \mathbf{g}_a - \mathbf{g}_b, \tag{13.3}$$

has a gradient wherever both have one, and its sign is the order of the pair: $a$ binds before $b$ exactly when $D_{ab} < 0$. Its variance budget says which entries of the state decide that order.

**Principle 13.1.** *Each element's binding load is its own scenario solved for the load, carried to the posterior by the implicit-function gradient. The binding element and the ranking have no gradient; they are read as probabilities over draws of the state. The order of a pair is the sign of their gap, and the gap's variance budget names the entries that decide it.*

### 13.1.2 The site and its four elements (step 0)

The plant is the site of Chapter 9: the two halls on one chilled-water plant, with the three existing chillers C1–C3 of 400 kW nameplate at the design day, the drawn economizer, one pump, the shared supply pipe and the power feed. Its physical relations are equation (9.3), and its state is the eight entries of equation (9.4),

$$\mathbf{z} = (\kappa_{C1},\ \kappa_{C2},\ \kappa_{C3},\ d,\ g_e,\ c,\ l_k,\ l_e)^\mathsf{T},$$

the chillers' as-built log-ratios, their derate per kelvin of wet bulb, the economizer's log-ratio to its handbook conductance, the log of the chillers' efficiency, and the logs of the pipe's flow coefficient and the pump's efficiency. Chapter 9 used the first element alone; this chapter sets all four against each other.

**The scenario node.** Three inputs are pinned, each a modeling choice. The wet bulb is pinned at the design day's $T_{\text{wb}}^* = 22°\text{C}$: this makes the four binding loads statements about the day the drawing is sized for, and it fails as a statement about any other day. Every unit runs: this is the plant with all three chillers up, and a plant with a unit out, or with the shared pipe out of service, is Chapter 15's question. The plant is the plant as drawn, the three existing chillers and the drawn economizer: this is what the design sheet ranks, and Section 13.4.1 sets the plant as built today, without the economizer, beside it. The site IT load $Q^*$ is free, and the loop's heat is $Q^* + Q_o$ with the halls' other heat $Q_o = 22\,\text{kW}$.

**The graph.** The site with its four elements is drawn in Figure 13.1. The halls on the left put their heat into the loop, which the pump drives through the shared pipe to the three chillers and the drawn economizer; the feed below carries the IT load, the chillers' power and the pump's. Each of the four elements is drawn with a red border and the limit that binds it, and each is labeled with the entries of the state its binding load depends on. The dashed arrows leave the graph for the operating record of Chapter 9; the economizer has none.

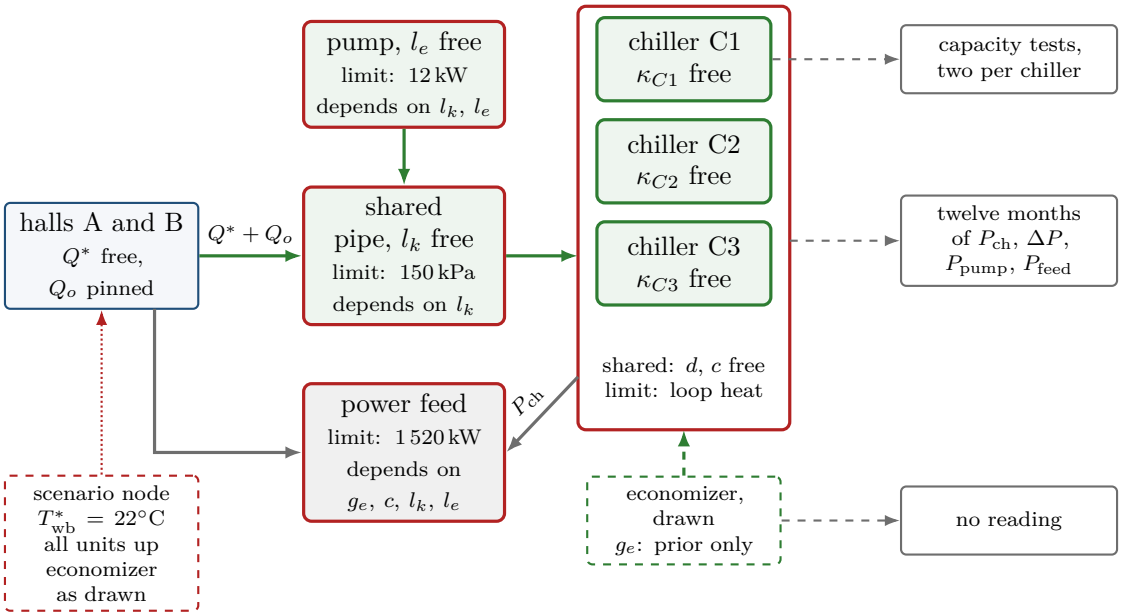



Figure 13.1: The site at the design day with its four elements (red borders), each labeled with its limit and with the entries of the state its binding load depends on. Free (green): the three chillers' as-built log-ratios with the derate $d$ and the efficiency $c$ they share, the pipe's log flow coefficient $l_k$ and the pump's log-efficiency $l_e$. The economizer (dashed) is drawn but not built; its log-ratio $g_e$ keeps the handbook prior, and no reading sees it. Gray: the power the feed carries, the IT load with its auxiliary load and the chillers' and the pump's power. Observed (dashed arrows): the capacity tests and the twelve monthly windows of Chapter 9. The dotted arrow marks the scenario node (red, dashed), which pins the wet bulb and the units and frees the site IT load $Q^*$.

**Modeling choices at the graph.** The plant is taken in steady state at the design day, as in Chapter 9: this makes each limit an algebraic relation in the load, and it fails in the minutes after a load step. The chilled-water supply and the loop's rise are pinned at 14°C and 10 K: this makes the loop's flow proportional to its heat, and it fails if the plant resets the supply or lets the rise grow as the load grows. The pump has one efficiency at every flow and the halls' own pressure drop is pinned at 100 kPa: this keeps the pump to one entry, and it fails where the pump's efficiency falls away from its best point or the halls' valves open as the load grows. The feed's power factor is pinned at 0.95: this makes its limit a real power, and it fails if the power factor falls with the load.

### 13.1.3 The four binding loads

At the scenario node the loop's flow is $V^* = (Q^* + Q_o)/(\rho c_p\,\Delta T_w)$ in L/s, the pipe's drop is $\Delta P^* = (V^*/k_p)^2$ and the pump's power is $P^*_{\text{pump}} = (\Delta P^* + \Delta P_{\text{h}})\,V^*/(1000\,\eta)$, by equation (9.3). At the design day the derate factor is one, the economizer's duty is $Q_e = 50\,\mathrm{e}^{g_e}$ kW by equation (9.5), and the chillers' efficiency is $\text{COP} = \mathrm{e}^{c}\,[1 - 0.02\,(22 - 18)] = 0.92\,\mathrm{e}^{c}$. Each element's limit, set equal to its value at the load, gives its binding load.

**The chillers.** The chillers with the economizer carry the loop's heat while $Q^* + Q_o \le \sum_i \text{Cap}_i + Q_e$. The limit is met at

$$Q^{\dagger}_{\text{ch}}(\mathbf{z}) = \sum_i 400\,\mathrm{e}^{\kappa_i} + 50\,\mathrm{e}^{g_e} - Q_o, \qquad \frac{\partial Q^{\dagger}_{\text{ch}}}{\partial \kappa_i} = 400\,\mathrm{e}^{\kappa_i}, \qquad \frac{\partial Q^{\dagger}_{\text{ch}}}{\partial g_e} = 50\,\mathrm{e}^{g_e} = Q_e, \tag{13.4}$$

explicit in the state: the heat side of the limit does not depend on the state, so the sensitivity in equation (13.1) is one. The gradient's entries are each unit's as-built capacity and the economizer's duty, in kW; a log-ratio moves the binding load by the capacity it scales.

**The pipe.** The pipe's drop meets its allowable $\Delta P_{\max} = 150$ kPa when the flow reaches $k_p\sqrt{\Delta P_{\max}}$, so that

$$Q^{\dagger}_{\text{pipe}}(\mathbf{z}) = \rho c_p\,\Delta T_w\,\sqrt{\Delta P_{\max}}\,\mathrm{e}^{l_k} - Q_o, \qquad \frac{\partial Q^{\dagger}_{\text{pipe}}}{\partial l_k} = Q^{\dagger}_{\text{pipe}} + Q_o, \tag{13.5}$$

the loop's heat at the limit, proportional to the flow coefficient, and every other entry zero. The pipe's binding load is a function of one entry, monotone in it.

**The pump.** The pump meets its motor rating $P_m = 12$ kW at the flow $V^{\dagger}$ that solves $h(V, \mathbf{z}) = P_m$ with $h = (V^3/k_p^2 + \Delta P_{\text{h}} V)/(1000\,\eta)$, a cubic in the flow with one positive root. Its sensitivities are $\partial h/\partial V = (3\Delta P + \Delta P_{\text{h}})/(1000\,\eta)$, $\partial h/\partial l_k = -2\,\Delta P\,V/(1000\,\eta)$ and $\partial h/\partial l_e = -h = -P_m$ at the root, with $\Delta P$ the pipe's drop at $V^{\dagger}$, so equation (13.1), with the load entering through the flow, gives

$$\begin{gathered} Q^{\dagger}_{\text{pump}} = \rho c_p\,\Delta T_w\,V^{\dagger} - Q_o, \\ \frac{\partial Q^{\dagger}_{\text{pump}}}{\partial l_k} = \rho c_p\,\Delta T_w\,\frac{2\,\Delta P\,V^{\dagger}}{3\,\Delta P + \Delta P_{\text{h}}}, \qquad \frac{\partial Q^{\dagger}_{\text{pump}}}{\partial l_e} = \rho c_p\,\Delta T_w\,\frac{(\Delta P + \Delta P_{\text{h}})\,V^{\dagger}}{3\,\Delta P + \Delta P_{\text{h}}}, \end{gathered} \tag{13.6}$$

both positive: a larger pipe or a better pump lets the pump drive more flow before its motor binds. The pump shares the pipe's entry $l_k$, since the pipe's drop is part of the head the pump supplies.

**The feed.** The feed meets its real power $S\cos\varphi = 1\,600 \times 0.95 = 1\,520\,\text{kW}$ where

$$F(Q^*, \mathbf{z}) = (1+a)\,Q^* + \frac{Q^* + Q_o - Q_e}{\text{COP}} + P^*_{\text{pump}} - S\cos\varphi = 0,$$

the IT load with its auxiliary fraction $a = 0.06$, the chillers' power for the heat the economizer leaves them, and the pump's power. The sensitivity to the load is $\partial F/\partial Q^* = 1 + a + 1/\text{COP} + \partial P^*_{\text{pump}}/\partial Q^*$, and equation (13.1) gives

$$\begin{aligned} \frac{\partial Q^\dagger_{\text{feed}}}{\partial c} &= \frac{P_{\text{ch}}}{\partial F/\partial Q^*}, \qquad \frac{\partial Q^\dagger_{\text{feed}}}{\partial g_e} = \frac{Q_e/\text{COP}}{\partial F/\partial Q^*}, \\ \frac{\partial Q^\dagger_{\text{feed}}}{\partial l_k} &= \frac{2\,\Delta P\, V/(1000\,\eta)}{\partial F/\partial Q^*}, \qquad \frac{\partial Q^\dagger_{\text{feed}}}{\partial l_e} = \frac{P_{\text{pump}}}{\partial F/\partial Q^*}, \end{aligned} \tag{13.7}$$

every quantity at the binding load. A more efficient chiller plant, a larger economizer, a larger pipe or a better pump each lowers the power drawn at a given load and raises the load at which the feed binds; each entry of the feed's gradient is the power that entry saves, divided by the power one more kilowatt of IT load costs.

**The pattern.** Read together, equations (13.4) to (13.7) fix which entries each binding load depends on before any number. With $*$ for an entry that can be nonzero, the four gradients have the pattern

$$\begin{bmatrix} \mathbf{g}^{\mathsf{T}}_{\text{ch}} \\ \mathbf{g}^{\mathsf{T}}_{\text{pipe}} \\ \mathbf{g}^{\mathsf{T}}_{\text{pump}} \\ \mathbf{g}^{\mathsf{T}}_{\text{feed}} \end{bmatrix} = \begin{bmatrix} * & * & * & 0 & * & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & * & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & * & * \\ 0 & 0 & 0 & 0 & * & * & * & * \end{bmatrix}, \tag{13.8}$$

columns in the order of equation (9.4). Three facts are read here. The derate $d$ has no entry: at the design day its factor is one. The chillers and the pipe share no column, so their gap's gradient $\mathbf{g}_{\text{pipe}} - \mathbf{g}_{\text{ch}}$ has the chillers' entries with their signs reversed and the pipe's entry beside them, and its variance is the sum of the two binding loads' variances and of the covariances between their entries. And the economizer's $g_e$, which no reading sees, sits in the chillers' row and the feed's: the drawn economizer moves two of the four binding loads.

### 13.1.4 The planner's questions (step 1)

The planner asks four questions of the site as drawn, and step 1 writes each as a probability or a variance over the state, with the evidence on the right of the bar. Table 13.1 lists them. The first column is the question in the planner's words, the drill question first; the second the general form of its category in the notation of Chapter 2; the third its mathematical form for the site; 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 drill question and the third are predictions under a threshold over a family of events, the binding element and the order; the second and fourth are estimation, the four binding loads with their bands and the split of the pipe-against-chillers gap. The answers are given in the order two, one, three, four, because the binding element is read from the binding loads.

### 13.1.5 The evidence as rows (step 2)

Step 2 attaches the operating record of Chapter 9 to the graph, unchanged: two capacity tests per chiller, twelve monthly windows of chiller power, pipe pressure drop, pump power and feed power, and eight prior rows, 62 rows on the 8 entries (Section 9.1.5). The outage record is not used: every unit runs at this scenario node, so the rates of failure and repair do not enter.

Table 13.1: The questions asked of the site at the design day, with all units up and the drawn economizer. $\mathbf{y}$ is the operating record and the priors of Chapter 9; $Q_e^\dagger$, $\pi$ and $D$ are the binding loads, the ranking and the pairwise gap of equations (13.1) to (13.3).

| Planner's question | General form | Mathematical form | Category | Computation | Answered in |
|---|---|---|---|---|---|
| Which element binds first as the load grows: the chiller, the pump, the pipe, or the power feed? | $\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y})$ over a family | $\Pr(e^\dagger(\mathbf{z}) = j \mid \mathbf{y})$, $j$ the chillers, the pipe, the pump, the feed | prediction under a threshold | posterior | 13.3.2 |
| At what load does each bind, and how sure? | $\Pr(f(\mathbf{z}) \leq f^* \mid \mathbf{y})$, quantiles | $Q_e^\dagger$ per element and $Q^\dagger = \min_e Q_e^\dagger$: median and 90 % band | estimation | posterior | 13.3.1 |
| What is the probability that the order flips? | $\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y})$ over a family | $\Pr(\pi(\mathbf{z}) \neq \pi(\hat{\mathbf{z}}) \mid \mathbf{y})$, and each order's probability | prediction under a threshold | posterior | 13.3.3 |
| What decides between the pipe and the chillers? | $\mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y})$, split by entry | the budget of $\mathbf{g}_D^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}_D$, $D = Q_{\text{pipe}}^\dagger - Q_{\text{ch}}^\dagger$ | estimation | posterior | 13.3.4 |

**Which row acts on which entry.** By the pattern (13.8) each binding load is read through a different part of the record. The pipe's entry $l_k$ is read by every month's pressure drop, twelve rows that see it and nearly nothing else. The pump's $l_e$ is read by the monthly pump power, beside $l_k$. The feed's entries are read by the monthly chiller and feed power, through the efficiency $c$. The chillers' log-ratios are read by the six capacity tests. The economizer's $g_e$ is read by nothing: the economizer is not built, no month ran it, and its row is the handbook prior alone, $g_e \sim \mathcal{N}(0,\ 0.35^2)$. Two of the four binding loads, the chillers' and the feed's, carry that prior through their gradients.

**Modeling choices at the rows.** The rows are those of Chapter 9, with its choices: each reading Gaussian, unbiased and independent given the state, and each prior independent. One choice matters more here than there. The economizer's prior is the whole of what the fit knows of it, so every answer that involves the chillers inherits whatever the handbook got wrong; the record cannot correct it, and Section 13.4.1 shows what the answers become without it.

## 13.2 The posterior of the state (step 3)

All four questions of Table 13.1 are read from one posterior, the site's one fit with its draws, as Chapter 9 builds them. No row is added and no fit is repeated; this section cites the posterior and the draws the answers use.

### 13.2.1 The fit of the site, cited

The MAP state $\hat{\mathbf{z}}$ of the eight entries and the Laplace covariance $\boldsymbol{\Sigma} = (\mathbf{J}^\mathsf{T}\mathbf{J})^{-1}$ at $\hat{\mathbf{z}}$ are those of Section 9.2.2, the Gauss–Newton fit of the operating record in the form of equation (4.38), tabulated with the stated truth in Table 9.2. Four of its facts carry this chapter. The chillers' log-ratios are $-0.0909$, $-0.0394$ and $-0.1157$, each to about 0.017, so that the units deliver 365.2, 384.5 and 356.3 kW at the design day. The pipe's $l_k$ is 0.8083 with posterior spread

0.0046, 0.015 of its prior spread, and above the handbook's $\ln 2 = 0.6931$: the pipe passes more flow than the handbook says. The pump's $l_e$ is $-0.2891$ with spread 0.0190, and the efficiency $c$ is 1.6176 with spread 0.0061. The economizer's $g_e$ stays at its prior, 0 with spread 0.35, uncorrelated with every other entry, since no row sees it.

### 13.2.2 The draws of the state

The answers read the posterior through the draws of Section 9.2.4: $N = 20\,000$ states $\mathbf{z}_i$ from the Laplace normal $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$, with the seed 9, and their weights $w_i$ of equation (5.16), which reweight them to the exact posterior of the operating record with an effective sample size of 19 913. Each draw is pushed through the four binding loads, the chillers' and the pipe's by equations (13.4) and (13.5), the pump's by solving its cubic for the flow, and the feed's by solving $F = 0$ for the load, which gives one row $(Q^\dagger_{\text{ch}}, Q^\dagger_{\text{pipe}}, Q^\dagger_{\text{pump}}, Q^\dagger_{\text{feed}})_i$ per draw, and with it the binding element and the ranking of that draw. Every answer of step 4 is a function of these rows.

## 13.3 The answers (step 4)

Every answer is read from the draws of Section 13.2.2. The reported value is the share or the quantile over the equal-weight draws, the joint reading over the four elements; the exact posterior stands beside it with the gap; and the delta form on each element, from its gradient at $\hat{\mathbf{z}}$, stands beside both. The binding loads come first, because the binding element is read from them.

### 13.3.1 Each element's binding load

**Planner's question.** At what load does each bind, and how sure? The question asks for the four binding loads, each with its band, and for the plant's binding load, the smallest of them.

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

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

summarized by its median and its central 90 % credible interval, equation (2.11), between the 5 % and the 95 % points.

**Specialized form.** The question is equation (2.9) with $f(\mathbf{z}) = Q^\dagger_e(\mathbf{z})$ for each element, by equations (13.4) to (13.7), and with $f(\mathbf{z}) = Q^\dagger(\mathbf{z}) = \min_e Q^\dagger_e(\mathbf{z})$ of equation (13.2) for the plant. The reported quantiles are those of the draws,

$$\hat{F}_e(q) = \frac{1}{N}\sum_{i=1}^{N} \mathbf{1}\big[Q^\dagger_e(\mathbf{z}_i) \le q\big], \qquad \hat{F}(q) = \frac{1}{N}\sum_{i=1}^{N} \mathbf{1}\Big[\min_e Q^\dagger_e(\mathbf{z}_i) \le q\Big], \tag{13.9}$$

and the exact posterior's are those of the same draws under the weights $w_i$. The delta form of each element is the normal $\mathcal{N}(Q^\dagger_e(\hat{\mathbf{z}}),\ \mathbf{g}_e^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_e)$ with the gradient of equation (13.1), and its band is $Q^\dagger_e(\hat{\mathbf{z}}) \pm 1.645\,(\mathbf{g}_e^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_e)^{1/2}$. For the pipe, a function of the one normal entry $l_k$ and monotone in it, the quantiles of the draws are the inversion band of $l_k$ carried through equation (13.5).

**Building blocks.** The draws and their weights (Section 13.2.2); the binding loads (13.4) to (13.7) at $\hat{\mathbf{z}}$ and in each draw; the gradients of the same equations; the covariance $\boldsymbol{\Sigma}$ of Section 13.2.1.

Table 13.2: The binding loads of the site's four elements at the design day, in site IT load: the design sheet's, the MAP value and the delta-form standard deviation, the median and 90 % band of the 20 000 draws, the exact posterior's median and the stated truth.

| Element | Sheet [kW] | $Q_e^{\dagger}(\hat{\mathbf{z}})$ [kW] | Delta sd [kW] | Median [kW] | 5 % point [kW] | 95 % point [kW] | Exact median [kW] | Truth [kW] |
|---|---|---|---|---|---|---|---|---|
| Chillers | 1 228.0 | 1 134.06 | 22.36 | 1 135.67 | 1 102.58 | 1 178.75 | 1 135.14 | 1 106.00 |
| Pipe | 1 001.9 | 1 126.81 | 5.29 | 1 126.82 | 1 118.12 | 1 135.68 | 1 126.86 | 1 127.31 |
| Feed | 1 202.6 | 1 188.23 | 3.15 | 1 188.31 | 1 184.05 | 1 194.90 | 1 188.32 | 1 184.66 |
| Pump | 1 155.2 | 1 272.00 | 10.66 | 1 272.08 | 1 254.48 | 1 289.70 | 1 272.31 | 1 250.37 |

**The four binding loads.** Table 13.2 gives, for each element, the design sheet's value, the binding load at $\hat{\mathbf{z}}$ with its delta-form standard deviation, the median and the ends of the 90 % band by the draws, the exact posterior's median, and the stated truth of the record, a check marker for Section 13.4.3. The four posteriors are drawn on one axis in Figure 13.2: the horizontal axis is the site IT load, the vertical axis the density of the delta form of each element; the bars above the curves are the 90 % bands of the draws, with the plant's joint band at the top; the dash-dotted vertical lines at the bottom are the design sheet's four values, and the black dots on the bars the stated truth.

Three findings stand out, and the first is the one the design sheet did not expect. The pipe, which the sheet put first by 153 kW at 1 001.9 kW, binds at a median of 1 126.82 kW with a band from 1 118.12 to 1 135.68 kW: the record's twelve months of pressure drop moved its flow coefficient above the handbook's, and with it the pipe's binding load up by 125 kW, into the chillers' band. The chillers, which the sheet put last at 1 228.0 kW, bind at a median of 1 135.67 kW: their as-built ratios lie below nameplate, and they move down by 92 kW. The two now sit within 9 kW of each other at their medians, the pipe's band entirely inside the chillers'. The feed and the pump follow at 1 188.31 and 1 272.08 kW; the pump, which the sheet put second, is last.

The second finding is the spreads. The pipe and the feed are narrow, 5.29 and 3.15 kW in the delta form, because the monthly pressure drops and powers read their entries closely; the chillers' standard deviation of 22.36 kW is four times the pipe's, because its entries are known less closely: the three as-built ratios to about 0.017 each and the economizer's to 0.35. The third is the shape of the chillers' band. Its delta form is $1\,134.06 \pm 1.645 \times 22.36$, the band $[1\,097.28,\ 1\,170.84]$ kW; the draws put it at $[1\,102.58,\ 1\,178.75]$ kW, with the median 1.61 kW above the MAP value. The economizer's duty $50\,\mathrm{e}^{g_e}$ is an exponential of a normal entry of spread 0.35, skewed to the right, and the symmetric normal of the delta form places both ends too low. The other three delta bands agree with the draws more closely: the pipe's, $1\,126.81 \pm 1.645 \times 5.29 = [1\,118.11,\ 1\,135.52]$ kW, to within 0.16 kW at each end; the pump's, $[1\,254.45,\ 1\,289.54]$ kW, to within 0.16 kW; and the feed's, $[1\,183.05,\ 1\,193.41]$ kW, to within 1.49 kW. The exact posterior moves no median by more than 0.53 kW, the chillers', from 1 135.67 to 1 135.14 kW, and no end of a band by more than 0.62 kW: its 90 % bands are $[1\,102.07,\ 1\,178.13]$ kW for the chillers, $[1\,118.21,\ 1\,135.79]$ kW for the pipe, $[1\,184.06,\ 1\,194.91]$ kW for the feed and $[1\,254.93,\ 1\,290.16]$ kW for the pump.

**The plant's binding load.** The plant's binding load is the smallest of the four in each draw, and the draws give the reported value,

$$\text{median } Q^{\dagger} = 1\,123.88\,\text{kW}, \qquad Q^{\dagger} \in [1\,102.58,\ 1\,133.97]\,\text{kW with } 90\,\%\text{ probability}, \tag{13.10}$$

in site IT load. The exact posterior gives the median 1 123.82 kW and the band from 1 102.07 to 1 133.99 kW, a gap of at most 0.51 kW. Table 13.3 sets beside them the routes on one element alone: each row is a route, the columns the median and the ends of the band. The element that

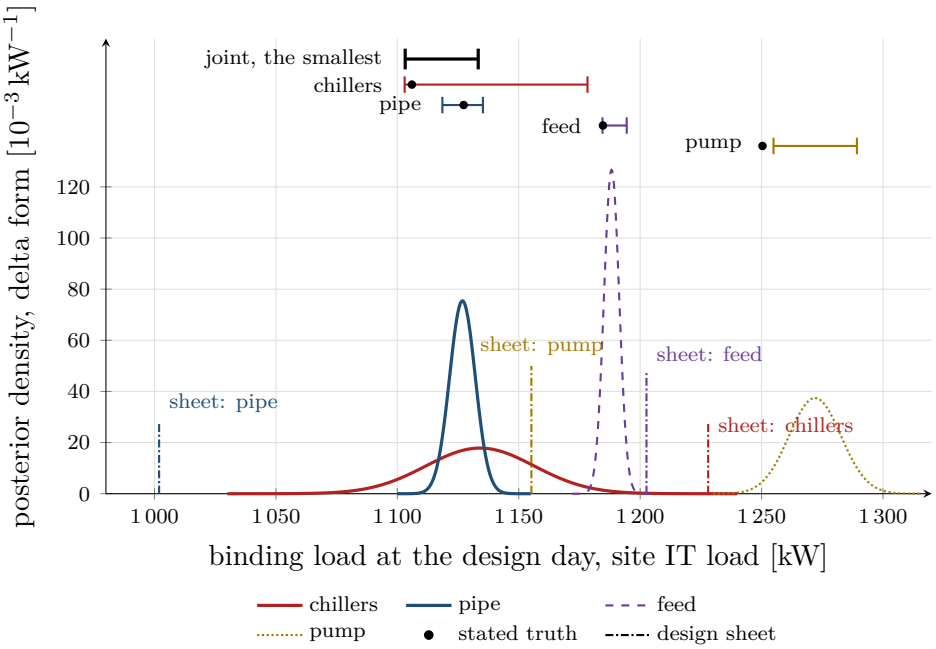



Figure 13.2: The binding loads of the site's four elements at the design day. Curves: the delta form of each element, $\mathcal{N}(Q_e^{\dagger}(\hat{\mathbf{z}}),\ \mathbf{g}_e^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_e)$, from the MAP values and standard deviations of Table 13.2. Bars: the 90 % bands of the 20 000 draws; the top bar is the plant's joint binding load, the smallest of the four in each draw, [1 102.58, 1 133.97] kW. Dots: the stated truth of the record, a check marker (Section 13.4.3). Dash-dotted lines at the bottom: the design sheet's values from nameplate and handbook ratings.

binds at the MAP state is the pipe, and its band alone, [1 118.12, 1 135.68] kW, misplaces the plant's at both ends. At the lower end the joint band reaches down to 1 102.58 kW, the chillers' own 5 % point to the hundredth of a kilowatt: below the pipe's lowest values only the chillers bind, so the lower tail of the minimum is the chillers' lower tail. At the upper end the joint band stops at 1 133.97 kW, short of the pipe's 1 135.68 kW, because in the draws where the pipe is high the chillers bind first and cut the tail. The gap between the pipe's band and the joint band is the chillers' share, at both ends.

### 13.3.2 The element that binds first

**Planner's question.** Which element binds first as the load grows: the chiller, the pump, the pipe, or the power feed? The question asks for the probability of each element being the one whose limit is met first.

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

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

asked here over a family of events, one per element.

Table 13.3: The plant's binding load at the design day by the joint draws, beside the routes on the pipe alone and on the chillers alone, each with the exact posterior, in site IT load.

| Route | Median [kW] | 5 % point [kW] | 95 % point [kW] |
|---|---|---|---|
| Joint draws, the smallest of the four (reported) | 1 123.88 | 1 102.58 | 1 133.97 |
| Joint, exact posterior | 1 123.82 | 1 102.07 | 1 133.99 |
| Pipe alone, draws (its inversion band) | 1 126.82 | 1 118.12 | 1 135.68 |
| Pipe alone, exact posterior | 1 126.86 | 1 118.21 | 1 135.79 |
| Pipe alone, delta form | 1 126.81 | 1 118.11 | 1 135.52 |
| Chillers alone, draws | 1 135.67 | 1 102.58 | 1 178.75 |
| Chillers alone, exact posterior | 1 135.14 | 1 102.07 | 1 178.13 |
| Chillers alone, delta form | 1 134.06 | 1 097.28 | 1 170.84 |

Table 13.4: The probability that each element of the site binds first at the design day, equation (13.11), from 20 000 draws of the Laplace posterior and from the same draws reweighted to the exact posterior.

| Element | $Q_e^\dagger(\hat{\mathbf{z}})$ [kW] | $P_j$, draws [–] | $P_j$, exact posterior [–] |
|---|---|---|---|
| Pipe | 1 126.81 | 0.6505 | 0.6411 |
| Chillers | 1 134.06 | 0.3494 | 0.3589 |
| Feed | 1 188.23 | $< 0.0001$ | $< 0.0001$ |
| Pump | 1 272.00 | $< 0.0001$ | $< 0.0001$ |

**Specialized form.** For element $j$ the question is equation (2.12) with $f(\mathbf{z}) = \min_{e \neq j} Q_e^\dagger(\mathbf{z}) - Q_j^\dagger(\mathbf{z})$ and $f^* = 0$, the event that $j$'s binding load is the smallest, as in equation (5.22),

$$P_j = \Pr(e^\dagger(\mathbf{z}) = j \mid \mathbf{y}) \approx \hat{P}_j = \frac{1}{N} \sum_{i=1}^{N} \mathbf{1}[e^\dagger(\mathbf{z}_i) = j], \qquad \mathrm{se}(\hat{P}_j) = \sqrt{\hat{P}_j\,(1 - \hat{P}_j)/N}, \quad (13.11)$$

with se the Monte Carlo standard error; the exact posterior sums the weights $w_i$ of the same draws. The event has no gradient, so the delta form cannot read it directly. It can read the one pair that matters: if only the pipe and the chillers ever bind first, the pipe binds first when the gap $D = Q_{\text{pipe}}^\dagger - Q_{\text{ch}}^\dagger$ of equation (13.3) is negative, and the delta form puts that probability at

$$\Pr(D < 0 \mid \mathbf{y}) \approx \Phi\left( \frac{-D(\hat{\mathbf{z}})}{\sqrt{\mathbf{g}_D^\mathsf{T} \boldsymbol{\Sigma} \mathbf{g}_D}} \right), \quad (13.12)$$

the single-pair route, which assumes the gap is normal.

**Building blocks.** The draws and their weights (Section 13.2.2); the binding element of each draw; the gap's MAP value and gradient, equations (13.3) and (13.1); the normal distribution function $\Phi$.

**The binding element.** Table 13.4 gives the probabilities: each row is an element, the columns its binding load at $\hat{\mathbf{z}}$ and the probability that it binds first by the draws and by the exact posterior.

The pipe binds first with probability

$$\Pr(e^\dagger = \text{pipe} \mid \mathbf{y}) = 0.6505, \quad (13.13)$$

with a Monte Carlo standard error of $\sqrt{0.6505 \times 0.3495/20\,000} = 0.0034$, and the chillers with probability 0.3494. The exact posterior gives 0.6411 and 0.3589, a gap of $-0.0094$ for the pipe. The feed and the pump bind first in fewer than one draw in ten thousand: the question is a contest between two elements. The single-pair route (13.12) gives $\Phi(7.25/22.98) = \Phi(0.32) = 0.62$, with $D(\hat{\mathbf{z}}) = 1\,126.81 - 1\,134.06 = -7.25$ kW and the delta-form standard deviation of the gap, 22.98 kW. It sits below the draws for the reason the chillers' band does: the gap carries the chillers' right skew as a left skew, so its median by the draws, $-8.84$ kW, lies further below zero than its MAP value, and a symmetric normal places less of it below zero.

The design sheet's answer was the pipe, and the analysis agrees on the element and disagrees on everything that would justify acting on it. The sheet's pipe at 1 001.9 kW is the handbook's pipe; the record moved it to 1 126.8 kW, beside the chillers, and the pipe binds first in two draws of three, not with the certainty a 153 kW lead suggests.

### 13.3.3 The order of the four

**Planner's question.** What is the probability that the order flips? The question asks how often the four elements bind in an order other than the order at the MAP state, and which orders the posterior gives.

**General form.** A prediction under a threshold over a family of events, equation (2.12),

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

one event per order of the four elements.

**Specialized form.** The question is equation (2.12) with the indicator of the event $\pi(\mathbf{z}) = \sigma$ for each order $\sigma$ of the four elements, by equation (13.2), and with the indicator of $\pi(\mathbf{z}) \neq \pi(\hat{\mathbf{z}})$ for the flip,

$$P_\sigma = \Pr\big(\pi(\mathbf{z}) = \sigma \mid \mathbf{y}\big), \qquad P_{\text{flip}} = \Pr\big(\pi(\mathbf{z}) \neq \pi(\hat{\mathbf{z}}) \mid \mathbf{y}\big) = 1 - P_{\pi(\hat{\mathbf{z}})}, \tag{13.14}$$

each read as the share of the draws in which its event holds, and by the exact posterior as the summed weights of those draws. A flip need not change the binding element: an exchange lower in the order, such as the feed passing the chillers, is a flip too.

**Building blocks.** The ranking of each draw (Section 13.2.2); the ranking at the MAP state; the weights.

**The orders.** At the MAP state the order is pipe, chillers, feed, pump, by the binding loads of Table 13.4. Table 13.5 lists every order the draws produce, with its probability, and the same probabilities are drawn in Figure 13.3, one bar per order, the MAP order marked.

The order flips with probability

$$P_{\text{flip}} = 0.3658, \tag{13.15}$$

with a Monte Carlo standard error of 0.0034; the exact posterior gives 0.3744. Four orders occur, and the bars tell their story. The pump is last in every order but the rarest. Almost the whole of the flip is the exchange of the first two: the chillers bind first in 0.3494 of the draws, the probability of the second bar, and in each of those the rest of the order is unchanged. The remaining $0.3658 - 0.3494 = 0.0164$ is the feed passing the chillers behind the pipe, in the draws where the economizer's guess is large: a large economizer raises the chillers' binding load by its duty and the feed's only by the chiller power that duty saves, so the chillers move past the feed. One question of the order is settled, the pump last, and one is open, the first.

Table 13.5: The orders in which the site's four elements bind at the design day, smallest binding load first, with their probabilities over the 20 000 draws. The first row is the order at the MAP state.

| Order [–] | Probability [–] |
|---|---|
| pipe, chillers, feed, pump (the MAP order) | 0.6342 |
| chillers, pipe, feed, pump | 0.3494 |
| pipe, feed, chillers, pump | 0.0163 |
| pipe, feed, pump, chillers | 0.0001 |
| any order other than the MAP order, draws (reported) | 0.3658 |
| any order other than the MAP order, exact posterior | 0.3744 |

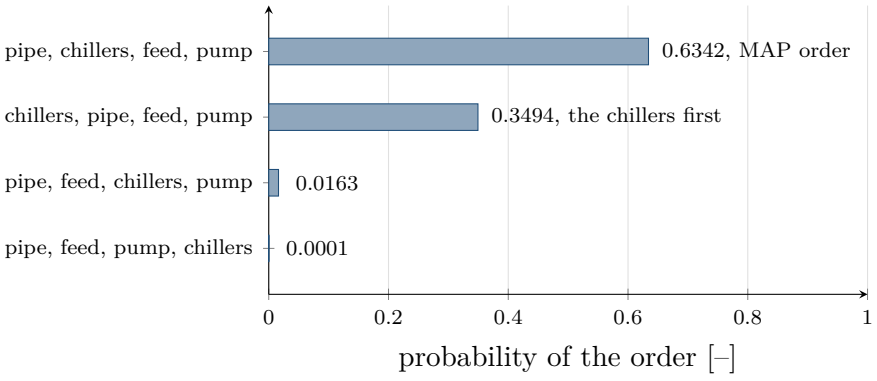



Figure 13.3: The probability of each order of the site's four elements at the design day, from the 20 000 draws (Table 13.5). The top bar is the order at the MAP state; every other bar is a flip, and together they make 0.3658.

Table 13.6: The variance budget of the gap between the pipe's and the chillers' binding loads at the design day, equation (5.25), at the MAP state, in shares of $\mathbf{g}_D^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_D = 22.98^2\,\mathrm{kW}^2$.

| Entry [unit] | Enters the binding load of [–] | Share [–] |
|---|---|---|
| $g_e$, the economizer's log-ratio [–] | chillers (and feed) | 0.580 |
| $\kappa_{C3}$ [–] | chillers | 0.136 |
| $\kappa_{C2}$ [–] | chillers | 0.135 |
| $\kappa_{C1}$ [–] | chillers | 0.095 |
| $l_k$, the pipe's log flow coefficient [–] | pipe (and pump, feed) | 0.053 |
| $d$, $c$, $l_e$ [$\mathrm{K}^{-1}$, –, –] | none of the two | 0.000 |

### 13.3.4 What decides between the pipe and the chillers

**Planner's question.** What decides between the pipe and the chillers? The question asks which entries of the state, and so which parts of the evidence, carry the band of the gap between the two elements that contest the first place.

**General form.** The variance of a derived quantity under the posterior, equation (2.10),

$$\mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y}) = \int \left(f(\mathbf{z}) - \mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]\right)^2 p(\mathbf{z} \mid \mathbf{y})\,\mathrm{d}\mathbf{z},$$

split into one share per entry.

**Specialized form.** The question is equation (2.10) with $f(\mathbf{z}) = D(\mathbf{z}) = Q^{\dagger}_{\mathrm{pipe}}(\mathbf{z}) - Q^{\dagger}_{\mathrm{ch}}(\mathbf{z})$, the gap of equation (13.3), on the Laplace normal, where its variance is $\mathbf{g}_D^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_D$ with $\mathbf{g}_D = \mathbf{g}_{\mathrm{pipe}} - \mathbf{g}_{\mathrm{ch}}$. It splits into one share per entry by the covariance split (5.25), $s_i = g_{D,i}\,(\boldsymbol{\Sigma}\mathbf{g}_D)_i / \mathbf{g}_D^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}_D$. By the pattern (13.8) the gradient has entries in the three log-ratios, in $g_e$ and in $l_k$, and none in $d$, $c$ or $l_e$, whose shares are zero. The entry $g_e$ has no covariance with any other, so its share is its diagonal term alone, $(Q_e(\hat{\mathbf{z}}) \times 0.35)^2$ over the gap's variance, with $Q_e(\hat{\mathbf{z}}) = 50\,\mathrm{e}^0 = 50\,\mathrm{kW}$ by equation (13.4).

**Building blocks.** The gradients (13.4) and (13.5) at $\hat{\mathbf{z}}$; the covariance $\boldsymbol{\Sigma}$; the split (5.25); the draws for the gap's quantiles.

**The gap and its budget.** The gap is $-7.25\,\mathrm{kW}$ at the MAP state, with a delta-form standard deviation of $22.98\,\mathrm{kW}$; by the draws its median is $-8.84\,\mathrm{kW}$ and its 90 % band $[-53.00,\ 25.29]\,\mathrm{kW}$, and its share below zero is 0.6505, the pipe's probability of binding first. The standard deviation is nearly the two elements' combined, $22.36^2 + 5.29^2 = 527.95\,\mathrm{kW}^2$ against $22.98^2 = 528.08\,\mathrm{kW}^2$: the covariances between the chillers' entries and the pipe's are near zero. Table 13.6 splits the variance: each row is an entry of the state, the columns the element whose binding load it enters and its share.

The economizer's handbook guess carries 0.580 of the variance, the diagonal term $(50 \times 0.35)^2 = 306.25\,\mathrm{kW}^2$ over $528.08\,\mathrm{kW}^2$. The three chillers' as-built ratios carry 0.366 together, and the pipe, the element the sheet would upgrade, 0.053. This is the surprise of the chapter: the order of the two elements that contest the first place is decided mostly by a guess about a third element that is not built. Neither the pipe nor the chillers are in doubt to that degree; the twelve months of pressure drop fix the pipe to $5.29\,\mathrm{kW}$, and the capacity tests fix each existing unit's log-ratio to under 0.02. What is in doubt is the duty the drawn economizer will add to the chillers, $50\,\mathrm{kW}$ at the handbook value with a spread of 35 %, and that alone is enough to put the chillers on either side of the pipe. The entries that carry a plant's band, and how much of it a measurement of each would remove, are the subject of Chapter 17.

## 13.4 The checks (step 5)

The answers of Section 13.3 rest on the drawn economizer, on the Laplace normal of the site's fit and on a record generated from a stated truth; each is tested here.

### 13.4.1 The plant as built today

The economizer is drawn but not built, so the plant as it stands has the chillers' binding load without the duty $50\,\mathrm{e}^{g_e}$. On the same draws, with the economizer's term removed from equation (13.4) and the other elements unchanged, the chillers bind at a median of 1 084.16 kW, with a 90 % band of [1 061.54, 1 107.32] kW, and they bind first with probability 0.9970; the pipe binds first in the remaining 0.0030, and the feed and the pump in fewer than one draw in ten thousand. Without the economizer the question has a clear answer, the chillers, and the pipe becomes a contender only through the duty the drawing adds. The design sheet's recommendation, read on the plant as built, would upgrade the element that binds first in three draws of a thousand.

### 13.4.2 The draws against the exact posterior

The reported answers average the 20 000 draws with equal weight; the exact posterior reweights the same draws, with an effective sample size of 19 913. The medians of the four binding loads move by −0.53, +0.04, +0.01 and +0.23 kW, exact minus reported, for the chillers, the pipe, the feed and the pump. The probabilities move more than the medians would suggest: the pipe's probability of binding first by −0.0094 and the flip by +0.0086, each between two and three Monte Carlo standard errors of 0.0034. The two estimates use the same states, so their difference is not sampling noise between independent sets: it is the departure of the posterior of the state from the Laplace normal, and this one has a visible cause. The exact posterior moves the chillers' median down by about half a kilowatt and the pipe's not at all, so the gap moves toward the chillers by about a fortieth of its 23 kW spread; at the center of the gap's distribution, where the probability of its sign is most sensitive, that is a shift of about a hundredth. Both readings put the pipe first in about two draws of three, and neither changes the report.

### 13.4.3 The answers against the stated truth

The operating record is generated from a stated truth, and the last check compares the answers with it. Table 13.7 collects the checks: each row is a quantity, the columns its reported value, the value beside it, the truth, and the reading. At the truth the chillers bind at 1 106.00 kW, the pipe at 1 127.31 kW, the feed at 1 184.66 kW and the pump at 1 250.37 kW, so the true order is chillers, pipe, feed, pump: not the MAP order, but the second most likely order of Table 13.5, to which the draws give 0.3494. The chillers' truth lies inside their band near its lower end: the true economizer is smaller than the handbook's, $g_e = -0.2231$ against the prior's 0, and the true units are a little weaker than their MAP ratios (Table 9.2). The pipe's truth lies inside its band, 0.49 kW above the median, and the feed's inside its band, 0.61 kW above its lower end. The pump's truth lies outside its band, 4.11 kW below the 5 % point: the true pump efficiency, $l_e = -0.3285$, lies $(-0.2891 + 0.3285)/0.0190 = 2.07$ posterior standard deviations below the MAP value, an excursion rare for one entry but not for one of eight. It does not touch the order, since the pump binds last by more than 60 kW either way.

**The years.** Mapped through the load path of equation (6.1) with today's site IT load of 660 kW and growth of 15 % a year, each binding load becomes a crossing year by equation (6.3), draw by draw: the pipe in year 3.83, band [3.77, 3.88]; the chillers in year 3.88, band [3.67, 4.15]; the feed in year 4.21, band [4.18, 4.25]; and the pump in year 4.69, band [4.60, 4.79], medians

Table 13.7: The checks of the binding-constraint answers. The reported values are from the 20 000 draws; beside them, the exact posterior or the route named; the truth is the stated truth of the record.

| Quantity | Reported [kW or –] | Beside [kW or –] | Truth [kW or –] | Reading |
|---|---|---|---|---|
| $Q^{\dagger}_{\mathrm{ch}}$, median | 1 135.67 | exact 1 135.14 | 1 106.00 | inside, near the lower end |
| $Q^{\dagger}_{\mathrm{pipe}}$, median | 1 126.82 | exact 1 126.86 | 1 127.31 | inside |
| $Q^{\dagger}_{\mathrm{feed}}$, median | 1 188.31 | exact 1 188.32 | 1 184.66 | inside, near the lower end |
| $Q^{\dagger}_{\mathrm{pump}}$, median | 1 272.08 | exact 1 272.31 | 1 250.37 | below the band; $l_e$ off by 2.07 sd |
| $Q^{\dagger}$, the plant, median | 1 123.88 | exact 1 123.82 | 1 106.00 | inside, near the lower end |
| Pr(pipe first) | 0.6505 | exact 0.6411 | – | truth: the chillers first |
| Pr(order flips) | 0.3658 | exact 0.3744 | – | true order has 0.3494 |
| Pipe first, single-pair delta form | – | 0.62 | – | below the draws; skew |
| *The plant as built, without the economizer* | | | | |
| $Q^{\dagger}_{\mathrm{ch}}$, median | 1 084.16 | – | – | band [1 061.54, 1 107.32] |
| Pr(chillers first) | 0.9970 | – | – | the pipe 0.0030 |
| *Crossing year at* 15 % *a year from* 660 kW*, median and* 90 % *band* | | | | |
| Pipe [yr] | 3.83, band [3.77, 3.88] | | | |
| Chillers [yr] | 3.88, band [3.67, 4.15] | | | |
| Feed [yr] | 4.21, band [4.18, 4.25] | | | |
| Pump [yr] | 4.69, band [4.60, 4.79] | | | |
| Plant, the smallest [yr] | 3.81, band [3.67, 3.87] | | | |

with 90 % bands. The crossing year is increasing in the binding load, so the plant's year has the quantiles of the joint binding load (13.10), mapped: a median of $\ln(1\,123.88/660)/\ln 1.15 = 3.81$, with the band [3.67, 3.87]. The two contenders cross within three weeks of each other at their medians, and the pipe's whole band lies inside the chillers'; which of them is crossed first is the question of Section 13.3.2, not a matter of the calendar.

## 13.5 The report (step 6)

The design sheet named the right element for the wrong reason. On the design day, with every unit up and the drawn economizer, the pipe binds first with probability 0.65 and the chillers with probability 0.35, at a plant binding load of 1 124 kW of site IT load, band 1 103 to 1 134 kW, reached in year 3.8 of the load path. The order of the two is decided mostly by the economizer's handbook conductance, and an upgrade committed on the sheet's ranking would rest on that guess.

### 13.5.1 The answers and their bands

Table 13.8 gives the answers in the order of Table 13.1. Each row is a question; the second column gives the reported answer with the exact posterior beside it, and the third its band or the routes beside it.

### 13.5.2 The design-review entry

The design review's entry replaces the sheet's ranking with a probability. On the design day the site's binding constraint is the shared pipe with probability 0.65 by the draws, 0.64 by the exact posterior, and the chillers with the drawn economizer otherwise; the feed and the pump bind later in every draw. The plant's binding load is 1 123.9 kW of site IT load at the median,

Table 13.8: The report for the site's four elements at the design day, with all units up and the drawn economizer, in site IT load. Reported values by 20 000 draws of the posterior of the state, joint over the four elements; the exact posterior beside each.

| Question | Answer | Band or routes beside |
|---|---|---|
| Which element binds first as the load grows: the chiller, the pump, the pipe, or the power feed? | the pipe, probability 0.6505 (exact: 0.6411); the chillers 0.3494 (exact: 0.3589); the feed and the pump below 0.0001 | Monte Carlo standard error 0.0034; single-pair delta form 0.62; without the economizer the chillers, 0.9970 |
| At what load does each bind, and how sure? | plant, the smallest: median 1 123.9 kW (exact: 1 123.8); pipe 1 126.8, chillers 1 135.7, feed 1 188.3, pump 1 272.1 kW (medians) | plant 1 102.6 to 1 134.0 kW (exact: 1 102.1 to 1 134.0); pipe 1 118.1 to 1 135.7; chillers 1 102.6 to 1 178.8; feed 1 184.1 to 1 194.9; pump 1 254.5 to 1 289.7 kW; central 90 % |
| What is the probability that the order flips? | 0.3658 (exact: 0.3744) | chillers, pipe, feed, pump 0.3494; pipe, feed, chillers, pump 0.0163 |
| What decides between the pipe and the chillers? | the economizer's handbook guess, 0.580 of the gap's variance | the three chillers' ratios 0.366; the pipe 0.053; gap median −8.8 kW, band −53.0 to 25.3 kW |

band 1 102.6 to 1 134.0 kW, crossed in year 3.81 of the load path, band 3.67 to 3.87. The entry records three further facts. The sheet's pipe at 1 001.9 kW was the handbook's; the record puts it at 1 126.8 kW, and an upgrade of the pipe alone would leave the chillers binding at a median of 1 135.7 kW, 9 kW above where the pipe binds now; what an upgrade is worth is Chapter 14's question. Without the economizer the chillers bind first with probability 0.997. And 58 % of the variance of the gap that decides the order is the economizer's handbook conductance, which no reading of the existing plant sees. The entry's decision follows: no upgrade is committed on the sheet's ranking, and the economizer's conductance is measured or specified before the order is relied on.

**Common mistake**

Ranking the bottlenecks by nameplate and handbook values. Each binding load moves by what the record knows of its own element, and the order of two close elements is decided by whichever entry carries the band of their gap, which may belong to neither.

## 13.6 The next measurement (step 7)

The report leaves the order of the two leading elements on a coin weighted about two to one, and the weight is set by one prior. Step 7 asks which test or survey would narrow the quantity that decides the build; this chapter names the entry and points forward, and computes nothing new.

### 13.6.1 The entry that decides the order

The gap's budget names the entry: the economizer's log-ratio $g_e$, with 0.580 of the variance. A measurement that fixed it would leave the gap with the remaining 0.420, carried by the chillers' as-built ratios and the pipe. No test of the existing plant can see the economizer, since it is not built: its conductance can be fixed only by the manufacturer's rating test of the selected unit, by a pilot, or by a specification that makes the handbook value a contractual minimum. The chillers' share, 0.366, is open to a capacity test of each unit at the design-day wet bulb. The pipe's share, 0.053, is already small: another year of pressure-drop readings would buy little.

Chapter 17 splits a plant's band by its entries and states what each costs the decision, and Chapter 18 ranks the tests and surveys a planner can commission by the width each takes off the quantity that decides the build; the economizer's conductance heads this chapter's list.

### 13.6.2 What the analysis showed

The design sheet put the pipe first by 153 kW and recommended an upgrade. The analysis puts the pipe first too, with probability 0.65, but 9 kW ahead of the chillers at the medians, not 153 kW, and with the order reversed in a third of the posterior. Three findings carry the answer. First, the record moved the two elements toward each other from opposite sides: the pipe passes more flow than its handbook coefficient, and the chillers deliver less than nameplate, so the sheet's comfortable lead was an artifact of two handbook values. Second, the joint binding load differs from the binding element's: the plant's band reaches down to the chillers' lower tail and stops short of the pipe's upper tail, because the chillers bind wherever the pipe is high. Third, and beyond this site, the order of two close elements can be decided by an element that is neither. Here the drawn economizer's guessed conductance carries 58 % of the gap's variance; without it the question has a clear answer, the chillers, and with it the answer turns on a number no reading of the plant contains.

### 13.6.3 Limits of the model

**The design day with all units up.** The four binding loads are read at the design day's wet bulb with every chiller running. With a unit out the chillers' binding load falls by that unit's capacity, some 350 kW or more, and the chillers bind long before the pipe; with the shared pipe itself out of service the plant carries nothing. Chapter 15 brings the outages and the halls into the same reading.

**One efficiency per pump, one coefficient per pipe.** The pump has one efficiency at every flow, and the pipe one flow coefficient with a drop that rises with the square of the flow. Near its motor rating a real pump runs away from its best point, and a pipe's drop includes fittings whose loss coefficients change with the flow; a pump curve and a measured system curve would replace both entries.

**Fixed supply, rise and power factor.** The supply temperature, the loop's rise and the feed's power factor are pinned. A plant that raises its supply temperature at high load moves heat into a smaller flow and relieves the pipe and the pump; a falling power factor would lower the feed's limit.

**The economizer at its handbook value.** The economizer enters every answer through a prior alone. Its truth in this record is smaller than the handbook's, which is why the true order puts the chillers first; a different guess would move the probabilities of Table 13.4 in either direction.

**The site's elements only.** The binding loads are those of the central plant's four elements. A rack inlet in a hall, the limit of Chapter 5, can bind before any of them, and the hall units' capacities are not in this chapter's state.

**The constructed record.** The operating record and the capacity tests are generated from a stated truth with the stated reading errors, not taken from a site, and the comparison with that truth in Section 13.4.3 is a check of the method on its own model, not of the model against a real plant.

**The next chapter.** Chapter 14 asks what raising each element's limit is worth: which upgrade moves the plant's binding load most, by the importance of each rating.

## 13.7 What to retain

**Definition 13.1** (Binding constraint and its ranking)**.** The binding load $Q_e^\dagger$ of each element $e$ of a plant is the scenario solved for the load at which its limit is met, $h_e(\mathbf{z}, Q_e^\dagger) = h_e^{\max}$, with the implicit-function gradient $\mathbf{g}_e = -\nabla_{\mathbf{z}} h_e / (\partial h_e / \partial Q^*)$ in kilowatts of load. The binding constraint is the smallest binding load, the binding element the one that attains it, and the ranking the order of all of them. The binding load is reported jointly, from draws of the posterior of the state pushed through every element; the binding element and the ranking have no gradient and are reported as probabilities over the same draws, with the exact posterior beside them. The order of two elements is the sign of their gap, whose variance budget names the entries that decide it.

Four conclusions carry beyond this site. A ranking of bottlenecks is a probability distribution over orders, not a list, and the probability that it flips is a number the review can read. Elements ranked on handbook values can trade places once the record is read: each binding load moves by what the record knows of its own element, and two that look far apart may be close. The plant's binding load is the joint minimum, whose tails belong to whichever element binds there, not to the element that binds at the MAP state. And the order of two close elements is decided by the entry that carries the band of their gap, which can be an unmeasured guess about a third element; that entry, not either contender, is what to measure before an upgrade is chosen.

Table 13.9 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 13.9: The terms of binding constraints and where each is defined. Units for the site at the design day.

| Term | Expression | Units | Meaning | Defined in |
|---|---|---|---|---|
| Element's binding load | $h_e(\mathbf{z}, Q_e^\dagger) = h_e^{\max}$; $\mathbf{g}_e = -\nabla_{\mathbf{z}} h_e / (\partial h_e / \partial Q^*)$ | kW | limit of one element | 13.1.1 |
| Binding element | $e^\dagger = \arg\min_e Q_e^\dagger$ | – | the element that binds first | 13.1.1 |
| Ranking | $\pi = \operatorname{argsort}_e Q_e^\dagger$ | – | the order of the elements | 13.1.1 |
| Pairwise gap | $D_{ab} = Q_a^\dagger - Q_b^\dagger$ | kW | its sign is the order | 13.1.1 |
| Chillers | $\sum_i 400\,\mathrm{e}^{\kappa_i} + 50\,\mathrm{e}^{g_e} - Q_o$ | kW | heat the plant can carry | 13.1.3 |
| Pipe | $\rho c_p \Delta T_w \sqrt{\Delta P_{\max}}\,\mathrm{e}^{l_k} - Q_o$ | kW | flow at the allowable drop | 13.1.3 |
| Pump, feed | roots of $h(V) = P_m$, $F(Q^*) = 0$ | kW | motor and feed limits | 13.1.3 |
| Gradient pattern | eq. (13.8) | – | the entries each element reads | 13.1.3 |
| *Questions and answers* | | | | |
| Joint binding load | quantiles of $\min_e Q_e^\dagger(\mathbf{z}_i)$ | kW | the reported plant value | 13.3.1 |
| Binding-element probability | $\Pr(e^\dagger = j \mid \mathbf{y})$ | – | which element binds first | 13.3.2 |
| Single-pair route | $\Phi(-D(\hat{\mathbf{z}}) / \sqrt{\mathbf{g}_D^\mathsf{T} \boldsymbol{\Sigma} \mathbf{g}_D})$ | – | delta form on the gap | 13.3.2 |
| Order flip | $\Pr(\pi(\mathbf{z}) \neq \pi(\hat{\mathbf{z}}) \mid \mathbf{y})$ | – | the order is not the MAP order | 13.3.3 |
| Gap budget | $s_i = g_{D,i} (\boldsymbol{\Sigma} \mathbf{g}_D)_i / \mathbf{g}_D^\mathsf{T} \boldsymbol{\Sigma} \mathbf{g}_D$ | – | what decides the order | 13.3.4 |

## 13.8 Exercises

**Exercise 13.1.** *Formulation.* The planner asks how likely it is that the feed binds before the chillers. Write the question as a probability over the state and state its answer from this chapter's draws. ($\Pr(Q^{\dagger}_{\text{feed}}(\mathbf{z}) < Q^{\dagger}_{\text{ch}}(\mathbf{z}) \mid \mathbf{y})$, equation (2.12) with the gap $D = Q^{\dagger}_{\text{feed}} - Q^{\dagger}_{\text{ch}}$ of equation (13.3) and $f^* = 0$. It is the summed probability of the orders in Table 13.5 in which the feed comes before the chillers: pipe, feed, chillers, pump at 0.0163 and pipe, feed, pump, chillers at 0.0001, together 0.0164.)

**Exercise 13.2.** *Derivation.* Derive the pipe's binding load and its gradient from the pressure-drop relation of equation (9.3), and check its delta-form standard deviation. (The drop $(V/k_p)^2$ meets $\Delta P_{\max}$ at $V = k_p\sqrt{\Delta P_{\max}}$; the loop's heat at that flow is $\rho c_p\,\Delta T_w\,V$, and the site IT load is that less $Q_o$, so $Q^{\dagger}_{\text{pipe}} = \rho c_p\,\Delta T_w\sqrt{\Delta P_{\max}}\,\mathrm{e}^{l_k} - Q_o$. Its only gradient entry is $\partial Q^{\dagger}_{\text{pipe}}/\partial l_k = Q^{\dagger}_{\text{pipe}} + Q_o = 1\,126.81 + 22 = 1\,148.81$ kW, and with $\mathrm{sd}(l_k) = 0.0046$ the standard deviation is $1\,148.81 \times 0.0046 = 5.28$ kW, against the printed 5.29 kW from the unrounded spread.)

**Exercise 13.3.** *Derivation.* Show that the pipe's binding load scales with the flow coefficient, and use the design sheet and the fit to check the MAP value. (By equation (13.5), $Q^{\dagger}_{\text{pipe}} + Q_o$ is proportional to $\mathrm{e}^{l_k}$. The sheet takes the handbook's $l_k = \ln 2 = 0.6931$ and finds $1\,001.9 + 22 = 1\,023.9$ kW of loop heat; the fit's $l_k = 0.8083$ multiplies that by $\mathrm{e}^{0.1152} = 1.122$, giving about 1 148.9 kW of loop heat and 1 126.9 kW of site IT load, against the printed 1 126.81 kW from the unrounded entry.)

**Exercise 13.4.** *Application.* Show from the gradient and the covariance why the economizer's share of the gap's budget is 0.580. (The economizer's entry $g_e$ has no row but its prior, so it is uncorrelated with every other entry and its share is its diagonal term alone. Its gradient entry in the gap is $-Q_e(\hat{\mathbf{z}}) = -50\,\mathrm{e}^0 = -50$ kW, by equation (13.4), and its variance $0.35^2$, so the term is $(50 \times 0.35)^2 = 306.25\,\mathrm{kW}^2$; over the gap's variance $22.98^2 = 528.08\,\mathrm{kW}^2$ it is 0.580.)

**Exercise 13.5.** *Application.* The facilities team proposes to drop the economizer from the drawing to save capital. Say which element then binds first, at about what load, and how sure. (Without the economizer the chillers' binding load loses $50\,\mathrm{e}^{g_e}$, 50 kW at the MAP state, so that $1\,134.06 - 50 = 1\,084.06$ kW there, below the pipe's 1 126.81 kW by more than 40 kW, with the economizer's share of the chillers' band removed with it. By the draws the chillers bind at a median of 1 084.16 kW, band 1 061.54 to 1 107.32 kW, and they bind first with probability 0.9970.)

# Chapter 14

# Capacity Importance Measures

## 14.1 Which upgrade buys the most headroom

The site's design review has a budget for one upgrade, and four are priced. The chiller plant can take a fourth 400 kW chiller for R$ 1.60 million; the shared supply pipe can be replaced one size up, which raises its flow coefficient by a quarter, for R$ 0.90 million; the pump's motor can go from 12 to 18 kW for R$ 0.15 million; and the power feed from 1 600 to 2 000 kVA for R$ 0.80 million. Each buys some rating of one element, and the review wants to know which buys the most headroom over today's 660 kW of site IT load.

The design sheet ranks the upgrades by its own binding order. Worked from the nameplates and the handbook values at the design day, it finds that the shared pipe reaches its allowable pressure drop first, at 1 001.9 kW of site IT load; the pump reaches its motor's power at 1 155.2 kW, the feed its rating at 1 202.6 kW, and the chillers their capacity at 1 228.0 kW. The pipe binds first, so the sheet says to upgrade the pipe. Its logic is the logic of a bottleneck: the element that binds first is the one whose rating the headroom moves with, and it moves with no other. The planner's question is the drill question, verbatim,

*If I could upgrade one element, which one buys the most headroom?*

It is answered with the method of Chapter 4 on the site of Chapter 9, with the one fit of its operating record. Chapter 13 asks which element binds first and finds the answer uncertain; this chapter asks what that uncertainty is worth when a rating is bought. What is new is the quantity asked for: not a binding load, but how the headroom, the smallest of four binding loads, changes when one of them is moved by a stated step.

**Jargon: importance measure, debottlenecking, rating, upgrade step, flow coefficient, power factor, binds first, derivative of a minimum, expected gain, complementary upgrades, gain per real**

An *importance measure* ranks the elements of a plant by how much a quantity of the whole plant changes with each of them; in the capacity sense used here, it is the change of the headroom per unit of an element's rating. *Debottlenecking* is the practice of raising the rating of the element that limits a plant. The *rating* $r_e$ of an element $e$ is the number on its nameplate that sets its limit: the chillers' total nameplate in kW, the pipe's flow coefficient, the pump motor's power in kW and the feed's rating in kVA. An *upgrade step* $\Delta_e$ is a stated change of one rating with a stated price. The *flow coefficient* $k_p$ of a pipe relates its flow to its pressure drop, $\Delta P = (V/k_p)^2$; a pipe one size up has a larger coefficient. The *power factor* of the feed is the ratio of the real power it delivers to its apparent power in kVA. An element *binds first* when its binding load is the smallest of the plant's. The *derivative of a minimum* of several smooth functions is, wherever one of them is strictly the smallest, the derivative of that one. The *expected gain* of an upgrade is the change of the headroom it buys, averaged over the posterior. Two upgrades are *complementary* when they buy more together than the sum of what each buys alone. The *gain per real* is the expected gain divided by the price, here in kW per million reais (R$ M).

**Key idea**

The headroom of a plant is the smallest of its elements' binding loads less today's load, and an importance measure is its derivative in one element's rating. That derivative is the element's own gradient where the element binds first and zero elsewhere, so over the posterior it is the probability that the element binds first times its gradient. It is local. A stated upgrade is a step, and the headroom gained by a step stops growing where the next element binds; the gain is read jointly from draws of the posterior with every element's binding load, before and after the step. The importance measure says where to look, and the joint gain says what to buy.

### 14.1.1 Importance read as the derivative of a minimum

The headroom of the site is set by its four elements together: the chillers, the shared pipe, the pump and the feed. Each element $e$ has a binding load $Q_e^\dagger$, the site IT load at which its limit binds, in the sense of Section 5.1.1; it depends on the plant's state $\mathbf{z}$ and on the element's own rating $r_e$. The headroom is the smallest binding load less today's load,

$$H(\mathbf{z}, \mathbf{r}) = \min_e Q_e^\dagger(\mathbf{z}, \mathbf{r}) - Q_0, \qquad e_1(\mathbf{z}, \mathbf{r}) = \arg\min_e Q_e^\dagger(\mathbf{z}, \mathbf{r}), \tag{14.1}$$

with $\mathbf{r} = (r_{\mathrm{ch}}, r_{\mathrm{p}}, r_{\mathrm{m}}, r_{\mathrm{f}})$ the four ratings, $Q_0 = 660\,\mathrm{kW}$ today's site IT load, a stated input, and $e_1$ the element that binds first. Write $g_e = \partial Q_e^\dagger / \partial r_e$ for the gradient of element $e$'s binding load in its own rating, in kW of site IT load per unit of the rating.

Three readings of an element's importance follow from equation (14.1). The first is the derivative of the headroom in the rating at one state. Wherever one element binds strictly first, the minimum moves with that element alone, so the derivative is that element's gradient, and the other elements' is zero. At the MAP state $\hat{\mathbf{z}}$ of the site's fit,

$$I_e^{\mathrm{MAP}} = \left.\frac{\partial H}{\partial r_e}\right|_{\hat{\mathbf{z}}} = g_e(\hat{\mathbf{z}})\,\mathbf{1}\big[e_1(\hat{\mathbf{z}}) = e\big]. \tag{14.2}$$

The second averages that derivative over the posterior. Ties between two binding loads have probability zero under a continuous posterior, so the derivative exists almost everywhere and

$$I_e = \mathbb{E}\Big[\frac{\partial H}{\partial r_e} \,\Big|\, \mathbf{y}\Big] = \int g_e(\mathbf{z})\,\mathbf{1}\big[e_1(\mathbf{z}) = e\big]\, p(\mathbf{z} \mid \mathbf{y})\,\mathrm{d}\mathbf{z} = \Pr(e_1 = e \mid \mathbf{y})\,\mathbb{E}\big[g_e \mid e_1 = e,\ \mathbf{y}\big], \tag{14.3}$$

the probability that the element binds first times its mean gradient where it does. This is the importance measure in the capacity sense: the headroom's analog of the Birnbaum importance, which *Reliability Engineering by Derivation*, Chapter 26, derives as the gradient of a system's probability of failure in one unit's unavailability, within the family of importance measures that chapter develops. The third reading asks of a stated step, not of a derivative: the change of the headroom when one rating is pinned at its upgraded value $r_e + \Delta_e$, averaged over the posterior,

$$G_e(\mathbf{z}) = H(\mathbf{z}, \mathbf{r} + \Delta_e\, \mathbf{e}_e) - H(\mathbf{z}, \mathbf{r}), \qquad \bar{G}_e = \mathbb{E}[G_e \mid \mathbf{y}], \tag{14.4}$$

with $\mathbf{e}_e$ the unit vector of rating $e$. The first two are in kW per unit of the element's rating, and so cannot be compared across elements whose ratings are in different units; multiplied by the step, $I_e\, \Delta_e$ is in kW, and it is what the importance measure promises for the step. The third is in kW directly, and it is what the step buys. They agree only if the element stays the one that binds first over the whole step.

**Principle 14.1.** *A capacity importance measure is the expected derivative of the headroom in one rating: the probability that the element binds first times its gradient. It ranks where the headroom is sensitive. What an upgrade buys is the expected change of the headroom over the stated step, read jointly over the elements, and it is smaller whenever another element binds within the step.*

### 14.1.2 The site and its four ratings (step 0)

The plant is the site of Chapter 9, drawn in Figure 9.1: the two halls on one chilled-water plant with three existing chillers C1–C3 of 400 kW nameplate, one pump with a 12 kW motor, the shared supply pipe with its allowable pressure drop of 150 kPa, the power feed of 1 600 kVA at power factor 0.95, and the drawn economizer. Its state is the eight entries of equation (9.4), and its physical relations are equations (9.3). The plant is taken at the design day, a wet bulb of 22°C, all units up and with the economizer as drawn.

**The four binding loads.** Each element's limit, set in the relations (9.3) at the design day, defines its binding load. Chapter 13 derives the four; they are restated here with each element's rating made explicit, since the rating is what this chapter moves. With $Q$ the site IT load and $V = (Q + Q_o)/(\rho c_p\, \Delta T_w)$ the loop flow,

$$\begin{aligned}
\text{chillers:} &\quad Q^{\dagger}_{\mathrm{ch}} = \textstyle\sum_i 400\, \mathrm{e}^{\kappa_i} + r_{\mathrm{ch}}\, \mathrm{e}^{\bar{\kappa}} + 50\, \mathrm{e}^{g_e} - Q_o, \\
\text{pipe:} &\quad Q^{\dagger}_{\mathrm{p}} = r_{\mathrm{p}}\, k_p \sqrt{150}\, \rho c_p\, \Delta T_w - Q_o, \\
\text{pump:} &\quad \left[(V/(r_{\mathrm{p}} k_p))^2 + \Delta P_{\mathrm{h}}\right] V/(1000\, \eta) = r_{\mathrm{m}} \quad \text{at } Q = Q^{\dagger}_{\mathrm{m}}, \\
\text{feed:} &\quad (1 + a)\, Q + P_{\mathrm{ch}}(Q) + P_{\mathrm{pump}}(Q) = 0.95\, r_{\mathrm{f}} \quad \text{at } Q = Q^{\dagger}_{\mathrm{f}},
\end{aligned} \tag{14.5}$$

with $\bar{\kappa} = (\kappa_{C1} + \kappa_{C2} + \kappa_{C3})/3$ the mean as-built log-ratio, the other symbols those of equation (9.3), and the ratings $r_{\mathrm{ch}}$ the chiller nameplate added to the existing 1 200 kW, $r_{\mathrm{p}}$ the factor on the pipe's flow coefficient, $r_{\mathrm{m}}$ the motor's power in kW and $r_{\mathrm{f}}$ the feed's rating in kVA. As drawn, $\mathbf{r} = (0,\ 1,\ 12,\ 1\,600)$. The pipe's upgrade also enters the pump's and the feed's limits, through the pump's power.

**The rating gradients.** The chillers' and the pipe's binding loads are linear in their own ratings, so their gradients are read off equation (14.5),

$$g_{\mathrm{ch}} = \mathrm{e}^{\bar{\kappa}}, \qquad g_{\mathrm{p}} = k_p \sqrt{150}\, \rho c_p\, \Delta T_w = Q^{\dagger}_{\mathrm{p}} + Q_o \quad (r_{\mathrm{p}} = 1), \tag{14.6}$$

the as-built ratio of an added unit, in kW of capacity per kW of nameplate, and the pipe's carrying capacity in heat per unit of the factor. The pump's and the feed's binding loads are defined by an equation, and their gradients are the implicit-function form of equation (5.1) with the rating in place of the state: the change of the limit per unit of rating, divided by the change of the element's demand per kilowatt of load,

$$g_{\mathrm{m}} = \left( \frac{\partial P_{\mathrm{pump}}}{\partial Q} \right)^{-1}, \qquad g_{\mathrm{f}} = 0.95 \left( \frac{\partial P_{\mathrm{feed}}}{\partial Q} \right)^{-1}, \tag{14.7}$$

each at its own binding load, with $P_{\mathrm{feed}}$ the left side of the feed's limit in equation (14.5).

**Modeling choices at the graph.** The added chiller is taken at the mean as-built ratio of the existing three, with no spread of its own: this keeps the chiller step a function of the fitted state alone, and it fails if the new unit is of another model or is commissioned differently (Chapter 9 gave a new unit a stated spread, equation (9.7), which would widen the chiller step's gain a little). The pipe one size up is taken as the stated factor 1.25 on its flow coefficient, with the same allowable pressure drop: a planning input, not a hydraulic calculation. All units are taken up, so the fourth chiller is read here as capacity, not as the redundancy that Chapter 9 reads it as. The plant is taken as drawn, with the economizer, which is not built; Chapter 13 shows that the binding order rests largely on it, and the Limits return to that.

### 14.1.3 The upgrade steps as a scenario

The question is asked at a scenario node of Section 4.3.8 with three things pinned: the wet bulb at the design day's 22°C, the load against which the headroom is measured at today's $Q_0 = 660\,\mathrm{kW}$, and the ratings. The ratings are pinned twice, as drawn and as upgraded, and the scenario quantities at the upgraded ratings are starred,

$$Q^{\dagger *}_{e'} = Q^{\dagger}_{e'}(\mathbf{z},\ \mathbf{r} + \Delta_e\, \mathbf{e}_e), \qquad H^* = \min_{e'} Q^{\dagger *}_{e'} - Q_0, \tag{14.8}$$

for every element $e'$ and each upgrade $e$. The four stated steps, planning inputs and not estimates, are

- the chillers: a fourth 400 kW chiller, $\Delta_{\mathrm{ch}} = 400\,\mathrm{kW}$ of nameplate, R$ 1.60 million;
- the pipe: one size up, $\Delta_{\mathrm{p}} = 0.25$ on the flow-coefficient factor, R$ 0.90 million;
- the pump: the motor from 12 to 18 kW, $\Delta_{\mathrm{m}} = 6\,\mathrm{kW}$, R$ 0.15 million;
- the feed: from 1 600 to 2 000 kVA, $\Delta_{\mathrm{f}} = 400\,\mathrm{kVA}$, R$ 0.80 million;

and two pairs of them, the pipe with the chillers and the pipe with the feed, are pinned the same way.

### 14.1.4 The planner's questions (step 1)

The planner asks four questions of the upgrades, and step 1 writes each over the site's state, with the evidence on the right of the bar. Table 14.1 lists them. The first column is the question in the planner's words, the drill question first; the second the general form of its category in the notation of Chapter 2; the third its mathematical form for the site; 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 drill question is a design question answered by an expected consequence, the headroom each step buys. The second and third are estimation: the importance measure is a moment of a derivative, and the disagreement between the measure and the gain is the moment of a minimum. The fourth is the drill question asked of two steps at once. The answers are given in the order two, one, three, four, because the gain is best read against the measure it departs from.

Table 14.1: The questions asked of the site's upgrades at the design day, the plant as drawn. $\mathbf{y}$ is the site's operating record and priors; $H$ the headroom (14.1); $G_e$ the gain of step $e$ (14.4); $p_e$ its stated price.

| Planner's question | General form | Mathematical form | Category | Computation | Answered in |
|---|---|---|---|---|---|
| If I could upgrade one element, which one buys the most headroom? | $\arg\min_a \mathbb{E}[L(a,\mathbf{z}) \mid \mathbf{y}]$ | $\arg\max_e \mathbb{E}[G_e \mid \mathbf{y}]$, and $\mathbb{E}[G_e \mid \mathbf{y}]/p_e$ | design | expected consequence | 14.3.2 |
| What is each element's importance measure? | $\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$ | $I_e = \mathbb{E}[\partial H/\partial r_e \mid \mathbf{y}]$: $\Pr(e_1 = e)$ times the gradient | estimation | posterior | 14.3.1 |
| Why do the gradient and the upgrade disagree? | moments of a minimum | $I_e\,\Delta_e - \mathbb{E}[G_e \mid \mathbf{y}]$ against the distance to the next element | estimation | posterior | 14.3.3 |
| What do two upgrades buy together? | $\arg\min_a \mathbb{E}[L(a,\mathbf{z}) \mid \mathbf{y}]$ | $\mathbb{E}[G_{e,e'} \mid \mathbf{y}]$ for the pipe with the chillers and with the feed | design | expected consequence | 14.3.4 |

### 14.1.5 The evidence as rows (step 2)

The evidence is the site's operating record of Section 9.1.5, unchanged: two capacity tests per chiller, twelve monthly windows of chiller power, pipe pressure drop, pump power and feed power, and the eight prior rows, 62 rows on the eight entries of the state. The outage record is not used: every unit is taken up. The upgrade steps and their prices are stated inputs, not rows: no reading of the existing plant sees a pipe one size up or a fourth chiller, and the chapter does not estimate a price. The economizer is not built, so no row sees its conductance, and its log-ratio keeps its prior, as in Chapter 9.

**Modeling choices at the rows.** The rows keep the choices of Section 9.1.5: Gaussian, unbiased and independent readings given the state. One choice is added: the stated steps are taken as delivered, so that a pipe one size up has exactly the factor 1.25 and a new chiller exactly the mean as-built ratio. This makes the gain a function of the fitted state alone, and it fails if an upgrade is installed to a different rating than the one priced.

## 14.2 The posterior of the state (step 3)

All four questions are read from one posterior, that of the site's state under its operating record. This section derives what the answers need from it, starting from the expected gain of the drill question, and names where each part is computed.

### 14.2.1 What the answer needs from the state

The drill question owes the expected gain of each step, equation (14.4), an integral over the posterior of the state,

$$\bar{G}_e = \int \Big[\min_{e'} Q^{\dagger *}_{e'}(\mathbf{z}) - \min_{e'} Q^{\dagger}_{e'}(\mathbf{z})\Big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \tag{14.9}$$

by equations (14.1) and (14.8), the load $Q_0$ cancelling. The integrand needs the four binding loads at two rating vectors, each by equation (14.5). It is a difference of two minima, and it has a kink wherever two binding loads cross, which here they do inside the posterior: Chapter 13

finds the pipe and the chillers within each other's band. A single gradient at $\hat{\mathbf{z}}$ therefore carries neither the integrand nor its mean, and the integral is taken by draws, as for every minimum since Chapter 5. The importance measure (14.3) needs the same posterior, with the rating gradient (14.6) or (14.7) of the element that binds first in each draw.

The posterior of the state is the Laplace normal of the site's fit,

$$p(\mathbf{z} \mid \mathbf{y}) \approx \mathcal{N}(\hat{\mathbf{z}},\ \boldsymbol{\Sigma}), \qquad \bar{G}_e \approx \frac{1}{N}\sum_{i=1}^{N} G_e(\mathbf{z}_i), \qquad \bar{G}_e^{\text{exact}} \approx \sum_{i=1}^{N} w_i\, G_e(\mathbf{z}_i), \tag{14.10}$$

with $\mathbf{z}_i$ the draws from it and $w_i$ the weights of equation (5.16), normalized over the draws used, which turn the same draws into the exact posterior. The MAP state and the covariance are the minimum of half the squared residual vector of the 62 rows and the inverse of $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ there, by equations (4.29) and (4.38); the quantiles of the gain are read from the distribution function of the $G_e(\mathbf{z}_i)$, and the probability that an element binds first is the share of the draws in which its binding load is the smallest.

### 14.2.2 The site's fit and draws

The fit is the one fit of the site, computed in Section 9.2.2 and listed in Table 9.2, and it is not repeated. Its draws are the 20 000 of Section 9.2.4, made with the seed 9, whose weights have an effective sample size of 19 913. Each draw here needs the four binding loads at seven rating vectors, as drawn and after each step and pair, those of the pump and the feed each a root of its limit, so this chapter uses the first 4 000 of those draws, with their weights renormalized over the 4 000, which have an effective sample size of 3 983. The subset reproduces Chapter 13's ranking to within its Monte Carlo error: the pipe binds first in 0.6405 of the 4 000 draws against 0.6505 of the 20 000, the chillers in 0.3595 against 0.3494, and the pump and the feed in none. Every number of Section 14.3 is read from the 4 000.

## 14.3 The answers (step 4)

Every answer is read from the draws of Section 14.2.2. The reported value is the average or quantile over the draws, the state route; the exact posterior stands beside it, and the readings at the MAP state, the single-state route, stand beside both.

### 14.3.1 Each element's importance measure

**Planner's question.** What is each element's importance measure? The question asks how fast the headroom grows with each element's rating, at the MAP state and over the posterior.

**General form.** An estimation question asks for the moments of a derived quantity under the posterior, equation (2.10),

$$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}] = \int f(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}.$$

**Specialized form.** The question is equation (2.10) with $f(\mathbf{z}) = \partial H/\partial r_e$, the derivative of the headroom (14.1) in rating $e$, which by equation (14.3) is $g_e(\mathbf{z})$ in the draws where $e$ binds first and zero in the others. It is reported as the average over the draws, and multiplied by the step $\Delta_e$ to put the four elements on one scale, in kW. Beside it stands the reading at the MAP state, equation (14.2).

**Building blocks.** The binding loads (14.5) in each draw and at $\hat{\mathbf{z}}$; the rating gradients (14.6) and (14.7); the draws (Section 14.2.2); the stated steps.

Table 14.2: The four elements' ratings as drawn and as upgraded, and the gradient of each element's binding load in its own rating at the MAP state of the site's fit, at the design day. The last column is the gradient times the step: what the step would buy if its element bound first over the whole step.

| Element | Rating [unit] | As drawn [rating's] | Upgraded [rating's] | Gradient $g_e(\hat{\mathbf{z}})$ [kW per rating's unit] | Gradient times step [kW] |
|---|---|---|---|---|---|
| Chillers | added nameplate [kW] | 0 | 400 | 0.9213 | 368.5 |
| Pipe | flow-coefficient factor [–] | 1.00 | 1.25 | 1 148.8 | 287.2 |
| Pump | motor [kW] | 12 | 18 | 46.66 | 280.0 |
| Feed | rating [kVA] | 1 600 | 2 000 | 0.7337 | 293.5 |

Table 14.3: The importance measures of the four elements: at the MAP state, equation (14.2), and expected over the posterior, equation (14.3), by the first 4 000 draws of the site. Per-unit values are in kW per unit of the element's rating (Table 14.2).

| Element | $I_e^{\mathrm{MAP}}$ [kW per rating's unit] | $I_e^{\mathrm{MAP}}\,\Delta_e$ [kW] | $\Pr(e_1 = e)$ [–] | $I_e$ [kW per rating's unit] | $I_e\,\Delta_e$ [kW] |
|---|---|---|---|---|---|
| Chillers | 0 | 0 | 0.3595 | 0.3286 | 131.4 |
| Pipe | 1 148.8 | 287.2 | 0.6405 | 735.36 | 183.8 |
| Pump | 0 | 0 | 0.0000 | 0 | 0.0 |
| Feed | 0 | 0 | 0.0000 | 0 | 0.0 |

**The gradients.** Table 14.2 gives each element's rating, its value as drawn and upgraded, and the gradient of its own binding load in that rating at the MAP state, with the gradient times the step, the headroom the step would buy if its element bound first and kept binding. The chillers' gradient is 0.9213 kW per kW of nameplate, the as-built ratio of an added unit; the pipe's 1 148.8 kW per unit of the factor, its whole carrying capacity in heat; the pump's 46.66 kW of load per kW of motor; and the feed's 0.7337 kW per kVA. Over its own step each would buy between 280.0 and 368.5 kW: steps of similar size, priced from R$ 0.15 to R$ 1.60 million.

**The importance measures.** Table 14.3 gives the two readings. Each row is an element; the columns give the reading at the MAP state, per unit of rating and times the step, the probability that the element binds first over the draws, the expected derivative (14.3) per unit of rating, and that times the step.

At the MAP state the pipe binds first, so only the pipe has an importance: 1 148.8 kW per unit of the factor, 287.2 kW over its step, and zero for the other three. That is the sheet's reading made with the fitted state instead of the handbook's, and it agrees with the sheet's ranking. Over the posterior the chillers enter. They bind first in 0.3595 of the draws, and their importance is 0.3286 kW per kW of nameplate, 131.4 kW over the 400 kW step; the pipe binds first in 0.6405 of the draws, and its importance falls to 735.36 per unit of the factor, 183.8 kW over its step. The pump and the feed never bind first and have no importance at either reading. The posterior ranking keeps the pipe first and puts the chillers a close second, a ranking the MAP reading cannot make, since it gives every element but one an importance of zero.

### 14.3.2 The headroom each upgrade buys

**Planner's question.** If I could upgrade one element, which one buys the most headroom? The question asks for the expected change of the headroom over each stated step, read jointly

Table 14.4: The headroom each stated step buys at the design day, the plant as drawn, joint over the four elements, by the first 4 000 draws of the site; the exact posterior's mean beside. Monte Carlo standard errors of the expected gains: pipe 0.29 kW, chillers 0.14 kW, pump and feed 0. The headroom now has a median of 463.8 kW, band 442.9 to 473.8 kW.

| Step | Price [R$ M] | Expected gain (exact) [kW] | 5 % point [kW] | 95 % point [kW] | Pr($G > 1$ kW) [–] | Gain per real [kW per R$ M] |
|---|---|---|---|---|---|---|
| Pipe, factor 1.25 | 0.90 | 15.5 (15.1) | 0.0 | 53.3 | 0.6222 | 17.2 |
| Chillers, a fourth 400 kW | 1.60 | 4.9 (5.1) | 0.0 | 24.9 | 0.3417 | 3.1 |
| Pump, 12 to 18 kW | 0.15 | 0.0 (0.0) | 0.0 | 0.0 | 0.0000 | 0.0 |
| Feed, to 2 000 kVA | 0.80 | 0.0 (0.0) | 0.0 | 0.0 | 0.0000 | 0.0 |

over the four elements, and for that change per real.

**General form.** A design question chooses the action of least expected loss, equation (2.23),

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

**Specialized form.** The question is equation (2.23) with the actions the four steps and the loss of a step the headroom it fails to buy, $L(e, \mathbf{z}) = -G_e(\mathbf{z})$ of equation (14.4), so that

$$a^* = \arg\max_e \bar{G}_e, \qquad \bar{G}_e / p_e, \tag{14.11}$$

with $\bar{G}_e$ by the draws (14.10) and the exact posterior beside it, the 90 % band of $G_e$ over the draws, and the probability $\Pr(G_e > 1\,\text{kW} \mid \mathbf{y})$ that the step buys anything at all. No value of a kilowatt of headroom in reais is stated, so the price $p_e$ is not netted against the gain: the gain per million reais is reported beside it.

**Building blocks.** The binding loads (14.5) at the drawn and the upgraded ratings in each draw, the scenario (14.8); the draws and their weights (Section 14.2.2); the stated prices.

**The gains.** Table 14.4 gives the four steps. Each row is a step; the columns give its price, the expected gain by the draws with the exact posterior's beside it, the 5 % and 95 % points of the gain, the probability that the gain exceeds 1 kW, and the gain per million reais. Over the same draws the headroom now has a median of 463.8 kW and a 90 % band of 442.9 to 473.8 kW. Figure 14.1 plots the distribution function of the gain for the pipe and the chiller steps through the deciles and the 95 % point the computation prints, with the pair of Section 14.3.4 beside them and the stated truth's gains dotted.

The pipe buys the most, and it buys little. Its expected gain is 15.5 kW, 15.1 kW by the exact posterior, with a 90 % band of 0 to 53.3 kW, and it buys more than 1 kW with probability 0.6222; per real, 17.2 kW per million. The chillers' step buys 4.9 kW, 5.1 kW by the exact posterior, band 0 to 24.9 kW, more than 1 kW with probability 0.3417, 3.1 kW per million. The pump and the feed buy nothing in any draw: neither binds first, so raising either moves only a binding load that is not the smallest. The answer to the drill question is the pipe, as the sheet said, but for 15.5 kW of a 463.8 kW headroom.

The distribution in Figure 14.1 shows why a mean is a poor summary of a single step. Each single step's gain is a spike at zero, the draws in which its element does not bind first, and a tail, the draws in which it does: 0.3778 of the pipe step's probability lies at or below 1 kW and 0.6583 of the chiller step's. The probability that the pipe buys more than 1 kW, 0.6222, is a

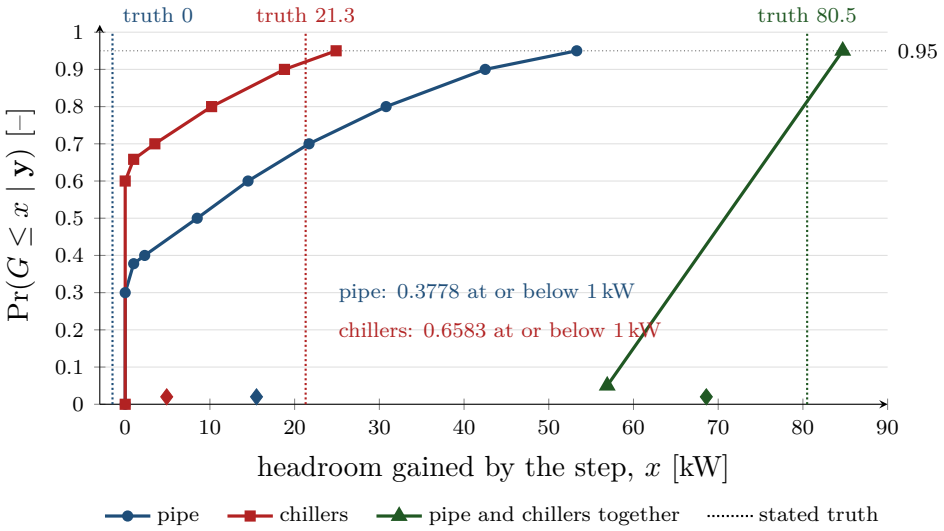



Figure 14.1: The distribution function of the headroom gained by the pipe step, the chiller step and the two together (Tables 14.4 and 14.6). For each single step the markers are the printed deciles, the probability of a gain of at most 1 kW, and the 95 % point; the vertical rise at zero is the spike, the draws in which the step buys nothing. For the pair the line joins the 5 % and 95 % points. Straight lines join printed points and do not claim the shape between them. Diamonds on the axis: the expected gains. Dotted verticals: the stated truth's gains, the pipe's 0 offset to the left of zero so that the rise there does not hide it.

little below the probability that it binds first, 0.6405, because in a few draws the next element binds within a kilowatt of it. The deciles put the median of the pipe step's gain at 8.5 kW, below its mean of 15.5 kW, and the chiller step's at zero, with its 70 % point at 3.5 kW: in half the posterior the pipe step buys less than 9 kW, and the mean is carried by the upper deciles. Even the tail is short: the 95 % point of the pipe's gain is 53.3 kW, a fifth of the 287.2 kW that its gradient promises over the step.

### 14.3.3 Why the gradient and the upgrade disagree

**Planner's question.** Why do the gradient and the upgrade disagree? The question asks where the 183.8 kW that the pipe's importance measure promises over its step goes, when the step buys 15.5 kW.

**General form.** The moments of a derived quantity under the posterior, equation (2.10),

$$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}] = \int f(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z},$$

here of the difference between two derived quantities, one of which is a minimum.

**Specialized form.** The question is equation (2.10) with $f = g_e\, \Delta_e\, \mathbf{1}[e_1 = e] - G_e$, the step's promise less its gain. For the chillers and the pipe, whose binding loads are linear in their own ratings by equation (14.5), the step raises the element's own binding load by exactly $g_e\, \Delta_e$. If the element does not bind first, raising it leaves the minimum where it is, since neither upgrade lowers another element's binding load and the ones it raises, the pump's and the feed's for the pipe, never bind first. If it does bind first, the minimum rises with it until the next element binds. Write $D_e^* = \min_{e' \neq e} Q_{e'}^{\dagger *} - Q_e^{\dagger}$ for the distance from the element's binding load to the

Table 14.5: What the pipe and the chiller steps buy, by five routes. The importance rows multiply equations (14.2) and (14.3) by the step; the gain at the MAP state is equation (14.12) at $\hat{\mathbf{z}}$, with the binding loads Chapter 13 prints there; the last two rows are the joint gain over the first 4 000 draws. The last column is the gap to the reported value.

| Route | Pipe step [kW] | Gap to reported [kW] | Chiller step [kW] | Gap to reported [kW] |
|---|---|---|---|---|
| Importance at the MAP times the step, $I_e^{\mathrm{MAP}}\,\Delta_e$ | 287.2 | 271.7 | 0 | −4.9 |
| Importance over the posterior times the step, $I_e\,\Delta_e$ | 183.8 | 168.3 | 131.4 | 126.5 |
| Gain at the MAP state | 7.25 | −8.25 | 0 | −4.9 |
| Joint gain by the draws (reported) | 15.5 | – | 4.9 | – |
| Joint gain, exact posterior | 15.1 | −0.4 | 5.1 | 0.2 |

nearest of the others at the upgraded ratings, in kW. Then, in every state,

$$G_e = \mathbf{1}[e_1 = e]\,\min(g_e\,\Delta_e,\ D_e^*), \qquad I_e\,\Delta_e - \bar{G}_e = \mathbb{E}\Big[\mathbf{1}[e_1 = e]\,\max(g_e\,\Delta_e - D_e^*,\ 0)\ \Big|\ \mathbf{y}\Big] \geq 0, \tag{14.12}$$

the second by taking the expectation of the first and of $g_e\,\Delta_e\,\mathbf{1}[e_1 = e]$, whose mean is $I_e\,\Delta_e$ by equation (14.3). The importance measure times the step overstates the gain by the expected part of the step that lies beyond the next element. It is reported as the difference of the two means already computed; beside it stand the single-state routes, the importance at the MAP times the step and the gain at the MAP state.

**Building blocks.** The importance measures (Table 14.3); the gains (Table 14.4); the binding loads at the MAP state, which Chapter 13 prints.

**The routes.** Table 14.5 sets the readings of the pipe and the chiller steps side by side. Each row is a route; the columns give what it says each step buys. The first two rows are the importance measures times the step, at the MAP state and over the posterior; the third is the gain at the MAP state alone; the last two are the reported gain by the draws and the exact posterior's. At the MAP state Chapter 13 puts the pipe's binding load at 1 126.81 kW and the chillers' at 1 134.06 kW, 7.25 kW apart, and the feed's above both; so at the MAP the pipe step buys 7.25 kW, the distance to the chillers, and the chiller step nothing. Figure 14.2 draws the two importance rows and the reported gain as bars for all four elements, each element's price under its name, with the stated truth's gain of each step dotted across its group.

The importance measure overstates the pipe step's gain by 168.3 kW, about twelvefold, and the chiller step's by 126.5 kW. Every route built on a gradient, Figure 14.2 shows, ranks the elements in the same order as the joint gain and is wrong about the size by an order of magnitude. The reason is equation (14.12): a step pays only up to the distance to the next element that binds, and that distance is short. Chapter 13 finds the pipe's binding load a median 8.8 kW below the chillers' and the feed's about 61 kW above the pipe's, while the pipe's step would raise its own binding load by 287.2 kW at the MAP. The step runs past the chillers, or past the feed, long before it runs its course, so nearly all of its promise lies beyond the next element. The chiller step is capped the same way by the pipe, which lies within a few tens of kilowatts of the chillers in the draws where the chillers bind first.

This is the chapter's surprise, and it holds beyond this site. An importance measure is a derivative and so a local statement: it assumes that the element stays the one that binds first over the whole step. Where the elements bind close together, as they do whenever a plant has been sized to balance its elements, no step of a useful size stays local. The gain at the MAP

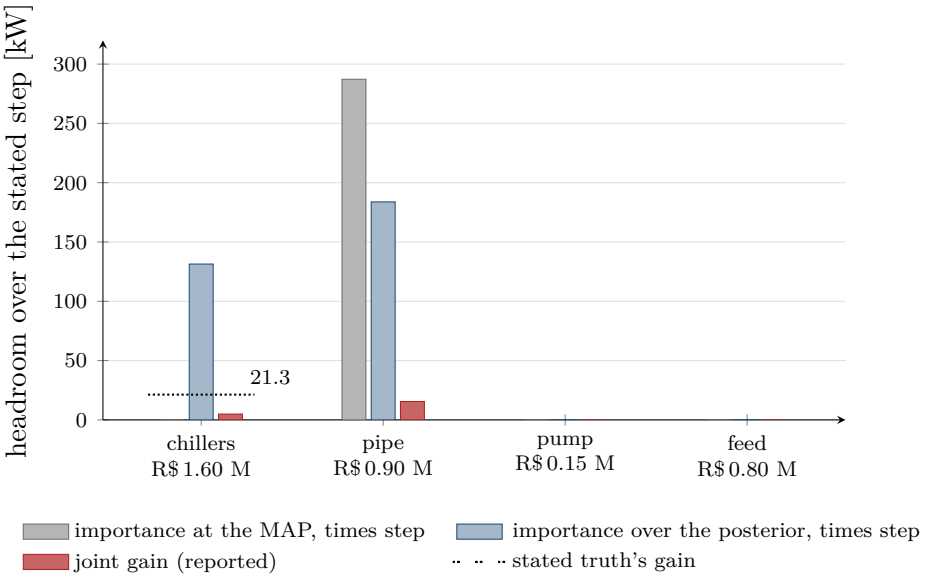



Figure 14.2: What each stated step promises and what it buys, per element: the importance at the MAP state and over the posterior, each times the step (Table 14.3), against the joint expected gain (Table 14.4). The pump and the feed have none of the three. Dotted: the stated truth's gain, 21.3 kW for the chiller step and zero for the other three.

state errs the other way: at 7.25 kW it understates the reported gain, because it sees the one distance to the chillers at their most likely value, where the draws see the pipe first by more in some states and the feed as the cap in others. The planner reads the importance measure as where to look, and the joint gain as what to buy.

### 14.3.4 Two upgrades together

**Planner's question.** What do two upgrades buy together? The question asks the drill question of two steps at once: the pipe with the chillers, the two elements that bind first, and the pipe with the feed, the element that caps the pipe's step in part of the posterior.

**General form.** The design form of Section 14.3.2, equation (2.23),

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

with an action now a pair of steps.

**Specialized form.** The question is equation (2.23) with the loss of a pair the headroom it fails to buy,

$$G_{e,e'}(\mathbf{z}) = H(\mathbf{z},\ \mathbf{r} + \Delta_e\,\mathbf{e}_e + \Delta_{e'}\,\mathbf{e}_{e'}) - H(\mathbf{z},\ \mathbf{r}), \qquad L = -G_{e,e'}, \tag{14.13}$$

read by the draws as in equation (14.10), with its band, its price as the sum of the two stated prices, and the element that binds first after the pair, read on the first 1 000 draws.

**Building blocks.** The binding loads (14.5) at the drawn ratings and with both steps; the draws; the stated prices.

Table 14.6: Two steps together at the design day, the plant as drawn, against the single pipe step and the sum of the two single steps; first 4 000 draws of the site. The element that binds next is read on the first 1 000 draws.

| Action | Price [R$ M] | Expected gain [kW] | 90 % band [kW] | Gain per real [kW per R$ M] | Binds first after [share of draws] |
|---|---|---|---|---|---|
| Pipe and chillers | 2.50 | 68.6 | 56.9 to 84.7 | 27.4 | feed 1.000 |
| Pipe and feed | 1.70 | 15.7 | 0.0 to 53.3 | 9.2 | chillers 1.000 |
| Pipe alone | 0.90 | 15.5 | 0.0 to 53.3 | 17.2 | – |
| Pipe alone plus chillers alone | 2.50 | 20.4 | – | – | – |

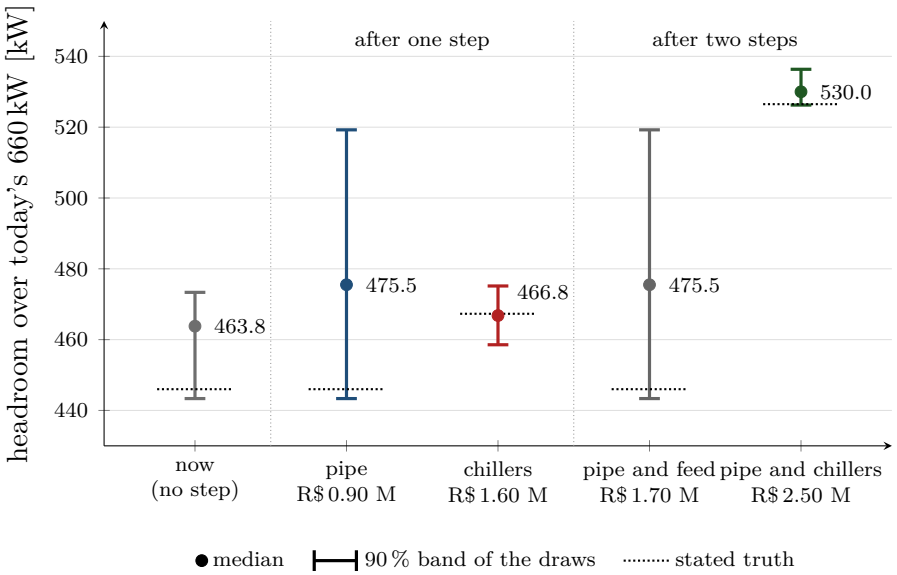



Figure 14.3: The site's headroom over today's 660 kW at the design day, now, after one step and after two, with each action's stated price under its name (Tables 14.4 and 14.6): dots at the median of the first 4 000 draws, bars over their 90 % band, labels at the median. The pump step and the feed step alone leave the headroom as it is now. Dotted: the stated truth's headroom, 446.0 kW now and after the pipe step alone or with the feed, 467.3 kW after the chiller step and 526.5 kW after the pipe with the chillers.

**The pairs.** Table 14.6 gives the two pairs, with the single pipe step and the sum of the two single steps beside them. Each row is an action; the columns give its price, the expected gain and its band, the gain per million reais, and the element that binds first after it; the Monte Carlo standard errors of the pairs' expected gains are 0.14 kW (the pipe with the chillers) and 0.31 kW (the pipe with the feed). Figure 14.3 plots the headroom itself, now, after one step and after two: for each action the median and the 90 % band of the headroom over the draws, with the stated truth's headroom dotted.

The pipe and the chillers together buy 68.6 kW, band 56.9 to 84.7 kW, for R$ 2.50 million: 27.4 kW per million reais, the most per real of any action, and more than three times the 20.4 kW that the two single steps buy between them. The two steps are complementary. Each single step is capped by the other element, the pipe's by the chillers in most of the draws where the pipe binds first and the chillers' by the pipe; raised together, neither caps the other, and the feed binds next in every one of the 1 000 draws checked. The pair's gain is then the distance from today's first element to the feed, and its band lies clear of zero: unlike either step alone, the pair buys headroom across the posterior, whichever of the two binds first. Figure 14.3 shows what that does to the headroom itself. After the pair its band runs from 525.8 to 536.8 kW,

wholly above the headroom now and after either single step, whose highest point is the pipe step's 519.7 kW; and it is 11.0 kW wide against the 30.9 kW of the headroom now, because the feed, whose binding load the record reads closely, sets it in every draw. A single step moves the median a little, to 475.5 kW after the pipe and 466.8 kW after the chillers, and stretches the band rather than lifting it: the pipe step leaves the band's lower end at 442.9 kW, where the chillers bind first.

The pipe and the feed buy 15.7 kW for R$ 1.70 million, 0.2 kW more than the pipe alone for R$ 0.80 million more. The feed caps the pipe's step only in the draws where the chillers lie further away than the feed; raising the feed there lets the chillers bind instead, and they do in every one of the 1 000 draws checked, a few kilowatts later. The importance measures do not add either: the pair's 68.6 kW is far below the $183.8 + 131.4$ kW the two measures promise over their steps and far above the sum of the single gains, and only the joint computation finds it.

## 14.4 The checks (step 5)

The answers of Section 14.3 rest on a subset of the site's draws, on the Laplace normal of its fit and on a record generated from a stated truth; each is tested here, and the reported upgrade is tested against the world the truth describes.

### 14.4.1 The draws against the exact posterior

The reported values average the first 4 000 draws with equal weight; the exact posterior reweights them by equation (5.16), renormalized over the 4 000. The exact mean gain of the pipe step is 15.1 kW against the reported 15.5 kW, a gap of $-0.4$ kW; that of the chiller step 5.1 kW against 4.9 kW, a gap of 0.2 kW; those of the pump and the feed are zero by both. The Monte Carlo standard errors of the reported means are 0.29 kW for the pipe step and 0.14 kW for the chiller step, so each gap is about one and a half standard errors; the gaps are a few percent of the gains and do not change their order, nor the order of the gain per real. The weights of the 4 000 draws have an effective sample size of 3 983, as the 19 913 of all 20 000 (Section 9.2.4) lead one to expect, so the Laplace normal loses little against the exact posterior. The subset's ranking differs from Chapter 13's on all the draws by 0.0100 in the probability that the pipe binds first, 0.6405 against 0.6505, the Monte Carlo error of using a fifth of the draws.

### 14.4.2 The answers against the stated truth

The operating record is generated from a stated truth, the last column of Table 9.2. At the truth the chillers bind first, at 1 106.0 kW of site IT load; the pipe at 1 127.3 kW, the feed at 1 184.7 kW and the pump at 1 250.4 kW. The truth's order is the second most likely of the posterior's, the chillers first, to which the draws give probability 0.3595. In the truth's world the chiller step buys 21.3 kW, the distance from the chillers to the pipe, and the pipe step buys nothing, nor do the pump and the feed. Table 14.7 collects the checks: each row is a quantity, the columns its reported value, the value beside it, the truth, and the reading.

The truth's headroom now, $1\,106.0 - 660 = 446.0$ kW, lies inside the reported band. Both true single-step gains lie inside their bands, the pipe's at the lower end, in the spike at zero, and the chillers' near the upper end, in the tail. The checks find no fault in the method, and they make one thing concrete that the report must carry. The reported upgrade, the pipe, buys nothing in the truth's world, and the step the report ranks second buys 21.3 kW. That is the 0.36 of the posterior in which the chillers bind first, made into a plant: an outcome the posterior gave about one chance in three, which the record cannot rule out. The reported answer is the best single step on the evidence, and the evidence leaves a one-in-three chance that it buys nothing; the report says both. The pair holds in the truth's world: the pipe with the chillers

Table 14.7: The checks of the importance and upgrade answers. Reported values by the first 4 000 draws of the site; beside them the exact posterior or the band; the truth is the stated truth of the record.

| Quantity | Reported [kW or –] | Beside [kW or –] | Truth [kW or –] | Reading |
|---|---|---|---|---|
| Headroom now, median | 463.8 | 442.9 to 473.8 | 446.0 | inside the band |
| Element that binds first | pipe, 0.6405 | chillers, 0.3595 | chillers | the second reading |
| Gain, pipe step | 15.5 | exact 15.1 | 0.0 | at the lower end of [0.0, 53.3] |
| Gain, chiller step | 4.9 | exact 5.1 | 21.3 | inside [0.0, 24.9] |
| Gain, pump step | 0.0 | exact 0.0 | 0.0 | no element moved |
| Gain, feed step | 0.0 | exact 0.0 | 0.0 | no element moved |
| Gain, pipe and chillers | 68.6 | 56.9 to 84.7 | 80.5 | inside, in the upper part |
| Gain, pipe and feed | 15.7 | 0.0 to 53.3 | 0.0 | at the lower end |

Table 14.8: The report for the site's upgrades at the design day, the plant as drawn, against today's 660 kW; first 4 000 draws of the site, the exact posterior beside. Headroom now: median 463.8 kW, band 442.9 to 473.8 kW.

| Question | Answer | Band or routes beside |
|---|---|---|
| If I could upgrade one element, which one buys the most headroom? | the pipe: 15.5 kW expected (exact: 15.1 kW), 17.2 kW per R$ M; the chillers 4.9 kW (exact: 5.1 kW); the pump and the feed nothing | pipe 0 to 53.3 kW, $\Pr(G > 1\,\text{kW}) = 0.6222$; chillers 0 to 24.9 kW, 0.3417 |
| What is each element's importance measure? | times the step: pipe 183.8 kW, chillers 131.4 kW, pump and feed 0 | at the MAP: pipe 287.2 kW, the others 0; $\Pr(e_1 = e)$: pipe 0.6405, chillers 0.3595 |
| Why do the gradient and the upgrade disagree? | the step runs past the next element: the measure overstates the pipe's gain by 168.3 kW, the chillers' by 126.5 kW | gain at the MAP state: pipe 7.25 kW, chillers 0 |
| What do two upgrades buy together? | pipe and chillers 68.6 kW, 27.4 kW per R$ M, the feed then binds; pipe and feed 15.7 kW | 56.9 to 84.7 kW; 0 to 53.3 kW; the two single steps sum to 20.4 kW |

buys 80.5 kW there, inside its band of 56.9 to 84.7 kW, while the pipe with the feed, like the pipe alone, buys nothing.

## 14.5 The report (step 6)

Upgrade the pipe if only one element can be touched, and the pipe with a fourth chiller if the budget allows. The pipe step buys an expected 15.5 kW of headroom, with a 0.62 chance of buying anything; the pair buys 68.6 kW, band 56.9 to 84.7 kW, whichever element binds first. Neither is sized from the importance measure, which promises twelve times what the pipe step buys.

### 14.5.1 The answers and their bands

Table 14.8 gives the answers in the order of Table 14.1. Each row is a question; the second column gives the reported answer with the exact posterior beside it where it differs, and the third its band or the routes beside it.

### 14.5.2 The design-review entry

The design review's entry replaces the sheet's bottleneck with the gain of each priced step. At the design day, with the plant as drawn, the site's headroom over today's 660 kW has a median of 463.8 kW, band 442.9 to 473.8 kW. If one element is upgraded, it is the shared pipe: one size up, for R$ 0.90 million, it buys an expected 15.5 kW (15.1 kW by the exact posterior), with a 90 % band of 0 to 53.3 kW and a probability of 0.62 of buying more than a kilowatt; a fourth chiller alone buys 4.9 kW for R$ 1.60 million, and the pump motor and the feed buy nothing. If the budget allows R$ 2.50 million, the pipe and a fourth chiller together buy 68.6 kW, band 56.9 to 84.7 kW, 27.4 kW per million reais, the best of any action priced, after which the feed binds. The entry records that the importance measure, 183.8 kW for the pipe's step and 131.4 kW for the chillers', is not a sizing figure: it assumes that the element stays the bottleneck over the step, and here the next element binds within tens of kilowatts.

**Common mistake**

Sizing an upgrade from its importance measure. The gradient of the headroom times the step is what the step would buy if its element kept binding; the step buys only up to the next element that binds, and the joint gain over every element's binding load is the figure to buy on.

## 14.6 The next measurement (step 7)

Step 7 asks which test or survey of the existing plant would narrow the quantity that decides the upgrade. The single-step choice and the pair answer this differently, and the difference is the chapter's guide to when to measure.

### 14.6.1 The economizer decides the single step

The single-step choice turns on which of the pipe and the chillers binds first: the pipe step buys only in the draws where the pipe binds first, and the chiller step only in the others. That question is Chapter 13's, and Chapter 13 finds that the economizer's handbook conductance carries 58 % of the variance of the gap between the pipe's and the chillers' binding loads; the economizer is drawn but not built, so no reading of the plant sees it. A test of the economizer's conductance, or of a comparable heat exchanger on the cooling-tower water, would move the probability that the pipe binds first toward zero or one, and with it the choice between the pipe and the chiller steps. The pair does not need that test: its gain is the distance from whichever element binds first to the feed, and its band lies clear of zero whichever it is. Chapter 18 ranks the tests and surveys a planner can commission on the site by the width each takes off the quantity that decides the build; for a single upgrade, the economizer's conductance is the measurement to price there first, and for the pair none is needed.

### 14.6.2 What the analysis showed

The design sheet ranked the upgrades by its binding order and chose the pipe. The analysis agrees with the choice and not with its reasons. The pipe is the best single step, at 15.5 kW of expected headroom for R$ 0.90 million, but it is the best by a margin the posterior holds with about two chances in three, and it buys about a twelfth of what its importance measure promises over the step.

Three findings carry the answer. First, and beyond this site, an importance measure is local: it is the derivative of a minimum, the element's gradient where it binds first, and a step of useful size runs past the next element that binds. Here the pipe's measure times its step is 183.8 kW against a gain of 15.5 kW, because the chillers and the feed bind within tens of kilowatts of the

pipe. Second, the MAP reading of importance gives every element but one a value of zero; over the posterior the chillers carry an importance of 131.4 kW over their step, close to the pipe's, because they bind first in a third of the posterior. Third, the steps are complementary: the pipe and a fourth chiller together buy 68.6 kW, more than three times what the two buy alone, and they buy it whichever of the two binds first. The checks make the second finding concrete: in the truth's world the chillers bind first, and the pipe step buys nothing, while the pair buys 80.5 kW, inside its band.

### 14.6.3 Limits of the model

**The plant as drawn.** The headroom and the gains are those of the plant with its drawn economizer, whose conductance keeps the handbook's guess. Chapter 13 finds that without the economizer, the plant as built today, the chillers bind first with probability 0.9970; the chiller step would then be the one that buys headroom, and the ranking of the single steps would reverse. This chapter does not compute the gains for the plant without the economizer.

**The design day and today's load.** The headroom is measured at the design day's wet bulb and against today's 660 kW. On cooler days the economizer does more and the chillers can carry more, which moves the binding order toward the pipe; the year's hours are Chapter 23's.

**The stated steps.** The steps and prices are planning inputs. The pipe one size up is taken as a factor of 1.25 on its flow coefficient with the same allowable drop, and the fourth chiller at the mean as-built ratio of the existing three, with no spread of its own. A different step would change the gains, though not the reason they fall short of the importance measures, which is the distance to the next element.

**All units up, no shared-path states.** Every unit is taken up and the shared pipe is taken as available. The fourth chiller is here capacity, not redundancy; what it buys as redundancy is Chapter 9's question, and the shared pipe as a state of its own is Chapter 15's.

**The subset of draws.** The answers read the first 4 000 of the site's 20 000 draws. The probabilities that an element binds first differ from Chapter 13's on all the draws by about a hundredth, and the Monte Carlo standard errors of the mean gains are at most 0.31 kW.

**The constructed record.** The operating record is generated from a stated truth with the stated reading errors, not taken from a site, and the comparison with that truth in Section 14.4.2 is a check of the method on its own model, not of the model against a real plant.

## 14.7 What to retain

**Definition 14.1** (Capacity importance measure and the gain of a step)**.** For a plant whose headroom is the smallest of its elements' binding loads less today's load, $H = \min_e Q_e^{\dagger} - Q_0$, the importance measure of element $e$ in the capacity sense is the expected derivative of the headroom in the element's rating, $I_e = \Pr(e \text{ binds first} \mid \mathbf{y})\, \mathbb{E}[g_e \mid e \text{ binds first}, \mathbf{y}]$, with $g_e$ the gradient of the element's own binding load. The gain of a stated step is the expected change of the headroom over the step, read jointly over every element's binding load. For an element whose binding load is linear in its rating, the gain in a state is the step's rise or the distance to the next element, whichever is smaller, and the importance measure times the step exceeds the gain by the expected part of the step beyond the next element.

Four conclusions carry beyond this site. An importance measure ranks where the headroom is sensitive, not what a step buys; it is a derivative, and a step of useful size leaves its range. A measure read at one state gives every element but the bottleneck a value of zero; read over the posterior, it weights each element by the probability that it binds first. Upgrades to elements that bind close together are complementary: each caps the other, so two together can buy more than twice what each buys alone. And the choice of a single upgrade inherits the uncertainty of the binding order, while an upgrade that clears every close element does not; the first may call for a measurement, the second may not.

Table 14.9 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 14.9: The terms of capacity importance measures and where each is defined. Units for the site's elements; $r_e$ in its own unit.

| Term | Expression | Units | Meaning | Defined in |
|---|---|---|---|---|
| Headroom | $H = \min_e Q_e^{\dagger}(\mathbf{z}, \mathbf{r}) - Q_0$ | kW | smallest binding load less today's | 14.1.1 |
| Rating gradient | $g_e = \partial Q_e^{\dagger} / \partial r_e$ | kW per unit of $r_e$ | one element's own slope | 14.1.2 |
| Importance at the MAP | $I_e^{\mathrm{MAP}} = g_e(\hat{\mathbf{z}})\,\mathbf{1}[e_1(\hat{\mathbf{z}}) = e]$ | kW per unit of $r_e$ | nonzero for one element | 14.1.1 |
| Importance measure | $I_e = \Pr(e_1 = e)\,\mathbb{E}[g_e \mid e_1 = e]$ | kW per unit of $r_e$ | expected derivative | 14.1.1 |
| Gain of a step | $\bar{G}_e = \mathbb{E}[H(\mathbf{r} + \Delta_e \mathbf{e}_e) - H(\mathbf{r})]$ | kW | what the step buys | 14.1.1 |
| *Questions and answers* | | | | |
| Gain per real | $\bar{G}_e / p_e$ | kW per R$ M | price beside, not netted | 14.3.2 |
| Gain in a state | $G_e = \mathbf{1}[e_1 = e] \min(g_e \Delta_e,\ D_e^*)$ | kW | capped by the next element | 14.3.3 |
| Overstatement | $I_e \Delta_e - \bar{G}_e = \mathbb{E}[\mathbf{1}[e_1 = e](g_e \Delta_e - D_e^*)^+]$ | kW | the step beyond the next element | 14.3.3 |
| Gain of a pair | $G_{e,e'} = H(\mathbf{r} + \Delta_e \mathbf{e}_e + \Delta_{e'} \mathbf{e}_{e'}) - H(\mathbf{r})$ | kW | complementary steps | 14.3.4 |

## 14.8 Exercises

**Exercise 14.1.** *Formulation.* The review asks how likely it is that the pipe step buys any headroom at all. Write the question over the site's state and give its answer. ($\Pr(G_{\mathrm{p}} > 1\,\mathrm{kW} \mid \mathbf{y}) = \int \mathbf{1}[G_{\mathrm{p}}(\mathbf{z}) > 1\,\mathrm{kW}]\, p(\mathbf{z} \mid \mathbf{y})\,\mathrm{d}\mathbf{z}$, with $G_{\mathrm{p}}$ of equation (14.4), the share of the draws in which the pipe step raises the headroom by more than a kilowatt; a little below the probability 0.6405 that the pipe binds first, it is 0.6222.)

**Exercise 14.2.** *Derivation.* Derive the pipe's rating gradient and check it against Chapter 13's binding load at the MAP state. (By equation (14.5) the pipe's binding load is $r_{\mathrm{p}} k_p \sqrt{150}\,\rho c_p\,\Delta T_w - Q_o$, linear in the factor $r_{\mathrm{p}}$, so its gradient is $k_p \sqrt{150}\,\rho c_p\,\Delta T_w$, which at $r_{\mathrm{p}} = 1$ is the binding load plus the other heat: $1\,126.81 + 22 = 1\,148.81\,\mathrm{kW}$ per unit of the factor, the computation's 1 148.8 kW.)

**Exercise 14.3.** *Derivation.* Show that the gain of a step on an element whose binding load is linear in its rating is $\mathbf{1}[e_1 = e] \min(g_e \Delta_e, D_e^*)$, and read the pipe step's overstatement from the tables. (Raising a rating that does not bind first leaves the minimum in place; raising the one that binds first lifts the minimum with it until it meets the nearest other binding load, a distance $D_e^*$ above; so the minimum rises by $\min(g_e \Delta_e, D_e^*)$. The mean of $\mathbf{1}[e_1 = e]\, g_e \Delta_e$ is $I_e \Delta_e = 183.8\,\mathrm{kW}$ for the pipe, and the overstatement is $183.8 - 15.5 = 168.3\,\mathrm{kW}$, leaving the reported gain of 15.5 kW.)

**Exercise 14.4.** *Application.* A colleague proposes to buy the pipe now and the fourth chiller next year, arguing that the two single steps add up. Compare the sum of the single gains with the pair's gain, per million reais. (The single steps buy 15.5 and 4.9 kW, 20.4 kW for R$ 2.50 million, about 8.2 kW per million; the pair buys 68.6 kW for the same R$ 2.50 million, because each step removes the other's cap and the feed binds next: $68.6/2.50 = 27.4$ kW per million reais.)

**Exercise 14.5.** *Application.* In the world of the stated truth, which single step would the planner want, and what would it buy? (At the truth the chillers bind first, at 1 106.0 kW, and the pipe next, at 1 127.3 kW. The pipe step raises a binding load that is not the smallest and buys nothing; the chiller step raises the chillers' by several hundred kilowatts, past the pipe, so it buys the distance to the pipe, $1\,127.3 - 1\,106.0 = 21.3$ kW.)

# Chapter 15

# Shared Paths and Common-Mode Limits

## 15.1 One pipe under two halls

The site's drawing shows the two halls side by side, each with its own three cooling units, and one supply pipe that carries the chilled water from the plant to both. A reviewer at the design review puts a finger on that pipe. Everything downstream of it is duplicated, hall by hall and unit by unit; the pipe is not. Its strainer can foul, and a fouled strainer throttles the flow to both halls at once. One element, one event, two halls short together.

The design sheet has an answer ready. It lists the pipe once, as an element of each hall's path, at its allowable drop of 150 kPa, and gives it a flow margin well above today's load. It lists each hall's three units at their nameplate and the three chillers at theirs. Read element by element, nothing on the sheet is short, and the pipe has the largest margin of all. What the sheet cannot say is which element actually limits the site when the site does fall short: the pipe, which rarely fails but takes everything when it does, or the halls' own units, which fail more often but one hall at a time. The planner's question is the drill question, verbatim,

*Two halls share one pipe; what is the probability that the shared pipe, not the halls' own units, is what limits them?*

It is answered with the method of Chapter 4 on the site of Chapter 9, with two changes. Both halls enter the calculation together, each with its own calibration, so hall A is calibrated here for the first time. And the shared pipe, which Chapter 9 took as always available, becomes a state of the enumeration: clean or fouled.

**Jargon: shared path, common-mode element, strainer, fouling, flow coefficient, air side, limiting element, cut set, minimal cut set, order of a cut set, importance, Fussell–Vesely importance**

A *shared path* is an element that carries what several otherwise separate parts of the plant need, here the supply pipe that feeds both halls. An element whose one state enters the capacity of several parts at once is a *common-mode element*; its failure is a *common-mode (shared-path) failure*. A *strainer* is the mesh in the pipe that catches debris; its *fouling* raises the pipe's resistance. The *flow coefficient* $k_p$ of the pipe relates its flow to its pressure drop, $V = k_p\sqrt{\Delta P}$. A hall's *air side* is its cooling units and the air they move, as distinct from the chilled-water plant. The *limiting element* of a state of the plant is the element whose capacity is the smallest in that state. A *cut set* is a set of elements whose failure together leaves the plant short; it is *minimal* when no smaller set within it does, and its *order* is the number of elements in it. The *importance* of an element ranks it by its part in the plant's failures; the *Fussell–Vesely importance* used here is the share of the shortfall probability that comes from states with that element down.

**Key idea**

A shared element is one state that sits in every path it serves. Its place in the answer is set by two things at once: how often it fails, and what is left when it does. The enumeration of the site's states carries both, and reading which element is smallest in each short state splits the shortfall probability among the elements. The minimal cut sets show where a shared element could take the site down alone; the split shows how often it actually does.

### 15.1.1 The shared path read as a cut set

Chapter 9 enumerated the states of the chillers alone. Here every element that can be down is a coordinate of the state: the three chillers C1–C3, the six hall cooling units A1–A3 and B1–B3, and the shared pipe, ten elements in all. Write the state as $\mathbf{s} = (s_e)$ with $s_e = 1$ when element $e$ is up (for the pipe: clean) and 0 when it is down (fouled). Each element is down a fraction $q_e$ of the time, independently, so the probability of a state is the product of equation (9.1) with one $q$ per element,

$$\begin{aligned} P(\mathbf{s} \mid \mathbf{q}) &= \prod_e q_e^{1-s_e}\,(1-q_e)^{s_e}, \\ \mathrm{Cap}(\mathbf{s};\ \boldsymbol{\theta}) = \min_{g\in\mathcal{G}} \mathrm{Cap}_g(\mathbf{s};\ \boldsymbol{\theta}), &\qquad \ell(\mathbf{s};\ \boldsymbol{\theta}) = \arg\min_{g\in\mathcal{G}} \mathrm{Cap}_g(\mathbf{s};\ \boldsymbol{\theta}), \end{aligned} \tag{15.1}$$

with $\mathbf{q}$ the ten unavailabilities, $\boldsymbol{\theta}$ the unknown of the plant, and $\mathcal{G}$ the six groups that can limit the site: hall A, hall B, the chillers, the pipe, the pump and the feed. The capacity of a state is no longer a sum, as it was for the chillers alone, but the smallest of the groups' capacities, each in site IT load (kW), since the load must pass through every group. The *limiting element* $\ell$ is the group that attains it. The site falls short at a site IT load $L$ in state $\mathbf{s}$ when $\mathrm{Cap}(\mathbf{s};\boldsymbol{\theta}) < L$.

Three readers of the enumeration follow, each a sum over the states,

$$\begin{aligned} P_L(\boldsymbol{\theta}) &= \sum_{\mathbf{s}} P(\mathbf{s} \mid \mathbf{q})\,\mathbf{1}\big[\mathrm{Cap}(\mathbf{s};\ \boldsymbol{\theta}) < L\big], \\ P_{L,g}(\boldsymbol{\theta}) &= \sum_{\mathbf{s}} P(\mathbf{s} \mid \mathbf{q})\,\mathbf{1}\big[\mathrm{Cap}(\mathbf{s};\ \boldsymbol{\theta}) < L\big]\,\mathbf{1}\big[\ell(\mathbf{s};\ \boldsymbol{\theta}) = g\big], \\ P_L^{(e)}(\boldsymbol{\theta}) &= \sum_{\mathbf{s}} P(\mathbf{s} \mid \mathbf{q})\,\mathbf{1}\big[\mathrm{Cap}(\mathbf{s};\ \boldsymbol{\theta}) < L\big]\,(1-s_e), \end{aligned} \tag{15.2}$$

all dimensionless: the shortfall probability given the unknown; the part of it from states that group $g$ limits; and the part from states with element $e$ down. The second splits the first,

$\sum_g P_{L,g} = P_L$, because every short state has one limiting group. The third does not: a state with two elements down counts toward both.

A set $S$ of elements is a *cut set* at the load $L$ when the state $\mathbf{s}_S$ with exactly the elements of $S$ down falls short,

$$\mathrm{Cap}(\mathbf{s}_S;\ \boldsymbol{\theta}) < L, \qquad S \text{ minimal if no } S' \subsetneq S \text{ is a cut set}, \qquad \text{order } |S|, \tag{15.3}$$

with $|S|$ the number of its elements. A cut set of order one is a single element whose failure alone leaves the site short. The derivation of minimal cut sets from a plant's structure, and of the probability of a common cause that enters several of them, is that of *Reliability Engineering by Derivation*, Chapters 14 and 15; this chapter reads the cut sets off the enumeration directly. A *common-mode element* is then one whose single state appears in the cut sets of several parts of the plant: the pipe's fouled state cuts the supply of hall A and of hall B together.

**Principle 15.1.** *A shared path is judged on the enumeration with the shared element as one state. Each short state has one limiting element, and the share of the shortfall probability from states that an element limits answers which element limits the plant; the minimal cut sets say where an element could take the plant down alone, and the shares say how often it does.*

### 15.1.2 The site with both halls and the pipe as a state (step 0)

The plant is the site of Chapter 9: the two halls on one chilled-water plant, with its three chillers, one pump, the shared supply pipe and the power feed, taken as built, without the drawn economizer. Each hall is the hall of Chapter 5: three cooling units drawing heat from the hall air into the chilled water, and five racks drawing their inlet air from the cold aisle with recirculation fractions $\beta_k$. Hall A's units are A1–A3 and its racks A1–A5, with the inlets of racks A1 and A3 measured; hall B's are those of Chapter 5.

**The halls' capacities.** A hall's capacity is its binding load, equation (5.8), at the drawing's supply of $T_s^* = 14°\mathrm{C}$ and the inlet limit $T_{\max} = 27°\mathrm{C}$, now with only its running units. A unit that is down passes no heat and moves no air, so both sums of equation (5.7) run over the running units $R$ alone,

$$a_k(\mathbf{z}_h;\ R) = \frac{1}{\sum_{u\in R} \bar{G}_u\, \mathrm{e}^{g_u}} - \frac{1-\beta_k}{\sum_{u\in R} C_u}, \qquad Q_h^{\dagger}(\mathbf{z}_h;\ R) = \min_{k\in h} \frac{T_{\max} - T_s^*}{a_k(\mathbf{z}_h;\ R)} - Q_{o,h}, \tag{15.4}$$

with $\mathbf{z}_h$ hall $h$'s state, $a_k$ the inlet rise of rack $k$ per kilowatt in K/kW, $\bar{G}_u$, $g_u$ and $C_u$ unit $u$'s rated conductance, log-derate and air capacity rate, and $Q_{o,h}$ the hall's other heat. With all three units running this is equation (5.8). Losing a unit lowers both sums; the first term rises and the second falls in magnitude, so every inlet rises faster with the load and the binding load falls. The hall's capacity in site IT load divides its binding load by its share of the site load,

$$\mathrm{Cap}_h(\mathbf{s};\ \boldsymbol{\theta}) = \frac{Q_h^{\dagger}(\mathbf{z}_h;\ R_h(\mathbf{s}))}{\varphi_h}, \qquad \varphi_A = 0.5455, \quad \varphi_B = 0.4545, \tag{15.5}$$

with $R_h(\mathbf{s})$ the units of hall $h$ up in state $\mathbf{s}$ and $\varphi_h$ the hall's share of today's 660 kW of site IT load: $0.5455 \times 660 = 360\,\mathrm{kW}$ in hall A and $0.4545 \times 660 = 300\,\mathrm{kW}$ in hall B.

**The plant's capacities.** The chillers, the pipe, the pump and the feed are those of equation (9.3) at the design day, a wet bulb of 22°C, each solved for the site IT load at which its limit binds, the binding load of Chapter 5. The chillers up carry the loop's heat, the IT load plus the halls' other heat $Q_o = 22\,\text{kW}$, with no economizer as built. The pipe's drop is $\Delta P = (V/k_p)^2$ with the loop flow $V = (Q_{\text{IT}} + Q_o)/(\rho c_p\,\Delta T_w)$, so the drop meets its allowable $\Delta P_{\max} = 150\,\text{kPa}$ when $V = k_p\sqrt{\Delta P_{\max}}$. A fouled strainer leaves the pipe a stated 60 % of its flow coefficient, so

$$\begin{aligned} \text{Cap}_{\text{C}}(\mathbf{s};\ \boldsymbol{\theta}) &= \sum_{i\in\{\text{C1,C2,C3}\}} s_i\,400\,\text{e}^{\kappa_i} - Q_o, \\ \text{Cap}_p(\mathbf{s};\ \boldsymbol{\theta}) &= \left[1 - 0.4\,(1 - s_p)\right]\rho c_p\,\Delta T_w\,\text{e}^{l_k}\sqrt{\Delta P_{\max}} - Q_o, \end{aligned} \tag{15.6}$$

in kW, with $s_p$ the pipe's state and $\rho c_p\,\Delta T_w$ the heat a liter of water carries over the loop's rise of 10 K. The pump's and the feed's capacities are their binding loads at the pump's 12 kW motor and the feed's 1 520 kW, by equation (9.3); neither can be down in this enumeration. Setting the six groups against each other in the healthy plant is the question of Chapter 13; here the halls' units and the pipe can fail.

**The outage states.** Each chiller and each hall unit fails and is repaired at its kind's rates, so that its unavailability is $q = \lambda/(\lambda + \mu)$ of equation (9.6). The pipe fouls at a stated rate $\lambda_p = 0.5$ a year and takes a stated $\tau_p = 12\,\text{h}$ to clear; with $\mu_p = 1/\tau_p$ the same relation gives

$$q_p = \frac{\lambda_p}{\lambda_p + 1/\tau_p} = \frac{x}{1+x}, \qquad x = \lambda_p\,\tau_p = \frac{0.5 \times 12}{8\,760}, \qquad q_p = 6.845 \times 10^{-4}, \tag{15.7}$$

dimensionless. The three inputs of the pipe, its rate, its clearing time and its fouled fraction, are stated, not estimated: no record of the strainer is kept.

**Modeling choices at the graph.** The pipe is one state, clean or fouled: this is what makes it a single element of every hall's path, and it fails if the strainer fouls by degrees, which would make its capacity a continuum rather than two values. The site load grows in today's proportions in the two halls: this is what turns each hall's binding load into a site IT load by equation (15.5), and it fails if one hall grows faster. A down unit's fan stops with its coil: this drops it from both sums of equation (15.4), and it fails if a down unit keeps blowing unchilled air into the cold aisle, which would make its loss worse. The halls' binding loads are taken at the pinned supply of 14°C whatever the chillers do; a chiller plant short of the load would warm the supply and lower the halls' binding loads together, a chain this chapter does not follow (Section 15.6.3).

**The graph.** The site is drawn in Figure 15.1. The two halls sit on the left, each with its three units; the shared pipe is one tall block between them and the chillers, drawn across both halls' paths, with the pump above it. The dashed arrows leave the graph for the records: each hall's night record, the outage records of the hall units and of the chillers. The pipe's fouling inputs, below it, are stated.

### 15.1.3 The planner's questions (step 1)

The planner asks four questions of the site, and step 1 writes each over the unknown $\boldsymbol{\theta}$ of the halls, the site and the outage rates, with the evidence on the right of the bar. Table 15.1 lists them. The first column is the question in the planner's words, the drill question first; the second the general form of its category in the notation of Chapter 2; the third its mathematical form for the site; the fourth its category; the fifth the computation of Section 2.3.6 that answers it; and

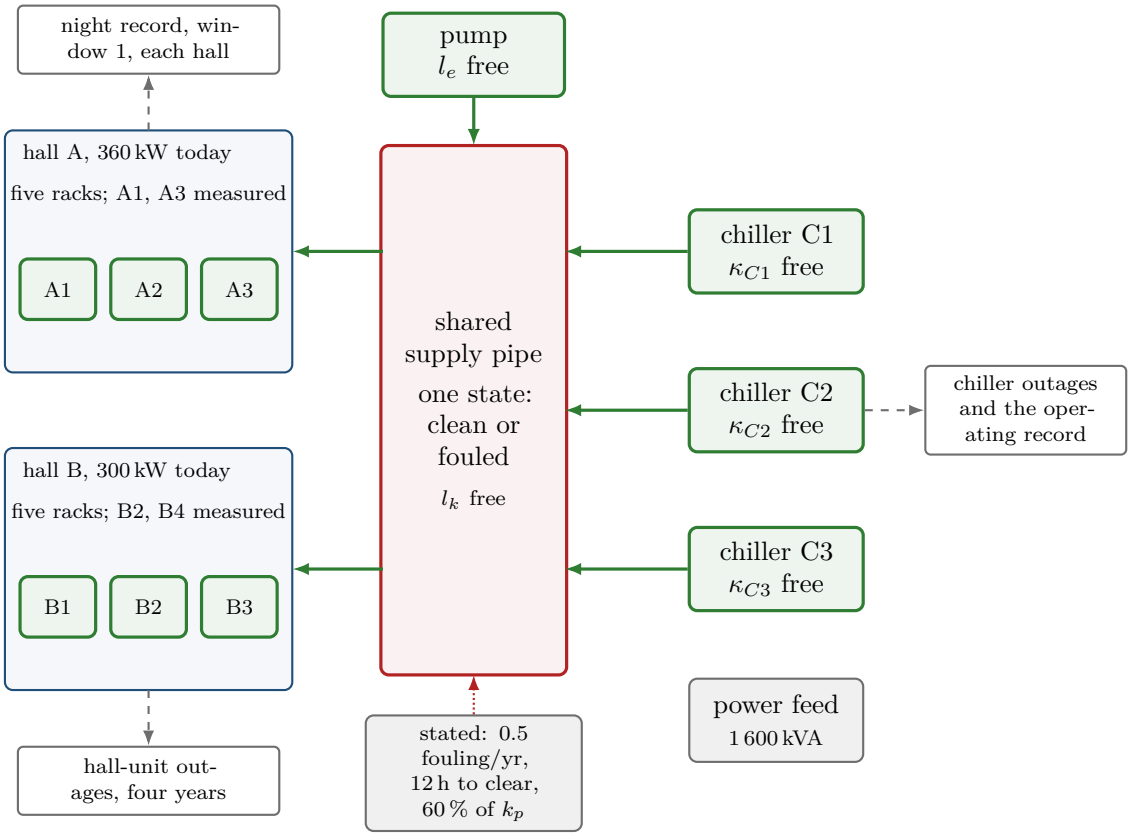



Figure 15.1: The site with both halls and the shared pipe as a state. Free (green): each hall's three units with their log-derates, and each hall's supply, load and five recirculation fractions (inside the hall blocks); the three chillers' as-built log-ratios, the pipe's and the pump's log-coefficients. The shared pipe (red border) is one block across both halls' supply paths: its single state, clean or fouled, enters both. Pinned: the feed's rating, and the pipe's stated fouling rate, clearing time and fouled fraction (dotted). Observed (dashed): each hall's readings in window 1 of the night record, the hall units' four-year outage record (both halls), and the site's operating record with the chillers' outage record.

Table 15.1: The questions asked of the site with the shared pipe as a state. $\mathbf{y}$ is both halls' night record, the site's operating record, the outage records and the priors; $P_L$, $P_{L,g}$ and $P_L^{(e)}$ are the readers of equation (15.2); $\ell$ is the limiting element; $\hat{\boldsymbol{\theta}}$ is the MAP of the plant.

| Planner's question | General form | Mathematical form | Category | Computation | Answered in |
|---|---|---|---|---|---|
| Two halls share one pipe; what is the probability that the shared pipe, not the halls' own units, is what limits them? | $\Pr(\ell = \text{pipe} \mid \text{shortfall}, \mathbf{y})$ | $\mathbb{E}[P_{L,\text{pipe}}]/\mathbb{E}[P_L]$, at $L = 660$, 759, 820, 860 kW | prediction under a threshold | posterior | 15.3.2 |
| What are the minimal cut sets, today and in a year? | the failing states of the enumeration | sets $S$ with $\text{Cap}(\mathbf{s}_S;\ \hat{\boldsymbol{\theta}}) < L$, orders one and two, at 660 and 759 kW | diagnosis | posterior | 15.3.3 |
| How likely is a shortfall at all? | $\Pr(\text{Cap} < L \mid \mathbf{y})$ | $\mathbb{E}[P_L(\boldsymbol{\theta}) \mid \mathbf{y}]$ by load, with its 90 % band | prediction under a threshold | posterior | 15.3.1 |
| How important is each element? | share of the shortfall from states with it down | $I_e = \mathbb{E}[P_L^{(e)}]/\mathbb{E}[P_L]$ per element, by load | estimation | posterior | 15.3.4 |

the last the subsection where it is answered. Each is asked at four site IT loads: today's 660 kW; year 1's 759 kW, the load path of Chapter 6 at 15 % a year; and 820 and 860 kW, near the load at which hall B is short with every unit up. The drill question is a conditional probability: of all the ways the site can fall short, the share in which the pipe is the smallest group. The answers come in the order three, one, two, four, since the drill question divides by the third.

### 15.1.4 The evidence as rows (step 2)

Step 2 attaches four records to the graph. They share no entry, so each acts on its own part of the unknown.

**The halls' night record.** Each hall is read in window 1 of the night record, at 19:00, by the rows of equation (5.11): the supply, the hall's return air, its three discharge temperatures, its load meter and the inlets of its two measured racks, with the priors of Chapter 5 on the supply, the three log-derates and the five fractions. Hall B's rows are those of Chapter 5 unchanged. Hall A's are the same rows with hall A's readings, ratings and other heat; its measured inlets are those of racks A1 and A3, and its unmeasured rack with the largest surveyed fraction is A5, surveyed at 0.45. Hall A's load meter reads 4 kW high throughout the record, a bias of the size of its noise that the rows do not model.

**The operating record and the outage records.** The site's state is read by the operating record of Chapter 9: two capacity tests per chiller and twelve monthly windows, sixty-two rows on the eight entries of equation (9.4). The chillers' outage record is that of Chapter 9. The hall units' outage record is new: the six units over the same four years, 24 unit-years of exposure, 21 failures and 460.8 h of repair in all. Its rows are of the same kind as the chillers': the failures a Poisson count in the exposure, with a gamma prior of shape 2 about a handbook rate of one failure a year, and the repair times exponential, with the weak prior of Chapter 9. The pipe has no record.

**Modeling choices at the rows.** Each hall is fitted on its own rows, and the supply reading, which both halls' rows contain, is used in both fits: this keeps each hall's calibration that of Chapter 5's form, and it fails to the extent that the two fits share that reading; the binding loads do not depend on the fitted supply, since the scenario pins it. Each reading is taken as Gaussian, unbiased and independent given the state: this fails for hall A's biased meter. The six hall units share one failure rate and one repair rate, as the three chillers do: this fails if one unit is worse than the others or if failures cluster.

## 15.2 The posterior of the state (step 3)

All four questions are read from one posterior, that of the plant's unknown under the four records of step 2. This section derives what the answers need from it and computes each part: hall A's fit, which is new; hall B's fit and the site's, which are cited; the outage rates; and the paired draws that carry them together.

### 15.2.1 What the answer needs from the state

The unknown of the chapter is

$$\boldsymbol{\theta} = (\mathbf{z}_A,\ \mathbf{z}_B,\ \mathbf{z}_{\text{site}},\ \lambda_c,\ \mu_c,\ \lambda_h,\ \mu_h), \tag{15.8}$$

each hall's reduced state of equation (5.4), the site's state of equation (9.4), and the failure and repair rates of the chillers (subscript $c$) and of the hall units (subscript $h$). By equations (15.4) to (15.6) the capacities of a state need the halls' log-derates and fractions and the site's as-built ratios and coefficients; by equation (15.1) its probability needs the unavailabilities, so the four rates and the stated $q_p$. The four records share no entry, so the posterior factors,

$$p(\boldsymbol{\theta} \mid \mathbf{y}) = p(\mathbf{z}_A \mid \mathbf{y}_A)\, p(\mathbf{z}_B \mid \mathbf{y}_B)\, p(\mathbf{z}_{\text{site}} \mid \mathbf{y}_{\text{op}})\, p(\lambda_c, \mu_c \mid \mathbf{y}_{\text{out}})\, p(\lambda_h, \mu_h \mid \mathbf{y}_{\text{out}}), \tag{15.9}$$

with $\mathbf{y}_A$ and $\mathbf{y}_B$ the halls' rows of window 1, $\mathbf{y}_{\text{op}}$ the operating record and $\mathbf{y}_{\text{out}}$ the outage records. Each hall's factor is the Laplace normal of its fit, equation (5.13); the site's is the one fit of Section 9.2.2; the rates' are the gamma posteriors of equation (9.12).

The answers are expectations over this posterior. The probability of a shortfall at load $L$ is, by the law of total probability of equation (3.2), the shortfall probability given the unknown averaged over its posterior,

$$\bar{P}_L = \Pr(\text{Cap} < L \mid \mathbf{y}) = \int P_L(\boldsymbol{\theta})\, p(\boldsymbol{\theta} \mid \mathbf{y})\, \mathrm{d}\boldsymbol{\theta}, \qquad \bar{P}_{L,g} = \int P_{L,g}(\boldsymbol{\theta})\, p(\boldsymbol{\theta} \mid \mathbf{y})\, \mathrm{d}\boldsymbol{\theta}, \tag{15.10}$$

and the same for $\bar{P}_L^{(e)}$. These three averages are all the answers need.

### 15.2.2 Hall A's fit

Hall A is fitted as hall B was in Section 5.2.2: its seventeen rows of window 1 form the residual vector of equation (5.11) with hall A's readings and ratings; the MAP state solves the normal equations (5.12) from the design sheet's state; and the Laplace covariance is the inverse of $\mathbf{J}^{\mathsf{T}}\mathbf{J}$ at the minimum, equation (5.13). The pattern of its rows is that of equation (5.10), with racks A1 and A3 in place of B2 and B4: no reading has a nonzero in the columns of the fractions of racks A2, A4 and A5, so their posterior is their survey prior, as for racks B1, B3 and B5.

Table 15.2 gives hall A's MAP state on window 1. Each row is an entry of the state; the columns give its MAP value, its posterior spread $\sqrt{\Sigma_{ii}}$ and the stated truth from which the record was generated, a check marker for Section 15.4.2, and the last column the rows that inform it.

Table 15.2: The MAP state of hall A in window 1 with the posterior spread of each entry and the stated truth.

| Entry [unit] | $\hat{z}_i$ [entry's] | $\sqrt{\Sigma_{ii}}$ [entry's] | Truth [entry's] | Informed by |
|---|---|---|---|---|
| $T_s$ [°C] | 14.1277 | 0.0978 | 14.0142 | supply sensor and set point |
| $Q_A$ [kW] | 341.5670 | 3.7278 | 337.1533 | meter and the temperature readings |
| $g_{A1}$ [–] | −0.0100 | 0.0242 | −0.0500 | discharge, return and inlet readings |
| $g_{A2}$ [–] | 0.0371 | 0.0233 | 0.0200 | discharge, return and inlet readings |
| $g_{A3}$ [–] | −0.0346 | 0.0247 | −0.0800 | discharge, return and inlet readings |
| $\beta_{A1}$ [–] | 0.1874 | 0.0267 | 0.1800 | rack-A1 inlet sensor and survey |
| $\beta_{A2}$ [–] | 0.2500 | 0.0800 | 0.2300 | survey only |
| $\beta_{A3}$ [–] | 0.3953 | 0.0268 | 0.4000 | rack-A3 inlet sensor and survey |
| $\beta_{A4}$ [–] | 0.3000 | 0.0800 | 0.2600 | survey only |
| $\beta_{A5}$ [–] | 0.4500 | 0.0800 | 0.5100 | survey only |

The readings see what they see in hall B. The supply is read to 0.0978 K and the three log-derates to about 0.024; the fractions of racks A1 and A3, whose inlets are measured, are read to 0.0267 and 0.0268, within a third of a spread of the truth; and the fractions of racks A2, A4 and A5 keep their survey's 0.08 and sit at its values. The load $Q_A$ lies 4.41 kW above the truth, the meter's bias, which the fit takes as load. The binding rack at the MAP state is A5, the unmeasured rack with the largest surveyed fraction, and the binding load there is 543.09 kW of hall load. Over the 2 000 draws of hall A's state of Section 15.2.4 the joint binding load, the smallest over the five racks in each draw, has median 537.76 kW and 90 % band 451.90 to 605.95 kW. Beside it, the exact posterior (the draws reweighted by hall A's own weights, effective sample size 1 990) gives median 537.56 kW and band 451.60 to 605.83 kW, and the delta form 543.09 kW with standard deviation 64.41 kW and band 437.13 to 649.05 kW, wider and symmetric where the binding load, a ratio, is skewed. The truth, 472.15 kW, lies inside the reported band, near its foot.

Table 15.3 sets the two halls' calibrations side by side. Each row is a quantity; the columns give it for hall A and hall B. The binding load at the MAP state is the hall's IT load at which its hottest inlet meets 27°C with all units running, equation (5.8); in site IT load it is divided by the hall's share, equation (15.5). The stated truth from which the record was generated is given beside it, a check marker for Section 15.4.2.

Hall A carries more than hall B: its binding load at the MAP state is 543.09 kW of hall load against hall B's 396.46 kW, and 995.7 against 872.2 kW of site load, so that with every unit up hall B is the hall that binds first. Hall B's fit reproduces Chapter 5's binding load to the last figure. Hall A's fit has the blind spot of hall B's, and a second one. Its unmeasured rack A5 is surveyed at 0.45 against a true fraction of 0.51, and no reading corrects a survey; its load meter reads 4 kW high. The MAP binding load lies $543.09 - 472.15 = 70.9$ kW of hall load above the truth, against $396.46 - 378.93 = 17.5$ kW for hall B, and hall A's band is wider than hall B's, 154.05 kW against 104.54 kW. Both truths lie inside their bands. The answers below carry both halls' posteriors through the draws.

Table 15.3: The two halls' calibrations on window 1 of the night record. Hall A's is new, its draws the 2 000 of this chapter; hall B's is that of Chapter 5, its draws the 50 000 of that chapter. The site IT load is the hall's load over its share.

| Quantity [unit] | Hall A [unit] | Hall B [unit] |
|---|---|---|
| Share of the site IT load $\varphi_h$ [–] | 0.5455 | 0.4545 |
| IT load today [kW] | 360 | 300 |
| Measured rack inlets [–] | A1, A3 | B2, B4 |
| Unmeasured rack with the largest survey fraction [–] | A5, 0.45 | B3, 0.50 |
| Binding rack at the MAP state [–] | A5 | B3 |
| *Binding load, all units, hall load [kW]* | | |
| MAP state | 543.09 | 396.46 |
| Median of the draws | 537.76 | 395.21 |
| 90 % band of the draws | 451.90 to 605.95 | 345.89 to 450.43 |
| Truth | 472.15 | 378.93 |
| *Binding load, all units, site IT load [kW]* | | |
| MAP state | 995.7 | 872.2 |
| Truth | 865.6 | 833.6 |

### 15.2.3 Hall B's fit and the site's, cited

Hall B's posterior is that of Section 5.2.2: its binding load at the MAP state is 396.46 kW at rack B3, whose fraction no reading sees, and the joint band of Section 5.3.1 runs from 345.89 to 450.43 kW of hall load (Table 5.7). The site's posterior is the one fit of Section 9.2.2, on the sixty-two rows of its operating record; its entries are listed in Table 9.2, and they are not repeated here.

### 15.2.4 The outage posteriors and the draws

The chillers' rates are the gamma posteriors of equation (9.12). The hall units' follow in the same way from their record: the failures add their count to the prior's shape and the exposure to its rate, and the repairs add their count and their hours,

$$\lambda_h \mid \mathbf{y}_{\text{out}} \sim \text{Gamma}(2{+}21,\ 2\,\text{yr}{+}24\,\text{yr}), \qquad \mu_h \mid \mathbf{y}_{\text{out}} \sim \text{Gamma}(1{+}21,\ 36\,\text{h}{+}460.8\,\text{h}), \tag{15.11}$$

exact, with $q_h = \lambda_h/(\lambda_h + \mu_h)$ computed in each draw. The prior's rate of 2 yr is its shape over its mean of one failure a year.

The answers are computed on $N = 2\,000$ paired draws of the unknown (15.8), by the factorization (15.9): in each, a draw of hall A's state and one of hall B's from their Laplace normals, made with the seed 15, paired with one of the first 2 000 draws of Section 9.2.4, which carry the site's state and both kinds' rates. Over these draws the mean unavailability of a chiller is 0.00551 and of a hall unit 0.00241; the pipe's is the stated $6.845 \times 10^{-4}$ in every draw. In each draw every state with at most two of the ten elements down is enumerated, its capacities computed by equations (15.4) to (15.6) and its limiting element read by equation (15.1), and the three readers (15.2) are summed. The states with three or more elements down are left out: at the mean unavailabilities their probability is $3.44 \times 10^{-6}$, under a three-thousandth of today's shortfall probability (Section 15.6.3). The reported answers are the plain averages over the draws, the state route, by equation (15.10).

The exact posterior reweights the same draws. Each draw carries the product of three weights of the form of equation (5.16): hall A's and hall B's, each the ratio of the hall's exact posterior to its Laplace normal at the draw, and the site's of Section 9.2.4; the rates' factors are

exact already. Normalized to $w_i$ with $\sum_i w_i = 1$, they give the exact route $\sum_i w_i\, P_L(\boldsymbol{\theta}_i)$ beside the plain average, with an effective sample size $1/\sum_i w_i^2$ of 1 967 of 2 000: the three Laplace normals together lose under two percent of the draws.

## 15.3 The answers (step 4)

Every answer is read from the paired draws of Section 15.2.4. The reported value is the average over the draws, joint over the six groups of every state; the band is the 5 % and 95 % points of the same quantity given the unknown, over the draws. The exact posterior, the same draws reweighted, stands beside it with the gap and the Monte Carlo standard error. The capacities at the MAP state stand beside both as the single-value route, and they are what the cut sets read.

### 15.3.1 The probability of a shortfall

**Planner's question.** How likely is a shortfall at all? The question asks for the probability that the site, in whatever state its elements are, carries less than the load.

**General form.** A prediction under a threshold is the posterior probability of an event, equation (2.12),

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

asked here, as in Chapter 9, of an event that is random even when the unknown is known.

**Specialized form.** The question is the complement of equation (2.12) with $f$ the capacity $\mathrm{Cap}(\mathbf{s};\ \boldsymbol{\theta})$ of equation (15.1), $f^* = L$, and the integral taken over the elements' states and the unknown together, which is $\bar{P}_L$ of equation (15.10). It is reported by the draws,

$$\bar{P}_L \approx \frac{1}{N}\sum_{i=1}^{N} P_L(\boldsymbol{\theta}_i), \qquad \bar{P}_{L,g} \approx \frac{1}{N}\sum_{i=1}^{N} P_{L,g}(\boldsymbol{\theta}_i), \tag{15.12}$$

with $N = 2\,000$; the band is that of $P_L(\boldsymbol{\theta}_i)$ over the draws, how far the shortfall probability itself is known. The exact route replaces each $1/N$ by the draw's weight $w_i$ of Section 15.2.4, and the Monte Carlo standard error of the plain average is the standard deviation of $P_L(\boldsymbol{\theta}_i)$ over the draws divided by $\sqrt{N}$. The capacity of each state at the MAP state is the single-value route.

**Building blocks.** The paired draws (Section 15.2.4); the capacities (15.4) to (15.6); the enumeration (15.1) and its readers (15.2).

**The capacities.** Table 15.4 gives the capacities at the MAP state. The upper block gives each group's capacity with every element up; the lower gives, for each element down alone, the smallest group's capacity and which group it is. All are in site IT load.

With every element up, hall B is the smallest group, at 872.2 kW; the pipe, the element the reviewer pointed at, is fourth, at 1 126.8 kW, behind hall A and the chillers. With one element down the order turns over. A hall-B unit out leaves the site 565.2 to 586.1 kW; a hall-A unit out, 632.6 to 676.1 kW; the pipe fouled, 667.3 kW; a chiller out, 699.5 to 727.8 kW. Hall B's air side is the weakest place in the site: with any one of its units out it is short of today's 660 kW. Hall B is N on its air side, not N+1. Hall A is N+1 for two of its units and not for A2, whose loss leaves 632.6 kW; the loss of A1 leaves 661.3 kW, 1.3 kW above today's load. The fouled pipe, at 667.3 kW, sits in the middle of the one-down states, above today's load and below year 1's.

Table 15.4: The capacities of the site's groups at the MAP state, at the design day, in site IT load: all elements up, and with each element down alone. The halls' capacities are their binding loads over their shares; the chillers carry the loop without the economizer.

| State | Capacity [kW] | Limiting group [–] |
|---|---|---|
| *All up, each group* | | |
| Hall B | 872.2 | – |
| Hall A | 995.7 | – |
| Chillers | 1 084.1 | – |
| Pipe, clean | 1 126.8 | – |
| Feed | 1 188.2 | – |
| Pump | 1 272.0 | – |
| *One element down, the smallest group* | | |
| B2 down | 565.2 | hall B |
| B1 down | 571.3 | hall B |
| B3 down | 586.1 | hall B |
| A2 down | 632.6 | hall A |
| A1 down | 661.3 | hall A |
| Pipe fouled | 667.3 | pipe |
| A3 down | 676.1 | hall A |
| C2 down | 699.5 | chillers |
| C1 down | 718.8 | chillers |
| C3 down | 727.8 | chillers |

**The shortfall probability.** Table 15.5 gives the answer at the four loads. Each row is a load; the columns give the shortfall probability by the draws with its 90 % band, and its split by the limiting group, $\bar{P}_{L,g}$ of equation (15.12).

Today the site falls short with probability $1.102 \times 10^{-2}$, band $4.611 \times 10^{-3}$ to $1.994 \times 10^{-2}$: the hall-B units down, each an order-one shortfall, and A2 down, make most of it. At year 1 it is $7.565 \times 10^{-2}$, and the mean lies above its own 95 % point of $5.821 \times 10^{-2}$. The band and the mean are read from different draws. In most draws the shortfall at 759 kW comes from the one-down states, which by Table 15.4 all fall short of it, and their probability is a few hundredths. In a few draws, those in which hall B's binding load with every unit up lies below the load, the all-up state is short and the shortfall probability is near one. Chapter 5's 5 % point of hall B's binding load, 345.89 kW of hall load, is $345.89/0.4545 = 761.0$ kW of site load, just above year 1's 759 kW: the draws below it are under one in twenty, and each counts for nearly one. At 820 and 860 kW those draws are the bulk: 860 kW of site load is $0.4545 \times 860 = 390.9$ kW in hall B, just under the median of 395.21 kW of Table 5.7, and the shortfall probability is $4.990 \times 10^{-1}$, with a band that reaches one. At these loads the question is no longer about failures but about hall B's capacity with every unit up.

**The exact route beside the draws.** Table 15.6 sets the plain average of the draws beside the exact posterior at each load. Each row is a load; the columns give the shortfall probability by both routes, their gap, the Monte Carlo standard error of the plain average, and the pipe's share given a shortfall of Section 15.3.2 by both routes.

The exact posterior raises the shortfall probability at every load, and by under half a Monte Carlo standard error each time: by $5 \times 10^{-5}$ today against $1.1 \times 10^{-4}$, and by $4.5 \times 10^{-3}$ at 860 kW against $1.1 \times 10^{-2}$. The pipe's share is the same to four places at the first three loads and 0.0013 against 0.0014 at 860 kW. The Laplace normals of the three fits change no answer of this chapter.

Table 15.5: The probability that the site falls short at four site IT loads, by 2 000 paired draws of the posterior, with its band over the draws and its split by the limiting group. The four parts sum to the shortfall probability.

| Load $L$ [kW] | $\bar{P}_L$ [–] | 90 % band [–] | Hall B limits [–] | Hall A limits [–] | Chillers limit [–] | Pipe limits [–] |
|---|---|---|---|---|---|---|
| 660 | $1.102 \times 10^{-2}$ | $4.611 \times 10^{-3}$ to $1.994 \times 10^{-2}$ | $6.877 \times 10^{-3}$ | $4.038 \times 10^{-3}$ | $1.029 \times 10^{-4}$ | $4.736 \times 10^{-6}$ |
| 759 | $7.565 \times 10^{-2}$ | $2.055 \times 10^{-2}$ to $5.821 \times 10^{-2}$ | $4.858 \times 10^{-2}$ | $1.032 \times 10^{-2}$ | $1.608 \times 10^{-2}$ | $6.766 \times 10^{-4}$ |
| 820 | $2.787 \times 10^{-1}$ | $2.114 \times 10^{-2}$ to $1.000$ | $2.224 \times 10^{-1}$ | $3.956 \times 10^{-2}$ | $1.609 \times 10^{-2}$ | $6.766 \times 10^{-4}$ |
| 860 | $4.990 \times 10^{-1}$ | $2.204 \times 10^{-2}$ to $1.000$ | $4.093 \times 10^{-1}$ | $7.292 \times 10^{-2}$ | $1.609 \times 10^{-2}$ | $6.766 \times 10^{-4}$ |

Table 15.6: The shortfall probability and the pipe's share given a shortfall, by the plain average of the 2 000 paired draws and by the exact posterior (the same draws reweighted, effective sample size 1 967), with the gap, exact less plain, and the Monte Carlo standard error of the plain average.

| Load $L$ [kW] | $\bar{P}_L$, draws [–] | $\bar{P}_L$, exact [–] | Gap [–] | MC se [–] | Pipe share, draws [–] | Pipe share, exact [–] |
|---|---|---|---|---|---|---|
| 660 | $1.102 \times 10^{-2}$ | $1.107 \times 10^{-2}$ | $5 \times 10^{-5}$ | $1.1 \times 10^{-4}$ | 0.0004 | 0.0004 |
| 759 | $7.565 \times 10^{-2}$ | $7.604 \times 10^{-2}$ | $3.9 \times 10^{-4}$ | $4.5 \times 10^{-3}$ | 0.0089 | 0.0089 |
| 820 | $2.787 \times 10^{-1}$ | $2.808 \times 10^{-1}$ | $2.1 \times 10^{-3}$ | $9.4 \times 10^{-3}$ | 0.0024 | 0.0024 |
| 860 | $4.990 \times 10^{-1}$ | $5.035 \times 10^{-1}$ | $4.5 \times 10^{-3}$ | $1.1 \times 10^{-2}$ | 0.0014 | 0.0013 |

Table 15.7: The answer to the drill question: the probability that the pipe limits the site given that the site falls short, the ratio of means of equation (15.13), with the band of the per-draw share $r_{\text{pipe}}(\boldsymbol{\theta}_i)$ beside it; and the shares of the halls' own units and of the chillers.

| Load $L$ [kW] | Pipe limits, given a shortfall [–] | Per-draw share, 90 % band [–] | Halls' own units [–] | Chillers [–] |
|---|---|---|---|---|
| 660 | 0.0004 | 0.0000 to 0.0000 | 0.9902 | 0.0093 |
| 759 | 0.0089 | 0.0115 to 0.0331 | 0.7786 | 0.2125 |
| 820 | 0.0024 | 0.0007 to 0.0321 | 0.9399 | 0.0577 |
| 860 | 0.0014 | 0.0007 to 0.0307 | 0.9664 | 0.0322 |

### 15.3.2 The share the shared pipe limits

**Planner's question.** Two halls share one pipe; what is the probability that the shared pipe, not the halls' own units, is what limits them? The question asks, of the states in which the site falls short, for the probability that the pipe is the smallest group.

**General form.** A prediction under a threshold is the posterior probability of an event, equation (2.12), here asked of the joint event that the site is short and the pipe limits it, and conditioned on the event that the site is short. A conditional probability is the probability of the joint event over that of the condition, equation (3.4),

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

**Specialized form.** The question is equation (3.4) with $A$ the event $\ell =$ pipe and $B$ the shortfall, each probability that of equation (2.12) averaged over the unknown by equation (15.10),

$$\Pr(\ell = g \mid \text{Cap} < L,\ \mathbf{y}) = \frac{\bar{P}_{L,g}}{\bar{P}_L} = \frac{\int P_{L,g}(\boldsymbol{\theta})\, p(\boldsymbol{\theta} \mid \mathbf{y})\, \mathrm{d}\boldsymbol{\theta}}{\int P_L(\boldsymbol{\theta})\, p(\boldsymbol{\theta} \mid \mathbf{y})\, \mathrm{d}\boldsymbol{\theta}}, \qquad g = \text{pipe}, \tag{15.13}$$

the ratio of the two averages of equation (15.12): a ratio of means. The ratio given the unknown, $r_g(\boldsymbol{\theta}) = P_{L,g}(\boldsymbol{\theta})/P_L(\boldsymbol{\theta})$, is the share the pipe would have if the plant were known. Its 5 % and 95 % points over the draws are the band, how far that share is known; its mean is not the answer, since by equation (15.13) each draw counts in proportion to its shortfall probability.

**Building blocks.** The averages $\bar{P}_L$ and $\bar{P}_{L,g}$ of Section 15.3.1; the per-draw shares $r_g(\boldsymbol{\theta}_i)$.

**The shares.** Figure 15.2 stacks the shares of the shortfall probability by limiting group at the four loads, each bar summing to one, with the shortfall probability and the pipe's share printed above it. Table 15.7 gives the pipe's share with its band over the draws, the halls' together and the chillers'.

The shared pipe is almost never what limits the two halls; their own units are. Given a shortfall, the probability that the pipe is the limiting group is 0.0004 today, 0.0089 at year 1, and 0.0024 and 0.0014 at 820 and 860 kW; the exact posterior gives the same four shares but the last, 0.0013 (Table 15.6). The halls' own units limit 0.9902 of today's shortfalls and 0.7786 of year 1's, the chillers most of the rest, 0.0093 and 0.2125. Hall B alone limits 0.6239 today and 0.6421 at year 1, by Table 15.5.

The pipe's share moves with the load for a reason the capacities give. Today, the fouled pipe leaves 667.3 kW at the MAP state, above the load, so the fouled state is short only in the draws whose pipe coefficient is low enough to take it below 660 kW; in more than nineteen draws in

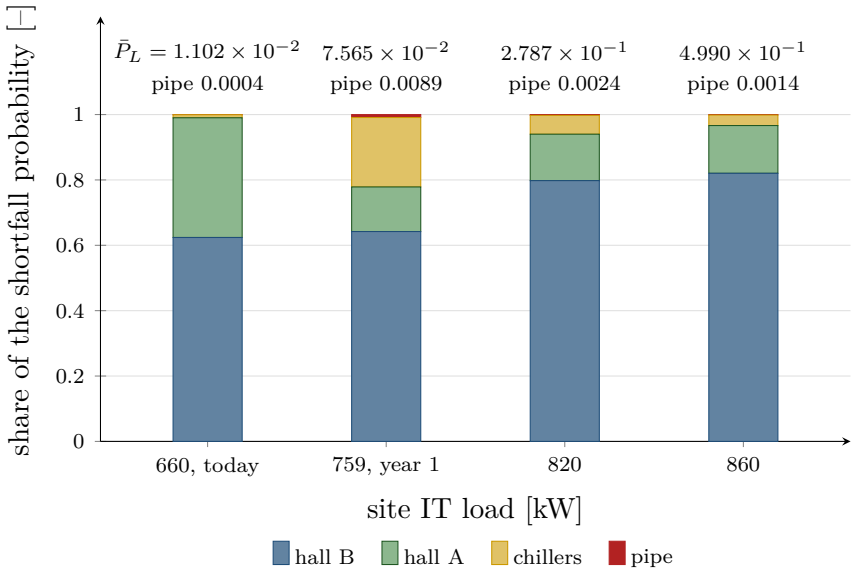



Figure 15.2: The share of the shortfall probability limited by each group, $\bar{P}_{L,g}/\bar{P}_L$ of equation (15.13), at today's load, year 1's and two higher loads, by 2 000 paired draws. Above each bar: the shortfall probability $\bar{P}_L$ and the pipe's share, whose segment (red) is too thin to see.

twenty it never is, and the per-draw band is 0.0000 at both ends. At year 1 the fouled state is short in every draw and the pipe limits it, since the halls and the chillers are all above 667.3 kW with every unit up; the pipe then limits with the probability of being fouled, $6.766 \times 10^{-4}$ in Table 15.5, a little under $q_p$ because the other elements must be up. That number stays the same at 820 and 860 kW, while the shortfall probability grows with hall B's all-up shortfalls, and the pipe's share falls.

At year 1 the ratio of means, 0.0089, lies below its own band over the draws, 0.0115 to 0.0331. Both are right; they answer different questions. In a typical draw the site is short only in its one-down states, a few hundredths of the time, and the pipe takes one to three hundredths of that. In the few draws in which hall B is short with every unit up, the shortfall probability is near one and the pipe's share in it near nothing; these draws are under the 5 % point of the band, so the band does not see them, but they carry most of $\bar{P}_L$, and the ratio of means weights them by it. The answer to the drill question is the ratio of means, because by equation (15.13) it is the conditional probability over everything not known; the band says that if hall B's capacity were pinned down, the pipe's share at year 1 would likely lie between one and three hundredths.

### 15.3.3 The minimal cut sets, today and in a year

**Planner's question.** What are the minimal cut sets, today and in a year? The question asks which elements, alone or in pairs, leave the site short when they are down.

**General form.** A diagnosis runs over a family of candidate faults generated from the plant graph, one for each component that could be at fault, and returns the candidates ranked by their probability, equation (2.15),

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

Here the candidates are asked not what the record shows but what each would do: whether it would leave the site short.

Table 15.8: The minimal cut sets of the site at the MAP state of the plant, orders one and two, at today's load and year 1's. Each is ranked by its leading-order probability, the product of its elements' mean unavailabilities.

| Load $L$ [kW] | Order one [–] | Minimal order two [–] |
|---|---|---|
| 660 | A2, B1, B2, B3 | {C1, C2}, {C1, C3}, {C2, C3}, {A1, A3} |
| 759 | C1, C2, C3, A1, A2, A3, B1, B2, B3, pipe | none: every element is a cut set alone |

**Specialized form.** The question is equation (2.15) with the family $M_S$ the sets $S$ of elements down, the rest up, generated from the ten elements of the graph. A candidate is kept when it is a minimal cut set of equation (15.3) at the MAP state $\hat{\boldsymbol{\theta}}$ of the plant, and it is ranked by its probability $\Pr(M_S) = P(\mathbf{s}_S \mid \bar{\mathbf{q}})$ of equation (15.1) at the mean unavailabilities, to leading order the product $\prod_{e \in S} \bar{q}_e$. Orders one and two are read at today's load and year 1's. The cut sets are a single-value reading: they test each state's capacity at one value of the plant, and the importances of Section 15.3.4 carry the rest of the posterior.

**Building blocks.** The capacities at the MAP state (Table 15.4) and of every two-down state; the mean unavailabilities (Section 15.2.4).

**The cut sets.** Table 15.8 lists the minimal cut sets at the two loads, and Figure 15.3 draws them as a fault tree: the site short at the top, an OR over the minimal cut sets below it, each order-two set an AND of its two elements.

Today the site has four order-one cut sets: the three hall-B units and A2, the four elements whose loss alone leaves less than 660 kW by Table 15.4. Each has the leading-order probability of a hall unit, $\bar{q}_h = 0.00241$. The minimal order-two sets are the three pairs of chillers and the pair A1 and A3 of hall A, each of the order of $\bar{q}^2$; a pair of chillers, for instance, $0.00551^2 = 3.0 \times 10^{-5}$. The pipe is in no cut set: fouled, it leaves 667.3 kW, and no pair that includes it falls below 660 kW at the MAP state except through an element that is a cut set on its own.

At year 1 every element is an order-one cut set, the pipe among them. No one-down state carries 759 kW, by Table 15.4, so every element alone takes the site down, and no pair is minimal. The pipe's place in the tree is the common mode the reviewer meant: its single state sits under the site's OR beside each hall's units, and when it fouls it takes the supply of both halls at once, while a hall unit takes one hall. It matters little only because it fouls rarely: its probability is $q_p = 6.845 \times 10^{-4}$, below a hall unit's 0.00241 and a chiller's 0.00551. The cut sets show where the pipe could take the site down alone; Section 15.3.2 showed how often it does.

### 15.3.4 The importance of each element

**Planner's question.** How important is each element? The question asks how much of the shortfall probability each element's failure is part of.

**General form.** An estimation question asks for the moments of a derived quantity under the posterior, equation (2.10),

$$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}] = \int f(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z},$$

here the expectations of two readers of the enumeration, one over the other.

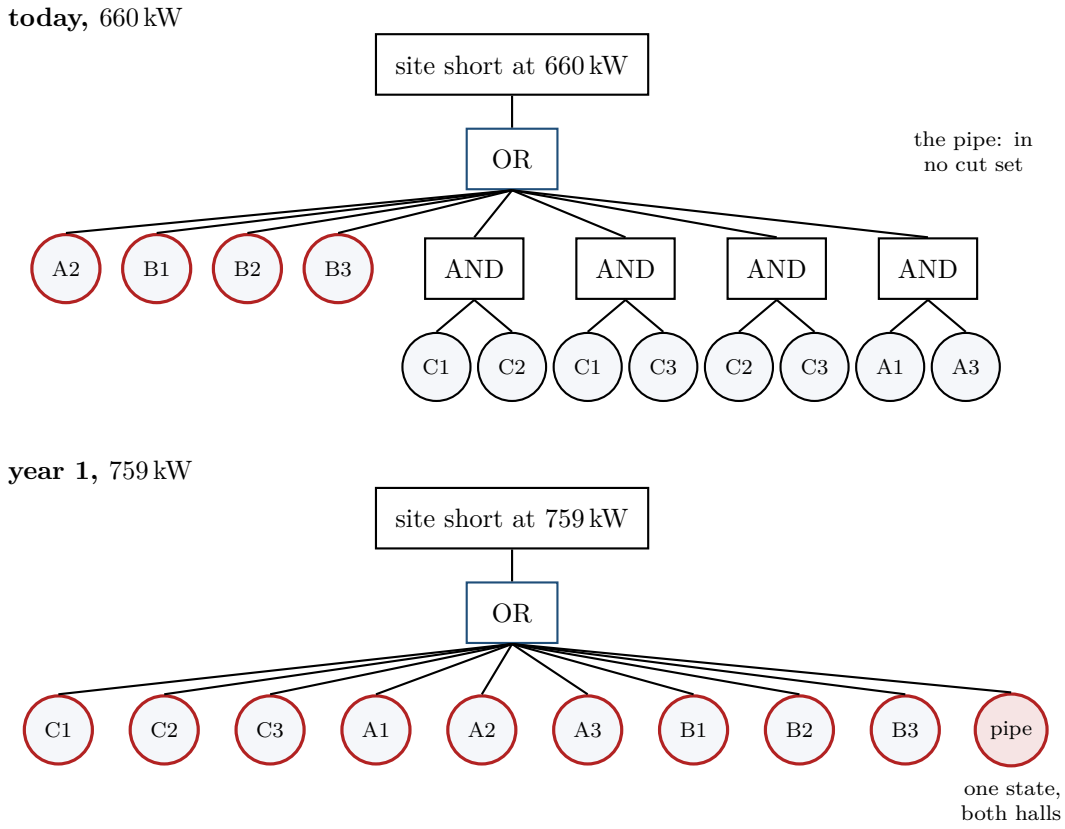



Figure 15.3: The minimal cut sets of Table 15.8 as a fault tree. Top: today, four order-one cut sets (red circles) and four order-two sets, each an AND of two elements; the pipe is in none. Bottom: year 1, every one of the ten elements is an order-one cut set, the shared pipe among them, whose single state cuts both halls' supply.

Table 15.9: The Fussell–Vesely importance of each element, equation (15.14): the share of the shortfall probability from states with the element down, by 2 000 paired draws.

| Element | 660 kW [–] | 759 kW [–] | 820 kW [–] | 860 kW [–] |
|---|---|---|---|---|
| C1 | 0.0116 | 0.0728 | 0.0198 | 0.0110 |
| C2 | 0.0116 | 0.0729 | 0.0198 | 0.0110 |
| C3 | 0.0116 | 0.0727 | 0.0198 | 0.0110 |
| A1 | 0.1159 | 0.0312 | 0.0086 | 0.0048 |
| A2 | 0.1572 | 0.0318 | 0.0086 | 0.0048 |
| A3 | 0.0990 | 0.0303 | 0.0086 | 0.0048 |
| B1 | 0.2108 | 0.0318 | 0.0086 | 0.0048 |
| B2 | 0.2132 | 0.0318 | 0.0086 | 0.0048 |
| B3 | 0.2025 | 0.0318 | 0.0086 | 0.0048 |
| Pipe | 0.0011 | 0.0090 | 0.0025 | 0.0014 |

**Specialized form.** The question is equation (2.10) with $f$ the readers $P_L^{(e)}$ and $P_L$ of equation (15.2), and the importance their ratio,

$$I_e = \frac{\mathbb{E}[P_L^{(e)}(\boldsymbol{\theta}) \mid \mathbf{y}]}{\mathbb{E}[P_L(\boldsymbol{\theta}) \mid \mathbf{y}]} = \Pr(s_e = 0 \mid \mathrm{Cap} < L,\ \mathbf{y}), \tag{15.14}$$

the share of the shortfall probability from states with element $e$ down, which by equation (3.4) is the probability that $e$ is down given that the site is short: the Fussell–Vesely importance of reliability engineering, read on the enumeration. It differs from the limiting share of equation (15.13): an element counts in every short state in which it is down, whatever limits that state, and the importances of the ten elements need not sum to one. The capacity-sense importance, the gain in headroom per unit of an element's rating, is the question of Chapter 14.

**Building blocks.** The averages of $P_L^{(e)}$ and $P_L$ over the paired draws (Section 15.2.4).

**The importances.** Table 15.9 gives the ten importances at the four loads. Each row is an element; the columns are the loads.

Today the hall-B units lead, about 0.21 each, then A2 at 0.1572; the pipe is last, at 0.0011. Two entries say more than the cut sets did. A1 and A3, which the MAP state leaves outside the order-one list by $661.3 - 660 = 1.3$ kW and $676.1 - 660 = 16.1$ kW, carry importances of 0.1159 and 0.0990, near A2's: across the posterior of hall A their loss alone leaves the site short in a large part of the draws, which a cut-set list at one value of the plant cannot show. And the pipe's importance today, 0.0011, is above its limiting share of 0.0004: in the short states with the pipe fouled, a hall unit is usually down too, and it is the hall that limits.

At year 1 every element is an order-one cut set, and the importances line up with the unavailabilities: the chillers about 0.073 each, the hall units about 0.03, the pipe 0.0090. When every element alone takes the site down, the share from states with it down is its unavailability times a factor common to all, so the ranking is the ranking of the $q$'s. The ten importances at year 1 sum to less than one half: more than half of year 1's shortfall probability comes from the state with every element up, the draws in which hall B is short on its own capacity. At 820 and 860 kW that share grows, and every element's importance falls with it.

## 15.4 The checks (step 5)

The answers rest on Laplace draws of three fits, on stated inputs for the pipe and on a record generated from a stated truth. Each is tested here as far as the computation allows, and the

Table 15.10: The checks of the shared-path answers. Reported values by 2 000 paired draws; the truth is the stated truth of the records with the outage rates at their true values; the fouling variation repeats the year-1 answer on every fourth draw with $q_p$ four times as large.

| Quantity | Reported [kW or –] | Truth [kW or –] | Reading |
|---|---|---|---|
| *Capacities, all up, site IT load [kW]* | | | |
| Hall B | 872.2 | 833.6 | MAP above the truth |
| Hall A | 995.7 | 865.6 | MAP above the truth |
| Chillers | 1 084.1 | 1 066.0 | close |
| Pipe | 1 126.8 | 1 127.3 | close |
| Feed | 1 188.2 | 1 184.7 | close |
| Pump | 1 272.0 | 1 250.4 | close |
| *Shortfall probability $\bar{P}_L$ [–]* | | | |
| 660 kW | $1.102 \times 10^{-2}$ | $1.312 \times 10^{-2}$ | inside the band |
| 759 kW | $7.565 \times 10^{-2}$ | $2.817 \times 10^{-2}$ | inside the band |
| 820 kW | $2.787 \times 10^{-1}$ | $2.817 \times 10^{-2}$ | inside the band |
| 860 kW | $4.990 \times 10^{-1}$ | 1.000 | at the band's top |
| *Pipe limits, given a shortfall [–]* | | | |
| 660 kW | 0.0004 | 0.0000 | both near zero |
| 759 kW | 0.0089 | 0.0240 | inside the per-draw band |
| 820 kW | 0.0024 | 0.0240 | inside the per-draw band |
| 860 kW | 0.0014 | 0.0007 | hall B short all up |
| *Hall A's calibration, hall load [kW]* | | | |
| Binding load, MAP state | 543.09 | 472.15 | MAP above the truth |
| Binding load, 90 % band of the draws | 451.90 to 605.95 | 472.15 | inside the band |
| *Variations* | | | |
| Pipe fouling ×4, $q_p = 2.732 \times 10^{-3}$, pipe limits at 759 kW | 0.0332 | – | against 0.0089 |
| Hall B binding load, Chapter 5 [kW] | 396.46 | 378.93 | the same fit |

answer is tested against the one input that is not estimated at all, the pipe's fouling rate.

### 15.4.1 The draws and the calibrations

Hall B's binding load at the MAP state is 396.46 kW, the value of Chapter 5: the fit written for both halls reproduces the template's. The exact posterior of the three fits together, the paired draws reweighted (Section 15.2.4), keeps an effective sample size of 1 967 of 2 000, beside the 49 638 of 50 000 found for hall B alone (Section 5.2.3) and the 19 913 of 20 000 for the site alone (Section 9.2.4); hall A's fit, new here, is checked only inside that joint weight. Every shortfall probability moves under the exact posterior by less than half its Monte Carlo standard error, and the pipe's share by at most 0.0001 (Table 15.6). Hall A's binding load, the one new calibration, has its truth of 472.15 kW inside its band of 451.90 to 605.95 kW (Table 15.3). The pipe's fouled capacity follows from its clean one by equation (15.6): $(1\,126.8 + 22) \times 0.6 - 22 = 667.3$ kW, the value of Table 15.4.

### 15.4.2 The answers against the stated truth

The halls' night record, the operating record and the outage records are generated from a stated truth. Table 15.10 collects the checks. Each row is a quantity; the columns give the reported value, the truth and the reading.

The truth's site is weaker than the MAP's in the halls and the same in the pipe. With every

unit up the true hall B carries 833.6 kW of site load and the true hall A 865.6 kW, against 872.2 and 995.7 kW at the MAP state; the chillers, the pipe, the feed and the pump are within a few kilowatts to a few tens. Hall B binds first in the truth too.

At today's load the true shortfall probability is $1.312 \times 10^{-2}$, inside the band of $4.611 \times 10^{-3}$ to $1.994 \times 10^{-2}$ and above the reported $1.102 \times 10^{-2}$, since the true hall A is weaker and its one-unit-out states fall shorter. The pipe limits none of the true shortfalls: 0.0000, against the reported 0.0004. At year 1 and at 820 kW the true shortfall probability is $2.817 \times 10^{-2}$ at both loads, since the true hall B, at 833.6 kW, carries both with every unit up; the shortfalls are the one-down states alone, and the pipe limits 0.0240 of them. That value lies inside the per-draw band, 0.0115 to 0.0331 at year 1, and above the ratio of means: it is the share in a plant whose hall B is known to carry the load, which is what the band of Section 15.3.2 describes. At 860 kW the true hall B is short with every unit up, the shortfall probability is one, and the pipe limits only when it is fouled, 0.0007.

The truth confirms the answer to the drill question at every load: the pipe limits under three hundredths of the shortfalls, and the halls' units, hall B's above all, limit almost all the rest. Hall A's new calibration places its MAP binding load 70.9 kW above the truth, at the rack whose survey lies below its true fraction, and its band still holds the truth.

### 15.4.3 A pipe that fouls four times as often

The pipe's fouling rate is stated, not estimated, and the answer depends on it directly: at year 1 the pipe limits only in its own fouled state. Repeating the year-1 answer with the rate four times as large, $q_p = 2.732 \times 10^{-3}$, on every fourth draw, the share rises from 0.0089 to 0.0332, about four times, the small remainder because the fouled states now add to the denominator too. The answer scales with the shared element's own rate. Even at four times the stated rate, the pipe limits about one shortfall in thirty at year 1, and the halls' units still limit most of the rest.

## 15.5 The report (step 6)

The shared pipe is not what limits the two halls. Given that the site falls short, the probability that the pipe is the limiting element is 0.0004 today and 0.0089 at year 1, by 2 000 draws of the posterior; the halls' own units limit 0.9902 and 0.7786 of the shortfalls, hall B's above all, and the chillers the rest. Hall B is N, not N+1, on its air side: any one of its units out leaves it short of today's load.

### 15.5.1 The answers and their bands

Table 15.11 gives the answers in the order of Table 15.1. Each row is a question; the second column gives the reported answer, and the third its band or the routes beside it.

### 15.5.2 The design-review entry

The design review's entry answers the reviewer. The shared supply pipe limits the two halls in under one percent of the site's shortfalls today and at year 1, 0.0004 and 0.0089; at four times its stated fouling rate it would limit 0.0332 at year 1. The halls' own cooling units limit almost all the rest, and hall B's first: it is N on its air side, each of its units an order-one cut set at today's load. The entry records two actions, in order. First, hall B's air side: a fourth unit in hall B, or another remedy for its lack of redundancy, comes before any work on the pipe; how much it buys with a unit out is the question of Chapter 20. Second, the strainer: keep it maintained and log its fouling, because at year 1 the pipe is an order-one cut set for both halls at once, and its share scales with its own rate, which no record yet supports.

Table 15.11: The report for the site with the shared pipe as a state, at the design day, by 2 000 paired draws of the posterior; the cut sets at the MAP state of the plant.

| Question | Answer | Band or routes beside |
|---|---|---|
| Two halls share one pipe; what is the probability that the shared pipe, not the halls' own units, is what limits them? | 0.0004 today, 0.0089 at year 1, 0.0024 at 820 kW, 0.0014 at 860 kW; halls' units 0.9902, 0.7786, 0.9399, 0.9664 | per-draw 90 % bands 0.0000 to 0.0000, 0.0115 to 0.0331, 0.0007 to 0.0321, 0.0007 to 0.0307; exact 0.0004, 0.0089, 0.0024, 0.0013; truth 0.0000, 0.0240, 0.0240, 0.0007 |
| What are the minimal cut sets, today and in a year? | today: A2, B1, B2, B3 of order one; the chiller pairs and {A1, A3} of order two; year 1: all ten elements of order one, the pipe among them | at the MAP state; A1 and A3 near order one across the posterior (importances 0.1159, 0.0990) |
| How likely is a shortfall at all? | $1.102 \times 10^{-2}$ today, $7.565 \times 10^{-2}$ at year 1, $2.787 \times 10^{-1}$ at 820 kW, $4.990 \times 10^{-1}$ at 860 kW | $4.611 \times 10^{-3}$ to $1.994 \times 10^{-2}$; $2.055 \times 10^{-2}$ to $5.821 \times 10^{-2}$; up to 1.000 at the two higher loads; exact $1.107 \times 10^{-2}$, $7.604 \times 10^{-2}$, $2.808 \times 10^{-1}$, $5.035 \times 10^{-1}$ |
| How important is each element? | today B1–B3 about 0.21 each, A2 0.1572, pipe 0.0011; year 1 chillers 0.073 each, hall units about 0.03, pipe 0.0090 | at year 1 in the order of the unavailabilities |

**Common mistake**

Judging a shared element by its margin alone. The pipe has the largest flow margin on the sheet and is still an order-one cut set for both halls at year 1; it limits little only because it fouls rarely. A shared element is judged on its state in the enumeration, by what is left when it fails and how often it does, not by its margin when it works.

## 15.6 The next measurement (step 7)

Step 7 asks which test or survey of the existing plant would narrow the quantity that decides the build. Here two quantities decide it: the capacity of hall B's air side, which sets most of the shortfall probability and the band of the pipe's share; and the pipe's fouling rate, which sets the pipe's share itself.

### 15.6.1 Hall B's air side

Hall B limits 0.6239 of today's shortfalls and 0.8203 of those at 860 kW, and the band of the shortfall probability at 820 and 860 kW reaches one because hall B's binding load with every unit up is known only to its band. That band is rack B3's, whose fraction no reading sees. Chapter 5 ranked the measurements that narrow it (Section 5.6.2): an inlet probe of 0.1 K at rack B3 removes 76.9 % of the band of hall B's binding load, a permanent 0.3 K sensor there 59.8 %, and a re-survey of rack B3's airflow 58.0 % (Table 5.9). The same measurement is the next one here: it sets hall B's binding load with all units up and with one out, and so the year at which every hall-B unit becomes, and stays, a cut set. Hall A's rack A5 is in the same position, unmeasured and surveyed below its true fraction, and an inlet reading or re-survey there would narrow hall A's band in the same way; its preposterior is not computed here.

The pipe's share needs a record, not a test: a log of the strainer's fouling and clearing would turn its stated rate into a posterior, and by Section 15.4.3 the share moves in proportion.

Chapter 18 ranks the tests and surveys a planner can commission on the site by the width each takes off the quantity that decides the build.

### 15.6.2 What the analysis showed

The reviewer feared the shared pipe, and the design sheet, counting it once with a large flow margin, could neither confirm nor dismiss the fear. The analysis answers the question asked: given a shortfall, the pipe is the limiting element with probability 0.0004 today and 0.0089 at year 1. The halls' own units limit almost all shortfalls, and hall B's air side is the weakest place on the site, N rather than N+1.

Three findings carry the answer. First, the cut sets and the shares answer different questions, and both are needed. At year 1 the pipe is an order-one cut set, a common-mode element under both halls, exactly as the reviewer feared; it limits under one percent of the shortfalls because its fouled state has probability $6.845 \times 10^{-4}$, under three tenths of a hall unit's, against nine other order-one cut sets. Second, and beyond this site, the answer to a conditional question is a ratio of means, and its band over the draws can lie on one side of it. At year 1 the ratio is 0.0089 and the per-draw band 0.0115 to 0.0331: the few draws in which hall B is short with every unit up carry most of the shortfall probability, and the pipe has almost no share in them. Third, a cut-set list at the MAP state hides the elements near the edge. A1 and A3 are not order-one cut sets at the MAP state, by 1.3 and 16.1 kW, but their importances, 0.1159 and 0.0990, show that across the posterior they nearly are.

### 15.6.3 Limits of the model

**The pipe's stated inputs.** The pipe fouls at a stated 0.5 a year, clears in a stated 12 h and keeps a stated 60 % of its flow coefficient. None is estimated, and the pipe's share scales with its rate (Section 15.4.3). A strainer that fouls by degrees would make the pipe's capacity a continuum, and its partly fouled states would lie between the clean and the fouled pipe.

**Independent elements and two states down at most.** Every element fails independently, with one rate pair per kind. A common cause outside the enumeration, a shared control fault or a shared electrical supply to several units, would raise the probability of the order-two states above the products of their unavailabilities. The states with three or more elements down are left out; at the mean unavailabilities they carry $3.44 \times 10^{-6}$. That is of the size of the part of today's shortfall probability that the pipe limits, $4.736 \times 10^{-6}$, so today's pipe share of 0.0004 is the one answer they could move; were every left-out state short and limited by the pipe, the share would still stay under one in a thousand.

**The halls at the pinned supply.** Each hall's binding load is taken at the drawing's 14°C supply, in every state. A chiller plant short of the loop's heat would warm the supply and lower both halls' binding loads together, a second common mode this chapter does not model; Chapter 21 follows such chains. The chillers carry the loop without the economizer, the plant as built.

**The load split.** The site load is split between the halls in today's proportions at every load. A hall that grows faster than the other would reach its binding load at a lower site load than its share says.

**Hall A's calibration.** Hall A's load meter reads 4 kW high, a bias the rows do not model, and its rack A5's fraction is known only from a survey that lies below the truth. Its binding load at the MAP state lies 70.9 kW of hall load above the truth, which lies inside its band of

451.90 to 605.95 kW. The exact posterior of hall A's fit is checked only jointly with hall B's and the site's (effective sample size 1 967 of 2 000), not on its own.

**One window and the design day.** Each hall is read in one window of the night record, and the site at the design day. The cost of a shortfall, which depends on the year's hours and loads, is not computed: Chapter 23 takes the expected unserved energy.

**The constructed record.** The halls' night record, the operating record and the outage records are generated from a stated truth with stated reading errors, not taken from a site; the comparison with that truth in Section 15.4.2 is a check of the method on its own model, not of the model against a real plant.

## 15.7 What to retain

**Definition 15.1** (Shared path in a capacity enumeration)**.** A shared path is an element whose single state enters the capacity of several parts of the plant. In the enumeration of the plant's states it is one coordinate, up or down with its own unavailability, and a state's capacity is the smallest of its groups' capacities. The share of the shortfall probability from states that the element limits, $\bar{P}_{L,g}/\bar{P}_L$, is the probability that it limits the plant given a shortfall: a ratio of means over the posterior. Its minimal cut sets say where it could take the plant down alone; its Fussell–Vesely importance, $\Pr(s_e = 0 \mid \text{shortfall}, \mathbf{y})$, counts every short state with it down.

Four conclusions carry beyond this site. A shared element can be an order-one cut set for every part it serves and still limit few of the plant's shortfalls, because its share is its own unavailability against everyone else's; it is judged by its state, not by its margin. When every element is an order-one cut set, the importances rank as the unavailabilities. A conditional probability over an uncertain plant is a ratio of means, and draws in which the plant is short on its own capacity can pull it below the band of the share given the plant. And a cut-set list read at one value of the plant leaves out the elements near the edge; the importances over the posterior find them.

Table 15.12 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 15.12: The terms of shared paths and common-mode limits, and where each is defined. Units for the site.

| Term | Expression | Units | Meaning | Defined in |
|---|---|---|---|---|
| State probability | $P(\mathbf{s} \mid \mathbf{q}) = \prod_e q_e^{1-s_e}(1-q_e)^{s_e}$ | – | ten elements, independent | 15.1.1 |
| State capacity | $\mathrm{Cap}(\mathbf{s}) = \min_g \mathrm{Cap}_g(\mathbf{s})$ | kW | smallest of six groups | 15.1.1 |
| Limiting element | $\ell(\mathbf{s}) = \arg\min_g \mathrm{Cap}_g(\mathbf{s})$ | – | the group that binds | 15.1.1 |
| Minimal cut set | $\mathrm{Cap}(\mathbf{s}_S) < L$, no smaller $S'$ | – | elements that cut together | 15.1.1 |
| Hall capacity | $Q_h^{\dagger}(\mathbf{z}_h;\ R)/\varphi_h$ | kW | running units only | 15.1.2 |
| Pipe capacity | $[1 - 0.4(1 - s_p)]\,\rho c_p \Delta T_w \mathrm{e}^{l_k}\sqrt{\Delta P_{\max}} - Q_o$ | kW | clean or fouled | 15.1.2 |
| Pipe unavailability | $q_p = \lambda_p\tau_p/(1+\lambda_p\tau_p)$ | – | stated rate and clearing time | 15.1.2 |
| *Rows and posterior* | | | | |
| Factored posterior | $p(\mathbf{z}_A)\,p(\mathbf{z}_B)\,p(\mathbf{z}_{\mathrm{site}})$ $p(\lambda_c,\mu_c)\,p(\lambda_h,\mu_h)$ | – | four records, no shared entry | 15.2.1 |
| Hall-unit rates | Gamma(23, 26 yr), Gamma(22, 496.8 h) | – | failures and repairs | 15.2.4 |
| *Questions and answers* | | | | |
| Shortfall probability | $\bar{P}_L = \mathbb{E}[P_L(\boldsymbol{\theta}) \mid \mathbf{y}]$ | – | over states and posterior | 15.3.1 |
| Limiting share | $\bar{P}_{L,g}/\bar{P}_L$ | – | a ratio of means | 15.3.2 |
| Per-draw share | $r_g(\boldsymbol{\theta}) = P_{L,g}/P_L$ | – | its band, not the answer | 15.3.2 |
| Fussell–Vesely importance | $I_e = \mathbb{E}[P_L^{(e)}]/\mathbb{E}[P_L]$ | – | share from states with $e$ down | 15.3.4 |

## 15.8 Exercises

**Exercise 15.1.** *Formulation.* The planner asks how likely it is that the chillers, rather than the halls or the pipe, limit the site when it falls short at year 1. Write the question over the enumeration and state its answer from this chapter's draws. ($\Pr(\ell = \text{chillers} \mid \mathrm{Cap} < 759\,\mathrm{kW},\ \mathbf{y}) = \bar{P}_{759,\mathrm{C}}/\bar{P}_{759}$, equation (15.13) with $g$ the chillers; by Table 15.5 it is $1.608 \times 10^{-2}$ over $7.565 \times 10^{-2}$, the share 0.2125.)

**Exercise 15.2.** *Derivation.* Derive the pipe's unavailability from its fouling rate and its clearing time, and compute it. (A fouling is a down spell of mean $\tau_p$, and the pipe is clean for a mean of $1/\lambda_p$ between foulings; by equation (9.6) with $\mu_p = 1/\tau_p$, $q_p = \lambda_p/(\lambda_p + 1/\tau_p) = x/(1+x)$ with $x = \lambda_p\tau_p$. With 0.5 a year and 12 h, $x = 6/8\,760 = 6.849 \times 10^{-4}$, and $q_p = 6.845 \times 10^{-4}$.)

**Exercise 15.3.** *Derivation.* Show that a fouled pipe keeping 60 % of its flow coefficient carries 60 % of the clean pipe's loop heat, and compute its capacity in site IT load at the MAP state. (At the allowable drop the flow is $V = k_p\sqrt{\Delta P_{\max}}$, so the loop heat $\rho c_p\,\Delta T_w V$ is proportional to $k_p$; the IT load is the loop heat less the other heat $Q_o = 22\,\mathrm{kW}$. The clean pipe carries $1\,126.8 + 22 = 1\,148.8\,\mathrm{kW}$ of loop heat, the fouled one $0.6 \times 1\,148.8 = 689.3\,\mathrm{kW}$, and its capacity is $689.3 - 22 = 667.3\,\mathrm{kW}$.)

**Exercise 15.4.** *Application.* At year 1 every element is an order-one cut set. Show that the importances of the pipe and of a hall unit then stand nearly in the ratio of their unavailabilities, and check it. (When every element alone takes the site down, the short states with element $e$ down are, to leading order, the state with $e$ alone down, of probability about $q_e$ times the probability that all others are up, a factor common to every element; so $I_e \propto q_e$. The unavailabilities stand in the ratio $6.845 \times 10^{-4}/0.00241 = 0.284$, and the importances in the ratio $0.0090/0.0318 = 0.283$; the pipe's importance is 0.0090.)

**Exercise 15.5.** *Application.* A colleague doubts the stated fouling rate and proposes four times as much. Predict the pipe's share of year 1's shortfalls to first order, and compare it with the computation. (The pipe limits at year 1 only in its fouled state, so the numerator grows about fourfold while the denominator, mostly hall B's all-up shortfalls and the units' one-down states, barely moves: $4 \times 0.0089 = 0.036$ to first order. The fouled states also add to the denominator, so the computed share is a little lower, 0.0332.)

# Chapter 16

# As-Built Margin

## 16.1 Twenty percent on the drawing

The drawing for the site's existing chilled-water plant carries one line that every later design review has repeated. Three chillers of 400 kW nameplate each, 1 200 kW in all, against a design load of 1 000 kW of heat to the chilled-water loop on the design day, a wet bulb of 22°C: a margin of 20 %, signed off at commissioning. The line has never been revisited. The chillers have since run for a year, each was tested for capacity twice at commissioning, and the growth of the halls has brought the design load within a few years of the site's load. The design sheet still says 20 %, and it says it without a band, because on the sheet the nameplate is the capacity.

The planner's question is the drill question, verbatim,

*The design says 20 % margin; what does the as-built model say the margin is?*

It is answered with the method of Chapter 4 on the site's one fit, the posterior of the plant that Chapter 9 built from the operating record and the capacity tests. What is new is the quantity asked for. It is not a binding load, an order date or an enumeration of unit states, but the drawing's own number, read on the plant as built instead of on its nameplates, with the band that the tests leave on it.

**Jargon: capacity margin, nameplate capacity, rating condition, as-built capacity, as-built log-ratio, capacity test, derate, extrapolation to the design day, monthly scenario chain, decision level, settled question**

The *capacity margin* $m$ of a plant is its capacity at the design day over the design load, less one; the drawing's margin $m_0$ is the same ratio taken on the nameplates. The *nameplate capacity* $N_i$ of a chiller is the capacity its maker rates it for at the *rating condition*, here a wet bulb of 22°C, which is also the design day. Its *as-built capacity* is what it delivers in place at that condition, $N_i\,\mathrm{e}^{\kappa_i}$, with $\kappa_i$ its *as-built log-ratio*, zero when the unit meets its nameplate. A *capacity test* runs one chiller at a measured duty and reads the capacity it reaches at the wet bulb of the day. The *derate* $d$ is the fraction of capacity lost per kelvin of wet bulb above the rating condition, and gained per kelvin below it. A test on a cooler day reads the capacity there, and the capacity at the design day is an *extrapolation to the design day* through the derate. A *monthly scenario chain* reads one quantity at a sequence of scenario nodes, one per month, each with that month's wet bulb pinned. A *decision level* is a probability beyond which the planner acts as if the event were certain, here 0.9 and, for its complement, 0.1; a question is *settled* by a measurement when the probability after it ends beyond one of them.

**Key idea**

A margin on a drawing is a function of the plant's capacities, and the as-built model gives it a posterior. At the rating condition the margin is linear in the units' as-built capacities, so its gradient is the units' capacities over the design load and its variance is the capacities' covariance summed. The tests that fix the capacities are taken on cooler days and read each unit's ratio through a derate the units share, so the ratios are correlated and the margin's band is wider than the units' bands added in quadrature. The drawing's number is read on that posterior, with the probability that the real margin lies below any level the review cares about.

### 16.1.1 The margin read on the as-built plant

The drawing states the margin as a ratio of two capacities. Written as a quantity of the plant, it is the as-built capacity of the existing units at the design day over the design load, less one,

$$m(\mathbf{z}) = \frac{1}{L^*}\sum_i \mathrm{Cap}_i^*(\mathbf{z}) - 1, \qquad \mathrm{Cap}_i^*(\mathbf{z}) = N_i\,\mathrm{e}^{\kappa_i}\left[1 - d\,(T_{\mathrm{wb}}^* - 22)\right], \tag{16.1}$$

with the sum over the three chillers, $L^*$ the design load in kW, $T_{\mathrm{wb}}^*$ the wet bulb of the scenario in °C and $\mathrm{Cap}_i^*$ the as-built capacity of chiller $i$ there, the capacity relation of equation (9.3). The drawing's own number is the same expression with every log-ratio at zero, the nameplate,

$$m_0 = \frac{1}{L^*}\sum_i N_i - 1 = \frac{3 \times 400}{1\,000} - 1 = 0.200.$$

The margin is a capacity margin, never an availability: it says how far the capacity of the units exceeds the load with every unit running, and nothing about how often a unit is down, which is the enumeration of Chapter 9.

The margin is linear in the capacities, and its gradient with respect to the state follows entry by entry,

$$\frac{\partial m}{\partial \kappa_i} = \frac{\mathrm{Cap}_i^*}{L^*}, \qquad \frac{\partial m}{\partial d} = -\,\frac{T_{\mathrm{wb}}^* - 22}{L^*}\sum_i N_i\,\mathrm{e}^{\kappa_i}, \qquad \mathrm{Var}(m \mid \mathbf{y}) \approx \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\,\mathbf{g}, \tag{16.2}$$

every other entry zero, the variance by the delta method of equation (4.53). A unit's entry is its share of the design load, so the largest unit carries the longest lever. At the design day $T_{\mathrm{wb}}^* = 22$°C, the derate's entry vanishes: the drawing states its margin at the condition the units are rated for, and the derate, which carries capacity from one wet bulb to another, does not move the margin there. It can still widen the band, through the posterior of the log-ratios, and Section 16.2.1 shows how.

> **Principle 16.1.** *A drawn margin is a derived quantity of the as-built state: the units' capacities at the design day over the design load, less one. Its gradient is each unit's capacity over the load, and its band is that gradient pushed through the covariance of the units' log-ratios, cross terms included.*

### 16.1.2 The site's chiller plant at the design day (step 0)

The plant is the site of Chapter 9, drawn in Figure 9.1: two halls on one chilled-water plant with three chillers C1, C2 and C3 of 400 kW nameplate, one pump, the shared supply pipe, the power feed, and a waterside economizer that is drawn but not built. Its relations are those of equation (9.3) and its state the eight entries of equation (9.4),

$$\mathbf{z} = (\kappa_{C1},\ \kappa_{C2},\ \kappa_{C3},\ d,\ g_e,\ c,\ l_k,\ l_e)^{\mathsf{T}},$$

the three as-built log-ratios, the derate, the economizer's log-ratio to its handbook conductance, and the logs of the chillers' efficiency, the pipe's flow coefficient and the pump's efficiency. Of the plant's relations the margin reads one, the capacity of a chiller, and at the design day that relation reduces to $\mathrm{Cap}_i = N_i\,\mathrm{e}^{\kappa_i}$ and the economizer's duty to $Q_e = 50\,\mathrm{e}^{g_e}$ kW, equation (9.5).

**Modeling choices at the graph.** The margin is taken with all three units running and nothing else in the way: this is the drawing's own definition, and it leaves out both a unit that is down and an element, the pipe, the pump or the feed, that limits the plant before the chillers do. Chapters 13 to 15 read those limits, and a unit out is Chapter 20's. The capacity of a unit is one log-ratio at the rating condition and one derate, linear in the wet bulb and shared by the three: this is what lets two tests at different wet bulbs fix both, and it fails if the units derate differently, or if the derate bends over the range from the tests to the design day. The supply temperature is pinned at the set point of 14°C: the capacity's dependence on the supply is pinned in the plant and does not enter the margin. The economizer is not counted in the drawing's margin, which names three chillers, and is read beside it as the drawing's other element.

### 16.1.3 The margin scenario

The questions are asked at a scenario node of Section 4.3.8, with three inputs pinned, each a modeling choice. The wet bulb is pinned at the design day, $T^*_{\mathrm{wb}} = 22°\mathrm{C}$, and the load at the design load, $L^* = 1\,000\,\mathrm{kW}$ to the loop: this makes the answer a statement on the drawing's terms. The plant is pinned at the existing three chillers, all running, with the economizer absent: this is the drawing's margin. At these pins equation (16.1) gives

$$m(\mathbf{z}) = \frac{400}{1\,000}\left(\mathrm{e}^{\kappa_{C1}} + \mathrm{e}^{\kappa_{C2}} + \mathrm{e}^{\kappa_{C3}}\right) - 1.$$

Two variants of the node answer the fourth question. The first adds the drawn economizer at its design-day duty,

$$m_e(\mathbf{z}) = m(\mathbf{z}) + \frac{50\,\mathrm{e}^{g_e}}{L^*}, \tag{16.3}$$

the margin the drawing would have if the economizer were built as the handbook rates it. The second keeps the three chillers and the design load and moves the wet bulb to each month's value $T_k$ of the record year, a monthly scenario chain. By equation (16.1) the capacity at $T_k$ is the design-day capacity times $1 + d\,(22 - T_k)$, so

$$m_k(\mathbf{z}) = \big(1 + m(\mathbf{z})\big)\big[1 + d\,(22 - T_k)\big] - 1, \qquad \frac{\partial m_k}{\partial d} = \big(1 + m(\mathbf{z})\big)(22 - T_k), \tag{16.4}$$

for $k$ from January to December. Away from the design day the derate has an entry of its own, which grows with the distance of the month's wet bulb from the rating condition.

### 16.1.4 The planner's questions (step 1)

The planner asks four questions of the drawing's margin, and step 1 writes each over the site's state, with the evidence on the right of the bar. Table 16.1 lists them. The first column is the question in the planner's words, the drill question first; the second the general form of its category in the notation of Chapter 2; the third its mathematical form for the site; 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, third and fourth are estimation: the margin with its band, what the band is made of, and the margin with the economizer and along the year. The second is a prediction under a threshold: whether the real margin lies under the level the review would act on.

Table 16.1: The questions asked of the drawing's margin. $\mathbf{y}$ is the site's operating record, its six capacity tests and the priors; $m$, $m_e$ and $m_k$ are the margins of equations (16.1), (16.3) and (16.4).

| Planner's question | General form | Mathematical form | Category | Computation | Answered in |
|---|---|---|---|---|---|
| The design says 20 % margin; what does the as-built model say the margin is? | $\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$, quantiles | $m$: median and 90 % band | estimation | posterior | 16.3.1 |
| What is the probability the real margin is under 10 %? | $\Pr(f(\mathbf{z}) < f^* \mid \mathbf{y})$ | $\Pr(m < 0.10 \mid \mathbf{y})$, with the levels 0.15 and 0 | prediction under a threshold | posterior | 16.3.2 |
| What is the margin made of? | $\mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y})$, split by entry | each chiller's $\mathrm{Cap}_i^*$ with its band; the budget of $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ | estimation | posterior | 16.3.3 |
| What does the drawn economizer add, and when in the year is the margin thinnest? | moments; a scenario chain over months | $m_e$, median and band; $m_k$ for each month | estimation | posterior | 16.3.4 |

### 16.1.5 The evidence as rows (step 2)

Step 2 attaches the site's operating record to the graph, the record of Chapter 9: twelve monthly windows of chiller power, pipe pressure drop, pump power and feed power, six capacity tests and eight priors, 62 rows on 8 entries. The margin is read from a part of it. Each chiller was tested twice for capacity at commissioning, at two different wet bulbs: C1 read 386.70 kW at 18°C and 363.56 kW at 22°C; C2 388.98 kW at 21°C and 425.84 kW at 15°C; C3 396.21 kW at 15°C and 361.82 kW at 20°C; each test of standard deviation 8 kW. The test of chiller $i$ at wet bulb $T_j$ is the row

$$r_j(\mathbf{z}) = \frac{N_i\,\mathrm{e}^{\kappa_i}\left[1 - d\,(T_j - 22)\right] - y_j}{8}, \qquad \frac{\partial r_j}{\partial \kappa_i} = \frac{\mathrm{Cap}_i(T_j)}{8}, \qquad \frac{\partial r_j}{\partial d} = \frac{N_i\,\mathrm{e}^{\kappa_i}\,(22 - T_j)}{8}, \tag{16.5}$$

with $y_j$ the reading and $\mathrm{Cap}_i(T_j)$ the capacity at the test's wet bulb. The priors put each log-ratio at zero with standard deviation 0.08, the nameplate within about 8 %, and the derate at $0.015\,\mathrm{K}^{-1}$ within 0.005.

**Which row acts on which entry.** The monthly windows read the chiller power, which depends on the duty and the efficiency but not on the capacity, and the pipe, pump and feed, which depend on the flow. None of them has an entry in the columns of the log-ratios or the derate. With $*$ for an entry that can be nonzero and the columns in the order of the state, the six test rows have the pattern

$$\mathbf{J}_{\text{test}} = \begin{bmatrix} * & 0 & 0 & * & 0 & 0 & 0 & 0 \\ * & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & * & 0 & * & 0 & 0 & 0 & 0 \\ 0 & * & 0 & * & 0 & 0 & 0 & 0 \\ 0 & 0 & * & * & 0 & 0 & 0 & 0 \\ 0 & 0 & * & * & 0 & 0 & 0 & 0 \end{bmatrix}, \tag{16.6}$$

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

rows in the order C1 at 18 and 22°C, C2 at 21 and 15°C, C3 at 15 and 20°C; the second line is not a row but the pattern of the margin's gradient (16.2) at the design day. Three facts are read here before any number. The margin's gradient lives in the three columns of the log-ratios, and the only rows with entries there are the six tests and three priors. Five of the six tests also have an entry in the derate's column, by equation (16.5), positive because they were taken below the rating condition; only C1's test at 22°C reads its log-ratio alone. And no reading touches the economizer, whose log-ratio keeps its prior.

**Modeling choices at the rows.** Each test is taken as Gaussian, unbiased and independent of the others given the state: this makes each test one normalized row, and it fails if the tests share a bias, such as one flow meter used for all six. The tests are taken to describe the units today: this lets commissioning tests stand for the as-built plant, and it fails if a unit has lost capacity since, to fouling or to a lost charge of refrigerant. The derate is one entry for the three units: this is what lets the tests at different wet bulbs inform it together, and it fails if one unit derates faster than the others.

## 16.2 The posterior of the state (step 3)

All four questions of Table 16.1 are read from one posterior, the site's one fit of Chapter 9. No row is added and no fit is repeated; this section derives what the margin needs from that posterior, cites the fit, and states the draws the answers use.

### 16.2.1 What the margin needs from the state

By the pattern (16.6) no row couples the log-ratios and the derate to the efficiency, the pipe, the pump or the economizer, so the information matrix of the site's fit is block diagonal, and so is its inverse,

$$\boldsymbol{\Lambda} = \begin{bmatrix} \boldsymbol{\Lambda}_{\kappa d} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & 0.35^{-2} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \boldsymbol{\Lambda}_{\mathrm{o}} \end{bmatrix}, \qquad \boldsymbol{\Sigma} = \begin{bmatrix} \boldsymbol{\Lambda}_{\kappa d}^{-1} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & 0.35^{2} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \boldsymbol{\Lambda}_{\mathrm{o}}^{-1} \end{bmatrix}, \tag{16.7}$$

with $\boldsymbol{\Lambda}_{\kappa d}$ the $4 \times 4$ block of the three log-ratios and the derate, built from the six tests and their four priors, and $\boldsymbol{\Lambda}_{\mathrm{o}}$ that of $(c, l_k, l_e)$, built from the monthly windows. The margin at the design day needs only the $3 \times 3$ covariance $\boldsymbol{\Sigma}_{\kappa\kappa}$ of the log-ratios, the upper left of $\boldsymbol{\Lambda}_{\kappa d}^{-1}$.

That block has a structure of its own. Each test row has an entry in one log-ratio's column, so the log-ratios' own information $\mathbf{D} = \boldsymbol{\Lambda}_{\kappa\kappa}$ is diagonal, and their information with the derate is the vector $\mathbf{a}$ with entries $a_i = \sum_j (\partial r_j / \partial \kappa_i)(\partial r_j / \partial d)$ over the tests of chiller $i$, nonnegative by equation (16.5). The covariance of the log-ratios with the derate integrated out is the inverse of the Schur complement of the derate's block, as in Chapter 3,

$$\boldsymbol{\Sigma}_{\kappa\kappa} = \Big(\mathbf{D} - \frac{\mathbf{a}\,\mathbf{a}^{\mathsf{T}}}{\Lambda_{dd}}\Big)^{-1} = \mathbf{D}^{-1} + \frac{\mathbf{D}^{-1}\mathbf{a}\,\mathbf{a}^{\mathsf{T}}\mathbf{D}^{-1}}{\Lambda_{dd} - \mathbf{a}^{\mathsf{T}}\mathbf{D}^{-1}\mathbf{a}}, \tag{16.8}$$

with $\Lambda_{dd}$ the derate's own information; the second form is the Sherman–Morrison identity. The correction is a positive multiple of the outer product of a vector with nonnegative entries, so every covariance between two log-ratios is positive or zero. The mechanism is the shared derate. A test on a cool day reads $N_i\, \mathrm{e}^{\kappa_i}[1 + d\,(22 - T_j)]$, and a larger derate explains the same reading with a smaller log-ratio: each log-ratio is negatively correlated with the derate, and so the three are positively correlated with each other. By equation (16.2) every entry of the margin's gradient is positive, so every cross term $g_i\, g_k\, \Sigma_{ik}$ adds to the variance,

$$\mathrm{Var}(m \mid \mathbf{y}) \approx \sum_i g_i^2\, \Sigma_{ii} + \sum_{i \neq k} g_i\, g_k\, \Sigma_{ik} \;\geq\; \sum_i g_i^2\, \Sigma_{ii}, \tag{16.9}$$

and the margin is less well known than the three units' spreads, added in quadrature, would say. The derate has no entry in the margin's gradient, and yet its uncertainty is inside the margin's band, carried by the log-ratios' covariance. Chiller C1, whose second test was taken at the rating condition, has the smallest coupling $a_{C1}$, from one test four kelvin below it; C2's and C3's tests sit eight and nine kelvin below it in all.

### 16.2.2 The site's fit, cited

The MAP state $\hat{\mathbf{z}}$ and the Laplace covariance $\boldsymbol{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1}$ at $\hat{\mathbf{z}}$ are those of Section 9.2.2, the Gauss–Newton fit of the operating record in the form of equation (4.38), with twice the objective 59.408 on 54 degrees of freedom and an upper tail probability of 0.285. Table 9.2 lists the entries. The ones the margin reads are the three log-ratios, $-0.0909$, $-0.0394$ and $-0.1157$ with posterior spreads 0.0158, 0.0175 and 0.0189, about a fifth of their prior spread, and the derate, at a MAP value of $0.0152\,\mathrm{K}^{-1}$ with spread 0.0029, a little more than half its prior's. All three log-ratios lie below zero: at the MAP state the chillers deliver 365.2, 384.5 and 356.3 kW at the design day, against their 400 kW nameplate. C1 has the narrowest spread, as equation (16.8) expects, since one of its tests reads it at the rating condition. The economizer's log-ratio keeps its prior, a spread of 0.35, uncorrelated with every other entry by equation (16.7).

### 16.2.3 The draws and the exact posterior

The answers read the posterior through the draws of Section 9.2.4: 20 000 states $\mathbf{z}_i$ from the Laplace normal $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$, drawn with the seed 9, each pushed through the exact margin (16.1) and the economizer's variant (16.3); the monthly chain (16.4), twelve margins per state, reads the first 4 000 of them. The reported values are the plain quantiles and shares over the draws, the state route. The exact posterior is the same draws under the weights of equation (5.16), with the objective of the site's fit; their effective sample size is 19 913 of the 20 000. The margin is nearly linear in the log-ratios over their posterior width, so the delta form of equation (16.2) at $\hat{\mathbf{z}}$ stands beside both as the single-gradient route.

## 16.3 The answers (step 4)

Every answer is read from the draws of Section 16.2.3. The reported value is the quantile or share over the Laplace draws, the exact posterior stands beside it with the gap, and the delta form at $\hat{\mathbf{z}}$ stands beside both.

### 16.3.1 The as-built margin

**Planner's question.** The design says 20 % margin; what does the as-built model say the margin is? The question asks for the drawing's margin read on the as-built units, with its 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 median and by the central credible interval of equation (2.11), here between the 5 % and 95 % points.

Table 16.2: The margin of the existing three chillers at the design day, by three routes, with the design sheet's number. Draws: 20 000 states from the Laplace posterior through the exact margin (16.1); exact: the same draws reweighted; delta form: the normal at $\hat{\mathbf{z}}$. Gaps from the printed values.

| Route | Median [–] | 5 % point [–] | 95 % point [–] | $\Pr(m < 0.10)$ [–] | $\Pr(m < 0.15)$ [–] |
|---|---|---|---|---|---|
| Laplace draws (reported) | 0.1062 | 0.0835 | 0.1293 | 0.3279 | 0.9988 |
| Exact posterior | 0.1057 | 0.0830 | 0.1287 | 0.3395 | 0.9991 |
| Gap, exact minus reported | −0.0005 | −0.0005 | −0.0006 | +0.0116 | +0.0003 |
| Delta form, $\mathcal{N}(0.1061,\ 0.0139^2)$ | 0.1061 | 0.0832 | 0.1290 | 0.3316 | 0.9992 |
| Design sheet, nameplate | 0.200 | – | – | – | – |

**Specialized form.** The question is equation (2.9) with $f(\mathbf{z}) = m(\mathbf{z})$ of equation (16.1) at the design-day scenario and $p(\mathbf{z} \mid \mathbf{y})$ the site's posterior. It is reported by the draws,

$$m_i = m(\mathbf{z}_i), \quad \mathbf{z}_i \sim \mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma}), \qquad \hat{F}(x) = \frac{1}{N} \sum_{i=1}^{N} \mathbf{1}[m_i \le x], \tag{16.10}$$

from whose distribution function the median and the band are read; the exact posterior weights the same $m_i$ by equation (5.16); and the delta form is the normal $\mathcal{N}(m(\hat{\mathbf{z}}),\ \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g})$ with $\mathbf{g}$ of equation (16.2).

**Building blocks.** The MAP state and the covariance of the site's fit (Section 16.2.2); the margin (16.1) in each draw; the gradient (16.2); the weights (5.16).

**The margin.** By the draws (16.10) the reported margin is

$$\text{median } m = 0.1062, \qquad m \in [0.0835,\ 0.1293] \text{ with } 90\,\% \text{ probability.} \tag{16.11}$$

The exact posterior gives a median of 0.1057 and the band [0.0830, 0.1287], gaps of −0.0005, −0.0005 and −0.0006, exact minus reported. The delta form puts the margin at $m(\hat{\mathbf{z}}) = 0.1061$ with standard deviation 0.0139, and its band, $0.1061 \pm 1.645 \times 0.0139$, is [0.0832, 0.1290]. Table 16.2 sets the three routes side by side, with the design sheet's number beneath them: each row is a route, the columns the median, the ends of the band, and the probabilities of the next subsection. In capacity, the margin's median is $1\,000 \times 1.1062 = 1\,106.2$ kW of the three units together, against the drawing's 1 200 kW.

The posterior of the margin is drawn in Figure 16.1: the horizontal axis is the margin, the vertical axis its posterior density. The solid curve is the delta form's normal for the existing three chillers, with the area below the 10 % level shaded; the dashed curve is the margin with the drawn economizer of Section 16.3.4; the red line is the drawing's 20 %, the dotted lines the stated truth of the record, check markers for Section 16.4.2; and the bars above the curves are the 90 % bands of the draws and of the exact posterior, with their medians.

The drawn 20 % lies far outside the posterior: half the drawn margin is not there. The three routes agree to within 0.0006 at every point, and the delta form's normal nearly coincides with the draws, because the margin is a sum of three exponentials whose arguments have posterior spreads under 0.02, and over that width an exponential is nearly a line. The exact posterior sits 0.0005 below the Laplace draws all along the band: the reweighting moves the margin down by a twenty-eighth of its standard deviation, the departure of the site's posterior from its Laplace normal in the direction of the margin's gradient.

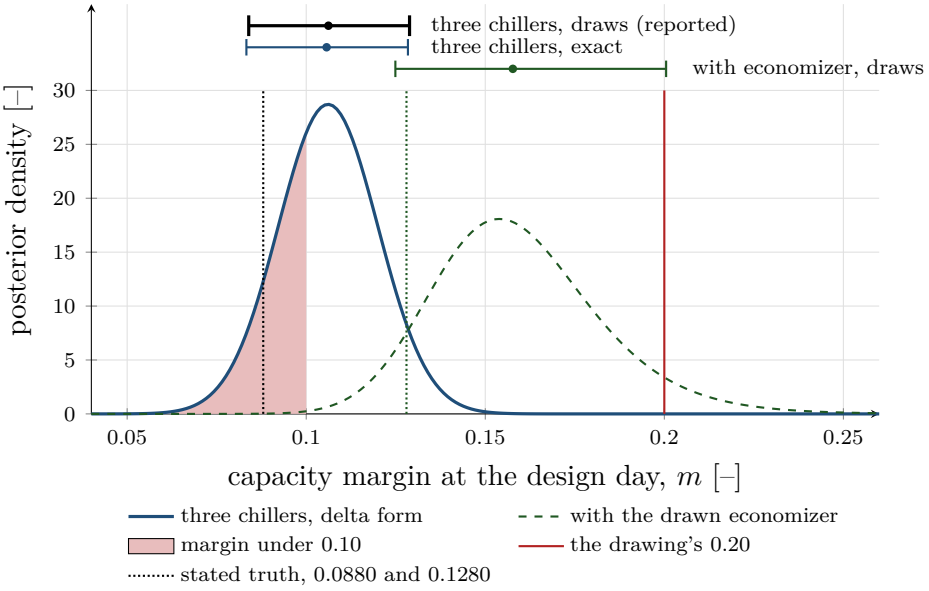



Figure 16.1: The posterior of the capacity margin at the design day. Solid: the delta form for the existing three chillers, $\mathcal{N}(0.1061,\ 0.0139^2)$, with the area below the 10 % level shaded. Dashed: the margin with the drawn economizer, $m + 0.05\,\mathrm{e}^{g_e}$, drawn from the same normal and the economizer's handbook prior $g_e \sim \mathcal{N}(0,\ 0.35^2)$, independent of it by equation (16.7). Red: the drawing's margin. Dotted: the stated truth of the record without and with the economizer, check markers (Section 16.4.2). Bars: the 90 % bands with their medians; the top bar is the reported band of the three chillers.

### 16.3.2 The probability that the margin is under ten percent

**Planner's question.** What is the probability the real margin is under 10 %? The question asks for a threshold probability on the margin, at the level below which the review would act.

**General form.** A prediction under a threshold is the posterior probability of an event on a derived quantity, equation (2.12),

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

**Specialized form.** The question is equation (2.12) with $f(\mathbf{z}) = m(\mathbf{z})$ and $f^* = 0.10$, and beside it the levels 0.15, 0 and the drawing's 0.20. The reported probability is the share of the draws (16.10) below the level,

$$\Pr(m < f^* \mid \mathbf{y}) \approx \hat{P} = \frac{1}{N}\sum_{i=1}^{N} \mathbf{1}\big[m_i < f^*\big], \qquad \mathrm{se}(\hat{P}) = \sqrt{\hat{P}\,(1-\hat{P})/N}, \tag{16.12}$$

with se its Monte Carlo standard error; the exact posterior sums the weights of the same draws; and the delta form is $\Phi\big((f^* - m(\hat{\mathbf{z}}))/\sqrt{\mathbf{g}^\top \boldsymbol{\Sigma} \mathbf{g}}\big)$.

**Building blocks.** The draws (16.10) and their weights (5.16); the delta form's mean and standard deviation of Section 16.3.1; the normal distribution function $\Phi$.

**The ten percent level.** By equation (16.12)

$$\Pr(m < 0.10 \mid \mathbf{y}) = 0.3279,$$

with a Monte Carlo standard error of 0.0033. The exact posterior gives 0.3395, and the delta form 0.3316. The three are stated as they stand: about one chance in three, 0.33 by the draws with the exact 0.34 beside it. The gap of +0.0116 between the exact posterior and the draws is three and a half standard errors, but it is not sampling noise: the two use the same states, and the gap is the shift of 0.0005 found in Section 16.3.1, read at a level that lies 0.44 of a standard deviation below the median, where the density is near its peak and a small shift moves much probability. Table 16.2 carries the other levels. The probability that the margin is under 15 % is 0.9988, with standard error 0.0002, exact 0.9991 and delta form 0.9992; under the drawing's 20 % it is 1.0000 by every route; and under zero it is 0.0000 by every route: in none of the 20 000 draws do the three units, running together, fail to carry the design load at the design day. The review can be practically sure the margin is under 15 %, and it cannot say on which side of 10 % it lies.

### 16.3.3 What the margin is made of

**Planner's question.** What is the margin made of? The question asks for each chiller's as-built capacity at the design day, against its nameplate, and for the entries of the state that carry the margin's band.

**General form.** The variance of a derived quantity under the posterior, equation (2.10),

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

split into one share per entry, with the moments of each unit's capacity beside it.

**Specialized form.** Each unit's capacity is equation (2.10) with $f(\mathbf{z}) = \mathrm{Cap}_i^*(\mathbf{z}) = 400\,\mathrm{e}^{\kappa_i}$, read from the draws as a median and a band, and with the ratio of the median to the nameplate. The margin's variance is split by the covariance split of equation (5.25),

$$s_i = \frac{g_i \, (\boldsymbol{\Sigma}\mathbf{g})_i}{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}}, \qquad \mathbf{g} = \Big(\frac{\mathrm{Cap}_{C1}^*}{L^*},\ \frac{\mathrm{Cap}_{C2}^*}{L^*},\ \frac{\mathrm{Cap}_{C3}^*}{L^*},\ 0,\ 0,\ 0,\ 0,\ 0\Big)^{\mathsf{T}}\Big|_{\hat{\mathbf{z}}}, \tag{16.13}$$

each entry's diagonal term with half of its cross terms. By equation (16.8) the cross terms among the log-ratios are positive, so each unit's share exceeds its own diagonal term over the variance; the derate's share is zero because its gradient entry is.

**Building blocks.** The draws (16.10); the gradient (16.2) at $\hat{\mathbf{z}}$; the covariance of the site's fit; the split (5.25).

**The units.** Table 16.3 gives each chiller's as-built capacity at the design day: each row is a unit, the columns its nameplate, its MAP log-ratio and posterior spread, the median and band of its capacity over the draws, and the ratio of the median to the nameplate. The same capacities are drawn in Figure 16.2, one bar per unit with its median, against the nameplate line, with C1's test at the rating condition marked beside its bar.

No unit reaches its nameplate, and no band does: the highest 95 % point, C2's 395.8 kW, stays below 400 kW. C3 is the furthest short, at 0.891 of nameplate, and C2 the closest, at 0.961. Together the medians come to 1 106.0 kW, about 94 kW short of the drawing's 1 200 kW: almost a whole tenth of the design load, which is the difference between the drawing's 20 % and the model's 10.6 %. C1's test at the rating condition, 363.56 kW, sits inside C1's band and below its median: the fit draws C1's capacity a little above its own direct test, through its other test at 18°C and the shared derate.

Table 16.3: The existing chillers as built at the design day. Log-ratios from the site's fit (Table 9.2); capacities over the 20 000 draws of the posterior.

| Chiller | Nameplate [kW] | $\hat{\kappa}_i$ [–] | $\sqrt{\Sigma_{ii}}$ [–] | Median [kW] | 5 % point [kW] | 95 % point [kW] | Ratio [–] |
|---|---|---|---|---|---|---|---|
| C1 | 400 | −0.0909 | 0.0158 | 365.3 | 356.1 | 375.0 | 0.9132 |
| C2 | 400 | −0.0394 | 0.0175 | 384.5 | 373.7 | 395.8 | 0.9612 |
| C3 | 400 | −0.1157 | 0.0189 | 356.2 | 345.4 | 367.5 | 0.8906 |

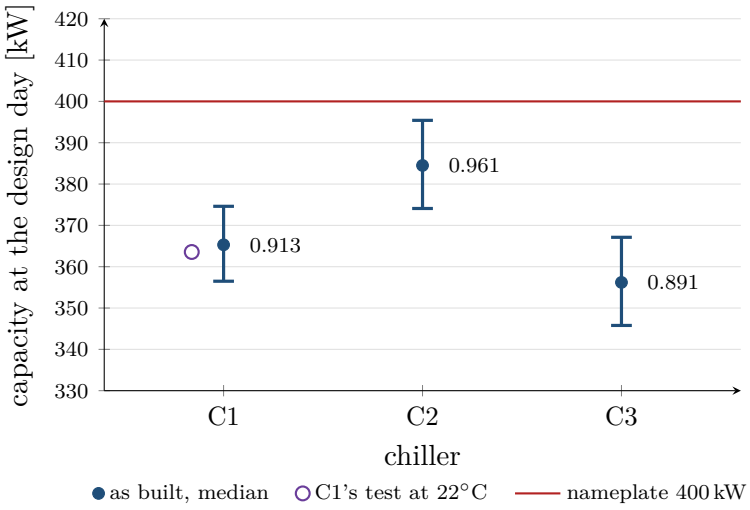



Figure 16.2: Each existing chiller's as-built capacity at the design day, median and 90 % band over the draws (Table 16.3), labeled with the ratio of the median to the nameplate. Red: the 400 kW nameplate. Open circle: C1's capacity test at the rating condition, 363.56 kW, offset to the left of C1's bar.

Table 16.4: The variance budget of the margin at the design day, equation (16.13), at the MAP state. Gradient entries are the MAP capacities over the 1 000 kW design load; the derate's vanishes at the rating condition.

| Entry [unit] | $\partial m/\partial z_i$ at $\hat{\mathbf{z}}$ [1/entry's unit] | $\sqrt{\Sigma_{ii}}$ [entry's unit] | Share $s_i$ [–] |
|---|---|---|---|
| $\kappa_{C1}$ [–] | 0.3652 | 0.0158 | 0.260 |
| $\kappa_{C2}$ [–] | 0.3845 | 0.0175 | 0.369 |
| $\kappa_{C3}$ [–] | 0.3563 | 0.0189 | 0.371 |
| $d$ [$\mathrm{K}^{-1}$] | 0 | 0.0029 | 0.000 |
| $g_e$ [–] | 0 | 0.3500 | 0.000 |
| $c$, $l_k$, $l_e$ [–] | 0 | 0.0061, 0.0046, 0.0190 | 0.000 |

**The budget.** Table 16.4 splits the margin's variance, $0.0139^2$ by the delta form: each row is an entry of the state, the columns its entry in the gradient (16.13) at $\hat{\mathbf{z}}$, its posterior spread and its share.

The band is the three log-ratios' and nothing else's: C3 carries 0.371, C2 0.369 and C1 0.260. C3's share comes from its wider spread, C2's from its larger lever, the largest of the three gradient entries, and C1's is the smallest because its test at the rating condition fixes it best. The derate has no share, as the gradient requires, but by equation (16.8) it is inside the three shares: the tests read each ratio through it, and its uncertainty correlates the three positively, so that every share is larger than the unit's own diagonal term over the variance.

### 16.3.4 The economizer and the months

**Planner's question.** What does the drawn economizer add, and when in the year is the margin thinnest? The question asks for the margin with the drawing's other element counted, and for the margin at each month's wet bulb at the design load.

**General form.** An estimation question asks for the moments and quantiles of a derived quantity, equation (2.10),

$$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}] = \int f(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z},$$

here of a family of quantities, one per scenario node of a chain.

**Specialized form.** The question is equation (2.10) with $f(\mathbf{z}) = m_e(\mathbf{z})$ of equation (16.3) for the economizer, and $f(\mathbf{z}) = m_k(\mathbf{z})$ of equation (16.4) for each month $k$ of the record year, each read from the draws as a median and a band. The economizer's log-ratio keeps its prior and, by equation (16.7), is independent of the chillers' entries, so the margin with it is the sum of two independent parts, the chillers' margin and the economizer's $0.05\,\mathrm{e}^{g_e}$, the second log-normal and skewed to the right. The monthly margins carry the derate through its own entry (16.4), which grows with the distance of the month's wet bulb from the rating condition.

**Building blocks.** The draws (16.10); the economizer's design-day duty (9.5); the months' wet bulbs of the record; the chain (16.4).

**With the economizer.** The economizer at the design day has a median duty of 50.2 kW over the draws, the handbook's 50 kW at the median of its prior. With it the margin has a median of 0.1577 and the band [0.1246, 0.2008], and the probability that it is under 10 % falls to 0.0015. The exact posterior gives a median of 0.1571 and a probability of 0.0017, gaps of

Table 16.5: The margin of the existing three chillers at the 1 000 kW design load and each month's wet bulb of the record year, equation (16.4), over the first 4 000 draws of the posterior; the design day over all 20 000 (Table 16.2).

| Month | $T_k$ [°C] | Median $m_k$ [–] | 5 % point [–] | 95 % point [–] |
|---|---|---|---|---|
| January | 19.83 | 0.1429 | 0.1263 | 0.1595 |
| February | 19.83 | 0.1429 | 0.1263 | 0.1595 |
| March | 18.54 | 0.1648 | 0.1492 | 0.1803 |
| April | 16.29 | 0.2023 | 0.1826 | 0.2215 |
| May | 13.71 | 0.2458 | 0.2155 | 0.2750 |
| June | 11.46 | 0.2836 | 0.2427 | 0.3223 |
| July | 10.17 | 0.3053 | 0.2577 | 0.3503 |
| August | 10.17 | 0.3053 | 0.2577 | 0.3503 |
| September | 11.46 | 0.2836 | 0.2427 | 0.3223 |
| October | 13.71 | 0.2458 | 0.2155 | 0.2750 |
| November | 16.29 | 0.2023 | 0.1826 | 0.2215 |
| December | 18.54 | 0.1648 | 0.1492 | 0.1803 |
| Design day | 22.00 | 0.1062 | 0.0835 | 0.1293 |

−0.0006 and +0.0002, the chillers' part of the shift found in Section 16.3.1. The band is skewed, 0.0331 below the median and 0.0431 above it, the economizer's log-normal guess; its width, 0.0762, is two thirds again the width of the chillers' band. The drawing's 20 % is only nearly recovered, at the band's upper end, and only by an element that is not built, on a handbook guess that no reading has tested. The dashed curve of Figure 16.1 is this margin. Chapter 17 splits the cooling headroom's band between the chillers and that guess.

**Month by month.** Table 16.5 gives the margin of the three chillers at the design load and each month's wet bulb: each row is a month, the columns its wet bulb and the median and band of $m_k$ over the first 4 000 draws; the last row is the design day of equation (16.11), over all 20 000. The same margins are drawn in Figure 16.3: the horizontal axis is the month, the vertical axis the margin; the shaded band is the 90 % band, the line its median; the point with its bar at the right is the design day; the red line is the drawing's 20 % and the dashed line the 10 % level.

At the design load the margin runs from 0.1429 in January and February to 0.3053 in July and August. The design day is the thinnest of all by construction: its wet bulb, 22°C, lies above every month's, the warmest being 19.83°C, and a chiller loses capacity as the wet bulb rises. The drawing's 20 % is met at the design load, by the median, only from April to November, in the months whose wet bulb is 16.29°C or lower. The bands are not equally wide. They are narrowest in March and December, 0.1492 to 0.1803, at a wet bulb of 18.54°C, near the middle of the wet bulbs at which the six tests were taken; they widen toward July, 0.2577 to 0.3503, where the month lies almost five kelvin below the coolest test, at 15°C, and the derate's entry in equation (16.4) is largest; and they widen toward the design day too, whose band, 0.0835 to 0.1293, is wider than January's, 0.1263 to 0.1595. The tests fix the capacity best where they were taken, and the design day lies at the warm edge of their range. The margin the drawing states is the one the record knows least well among the conditions near the tests.

## 16.4 The checks (step 5)

The answers of Section 16.3 rest on the Laplace normal of the site's fit, on its draws, and on a record generated from a stated truth; each is tested here.

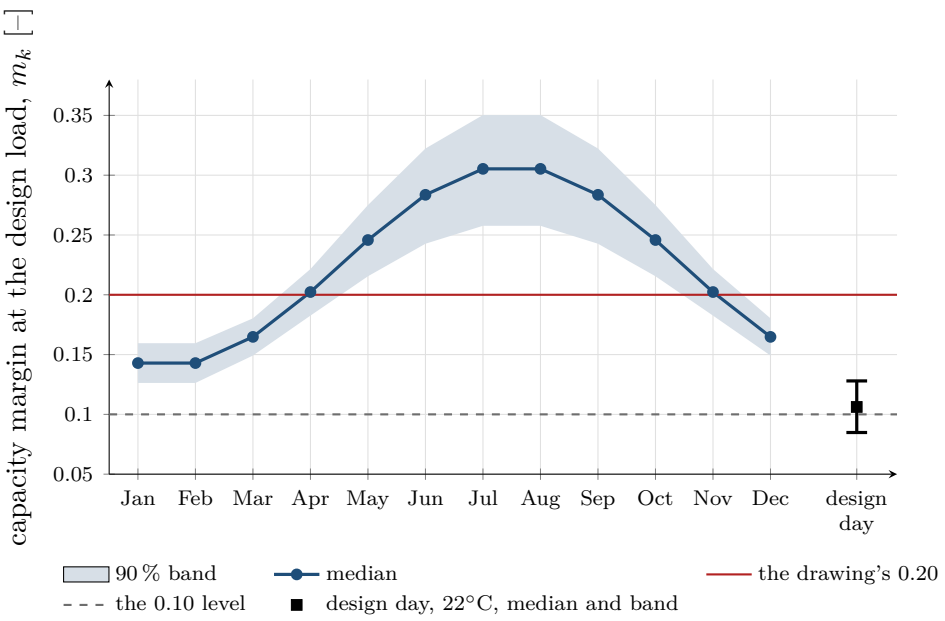



Figure 16.3: The margin of the existing three chillers at the 1 000 kW design load and each month's wet bulb of the record year (Table 16.5), by the monthly scenario chain (16.4). Shaded: the 90 % band over the first 4 000 draws; line: the median. At the right, the design day over all 20 000 draws: median 0.1062, below every month's band, and band 0.0835 to 0.1293.

### 16.4.1 The draws against the exact posterior and the delta form

The reported values are quantiles and shares of 20 000 draws; the exact posterior reweights the same draws, with an effective sample size of 19 913. Table 16.2 gives the gaps: −0.0005 at the median and the 5 % point and −0.0006 at the 95 % point, and +0.0116 and +0.0003 in the probabilities of the 10 and 15 % levels, against Monte Carlo standard errors of 0.0033 and 0.0002. The draws and the exact posterior use the same states, so the gaps are not sampling noise: they are the small departure of the site's posterior from its Laplace normal, a downward shift of the margin of 0.0005, which matters only where a level falls near the center of the distribution, as 10 % does. The delta form sits between the two at the 10 % level, 0.3316, and within 0.0003 of the draws at every quantile: the margin's nonlinearity in the log-ratios is negligible over their posterior width. With the economizer the exact posterior gives a median of 0.1571 against 0.1577 and a probability under 10 % of 0.0017 against 0.0015: the economizer's entry is drawn from its exact normal prior, so the gap is the chillers' part alone. For the monthly margins, read on the first 4 000 draws, the reweighting is not computed.

### 16.4.2 The answers against the stated truth

The operating record is generated from a stated truth, the last column of Table 9.2, and the last check compares the answers with it. At the true log-ratios −0.1054, −0.0619 and −0.1278, the units deliver $400\,\mathrm{e}^{-0.1054} = 360.0$ kW, 376.0 kW and 352.0 kW at the design day, and the true margin is

$$m_{\text{true}} = \frac{360.0 + 376.0 + 352.0}{1\,000} - 1 = 0.0880,$$

and 0.1280 with the economizer at its true conductance. Table 16.6 collects the checks: each row is a quantity, the columns its reported value, the route beside it, the truth and the reading. The true margin lies inside the reported band, 0.0835 to 0.1293, near its lower end; the event of the second question holds of it, since 0.0880 is under 10 %; and the true margin with the

Table 16.6: The checks of the margin answers. Reported values over 20 000 draws of the Laplace posterior; beside them the exact posterior or the route named; the truth is the stated truth of the record.

| Quantity | Reported [– or kW] | Beside [– or kW] | Truth [– or kW] | Reading |
|---|---|---|---|---|
| Margin $m$, median | 0.1062 | exact 0.1057 | 0.0880 | inside [0.0835, 0.1293] |
| Margin $m$ at $\hat{\mathbf{z}}$, delta form | 0.1061 | sd 0.0139 | 0.0880 | MAP 0.0181 high, 1.30 sd |
| $\Pr(m < 0.10)$ | 0.3279 | exact 0.3395 | event holds | delta form 0.3316 |
| $\Pr(m < 0.15)$ | 0.9988 | exact 0.9991 | event holds | delta form 0.9992 |
| Margin with the economizer, median | 0.1577 | exact 0.1571 | 0.1280 | inside [0.1246, 0.2008] |
| C1 at the design day [kW] | 365.3 | – | 360.0 | inside [356.1, 375.0] |
| C2 at the design day [kW] | 384.5 | – | 376.0 | inside [373.7, 395.8] |
| C3 at the design day [kW] | 356.2 | – | 352.0 | inside [345.4, 367.5] |
| Derate $d$, MAP [$\mathrm{K}^{-1}$] | 0.0152 | sd 0.0029 | 0.0180 | read low, by 0.97 sd |
| Design sheet, nameplate | 0.200 | – | 0.0880 | outside every band |
| Draws | $N = 20\,000$, seed 9, effective sample size 19 913 | | | |

economizer lies inside its band, 0.1246 to 0.2008. Each unit's true capacity lies inside its own band.

The MAP margin lies 0.0181 above the truth, 1.30 of the delta form's standard deviation: the tests read the chillers slightly generous. Unit by unit, the MAP capacities exceed the truth by 5.2, 8.5 and 4.3 kW, C2 the most. Two causes combine. C1's and C2's tests read above the capacities they measured; C1's direct test at the rating condition read 363.56 kW against a true 360.0 kW. And the derate is read low, 0.0152 against a true 0.0180 $\mathrm{K}^{-1}$: a test on a cool day reads the design-day capacity times $1 + d\,(22 - T_j)$, so a derate read too small leaves too large a capacity at the design day. That second cause raises C3 too, whose two tests were taken at 15 and 20°C. Both causes act on all three units in the same direction, which is what the positive covariances of equation (16.8) say they can do. The design sheet's 20 % lies more than eight delta-form standard deviations above the truth.

## 16.5 The report (step 6)

The as-built margin of the existing chillers at the design day is 10.6 %, with a 90 % band of 8.4 to 12.9 %, not the drawing's 20 %. There is about one chance in three that it is under 10 %, 0.33 by the draws and 0.34 by the exact posterior, and the stated truth of the record, 8.8 %, lies under it.

### 16.5.1 The answers and their bands

Table 16.7 gives the answers in the order of Table 16.1. Each row is a question; the second column gives the reported answer with the exact posterior beside it where it is computed, and the third its band or the routes beside it.

### 16.5.2 The design-review entry

The design review's entry corrects the drawing. The margin of 20 % was a nameplate margin: three 400 kW units over a 1 000 kW design load. On the as-built model, fitted to the six

Table 16.7: The report on the drawing's margin, at the design day of 22°C wet bulb and the 1 000 kW design load, with the three existing chillers running. Reported values by 20 000 draws of the site's posterior; the exact posterior beside each where computed.

| Question | Answer | Band or routes beside |
|---|---|---|
| The design says 20 % margin; what does the as-built model say the margin is? | median 0.1062 (exact: 0.1057), against the drawing's 0.200 | 0.0835 to 0.1293 (exact: 0.0830 to 0.1287), central 90 %; delta form 0.1061, sd 0.0139 |
| What is the probability the real margin is under 10 %? | 0.3279 (exact: 0.3395) | delta form 0.3316; under 15 % 0.9988; under 0 0.0000 |
| What is the margin made of? | C1 365.3, C2 384.5, C3 356.2 kW at the design day, 0.913, 0.961 and 0.891 of nameplate | variance shares 0.260, 0.369, 0.371; the derate 0 |
| What does the drawn economizer add, and when in the year is the margin thinnest? | with it, median 0.1577 (exact: 0.1571); month by month at the design load, 0.1429 (January) to 0.3053 (July) | 0.1246 to 0.2008, $\Pr(m_e < 0.10) = 0.0015$ (exact: 0.0017); the design day thinnest, 0.1062 |

commissioning tests and the year of operation, the three units carry a median of 1 106.2 kW at the design day, a margin of 10.6 % with a 90 % band of 8.4 to 12.9 %; the probability that it is under 10 % is 0.33 by the draws, 0.34 by the exact posterior and 0.33 by the delta form, and the probability that it is under 15 % is 0.999 by every route. The entry restates the margin on the as-built model and records what the band rests on: the three units' log-ratios, read through a derate they share, in nearly equal parts. It records too that the drawn economizer would raise the margin to a median of 15.8 %, band 12.5 to 20.1 %, and that it is not counted: an element enters the margin when it is built and tested, not when it is drawn, and its conductance is a handbook guess. Every later review that quotes the margin quotes 10.6 % with its band.

**Common mistake**

Carrying a nameplate margin forward as the plant's margin. The drawing's number assumes every unit meets its rating; the tests say none does, and the margin falls by the units' shortfall summed, about a tenth of the design load here. Nor is a drawn element part of the margin before it exists.

## 16.6 The next measurement (step 7)

The margin's shortfall against the drawing is settled already: the probability that it is under 15 % is 0.9988, and no measurement of the existing plant would bring back the 20 %. What is open is the 10 % question, at a probability of one in three. A further capacity test is one more row of the residual vector, and what it would do to that question is known before its value is.

### 16.6.1 A capacity test at the design day

A test of chiller $i$ at the rating condition, $T_j = 22°\mathrm{C}$, is the row (16.5) with the derate's entry zero,

$$\mathbf{J}_a = \frac{\mathrm{Cap}_i(\hat{\mathbf{z}})}{8}\,\mathbf{e}_{\kappa_i}^{\mathsf{T}}, \qquad \boldsymbol{\Sigma}_a = (\boldsymbol{\Sigma}^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a)^{-1}, \qquad \mathrm{sd}_a(m) = \sqrt{\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_a\,\mathbf{g}}, \tag{16.14}$$

with $\mathbf{e}_{\kappa_i}$ the unit vector of the log-ratio and $\boldsymbol{\Sigma}_a$ the preposterior covariance of equation (5.26); a campaign of one test of each unit stacks three such rows. A test at the rating condition reads

Table 16.8: Capacity tests at the rating condition, 22°C, of standard deviation 8 kW, before the 10 % question; delta form, by the preposterior covariance (16.14) and draws of the shift of the posterior mean (16.15).

| Campaign | $\mathrm{sd}_a(m)$ [–] | Width removed [%] | $P_a$, 5 % [–] | $P_a$, 50 % [–] | $P_a$, 95 % [–] | Settles [–] |
|---|---|---|---|---|---|---|
| None (today) | 0.0139 | 0.0 | – | 0.3316 | – | – |
| One test of C2 | 0.0121 | 12.9 | 0.0779 | 0.3081 | 0.6655 | 0.0825 |
| One test of each, C1 to C3 | 0.0098 | 29.7 | 0.0117 | 0.2671 | 0.8517 | 0.2861 |

the log-ratio alone, without the derate that couples the units and that the record reads low. It narrows the margin by $1 - \mathrm{sd}_a(m)/\mathrm{sd}(m)$. The question asks more, as in Chapter 6: where the probability of a margin under 10 % would land once the test is read. On the delta form the test moves the margin's posterior mean by an amount not known until it is read, normal around today's value with the variance the test removes, by equation (6.11),

$$\hat{m}_a \sim \mathcal{N}\big(m(\hat{\mathbf{z}}),\ \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} - \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}_a\mathbf{g}\big), \qquad P_a = \Phi\left(\frac{0.10 - \hat{m}_a}{\mathrm{sd}_a(m)}\right), \qquad \Pr(P_a > 0.9) + \Pr(P_a < 0.1), \tag{16.15}$$

the last the probability that the test settles the 10 % question at the decision levels 0.9 and 0.1, the form of equation (6.12). It is read by draws of the shift, each pushed through $P_a$.

### 16.6.2 The candidates

Table 16.8 gives two candidates against today's posterior. Each row is a campaign of tests at 22°C of standard deviation 8 kW; the columns give the margin's standard deviation after it by equation (16.14), the share of today's 0.0139 that it removes, the 5, 50 and 95 % points of the probability $P_a$ after it, and the probability that it settles the question, by equation (16.15). The single test is of C2, the unit with the largest gradient entry and, with C3, the largest share of the budget.

One test of C2 takes 12.9 % off the margin's width and settles the 10 % question in 8 % of its outcomes. It removes little because C2's capacity is already known to about $384.5 \times 0.0175 = 6.7$ kW, better than the test's 8 kW, and because C2 carries only 0.369 of the variance. A test of each unit takes 29.7 % off and settles the question in 29 % of outcomes; in the other 71 % the probability lands between 0.1 and 0.9, and the question stays open. The median outcome moves the probability down in both campaigns, to 0.3081 and 0.2671: an outcome that leaves the mean where it is still narrows the band, and with the 10 % level below the center a narrower band puts less probability under it.

### 16.6.3 The recommendation

The review needs no test to restate the margin: under every route the margin is under 15 % with probability above 0.998, and the correction from 20 % stands on the record as it is. A test campaign at the rating condition is worth commissioning only if a decision turns on the 10 % level, and then it should be the campaign of three, one test of each unit on a design-condition day, which reads the log-ratios without the derate and removes the coupling through which the record's error reached all three units. Even so it settles the 10 % question in fewer than one outcome in three. Chapter 18 ranks it against the other tests and surveys a planner can commission on the site, by the width each takes off the quantity that decides the build.

### 16.6.4 What the analysis showed

The drawing's 20 % was a nameplate margin, and the as-built model puts the margin at 10.6 %, with a band of 8.4 to 12.9 %: half the drawn margin is not there. The missing half is the three units' shortfall from their ratings, 94 kW over the three, and no unit's band reaches its nameplate. The margin is a linear function of the three capacities, and its band is theirs, but not their bands added in quadrature: the tests were taken on cooler days and read each unit's ratio through a shared derate, which correlates the three positively and widens the margin's band beyond the units' own. The same coupling explains the error at the truth: the derate was read low, and every unit's design-day capacity, extrapolated from a cooler test, came out high together, 1.30 standard deviations in all. The design day is the thinnest condition of the year, and also the warm edge of the tests' range, where the record knows the capacity less well than in the months the tests resemble. And the drawing's 20 % comes back only with an element that does not exist, at the upper end of a band set by a handbook guess. The margin is the planner's number, and its posterior says both what it is and where the record leaves it open.

### 16.6.5 Limits of the model

**All units running, chillers only.** The margin is the drawing's quantity: the three chillers' capacity with all three running. It says nothing of a unit out, of the outage rates, or of the pipe, pump and feed limits that may bind first; those are the enumeration of Chapter 9, the binding constraints of Chapter 13 and the unit-out headroom of Chapter 20.

**Commissioning tests as today's capacity.** The six tests are the only rows on the log-ratios, and they were taken at commissioning. A unit that has fouled or lost refrigerant since has a lower capacity than the tests say, and no row of the operating record would see it, since the monthly windows read power and flow, not capacity.

**One linear derate.** The three units share one derate, linear in the wet bulb. A unit that derates faster, or a derate that steepens toward the design day, would move the design-day capacities beyond the band, which is an extrapolation from tests at 15 to 22°C.

**The pinned supply and load.** The supply temperature is held at its set point and the design load at 1 000 kW; a different supply moves every unit's capacity, and a different design load moves the margin directly.

**The economizer's guess.** The margin with the economizer rests on a handbook conductance that no reading has tested, and its band is the handbook's assumed spread.

**The constructed record.** The capacity tests and the monthly windows are generated from a stated truth with stated errors, not taken from a site, and the comparison with that truth in Section 16.4.2 is a check of the method on its own model, not of the model against a real plant.

**The next part.** Part IV read where the existing plant's limits lie: the constraint that binds first (Chapter 13), the elements whose capacity matters most (Chapter 14), the shared path (Chapter 15) and, here, the drawing's margin as built. Part V, Chapters 17 to 19, turns to what the uncertainty in the design inputs costs, and Chapter 17 splits the band of the site's cooling headroom among the design parameters, among them the economizer's guess.

## 16.7 What to retain

**Definition 16.1** (As-built capacity margin)**.** The capacity margin of a set of units at a scenario is their as-built capacity there over the design load, less one, $m(\mathbf{z}) = \sum_i N_i \, \mathrm{e}^{\kappa_i}[1 - d\,(T^*_{\mathrm{wb}} - 22)]/L^* - 1$; the drawing's margin is the same expression on the nameplates. Its gradient is each unit's capacity over the load in the log-ratios' columns, and in the derate's column a term that vanishes at the rating condition; its band is the gradient pushed through the covariance of the log-ratios, whose cross terms are positive when the units' tests read them through a shared derate.

Four conclusions carry beyond the site. A drawn margin is a derived quantity of the as-built state, and it is read with a band and with the probability of any level the review cares about; the nameplate value is one point of that state, usually not a likely one. The margin's band is the units' capacities' covariance, not their variances alone: a shared parameter that the tests read the units through correlates them and widens the band, even where that parameter has no entry of its own in the margin. A margin stated at the rating condition is known best near the conditions at which the units were tested, and the design day may lie at the edge of them. And a drawn element enters the margin when it is built and tested; until then it is a scenario, carried with its guess.

Table 16.9 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 16.9: The terms of the as-built margin and where each is defined. Units for the site's chiller plant.

| Term | Expression | Units | Meaning | Defined in |
|---|---|---|---|---|
| Capacity margin | $m = \sum_i \mathrm{Cap}^*_i / L^* - 1$ | – | the drawing's quantity, as built | 16.1.1 |
| Drawing's margin | $m_0 = \sum_i N_i / L^* - 1 = 0.200$ | – | the same on the nameplates | 16.1.1 |
| Gradient of the margin | $\partial m/\partial \kappa_i = \mathrm{Cap}^*_i / L^*$; $\partial m/\partial d = 0$ at 22°C | – | each unit's lever | 16.1.1 |
| Margin with the economizer | $m_e = m + 50\,\mathrm{e}^{g_e}/L^*$ | – | the drawn element counted | 16.1.3 |
| Monthly margin | $m_k = (1+m)[1 + d\,(22 - T_k)] - 1$ | – | a scenario chain over the year | 16.1.3 |
| *Rows and posterior* | | | | |
| Capacity test row | $[N_i \mathrm{e}^{\kappa_i}(1 - d(T_j - 22)) - y_j]/8$ | – | reads a ratio through the derate | 16.1.5 |
| Block structure | eq. (16.7) | – | the tests alone inform the ratios | 16.2.1 |
| Ratios with the derate integrated out | $(\mathbf{D} - \mathbf{a}\mathbf{a}^{\mathsf{T}}/\Lambda_{dd})^{-1}$ | – | positive covariances | 16.2.1 |
| *Questions and answers* | | | | |
| Margin by the draws | quantiles of $m(\mathbf{z}_i)$, $\mathbf{z}_i \sim \mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ | – | the reported margin | 16.3.1 |
| Threshold probability | $N^{-1} \sum_i \mathbf{1}[m_i < f^*]$ | – | margin under a level | 16.3.2 |
| Variance budget | $s_i = g_i(\boldsymbol{\Sigma}\mathbf{g})_i / \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ | – | entries behind the band | 16.3.3 |
| Test at the rating condition | $\mathbf{J}_a = (\mathrm{Cap}_i/8)\,\mathbf{e}^{\mathsf{T}}_{\kappa_i}$ | – | a row with no derate entry | 16.6.1 |
| Settling probability | $\Pr(P_a > 0.9) + \Pr(P_a < 0.1)$ | – | the test decides the level | 16.6.1 |

## 16.8 Exercises

**Exercise 16.1.** *Formulation.* Write the planner's question *is the margin under 15 %?* as a probability over the site's state, say which entries it depends on at the design day, and give its value by the draws, the exact posterior and the delta form. ($\Pr(m < 0.15 \mid \mathbf{y})$ with $m(\mathbf{z}) = 0.4\,(\mathrm{e}^{\kappa_{C1}} + \mathrm{e}^{\kappa_{C2}} + \mathrm{e}^{\kappa_{C3}}) - 1$; at the design day it depends on the three log-ratios alone, the derate's entry vanishing at the rating condition. Its value is 0.9988 by the draws, with a Monte Carlo standard error of 0.0002, 0.9991 by the exact posterior and 0.9992 by the delta form.)

**Exercise 16.2.** *Derivation.* Derive the gradient of the margin (16.1) with respect to the state, show that its derate entry vanishes at the design day, and give C2's entry at the MAP state. (Differentiating $m = \sum_i N_i \mathrm{e}^{\kappa_i}[1 - d\,(T^* - 22)]/L^* - 1$ gives $\partial m/\partial \kappa_i = N_i \mathrm{e}^{\kappa_i}[1 - d\,(T^* - 22)]/L^* = \mathrm{Cap}_i^*/L^*$ and $\partial m/\partial d = -(T^* - 22)\sum_i N_i \mathrm{e}^{\kappa_i}/L^*$, zero at $T^* = 22°\mathrm{C}$; every other entry is zero, since no other entry enters the capacity. C2's entry is its MAP capacity over the load, $384.5/1\,000 = 0.3845$.)

**Exercise 16.3.** *Derivation.* Show that the margin at a month's wet bulb is $m_k = (1 + m)[1 + d\,(22 - T_k)] - 1$, and check July at the MAP state against the draws. (By equation (16.1) each capacity at $T_k$ is its design-day value times $1 - d\,(T_k - 22) = 1 + d\,(22 - T_k)$, a factor common to the three units; the sum of the design-day capacities over the load is $1 + m$. July's wet bulb is 10.17°C, so at $m(\hat{\mathbf{z}}) = 0.1061$ and $d = 0.0152$ the margin is $1.1061 \times (1 + 0.0152 \times 11.83) - 1 = 0.3050$, against the draws' median of 0.3053.)

**Exercise 16.4.** *Application.* A commissioning engineer proposes one capacity test of C2 at 22°C to settle whether the margin is under 10 %. Explain from the preposterior why it rarely does. (The test is a row on $\kappa_{C2}$ alone of standard deviation 8 kW, but C2's capacity is already known to about 6.7 kW, and C2 carries 0.369 of the margin's variance; by equation (16.14) the margin's standard deviation falls only from 0.0139 to 0.0121. The posterior mean can then move only a little, and by equation (16.15) the probability after the test lands between 0.0779 and 0.6655 in nine outcomes of ten; it ends beyond 0.9 or below 0.1 with probability 0.0825.)

**Exercise 16.5.** *Application.* The design review is asked to count the drawn economizer in the margin. Give the margin with it and the probability that it is under 10 %, and say why the entry still does not count it. (By equation (16.3) the margin with the economizer has a median of 0.1577 and the band 0.1246 to 0.2008, skewed to the right by the log-normal guess on its conductance; the drawing's 0.20 is reached only at the band's upper end. The economizer is not built, and its conductance is a handbook prior that no reading has tested, so the entry restates the margin without it. With it, the probability of a margin under 10 % would be 0.0015, and 0.0017 by the exact posterior.)

# Part V

# The Cost of Input Uncertainty

# Chapter 17

# Variance Budgets for Design Parameters

## 17.1 How much of the band one handbook guess carries

This chapter opens Part V. The parts before it asked questions of the plant: how much load it can take, how much to build, where it binds and how much margin it has. Part V asks questions a spreadsheet cannot ask, because they are about the uncertainty of the inputs, not about the inputs. A spreadsheet holds one value per input and returns one number; it has no place for how well an input is known, and so it cannot say which input its answer rests on. This chapter splits a headroom's band among the inputs behind it, Chapter 18 ranks the measurements that would narrow it before a commitment, and Chapter 19 asks what a pilot is worth.

The design sheet for the site's chilled-water plant takes the drawn waterside economizer at its handbook conductance, 25 kW/K, and writes its contribution on the design day as 50.0 kW: the conductance times the 2 K approach the exchanger sees at a wet bulb of 22°C. With the three chillers at their 400 kW nameplates, the sheet's cooling headroom over today's 660 kW of site IT load comes to 568.0 kW. It carries no band. The handbook that supplied the 25 kW/K says its values are right to a factor of about 1.4, and the economizer is not built, so no reading of the plant has ever seen it. The planner's question is the drill question, verbatim,

*The conductance of the economizer path is a handbook guess; how much of the headroom band does that one guess carry?*

It is answered on the site's one fit of Chapter 9, with the binding loads of Chapter 13. What is new is the quantity asked for. It is not a headroom but a split of the headroom's variance, one share per entry of the state, and the split is asked of two headrooms: that of the chillers with the economizer, which is a sum, and that of the site, which is a minimum over four elements. The two give different answers, and the reason is the subject of the chapter.

**Jargon: parameter risk (variance budget), handbook value, handbook factor, cooling headroom, site headroom, conditional mean, conditional variance, law of total variance, first-order share, equal-count binning, conditional normal, quantile points, factory rating test**

The *parameter risk* of a derived quantity, its *variance budget* in the framework, splits the quantity's posterior variance into one share per entry of the state. A *handbook value* is a rating taken from a catalog or a design guide rather than from the plant; its *handbook factor* is the factor within which the guide says it is right, so that a log-ratio to the handbook value has a prior standard deviation of about the logarithm of the factor. The *cooling headroom* $H_c$ is the load the chillers with the drawn economizer can still take over today's site IT load; the *site headroom* $H$ is the load the whole plant can still take, the smallest over its four elements. The *conditional mean* $\mathbb{E}[H \mid z_i]$ of a quantity given one entry is its posterior mean with that entry held at a value, read as a function of the value; its *conditional variance* $\mathrm{Var}(H \mid z_i)$ is the posterior variance with the entry held. The *law of total variance* splits a variance into the mean of the conditional variance and the variance of the conditional mean. The *first-order share* of an entry is the variance of the conditional mean over the variance, the part of the band that knowing the entry alone would remove on average. *Equal-count binning* sorts posterior draws by one entry into groups of equal size and reads the conditional mean as each group's average. The *conditional normal* is a normal distribution of the other entries with one entry held at a value. *Quantile points* of a normal entry are values that split its distribution into slices of equal probability. A *factory rating test* is the manufacturer's test of a selected unit at stated conditions before it ships.

**Key idea**

A variance budget says which inputs a band rests on. For a quantity that is a sum of the entries' effects, each entry's share is the familiar covariance split, and knowing the entry narrows the band by the same amount whatever value it turns out to have. For a quantity that is a minimum, the budget is read by the law of total variance on draws of the posterior: the share of an entry is the variance of the headroom's conditional mean given that entry, and the band with the entry known depends on the value it turns out to have. An input can carry most of one element's band and much less of the plant's, and learning its value can widen the plant's band as well as narrow it.

### 17.1.1 A variance budget read by the law of total variance

Chapter 5 split the variance of a binding load by its gradient, the covariance split of equation (5.25): each entry's diagonal term with half of every cross term it takes part in. That split needs one gradient. A headroom that is the smallest of several binding loads has a gradient that changes with the element that binds, so no single gradient carries its variance, and the split must be read from the posterior itself. The question owes, for each entry $z_i$ of the state, the part of $\mathrm{Var}(H \mid \mathbf{y})$ that entry carries. In this subsection every expectation and variance is under the posterior, the evidence $\mathbf{y}$ on the right of every bar is dropped from the notation, and $\mu = \mathbb{E}[H]$ is the headroom's posterior mean.

Start from the variance as the mean squared deviation from the mean, $\mathrm{Var}(H) = \mathbb{E}[(H-\mu)^2]$. Write $m_i(z_i) = \mathbb{E}[H \mid z_i]$ for the conditional mean, the posterior mean of the headroom over every other entry with $z_i$ held, a function of the value of $z_i$. Adding and subtracting it splits each deviation into two parts,

$$H - \mu = (H - m_i(z_i)) + (m_i(z_i) - \mu),$$

the deviation of the headroom from its conditional mean, and the deviation of the conditional

mean from the overall mean. The square of the sum is the two squares and twice their product, and its expectation is taken in two stages, first over the other entries with $z_i$ held and then over $z_i$, by the law of total expectation of Section 3.6.1,

$$\mathrm{Var}(H) = \mathbb{E}\Big[\mathbb{E}\big[(H - m_i)^2 \mid z_i\big]\Big] + 2\,\mathbb{E}\Big[(m_i - \mu)\,\mathbb{E}\big[H - m_i \mid z_i\big]\Big] + \mathbb{E}\big[(m_i - \mu)^2\big].$$

The middle term vanishes: with $z_i$ held, $m_i - \mu$ is a constant, and the inner expectation $\mathbb{E}[H - m_i \mid z_i]$ is zero by the definition of $m_i$. The first term is the mean of the conditional variance. The third is the variance of the conditional mean, since $\mathbb{E}[m_i] = \mu$ by the same law. So

$$\mathrm{Var}(H) = \mathbb{E}\big[\mathrm{Var}(H \mid z_i)\big] + \mathrm{Var}\big(\mathbb{E}[H \mid z_i]\big), \tag{17.1}$$

the law of total variance of equation (3.32), with an entry of the state in place of the rate. The first term is the variance the headroom keeps on average when $z_i$ is known; the second is the variance that knowing $z_i$ removes on average. The share an entry carries is the second over the whole,

$$S_i = \frac{\mathrm{Var}\big(\mathbb{E}[H \mid z_i]\big)}{\mathrm{Var}(H)} = 1 - \frac{\mathbb{E}\big[\mathrm{Var}(H \mid z_i)\big]}{\mathrm{Var}(H)}, \tag{17.2}$$

the *first-order share* of $z_i$. Its two forms are the two terms of equation (17.1), and each leads to its own computation.

**The variance of the conditional mean.** The first form needs $m_i(z_i)$ at the values $z_i$ takes in the posterior. On $N$ draws $\mathbf{z}_j$ of the posterior with headroom values $H_j$, sort the draws by their entry $z_{i,j}$ and cut them into $B$ bins of equal count; the average $\bar{H}_b$ of the headroom in bin $b$ stands for the conditional mean over that bin's range of $z_i$, and with $\bar{H}$ the average over all draws and $s_H^2$ their variance,

$$S_i \approx \frac{1}{s_H^2} \sum_{b=1}^{B} \frac{n_b}{N} \left(\bar{H}_b - \bar{H}\right)^2, \tag{17.3}$$

with $n_b$ the count of bin $b$. Every entry is read from the same draws, so the binning route gives all eight shares at once.

**The mean of the conditional variance.** The second form needs $\mathrm{Var}(H \mid z_i = v)$ at the values $v$ that $z_i$ takes, averaged over them. Draws of the state with $z_i$ held at $v$ come from the conditional normal, derived in Section 17.2.1; the variance of the headroom over them is $\mathrm{Var}(H \mid z_i = v)$. The average over $v$ is taken on $K$ quantile points of $z_i$'s posterior normal,

$$v_k = \hat{z}_i + \sqrt{\Sigma_{ii}}\;\Phi^{-1}\!\left(\frac{k - \frac{1}{2}}{K}\right), \qquad \mathbb{E}\big[\mathrm{Var}(H \mid z_i)\big] \approx \frac{1}{K}\sum_{k=1}^{K} \mathrm{Var}(H \mid z_i = v_k), \tag{17.4}$$

with $\Phi^{-1}$ the inverse of the standard normal distribution function. Each point stands at the middle, in probability, of a slice of probability $1/K$, so the points take equal weights. The conditioning route reads one entry at a time, and reads more than the share: the conditional variance at each point is itself an answer, the band the headroom would have if $z_i$ turned out to be $v_k$.

**The delta form as a special case.** For a quantity linear in the state near $\hat{\mathbf{z}}$, $H \approx H(\hat{\mathbf{z}}) + \mathbf{g}^{\mathsf{T}}(\mathbf{z} - \hat{\mathbf{z}})$, on the Laplace normal, the conditional mean is linear in $z_i$ too: holding $z_i$ moves each other entry $z_j$ by $\Sigma_{ji}/\Sigma_{ii}$ times the deviation of $z_i$, so $m_i - \mu = (\boldsymbol{\Sigma}\mathbf{g})_i\,(z_i - \hat{z}_i)/\Sigma_{ii}$, and

$$S_i = \frac{(\boldsymbol{\Sigma}\mathbf{g})_i^2}{\Sigma_{ii}\;\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}}, \qquad \mathrm{Var}(H \mid z_i = v) = \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}\,(1 - S_i) \quad \text{for every } v. \tag{17.5}$$

Two facts follow. For an entry uncorrelated with every other, $(\boldsymbol{\Sigma}\mathbf{g})_i = \nabla_i \Sigma_{ii}$ with $\nabla_i$ the gradient's entry, and $S_i = \nabla_i^2 \Sigma_{ii} / \mathbf{g}^\mathsf{T} \boldsymbol{\Sigma} \mathbf{g}$, the entry's term of the covariance split (5.25): on such an entry the two budgets agree. For correlated entries they differ: the first-order share of an entry counts what the others carry through their covariance with it, so an entry with no gradient of its own can have a share, the shares of correlated entries overlap, and they need not sum to one. And on a linear quantity the conditional variance does not depend on the value at which the entry is held: knowing the entry narrows the band by the same amount whatever it turns out to be. A minimum is not linear, and Section 17.3.3 shows what that costs.

**Principle 17.1.** *The variance of a derived quantity is the mean of its conditional variance given an entry plus the variance of its conditional mean. The second term over the whole is the entry's first-order share: read from the posterior draws by binning on the entry, or as one less the mean conditional variance, read by holding the entry at quantile points of its posterior. On a linear quantity and an uncorrelated entry it is the entry's term of the covariance split.*

### 17.1.2 The site and its two headrooms (step 0)

The plant is the site of Chapter 9, with the four elements of Chapter 13 drawn in Figure 13.1: the three existing chillers C1–C3 of 400 kW nameplate with the drawn waterside economizer, the shared supply pipe, the pump and the power feed. Its relations are those of equation (9.3), and its state the eight entries of equation (9.4),

$$\mathbf{z} = \left(\kappa_{C1},\ \kappa_{C2},\ \kappa_{C3},\ d,\ g_e,\ c,\ l_k,\ l_e\right)^\mathsf{T},$$

the chillers' as-built log-ratios, their derate per kelvin of wet bulb, the economizer's log-ratio to its handbook conductance, the log of the chillers' efficiency, and the logs of the pipe's flow coefficient and the pump's efficiency. The entry this chapter is about is $g_e$: the economizer's conductance is $G_e = 25\,\mathrm{e}^{g_e}$ kW/K, and at the design day its duty is $Q_e = 50\,\mathrm{e}^{g_e}$ kW by equation (9.5).

**Modeling choices at the graph.** The choices are those of Chapter 13: steady state at the design day; the chilled-water supply and the loop's rise pinned, so that the loop's flow is proportional to its heat; one efficiency for the pump and one flow coefficient for the pipe; the feed's power factor pinned. One choice is this chapter's own. The economizer's conductance is one log-ratio to the handbook, with a normal prior: this makes the duty lognormal, skewed to the right, and it fails if the selected exchanger's conductance has a distribution the handbook's factor does not describe, such as a fouled exchanger that loses conductance over its life. The prior's width is varied in Section 17.3.4.

### 17.1.3 The scenario node of the two headrooms

The questions are asked at the scenario node of Section 13.1.2, with its three pins: the wet bulb at the design day's 22°C, every unit running, and the plant as drawn, the existing chillers with the drawn economizer. The site IT load is free, and each element's binding load is the load at which its limit is met, by equation (13.1). The headroom is read over today's site IT load of 660 kW.

**The cooling headroom.** The chillers with the economizer bind at the load (13.4), explicit in the state, so the cooling headroom is

$$H_c(\mathbf{z}) = Q_{\mathrm{ch}}^{\dagger}(\mathbf{z}) - 660 = \sum_i 400\,\mathrm{e}^{\kappa_i} + 50\,\mathrm{e}^{g_e} - Q_o - 660, \qquad \frac{\partial H_c}{\partial \kappa_i} = 400\,\mathrm{e}^{\kappa_i}, \qquad \frac{\partial H_c}{\partial g_e} = 50\,\mathrm{e}^{g_e}, \tag{17.6}$$

Table 17.1: The questions asked of the site's two headrooms at the design day, over today's 660 kW of site IT load. $\mathbf{y}$ is the operating record and the priors of Chapter 9; $H_c$ and $H$ are the headrooms of equations (17.6) and (17.7); $S_i$ is the first-order share (17.2).

| Planner's question | General form | Mathematical form | Category | Computation | Answered in |
|---|---|---|---|---|---|
| The conductance of the economizer path is a handbook guess; how much of the headroom band does that one guess carry? | $\mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y})$, split by entry | $g_e$'s share of $\mathrm{Var}(H_c \mid \mathbf{y})$ and of $\mathrm{Var}(H \mid \mathbf{y})$ | estimation | posterior | 17.3.1 |
| What does each entry carry? | $\mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y})$, split by entry | $S_i = \mathrm{Var}(\mathbb{E}[H \mid z_i])/\mathrm{Var}(H)$, all eight entries | estimation | posterior | 17.3.2 |
| What would the band be with the guess pinned down? | $\mathrm{Var}(f(\mathbf{z}) \mid z_i, \mathbf{y})$ | $H$'s band with $g_e$ known; $\mathbb{E}[\mathrm{Var}(H \mid g_e)]$ | estimation | posterior | 17.3.3 |
| How much does a better handbook buy? | $\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$, quantiles | the 90 % bands of $H$ and $H_c$ at handbook factors 1.4, 1.2, 1.1 | estimation | posterior | 17.3.4 |

in kW, with the halls' other heat $Q_o = 22$ kW pinned and every other entry of the gradient zero. It is a sum of one term per unit and one for the economizer, each a function of its own entry. Over the posterior widths of the chillers' log-ratios each term $400\,\mathrm{e}^{\kappa_i}$ is close to linear; the economizer's term $50\,\mathrm{e}^{g_e}$ is an exponential of an entry of wide spread and bends visibly over it. The delta form of $H_c$ is therefore exact to first order, and its covariance split (5.25) is the cooling headroom's budget; the draws stand beside it for the skew.

**The site headroom.** The site binds at the smallest of the four binding loads, equation (13.2), so the site headroom is

$$H(\mathbf{z}) = \min_e Q_e^\dagger(\mathbf{z}) - 660, \qquad \nabla_{\mathbf{z}} H = \nabla_{\mathbf{z}} Q^\dagger_{e^\dagger(\mathbf{z})}, \tag{17.7}$$

with $e^\dagger(\mathbf{z})$ the element that binds at $\mathbf{z}$ and the gradient that element's, wherever one element binds alone. By the pattern (13.8), the economizer's entry $g_e$ sits in the gradients of the chillers and the feed and not in those of the pipe and the pump. The headroom's gradient therefore has an entry in $g_e$ in the states where the chillers bind, and none in the states where the pipe binds; at $\hat{\mathbf{z}}$ it is the pipe's (Section 13.3.1), and a delta form at $\hat{\mathbf{z}}$ would give the guess no share at all. The budget of $H$ is read by the law of total variance, equations (17.3) and (17.4), on draws of the posterior pushed through all four binding loads.

### 17.1.4 The planner's questions (step 1)

The planner asks four questions of the two headrooms, and step 1 writes each over the site's state, with the evidence on the right of the bar. Table 17.1 lists them. The first column is the question in the planner's words, the drill question first; the second the general form of its category in the notation of Chapter 2; the third its mathematical form for the site; 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. All four are estimation: the guess's share of each headroom's variance, every entry's share, the band with the guess known, and the band with the guess known better.

### 17.1.5 The evidence as rows (step 2)

Step 2 attaches the site's operating record of Chapter 9, unchanged: two capacity tests per chiller, twelve monthly windows of chiller power, pipe pressure drop, pump power and feed power, and eight prior rows, 62 rows on the 8 entries (Section 9.1.5). Every unit runs at the scenario node, so the outage record does not enter.

**Which row acts on which entry.** The pattern of the rows is that of Section 13.1.5. The capacity tests read the chillers' log-ratios and the derate; the monthly windows read the efficiency, the pipe and the pump. The economizer is not built, no month ran it, and no test touched it, so its column of the Jacobian holds one nonzero entry, that of its prior row,

$$r_{g_e}(\mathbf{z}) = \frac{g_e - 0}{0.35}, \qquad \frac{\partial r_{g_e}}{\partial g_e} = \frac{1}{0.35},$$

the handbook value at a log-ratio of zero with the prior standard deviation 0.35, a factor of about $e^{0.35}$. No other row has an entry in that column, and that row has no entry in any other. Read as observability, the economizer's conductance is unobservable from the record: whatever the fit knows of it, it knows from the handbook.

**Modeling choices at the rows.** The rows are those of Chapter 9, with its choices: each reading Gaussian, unbiased and independent given the state, and each prior independent. The one that matters here is the economizer's prior. Its center, the handbook value, and its width, the handbook's factor, are the whole of what the fit knows of the economizer, and every answer of this chapter is an answer about that prior as much as about the plant. The center is tested against the stated truth in Section 17.4.3, and the width is varied in Section 17.3.4.

## 17.2 The posterior of the state (step 3)

All four questions of Table 17.1 are read from one posterior, the site's one fit of Chapter 9. No row is added and no fit is repeated. This section derives what the budgets need from that posterior, the normal of the other entries with one entry held, cites the fit, and states the draws the answers use.

### 17.2.1 What the budgets need from the state

The conditioning route (17.4) owes draws of the state with one entry held at a value $v$: the conditional distribution of the other entries, $p(\mathbf{z}_{-i} \mid z_i = v, \mathbf{y})$, with $\mathbf{z}_{-i}$ the state without $z_i$. On the Laplace normal $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of equation (4.38), write the information matrix $\boldsymbol{\Lambda} = \boldsymbol{\Sigma}^{-1}$ in blocks for $z_i$ and the rest, and $\boldsymbol{\delta} = \mathbf{z} - \hat{\mathbf{z}}$ for the deviation from the MAP state. The conditional density is the joint density with $z_i$ held, read as a function of the rest,

$$p(\mathbf{z}_{-i} \mid z_i = v, \mathbf{y}) \propto \exp\Big[-\tfrac{1}{2}\big(\boldsymbol{\delta}_{-i}^{\mathsf{T}}\boldsymbol{\Lambda}_{-i,-i}\,\boldsymbol{\delta}_{-i} + 2\,\boldsymbol{\delta}_{-i}^{\mathsf{T}}\boldsymbol{\Lambda}_{-i,i}\,(v - \hat{z}_i) + \Lambda_{ii}\,(v - \hat{z}_i)^2\big)\Big],$$

with $\boldsymbol{\Lambda}_{-i,-i}$ the block of the other entries, $\boldsymbol{\Lambda}_{-i,i}$ their column against $z_i$ and $\Lambda_{ii}$ the diagonal entry of $z_i$. The last term does not depend on the other entries and goes into the constant. Completing the square in $\boldsymbol{\delta}_{-i}$ leaves a normal with information $\boldsymbol{\Lambda}_{-i,-i}$, centered where the linear term is absorbed, $\boldsymbol{\delta}_{-i} = -\boldsymbol{\Lambda}_{-i,-i}^{-1}\boldsymbol{\Lambda}_{-i,i}\,(v - \hat{z}_i)$. The inverse of a block of $\boldsymbol{\Lambda}$ is the Schur complement of $\boldsymbol{\Sigma}$ of Chapter 3, and the same block identity gives $-\boldsymbol{\Lambda}_{-i,-i}^{-1}\boldsymbol{\Lambda}_{-i,i} = \boldsymbol{\Sigma}_{-i,i}/\Sigma_{ii}$, so

$$\mathbf{z}_{-i} \mid z_i = v, \mathbf{y} \sim \mathcal{N}\left(\hat{\mathbf{z}}_{-i} + \frac{\boldsymbol{\Sigma}_{-i,i}}{\Sigma_{ii}}\,(v - \hat{z}_i),\ \boldsymbol{\Sigma}_{-i,-i} - \frac{\boldsymbol{\Sigma}_{-i,i}\,\boldsymbol{\Sigma}_{i,-i}}{\Sigma_{ii}}\right), \tag{17.8}$$

the conditional normal: holding $z_i$ moves each other entry by its covariance with $z_i$ over $z_i$'s variance, and removes from the other entries' covariance what $z_i$ explained.

**The economizer's entry.** By step 2, no row couples $g_e$ to another entry, so its row and column of $\mathbf{\Lambda}$ hold the prior's $0.35^{-2}$ on the diagonal and zeros elsewhere, the block form of equation (16.7), and so do its row and column of $\mathbf{\Sigma}$,

$$\Sigma_{g_e g_e} = 0.35^2, \qquad \mathbf{\Sigma}_{-g_e,\, g_e} = \mathbf{0}, \qquad \mathbf{z}_{-g_e} \mid g_e = v, \mathbf{y} \sim \mathcal{N}\big(\hat{\mathbf{z}}_{-g_e},\ \mathbf{\Sigma}_{-g_e,-g_e}\big). \tag{17.9}$$

Holding the economizer's entry at any value leaves the other seven at their posterior: conditioning on $g_e$ is drawing the seven from their posterior and setting $g_e = v$. The same structure answers the fourth question. A handbook right to a factor $F$ is the prior $g_e \sim \mathcal{N}(0,\ (\ln F)^2)$, and since no other row sees $g_e$, its posterior is its prior and only one entry of $\mathbf{\Sigma}$ changes,

$$\Sigma_{g_e g_e}(F) = (\ln F)^2, \qquad \text{every other entry of } \mathbf{\Sigma} \text{ unchanged.} \tag{17.10}$$

### 17.2.2 The site's fit, cited

The MAP state $\hat{\mathbf{z}}$ and the Laplace covariance $\mathbf{\Sigma} = (\mathbf{J}^\mathsf{T}\mathbf{J})^{-1}$ at $\hat{\mathbf{z}}$ are those of Section 9.2.2, the Gauss–Newton fit of the operating record, tabulated with the stated truth in Table 9.2; Section 13.2.1 lists the entries the binding loads read. One fact of the fit carries this chapter. The economizer's entry stays where the handbook put it, at 0, and its posterior standard deviation is 0.3500 against a prior standard deviation of 0.3500: the ratio is one, as the block structure (17.9) requires, and no reading of the plant informs it. Every other entry was moved and narrowed by the record.

### 17.2.3 The draws and the conditional draws

The answers read the posterior through the draws of Section 9.2.4: $N = 20\,000$ states from the Laplace normal, with the seed 9, each pushed through the cooling headroom (17.6) and through the four binding loads of Section 13.2.2 for the site headroom (17.7). The exact posterior is the same draws under the weights of equation (5.16), with an effective sample size of 19 913. The binning route (17.3) reads these draws, sorted on each entry into $B = 40$ equal-count bins of 500 draws. The conditioning route (17.4) reads conditional draws from equation (17.9), with the seed 17: 4 000 states at each of $K = 9$ quantile points of $g_e$'s normal, and two further sets of 20 000 with $g_e$ held at the handbook value and at the stated truth's value. The fourth question draws a fresh set of 20 000 from the Laplace normal at each handbook factor, with $\mathbf{\Sigma}$ changed by equation (17.10), on a random stream of its own; the probability that the pipe binds first is read on the first 4 000 of each set.

## 17.3 The answers (step 4)

Every answer is read from the posterior of Section 17.2.2 through the draws of Section 17.2.3. The cooling headroom's budget is the delta form's covariance split, with the draws and the exact posterior beside it for its band; the site headroom's budget is read by the law of total variance on the draws, with the two single-element routes beside it.

### 17.3.1 The guess's share of the two headrooms

**Planner's question.** The conductance of the economizer path is a handbook guess; how much of the headroom band does that one guess carry? The question asks for the economizer entry's share of the variance of the headroom, for the chillers' headroom and for the site's.

Table 17.2: The cooling headroom $H_c$ at the design day, the chillers with the drawn economizer over today's 660 kW, by route. The delta form is $H_c(\hat{\mathbf{z}}) \pm 1.645\,\mathrm{sd}$; the draws are 20 000 states of the Laplace posterior, and the exact posterior is the same draws reweighted; the last row is 20 000 conditional draws with $g_e$ held at the handbook value.

| Route | Location [kW] | 5 % point [kW] | 95 % point [kW] | sd [kW] |
|---|---|---|---|---|
| Design sheet, nameplate and handbook | 568.0 | – | – | – |
| Delta form at $\hat{\mathbf{z}}$ | 474.06 | 437.3 | 510.8 | 22.36 |
| Draws, median (reported band) | 475.67 | 442.58 | 518.75 | – |
| Exact posterior, median | 475.14 | – | – | – |
| Delta form with $g_e$ known | 474.06 | 451.2 | 497.0 | 13.92 |
| Draws with $g_e$ known at 0, median | 474.18 | 451.67 | 496.94 | – |

**General form.** The variance of a derived quantity under the posterior, equation (2.10),

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

split into one share per entry.

**Specialized form.** For the cooling headroom the question is equation (2.10) with $f(\mathbf{z}) = H_c(\mathbf{z})$ of equation (17.6) on the Laplace normal, where the variance is $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$ with the gradient (17.6) at $\hat{\mathbf{z}}$, split by the covariance split (5.25). The economizer's entry is uncorrelated with every other, equation (17.9), so its share is its diagonal term alone,

$$s_{g_e} = \frac{\big(50\,\mathrm{e}^{\hat{g}_e} \times 0.35\big)^2}{\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}}, \qquad \mathrm{sd}(H_c \mid g_e, \mathbf{y}) = \sqrt{\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g} - \big(50\,\mathrm{e}^{\hat{g}_e} \times 0.35\big)^2}, \tag{17.11}$$

the second the standard deviation with the guess known, by equation (17.5), the same at every value of $g_e$. For the site headroom the question is equation (2.10) with $f(\mathbf{z}) = H(\mathbf{z})$ of equation (17.7), and the share is the first-order share $S_{g_e}$ of equation (17.2), read by both of its routes: binning the draws, equation (17.3), and one less the mean conditional variance on nine quantile points, equation (17.4). Two single-element routes stand beside it. The element that binds at $\hat{\mathbf{z}}$, the pipe, has no entry in $g_e$ by the pattern (13.8) and gives the guess a share of zero; the chillers alone are the cooling headroom and give it $s_{g_e}$.

**Building blocks.** The covariance $\boldsymbol{\Sigma}$ and its block structure (17.9); the gradient (17.6); the split (5.25); the draws and their weights; the binding loads of Chapter 13 in each draw; the two routes (17.3) and (17.4).

**The cooling headroom.** Table 17.2 gives the cooling headroom by each route: each row is a route, the columns its location, the ends of its 90 % band and its standard deviation. Table 17.3 gives its budget: each row is an entry of the state, the columns its gradient entry at $\hat{\mathbf{z}}$ and its share by equation (5.25).

The cooling headroom is 474.06 kW at the MAP state, with a delta-form standard deviation of 22.36 kW; the draws put its median at 475.67 kW and its 90 % band at $[442.58,\ 518.75]$ kW, and the exact posterior its median at 475.14 kW. The sheet's 568.0 kW lies far above the band: the chillers deliver less than their nameplates, the shortfall of Chapter 16. The draws' band sits to the right of the delta form's, $[437.3,\ 510.8]$ kW, and their median above the MAP value: the

Table 17.3: The variance budget of the cooling headroom $H_c$ at the MAP state, equation (5.25), in shares of $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} = 22.36^2\,\mathrm{kW}^2$; the gradient entries by equation (17.6).

| Entry [unit] | $\partial H_c/\partial z_i$ at $\hat{\mathbf{z}}$ [kW] | Share $s_i$ [–] |
|---|---|---|
| $g_e$, the economizer's log-ratio [–] | 50.00 | 0.613 |
| $\kappa_{C3}$ [–] | 356.28 | 0.144 |
| $\kappa_{C2}$ [–] | 384.54 | 0.143 |
| $\kappa_{C1}$ [–] | 365.24 | 0.101 |
| $d$, $c$, $l_k$, $l_e$ [$\mathrm{K}^{-1}$, –, –, –] | 0 | 0.000 |

economizer's duty is an exponential of its entry, skewed to the right, as Section 13.3.1 found for the chillers' binding load.

The economizer's guess carries 0.613 of the variance, the three chillers' log-ratios 0.101, 0.143 and 0.144, and every other entry nothing: none enters $H_c$, and the covariance split gives an entry with no gradient entry no share, whatever its correlations. The guess's term is its gradient entry times its spread, 50.00 kW per unit of $g_e$ times 0.350, or 17.50 kW of the 22.36 kW standard deviation. A share is a share of variance, not of standard deviation: the guess's 17.50 kW is 78 % of the standard deviation and 61 % of the variance, and with the guess known the remaining standard deviation is 13.92 kW, not $22.36 - 17.50$. The half-width of the delta band falls from 36.8 to 22.9 kW. The draws with the guess held at the handbook value agree: their band is [451.67, 496.94] kW, 45.27 kW wide against today's 76.17 kW, beside the delta form's [451.2, 497.0] kW, since with $g_e$ fixed its skew goes too. C1 has the smallest share of the three units, although its gradient entry lies between the other two: one of its capacity tests was taken at the rating condition, and its log-ratio is known best (Section 16.2.1).

**The site headroom.** The site headroom, the smallest of the four binding loads less 660 kW, has by the draws a median of 463.88 kW, a 90 % band of [442.58, 473.97] kW and a standard deviation of 9.62 kW; the exact posterior gives the median 463.82 kW and the band [442.07, 473.99] kW. Its band is less than half as wide as the cooling headroom's and shares its lower end, to the hundredth of a kilowatt: the lower tail of the minimum is the chillers' lower tail, and the upper tail is cut by the pipe, as Section 13.3.1 found for the binding loads. On this band the economizer's first-order share is

$$\begin{aligned} S_{g_e} &= 0.305 && \text{by binning, equation (17.3),} \\ S_{g_e} &= 1 - 0.677 = 0.323 && \text{by conditioning, equation (17.4),} \end{aligned} \tag{17.12}$$

two routes to one quantity, compared in Section 17.4.2. The two single-element routes bracket them: the pipe alone gives the guess no share, the chillers alone 0.613.

The guess carries 61 % of the chillers' headroom band but only about a third of the site's, the first surprise of the chapter. The pipe binds first in 0.6505 of the 20 000 draws, and whenever it binds, the economizer does not matter: the site headroom is the pipe's, and the pipe's binding load has no entry in $g_e$. In the draws where the chillers bind, the guess moves the headroom with its full lever of $50\,\mathrm{e}^{g_e}$ kW per unit; in the draws where the pipe binds, not at all. The minimum dilutes the guess. The planner who reads the chillers' budget and concludes that the site's band is mostly the economizer has read the budget of the wrong quantity.

### 17.3.2 What each entry carries

**Planner's question.** What does each entry carry? The question asks for the first-order share of every entry of the state in the site headroom's variance.

Table 17.4: The variance budgets of the two headrooms at the design day: the cooling headroom's by the covariance split (5.25) at $\hat{\mathbf{z}}$, and the site headroom's first-order shares (17.2) by binning the 20 000 draws, equation (17.3).

| Entry [unit] | Enters the binding load of [–] | Share of Var($H_c$) [–] | $S_i$, share of Var($H$) [–] |
|---|---|---|---|
| $\kappa_{C1}$ [–] | chillers | 0.101 | 0.106 |
| $\kappa_{C2}$ [–] | chillers | 0.143 | 0.169 |
| $\kappa_{C3}$ [–] | chillers | 0.144 | 0.172 |
| $d$ [K$^{-1}$] | none at the design day | 0.000 | 0.165 |
| $g_e$ [–] | chillers, feed | 0.613 | 0.305 |
| $c$ [–] | feed | 0.000 | 0.001 |
| $l_k$ [–] | pipe, pump, feed | 0.000 | 0.139 |
| $l_e$ [–] | pump, feed | 0.000 | 0.007 |
| Sum | | 1 | 1.065 |

**General form.** The variance of a derived quantity under the posterior, equation (2.10), split into one share per entry.

**Specialized form.** The question is equation (2.10) with $f(\mathbf{z}) = H(\mathbf{z})$ of equation (17.7), split by the first-order share (17.2) of each of the eight entries, read by the binning route (17.3) on the 20 000 draws. Beside each stands the entry's share of the cooling headroom by the covariance split. By equation (17.5), the first-order shares of correlated entries overlap, and their sum is not bound to one.

**Building blocks.** The draws and the site headroom in each (Section 17.2.3); the binning route (17.3); the cooling budget of Table 17.3.

**The shares.** Table 17.4 gives them: each row is an entry, the columns the binding loads it enters by the pattern (13.8), its share of the cooling headroom and its first-order share of the site headroom. The two budgets are drawn side by side in Figure 17.1, one pair of bars per entry, the horizontal axis the share of the variance.

The bars tell two stories. In the cooling headroom one entry dominates and four have nothing. In the site headroom the guess still has the largest share, 0.305, but six entries have a share above a tenth, and the shares sum to 1.065. The excess over one is the overlap of correlated entries, which equation (17.5) predicts: the entries are correlated in the posterior, so holding one also moves others, and the part of the variance two entries carry together is counted in both shares. The derate is the plainest case, and both facts about it are stated plainly. It has no direct effect at the design day: its entry in every binding load's gradient is zero there, equation (13.8). It has a share of 0.165 nonetheless, through its correlation with the chillers' log-ratios. The capacity tests read each log-ratio through the derate, so the log-ratios are correlated with it and with each other (Section 16.2.1), and the derate's conditional mean of the headroom moves with the log-ratios it carries. Its share is not a lever the derate has, but a part of the log-ratios' band seen from the derate's side.

The chillers' three log-ratios together carry more of the site's band than of the cooling headroom's, each a little more than its cooling share, for the same reason: their correlations with each other and with the derate add to each first-order share. The pipe's flow coefficient $l_k$ carries 0.139 of the site's band and nothing of the chillers': it decides the headroom wherever the pipe binds. The efficiency $c$ and the pump's $l_e$, which enter only the feed and the pump, two elements that bind first in fewer than one draw in ten thousand (Section 13.3.2), have shares of

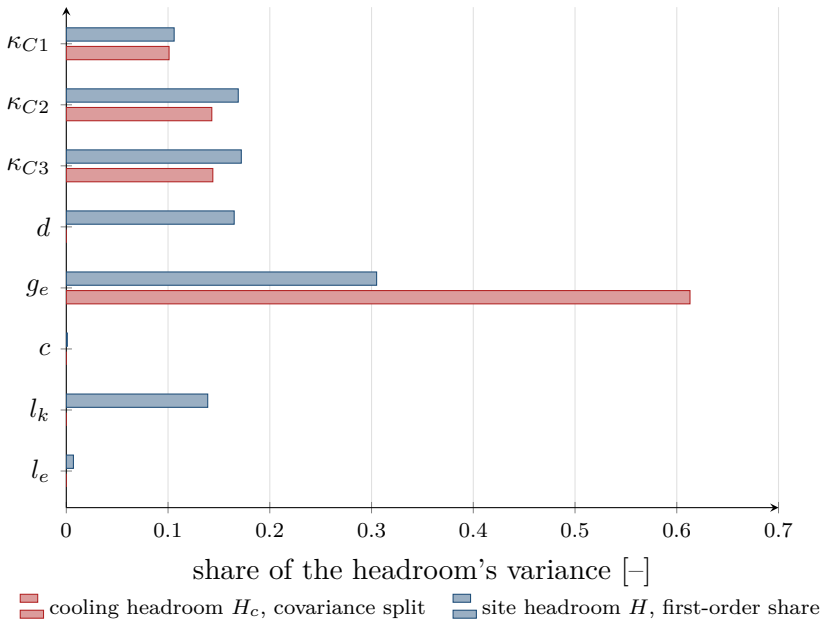



Figure 17.1: The variance budgets of the two headrooms at the design day, from Table 17.4. Red: the cooling headroom's shares by the covariance split at $\hat{\mathbf{z}}$, which sum to one. Blue: the site headroom's first-order shares by binning the posterior draws, which overlap and sum to 1.065.

0.001 and 0.007, at the level of what binning resolves and of what their correlation with $l_k$ in the monthly rows carries.

### 17.3.3 The band with the guess pinned down

**Planner's question.** What would the band be with the guess pinned down? The question asks for the site headroom's band if the economizer's conductance were known, at the value it might turn out to have, and on average over those values.

**General form.** The variance of a derived quantity under the posterior with one entry held, equation (2.10) applied to the conditional posterior,

$$\operatorname{Var}(f(\mathbf{z}) \mid z_i = v, \mathbf{y}) = \int \left(f(\mathbf{z}) - \mathbb{E}[f(\mathbf{z}) \mid z_i = v, \mathbf{y}]\right)^2 p(\mathbf{z} \mid z_i = v, \mathbf{y}) \, \mathrm{d}\mathbf{z}_{-i},$$

and its average over the values $v$ that $z_i$ takes.

**Specialized form.** The question is equation (2.10) with $f(\mathbf{z}) = H(\mathbf{z})$ of equation (17.7), $z_i = g_e$, and $p(\mathbf{z} \mid g_e = v, \mathbf{y})$ the conditional normal (17.9): the seven other entries at their posterior, the economizer's held. It is read at three kinds of value: the handbook's, $v = 0$, the value the review would assume if told the handbook was right; the stated truth's, $v = -0.2231$, the value a test would find; and the nine quantile points of equation (17.4), $v_k = 0.35\,\Phi^{-1}((2k-1)/18)$, from $-0.5576$ to $0.5576$, whose average conditional variance is the first term of the law of total variance (17.1). Each conditional variance is reported against today's, $\operatorname{Var}(H \mid g_e = v)/\operatorname{Var}(H)$.

**Building blocks.** The conditional normal (17.9); the conditional draws of Section 17.2.3, each pushed through the four binding loads; the quantile points (17.4).

Table 17.5: The site headroom $H$ with the economizer's entry known, from the conditional normal (17.9). Upper block: $g_e$ held at the handbook value and at the stated truth's value, beside today's posterior. Lower block: the nine quantile points of equation (17.4), 4 000 conditional draws each.

| $g_e$ held at [–] | Median [kW] | 5 % point [kW] | 95 % point [kW] | sd [kW] | $\mathrm{Var}(H \mid g_e)/\mathrm{Var}(H)$ [–] |
|---|---|---|---|---|---|
| not held (today) | 463.88 | 442.58 | 473.97 | 9.62 | 1 |
| the handbook value, 0.0000 | 464.67 | 451.40 | 473.72 | 6.86 | 0.509 |
| the stated truth, −0.2231 | 461.12 | 441.64 | 471.92 | 9.33 | 0.940 |
| *Quantile points, level* $(2k-1)/18$ *of* $g_e$*'s normal* | | | | | |
| $k = 1$, level 1/18, $v_1 = -0.5576$ | | | | | 1.534 |
| $k = 2$, level 3/18, $v_2 = -0.3386$ | | | | | 1.153 |
| $k = 3$, level 5/18, $v_3 = -0.2063$ | | | | | 0.898 |
| $k = 4$, level 7/18, $v_4 = -0.0988$ | | | | | 0.685 |
| $k = 5$, level 9/18, $v_5 = 0.0000$, the handbook value | | | | | 0.505 |
| $k = 6$, level 11/18, $v_6 = 0.0988$ | | | | | 0.401 |
| $k = 7$, level 13/18, $v_7 = 0.2063$ | | | | | 0.329 |
| $k = 8$, level 15/18, $v_8 = 0.3386$ | | | | | 0.287 |
| $k = 9$, level 17/18, $v_9 = 0.5576$ | | | | | 0.298 |
| Average over the points, $\mathbb{E}[\mathrm{Var}(H \mid g_e)]/\mathrm{Var}(H)$ | | 0.677; the share by conditioning $1 - 0.677 = 0.323$ | | | |

**The band with the guess known.** Table 17.5 gives the results: in its upper block each row is a value at which the guess is held, the columns the site headroom's median, the ends of its 90 % band, its standard deviation and its variance against today's; in its lower block each row is a quantile point, the column its variance ratio. The ratios at the nine points are drawn in Figure 17.2: the horizontal axis is the probability level of the point in $g_e$'s normal, low conductances on the left; the vertical axis the variance with $g_e$ known over today's; the solid horizontal line marks one, today's band, the dashed line the average over the points, and the dotted line the cooling headroom's ratio, the same at every value.

With the guess known at the handbook value the site headroom's variance falls to 50.9 % of today's, its standard deviation from 9.62 to 6.86 kW, and its band from [442.58, 473.97] to [451.40, 473.72] kW: the lower end comes up by 8.8 kW, and the upper end barely moves. Known at the stated truth's value, the variance falls only to 94.0 % of today's, and the band moves down to [441.64, 471.92] kW. The value the guess turns out to have decides how much knowing it is worth.

Across the nine points the ratio runs from 1.534 at a low conductance down to 0.29 at the high ones, and the two points furthest to the left lie above one. Here is the second surprise of the chapter: learning that the guess is low can widen the band. A low economizer pulls the chillers' binding load down, below the pipe's, and the chillers bind in nearly every draw; the site headroom then inherits the chillers' band with the guess known, a standard deviation of 13.92 kW by Table 17.2, whose variance is $(13.92/9.62)^2 = 2.09$ times the site's today. The chillers' band is wider than the pipe's, and today's site band is narrow only because the pipe cuts it wherever the chillers would be high. Learning that the economizer is weak removes that cut. At the high points the pipe binds in nearly every draw, and the ratio settles near 0.29, the pipe's own narrow band, Table 13.2; the last two points, 0.287 and 0.298, agree within the noise of 4 000 draws.

On average over the points the variance with the guess known is 0.677 of today's, so knowing the guess removes 0.323 of the site headroom's variance on average. The average hides the spread: the cooling headroom, a sum, loses 0.613 of its variance at every value of the guess, the

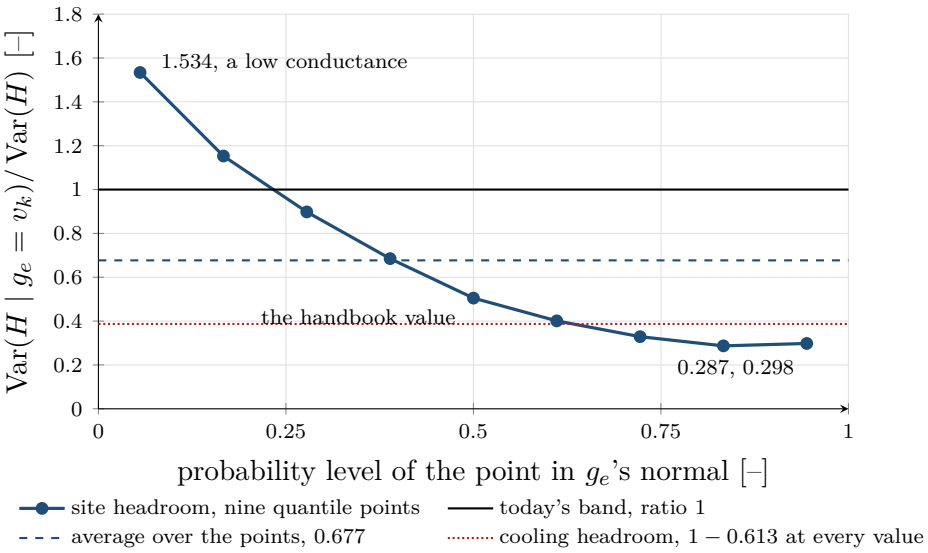



Figure 17.2: The site headroom's variance with the economizer's entry held at each of the nine quantile points of its normal, over today's variance, from Table 17.5. Left: low conductances, where the chillers bind; right: high conductances, where the pipe binds. The dotted line is the cooling headroom's ratio by the delta form, the same at every value because $H_c$ is a sum.

dotted line of the figure, while the site headroom loses anything from nothing to 0.71 of it, or gains up to half again.

### 17.3.4 What a better handbook buys

**Planner's question.** How much does a better handbook buy? The question asks how the bands of the two headrooms narrow if the economizer's conductance were known to a tighter factor before the design is committed, with its center still at the handbook value.

**General form.** The posterior distribution of a derived quantity, equation (2.9), summarized by its central credible interval, equation (2.11), here the central 90 % interval.

**Specialized form.** The question is equation (2.11) with $f(\mathbf{z}) = H(\mathbf{z})$ and $f(\mathbf{z}) = H_c(\mathbf{z})$, under the Laplace normal with the economizer's variance set by equation (17.10) for the handbook factors $F = 1.4$, 1.2 and 1.1, the prior standard deviations $\ln F = 0.336$, 0.182 and 0.095. Each band is read as the 5 % and 95 % points of a fresh set of draws pushed through the four binding loads, and with it the probability that the pipe binds first, the share of the draws in which the pipe's binding load is the smallest, by equation (13.11).

**Building blocks.** The covariance $\boldsymbol{\Sigma}$ with one entry changed, equation (17.10); the draws at each factor; the binding loads of Chapter 13.

**The bands against the factor.** Table 17.6 gives them: each row is a handbook factor, the columns the economizer's prior standard deviation, the site headroom's band and its width, the cooling headroom's width and the probability that the pipe binds first. Today's prior is the first row, and the row with the guess known at the handbook value, from Table 17.5, is the last. The widths are drawn in Figure 17.3: the horizontal axis is the economizer's prior standard deviation, the vertical axis the width of the 90 % band, one line per headroom.

Table 17.6: The 90 % bands of the two headrooms at the design day as the handbook's factor narrows, with the center at the handbook value. Today's row is the 20 000 draws of the site's posterior, each factor's row a fresh set of 20 000 (Pr(pipe first) on its first 4 000), and the known row the 20 000 conditional draws with $g_e$ held at the handbook value. The site headroom's width in today's row and in the known row is the difference of the printed ends.

| Handbook factor [–] | sd of $g_e$ [–] | $H$, 5 % [kW] | $H$, 95 % [kW] | $H$ width [kW] | $H_c$ width [kW] | Pr(pipe first) [–] |
|---|---|---|---|---|---|---|
| today's prior | 0.350 | 442.58 | 473.97 | 31.39 | 76.17 | 0.6505 |
| 1.4 | 0.336 | 442.84 | 473.70 | 30.86 | 73.78 | 0.6340 |
| 1.2 | 0.182 | 448.23 | 473.89 | 25.66 | 54.91 | 0.6763 |
| 1.1 | 0.095 | 450.35 | 473.76 | 23.40 | 48.23 | 0.6970 |
| known at the handbook value | 0 | 451.40 | 473.72 | 22.32 | 45.27 | – |

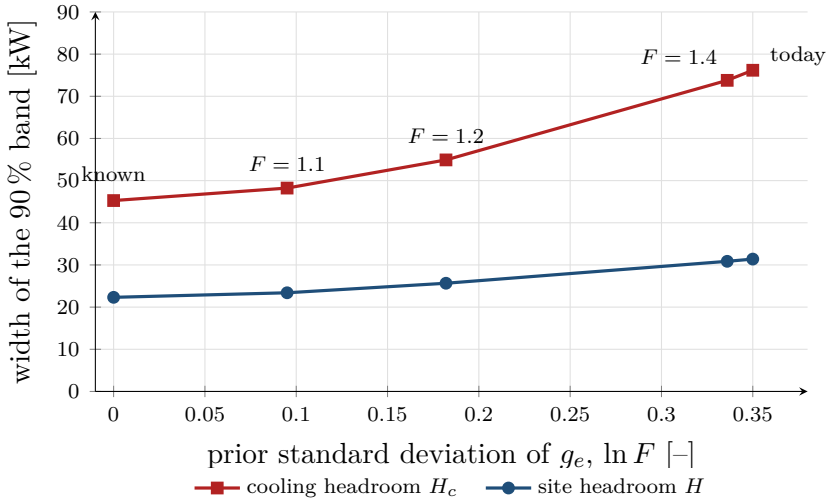



Figure 17.3: The width of the 90 % band of each headroom against the economizer's prior standard deviation, from Table 17.6. The points at zero are the bands with the guess known at the handbook value, read on the conditional draws; the points at 0.350 are today's.

The two lines fall at different rates. The cooling headroom's band narrows from 76.17 kW today to 54.91 kW at a factor of 1.2 and 48.23 kW at 1.1, close to the 45.27 kW it has with the guess known. The site headroom's band narrows from 31.39 to 25.66 and 23.40 kW, within 1.1 kW of the 22.32 kW it would have with the guess known at the handbook value. Every change is at the lower end: the site band's 5 % point rises from 442.58 to 450.35 kW while its 95 % point stays between 473.70 and 473.97 kW, the pipe's. The guess lives in the lower tail, where the chillers bind. As the guess narrows, the pipe binds first more surely, from 0.6505 today to 0.6763 at a factor of 1.2 and 0.6970 at 1.1: at the handbook value the chillers' binding load lies above the pipe's (Section 13.3.1), and a narrower guess leaves fewer draws in which a weak economizer pulls the chillers under the pipe. The factor of 1.4, at 0.6340, sits below today's although its prior is slightly narrower; it is read on 4 000 draws, whose standard error is about 0.008, so the gap of 0.017 is about two standard errors, and the step from a prior standard deviation of 0.350 to 0.336 is too small to show through that noise. A handbook right to a factor of 1.2 buys most of what knowing the conductance would, on these terms; the terms are that the handbook's center is right, which Section 17.4.3 tests.

## 17.4 The checks (step 5)

The answers of Section 17.3 rest on the Laplace normal of the site's fit, on two numerical routes to one share, and on a record generated from a stated truth; each is tested here.

### 17.4.1 The draws against the exact posterior

The reported bands are quantiles of 20 000 Laplace draws; the exact posterior reweights the same draws, with an effective sample size of 19 913. The site headroom's median moves from 463.88 to 463.82 kW, its 5 % point from 442.58 to 442.07 kW and its 95 % point from 473.97 to 473.99 kW; the cooling headroom's median from 475.67 to 475.14 kW. Every gap is under 0.6 kW, and the largest sit where the chillers set the headroom, the cooling headroom's median and the site's lower end, as in Section 13.4.2. The economizer's entry is drawn from its exact normal prior, so none of the gap is the guess's. The posterior standard deviation of $g_e$ is 0.3500 against its prior's 0.3500: the record is silent on it, as the block structure (17.9) says it must be.

### 17.4.2 The two routes to the guess's share

The law of total variance (17.1) makes the binning route and the conditioning route two readings of one quantity, since $g_e$ is uncorrelated with every other entry and the two forms of equation (17.2) are equal. They read 0.305 and 0.323, a gap of 0.018. Each has its own error. Binning replaces the conditional mean by one value per bin, and averaging within a bin can only lower the variance of the conditional mean; the conditioning route reads a nine-point rule, each point's variance from 4 000 draws. Both put the guess at about a third of the site's variance, against 0.613 of the chillers'. A second agreement checks the conditional draws. The fifth quantile point, at the level 9/18, is the handbook value itself, and its 4 000 draws give a ratio of 0.505; the separate set held at the handbook value gives 0.509, the 50.9 % of Table 17.5.

The cooling headroom's routes differ in the band, not in the budget. The delta form's band, $[437.3,\ 510.8]$ kW, lies 5 kW and 8 kW to the left of the draws' at its two ends, the skew of $50\,\mathrm{e}^{g_e}$; its budget is the variance split at $\hat{\mathbf{z}}$, which the delta form owns.

### 17.4.3 The answers against the stated truth

The operating record is generated from a stated truth, and the last check compares the answers with it. Table 17.7 collects the checks: each row is a quantity, the columns its reported value, the value beside it, the truth and the reading. At the truth the cooling headroom is 446.00 kW

Table 17.7: The checks of the variance-budget answers. Reported values from the 20 000 draws; beside them the exact posterior or the route named; the truth is the stated truth of the record.

| Quantity | Reported [kW or –] | Beside [kW or –] | Truth [kW or –] | Reading |
|---|---|---|---|---|
| $H_c$, median | 475.67 | exact 475.14 | 446.00 | inside [442.58, 518.75], near the lower end |
| $H_c$ at $\hat{\mathbf{z}}$, delta form | 474.06 | sd 22.36 | 446.00 | 1.25 sd above the truth |
| $H$, median | 463.88 | exact 463.82 | 446.00 | inside [442.58, 473.97], near the lower end |
| $H$, 5 % and 95 % points | 442.58, 473.97 | exact 442.07, 473.99 | – | gaps under 0.6 kW |
| $g_e$, posterior sd | 0.3500 | prior 0.3500 | −0.2231 | the record is silent on it |
| $g_e$'s share of Var($H$) | 0.305 | 0.323 | – | binning; conditioning |
| Var($H \mid g_e = 0$)/ Var($H$) | 0.509 | 0.505 | – | separate set; fifth point |
| Var($H \mid g_e =$ −0.2231)/ Var($H$) | 0.940 | – | – | the band barely narrows |
| Design sheet, nameplate and handbook | 568.0 | – | 446.00 | outside every band |
| Draws | $N = 20\,000$, seed 9, effective sample size 19 913; conditional draws seed 17 | | | |

and the site headroom 446.00 kW: with the truth's low economizer, $g_e = -0.2231$, the chillers bind. Both lie inside their bands, near the lower ends, which the chillers set. The true duty on the design day is $50\,\mathrm{e}^{-0.2231} = 40.0$ kW, 10.0 kW below the handbook's. The MAP cooling headroom lies $474.06 - 446.00 = 28.06$ kW above the truth, 1.25 delta-form standard deviations: 10.0 kW of it is the guess, and the rest the chillers' as-built capacities, which the record reads a little high (Section 16.4.2). The truth's $g_e$ lies $0.2231/0.35 = 0.64$ prior standard deviations below the handbook value, and held there, the site headroom keeps 94.0 % of its variance: the true economizer sits where the chillers and the pipe contest the first place, and knowing it leaves both in play.

## 17.5 The report (step 6)

The handbook guess on the economizer's conductance carries 61 % of the variance of the chiller plant's headroom and about a third of the site's, 0.305 by binning and 0.323 by conditioning. Pinning the guess down narrows the site's band on average, to 0.677 of its variance, but a low reading can widen it, to 1.534 of it.

### 17.5.1 The answers and their bands

Table 17.8 gives the answers in the order of Table 17.1. Each row is a question; the second column gives the reported answer with the exact posterior or the second route beside it, and the third its band or the values beside it.

### 17.5.2 The design-review entry

The design review's entry names what the headroom's band rests on. On the design day, with every unit up and the drawn economizer, the site can take a median of 463.9 kW more over today's 660 kW, band 442.6 to 474.0 kW, and the economizer's handbook conductance carries about a third of that band's variance. The chiller plant's own headroom, the number the sheet writes as 568.0 kW, is 475.7 kW at the median with a band of 442.6 to 518.8 kW, and the same guess carries 61 % of its variance. The entry records that the two shares differ because the site binds at the pipe in most of the posterior, and that knowing the guess does not always narrow the site's band: if the exchanger proves weak, the chillers bind and the band widens.

Table 17.8: The report on the economizer's guess, at the design day with all units up and the drawn economizer, over today's 660 kW of site IT load. Reported values by the site's posterior and its draws.

| Question | Answer | Band or values beside |
|---|---|---|
| The conductance of the economizer path is a handbook guess; how much of the headroom band does that one guess carry? | chillers' headroom: 0.613 of the variance; site headroom: 0.305 (by conditioning: 0.323) | $H_c$ median 475.67 kW (exact: 475.14), band 442.58 to 518.75 kW; $H$ median 463.88 kW (exact: 463.82), band 442.58 to 473.97 kW, central 90 % |
| What does each entry carry? | $g_e$ 0.305; $\kappa_{C1}$, $\kappa_{C2}$, $\kappa_{C3}$ 0.106, 0.169, 0.172; $d$ 0.165; $l_k$ 0.139 | $c$ 0.001, $l_e$ 0.007; the shares sum to 1.065, overlapping through the posterior correlations |
| What would the band be with the guess pinned down? | on average 0.677 of today's variance; at the handbook value 0.509, sd 6.86 kW | at the truth's value 0.940; across the quantile points 0.29 to 1.534 |
| How much does a better handbook buy? | site band 31.39 kW wide today; 30.86, 25.66, 23.40 kW at factors 1.4, 1.2, 1.1; 22.32 kW with the guess known | cooling band 76.17 today; 73.78, 54.91, 48.23 kW; 45.27 kW known; Pr(pipe first) 0.6505 today; 0.6340, 0.6763, 0.6970 |

Its decision follows. The economizer is carried as uncertain, at the handbook value with the handbook's factor, until it is measured; no headroom that depends on the chillers binding, and no build that would use that headroom, is quoted without its band. Chapter 18 ranks the measurement against the others the review can buy.

**Common mistake**

Quoting one element's variance budget as the plant's. The chillers' budget says the guess carries 61 %; the site's headroom is a minimum, and the guess carries its share only in the states where the chillers bind. A budget belongs to the quantity it was computed for.

## 17.6 The next measurement (step 7)

The report leaves a third of the site's band, and most of the chillers', on a number no reading of the plant contains. A measurement of the economizer cannot be taken on the plant, since it is not built; it can be taken on the exchanger before it is installed. This step names that measurement and states its ceiling; Chapter 18 computes its value.

### 17.6.1 The factory test of the exchanger

The manufacturer's rating test of the selected exchanger, at stated flows and temperatures, reads its conductance, and with it $g_e$. It is one row on that entry alone, with $\sigma_a$ the test's standard deviation on the log-ratio,

$$\mathbf{J}_a = \frac{1}{\sigma_a}\,\mathbf{e}_{g_e}^{\mathsf{T}}, \qquad \Sigma_{a,\,g_e g_e} = \left(0.35^{-2} + \sigma_a^{-2}\right)^{-1},$$

the preposterior of equation (5.26), with $\mathbf{e}_{g_e}$ the unit vector of the economizer's entry. By the block structure (17.9) the row changes that one entry of $\mathbf{\Sigma}$ and no other, and moves $g_e$'s posterior mean to the test's reading. Its ceiling is a test so precise that it fixes $g_e$, the conditioning of Section 17.3.3. For the cooling headroom that ceiling is known in advance: the standard deviation falls from 22.36 to 13.92 kW, a width removed of $1 - 13.92/22.36 = 37.7$ %, whatever

the test reads. For the site headroom it is not: it removes 0.323 of the variance on average, and the reading decides the rest, from about half of today's variance at the handbook value to more than today's at a low reading. The factor rows of Table 17.6 are a test of finite precision whose reading happens to equal the handbook value.

The test does not touch the chillers' log-ratios, the derate or the pipe's coefficient, which together carry the other two thirds of the site's band (Table 17.4). Chapter 18 sets the factory test beside the capacity tests of the existing chillers and a high-flow pressure-drop test of the pipe, and ranks them by how often each settles the build decision, with the width each takes off the deciding quantity beside; Chapter 19 asks what a pilot of the economizer is worth.

### 17.6.2 What the analysis showed

The design sheet wrote the economizer as 50 kW and the cooling headroom as 568 kW, with no band. The analysis puts the chiller plant's headroom at a median of 476 kW with a band of 443 to 519 kW, and the site's at 464 kW with a band of 443 to 474 kW, and asks what the bands rest on. Three findings carry the answer. The guess carries 61 % of the chillers' band, a single handbook number outweighing the six capacity tests that fix the existing units. It carries about a third of the site's, because the site's headroom is a minimum, and wherever the pipe binds the economizer does not matter: a budget of one element is not the budget of the plant. And knowing the guess is not the same as narrowing the band. The chillers' headroom narrows by the same amount whatever the exchanger turns out to be; the site's narrows to half if the handbook is right and widens if the exchanger is weak, because a weak exchanger makes the chillers bind and the chillers' band is wider than the pipe's. The law of total variance says why: the average conditional variance is fixed by the share, and the conditional variance at each value is free to lie above it or below. A budget on a minimum is read on the draws, and the value of a measurement on a minimum depends on what the measurement may read.

### 17.6.3 Limits of the model

**The design day with all units up.** The two headrooms are read at the design day's wet bulb with every chiller running and the economizer at its design-day approach. On a cooler day the economizer's duty grows with its approach and the guess's share with it; with a unit out the chillers bind in every state, and the site's budget becomes the chillers'.

**One log-ratio for the exchanger.** The economizer is one conductance with a lognormal prior, constant over its life and over the flows it sees. A fouling exchanger loses conductance over time; an exchanger whose conductance depends on the tower-water flow has more than one entry. Either would change the guess's prior and, through it, every share above.

**Shares of variance, not of a decision.** The budget splits the headroom's variance, not the risk of a decision taken on it. An entry with a small share can still move a probability near a threshold, and an entry with a large share can leave a decision unchanged; which entries change a build decision is the question of Chapters 18 and 19.

**Numerical routes.** The first-order shares are read by binning on the draws, and the mean conditional variance by a nine-point rule with 4 000 draws at each point; the two routes to the guess's share differ by 0.018. The conditional draws use the Laplace normal; for the economizer's entry, which has no row but its prior, the conditional normal is exact, and for the others it carries the small departure of Section 17.4.1.

**The constructed record.** The operating record and the capacity tests are generated from a stated truth with stated reading errors, not taken from a site, and the comparison with that truth in Section 17.4.3 is a check of the method on its own model, not of the model against a real plant.

**The next chapter.** Chapter 18 ranks the tests and surveys the review can commission before the build, the factory test of the exchanger among them, by how often each settles the build decision.

## 17.7 What to retain

**Definition 17.1** (First-order share and the band with an entry known)**.** The variance of a derived quantity under the posterior splits, for any entry $z_i$, as $\mathrm{Var}(H) = \mathbb{E}[\mathrm{Var}(H \mid z_i)] + \mathrm{Var}(\mathbb{E}[H \mid z_i])$, the law of total variance. The first-order share of $z_i$ is $S_i = \mathrm{Var}(\mathbb{E}[H \mid z_i])/\mathrm{Var}(H)$, read by binning the posterior draws on $z_i$, or as $1 - \mathbb{E}[\mathrm{Var}(H \mid z_i)]/\mathrm{Var}(H)$ by holding $z_i$ at quantile points of its posterior through the conditional normal. On a quantity linear in the state it is $(\boldsymbol{\Sigma}\mathbf{g})_i^2/(\Sigma_{ii}\,\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g})$, the entry's term of the covariance split when the entry is uncorrelated, and the band with $z_i$ known is the same at every value. On a minimum the band with $z_i$ known depends on the value, and can be wider than today's.

Four conclusions carry beyond this site. A variance budget names the inputs a band rests on, and an unmeasured handbook value can outweigh every test of the existing plant. A budget belongs to the quantity it is computed for: on a minimum over elements an input carries its share only where its element binds, and the plant's budget can be far smaller than the element's. First-order shares of correlated entries overlap, and an entry with no lever of its own can carry a share through its correlations; their sum is not one. And the band with an input known is a function of the value it turns out to have: on a sum it narrows by the same amount at every value, on a minimum it can widen, so the worth of a measurement depends on what it may read, which is the preposterior of Chapter 18.

Table 17.9 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 17.9: The terms of variance budgets and where each is defined. Units for the site at the design day.

| Term | Expression | Units | Meaning | Defined in |
|---|---|---|---|---|
| Law of total variance | $\mathrm{Var}(H) = \mathbb{E}[\mathrm{Var}(H \mid z_i)] + \mathrm{Var}(\mathbb{E}[H \mid z_i])$ | $\mathrm{kW}^2$ | the split by one entry | 17.1.1 |
| First-order share | $S_i = \mathrm{Var}(\mathbb{E}[H \mid z_i])/\mathrm{Var}(H)$ | – | variance removed on average | 17.1.1 |
| Binning route | $s_H^{-2}\sum_b (n_b/N)(\bar{H}_b - \bar{H})^2$ | – | all shares from one set of draws | 17.1.1 |
| Conditioning route | $K^{-1}\sum_k \mathrm{Var}(H \mid z_i = v_k)$ | $\mathrm{kW}^2$ | mean conditional variance | 17.1.1 |
| Linear case | $S_i = (\boldsymbol{\Sigma}\mathbf{g})_i^2/(\Sigma_{ii}\,\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g})$ | – | overlaps when correlated | 17.1.1 |
| Cooling headroom | $\sum_i 400\,\mathrm{e}^{\kappa_i} + 50\,\mathrm{e}^{g_e} - Q_o - 660$ | kW | the chillers with the economizer | 17.1.3 |
| Site headroom | $\min_e Q_e^\dagger - 660$ | kW | the plant's | 17.1.3 |
| *Posterior* | | | | |
| Conditional normal | eq. (17.8) | mixed | the rest with one entry held | 17.2.1 |
| Unobserved entry | $\boldsymbol{\Sigma}_{-g_e,g_e} = \mathbf{0}$, $\Sigma_{g_e g_e} = 0.35^2$ | – | holding it leaves the rest | 17.2.1 |
| Handbook factor | $\Sigma_{g_e g_e}(F) = (\ln F)^2$ | – | the guess's width | 17.2.1 |
| *Questions and answers* | | | | |
| Guess's share of a sum | $(50\,\mathrm{e}^{\hat{g}_e} \times 0.35)^2/\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$ | – | the covariance split | 17.3.1 |
| Band with the guess known | $\mathrm{Var}(H \mid g_e = v)/\mathrm{Var}(H)$ | – | depends on $v$ for a minimum | 17.3.3 |
| Factory test | $\mathbf{J}_a = \mathbf{e}_{g_e}^\mathsf{T}/\sigma_a$ | – | one row on the guess | 17.6.1 |

## 17.8 Exercises

**Exercise 17.1.** *Formulation.* The planner asks how much of the site headroom's band the pipe's flow coefficient carries. Write the question over the site's state, say why the cooling headroom's budget gives it nothing, and state its answer from this chapter's draws. ($S_{l_k} = \mathrm{Var}(\mathbb{E}[H \mid l_k, \mathbf{y}] \mid \mathbf{y})/\mathrm{Var}(H \mid \mathbf{y})$, the first-order share (17.2) with $H$ of equation (17.7), read by binning the draws on $l_k$, equation (17.3). The cooling headroom (17.6) has no entry in $l_k$, and the covariance split (5.25) gives an entry with no gradient entry no share, so its share there is 0.000. In the site headroom the pipe binds in most draws, and its coefficient's first-order share is 0.139.)

**Exercise 17.2.** *Derivation.* Derive the law of total variance for the site headroom and one entry, and use it to give the economizer's share from the mean conditional variance. (Write $H - \mu = (H - m_i) + (m_i - \mu)$ with $m_i = \mathbb{E}[H \mid z_i]$; square, and take the expectation first with $z_i$ held and then over $z_i$. The cross term vanishes because $\mathbb{E}[H - m_i \mid z_i] = 0$, the first square averages to $\mathbb{E}[\mathrm{Var}(H \mid z_i)]$, and the second is $\mathrm{Var}(m_i)$, so $\mathrm{Var}(H) = \mathbb{E}[\mathrm{Var}(H \mid z_i)] + \mathrm{Var}(\mathbb{E}[H \mid z_i])$, equation (17.1). Dividing by $\mathrm{Var}(H)$, $S_i = 1 - \mathbb{E}[\mathrm{Var}(H \mid z_i)]/\mathrm{Var}(H)$. The nine quantile points give $\mathbb{E}[\mathrm{Var}(H \mid g_e)]/\mathrm{Var}(H) = 0.677$, so $S_{g_e} = 0.323$.)

**Exercise 17.3.** *Derivation.* Show that the cooling headroom's standard deviation with the economizer's conductance known is the same at every value of the conductance, and give it. ($H_c = \sum_i 400\,\mathrm{e}^{\kappa_i} + 50\,\mathrm{e}^{g_e} - Q_o - 660$ is a function of the chillers' entries plus a function of $g_e$, and $g_e$ is uncorrelated with every other entry, equation (17.9). Holding $g_e$ at $v$ fixes the second term at $50\,\mathrm{e}^v$ and leaves the first with its posterior, so the conditional variance is that of the first term at every $v$. By the delta form it is $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g} - (50.00 \times 0.350)^2 = 22.36^2 - 17.50^2$, a standard deviation of 13.92 kW.)

**Exercise 17.4.** *Application.* A factory test reports that the selected exchanger is weak, at the lowest of the nine quantile points of the handbook's normal. The review expected the test to narrow the site headroom's band. Explain what happens, and give the variance against today's. (With $g_e$ low, the economizer's duty pulls the chillers' binding load below the pipe's, and the chillers bind in nearly every draw. The site headroom then inherits the chillers' band with the guess known, a standard deviation of 13.92 kW, whose variance is $(13.92/9.62)^2 = 2.09$ times the site's today; today's site band is narrow only because the pipe cuts the chillers' upper tail. At the lowest point the pipe still binds first in some draws, and the variance with $g_e$ known is 1.534 times today's.)

**Exercise 17.5.** *Application.* The economizer's vendor offers a catalog value right to a factor of 1.2 instead of the handbook's. Give the site and cooling headrooms' bands and the probability that the pipe binds first, and say which end of the site's band moves. (By equation (17.10) the economizer's prior standard deviation becomes $\ln 1.2 = 0.182$, and every other entry of $\mathbf{\Sigma}$ is unchanged. The site headroom's 90 % band is $[448.23,\ 473.89]$ kW, 25.66 kW wide against 31.39 kW today: its lower end rises, set by the chillers, and its upper end, the pipe's, stays. The cooling headroom's band is 54.91 kW wide against 76.17 kW, and the pipe binds first with probability 0.6763.)

# Chapter 18

# Measurements Before Commitment

## 18.1 Which test narrows the build decision

The design sheet of the site was approved on one line: three 400 kW chillers against a 1 000 kW design load, a 20 % margin. Since then the review has learned that the line hides a contest. The chillers with their drawn economizer and the shared supply pipe bind within each other's bands at the design day, and which of the two binds first decides both the element to upgrade and the date by which the next order must be placed. Before either is committed the review can buy one test or survey, and six are quoted.

The economizer's maker offers a factory performance test of the exchanger the drawing selects, for R$ 40 thousand. A test crew offers a capacity test of each existing chiller at the design wet bulb of 22°C, for R$ 60 thousand, or of chiller C2 alone, for R$ 20 thousand. The same crew offers a high-flow pressure-drop test of the shared pipe, four readings between 25 and 30 L/s, for R$ 15 thousand, and a pump power test at the same four flows, for R$ 10 thousand. The economizer test with the three chiller tests together costs R$ 100 thousand. The design sheet knows which it would buy. It doubts the chillers most, because their nameplate is a catalog figure for machines that have run for years, and it doubts the pipe least, because its handbook coefficient puts the pipe far ahead as the first element to bind. It orders the chiller tests.

The sheet's choice rests on a reading of doubt, not of the decision. A test that narrows an element the review already knows well buys little; a test that narrows an element in doubt buys much only if the doubt is what holds the decision open. The question of this chapter is the drill question, verbatim,

*Which test or survey on the existing plant would most narrow the build decision?*

The economizer is drawn but not built, so its maker's test is the only test of it before the build, and the review lists it with the tests of the existing plant because the decision turns on it. The question is answered on the site's one fit of Chapter 9, with the binding loads of Chapter 13 as the quantity that decides the build. What is new is that no reading is taken: each candidate is a row whose value is not yet known, and the chapter reads what each would do to the decision before it is bought.

**Jargon: deciding quantity, decision level, settled decision, candidate measurement, factory performance test, capacity test, pressure-drop test, pump power test, preposterior covariance, shift of the posterior mean, width removed, value of information, common random numbers, Monte Carlo standard error, push-through identity**

The *deciding quantity* of a build decision is the derived quantity of the state whose value would change the decision; here it is the site's binding load at the design day, the smaller of the chillers' and the pipe's. The decision it carries is read as a probability, here the probability that the pipe binds first. A *decision level* is a value of that probability beyond which the review acts without further evidence, here 0.1 and 0.9; the decision is *settled* when the probability lies outside $[0.1,\ 0.9]$ and open when it lies inside. A *candidate measurement* is a test or survey not yet bought, written as one or more rows whose values are not yet known. A *factory performance test* is a test of a unit at its maker's works before delivery. A *capacity test* reads the cooling a chiller delivers at a stated wet bulb; a *pressure-drop test* reads the drop across the pipe at stated flows; a *pump power test* reads the pump's input power at stated flows. The *preposterior covariance* of a candidate is the covariance of the state that the candidate would leave, known before the candidate's value. The *shift of the posterior mean* is the change the candidate's value would make to the MAP state, a random quantity before the reading. The *width removed* by a candidate is the share of the deciding quantity's 90 % band that it would take off. The *value of information* of a candidate is what its outcome is worth to the decision, averaged over the outcomes before it is taken; here it is read as the probability of settling the decision. *Common random numbers* are one set of random draws used for every candidate, so that differences between candidates are not differences of sampling. The *Monte Carlo standard error* of a share of $N$ draws is the standard deviation of that share between independent sets of draws. The *push-through identity* moves a matrix through an inverse, $(\mathbf{I}+\mathbf{AB})^{-1}\mathbf{A}=\mathbf{A}(\mathbf{I}+\mathbf{BA})^{-1}$.

**Key idea**

A measurement not yet taken is a row whose value is unknown and whose Jacobian is known. The Jacobian alone gives the covariance the measurement would leave; the unknown value moves the posterior mean, and before the reading that move is normal with the covariance the measurement removes. Each move gives the probability that decides the build, so the measurement's worth to the decision is the share of its outcomes that would push that probability past a decision level. That share and the width the measurement takes off the band are different measures, and they need not rank the candidates alike.

### 18.1.1 A measurement read before it is taken

Every chapter of this book has closed its step 7 on the same question: which further reading would narrow the quantity that decides, and by how much. Each answered it for its own question, from the covariance a candidate would leave. This chapter derives that step once, for any plant, any candidate and any decision read as a probability.

The quantity owed is the probability that a candidate $a$ settles the decision. Let $f(\mathbf{z})$ be the deciding quantity and let the decision be read as the threshold probability $p=\Pr(f(\mathbf{z})<f^{*}\mid\mathbf{y})$, with the decision levels 0.1 and 0.9. After the candidate's value $\tilde{y}_a$ is read, the same probability becomes $p_a(\tilde{y}_a)$, and the candidate settles the decision when $p_a$ leaves $[0.1,\ 0.9]$. Before the reading $\tilde{y}_a$ is unknown, and the probability of settling averages over its predictive,

$$\begin{aligned} S_a &= \Pr(p_a(\tilde{y}_a)\notin[0.1,\ 0.9]\mid\mathbf{y}) = \int \mathbf{1}[p_a(\tilde{y}_a)\notin[0.1,\ 0.9]]\,p(\tilde{y}_a\mid\mathbf{y})\,\mathrm{d}\tilde{y}_a, \\ p_a(\tilde{y}_a) &= \Pr(f(\mathbf{z})<f^{*}\mid\mathbf{y},\ \tilde{y}_a), \end{aligned} \tag{18.1}$$

the design form of Chapter 2 with the indicator of a settled decision as the quantity averaged. Two objects are owed: the posterior after an outcome, which gives $p_a$, and the predictive of the outcome, which weights it.

**The posterior after an outcome.** The candidate is a stack of rows, each a function $h_a(\mathbf{z})$ of the state read with noise of standard deviation $\sigma_a$, $r_a(\mathbf{z}) = (h_a(\mathbf{z}) - \tilde{y}_a)/\sigma_a$. Near the MAP state $\hat{\mathbf{z}}$ each row is linear in the state,

$$\mathbf{r}_a(\mathbf{z}) \approx \mathbf{J}_a\,(\mathbf{z} - \hat{\mathbf{z}}) - \boldsymbol{\epsilon}_a, \qquad \mathbf{J}_a = \frac{\nabla_{\mathbf{z}} h_a^{\mathsf{T}}}{\sigma_a}\bigg|_{\hat{\mathbf{z}}}, \qquad \boldsymbol{\epsilon}_a = \frac{\tilde{y}_a - h_a(\hat{\mathbf{z}})}{\sigma_a}, \tag{18.2}$$

with $\mathbf{J}_a$ the candidate's Jacobian, one line per row, and $\boldsymbol{\epsilon}_a$ its normalized innovation, the outcome less what the fit expects, in standard deviations. The present posterior $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ enters the objective as the quadratic $\frac{1}{2}(\mathbf{z}-\hat{\mathbf{z}})^{\mathsf{T}}\boldsymbol{\Sigma}^{-1}(\mathbf{z}-\hat{\mathbf{z}})$, and the candidate's rows add $\frac{1}{2}\|\mathbf{J}_a(\mathbf{z}-\hat{\mathbf{z}})-\boldsymbol{\epsilon}_a\|^2$. Setting the gradient of the sum to zero gives the normal equations of the posterior after the outcome, $(\boldsymbol{\Sigma}^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a)(\mathbf{m}_a - \hat{\mathbf{z}}) = \mathbf{J}_a^{\mathsf{T}}\boldsymbol{\epsilon}_a$, so that

$$\boldsymbol{\Sigma}_a = \left(\boldsymbol{\Sigma}^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a\right)^{-1}, \qquad \mathbf{m}_a = \hat{\mathbf{z}} + \boldsymbol{\Sigma}_a\mathbf{J}_a^{\mathsf{T}}\boldsymbol{\epsilon}_a, \qquad \mathbf{z} \mid \mathbf{y},\ \tilde{y}_a \approx \mathcal{N}(\mathbf{m}_a,\ \boldsymbol{\Sigma}_a), \tag{18.3}$$

with $\mathbf{m}_a$ the posterior mean after the outcome and $\boldsymbol{\Sigma}_a$ the preposterior covariance, the covariance of equation (4.67) for a stack of rows. *Reliability Engineering by Derivation*, Chapter 31, derives the same covariance for a reliability test. The covariance holds no outcome: it is known from the Jacobian before the candidate is bought. The mean holds the outcome linearly, through $\boldsymbol{\epsilon}_a$.

**The predictive of the outcome.** Before the reading the innovation is the candidate's rows at a state drawn from the present posterior, plus their noise: $\boldsymbol{\epsilon}_a = \mathbf{J}_a(\mathbf{z} - \hat{\mathbf{z}}) + \mathbf{e}_a$ with $\mathbf{z} - \hat{\mathbf{z}} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$ and $\mathbf{e}_a \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ in normalized units. It is normal,

$$\boldsymbol{\epsilon}_a \mid \mathbf{y} \sim \mathcal{N}(\mathbf{0},\ \mathbf{S}_a), \qquad \mathbf{S}_a = \mathbf{I} + \mathbf{J}_a\boldsymbol{\Sigma}\mathbf{J}_a^{\mathsf{T}}, \tag{18.4}$$

the predictive of equation (2.13) for linear rows, with $\mathbf{S}_a$ the covariance of the innovation: the noise of the rows and the spread the present posterior puts on what they read.

**The shift of the posterior mean.** By equation (18.3) the shift $\mathbf{m}_a - \hat{\mathbf{z}} = \boldsymbol{\Sigma}_a\mathbf{J}_a^{\mathsf{T}}\boldsymbol{\epsilon}_a$ is linear in a normal innovation, so it is normal with covariance $\boldsymbol{\Sigma}_a\mathbf{J}_a^{\mathsf{T}}\mathbf{S}_a\mathbf{J}_a\boldsymbol{\Sigma}_a$. Multiplying out $(\boldsymbol{\Sigma}^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a)\boldsymbol{\Sigma}\mathbf{J}_a^{\mathsf{T}} = \mathbf{J}_a^{\mathsf{T}}(\mathbf{I} + \mathbf{J}_a\boldsymbol{\Sigma}\mathbf{J}_a^{\mathsf{T}}) = \mathbf{J}_a^{\mathsf{T}}\mathbf{S}_a$ gives the push-through form $\boldsymbol{\Sigma}_a\mathbf{J}_a^{\mathsf{T}} = \boldsymbol{\Sigma}\mathbf{J}_a^{\mathsf{T}}\mathbf{S}_a^{-1}$, and with it the covariance becomes $\boldsymbol{\Sigma}\mathbf{J}_a^{\mathsf{T}}\mathbf{S}_a^{-1}\mathbf{J}_a\boldsymbol{\Sigma}$, which the matrix-inversion lemma identifies as $\boldsymbol{\Sigma} - \boldsymbol{\Sigma}_a$:

$$\mathbf{m}_a \mid \mathbf{y} \sim \mathcal{N}(\hat{\mathbf{z}},\ \boldsymbol{\Sigma} - \boldsymbol{\Sigma}_a), \qquad \boldsymbol{\Sigma} = \mathbb{E}[\boldsymbol{\Sigma}_a] + \mathrm{Cov}(\mathbf{m}_a \mid \mathbf{y}). \tag{18.5}$$

The covariance the candidate removes from the state is the covariance of the move it makes to the mean, the second relation, which is the law of total variance with an outcome-free $\boldsymbol{\Sigma}_a$. A candidate that removes nothing cannot move the mean, and one that would pin an entry moves the mean by that entry's whole spread. The scalar form of this relation on one delta-form quantity is equation (6.11); equation (18.5) carries the whole state, so that the decision can be read jointly over every element.

**Down to a computation.** Each outcome gives the probability that decides,

$$p_a(\mathbf{m}_a) = \int \mathbf{1}\left[f(\mathbf{z}) < f^*\right] \mathcal{N}(\mathbf{z};\ \mathbf{m}_a,\ \boldsymbol{\Sigma}_a)\,\mathrm{d}\mathbf{z}, \tag{18.6}$$

a function of the outcome through $\mathbf{m}_a$ alone. Substituting the shift (18.5) for the predictive of the outcome turns equation (18.1) into two nested sets of draws: means $\mathbf{m}_j$ from $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma} - \boldsymbol{\Sigma}_a)$,

one per outcome; and for each, states from $\mathcal{N}(\mathbf{m}_j, \mathbf{\Sigma}_a)$ pushed through $f$, whose share below $f^*$ is $p_a(\mathbf{m}_j)$. The share of the outcomes whose $p_a$ lies outside $[0.1,\ 0.9]$ is $S_a$. The width removed needs only the inner draws: the 90 % band of $f$ under $\mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma}_a)$ against its band under $\mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$.

**Principle 18.1.** *A measurement not yet taken leaves the covariance* $\mathbf{\Sigma}_a = (\mathbf{\Sigma}^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a)^{-1}$*, known from its Jacobian, and moves the posterior mean by a normal amount of covariance* $\mathbf{\Sigma} - \mathbf{\Sigma}_a$*, known before its value. Its worth to a decision read as a probability is the share of those moves that push the probability past a decision level, read by draws of the move and, for each, draws of the state.*

### 18.1.2 The site, its deciding quantity and six candidates (step 0)

The plant is the site of Chapter 13, drawn with its four elements in Figure 13.1: three existing chillers C1–C3 with the drawn economizer, the shared supply pipe, the pump and the power feed, with the state of equation (9.4),

$$\mathbf{z} = \left(\kappa_{C1},\ \kappa_{C2},\ \kappa_{C3},\ d,\ g_e,\ c,\ l_k,\ l_e\right)^{\mathsf{T}}.$$

The scenario node is Chapter 13's: the wet bulb pinned at the design day's 22°C, every unit up, the economizer as drawn, and the site IT load free.

**The deciding quantity.** Chapter 13 found that of the four elements only the chillers and the pipe ever bind first; the pump and the feed bind later in every draw. The site's binding load is therefore the smaller of two, and the decision is their order. With the chillers' and the pipe's binding loads of equations (13.4) and (13.5),

$$Q^{\dagger}(\mathbf{z}) = \min\left(Q^{\dagger}_{\text{ch}}(\mathbf{z}),\ Q^{\dagger}_{\text{pipe}}(\mathbf{z})\right), \qquad D(\mathbf{z}) = Q^{\dagger}_{\text{pipe}}(\mathbf{z}) - Q^{\dagger}_{\text{ch}}(\mathbf{z}), \qquad p = \Pr\left(D(\mathbf{z}) < 0 \mid \mathbf{y}\right), \tag{18.7}$$

in kW of site IT load, with $D$ the pairwise gap of equation (13.3) and $p$ the probability that the pipe binds first. The binding load $Q^{\dagger}$ sets the order date of Chapter 11; the sign of $D$ sets which element the upgrade of Chapter 14 should raise, the decision "pipe or chillers first". It is the $f$ and the $f^* = 0$ of Section 18.1.1 with $f = D$, and it is settled when $p$ leaves $[0.1,\ 0.9]$.

**The candidates as rows.** Each candidate is a stack of rows of the site's physics, equation (9.3), at stated conditions and with a stated standard deviation per row. The factory test of the economizer reads its log-ratio $g_e$ directly, with standard deviation 0.10. A capacity test of chiller $i$ at 22°C reads $\text{Cap}_i = 400\,\text{e}^{\kappa_i}$, since the derate factor is one at the rating condition, with standard deviation 8 kW, the row of equation (16.14). The pressure-drop test reads the pipe's drop $\Delta P = (V_k/k_p)^2$ at four flows $V_k$ between 25 and 30 L/s, with $k_p = \text{e}^{l_k}$, each of standard deviation 1.5 kPa. The pump power test reads $P_{\text{pump}} = (\Delta P + \Delta P_{\text{h}})\, V_k/(1000\,\text{e}^{l_e})$ at the same four flows, with the halls' pinned $\Delta P_{\text{h}} = 100$ kPa, each of standard deviation 0.2 kW. By equation (18.2) their Jacobian lines are

$$\begin{aligned} \mathbf{J}_E &= \frac{\mathbf{e}_{g_e}^{\mathsf{T}}}{0.10}, \qquad \mathbf{J}_{C,i} = \frac{400\,\text{e}^{\hat{\kappa}_i}}{8}\,\mathbf{e}_{\kappa_i}^{\mathsf{T}}, \quad i = \text{C1}, \text{C2}, \text{C3}, \\ \mathbf{J}_{P,k} &= -\frac{2\,\Delta P_k}{1.5}\,\mathbf{e}_{l_k}^{\mathsf{T}}, \qquad \mathbf{J}_{U,k} = -\frac{1}{0.2}\left(\frac{2\,\Delta P_k\, V_k}{1000\,\eta}\,\mathbf{e}_{l_k}^{\mathsf{T}} + P_{\text{pump},k}\,\mathbf{e}_{l_e}^{\mathsf{T}}\right), \quad k = 1, \ldots, 4, \end{aligned} \tag{18.8}$$

with $\mathbf{e}_x$ the unit vector of entry $x$, $\eta = \text{e}^{l_e}$ the pump's efficiency, and $\Delta P_k$ and $P_{\text{pump},k}$ the drop and the pump power at flow $V_k$, all at $\hat{\mathbf{z}}$. A larger flow coefficient lowers the drop and the pump's power; a better pump lowers its power: both lines are negative. The candidate C2 is the line $\mathbf{J}_{C,\text{C2}}$ alone, and the candidate EC stacks $\mathbf{J}_E$ on the three lines of C. Table 18.1 lists the six candidates: each row is a candidate, the columns what it reads, its number of rows, the standard deviation of each row, the entries its lines touch and its stated price.

Table 18.1: The six candidate measurements before the build, with the rows of equation (18.8). Precisions and prices are stated inputs of this chapter, not estimated.

| Candidate | Reads | Rows [–] | sd per row [row's unit] | Entries its lines touch | Price [R$ thousand] |
|---|---|---|---|---|---|
| E | factory performance test of the economizer exchanger: $g_e$ | 1 | 0.10 | $g_e$ | 40 |
| C | capacity test of each existing chiller at 22°C | 3 | 8 kW | $\kappa_{C1}$, $\kappa_{C2}$, $\kappa_{C3}$ | 60 |
| C2 | capacity test of chiller C2 at 22°C | 1 | 8 kW | $\kappa_{C2}$ | 20 |
| P | pressure drop of the shared pipe at 25 to 30 L/s | 4 | 1.5 kPa | $l_k$ | 15 |
| U | pump power at the same four flows | 4 | 0.2 kW | $l_k$, $l_e$ | 10 |
| EC | E and C together | 4 | as E and C | $g_e$, $\kappa_{C1}$, $\kappa_{C2}$, $\kappa_{C3}$ | 100 |

**The pattern.** Read beside the gradients of the deciding quantity, the lines fix which candidate can reach the decision before any number. With $*$ for an entry that can be nonzero, in the columns of equation (9.4),

$$\begin{aligned} \mathbf{J}_E &= (0,\ 0,\ 0,\ 0,\ *,\ 0,\ 0,\ 0), & \mathbf{J}_{C,\mathrm{C2}} &= (0,\ *,\ 0,\ 0,\ 0,\ 0,\ 0,\ 0), \\ \mathbf{J}_{P,k} &= (0,\ 0,\ 0,\ 0,\ 0,\ 0,\ *,\ 0), & \mathbf{J}_{U,k} &= (0,\ 0,\ 0,\ 0,\ 0,\ 0,\ *,\ *), \\ \mathbf{g}_{\mathrm{ch}}^{\mathsf{T}} &= (*,\ *,\ *,\ 0,\ *,\ 0,\ 0,\ 0), & \mathbf{g}_{\mathrm{pipe}}^{\mathsf{T}} &= (0,\ 0,\ 0,\ 0,\ 0,\ 0,\ *,\ 0), \end{aligned} \tag{18.9}$$

and the lines of C1 and C3 sit in their own columns as C2's does. Four facts are read here. The economizer test sits in one column, $g_e$, which no row of the record sees and which has no covariance with any other entry; what it does to the state is what it does to $g_e$. The pressure-drop test sits in the pipe's column alone, so it acts on the pipe's binding load and on nothing else, and its reach over the decision is bounded by $l_k$'s part in the gap. The pump test reaches the decision only through $l_k$: its other column, $l_e$, is in neither gradient. And the capacity tests at 22°C have no entry in the derate's column, yet the record read each unit's log-ratio together with the derate in tests at other wet bulbs, so the posterior couples them, and a test of $\kappa_i$ narrows $d$ through that coupling, an entry neither gradient holds at the design day.

### 18.1.3 The planner's questions (step 1)

The planner asks four questions of the six candidates, and step 1 writes each over the state and the unknown outcome of each candidate. Table 18.2 lists them. The first column is the question in the planner's words, the drill question first; the second the general form of its category in the notation of Chapter 2; the third its mathematical form for the site; 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. Three are design questions, read before any outcome exists; the second is an estimation question about the band each candidate would leave. The answers are given in the order two, one, three, four, because settling is read beside the width.

### 18.1.4 The evidence as rows (step 2)

Step 2 attaches the operating record of Chapter 9 to the graph, unchanged: two capacity tests per chiller, twelve monthly windows of chiller power, pipe pressure drop, pump power and feed power, and eight prior rows, 62 rows on the 8 entries (Section 9.1.5). The outage record is not used: every unit runs at this scenario node. The candidates add no evidence. Their rows, equation (18.8), enter the graph with their Jacobians and without their values, which is all the preposterior (18.3) needs.

Table 18.2: The questions asked of the six candidates before the build. $\mathbf{y}$ is the site's operating record and priors; $\tilde{y}_a$ the unknown outcome of candidate $a$; $Q^{\dagger}$, $D$ and $p$ the deciding quantity of equation (18.7); $S_a$ the probability of settling, equation (18.1); $c_a$ the stated price.

| Planner's question | General form | Mathematical form | Category | Computation | Answered in |
|---|---|---|---|---|---|
| Which test or survey on the existing plant would most narrow the build decision? | $\mathbb{E}_{\tilde{y}_a \mid \mathbf{y}}[u]$, $u$ an acceptance indicator | $S_a = \Pr(p_a \notin [0.1,\ 0.9] \mid \mathbf{y})$, ranked, the width beside | design | preposterior | 18.3.2 |
| How much does each narrow the deciding quantity? | $\mathbb{E}_{\tilde{y}_a \mid \mathbf{y}}[u]$, $u$ a spread | width of the 90 % band of $Q^{\dagger}$ under $\mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma}_a)$, and the share removed | estimation | preposterior | 18.3.1 |
| Is narrowing the band the same as settling the decision? | the two measures of one candidate side by side | width removed against $S_a$, by candidate | design | preposterior | 18.3.3 |
| What is each worth per real? | $\mathcal{R}(a) = c_a + \sum_e C_e \Pr(e \mid \mathbf{y}, a)$ | $S_a$ per R$ 10 thousand of $c_a$ | design | expected consequence | 18.3.4 |

**Which row acts on which entry.** By Section 13.1.5 the record reads the pipe's $l_k$ by twelve monthly drops, the pump's $l_e$ by twelve monthly powers, the efficiency $c$ by the chiller and feed powers, and each chiller's log-ratio with the derate by its two capacity tests. The economizer's $g_e$ is read by nothing; its row is the handbook prior. The candidates fall on this pattern unevenly: E on the one entry no row sees, C and C2 on entries two tests each already read, P and U on entries twelve months already read.

**Modeling choices at the rows.** Five choices make the candidates' rows, and each is stated with what it gives and where it fails. The precisions and prices of Table 18.1 are stated inputs: this makes the ranking a statement about these offers, and it fails if a crew cannot reach the stated precision in the field. Each candidate row is Gaussian, unbiased and independent of the record given the state: this makes the candidate one more stack of normalized rows, and it fails if the test crew's instrument shares a bias with the record's meters. The rows are linear at $\hat{\mathbf{z}}$, equation (18.2): this makes $\mathbf{\Sigma}_a$ free of the outcome and the shift exactly normal, and it fails where an outcome would move the state far enough to change the Jacobian; over the spreads of the fit the lines change little. The factory test reads the log-ratio of the economizer as it will run at the site: this lets a reading at the maker's works stand for the installed exchanger, and it fails if the site's water quality or approach differs from the test's. And the decision levels are 0.1 and 0.9: this is the review's policy, the levels Chapter 16 used, and a stricter pair would settle less often under every candidate.

## 18.2 The posterior of the state and the deciding quantity (step 3)

All four questions of Table 18.2 are read from one posterior, the site's one fit, and from the preposterior covariance each candidate would leave. No fit is repeated; this section cites the posterior and reads the deciding quantity from it as it stands, the baseline every candidate is measured against.

Table 18.3: The deciding quantity at the design day before any candidate, by this chapter's 20 000 draws (seed 18) and by Chapter 13's (seed 9, the smallest of four elements, with its exact posterior). The gaps are Chapter 13's values less this chapter's.

| Route | Median [kW] | 5 % point [kW] | 95 % point [kW] | Width [kW] | $p$ [–] |
|---|---|---|---|---|---|
| This chapter's draws (reported) | 1 123.92 | 1 102.95 | 1 133.83 | 30.88 | 0.6520 |
| Chapter 13's draws | 1 123.88 | 1 102.58 | 1 133.97 | 31.39 | 0.6505 |
| Gap | −0.04 | −0.37 | +0.14 | +0.51 | −0.0015 |
| Chapter 13's exact posterior | 1 123.82 | 1 102.07 | 1 133.99 | 31.92 | 0.6411 |
| Gap | −0.10 | −0.88 | +0.16 | +1.04 | −0.0109 |

### 18.2.1 The fit of the site, cited

The MAP state $\hat{\mathbf{z}}$ and the Laplace covariance $\boldsymbol{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1}$ at $\hat{\mathbf{z}}$ are those of Section 9.2.2, tabulated with the stated truth in Table 9.2. The posterior standard deviations of the eight entries are 0.0158, 0.0175 and 0.0189 for the three chillers' log-ratios, $0.0029\,\mathrm{K}^{-1}$ for the derate, 0.3500 for the economizer's log-ratio, 0.0061 for the efficiency, 0.0046 for the pipe's log flow coefficient and 0.0190 for the pump's log-efficiency. The economizer's spread is its prior's, untouched by the record; every other entry is known to a few hundredths or better. The preposterior covariance of each candidate is formed from this $\boldsymbol{\Sigma}$ by equation (18.3) with the lines of equation (18.8) at $\hat{\mathbf{z}}$.

### 18.2.2 The deciding quantity now

The deciding quantity is read from $N = 20\,000$ draws $\mathbf{z}_i = \hat{\mathbf{z}} + \mathbf{L}\boldsymbol{\xi}_i$ of the Laplace normal, with $\mathbf{L}$ the Cholesky factor of $\boldsymbol{\Sigma}$ and the seed 18, each pushed through equations (13.4) and (13.5). The same standard normal vectors $\boldsymbol{\xi}_i$ are used for every candidate in step 4, with $\mathbf{L}$ replaced by the Cholesky factor of $\boldsymbol{\Sigma}_a$: the common random numbers that make the candidates' differences free of sampling noise. The draws give

$$\begin{aligned} \text{median } Q^{\dagger} = 1\,123.92\,\text{kW}, \qquad & Q^{\dagger} \in [1\,102.95,\ 1\,133.83]\,\text{kW}, \\ W_0 = 30.88\,\text{kW}, \qquad & p = 0.6520, \end{aligned} \tag{18.10}$$

in site IT load: the median and the 90 % band of the binding load, the band's width $W_0$, and the probability that the pipe binds first. Table 18.3 sets the same quantities from Chapter 13's draws beside them: each row is a route, the columns the median, the ends of the band, the width and $p$, with the gaps to this chapter's draws.

The two sets of draws agree to under 0.4 kW at every point of the band and to 0.0015 in $p$, inside the Monte Carlo standard error of 0.0034 that Chapter 13 gives for $p$; restricting the minimum to the chillers and the pipe drops nothing, since the pump and the feed never bind first. The exact posterior moves $p$ by −0.0109 and the lower end of the band by −0.88 kW, the departure of the posterior from the Laplace normal that Chapter 13 traces to the chillers' skew. The decision today is open: the pipe binds first in about two draws of three, inside [0.1, 0.9] by a wide margin on every route.

## 18.3 The answers (step 4)

Every answer is read from the pair $(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$ of Section 18.2.1 and the preposterior covariance $\boldsymbol{\Sigma}_a$ of each candidate, by draws with common random numbers. The width comes first, because the planner's instinct reads a measurement by the band it narrows; the settling of the decision is read beside it.

Table 18.4: The width of the 90 % band of the site's binding load after each candidate, equation (18.11), against today's 30.88 kW; 20 000 draws with common random numbers. The entries are those whose standard deviation falls by more than 1 %, now and after.

| Candidate | Width after [kW] | Removed [%] | Entries narrowed, sd now → after [entry's unit] |
|---|---|---|---|
| EC | 18.49 | 40.1 | $g_e$ 0.3500 → 0.0962; $\kappa_{C1}$ 0.0158 → 0.0127, $\kappa_{C2}$ 0.0175 → 0.0131, $\kappa_{C3}$ 0.0189 → 0.0141; $d$ 0.0029 → 0.0024 |
| E | 23.08 | 25.2 | $g_e$ 0.3500 → 0.0962 |
| C | 26.59 | 13.9 | $\kappa_{C1}$ 0.0158 → 0.0127, $\kappa_{C2}$ 0.0175 → 0.0131, $\kappa_{C3}$ 0.0189 → 0.0141; $d$ 0.0029 → 0.0024 |
| P | 27.30 | 11.6 | $l_k$ 0.0046 → 0.0022 |
| C2 | 28.90 | 6.4 | $\kappa_{C2}$ 0.0175 → 0.0134; $\kappa_{C1}$ 0.0158 → 0.0156, $\kappa_{C3}$ 0.0189 → 0.0183; $d$ 0.0029 → 0.0026 |
| U | 30.81 | 0.2 | $l_e$ 0.0190 → 0.0105 |

### 18.3.1 The width each candidate removes

**Planner's question.** How much does each narrow the deciding quantity? The question asks for the 90 % band of the site's binding load that each candidate would leave, against today's.

**General form.** A design question values a posterior that does not exist yet by the preposterior expectation of a quantity $u$ of it, equation (2.16),

$$U(a) = \int u\big(p(\mathbf{z} \mid \mathbf{y}, \tilde{y}_a)\big)\, p(\tilde{y}_a \mid \mathbf{y})\, \mathrm{d}\tilde{y}_a,$$

here with $u$ the spread of a derived quantity after the outcome.

**Specialized form.** The question is equation (2.16) with $u$ the width of the central 90 % band of $Q^\dagger$ of equation (18.7). On linear rows the covariance $\mathbf{\Sigma}_a$ of equation (18.3) holds no outcome, and the width is read about $\hat{\mathbf{z}}$, so the average over the outcomes is the width itself,

$$W_a = \hat{F}_a^{-1}(0.95) - \hat{F}_a^{-1}(0.05), \qquad \hat{F}_a(q) = \frac{1}{N}\sum_{i=1}^{N} \mathbf{1}\big[Q^\dagger(\hat{\mathbf{z}} + \mathbf{L}_a \boldsymbol{\xi}_i) \le q\big], \qquad \text{removed} = 1 - \frac{W_a}{W_0}, \tag{18.11}$$

with $\mathbf{L}_a$ the Cholesky factor of $\mathbf{\Sigma}_a$, the $\boldsymbol{\xi}_i$ the common normal vectors of Section 18.2.2 and $W_0$ today's width, equation (18.10). The entries a candidate narrows are read from the diagonal of $\mathbf{\Sigma}_a$ against that of $\mathbf{\Sigma}$.

**Building blocks.** The covariance $\mathbf{\Sigma}$ (Section 18.2.1); the lines (18.8) at $\hat{\mathbf{z}}$; the preposterior covariance (18.3); the binding loads (13.4) and (13.5) in each draw; the 20 000 common normal vectors.

**The widths.** Table 18.4 gives, for each candidate, the width of the band after it, the share of today's 30.88 kW it removes, and the posterior standard deviation of every entry it narrows by more than a hundredth of its spread, before and after.

The economizer test removes the most of any single candidate, 25.2 %: it takes $g_e$ from the handbook's 0.3500 to 0.0962, and $g_e$ moves the chillers' binding load by the economizer's duty. The chiller tests come second, at 13.9 %. Each narrows its unit's log-ratio by about a fifth to a quarter, because a test of standard deviation 8 kW reads a capacity the record already knows to a few kilowatts, and through the coupling of the log-ratios with the derate in the record they narrow the derate too, from 0.0029 to 0.0024 $\mathrm{K}^{-1}$, an entry that does not move the binding load

at the design day. The test of C2 alone removes 6.4 %. The pressure-drop test removes 11.6 %, almost as much as the three chiller tests: it halves the spread of $l_k$, from 0.0046 to 0.0022, and the pipe's binding load is the upper end of the joint band, where the pipe binds in the draws in which the chillers are high. The pump test narrows its own entry from 0.0190 to 0.0105 and removes 0.2 %, since the pump never binds first. The two tests together, EC, remove 40.1 %, a little more than the sum of their single shares, $25.2 + 13.9 = 39.1$ %: once the economizer is fixed, the chillers' own log-ratios carry more of what is left.

### 18.3.2 The probability each settles the decision

**Planner's question.** Which test or survey on the existing plant would most narrow the build decision? The question asks, for each candidate, how likely it is that its outcome would leave the order of the pipe and the chillers decided.

**General form.** The preposterior expectation, equation (2.16),

$$U(a) = \int u\big(p(\mathbf{z} \mid \mathbf{y}, \tilde{y}_a)\big)\, p(\tilde{y}_a \mid \mathbf{y})\, \mathrm{d}\tilde{y}_a,$$

here with $u$ the indicator that an acceptance rule is met after the outcome, so that $U(a)$ is the probability that the measurement ends with the question decided.

**Specialized form.** The question is equation (2.16) with $u = \mathbf{1}[p_a \notin [0.1,\ 0.9]]$ and $p_a$ the probability that the pipe binds first after the outcome, equation (18.6) with $f = D$ and $f^* = 0$; that is the probability $S_a$ of equation (18.1). By the shift (18.5) it is read by nested draws,

$$\begin{aligned} \mathbf{m}_j &= \hat{\mathbf{z}} + \mathbf{L}_\Delta \boldsymbol{\zeta}_j, \qquad \hat{p}_a(\mathbf{m}_j) = \frac{1}{K} \sum_{k=1}^{K} \mathbf{1}\big[D(\mathbf{m}_j + \mathbf{L}_a \boldsymbol{\eta}_k) < 0\big], \\ \hat{S}_a &= \frac{1}{J} \sum_{j=1}^{J} \mathbf{1}\big[\hat{p}_a(\mathbf{m}_j) \notin [0.1,\ 0.9]\big], \qquad \mathrm{se}(\hat{S}_a) = \sqrt{\hat{S}_a\,(1 - \hat{S}_a)/J}, \end{aligned} \tag{18.12}$$

with $\mathbf{L}_\Delta \mathbf{L}_\Delta^{\mathsf{T}} = \boldsymbol{\Sigma} - \boldsymbol{\Sigma}_a$, $J = 600$ outcomes and $K = 4\,000$ states per outcome, the normal vectors $\boldsymbol{\zeta}_j$ and $\boldsymbol{\eta}_k$ common to every candidate. The probability of the pipe binding first after each outcome is joint over the chillers' and the pipe's entries, the state route; the quantiles of the 600 values $\hat{p}_a(\mathbf{m}_j)$ give its preposterior distribution.

**Building blocks.** The preposterior covariance (18.3) and the difference $\boldsymbol{\Sigma} - \boldsymbol{\Sigma}_a$; the shift (18.5); the gap $D$ of equation (18.7) in each draw; the decision levels.

**The settlement.** Table 18.5 gives, for each candidate, the 5, 50 and 95 % points of the probability that the pipe binds first after it, the mean of that probability over the outcomes, and the probability that the candidate settles the decision, with its Monte Carlo standard error. Read on 4 000 settlement draws of today's posterior, today's probability is 0.6380, the common baseline of these columns. The same distribution is drawn in Figure 18.1: one bar per candidate from the 5 % to the 95 % point, short ticks at the deciles, the median as a dot, the horizontal axis the probability that the pipe binds first after the outcome, the shaded region the open interval [0.1, 0.9] and the dashed line today's 0.6520. A dotted tick on each bar marks the median of $p_a$ when the readings are drawn from the stated truth instead of from today's predictive, a check marker for Section 18.4.3.

The economizer test with the chiller tests settles the decision with probability

$$S_{EC} = 0.4450, \qquad S_E = 0.2500, \tag{18.13}$$

Table 18.5: The preposterior distribution of the probability that the pipe binds first, $p_a$ of equation (18.6), and the probability that each candidate settles the decision, equation (18.12); 600 outcomes of 4 000 draws each, common random numbers across the candidates.

| Candidate | $p_a$, 5 % [–] | $p_a$, 50 % [–] | $p_a$, 95 % [–] | Mean of $p_a$ [–] | $S_a$ [–] | se of $S_a$ [–] |
|---|---|---|---|---|---|---|
| EC | 0.0377 | 0.7229 | 0.9998 | 0.6300 | 0.4450 | 0.0203 |
| E | 0.1699 | 0.6647 | 0.9953 | 0.6325 | 0.2500 | 0.0177 |
| C | 0.3252 | 0.6509 | 0.9103 | 0.6437 | 0.0550 | 0.0093 |
| C2 | 0.4291 | 0.6435 | 0.8333 | 0.6370 | 0.0067 | 0.0033 |
| P | 0.5012 | 0.6435 | 0.7676 | 0.6379 | 0.0000 | 0.0000 |
| U | 0.6192 | 0.6382 | 0.6590 | 0.6388 | 0.0000 | 0.0000 |

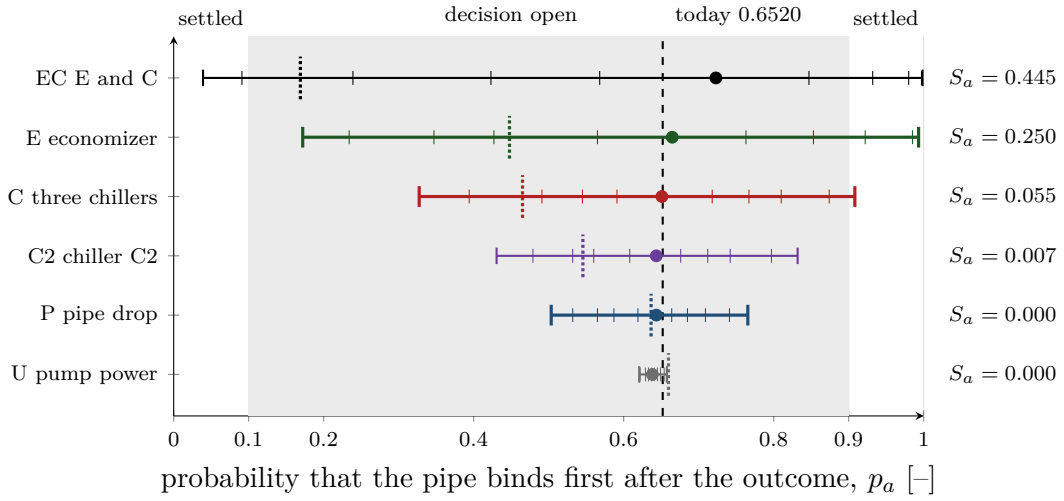



Figure 18.1: The preposterior distribution of the probability that the pipe binds first, for each candidate (Table 18.5). Bars: from the 5 % to the 95 % point of $p_a$ over the 600 outcomes; short ticks: the deciles; dots: the medians. Dotted ticks: the median of $p_a$ with the readings drawn from the stated truth (Section 18.4.3). Shaded: the open interval [0.1, 0.9]; an outcome outside it settles the decision. Dashed: today's probability 0.6520. Right: the probability $S_a$ that the candidate settles the decision.

with Monte Carlo standard errors of 0.0203 and 0.0177, and the economizer test alone settles it a quarter of the time: of the single candidates it is the one. The chiller tests settle it with probability 0.0550, the test of C2 alone with 0.0067, and the pressure-drop and pump tests in none of the 600 outcomes.

The bars of Figure 18.1 show how. The economizer test spreads the probability from 0.1699 to 0.9953 between its 5 and 95 % points: its outcomes carry the pipe's probability past 0.9 often and below 0.1 in fewer than one outcome in twenty. The asymmetry is the economizer's duty, $50\,\mathrm{e}^{g_e}$ kW by equation (13.4): a high reading of $g_e$ raises the chillers' binding load by a growing amount and leaves the pipe clearly first, while a low reading cannot take more than the duty itself off the chillers. With the chiller tests as well the lower tail reaches 0.0377, since the chillers' own log-ratios, now read too, can carry them below the pipe. The chiller tests alone stretch the probability from 0.3252 to 0.9103, past 0.9 only at their upper tail. The pressure-drop test keeps it between 0.5012 and 0.7676, and the pump test between 0.6192 and 0.6590: neither moves it near a decision level.

### 18.3.3 Width against settlement

**Planner's question.** Is narrowing the band the same as settling the decision? The question asks whether the two measures of a candidate, the width it removes and the probability that it settles, order the six candidates alike.

**General form.** Two values of the preposterior expectation of one candidate, equation (2.16),

$$U(a) = \int u(p(\mathbf{z} \mid \mathbf{y}, \tilde{y}_a))\, p(\tilde{y}_a \mid \mathbf{y})\, \mathrm{d}\tilde{y}_a,$$

with $u$ a spread for the one and an acceptance indicator for the other, set side by side across the candidates.

**Specialized form.** The question is equation (2.16) read twice for each candidate: with $u$ the width, giving the share removed of equation (18.11), and with $u$ the indicator of a settled decision, giving $S_a$ of equation (18.12). The two order the candidates alike only if the width a candidate removes from $Q^\dagger$ is the variance it removes from the gap $D$. It need not be: the band of $Q^\dagger$ is set at its lower end by the chillers and at its upper end by the pipe, while the sign of $D$ is set by whatever carries the band of the gap, the economizer's $g_e$ by Section 13.3.4.

**Building blocks.** The widths of Table 18.4; the settlement of Table 18.5; the gap's budget of Table 13.6.

**The two measures.** The two measures of each candidate are plotted against each other in Figure 18.2: the horizontal axis is the width removed, the vertical axis the probability of settling, one labeled point per candidate with a bar of one Monte Carlo standard error, and the dashed vertical line the width that an exact knowledge of $g_e$ would remove, the bound of Section 18.4.1.

The two measures agree at the top and part at the bottom. The economizer test, alone or with the chiller tests, leads on both. Below them the order breaks: by width the pressure-drop test is fourth, at 11.6 %, ahead of the test of C2 at 6.4 %; by settlement it is last with the pump test, at zero, behind C2's 0.0067. This is the surprise of the chapter: the pipe test narrows the band by 11.6 % and never settles the decision. It sharpens an entry that is not in doubt. The record already fixes the pipe's flow coefficient to 0.0046, and $l_k$ carries 0.053 of the variance of the gap (Table 13.6); halving its spread moves the pipe's end of the band and leaves the gap's sign where the economizer and the chillers put it. Width and settlement are not the same

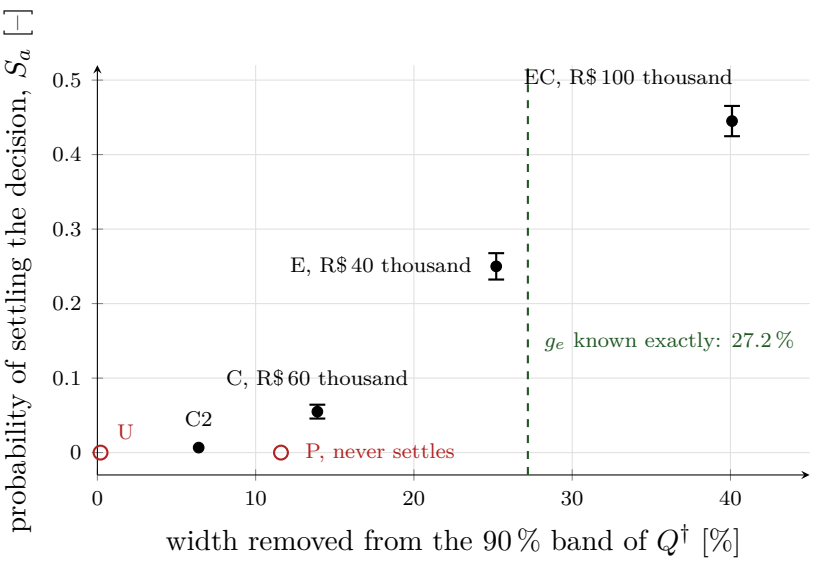



Figure 18.2: The two measures of each candidate. Horizontal: the share of the deciding quantity's 90 % band that it removes (Table 18.4); vertical: the probability that its outcome settles the decision (Table 18.5), with a bar of one Monte Carlo standard error. Open red circles: candidates that settle the decision in none of the 600 outcomes. Dashed: the width removed if $g_e$ were known exactly.

measure, and the planner wants settlement. A narrower band is worth buying only where the narrowing falls on what holds the decision open.

### 18.3.4 The worth of each candidate per real

**Planner's question.** What is each worth per real? The question asks how much settling each candidate buys for its price.

**General form.** The expected consequence of an action, with its certain cost and the events it leaves, equation (2.23),

$$\mathcal{R}(a) = c_a + \sum_e C_e \Pr(e \mid \mathbf{y}, a),$$

here with the price $c_a$ of the candidate as its certain cost and the open decision as the event it leaves.

**Specialized form.** The question is equation (2.23) with $c_a$ the stated price of Table 18.1 and the event an open decision, of probability $1 - S_a$. The consequence of an open decision is not priced in this chapter: what the review loses by committing the wrong upgrade is a question of the build's value, not of the test. The comparison therefore reports the probability of settling against the price, not netted,

$$s_a = \frac{10\, S_a}{c_a}, \tag{18.14}$$

the probability of settling per R$ 10 thousand, with $c_a$ in thousands of reais and $S_a$ from equation (18.12).

**Building blocks.** The probabilities $S_a$ (Section 18.3.2); the stated prices.

Table 18.6: The six candidates ranked by the probability of settling the decision, with the width removed, the stated price and the probability of settling per R$ 10 thousand, equation (18.14); the last three columns rank the candidates by width, by settlement and by settlement per real.

| Candidate | $S_a$ [–] | Width removed [%] | Price [R$ thousand] | $s_a$ [per R$ 10 thousand] | Rank, width [–] | Rank, $S_a$ [–] | Rank, $s_a$ [–] |
|---|---|---|---|---|---|---|---|
| EC | 0.4450 | 40.1 | 100 | 0.0445 | 1 | 1 | 2 |
| E | 0.2500 | 25.2 | 40 | 0.0625 | 2 | 2 | 1 |
| C | 0.0550 | 13.9 | 60 | 0.0092 | 3 | 3 | 3 |
| C2 | 0.0067 | 6.4 | 20 | 0.0033 | 5 | 4 | 4 |
| P | 0.0000 | 11.6 | 15 | 0.0000 | 4 | 5 | 5 |
| U | 0.0000 | 0.2 | 10 | 0.0000 | 6 | 5 | 5 |

**The ranking.** Table 18.6 ranks the six candidates by their probability of settling: each row is a candidate, the columns its probability of settling, the width it removes, its price, its probability of settling per R$ 10 thousand and its rank by each measure.

Per real the economizer test leads, at 0.0625 per R$ 10 thousand, ahead of the economizer test with the chiller tests at 0.0445, which settles more but costs two and a half times as much. The design sheet's instinct, the chiller tests, is third at 0.0092, and the economizer test buys $0.0625/0.0092 \approx 7$ times as much settling per real. The two together settle more than the sum of their single probabilities, 0.4450 against $0.2500 + 0.0550 = 0.3050$: once the economizer's duty is read, the chillers' own log-ratios carry most of what keeps the gap near zero, and reading them as well pushes more outcomes past a level. The pressure-drop and pump tests, the cheapest on the list, buy nothing for this decision.

## 18.4 The checks (step 5)

The answers of Section 18.3 rest on the preposterior of linear rows, on Monte Carlo draws and on a record generated from a stated truth; each is tested here.

### 18.4.1 The economizer known exactly

No reading of $g_e$ can remove more of the band than knowing it exactly. A row on $g_e$ of vanishing standard deviation, through the same preposterior (18.3) and the same common draws, leaves a width of 22.48 kW, a share removed of 27.2 %. The factory test of standard deviation 0.10 removes 25.2 %, close to the bound: its row takes $g_e$'s spread from 0.3500 to 0.0962, and what remains of the band after exact knowledge belongs to the chillers' log-ratios and the pipe. A more precise factory test would buy at most two more points of width. A sharper reading of the economizer is not where more width lies; the chillers' log-ratios are, which is what the pair EC adds.

### 18.4.2 The Monte Carlo errors and the mean of the outcomes

The probabilities of settling carry the Monte Carlo standard errors of 600 outcomes: 0.0203 for EC, 0.0177 for E, 0.0093 for C and 0.0033 for C2. The gaps between the four are many times their errors, so the ranking of Table 18.6 is not a sampling effect; and because the outcomes and the states are the same normal vectors for every candidate, the differences between candidates carry less noise than these errors suggest. The pressure-drop and pump tests settle in none of the 600 outcomes, so their probability is below about one in two hundred; their 95 % points, 0.7676 and 0.6590, lie far inside the upper level.

A second check is exact. Before an outcome the posterior probability averages to today's, since the shift (18.5) has mean zero and the probability after an outcome integrates back to the probability before. The baseline is today's probability on the same 4 000 settlement draws, 0.6380. The means of $p_a$ in Table 18.5 run from 0.6300 to 0.6437, within 0.008 of it, small beside the spread of each candidate's 600 outcomes; the pump test, which barely moves the state, has the mean 0.6388. The baseline itself lies 0.0140 below the 0.6520 of the 20 000 draws of Section 18.2.2, inside two Monte Carlo standard errors of a share of 4 000 draws, about 0.0076 each. The offset is common to every candidate, since all read the same draws, and the probability of settling is a count of outcomes beyond the levels, which an offset of this size moves by a few outcomes at most.

### 18.4.3 The answers against the stated truth

The operating record is generated from a stated truth, and the last check compares the answers with it. Table 18.7 collects the checks: each row is a quantity, the columns its reported value, the value beside it and the reading. At the truth the chillers bind at 1 106.00 kW and the pipe at 1 127.31 kW, so the chillers bind first, by 21.31 kW, and the site's binding load is 1 106.00 kW, inside today's band [1 102.95, 1 133.83] kW near its lower end. Today's posterior gives the true order a probability of $1 - 0.6520 = 0.3480$; the decision would be wrong if it were made now on the more likely order. The truth's economizer is smaller than the handbook's, $g_e = -0.2231$ (Section 13.4.3): a reading near it would lower the chillers' binding load and move the probability toward the true order. The preposterior averages over the outcomes that today's posterior expects, not over those the truth would produce; it ranks the candidates by what they can do. The check draws each candidate's readings from the stated truth instead, with the stated precision, and reads $p_a$ after each, on the same draws. With the economizer test the median of $p_a$ falls to 0.4477; the probability then points to the chillers, below 0.5, in 0.7000 of the readings, and settles on them, below 0.1, in none. The chiller tests point to the chillers in 0.5967 of the readings, the test of C2 in 0.4000, and the pipe and pump tests in none, with medians of 0.6365 and 0.6597 that stay on today's wrong side. Only the pair EC settles the decision on the true side, in 0.3033 of the readings, with the median at 0.1689 and the order pointed right in 0.8867. The dotted ticks of Figure 18.1 mark these medians. The check confirms the ranking's order, the economizer test first among the single candidates and the pipe and pump tests last; it also shows that at this truth the economizer test alone would turn the decision toward the true order but leave it open, and the chillers' own log-ratios are needed to close it.

## 18.5 The report (step 6)

Buy the economizer's factory test before committing the upgrade or the order. For R$ 40 thousand it settles the order of the pipe and the chillers in a quarter of its outcomes, 0.250, and takes a quarter, 25.2 %, off the deciding band, the most of any single candidate and the most per real. With the chiller tests as well, for R$ 100 thousand, it settles 0.445. The pipe and pump tests should not be bought for this decision, and the chiller tests the design sheet would order buy about a seventh as much settling per real. In the world the record was built from, where the chillers bind first, the factory test alone points the review to the chillers in 0.7000 of its outcomes but never settles it there; the two together point there in 0.8867 and settle it in 0.3033. If the review must settle the order rather than lean it, the pair is the purchase.

### 18.5.1 The answers and their bands

Table 18.8 gives the answers in the order of Table 18.2. Each row is a question; the second column gives the reported answer, and the third its Monte Carlo error, its bound or the values

Table 18.7: The checks of the measurement ranking. The reported values are from this chapter's draws; beside them, the bound, the route or the value named; the truth is the stated truth of the record.

| Quantity | Reported [kW or –] | Beside [kW or –] | Reading |
|---|---|---|---|
| Width after E, against $g_e$ known exactly | 23.08 | 22.48 | E removes 25.2 % of the bound's 27.2 % |
| Pr(pipe first) today, against Chapter 13's draws | 0.6520 | 0.6505 | inside a standard error of 0.0034 |
| Pr(pipe first) today, against the exact posterior | 0.6520 | 0.6411 | the Laplace normal's departure |
| Today's $p$ on the settlement draws, against the 20 000 draws | 0.6380 | 0.6520 | the settlement draws' offset, common to all |
| Mean of $p_a$ after E and EC | 0.6325, 0.6300 | 0.6380 | within 0.008 of the common baseline |
| $S_a$ of E, EC, Monte Carlo standard error | 0.0177, 0.0203 | – | the ranking holds by many errors |
| *The stated truth* | | | |
| Chillers' binding load [kW] | – | 1 106.00 | binds first |
| Pipe's binding load [kW] | – | 1 127.31 | second, by 21.31 kW |
| Site's binding load [kW] | 1 123.92 | 1 106.00 | inside [1 102.95, 1 133.83], near the lower end |
| Probability of the true order today [–] | 0.3480 | – | the decision open, leaning the wrong way |
| Readings from the truth: share with $p_a < 0.5$, pointing to the chillers [–] | E 0.7000 | EC 0.8867 | C 0.5967, C2 0.4000, P and U 0 |
| Readings from the truth: share with $p_a < 0.1$, settled on the chillers [–] | E 0.0000 | EC 0.3033 | C, C2, P and U 0: only the pair settles on the true side |

Table 18.8: The report on the six candidates before the build, at the design day with the plant as drawn. Today: the site's binding load median 1 123.92 kW, band 1 102.95 to 1 133.83 kW (width 30.88 kW); the pipe binds first with probability 0.6520.

| Question | Answer | Error, bound or beside |
|---|---|---|
| Which test or survey on the existing plant would most narrow the build decision? | the economizer's factory test, settles with probability 0.250; with the chiller tests 0.445 | se 0.018 and 0.020; the chiller tests 0.055, C2 0.007, the pipe and pump tests 0.000 |
| How much does each narrow the deciding quantity? | EC 40.1 %, E 25.2 %, C 13.9 %, P 11.6 %, C2 6.4 %, U 0.2 % of the width | $g_e$ known exactly 27.2 % |
| Is narrowing the band the same as settling the decision? | no: the pipe test removes 11.6 % and settles in no outcome | by width P is fourth, by settlement last |
| What is each worth per real? | per R$ 10 thousand: E 0.0625, EC 0.0445, C 0.0092, C2 0.0033, P and U 0 | prices R$ 40, 100, 60, 20, 15 and 10 thousand |

beside it.

### 18.5.2 The design-review entry

The design review's entry replaces the sheet's test with the test that moves the decision. Before the upgrade of Chapter 14 or the order of Chapter 11 is committed, the review buys the economizer's factory performance test, at R$ 40 thousand. Today the pipe binds first with probability 0.652, an open decision; the test settles it, past 0.9 or below 0.1, with probability 0.250 (Monte Carlo standard error 0.018), and it takes 25.2 % off the band of the site's binding load, against the 27.2 % that exact knowledge of the economizer would take. If the budget allows R$ 100 thousand, the chiller tests at 22°C are added, and the pair settles the decision with probability 0.445. The entry records what is not bought: the pipe's pressure-drop test, which would narrow the band by 11.6 % and leaves the decision open in every outcome, and the pump test, which narrows an entry the decision does not read. It also records the limit of the purchase: in three outcomes of four the decision stays open after the economizer test, and the review then commits on a probability, not a certainty.

**Common mistake**

Ranking measurements by the width they take off the band. A test can narrow the deciding quantity where it is already known and leave the decision where it was; the measurement to buy is the one whose outcomes push the probability that decides past a decision level.

## 18.6 The next measurement (step 7)

Every earlier chapter's step 7 pointed to this chapter, and this chapter is that step made general: the preposterior covariance of a stack of rows, the shift it allows the posterior mean, and the share of shifts that settle the decision. Its own next measurement is the one the ranking leaves.

### 18.6.1 The step made general

Each earlier step 7 read one part of the derivation of Section 18.1.1. Chapter 5 ranked candidates by the width each removes, equation (5.26), the first half of equation (18.3) on one

gradient. Chapters 6 and 16 read the shift of the posterior mean on one delta-form quantity, equations (6.11) and (16.15), the scalar form of equation (18.5), and asked whether it would settle a question. Chapter 17 named the factory test of the economizer and its ceiling, $g_e$ known exactly, the bound of Section 18.4.1. This chapter carries the shift on the whole state, reads the decision jointly over the elements that contend for it, and sets the width and the settlement side by side for every candidate. The general step needs three things a chapter can always supply: the candidate's Jacobian lines and precision, the decision read as a probability with its levels, and the draws that carry the state through the deciding quantity.

### 18.6.2 The measurement after the factory test

The factory test reads the economizer before it is built, at its maker's works and under the maker's conditions. Its limit is the bound of Section 18.4.1: even exact, it removes 27.2 % of the band, and its reading at the works stands for the exchanger as it will run at the site only by the modeling choice of Section 18.1.4. The next measurement is the economizer measured where it will run, on one hall, which also bears on whether to build it at all; Chapter 19 values it.

### 18.6.3 What the analysis showed

The design sheet would buy the chiller tests, the element it doubts most. The analysis ranks them third of six, at a probability of settling the decision of 0.055 and about a seventh of the economizer test's settling per real. Three findings carry the answer. First, a test is worth what it does to the entry that holds the decision open, and here that entry is the economizer's handbook conductance, which no reading of the plant sees; its maker's test settles the order of the pipe and the chillers in a quarter of its outcomes. Second, and beyond this site, width and settlement are different measures: the pipe's pressure-drop test takes 11.6 % off the deciding band and settles the decision in none of 600 outcomes, because it sharpens the pipe's end of the band, an entry the record already knows, and leaves the sign of the gap to the economizer and the chillers. Third, tests can be complementary in what they settle: the economizer and chiller tests together settle 0.445, more than the 0.305 of the two apart, because each removes the doubt that would otherwise keep the other's outcomes inside the levels.

### 18.6.4 Limits of the model

**Linear rows at the MAP state.** The preposterior covariance is formed from the candidates' lines at $\hat{\mathbf{z}}$, so it holds no outcome, and the band's width is read about $\hat{\mathbf{z}}$. An outcome far in a tail would move the state enough to change the lines, and the chillers' binding load is an exponential in $g_e$; the economizer test's tail outcomes are where this approximation is weakest. Chapter 17 found the same for the site's headroom with $g_e$ known exactly: the variance left depends on the value read, so the widths of Table 18.4 are the widths at a reading equal to the MAP value, not their average over the readings.

**The Laplace normal throughout.** The settlement is read on draws of the Laplace normal before and after each outcome. The exact posterior moves today's probability by $-0.0109$ (Table 18.3); the preposterior on the exact posterior, which would reweight every outcome's draws, is not computed.

**One decision, two levels.** The candidates are ranked for one decision, the order of the pipe and the chillers, at the levels 0.1 and 0.9. The same tests bear differently on other decisions: the order date of Chapter 11 reads the chillers' log-ratios, which the chiller tests narrow, and Chapter 16's 10 % question reads them alone. Stricter levels settle less often under every candidate.

**Stated precisions and prices.** The candidates' precisions and prices are stated inputs of this chapter, not estimated, and the consequence of an open decision is not priced, so the comparison per real is a ratio of a probability to a price, not a value netted against a loss.

**The constructed record.** The operating record is generated from a stated truth with the stated reading errors, not taken from a site, and the comparison with that truth in Section 18.4.3 is a check of the method on its own model, not of the model against a real plant.

**The next chapter.** Chapter 19 asks whether a pilot of the economizer on one hall would change the decision to build it.

## 18.7 What to retain

**Definition 18.1** (Preposterior of a decision)**.** A candidate measurement with Jacobian lines $\mathbf{J}_a$ leaves the preposterior covariance $\boldsymbol{\Sigma}_a = (\boldsymbol{\Sigma}^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a)^{-1}$ and moves the posterior mean by a shift $\mathbf{m}_a - \hat{\mathbf{z}} \sim \mathcal{N}(\mathbf{0},\ \boldsymbol{\Sigma} - \boldsymbol{\Sigma}_a)$, both known before its value. For a decision read as the probability $p = \Pr(f(\mathbf{z}) < f^* \mid \mathbf{y})$, each shift gives the probability after the outcome under $\mathcal{N}(\mathbf{m}_a, \boldsymbol{\Sigma}_a)$, and the candidate settles the decision with the probability $S_a$ that this leaves the decision levels' interval. The width removed, $1 - W_a/W_0$, is read beside it, and the price beside both.

Four conclusions carry beyond this site. A measurement not yet taken is known through its Jacobian: the covariance it leaves needs no outcome, and the move it makes to the mean is normal with the covariance it removes. A measurement's worth to a decision is the share of its outcomes that would decide, and that is a different number from the width it takes off a band; a test can narrow the band where the band is already known and leave the decision untouched. The test to buy acts on the entry that carries the decision's gap, which can belong to an element not yet built. And tests can be complementary in what they settle, because each removes the doubt that keeps the other's outcomes short of a level.

Table 18.9 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 18.9: The terms of measurements before commitment and where each is defined. Units for the site at the design day.

| Term | Expression | Units | Meaning | Defined in |
|---|---|---|---|---|
| Probability of settling | $S_a = \Pr(p_a \notin [0.1,\ 0.9] \mid \mathbf{y})$ | – | the decision decided by the outcome | 18.1.1 |
| Linearized candidate rows | $\mathbf{J}_a(\mathbf{z} - \hat{\mathbf{z}}) - \boldsymbol{\epsilon}_a$ | – | a row whose value is unknown | 18.1.1 |
| Preposterior covariance | $\boldsymbol{\Sigma}_a = (\boldsymbol{\Sigma}^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a)^{-1}$ | entries' | the covariance left | 18.1.1 |
| Predictive of the innovation | $\boldsymbol{\epsilon}_a \sim \mathcal{N}(\mathbf{0},\ \mathbf{I} + \mathbf{J}_a\boldsymbol{\Sigma}\mathbf{J}_a^{\mathsf{T}})$ | – | the outcome before it is read | 18.1.1 |
| Shift of the mean | $\mathbf{m}_a \sim \mathcal{N}(\hat{\mathbf{z}},\ \boldsymbol{\Sigma} - \boldsymbol{\Sigma}_a)$ | entries' | the move the outcome makes | 18.1.1 |
| Deciding quantity | $Q^{\dagger} = \min(Q^{\dagger}_{\text{ch}}, Q^{\dagger}_{\text{pipe}})$; $D = Q^{\dagger}_{\text{pipe}} - Q^{\dagger}_{\text{ch}}$ | kW | the binding load and the order | 18.1.2 |
| Candidate lines | eq. (18.8) | – | E, C, C2, P, U, EC as rows | 18.1.2 |
| *Questions and answers* | | | | |
| Width removed | $1 - W_a/W_0$, band under $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma}_a)$ | % | the band a candidate takes off | 18.3.1 |
| Nested draws | means from $\boldsymbol{\Sigma} - \boldsymbol{\Sigma}_a$, states from $\boldsymbol{\Sigma}_a$ | – | the settlement read by draws | 18.3.2 |
| Common random numbers | one set of normal vectors for every candidate | – | differences free of sampling | 18.3.2 |
| Settlement per real | $s_a = 10\,S_a/c_a$ | per R$ 10 thousand | settling bought per price | 18.3.4 |
| Bound of a reading | the entry known exactly | % | most a reading of it can remove | 18.4.1 |

## 18.8 Exercises

**Exercise 18.1.** *Formulation.* The facilities team asks whether the pipe's pressure-drop test would settle which element binds first. Write the question over the state and the test's unknown outcome, and state its answer from this chapter. ($S_P = \Pr(p_P \notin [0.1,\ 0.9] \mid \mathbf{y})$, equation (18.1) with $p_P = \Pr(D(\mathbf{z}) < 0 \mid \mathbf{y},\ \tilde{y}_P)$, the four rows $\mathbf{J}_{P,k}$ of equation (18.8) and the shift (18.5). Over the 600 outcomes the probability that the pipe binds first stays between 0.5012 and 0.7676 at its 5 and 95 % points, and the test settles the decision with probability 0.0000.)

**Exercise 18.2.** *Derivation.* Show that the economizer test leaves $g_e$ with the standard deviation 0.0962. (The entry $g_e$ has no row but its prior, so it has no covariance with any other entry, and the line $\mathbf{J}_E = \mathbf{e}_{g_e}^{\mathsf{T}}/0.10$ touches its column alone. The preposterior covariance (18.3) then changes only that diagonal entry, whose precision becomes $1/0.35^2 + 1/0.10^2 = 8.16 + 100 = 108.16$, so that the standard deviation is $1/\sqrt{108.16} = 0.0962$.)

**Exercise 18.3.** *Derivation.* Show from equations (18.3) and (18.4) that the shift of the posterior mean has the covariance $\boldsymbol{\Sigma} - \boldsymbol{\Sigma}_a$, and say why the mean of the probability after the outcome should equal today's. (The shift is $\boldsymbol{\Sigma}_a\mathbf{J}_a^{\mathsf{T}}\boldsymbol{\epsilon}_a$ with $\boldsymbol{\epsilon}_a \sim \mathcal{N}(\mathbf{0}, \mathbf{S}_a)$, so its covariance is $\boldsymbol{\Sigma}_a\mathbf{J}_a^{\mathsf{T}}\mathbf{S}_a\mathbf{J}_a\boldsymbol{\Sigma}_a$. Since $(\boldsymbol{\Sigma}^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a)\boldsymbol{\Sigma}\mathbf{J}_a^{\mathsf{T}} = \mathbf{J}_a^{\mathsf{T}}\mathbf{S}_a$, $\boldsymbol{\Sigma}_a\mathbf{J}_a^{\mathsf{T}} = \boldsymbol{\Sigma}\mathbf{J}_a^{\mathsf{T}}\mathbf{S}_a^{-1}$, and the covariance is $\boldsymbol{\Sigma}\mathbf{J}_a^{\mathsf{T}}\mathbf{S}_a^{-1}\mathbf{J}_a\boldsymbol{\Sigma} = \boldsymbol{\Sigma} - \boldsymbol{\Sigma}_a$ by the matrix-inversion lemma. Averaging $\mathcal{N}(\mathbf{z}; \mathbf{m}_a, \boldsymbol{\Sigma}_a)$ over $\mathbf{m}_a \sim \mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma} - \boldsymbol{\Sigma}_a)$ gives $\mathcal{N}(\mathbf{z}; \hat{\mathbf{z}}, \boldsymbol{\Sigma})$, so the probability after the outcome averages to today's. On the settlement draws the means run from 0.6300 to 0.6437, within 0.008 of 0.6380, today's probability on those draws.)

**Exercise 18.4.** *Application.* The design sheet orders the chiller tests at R$ 60 thousand. Compare them per real with the economizer's factory test. (By equation (18.14) the chiller tests

settle $10 \times 0.0550/60 = 0.0092$ per R$ 10 thousand and the economizer test $10 \times 0.2500/40 = 0.0625$, about seven times as much; the chiller tests are third of six, behind the economizer test with the chiller tests at 0.0445.)

**Exercise 18.5.** *Application.* The test crew points out that its pressure-drop test narrows the deciding band more than the test of chiller C2 and costs less. Explain from the gap's budget why it is still the worse buy for this decision. (The pressure-drop test's lines sit in the column of $l_k$ alone, equation (18.9), and $l_k$ carries 0.053 of the variance of the gap $D$ (Table 13.6); halving its spread narrows the pipe's end of the band, 11.6 % of the width, but cannot move the sign of the gap, which the economizer and the chillers carry. The test of C2 reads a log-ratio inside the gap's budget and settles the decision with probability 0.0067; the pressure-drop test settles it with probability 0.0000.)

# Chapter 19

# The Value of a Pilot

## 19.1 A pilot before the full build

The site's drawing carries a waterside economizer that was never built. The full build is that economizer for the whole site, a stated R$ 1 800 000 installed, which the capital recovery factor of 0.131 a year turns into R$ 235 800 a year. The design sheet values it the way it values everything, at the handbook. With the exchanger at its handbook conductance and the chillers at their fitted efficiency, the economizer takes over enough of the chillers' duty in the record year's months to save R$ 270 586 a year of chiller energy. The savings exceed the annualized cost, and the sheet's single line reads "build".

A vendor offers another way in: a one-hall pilot first, the economizer on one hall for a season, for a stated R$ 500 000, of which R$ 350 000 is hardware the full build would reuse. A season of the pilot's metered duty would read the exchanger's conductance, the one number on the sheet that no reading of the site has ever touched. The planner's question is the drill question, verbatim,

*Would a one-hall pilot change the decision on the full build?*

It is answered with the method of Chapter 4 on the site's one fit, the posterior that Chapter 9 built from the operating record, and with the value of information of Section 2.6.2. What is new is the question asked of the posterior. It is not a band, a binding load or a margin, but a decision, build or not, and the probability that a measurement not yet taken would reverse it, set against what that measurement costs.

**Jargon: waterside economizer, full build, pilot, approach, chiller energy saved, annualized cost, capital recovery factor, net annual value, breakeven conductance, factory test, normal–normal update, preposterior, decision flip, value of information (expected value of sample information, EVSI), expected value of perfect information (EVPI), capitalized value**

A *waterside economizer* cools the chilled-water loop's return against the outdoor wet bulb through a heat exchanger and a cooling tower, taking part of the loop's heat off the chillers without their compressors. The *full build* is the economizer for the whole site; the *pilot* is the same economizer on one hall for a season, whose metered duty reads the exchanger's conductance. The *approach* of the economizer is the difference between the loop's return temperature and the wet bulb, the temperature difference across which it works. The *chiller energy saved* is the electricity the chillers no longer draw for the duty the economizer takes. The *annualized cost* of an installed price is that price times the *capital recovery factor* $\rho$, the constant payment a year that repays it over its life at the stated rate. The *net annual value* of the build is the savings less the annualized cost. The *breakeven conductance* is the conductance at which the net annual value is zero. A *factory test* reads the exchanger's conductance on its maker's rig before it is installed. A *normal–normal update* is the posterior of a normal entry after a reading of it with normal error, itself normal, in closed form. The *preposterior* of a measurement is the distribution, before the measurement is taken, of the posterior it would leave. A *decision flip* is an outcome of the measurement after which the better action differs from the better action now. The *value of information*, the *expected value of sample information* (EVSI) in the framework, is the expected improvement of the decision that a measurement buys; the *expected value of perfect information* (EVPI) is the same for a measurement that reads the quantity exactly, an upper bound on the value of any measurement of it. The *capitalized value* of a value per year is its value once, the amount whose annualized cost it would repay: the value per year over $\rho$.

**Key idea**

A measurement is worth what it changes in a decision, not what it narrows in a band. A build decision is the sign of an expected net value, and a measurement moves that expectation by an amount not known until it is read; before it is read, the amount is normal, with the variance the measurement removes. The probability that the decision flips is the share of those outcomes that carry the expected net across zero, and the value of the measurement is the expected loss it avoids in them. When the quantity that decides the build carries no reading at all, both follow in closed form, and a perfect reading of it bounds what any measurement can be worth.

### 19.1.1 A decision read through its preposterior

The full build is a choice between two actions, build and do not build, and Chapter 9 chose between drawings by their expected loss, equation (2.23). Here the loss of not building is zero and the loss of building is the annualized cost less the savings, so the expected-loss rule becomes a sign. Write $W(\mathbf{z})$ for the net annual value of the build at the state $\mathbf{z}$, the savings $S(\mathbf{z})$ less the annualized cost $c_{\text{ann}}$. The rule is

$$W(\mathbf{z}) = S(\mathbf{z}) - c_{\text{ann}}, \qquad \bar{W} = \mathbb{E}[W(\mathbf{z}) \mid \mathbf{y}], \qquad \text{build if } \bar{W} > 0, \tag{19.1}$$

with $\bar{W}$ the expected net annual value on today's evidence, in R$ a year. The decision is not $W > 0$, which is unknown, but $\bar{W} > 0$, which is computed; the probability $\Pr(W > 0 \mid \mathbf{y})$ says how often the chosen action would turn out right.

A measurement $a$ with outcome $\tilde{y}_a$ adds a row, and after it the rule is applied again with the posterior that row leaves, $\bar{W}_a(\tilde{y}_a) = \mathbb{E}[W(\mathbf{z}) \mid \mathbf{y},\ \tilde{y}_a]$. Before the measurement is taken the

outcome is unknown, and so is $\bar{W}_a$; its distribution over the predictive of the outcome is the preposterior of the decision, the general form (2.16) with $u$ the expected net annual value after the outcome. Two quantities are read from it. The decision flips when $\bar{W}_a$ falls on the other side of zero from $\bar{W}$; and the value of the measurement is the expected value of the better action after it less the value of the better action now, the EVSI of equation (2.53) written for a gain instead of a loss,

$$P_{\text{flip}}(a) = \Pr\big(\operatorname{sign}\bar{W}_a \neq \operatorname{sign}\bar{W}\big), \qquad \text{EVSI}_a = \mathbb{E}\big[\max(\bar{W}_a,\ 0)\big] - \max(\bar{W},\ 0), \tag{19.2}$$

both expectations over the predictive of $\tilde{y}_a$. By the law of total expectation the average of $\bar{W}_a$ over the outcomes is $\bar{W}$. When the build is chosen now, $\bar{W} > 0$, the second relation becomes

$$\text{EVSI}_a = \mathbb{E}\big[\max(\bar{W}_a,\ 0) - \bar{W}_a\big] = \mathbb{E}\big[\max(-\bar{W}_a,\ 0)\big], \tag{19.3}$$

the expected size of the loss that the measurement avoids, counted only in the outcomes that flip the decision. An outcome that leaves the decision where it is adds nothing to the value, however much it narrows the band; and a measurement whose outcomes can never flip the decision is worth nothing for it. The value is a value a year; set against a price paid once, it is capitalized by the capital recovery factor,

$$V_a = \frac{\text{EVSI}_a}{\rho}, \qquad \text{the measurement pays if } V_a > K_a, \tag{19.4}$$

with $K_a$ its price. The same three relations, at a measurement that reads the deciding quantity exactly, give the expected value of perfect information, which no measurement of that quantity can exceed. The value of information is derived in general in *Reliability Engineering by Derivation*, Chapter 31; this chapter writes it for one build decision on the site.

**Principle 19.1.** *A build decision is the sign of the expected net annual value. A measurement's preposterior moves that expectation; the probability the decision flips is the share of the outcomes that move it across zero, and the measurement's value is the expected loss avoided in those outcomes, capitalized against its price.*

### 19.1.2 The site's economizer and its savings (step 0)

The plant is the site of Chapter 9, drawn in Figure 9.1: two halls on one chilled-water plant with three chillers, one pump, the shared supply pipe, the power feed, and the waterside economizer that is drawn but not built. Its relations are those of equation (9.3) and its state the eight entries of equation (9.4),

$$\mathbf{z} = (\kappa_{C1},\ \kappa_{C2},\ \kappa_{C3},\ d,\ g_e,\ c,\ l_k,\ l_e)^{\mathsf{T}}.$$

Of these relations the savings read two: the economizer's duty and the chillers' efficiency.

**The physical relations.** The savings are the record year's chiller energy saved, priced. Over the year they are the sum over the twelve months of the energy the chillers no longer draw, at the energy price $p$,

$$S = p \sum_{k=1}^{12} h\, \frac{Q_{e,k}}{\text{COP}_k},$$

with $h$ the hours in a month, $Q_{e,k}$ the economizer's duty in month $k$ in kW, and $\text{COP}_k$ the chillers' efficiency that month: by equation (9.3) the chiller power is the duty left to the chillers over their efficiency, so taking the duty $Q_{e,k}$ off them saves $Q_{e,k}/\text{COP}_k$ of power. Both terms

are relations of the site. The economizer's duty is its conductance times its approach, the chilled-water return $T_s + \Delta T_w$ less the month's wet bulb $T_k$, and the efficiency falls by $c_T$ per kelvin of wet bulb above 18°C,

$$Q_{e,k} = G_e^{\mathrm{hb}}\, \mathrm{e}^{g_e}\, (T_s + \Delta T_w - T_k), \qquad \mathrm{COP}_k = \mathrm{e}^{c}\big[1 - c_T\,(T_k - 18)\big],$$

the approach positive in every month, since the warmest month's wet bulb lies below the 24°C return. Substituting both into the sum, the two entries of the state leave it as one factor,

$$S(g_e, c) = \mathrm{e}^{g_e}\, S_0(c), \qquad S_0(c) = \mathrm{e}^{-c}\, p\, h\, G_e^{\mathrm{hb}} \sum_{k=1}^{12} \frac{T_s + \Delta T_w - T_k}{1 - c_T\,(T_k - 18)}, \tag{19.5}$$

in R$ a year, with $S_0(c)$ the savings at the handbook conductance. The pinned quantities are those of equation (9.3), $T_s + \Delta T_w = 24$°C, $G_e^{\mathrm{hb}} = 25\,\mathrm{kW/K}$ and $c_T = 0.02\,\mathrm{K}^{-1}$, with the stated $h = 730\,\mathrm{h}$ a month and $p = \mathrm{R\$}\,0.75$ per kWh, and the twelve wet bulbs $T_k$ of the record year, from 10.17°C in July and August to 19.83°C in January and February (Table 16.5). The savings are monotone in the conductance and multiplicative in it: the economizer at a fraction $\mathrm{e}^{g_e}$ of its handbook conductance saves that fraction of the handbook savings, in every month alike. Their gradient with respect to the state follows,

$$\frac{\partial S}{\partial g_e} = S, \qquad \frac{\partial S}{\partial c} = -S, \tag{19.6}$$

every other entry zero.

**Modeling choices at the graph.** The economizer is taken to run every hour of every month at the month's mean wet bulb, with its duty taken whole off the chillers: this makes each month one term of equation (19.5), and it fails where the wet bulb swings within the month, since the duty is linear in the approach but the economizer cannot take more than the loop's heat, or where the economizer is shut for maintenance. The duty does not depend on the halls' load as long as the loop's heat exceeds it, as it does at today's loads in every month: the savings are those of the record year at today's loads, and the growth of later chapters does not raise them. The record year stands for every year of the build's life: this is what lets one year's savings be set against the annualized cost, and it fails if the energy price or the climate moves over the fifteen years. And the savings are the chiller energy alone: the economizer's fans and pumps, its maintenance, and the capacity it adds on the design day are not counted (Section 19.6.5).

### 19.1.3 The build decision and the measurements

The questions are asked at a scenario node of Section 4.3.8 on which the economizer is built at full size, with the record year's months and the chilled-water return pinned. The build's annualized cost is the stated installed price times the capital recovery factor,

$$c_{\mathrm{ann}} = \rho\, K_{\mathrm{build}} = 0.131\,\mathrm{yr}^{-1} \times \mathrm{R\$}\,1\,800\,000 = \mathrm{R\$}\,235\,800\,\mathrm{yr}^{-1},$$

fifteen years at ten percent, as in equation (9.16). By equations (19.1) and (19.5) the net annual value is $W(g_e, c) = \mathrm{e}^{g_e} S_0(c) - c_{\mathrm{ann}}$, which is zero at the breakeven conductance $g_e^{\dagger}$; written with the savings averaged over the efficiency's posterior, $\bar{S}_0 = \mathbb{E}[S_0(c) \mid \mathbf{y}]$,

$$g_e^{\dagger} = \ln \frac{c_{\mathrm{ann}}}{\bar{S}_0}, \qquad \bar{W} = \bar{S}_0\, \mathbb{E}\big[\mathrm{e}^{g_e} \mid \mathbf{y}\big] - c_{\mathrm{ann}}, \tag{19.7}$$

the second because, by equation (16.7), no row couples the economizer's entry to the efficiency, so the expectation of the product is the product of the expectations. If $g_e$ is normal with mean $\mu$ and spread $s$, its exponential is log-normal and $\mathbb{E}[\mathrm{e}^{g_e}] = \mathrm{e}^{\mu + s^2/2}$, so

$$\bar{W}(\mu, s) = \bar{S}_0\, \mathrm{e}^{\mu + s^2/2} - c_{\mathrm{ann}}, \qquad \bar{W}(\mu, s) > 0 \iff \mu > g_e^{\dagger} - s^2/2. \tag{19.8}$$

The build is chosen when the mean of the conductance clears the breakeven, less a correction for its spread that the log-normal mean adds.

**The pilot and the factory test.** Each measurement is a reading of $g_e$ with normal error of stated spread $\sigma_a$: the pilot, a season of its metered duty with $\sigma_a = 0.08$ at a net cost of R\$ 500 000 − R\$ 350 000 = R\$ 150 000; the factory test of the exchanger of Chapter 18, with $\sigma_a = 0.10$ at R\$ 40 000; the two together, whose readings combine to one of spread $(0.08^{-2} + 0.10^{-2})^{-1/2} = 0.0625$; and a perfect reading, $\sigma_a \to 0$. With the prior $g_e \sim \mathcal{N}(0,\ \sigma_0^2)$ and the reading $\tilde{y}_a = g_e + \varepsilon_a$, the posterior is the normal–normal update,

$$g_e \mid \mathbf{y},\ \tilde{y}_a \sim \mathcal{N}(\mu_a,\ s_a^2), \qquad s_a^2 = \left(\sigma_0^{-2} + \sigma_a^{-2}\right)^{-1}, \qquad \mu_a = \frac{s_a^2}{\sigma_a^2}\,\tilde{y}_a, \tag{19.9}$$

the one-entry case of the preposterior covariance of equation (5.26), since the reading's row has its one entry in the column of $g_e$, uncorrelated with every other entry. The posterior spread $s_a$ is known before the reading; the posterior mean is not. Its predictive follows as in equation (6.11): the reading is $\tilde{y}_a \sim \mathcal{N}(0,\ \sigma_0^2 + \sigma_a^2)$ before it is taken, and $\mu_a$ is a fixed multiple of it, so

$$\mu_a \sim \mathcal{N}(0,\ \tau_a^2), \qquad \tau_a^2 = \sigma_0^2 - s_a^2, \tag{19.10}$$

the variance the reading removes, by the law of total variance. Each outcome gives the expected net annual value after it, $\bar{W}_a = \bar{W}(\mu_a, s_a)$ by equation (19.8), and with the build chosen now the decision flips when $\mu_a$ falls below $g_e^\dagger - s_a^2/2$,

$$P_{\text{flip}}(a) = \Phi\left(\frac{g_e^\dagger - s_a^2/2}{\tau_a}\right), \qquad \text{EVSI}_a = \mathbb{E}\Big[\max\big(c_{\text{ann}} - \bar{S}_0\,\mathrm{e}^{\mu_a + s_a^2/2},\ 0\big)\Big], \tag{19.11}$$

the second equation (19.3) with equation (19.8) inside it. At $\sigma_a \to 0$ the posterior spread vanishes, $\tau_a \to \sigma_0$, and $\mu_a$ is $g_e$ itself: the flip probability becomes $\Pr(g_e < g_e^\dagger \mid \mathbf{y})$, the probability that the build is the wrong call, and the value becomes the EVPI,

$$\text{EVPI} = \mathbb{E}\big[\max\big(c_{\text{ann}} - \mathrm{e}^{g_e}\,\bar{S}_0,\ 0\big) \mid \mathbf{y}\big], \tag{19.12}$$

the expected loss of building over the conductances below the breakeven.

### 19.1.4 The planner's questions (step 1)

The planner asks four questions of the build decision, and step 1 writes each over the site's state, with the evidence on the right of the bar. Table 19.1 lists them. The first column is the question in the planner's words, the drill question first; the second the general form of its category in the notation of Chapter 2; the third its mathematical form for the site; 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. All four are design questions. The second, the decision now, is an expected consequence; the first and the fourth are preposteriors of that decision; the third sets the value the first finds against a price. The atlas of Appendix A collects them with the other questions of Volume II.

### 19.1.5 The evidence as rows (step 2)

Step 2 attaches the site's operating record to the graph, the record of Chapter 9: twelve monthly windows of chiller power, pipe pressure drop, pump power and feed power, six capacity tests and eight priors, 62 rows on 8 entries. The savings read two entries of it. The efficiency $c$ is read by every month's chiller power, which by equation (9.3) is the loop's heat over the efficiency, since the economizer takes none of it. The economizer's log-ratio $g_e$ is read by no row: the economizer is not built, so no window sees its duty, and $g_e$ keeps the handbook's prior, $g_e \sim \mathcal{N}(0,\ 0.35^2)$, a conductance right to about a factor of $\mathrm{e}^{0.35}$.

Table 19.1: The questions asked of the economizer's full build. $\mathbf{y}$ is the site's operating record, its capacity tests and the priors; $W$ is the net annual value (19.1), $\bar{W}_a$ its expectation after the measurement $a$, $g_e^\dagger$ the breakeven conductance (19.7), and $V_a$ the capitalized value (19.4).

| Planner's question | General form | Mathematical form | Category | Computation | Answered in |
|---|---|---|---|---|---|
| Would a one-hall pilot change the decision on the full build? | preposterior of a decision, $\mathbb{E}_{\tilde{y}_a \mid \mathbf{y}}[u]$ | $\Pr(\operatorname{sign}\bar{W}_a \neq \operatorname{sign}\bar{W})$ for the pilot | design | preposterior | 19.3.2 |
| What is the decision now, and how sure? | $\arg\min_a \mathbb{E}[L(a,\mathbf{z}) \mid \mathbf{y}]$ | $\bar{W} = \mathbb{E}[W \mid \mathbf{y}]$; $\Pr(W>0 \mid \mathbf{y})$; $g_e^\dagger$ | design | expected consequence | 19.3.1 |
| Is the pilot worth its cost? | EVSI, $\bar{L}_{\text{now}} - \bar{L}_{\text{after}}$ | $V_a = \text{EVSI}_a/\rho$ against the pilot's net cost | design | expected consequence | 19.3.3 |
| Is there a cheaper way to the same answer? | preposterior, over candidates | flips and $V_a$ for the factory test, both, and a perfect reading (EVPI) | design | preposterior | 19.3.4 |

**Which row acts on which entry.** With $*$ for an entry that can be nonzero and the columns in the order of the state, the chiller-power rows of the record, the prior row of $g_e$ and the two candidate readings have the pattern

$$\begin{bmatrix} \mathbf{J}_{\text{ch}} \\ \mathbf{J}_{g_e,\text{prior}} \\ \mathbf{J}_{\text{pilot}} \\ \mathbf{J}_{\text{factory}} \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & * & 0 & 0 \\ 0 & 0 & 0 & 0 & 1/0.35 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1/0.08 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1/0.10 & 0 & 0 & 0 \end{bmatrix}, \qquad \mathbf{g}_S^{\mathsf{T}} = (0,\ 0,\ 0,\ 0,\ S,\ -S,\ 0,\ 0), \tag{19.13}$$

with the first line standing for the twelve chiller-power rows, each with an entry in the efficiency's column; the second line is not a row but the gradient (19.6) of the savings. Three facts are read here before any number. The savings' gradient lives in two columns, $g_e$ and $c$, with entries of equal size. The column of $g_e$ holds one row of today's evidence, its prior, so its posterior is its prior exactly, and the decision's uncertainty comes from it unless the record has left the efficiency nearly as uncertain. And each candidate adds one row in that column alone: the pilot and the factory test read nothing else, and they differ only in their spread.

**Modeling choices at the rows.** The two candidates are taken as unbiased normal readings of $g_e$ with the stated spreads: this is what makes the update (19.9) closed, and it fails if the pilot's hall differs from the site in its approach or its flow, so that one hall's conductance is not the site's, or if the maker's rig reads the exchanger at conditions the site never sees. The two readings are independent given $g_e$: this lets them combine to one of spread 0.0625, and it fails if both use the maker's own rating method. The prices are stated inputs: the pilot's R$ 500 000 with its R$ 350 000 of reused hardware, and the factory test's R$ 40 000; the book does not estimate them. The pilot is priced at its net cost on the understanding that its hardware becomes the first hall of the full build.

## 19.2 The posterior of the state (step 3)

All four questions of Table 19.1 are read from one posterior, the site's one fit of Chapter 9. No row is added and no fit is repeated; this section derives what the decision needs from that

posterior, cites the fit, and states the draws the answers use.

### 19.2.1 What the decision needs from the state

By equations (19.7) and (19.8) the decision needs two things: the posterior of $g_e$, and the average of $S_0(c)$ over the posterior of $c$. By the block structure (16.7) of the site's information matrix, $g_e$ sits in a block of its own,

$$\mathbf{\Lambda}_{g_e g_e} = 0.35^{-2}, \qquad \Sigma_{g_e g_e} = 0.35^2, \qquad \Sigma_{g_e j} = 0 \quad \text{for every other entry } j,$$

so its posterior is the normal prior exactly, with no Laplace step in between, and it is independent of $c$. The efficiency enters the savings through $\mathrm{e}^{-c}$, so by the gradient (19.6) the variance of $\ln S$ is $\Sigma_{g_e g_e} + \Sigma_{cc}$, the two entries' variances added, and each entry's share of it is its own variance over the sum. The decision itself is a function of $g_e$ through $\mathrm{e}^{g_e}$ and of $c$ only through the number $\bar{S}_0$.

### 19.2.2 The site's fit, cited

The MAP state $\hat{\mathbf{z}}$ and the Laplace covariance $\mathbf{\Sigma} = (\mathbf{J}^\mathsf{T}\mathbf{J})^{-1}$ at $\hat{\mathbf{z}}$ are those of Section 9.2.2, the Gauss–Newton fit of the operating record in the form of equation (4.38), with twice the objective 59.408 on 54 degrees of freedom. Table 9.2 lists the entries. The two the savings read are the efficiency, at a MAP value of 1.6176 with posterior spread 0.0061, about a twenty-fifth of its prior spread of 0.15, read by every month's chiller power; and the economizer's log-ratio, at 0 with spread 0.35, its prior, untouched by the record. By the variance split of Section 19.2.1 the efficiency's share of the variance of $\ln S$ is $0.0061^2/(0.35^2 + 0.0061^2)$, about three parts in ten thousand: the record has fixed the chillers' efficiency, and the savings' uncertainty is the economizer's guess alone.

### 19.2.3 The draws

The answers read the efficiency's posterior through the first 2 000 of the site's draws of Section 9.2.4, states from the Laplace normal $\mathcal{N}(\hat{\mathbf{z}}, \mathbf{\Sigma})$ drawn with the seed 9, whose weights toward the exact posterior have an effective sample size of 19 913 of 20 000; $\bar{S}_0$ and every savings figure below is the mean over those draws of $c$. The conductance's posterior is its exact normal, and the expectations over it are read by Gauss–Hermite quadrature, which integrates a smooth function against a normal by a weighted sum at fixed nodes; the probability and the band of the decision now are read by 20 000 draws of $g_e$ from the same normal, drawn with the seed 190, at each of which the savings are averaged over the first 200 draws of $c$. The preposterior of each candidate reads 4 000 outcomes of its posterior mean from the normal (19.10), with the seeds 19 for the pilot and 191, 192 and 193 for the factory test, the two together and the perfect reading; each outcome's expected net annual value is read by quadrature over its posterior (19.9). The reported values are those of the state route, the expectations and shares over the posterior of the state, $g_e$ exact and $c$ by its draws.

## 19.3 The answers (step 4)

Every answer is read from the posterior of Section 19.2.3. The decision now is answered first, since the other three ask what a measurement would do to it; the drill question follows.

### 19.3.1 The decision now

**Planner's question.** What is the decision now, and how sure? The question asks whether to build on today's evidence, and how often that choice would turn out right.

Table 19.2: The economizer's savings against its conductance, by equation (19.5), averaged over the first 2 000 draws of the site's efficiency; the net annual value against the annualized cost $c_{\text{ann}} = \text{R\$ } 235\,800$ a year. Net values from the printed savings.

| $g_e$ [–] | Factor $e^{g_e}$ [–] | Savings $S$ [R$/yr] | Net $W$ [R$/yr] |
|---|---|---|---|
| −0.50 | 0.607 | 164 163 | −71 637 |
| −0.25 | 0.779 | 210 790 | −25 010 |
| −0.10 | 0.905 | 244 903 | +9 103 |
| 0.00 | 1.000 | 270 659 | +34 859 |
| +0.10 | 1.105 | 299 125 | +63 325 |
| +0.25 | 1.284 | 347 533 | +111 733 |
| +0.50 | 1.649 | 446 242 | +210 442 |

**General form.** A design question chooses the action of least expected loss, equation (2.23),

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

with the loss of an action its cost less what it returns.

**Specialized form.** The question is equation (2.23) with two actions, the loss of not building zero and the loss of building $-W(\mathbf{z})$, so that the build is chosen when $\mathcal{R}(\text{build}) < 0$, that is when $\bar{W} > 0$, equation (19.1). The expected net annual value is equation (19.8) at the prior, $\bar{W}(0,\ 0.35)$; the probability that the build is right is

$$\Pr(W > 0 \mid \mathbf{y}) = \Pr(g_e > g_e^\dagger \mid \mathbf{y}) = \Phi\left(\frac{-g_e^\dagger}{0.35}\right), \tag{19.14}$$

with the breakeven of equation (19.7); and the distribution of $W$ itself is read from the draws of $g_e$, with $c$ averaged over its draws.

**Building blocks.** The savings (19.5) at the draws of $c$; the prior of $g_e$; the stated price, the capital recovery factor and the energy price; the normal distribution function $\Phi$.

**The savings against the conductance.** Table 19.2 gives the savings at seven conductances. Each row is a value of $g_e$; the columns give the factor $e^{g_e}$ on the handbook conductance, the savings averaged over the draws of $c$, and the net annual value against the annualized cost of R\$ 235 800. At the handbook conductance the savings are R\$ 270 659 a year, the mean $\bar{S}_0$, against the sheet's R\$ 270 586 at the MAP efficiency alone; the two differ by the average of $e^{-c}$ over a spread of 0.0061. Down the column the savings are the handbook savings times the factor, as equation (19.5) requires: $e^{-0.5} \times 270\,659$ is the 164 163 of the first row.

The same savings are drawn in Figure 19.1. The upper panel puts the conductance's log-ratio on the horizontal axis and the savings on the vertical, in thousands of reais a year: the curve is $\bar{S}_0\, e^{g_e}$, the points are the seven rows of the table, the dashed line is the annualized cost and the red line marks the breakeven. The lower panel draws the prior of $g_e$ on the same horizontal axis, shaded red below the breakeven and green above it.

**The breakeven.** By equation (19.7) the breakeven conductance is

$$g_e^\dagger = \ln \frac{235\,800}{270\,659} = -0.1379, \qquad e^{g_e^\dagger} = 0.8712,$$

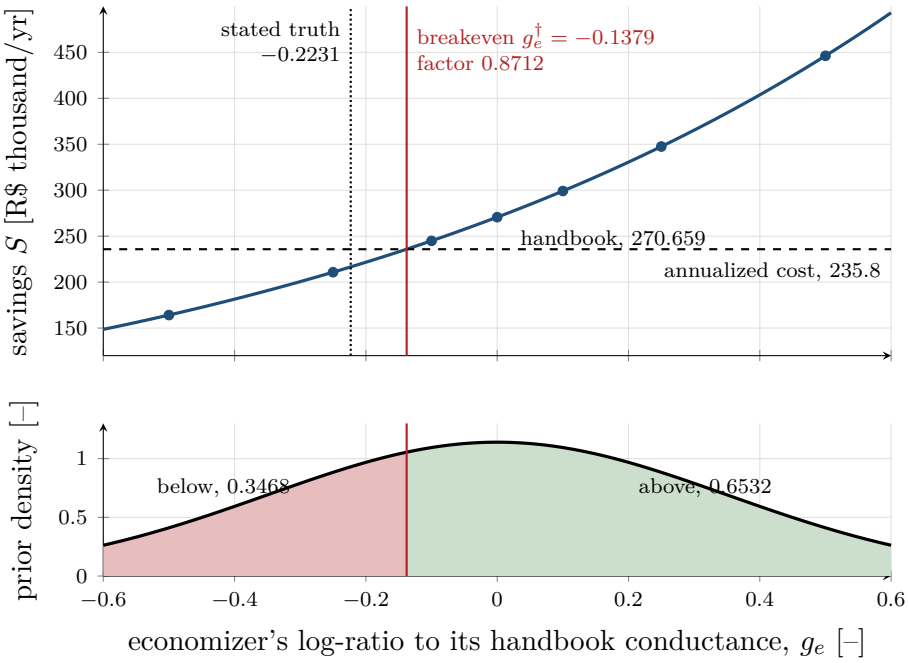



Figure 19.1: The economizer's savings and the conductance that decides them. Upper panel: the savings $\bar{S}_0\,\mathrm{e}^{g_e}$ of equation (19.5), averaged over the draws of the efficiency (curve), with the seven conductances of Table 19.2 (points), against the annualized cost of R$ 235 800 a year (dashed); the red line is the breakeven $g_e^\dagger = -0.1379$, a factor 0.8712 on the handbook conductance, and the dotted line the stated truth of the record, a check marker for Section 19.4.2. Lower panel: the handbook's prior $g_e \sim \mathcal{N}(0,\ 0.35^2)$, the posterior of $g_e$ today, with its mass below the breakeven (red) and above it (green).

Table 19.3: The decision now on the economizer's full build, by three routes. Quadrature: equation (19.8) at the prior, exact in $g_e$, over the draws of $c$; draws: 20 000 draws of $g_e$ from its prior (seed 190), each with the savings averaged over the first 200 draws of $c$, Monte Carlo standard error of $\Pr(W > 0)$ 0.0034; design sheet: the handbook conductance at the MAP efficiency. Gaps from the printed values.

| Route | $\bar{W}$ [R$/yr] | 5 % point [R$/yr] | 95 % point [R$/yr] | $\Pr(W > 0)$ [–] | Decision [–] |
|---|---|---|---|---|---|
| Quadrature, exact in $g_e$ (reported) | 51 955 | – | – | 0.6532 | build |
| Draws of $g_e$, $c$ averaged | 53 268 | −83 596 | 246 435 | 0.6586 | build |
| Gap, draws minus reported | +1 313 | – | – | +0.0054 | – |
| Design sheet, $g_e = 0$ | 34 786 | – | – | – | build |

and by equation (19.14) the probability that the conductance lies above it, so that the build is right, is

$$\Pr(W > 0 \mid \mathbf{y}) = 0.6532.$$

The economizer pays if its exchanger delivers at least 87 % of the handbook's conductance, and the handbook's own spread puts it below that about one time in three. In Figure 19.1 the breakeven lies 0.39 prior standard deviations below the handbook's guess: the build is favored, but not by much.

**The decision.** By equation (19.8) at the prior, read by quadrature, the expected net annual value is

$$\bar{W} = \text{R\$}\,51\,955\,\text{yr}^{-1} > 0 \quad \Longrightarrow \quad \text{build.} \tag{19.15}$$

Table 19.3 sets the routes side by side. Each row is a route; the columns give the expected net annual value, the ends of the 90 % band of the net value, the probability that it is positive, and the decision. The 20 000 draws of $g_e$, each with the savings averaged over 200 draws of $c$, put the mean at R$ 53 268, R$ 1 313 above the quadrature, the band at −R$ 83 596 to R$ 246 435, and the probability of a positive net at 0.6586, 0.0054 above equation (19.14), with a Monte Carlo standard error of 0.0034: the gap is 1.6 standard errors. The design sheet, at $g_e = 0$ and the MAP efficiency, has R$ 270 586 − R$ 235 800 = R$ 34 786 a year and no band.

The decision is to build, and it is a decision made with a one-in-three chance of losing. The band of the net value runs from a loss of R$ 83 596 a year to a gain of R$ 246 435: the build's outcome is mostly the economizer's conductance, which no reading has seen. The expected value, R$ 51 955, is larger than the sheet's R$ 34 786 at the same handbook guess, because the savings are convex in $g_e$: a conductance above the guess gains more than the same distance below it loses, and the log-normal mean $\mathrm{e}^{s^2/2}$ of equation (19.8) is that asymmetry. The draws and the quadrature agree on the decision and differ by their sampling error; the quadrature is exact in $g_e$ and is the reported value.

### 19.3.2 The pilot's preposterior

**Planner's question.** Would a one-hall pilot change the decision on the full build? The question asks how likely it is that the pilot's reading would turn today's "build" into "do not build".

**General form.** A design question about a measurement not yet taken is the preposterior expectation of a quantity of the posterior it would leave, equation (2.16),

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

here with $u$ the indicator that the decision after the outcome differs from the decision now.

**Specialized form.** The question is equation (2.16) with $a$ the pilot, $\sigma_a = 0.08$, and $u = \mathbf{1}[\bar{W}_a < 0]$, since the build is chosen now: the flip probability of equation (19.2), in closed form equation (19.11). It is read by 4 000 outcomes $\mu_a$ drawn from the predictive (19.10), each pushed through equation (19.8) by quadrature over the posterior spread (19.9),

$$\hat{P}_{\text{flip}} = \frac{1}{4\,000}\sum_{i=1}^{4\,000} \mathbf{1}[\bar{W}(\mu_{a,i},\ s_a) < 0], \qquad \text{se}(\hat{P}_{\text{flip}}) = \sqrt{\hat{P}_{\text{flip}}\,(1-\hat{P}_{\text{flip}})/4\,000}, \tag{19.16}$$

with se the Monte Carlo standard error.

**Building blocks.** The prior spread 0.35 and the pilot's 0.08; the update (19.9) and the predictive (19.10); $\bar{S}_0$ and $c_{\text{ann}}$ of Section 19.3.1.

**The pilot.** By equation (19.9) the pilot would leave the conductance with a posterior spread of

$$s_{\text{pilot}} = \big(0.35^{-2} + 0.08^{-2}\big)^{-1/2} = 0.0780,$$

under a quarter of today's 0.35: the pilot nearly settles the conductance. By equation (19.16) the probability that it flips the decision is

$$P_{\text{flip}}(\text{pilot}) = 0.3445, \tag{19.17}$$

with a Monte Carlo standard error of $\sqrt{0.3445 \times 0.6555/4\,000} = 0.0075$. Beside it, the closed form (19.11), with its correction $s_a^2/2$, gives 0.3398, a gap of +0.0047, under two thirds of the standard error. The answer is yes, about a third of the time.

The distribution of the expected net annual value after the pilot is drawn in Figure 19.2: the horizontal axis is $\bar{W}_{\text{pilot}}$ in thousands of reais a year, the vertical axis its preposterior density. The curve is the density of $\bar{S}_0\,\mathrm{e}^{\mu_a + s_a^2/2} - c_{\text{ann}}$ with $\mu_a$ from the normal (19.10), $\tau_a^2 = 0.35^2 - 0.0780^2$; the red area left of zero is the flip region, the outcomes after which the build is dropped; the dotted lines are the 5, 50 and 95 % points of the 4 000 outcomes, and the blue line today's $\bar{W}$.

Over the 4 000 outcomes the expected net annual value after the pilot has its 5, 50 and 95 % points at −R$ 80 773, R$ 33 831 and R$ 236 856 a year. The median outcome still builds, and it builds with a smaller expected net than today's R$ 51 955: the distribution is skewed to the right, so its median lies below its mean, and the mean is today's value, as the law of total expectation requires. The pilot does not change the expected net value; it spreads it out, and a third of that spread lies below zero. In those outcomes the pilot reads the exchanger low enough that the full build would lose money, and the planner who ran the pilot would not build.

### 19.3.3 The value of the pilot

**Planner's question.** Is the pilot worth its cost? The question asks whether what the pilot does to the decision is worth its net cost of R$ 150 000.

**General form.** The value of information is the reduction of the minimum expected loss that a measurement buys, the EVSI of equation (2.53), an instance of equation (2.16) with $u$ the minimum expected loss after the outcome,

$$\text{EVSI} = \bar{L}_{\text{now}} - \bar{L}_{\text{after}}.$$

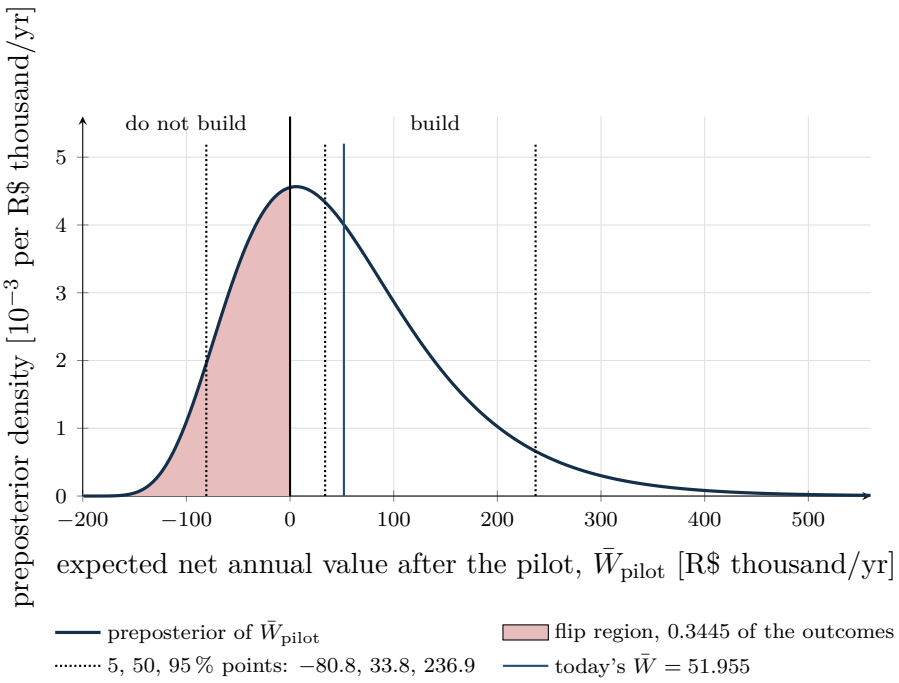



Figure 19.2: The preposterior of the expected net annual value of the full build after the one-hall pilot, by equations (19.8) and (19.10) with the pilot's posterior spread 0.0780. The density is that of a shifted log-normal, skewed to the right; the area left of zero (red) is the share of outcomes after which the build is dropped, 0.3445 by 4 000 outcomes. Dotted: the 5, 50 and 95 % points of the outcomes. Blue: the expected net annual value today, the mean of the distribution by the law of total expectation.

**Specialized form.** With the loss of building $-W$ and of not building zero, the minimum expected loss is $-\max(\bar{W}, 0)$, and the general form becomes the value of equation (19.2); it is a case of equation (2.16) with $u = \max(\bar{W}_a, 0)$, less its value now. With the build chosen now it is the expected loss avoided (19.3), read over the same 4 000 outcomes as the flip, and capitalized by equation (19.4) against $K_{\text{pilot}} = \text{R\$ } 150\,000$.

**Building blocks.** The outcomes $\bar{W}(\mu_{a,i},\ s_a)$ of Section 19.3.2; $\bar{W}$ of Section 19.3.1; the capital recovery factor $\rho = 0.131\,\text{yr}^{-1}$ and the pilot's stated net cost.

**The value.** By equation (19.3) the pilot's value of information is

$$\text{EVSI}_{\text{pilot}} = \text{R\$ } 14\,241\,\text{yr}^{-1}, \qquad V_{\text{pilot}} = \frac{14\,241}{0.131} = \text{R\$ } 108\,713 \ < \ \text{R\$ } 150\,000, \tag{19.18}$$

a ratio of 0.72 of its net cost. The pilot is not worth its net cost for this decision alone. Its value is the decision it changes, not the measurement it buys: in the third of the outcomes that flip, it saves the planner from a build that loses money, and the expected size of that saving, averaged over every outcome, is R\$ 14 241 a year. In the other two thirds it confirms a decision already taken, and confirmation is worth nothing to the decision. The pilot buys a conductance known to 0.0780 instead of 0.35; the decision needs far less than that.

### 19.3.4 A cheaper way to the same answer

**Planner's question.** Is there a cheaper way to the same answer? The question asks what the factory test of Chapter 18 does to the same decision, what the two together do, and what any measurement of the conductance could do at best.

**General form.** The preposterior expectation of equation (2.16), over a family of candidates,

$$U(a) = \mathbb{E}_{\tilde{y}_a|\mathbf{y}}\Big[\, u\big(p(\mathbf{z} \mid \mathbf{y}, \tilde{y}_a)\big)\Big], \qquad a \in \{\text{factory test},\ \text{pilot},\ \text{both},\ \text{perfect}\}.$$

**Specialized form.** The question is equation (2.16) for each candidate with $u$ the flip indicator and $u = \max(\bar{W}_a, 0)$, by equations (19.2) to (19.4), at $\sigma_a = 0.10$, 0.08, 0.0625 and $\sigma_a \to 0$; the last is the EVPI of equation (19.12). Each reads 4 000 outcomes of its own.

**Building blocks.** As in Section 19.3.3, with the candidates' spreads and stated prices: the factory test R\$ 40 000, the pilot R\$ 150 000 net, the two together R\$ 150 000 + R\$ 40 000 = R\$ 190 000, and no price for a perfect reading, which no measurement provides.

**The candidates.** Table 19.4 gives the four. Each row is a candidate; the columns give the spread of its reading, its price, the probability that it flips the decision, its value of information a year and capitalized, and the ratio of the capitalized value to the price. The same values are drawn in Figure 19.3, each candidate's capitalized value as a bar with its price beside it as a red diamond, on one axis in thousands of reais.

This is the surprise of the chapter. A perfect reading of the conductance, the most any measurement of it could do, is worth R\$ 19 030 a year, capitalized R\$ 145 265: still below the pilot's net cost of R\$ 150 000. No measurement of $g_e$ could be worth the pilot's price on this decision, however precise. In Figure 19.3 the pilot's diamond lies to the right of the dashed bound, and no bar can reach it. The factory test, at R\$ 40 000, flips the decision with probability 0.3407, against the pilot's 0.3445, and its value, R\$ 11 703 a year or R\$ 89 333 capitalized, is 2.23 times its price and about three fifths of what a perfect reading would be worth, $89\,333/145\,265 = 0.61$.

Table 19.4: The candidate measurements of the economizer's conductance before the full build: flip probability and value of information by 4 000 outcomes each, equations (19.2) to (19.4); capitalized at $\rho = 0.131\,\mathrm{yr}^{-1}$. Prices are stated inputs.

| Candidate | Spread $\sigma_a$ [–] | Price $K_a$ [R$] | $P_{\text{flip}}$ [–] | EVSI [R$/yr] | $V_a$ [R$] | $V_a/K_a$ [–] |
|---|---|---|---|---|---|---|
| Factory test | 0.10 | 40 000 | 0.3407 | 11 703 | 89 333 | 2.23 |
| One-hall pilot | 0.08 | 150 000 | 0.3445 | 14 241 | 108 713 | 0.72 |
| Both | 0.0625 | 190 000 | 0.3533 | 14 631 | 111 691 | – |
| Perfect reading (EVPI) | $\to 0$ | – | 0.3392 | 19 030 | 145 265 | – |

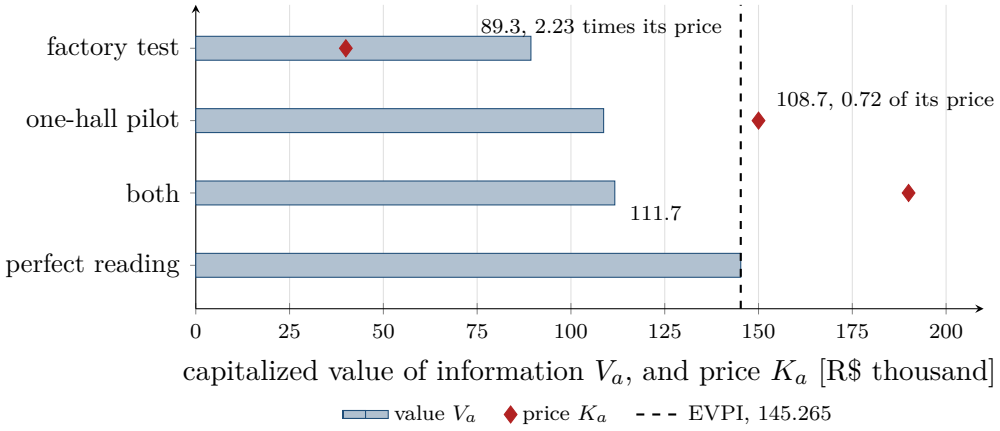



Figure 19.3: The capitalized value of information of each candidate measurement of the economizer's conductance (bars), against its price (diamonds), from Table 19.4. The dashed line is the value of a perfect reading, the EVPI, which bounds every bar. The factory test's bar lies far to the right of its price; the pilot's and the pair's fall short of theirs, and the pilot's price lies to the right of the bound itself.

Adding the pilot to the factory test raises the capitalized value from R$ 89 333 to R$ 111 691, by R$ 22 358 for R$ 150 000 more.

The flip probabilities barely move from candidate to candidate. The 4 000 outcomes give 0.3407 for the factory test, 0.3445 for the pilot, 0.3533 for the two, and 0.3392 for a perfect reading. Beside them, the closed form (19.11) gives 0.3360, 0.3398, 0.3425 and 0.3468, rising as the spread falls; the gaps, +0.0047, +0.0047, +0.0108 and −0.0076, are each within about one and a half Monte Carlo standard errors of 0.0075, and the outcomes' departure from the closed forms' order is that sampling error. Every candidate reads the conductance well enough to put it on the right side of the breakeven in most outcomes, since the breakeven lies 0.39 prior spreads from the guess while even the factory test's spread is 0.10. What a more precise reading adds is not more flips but better ones: the flips it adds or removes are near the breakeven, where the loss avoided is small, and the value climbs slowly, from R$ 11 703 to R$ 19 030 a year, as the spread falls from 0.10 to zero. The pilot pays only if it is built anyway, as the first hall of the full build.

## 19.4 The checks (step 5)

The answers of Section 19.3 rest on quadrature and draws over the site's posterior, on the closed forms of Section 19.1.3, and on a record generated from a stated truth; each is tested here.

### 19.4.1 The routes against each other

Four checks bear on the computation. First, the decision now by quadrature and by draws: R$ 51 955 against R$ 53 268 for the expected net, and 0.6532 against 0.6586 for the probability of a positive net (Table 19.3); the quadrature is exact in $g_e$, and the probability's gap of 0.0054 is 1.6 Monte Carlo standard errors of 0.0034, without disagreeing on the decision. Second, each candidate's flip probability against its closed form (Table 19.5). As $\sigma_a \to 0$, equation (19.11) gives exactly the probability that $g_e$ lies below the breakeven, $1 - 0.6532 = 0.3468$, and the 4 000 outcomes give 0.3392, a gap of 0.0076, about one Monte Carlo standard error of $\sqrt{0.3392 \times 0.6608/4\,000} = 0.0075$. At the other three spreads the closed form (19.11), with its correction $s_a^2/2$, gives 0.3360, 0.3398 and 0.3425, and the gaps are +0.0047, +0.0047 and +0.0108, the largest 1.4 standard errors. Third, the values must grow as the spread falls, since a reading of smaller spread is the same reading with more information added: R$ 11 703, R$ 14 241, R$ 14 631 and R$ 19 030 a year at the spreads 0.10, 0.08, 0.0625 and zero, in that order, and every value lies below the EVPI. Fourth, the pilot's outcomes average to today's value by the law of total expectation, which equations (19.8) and (19.10) give exactly, $\bar{S}_0\,\mathbb{E}[\mathrm{e}^{\mu_a + s_a^2/2}] = \bar{S}_0\,\mathrm{e}^{(\tau_a^2 + s_a^2)/2} = \bar{S}_0\,\mathrm{e}^{\sigma_0^2/2}$; the outcomes' median, R$ 33 831, lies below that mean of R$ 51 955 in Figure 19.2, as the right skew requires. The efficiency's draws are not reweighted toward the exact posterior here: by Section 19.2.2 the efficiency carries about three parts in ten thousand of the variance of the savings, and the site's draws lose almost nothing to the exact posterior, with an effective sample size of 19 913 of 20 000.

### 19.4.2 The answers against the stated truth

The operating record is generated from a stated truth, the last column of Table 9.2, and the last check compares the answers with it. The true economizer has $g_e = -0.2231$, a factor 0.8000 on the handbook conductance, below the breakeven factor of 0.8712; at that conductance the net annual value of the full build is −R$ 19 273, and the right decision is *do not build*. The handbook's "build" is the wrong call, and the decision now, which builds on the handbook's prior, is wrong with it: the truth lies in the red third of the lower panel of Figure 19.1.

Table 19.5: The checks of the build-decision answers. Reported values by quadrature in $g_e$ and draws of $c$; the truth is the stated truth of the record. The closed forms are equation (19.11), with its correction $s_a^2/2$.

| Quantity | Reported [– or R$/yr] | Beside [– or R$/yr] | Truth [– or R$/yr] | Reading |
|---|---|---|---|---|
| Economizer log-ratio $g_e$ | 0 | sd 0.35 | −0.2231 | 0.64 prior sd low |
| Breakeven $g_e^\dagger$ | −0.1379 | factor 0.8712 | factor 0.8000 | truth below the breakeven |
| Expected net annual value | 51 955 | draws 53 268 | −19 273 | inside [−83 596, 246 435] |
| $\Pr(W > 0)$ | 0.6532 | draws 0.6586 | event fails | gap 1.6 se of 0.0034; the decision now is wrong |
| Decision now | build | sheet: build | do not build | wrong call, one in three |
| $P_{\text{flip}}$, factory test | 0.3407 | closed form 0.3360 | – | gap +0.0047 |
| $P_{\text{flip}}$, pilot | 0.3445 | closed form 0.3398 | – | gap +0.0047 |
| $P_{\text{flip}}$, both | 0.3533 | closed form 0.3425 | – | gap +0.0108, 1.4 se |
| $P_{\text{flip}}$, perfect reading | 0.3392 | closed form 0.3468 | – | gap −0.0076, one se |
| Pilot flips, given the truth | – | – | about 0.835 | pilot changes the call |
| Draws | $c$: first 2 000 site draws, seed 9; decision now: 20 000 of $g_e$, seed 190, each over 200 of $c$; outcomes: 4 000 each, seeds 19, 191–193, se about 0.0075 | | | |

What would the pilot have done in that world? The pilot's reading is $\tilde{y} = g_e + \varepsilon$ with $g_e$ the truth, and the decision flips when the posterior mean $\mu_a = (s_a^2/\sigma_a^2)\,\tilde{y}$ falls below the breakeven. Since $s_a^2/\sigma_a^2 = 1/(1+\sigma_a^2/\sigma_0^2)$ by equation (19.9), that is when the reading falls below $g_e^\dagger(1+\sigma_a^2/\sigma_0^2)$, so

$$\Pr(\mu_a < g_e^\dagger \mid g_e) = \Phi\left(\frac{g_e^\dagger\,(1+\sigma_a^2/\sigma_0^2) - g_e}{\sigma_a}\right), \tag{19.19}$$

which at the truth gives about 0.835. The figure is read from the posterior mean against the breakeven and leaves out the small correction $s_a^2/2$ of equation (19.8), hence "about". In the world the record was built from, the pilot would change the decision about five times in six, and change it to the right one. Table 19.5 collects the checks: each row is a quantity, the columns its reported value, the route or reference beside it, the truth and the reading.

The answers hold where they can be checked. The truth's net value lies inside the band of the decision now, 0.64 of a prior spread below the guess; the decision now loses, which a decision made with probability 0.6532 of being right does one time in three. And the pilot, which in the posterior's average flips the decision a third of the time, in this world flips it about five times in six, because this world is one of the third in which the build is wrong. The probability that a measurement changes a decision is an average over worlds, and the value of the measurement lives in the worlds where it does.

## 19.5 The report (step 6)

The decision on today's evidence is to build the economizer, with an expected net of R$ 51 955 a year and a one-in-three chance that the build loses. The one-hall pilot would change that decision with probability 0.3445, and in the world the record was built from it would change it. But its value for the decision is R$ 108 713 capitalized, against a net cost of R$ 150 000, and even a perfect reading of the conductance would be worth only R$ 145 265. The R$ 40 000 factory test is worth 2.23 times its price.

Table 19.6: The report on the economizer's full build, at the record year's months and today's loads, priced at the stated R$ 1 800 000 and capital recovery factor 0.131 a year. Reported values by quadrature in $g_e$ and 4 000 outcomes per candidate.

| Question | Answer | Band or routes beside |
|---|---|---|
| Would a one-hall pilot change the decision on the full build? | yes, about a third of the time: $P_{\text{flip}} = 0.3445$ | se 0.0075; posterior sd of $g_e$ after it 0.0780; $\bar{W}$ after, 5/50/95 %: −80 773, 33 831, 236 856 R$/yr |
| What is the decision now, and how sure? | build: $\bar{W}$ = R$ 51 955 a year; $\Pr(W > 0) = 0.6532$ | draws 53 268, band −83 596 to 246 435 R$/yr, $\Pr(W > 0) = 0.6586$; breakeven factor 0.8712 |
| Is the pilot worth its cost? | no: EVSI R$ 14 241 a year, R$ 108 713 capitalized, against R$ 150 000 | ratio 0.72 |
| Is there a cheaper way to the same answer? | the factory test: flips 0.3407, R$ 89 333 capitalized, 2.23 times its R$ 40 000 | both: 0.3533, R$ 111 691; perfect (EVPI): 0.3392, R$ 145 265, below the pilot's net cost |

### 19.5.1 The answers and their bands

Table 19.6 gives the answers in the order of Table 19.1. Each row is a question; the second column gives the reported answer, and the third its band or the routes and candidates beside it.

### 19.5.2 The design-review entry

The design review's entry does not approve the full build on the sheet, and it does not buy the pilot. The sheet's "build" rests on a handbook conductance that no reading has touched; on the site's posterior the build is still favored, with an expected net of R$ 51 955 a year, but it loses money in a third of the conductances the handbook allows. A measurement of the conductance would change the decision about a third of the time, and the review buys the cheapest one that does so: the R$ 40 000 factory test of the exchanger, which flips the decision with probability 0.3407 and is worth R$ 89 333 capitalized, 2.23 times its price. The full build is approved only if the test's posterior puts the expected net above zero, which, to within the small spread correction of equation (19.8), is to say only if the test reads the exchanger above the breakeven conductance, a factor 0.871 on the handbook's. The one-hall pilot is recorded as not justified by this decision: its value for it, R$ 108 713 capitalized, falls short of its R$ 150 000 net cost, and no measurement of the conductance could reach that cost, since a perfect one is worth R$ 145 265. If the build goes ahead, the first hall may be built as a pilot, at no cost to the decision, since that hardware is spent either way.

**Common mistake**

Valuing a measurement by its precision. The pilot reads the conductance with a spread under a quarter of the handbook's, and its precision is not what the decision buys: it pays only in the outcomes that change the build, and there it is worth the loss it avoids. A measurement priced above the value of perfect information cannot pay for that decision, however good it is.

## 19.6 The next measurement (step 7)

The report leaves a build decision whose whole uncertainty is one handbook number, and Section 19.3.4 has already ranked the candidates that read it. Step 7 here is the order in which

they are bought.

### 19.6.1 The factory test first

The factory test is the first and, on this decision, the last measurement. It is the row $\mathbf{J}_{\text{factory}} = \mathbf{e}_{g_e}^{\mathsf{T}}/0.10$ of the pattern (19.13), and by Table 19.4 it flips the decision with probability 0.3407, for a value of R$ 89 333 capitalized against its R$ 40 000. After it, the full build is decided by the sign of its posterior's expected net, equation (19.8). Adding the pilot after the test would raise the capitalized value of the pair to R$ 111 691, R$ 22 358 above the test alone, for R$ 150 000 more: it does not pay. Nor does any measurement of the efficiency or of the chillers, which carry about three parts in ten thousand of the savings' variance; the other tests and surveys that Chapter 18 ranks on the existing plant decide other questions, not this one.

### 19.6.2 The pilot as the first hall

The pilot has one use that the decision does not price. If the factory test reads above the breakeven and the full build is approved, the first hall can be built and metered before the rest, at no added cost beyond the build's own, and a season of its duty reads the conductance with the pilot's spread of 0.08 before the other halls are committed. That is not a measurement before the decision but a staged build, and its value would be a question of its own: the value of stopping after one hall. The chapter does not compute it.

### 19.6.3 The recommendation

Buy the factory test of the exchanger for R$ 40 000. Build the full economizer only if the test's posterior puts the expected net above zero, the test reading above the breakeven factor of 0.871 to within the spread correction; do not build otherwise. Do not buy the one-hall pilot as a measurement: it is worth less than its net cost to this decision, and so would any reading of the conductance be. If the build proceeds, stage it hall by hall if the vendor allows, so that the pilot's information comes free.

### 19.6.4 What the analysis showed

The sheet's "build" is right on average and wrong a third of the time. The economizer saves its handbook savings times the ratio of its conductance to the handbook's, and that ratio is the one quantity on the site that no reading has seen. The breakeven lies at 87 % of the handbook, and the handbook's own spread puts the exchanger below it with probability 0.3468; the build is favored at an expected net of R$ 51 955 a year, more than the sheet's plug-in, because the savings are convex in the log-ratio. A one-hall pilot would flip the decision a third of the time, and in the world the record was built from it would flip it about five times in six. Yet the pilot is not worth its net cost: it pays only in the outcomes that flip the decision, and there it is worth the loss it avoids, which, averaged over every outcome, is R$ 14 241 a year. The bound on every such value is the value of a perfect reading, R$ 145 265 capitalized, and the pilot's R$ 150 000 lies above it. A factory test whose spread is a quarter larger than the pilot's does three fifths of what a perfect reading would, at about a quarter of the pilot's net cost. Precision is what a measurement sells; a change of decision is what a planner buys.

### 19.6.5 Limits of the model

**Chiller energy alone.** The savings are the chillers' electricity at a fixed price. The economizer's own fans and pumps, its maintenance, the water its tower uses, and the capacity it adds on the design day, $50\,\mathrm{e}^{g_e}$ kW by equation (9.5), 50 kW at the handbook, are not counted.

Chapter 9 counts that capacity in every drawing; it is not monetized here, and a build decision that counted it would move the breakeven down and the decision further toward building.

**One year for every year.** The record year's months stand for every year of the build's fifteen-year life, at today's energy price. A change in price or climate scales the savings, and with them the breakeven; neither is carried as an uncertainty.

**Monthly means.** Each month runs at its mean wet bulb for 730 h. The duty is linear in the approach, so the monthly mean gives the month's mean duty while the economizer never saturates; hours in which the approach is negative, or in which the economizer would carry the whole loop, are not modeled.

**Stated prices and precisions.** The installed price, the pilot's price and its reuse, the factory test's price, and the spreads 0.08 and 0.10 of the two readings are stated inputs, not estimates. The pilot's net cost takes the reuse as given; in the outcomes in which the pilot stops the full build, its hardware is not reused, and the one-hall economizer it leaves would save a share of the site's savings that the chapter does not count. The two move the pilot's true cost in opposite directions.

**One hall for the site.** The pilot's reading is taken as a reading of the site's conductance. A hall whose flow or approach differs from the site's would read a different conductance, and the spread 0.08 would have to carry that difference as well as the meter's error.

**The constructed record.** The operating record is generated from a stated truth with stated errors, not taken from a site, and the comparison with that truth in Section 19.4.2 is a check of the method on its own model, not of the model against a real plant. The truth was set so that the handbook's decision is wrong; a different truth would make the same decision right, and the pilot's value, an average over the posterior, would not change.

**The next part.** Part V has read what the uncertainty in the design inputs costs: the band of the site's cooling headroom split among the design parameters (Chapter 17), the measurements on the existing plant that would most narrow the build decision (Chapter 18), and here the value of a measurement read through the decision it can change. Part VI turns from the plant as built to the plant under failure, and Chapter 20 reads the headroom left when a unit is out.

## 19.7 What to retain

**Definition 19.1** (Value of a measurement before a build)**.** For a build decision taken on the sign of the expected net annual value $\bar{W} = \mathbb{E}[W(\mathbf{z}) \mid \mathbf{y}]$, a measurement $a$ flips the decision with probability $\Pr(\operatorname{sign}\bar{W}_a \neq \operatorname{sign}\bar{W})$ over the predictive of its outcome, and its value of information is $\mathrm{EVSI}_a = \mathbb{E}[\max(\bar{W}_a, 0)] - \max(\bar{W}, 0)$, the expected loss it avoids in the outcomes that flip the decision. Capitalized by the capital recovery factor, $V_a = \mathrm{EVSI}_a/\rho$, it is set against the measurement's price, and it never exceeds the value of a perfect reading of the deciding quantity.

Four conclusions carry beyond the economizer. A build decision is the sign of an expected net value, and its posterior says both the decision and how often it would lose; a design sheet that prints one number at the handbook hides the second. A measurement changes the decision only in the outcomes that carry the expected net across zero, and its value is the loss it avoids there: confirmation is worth nothing to the decision, and precision beyond what moves it is

paid for and not used. The value of a perfect reading bounds every measurement of the same quantity, and a measurement priced above it cannot pay for that decision, however precise. And when the deciding quantity carries no reading, the normal–normal update gives the preposterior in closed form: the posterior spread is known before the reading, and the posterior mean is normal with the variance the reading removes.

Table 19.7 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 19.7: The terms of the value of a pilot and where each is defined. Units for the site's economizer.

| Term | Expression | Units | Meaning | Defined in |
|---|---|---|---|---|
| Net annual value | $W = S - c_{\text{ann}}$; build if $\bar{W} > 0$ | R$/yr | the build decision | 19.1.1 |
| Flip probability | $\Pr(\operatorname{sign}\bar{W}_a \neq \operatorname{sign}\bar{W})$ | – | the measurement changes the call | 19.1.1 |
| Value of information | $\mathbb{E}[\max(\bar{W}_a, 0)] - \max(\bar{W}, 0)$ | R$/yr | expected loss avoided | 19.1.1 |
| Capitalized value | $V_a = \text{EVSI}_a/\rho$ | R$ | against a price paid once | 19.1.1 |
| *The site's economizer* | | | | |
| Savings | $S = \text{e}^{g_e} S_0(c)$ | R$/yr | chiller energy saved | 19.1.2 |
| Breakeven conductance | $g_e^\dagger = \ln(c_{\text{ann}}/\bar{S}_0)$ | – | net value zero | 19.1.3 |
| Expected net | $\bar{S}_0\,\text{e}^{\mu+s^2/2} - c_{\text{ann}}$ | R$/yr | log-normal mean | 19.1.3 |
| Normal–normal update | $s_a^2 = (\sigma_0^{-2} + \sigma_a^{-2})^{-1}$, $\mu_a = s_a^2 \tilde{y}_a/\sigma_a^2$ | – | posterior after a reading | 19.1.3 |
| Shift of the mean | $\mu_a \sim \mathcal{N}(0,\ \sigma_0^2 - s_a^2)$ | – | the preposterior | 19.1.3 |
| EVPI | $\mathbb{E}[\max(c_{\text{ann}} - \text{e}^{g_e}\bar{S}_0,\ 0)]$ | R$/yr | bound on every reading of $g_e$ | 19.1.3 |
| *Questions and answers* | | | | |
| Decision now | $\bar{W}$, $\Pr(W > 0) = \Phi(-g_e^\dagger/\sigma_0)$ | R$/yr, – | build, and how sure | 19.3.1 |
| Pilot's flips | share of outcomes with $\bar{W}_a < 0$ | – | the drill answer | 19.3.2 |
| Flip given the truth | $\Phi\big((g_e^\dagger(1 + \sigma_a^2/\sigma_0^2) - g_e)/\sigma_a\big)$ | – | from the posterior mean | 19.4.2 |

## 19.8 Exercises

**Exercise 19.1.** *Formulation.* Write the planner's question *should the economizer be built on today's evidence, and how often would that be the wrong call?* as an expected consequence and a probability over the site's state, say which entries they depend on, and give both values. (The build is chosen if $\bar{W} = \mathbb{E}[\text{e}^{g_e} S_0(c) - c_{\text{ann}} \mid \mathbf{y}] > 0$, and it is wrong with probability $\Pr(W < 0 \mid \mathbf{y}) = \Pr(g_e < g_e^\dagger \mid \mathbf{y})$. They depend on $g_e$, whose posterior is its prior, and on $c$ through $\bar{S}_0$ alone. The expected net is R$ 51 955 a year, so the build is chosen, and it is the wrong call with probability $1 - 0.6532 = 0.3468$.)

**Exercise 19.2.** *Derivation.* Show that the savings (19.5) are the handbook savings times $\text{e}^{g_e}$, derive the breakeven conductance, and give its value and factor. (The duty is $G_e^{\text{hb}}\text{e}^{g_e}$ times the approach in every month, and the efficiency does not depend on $g_e$, so $\text{e}^{g_e}$ factors out of the sum; the net value $\text{e}^{g_e}\bar{S}_0 - c_{\text{ann}}$ is zero at $g_e^\dagger = \ln(c_{\text{ann}}/\bar{S}_0) = \ln(235\,800/270\,659) = -0.1379$, a factor 0.8712 on the handbook conductance.)

**Exercise 19.3.** *Derivation.* Derive the predictive of the posterior mean after a reading of spread $\sigma_a$ of a normal entry with prior $\mathcal{N}(0,\ \sigma_0^2)$, and give the pilot's posterior spread. (The posterior mean is $\mu_a = (s_a^2/\sigma_a^2)\,\tilde{y}_a$ with $s_a^2 = (\sigma_0^{-2} + \sigma_a^{-2})^{-1}$, and before the reading $\tilde{y}_a \sim \mathcal{N}(0,\ \sigma_0^2 + \sigma_a^2)$; so $\mu_a$ is normal with mean zero. From $s_a^2 = \sigma_0^2\sigma_a^2/(\sigma_0^2 + \sigma_a^2)$ the multiplier is $s_a^2/\sigma_a^2 = \sigma_0^2/(\sigma_0^2 + \sigma_a^2)$,

and the variance is $\sigma_0^4/(\sigma_0^2+\sigma_a^2) = \sigma_0^2 - \sigma_0^2\sigma_a^2/(\sigma_0^2+\sigma_a^2) = \sigma_0^2 - s_a^2$, the variance the reading removes. For the pilot, $s_a = (0.35^{-2} + 0.08^{-2})^{-1/2} = 0.0780$.)

**Exercise 19.4.** *Application.* The vendor argues that the pilot is worth buying because it changes the decision a third of the time. Explain from the value of information why that is not enough, and give the pilot's value against its net cost. (Changing the decision is necessary but not sufficient: the pilot is worth the loss it avoids in the outcomes that flip, averaged over all outcomes, and the flips it causes are mostly near the breakeven, where the loss is small. By equation (19.3) its value is R$ 14 241 a year, R$ 108 713 capitalized at 0.131 a year, against a net cost of R$ 150 000: a ratio of 0.72.)

**Exercise 19.5.** *Application.* A reviewer asks for a more precise pilot, a full year of metering that would read the conductance almost exactly, at the same net cost. Say whether it could pay for the decision, and name the measurement to buy instead. (No reading of the conductance can be worth more than a perfect one, whose value is R$ 19 030 a year, R$ 145 265 capitalized, below the R$ 150 000 net cost; so no pilot at that cost can pay for this decision, however precise. The factory test of the exchanger is worth R$ 89 333 capitalized, 2.23 times its R$ 40 000.)

# Part VI

# Design Under Failure

# Chapter 20

# Headroom with the Largest Unit Out

## 20.1 Hall B with its largest unit out

Hall B's drawing has three cooling units, and the design sheet gives each a duty of 180 kW. The hall's design load is 340 kW of IT load, and its fans and lighting add 10 kW of other heat, 350 kW in all. The sheet's redundancy check takes one unit away. Two units leave $2 \times 180 = 360$ kW of duty against the 350 kW, a ratio of 1.029, and the sheet marks the hall N+1. It is the check Chapter 9's sheet made for the chillers: count the units, multiply by the duty, compare with the load.

The sheet reads coil duty. The limit the hall is held to is 27°C at every rack inlet, and that limit is not in the sheet's arithmetic. A coil's duty is a rating at a stated return-air temperature; an inlet's temperature depends on how warm the return air becomes when fewer coils carry the heat, on how much air the running fans move, and on how much of that return air each rack draws back. The planner's question is the drill question, verbatim,

*With the largest unit out, what is the probability the remaining plant holds the inlets at the design load?*

It is answered with the method of Chapter 4 on hall B's fit of Chapter 5, at the scenario node of Section 4.3.8 with one unit removed. Two things are new. The scenario removes a unit whose identity is itself uncertain: the largest unit is the one with the largest as-built conductance, and the fit knows the three conductances only to a band. And the answer is an event over five inlets at once, whose probability is not the product of the five racks' probabilities.

**Jargon: sheet duty, N+1 redundancy (single-unit fault tolerance), unit out, largest unit, worst unit, unit-out binding load, joint tail, marginal tail, product of the marginals, supply reset, chillers' capacity factor, crossing year**

The *sheet duty* of a cooling unit is the heat its coil is rated to pass at the rating's air and water temperatures. A plant is *N+1* on an element when it carries its load with any one of its units of that element out; in reliability terms it is tolerant of a single unit's fault, *single-unit fault tolerance*. A *unit out* is a unit whose coil passes no heat and whose fan moves no air. The *largest unit* is the unit with the largest as-built conductance; the *worst unit* is the unit whose loss leaves the smallest binding load. The *unit-out binding load* is the binding load of Chapter 5 with that unit out. The *joint tail* of several limits is the probability that every one of them holds at once; a *marginal tail* is that probability for one limit alone, and the *product of the marginals* is the joint tail the limits would have if they were independent. A *supply reset* lowers the chilled-water supply temperature to restore the inlets; the *chillers' capacity factor* is the fraction of their design-day capacity the chillers keep at the lower supply. The *crossing year* is the year in which a growing load meets a binding load.

**Key idea**

Redundancy is a property of the inlets, not of the duties. Removing a unit removes its conductance from the hall balance and its air from the mixing, and every inlet then rises faster with the load by a factor the hall's equations give in closed form. The largest unit is a posterior quantity read on each draw of the state, and because the three units move the same air, the largest unit out is the worst unit out. The answer is a joint tail over the five inlets, which share the running units' conductance: the product of the racks' tails understates it, the smallest of them bounds it, and near the answer one rack decides it.

### 20.1.1 Redundancy read at the inlets

Chapter 9 counted redundancy by state enumeration: each unit up or down, each state with its probability and its capacity. The planner's question here fixes the state, one unit down, and asks for the probability over the plant's parameters that the state carries the load. Write $R$ for the set of units running and $Q^{\dagger}(\mathbf{z};\ R)$ for the hall's binding load with those units, the smallest over the racks of the binding loads of Section 5.1.1. With unit $u$ removed from the full set $U$, the unit-out binding load is $Q^{\dagger}(\mathbf{z};\ U \setminus \{u\})$, and the hall is N+1 at the design load $Q_d$, at the state $\mathbf{z}$, when every unit's removal leaves it at or above that load. The largest unit is the one with the largest as-built conductance $G_u(\mathbf{z})$, and the worst unit the one whose removal leaves the smallest binding load,

$$u^{\star}(\mathbf{z}) = \arg\max_{u \in U} G_u(\mathbf{z}), \qquad u_{\mathrm{w}}(\mathbf{z}) = \arg\min_{u \in U} Q^{\dagger}(\mathbf{z};\ U \setminus \{u\}), \qquad Q_L^{\dagger}(\mathbf{z}) = Q^{\dagger}(\mathbf{z};\ U \setminus \{u^{\star}(\mathbf{z})\}), \tag{20.1}$$

with $Q_L^{\dagger}$ the binding load with the largest unit out, in kW. Both $u^{\star}$ and $u_{\mathrm{w}}$ are functions of the state, so on the posterior each is a probability over the units, not a label. The planner's question asks for the posterior probability that $Q_L^{\dagger}$ is at least the design load. Since every inlet at the scenario rises with the load, the event that the binding load is at least $Q_d$ is the event that every inlet holds at $Q_d$, and the probability is a joint tail over the racks.

> **Principle 20.1.** *A plant is N+1 at a load when the scenario with any one unit out holds every limit at that load. With the units' ratings uncertain, which unit is largest is read on each draw of the state, and the probability that the plant is N+1 is the posterior mass of the states in which the scenario holds.*

### 20.1.2 Hall B and its units (step 0)

The plant is hall B of Chapter 5, drawn in Figure 5.1: three cooling coils B1–B3 fed by the chilled water at the supply temperature $T_s$, five racks B1–B5 drawing their inlet air from the cold aisle $T_c$ with recirculation fractions $\beta_k$, the return air $T_r$, and the inlets of racks B2 and B4 measured. Its physical rows are the ten of equation (5.2) and its reduced state the ten free entries of equation (5.4),

$$\mathbf{z} = \left(T_s,\ Q_B,\ g_{B1},\ g_{B2},\ g_{B3},\ \beta_{B1},\ \beta_{B2},\ \beta_{B3},\ \beta_{B4},\ \beta_{B5}\right)^{\mathsf{T}},$$

with the pinned ratings $\bar{G}_u = 7\,\mathrm{kW/K}$, the air capacity rates $C_u = 10\,\mathrm{kW/K}$ and the other heat $Q_o = 10\,\mathrm{kW}$. Unit $u$'s as-built conductance is $G_u = \bar{G}_u\,\mathrm{e}^{g_u}$, and since the three ratings are equal, the largest unit is the one with the largest log-derate. Chapter 15 put these units into the site's enumeration of outage states at the MAP state of this fit (Table 15.4); here a unit out is read on the whole posterior, at the inlets, with the identity of the largest unit uncertain.

**Modeling choices at the graph.** A unit out loses its coil and its fan together: this removes its conductance from the hall balance and its air from the cold-aisle mixing, and it fails if the stopped unit's fan keeps turning, or if return air flows backward through the stopped unit into the cold aisle. The running units keep their fixed fan speed: this keeps their air capacity rates at 10 kW/K each, and it fails if a controller speeds the remaining fans up when a unit stops. The racks' fractions keep their calibrated values with a unit out: the cold aisle is one mixed volume, so which unit is out matters only through its conductance and its air, and this fails if the loss of one unit changes the airflow pattern at the racks nearest it. These three choices are taken up in Section 20.6.4.

### 20.1.3 The scenario node with a unit out

The questions are asked at the scenario node of Section 4.3.8. Three inputs are pinned. The supply is pinned at the drawing's design-day value $T_s^* = 14.0^\circ\mathrm{C}$; the load is pinned at the design load $Q^* = Q_d = 340\,\mathrm{kW}$, or left free when the binding load is solved for; and the running units are the set $R$, all three less the one out. At these pins the rows of equation (5.2) keep their form, with every sum over the running units only and the down unit's air-side row and discharge temperature removed,

$$\begin{aligned}
r_{\mathrm{h}}^* &= \textstyle\sum_{u\in R} G_u\,(T_r^* - T_s^*) - Q^* - Q_o && && \text{hall balance,}\\
r_{\mathrm{a},u}^* &= C_u\,(T_r^* - T_{d,u}^*) - G_u\,(T_r^* - T_s^*), && u \in R && \text{air side of each running unit,}\\
r_{\mathrm{m}}^* &= \textstyle\sum_{u\in R} C_u\,T_c^* - \sum_{u\in R} C_u\,T_{d,u}^* && && \text{cold-aisle mixing,}\\
r_{\mathrm{i},k}^* &= T_k^* - T_c^* - \beta_k\,(T_r^* - T_c^*), && k = \mathrm{B1},\ldots,\mathrm{B5} && \text{rack inlet } k,
\end{aligned} \tag{20.2}$$

nine rows on the nine scenario temperatures, each zero when it holds. Write the running units' summed conductance and summed air capacity rate as

$$S_G(\mathbf{z};\ R) = \sum_{u\in R} \bar{G}_u\,\mathrm{e}^{g_u}, \qquad S_C(R) = \sum_{u\in R} C_u,$$

in kW/K; with one unit out, $S_C = 20\,\mathrm{kW/K}$ whichever unit it is.

**The inlet rise per kilowatt.** The quantity the questions need is the rise of rack $k$'s inlet per kilowatt of hall heat at the scenario, $a_k(\mathbf{z};\ R) = \partial T_k^*/\partial Q^*$, in K/kW. The inlet row of equation (20.2) writes the inlet's rise over the supply through two temperature differences, the return air's rise over the supply and its rise over the cold aisle,

$$T_k^* - T_s^* = (T_r^* - T_s^*) - (1 - \beta_k)\,(T_r^* - T_c^*).$$

The first difference is set by the hall balance, whose sum now runs over the running coils alone,

$$S_G\,(T_r^* - T_s^*) = Q^* + Q_o \quad \Longrightarrow \quad T_r^* - T_s^* = \frac{Q^* + Q_o}{S_G}.$$

The second is set by the air. Summing the running units' air-side rows and substituting the hall balance gives $\sum_{u\in R} C_u\,(T_r^* - T_{d,u}^*) = Q^* + Q_o$, and the mixing row replaces the discharge temperatures' weighted sum by the cold aisle's,

$$S_C\,(T_r^* - T_c^*) = Q^* + Q_o \quad \Longrightarrow \quad T_r^* - T_c^* = \frac{Q^* + Q_o}{S_C},$$

the air balance (5.6) over the running units. Substituting the two differences into the inlet's rise,

$$T_k^*(\mathbf{z},\ Q^*;\ R) = T_s^* + a_k(\mathbf{z};\ R)\,(Q^* + Q_o), \qquad a_k(\mathbf{z};\ R) = \frac{1}{S_G} - \frac{1-\beta_k}{S_C} = \frac{1-(1-\beta_k)\,S_G/S_C}{S_G}, \tag{20.3}$$

every inlet linear in the heat, with $a_k$ the rise per kilowatt of equation (5.7) written over the running units. The second form reads the loss of a unit. The bracket depends on the ratio of the running units' conductance to their air, $S_G/S_C$, which is 0.7 at the rating whether two units run or three, since every unit has $\bar{G}_u/C_u = 0.7$; the factor in front is $1/S_G$, which grows by a half when one of three equal coils stops. At equal units, then, every inlet rises half again as fast with the load when one unit is out, and the binding load plus the other heat falls to two thirds of its value with all units running.

**The binding load and its gradient.** The binding load of rack $k$ with the units $R$ running solves $T_k^* = T_{\max}$ with $T_{\max} = 27°\mathrm{C}$, as in equation (5.8),

$$Q_k^\dagger(\mathbf{z};\ R) = \frac{T_{\max} - T_s^*}{a_k(\mathbf{z};\ R)} - Q_o, \qquad Q^\dagger(\mathbf{z};\ R) = \min_k Q_k^\dagger(\mathbf{z};\ R), \tag{20.4}$$

in kW, and by the implicit-function form (5.1) its gradient is that of equation (5.9) with the sums over $R$,

$$\frac{\partial Q_k^\dagger}{\partial g_u} = \frac{(Q_k^\dagger + Q_o)\,G_u}{a_k\,S_G^2} \quad (u \in R), \qquad \frac{\partial Q_k^\dagger}{\partial g_u} = 0 \quad (u \notin R), \qquad \frac{\partial Q_k^\dagger}{\partial \beta_k} = -\frac{Q_k^\dagger + Q_o}{a_k\,S_C}, \tag{20.5}$$

every other entry zero. The down unit's log-derate leaves the gradient: with unit B2 out, the binding load does not depend on how good B2's coil is.

**The largest unit out is the worst.** Two facts of equation (20.3) settle the order of the units and of the racks before any number. First, with one unit out $S_C = 20\,\mathrm{kW/K}$ whichever unit it is, so $a_k$ depends on which unit is out only through $S_G$, and $\partial a_k/\partial S_G = -1/S_G^2 < 0$: the unit whose loss leaves the smallest $S_G$, the largest unit, raises every rack's rise the most, and so $u_{\mathrm{w}}(\mathbf{z}) = u^\star(\mathbf{z})$ at every state. Second, at any one state all five racks share $S_G$ and $S_C$, so they bind in the order of their fractions whichever units run, the largest fraction first: the binding rack with a unit out is the binding rack with all units running.

### 20.1.4 The planner's questions (step 1)

The planner asks four questions of hall B with a unit out, and step 1 writes each as a probability over the reduced state, with the readings and the priors of window 1 on the right of the bar. Table 20.1 lists them. The first column is the question in the planner's words, the drill question first; the second the general form of its category in the notation of Chapter 2; the third its

Table 20.1: The questions asked of hall B with a unit out. $\mathbf{y}$ is the eight readings and nine priors of window 1; $Q_L^{\dagger}$ is the binding load with the largest unit out, equation (20.1); $T_k^*$ the scenario inlets of equation (20.3); $Q_d = 340\,\mathrm{kW}$.

| Planner's question | General form | Mathematical form | Category | Computation | Answered in |
|---|---|---|---|---|---|
| With the largest unit out, what is the probability the remaining plant holds the inlets at the design load? | $\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y})$ | $\Pr(Q_L^{\dagger} \geq Q_d \mid \mathbf{y})$, $u^{\star} = \arg\max_u G_u$ out, $T_s^* = 14^{\circ}\mathrm{C}$ | prediction under a threshold | posterior | 20.3.1 |
| Is the joint tail the product of the racks' tails? | joint against marginal tails | $\Pr(\cap_k\{T_k^* \leq 27^{\circ}\mathrm{C}\})$ against $\prod_k \Pr(T_k^* \leq 27^{\circ}\mathrm{C})$ | prediction under a threshold | posterior | 20.3.2 |
| What supply would restore it, and what does that cost the chillers? | $\Pr(f(\mathbf{z}) \leq f^* \mid \mathbf{y})$, quantiles | $T_s^{\dagger} = T_{\max} - \max_k a_k\,(Q_d + Q_o)$; the chillers' binding load at $T_s^{\dagger}$ | design | posterior | 20.3.3 |
| What does a fourth unit buy? | $\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y})$ under an alternative | the first question with four units, the largest of four out | design | posterior | 20.3.4 |

mathematical form for hall B; 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 predictions under a threshold: the drill question itself, and how its joint tail relates to the racks' own. The last two are design questions answered at a scenario: the supply that would restore the inlets, read with its cost to the chillers, and a fourth unit. The atlas of Appendix A collects them with the other questions of Volume II.

### 20.1.5 The evidence as rows (step 2)

The evidence is that of Chapter 5: window 1 of the night record at 19:00, the healthy plant, with its eight readings (the supply, the return air, the three discharge temperatures, the hall load meter and the inlets of racks B2 and B4) and its nine priors (the supply's set point, the commissioning record's $g_u \sim \mathcal{N}(0,\ 0.15^2)$ for each coil, and the airflow survey's fractions). The seventeen rows are those of the residual vector (5.11). No row is added: the unit out is a scenario, not a reading, and the record holds no window with a unit stopped.

**Which row acts on which entry.** With $*$ for an entry that can be nonzero and the columns in the order of the state, the gradient (20.5) of rack B3's binding load with unit B2 out, and the row a 0.3 K inlet sensor at rack B3 would add, have the patterns

$$\begin{aligned} \mathbf{g}_{\mathrm{B3},-\mathrm{B2}}^{\mathsf{T}} &= (0,\ 0,\ *,\ 0,\ *,\ 0,\ 0,\ *,\ 0,\ 0), \\ \mathbf{J}_{\mathrm{B3}} &= (*,\ *,\ *,\ *,\ *,\ 0,\ 0,\ *,\ 0,\ 0)/0.3, \end{aligned} \tag{20.6}$$

neither of them a row of today's evidence. Three facts are read here before any number. The unit-out gradient has no entry in the down unit's column, nor in the supply's or the load's, which the scenario pins. Its one fraction entry is $\beta_{B3}$'s, whose column holds only its survey prior (equation (5.10)), so the band of the unit-out binding load rests on the same unseen airflow number as the band with all units running. And a sensor at rack B3 would be the only row with an entry in that column.

Table 20.2: The entries of hall B's state that the unit-out answers read, from Chapter 5's fit of window 1, with the as-built conductances at the MAP state and the probability that each unit is the largest, by the Laplace draws and by the exact posterior.

| Entry | $\hat{z}_i$ [–] | $\sqrt{\Sigma_{ii}}$ [–] | Truth [–] | $G_u(\hat{\mathbf{z}})$ [kW/K] | $\Pr(u^\star = u)$, draws [–] | $\Pr(u^\star = u)$, exact [–] |
|---|---|---|---|---|---|---|
| $g_{B1}$ | −0.0151 | 0.0285 | 0.0300 | 6.895 | 0.3468 | 0.3472 |
| $g_{B2}$ | −0.0018 | 0.0282 | 0.0000 | 6.987 | 0.5761 | 0.5776 |
| $g_{B3}$ | −0.0479 | 0.0293 | −0.0400 | 6.672 | 0.0771 | 0.0752 |
| $\beta_{B1}$ | 0.2000 | 0.0800 | 0.1700 | – | – | – |
| $\beta_{B2}$ | 0.3168 | 0.0320 | 0.3400 | – | – | – |
| $\beta_{B3}$ | 0.5000 | 0.0800 | 0.5700 | – | – | – |
| $\beta_{B4}$ | 0.3670 | 0.0320 | 0.3800 | – | – | – |
| $\beta_{B5}$ | 0.3500 | 0.0800 | 0.4000 | – | – | – |

**Modeling choices at the rows.** Three stated design inputs enter the questions and are not rows. The sheet's duty of 180 kW per unit is printed to state the sheet's check and enters no computation. A fourth hall-B unit is the existing model, rated 7 kW/K with an air capacity rate of 10 kW/K, and its as-built log-ratio $g_{\text{new}}$ carries the commissioning prior of an unmeasured coil, $\mathcal{N}(0,\ 0.15^2)$, independent of the hall's state. And the chillers' capacity changes by a stated 3 % per kelvin of supply below or above 14°C, the slope pinned in the site's plant model. The chillers' as-built capacities are read from the site's fit of Chapter 9, which shares no row with hall B's window, so the two posteriors are taken as independent.

## 20.2 The posterior of the state (step 3)

All four questions of Table 20.1 are read from one posterior, that of hall B's reduced state under the seventeen rows of window 1. No row is added and no fit is repeated: this section cites the fit, reads from it which unit is the largest, and states the draws the answers use.

### 20.2.1 Chapter 5's fit, cited

The MAP state $\hat{\mathbf{z}}$ minimizes the objective $\mathcal{L}(\mathbf{z}) = \frac{1}{2}\,\mathbf{r}(\mathbf{z})^{\mathsf{T}}\mathbf{r}(\mathbf{z})$ of the residual vector (5.11), by the normal equations (5.12) $(\mathbf{J}^{\mathsf{T}}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^{\mathsf{T}}\mathbf{r}$ with the reduced Jacobian of equation (5.15), and the Laplace covariance is $\mathbf{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1}$ at $\hat{\mathbf{z}}$, equation (5.13). Table 20.2 lists the eight entries the unit-out answers depend on. Each row is an entry; the columns give its MAP value, its posterior spread $\sqrt{\Sigma_{ii}}$, the stated truth from which the record was generated, a check marker for Section 20.4.2, and for the three coils the as-built conductance at the MAP state and the probability that the unit is the largest.

The coils are known to about a fifth of their commissioning spread, and the fractions of racks B1, B3 and B5 keep the survey's spread of 0.08, as Chapter 5 found. The supply and the load, which the scenario pins, are not listed.

### 20.2.2 Which unit is the largest

At the MAP state the as-built conductances are 6.895, 6.987 and 6.672 kW/K, and unit B2 is the largest by 0.09 kW/K over B1. That margin is small against the spread: a log-derate's spread of about 0.028 is about 0.2 kW/K of conductance, so the three units lie within each other's band and which one is largest is a probability, not a fact. On each draw $\mathbf{z}_i$ of the

posterior the largest unit is read as $u^\star(\mathbf{z}_i)$ of equation (20.1), and its probability is

$$\Pr(u^\star = u \mid \mathbf{y}) \approx \frac{1}{N}\sum_{i=1}^{N} \mathbf{1}\big[u^\star(\mathbf{z}_i) = u\big], \tag{20.7}$$

by the Laplace draws, and the summed weights of the same draws by the exact posterior. By Table 20.2, unit B2 is the largest with probability 0.578 by the exact posterior, unit B1 with 0.347 and unit B3 with 0.075. The readings rank B3's coil lowest with some confidence and cannot separate B1 from B2: the return-air reading fixes the sum of the three conductances, and the three discharge readings split it with errors of the size of the gap.

### 20.2.3 The draws

The answers read the 50 000 Laplace draws of Section 5.2.3, states from $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$, with their exact weights (5.16), of effective sample size 49 638. Each draw is pushed through the unit-out scenario (20.3) with its own largest unit removed. The reported values are those of the state route, the Laplace draws; the exact posterior, the same draws under the weights, stands beside each. The question of the chillers reads 50 000 of the site's draws of Section 9.2.4, and hall A, set beside hall B in Section 20.3.4, reads 20 000 Laplace draws of its own fit of Section 15.2.2.

## 20.3 The answers (step 4)

Every answer is read from the posterior of Section 20.2.3, pushed through the scenario with a unit out. The reported value is joint over the five racks, by the draws; the exact posterior stands beside it, and for a named unit out the inversion route and the delta form on the binding rack alone stand beside both.

### 20.3.1 The design load with the largest unit out

**Planner's question.** With the largest unit out, what is the probability the remaining plant holds the inlets at the design load? The question asks for the posterior probability that every rack inlet stays at or below 27°C at 340 kW when the largest unit is lost.

**General form.** A requirement is accepted when its posterior probability reaches the assurance level, equation (2.12),

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

**Specialized form.** The question is equation (2.12) with $f(\mathbf{z}) = Q_L^\dagger(\mathbf{z})$ of equation (20.1), the binding load with the largest unit out, and $f^* = Q_d = 340\,\mathrm{kW}$. Because each inlet of equation (20.3) rises with the load, the event is the intersection of the five inlets' events at the design load,

$$\Pr(Q_L^\dagger \geq Q_d \mid \mathbf{y}) = \Pr\Big(\textstyle\bigcap_k \{T_k^*(\mathbf{z},\ Q_d;\ U \setminus \{u^\star(\mathbf{z})\}) \leq T_{\max}\} \;\Big|\; \mathbf{y}\Big) \approx \tfrac{1}{N}\textstyle\sum_{i=1}^{N} \mathbf{1}\big[Q_L^\dagger(\mathbf{z}_i) \geq Q_d\big], \tag{20.8}$$

read by the draws, each with its own largest unit removed; the exact posterior sums the weights (5.16) of the same draws. The binding load $Q_L^\dagger$ itself is read from the same draws, its median and its 5 % and 95 % points. Beside it, for each named unit out, two routes read the binding rack, B3, alone at $\hat{\mathbf{z}}$. The inversion route carries rack B3's rise per kilowatt $a_{\mathrm{B3}}$ of equation (20.3) as a normal by the delta method, equation (5.18) with the sums over $R$, and maps its quantiles and tails through the decreasing ratio (20.4), as equation (5.19) does with all

Table 20.3: The binding load of hall B at the drawing's supply of 14°C, with all units running and with each unit out. Draws: the 50 000 Laplace draws, joint over the five racks; exact: the same draws reweighted; delta: the binding rack at $\hat{\mathbf{z}}$ with the gradient (20.5). Gaps are exact minus draws, from the printed values.

| Units running | 5 % [kW] | Median [kW] | 95 % [kW] | Exact median [kW] | Gap [kW] | Delta at $\hat{\mathbf{z}}$ [kW] | Delta sd [kW] | Holds 300 kW [–] |
|---|---|---|---|---|---|---|---|---|
| All three | 345.89 | 395.21 | 450.43 | 394.92 | −0.29 | – | – | – |
| B1 out | 225.45 | 258.78 | 296.53 | 258.57 | −0.21 | 259.67 | 23.66 | 0.0353 |
| B2 out | 223.55 | 256.02 | 292.89 | 255.81 | −0.21 | 256.91 | 23.20 | 0.0239 |
| B3 out | 230.80 | 265.54 | 305.28 | 265.39 | −0.15 | 266.41 | 24.79 | 0.0788 |
| Largest out (reported) | 221.90 | 253.55 | 288.86 | 253.39 | −0.16 | none | – | 0.0125 |
| Worst out | 221.90 | 253.55 | 288.86 | 253.39 | −0.16 | none | – | 0.0125 |

units running; its median is therefore $Q^{\dagger}(\hat{\mathbf{z}};\ R)$. The delta form reads the same rack with the gradient (20.5), $\mathcal{N}(Q^{\dagger}(\hat{\mathbf{z}};\ R),\ \mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g})$. The largest unit has neither: $u^{\star}$ jumps from one unit to another across the posterior, and a minimum over units has no gradient where two units tie.

**Building blocks.** The draws and their weights (5.16); the largest unit (20.1) in each draw; the rise per kilowatt (20.3) and the binding loads (20.4) over the running units; the gradient (20.5), the inversion (5.19) and the covariance $\boldsymbol{\Sigma}$.

**The binding load with each unit out.** Table 20.3 gives the binding load of hall B with all units running and with each unit out. Each row is a set of running units; the columns give the 5 % point, the median and the 95 % point by the draws, the exact median and its gap from the draws', the delta form's value at $\hat{\mathbf{z}}$ and its standard deviation, and the probability that the binding load reaches today's 300 kW; the inversion route on rack B3 alone is set beside the draws in Table 20.4. The densities are drawn in Figure 20.1: the horizontal axis is the binding load, the vertical axis its density; the three curves are the delta form's normals for B1, B2 and B3 out, and the bars above them the 90 % bands of the draws with their medians, for the largest unit out, each named unit out and all three running. The vertical lines mark today's 300 kW and the 340 kW design load, and a dotted purple tick on the largest-out bar marks the stated truth.

With all three units running the draws reproduce Chapter 5's reported binding load, a median of 395.21 kW and the band [345.89, 450.43] kW, and its exact probability of the design load, 0.9692. With one unit out the whole band moves below the design load. Losing B1 leaves a median of 258.78 kW, losing B2 256.02 kW, losing B3 265.54 kW, and every 95 % point lies below 340 kW; the shift is close to the two thirds that equation (20.3) gives for equal units, $(395.21 + 10) \times 2/3 - 10 = 260.1$ kW. Losing B3, the unit the fit puts lowest, costs least, and losing B2, the unit most likely to be the largest, costs most. The draws with their own largest unit removed give the reported binding load,

$$\text{median } Q_L^{\dagger} = 253.55\,\text{kW}, \qquad Q_L^{\dagger} \in [221.90,\ 288.86]\,\text{kW with 90 \% probability}, \tag{20.9}$$

below every named unit's median, because in each draw it removes that draw's largest coil. The exact posterior gives a median of 253.39 kW, a gap of −0.16 kW. In every one of the 50 000 draws the largest unit out is also the worst unit out (Pr = 1.0000), as the equal air rates of Section 20.1.3 require. The delta form for a named unit out sits 0.87 to 0.89 kW above the draws' median (for B2, $256.91 - 256.02 = 0.89$ kW): the binding load is a ratio over $a_k$, skewed to the right, and its median lies below its value at $\hat{\mathbf{z}}$.

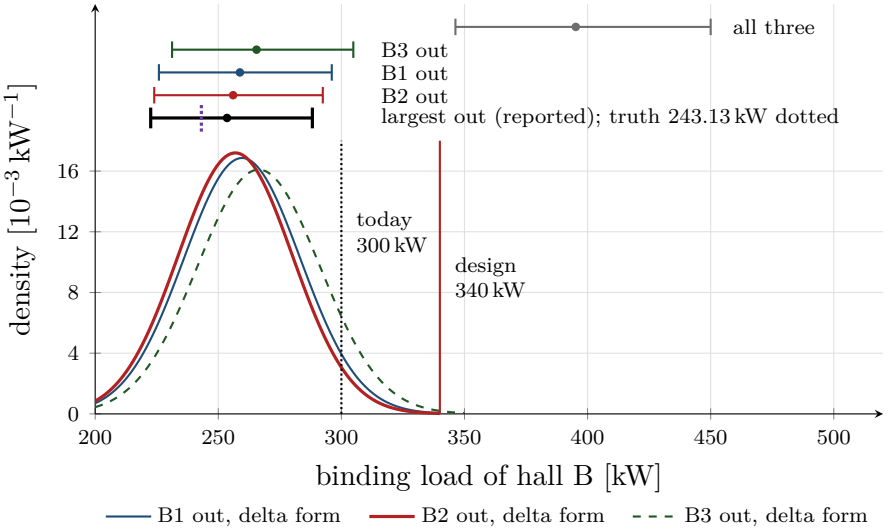



Figure 20.1: The binding load of hall B with a unit out, at the drawing's supply of 14°C. Curves: the delta form for each named unit out, $\mathcal{N}(259.67,\ 23.66^2)$, $\mathcal{N}(256.91,\ 23.20^2)$ and $\mathcal{N}(266.41,\ 24.79^2)$ kW for B1, B2 and B3 out. Bars: the 90 % bands of the 50 000 draws, joint over the five racks, with their medians as dots; the top bar is hall B with all three units running, [345.89, 450.43] kW. The dotted purple tick on the largest-out bar is the stated truth, 243.13 kW with the true largest unit, B1, out. The largest unit has no delta form, since it changes from draw to draw.

Table 20.4: The binding load of hall B with each named unit out, at the drawing's supply of 14°C: the 50 000 draws, joint over the five racks, beside the inversion route on rack B3 alone, equation (5.19) with the sums over the running units. Gaps are draws minus inversion, from the printed values.

| | Median | | | 95 % point | | | Pr(≥ 300 kW) | | | Pr(≥ 340 kW) |
|---|---|---|---|---|---|---|---|---|---|---|
| Unit out | Draws [kW] | Inversion [kW] | Gap [kW] | Draws [kW] | Inversion [kW] | Gap [kW] | Draws [–] | Inversion [–] | Gap [–] | Inversion [–] |
| B1 | 258.78 | 259.67 | −0.89 | 296.53 | 305.14 | −8.61 | 0.0353 | 0.0690 | −0.0337 | $4.4 \times 10^{-3}$ |
| B2 | 256.02 | 256.91 | −0.89 | 292.89 | 301.45 | −8.56 | 0.0239 | 0.0549 | −0.0310 | $3.2 \times 10^{-3}$ |
| B3 | 265.54 | 266.41 | −0.87 | 305.28 | 314.23 | −8.95 | 0.0788 | 0.1134 | −0.0346 | $9.5 \times 10^{-3}$ |

**Rack B3 alone beside the joint.** Table 20.4 sets the inversion route on rack B3 alone beside the joint draws for each named unit out. Each row is a unit out; the column groups give the median, the 95 % point and the probability of reaching today's 300 kW, each by the draws, joint over the five racks, and by the inversion route, with the gap, draws minus inversion; the last column is the inversion route's probability of reaching the 340 kW design load.

The two routes agree at the lower end of the band, where rack B3 binds: the 5 % points are 225.45 against 225.66 kW with B1 out, 223.55 against 223.52 kW with B2 out and 230.80 against 230.88 kW with B3 out. They part in the upper tail. Rack B3 alone puts the 95 % point 8.56 to 8.95 kW above the joint's, and with B2 out its probability of reaching 300 kW is 0.0549 against the joint's 0.0239, more than twice. The gap is the other racks' share: in the draws where rack B3 alone would carry a high load, rack B5 or rack B4 binds first instead, as they do in 0.0813 and 0.0400 of the draws with the largest unit out (Table 20.5). At the design load the gap decides the reading. With the MAP state's largest unit, B2, out, the inversion route on rack B3 alone gives a probability of $3.2 \times 10^{-3}$ of reaching 340 kW, and $4.4 \times 10^{-3}$ and $9.5 \times 10^{-3}$ with B1 or B3 out; the joint answer read next, with each draw's own largest unit

Table 20.5: The rack inlets of hall B at the 340 kW design load with the largest unit out and the supply at 14°C, by the 50 000 draws. The probability of binding first is the share of the draws in which the rack's binding load is the smallest.

| Rack | Median [°C] | 5 % [°C] | 95 % [°C] | $\Pr(T_k^* \le 27°\mathrm{C})$ [–] | Binds first [–] | Truth [°C] |
|---|---|---|---|---|---|---|
| B1 | 25.97 | 23.57 | 28.36 | 0.7605 | 0.0013 | 24.97 |
| B2 (measured) | 28.00 | 27.00 | 29.03 | 0.0498 | 0.0022 | 27.95 |
| B3 | 31.21 | 28.80 | 33.62 | 0.0018 | 0.8752 | 31.97 |
| B4 (measured) | 28.88 | 27.87 | 29.90 | 0.0011 | 0.0400 | 28.65 |
| B5 | 28.58 | 26.19 | 30.99 | 0.1391 | 0.0813 | 29.00 |

out, is below $10^{-4}$ by the exact posterior. One rack's normal tail overstates the joint by more than an order of magnitude, and the reported answer is the joint.

**The drill question.** By equation (20.8), at the design load of 340 kW with the largest unit out and the supply at 14°C,

$$\Pr(Q_L^{\dagger} \ge 340\,\mathrm{kW} \mid \mathbf{y}): \quad 0 \text{ of } 50\,000 \text{ draws; exact below } 10^{-4}. \tag{20.10}$$

The probability is below what the draws can resolve, and the exact posterior puts it below $10^{-4}$ too. Not one state of the posterior holds the design load with its largest unit out. At today's 300 kW the same scenario holds with probability 0.0125 by the draws, with a Monte Carlo standard error of $\sqrt{0.0125 \times 0.9875/50\,000} = 0.0005$, and 0.0122 by the exact posterior. The drawing's N+1 is N: one unit out does not hold today's load, let alone the design load.

**The inlets at the design load.** Table 20.5 gives the five inlets at 340 kW with the largest unit out. Each row is a rack; the columns give its inlet's median and 5 % and 95 % points by the draws, the probability that it stays at or below 27°C, the probability that the rack binds first, and the stated truth, a check marker for Section 20.4.2. The hall with those inlets is drawn in Figure 20.2: the cold aisle on the left feeds the racks, the return air in the middle feeds the units, and the largest unit, B2 at the MAP state, is grayed out; each rack carries its inlet's median and its probability of holding.

Four of the five inlets have medians above the limit. Rack B3, the unmeasured rack with the largest surveyed fraction, sits at a median of 31.21°C, more than 4 K over, and binds first in 0.875 of the draws, the probability Chapter 5 found with all units running: by Section 20.1.3 the racks bind in the order of their fractions whichever unit is out. The rack least likely to hold is not B3 but the measured rack B4, at 0.0011 against 0.0018: its median is lower, 28.88°C, but its sensor narrows its band to about ±1 K, all of it above the limit, while rack B3's band of about ±2.4 K reaches down toward it. The rack most likely to bind first and the rack least likely to hold are answers to different questions.

### 20.3.2 The joint tail and the racks' tails

**Planner's question.** Is the joint tail the product of the racks' tails? The question asks whether the probability that all five inlets hold can be built from each rack's probability alone, and which rack the joint tail follows.

**General form.** The joint tail is a threshold question, equation (2.12), on the intersection of the racks' events, set beside the racks' marginal tails, each the same form on one event,

$$P_{\cap} = \Pr\Big(\textstyle\bigcap_k E_k \mid \mathbf{y}\Big), \qquad P_k = \Pr(E_k \mid \mathbf{y}), \qquad \textstyle\prod_k P_k \;\le\; P_{\cap} \;\le\; \min_k P_k,$$

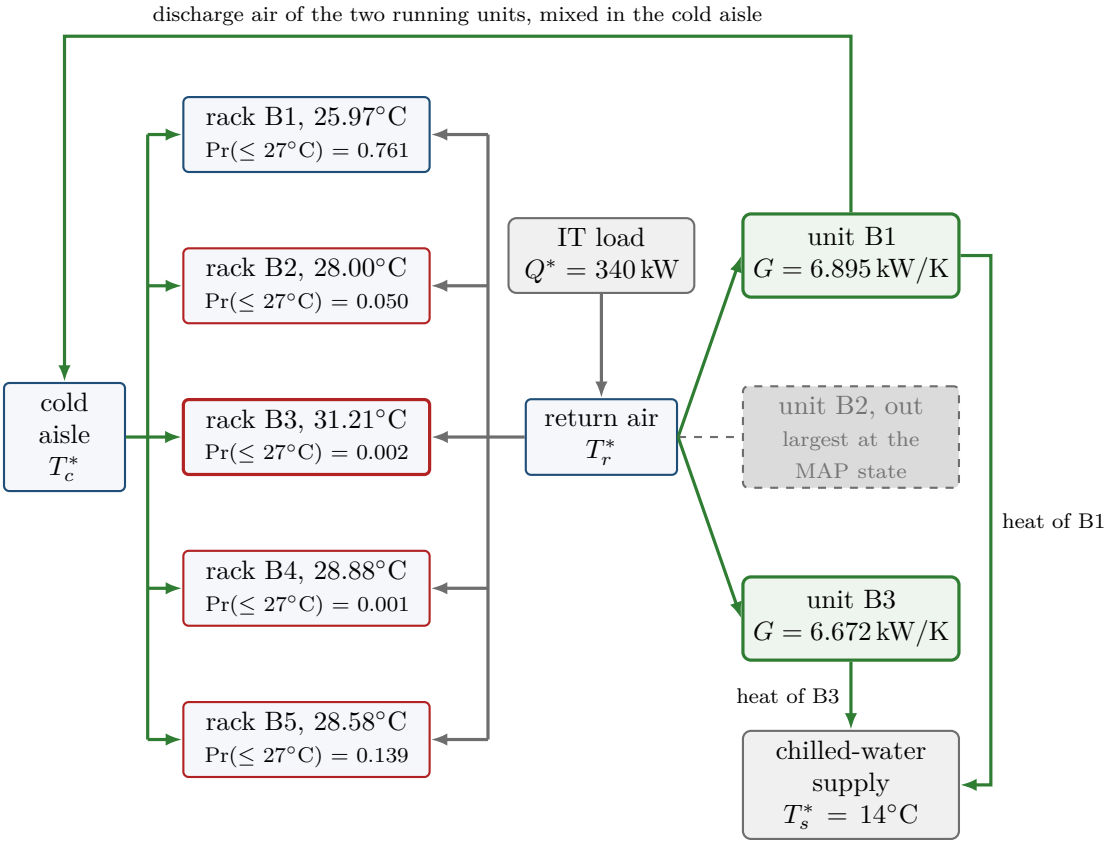



Figure 20.2: Hall B at the scenario node with its largest unit out, at the 340 kW design load and the supply of 14°C. Unit B2, the largest at the MAP state (probability 0.578), is drawn out, gray and dashed; on each draw of the posterior the unit removed is that draw's largest. Each rack carries its inlet's median over the draws and the probability that it stays at or below 27°C (Table 20.5); the four racks outlined in red have medians above the limit, and rack B3, the thickest outline, binds first with probability 0.875. The running coils' conductances are those at the MAP state.

Table 20.6: The joint tail of the five inlets at the 340 kW design load with the largest unit out, against the pinned supply, beside the product of the five marginals and the smallest marginal. 50 000 draws; gaps are exact minus draws, from the printed values.

| Supply [°C] | Joint, draws [–] | Joint, exact [–] | Gap [–] | Product of marginals [–] | Smallest marginal [–] | Its rack [–] |
|---|---|---|---|---|---|---|
| 14.0 | $< 10^{-4}$ | $< 10^{-4}$ | – | $< 10^{-4}$ | 0.0011 | B4 |
| 13.0 | 0.0005 | 0.0004 | −0.0001 | 0.0002 | 0.0134 | B3 |
| 12.0 | 0.0303 | 0.0297 | −0.0006 | 0.0215 | 0.0652 | B3 |
| 11.0 | 0.1711 | 0.1686 | −0.0025 | 0.1623 | 0.2024 | B3 |
| 10.0 | 0.4248 | 0.4214 | −0.0034 | 0.4216 | 0.4438 | B3 |
| 9.0 | 0.7001 | 0.6970 | −0.0031 | 0.6994 | 0.7065 | B3 |

the right inequality for any events and the left for events that rise and fall together over the posterior, which for positively correlated normal quantities is Slepian's inequality of Chapter 7. The smallest marginal is that chapter's worst-marginal bound, read for holding rather than for violation; the product of the marginals is its independence product, the joint tail the events would have if they were independent.

**Specialized form.** The question is equation (2.12) with $E_k = \{T_k^*(\mathbf{z},\ Q_d;\ U \setminus \{u^\star\}) \le T_{\max}\}$, each inlet of equation (20.3) at the design load with the largest unit out, asked over a family of pinned supplies $T_s^*$ from 14 down to 9°C. At each supply the draws give the joint tail and the five marginals,

$$\hat{P}_\cap(T_s^*) = \frac{1}{N}\sum_{i=1}^{N}\prod_k \mathbf{1}\big[E_k(\mathbf{z}_i)\big], \qquad \hat{P}_k(T_s^*) = \frac{1}{N}\sum_{i=1}^{N}\mathbf{1}\big[E_k(\mathbf{z}_i)\big], \tag{20.11}$$

and the exact posterior weights the first by (5.16). The five events rise and fall together: by equation (20.3) every $a_k$ carries the common term $1/S_G$, so a draw with weak coils raises every inlet at once, and the left inequality holds.

**Building blocks.** The draws with their own largest unit out; the inlets (20.3) at each pinned supply; the weights (5.16).

**The joint tail against the supply.** Table 20.6 gives the answer at six supplies. Each row is a supply; the columns give the joint tail by the draws and by the exact posterior with their gap, the product of the five marginals, and the smallest marginal with its rack. The upper panel of Figure 20.3 plots the same rows against the supply: the horizontal axis is the pinned supply, the vertical axis a probability, with the joint tail solid, the product of the marginals dashed, the smallest marginal dotted and the 0.95 assurance level as a horizontal line; the two dots mark the median supply that holds every inlet and the supply that holds them with probability 0.95, both read in Section 20.3.3.

At the drawing's supply of 14°C none of the three readings resolves a nonzero joint tail, and the smallest marginal, rack B4's 0.0011, bounds it from above. As the supply falls the three separate. The product of the marginals understates the joint tail throughout: at 12°C it is 0.0215 against the joint 0.0303, short by 0.0088, close to a third of the joint, and at 13°C it is 0.0002 against 0.0005, short by more than half. The cause is the shared term of equation (20.3): the five inlets share the posterior of $S_G$, so the draws in which one inlet holds are the draws in which the coils are good and the others hold too, and multiplying the marginals counts that good fortune five times over. Near the answer the joint tail collapses onto rack B3's marginal. From 13°C down, B3's is the smallest marginal, and by 10°C the joint is 0.4248 against B3's

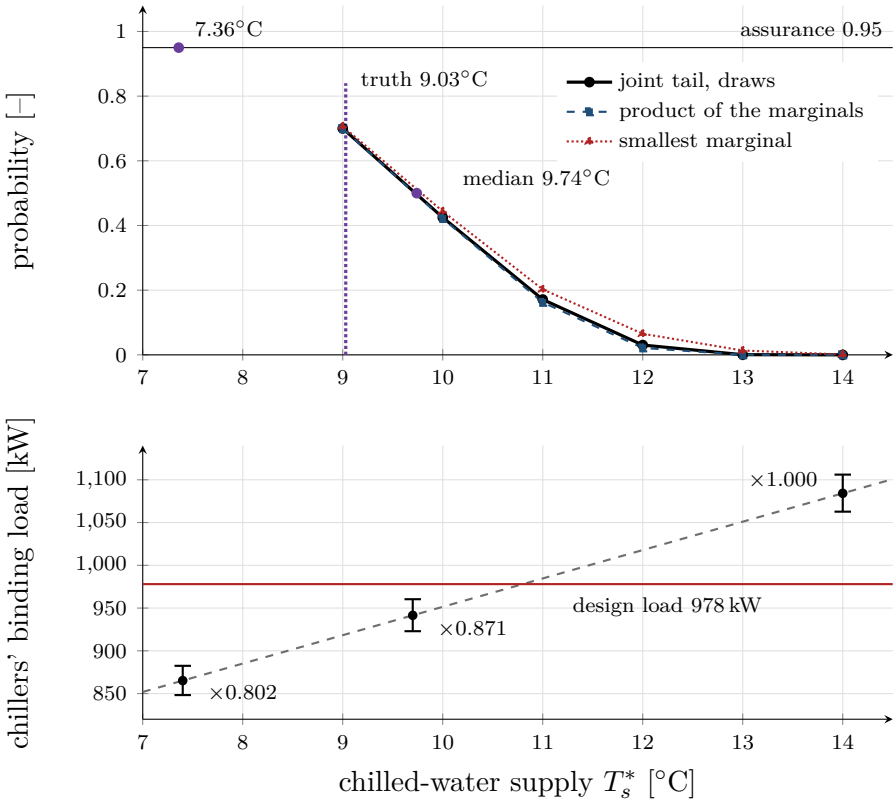



Figure 20.3: The supply reset. Upper panel: the joint tail of the five inlets at 340 kW with the largest unit out, the product of the five marginals and the smallest marginal, against the pinned supply (Table 20.6); the dots mark the median supply that holds every inlet, 9.74°C, and the supply that holds them with probability 0.95, 7.36°C; the dotted purple line is the stated truth's supply that holds, 9.03°C. Lower panel: the chillers' binding load in site IT load at the design day, median with its 90 % band (Table 20.7), at 14, 9.7 and 7.4°C, each labeled with the capacity factor; the dashed line is the median at 14°C scaled by the stated 3 % per kelvin. Today's site IT load, 660 kW, lies below the panel.

0.4438, by 9°C 0.7001 against 0.7065: the other racks hold in almost every draw in which B3 holds, and one rack decides. The exact posterior lies at most 0.0034 below the draws, at 10°C, where the Monte Carlo standard error of the draws is $\sqrt{0.4248 \times 0.5752/50\,000} = 0.0022$; the two use the same states, and the gap is the posterior's small departure from the Laplace normal.

### 20.3.3 The supply that would restore it

**Planner's question.** What supply would restore it, and what does that cost the chillers? The question asks how far the chilled-water supply must be reset for every inlet to hold at the design load with the largest unit out, and whether the chillers can carry the site at that supply.

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

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

read by its median, its central band and a quantile at the assurance level.

**Specialized form.** The question is equation (2.9) with $f(\mathbf{z}) = T_s^\dagger(\mathbf{z})$, the highest supply at which every inlet of equation (20.3) holds at the design load with the largest unit out. Setting the hottest inlet equal to the limit and solving for the supply,

$$\begin{aligned} T_s^\dagger(\mathbf{z}) &= T_{\max} - \max_k a_k\big(\mathbf{z};\ U \setminus \{u^\star\}\big)\,(Q_d + Q_o), \\ \Pr\big(\text{every inlet holds at } T_s^* \mid \mathbf{y}\big) &= \Pr\big(T_s^\dagger \ge T_s^* \mid \mathbf{y}\big), \end{aligned} \tag{20.12}$$

in °C. The second relation says that the joint tail of Table 20.6, read as a function of the supply, is the upper tail of $T_s^\dagger$; the supply that holds with probability 0.95 is the 5 % point of $T_s^\dagger$. The supply is the site's: hall A and hall B are fed by the same chilled water, so a reset for hall B is a reset of the loop. The chillers' capacity at the design day falls with the supply by the stated factor $f(T_s) = 1 + 0.03\,(T_s - 14)$, and their binding load in site IT load is that of equation (15.6) with all three chillers up and the factor applied,

$$Q_{\mathrm{ch}}^\dagger(\boldsymbol{\kappa};\ T_s) = f(T_s) \sum_i 400\,\mathrm{e}^{\kappa_i} - Q_{o,\mathrm{site}}, \qquad Q_{o,\mathrm{site}} = 22\,\mathrm{kW}, \tag{20.13}$$

read by the site's draws at three supplies, each with the probability that it carries today's 660 kW and the design load of 1 000 kW to the loop, 978 kW of site IT load.

**Building blocks.** The hall's draws with their own largest unit out; the rise per kilowatt (20.3); the site's draws of the chillers' log-ratios $\kappa_i$ (Section 9.2.4); the stated capacity slope.

**The supply.** By equation (20.12) the supply that holds every inlet at 340 kW with the largest unit out has

$$\text{median } T_s^\dagger = 9.74°\mathrm{C}, \qquad T_s^\dagger \in [7.38,\ 11.78]°\mathrm{C} \text{ with } 90\,\% \text{ probability}, \tag{20.14}$$

and 9.73°C by the exact posterior. To hold with probability 0.95 the supply must be 7.36°C, more than 6.6 K below the drawing's 14°C.

Table 20.7: The site's chillers at the design day, all three up, at the drawing's supply and at two resets, from the site's draws. The binding load is in site IT load, equation (20.13); the design load is 1 000 kW to the loop, 978 kW of site IT load.

| Supply [°C] | Factor $f(T_s)$ [–] | Median [kW] | 5 % [kW] | 95 % [kW] | Pr(≥ 660 kW) [–] | Pr(≥ 978 kW) [–] |
|---|---|---|---|---|---|---|
| 14.0 | 1.0000 | 1 084.2 | 1 061.5 | 1 107.3 | 1.0000 | 1.0000 |
| 9.7 | 0.8710 | 941.5 | 921.7 | 961.6 | 1.0000 | 0.0017 |
| 7.4 | 0.8020 | 865.2 | 847.0 | 883.7 | 1.0000 | $< 10^{-4}$ |

**The cost to the chillers.** Table 20.7 gives the chillers at three supplies: the drawing's, the median reset and the reset that holds with probability 0.95, each rounded to a tenth of a kelvin. Each row is a supply; the columns give the capacity factor, the chillers' binding load in site IT load with its band, and the probabilities that it carries today's load and the design load. The lower panel of Figure 20.3 plots the binding loads with their bands against the supply, beside the 978 kW design load.

At the drawing's supply the chillers carry the design load with room to spare, a median of 1 084.2 kW against 978 kW. At the median reset of 9.7°C they lose 13 % of their capacity, a factor of 0.871; they still carry today's 660 kW in every draw, but carry the design load with probability only 0.0017. At 7.4°C, the reset that holds hall B's inlets with probability 0.95, they lose about a fifth, a factor of 0.802, and carry the design load in no draw that the table resolves, below $10^{-4}$. The reset borrows from the chillers' margin: it holds the hall today and fails the site at the design load. Hall B's units and the site's chillers are separate elements on the drawing, but they share one resource, the chilled-water supply, and a remedy taken from the shared resource moves the shortfall from one element to the other.

### 20.3.4 A fourth unit

**Planner's question.** What does a fourth unit buy? The question asks for the probability of the drill question with a fourth hall-B unit installed and the largest of the four out, and for how long it holds as the load grows.

**General form.** The threshold form, equation (2.12), asked of a design alternative: the state is extended by the alternative's own uncertain ratings, given their stated distribution, and the same event is read over the extended state,

$$\Pr(f(\mathbf{z},\ \boldsymbol{\theta}_a) \geq f^* \mid \mathbf{y}) = \iint \mathbf{1}\left[f(\mathbf{z},\ \boldsymbol{\theta}_a) \geq f^*\right] p(\boldsymbol{\theta}_a)\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\boldsymbol{\theta}_a\, \mathrm{d}\mathbf{z}.$$

**Specialized form.** The question is equation (2.12) with $\boldsymbol{\theta}_a = g_{\text{new}} \sim \mathcal{N}(0,\ 0.15^2)$, the new unit's as-built log-ratio, the units $U_4 = \{\text{B1}, \text{B2}, \text{B3}, \text{new}\}$, and $f$ the binding load with the largest of the four out,

$$Q^{\dagger}_{L,4}(\mathbf{z},\ g_{\text{new}}) = Q^{\dagger}(\mathbf{z};\ U_4 \setminus \{u_4^{\star}\}), \qquad u_4^{\star} = \arg\max_{u \in U_4} G_u, \qquad t^{\dagger} = \frac{\ln(Q^{\dagger}_{L,4}/300)}{\ln 1.15}, \tag{20.15}$$

by equations (20.3) and (20.4) with three units running, $S_C = 30\,\text{kW/K}$; and $t^{\dagger}$ the crossing year of equation (6.3) at which hall B's load, 300 kW today and growing 15 % a year, meets the unit-out binding load. Each draw of the hall's state is paired with one draw of $g_{\text{new}}$. The new unit's spread is a stated input, varied from 0 to 0.30. Hall A is read beside, by the same event on its own fit at its own today's load of 360 kW.

Table 20.8: The binding load with the largest unit out, and the probability of holding the load asked, for hall B as drawn, hall B with a fourth unit (its as-built log-ratio drawn from its stated prior of spread 0.15, or of the spread given), and hall A beside. Supply 14°C.

| Plant, largest unit out | Load [kW] | Pr(hold), draws [–] | Pr(hold), exact [–] | Median [kW] | 5 % [kW] | 95 % [kW] |
|---|---|---|---|---|---|---|
| Hall B, three units | 340 | 0 of 50 000 | $< 10^{-4}$ | 253.55 | 221.90 | 288.86 |
| Hall B, three units | 300 | 0.0125 | 0.0122 | 253.55 | 221.90 | 288.86 |
| Hall B, four units (reported) | 340 | 0.9034 | 0.9024 | 382.97 | 329.48 | 441.63 |
| Hall B, four units, new spread 0 | 340 | 0.9641 | – | – | – | – |
| Hall B, four units, new spread 0.30 | 340 | 0.8031 | – | – | – | – |
| Hall A, three units | 360 | 0.2713 | – | 341.79 | 287.83 | 383.87 |

**Building blocks.** The hall's draws and weights; draws of $g_{\text{new}}$ from its stated prior; the largest unit of four in each pair; the crossing-year map (6.3); hall A's draws.

**The fourth unit.** Table 20.8 sets the fourth unit beside the hall as drawn and beside hall A. Each row is a plant with its largest unit out; the columns give the load asked, the probability that every inlet holds there by the draws and by the exact posterior, and the unit-out binding load's median and band.

With a fourth unit and the largest of the four out, every inlet holds at the 340 kW design load with probability 0.9034, and 0.9024 by the exact posterior; the unit-out binding load has a median of 382.97 kW and the band [329.48, 441.63] kW. The new unit is the largest of the four with probability 0.4732, nearly one half: its prior spread of 0.15 is five times the fitted spread of the existing coils, so its conductance lands above theirs about as often as below. With all four units running the binding load's median is 535.07 kW. The spread of the new unit's as-built ratio moves the answer: at zero spread, a unit exactly at its rating, the probability is 0.9641; at a spread of 0.30 it is 0.8031. A wide spread makes the new unit the largest more often and, when it is the largest, removes more. The fourth unit restores N+1 at the design load with probability 0.90, short of the assurance level of 0.95 that Chapter 5's review applied with all units running.

It restores it for a short time. By equation (20.15) the load meets the four-unit hall's unit-out binding load in year

$$\text{median } t^{\dagger} = 1.75, \qquad t^{\dagger} \in [0.67,\ 2.77] \text{ with } 90\,\% \text{ probability},$$

the quantiles of the binding load carried through the monotone map of Chapter 6. The fourth unit buys N+1 for under two years at the median.

**Hall A beside.** Hall A, on its own fit of Section 15.2.2, has a binding load of 539.27 kW at the median with all units running, band [453.55, 604.17] kW, and of 341.79 kW with its largest unit out, band [287.83, 383.87] kW, against its 360 kW today. It holds today's load with its largest unit out with probability 0.2713. Hall A is N as well, by the same reading, though less deeply than hall B: its median falls short of today's load by about 18 kW, hall B's by about 46 kW.

## 20.4 The checks (step 5)

The answers of Section 20.3 rest on the closed form of the unit-out inlets, on the Laplace draws of the state, and on a record generated from a stated truth; each is tested here.

### 20.4.1 The closed form and the routes

Three checks bear on the computation. First, the closed form (20.3) against the hall's own equations: at the MAP state, with each unit out in turn, the nine rows of equation (20.2) were solved for the five inlets at 340 kW, and the largest difference from $T_s^* + a_k(Q_d + Q_o)$ is below $10^{-4}$ K. Second, the draws reproduce Chapter 5: with all units running the binding load's median and band are that chapter's reported 395.21 kW and [345.89, 450.43] kW, and the exact probability of the design load is its 0.9692. Third, the exact posterior against the draws: the drill question's probability is below $10^{-4}$ by both, the exact medians of Table 20.3 lie within 0.29 kW of the draws' (394.92 against 395.21 kW with all units running, within 0.21 kW with a unit out), the probability of holding today's load with the largest unit out moves by 0.0003, within the draws' standard error of 0.0005, and the joint tails of Table 20.6 move by at most 0.0034. With an effective sample size of 49 638 of 50 000 the posterior of the state is nearly the Laplace normal, and the routes differ in the derived quantity: the delta form's centers sit about 0.9 kW above the draws' medians, the skew of a ratio, and the inversion route on rack B3 alone, which ignores the other racks, puts the 95 % point about 9 kW higher and the probability of the design load at $3.2 \times 10^{-3}$ with B2 out, against the joint's below $10^{-4}$ (Table 20.4).

### 20.4.2 The answers against the stated truth

The record of hall B is generated from a stated truth, the truth column of Table 20.2, and the last check compares the answers with it. In the truth the largest unit is B1, with a log-derate of 0.03: the unit the posterior ranks second, at probability 0.347. With B1 out the true unit-out binding load is 243.13 kW, inside the reported band [221.90, 288.86] kW and below today's 300 kW: the truth agrees with the answer that the hall is N. The true inlets at 340 kW are 24.97, 27.95, 31.97, 28.65 and 29.00°C for racks B1 to B5, each inside its band of Table 20.5; four of the five exceed the limit, as the medians do. The true supply that would hold every inlet is 9.03°C, inside the reported band [7.38, 11.78]°C. Later in the night record unit B2's coil derates; with B2 at 0.70 of its rating and the largest unit, B1, out, the true unit-out binding load falls to 191.50 kW. The derate makes a bad answer worse. Table 20.9 collects the checks: each row is a quantity, the columns its reported value, the route beside it, the truth and the reading.

The answers hold where they can be checked. The posterior did not name the true largest unit, and did not need to: the question is asked of whichever unit is largest, and the answer, read over the posterior's uncertainty about which that is, brackets the truth. The truth lies in the lower half of the band, as Chapter 5 found with all units running, for the same reason: the survey's 0.50 for rack B3 is below the true 0.57, and no reading corrects it.

## 20.5 The report (step 6)

Hall B is N on its air side, not the N+1 the sheet says. With the largest unit out, the probability that every inlet holds at the 340 kW design load is below $10^{-4}$, and at today's 300 kW it is 0.012. Rack B3 decides the joint tail. A supply reset to about 7.4°C would hold the hall, but it costs the chillers a fifth of their capacity and fails the site at the design load. A fourth hall-B unit gives a probability of 0.90 at 340 kW, and at the forecast growth it holds for about two years.

### 20.5.1 The answers and their bands

Table 20.10 gives the answers in the order of Table 20.1. Each row is a question; the second column gives the reported answer with the exact posterior beside it, and the third its band or the routes beside it.

Table 20.9: The checks of the unit-out answers. Reported values by the 50 000 Laplace draws, joint over the five racks; the truth is the stated truth of the record.

| Quantity [unit] | Reported | Beside | Truth | Reading |
|---|---|---|---|---|
| Largest unit [–] | B2, 0.5761 | exact 0.5776; B1 0.3472 | B1 | the posterior's second unit |
| Binding load, largest out [kW] | 253.55 | exact 253.39 | 243.13 | inside [221.90, 288.86]; below 300 |
| The same after B2's derate [kW] | – | – | 191.50 | the derate lowers it further |
| Pr(hold at 340 kW), largest out [–] | 0 of 50 000 | exact $< 10^{-4}$ | fails | agrees |
| Pr(hold at 300 kW), largest out [–] | 0.0125 | exact 0.0122 | fails | gap within se 0.0005 |
| Pr($\geq$ 340 kW), B2 out, rack B3 alone [–] | – | inversion $3.2 \times 10^{-3}$ | fails | one rack's tail overstates the joint |
| Inlet B3 at 340 kW [°C] | 31.21 | – | 31.97 | inside [28.80, 33.62] |
| Inlet B5 at 340 kW [°C] | 28.58 | – | 29.00 | inside [26.19, 30.99] |
| Supply that holds [°C] | 9.74 | exact 9.73 | 9.03 | inside [7.38, 11.78] |
| Largest out is worst out [–] | 1.0000 | equal air rates | – | as derived |
| Closed form (20.3) [K] | – | rows, $< 10^{-4}$ | – | agrees |
| All three running [kW] | 395.21 | Chapter 5's | – | reproduced |

Table 20.10: The report for hall B with a unit out, in window 1, at the drawing's supply of 14°C unless stated. Reported values by the 50 000 Laplace draws, joint over the five racks, each draw with its own largest unit out; the exact posterior beside each.

| Question | Answer | Band or routes beside |
|---|---|---|
| With the largest unit out, what is the probability the remaining plant holds the inlets at the design load? | 0 of 50 000 draws (exact: below $10^{-4}$); at today's 300 kW, 0.0125 (exact: 0.0122) | unit-out binding load, median 253.55 kW, band 221.90 to 288.86 kW (exact median 253.39 kW); largest unit B2 with probability 0.578; rack B3 alone with B2 out, inversion route, $3.2 \times 10^{-3}$ at 340 kW |
| Is the joint tail the product of the racks' tails? | no: the product understates it, 0.0215 against 0.0303 at 12°C | near the answer the joint follows rack B3's marginal: 0.4248 against 0.4438 at 10°C |
| What supply would restore it, and what does that cost the chillers? | median 9.74°C (exact: 9.73°C); 7.36°C to hold with probability 0.95 | band 7.38 to 11.78°C; chillers $\times$0.871 at 9.7°C, design load carried with probability 0.0017; $\times$0.802 at 7.4°C, below $10^{-4}$ |
| What does a fourth unit buy? | 0.9034 at 340 kW (exact: 0.9024) | unit-out binding load 329.48 to 441.63 kW; met by growth in year 1.75, band 0.67 to 2.77; 0.9641 to 0.8031 over the new unit's spread |
| Hall A beside | 0.2713 at today's 360 kW | unit-out binding load 287.83 to 383.87 kW, median 341.79 kW |

### 20.5.2 The design-review entry

The design review's entry counts redundancy in inlet limits, not in coil duty. Hall B is recorded as N on its air side. With its largest unit out it does not hold the 340 kW design load in any state of the posterior (below $10^{-4}$), and it holds today's 300 kW with probability 0.012; the sheet's N+1, 360 kW of duty against 350 kW, is withdrawn. The entry records that the shortfall is decided at rack B3, the unmeasured rack, which binds first in 0.875 of the posterior, and that the largest unit, read on the posterior, is B2 with probability 0.578 and B1 with 0.347; since the units move the same air, the largest unit out is the worst in every state, and the entry need not name it. Two remedies are recorded with their costs. A reset of the site's chilled-water supply to 7.4°C holds hall B with probability 0.95 but leaves the chillers unable to carry the site's design load; it is not an acceptable remedy, and a reset to the median 9.7°C fails the site's design load nearly as surely. A fourth hall-B unit of the existing model restores N+1 at the design load with probability 0.90, below the review's 0.95, and the load reaches its unit-out binding load in about two years. The entry asks for the fourth unit's price and for a remedy of the air side that lasts longer, and it records hall A as N by the same reading.

**Common mistake**

Counting redundancy in coil duty. Two units of 180 kW against 350 kW of heat says nothing about the inlets: losing a unit raises the return air, every inlet rises with it, and the hall's binding load with all three units running, 405 kW of heat, is already far below the 540 kW the three duties add to. The check that counts is the inlets' joint tail at the scenario with the unit out.

## 20.6 The next measurement (step 7)

The report leaves an answer that no plausible state of the posterior overturns, and a remedy, the fourth unit, whose value rests on the same unit-out binding load. Step 7 asks which test of the existing hall would narrow that binding load, the quantity that sizes the remedy.

### 20.6.1 An inlet sensor at rack B3

A candidate measurement adds a row $\mathbf{J}_a$ to the residual vector, and before its value is read its effect on the covariance is known: the preposterior covariance $\boldsymbol{\Sigma}_a = (\boldsymbol{\Sigma}^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a)^{-1}$ of equation (5.26), the form whose ranking of candidates Chapter 18 derives in general, equation (18.3). The quantity is the binding load with unit B2 out, the MAP state's largest unit, whose delta-form standard deviation is 23.20 kW. A permanent inlet sensor at rack B3 of standard deviation 0.3 K is the row $\mathbf{J}_{\mathrm{B3}}$ of the pattern (20.6), read in window 1 with all units running. It takes the posterior spread of $\beta_{B3}$ from 0.0800 to 0.0320 and the standard deviation of the B2-out binding load from 23.20 to 10.47 kW: it removes 54.9 % of the width. Sensors of the same standard deviation at racks B1 or B5, the other unmeasured racks, take it from 23.20 to 23.19 kW, 0.0 %. The ranking is that of Chapter 5 with all units running, and for the same reason: by the pattern (20.6) the unit-out binding load has one fraction entry, $\beta_{B3}$'s, and only a reading at rack B3 has an entry in that column.

### 20.6.2 The recommendation

Fit the inlet sensor at rack B3. It is the measurement Chapter 5 recommended for the headroom with all units running, and it removes more than half of the band of the unit-out binding load as well. It will not change whether hall B is N: averaged over the outcomes of the reading, the probability of holding today's load with the largest unit out stays at today's 0.0125, the law of total probability, and every state of the posterior fails the design load. What it will change is

the size of the remedy: the band of the unit-out binding load, and with it the year at which a fourth unit's redundancy is spent, narrows by about half. A sensor at rack B1 or B5 is not worth fitting for this question.

### 20.6.3 What the analysis showed

The sheet's N+1 is N. Two units of 180 kW against 350 kW of heat pass the sheet's check, but the inlets with the largest unit out do not hold the design load in any state of the posterior, and hold today's load with probability 0.012. The reason is the hall's own equations. Removing one of three nearly equal coils leaves the ratio of conductance to air unchanged and multiplies every inlet's rise per kilowatt by about three halves, so the binding load plus the other heat falls to about two thirds of its value with all units running: from 405 kW of heat to about 270 kW. The sheet's duty never meets this arithmetic, because the duty is a coil rating and the limit is an inlet.

Three further findings carry the answer. The largest unit is a posterior quantity: the fit cannot tell B1 from B2, and in the truth the largest unit is the one the posterior ranks second. That does not matter here, because with equal air rates the largest unit out is the worst unit out in every state, and the question asked of "the largest" is the question asked of the worst. The joint tail is not the product of the racks' tails: the inlets share the running coils' conductance, the product understates the joint by a third or more where the answer is small, and near the answer one rack, B3, decides it. And the remedy taken from a shared resource moves the shortfall instead of removing it: a chilled-water reset deep enough to hold hall B costs the chillers a fifth of their capacity and fails the site at its design load. The fourth unit restores N+1 with probability 0.90, and the load's growth spends that redundancy in under two years at the median.

### 20.6.4 Limits of the model

**A unit out, as modeled.** The down unit loses its coil and its fan together, and the running fans keep their speed. A controller that sped the remaining fans up would raise $S_C$ and lower every inlet; a stopped unit through which return air flows backward into the cold aisle would raise them. Neither is in equation (20.2).

**Fractions unchanged by the loss.** The racks' recirculation fractions keep their calibrated values with a unit out, and the cold aisle is one mixed volume. In a real hall the racks nearest a stopped unit lose their supply first and draw more return air; which unit is out would then matter through its position, not only through its conductance, and the fractions would need rows read with a unit stopped, which the record does not hold.

**A pinned supply, and two separate fits.** The hall's answers are read at the drawing's supply of 14°C with all three chillers up. The reset's cost reads the chillers at three supplies with the stated 3 % per kelvin, not a fitted slope, and the site's posterior is taken as independent of the hall's. A chiller out at the same time as a hall unit is a question of the enumeration of Chapter 15, not of this scenario.

**One window, the healthy plant.** The fit reads window 1, the plant before unit B2's derate. The check of Section 20.4.2 shows that the derate lowers the true unit-out binding load further; a fit of a later window would describe that plant.

**Stated inputs.** The fourth unit's rating and its prior spread are stated, and the answer moves with the spread from 0.9641 to 0.8031. The growth of 15 % a year is the planning figure, without its band.

**The constructed record.** The readings of window 1 are generated from a stated truth with stated sensor errors, not taken from a hall, and the comparison with that truth in Section 20.4.2 is a check of the method on its own model, not of the model against a real hall.

**The next chapter.** A unit out here is one unit, lost alone. Chapter 21 asks what follows when the load a lost unit carried moves onto the units that remain, and whether one loss can bring on another.

## 20.7 What to retain

**Definition 20.1** (Redundancy at the inlets)**.** A plant is N+1 at a load $Q_d$ when, with any one unit out, the scenario node holds every constrained entry at $Q_d$. With units $R$ running, each inlet is $T_k^* = T_s^* + a_k(Q^* + Q_o)$ with $a_k = (1 - (1 - \beta_k)\, S_G/S_C)/S_G$, $S_G$ and $S_C$ the running units' summed conductance and air capacity rate. The largest unit $u^\star = \arg\max_u G_u$ is read on each draw of the posterior, and the probability that the plant holds with it out is the joint tail $\Pr(\cap_k\{T_k^* \le T_{\max}\} \mid \mathbf{y})$ over the draws, with the exact posterior beside it, the product of the marginals below it and the smallest marginal above it.

Four conclusions carry beyond hall B. Redundancy is counted in the limits the plant is held to, not in the ratings of its units, and a unit out changes the plant's equations, not only its sum of duties: at equal units, losing one of $n$ scales every inlet's rise per kilowatt by $n/(n-1)$. Which unit is largest is a probability when the ratings are fitted, and the question asked of "the largest" is answered on each draw; when the units move the same air it is the question of the worst. A joint tail over entries that share a term is not the product of their tails: the product understates it, the smallest tail bounds it, and near the answer the entry with the smallest tail decides it. And a remedy drawn from a shared resource, such as the chilled-water supply, is checked against every element that shares it before it is counted.

Table 20.11 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 20.11: The terms of headroom with the largest unit out and where each is defined. Units for hall B.

| Term | Expression | Units | Meaning | Defined in |
|---|---|---|---|---|
| Largest and worst unit | $u^\star = \arg\max_u G_u$; $u_{\mathrm{w}} = \arg\min_u Q^\dagger(\mathbf{z};\ U \setminus \{u\})$ | – | read on each draw | 20.1.1 |
| Unit-out binding load | $Q_L^\dagger = Q^\dagger(\mathbf{z};\ U \setminus \{u^\star\})$ | kW | the N+1 check | 20.1.1 |
| Unit-out rows | eq. (20.2), sums over the running units | kW, K | the hall with a unit out | 20.1.3 |
| Rise per kilowatt | $a_k = \big(1 - (1 - \beta_k)\, S_G/S_C\big)/S_G$ | K/kW | inlet slope with $R$ running | 20.1.3 |
| Unit-out gradient | eq. (20.5), no entry for the down unit | kW | delta form for a named unit | 20.1.3 |
| *Posterior and questions* | | | | |
| Probability of the largest | $\Pr(u^\star = u \mid \mathbf{y})$ by the draws | – | which unit is largest | 20.2.2 |
| Hold probability | $\Pr(Q_L^\dagger \ge Q_d \mid \mathbf{y})$ | – | the drill answer | 20.3.1 |
| Joint against marginal | $\prod_k P_k \le P_\cap \le \min_k P_k$ | – | shared term, one rack decides | 20.3.2 |
| Supply that holds | $T_s^\dagger = T_{\max} - \max_k a_k(Q_d + Q_o)$ | °C | the reset | 20.3.3 |
| Chillers at a reset | $f(T_s)\sum_i 400\,\mathrm{e}^{\kappa_i} - Q_{o,\mathrm{site}}$ | kW | cost of the reset | 20.3.3 |
| Fourth unit | $Q_{L,4}^\dagger$; crossing year $\ln(Q_{L,4}^\dagger/300)/\ln 1.15$ | kW, yr | what a unit buys, and for how long | 20.3.4 |

## 20.8 Exercises

**Exercise 20.1.** *Formulation.* Write the planner's question *with unit B3 out, can hall B carry today's* 300 kW*?* as a probability over the reduced state, say which entries it depends on, and give its value. ($\Pr(Q^{\dagger}(\mathbf{z};\ \{\mathrm{B1},\mathrm{B2}\}) \geq 300\,\mathrm{kW} \mid \mathbf{y})$, with $Q^{\dagger}$ the smallest over the racks of $13/a_k - 10$ and $a_k = 1/S_G - (1-\beta_k)/20$, $S_G = 7\,\mathrm{e}^{g_{B1}} + 7\,\mathrm{e}^{g_{B2}}$. It depends on $g_{B1}$, $g_{B2}$ and the fractions, not on $g_{B3}$, the supply or the load. By the draws it is 0.0788.)

**Exercise 20.2.** *Derivation.* Derive the rise per kilowatt with the units $R$ running from the rows (20.2), and show that at equal units the binding load plus the other heat falls to two thirds of its value when one of three units stops. Estimate hall B's unit-out binding load from Chapter 5's. (The hall balance over $R$ gives $T_r^* - T_s^* = (Q^* + Q_o)/S_G$; the air-side rows over $R$ summed, with the hall balance and the mixing row, give $T_r^* - T_c^* = (Q^* + Q_o)/S_C$; the inlet row then gives $a_k = 1/S_G - (1-\beta_k)/S_C = (1 - (1-\beta_k)S_G/S_C)/S_G$. With $\bar{G}_u/C_u = 0.7$ for every unit, $S_G/S_C$ is the same with two units or three, and $1/S_G$ grows by 3/2; so $Q^{\dagger} + Q_o = 13/a_k$ falls to 2/3. From the all-running median, $(395.21 + 10) \times 2/3 - 10 = 260.1$ kW, between the B2-out median of 256.02 kW and the B3-out median of 265.54 kW.)

**Exercise 20.3.** *Derivation.* Show that, with one unit out of three equal-airflow units, the largest unit out is the worst unit out at every state, and that the binding rack does not depend on which unit is out. Give the probability the draws find. (With any one unit out $S_C = 20$, so $a_k$ depends on the unit only through $S_G$, and $\partial a_k/\partial S_G = -1/S_G^2 < 0$: removing the largest conductance leaves the smallest $S_G$ and the largest $a_k$ for every rack, the smallest binding load. At one state every rack shares $S_G$ and $S_C$, so $a_k$ is largest for the largest $\beta_k$ whichever unit is out. The draws find the largest unit out to be the worst in every draw, Pr = 1.0000.)

**Exercise 20.4.** *Application.* A reviewer multiplies the five racks' probabilities of holding at 340 kW, with the largest unit out and the supply reset to 12°C, and reports the result as the hall's. Say what is wrong, and give both numbers. (The racks' events share the posterior of $S_G$, so they hold and fail together, and their product understates the probability that all five hold; the smallest marginal, rack B3's 0.0652, bounds it from above. The product of the marginals is 0.0215; the joint tail by the draws is 0.0303.)

**Exercise 20.5.** *Application.* A vendor proposes a fourth hall-B unit to make the hall N+1. Give the probability that it holds the design load with the largest of the four out, say whether it meets an assurance of 0.95, and say how long it lasts at 15 % a year from today's 300 kW. (With the new unit's log-ratio on its prior $\mathcal{N}(0,\ 0.15^2)$, the probability is 0.9034 by the draws and 0.9024 by the exact posterior, below 0.95; it would be 0.9641 for a unit exactly at its rating. The load meets the unit-out binding load, median 382.97 kW, in the crossing year $\ln(382.97/300)/\ln 1.15$, a median of year 1.75.)

# Chapter 21

# Cascading Overloads

## 21.1 A chiller trip on the design day

The site's chilled-water plant has three chillers of 400 kW nameplate, and the design sheet checks what happens when one of them trips. Two chillers are left, 800 kW of nameplate, against the heat of the loop: today's 660 kW of IT load and the halls' other heat, 682 kW in all, a ratio of $800/682 = 1.173$; and at year 1, when the IT load has grown to 759 kW, 781 kW in all, a ratio of 1.024. The sheet's rule is that the two running chillers take the tripped unit's share as long as the ratio stays above one, so that neither is pushed past its rating, and it records the plant as tolerating a single trip through the plan year. A second trip, on the sheet, needs a ratio below one.

The ratio counts nameplate, and it counts only the chillers. A chiller that trips on a hot day leaves its share to units that were built a little short of their nameplate, and what happens to the rest of the site when they cannot carry it is not on the sheet. The planner's question is the drill question, verbatim,

*If one unit trips, does the load shift onto the others push a second one past its envelope?*

It is answered with the method of Chapter 4, on the site's one fit of Chapter 9 and the two halls' fits of Chapters 5 and 15. What is new is the device. A single scenario node pins a load and reads the plant there; here the planner needs a chain of them, each pinned by the output of the one before: the trip removes a chiller, the shortfall it leaves floats the chilled-water supply, and the floated supply meets two envelopes, the running chillers' motors and the racks' inlets. The answer is the probability of the second event given the first.

**Jargon: cascading failure (cascading overload), first trip, second excursion, chain of scenario nodes, set point, supply float, capacity–supply slope, motor envelope, inlet margin, fault tolerance, hall-unit chain, plan year, design hours, expected count**

A *cascading failure*, or cascading overload (in the framework, a scenario chain), is a failure whose shifted load drives a second element past its envelope. The *first trip* is the outage that starts it; the *second excursion* is any element pushed past its limit by the load the first trip shifts. A *chain of scenario nodes* is a sequence of scenario nodes of Section 4.3.8 in which each node's pinned input is the solved output of the node before. The *set point* of the chilled-water supply is the temperature the chillers hold when they can, here the drawing's 14°C; the *supply float* $\Delta T_s$ is the rise of the supply above it when the running chillers cannot carry the loop's heat at the set point. The *capacity–supply slope* $\gamma$ is the fraction by which a chiller's capacity grows per kelvin of warmer supply. A chiller's *motor envelope* is the largest duty its compressor motor carries, as a multiple of the unit's as-built duty at the set point. A hall's *inlet margin* $m_h$ is the 27°C inlet limit less the hottest inlet of the hall at the set point: the supply rise the hall can absorb before an inlet crosses the limit. *Fault tolerance* to a single trip is the property that no second excursion follows any one first trip. The *hall-unit chain* is the same chain started by a hall's cooling unit instead of a chiller. The *plan year* is the twelve months from today, in which the load grows from 660 kW to 759 kW. The *design hours* are the stated 200 h a year spent at or near the design day. An *expected count* is the expected number of events in a stated period; it is not a rate of failure of any element.

**Key idea**

A cascade is a chain of scenario nodes. The first node pins the trip and solves the plant for what the trip shifts; the next pins that shift and solves for the second element's load against its envelope. Every link is a function of the posterior of the state, so the probability of the second excursion given the first is one integral over that posterior, joint over every envelope the shifted load can cross. The envelope the design sheet checks need not be the one that breaks first, and the chain finds which does. Multiplied by the rate of first trips and summed over the months, the conditional probability becomes an expected count of cascades in the plan year.

### 21.1.1 A cascade read as a chain of scenario nodes

Chapter 9 enumerated the states of the chillers, up or down, and read the capacity of each. A cascade asks a question that no single state answers: given that the plant has just entered a state, does the load of the element that left push the plant into a worse one? The answer is a chain. Write $F$ for the first event, the trip of one unit, and $S$ for the second, any element of the plant past its limit after the load has moved. The first event pins a scenario node, which solves the plant for the redistributed load $x^*(\mathbf{z})$; the second event is a threshold on that load, $S = \{x^*(\mathbf{z}) > \vartheta(\mathbf{z})\}$, with $\vartheta$ the smallest of the envelopes it meets. Both are functions of the state, so the probability of the second given the first is the law of total probability of equation (3.2) over the posterior, and the expected number of cascades over a period $T$ is the rate $\nu_F$ of first events times that conditional probability,

$$\Pr(S \mid F,\, \mathbf{y}) = \int \mathbf{1}\big[x^*(\mathbf{z}) > \vartheta(\mathbf{z})\big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}, \qquad \mathbb{E}[N_S \mid \mathbf{y}] = \mathbb{E}\big[\nu_F\, T\, \Pr(S \mid F,\, \mathbf{z}) \mid \mathbf{y}\big], \tag{21.1}$$

with $N_S$ the number of second excursions in the period. The first relation is the threshold form of a single node with the threshold taken after the load has moved; the second is a count, never a rate of failure of the second element, since it counts first events that happen to be followed by a second.

**Principle 21.1.** *A cascade is the probability of a second excursion given the first, read along a chain of scenario nodes from the posterior of the state. The second link is the smallest of every envelope the shifted load meets, so the probability is joint over all of them. Its expected count in a period is the rate of first events times that probability.*

### 21.1.2 The site's chillers and the halls' inlets (step 0)

The plant is the site of Chapter 15, drawn in Figure 15.1: two halls on one chilled-water plant with three chillers C1–C3, one pump, the shared supply pipe and the power feed, taken as built, without the drawn economizer. Each hall has three cooling units and five racks. The chain reads three parts of it: the chillers' capacities, through the site's state (9.4),

$$\mathbf{z} = (\kappa_{C1},\ \kappa_{C2},\ \kappa_{C3},\ d,\ g_e,\ c,\ l_k,\ l_e)^{\mathsf{T}},$$

and each hall's inlets, through its reduced state $\mathbf{z}_h$ of equation (5.4).

**The quantity the chain owes.** After chiller $i$ trips, the chain owes the probability of a second excursion, the event that the floated supply $\Delta T^*_{s,i}$ crosses at least one envelope: a running chiller's motor, which binds at a float $\Delta T_m$, or a hall's inlets, which bind at that hall's inlet margin $m_A$ or $m_B$. Each envelope is a threshold on the same float, so the union of the three events is one event at the smallest threshold,

$$E_i = \{\Delta T^*_{s,i} > \Delta T_m\} \cup \{\Delta T^*_{s,i} > m_A\} \cup \{\Delta T^*_{s,i} > m_B\} = \{\Delta T^*_{s,i} > \min(\Delta T_m,\ m_A,\ m_B)\}, \tag{21.2}$$

dimensionless as an event, with every term in kelvin. The rest of this subsection derives the three terms in turn: the float, the motor's threshold and the halls' margins.

**The float.** With chiller $i$ down, the two running chillers $j \neq i$ go to their limit, and the supply rises until their capacity, which grows with warmer supply, meets the loop's heat. Each running chiller's capacity at the set point is that of equation (9.3) at the wet bulb $T_{\mathrm{wb}}$,

$$\mathrm{Cap}_j(\mathbf{z}) = N_j\, \mathrm{e}^{\kappa_j}\,[1 - d\,(T_{\mathrm{wb}} - 22)],$$

with $N_j = 400\,\mathrm{kW}$ its nameplate. At a supply $\Delta T_s$ above the set point the capacity is larger by the pinned capacity–supply slope $\gamma = 0.03\,\mathrm{K}^{-1}$,

$$\mathrm{Cap}_j(\mathbf{z},\ \Delta T_s) = \mathrm{Cap}_j(\mathbf{z})\,(1 + \gamma\,\Delta T_s).$$

The heat the running chillers must carry is the loop's, the pinned IT load $L^*$ plus the halls' other heat $Q_o = 22\,\mathrm{kW}$,

$$Q_{\mathrm{ch}} = L^* + Q_o,$$

682 kW at today's load and 781 kW at year 1's. The float is the supply at which the running capacity meets that heat, a row that vanishes when it holds,

$$r_i(\Delta T_s;\ \mathbf{z}) = \mathrm{Cap}_{\mathrm{rem},i}(\mathbf{z})\,(1 + \gamma\,\Delta T_s) - Q_{\mathrm{ch}} = 0, \qquad \mathrm{Cap}_{\mathrm{rem},i}(\mathbf{z}) = \sum_{j \neq i} \mathrm{Cap}_j(\mathbf{z}), \tag{21.3}$$

in kW, with $\mathrm{Cap}_{\mathrm{rem},i}$ the capacity left when $i$ trips. The row is linear in the float, with slope $\partial r_i / \partial \Delta T_s = \gamma\,\mathrm{Cap}_{\mathrm{rem},i} > 0$, so it has one root, and the supply stays at its set point whenever the running chillers carry the heat there,

$$\Delta T^*_{s,i}(\mathbf{z}) = \frac{1}{\gamma}\,\max\!\left(\frac{Q_{\mathrm{ch}}}{\mathrm{Cap}_{\mathrm{rem},i}(\mathbf{z})} - 1,\ 0\right), \tag{21.4}$$

in kelvin. Where the supply floats, its gradient with respect to the state follows from the row by the implicit function theorem of equation (5.1), the row's gradient in the state over its slope in the float,

$$\nabla_{\mathbf{z}}\,\Delta T^{*}_{s,i} = -\,\frac{(1+\gamma\,\Delta T^{*}_{s,i})\,\nabla_{\mathbf{z}}\mathrm{Cap}_{\mathrm{rem},i}}{\gamma\,\mathrm{Cap}_{\mathrm{rem},i}} = -\,\frac{Q_{\mathrm{ch}}}{\gamma\,\mathrm{Cap}^{2}_{\mathrm{rem},i}}\,\nabla_{\mathbf{z}}\mathrm{Cap}_{\mathrm{rem},i}, \tag{21.5}$$

in K per unit of each entry. A larger remaining capacity lowers the float, and the float is steepest where the remaining capacity is smallest.

**The motor's threshold.** A running chiller $j$ at the floated supply carries the capacity it reaches there, $(1+\gamma\,\Delta T_s)\,\mathrm{Cap}_j$, and its motor is sized for the stated envelope $e_m = 1.10$ of its as-built duty at the set point, $e_m\,\mathrm{Cap}_j$. The motor trips when the first exceeds the second; $\mathrm{Cap}_j$ cancels, so both running chillers reach their envelope at the same float,

$$\begin{aligned}(1+\gamma\,\Delta T_s)\,\mathrm{Cap}_j > e_m\,\mathrm{Cap}_j \;&\Longleftrightarrow\; \Delta T_s > \Delta T_m = \frac{e_m-1}{\gamma} = \frac{0.10}{0.03} = 3.33\,\mathrm{K}\\ &\Longleftrightarrow\; \mathrm{Cap}_{\mathrm{rem},i} < \frac{Q_{\mathrm{ch}}}{e_m},\end{aligned} \tag{21.6}$$

the last form by equation (21.4). The motor's threshold is a fixed float, the same in every state; whether the float reaches it is a question about the remaining capacity alone.

**The halls' margins.** Each hall's inlets at the scenario node are those of equation (5.7), now with the supply at the set point plus the float. Every inlet of hall $h$ is the supply plus a rise that does not depend on the supply,

$$T^{*}_{k}(\mathbf{z}_h,\ Q^{*}_h,\ \Delta T_s) = T^{*}_s + \Delta T_s + a_k(\mathbf{z}_h)\,(Q^{*}_h + Q_{o,h}),$$

with $T^{*}_s = 14^{\circ}\mathrm{C}$, $a_k$ the rise per kilowatt of rack $k$, in K/kW, $Q^{*}_h$ the hall's IT load and $Q_{o,h}$ its other heat. Every inlet rises one for one with the supply. The hottest inlet meets the limit $T_{\max} = 27^{\circ}\mathrm{C}$ when the float equals the hall's margin,

$$m_h(\mathbf{z}_h) = T_{\max} - T^{*}_s - \max_{k\in h}\,a_k(\mathbf{z}_h)\,(Q^{*}_h + Q_{o,h}), \qquad T^{*}_k > T_{\max}\text{ for some }k \;\Longleftrightarrow\; \Delta T_s > m_h, \tag{21.7}$$

in kelvin. The hall's IT load is its share of the site's, $Q^{*}_h = \varphi_h\,L^{*}$ with $\varphi_A = 0.5455$ and $\varphi_B = 0.4545$ of equation (15.5). A margin below zero is a hall already past its limit with every chiller running, before any trip. Substituting equations (21.4), (21.6) and (21.7) into equation (21.2) gives the second excursion as one inequality on the three states,

$$E_i = \Big\{\tfrac{1}{\gamma}\max(Q_{\mathrm{ch}}/\mathrm{Cap}_{\mathrm{rem},i}(\mathbf{z}) - 1,\ 0) > \min(\Delta T_m,\ m_A(\mathbf{z}_A),\ m_B(\mathbf{z}_B))\Big\}.$$

**The hall-unit chain.** The same chain can start in a hall. When unit $u$ of hall $h$ trips, its coil and fan stop, and the running units $R = h\setminus\{u\}$ take its heat at the set point. By equation (5.5) with the sums over the running units, the return air sits above the supply by the hall's heat over their conductance, and each running coil passes its conductance's share,

$$T^{*}_r - T^{*}_s = \frac{Q^{*}_h + Q_{o,h}}{\sum_{y\in R} G_y}, \qquad Q^{*}_x = G_x\,(T^{*}_r - T^{*}_s) = \frac{G_x}{\sum_{y\in R} G_y}\,(Q^{*}_h + Q_{o,h}), \quad x\in R, \tag{21.8}$$

in K and kW, with $G_x = \bar{G}_x\,\mathrm{e}^{g_x}$ the as-built conductance of unit $x$. The coil duty $Q^{*}_x$ is set against the sheet's duty of one unit, the stated 180 kW in hall B and 210 kW in hall A; the inlets with unit $u$ out are those of equation (15.4). Chapter 20 reads hall B's inlets with its largest unit out; here the hall-unit chain stands beside the chillers' as the second place a first trip can start.

**Modeling choices at the graph.** The float is the steady state after the trip: this makes it the root of one row, equation (21.3), and it fails in the minutes in which the loop's water and the halls' air store heat, which delay the float, and in which a tripped chiller restarts, which ends it. The running chillers' capacity grows linearly with the supply at the pinned slope of 0.03 per kelvin: this gives the closed form (21.4), and it fails far from the set point, where the compressor's lift changes the slope. The motor envelope is a stated multiple of each unit's as-built duty: this makes both running units reach it together, at a float that does not depend on the state, and it fails if a unit's motor was sized to its nameplate instead, which would give each unit its own threshold. Every inlet rises one for one with the supply: by equation (5.7) this holds with the units' airflow fixed and no controller moving their valves, and it fails if a unit opens its valve to hold its discharge air as the supply warms. The chain stops at the second excursion: what the plant does after a second chiller trips or a hall crosses its limit is a third link that this chapter does not follow. The other elements of Chapter 13, the pipe, the pump and the feed, are held at their healthy state.

### 21.1.3 The chain's scenario nodes

The chain is asked at three kinds of node. On the design day the wet bulb is pinned at 22°C, the chillers' rating condition, at which the derate factor of $\mathrm{Cap}_j$ is one; the load is pinned at three values, today's $L^* = 660\,\mathrm{kW}$, 700 kW, and year 1's 759 kW on the load path of Chapter 6. Through the plan year the wet bulb of month $m$ is the record year's monthly mean and the load grows month by month from today's at 15 % a year,

$$L^*_m = 660 \times 1.15^{\,m/12}\,\mathrm{kW}, \qquad m = 1, \ldots, 12, \tag{21.9}$$

from 667.7 kW in January to 759.0 kW in December; the design day, at year 1's load, bounds it. At every node the first trip is pinned to one chiller $i$ in turn, the supply is free above its set point, and both halls run all three units. A design alternative adds a fourth chiller, the 400 kW unit of the existing model of Chapter 9, with its as-built log-ratio $\kappa_n$ of equation (9.7).

**The rate of first trips.** Each chiller trips at the rate $\lambda$ of the outage record, the same for all three, so the first trip is equally likely to be any of them, and the probability of a second excursion averaged over the first trip is the plain mean,

$$\bar{P}(\mathbf{z},\ \mathbf{z}_A,\ \mathbf{z}_B) = \frac{1}{3}\sum_{i=1}^{3} \mathbf{1}[E_i], \qquad N_2 = \sum_{m=1}^{12} 3\,\lambda\,h\,\bar{P}_m, \tag{21.10}$$

with $\bar{P}_m$ the mean at month $m$'s node, $h = 730\,\mathrm{h}$ the hours in a month and $3\lambda h$ the expected number of first trips in the month, so that $N_2$ is the expected number of second excursions in the plan year given the state and the rate; by equation (21.1) its expectation over the posterior is the answer. The design hours enter the same way, $3\lambda \times 200\,\mathrm{h} \times \bar{P}$ at the design day's node.

**The sheet's arithmetic as a node.** The design sheet's ratio is equation (21.4) with nameplate in place of the as-built capacity, $\mathrm{Cap}_{\mathrm{rem},i} = 800\,\mathrm{kW}$: at $Q_{\mathrm{ch}} = 781\,\mathrm{kW}$ the ratio 781/800 is below one, the float is zero, and the sheet sees no chain. It does not check the halls at all.

### 21.1.4 The planner's questions (step 1)

The planner asks five questions of the chain, and step 1 writes each over the unknown $\boldsymbol{\theta} = (\mathbf{z},\ \mathbf{z}_A,\ \mathbf{z}_B,\ \lambda)$ of the site, the two halls and the chillers' trip rate, with the evidence on the right of the bar. Table 21.1 lists them. The first column is the question in the planner's words,

Table 21.1: The questions asked of a trip on the site. $\mathbf{y}$ is the site's operating record, the chillers' outage record, both halls' window 1 and the priors; $E_i$ is the second excursion after chiller $i$ trips, equation (21.2); $\bar{P}$ its mean over the first trip and $N_2$ the expected count, equation (21.10); $Q_x^*$ a running hall unit's duty, equation (21.8).

| Planner's question | General form | Mathematical form | Category | Computation | Answered in |
|---|---|---|---|---|---|
| If one unit trips, does the load shift onto the others push a second one past its envelope? | $\Pr(f \geq f^* \mid \mathbf{y})$ along a scenario chain, a conditional probability | $\mathbb{E}[\bar{P} \mid \mathbf{y}]$, and the chiller term alone, at 660, 700, 759 kW on the design day | prediction under a threshold | posterior | 21.3.1 |
| Which first trip matters? | the same, for each first event | $\Pr(E_i \mid \mathbf{y})$ for each $i$; $\Pr(\mathrm{Cap}_{C2} = \max_j \mathrm{Cap}_j \mid \mathbf{y})$ | prediction under a threshold | posterior | 21.3.2 |
| How often a year? | moments, $\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$ | $\mathbb{E}[N_2 \mid \mathbf{y}]$ over the plan year, with its band; the design hours beside | estimation | posterior | 21.3.3 |
| Does the hall side cascade? | $\Pr(f \geq f^* \mid \mathbf{y})$ along the hall-unit chain | $\Pr(\max_x Q_x^* > Q_{\mathrm{sheet}})$ and $\Pr(\max_k T_k^* > 27^\circ\mathrm{C})$, a unit out | prediction under a threshold | posterior | 21.3.4 |
| What does Chapter 9's fourth chiller do to the chain? | the same under a design alternative | the first question with four chillers; $\Pr(m_B < 0 \mid \mathbf{y})$ | prediction under a threshold | posterior | 21.3.5 |

the drill question first; the second the general form of its category in the notation of Chapter 2; the third its mathematical form for the site; 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 drill question is a prediction under a threshold taken along the chain, a conditional probability; the second splits it by the first trip; the third turns it into an expected count, an estimation; the fourth asks it of the halls' own units; and the fifth asks it of the drawing with a fourth chiller. The atlas of Appendix A collects them with the other questions of Volume II.

### 21.1.5 The evidence as rows (step 2)

Step 2 attaches four records to the graph, and they share no entry. The site's operating record of Chapter 9, two capacity tests per chiller, twelve monthly windows and eight priors, 62 rows on the 8 entries of $\mathbf{z}$, reads the chillers' log-ratios and the derate. The chillers' outage record, 13 trips in 12 unit-years, reads the trip rate $\lambda$ through the gamma posterior of equation (9.12). Hall B's window 1 of the night record, the rows of equation (5.11), reads $\mathbf{z}_B$; hall A's window 1, the same rows with its own readings (Section 15.2.2), reads $\mathbf{z}_A$.

**Which row acts on which entry.** The float reads the site's state through the remaining capacity alone. By equation (21.5), with $\mathrm{Cap}_j$ of the running units and $c_i = Q_{\mathrm{ch}}/(\gamma\,\mathrm{Cap}_{\mathrm{rem},i}^2)$, the float after C2 trips has the gradient

$$\nabla_{\mathbf{z}}\,\Delta T_{s,2}^{*\mathsf{T}} = -\,c_2\,\Big(\mathrm{Cap}_{C1},\ 0,\ \mathrm{Cap}_{C3},\ -\textstyle\sum_{j\neq 2} N_j\,\mathrm{e}^{\kappa_j}\,(T_{\mathrm{wb}}-22),\ 0,\ 0,\ 0,\ 0\Big), \tag{21.11}$$

in the order of $\mathbf{z}$. Three facts are read here before any number. The float after a trip reads the other two units' log-ratios and not the tripped unit's: what is lost matters only through what is left. The derate's entry carries the factor $T_{\mathrm{wb}} - 22$, which vanishes on the design day, so the

design-day chain reads two log-ratios and nothing else of the site; through the plan year's cooler months it reads the derate too. And the margins read the halls' states alone, the float the site's alone, and the expected count the outage record alone, so the four records enter the answer through separate factors.

**Modeling choices at the rows.** The rows are those of the cited fits, with their modeling choices. The trip rate is the outage record's single rate for the three units: this is what makes the first trip uniform over them, and it fails if one unit trips more often, for instance the strongest, which carries the most load. The motor envelope $e_m = 1.10$ and the capacity–supply slope $\gamma = 0.03$ are stated inputs, not rows: no record of motor current or of capacity at a warm supply is kept, so neither has a posterior. The sheet's unit duties of 180 kW in hall B and 210 kW in hall A are stated.

## 21.2 The posterior of the state (step 3)

All five questions are read from one posterior, that of the unknown $\boldsymbol{\theta}$ under the four records of step 2. No row is added and no fit is repeated; this section states what the chain needs from the posterior, cites the fits, and gives the draws the answers use.

### 21.2.1 What the chain needs from the state

By equation (21.2) the second excursion needs the remaining capacity of each pair of chillers, read from $\mathbf{z}$; the two margins, read from $\mathbf{z}_A$ and $\mathbf{z}_B$; and, for the count, the rate $\lambda$. The four records share no entry, so the posterior factors as in equation (15.9),

$$p(\boldsymbol{\theta} \mid \mathbf{y}) = p(\mathbf{z} \mid \mathbf{y}_{\mathrm{op}})\, p(\mathbf{z}_A \mid \mathbf{y}_A)\, p(\mathbf{z}_B \mid \mathbf{y}_B)\, p(\lambda \mid \mathbf{y}_{\mathrm{out}}), \tag{21.12}$$

and the probability of a second excursion after chiller $i$ trips is the threshold of equation (21.1) over it,

$$\begin{aligned} \Pr(E_i \mid \mathbf{y}) = \int \mathbf{1}\Big[\Delta T^{*}_{s,i}(\mathbf{z}) > \min\big(\Delta T_m,\ m_A(\mathbf{z}_A),\ m_B(\mathbf{z}_B)\big)\Big] \\ \times\, p(\mathbf{z} \mid \mathbf{y}_{\mathrm{op}})\, p(\mathbf{z}_A \mid \mathbf{y}_A)\, p(\mathbf{z}_B \mid \mathbf{y}_B)\, \mathrm{d}\mathbf{z}\, \mathrm{d}\mathbf{z}_A\, \mathrm{d}\mathbf{z}_B. \end{aligned} \tag{21.13}$$

The rate does not enter it: the conditional probability is a property of the plant given the trip, and the rate enters only the count.

### 21.2.2 The fits, cited

The site's state is the one fit of Section 9.2.2: the MAP state $\hat{\mathbf{z}}$ minimizes half the squared length of the residual vector of the 62 rows by the normal equations of equation (4.29), and its Laplace covariance is $\boldsymbol{\Sigma} = (\mathbf{J}^{\mathsf{T}}\mathbf{J})^{-1}$ at $\hat{\mathbf{z}}$, equation (4.38), with the entries of Table 9.2. Hall B's is the fit of Section 5.2.2, $(\hat{\mathbf{z}}_B, \boldsymbol{\Sigma}_B)$ on its window 1, and hall A's that of Section 15.2.2. The trip rate's posterior is the gamma of equation (9.12), exact.

Table 21.2 lists what the chain reads from them. Each row is a quantity; the columns give its median, or its mean for the rate, and its 90 % band. The as-built capacities at the design day are those whose margin Chapter 16 set against the drawing's: all three below their 400 kW nameplate, C2 the strongest. Over the draws C2 is the strongest unit with probability 0.9924, so that losing C2 would leave the smallest remainder. The trip rate's mean is 1.124 trips a year per unit, from 13 trips in 12 unit-years. The halls' inlet margins at the set point fall with the load: hall B's from a median of 3.07 K at 660 kW to 1.63 K at 759 kW, with a 5 % point of 0.06 K there, so that in a twentieth of the draws hall B has almost no margin left at year 1's load before any trip. Hall A, the larger hall, keeps more.

Table 21.2: What the chain reads from the posterior: the chillers' as-built capacities at the design day (the site's fit), the trip rate (the outage record's gamma posterior, mean) and the halls' inlet margins (21.7) at the set point of 14°C with every unit running (each hall's fit), by site IT load.

| Quantity [unit] | Median [unit] | 5 % point [unit] | 95 % point [unit] |
|---|---|---|---|
| $\mathrm{Cap}_{C1}$ at 22°C [kW] | 365.3 | 356.1 | 375.0 |
| $\mathrm{Cap}_{C2}$ at 22°C [kW] | 384.5 | 373.7 | 395.8 |
| $\mathrm{Cap}_{C3}$ at 22°C [kW] | 356.2 | 345.4 | 367.5 |
| Trip rate $\lambda$ per unit, mean [1/yr] | 1.124 | 0.693 | 1.636 |
| *Inlet margin $m_h$ at 14°C, every unit running [K]* | | | |
| Hall A, site 660 kW | 4.26 | 2.71 | 5.17 |
| Hall B, site 660 kW | 3.07 | 1.70 | 4.25 |
| Hall A, site 700 kW | 3.74 | 2.10 | 4.71 |
| Hall B, site 700 kW | 2.49 | 1.04 | 3.74 |
| Hall A, site 759 kW | 2.99 | 1.21 | 4.03 |
| Hall B, site 759 kW | 1.63 | 0.06 | 2.98 |

### 21.2.3 The draws

The answers read the posterior by draws, the state route. The site's state is drawn 20 000 times from the Laplace normal $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma})$, the draws of Section 9.2.4 with the seed 9, each with its draw of the trip rate from the gamma posterior. Their weights toward the exact posterior, equation (5.16) with the objective of the site's fit, have an effective sample size of 19 913 of 20 000. Each hall's state is drawn 4 000 times from its Laplace normal, equation (5.13), with the seed 21, and draw $n$ of each hall is paired with draw $n$ of the site. Every quantity of the site alone, the float, its probability and the motor's link, is read over the 20 000; every quantity that meets a margin, the halls' links and the second excursion itself, over the first 4 000, paired. In each draw the float is the closed form (21.4) and the margins are equation (21.7) at the node's load, and the indicator of equation (21.13) is evaluated exactly; the reported probabilities are the shares of the draws in which it is one.

## 21.3 The answers (step 4)

Every answer is read from the draws of Section 21.2.3. The drill question comes first, on the design day; the second splits it by the first trip; the third carries it through the plan year; the fourth and fifth ask it of the halls' units and of a fourth chiller.

### 21.3.1 The second excursion after a first trip

**Planner's question.** If one unit trips, does the load shift onto the others push a second one past its envelope? The question asks for the probability that a first chiller trip on the design day is followed by a second excursion, and for the share of it that is a second chiller.

**General form.** A prediction under a threshold is the posterior probability of an event, equation (2.12),

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

here taken after the first event has moved the load, the conditional probability of equation (21.1).

Table 21.3: The chain on the design day, by site IT load and first trip: the float (21.4) and the second excursion of each envelope, equations (21.6) and (21.7), and of any of them, equation (21.2). The float's four columns and the motor's over the 20 000 draws of the site, where $< 10^{-4}$ marks a share below one in ten thousand; the halls' columns and the last over the 4 000 paired draws, where "none" marks no draw.

| Load [kW] | First trip [–] | Pr (float) [–] | $\Delta T_s^*$ 5 % [K] | $\Delta T_s^*$ median [K] | $\Delta T_s^*$ 95 % [K] | Pr (motor) [–] | Pr (hall B) [–] | Pr (hall A) [–] | Pr (any) [–] |
|---|---|---|---|---|---|---|---|---|---|
| 660 | C1 | $< 10^{-4}$ | 0.00 | 0.00 | 0.00 | $< 10^{-4}$ | none | none | none |
| 660 | C2 | 0.0001 | 0.00 | 0.00 | 0.00 | $< 10^{-4}$ | none | none | none |
| 660 | C3 | $< 10^{-4}$ | 0.00 | 0.00 | 0.00 | $< 10^{-4}$ | none | none | none |
| 700 | C1 | 0.0444 | 0.00 | 0.00 | 0.00 | $< 10^{-4}$ | 0.0018 | none | 0.0018 |
| 700 | C2 | 0.5147 | 0.00 | 0.02 | 0.77 | $< 10^{-4}$ | 0.0040 | 0.0005 | 0.0045 |
| 700 | C3 | 0.0017 | 0.00 | 0.00 | 0.00 | $< 10^{-4}$ | 0.0015 | none | 0.0015 |
| 759 | C1 | 0.9997 | 0.94 | 1.81 | 2.69 | 0.0030 | 0.5685 | 0.1635 | 0.6212 |
| 759 | C2 | 1.0000 | 1.95 | 2.74 | 3.56 | 0.1145 | 0.8700 | 0.4258 | 0.9050 |
| 759 | C3 | 0.9991 | 0.66 | 1.38 | 2.14 | $< 10^{-4}$ | 0.4068 | 0.0848 | 0.4500 |

**Specialized form.** The question is equation (2.12) with the event $E_i$ of equation (21.2), the float $f = \Delta T_{s,i}^*$ against the threshold $f^* = \min(\Delta T_m, m_A, m_B)$, a threshold that is itself a function of the halls' states, over the factored posterior (21.12): equation (21.13) for each first trip, averaged over the three by equation (21.10). The reported value is joint over the three envelopes. Beside it stand the worst marginal, the event on hall B's inlets alone, $\Delta T_s > m_B$, whose gap from the joint is the share of the other two envelopes; and, for the joint and for the motor's link, the exact posterior, the same draws under the site's weights, with the gap. The weights correct the site's factor of the posterior only; the halls' factors stay on their Laplace normals.

**Building blocks.** The float (21.4); the motor's threshold (21.6); the margins (21.7); the paired draws of Section 21.2.3 and their weights.

**The chain on the design day.** Table 21.3 gives the chain at three loads. Each row is a load and a first trip; the columns give the probability that the supply floats at all, the 5 % point, median and 95 % point of the float, and the probability of each envelope's excursion and of any of them. Today, at 660 kW, the loop's 682 kW sits below what any two chillers carry: the supply floats in at most one draw in ten thousand whichever unit trips, and no draw of the 4 000 crosses any envelope. At 700 kW the loop's $700 + 22 = 722$ kW sits at the capacity C1 and C3 leave, the sum of their medians $365.3 + 356.2 = 721.5$ kW, and the supply floats after a C2 trip in about half the draws, by a median of 0.02 K; the second excursion stays below 0.005 for every first trip. At year 1's 759 kW every first trip floats the supply in essentially every draw.

**The answer.** Averaged over the first trip at year 1's load, the probability that a first trip is followed by a second excursion is

$$\begin{aligned} \Pr(\text{second excursion} \mid \text{first trip},\ \mathbf{y}) &= 0.6587 \quad (\text{exact } 0.6642), \\ \Pr(\text{second chiller} \mid \text{first trip},\ \mathbf{y}) &= 0.0392 \quad (\text{exact } 0.0414), \end{aligned} \tag{21.14}$$

at 759 kW on the design day. Today the answer is no: at 660 kW no draw shows a second excursion after any first trip, and the site is single-fault tolerant on its chillers. At 700 kW the averaged probability is 0.0026 (exact 0.0027). At 759 kW it is yes, about two first trips in three, and the second excursion is almost never a second chiller.

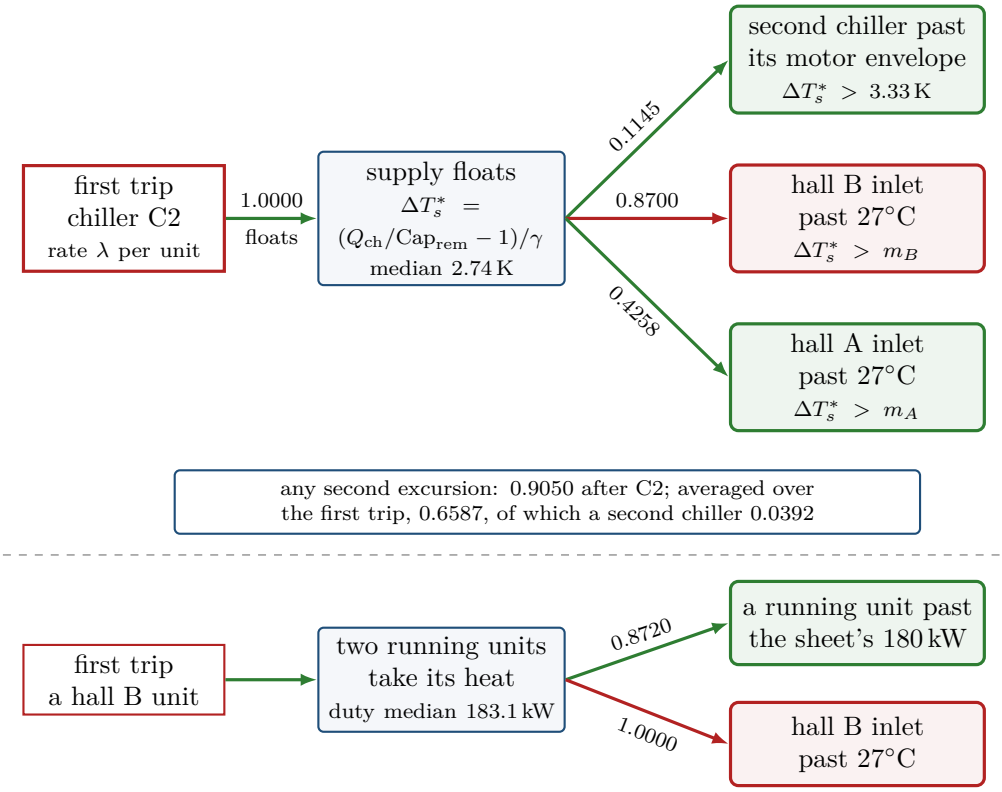



Figure 21.1: The cascade as a chain of scenario nodes at year 1's site IT load of 759 kW on the design day. Upper rows: chiller C2 trips, the two running chillers go to their limit and the supply floats by equation (21.4); the float meets three envelopes, each arrow labeled with the probability of crossing it (Table 21.3); the red path is the one the posterior favors. Lower row, beside: a hall B unit trips and the two running units take its heat, equation (21.8); their duty is set against the sheet's and the inlets against 27°C (Table 21.6).

The chain is drawn in Figure 21.1, with C2 as the first trip. The boxes are the nodes: the trip, the float it causes, and the three envelopes the float meets; each arrow carries the probability that the chain passes along it at 759 kW. A second row draws the hall-unit chain beside it, answered in Section 21.3.4.

This is the finding of the chapter. The cascade's second link is thermal, at the racks, and not a second chiller. After C2 trips at 759 kW the supply floats by a median of 2.74 K, below the motors' 3.33 K, which it passes in 0.1145 of the draws; but hall B's margin at that load is 1.63 K at the median, and the float crosses it in 0.8700 of the draws. The design sheet checks the chillers' ratio, which is the envelope the float reaches last. The racks, which the sheet never checks, are the envelope it reaches first.

The float and the margin it crosses are drawn in Figure 21.2. The horizontal axis is in kelvin, the float above the set point for the one curve and the margin for the other; the vertical axis is the posterior density. The red curve is the float after C2 trips, the blue curve hall B's margin at 759 kW, and the dashed line the motors' 3.33 K, with the share of the float beyond it shaded. The dots mark the points the draws print, and two dotted lines the stated truth of Section 21.4.2.

The two curves overlap only in part. The float's 90 % band runs from 1.95 to 3.56 K, the margin's from 0.06 to 2.98 K, and the event of the second link, the float to the right of the margin, holds in 0.8700 of the paired draws. The float crosses the motors' line only in its upper tail. The margin's band is the wider of the two: its ends lie 1.57 and 1.35 K from its median,

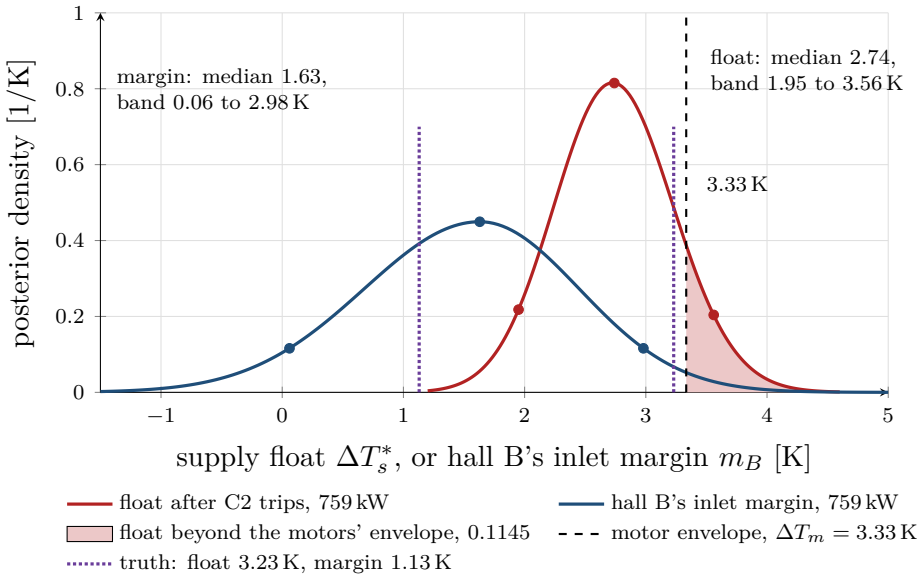



Figure 21.2: The supply float after chiller C2 trips at 759 kW on the design day (red) against hall B's inlet margin at the same load (blue). The float's density is equation (21.4) applied to a normal remaining capacity $\mathrm{Cap}_{C1} + \mathrm{Cap}_{C3}$ of standard deviation 9.79 kW (Section 21.6), centered so that its median is the draws' 2.74 K; its 5 and 95 % points are then the draws' 1.95 and 3.56 K (dots). The margin's curve is a sketch, not a density the draws print: a two-piece normal peaked at the draws' median, each half's spread set so that it reaches the draws' 5 or 95 % point (dots) at 1.645 of its standard deviations. The dashed line is the motors' threshold (21.6); the shaded tail is the float beyond it. The dotted lines are the stated truth: the true float after C2's trip and hall B's true margin, both at 759 kW. Most of the float lies to the right of most of the margin.

Table 21.4: The routes to the second excursion at 759 kW on the design day. Joint: the event (21.2) over the three envelopes, reported; worst marginal: hall B's inlets alone; exact: the site's draws under their weights (ESS 19 913), on the paired draws for the joint, with the halls' factors unweighted; the exact route is printed for the average over the first trip only. Gaps from the printed values; the design sheet uses nameplate.

| Route | After C1 [–] | After C2 [–] | After C3 [–] | First trip averaged [–] |
|---|---|---|---|---|
| Any second excursion, joint (reported) | 0.6212 | 0.9050 | 0.4500 | 0.6587 |
| Any second excursion, exact posterior | – | – | – | 0.6642 |
| Gap, exact minus joint | – | – | – | +0.0055 |
| Hall B's inlets alone, worst marginal | 0.5685 | 0.8700 | 0.4068 | – |
| Gap, joint minus worst marginal | +0.0527 | +0.0350 | +0.0432 | – |
| Second chiller, draws | 0.0030 | 0.1145 | $< 10^{-4}$ | 0.0392 |
| Second chiller, exact posterior | – | – | – | 0.0414 |
| Gap, exact minus draws | – | – | – | +0.0022 |
| Design sheet, ratio 800/781 = 1.024 | no float | no float | no float | no second trip |

against 0.79 and 0.82 K for the float's, so that most of what is not known about the second link is the hall's margin, not the chillers. The truth sits on the side the event favors: its float lies right of the float's median and its margin left of the margin's.

**The routes.** Table 21.4 sets the routes side by side. Each column is a first trip, or the average over them; the rows give the joint answer, its exact value and their gap, the worst marginal and its gap from the joint, then the motor's link by the draws and by the exact posterior with their gap, and the design sheet. The joint exceeds hall B's marginal by 0.0527, 0.0350 and 0.0432 after C1, C2 and C3: the share of hall A and the motors, the draws in which one of them binds below hall B's margin. That share is small because the three envelopes are thresholds on the same float, and a float that crosses hall A's margin or the motors' threshold has, in most draws, already crossed hall B's. After C2 the Monte Carlo standard error of the joint is $\sqrt{0.905 \times 0.095 / 4\,000} = 0.0046$, and the gap is more than seven of them. For the joint, averaged over the first trip, the site's weights on the paired draws move the answer from 0.6587 to 0.6642, a gap of +0.0055, in the same direction as the motor's link. For the motor's link the exact posterior moves the average from 0.0392 to 0.0414, a gap of +0.0022, against a Monte Carlo standard error of at most $\sqrt{0.0392 \times 0.9608 / 20\,000} = 0.0014$; the two use the same draws, so the gap is the departure of the posterior from the Laplace normal in the tail where C1 and C3 are both weak. The motor's link has a closed threshold beside it: by equation (21.6) a running chiller trips after C2's trip when $\mathrm{Cap}_{C1} + \mathrm{Cap}_{C3} < 781/1.10 = 710.0$ kW, and the sum of their medians, 721.5 kW, lies 11.5 kW above that, a little over one standard deviation of 9.79 kW: about one draw in nine falls below it, the 0.1145 of the table.

### 21.3.2 The first trip that matters

**Planner's question.** Which first trip matters? The question asks which chiller's trip carries the chain, and why.

**General form.** The threshold form of equation (2.12),

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

asked once for each first event.

**Specialized form.** The question is equation (2.12) with the event $E_i$ for each $i$ separately, equation (21.13), and with the motor's link $\{\Delta T^{*}_{s,i} > \Delta T_m\}$ beside it; and the probability that C2 is the strongest unit, the same form with the event $\mathrm{Cap}_{C2} = \max_j \mathrm{Cap}_j$, over the site's draws.

**Building blocks.** The rows of Table 21.3 at 759 kW; the capacities of Table 21.2.

**C2's trip.** C2's trip is the one that matters. After it a second chiller's motor trips with probability 0.1145, against 0.0030 after C1's and below $10^{-4}$ after C3's; any second excursion follows with probability 0.9050, against 0.6212 and 0.4500. The reason is in the fit: C2 is the strongest unit with probability 0.9924, so its trip leaves the smallest remainder, the two weakest units, and by equation (21.4) the float is largest where the remainder is smallest. The motor's link is concentrated almost wholly in C2's trip: of the averaged 0.0392, a third of 0.1145 is C2's share and a third of 0.0030 the rest. The halls' links are shared more evenly, because hall B's margin at 759 kW is small enough that even the smaller float after C3's trip, a median of 1.38 K, crosses it in four draws of ten.

### 21.3.3 The plan year

**Planner's question.** How often a year? The question asks for the expected number of second excursions in the plan year, as the load grows and the wet bulb moves through the months.

**General form.** An estimation question asks for the moments of a derived quantity, equation (2.10),

$$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}] = \int f(\mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z},$$

here of a count summed over a chain of twelve monthly nodes.

**Specialized form.** The question is equation (2.10) with $f = N_2$ of equation (21.10), the unknown including the trip rate, over the factored posterior (21.12): in each paired draw the twelve nodes are solved at the month's wet bulb and the load (21.9), the averaged indicator $\bar{P}_m$ is multiplied by the draw's expected first trips $3\lambda h$, and the twelve products are summed. The mean over the draws is the answer, and the 5 and 95 % points of the per-draw sums its band. It is an expected count, not a rate of failure of any chiller or hall.

**Building blocks.** The month's wet bulbs of the record year (Table 16.5); the load path (21.9); the trip rate's draws; $h = 730$ h and the stated 200 h a year of design hours.

**By month.** Table 21.5 gives the twelve nodes. Each row is a month; the columns give its mean wet bulb, the site IT load at its end, the probability of a second chiller and of any second excursion given a first trip, and the expected number of second excursions in the month. The same probabilities and counts are drawn in Figure 21.3 on a logarithmic scale, the month on the horizontal axis.

The probability rises by more than two orders of magnitude through the year, from 0.0005 in January to 0.1292 in December. The load drives it. The wet bulb is highest in January and February, which lowers the chillers' capacity through the derate, but the load is then lowest; from July the wet bulb rises again while the load grows, and the year's last month combines the largest load with a warm wet bulb. No month's mean wet bulb is near the design day's 22°C, and in no month does any draw float the supply past the motors' threshold. Each month has

Table 21.5: The chain through the plan year, one scenario node per month at the month's mean wet bulb and the load (21.9): the probability of a second excursion given a first trip, averaged over the first trip, and the expected number in the month, $3\lambda h\,\bar{P}_m$ with $h = 730\,\text{h}$. No draw of the 20 000 shows a second chiller trip in any month.

| Month | $T_{\text{wb}}$ [°C] | IT load [kW] | Pr(second chiller) [–] | Pr(any) [–] | Expected in month [–] |
|---|---|---|---|---|---|
| January | 19.83 | 667.7 | $< 10^{-4}$ | 0.0005 | $1.32 \times 10^{-4}$ |
| February | 19.83 | 675.6 | $< 10^{-4}$ | 0.0005 | $1.32 \times 10^{-4}$ |
| March | 18.54 | 683.5 | $< 10^{-4}$ | 0.0008 | $2.20 \times 10^{-4}$ |
| April | 16.29 | 691.5 | $< 10^{-4}$ | 0.0010 | $2.97 \times 10^{-4}$ |
| May | 13.71 | 699.6 | $< 10^{-4}$ | 0.0013 | $3.99 \times 10^{-4}$ |
| June | 11.46 | 707.8 | $< 10^{-4}$ | 0.0032 | $9.13 \times 10^{-4}$ |
| July | 10.17 | 716.1 | $< 10^{-4}$ | 0.0060 | $1.66 \times 10^{-3}$ |
| August | 10.17 | 724.5 | $< 10^{-4}$ | 0.0105 | $2.89 \times 10^{-3}$ |
| September | 11.46 | 732.9 | $< 10^{-4}$ | 0.0155 | $4.38 \times 10^{-3}$ |
| October | 13.71 | 741.5 | $< 10^{-4}$ | 0.0260 | $7.18 \times 10^{-3}$ |
| November | 16.29 | 750.2 | $< 10^{-4}$ | 0.0381 | $1.06 \times 10^{-2}$ |
| December | 18.54 | 759.0 | $< 10^{-4}$ | 0.1292 | $3.61 \times 10^{-2}$ |

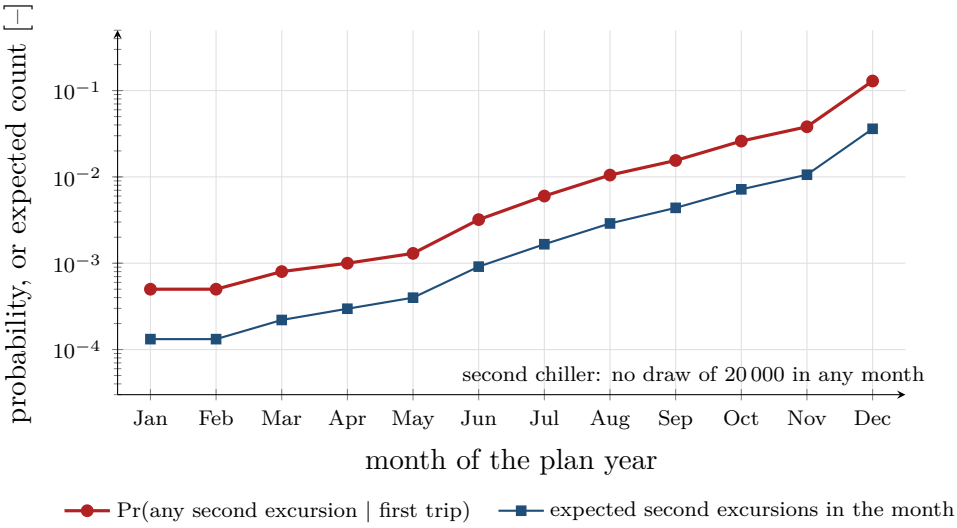



Figure 21.3: The chain through the plan year: the probability of any second excursion given a first trip (red) and the expected number of second excursions in the month (blue), from Table 21.5, on a logarithmic scale. The two curves are a factor of the month's expected first trips apart. A second chiller trip appears in no month.

$3 \times 1.124/12 = 0.281$ first trips expected, so the expected count in a month is about a quarter of its probability, as the two curves of Figure 21.3 show.

**The answer.** Summed over the plan year, the expected number of second excursions is

$$\mathbb{E}[N_2 \mid \mathbf{y}] = 0.065 \text{ in the plan year}, \qquad 90\,\% \text{ band from 0 to } 0.231, \tag{21.15}$$

and the expected number of second chiller trips is zero to the draws' resolution: no draw shows one in any month. The lower end of the band is zero because at least one draw in twenty has a margin wide enough in both halls that no month's float crosses it. December alone carries 0.0361, more than half the year's total. At the monthly mean wet bulb the chain is an event of the year's end.

The months' means leave out the hot days. The design hours, the stated 200 h a year at the design day, bound them: at today's load they add no second excursion in any draw, and at year 1's 759 kW they add an expected 0.0506 second excursions, of which 0.0030 are a second chiller trip, by equation (21.10) with the design day's 0.6587 and 0.0392. Reported as an expected count, a first trip followed by a second excursion happens about 0.07 times in the plan year at mean weather, about 0.12 times, $0.065 + 0.051$, with the design hours added at year 1's load as a bound, and the count grows with the load. It is not a failure rate of any chiller: it counts first trips, at the outage record's rate, that the plant fails to tolerate.

### 21.3.4 The hall-unit chain

**Planner's question.** Does the hall side cascade? The question asks whether a hall unit's trip pushes the running units past the sheet's duty, and whether it pushes an inlet past the limit.

**General form.** The threshold form of equation (2.12),

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

along the hall-unit chain.

**Specialized form.** The question is equation (2.12) over each hall's state, with two events: $f = \max_u \max_{x \in R_u} Q_x^*$ of equation (21.8) against the sheet's duty, and $f = \max_u \max_k T_k^*$ of equation (15.4) against 27°C, each the worst over the three units that could trip; both at the set point of 14°C, over the hall's 4 000 draws.

**Building blocks.** Each hall's draws; the duty split (21.8); the inlets with a unit out (15.4); the sheet's stated duties.

**The two halls.** Table 21.6 gives the chain in each hall at today's and year 1's site load. Each row is a hall and a load; the columns give the hall's IT load, the median and band of the largest running unit's duty, the sheet's duty, and the probabilities that the duty exceeds it and that an inlet exceeds 27°C.

At today's load neither hall's running units exceed the sheet's duty in any draw: hall B's carry 159.9 kW at the median against the sheet's 180 kW, hall A's 193.2 kW against 210 kW. The coils, by the sheet's measure, have duty to spare. Yet an inlet crosses 27°C with probability 0.9895 in hall B and 0.7177 in hall A. At year 1's load the duties cross the sheet too, with probability 0.8720 in hall B and in every draw in hall A, and the inlets cross in essentially every draw. On the air side the cascade never needs a second unit: the first trip already breaks the inlets, at today's load, before any unit's duty reaches the sheet's figure. This is the finding of Chapter 20 for hall B with its largest unit out, read here as a chain beside the chillers'.

Table 21.6: The hall-unit chain beside the chillers': a hall unit trips and the two running units take its heat at the set point of 14°C, the worst of the three units out in each draw, over each hall's 4 000 draws. The sheet's duties are stated; "none" marks no draw, and 1.0000 every draw.

| Hall | Site load [kW] | Hall load [kW] | Duty median [kW] | Duty band [kW] | Sheet [kW] | Pr(duty > sheet) [–] | Pr(inlet > 27°C) [–] |
|---|---|---|---|---|---|---|---|
| A | 660 | 360 | 193.2 | 188.7 to 198.4 | 210 | none | 0.7177 |
| A | 759 | 414 | 221.3 | 216.1 to 227.2 | 210 | 1.0000 | 0.9985 |
| B | 660 | 300 | 159.9 | 156.3 to 164.8 | 180 | none | 0.9895 |
| B | 759 | 345 | 183.1 | 179.0 to 188.7 | 180 | 0.8720 | 1.0000 |

Table 21.7: The chain with Chapter 9's fourth chiller on the design day, the first trip uniform over the four. The first row is the three existing chillers at 759 kW. Shares over the 20 000 draws of the site, where $< 10^{-4}$ marks a share below one in ten thousand, and over the 4 000 paired draws for the halls.

| Drawing | Site load [kW] | Pr(second chiller) [–] | Pr(any) [–] | Pr(float) [–] | Pr($m_B < 0$) [–] |
|---|---|---|---|---|---|
| Three chillers | 759 | 0.0392 | 0.6587 | – | – |
| Four chillers | 759 | $< 10^{-4}$ | 0.0503 | – | – |
| Four chillers | 860 | $< 10^{-4}$ | – | $< 10^{-4}$ | 0.4442 |
| Four chillers | 978 | $< 10^{-4}$ | – | $< 10^{-4}$ | 0.9273 |

### 21.3.5 A fourth chiller

**Planner's question.** What does Chapter 9's fourth chiller do to the chain? The question asks for the second excursion with four chillers, one tripped and three running.

**General form.** The threshold form of equation (2.12),

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

at a scenario node of a design alternative.

**Specialized form.** The question is equation (2.12) with the event $E_i$ of equation (21.2) on four chillers, the fourth's capacity $400\,\mathrm{e}^{\kappa_n}$ with $\kappa_n$ of equation (9.7), $\mathrm{Cap}_{\mathrm{rem},i}$ the sum of the three running, and the first trip uniform over the four; and, at loads beyond year 1's, the probability that the supply floats at all and that hall B's margin is negative with nothing tripped, $\Pr(m_B < 0 \mid \mathbf{y})$.

**Building blocks.** The site's draws with the new unit's deviation; the float (21.4); hall B's margin (21.7) at each load.

**Four chillers.** Table 21.7 gives the chain with four chillers. Each row is a site load; the columns give the probability of a second chiller trip and of any second excursion given a first trip, the probability that the supply floats at all, and the probability that hall B is past its limit with nothing tripped. The three-chiller answer of equation (21.14) is the first row, for comparison.

The fourth chiller removes the float. With three units running after any first trip, the supply floats in fewer than one draw in ten thousand even at 860 kW and at the 978 kW of the design load, and a second chiller trip falls below $10^{-4}$ at every load. What remains at 759 kW, a

probability of 0.0503, is not a cascade: with no float, the second excursion of equation (21.2) holds only where a margin is already negative, hall B past its limit with every unit running and nothing tripped. That share is also inside the three-chiller 0.6587. At higher loads it grows into hall B's own limit, 0.4442 at 860 kW, near hall B's binding load with every unit up (Table 15.3), and 0.9273 at the design load. The fourth chiller makes the chillers fault tolerant through the design load; it does nothing for hall B's air side, which Chapter 5 and Chapter 20 read.

## 21.4 The checks (step 5)

The answers rest on a closed form for the float, on draws of the Laplace posteriors, and on a record generated from a stated truth; each is tested here.

### 21.4.1 The closed form and the routes

The float of equation (21.4) was set against a root solve of the row (21.3) by bracketing, on the first draw of the site with C2 tripped at 759 kW: the two differ by less than $10^{-4}$ K. The motor's link by the exact posterior, the site's draws under their weights, is 0.0414 against the draws' 0.0392, and the joint answer under the same weights on the paired draws is 0.6642 against 0.6587, the gaps of Table 21.4. The paired draws of the halls are read on their Laplace normals without reweighting: the halls' own fits lose little to their exact posteriors (Sections 5.4.2 and 15.2.2), so the exact joint corrects the site's factor only. Three consistency checks need no draws. At 700 kW the loop's 722 kW sits within 0.5 kW of the sum of C1's and C3's medians, and the supply floats after C2's trip in 0.5147 of the draws, as a remainder centered on the loop's heat requires. At 759 kW the threshold $781/1.10 = 710.0$ kW lies a little over one standard deviation below the remainder's center, and 0.1145 of the draws fall below it. And with four chillers the probability of any second excursion, 0.0503, matches a margin already negative: hall B's 5 % point at 759 kW is 0.06 K, so about one draw in twenty sits at or below zero.

### 21.4.2 The answers against the stated truth

The operating record is generated from a stated truth, the last column of Table 9.2, and the last check compares the chain with it. Table 21.8 collects the checks: each row is a quantity, the columns its reported value, the route or reference beside it, the truth, and the reading. The true chillers are weaker than the fit's medians, all three log-ratios a little lower than their MAP values, so the true floats lie above the posterior medians: 2.43, 3.23 and 2.04 K after C1's, C2's and C3's trips at 759 kW, each between its median and its 95 % point in Table 21.3. No true float reaches the motors' 3.33 K: a second chiller trips after none of the three first trips, and C2's 3.23 K sits just under the envelope. At 700 kW the true float after C2's trip is 0.47 K, inside its band of 0 to 0.77 K, and at 660 kW the true supply holds after every trip. The true halls' inlet margins at 759 kW are 1.56 K in hall A and 1.13 K in hall B, each inside its 90 % band of Table 21.2 and below its median; hall B's is positive, so the true hall B carries its 345 kW at year 1's load with every unit up, as its true binding load requires (Table 5.7). Every true float at 759 kW, the smallest 2.04 K after C3's trip, exceeds both true margins, so in the true plant every first trip breaks the inlets of both halls. At 700 kW the true margins are 2.43 and 2.03 K and at 660 kW 3.01 and 2.64 K, in hall A and hall B, and no true float crosses them.

The chain holds against the truth: the floats and the margins fall inside their bands, and the motors' link, which the draws put at about one first trip in twenty-five, does not fire in the true plant, whose worst float stops 0.10 K short of it. The inlet link does fire, after every first trip and in both halls. The truth lies on the unfavorable side of both halves of the chain at once, chillers weaker than the fit's medians and margins thinner, so its second excursion follows

Table 21.8: The checks of the chain. Reported values by the paired draws; the exact route by the site's draws under their weights; the truth is the stated truth of the record, on the design day.

| Quantity | Reported [K or –] | Beside [K or –] | Truth [K or –] | Reading |
|---|---|---|---|---|
| Float after C1, 759 kW | 1.81 | 95 %: 2.69 | 2.43 | inside, above the median |
| Float after C2, 759 kW | 2.74 | 95 %: 3.56 | 3.23 | inside, under 3.33 |
| Float after C3, 759 kW | 1.38 | 95 %: 2.14 | 2.04 | inside, above the median |
| Float after C2, 700 kW | 0.02 | 95 %: 0.77 | 0.47 | inside |
| Float after any trip, 660 kW | 0.00 | 95 %: 0.00 | 0.00 | supply holds |
| Hall A's margin, 759 kW | 2.99 | 5 %: 1.21 | 1.56 | inside, below the median |
| Hall B's margin, 759 kW | 1.63 | 5 %: 0.06 | 1.13 | inside, below the median |
| Second chiller, 759 kW | 0.0392 | exact 0.0414 | after none | gap +0.0022 |
| Any second excursion, 759 kW | 0.6587 | exact 0.6642 | after each | both halls' inlets, after every first trip |
| Any second excursion, 700 kW | 0.0026 | exact 0.0027 | after none | true floats below both margins |
| Any second excursion, 660 kW | none | – | after none | supply holds |
| Float, closed form against root solve | – | – | – | below $10^{-4}$ K |
| Remainder against $Q_{\mathrm{ch}}/1.10$, after C2 | 721.5 kW | 710.0 kW | – | one sd of 9.79 kW; 0.1145 below |
| Draws | site: 20 000, seed 9, ESS 19 913; halls: 4 000 each, seed 21, paired with the first 4 000 of the site | | | |

every trip, where the posterior gives about two trips in three. The inlets need far less float to cross than the motors.

## 21.5 The report (step 6)

Today a single chiller trip does not cascade: at 660 kW the two running chillers carry the loop at the set point after any trip. At next year's 759 kW a trip of C2, the strongest unit, on a design day floats the supply by about 2.7 K, which breaks hall B's inlets with probability 0.87 and a second chiller's motor with probability 0.11. The cascade's second link is thermal, at the racks. Chapter 9's fourth chiller removes the float; what remains is hall B's own capacity, the question of Chapter 5 and of Chapter 20.

### 21.5.1 The answers and their bands

Table 21.9 gives the answers in the order of Table 21.1. Each row is a question; the second column gives the reported answer, and the third its routes or the values beside it.

### 21.5.2 The design-review entry

The design review's entry does not accept the sheet's "single trip tolerated" for the plan year. The sheet's ratio of 1.024 at year 1 counts nameplate; the as-built chillers left after C2's trip carry less than the loop's heat, the supply floats, and the float reaches the racks before it reaches any motor. The entry records the site as single-fault tolerant on its chillers at today's load and not at year 1's, with the second excursion at hall B's inlets with probability 0.87 after C2's trip on a design day, and an expected 0.065 second excursions in the plan year at mean weather, about 0.12 with the design hours, rising with the load. It records the second chiller trip, the event the sheet checks, at 0.0392 on a design day and in no draw at the months' mean weather.

Table 21.9: The report on a single trip on the site, on the design day unless stated, with the set point at 14°C. Reported values by the paired draws, joint over the envelopes.

| Question | Answer | Routes or values beside |
|---|---|---|
| If one unit trips, does the load shift onto the others push a second one past its envelope? | not today; at year 1's 759 kW, a second excursion follows 0.6587 of first trips, a second chiller 0.0392 | exact 0.6642, second chiller exact 0.0414; at 700 kW, 0.0026; the excursion is at the racks: hall B alone 0.8700 after C2 |
| Which first trip matters? | C2's: second chiller 0.1145, any 0.9050 | C1: 0.0030, 0.6212; C3: below $10^{-4}$, 0.4500; C2 strongest with probability 0.9924 |
| How often a year? | 0.065 second excursions expected in the plan year; no second chiller trip in any draw | band 0 to 0.231; December 0.0361; design hours at 759 kW add 0.0506, of them 0.0030 a second chiller |
| Does the hall side cascade? | yes, at the inlets, today: hall B 0.9895, hall A 0.7177, duty within the sheet in every draw | at 759 kW duty above the sheet: hall B 0.8720, hall A every draw |
| What does the fourth chiller do to the chain? | removes the float: second chiller below $10^{-4}$; any second excursion 0.0503 at 759 kW | that 0.0503 is hall B past its limit with nothing tripped; 0.4442 at 860 kW, 0.9273 at 978 kW |

The trip check of every later drawing is required to set the float against the halls' inlet margins as well as against the chillers' motors. The fourth chiller of Chapter 9 is recorded as removing the float through the design load, and hall B's air side as the limit it leaves, for the air-side review of Chapter 20.

**Common mistake**

Checking a trip against the envelope of the units that take the load. The running chillers absorb the trip by letting the supply warm, and the warmer supply goes to every rack on the site. The envelope that a cascade crosses first is the one with the least margin to the shifted quantity, and here that is a rack inlet, which no chiller ratio sees.

## 21.6 The next measurement (step 7)

The chain has two halves, the float and the margin it meets, and each is read from a different fit. Step 7 asks which test or survey of the existing plant narrows each, by the preposterior covariance of equation (5.26), the ranking form of Chapter 18, equation (18.3).

### 21.6.1 The float's half

The float after C2's trip is a function of the remaining capacity $\mathrm{Cap}_{C1} + \mathrm{Cap}_{C3}$ alone, by equation (21.4), whose gradient is the pattern of equation (21.11); on the design day the remaining capacity's own gradient is

$$\mathbf{g} = \nabla_{\mathbf{z}} \mathrm{Cap}_{\mathrm{rem},2} = (\mathrm{Cap}_{C1},\ 0,\ \mathrm{Cap}_{C3},\ 0,\ 0,\ 0,\ 0,\ 0)^{\mathsf{T}}.$$

A capacity test of chiller $j$ at the design wet bulb is one row, $\mathbf{J}_a = \nabla_{\mathbf{z}} \mathrm{Cap}_j^{\mathsf{T}} / 8$ with the stated standard deviation of 8 kW; at 22°C, the rating condition, the row has no entry in the derate's column and reads the unit's log-ratio alone. Its preposterior covariance is

$$\boldsymbol{\Sigma}_a = \left(\boldsymbol{\Sigma}^{-1} + \mathbf{J}_a^{\mathsf{T}} \mathbf{J}_a\right)^{-1}, \qquad \mathrm{sd}_a(\mathrm{Cap}_{\mathrm{rem},2}) = \sqrt{\mathbf{g}^{\mathsf{T}} \boldsymbol{\Sigma}_a \, \mathbf{g}}.$$

Table 21.10: Candidate measurements for the chain, by the preposterior covariance (5.26): capacity tests at the design wet bulb of 22°C, standard deviation 8 kW each, read on the capacity left when C2 trips. Shares removed from the printed standard deviations.

| Candidate | sd now [kW] | sd after [kW] | Removed [%] |
|---|---|---|---|
| Tests of C1 and C3 | 9.79 | 7.39 | 24.6 |
| Test of C2 | 9.79 | 9.46 | 3.4 |

Table 21.10 ranks the candidates. Each row is a measurement; the columns give the standard deviation of the remaining capacity after C2's trip before and after it, and the share of it removed.

Testing the two units that would be left removes 24.6 % of the remaining capacity's spread, from 9.79 to 7.39 kW. Testing C2 itself, the unit whose trip matters, removes $1 - 9.46/9.79$, about 3.4 %: its row has no entry in the gradient, and it narrows the remainder only through the positive covariance the shared derate leaves among the three log-ratios, equation (16.8). The test that matters for a trip is a test of what is left after it.

### 21.6.2 The margin's half

The float's half is the narrower one. The other half of the chain is hall B's inlet margin at 759 kW, a median of 1.63 K with a 90 % band of 0.06 to 2.98 K, whose ends lie further from its median than the float's 95 % point lies from the float's median (Section 21.3.1). The margin's band is that of hall B's binding rack, rack B3, whose recirculation fraction no reading sees; Chapter 5 ranked the measurements that narrow it (Section 5.6.2), and the inlet sensor at rack B3 that Chapter 20 reads is the same measurement here. Its preposterior on the chain is not computed in this chapter.

### 21.6.3 The recommendation

Commission capacity tests of C1 and C3 at the design wet bulb: they narrow the float that a C2 trip causes, and with it the motor's link. Read hall B's margin by the rack B3 sensor that Chapter 20 recommends: it narrows the half of the chain that decides. Neither measurement changes the finding that the second link is at the racks; both narrow when, in the load's growth, the chain becomes likely, and the design review's trip check reads both.

### 21.6.4 What the analysis showed

The design sheet checks a trip against the chillers' ratio, and on its nameplate arithmetic the site tolerates a single trip through the plan year. On the site's posterior it does today, and does not at year 1's load. The as-built chillers are a little short of their nameplate, so that after C2's trip the two left carry less than the loop's heat and the supply floats, by a median of 2.74 K on a design day. The float would have to exceed 3.33 K to push a second chiller past its motor envelope, which it does in about one draw in nine; but every rack inlet on the site rises with the supply, and hall B's margin at that load is 1.63 K, which the float crosses in 0.87 of the draws. The cascade's second link is thermal, at the racks, not a second chiller. Averaged over the first trip, a second excursion follows two first trips in three on a design day at year 1's load; through the plan year, at the months' mean weather, the expected count is 0.065, mostly in December, and the design hours raise it to about 0.12. The hall units' own trips break the inlets already at today's load, with every coil inside the sheet's duty. A fourth chiller removes the float and

leaves hall B's own limit, which the chain cannot fix. The chain found the envelope the sheet does not check, and found it first.

### 21.6.5 Limits of the model

**The motor envelope is stated.** The envelope $e_m = 1.10$ of each chiller's as-built duty is a stated input, not inferred: no record of motor current or of a chiller's duty at a warm supply is kept, and the motor's threshold of 3.33 K inherits it whole. A different envelope moves the threshold to $(e_m - 1)/\gamma$ and the motor's link with it; the halls' link does not depend on it, and the finding that the racks cross first holds for any envelope whose threshold lies above hall B's margin.

**The capacity–supply slope is pinned.** The 3 % per kelvin by which a chiller's capacity grows with the supply is pinned, not fitted, for the same reason. A smaller slope makes the float larger for the same shortfall, by equation (21.4), and lowers the motor's threshold; both make the chain more likely.

**Steady state, one link.** The float is the steady state after the trip. The minutes in which the loop and the halls store heat, and a tripped chiller's restart, are not modeled, so the chain says whether a second excursion follows, not when or for how long. Its duration and the energy not served are the question of Chapter 23. The chain stops at the second excursion: a second chiller trip leaves one chiller, and what follows is not followed.

**Two chains, not combined.** The chillers' chain runs with every hall unit up, and the hall-unit chain at the set point with every chiller up. A chiller trip during a hall unit's outage combines a float with a smaller margin and is not computed. The pipe, the pump and the feed are held healthy.

**Mean months and a bound.** Each month runs at its mean wet bulb, and the design hours at the design day bound the hot days; the truth of the year lies between the two counts. The load grows in today's proportions in the two halls and month by month at 15 % a year.

**The halls' draws are not reweighted.** The exact route reweights the site's factor of the posterior, for the motor's link and for the joint answer; the halls' factors, which carry most of the answer's spread, are read on their Laplace normals.

**The constructed record.** The operating record, the outage record and the halls' night record are generated from a stated truth with stated errors, not taken from a site, and the comparison in Section 21.4.2 is a check of the method on its own model, not of the model against a real plant.

**The next chapter.** The chain asked what an unplanned trip does. A planned outage, a unit taken down for maintenance while the rest must carry the load, is the same scenario node held for days instead of minutes; Chapter 22 reads the site's concurrent maintainability.

## 21.7 What to retain

**Definition 21.1** (Cascade probability and its expected count)**.** For a first event $F$ that pins a scenario node and a second event $S = \{x^*(\mathbf{z}) > \vartheta(\mathbf{z})\}$ on the load that node shifts, with $\vartheta$ the smallest of every envelope the shifted load meets, the cascade probability is $\Pr(S \mid F, \mathbf{y}) =$

$\int \mathbf{1}[x^*(\mathbf{z}) > \vartheta(\mathbf{z})]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$, joint over the envelopes. Its expected count over a period $T$ is $\mathbb{E}[\nu_F\, T\, \Pr(S \mid F, \mathbf{z}) \mid \mathbf{y}]$ with $\nu_F$ the rate of first events: an expected count of cascades, never a failure rate of the second element.

Four conclusions carry beyond the site. A cascade is a chain of scenario nodes, each pinned by the solved output of the one before, and its probability is one integral over the posterior of every state the chain touches. Its second link is the smallest of several envelopes on the same shifted quantity, so the reported probability is joint over them, and the worst marginal stands beside it with the gap. The envelope that breaks first need not belong to the units that take the load: a chiller trip shifts heat to the running chillers and temperature to every rack, and the racks have the least margin. And the measurement that narrows a trip is a test of what the trip leaves, not of the unit that trips.

Table 21.11 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 21.11: The terms of a cascade and where each is defined. Units for the site.

| Term | Expression | Units | Meaning | Defined in |
|---|---|---|---|---|
| Cascade probability | $\Pr(S \mid F, \mathbf{y})$, joint over the envelopes | – | second event given the first | 21.1.1 |
| Expected count | $\mathbb{E}[\nu_F\, T\, \Pr(S \mid F, \mathbf{z})]$ | – | cascades in a period | 21.1.1 |
| *The site's chain* | | | | |
| Second excursion | $\Delta T^*_{s,i} > \min(\Delta T_m, m_A, m_B)$ | – | any envelope crossed | 21.1.2 |
| Supply float | $\max(Q_{\mathrm{ch}}/\mathrm{Cap}_{\mathrm{rem},i} - 1,\ 0)/\gamma$ | K | supply rise after a trip | 21.1.2 |
| Float's gradient | $-Q_{\mathrm{ch}} \nabla \mathrm{Cap}_{\mathrm{rem},i} / (\gamma\, \mathrm{Cap}^2_{\mathrm{rem},i})$ | K | implicit function theorem | 21.1.2 |
| Motor's threshold | $\Delta T_m = (e_m - 1)/\gamma = 3.33$ | K | second chiller trips | 21.1.2 |
| Inlet margin | $T_{\max} - T^*_s - \max_k a_k (Q^*_h + Q_{o,h})$ | K | supply rise a hall absorbs | 21.1.2 |
| Hall-unit duty | $G_x (Q^*_h + Q_{o,h}) / \sum_{y \in R} G_y$ | kW | running unit after a unit trip | 21.1.2 |
| Plan-year count | $\sum_m 3 \lambda h\, \bar{P}_m$ | – | second excursions in the year | 21.1.3 |
| *Questions and answers* | | | | |
| Drill answer | 0.6587 any, 0.0392 a second chiller | – | at 759 kW, design day | 21.3.1 |
| Plan year | 0.065, band 0 to 0.231 | – | mean weather | 21.3.3 |

## 21.8 Exercises

**Exercise 21.1.** *Formulation.* Write the planner's question *if one chiller trips at year 1's load on a design day, does the load shift push a second element past its envelope?* as a probability over the site's and the halls' states, say which entries it depends on, and give its value and the share that is a second chiller. (It is $\mathbb{E}[\bar{P} \mid \mathbf{y}]$ with $\bar{P} = \frac{1}{3} \sum_i \mathbf{1}[\Delta T^*_{s,i} > \min(\Delta T_m, m_A, m_B)]$, equation (21.13) averaged over the first trip. On the design day the floats depend on the three log-ratios only, each trip on the other two, and the margins on each hall's log-derates and fractions. At 759 kW the value is 0.6587, and a second chiller 0.0392.)

**Exercise 21.2.** *Derivation.* Derive the float after chiller $i$ trips from the running chillers' balance, and show that a running chiller passes its motor envelope when the float exceeds $(e_m - 1)/\gamma$, whatever its own capacity. Give the threshold on the remaining capacity after C2's trip at 759 kW and the probability of crossing it. (The row $\mathrm{Cap}_{\mathrm{rem},i}(1 + \gamma\, \Delta T_s) - Q_{\mathrm{ch}} = 0$ gives $\Delta T_s = (Q_{\mathrm{ch}}/\mathrm{Cap}_{\mathrm{rem},i} - 1)/\gamma$ where positive. Running unit $j$ carries $(1 + \gamma \Delta T_s)\mathrm{Cap}_j$ against $e_m \mathrm{Cap}_j$, and $\mathrm{Cap}_j$ cancels, so the threshold is $0.10/0.03 = 3.33$ K, or $\mathrm{Cap}_{\mathrm{rem}} < Q_{\mathrm{ch}}/e_m = 781/1.10 = 710.0$ kW. After C2's trip the probability is 0.1145.)

**Exercise 21.3.** *Derivation.* Show from the scenario inlets of hall B that every inlet rises one for one with the supply, and so that the hall's inlet margin at the set point is the supply rise it can absorb. Give hall B's margin at 759 kW. (By equation (5.7), $T_k^* = T_s + a_k(\mathbf{z}_B)(Q_B^* + Q_o)$, and $a_k$ does not depend on the supply, so $\partial T_k^*/\partial T_s = 1$ for every rack. The hottest inlet at $T_s^* = 14$°C is 27°C less the margin, so a float larger than the margin carries it past the limit. Hall B's margin at 759 kW has median 1.63 K and 90 % band 0.06 to 2.98 K.)

**Exercise 21.4.** *Application.* From the monthly probabilities, check December's expected count and give the plan year's, and say how to report it. (December's first trips number $3 \times 1.124/12 = 0.281$ in expectation, and $0.281 \times 0.1292 = 0.0363$, the printed 0.0361 to within the product of means. Summed over the months the expected count is 0.065 second excursions in the plan year, band 0 to 0.231, about 0.12 with the design hours added as a bound; it is reported as an expected count of cascades, never as a failure rate of a chiller. The year's count is 0.065.)

**Exercise 21.5.** *Application.* A reviewer proposes Chapter 9's fourth chiller as the cure for the cascade. Say what it removes and what it leaves, with the numbers at 759 kW and at the design load. (With three units running after any first trip the supply does not float in any draw even at 978 kW, and a second chiller trip falls below $10^{-4}$: the chillers become fault tolerant through the design load. The second excursion at 759 kW falls from 0.6587 to 0.0503, which is hall B already past its limit with nothing tripped; at the 978 kW design load hall B is past its limit with nothing tripped with probability 0.9273.)

# Chapter 22

# Concurrent Maintainability

## 22.1 A maintenance calendar for the site

The site's maintenance planner has a calendar to fill. Each of the three chillers and each of the six hall cooling units is due its annual service, a planned outage of 72 h, and the planner wants a month for each. The design sheet answers the question one unit at a time, with nameplates. With one chiller out, two of 400 kW leave 800 kW of nameplate against a loop heat that grows from 690 kW in January to 781 kW in December of the plan year. With one hall unit out, hall B keeps two units of 180 kW of duty, 360 kW against a hall load of 314 to 355 kW, and hall A two of 210 kW, 420 kW against 376 to 426 kW. The sheet passes the chillers in every month and hall B in every month, but hall A only from January to October: in November and December the hall's load exceeds its 420 kW, by 6 kW in December at 426 kW. The sheet does not list the pump or the strainer in the shared supply pipe: there is one of each, and the drawing gives neither a second path.

The planner's question is the drill question, verbatim,

*Can any one unit be taken out for service in any month of the year without a violation?*

It is answered with the method of Chapter 4 on the site's one fit of Chapter 9 and the two halls' fits of Chapters 5 and 15. What is new is the shape of the question. It names no single scenario but every pair of a unit and a month, $11 \times 12 = 132$ of them, and its answer is not a probability but the smallest of 132 probabilities, the pair at which it falls, and the months in which each unit may come out.

**Jargon: concurrent maintainability, fault tolerance, planned outage, service window, scenario grid, staying inside, capability probability under a planned outage, assurance level, limiting element, flow path, standby pump, strainer bypass, second failure during service**

A plant is *concurrently maintainable* when any one of its units can be taken out for planned service, at any time of the year, without loss of the load; it is *fault tolerant* when the load also survives an unplanned failure, here a second unit failing while the first is out. In the tier classification used for data centers the first is the mark of Tier III and the second of Tier IV. A *planned outage* is a unit taken out on purpose; its *service window* is the time it is out, here 72 h, and a unit's service windows in the year are the months in which it may come out. The *scenario grid* is a family of scenario nodes, one per pair of a unit out and a month, each with that month's load and weather. *Staying inside* at a node means that every element of the plant carries the load there; its probability is a *capability probability under a planned outage*, the probability that the plant as it is carries a stated load with a stated unit out, not the fraction of time the plant is up. The *assurance level* is the probability a requirement must reach to be accepted, here 0.95. The *limiting element* of a node is the element whose binding load is the smallest. A *flow path* is a route the chilled water can take from the plant to the halls; a *standby pump* is a second pump on its own path, and a *strainer bypass* a valved pipe around the strainer, either of which keeps a path open while the other is serviced. A *second failure during service* is an unplanned failure of another unit inside the first unit's service window.

**Key idea**

A maintenance calendar is a capacity result. Each unit out in each month is a scenario node: the unit's terms leave the plant's relations, the month's load and weather are pinned, and every element's binding load is solved again. The probability of staying inside is read at each node jointly over the elements, from one posterior; the plant is concurrently maintainable when the smallest of those probabilities clears the assurance level, and each unit's service windows are the months in which its own row clears it. The question of a second failure during the service is a different question, fault tolerance, and is asked of the same grid with the failure rates attached.

### 22.1.1 A grid of scenario nodes

The quantity the question owes is the smallest probability of staying inside over every unit and every month, and where it falls. Write $\mathcal{U}$ for the units that can be taken out and $m = 1, \dots, 12$ for the months of the plan year. Each pair $(u, m)$ is a scenario node of Section 4.3.8 with unit $u$ out, every other element up, and the month's site IT load $Q^*_m$ and wet bulb $T^*_m$ pinned. The answer is

$$P_{\min} = \min_{u \in \mathcal{U},\ m} P_{um}, \qquad (u^\dagger, m^\dagger) = \arg\min_{u,\ m} P_{um},$$
$$\text{concurrently maintainable if } P_{\min} \geq \kappa, \tag{22.1}$$

with $\kappa = 0.95$ the assurance level. The pair $(u^\dagger, m^\dagger)$ names the unit and the month that fail the plant first. Unit $u$'s service windows are the months of its row that clear the level,

$$W_u = \{\, m : P_{um} \geq \kappa \,\}, \tag{22.2}$$

so that the plant is concurrently maintainable exactly when every $W_u$ holds at least one month and the calendar can give each unit one of its own; the drill question asks the stronger thing, every month for every unit, which is $P_{\min} \geq \kappa$.

Each $P_{um}$ is a probability over the posterior. At node $(u, m)$ each element $e$ of the plant has a binding load $Q_e^{\dagger}(\boldsymbol{\theta};\ u, m)$, the site IT load at which its limit binds with unit $u$ out at the wet bulb $T_m^{*}$, a function of the unknown $\boldsymbol{\theta}$ of the plant. The plant stays inside when every element carries the month's load, that is when the smallest of the binding loads is at least $Q_m^{*}$, and the probability of that event is its posterior mass,

$$\begin{aligned} P_{um} &= \Pr\Big(\min_e Q_e^{\dagger}(\boldsymbol{\theta};\ u, m) \geq Q_m^{*} \;\Big|\; \mathbf{y}\Big) \\ &= \int \mathbf{1}\Big[\min_e Q_e^{\dagger}(\boldsymbol{\theta};\ u, m) \geq Q_m^{*}\Big]\, p(\boldsymbol{\theta} \mid \mathbf{y})\, \mathrm{d}\boldsymbol{\theta}, \end{aligned} \tag{22.3}$$

joint over the elements: one indicator of an intersection, not a product of the elements' own probabilities. The limiting element of the node is the element that attains the minimum most often over the posterior. The integral is read by draws $\boldsymbol{\theta}_i$ from the posterior, each pushed through every element's binding load,

$$\hat{P}_{um} = \frac{1}{N}\sum_{i=1}^{N} \mathbf{1}\Big[\min_e Q_e^{\dagger}(\boldsymbol{\theta}_i;\ u, m) \geq Q_m^{*}\Big], \qquad \mathrm{se}(\hat{P}_{um}) = \sqrt{\hat{P}_{um}\,(1 - \hat{P}_{um})/N}, \tag{22.4}$$

with se the Monte Carlo standard error. A node at which no draw stays inside has $\hat{P}_{um} = 0$, which says only that the probability lies below about $1/N$; a node at which the model leaves no flow path at all has $P_{um} = 0$ exactly, with no draw needed.

The minimum in equation (22.1) is a minimum of probabilities, not the probability that every node stays inside. The nodes are not states that occur together: each unit is out in one month of the planner's choosing, and the question is whether some month serves it. Each node is therefore its own event, read on its own, and the calendar is the table of all of them.

Fault tolerance asks a second question of the same nodes. While unit $u$ is out in month $m$, another unit $v$ may fail. With $\lambda_v$ its failure rate and $\tau = 72\,\mathrm{h}$ the service window, the chance that $v$ fails inside the window is $1 - \mathrm{e}^{-\lambda_v \tau}$; the plant with both out is the node $(u, v, m)$, solved like the others. To first order in the rates, so that two further failures in one window are neglected, the probability of a second failure that leaves the plant outside is

$$P_{um}^{(2)} = \sum_{v \neq u} \mathbb{E}\Big[(1 - \mathrm{e}^{-\lambda_v \tau})\, \mathbf{1}\Big[\min_e Q_e^{\dagger}(\boldsymbol{\theta};\ u, v, m) < Q_m^{*}\Big] \;\Big|\; \mathbf{y}\Big], \tag{22.5}$$

the expectation over the posterior of the plant and of the rates together. Its first factor alone, summed, is the probability that a second unit fails at all, $\sum_{v \neq u} \mathbb{E}[1 - \mathrm{e}^{-\lambda_v \tau} \mid \mathbf{y}]$. The indicator counts the draws outside with both units out, including any already outside with $u$ alone.

**Principle 22.1.** *Concurrent maintainability is a minimum over a grid of scenario nodes, one per unit out and month. Each node's probability of staying inside is joint over the plant's elements and read from one posterior; the smallest of them, where it falls, and each unit's months above the assurance level are the answer. A second failure during the service is fault tolerance, read on the same nodes with the failure rates attached.*

### 22.1.2 The site with its units out of service (step 0)

The plant is the site of Chapter 9, drawn in Figure 9.1, with both halls as Chapter 15 draws them in Figure 15.1: the three chillers C1–C3, one pump, the shared supply pipe with its strainer, the power feed, and the two halls of three cooling units and five racks each, taken as built, without the drawn economizer. The unknown is that of equation (15.8): each hall's reduced state $\mathbf{z}_A$ and $\mathbf{z}_B$ of equation (5.4), the site's eight entries $\mathbf{z}_{\mathrm{site}}$ of equation (9.4), and the failure rates of the chillers and of the hall units. The units that can be taken out are eleven,

$$\mathcal{U} = \{\mathrm{C1},\ \mathrm{C2},\ \mathrm{C3},\ \mathrm{A1},\ \mathrm{A2},\ \mathrm{A3},\ \mathrm{B1},\ \mathrm{B2},\ \mathrm{B3},\ \mathrm{pump},\ \mathrm{strainer}\}.$$

**The physical relations.** At node $(u, m)$ the six elements of equation (15.1) each have their binding load in site IT load, kW. The chillers up carry the loop's heat, the IT load plus the halls' other heat $Q_o = 22\,\text{kW}$, at the month's wet bulb, where by equation (9.3) each unit's capacity is its design-day capacity times the derate factor; each hall carries its share of the load with its running units, by equations (15.4) and (15.5); the pipe and the pump are those of equations (13.5) and (13.6); and the feed solves its limit at the month's wet bulb, where the chillers' efficiency is that month's. Writing $R_h(u)$ for the units of hall $h$ other than $u$,

$$\begin{aligned}
Q^{\dagger}_{\text{ch}}(\boldsymbol{\theta};\ u, m) &= \Big[1 - d\,(T^*_m - 22)\Big] \sum_{i \neq u} 400\,\mathrm{e}^{\kappa_i} - Q_o, \\
Q^{\dagger}_{h}(\boldsymbol{\theta};\ u, m) &= \frac{1}{\varphi_h}\Big[\min_{k \in h} \frac{T_{\max} - T^*_s}{a_k(\mathbf{z}_h;\ R_h(u))} - Q_{o,h}\Big], \qquad h = A,\ B, \\
Q^{\dagger}_{\text{pipe}}(\boldsymbol{\theta}) &= \rho c_p\,\Delta T_w\,\sqrt{\Delta P_{\max}}\,\mathrm{e}^{l_k} - Q_o, \\
Q^{\dagger}_{\text{pump}}(\boldsymbol{\theta}) &= \rho c_p\,\Delta T_w\,V^{\dagger}(l_k,\ l_e) - Q_o, \\
Q^{\dagger}_{\text{feed}}(\boldsymbol{\theta};\ m): &\quad (1 + a)\,Q + \frac{Q + Q_o}{\text{COP}(T^*_m)} + P_{\text{pump}}(Q) = 1\,520\,\text{kW},
\end{aligned} \tag{22.6}$$

the first with the sum over the chillers other than $u$, the second with the inlet rise $a_k$ of equation (15.4) at the drawing's supply $T^*_s = 14°\text{C}$ and the limit $T_{\max} = 27°\text{C}$, the pipe's and the pump's with the pump and the strainer both in service, and the last solved for the load $Q$. The hall's closed form is the one Chapter 20 writes for the largest unit out; here it is written for every unit in turn. Only the chillers and the feed depend on the month's wet bulb; the halls, the pipe and the pump depend on the month only through the load they are asked to carry. With the pump or the strainer out there is no flow path: no water reaches either hall, and the pipe's and the pump's binding loads are zero,

$$Q^{\dagger}_{\text{pipe}} = Q^{\dagger}_{\text{pump}} = 0, \qquad u \in \{\text{pump},\ \text{strainer}\},$$

so the plant is outside at every load and every state. The pinned quantities are those of equations (9.3) and (15.4); the derate $d$ is per kelvin of wet bulb above the 22°C rating, so in the cooler months the factor exceeds one and the chillers carry more than at the design day.

**Modeling choices at the graph.** The plant is taken in steady state at the month's hottest days: this makes each node one algebraic check, and it fails if the hottest hours of a month are hotter than its hottest days' average. A unit out passes no heat and moves no air, as in Chapter 15: this drops it from every sum it enters, and it fails if a hall unit kept blowing unchilled air. The supply is pinned at 14°C whatever the chillers do: a chiller plant short of the load would warm the supply and lower both halls' binding loads at once, a chain that Chapter 21 follows and this chapter does not. The pump and the strainer are taken as single points of the flow path, as the drawing shows them; a site whose pump had a cross-connection to another loop would not be modeled by the zeros above. And the halls' shares of the site load are today's, $\varphi_A = 0.5455$ and $\varphi_B = 0.4545$: this fails if one hall grows faster.

### 22.1.3 The grid's nodes and the remedies

The months pin two inputs, the load and the weather. The load is the plan year's, the load path of equation (6.1) with today's site IT load of 660 kW and growth of 15 % a year, read at the month's end,

$$Q^*_m = 660\,(1.15)^{m/12}\,\text{kW}, \qquad T^*_m = T_m + 2\,\text{K}, \qquad m = 1, \ldots, 12, \tag{22.7}$$

and the wet bulb is the month's hottest days, the record year's monthly mean $T_m$ of Table 16.5 plus a stated 2 K. The 2 K is chosen so that the hottest month's hottest days fall at the design

Table 22.1: The inputs pinned at each month's nodes: the hottest days' wet bulb, the record year's monthly mean plus 2 K; the plan year's site IT load at 15 % a year from today's 660 kW; and the loop heat, the load plus $Q_o = 22$ kW. The sheet's chillers with one out are 800 kW of nameplate in every month.

| Month | Wet bulb $T_m^*$ [°C] | Site IT load $Q_m^*$ [kW] | Loop heat $Q_m^* + Q_o$ [kW] |
|---|---|---|---|
| January | 21.83 | 668 | 690 |
| February | 21.83 | 676 | 698 |
| March | 20.54 | 683 | 705 |
| April | 18.29 | 691 | 713 |
| May | 15.71 | 700 | 722 |
| June | 13.46 | 708 | 730 |
| July | 12.17 | 716 | 738 |
| August | 12.17 | 724 | 746 |
| September | 13.46 | 733 | 755 |
| October | 15.71 | 742 | 764 |
| November | 18.29 | 750 | 772 |
| December | 20.54 | 759 | 781 |

day: 22°C is the warmest monthly mean plus about 2 K, a check of Section 22.4.2. Table 22.1 gives the twelve nodes' inputs. Each row is a month; the columns give the hottest days' wet bulb, the site IT load and the loop heat the chillers must carry, the load plus the halls' other 22 kW.

The two inputs move against each other. The site lies in the southern hemisphere, and its warmest months are January and February, at 21.83°C; but the load grows through the year, and December, at 20.54°C and 759 kW, carries the most load at the second-warmest wet bulb. Which month binds is not read off either column alone; it is the month at which the elements' binding loads, moved by the wet bulb, fall furthest below the load.

**The remedies.** The earlier chapters name four design alternatives that bear on a unit out of service, and the grid is solved again with all four built: a fourth cooling unit in each hall, of that hall's model, the design alternative Chapter 20 weighs for hall B; the fourth chiller of Chapter 9; a standby pump; and a bypass around the strainer. The fourth hall unit is not yet built, so its log-derate is stated, and the fourth chiller's log-ratio is that of equation (9.7),

$$g_{h4} \sim \mathcal{N}(0,\ 0.15^2), \qquad \kappa_n = \tfrac{1}{3}(\kappa_{C1} + \kappa_{C2} + \kappa_{C3}) + 0.04\,\varepsilon_n, \tag{22.8}$$

with the fourth hall unit at its hall's rated conductance and air capacity rate. The fourth hall unit enters both sums of equation (15.4); the fourth chiller enters the sum of equation (22.6); the standby pump and the bypass give the water a second path, so that with the pump or the strainer out the pipe and the pump keep their binding loads. The remedies are stated design alternatives, not estimates; this chapter reads what they buy in probability and leaves their price to Chapters 23 and 24.

### 22.1.4 The planner's questions (step 1)

The planner asks four questions of the calendar, and step 1 writes each over the plant's unknown, with the evidence on the right of the bar. Table 22.2 lists them. The first column is the question in the planner's words, the drill question first; the second the general form of its category in the notation of Chapter 2; the third its mathematical form for the site; the fourth its category; the fifth the computation of Section 2.3.6 that answers it; and the last the subsection where it

Table 22.2: The questions asked of the site's maintenance calendar. $\mathbf{y}$ is both halls' night record, the site's operating record, the outage records and the priors; $P_{um}$ is the probability of staying inside at node $(u, m)$, equation (22.3), $W_u$ unit $u$'s service windows (22.2), and $P^{(2)}_{um}$ the probability of a second failure that leaves the site outside (22.5).

| Planner's question | General form | Mathematical form | Category | Computation | Answered in |
|---|---|---|---|---|---|
| Can any one unit be taken out for service in any month of the year without a violation? | $\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y})$, the minimum over scenarios | $P_{\min} = \min_{u,m} P_{um}$ and its $\arg\min$ | prediction under a threshold | posterior | 22.3.1 |
| When can each unit come out? | $\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}) \geq \kappa$, over a family | $W_u = \{m : P_{um} \geq 0.95\}$ for each unit | prediction under a threshold | posterior | 22.3.2 |
| What do the named remedies buy? | the same, at a design alternative | $P_{um}$ and $W_u$ with the four remedies built | design | posterior | 22.3.3 |
| What if a second unit fails during the service? | $p(y_{\text{new}} \mid \mathbf{y})$, the predictive | $P^{(2)}_{um}$ with the rates' posterior, $\tau = 72\,\text{h}$ | prediction under a threshold | posterior | 22.3.4 |

is answered. The drill question and the second are predictions under a threshold, the first a minimum over the grid and the second a set of months; the third is the same reading under a design alternative; the fourth is a prediction of a future event, a second failure, which stays random even when the plant is known. The atlas of Appendix A collects them with the other questions of Volume II.

### 22.1.5 The evidence as rows (step 2)

Step 2 attaches the records of Chapter 15, unchanged. Each hall is read in window 1 of the night record, its seventeen rows of equation (5.11): hall B's those of Chapter 5, hall A's those of Section 15.2.2. The site is read by the operating record of Chapter 9, sixty-two rows on its eight entries. The outage records of the chillers and of the hall units give the failure rates, and only the fourth question reads them. No row is added: the grid asks new questions of the same evidence.

**Which row acts on which entry.** By equation (22.6) each element's binding load reads its own block of the unknown. With $*$ for a block that can be nonzero, the six binding loads' gradients have the pattern

$$\begin{bmatrix} \mathbf{g}_{\text{ch}}^{\mathsf{T}} \\ \mathbf{g}_{A}^{\mathsf{T}} \\ \mathbf{g}_{B}^{\mathsf{T}} \\ \mathbf{g}_{\text{pipe}}^{\mathsf{T}} \\ \mathbf{g}_{\text{pump}}^{\mathsf{T}} \\ \mathbf{g}_{\text{feed}}^{\mathsf{T}} \end{bmatrix} = \begin{bmatrix} 0 & 0 & * & * & 0 & 0 & 0 \\ * & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & * & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & * & 0 \\ 0 & 0 & 0 & 0 & 0 & * & * \\ 0 & 0 & 0 & 0 & * & * & * \end{bmatrix}, \tag{22.9}$$

with the columns, in order, hall A's state, hall B's state, the chillers' log-ratios, the derate $d$, the efficiency $c$, the pipe's $l_k$ and the pump's $l_e$. Three facts are read here. Each hall's binding load reads its own state alone, and the two halls' states are read by different records, so hall A's and hall B's limits are independent given the evidence. The chillers' binding load reads the log-ratios at every node, and the derate's column grows with the distance of the month's wet

bulb from 22°C. And no element's row changes when a hall unit is out except that hall's: taking out A2 moves $\mathbf{g}_A$ alone.

**Modeling choices at the rows.** The rows are those of Chapter 15, with its choices: each reading Gaussian, unbiased and independent given the state, each hall fitted on window 1 of its own record, and the hall units sharing one failure rate and the chillers another. One choice matters more here than there. The halls' binding loads with a unit out lean on the fractions of the racks no reading sees, B3 in hall B and A5 in hall A, whose posteriors are their survey priors; every hall row of the grid inherits what those surveys got wrong.

## 22.2 The posterior of the state (step 3)

All four questions of Table 22.2 are read from one posterior, the product of the three fits that Chapters 5, 9 and 15 compute, with the rates' posteriors beside them. No row is added and no fit is repeated; this section derives what the grid needs from that posterior, cites the fits, and states the draws the answers use.

### 22.2.1 What the grid needs from the state

By equation (22.3) each node needs the joint distribution of the six binding loads at that node, which by the pattern (22.9) needs the three blocks of the state and nothing of their cross-covariances: no record reads two blocks at once. The posterior factors as in equation (15.9),

$$p(\boldsymbol{\theta} \mid \mathbf{y}) = p(\mathbf{z}_A \mid \mathbf{y}_A)\, p(\mathbf{z}_B \mid \mathbf{y}_B)\, p(\mathbf{z}_{\text{site}} \mid \mathbf{y}_{\text{op}})\, p(\lambda_c \mid \mathbf{y}_{\text{out}})\, p(\lambda_h \mid \mathbf{y}_{\text{out}}), \tag{22.10}$$

with the repair rates dropped, since a service window is planned and its length stated. Within a node the six binding loads are not independent: the pipe, the pump and the feed share $l_k$, and the feed and the chillers both move with the month. The joint indicator of equation (22.3) is therefore read on draws of the whole unknown, never as a product of the elements' own probabilities.

### 22.2.2 The three fits, cited

Each block's posterior is the Laplace normal of its own fit. For block $b$, one of hall A, hall B and the site, the residual vector $\mathbf{r}_b(\mathbf{z}_b)$ stacks that block's normalized rows, its Jacobian is $\mathbf{J}_b = \partial \mathbf{r}_b / \partial \mathbf{z}_b$, and the MAP state solves the normal equations of equation (5.12) from the design sheet's state, with the Laplace covariance the inverse of the information at the minimum, equation (5.13),

$$\begin{aligned} (\mathbf{J}_b^{\mathsf{T}} \mathbf{J}_b)\, \Delta \mathbf{z}_b &= -\mathbf{J}_b^{\mathsf{T}} \mathbf{r}_b(\mathbf{z}_b), \\ \boldsymbol{\Sigma}_b = (\mathbf{J}_b^{\mathsf{T}} \mathbf{J}_b)^{-1}\big|_{\hat{\mathbf{z}}_b}, &\qquad \mathbf{z}_b \mid \mathbf{y}_b \approx \mathcal{N}(\hat{\mathbf{z}}_b,\ \boldsymbol{\Sigma}_b). \end{aligned} \tag{22.11}$$

Hall B's fit is that of Section 5.2.2 and hall A's that of Section 15.2.2, both on window 1 of the night record, both with their binding rack an unmeasured one, B3 and A5 (Table 15.3). The site's is the one fit of Section 9.2.2, tabulated with the stated truth in Table 9.2. The failure rates are the exact gamma posteriors of equations (9.12) and (15.11).

### 22.2.3 The draws

The answers read the posterior through $N = 4\,000$ paired draws. The site's are the first 4 000 of the draws of Section 9.2.4, from the Laplace normal of the site's fit, whose weights (5.16) toward the exact posterior have an effective sample size of 3 983 of the 4 000. Each hall's are 4 000 draws from the Laplace normal (22.11) of its fit, paired with the site's by index, which

by the factorization (22.10) makes each pair a draw of the whole unknown. The fourth units of the remedies are drawn from their stated normals (22.8), once per pair, and the failure rates of the fourth question from their gamma posteriors. Each pair is pushed through the six binding loads (22.6) at all 132 nodes, and each node's probability is the share of pairs inside, equation (22.4): the state route, joint over the elements. Written out,

$$\boldsymbol{\theta}_i = (\mathbf{z}_{A,i},\ \mathbf{z}_{B,i},\ \mathbf{z}_{\text{site},i}), \qquad \hat{P}_{um} = \frac{1}{4\,000} \sum_{i=1}^{4\,000} \mathbf{1}\Big[\min_{e} Q_e^{\dagger}(\boldsymbol{\theta}_i;\ u, m) \geq Q_m^{*}\Big], \tag{22.12}$$

with $i$ the index that pairs the three blocks. For the chillers' own rows the exact posterior, the same draws under the site's weights, stands beside the draws. At the assurance level the Monte Carlo standard error of a node is $\sqrt{0.95 \times 0.05/4\,000} = 0.0034$, the precision with which the grid places a unit's window edge.

## 22.3 The answers (step 4)

Every answer is read from the draws of Section 22.2.3. The grid as built answers the drill question and fixes the windows; the grid with the remedies and the second failure follow.

### 22.3.1 The grid as built

**Planner's question.** Can any one unit be taken out for service in any month of the year without a violation? The question asks for the smallest probability of staying inside over every unit and month, and for the unit and month at which it falls.

**General form.** A prediction under a threshold is the posterior probability of an event, compared with an assurance level, equation (2.12),

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

here taken over a family of scenarios and minimized.

**Specialized form.** The question is equation (2.12) at each node, with $f$ the smallest binding load less the month's load and $f^{*} = 0$, which is equation (22.3); the minimum and its place are equation (22.1), each node read by the draws of equation (22.12).

**Building blocks.** The six binding loads (22.6); the month's inputs of Table 22.1; the paired draws of the three fits.

**The grid.** Table 22.3 gives the whole grid. Each row is a unit out; the twelve columns give the probability of staying inside in each month, to three decimals; the last column names the limiting element at the row's smallest value. A circle marks a node at which none of the 4 000 pairs stays inside, a probability below 1/4 000; a zero marks a node with no flow path, outside in every state.

The same grid is drawn as a map in Figure 22.1. The months run across and the units down, chillers, hall A, hall B and the flow path; each cell is shaded by its probability in five bands, from green at or above the assurance level to dark red below 0.01, and carries its value. The thick line encloses the cells at or above 0.95, the units' service windows.

Table 22.3: The probability of staying inside with each unit out, by month of the plan year, as built, equation (22.3), by 4 000 paired draws of the site's and the halls' fits; Monte Carlo standard error 0.0034 at the assurance level 0.95. ○: no draw inside, below 1/4 000; 0.000: some draw inside, below 0.0005. 0: no flow path, outside in every state. The last column is the limiting element at the row's smallest value.

| Unit out [–] | Jan [–] | Feb [–] | Mar [–] | Apr [–] | May [–] | Jun [–] | Jul [–] | Aug [–] | Sep [–] | Oct [–] | Nov [–] | Dec [–] | Limiting [–] |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| C1 | 1.000 | 1.000 | 1.000 | 0.999 | 0.998 | 0.996 | 0.992 | 0.989 | 0.983 | 0.974 | 0.878 | 0.007 | chillers |
| C2 | 1.000 | 0.997 | 1.000 | 0.999 | 0.998 | 0.996 | 0.992 | 0.989 | 0.983 | 0.972 | 0.105 | ○ | chillers |
| C3 | 1.000 | 1.000 | 1.000 | 0.999 | 0.998 | 0.996 | 0.992 | 0.989 | 0.983 | 0.974 | 0.957 | 0.045 | chillers |
| A1 | 0.414 | 0.364 | 0.317 | 0.271 | 0.224 | 0.172 | 0.132 | 0.094 | 0.068 | 0.046 | 0.030 | 0.017 | hall A |
| A2 | 0.227 | 0.179 | 0.135 | 0.101 | 0.070 | 0.043 | 0.029 | 0.017 | 0.009 | 0.004 | 0.003 | 0.002 | hall A |
| A3 | 0.509 | 0.471 | 0.423 | 0.367 | 0.316 | 0.271 | 0.219 | 0.174 | 0.135 | 0.099 | 0.069 | 0.045 | hall A |
| B1 | 0.022 | 0.014 | 0.007 | 0.005 | 0.004 | 0.002 | 0.001 | 0.001 | ○ | ○ | ○ | ○ | hall B |
| B2 | 0.017 | 0.011 | 0.005 | 0.004 | 0.002 | 0.002 | 0.001 | 0.000 | ○ | ○ | ○ | ○ | hall B |
| B3 | 0.061 | 0.043 | 0.030 | 0.019 | 0.012 | 0.009 | 0.005 | 0.003 | 0.002 | 0.001 | 0.001 | 0.001 | hall B |
| pump | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | no flow path |
| strainer | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | no flow path |

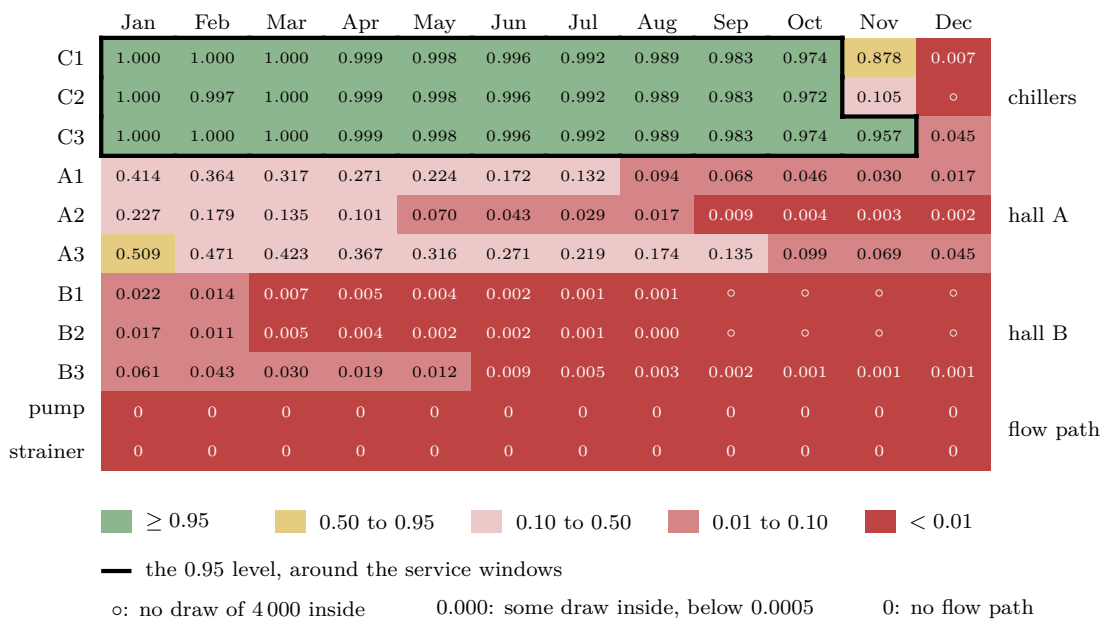



Figure 22.1: The scenario grid as built: the probability of staying inside with each unit out (rows) in each month of the plan year (columns), from Table 22.3. The thick line encloses the nodes at or above the assurance level of 0.95: the chillers from January to October, and C3 in November as well. No hall unit, and neither the pump nor the strainer, has a single cell inside it.

**The answer.** By equation (22.1) the smallest probability over the grid is

$$P_{\min} = 0, \qquad (u^{\dagger}, m^{\dagger}) \in \{(\text{pump},\ m),\ (\text{strainer},\ m) : \text{every } m\} \cup \{(\text{B1},\ m),\ (\text{B2},\ m) : m \geq 9\}, \tag{22.13}$$

exactly zero at the pump and the strainer, where the model leaves no flow path, and with no draw of 4 000 inside, a probability below 1/4 000, for B1 and B2 from September on. The answer to the drill question is no. *The site is not concurrently maintainable.* No unit of either hall can come out in any month, and neither can the pump or the strainer; the chillers can, from January to October.

Read row by row, Figure 22.1 has three bands. The chillers' rows are green to October: with one chiller out the other two, raised by the derate factor in the cooler months, carry the loop with probability 0.974 or more for C1 and C3 and 0.97 or more for C2. In November they part, 0.878 for C1, 0.105 for C2 and 0.957 for C3, and in December all three fall, to 0.007, below 1/4 000 and 0.045; the limiting element is the chillers in every one of those cells. Hall A's rows never reach the level: A3 out in January is the best of them at 0.51, and A2 out in December the worst at 0.0015. Hall B's rows are worse: B3 out in January is the best at 0.061, and B1 and B2 out have no draw inside from September. The pump and the strainer are zero in every month.

Hall B's own limit shows in the chillers' rows too. From March to October the rows of C1 and C3 are equal to the fourth decimal to the grid's row with every element up, 0.9998 in March down to 0.9740 in October, and C2's row equals it from April to September: in those months the chiller out costs nothing, and what keeps the row below one is hall B's binding load with all its units running, which no reading of rack B3 has fixed. With every unit up the plant stays inside with probability 0.9457 in December, already below the level before any unit is taken out.

The sheet passes all three of the places that fail. It passes hall B because it compares two units' duty with the hall's load, while the inlets are limited by the air the units move and by the recirculation at rack B3: Chapter 15 found hall B N on its air side, not N+1, and the grid shows the same fact as twelve months of red. It passes hall A on duty from January to October for the same reason, though hall A, with the larger binding load, fails less badly. And it passes the pump and the strainer by not listing them.

### 22.3.2 The service windows

**Planner's question.** When can each unit come out? The question asks, for each unit, the months in which its service can be scheduled at the assurance level.

**General form.** The prediction under a threshold of equation (2.12), accepted where the probability reaches the assurance level, here over a family of events, one per month.

**Specialized form.** The question is equation (2.12) at each node of a unit's row with $\kappa = 0.95$, which is the set $W_u$ of equation (22.2), each node read jointly by equation (22.12). Beside it stand the single-element routes of the chillers: the chillers' binding load alone at the same draws, its exact posterior under the site's weights, and a normal with the draws' median margin and the delta form's spread.

**Building blocks.** The grid of Section 22.3.1; the site's weights; the delta form of the chillers' margin, the gradient of equation (22.6) pushed through the site's $\mathbf{\Sigma}$ of equation (22.11).

Table 22.4: Each unit's service windows, the months whose probability of staying inside reaches 0.95, equation (22.2), and its lowest probability, as built (Table 22.3) and with the four remedies (Table 22.6). "No draw": none of the 4 000 draws inside, below 1/4 000.

| Unit out [–] | Window, as built [–] | Lowest, as built [–] | Window, remedies [–] | Lowest, remedies [–] |
|---|---|---|---|---|
| C1 | Jan–Oct | 0.0067, December | every month | 1.000, every draw |
| C2 | Jan–Oct | no draw, December | every month | 1.000, every draw |
| C3 | Jan–Nov | 0.0447, December | every month | 1.000, every draw |
| A1 | none | 0.0168, December | every month | 0.9815, December |
| A2 | none | 0.0015, December | every month | 0.9675, December |
| A3 | none | 0.0447, December | every month | 0.9860, December |
| B1 | none | no draw, Sep–Dec | Jan–Sep | 0.9120, December |
| B2 | none | no draw, Sep–Dec | Jan–Aug | 0.9022, December |
| B3 | none | 0.0005, December | Jan–Oct | 0.9343, December |
| Pump | none | 0, every month | every month | 1.000, every draw |
| Strainer | none | 0, every month | every month | 1.000, every draw |

Table 22.5: The routes at the edge of the chillers' service window. Joint: the grid, every element, Table 22.3. Chillers alone: the chillers' binding load against the load, the same 4 000 draws; exact: the same under the site's weights (effective sample size 3 983). Delta-route arithmetic: $\Phi$(median/sd) on two printed numbers of Table 22.10, the draws' median margin and the delta form's spread today; a beside-route only, never the answer. Gaps from the printed values. "No draw": none of 4 000 inside.

| Node [–] | Joint, reported [–] | Chillers alone, draws [–] | Chillers alone, exact [–] | Delta-route arithmetic [–] | Gap, joint less chillers alone [–] |
|---|---|---|---|---|---|
| C1 out, November | 0.8775 | 0.9117 | 0.9069 | 0.907 | −0.0342 |
| C2 out, November | 0.1047 | 0.1087 | 0.1035 | 0.110 | −0.0040 |
| C3 out, November | 0.9573 | 0.9938 | 0.9930 | 0.994 | −0.0365 |
| C1 out, December | 0.0067 | 0.0070 | 0.0059 | 0.0056 | −0.0003 |
| C2 out, December | no draw | no draw | no draw | below $10^{-6}$ | – |
| C3 out, December | 0.0447 | 0.0467 | 0.0421 | 0.0419 | −0.0020 |

**The windows.** Table 22.4 gives the windows. Each row is a unit; the columns give its months at or above 0.95 and its lowest probability with the month, as built and with the remedies of Section 22.3.3.

As built, by equation (22.2),

$$W_{\mathrm{C1}} = W_{\mathrm{C2}} = \{\text{January to October}\}, \qquad W_{\mathrm{C3}} = \{\text{January to November}\},$$
$$W_u = \varnothing \text{ for every other unit.} \tag{22.14}$$

The calendar can give each chiller its 72 h in any month to October, and C3 in November as well; it can give no other unit any month. November is open to C3 at 0.9573, 2.1 standard errors of 0.0034 above the level, and nearly open to C1, whose chillers alone carry it with probability 0.9117 by the draws and 0.9069 by the exact posterior. December is closed to all three.

**The edge of the chillers' window.** Table 22.5 sets the routes side by side at the two months where the chillers' windows end. Each row is a chiller out in November or December; the columns give the reported joint probability of the grid, the chillers' binding load alone by the same draws, the chillers alone by the exact posterior, the delta-route arithmetic, a normal with the draws' median margin and the delta form's spread, and the gap of the joint value to the chillers alone.

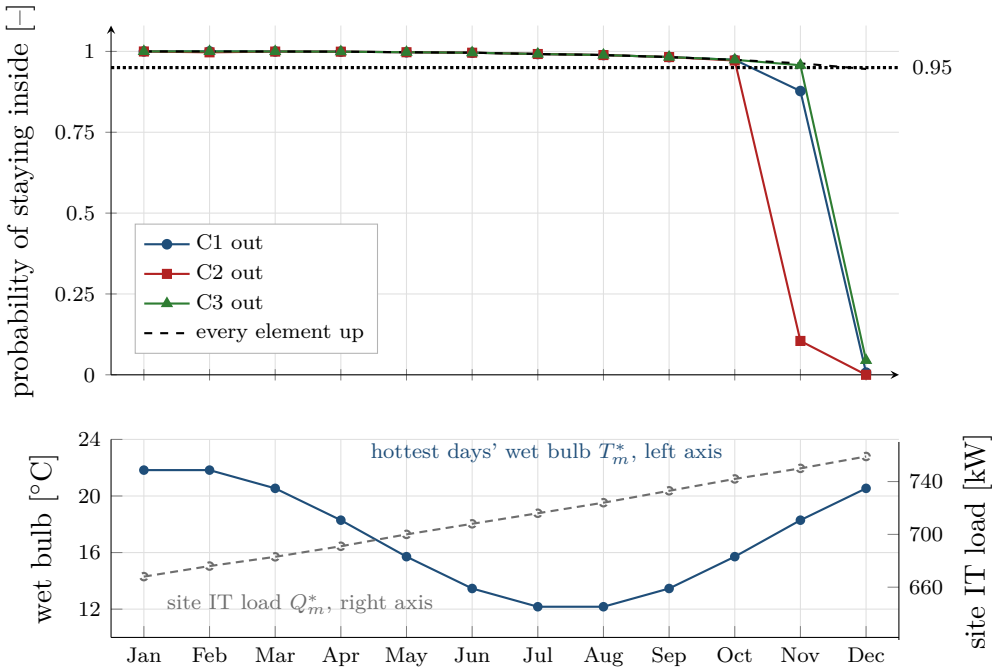



Figure 22.2: The chillers' service windows. Upper panel: the probability of staying inside with each chiller out, by month, from the grid of Table 22.3; dashed, the grid with every element up, the ceiling hall B sets; dotted, the assurance level of 0.95. The three lines lie together to October and part in November. Lower panel: the hottest days' wet bulb (left axis, solid) and the plan year's site IT load (right axis, dashed), the two inputs of Table 22.1.

The chillers alone are the worst marginal, the naive baseline, and the joint lies below it in every row. The gap is the share of the other elements, and it is hall B's: in November the grid's row with every element up is 0.9630, and the joint probability with a chiller out cannot exceed it. For C1 the gap of 0.0342 is the difference between a month nearly open and a month clearly closed; for C3 it takes 0.9938 down to 0.9573, still above the level. In December the chillers alone decide, and the gap is a few thousandths. The single-entry routes agree with each other. The exact posterior sits below the draws by 0.0048, 0.0052 and 0.0008 in November, each about one Monte Carlo standard error (0.0045, 0.0049 and 0.0012), and by 0.0011 and 0.0046 in December. The delta-route arithmetic lands within 0.007 of the exact posterior in every row, because the chillers' margin is a sum of exponentials of small, nearly normal log-ratios, close to normal itself; it is a beside-route and decides nothing here, since no threshold falls between it and the others.

The chillers' rows are drawn against the month in Figure 22.2. In the upper panel the horizontal axis is the month and the vertical axis the probability of staying inside: one line per chiller out, from the grid, the dashed line the grid's row with every element up, and the dotted line the assurance level. The lower panel puts the month's two inputs on the same months, the hottest days' wet bulb on the left axis and the site IT load on the right.

The upper panel shows the chillers' window as a cliff, not a slope. The three lines sit within a few hundredths of one through October and fall together in two months, because the load rises through the year while the wet bulb, after its winter low, climbs back toward the summer value: in the lower panel the two inputs meet the year's end together. January and February are warmer than December, by 1.29 K, and carry 91 and 83 kW less load; the load wins, and December is the binding month. The dashed line, every element up, is the ceiling that hall B lays over every row; in November it already sits at 0.9630, so C1's November cannot open unless

Table 22.6: The probability of staying inside with each unit out, by month, with the four remedies built: a fourth unit in each hall, the fourth chiller, a standby pump and a strainer bypass. 4 000 paired draws, the fourth units drawn from their stated normals; Monte Carlo standard error 0.0034 at 0.95. Below 0.95 in bold.

| Unit out [–] | Jan [–] | Feb [–] | Mar [–] | Apr [–] | May [–] | Jun [–] | Jul [–] | Aug [–] | Sep [–] | Oct [–] | Nov [–] | Dec [–] |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| C1–C3, pump, strainer | 1.000 in every month: every draw inside | | | | | | | | | | | |
| A1 | 0.999 | 0.999 | 0.999 | 0.999 | 0.998 | 0.997 | 0.996 | 0.994 | 0.991 | 0.989 | 0.986 | 0.982 |
| A2 | 0.999 | 0.999 | 0.999 | 0.998 | 0.996 | 0.995 | 0.992 | 0.990 | 0.986 | 0.982 | 0.977 | 0.968 |
| A3 | 0.999 | 0.999 | 0.999 | 0.999 | 0.999 | 0.999 | 0.998 | 0.995 | 0.993 | 0.990 | 0.989 | 0.986 |
| B1 | 0.997 | 0.996 | 0.994 | 0.990 | 0.987 | 0.981 | 0.973 | 0.964 | 0.955 | **0.944** | **0.929** | **0.912** |
| B2 | 0.996 | 0.993 | 0.991 | 0.989 | 0.985 | 0.976 | 0.967 | 0.959 | **0.9495** | **0.936** | **0.919** | **0.902** |
| B3 | 0.999 | 0.997 | 0.996 | 0.995 | 0.993 | 0.989 | 0.983 | 0.976 | 0.969 | 0.958 | **0.947** | **0.934** |

both the chillers and hall B are nearly certain.

### 22.3.3 The grid with the remedies

**Planner's question.** What do the named remedies buy? The question asks for the grid and the windows with a fourth unit in each hall, the fourth chiller, a standby pump and a strainer bypass built.

**General form.** A design question compares the posterior quantity of interest under the design alternatives; here the quantity is the prediction under a threshold of equation (2.12), read at a scenario node whose pinned ratings are those of the alternative.

**Specialized form.** The question is equation (2.12) at each node of the grid with the remedies' terms (22.8) added to the binding loads (22.6), read by the same draws, equation (22.12), with the fourth units drawn once per pair; the minimum and the windows are equations (22.1) and (22.2).

**Building blocks.** The grid as built; the fourth units' stated normals; the standby pump and the bypass as a second flow path.

**The grid with the remedies.** Table 22.6 gives it. The chillers, the pump and the strainer occupy one row, since every draw stays inside with any of them out in every month; each hall unit has its row, its twelve months to three decimals.

The same grid is mapped in Figure 22.3, in the layout, bands and contour of Figure 22.1.

The remedies close most of the grid. The chillers, the pump and the strainer can come out in every month: the fourth chiller makes the chiller out one of four, and the standby pump and the bypass give the water its second path. Hall A's units can come out in every month, the worst at 0.9675, A2 in December. Hall B's cannot. With its fourth unit built, hall B's rows fall through the year to 0.9120, 0.9022 and 0.9343 in December for B1, B2 and B3, and drop below the level from October for B1, from September for B2 and from November for B3. By equation (22.1) the grid's minimum with the remedies is

$$P_{\min}^{\text{rem}} = 0.9022, \qquad (u^{\dagger}, m^{\dagger}) = (\text{B2},\ \text{December}), \tag{22.15}$$

below the assurance level. *Even with a fourth unit, hall B cannot release a unit in the last months of the plan year.* Chapter 20 finds that a fourth unit restores hall B's N+1 for under

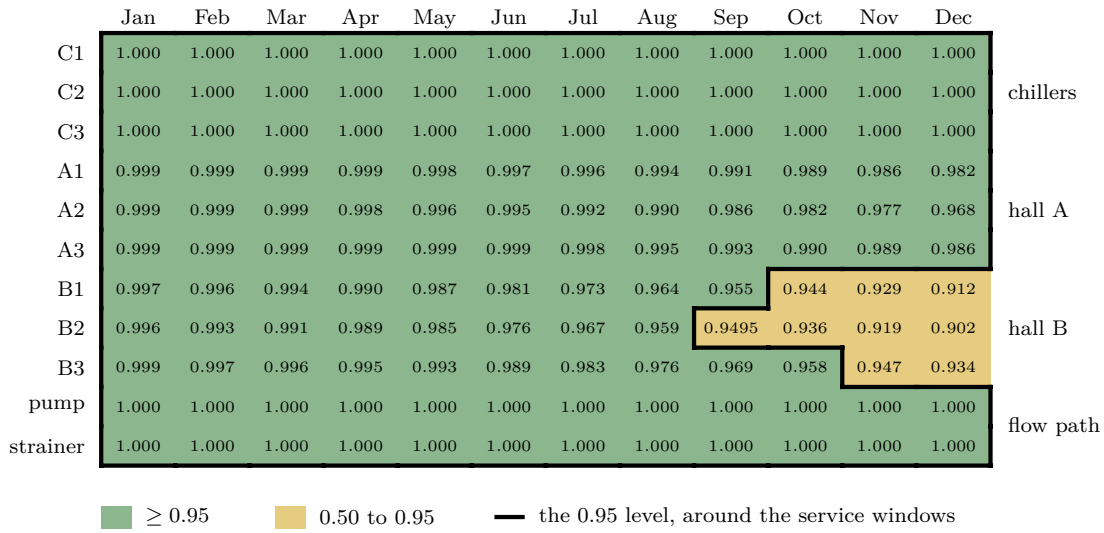



Figure 22.3: The scenario grid with the four remedies built, from Table 22.6, in the layout of Figure 22.1. The thick line now encloses every node but the lower right of hall B's rows: B1 from October, B2 from September and B3 from November fall below the assurance level.

two years of growth; the grid shows the same erosion inside the first year, as months. With every element up and the remedies built the site stays inside in every draw of every month.

Two of the window edges lie within the grid's precision. B2's September is 0.9495, below the level by 0.0005, a seventh of a standard error; B1's September is 0.9547, above it by 1.4 standard errors. The months themselves are not in doubt: by October every hall-B row is below 0.96 and falling by one to two hundredths a month.

### 22.3.4 A second failure during the service

**Planner's question.** What if a second unit fails during the service? The question asks how likely it is that another unit fails inside the 72 h window and leaves the site outside, with the remedies built.

**General form.** When the event lies in the future and stays random even when the state is known, the probability is taken of the predictive distribution, equation (2.13),

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

here with $y_{\text{new}}$ the event that a second unit fails in the window and the site is then outside.

**Specialized form.** The question is equation (2.13) with the event's probability given the unknown the exponential chance of each unit's failure in $\tau = 72$ h times the indicator of the node with both units out, summed over the second unit, which is equation (22.5), read over the draws (22.12) with each draw's failure rates from the gamma posteriors.

**Building blocks.** The grid with the remedies; the nodes with two units out; the posteriors of the chillers' and the hall units' failure rates, equations (9.12) and (15.11).

**The second failure.** Table 22.7 gives it at four nodes. Each row is a unit out and a month; the columns give the node's probability of staying inside with the remedies, the probability that

Table 22.7: A second failure during the 72 h service window, with the four remedies built, equation (22.5): the eight other chillers and hall units can fail, at the rates' posteriors; 4 000 paired draws. Inside: the node with the one unit out, Table 22.6.

| Node [–] | Inside, one out [–] | A second unit fails [–] | A second failure, outside [–] |
|---|---|---|---|
| C2 out, July | 1.000 | 0.0619 | $6.55 \times 10^{-4}$ |
| C2 out, December | 1.000 | 0.0619 | 0.0201 |
| A2 out, December | 0.9675 | 0.0638 | 0.0156 |
| B2 out, December | 0.9022 | 0.0638 | 0.0191 |

a second unit fails in the window, and the probability of a second failure that leaves the site outside.

A second failure is about a one-in-sixteen event in any window, 0.0619 with a chiller out and 0.0638 with a hall unit out; the rates do not depend on the month. Whether it matters does. With C2 out in July, a second failure leaves the site outside with probability $6.55 \times 10^{-4}$: the cool month's chillers and halls have room for two units out. With C2 out in December the same failure leaves it outside with probability 0.0201, about a third of the second failures, and with A2 or B2 out in December 0.0156 and 0.0191. In December, with the remedies built, a planned outage carries roughly a two-percent chance that an unplanned failure during it takes the site outside.

This is the fault-tolerance question, and concurrent maintainability does not answer it. A plant can be concurrently maintainable, every node of the grid above the level, and still fall short when a second unit fails during the service; the July and December rows of C2 have the same probability of staying inside with one unit out, one, and differ thirtyfold on a second failure. The grid answers the planned outage alone.

## 22.4 The checks (step 5)

The answers of Section 22.3 rest on draws of three fits, on the scenario's choice of the month's hottest days, and on records generated from a stated truth; each is tested here.

### 22.4.1 The routes against each other

Three checks bear on the computation. First, the exact posterior against the draws, on the chillers' own rows in November and December (Table 22.5). In November the exact values are 0.9069, 0.1035 and 0.9930 against 0.9117, 0.1087 and 0.9938 by the draws with C1, C2 and C3 out, gaps of 0.0048, 0.0052 and 0.0008 against Monte Carlo standard errors of 0.0045, 0.0049 and 0.0012. In December they are 0.0059 against 0.0070 with C1 out, a gap of 0.0011 within one standard error of 0.0013; 0.0421 against 0.0467 with C3 out, a gap of 0.0046, 1.4 standard errors of 0.0033; and no draw inside, by either, with C2 out, both below 1/4 000. The site's draws lose little to the exact posterior, with an effective sample size of 3 983 of 4 000, and the halls' Laplace normals, which are not reweighted here, lose as little in Chapter 15's check of hall A (Section 15.2.2). Second, the joint against the worst marginal: the joint is at or below the chillers alone at every node of Table 22.5, as an intersection must be, and the gap is no larger than the room hall B leaves, one less the grid's row with every element up. Third, the structure of the grid: every row is at or below the row with every element up, the chillers' rows equal it where the chillers have room, and the remedies raise every node.

Table 22.8: The checks of the maintenance-calendar answers. Reported values by 4 000 paired draws; the truth is the stated truth of the record, for the chillers.

| Quantity [unit] | Reported [–] | Beside [–] | Truth [–] | Reading |
|---|---|---|---|---|
| C1's window | January–October | chillers alone, November 0.9117 | January–November | one month short, on the side of the truth |
| C2's window | January–October | – | January–October | agrees |
| C3's window | January–November | – | January–November | agrees |
| C1 out, November | 0.9117 chillers alone | exact 0.9069 | – | gap 0.0048, se 0.0045 |
| C2 out, November | 0.1087 chillers alone | exact 0.1035 | – | gap 0.0052, se 0.0049 |
| C3 out, November | 0.9938 chillers alone | exact 0.9930 | – | gap 0.0008, se 0.0012 |
| C1 out, December | 0.0070 chillers alone | exact 0.0059 | – | gap 0.0011, se 0.0013 |
| C3 out, December | 0.0467 chillers alone | exact 0.0421 | – | gap 0.0046, se 0.0033 |
| C2 out, December | no draw | exact: no draw | – | both below 1/4 000 |
| Joint against chillers alone | at or below, every node | gap at most 0.0365 | – | hall B's share |
| Hottest month's wet bulb [°C] | 21.83 | design day 22.0 | – | stated 2 K consistent |
| Draws | 4 000 paired: the site's first 4 000 of Chapter 9 (effective sample size 3 983), each hall's 4 000 from its Laplace normal; se 0.0034 at 0.95 | | | |

### 22.4.2 The answers against the stated truth

The records are generated from a stated truth, and the last checks compare the answers with it. With the true chillers, the two left with C1 or C3 out carry the load from January to November, and with C2 out from January to October. The posterior's windows open every one of those months but one: C1 in November, where the truth carries the load and the posterior gives the chillers alone 0.9117 and the grid 0.8775, below the level but on the side of the truth. The windows never open a month the truth closes. The month's hottest days are checked against the drawing: the hottest month's are 21.83°C, against the design day's 22.0°C, so the stated 2 K above the monthly mean reproduces the design day to under two tenths of a kelvin. Table 22.8 collects the checks: each row is a quantity, the columns its reported value, the route or reference beside it, the truth and the reading.

The answers hold where they can be checked, and they lean the safe way where they differ. The truth does not reach the halls' rows here: the grid's red for hall B rests on the posterior of its binding load, whose blind spot is rack B3, and Chapter 15 found the truth's binding load for hall B below the posterior's (Table 15.3), so the true hall B is, if anything, worse than the grid shows.

## 22.5 The report (step 6)

The site is not concurrently maintainable today. The chillers can be serviced from January to October, C3 in November as well, and none in December. No hall unit can be serviced in any month, and neither can the pump or the strainer. With the remedies built, hall B's units still cannot come out from September, October or November on. The maintenance calendar is a capacity result, and a calendar written from the sheet schedules violations.

Table 22.9: The report on the site's maintenance calendar, plan year at 15 % a year from 660 kW, the month's hottest days, a 72 h service window and the assurance level 0.95. Reported values by 4 000 paired draws, joint over the elements.

| Question | Answer | Band or routes beside |
|---|---|---|
| Can any one unit be taken out for service in any month of the year without a violation? | no: $P_{\min} = 0$, at the pump and the strainer (no flow path) and at B1 and B2 from September (no draw inside) | hall A's best 0.51 (A3, January); hall B's best 0.061 (B3, January); every element up, 0.9457 in December |
| When can each unit come out? | C1, C2 January to October; C3 January to November; every other unit never | se 0.0034 at 0.95; C1 in November 0.8775, chillers alone 0.9117, exact 0.9069; December closed, chillers alone exact 0.0059 (C1), below 1/4 000 (C2), 0.0421 (C3) |
| What do the named remedies buy? | every month for the chillers, the pump, the strainer and hall A ($\geq 0.9675$); hall B to September (B1), August (B2), October (B3) | minimum 0.9022, B2 in December |
| What if a second unit fails during the service? | with the remedies, C2 out in December: 0.0201; B2: 0.0191; A2: 0.0156 | a second unit fails in the window 0.0619 to 0.0638; C2 out in July $6.55 \times 10^{-4}$ |

### 22.5.1 The answers and their bands

Table 22.9 gives the answers in the order of Table 22.2. Each row is a question; the second column gives the reported answer, and the third its band or the routes beside it.

### 22.5.2 The design-review entry

The design review's entry does not accept the sheet's maintenance calendar. The sheet checks each unit out against duty and nameplate, passes the chillers and the halls in nearly every month, and does not list the pump or the strainer; on the site's posterior, read jointly over every element at each month's hottest days and the plan year's load, three places fail in every month. The pump and the strainer have no second path, and either out stops the flow. Hall B's air side is N, not N+1, as Chapter 15 found; with any of its units out the probability of staying inside is at most 0.061. Hall A's is N as well, at most 0.51. The chillers alone are concurrently maintainable for ten months, January to October, with C3 also in November; December is closed to them.

The entry records the calendar the plant allows today, chillers only, January to October, and records that no hall unit, the pump or the strainer can be serviced without a violation at the assurance level. It records the four remedies as design alternatives with what each buys: the standby pump and the strainer bypass open the flow path in every month; the fourth chiller opens the chillers' December; a fourth unit in hall A opens hall A's every month; a fourth unit in hall B opens hall B to August, September or October of the plan year, by unit, and not after. It records that concurrent maintainability does not answer a second failure during the service, which with the remedies is about two percent in December. Their prices are the subject of Chapters 23 and 24.

Table 22.10: The chillers' margin with one unit out at the edge of their service window: the binding load of the two left running less the month's load, by the first 4 000 draws of the site's fit; the spread by the delta form today and after capacity tests of the two running chillers at 22°C, equation (5.26).

| Node [–] | Median [kW] | 5 % point [kW] | 95 % point [kW] | Pr(> 0) [–] | sd today [kW] | sd after tests [kW] |
|---|---|---|---|---|---|---|
| C1 out, November | 10.4 | −2.3 | 23.6 | 0.9117 | 7.87 | 6.69 |
| C2 out, November | −9.8 | −22.6 | 3.2 | 0.1087 | 7.99 | 7.15 |
| C3 out, November | 20.1 | 6.9 | 33.4 | 0.9938 | 8.08 | 7.31 |
| C1 out, December | −23.7 | −38.6 | −7.9 | 0.0070 | 9.35 | 6.79 |
| C2 out, December | −43.2 | −56.7 | −29.4 | no draw | 8.38 | 6.52 |
| C3 out, December | −14.4 | −27.7 | −0.4 | 0.0467 | 8.33 | 6.50 |

**Common mistake**

Writing the maintenance calendar from the design sheet. The sheet checks a unit out against nameplate duty, one element at a time, and passes nearly every month; the plant stays inside only if every element carries the load at once, at the month's weather and load, and the elements the sheet does not list, a single pump and a single strainer, fail the calendar in every month. A calendar is a capacity result and must be read jointly from the plant's posterior.

## 22.6 The next measurement (step 7)

The report leaves two kinds of edge. The hall rows and the flow path fail by wide margins that no measurement on the existing plant can close: they need equipment. The chillers' window ends on a measurable edge, November, and step 7 asks what would move it.

### 22.6.1 The edge of the chillers' window

The chillers' margin with one unit out, the chillers' binding load less the month's load, decides their November and December. Table 22.10 gives it. Each row is a chiller out in November or December; the columns give the margin's median and 90 % band over the draws, its probability of being positive, and its delta form's spread today and after capacity tests of the two chillers left running, read at the design day's 22°C, by the preposterior covariance of equation (5.26).

November with C1 out is the one node on the edge. Its margin is 10.4 kW, with a band from −2.3 to 23.6 kW and a probability of 0.912 of being positive; capacity tests of C2 and C3 at 22°C would take its spread from 7.87 to 6.69 kW, about fifteen percent. As delta-route arithmetic on those two printed numbers, a normal at today's median with the narrower spread gives $\Phi(10.4/6.69) = 0.94$, a beside-route rather than an answer, and still short of the level even before hall B's ceiling below is applied: the tests open November to C1 only if they read C2 and C3 above their fitted ratios, and they close it for good if they read them below. They decide the month; they do not open it by narrowing. And the joint grid adds hall B's ceiling of 0.9630 in November, so a November service of C1 needs the tests to read high and hall B's band to narrow as well, the rack B3 measurement that Chapter 5 ranks first. December is beyond any test: its margins run from −14.4 to −43.2 kW at the median, and a test that narrows a band does not move a margin that far.

### 22.6.2 The recommendation

Schedule the chillers' services from January to October, one at a time, and none in November or December until a test has read C2 and C3. Do not schedule any hall unit, the pump or the strainer: as built, the site cannot release them at the assurance level in any month. Take the four remedies to the design review as stated alternatives, with what each buys from Table 22.4: the standby pump and the strainer bypass first, since without them the flow path has no window at all; hall A's fourth unit, which opens every month; and hall B's fourth unit, which opens only the first eight to ten months of the plan year and needs the rack B3 measurement and a second step after it. Chapter 18 ranks the tests on the existing plant that would narrow the build decision; the capacity tests of C2 and C3 are the one that bears on this calendar.

### 22.6.3 What the analysis showed

The design sheet passes nearly all of the maintenance calendar, and the plant fails it. Read as 132 scenario nodes, one per unit out and month, each joint over the six elements at the month's hottest days and the plan year's load, the site stays inside with each chiller out from January to October and with no other unit out in any month. The smallest probability on the grid is zero, at the pump and the strainer, which the sheet does not list and which have no second path; hall B's units, which the sheet passes on duty, have no draw inside from September. The chillers' window ends in a cliff two months wide, because the load grows through the year while the wet bulb climbs back toward its summer value; December binds, not the summer. Hall B's own limit lies over every row as a ceiling, so that even the chillers' rows are capped by the hall they serve. The four remedies close the flow path, the chillers and hall A, and leave hall B's units closed from September, October or November: one unit of margin in hall B is used up by the year's growth. And the second failure during a service, fault tolerance, is a separate answer: about two percent in December with the remedies, thirty times its July value.

### 22.6.4 Limits of the model

**The month's hottest days.** Each month is checked at its mean wet bulb plus a stated 2 K, the same for every month. A month whose hottest hours lie further above its mean, or a service window that falls in a cooler week of the month, would move its cells; the stated 2 K reproduces the design day in the hottest month and is not otherwise estimated.

**The month's end.** Each month carries its end-of-month load, the largest it sees. A service early in the month meets less load, so the windows are slightly conservative at their edges.

**One unit out, no repair of the plan.** A unit out is taken out whole and the plant otherwise unchanged: no temporary chiller, no reset of the supply, no load moved between halls. Each of these is a design alternative that the grid could read as another node; none is read here.

**The flow path as drawn.** The pump and the strainer are single points because the drawing shows them so. The zeros of their rows are the drawing's, not the record's, and a site with a cross-connection or a hand bypass would need that path drawn and read.

**The supply pinned.** The halls' binding loads are taken at the supply of 14°C whatever the chillers do, so the chillers' shortfall and the halls' are read as separate events at each node. A chiller plant short of the load would warm the supply and pull both halls' binding loads down at once, the chain Chapter 21 follows.

**The constructed record.** The night records, the operating record and the outage records are generated from a stated truth with stated errors, not taken from a site, and the comparison with that truth in Section 22.4.2 is a check of the method on its own model, not of the model against a real plant.

**The next part.** Part VI has read the plant under failure: the headroom left with the largest unit out (Chapter 20), the overloads that one unit's trip can carry to another (Chapter 21), and here the calendar on which the plant can be serviced. Part VII prices what those failures cost: the expected unserved energy of Chapter 23 and the risk–cost curve of Chapter 24.

## 22.7 What to retain

**Definition 22.1** (Concurrent maintainability as a scenario grid)**.** For a plant with units $\mathcal{U}$ and months $m = 1, \dots, 12$, the probability of staying inside with unit $u$ out in month $m$ is $P_{um} = \Pr(\min_e Q_e^{\dagger}(\boldsymbol{\theta};\ u, m) \geq Q_m^{*} \mid \mathbf{y})$, joint over the plant's elements at the month's load and weather. The plant is concurrently maintainable at the assurance level $\kappa$ when $P_{\min} = \min_{u,m} P_{um} \geq \kappa$; unit $u$'s service windows are $W_u = \{m : P_{um} \geq \kappa\}$. A second failure during the service, $\sum_{v \neq u} \mathbb{E}[(1 - \mathrm{e}^{-\lambda_v \tau})\, \mathbf{1}[\text{outside with } u, v \text{ out}] \mid \mathbf{y}]$, is fault tolerance and is reported beside, never in place of, the grid.

Four conclusions carry beyond the site. A maintenance calendar is a capacity question asked at many scenario nodes, and its answer is a minimum, a place and a set of windows, not a single margin. Each node is read jointly over every element of the plant; the element the sheet checks is the worst marginal, and the gap to the joint is the share of the elements it does not check, here hall B's. The binding month is set by the load and the weather together, and on a growing load it can fall far from the hottest month. And concurrent maintainability and fault tolerance are two answers on the same grid: the first asks whether the plant carries the load with one unit out on purpose, the second whether it survives a second unit out by accident during that outage, and a plant can pass the first and not the second. Neither is an availability: each is a capability probability, the probability that the plant as it is carries a stated load in a stated state.

Table 22.11 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 22.11: The terms of concurrent maintainability and where each is defined. Units for the site.

| Term | Expression | Units | Meaning | Defined in |
|---|---|---|---|---|
| Staying inside at a node | $P_{um} = \Pr(\min_e Q_e^{\dagger} \geq Q_m^{*} \mid \mathbf{y})$ | – | capability with unit $u$ out in month $m$ | 22.1.1 |
| Grid minimum | $P_{\min} = \min_{u,m} P_{um}$, and its arg min | – | concurrently maintainable if $\geq \kappa$ | 22.1.1 |
| Service windows | $W_u = \{m : P_{um} \geq \kappa\}$ | – | months unit $u$ may come out | 22.1.1 |
| Second failure | $\sum_{v \neq u} \mathbb{E}[(1 - \mathrm{e}^{-\lambda_v \tau})\, \mathbf{1}[\text{outside}]]$ | – | fault tolerance during the service | 22.1.1 |
| Monte Carlo error | $\sqrt{P(1-P)/N}$ | – | 0.0034 at 0.95, $N = 4\,000$ | 22.1.1 |
| *The site's nodes* | | | | |
| Chillers, unit out | $[1 - d(T_m^{*} - 22)] \sum_{i \neq u} 400\,\mathrm{e}^{\kappa_i} - Q_o$ | kW | binding load at the month's wet bulb | 22.1.2 |
| Hall, unit out | equation (15.4) over $R_h(u)$, over $\varphi_h$ | kW | binding load with its running units | 22.1.2 |
| No flow path | $Q_{\text{pipe}}^{\dagger} = Q_{\text{pump}}^{\dagger} = 0$ | kW | pump or strainer out, no second path | 22.1.2 |
| Month's inputs | $Q_m^{*} = 660\,(1.15)^{m/12}$, $T_m^{*} = T_m + 2$ | kW, °C | load at month's end, hottest days | 22.1.3 |
| *Questions and answers* | | | | |
| As built | $P_{\min} = 0$; windows: chillers to October, C3 to November | – | not concurrently maintainable | 22.3.1 |
| With the remedies | $P_{\min} = 0.9022$, B2 in December | – | hall B closed late in the year | 22.3.3 |

## 22.8 Exercises

**Exercise 22.1.** *Formulation.* Write the planner's question *can C1 be serviced in November?* as a probability over the site's unknown, say which blocks of the state it depends on, and give its value and the decision at the assurance level of 0.95. (It is $P_{\mathrm{C1},11} = \Pr(\min_e Q_e^{\dagger}(\boldsymbol{\theta};\ \mathrm{C1},\ 11) \geq 750 \mid \mathbf{y})$, joint over the six elements at $T_{11}^{*} = 18.29$°C. By the pattern (22.9) it depends on the chillers' log-ratios and derate, hall A's and hall B's states, and the pipe's, the pump's and the feed's entries. November is not in C1's window: the value is below 0.95, at 0.8775.)

**Exercise 22.2.** *Derivation.* Show that the probability of staying inside with a chiller out cannot exceed the probability with every element up, and use it to bound C1's November from the grid. (Taking a chiller out removes a nonnegative term from the chillers' sum in equation (22.6) and changes no other element, so each binding load with the chiller out is at most its value with every element up; the event of staying inside with C1 out is contained in the event with every element up, and its probability is at most that one's. In November the grid with every element up is 0.9630, so $P_{\mathrm{C1},11} \leq 0.9630$; the reported value is 0.8775.)

**Exercise 22.3.** *Derivation.* From equation (22.5), derive the probability that a second unit fails at all during a 72 h window with a chiller out, and say why it does not depend on the month while the probability of a violation does. (Summing the first factor over the two other chillers and the six hall units gives $\sum_{v \neq u} \mathbb{E}[1 - \mathrm{e}^{-\lambda_v \tau} \mid \mathbf{y}]$, which reads only the failure rates' posteriors, and they do not depend on the month; the indicator reads the node with both units out, whose load and wet bulb are the month's. With C2 out it is 0.0619 in every month.)

**Exercise 22.4.** *Application.* The maintenance contractor proposes to service hall B's units in March, the month after the hottest, arguing that the sheet's two units of 180 kW carry hall B's

load with room to spare. Say what the grid answers, as built and with hall B's fourth unit. (As built, B1, B2 and B3 out in March stay inside with probability 0.0067, 0.0053 and 0.0295: hall B is N on its air side, and the sheet's duty does not see the inlets. With the four remedies built, March is open: 0.9935, 0.9912 and 0.9962.)

**Exercise 22.5.** *Application.* With the remedies built, the planner wants a single month in which every unit of the site can be serviced, one after another. Name the latest month that serves, and the probability that a second failure during B2's service in December takes the site outside. (Every row of the grid with the remedies is at or above 0.95 through August, and B2 falls below the level in September at 0.9495, so August is the latest month that serves every unit. In December a second failure during B2's service takes the site outside with probability 0.0191.)

# Part VII

# Expected Cost and Risk

# Chapter 23

# Expected Unserved Energy

## 23.1 The site's shed energy over a year

The design sheet closes with a risk line. It takes the site's ten elements that can fail, the three chillers, the six hall cooling units and the strainer in the shared pipe, gives each chiller its nameplate of 400 kW and each hall unit the sheet's duty, 210 kW in hall A and 180 kW in hall B, and gives each kind the handbook's outage rates: 1.5 failures a year and 36 h to repair for a chiller, 1.0 a year and 24 h for a hall unit. It enumerates the states of those units over the plan year, month by month, as the site's IT load grows from 668 kW in January to 759 kW in December, and prints three numbers: an expected unserved energy of 532.2 kWh/yr, 13.24 h a year short, and an "availability" of 0.998488, the hours short taken from the hours of the year. The sheet does not list the strainer; it has no state for it.

The planner's question is the drill question, verbatim,

*How many kWh of IT load would be shed per year, in expectation, under this design?*

It is answered with the method of Chapter 4 on the site's one fit of Chapter 9, the two halls' fits of Chapters 5 and 15, and the outage posteriors of both. Three things are new. The question is an expectation, not a probability: it asks how much energy, not how often. It runs over the year, month by month, at each month's load and weather; this chapter replaces Chapter 9's design-day basis with the year. And its answer splits into two parts of different kinds, one from the plant short with nothing down, a capacity problem, and one from the plant short because a unit is down, an availability problem, and the planner needs to know which part is which.

**Jargon: expected unserved energy, loss-of-load expectation (hours short), plan year, all-up state, capacity term, availability term, limiting element, share from states with an element down, residual mass, shed value**

The *expected unserved energy* (EUE) of a year is the IT energy the site cannot carry, summed over the year, in expectation over the states of its units: kWh a year. The *loss-of-load expectation* (LOLE) is the expected number of hours in the year during which the site is short, whatever the depth: hours a year; this chapter calls it the *hours short*. The *plan year* is the coming twelve months, each with its own load and weather. The *all-up state* is the state with every element up. The *capacity term* of the EUE is the part from the all-up state, the plant short with nothing down; the *availability term* is the part from the states with one or more elements down. The *limiting element* of a state is the group of smallest capacity, as in Chapter 15. The *share from states with an element down* counts every state with that element down, whatever limits it, so two elements down count toward both. The *residual mass* is the probability of the states the enumeration leaves out. The *shed value* is the stated value of one kWh of IT load not served, in reais.

**Key idea**

Expected unserved energy is a sum over the months of the year and over the states of the plant's units, each term a state's probability times its shortfall at that month's load. The posterior enters twice: through the fits, in the capacity of every state, and through the outage posteriors, in the probability of every state. One term of the sum, the all-up state's, measures capacity: the plant short with nothing broken. Every other term measures availability: the plant short because a unit is down. The two are reported side by side, with the hours short beside the energy, and neither is ever converted into the other.

### 23.1.1 The shed energy as a sum over months and states

The quantity the question owes is the posterior expectation of the year's shed energy, with its band. Write $\boldsymbol{\theta}$ for the unknown of the plant and its outage rates, $\mathbf{s} = (s_e)$ for a state of the ten elements with $s_e = 1$ when element $e$ is up (for the strainer: clean) and 0 when it is down (fouled), $m = 1, \ldots, 12$ for the months of the plan year, $L_m$ for the month's site IT load and $h_m$ for its hours. The answer is the expectation, and each line below writes the term above it in terms of the next,

$$\begin{aligned}
\mathbb{E}[\mathrm{EUE} \mid \mathbf{y}] &= \int \mathrm{EUE}(\boldsymbol{\theta})\, p(\boldsymbol{\theta} \mid \mathbf{y})\, \mathrm{d}\boldsymbol{\theta}, \\
\mathrm{EUE}(\boldsymbol{\theta}) &= \sum_{m=1}^{12} h_m\, S_m(\boldsymbol{\theta}), \\
S_m(\boldsymbol{\theta}) &= \sum_{\mathbf{s}} P(\mathbf{s} \mid \mathbf{q})\, \max(L_m - \mathrm{Cap}(\mathbf{s},\ m;\ \boldsymbol{\theta}),\ 0), \\
\mathrm{Cap}(\mathbf{s},\ m;\ \boldsymbol{\theta}) &= \min_{g \in \mathcal{G}} \mathrm{Cap}_g(\mathbf{s},\ m;\ \boldsymbol{\theta}), \\
P(\mathbf{s} \mid \mathbf{q}) &= \prod_e q_e^{1-s_e}\, (1 - q_e)^{s_e}, \qquad q_e = \frac{\lambda_e}{\lambda_e + \mu_e}.
\end{aligned} \tag{23.1}$$

The first line is the expected-consequence form of equation (2.23) with the year's shed energy as the loss and no action yet taken. The second sums the year's months, each its hours $h_m$ times $S_m$, the expected shortfall in kW at the month's load, so that $\mathrm{EUE}(\boldsymbol{\theta})$ is in kWh a year. The third is the enumeration's expectation over the unit states: each state's probability times the load it cannot carry, zero in a state that carries the load. The fourth is the capacity of a state in site IT load, the smallest of the groups' capacities, since the load must pass through every group; $\mathcal{G}$ holds the six of Chapter 15, the chillers, hall A, hall B, the pipe, the pump and the feed. The last is the probability of a state, the product of its elements' probabilities, with each element down the fraction $q_e$ of the time, its failure rate $\lambda_e$ over the sum of its failure and repair rates $\mu_e$. The enumeration and its derivation from the units' up and down processes are those of *Reliability Engineering by Derivation*, Chapters 13 to 15, which Chapter 9 introduced and equation (15.1) extended to ten elements; this chapter uses the result.

The same sum with the indicator of a shortfall in place of its depth counts hours rather than energy,

$$\mathrm{LOLE}(\boldsymbol{\theta}) = \sum_{m=1}^{12} h_m \sum_{\mathbf{s}} P(\mathbf{s} \mid \mathbf{q})\, \mathbf{1}\big[\mathrm{Cap}(\mathbf{s},\ m;\ \boldsymbol{\theta}) < L_m\big], \tag{23.2}$$

the hours short, in hours a year. The two are readers of one enumeration, as Chapter 9's availability and headroom are: the hours short count how long, the shed energy measures how deep. Their ratio, EUE/LOLE, is the mean shortfall when short, in kW.

**Capacity against availability.** The sum over states in (23.1) has one term that is not like the others. With $\mathbf{s}_0$ the all-up state, splitting it off gives

$$\begin{aligned}\mathrm{EUE}(\boldsymbol{\theta}) = &\underbrace{\sum_m h_m\, P(\mathbf{s}_0 \mid \mathbf{q})\, \max(L_m - \mathrm{Cap}(\mathbf{s}_0,\ m),\ 0)}_{\mathrm{EUE}_{\mathrm{cap}}\text{: nothing down}} \\ &+ \underbrace{\sum_m h_m \sum_{\mathbf{s}\neq\mathbf{s}_0} P(\mathbf{s} \mid \mathbf{q})\, \max(L_m - \mathrm{Cap}(\mathbf{s},\ m),\ 0)}_{\mathrm{EUE}_{\mathrm{av}}\text{: units down}},\end{aligned} \tag{23.3}$$

with $\boldsymbol{\theta}$ dropped from the capacities. The first term is zero whenever the plant with every element up carries every month's load: it measures *capacity*, and no amount of redundancy changes it, since it counts no failure. The second is zero whenever every state with units down still carries the load: it measures *availability*, and no amount of capacity in the all-up state changes it except through the states with units down. The hours short split the same way, $\mathrm{LOLE} = \mathrm{LOLE}_{\mathrm{cap}} + \mathrm{LOLE}_{\mathrm{av}}$. Two further splits are read from the same sum: by limiting element, with the indicator $\mathbf{1}[\ell(\mathbf{s}, m) = g]$ inside the sum, whose parts add to the whole because every short state has one limiting group; and by element down, with the factor $(1 - s_e)$, whose parts do not, because a state with two elements down counts toward both.

**The states left out.** The enumeration takes every state with at most two of the ten elements down. Write $r(\mathbf{q})$ for the probability of the rest. A state's shortfall cannot exceed the month's load, so the energy those states could add is bounded,

$$0 \le \mathrm{EUE}_{\ge 3}(\boldsymbol{\theta}) \le r(\mathbf{q}) \sum_{m=1}^{12} h_m\, L_m, \tag{23.4}$$

the residual mass times the year's IT energy.

> **Principle 23.1.** *Expected unserved energy is a sum over the months of the year and the states of the plant's units of each state's probability times its shortfall at the month's load. The all-up state's term is the capacity part, the rest the availability part; the hours short are the same sum counting hours. Each is a function of the posterior, read with its band, and they are reported side by side.*

### 23.1.2 The site as an enumeration of element states (step 0)

The plant is the site of Chapter 9, drawn with both halls and the shared pipe as a state in Figure 15.1: the three chillers C1–C3, the pump, the shared supply pipe with its strainer, the power feed, and the two halls of three cooling units and five racks each, taken as built, without the drawn economizer. The unknown is that of equation (15.8),

$$\boldsymbol{\theta} = (\mathbf{z}_A,\ \mathbf{z}_B,\ \mathbf{z}_{\mathrm{site}},\ \lambda_c,\ \mu_c,\ \lambda_h,\ \mu_h),$$

each hall's reduced state of equation (5.4), the site's eight entries of equation (9.4), and the failure and repair rates of the chillers and of the hall units. The ten elements that can be down are the three chillers, the six hall units and the strainer.

**The physical relations.** In state $\mathbf{s}$ and month $m$ each group's capacity is its binding load in site IT load, kW. The chillers up carry the loop's heat, the IT load plus the halls' other heat $Q_o = 22\,\mathrm{kW}$, at the month's mean wet bulb $T_m$; each hall carries its share of the load with its running units, by the closed form of Chapter 20, equations (20.3) and (20.4), written over the hall's running units in equation (15.4); the pipe carries its flow clean or fouled; and the pump and the feed carry theirs,

$$\begin{aligned} \mathrm{Cap_C}(\mathbf{s},\ m) &= \left[1 - d\,(T_m - 22)\right] \sum_i s_i\, 400\,\mathrm{e}^{\kappa_i} - Q_o, \\ \mathrm{Cap}_h(\mathbf{s}) &= Q_h^{\dagger}(\mathbf{z}_h;\ R_h(\mathbf{s}))/\varphi_h, \qquad h = A,\ B, \\ \mathrm{Cap}_p(\mathbf{s}) &= \left[1 - 0.4\,(1 - s_p)\right] \rho c_p\, \Delta T_w\, \mathrm{e}^{l_k} \sqrt{\Delta P_{\max}} - Q_o, \end{aligned} \tag{23.5}$$

with the pump's and the feed's capacities the binding loads of equation (22.6), the feed's at the month's wet bulb. Here $R_h(\mathbf{s})$ is the set of hall $h$'s units up in state $\mathbf{s}$, $\varphi_A = 0.5455$ and $\varphi_B = 0.4545$ the halls' shares of today's site load, equation (15.5), and $s_p$ the strainer's state; a fouled strainer leaves the pipe a stated 60 % of its flow coefficient, equation (15.6). A hall with no unit running carries nothing. Only the chillers and the feed depend on the month's weather; the halls, at the pinned supply of 14°C, and the pipe depend on the month only through the load they are asked to carry.

**The outage states.** Each chiller and each hall unit fails and is repaired at its kind's rates, so that $q_e = \lambda/(\lambda + \mu)$ of equation (9.6) for its kind. The strainer fouls at a stated 0.5 a year and takes a stated 12 h to clear, so that its $q_p = 6.845 \times 10^{-4}$ by equation (15.7). The pump and the feed are not states of this enumeration: they do not fail in it.

**Modeling choices at the graph.** Each month is taken in steady state at its mean wet bulb for all its hours: this makes each month one check per state, and it fails if the month's hottest hours carry most of its shortfall, a check of Section 23.4.2. The elements fail and are repaired independently: this makes the probability of a state the product of equation (23.1), and it fails under a common cause, a power event that stops several units at once. A state's probability is its long-run fraction of time, so the shortfall of a state is shed for that fraction of each month's hours: this fails for a month whose outages bunch. And the halls are read at the pinned supply of 14°C whatever the chillers do: a chiller plant short of the load would warm the supply and lower both halls at once, the chain of Chapter 21, which this sum does not follow.

### 23.1.3 The plan year's months

The months pin two inputs. The load is the plan year's, the load path of equation (6.1) from today's 660 kW at 15 % a year, read at each month's end as in equation (22.7), from 668 kW in January to 759 kW in December (Table 22.1). The weather is the record year's monthly mean wet bulb $T_m$ of Table 16.5, from 10.17°C in July and August to 19.83°C in January and February. Each month holds a twelfth of the year's hours,

$$L_m = 660\,(1.15)^{m/12}\,\mathrm{kW}, \qquad h_m = \frac{8\,760}{12} = 730\,\mathrm{h}, \qquad \sum_{m=1}^{12} h_m\, L_m = 6\,241\,\mathrm{MWh}, \tag{23.6}$$

the last the year's IT energy, the scale against which the shed energy and the bound (23.4) are read.

**The sheet's version.** The sheet runs the same sum with three substitutions: each chiller at its 400 kW nameplate in every month; each hall's capacity its running units times the sheet's

Table 23.1: The design sheet's expected unserved energy over the plan year: its inputs, its three outputs, and what this chapter reads in their place.

| Quantity [unit] | Sheet [unit] | Read here instead [–] |
|---|---|---|
| Chiller capacity [kW] | 400, nameplate | as-built, $400\,e^{\kappa_i}$ at the month's wet bulb, site's fit |
| Hall unit capacity [kW] | duty: A 210, B 180 | the halls' binding loads at the inlets, equation (20.4) |
| Chiller unavailability [–] | 0.00613 (handbook) | gamma posterior of the chiller outage record |
| Hall-unit unavailability [–] | 0.00273 (handbook) | gamma posterior of the hall-unit outage record |
| Strainer, pipe, pump, feed [–] | not listed | strainer a state; pipe, pump, feed as groups |
| Expected unserved energy [kWh/yr] | 532.2 | Section 23.3.1 |
| Hours short [h/yr] | 13.24 | Section 23.3.2 |
| "Availability" [–] | 0.998488, from the hours | not reported: hours and energy side by side |

duty, less the hall's other heat, over its share; and each kind's unavailability from the handbook's rate and repair time, 0.00613 for a chiller and 0.00273 for a hall unit. It has no strainer, no pipe, no pump and no feed, and no posterior. Table 23.1 sets its inputs and its result beside the quantities this chapter reads instead. Each row is an input or an output; the second column gives the sheet's value, the third what replaces it here.

The sheet's two hall duties settle where its shed energy comes from. With one hall B unit down, two units of 180 kW give 360 kW, more than the hall's heat in every month of the plan year, so the sheet sees no shortfall in hall B; with one hall A unit down, two units of 210 kW give 420 kW, short of hall A's heat in November and December by a few kilowatts, as Chapter 22 found (Section 22.1). The sheet's 532.2 kWh/yr is those two months of hall A with a unit down, and the rarer states with two units down.

### 23.1.4 The planner's questions (step 1)

The planner asks four questions of the year, and step 1 writes each over the plant's unknown, with the evidence on the right of the bar. Table 23.2 lists them. The first column is the question in the planner's words, the drill question first; the second the general form of its category in the notation of Chapter 2; the third its mathematical form for the site; 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 estimation questions about a loss, each the mean and band of a derived quantity: the energy, the hours, and the split of the energy into its capacity and availability parts. The fourth prices the energy at the stated shed value, the loss term a design comparison needs, and Chapter 24 uses it to compare designs. All four are answered by the expected consequence, the enumeration's sum averaged over the posterior. The atlas of Appendix A collects them with the other questions of Volume II.

### 23.1.5 The evidence as rows (step 2)

Step 2 attaches the records of Chapter 15, unchanged. Each hall is read in window 1 of its night record, the seventeen rows of equation (5.11): hall B's those of Chapter 5, hall A's those of Section 15.2.2. The site is read by the operating record of Chapter 9, sixty-two rows on its eight entries. The outage records of the chillers and of the hall units, four years each, give the failure counts with their exposures and the repair times. No row is added: the year asks a new

Table 23.2: The questions asked of the site's year. $\mathbf{y}$ is both halls' night record, the site's operating record, the outage records and the priors; EUE is equation (23.1), LOLE equation (23.2), and the split equation (23.3); $v$ is the stated shed value.

| Planner's question | General form | Mathematical form | Category | Computation | Answered in |
|---|---|---|---|---|---|
| How many kWh of IT load would be shed per year, in expectation, under this design? | $\mathbb{E}[L(\mathbf{z}) \mid \mathbf{y}]$ and its quantiles | $\mathbb{E}[\text{EUE} \mid \mathbf{y}]$, its median and 90 % band | estimation | expected consequence | 23.3.1 |
| How many hours is the site short? | the same, counting hours | $\mathbb{E}[\text{LOLE} \mid \mathbf{y}]$, its median and band | estimation | expected consequence | 23.3.2 |
| Is it capacity or availability? | the expectation of each part of a sum | $\mathbb{E}[\text{EUE}_{\text{cap}}]$ against $\mathbb{E}[\text{EUE}_{\text{av}}]$; by limiting element, element down and month | estimation | expected consequence | 23.3.3 |
| What is it worth? | $\mathcal{R} = \sum_e C_e \Pr(e \mid \mathbf{y})$ | $v\,\mathbb{E}[\text{EUE} \mid \mathbf{y}]$, $v = \text{R\$ } 1\,500/\text{kWh}$ | design | expected consequence | 23.3.4 |

question of the same evidence.

**Which row acts on which entry.** By equation (23.1) the shed energy reads the unknown in two places, and the two are read by different records. The capacities of equation (23.5) read the halls' states and the site's entries, through the night records and the operating record. The probabilities $P(\mathbf{s} \mid \mathbf{q})$ read the rates alone, through the outage records, and the strainer's stated $q_p$. No record reads both, so the capacity of a state and its probability are independent given the evidence. The pattern of each capacity's gradient is that of equation (22.9): each hall's binding load reads its own state alone, the chillers' read the log-ratios and the derate, the pipe's the flow coefficient. The entries that move the shed energy most are therefore the ones that move the halls' binding loads with a unit down, and of those the fractions of the racks no reading sees, B3 in hall B and A5 in hall A, whose posteriors are their survey priors.

**Modeling choices at the rows.** The rows keep the choices of Chapter 15: each reading Gaussian, unbiased and independent given the state; each hall fitted on window 1; the hall units sharing one pair of rates and the chillers another. The repair rates enter here, as they did not in Chapter 22: an unplanned outage lasts its repair time, and $q$ reads it. The shed value of R\$ 1 500 per kWh is a stated planning input, not an estimate.

## 23.2 The posterior of the state (step 3)

All four questions of Table 23.2 are read from one posterior, the product of the three fits of Chapters 5, 9 and 15 with the outage posteriors beside them. No row is added and no fit is repeated; this section derives what the sum needs from that posterior, cites the fits, and states the draws the answers use.

### 23.2.1 What the sum needs from the state

The posterior enters equation (23.1) twice: in the capacities, through the fits, and in the probabilities, through the rates. By the independence of Section 23.1.5 it factors as equation (15.9),

$$p(\boldsymbol{\theta} \mid \mathbf{y}) = p(\mathbf{z}_A \mid \mathbf{y}_A)\, p(\mathbf{z}_B \mid \mathbf{y}_B)\, p(\mathbf{z}_{\text{site}} \mid \mathbf{y}_{\text{op}})\, p(\lambda_c, \mu_c \mid \mathbf{y}_{\text{out}})\, p(\lambda_h, \mu_h \mid \mathbf{y}_{\text{out}}), \tag{23.7}$$

Table 23.3: The outage posteriors over the 4 000 draws the answers read, with the sheet's handbook values beside them, and the strainer's stated unavailability.

| Quantity [unit] | Mean [unit] | 5 % point [unit] | 95 % point [unit] | Sheet [unit] |
|---|---|---|---|---|
| Chiller unavailability $q_c$ [–] | 0.00552 | 0.00271 | 0.00967 | 0.00613 |
| Chiller failure rate $\lambda_c$ [1/yr] | 1.122 | – | – | 1.5 |
| Hall-unit unavailability $q_h$ [–] | 0.00239 | 0.00139 | 0.00375 | 0.00273 |
| Hall-unit failure rate $\lambda_h$ [1/yr] | 0.885 | – | – | 1.0 |
| Strainer unavailability $q_p$ [–] | $6.845 \times 10^{-4}$ | stated | stated | not listed |

now with the repair rates kept. The factorization does not make the shed energy a product of its parts: a state's shortfall is a minimum over the groups, and the sum is over states that read every block at once. The expectation is therefore read on draws of the whole unknown.

### 23.2.2 The fits and the outage posteriors, cited

Each block of the state has the Laplace posterior of its own fit. For block $b$, one of hall A, hall B and the site, the residual vector $\mathbf{r}_b(\mathbf{z}_b)$ stacks that block's rows divided by their standard deviations, its Jacobian is $\mathbf{J}_b = \partial \mathbf{r}_b / \partial \mathbf{z}_b$, the MAP state solves the normal equations (5.12), and the covariance is the inverse of the information at the minimum, equation (5.13),

$$(\mathbf{J}_b^{\mathsf{T}} \mathbf{J}_b)\, \Delta \mathbf{z}_b = -\mathbf{J}_b^{\mathsf{T}} \mathbf{r}_b(\mathbf{z}_b), \qquad \mathbf{\Sigma}_b = (\mathbf{J}_b^{\mathsf{T}} \mathbf{J}_b)^{-1}\big|_{\hat{\mathbf{z}}_b}, \qquad \mathbf{z}_b \mid \mathbf{y}_b \approx \mathcal{N}(\hat{\mathbf{z}}_b,\ \mathbf{\Sigma}_b). \tag{23.8}$$

Hall B's fit is that of Section 5.2.2, hall A's that of Section 15.2.2, and the site's the one fit of Section 9.2.2; the halls' binding loads with all units running are tabulated with their stated truth in Table 15.3. The rates' posteriors are the exact gamma posteriors of equations (9.12) and (15.11), and each draw of a kind's rates gives its $q$ by equation (9.6).

Table 23.3 gives the posteriors of the quantities the probabilities read. Each row is a quantity; the columns give its mean, its 5 % and 95 % points over the draws, and the sheet's handbook value beside it.

The records lower both kinds' failure rates from the handbook's, and both unavailabilities with them: 0.00552 against 0.00613 for a chiller and 0.00239 against 0.00273 for a hall unit. Each band is wide, a factor of about three and a half for a chiller and nearly three for a hall unit, because thirteen chiller failures and twenty-one hall-unit failures fix a rate only to a fraction of itself. The sheet's rates are not what makes it low: they are higher than the posterior's.

### 23.2.3 The draws

The answers read the posterior through $N = 4\,000$ paired draws, those of Chapter 22: the site's first 4 000 draws of Section 9.2.4, which carry the rates of both kinds, each paired by index with a draw of each hall's state from the Laplace normal (23.8) of its fit. By the factorization (23.7) each pair is a draw $\boldsymbol{\theta}_i$ of the whole unknown. In each, the 56 states with at most two of the ten elements down, one with none down, ten with one and forty-five with two, are enumerated at each of the twelve months, and the sums (23.1) and (23.2) and their splits are formed. The reported mean and band are those of the draws, the state route,

$$\mathbb{E}[\mathrm{EUE} \mid \mathbf{y}] \approx \frac{1}{N} \sum_{i=1}^{N} \mathrm{EUE}(\boldsymbol{\theta}_i), \qquad \mathrm{se} = \frac{s_{\mathrm{EUE}}}{\sqrt{N}}, \tag{23.9}$$

with $s_{\mathrm{EUE}}$ the standard deviation of the shed energy over the draws and se the Monte Carlo standard error of the mean; the band is the 5 % and 95 % points of the same draws. The exact

Table 23.4: The shed energy and the hours short over the plan year: the 4 000 paired draws (reported), the exact posterior (the same draws reweighted), and the design sheet. Gaps are exact minus draws, from the printed values.

| Quantity [unit] | Draws, mean [unit] | Exact, mean [unit] | Gap [unit] | Monte Carlo se [unit] | Median [unit] | 90 % band [unit] | Sheet [unit] |
|---|---|---|---|---|---|---|---|
| EUE [kWh/yr] | 16 562.4 | 16 590.8 | 28.4 | 362.3 | 13 008.1 | 6 049 to 29 353 | 532.2 |
| Hours short [h/yr] | 255.7 | 256.2 | 0.5 | – | 124.5 | 74.7 to 817.1 | 13.24 |

posterior is the same draws under the site's weights (5.16), with an effective sample size of 3 983 of the 4 000; the rates are drawn from their exact posteriors already, and the halls' Laplace normals are not reweighted, as in Chapter 22.

At the posterior's draws the residual mass of equation (23.4) has a mean of $3.92 \times 10^{-6}$ and a largest value of $3.47 \times 10^{-5}$. Against the year's 6 241 MWh the bound is 24.47 kWh/yr at the mean residual mass and 216.44 kWh/yr at the largest. Both are small against the answers below and the first is well inside their Monte Carlo error.

## 23.3 The answers (step 4)

Every answer is read from the paired draws of Section 23.2.3. The capacity of each state is joint over the six groups by construction, the minimum of equation (23.1), so every reported value is joint; the exact posterior stands beside it with the gap and the Monte Carlo standard error, and the sheet beside both. No delta form is reported: the shed energy is a sum of terms $\max(L_m - \text{Cap}, 0)$, each with a kink where a state's capacity meets the month's load, and its all-up term is zero over most of the posterior and positive in a tail, so no single gradient carries it.

### 23.3.1 The expected unserved energy

**Planner's question.** How many kWh of IT load would be shed per year, in expectation, under this design? The question asks for the posterior mean of the year's shed energy at the site as built, with its band.

**General form.** An expected consequence is the posterior expectation of a loss, equation (2.23), here with no action taken and no certain cost,

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

read with the quantiles of the loss's distribution, equation (2.11), as its band.

**Specialized form.** The question is equation (2.23) with the unknown $\boldsymbol{\theta}$ in place of $\mathbf{z}$ and the loss $\text{EUE}(\boldsymbol{\theta})$ of equation (23.1), which is the first line of that equation, read by the draws (23.9). Its median and 90 % band are the 50 %, 5 % and 95 % points of $\text{EUE}(\boldsymbol{\theta}_i)$ over the draws.

**Building blocks.** The enumeration of the 56 states (23.1); the capacities (23.5) with the halls' closed form (20.4); the months' inputs (23.6); the paired draws and the site's weights.

**The routes.** Table 23.4 sets the answer beside its routes. Each row is a quantity; the columns give the mean by the draws, the mean by the exact posterior and the gap between them, the Monte Carlo standard error, the median and 90 % band of the draws, and the sheet's value.

The exact posterior moves the mean by 28.4 kWh/yr, under a tenth of the Monte Carlo standard error of 362.3: the Laplace normals lose nothing that matters here. The distribution is skewed to the right, its mean above its median by about a quarter and its 95 % point more than twice the median, because a tail of the posterior puts the plant short with nothing down (Section 23.3.3).

**The drill question.** By equation (23.1), over the plan year at the site as built,

$$\mathbb{E}[\text{EUE} \mid \mathbf{y}] = 16\,562\,\text{kWh/yr}, \qquad \text{EUE} \in [6\,049,\ 29\,353]\,\text{kWh/yr with 90 \% probability}, \tag{23.10}$$

with the exact 16 591 kWh/yr beside it and a median of 13 008 kWh/yr. That is 31 times the sheet's 532.2 kWh/yr, the ratio of the two means. It is $2.65 \times 10^{-3}$ of the year's IT energy of 6 241 MWh. The probability that the year sheds more than 10 MWh is 0.7225. Even the 5 % point of the posterior, 6 049 kWh/yr, is more than eleven times the sheet's figure.

The sheet is low by a factor of 31, and not because of its rates, which are higher than the posterior's (Table 23.3). It is low because of hall B. Its duties of 180 kW make hall B N+1: two units carry the hall's heat, so a hall B unit down sheds nothing. At the inlets hall B is N: Chapter 20 found that with its largest unit out the hall holds today's 300 kW with probability 0.0125 and the design load in none of its draws. Every state with a hall B unit down sheds load in every month of the plan year, and the sheet counts none of them.

### 23.3.2 The hours short beside it

**Planner's question.** How many hours is the site short? The question asks for the posterior mean of the hours in the plan year during which the site cannot carry its load, the loss-of-load expectation, beside the energy.

**General form.** The expected consequence of equation (2.23) with the hours short as the loss.

**Specialized form.** The question is equation (2.23) with the loss $\text{LOLE}(\boldsymbol{\theta})$ of equation (23.2), read by the same draws as the energy.

**Building blocks.** Those of Section 23.3.1, with the indicator of a shortfall in place of its depth.

**The hours.** By Table 23.4 the site is short for

$$\mathbb{E}[\text{LOLE} \mid \mathbf{y}] = 255.7\,\text{h/yr}, \qquad \text{LOLE} \in [74.7,\ 817.1]\,\text{h/yr with 90 \% probability}, \tag{23.11}$$

with the exact 256.2 h beside it and a median of 124.5 h. The mean shortfall when short, the ratio of the two means, is $16\,562.4/255.7 = 64.8$ kW. The hours band is wider than the energy's in proportion, its 95 % point more than six times its median: the tail of the posterior that puts the plant short with nothing down adds whole months of shallow shortfall, many hours and comparatively little energy.

**Hours are not an availability.** The sheet turns its 13.24 h into an "availability" of 0.998488, one less the hours short over the year's hours. The same arithmetic on the posterior's 255.7 h would print 0.971, a number that reads as a property of the units' failures and repairs and hides two things the planner needs. It hides the depth: 255.7 h at 64.8 kW is the energy of equation (23.10), and the same hours at a few kilowatts would be a nuisance. And it hides the kind: half of these hours, as the next section shows, come from the plant short with nothing down, a capacity shortfall that no repair would end. The chapter reports the energy and the hours side by side and converts neither into the other, nor either into an availability.

Table 23.5: The shed energy and the hours short of the plan year, split into the all-up state's term (capacity) and the states with units down (availability), equation (23.3); 4 000 paired draws.

| Part [–] | Mean EUE [kWh/yr] | Share [–] | 5 % point [kWh/yr] | 95 % point [kWh/yr] | Hours short [h/yr] |
|---|---|---|---|---|---|
| All-up state (capacity) | 2 728.4 | 0.165 | 0.0 | 918.4 | 127.1 |
| Units down (availability) | 13 834.0 | 0.835 | – | – | 128.6 |
| Whole year | 16 562.4 | 1.000 | 6 049.2 | 29 352.5 | 255.7 |

### 23.3.3 Capacity against availability

**Planner's question.** Is it capacity or availability? The question asks how much of the shed energy and of the hours short come from the plant short with nothing down, and how much from states with units down; and, within each, which element limits, which element is down, and in which months.

**General form.** The expectation of each part of a sum is the part of the expectation, by the linearity of equation (2.10),

$$\mathbb{E}\Big[\textstyle\sum_j f_j(\mathbf{z}) \mid \mathbf{y}\Big] = \textstyle\sum_j \mathbb{E}\big[f_j(\mathbf{z}) \mid \mathbf{y}\big],$$

so each part's mean is a share of the whole mean, read from the same draws.

**Specialized form.** The question is equation (2.10), the mean, taken part by part with the two parts of equation (23.3), the capacity term $\mathrm{EUE}_{\mathrm{cap}}$ and the availability term $\mathrm{EUE}_{\mathrm{av}}$, and their counterparts in hours; then with the parts by limiting element and by month, which add to the whole; and with the parts by element down, which do not.

**Building blocks.** The sums of Section 23.3.1, kept term by term; the limiting element $\ell(\mathbf{s}, m)$ of each state, equation (15.1).

**The two terms.** Table 23.5 gives the split. Each row is a part of the sum; the columns give its mean energy, its share of the mean energy, the 5 % and 95 % points of its energy over the draws, and its mean hours short.

The hours split evenly and the energy does not. The states with units down carry 0.835 of the energy, 13 834 kWh/yr, in 128.6 h; the all-up state carries 0.165, 2 728 kWh/yr, in 127.1 h. A unit down leaves a hall far short of its load, for a short time; the plant short with nothing down is short by little, for long. The all-up term also sits in a tail: its mean of 2 728 kWh/yr lies far above its 95 % point of 918.4 kWh/yr, and its 5 % point is zero. In most draws the all-up plant carries every month's load; in a few, those in which hall B's binding load with all units running is low, it is short for months at a time. The 5 % point of that binding load, 345.89 kW of hall load (Table 15.3), is 345.89/0.4545 = 761.0 kW of site load, just above December's 759 kW: at the foot of hall B's band the plant runs out of capacity within the plan year with nothing broken.

**By group of states.** Table 23.6 gives the enumeration group by group, each group the states with the same kinds of element down. Each row is a group; the columns give its number of states, its probability summed over its states (the mean over the draws), its share of the mean energy, its probability-weighted shortfall averaged over the months, and its shortfall given

Table 23.6: The shed energy of the plan year by group of states, each group the states with the same kinds of element down; probabilities summed over the group and averaged over the 4 000 paired draws, shortfalls averaged over the months. Each state lies in one group.

| Group, elements down [–] | States [–] | Probability [–] | EUE share [–] | Probability × shortfall [kW] | Shortfall given the group [kW] |
|---|---|---|---|---|---|
| None (all up) | 1 | $9.689 \times 10^{-1}$ | 0.1647 | 0.3115 | 0.3 |
| A chiller | 3 | $1.611 \times 10^{-2}$ | 0.0088 | 0.0167 | 1.0 |
| A hall A unit | 3 | $6.943 \times 10^{-3}$ | 0.2576 | 0.4871 | 70.2 |
| A hall B unit | 3 | $6.943 \times 10^{-3}$ | 0.5092 | 0.9628 | 138.7 |
| The strainer | 1 | $6.636 \times 10^{-4}$ | 0.0159 | 0.0300 | 45.3 |
| Two chillers | 3 | $1.036 \times 10^{-4}$ | 0.0179 | 0.0338 | 326.5 |
| Chiller, hall A unit | 9 | $1.155 \times 10^{-4}$ | 0.0043 | 0.0081 | 69.8 |
| Chiller, hall B unit | 9 | $1.155 \times 10^{-4}$ | 0.0085 | 0.0161 | 139.0 |
| Chiller, strainer | 3 | $1.104 \times 10^{-5}$ | 0.0003 | 0.0005 | 45.2 |
| Two hall A units | 3 | $1.818 \times 10^{-5}$ | 0.0038 | 0.0072 | 398.7 |
| Hall A unit, hall B unit | 9 | $5.453 \times 10^{-5}$ | 0.0042 | 0.0079 | 145.7 |
| Hall A unit, strainer | 3 | $4.755 \times 10^{-6}$ | 0.0002 | 0.0004 | 80.6 |
| Two hall B units | 3 | $1.818 \times 10^{-5}$ | 0.0042 | 0.0079 | 436.2 |
| Hall B unit, strainer | 3 | $4.755 \times 10^{-6}$ | 0.0003 | 0.0007 | 139.1 |

the group, the mean depth when one of its states holds. Each state lies in one group, so the probabilities and the shares add to one.

The table shows why the hours and the energy split differently. The all-up state holds 0.969 of the probability and, given it, the site is short by 0.3 kW on average over the months and the draws: it carries 0.165 of the energy. The three states with a hall B unit down hold $6.9 \times 10^{-3}$ of the probability, about a hundred and forty times less, and given them the site is short by 138.7 kW: they carry 0.509 of the energy, about three times the all-up state's share. A hall A unit down is as likely and half as deep, 70.2 kW, and carries 0.258. A chiller down is the likeliest group with anything down, $1.611 \times 10^{-2}$, and nearly harmless, 1.0 kW given it; two chillers down are short by 326.5 kW but rare enough to carry 0.018. The hall B group's 0.509 is the share of the states with only a hall B unit down; the 0.530 of the hall B units read by element down below adds the two-down states that include one.

The availability term's hours follow from the same table. The states with a hall A unit down and those with a hall B unit down hold $6.943 \times 10^{-3}$ of the probability each, together $1.39 \times 10^{-2}$ of the year, about 122 h. With hall B N at its inlets and hall A nearly so (Chapter 22 found a hall A unit's service at best 0.51 likely to stay inside), those states put the site short in nearly every month, which accounts for most of the 128.6 h of Table 23.5; the rest comes from the strainer and from the chillers late in the year.

**By limiting element, element down and month.** Table 23.7 gives the two finer splits. Each row is an element; the first column of numbers is its share of the mean energy from the states it limits, the second its share from the states with it down. The same shares are drawn in Figure 23.2: the upper panel stacks the energy by limiting element and the two terms of the split, with the hours split beside them as a bar of its own; the lower panel gives one bar per element down. The enumeration behind both is sketched in Figure 23.1: the all-up state on the left, set apart as the capacity term, and the states with one or two elements down fanned out to the right by the number down, each group with its probability, its share of the energy and its shortfall given the group from Table 23.6.

Hall B limits 0.688 of the shed energy and hall A 0.272; the chillers limit 0.024 and the pipe 0.016, and the pump and the feed none. Read by element down, the three hall B units carry

Table 23.7: The shed energy of the plan year by limiting element (shares add to one) and by element down (shares from the states with that element down; a state with two elements down counts toward both, so they do not add to one); 4 000 paired draws.

| Element [–] | Share, as limiting element [–] | Share, from states with it down [–] |
|---|---|---|
| Hall B | 0.688 | B1 0.180, B2 0.188, B3 0.162 |
| Hall A | 0.272 | A1 0.086, A2 0.115, A3 0.073 |
| Chillers | 0.024 | C1 0.018, C2 0.022, C3 0.017 |
| Pipe (strainer) | 0.016 | 0.017 |
| Pump | 0.000 | – |
| Feed | 0.000 | – |

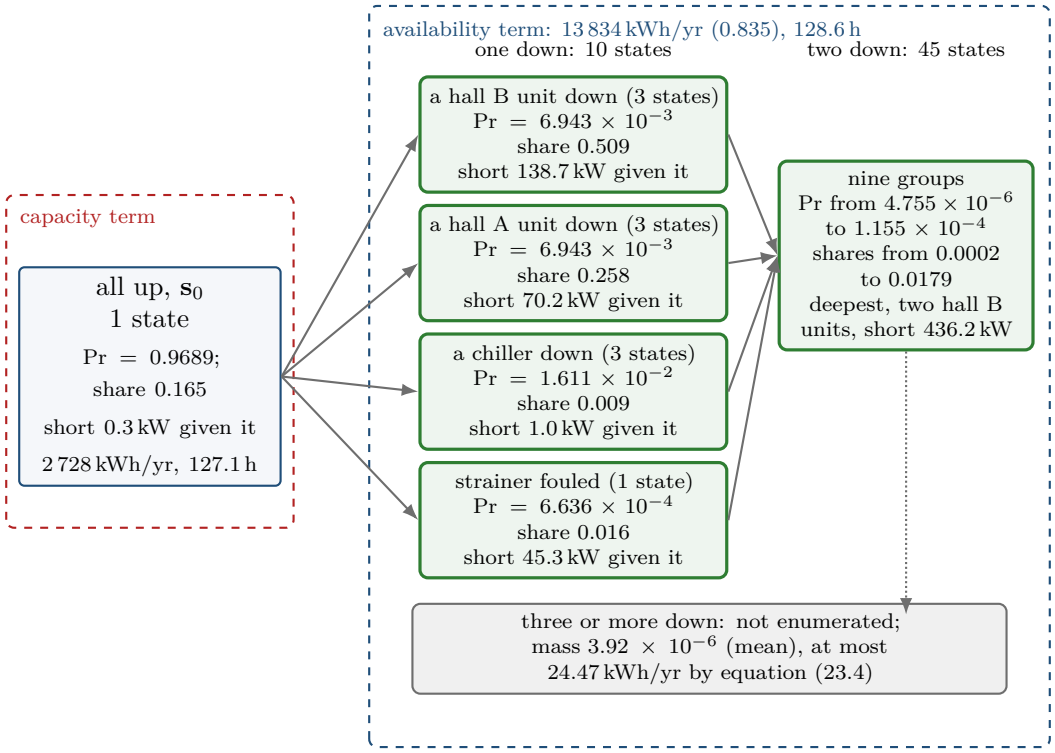



Figure 23.1: The enumeration of the site's ten elements over the plan year, fanned by the number of elements down. Left, in the red frame: the all-up state, the capacity term of equation (23.3), with its probability, share of the mean energy, shortfall given it, mean shed energy and hours short. Right, in the blue frame: the 55 states with one or two elements down, the availability term, grouped by the kinds of element down; each one-down group carries its probability summed over its states, its share of the mean energy and its shortfall given the group, and the nine two-down groups are summarized by their ranges (Table 23.6). Each state lies in one group. Bottom: the states left out, with the bound on their energy.

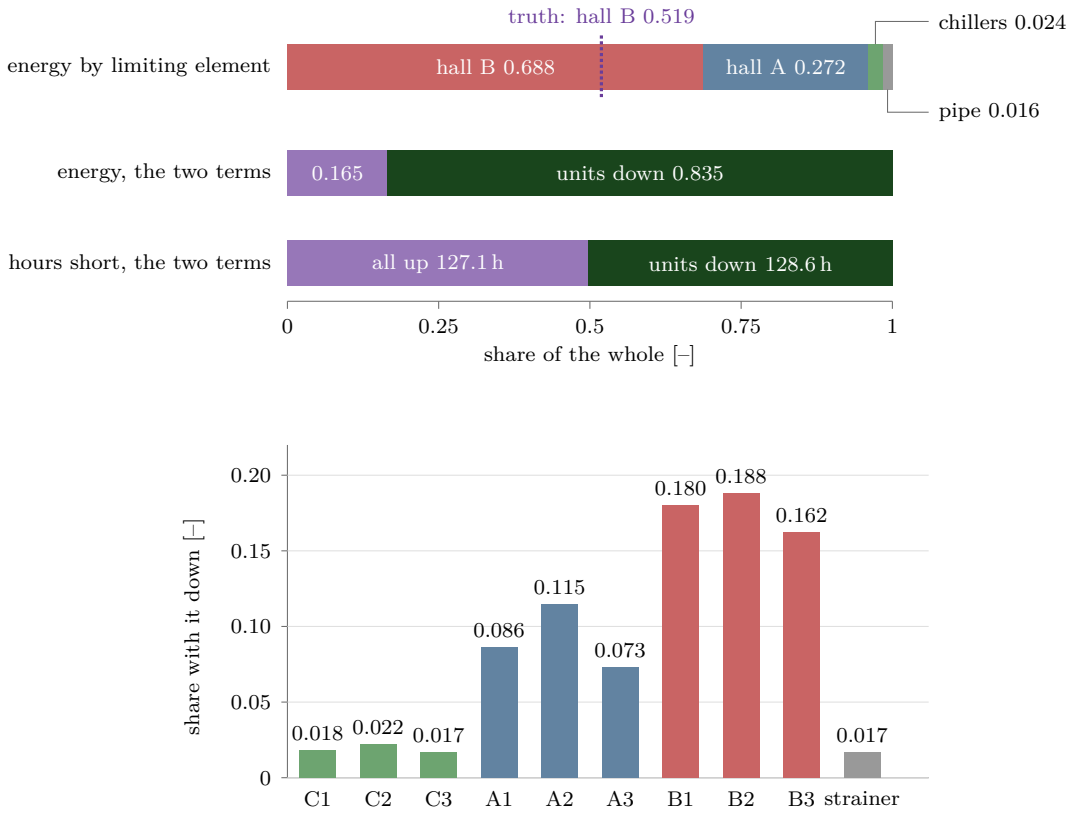



Figure 23.2: The shed energy of the plan year split three ways, from Tables 23.5 and 23.7. Upper panel, as shares of the whole: the energy by limiting element (hall B red, hall A blue, chillers green, pipe gray), with the stated truth's hall B share, 9 423.8 of 18 171.7 kWh or 0.519, dotted; the energy in its two terms, the all-up state (capacity, purple) and the states with units down (availability, dark green); and the hours short in the same two terms, 127.1 and 128.6 h of 255.7 h. Lower panel: the share of the mean energy from the states with each element down, chillers green, hall A blue, hall B red, the strainer gray; a state with two elements down counts toward both bars.

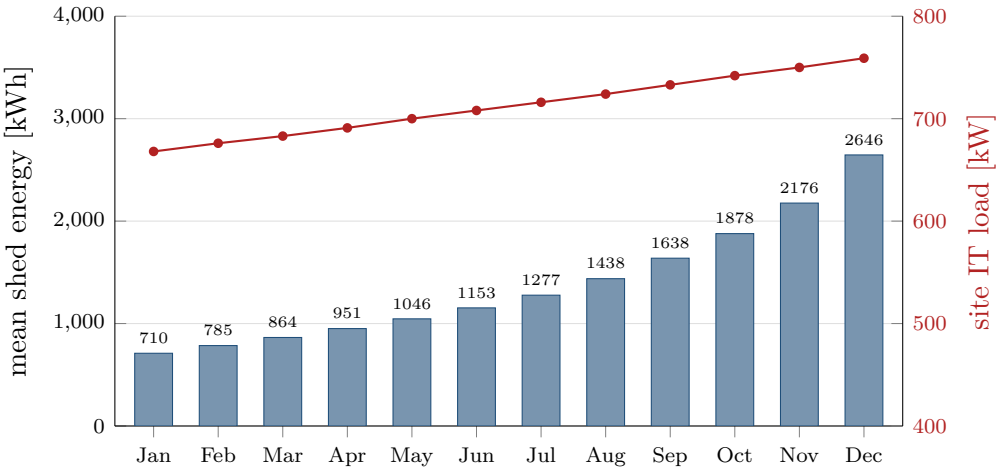



Figure 23.3: The mean shed energy of each month of the plan year (blue bars, left axis, kWh in the month), over the 4 000 paired draws, and the month's site IT load at its end (red line, right axis), 660 kW grown at 15 % a year (Table 22.1). October to December carry 0.404 of the year's 16 562 kWh.

$0.180 + 0.188 + 0.162 = 0.530$ of the energy, the three hall A units $0.086 + 0.115 + 0.073 = 0.274$, the three chillers 0.057 and the strainer 0.017. The two readings agree on the order and differ on hall B, whose share as the limiting element is larger than its units' share down by $0.688 - 0.530 = 0.158$, close to the all-up term's 0.165: hall B limits the plant short with nothing down, and that term carries no unit's share. Among hall B's units B3 carries the least, the unit the fit puts lowest (Chapter 20), whose loss costs the hall least. The ten shares by element down sum to 0.878, above the availability term's 0.835, the two-down states counted twice.

**The chillers.** The chillers limit 0.024 of the shed energy, 395.0 kWh/yr. They are the subject of Chapter 9's N+1, its fourth chiller weighed at the design day. Over the year the chillers meet a lower wet bulb than the design day's in every month, and Chapter 22 found that with one chiller out the other two carry the load until late in the plan year, when the load is highest and the wet bulb climbs back; a chiller down is short only in those months, and only for its repair time. The year's shed energy is an air-side availability problem today: the hall units, not the chillers, and units down, not the plant with nothing down.

**By month.** Figure 23.3 gives the shed energy by month, as bars on the left axis, with the month's site IT load as a line on the right axis. The energy grows every month, from 710 kWh in January to 2 646 kWh in December, and October to December carry 0.404 of the year. The weather is nearly absent from the shape: January and February are the warmest months and among the cheapest, because the halls, which carry almost all of the energy, do not see the wet bulb at the pinned supply, while the load they must carry grows by 15 % a year.

### 23.3.4 The value of the shed energy

**Planner's question.** What is it worth? The question asks for the shed energy priced at the stated shed value, the loss term against which a design that sheds less is compared.

**General form.** The expected consequence of equation (2.23), risk as probability times consequence summed over the events, $\mathcal{R} = \sum_e C_e \Pr(e \mid \mathbf{y})$, with no certain cost while no action is taken.

**Specialized form.** The question is equation (2.23) with each state and month an event, its consequence the energy it sheds times the shed value $v$, so that

$$\mathcal{R}_{\text{shed}} = v\,\mathbb{E}[\text{EUE} \mid \mathbf{y}], \qquad v = \text{R\$ } 1\,500/\text{kWh}, \tag{23.12}$$

with the band $v$ times the band of the energy, since $v$ is stated.

**Building blocks.** The shed energy of Section 23.3.1; the stated shed value.

**The value.** At the stated R$ 1 500 per kWh, the shed energy of the plan year is worth

$$\mathcal{R}_{\text{shed}} = \text{R\$ } 24.8 \text{ million a year}, \qquad [\text{R\$ } 9.1,\ \text{R\$ } 43.9] \text{ million with } 90\,\% \text{ probability}, \tag{23.13}$$

R$ 24 843 596 at the mean, from R$ 9 075 104 at the 5 % point to R$ 43 850 460 at the 95 %. The sheet's 532.2 kWh/yr at the same price is under a million. The value is a loss term, not a decision: which design removes how much of it, and at what price, is the comparison of Chapter 24.

## 23.4 The checks (step 5)

The answers rest on draws of three fits and two outage posteriors, on an enumeration cut at two elements down, on each month at its mean wet bulb, and on records generated from a stated truth; each is tested here.

### 23.4.1 The routes against each other

Three checks bear on the computation. First, the exact posterior against the draws: 16 590.8 against 16 562.4 kWh/yr, a gap of 28.4, under a tenth of the Monte Carlo standard error of 362.3; and 256.2 against 255.7 h for the hours short. Second, the states left out: by the bound (23.4) they add at most 24.47 kWh/yr at the mean residual mass and 216.44 at the largest, the first under a tenth and the second under two thirds of the Monte Carlo standard error. Third, the structure of the sums: the shares by limiting element add to one, the capacity and availability terms add to the whole in energy and in hours, and the shares by element down exceed the availability term by the two-down states, as equation (23.3) requires.

### 23.4.2 The weather and the load's growth

Each month is taken at its mean wet bulb for all its hours. Taking every hour of every month at the month's hottest days instead, its mean wet bulb plus 2 K as in Chapter 22, gives 16 822.4 kWh/yr, 260 kWh/yr or about 1.6 % above the reported mean, on the same draws. This bounds the weather's part from above: a month's hours lie between its mean and its hottest days. The weather's part is small because the halls, which carry nearly all of the energy, are read at the pinned supply and do not see the wet bulb.

The load's growth is not small. With today's 660 kW held flat through the year, the shed energy is 7 711.8 kWh/yr, with an all-up part of 2.8 kWh/yr. The year's growth from 668 to 759 kW doubles the shed energy and creates nearly all of the capacity term: at today's load the plant with nothing down is almost never short.

### 23.4.3 The answers against the stated truth

The records are generated from a stated truth, and the last checks compare the answers with it. At the true plant and the true rates the plan year sheds 18 171.7 kWh/yr in 127.8 h, with an

Table 23.8: The checks of the shed-energy answers. Reported values by 4 000 paired draws; the truth is the stated truth of the record, at the true plant and the true outage rates.

| Quantity [unit] | Reported [unit] | Beside [unit] | Truth [unit] | Reading |
|---|---|---|---|---|
| EUE [kWh/yr] | 16 562.4 | exact 16 590.8 | 18 171.7 | gap 28.4, se 362.3; truth inside the band, $P(\leq) = 0.7845$ |
| Hours short [h/yr] | 255.7 | exact 256.2 | 127.8 | truth's all from units down |
| All-up part [kWh/yr] | 2 728.4 | 95 % point 918.4 | 0.0 | a tail of the posterior |
| Units-down hours [h/yr] | 128.6 | – | 127.8 | agrees |
| Hall A, limiting [kWh/yr] | share 0.272 | – | 8 167.1 | posterior gives hall A too much room |
| Hall B, limiting [kWh/yr] | share 0.688 | – | 9 423.8 | largest in both |
| States left out [kWh/yr] | $\leq$ 24.47 | $\leq$ 216.44 at the largest mass | – | inside the se |
| Hottest days [kWh/yr] | 16 562.4 | 16 822.4 | – | weather's part small |
| Load held at 660 kW [kWh/yr] | 16 562.4 | 7 711.8, all-up 2.8 | – | growth doubles it |
| Draws | 4 000 paired: the site's first 4 000 of Chapter 9 (effective sample size 3 983), each hall's 4 000 from its Laplace normal, the rates from their gamma posteriors | | | |

all-up part of zero: the true plant with nothing down carries every month's load, since hall B's true binding load with all units running is 833.6 kW of site load (Table 15.3). The truth lies inside the posterior's band, at $P(\text{EUE} \leq \text{truth}) = 0.7845$. Its hours, all from states with units down, sit beside the posterior's 128.6 h of the availability term. By limiting element the truth's energy is 9 423.8 kWh/yr in hall B, 8 167.1 in hall A, 315.2 in the chillers and 265.6 in the pipe, none in the pump or the feed. Hall A carries nearly half of the truth's energy, more than the posterior's 0.272 share: hall A's fit places its binding load above the truth, at the unmeasured rack A5 whose survey reads 0.45 against a true 0.51 (Section 15.2.2), so the posterior gives hall A more room than it has. Table 23.8 collects the checks: each row is a quantity, the columns its reported value, the route or reference beside it, the truth and the reading.

The answers hold where they can be checked. The one place they part from the truth is the capacity term, which the truth does not have and the posterior carries in a tail: the posterior cannot rule out a hall B whose binding load with all units running falls inside the plan year, and the truth is not such a hall. The energy from states with units down, the part that matters most, the posterior gets about right in hours and somewhat low in energy, because it places hall A higher than it is.

## 23.5 The report (step 6)

The site sheds an expected 16.6 MWh of IT load in the plan year, with a band from 6.0 to 29.4 MWh, over about 256 h. That is 31 times the sheet's figure. Of the energy, 84 % comes from states with a unit down, most of them a hall unit: hall B is N on its air side, and hall A nearly so. The chillers carry 2 %.

### 23.5.1 The answers and their bands

Table 23.9 gives the answers in the order of Table 23.2. Each row is a question; the second column gives the reported answer, and the third its band or the routes beside it.

Table 23.9: The report on the site's shed energy over the plan year, at 15 % a year from 660 kW, each month at its mean wet bulb, the halls at the supply of 14°C. Reported values by 4 000 paired draws.

| Question | Answer | Band or routes beside |
|---|---|---|
| How many kWh of IT load would be shed per year, in expectation, under this design? | 16 562 kWh/yr, 31 times the sheet's 532.2; $2.65 \times 10^{-3}$ of the year's IT energy | 90 % band 6 049 to 29 353; exact 16 591; se 362; $P(> 10\,\text{MWh}) = 0.7225$ |
| How many hours is the site short? | 255.7 h/yr, beside the energy; mean shortfall 64.8 kW | band 74.7 to 817.1 h; exact 256.2; the sheet's 13.24 |
| Is it capacity or availability? | availability: units down 0.835 of the energy, 128.6 h; all up 0.165, 127.1 h | by limiting element hall B 0.688, hall A 0.272, chillers 0.024, pipe 0.016; Oct–Dec 0.404 |
| What is it worth? | R$ 24.8 million a year at R$ 1 500/kWh | band R$ 9.1 to 43.9 million |

### 23.5.2 The design-review entry

The design review's entry does not accept the sheet's risk line. The sheet prices the site's shed energy at 532.2 kWh/yr with nameplate chillers, duty-rated hall units and handbook rates; on the site's posterior, enumerated over the plan year at each month's load and weather and joint over every group in every state, the expected shed energy is 16 562 kWh/yr, with a band of 6 049 to 29 353, beside 255.7 h short. The sheet's rates are not the cause; its duties are. They make hall B N+1, and at the inlets it is N.

The entry records the shed energy and the hours short side by side, and records no availability made from the hours. It records which part is which: 13 834 kWh/yr, 0.835 of the energy, from states with units down, an availability problem of the halls' air side; 2 728 kWh/yr, 0.165, from the plant short with nothing down, a capacity problem that sits in the tail of hall B's binding load and grows with the load. It records that the chillers, whose redundancy Chapter 9 weighed, carry 395.0 kWh/yr of it. And it records the value at the stated shed price, R$ 24.8 million a year, as the loss term against which Chapter 24 compares the remedies.

**Common mistake**

Turning hours short into an availability. Dividing the hours short by the hours of the year and subtracting from one gives a number that reads as a property of the units' outages, 0.971 here, and it folds into that number the hours the plant is short with nothing down, which no repair ends, while dropping the depth of each shortfall. The hours short and the shed energy are two readers of one enumeration; report both, side by side, and say which part of each is capacity and which is availability.

## 23.6 The next measurement (step 7)

The report leaves a shed energy whose band spans nearly a factor of five. Step 7 asks which entry of the state carries that band, and which measurement on the existing plant would narrow it.

Table 23.10: The first-order shares of the variance of the plan year's shed energy, equation (23.14), by 20 equal-count bins of the 4 000 paired draws; the eight largest of 26 entries.

| Entry [–] | First-order share [–] |
|---|---|
| Hall B, rack B3's fraction $\beta_{B3}$ | 0.319 |
| Hall-unit unavailability $q_h$ | 0.034 |
| Hall A, rack A5's fraction $\beta_{A5}$ | 0.025 |
| Site, chiller C1's log-ratio $\kappa_{C1}$ | 0.009 |
| Hall B, unit B3's log-derate $g_{B3}$ | 0.008 |
| Hall A, rack A3's fraction $\beta_{A3}$ | 0.007 |
| Site, chiller C3's log-ratio $\kappa_{C3}$ | 0.006 |
| Hall A, rack A2's fraction $\beta_{A2}$ | 0.006 |
| Sum over all 26 entries | 0.496 |

### 23.6.1 The band of the shed energy by entry

The first-order share of an entry is the part of the variance that knowing it would remove on average, equation (17.2),

$$S_i = \frac{\mathrm{Var}(\mathbb{E}[\mathrm{EUE} \mid z_i])}{\mathrm{Var}(\mathrm{EUE})}, \tag{23.14}$$

read by sorting the draws on $z_i$ into 20 bins of equal count, equation (17.3), over the 26 entries: each hall's log-derates and fractions, the site's eight entries, and the two kinds' unavailabilities. Table 23.10 gives the eight largest. Each row is an entry; the column gives its first-order share.

One entry carries most of what any single entry carries: rack B3's recirculation fraction, 0.319 of the variance, the unmeasured rack of hall B whose posterior is its survey prior. It moves hall B's binding load with every unit running and with each unit down, so it moves both terms of the split at once, and it alone decides whether a draw falls in the tail where the plant is short with nothing down. The hall units' unavailability follows at 0.034 and rack A5's fraction at 0.025; every other entry carries under 0.01. The shares over all 26 entries sum to 0.496. The other half of the variance is carried by interactions: the shed energy is a product of a state's probability and its shortfall, summed over states, and an entry that sets the depth and an entry that sets the probability move it together, which no single entry's share counts. Knowing $\beta_{B3}$ exactly would remove about a third of the variance, not all of it.

### 23.6.2 The recommendation

Fit the inlet sensor at rack B3. It is the measurement Chapter 20 ranked first for hall B's binding load with a unit out (Section 20.6.1), which takes the posterior spread of $\beta_{B3}$ from 0.0800 to 0.0320, and it bears on the one entry that carries a third of the shed energy's variance. It will not change the expected shed energy on average: averaged over the reading's outcomes the posterior mean stays where it is, by the law of total expectation. What it will change is the width of the band, and with it how firmly the capacity term can be ruled in or out. Chapter 18 ranks such measurements by the width they remove from the quantity that decides; whether any measurement would change the design itself is the question of Chapter 24.

### 23.6.3 What the analysis showed

The sheet's risk line is low by a factor of 31. Read as an enumeration of the site's ten elements over the plan year, month by month, each state's capacity joint over the six groups and each state's probability from the outage posteriors, the site sheds an expected 16 562 kWh/yr, with a band of 6 049 to 29 353, beside 255.7 h short. The sheet's rates are higher than the posterior's;

its error is its duties, which make hall B N+1 where the inlets make it N, so that every state with a hall B unit down sheds load and the sheet counts none. The shed energy is an air-side availability problem today: 0.835 of it comes from states with units down, hall B's above all, and the chillers, the subject of Chapter 9's N+1, carry 0.024. The hours split evenly between the all-up state and the states with units down while the energy does not, which is why the two are reported side by side: a unit down is deep and short, the plant short with nothing down shallow and long. The capacity term sits in a tail of the posterior, set by hall B's unseen rack B3, and the load's growth through the year creates it; with the load held at today's 660 kW it nearly vanishes and the shed energy halves.

### 23.6.4 Limits of the model

**The month at its mean wet bulb.** Each month is taken at its mean wet bulb for all its hours. The hottest-days check of Section 23.4.2 bounds what this misses at about 260 kWh/yr; a site whose halls saw the weather would need the hours of each month resolved.

**The month's end.** Each month carries its end-of-month load, the largest it sees, so the sum leans high by part of a month's growth.

**Independent outages, no common cause.** The probability of a state is the product of its elements' unavailabilities. A common cause, a power event or a control fault that stops several units at once, would add mass to the states with two or more down, which the enumeration weights at the product of two small numbers.

**The supply pinned, no cascade.** The halls are read at the supply of 14°C whatever the chillers do. A chiller short of the load would warm the supply and pull both halls down at once, the chain of Chapter 21; the chillers' share here is small, so the energy that chain would add is small too, but it is not in this sum.

**The strainer and the shares.** The strainer's fouling rate, clearing time and fouled fraction are stated, not estimated, and the halls grow in today's proportions; both are as in Chapter 15.

**The shed value.** The R$ 1 500 per kWh is a stated planning input; the value of equation (23.13) scales with it.

**The constructed record.** The night records, the operating record and the outage records are generated from a stated truth with stated errors, not taken from a site, and the comparison with that truth in Section 23.4.3 is a check of the method on its own model, not of the model against a real plant.

## 23.7 What to retain

**Definition 23.1** (Expected unserved energy by enumeration)**.** For a plant whose elements are up or down independently with unavailabilities $q_e = \lambda_e/(\lambda_e + \mu_e)$, the expected unserved energy of a year of months $m$ with hours $h_m$ and loads $L_m$ is $\mathbb{E}[\mathrm{EUE} \mid \mathbf{y}]$ with $\mathrm{EUE} = \sum_m h_m \sum_{\mathbf{s}} P(\mathbf{s} \mid \mathbf{q})\, \max(L_m - \mathrm{Cap}(\mathbf{s}, m), 0)$, each state's capacity the smallest of its groups'. The all-up state's term is the capacity part, the states with units down the availability part. The hours short are the same sum with the indicator of a shortfall, reported beside the energy and never converted into an availability.

Four conclusions carry beyond the site. Expected unserved energy is an enumeration read twice from the posterior, through the capacity of each state and through its probability, and the band of the answer comes from both. A sheet that rates units by duty can miss a whole class of states that shed load: a unit the sheet counts as redundant may not be redundant at the limit the plant is held to. The sum splits by construction into a capacity term and an availability term, and the split says what kind of remedy can reduce it: units for availability, capacity for the all-up term. And the hours short and the energy are two readers of one sum, which can split differently, as here, where the hours divide evenly and the energy five to one; neither is an availability, and neither stands for the other.

Table 23.11 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 23.11: The terms of expected unserved energy and where each is defined. Units for the site.

| Term | Expression | Units | Meaning | Defined in |
|---|---|---|---|---|
| Expected unserved energy | $\sum_m h_m \sum_{\mathbf{s}} P(\mathbf{s}) \max(L_m - \text{Cap}, 0)$, averaged over the posterior | kWh/yr | IT energy shed in a year | 23.1.1 |
| Hours short (LOLE) | $\sum_m h_m \sum_{\mathbf{s}} P(\mathbf{s})\, \mathbf{1}[\text{Cap} < L_m]$ | h/yr | time short, any depth | 23.1.1 |
| Capacity term | the all-up state's part | kWh/yr | short with nothing down | 23.1.1 |
| Availability term | the part from states with units down | kWh/yr | short because a unit is down | 23.1.1 |
| Residual bound | $r(\mathbf{q}) \sum_m h_m L_m$ | kWh/yr | states left out | 23.1.1 |
| Shed value | $v\, \mathbb{E}[\text{EUE} \mid \mathbf{y}]$ | R$/yr | the loss term | 23.3.4 |
| *The site* | | | | |
| State probability | $\prod_e q_e^{1-s_e}(1-q_e)^{s_e}$ | – | ten elements, at most two down | 23.1.2 |
| State capacity | $\min_g \text{Cap}_g(\mathbf{s}, m)$, equation (23.5) | kW | six groups, month's wet bulb | 23.1.2 |
| Months | $L_m = 660(1.15)^{m/12}$, $h_m = 730$ | kW, h | plan year | 23.1.3 |
| *Answers* | | | | |
| Shed energy | 16 562, band 6 049 to 29 353 | kWh/yr | 31 times the sheet | 23.3.1 |
| Hours short | 255.7, band 74.7 to 817.1 | h/yr | beside the energy | 23.3.2 |
| Split | units down 0.835; all up 0.165 | – | availability today | 23.3.3 |

## 23.8 Exercises

**Exercise 23.1.** *Formulation.* Write the planner's question *how much of the year's shed energy comes from the plant short with nothing down?* as an expectation over the site's unknown, say which blocks of the unknown it depends on, and give its value. (It is $\mathbb{E}[\text{EUE}_{\text{cap}} \mid \mathbf{y}]$ with $\text{EUE}_{\text{cap}} = \sum_m h_m\, P(\mathbf{s}_0 \mid \mathbf{q}) \max(L_m - \text{Cap}(\mathbf{s}_0, m), 0)$, the first term of equation (23.3). Its capacity reads every block of the plant's state at the all-up state, its probability the rates of both kinds and the strainer's $q_p$. Its value is 2 728.4 kWh/yr, a share of 0.165.)

**Exercise 23.2.** *Derivation.* From equation (23.1), derive the bound on the energy of the states the enumeration leaves out, and evaluate it at the mean residual mass. (In each left-out state the shortfall is $\max(L_m - \text{Cap}, 0) \le L_m$, since a capacity is not negative; summing over those states, their probabilities add to $r(\mathbf{q})$, so $\text{EUE}_{\ge 3} \le r(\mathbf{q}) \sum_m h_m L_m$, equation (23.4). At the mean residual mass of $3.92 \times 10^{-6}$ and the year's 6 241 MWh the bound is 24.47 kWh/yr.)

**Exercise 23.3.** *Derivation.* Show that the shares of the shed energy by limiting element add to one while the shares by element down do not, and say which is larger than the availability term. (Every short state has exactly one limiting group, so the indicators $\mathbf{1}[\ell = g]$ add to one

in each state and the parts add to the whole; a state with two elements down carries the factor $(1 - s_e)$ for each, so it is counted in both elements' shares. The ten shares by element down sum to 0.878, above the availability term's share of 0.835.)

**Exercise 23.4.** *Application.* A facilities report takes the posterior's hours short and prints "service availability 0.971". Say what is wrong with the figure and what the report should carry instead. (It turns hours into an availability, folding into it the 127.1 h the plant is short with nothing down, which no repair ends, and dropping the depth of each shortfall. The report should carry the shed energy and the hours side by side, each with its split: 16 562.4 kWh/yr beside 255.7 h/yr.)

**Exercise 23.5.** *Application.* A planner proposes to defer the year's new IT load, holding the site at today's 660 kW for the whole plan year. Give the shed energy and its capacity term at that load. (With the load flat at 660 kW the plant with nothing down is almost never short, and what remains is mostly the availability term of the hall units: the shed energy is 7 711.8 kWh/yr, about half of the reported 16 562.4, and its all-up part is 2.8 kWh/yr.)

# Chapter 24

# Risk–Cost Tradeoffs in Design

## 24.1 The design review's list of additions

The design review has a list. Four additions to the site have been named in the chapters before this one, each for its own reason: a fourth cooling unit in hall B, to give the hall the N+1 its sheet claims (Chapter 20); a fourth unit in hall A, its twin (Chapter 22); a fourth chiller, the N+1 of the chilled-water plant (Chapter 9); and a bypass around the strainer of the shared supply pipe, whose fouling cuts both halls at once (Chapter 15). The prices are stated: a hall cooling unit installed at R$ 400 000, a strainer bypass at R$ 120 000, a 400 kW chiller at R$ 1 600 000.

The design sheet has one number to set against them. With nameplate chillers, the sheet's unit duties and handbook failure rates, its expected unserved energy over the plan year is 532.2 kWh, and at the stated shed value of R$ 1 500 per kWh that energy is worth R$ 798 357 a year. The sheet cannot say which addition removes which part of that number, because it does not split the number by element, and it cannot say how sure the number is, because it has no band. The planner's question is the drill question, verbatim,

*Where on the cost-versus-expected-shortfall curve does the design sit, and what is the band on that position?*

It is answered with the method of Chapter 4 on the enumeration of Chapter 23, which reads the site's one fit of Chapter 9, the two halls' fits of Chapters 5 and 15, and the outage records. What is new is that the design itself is the variable. Each build-out is a scenario node with its own units, each node has its own expected unserved energy with its own band, and the curve through them is read on one set of draws, so that the differences between designs are sharper than the designs' own bands.

**Jargon: expected unserved energy, loss-of-load expectation, shed value, design variable, annualized price, capital recovery factor, expected consequence, cost–EUE curve, frontier, dominated design, iso-consequence line, common random numbers, breakeven shed value, value of knowing an entry**

The *expected unserved energy* (EUE) of a design is the IT load it is expected to shed over the plan year, in kWh a year, by the enumeration of Chapter 23; the *loss-of-load expectation* beside it counts the hours short. The *shed value* $v$ is the stated price of one kWh of IT load shed, in reais; it is an input, not an estimate. The *design variable* $a$ is the build-out, one of a list of additions to the existing site. Its *annualized price* $c_a$ is the installed price times the *capital recovery factor*, the fraction of a price paid once that a loan over the equipment's life repays each year. The *expected consequence* of a design is its annualized price plus the shed value times its EUE. The *cost–EUE curve* places each design at its annualized price and its EUE. A design is *dominated* when another costs no more and sheds less, or costs less and sheds no more; the *frontier* (the efficient set) is the designs that are not dominated. An *iso-consequence line* joins the points of the plane with the same expected consequence at a stated shed value. *Common random numbers* are the same draws of the plant's unknown used for every design, so that a difference between designs is read draw by draw. The *breakeven shed value* of a step from one design to the next is the shed value at which the step's price equals the shed energy it saves. The *value of knowing an entry* is the expected value of perfect information about that one entry: how much less the best choice is expected to cost if the entry were known before choosing.

**Key idea**

A design choice under uncertainty is an expected consequence minimized over the design variable. With the price certain and the shed energy uncertain, the expected consequence is the annualized price plus the shed value times the expected unserved energy, and the uncertainty enters only through the EUE's posterior. Every design is read on the same draws of the plant's unknown, so that the difference between two designs, the quantity a choice turns on, carries the correlation between them. The shed value is stated, not known; the breakeven shed value of each step, read on the draws, shows which choices depend on it and which do not.

### 24.1.1 Expected consequence over a design variable

The quantity the question owes is the expected consequence of each design, the design that makes it least, and how sure that choice is. Chapter 2 gives its general form, equation (2.23): an action $a$ that costs $c_a$ for certain and leaves events $e$ of consequence $C_e$ has the expected consequence

$$\mathcal{R}(a) = c_a + \sum_e C_e \Pr(e \mid \mathbf{y}, a), \qquad a^* = \arg\min_a \mathcal{R}(a). \tag{24.1}$$

Here the actions are the build-outs, and each term is derived in turn.

**The events and their consequence.** A shortfall event is a state $s$ of the site's elements, each unit up or down, in a month $m$ of the plan year. In that state and month the site carries an IT load of at most its capacity $C_{s,m}(\boldsymbol{\theta};\ a)$, the smallest binding load over its elements with the units of $s$ running, a function of the plant's unknown $\boldsymbol{\theta}$ and of the design. Against the month's load $L_m$ it sheds $\max(L_m - C_{s,m}, 0)$ kW for as many of the month's $h_m$ hours as the state lasts. The consequence of the event is the shed energy priced at $v$, and its probability given the plant is the state's probability $\Pr(s \mid \mathbf{q},\ a)$, the product over the design's elements of each element's unavailability $q$ when down and $1 - q$ when up. Given the unknown, the sum

over the events of equation (24.1) is therefore $v$ times the shed energy,

$$\begin{aligned}\sum_{e} C_e \Pr(e \mid \boldsymbol{\theta},\ a) &= \sum_{m} \sum_{s \in \mathcal{S}_a} v\, h_m \max\bigl(L_m - C_{s,m}(\boldsymbol{\theta};\ a),\ 0\bigr) \Pr(s \mid \mathbf{q},\ a) \\ &= v\, \mathrm{EUE}(a;\ \boldsymbol{\theta}),\end{aligned} \tag{24.2}$$

with $\mathcal{S}_a$ the enumerated states of design $a$ and

$$\mathrm{EUE}(a;\ \boldsymbol{\theta}) = \sum_{m=1}^{12} h_m \sum_{s \in \mathcal{S}_a} \Pr(s \mid \mathbf{q},\ a) \max\bigl(L_m - C_{s,m}(\boldsymbol{\theta};\ a),\ 0\bigr) \tag{24.3}$$

the expected unserved energy of Chapter 23 with the design's own elements. The shed value is a constant, so it leaves the sum.

**The certain cost.** The price of a design is paid once; the EUE is a yearly quantity. The price is put on the same footing by the capital recovery factor $\rho$, fifteen years at ten percent, the factor Chapter 9 uses,

$$\begin{gathered} c_a = \rho\, P_a, \qquad \rho = 0.131\,\mathrm{yr}^{-1}, \\ P_a = \mathrm{R\$}\,400\,000\,(n_{A4} + n_{B4}) + \mathrm{R\$}\,1\,600\,000\,n_{C4} + \mathrm{R\$}\,120\,000\,n_{\mathrm{byp}}, \end{gathered} \tag{24.4}$$

with each $n$ one when the design includes that addition and zero otherwise. Operating and maintenance costs of the additions are left out; a design review that has them adds them to $c_a$.

**The expected consequence of a design.** The loss given the unknown is $L(a, \boldsymbol{\theta}) = c_a + v\,\mathrm{EUE}(a;\ \boldsymbol{\theta})$. Its posterior expectation, by equation (24.1), is

$$\mathcal{R}(a) = \mathbb{E}\bigl[L(a,\boldsymbol{\theta}) \mid \mathbf{y}\bigr] = c_a + v\,\mathbb{E}\bigl[\mathrm{EUE}(a;\ \boldsymbol{\theta}) \mid \mathbf{y}\bigr], \qquad a^* = \arg\min_a \mathcal{R}(a). \tag{24.5}$$

The loss is linear in the EUE, so the expected consequence needs only the EUE's posterior mean; the band of the loss, and the probability that one design costs less than another at the true state, need its draws. Two consequences of the linearity are used below. For a fixed design, $\mathcal{R}(a)$ is a straight line in $v$ with intercept $c_a$ and slope $\mathbb{E}[\mathrm{EUE}(a) \mid \mathbf{y}]$, so the design of least expected consequence changes only where two lines cross. And the points of the cost–EUE plane with the same expected consequence at a given $v$ lie on the line $c + v\,E = \text{constant}$, the iso-consequence line.

**The frontier.** Before any shed value is chosen, some designs can be set aside. Design $b$ dominates design $a$ when it costs no more and sheds no more in expectation, one of the two strictly,

$$c_b \le c_a, \quad \mathbb{E}[\mathrm{EUE}(b) \mid \mathbf{y}] \le \mathbb{E}[\mathrm{EUE}(a) \mid \mathbf{y}], \quad \text{one strictly} \quad \Longrightarrow \quad \mathcal{R}(b) < \mathcal{R}(a) \text{ for every } v > 0, \tag{24.6}$$

so a dominated design is never the choice at any shed value. The frontier is the designs that no other dominates.

**The breakeven of a step.** A step from design $a$ to a costlier design $b$ pays at shed value $v$ on a draw $i$ of the unknown when the shed energy it saves, priced at $v$, exceeds its price. On that draw the shed value at which it pays is

$$v^*_{a\to b}(\boldsymbol{\theta}_i) = \frac{c_b - c_a}{\mathrm{EUE}(a;\ \boldsymbol{\theta}_i) - \mathrm{EUE}(b;\ \boldsymbol{\theta}_i)}, \qquad \Pr(\text{step pays at } v \mid \mathbf{y}) = \Pr\bigl(v^*_{a\to b} < v \mid \mathbf{y}\bigr), \tag{24.7}$$

a ratio whose denominator is the step's saving in kWh a year. The breakeven is read on the draws, its median and its 90 % band the quantiles of $v^*$ over them. Its delta form, the ratio of the mean price to the mean saving with a band from the gradient, is a ratio's delta form and is not reported: the saving is skewed, and the ratio's quantiles are not those of a normal.

**Common random numbers.** Every design's EUE is read on the same draws $\boldsymbol{\theta}_i$, $i = 1, \dots, N$. A difference between two designs is then a difference on each draw, and its variance is

$$\operatorname{Var}\big(\mathrm{EUE}(a) - \mathrm{EUE}(b)\big) = \operatorname{Var}\big(\mathrm{EUE}(a)\big) + \operatorname{Var}\big(\mathrm{EUE}(b)\big) - 2\,\operatorname{Cov}\big(\mathrm{EUE}(a),\ \mathrm{EUE}(b)\big), \quad (24.8)$$

the covariance taken over the posterior. Two designs share most of the plant: a hall B that sheds much as built sheds much with a fourth chiller too, so the covariance is large and positive and the difference is far narrower than either design's band. Read on independent draws the covariance term would be lost, the variance of a difference would be the sum of the two variances, and a step whose saving is clearly positive could look uncertain. The check of the paired against the unpaired spread is in Section 24.4.2.

> **Principle 24.1.** *The expected consequence of a design is its annualized price plus the shed value times its expected unserved energy, the expected-consequence form of Chapter 2 with the events the enumerated states and months. Designs that another dominates in price and EUE are never the choice. The others are compared on one set of draws, so that each difference carries the correlation between designs, and the shed value at which each step pays is read on those draws as a distribution, not as a ratio of means.*

### 24.1.2 The site and its additions (step 0)

The plant is the site of Chapter 9, drawn in Figure 9.1, with both halls as Chapter 15 draws them in Figure 15.1, and its elements as Chapter 23 enumerates them: the chillers C1–C3, the hall units A1–A3 and B1–B3, and the strainer of the shared pipe, clean or fouled, each up or down independently with its kind's unavailability. The pump and the feed are taken as up. The unknown is that of equation (15.8) with the two unavailabilities beside it,

$$\boldsymbol{\theta} = \big(\mathbf{z}_A,\ \mathbf{z}_B,\ \mathbf{z}_{\mathrm{site}},\ q_c,\ q_h\big), \qquad q = \frac{\lambda}{\lambda + \mu},$$

each hall's reduced state of equation (5.4), the site's eight entries of equation (9.4), and the unavailabilities of the chillers and the hall units from their failure and repair rates, equation (9.6). The strainer's unavailability is stated, $6.845 \times 10^{-4}$, from Chapter 15's fouling rate of 0.5 a year and twelve hours to clear.

**The physical relations.** A state's capacity in site IT load is the smallest of the elements' binding loads with the state's units running, equation (22.6), with every unit of the state that is down taken out of its sum,

$$C_{s,m}(\boldsymbol{\theta};\ a) = \min\Big\{Q^{\dagger}_{\mathrm{ch}}(\boldsymbol{\theta};\ s, m),\ Q^{\dagger}_{A}(\boldsymbol{\theta};\ s),\ Q^{\dagger}_{B}(\boldsymbol{\theta};\ s),\ Q^{\dagger}_{\mathrm{pipe}}(\boldsymbol{\theta};\ s),\ Q^{\dagger}_{\mathrm{pump}}(\boldsymbol{\theta}),\ Q^{\dagger}_{\mathrm{feed}}(\boldsymbol{\theta};\ m)\Big\}, \quad (24.9)$$

the chillers' at the month's mean wet bulb, each hall's by the closed form of Chapter 20 over its running units and its share of the site load, and the pipe's at 60 % of its flow coefficient when the strainer is fouled. The month's load is the plan year's, from 668 kW in January to 759 kW in December, each month of $h_m = 8\,760/12 = 730$ h. The states enumerated are every state with at most two elements down; the rest carry a mass whose mean is $3.92 \times 10^{-6}$, and Chapter 23 bounds their part of the EUE by 24.47 kWh a year at that mass.

Table 24.1: The design variable and its prices: installed prices stated, a hall cooling unit R$ 400 000, a strainer bypass R$ 120 000, a 400 kW chiller R$ 1 600 000; annualized at the capital recovery factor 0.131 a year, equation (24.4).

| Design [–] | Additions [–] | Installed price $P_a$ [R$] | Annualized $c_a$ [R$/yr] |
|---|---|---|---|
| as built | none | 0 | 0 |
| +B4 | hall B unit | 400 000 | 52 400 |
| +A4 | hall A unit | 400 000 | 52 400 |
| +A4+B4 | both hall units | 800 000 | 104 800 |
| +A4+B4+bypass | both, strainer bypass | 920 000 | 120 520 |
| +A4+B4+C4 | both, fourth chiller | 2 400 000 | 314 400 |
| +A4+B4+C4+bypass | all four | 2 520 000 | 330 120 |
| +C4 | fourth chiller alone | 1 600 000 | 209 600 |

**Modeling choices at the graph.** Each choice is Chapter 23's. The elements fail independently: this makes a state's probability a product, and it fails if a shared cause, a power event or a control fault, takes several units at once. Each month is read at its mean wet bulb and its end-of-month load: this makes the year twelve steady states, and it fails if a month's shortfall hides in its hottest hours. The supply is pinned at 14°C in every state, so a short chiller plant does not also lower the halls' binding loads, the chain Chapter 21 follows. A shed is taken as the load above capacity for the state's whole duration: no load is moved between halls and no temporary chiller is brought in.

### 24.1.3 The design variable and its prices

A design is a build-out of the four additions. The design variable takes eight values: the site as built; a fourth unit in hall B; a fourth unit in hall A; both; both and the strainer bypass; both and the fourth chiller; all four; and, beside the build-out, the fourth chiller alone, the addition Chapter 9 weighs first. Each is a scenario node of Section 4.3.8 whose pinned inputs are the ratings of the units added. The additions are not yet built, so their ratings are stated, as Chapter 22 states them in equation (22.8): a fourth hall unit at its hall's rated conductance times $e^{g_{h4}}$ with $g_{h4} \sim \mathcal{N}(0,\ 0.15^2)$, and a fourth chiller with the log-ratio of equation (9.7). Each added unit takes its kind's unavailability, and the strainer bypass sets the strainer's to zero: the pipe runs clean while its strainer is cleared.

Table 24.1 gives the design variable. Each row is a design; the columns name the additions it includes, its installed price $P_a$ and its annualized price $c_a$ of equation (24.4).

A hall unit costs a quarter of a chiller, and the bypass less than a third of a hall unit. The fourth chiller is the costliest single addition by a factor of four; the full package is a little over three times the price of the two hall units together.

### 24.1.4 The planner's questions (step 1)

The planner asks six questions of the design list, and step 1 writes each over the plant's unknown, with the evidence on the right of the bar. Table 24.2 lists them. The first column is the question in the planner's words, the drill question first; the second the general form in the notation of Chapter 2; the third its mathematical form for the site; 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. Four are design questions answered by the expected consequence; the fifth estimates where the curve sits a year later; the sixth asks which measurement could change the choice, a design question that the preposterior answers. The atlas of Appendix A collects them with the other questions of Volume II.

Table 24.2: The questions asked of the design list. $\mathbf{y}$ is both halls' night record, the site's operating record, the outage records and the priors; $\mathcal{R}(a)$ is the expected consequence of design $a$, equation (24.5), $v^*$ the breakeven shed value of a step (24.7), and $v$ the stated shed value.

| Planner's question | General form | Mathematical form | Category | Computation | Answered in |
|---|---|---|---|---|---|
| Where on the cost-versus-expected-shortfall curve does the design sit, and what is the band on that position? | expected consequence over the design variable, eq. (2.23) | $(c_a,\ \mathbb{E}[\mathrm{EUE}(a) \mid \mathbf{y}])$ with its band; the frontier | design | expected consequence | 24.3.1 |
| Which design has the least expected consequence, and how sure? | $\arg\min_a \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]$ | $a^* = \arg\min_a \mathcal{R}(a)$; $\Pr(a^* \text{ least on a draw})$ | design | expected consequence | 24.3.2 |
| At what shed value does each step pay? | the breakeven of an increment | $v^*_{a\to b}$ on each draw, its median and band | design | expected consequence | 24.3.3 |
| How does the choice move with the shed value? | the minimizer against a stated input | $\arg\min_a \mathcal{R}(a)$ for $v$ from 1 to 3 000 R\$/kWh | design | expected consequence | 24.3.4 |
| Where does the curve sit at next year's load? | the same, at a later scenario node | $\mathbb{E}[\mathrm{EUE}(a) \mid \mathbf{y}]$ at 768 to 873 kW | estimation | expected consequence | 24.3.5 |
| Which measurement could change the choice? | expected value of perfect information, per entry, eq. (2.53) | $\min_a \mathcal{R}(a) - \mathbb{E}_{z_k}[\min_a \mathcal{R}(a \mid z_k)]$ | design | preposterior | 24.6.1 |

### 24.1.5 The evidence as rows (step 2)

Step 2 attaches the records of Chapter 23, unchanged. Each hall is read in window 1 of the night record, its seventeen rows of equation (5.11); the site is read by its operating record of Chapter 9, sixty-two rows on its eight entries; and the outage records of the chillers and of the hall units give their failure and repair rates. No row is added. The design variable adds units, and a unit not yet built has no record: its rating enters through its stated normal, not through a row.

**Which row acts on which part of the EUE.** The EUE of equation (24.3) splits by the element that limits each short state, and each part reads its own blocks of the unknown through the capacity (24.9) and all of the unavailabilities through the state's probability. With $*$ for a block that can be nonzero,

$$\begin{bmatrix} \partial \mathrm{EUE}_{\mathrm{ch}} \\ \partial \mathrm{EUE}_{A} \\ \partial \mathrm{EUE}_{B} \\ \partial \mathrm{EUE}_{\mathrm{pipe}} \end{bmatrix} = \begin{bmatrix} 0 & 0 & * & * & * \\ * & 0 & 0 & * & * \\ 0 & * & 0 & * & * \\ 0 & 0 & * & * & * \end{bmatrix}, \tag{24.10}$$

with the columns, in order, hall A's state, hall B's state, the site's entries, $q_c$ and $q_h$. Two facts are read here. Each hall's part reads its own state and no other, so an addition to hall B moves the hall-B part and hands the states it no longer leaves short to whatever element limits them next. And every part reads both unavailabilities, since every state's probability is a product over all the elements; the chillers' part reads the chillers' capacity entries and $q_c$ most strongly, the halls' parts their states and $q_h$.

**Modeling choices at the rows.** The rows are those of Chapter 23, with its choices: each reading Gaussian, unbiased and independent given the state; each hall fitted on window 1 of its own record; the chillers sharing one failure rate and one repair rate, the hall units another pair. A fourth unit takes its kind's rates, which assumes that a new unit fails like the fleet it joins; a new unit with a burn-in period would fail more often in its first year.

## 24.2 The posterior of the state (step 3)

All six questions of Table 24.2 are read from one posterior, Chapter 23's, the product of the three fits of Chapters 5, 9 and 15 with the outage posteriors beside them. No row is added and no fit is repeated; this section states what the curve needs from that posterior, cites the fits, and states the draws on which every design is read.

### 24.2.1 What the curve needs from the state

By equation (24.5) each design needs the posterior mean of its EUE, and the band, the probability of being least and the breakevens need the EUE's joint distribution over the designs. By the pattern (24.10) that needs the three blocks of the state and the two unavailabilities, and nothing of their cross-covariances, since no record reads two blocks at once. The posterior factors as in equation (22.10), with the repair rates kept,

$$p(\boldsymbol{\theta} \mid \mathbf{y}) = p(\mathbf{z}_A \mid \mathbf{y}_A)\, p(\mathbf{z}_B \mid \mathbf{y}_B)\, p(\mathbf{z}_{\mathrm{site}} \mid \mathbf{y}_{\mathrm{op}})\, p(q_c \mid \mathbf{y}_{\mathrm{out}})\, p(q_h \mid \mathbf{y}_{\mathrm{out}}), \tag{24.11}$$

each $q$ carried by the gamma posteriors of its failure and repair rates. The EUE is a sum of maxima over states and months, not a smooth function of the unknown, so it is read on draws of the whole unknown and never through a gradient.

### 24.2.2 The fits, cited

Each block's posterior is the Laplace normal of its own fit. For block $b$, one of hall A, hall B and the site, the residual vector $\mathbf{r}_b(\mathbf{z}_b)$ stacks that block's normalized rows, its Jacobian is $\mathbf{J}_b = \partial \mathbf{r}_b / \partial \mathbf{z}_b$, and the MAP state solves the normal equations of equation (5.12), with the Laplace covariance of equation (5.13),

$$\begin{aligned} (\mathbf{J}_b^{\mathsf{T}} \mathbf{J}_b)\, \Delta \mathbf{z}_b &= -\mathbf{J}_b^{\mathsf{T}} \mathbf{r}_b(\mathbf{z}_b), \\ \boldsymbol{\Sigma}_b = (\mathbf{J}_b^{\mathsf{T}} \mathbf{J}_b)^{-1}\big|_{\hat{\mathbf{z}}_b}, \qquad & \mathbf{z}_b \mid \mathbf{y}_b \approx \mathcal{N}(\hat{\mathbf{z}}_b,\ \boldsymbol{\Sigma}_b). \end{aligned} \tag{24.12}$$

Hall B's fit is that of Section 5.2.2 and hall A's that of Section 15.2.2, both on window 1. The site's is the one fit of Section 9.2.2. The rates are the exact gamma posteriors of equations (9.12) and (15.11); Chapter 23 reads from them the chillers' unavailability with mean 0.00552 and 90 % band from 0.00271 to 0.00967, and the hall units' with mean 0.00239, band 0.00139 to 0.00375.

### 24.2.3 The draws, common to every design

The answers read the posterior through Chapter 23's $N = 4\,000$ paired draws. The site's are the first 4 000 of the draws of Section 9.2.4, whose weights (5.16) toward the exact posterior have an effective sample size of 3 983; each hall's are 4 000 draws from the Laplace normal of its fit, paired with the site's by index; each draw carries its own $q_c$ and $q_h$ from the rates' posteriors and its own ratings of the units not yet built. Written out, with $i$ the index that pairs them,

$$\boldsymbol{\theta}_i = (\mathbf{z}_{A,i},\ \mathbf{z}_{B,i},\ \mathbf{z}_{\text{site},i},\ q_{c,i},\ q_{h,i}), \qquad \hat{\mathcal{R}}(a) = c_a + \frac{v}{4\,000} \sum_{i=1}^{4\,000} \text{EUE}(a;\ \boldsymbol{\theta}_i), \tag{24.13}$$

and the exact posterior's value is the same sum under the site's weights. These are the common random numbers of equation (24.8): the eight designs are eight functions of one set of 4 000 draws, so the design whose loss is least on draw $i$, and the breakeven of each step on draw $i$, are read with the correlation between designs intact.

## 24.3 The answers (step 4)

Every answer is read from the draws of Section 24.2.3 at the stated shed value of R$ 1 500 per kWh unless the answer varies it. The curve answers the drill question; the least design, the breakevens and the choice against the shed value follow; the year after closes the step.

### 24.3.1 The cost–EUE curve

**Planner's question.** Where on the cost-versus-expected-shortfall curve does the design sit, and what is the band on that position? The question asks for each design's annualized price and EUE with its band, for the frontier through them, and for the as-built design's place on it.

**General form.** A design question chooses among actions by their expected consequence, equation (2.23),

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

read here at every value of the design variable, not only at the minimum.

Table 24.3: The cost–EUE curve at the stated shed value of R$ 1 500 per kWh: each design's expected unserved energy over the plan year and its expected consequence $\mathcal{R}(a)$, equation (24.5), by 4 000 paired draws, common to every design; exact: the same under the site's weights (effective sample size 3 983). Gap: exact less draws, as a share of the exact. Prices in Table 24.1.

| Design [–] | EUE mean [kWh/yr] | 5 % [kWh/yr] | 95 % [kWh/yr] | $\mathcal{R}(a)$, draws [R$ M/yr] | 90 % band [R$ M/yr] | $\mathcal{R}(a)$, exact [R$ M/yr] | Gap [%] |
|---|---|---|---|---|---|---|---|
| as built | 16 562.4 | 6 050.1 | 29 233.6 | 24.844 | 9.075 to 43.850 | 24.886 | 0.17 |
| +B4 | 5 304.2 | 1 131.0 | 11 983.5 | 8.009 | 1.749 to 18.028 | 8.011 | 0.03 |
| +A4 | 12 110.8 | 3 889.4 | 20 067.9 | 18.219 | 5.886 to 30.154 | 18.262 | 0.24 |
| +A4+B4 | 806.2 | 428.0 | 1 532.7 | 1.314 | 0.747 to 2.404 | 1.317 | 0.24 |
| +A4+B4+bypass | 535.4 | 159.2 | 1 263.9 | 0.924 | 0.359 to 2.016 | 0.927 | 0.38 |
| +A4+B4+C4 | 409.6 | 276.9 | 713.9 | 0.929 | 0.730 to 1.385 | 0.929 | 0.02 |
| +A4+B4+C4+bypass | 138.5 | 17.0 | 449.8 | 0.538 | 0.356 to 1.005 | 0.538 | 0.09 |
| +C4 | 16 167.8 | 5 747.5 | 29 079.2 | 24.461 | 8.831 to 43.828 | 24.501 | 0.16 |

**Specialized form.** The question is equation (2.23) with the loss the annualized price plus the shed energy priced at $v$, which is equation (24.5), with the EUE of equation (24.3) read on the draws (24.13); the position of a design is the pair $(c_a,\ \mathbb{E}[\mathrm{EUE}(a) \mid \mathbf{y}])$ with the 90 % band of EUE$(a)$ over the draws, and the frontier is the set of designs not dominated in the sense of equation (24.6).

**Building blocks.** The enumeration (24.3) under each design; the capacity (24.9) with the added units; the prices of Table 24.1; the paired draws and the site's weights.

**The curve.** Table 24.3 gives it. Each row is a design; the columns give its EUE's mean over the draws and its 90 % band, its expected consequence at R$ 1 500 per kWh by the draws with its band, the same by the exact posterior, and the gap between the two routes.

The curve is drawn in Figure 24.1. The horizontal axis is the annualized price and the vertical axis the EUE on a logarithmic scale; each design is a point at its mean with a bar over its 90 % band, the frontier is the solid line, the stated truth's points are the open circles, joined by a dotted line along the frontier, and the two thin lines are iso-consequence lines, at R$ 1 500 per kWh through the least design and at R$ 663 per kWh, the fourth chiller's median breakeven of Section 24.3.3, through the least design there.

**The answer.** The as-built design sits at the far end of the curve,

$$\begin{aligned} \mathbb{E}[\mathrm{EUE}(\text{as built}) \mid \mathbf{y}] &= 16\,562.4 \text{ kWh/yr}, \quad 90\,\% \text{ band } 6\,050.1 \text{ to } 29\,233.6, \\ \mathcal{R}(\text{as built}) &= \text{R\$ } 24.84 \text{ M/yr}, \quad 90\,\% \text{ band R\$ } 9.08 \text{ M to R\$ } 43.85 \text{ M}, \end{aligned} \tag{24.14}$$

at zero price and the largest shed energy of any design. The frontier runs as built, +B4, +A4+B4, +A4+B4+bypass, +A4+B4+C4, +A4+B4+C4+bypass; +A4 and +C4 are dominated, +A4 by +B4 at the same price and +C4 by every design from +B4 on, which cost less and shed less.

*The fourth chiller alone is off the frontier.* It is the addition Chapter 9 found worth building, and on that chapter's basis, the chillers alone at the design condition, it is. On the year, with both halls in the enumeration, it buys $16\,562.4 - 16\,167.8 = 394.6$ kWh a year, about the chillers' part of the as-built EUE, while hall B goes on shedding thousands. A fourth hall-B unit at a quarter of its price removes 11 258 kWh a year. The fourth chiller pays for itself against the as-built design, R$ 24.46 M against R$ 24.84 M a year, and still is not the addition to make first. Measured as EUE removed per real a year, the hall-B unit gives 0.2149 kWh, the hall-A unit 0.0850 and the fourth chiller 0.0019, a ratio of more than a hundred to one between the first and the last.

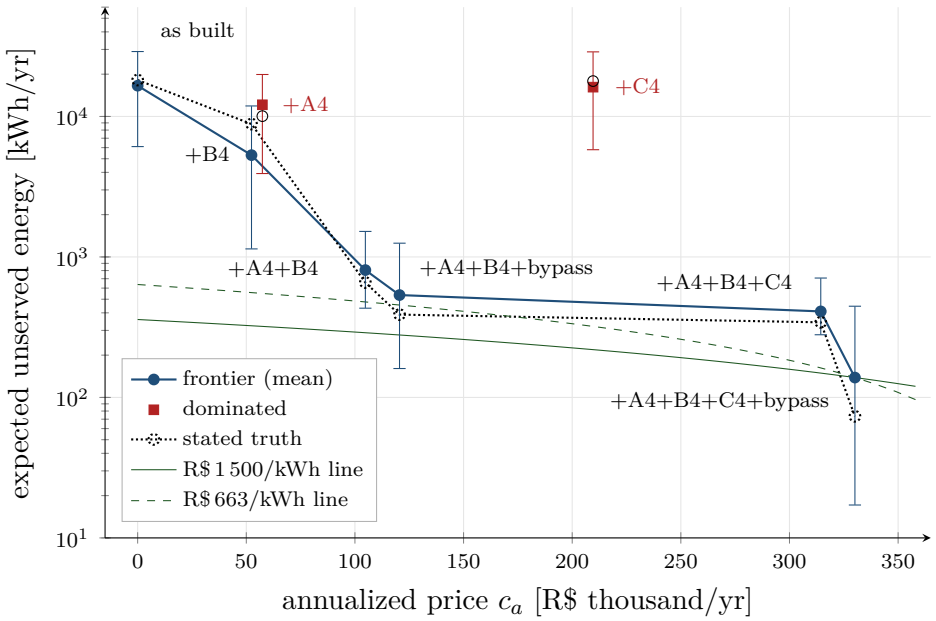



Figure 24.1: The cost–EUE plane: each design at its annualized price and its posterior mean EUE (logarithmic axis), with the 90 % band over the paired draws as a bar, from Table 24.3. Solid line: the frontier. Squares: the dominated designs, +A4 drawn R$ 5 thousand right of its price so that its bar does not cover that of +B4. Open circles: the stated truth's EUE for each design, joined by the dotted line along the frontier. Thin solid line: the iso-consequence line at R$ 1 500 per kWh through the least design; dashed, the same at R$ 663 per kWh, the fourth chiller's median breakeven.

**The limiting element along the curve.** Table 24.4 says why the frontier has this order. Each row is a design; the columns give its hours short, the loss-of-load expectation beside the EUE, and the shares of its mean EUE by the element that limits each short state.

*Each addition hands the shortfall to the next element.* As built, hall B carries 0.688 of the EUE. With its fourth unit, hall A carries 0.854; with both hall units, the chillers carry 0.493 and the pipe 0.336. The bypass takes the pipe out and leaves the chillers 0.743; the fourth chiller instead takes the chillers out and leaves the pipe 0.662; with all four, the halls are back in front, hall B at 0.756, now on states with two of its four units down. The order of the frontier is the order in which the elements bind, the binding constraints of Chapter 13 read in energy rather than in load. It also explains the dominated designs. The fourth chiller alone and the fourth hall-A unit alone each relieve an element that is not the one limiting most short states: with them built, hall B still carries 0.704 and 0.942 of what is left.

The hours fall less than the energy along the first steps and more along the last. As built the site is short 255.7 hours a year, which the fourth hall-B unit takes to 73.1; with all four additions it is short 3.2 hours. The hours and the energy are two readers of the same enumeration and neither is converted to the other: 3.2 hours short is not an availability.

### 24.3.2 The design of least expected consequence

**Planner's question.** Which design has the least expected consequence, and how sure? The question asks for the minimizer of equation (24.5) at the stated shed value, and for the probability that it is least at the true state.

Table 24.4: Hours short and the limiting element along the curve: each design's loss-of-load expectation and the shares of its mean EUE by the element that limits the short state, by 4 000 paired draws. A dash: no short state limited by that element (the pipe, with the bypass built). The pump and the feed limit no state.

| Design [–] | Hours short [h/yr] | Chillers [–] | Hall A [–] | Hall B [–] | Pipe [–] |
|---|---|---|---|---|---|
| as built | 255.7 | 0.024 | 0.272 | **0.688** | 0.016 |
| +B4 | 73.1 | 0.075 | **0.854** | 0.020 | 0.051 |
| +A4 | 200.7 | 0.033 | 0.003 | **0.942** | 0.022 |
| +A4+B4 | 16.9 | **0.493** | 0.040 | 0.130 | 0.336 |
| +A4+B4+bypass | 11.1 | **0.743** | 0.061 | 0.196 | – |
| +A4+B4+C4 | 9.0 | 0.004 | 0.079 | 0.256 | **0.662** |
| +A4+B4+C4+bypass | 3.2 | 0.011 | 0.233 | **0.756** | – |
| +C4 | 248.2 | 0.000 | 0.279 | **0.704** | 0.017 |

**General form.** The action of least expected loss, equation (2.23), with the firmness of a ranking read as the probability that one action costs less than the others at the true state, the comparison of equation (2.24),

$$a^{*} = \arg\min_{a} \mathcal{R}(a), \qquad \Pr\big(L(a^{*}, \mathbf{z}) \le L(a, \mathbf{z}) \text{ for every } a \mid \mathbf{y}\big).$$

**Specialized form.** The question is equation (2.23) with $\mathcal{R}(a)$ of equation (24.5) at $v =$ R\$ 1 500 per kWh, the minimum over the eight designs. The probability that a design is least is read on the common draws (24.13), as the share of draws on which its loss is the smallest of the eight, rather than by the normal comparison of equation (2.24): the EUE is a sum of maxima and its distribution is skewed.

**Building blocks.** The curve of Table 24.3; the loss of every design on each draw.

**The answer.** The least design is the full package,

$$\begin{aligned} a^{*} &= \text{+A4+B4+C4+bypass}, \qquad \mathcal{R}(a^{*}) = \text{R\$ } 537\,944/\text{yr} \ \ (\text{exact R\$ } 538\,450), \\ \Pr(a^{*} \text{ least on a draw} \mid \mathbf{y}) &= 0.887, \qquad \Pr(\text{+A4+B4+bypass least} \mid \mathbf{y}) = 0.113, \end{aligned} \tag{24.15}$$

and on no draw is any other design least. Against the as-built design it saves R\$ 24.31 M a year in expectation, with a 90 % band from R\$ 8.72 M to R\$ 42.62 M, and the saving is positive on every draw, a probability of 1.000 at the resolution of 4 000 draws.

The two designs that share the probability differ by the fourth chiller alone. Whether to build the hall units and the bypass is not in doubt at this shed value; whether to add the chiller is, on about one draw in nine. The next two subsections say why, and how the answer moves when the stated shed value does. In expectation the two contenders cost R\$ 537 944 and R\$ 923 580 a year, a difference of R\$ 385 636 and a wide margin for a choice that fails on one draw in nine: on the draws on which the chiller does not pay it loses at most its price, while on the draws on which it pays it can save several times its price, and the mean weighs those draws by what they save.

### 24.3.3 The shed value at which each step pays

**Planner's question.** At what shed value does each step pay? The question asks, for each step along the frontier, the shed value above which the step's saving exceeds its price, and how sure that value is.

Table 24.5: The four steps along the frontier: the price of the step a year, its saving in EUE on the paired draws, and the breakeven shed value of equation (24.7) on each draw, its median and 90 % band. $\Pr(v^* < 1\,500)$: the step pays at the stated shed value.

| Step [–] | Price [R$/yr] | Saving, mean [kWh/yr] | 90 % band [kWh/yr] | $v^*$ median [R$/kWh] | 90 % band [R$/kWh] | $\Pr(v^* < 1\,500)$ [–] |
|---|---|---|---|---|---|---|
| as built to +B4 | 52 400 | 11 258.2 | 3 175.4 to 18 891.1 | 6.48 | 2.77 to 16.50 | 1.0000 |
| +B4 to +A4+B4 | 52 400 | 4 498.0 | 499.1 to 10 980.8 | 14.59 | 4.77 to 105.00 | 0.9995 |
| +A4+B4 to +bypass | 15 720 | 270.8 | 241.3 to 301.4 | 58.13 | 52.15 to 65.14 | 1.0000 |
| +bypass to +C4 | 209 600 | 396.8 | 102.5 to 961.3 | 662.61 | 218.05 to 2 044.90 | 0.8872 |

**General form.** The comparison of two actions by expected consequence, equation (2.23), written for an increment: the step from $a$ to $b$ pays when

$$L(b, \mathbf{z}) < L(a, \mathbf{z}) \iff c_b - c_a < \sum_e C_e \left[\Pr(e \mid \mathbf{z}, a) - \Pr(e \mid \mathbf{z}, b)\right],$$

and its breakeven is the value of the stated consequence at which the two sides are equal.

**Specialized form.** The question is equation (2.23) for the increment, with the right-hand side $v$ times the step's saving in EUE by equation (24.2), solved for $v$ on each draw, which is equation (24.7). The steps are those of the frontier that the least design passes through, as built to +B4, to +A4+B4, to +A4+B4+bypass, to the full package; the step from the bypass to the full package is the fourth chiller.

**Building blocks.** The paired EUEs of consecutive designs on each draw; the step's price $c_b - c_a$ from Table 24.1.

**The breakevens.** Table 24.5 gives the four steps. Each row is a step; the columns give its price a year, its mean saving in EUE with the 90 % band, the median breakeven with its band, and the probability that it pays at the stated R$ 1 500 per kWh.

The breakevens are set against the stated shed value in Figure 24.2, on a logarithmic axis in reais per kWh: each step is a bar over its 90 % band with its median marked, and the dashed line is the stated R$ 1 500.

*Only the fourth chiller's decision depends on the stated shed value.* The two hall units pay at any shed value above a few reais per kWh: their medians are R$ 6.48 and R$ 14.59, and the upper ends of their bands R$ 16.50 and R$ 105.00, one to two orders of magnitude below the stated value. The bypass pays above about R$ 58 per kWh, with a band from R$ 52.15 to R$ 65.14, the narrowest of the four, because its saving is fixed by the strainer's stated fouling rate and its price is small; its saving's band runs only from 241.3 to 301.4 kWh. The fourth chiller's median breakeven is R$ 662.61 per kWh, with a band from R$ 218.05 to R$ 2 044.90, and the stated R$ 1 500 falls inside it: the step pays on 0.8872 of the draws, the same share on which the full package is least in equation (24.15), since the two contenders differ by that step alone.

The band of the fourth chiller's breakeven is the band of its saving turned over. The saving runs from 102.5 to 961.3 kWh a year, a factor of nine, because the chillers' part of the EUE scales with the chance of a chiller down, and the chillers' unavailability is known from a short record to about a factor of three and a half. The breakeven inherits the factor of nine, reversed, and with the price fixed it is the saving's distribution read through a reciprocal. That reciprocal is why the median breakeven, R$ 662.61, lies above the ratio of the step's price to its mean

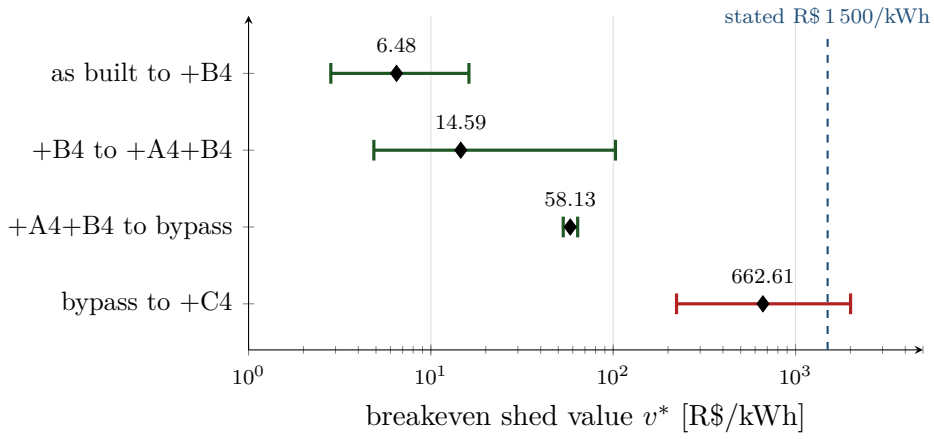



Figure 24.2: The breakeven shed value of each step along the frontier, equation (24.7): the 90 % band over the paired draws (bar) and the median (diamond, value above), on a logarithmic axis, from Table 24.5; dashed, the stated shed value. Three bands lie wholly to the left of it; the fourth chiller's straddles it.

saving: the saving is skewed to the right, its mean lies above its median, and a ratio built on the mean moves the breakeven toward the cheap side. The breakeven is reported on the draws.

### 24.3.4 The choice against the shed value

**Planner's question.** How does the choice move with the shed value? The question asks for the design of least expected consequence over a range of shed values, from one real to three thousand per kWh, and for how sure each choice is.

**General form.** The minimizer of the expected consequence (2.23) taken as a function of a stated input of the consequence schedule,

$$a^*(v) = \arg\min_a \mathcal{R}(a;\ v),$$

the stated input varied and nothing else.

**Specialized form.** The question is equation (2.23) with $\mathcal{R}(a;\ v)$ of equation (24.5), a straight line in $v$ for each design, minimized over the frontier at each of nine values of $v$, with the probability of being least read on the common draws as in Section 24.3.2. Because each line is $c_a + v\,\mathbb{E}[\mathrm{EUE}(a) \mid \mathbf{y}]$, the least design changes from $a$ to $b$ exactly where the two lines cross,

$$\bar{v}_{a\to b} = \frac{c_b - c_a}{\mathbb{E}[\mathrm{EUE}(a) \mid \mathbf{y}] - \mathbb{E}[\mathrm{EUE}(b) \mid \mathbf{y}]},$$

the shed value at which the two expected consequences are equal. This is not the breakeven of equation (24.7): it is the switch point of the expectation, exact for a loss linear in the EUE, and the breakeven's median is a different quantity that answers a different question.

**Building blocks.** The mean EUE and the price of each frontier design; the loss of every design on each draw.

Table 24.6: The design of least expected consequence against the shed value $v$, over the frontier designs, equation (24.5); the share of the 4 000 paired draws on which it is least beside.

| Shed value $v$ [R$/kWh] | Least design [–] | $\mathcal{R}(a^*)$ [R$/yr] | Pr(least on a draw) [–] |
|---:|---|---:|---:|
| 1 | as built | 16 562 | 0.9798 |
| 5 | +B4 | 78 921 | 0.2445 |
| 10 | +B4 | 105 442 | 0.5222 |
| 25 | +A4+B4 | 124 955 | 0.6887 |
| 100 | +A4+B4+bypass | 174 057 | 0.9593 |
| 300 | +A4+B4+bypass | 281 132 | 0.8795 |
| 1 000 | +A4+B4+C4+bypass | 468 669 | 0.7335 |
| 1 500 | +A4+B4+C4+bypass | 537 944 | 0.8872 |
| 3 000 | +A4+B4+C4+bypass | 745 768 | 0.9890 |

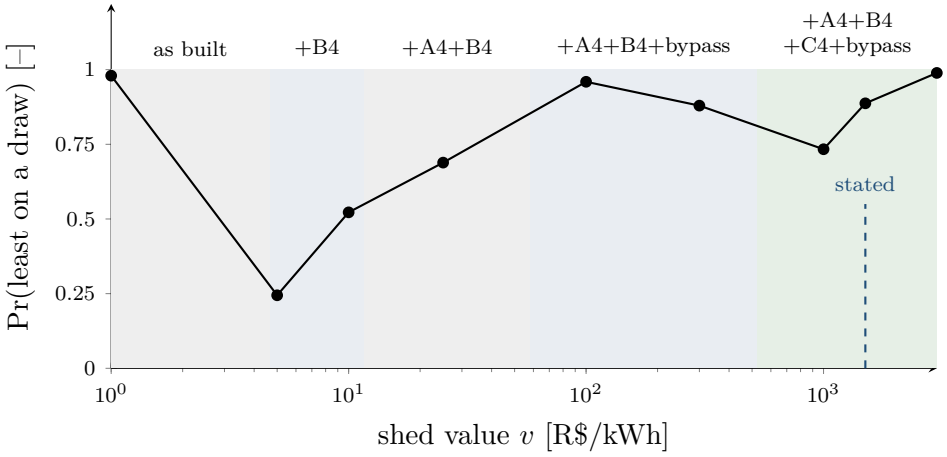



Figure 24.3: The design of least expected consequence against the shed value, from Table 24.6. Shaded bands: the least design between the switch points of the expected consequences, R$ 4.65, R$ 11.65, R$ 58.04 and R$ 528.19 per kWh. Line: the share of the paired draws on which that design is least, lowest at the values just past a switch, R$ 5 and R$ 1 000. Dashed: the stated shed value.

**The choice.** Table 24.6 gives it. Each row is a shed value; the columns give the least design, its expected consequence, and the share of draws on which it is least.

The same choice is plotted against the shed value in Figure 24.3. The horizontal axis is the shed value on a logarithmic scale; the shaded bands mark the least design between the switch points $\bar{v}$, which fall at R$ 4.65, R$ 11.65, R$ 58.04 and R$ 528.19 per kWh, and the line gives the share of draws on which the least design is least at each of the nine values of Table 24.6.

The choice walks the frontier as the shed value grows: the as-built design below about R$ 5 per kWh, the hall-B unit to about R$ 12, both hall units to about R$ 58, the bypass with them to about R$ 528, and the full package above. The share of draws on which the least design is least dips where a value of the table falls just past a switch, to 0.2445 at R$ 5 past the first and to 0.7335 at R$ 1 000 past the last, and recovers away from the switches, to 0.9890 at R$ 3 000. Every switch but the last lies below R$ 60 per kWh. A shed value for a data center's IT load is a matter of contracts and reputation and is not estimated here, but no plausible value for a site whose sheet prices it at R$ 1 500 falls below R$ 60: the hall units and the bypass are chosen at every plausible shed value, and the stated value matters only to the fourth chiller.

The last switch, R$ 528.19, lies below the fourth chiller's median breakeven of R$ 662.61,

Table 24.7: The curve at next year's load, year 2 of the plan, site IT load 768 to 873 kW: EUE by the same 4 000 paired draws, its part from the state with every element up (capacity alone), and the hours short.

| Design [–] | EUE mean [kWh/yr] | 90 % band [kWh/yr] | All-up part [kWh/yr] | Hours short [h/yr] |
|---|---|---|---|---|
| as built | 137 567.2 | 19 885.4 to 564 902.1 | 104 845.9 | 2 642.0 |
| +A4+B4+C4+bypass | 1 584.1 | 86.7 to 6 666.3 | 16.0 | 30.1 |

and between them the full package is the least design in expectation while the chiller step pays on fewer than half of the draws. The two statements do not conflict: the expectation weighs the draws on which the chiller saves much, and the median counts draws. A design review that wants the chiller to pay on most draws, not just in expectation, needs a shed value above about R$ 663; one that minimizes the expected consequence needs about R$ 528.

### 24.3.5 The curve at next year's load

**Planner's question.** Where does the curve sit at next year's load? The question asks for the EUE of the as-built design and of the least design in year 2 of the plan, when the site IT load grows from 768 to 873 kW.

**General form.** The expected consequence of equation (2.23) at a later scenario node, its event probabilities those of the later load,

$$\mathcal{R}(a) = c_a + \sum_{e} C_e \, \Pr(e \mid \mathbf{y}, a),$$

with every event a state and month of year 2.

**Specialized form.** The question is equation (2.23) with the events of equation (24.2) at year 2's monthly loads $L_m$, the EUE of equation (24.3) on the same draws, its all-up part the term of the state with every element up.

**Building blocks.** The enumeration with year 2's loads; the same draws and designs.

**The year after.** Table 24.7 gives it. Each row is a design; the columns give the EUE's mean and 90 % band in year 2, the part of the mean from the all-up state, and the hours short.

*In year 2 the shortfall turns from availability to capacity.* As built, the site sheds an expected 137 567 kWh, of which 104 846, three quarters, comes from the state with every element up: the plant is short with nothing broken, and it is short 2 642 hours of the year in all. In the plan year, by Chapter 23, the all-up state carried 0.165 of the EUE. With the full package the EUE is 1 584.1 kWh, about eleven times its plan-year value of 138.5, and its all-up part is 16.0: the additions still buy availability, but the load is approaching what even the full package carries with everything running. The curve of Figure 24.1 is a plan-year curve. Every point on it moves up as the load grows, and the order of the frontier can change once capacity, not the loss of a unit, limits the short states.

## 24.4 The checks (step 5)

The answers of Section 24.3 rest on draws of three fits and two rate posteriors, on reading every design on the same draws, and on records generated from a stated truth; each is tested here.

### 24.4.1 The routes against each other

The exact posterior against the draws is the last pair of columns of Table 24.3. The expected consequence by the site's weights lies within 0.4 % of the draws' for every design: from 0.02 % for +A4+B4+C4 to 0.38 % for +A4+B4+bypass. For the as-built design the gap is R$ 42 654 a year against a Monte Carlo standard error of 362.3 kWh in the EUE, R$ 543 450 at the shed value, so under a twelfth of a standard error. For the least design the exact value is R$ 538 450 against R$ 537 944. No conclusion of Section 24.3 moves by more than the draws' own error when the exact posterior replaces the Laplace draws of the site. The halls' Laplace normals are not reweighted here, as in Chapter 23.

### 24.4.2 Common random numbers against independent draws

The case for reading every design on one set of draws is equation (24.8), and it is checked on the step that decides, the fourth chiller after both hall units, EUE(+A4+B4) − EUE(+A4+B4+C4). On the paired draws the difference has a standard deviation of 301.3 kWh a year. With the second design's draws shuffled, so that the two are read on independent draws, it has 690.6. The Monte Carlo standard error of the mean saving is 4.8 kWh a year paired against 10.9 unpaired: pairing is worth the same precision as about five times as many draws. The breakevens of Table 24.5 and the shares of draws of Table 24.6 are per-draw quantities of a difference and exist only on paired draws; on independent draws they would describe a comparison of two different plants.

### 24.4.3 The answers against the stated truth

The records are generated from a stated truth, and the last checks read every design at it. Table 24.8 collects them: each row is a design or a quantity, the columns its reported value, the band or route beside it, the truth, and the reading.

The truth falls inside the 90 % band of every design's EUE, and it places the designs in the same order. Its least design is the same full package, at R$ 440 267 a year against R$ 703 500 for the bypass without the chiller, and its frontier is the same six designs, with +A4 and +C4 dominated, as in the posterior. The decisive step agrees most closely: at the truth, the fourth chiller after the bypass saves $388.7 - 73.4 = 315.3$ kWh a year for its R$ 209 600, a breakeven of about R$ 665 per kWh, against the posterior's median of R$ 662.61. The posterior's mean EUE lies below the truth as built, 16 562 against 18 172, and above it for every design with both hall units; the truth sits at the 0.78 quantile of the as-built EUE by Chapter 23's check, inside its 90 % band.

## 24.5 The report (step 6)

As built, the site sits at the far, expensive end of the cost–EUE curve: an expected 16.6 MWh a year shed, worth R$ 24.8 M a year at the stated shed value, with a band from R$ 9.1 M to R$ 43.9 M. Two hall units, R$ 0.8 M installed, take the expected shed energy to 0.8 MWh a year, with median breakevens of about R$ 6 and R$ 15 per kWh. The strainer bypass pays above about R$ 58. The fourth chiller is the only marginal call: it pays above a median breakeven of about R$ 663 per kWh, band R$ 218 to R$ 2 045, so at the stated R$ 1 500 it is the right addition with probability 0.89. Built alone, it is off the frontier.

### 24.5.1 The answers and their bands

Table 24.9 gives the answers in the order of Table 24.2. Each row is a question; the second column gives the reported answer, and the third its band or the routes beside it.

Table 24.8: The checks of the risk–cost answers. EUE in kWh a year; reported values by 4 000 paired draws, bands 90 %; the truth is the enumeration at the stated truth of the records, with the same designs and prices.

| Quantity [unit] | Reported [–] | Band or beside [–] | Truth [–] | Reading |
|---|---|---|---|---|
| EUE, as built [kWh/yr] | 16 562.4 | 6 050.1 to 29 233.6 | 18 171.7 | inside the band |
| EUE, +B4 [kWh/yr] | 5 304.2 | 1 131.0 to 11 983.5 | 8 828.0 | inside |
| EUE, +A4 [kWh/yr] | 12 110.8 | 3 889.4 to 20 067.9 | 10 051.7 | inside |
| EUE, +A4+B4 [kWh/yr] | 806.2 | 428.0 to 1 532.7 | 657.4 | inside |
| EUE, +A4+B4+bypass [kWh/yr] | 535.4 | 159.2 to 1 263.9 | 388.7 | inside |
| EUE, +A4+B4+C4 [kWh/yr] | 409.6 | 276.9 to 713.9 | 342.4 | inside |
| EUE, full package [kWh/yr] | 138.5 | 17.0 to 449.8 | 73.4 | inside |
| EUE, +C4 [kWh/yr] | 16 167.8 | 5 747.5 to 29 079.2 | 17 855.5 | inside |
| Least design | full package | Pr = 0.887 | full package | agrees |
| $\mathcal{R}$ of the least [R$/yr] | 537 944 | exact 538 450 | 440 267 | truth below, inside the band |
| Frontier | six designs | +A4, +C4 dominated | same six | agrees |
| Fourth chiller's breakeven [R$/kWh] | 662.61 | 218.05 to 2 044.90 | 665 | agrees |
| Exact against draws, $\mathcal{R}$ | every design | within 0.4 % | – | below the draws' error |
| Paired against unpaired, sd | 301.3 | 690.6 unpaired | – | pairing narrows |

Table 24.9: The report on the site's design list, plan year at 15 % a year from 660 kW, stated prices annualized at 0.131 a year, stated shed value R$ 1 500 per kWh. Reported values by 4 000 paired draws, common to every design.

| Question | Answer | Band or routes beside |
|---|---|---|
| Where on the cost-versus-expected-shortfall curve does the design sit, and what is the band on that position? | at the far end: EUE 16 562 kWh/yr at zero price, $\mathcal{R}$ = R$ 24.84 M/yr; frontier as built, +B4, +A4+B4, +bypass, +A4+B4+C4, full package | EUE 6 050 to 29 234; $\mathcal{R}$ R$ 9.08 M to R$ 43.85 M, exact R$ 24.89 M; +A4, +C4 dominated |
| Which design has the least expected consequence, and how sure? | +A4+B4+C4+bypass, R$ 537 944/yr | exact R$ 538 450; least on a draw 0.887, +A4+B4+bypass 0.113; saving against as built R$ 24.31 M (8.72 to 42.62 M) |
| At what shed value does each step pay? | hall B unit R$ 6.48, hall A unit R$ 14.59, bypass R$ 58.13, fourth chiller R$ 662.61 per kWh (medians) | bands 2.77–16.50, 4.77–105.00, 52.15–65.14, 218.05–2 044.90; pays at R$ 1 500: 1.0000, 0.9995, 1.0000, 0.8872 |
| How does the choice move with the shed value? | as built, +B4, +A4+B4, +bypass, full package, switching at R$ 4.65, R$ 11.65, R$ 58.04, R$ 528.19 | least on a draw from 0.2445 (at R$ 5) to 0.9890 (at R$ 3 000) |
| Where does the curve sit at next year's load? | as built 137 567 kWh/yr, three quarters from the all-up state; full package 1 584 | as built 19 885 to 564 902; full package 86.7 to 6 666.3 |
| Which measurement could change the choice? | the chillers' unavailability: R$ 4 722/yr of R$ 5 901 at R$ 1 500 | at R$ 663: R$ 37 274 of R$ 38 571; every capacity entry 0 |

### 24.5.2 The design-review entry

The design review's entry ranks the additions on the cost–EUE frontier, not by their N+k counts. Counted by redundancy, the fourth chiller is the plant's N+1 and stands first on the list; read on the year's shed energy, it stands last, and alone it is dominated by a hall unit at a quarter of its price. The entry records the frontier in its order: hall B's fourth unit, hall A's fourth unit, the strainer bypass, the fourth chiller. It records that the first three are chosen in expectation at any shed value above R$ 60 per kWh, so that their choice does not rest on the stated R$ 1 500, and that the fourth chiller's does: it enters the least design in expectation above the switch point of R$ 528 per kWh, it pays on most draws only above its median breakeven of R$ 663, and at the stated value it is right with probability 0.89. It records that the curve is the plan year's, and that in year 2 the as-built site is short with nothing broken, so that the list must be read again against the load each year of growth.

**Common mistake**

Ranking additions by the redundancy count they restore. An N+1 count says how many units can be lost before an element is short; it does not say how much load is shed, by which element, or what removing it costs. The fourth chiller restores the count the sheet checks first, and buys under 400 kWh a year while hall B sheds thousands. Additions are ranked by the expected consequence they remove per real, on the frontier of price against expected unserved energy, and each step is read on the same draws as the design it improves.

## 24.6 The next measurement (step 7)

The report leaves one decision open, the fourth chiller, and asks of the plant's unknown which entry, if known before choosing, would change it. The question is a value of information: not which measurement narrows a band most, which Chapter 18 ranks, but which could change the design, and what that is worth.

### 24.6.1 The value of knowing one entry before choosing

**Planner's question.** Which measurement could change the choice? The question asks, for each entry of the unknown, how much less the best design would be expected to cost if that entry were known before the design is chosen.

**General form.** The design form of equation (2.16) with $u$ the minimum expected loss after the outcome: the expected value of sample information of Section 2.6.2, equation (2.53), the minimum expected loss now less its expectation after the outcome,

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

which for an exact reading of the quantity is the expected value of perfect information, an upper bound on what any measurement of it could be worth.

**Specialized form.** The question is equation (2.16), in its value-of-information reading (2.53), with the loss the expected consequence of equation (24.5) and the outcome an exact reading of one entry $z_k$ of $\boldsymbol{\theta}$,

$$\mathrm{EVPI}_k = \min_a \mathcal{R}(a) - \mathbb{E}_{z_k|\mathbf{y}}\Big[\min_a \mathbb{E}\big[L(a,\boldsymbol{\theta}) \mid \mathbf{y},\ z_k\big]\Big], \tag{24.16}$$

read on the common draws by sorting them on $z_k$ into 20 bins of equal count, taking the least design in each bin, and averaging the bins' least expected consequences. The full value, every entry known, is $\min_a \mathcal{R}(a) - \mathbb{E}[\min_a L(a,\boldsymbol{\theta}) \mid \mathbf{y}]$, the least design chosen on each draw.

Table 24.10: The value of knowing one entry before choosing, equation (24.16), over the frontier designs, by 20 equal-count bins of the 4 000 paired draws; at the stated shed value, where the least design's $\mathcal{R}$ is R$ 537 944, and at the fourth chiller's median breakeven, where it is R$ 421 978. Bins: those in which knowing the entry changes the least design.

| | $v$ = R$ 1 500/kWh | | $v$ = R$ 663/kWh | |
|---|---|---|---|---|
| Entry known [–] | Value [R$/yr] | Bins [of 20] | Value [R$/yr] | Bins [of 20] |
| hall B fraction $\beta_{B3}$ | 0 | 0 | 0 | 0 |
| hall A fraction $\beta_{A5}$ | 0 | 0 | 0 | 0 |
| chiller log-ratio $\kappa_{C2}$ | 0 | 0 | 0 | 0 |
| derate $d$ | 0 | 0 | 0 | 0 |
| hall units' unavailability $q_h$ | 0 | 0 | 0 | 0 |
| chillers' unavailability $q_c$ | **4 722** | 2 | **37 274** | 10 |
| every entry (full value) | 5 901 | – | 38 571 | – |

**Building blocks.** The loss of every frontier design on each draw; the draws' values of six candidate entries: the fractions of the unmeasured racks B3 and A5, the second chiller's log-ratio, the derate, and the two unavailabilities.

**The values.** Table 24.10 gives them at the stated shed value and at the fourth chiller's median breakeven. Each row is an entry; the columns give the value of knowing it a year and the number of the 20 bins in which knowing it changes the choice.

*Nothing about the halls changes the choice.* The fraction at rack B3, which Chapter 23 found to carry 0.319 of the variance of the as-built EUE, is worth nothing here: in every bin of its values the least design includes both hall units, because their breakevens lie two orders of magnitude below any plausible shed value. Neither is any capacity entry of the chillers, their log-ratio or their derate, worth anything to the choice. One entry is. Knowing the chillers' unavailability is worth R$ 4 722 a year at the stated shed value, 0.80 of the full value of R$ 5 901, and changes the choice in 2 of the 20 bins, those of lowest unavailability, where the chiller is not worth adding. At the fourth chiller's median breakeven it is worth R$ 37 274 of R$ 38 571 and changes the choice in half the bins. Capitalized at the recovery factor, the value at the stated shed value is about $4\,722/0.131 \approx$ R$ 36 000 paid once; at the breakeven, about R$ 285 000.

The contrast with Chapter 23's ranking is the point of the section. Chapter 23 ranks the entries by their share of the EUE's variance, and rack B3's fraction comes first; that is the measurement that would narrow the as-built EUE most. Here the question is the decision, and the decision among designs that all include the hall units does not turn on the hall's shed energy at all. A measurement can take a third of a band away and be worth nothing to the choice, as Section 2.6.2 warns.

### 24.6.2 The chillers' outage record

The chillers' unavailability is read from their outage record, four years of the three units, thirteen failures in twelve unit-years and 527.4 repair hours, through the gamma posteriors of equation (9.12). Its 90 % band, 0.00271 to 0.00967, spans a factor of three and a half, and the fourth chiller's saving spans it with it. The measurement that bears on the choice is therefore more of that record: further years of the site's own chillers, or failure and repair records of the same chiller model pooled from the manufacturer or from other sites, with a stated judgment of how similar those units are to these. Each narrows the rates' posteriors and moves the unavailability's band toward one side or the other of the fourth chiller's breakeven.

The estimation of failure and repair rates from such records, and the availability that follows from them, are derived in *Reliability Engineering by Derivation*, Chapters 6 and 8. The value of perfect information, R$ 4 722 a year at the stated shed value, bounds what any such record can be worth to this choice; a record that costs more than its capitalized value, about R$ 36 000, is not worth collecting for the choice alone.

### 24.6.3 The recommendation

Build the fourth units in hall B and hall A and the strainer bypass, in that order: they pay at any plausible shed value, and no measurement on the existing plant would change them. Decide the fourth chiller on the stated shed value and the chillers' outage record. At the stated R$ 1 500 per kWh it is the least-consequence choice and right on 0.89 of the draws; if the design review's shed value lies near or below R$ 663, collect more of the chillers' outage record before committing, since there the record decides the choice in half its outcomes. Do not build the fourth chiller ahead of the hall units. And read the curve again before year 2, when the site begins to be short with nothing broken.

### 24.6.4 What the analysis showed

The design sheet prices the site's shed energy at R$ 0.80 M a year and cannot place the additions; the plant's posterior prices it at R$ 24.8 M and places every one. Read as eight scenario nodes on one set of draws, the build-out traces a frontier whose order is the order in which the elements bind: hall B, then hall A, then the chillers and the pipe, then hall B again on two units down. Each addition hands the shortfall to the next element. The hall units are cheap and remove almost all of the shed energy; their breakevens lie below R$ 20 per kWh at the median, and their choice survives any plausible shed value. The fourth chiller, the plant's N+1 and the first addition a redundancy count would name, is off the frontier when built alone and marginal when built last, with a breakeven of R$ 663 per kWh whose band straddles the stated value. Reading every design on the same draws narrows the decisive difference by more than half. And the one measurement that could change the design is the chillers' outage record; the hall measurement that would narrow the EUE most is worth nothing to the choice.

### 24.6.5 Limits of the model

**The stated inputs.** The installed prices, the capital recovery factor and the shed value are stated, not estimated, and the answer's position depends on them. The hall units' choice does not: it holds for any shed value above about R$ 60 per kWh, and at the stated shed value for a hall unit price up to about ten times the stated one, since the upper ends of the hall steps' breakeven bands lie below R$ 105. The fourth chiller's does, and Sections 24.3.3 and 24.3.4 give its dependence on the shed value; a different chiller price moves its breakeven in proportion.

**One shed value for every kWh.** The loss is linear in the shed energy: a kWh shed in a two-hour event is priced like a kWh in a two-minute one. A shed value that grew with the event's length or depth, or a fixed cost per event, would weigh the deep, rare states more, and would favor the additions that remove them.

**The plan year.** The curve is read at the plan year's load. Section 24.3.5 shows that in year 2 three quarters of the as-built site's shed energy comes from the state with nothing broken, so the curve moves with the load, and a build-out chosen for the plan year is a first step of a longer plan whose later steps add capacity, not redundancy.

**Independent failures and a pinned supply.** The enumeration takes the elements to fail independently and pins the supply at 14°C. A common cause would add short states with several units down at once, which the additions treat differently; a short chiller plant that warmed the supply would join the chillers' and the halls' parts of the EUE, the chain of Chapter 21.

**New units like old ones.** Each addition takes its kind's failure and repair rates and a stated rating. A new unit with a burn-in period, or a fourth hall unit whose air path differs from its neighbors', would change the step it belongs to.

**The constructed record.** The night records, the operating record and the outage records are generated from a stated truth with stated errors, not taken from a site, and the comparison with that truth in Section 24.4.3 is a check of the method on its own model, not of the model against a real plant.

## 24.7 What to retain

**Definition 24.1** (Expected consequence of a design)**.** For a design variable $a$ with annualized price $c_a = \rho P_a$ and a stated shed value $v$, the expected consequence is $\mathcal{R}(a) = c_a + v\,\mathbb{E}[\mathrm{EUE}(a) \mid \mathbf{y}]$, the expected-consequence form of Chapter 2 with the events the enumerated states and months. A design is dominated when another costs no more and sheds no more, one strictly; the others form the frontier. Every design is read on common random numbers, and the breakeven shed value of a step, $v^* = (c_b - c_a)/(\mathrm{EUE}(a) - \mathrm{EUE}(b))$, is read on each draw and reported by its median and band, never by its delta form. The value of knowing an entry is $\min_a \mathcal{R}(a) - \mathbb{E}_{z_k}[\min_a \mathcal{R}(a \mid z_k)]$.

Four conclusions carry beyond the site. A design choice under uncertainty is a minimum of expected consequence over the design variable, and with a certain price and an uncertain shortfall it needs the shortfall's posterior and a stated value of the shortfall, nothing more. Additions are ranked on the frontier of price against expected unserved energy, and the frontier's order is the order in which the plant's elements bind; a redundancy count can name first an addition that a cheaper one dominates. The comparisons that decide are differences between designs, and they are read on common random numbers, or their variances are the sum of two large variances instead of one small one. And a stated input that cannot be estimated is handled by asking where it matters: the breakeven of each step, read on the draws, separates the choices that hold at any plausible value from the one that turns on it, and the value of information then names the entry whose measurement could turn it, which need not be the entry that narrows the shortfall most.

Table 24.11 lists the terms of the chapter with their expressions, units and the subsection that defines each.

Table 24.11: The terms of risk–cost tradeoffs in design and where each is defined. Units for the site.

| Term | Expression | Units | Meaning | Defined in |
|---|---|---|---|---|
| Expected unserved energy | $\sum_m h_m \sum_s \Pr(s \mid \mathbf{q}, a) \max(L_m - C_{s,m}, 0)$ | kWh/yr | shed energy of design $a$ | 24.1.1 |
| Annualized price | $c_a = \rho P_a$, $\rho = 0.131$ | R$/yr | price paid once, per year | 24.1.1 |
| Expected consequence | $\mathcal{R}(a) = c_a + v\,\mathbb{E}[\mathrm{EUE}(a) \mid \mathbf{y}]$ | R$/yr | what the choice minimizes | 24.1.1 |
| Dominated design | $c_b \le c_a$, $\mathbb{E}[\mathrm{EUE}(b)] \le \mathbb{E}[\mathrm{EUE}(a)]$, one strict | – | never the choice | 24.1.1 |
| Breakeven shed value | $v^* = \Delta c / \Delta\mathrm{EUE}$ on each draw | R$/kWh | step pays above it | 24.1.1 |
| Common random numbers | $\mathrm{Var}(D) = \mathrm{Var}_a + \mathrm{Var}_b - 2\,\mathrm{Cov}_{ab}$ | (kWh/yr)$^2$ | differences read draw by draw | 24.1.1 |
| *The site's designs* | | | | |
| State capacity | eq. (24.9), minimum over elements | kW | with the state's units running | 24.1.2 |
| Added units | $g_{h4} \sim \mathcal{N}(0, 0.15^2)$; $\kappa_n$ of eq. (9.7) | – | stated ratings | 24.1.3 |
| *Questions and answers* | | | | |
| Switch point | $\bar{v} = \Delta c / \Delta\mathbb{E}[\mathrm{EUE}]$ | R$/kWh | least design changes | 24.3.4 |
| Value of knowing an entry | $\min_a \mathcal{R}(a) - \mathbb{E}_{z_k}[\min_a \mathcal{R}(a \mid z_k)]$ | R$/yr | entry that could change the choice | 24.6.1 |
| Least design | full package, R$ 537 944/yr, $\Pr = 0.887$ | R$/yr | at R$ 1 500/kWh | 24.3.2 |
| Fourth chiller's breakeven | median 662.61, band 218.05 to 2 044.90 | R$/kWh | the one marginal step | 24.3.3 |

## 24.8 Exercises

**Exercise 24.1.** *Formulation.* Write the planner's question *should the fourth chiller be added to the two hall units and the bypass?* as an expected consequence over the site's unknown, say which entries it depends on most, and give the probability that the addition is right at the stated shed value. (It is the comparison of $\mathcal{R}$ for +A4+B4+bypass and +A4+B4+C4+bypass by equation (24.5): the step pays on a draw when $v\,[\mathrm{EUE}(+\mathrm{A4}+\mathrm{B4}+\mathrm{bypass}) - \mathrm{EUE}(\mathrm{full})] >$ R$ 209 600. By the pattern (24.10) the saving is the chillers' part, which reads the site's capacity entries and the chillers' unavailability; by Table 24.10 only the unavailability changes the choice. At R$ 1 500 per kWh the step pays on a share of the draws of 0.8872.)

**Exercise 24.2.** *Derivation.* Show that the shed value at which the least design switches from $a$ to $b$ is the ratio of the price difference to the difference of the mean EUEs, and explain why the fourth chiller's switch point lies below its median breakeven. (Each $\mathcal{R}(a;\ v) = c_a + v\,\mathbb{E}[\mathrm{EUE}(a)]$ is a line in $v$; two lines are equal at $v = (c_b - c_a)/(\mathbb{E}[\mathrm{EUE}(a)] - \mathbb{E}[\mathrm{EUE}(b)])$. For the fourth chiller that is $209\,600/396.8 \approx 528$. The breakeven of equation (24.7) is the same price over the saving on each draw; the saving is skewed to the right, its mean above its median, so the reciprocal's median lies above the price over the mean saving. The median breakeven is R$ 662.61 per kWh.)

**Exercise 24.3.** *Derivation.* From equation (24.8), show that pairing two designs on the same draws narrows the standard error of their mean difference whenever their EUEs are positively correlated, and give the standard errors for the step from +A4+B4 to +A4+B4+C4. (The variance of the difference is the sum of the variances less twice the covariance; on independent draws the covariance term is zero, so a positive covariance makes the paired variance the smaller, and the standard error of a mean of $N$ differences is the standard deviation over $\sqrt{N}$. The

paired standard deviation is 301.3 kWh a year against 690.6 unpaired, and the standard errors of the mean saving are 4.8 against 10.9.)

**Exercise 24.4.** *Application.* A reviewer proposes to build the fourth chiller first, as the plant's N+1, and to defer the hall units. Say what the curve answers. (The fourth chiller alone is dominated: +B4 costs R$ 52 400 a year against R$ 209 600 and sheds 5 304.2 kWh a year against 16 167.8. Alone the chiller removes 394.6 kWh a year, because hall B still limits 0.704 of the shed energy. The expected consequence of the chiller alone is R$ 24.46 M a year, and that of the hall-B unit R$ 8.01 M.)

**Exercise 24.5.** *Application.* The design review revises the shed value to R$ 300 per kWh. Name the least design, its expected consequence, and how sure it is, and say which measurement would then matter. (Below the fourth chiller's switch point of about R$ 528, the least design is +A4+B4+bypass, without the chiller; the chiller step's saving reads the chillers' unavailability, so the chillers' outage record remains the measurement that bears on it. The expected consequence is R$ 281 132 a year, least on a share of the draws of 0.8795.)

# Appendix A

# Atlas of planning questions

This appendix can be read independently of the chapters. Each row is a question a planner asks. The first column gives the question in the planner'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. The rows grow as chapters are added.

| Planner's question | General form | Mathematical form | Category | Computation | Answered in |
|---|---|---|---|---|---|
| *Chapter 1. Planned Plants as Graphs* | | | | | |
| What is the hall air temperature at today's load? | $\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$, $\mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y})$ | $p(T \mid \mathbf{y})$, $T = z_1$ | 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 |
| Will the hall carry its 60 kW design load inside 27.5°C? | $\Pr(f(\mathbf{z}) < f^* \mid \mathbf{y})$ | $\Pr(T_2 < 27.5^\circ\mathrm{C} \mid \mathbf{y})$, $T_2 = T_c + (Q_2/G_{\mathrm{ref}})\, \mathrm{e}^{-z_2}$ | prediction under a threshold | posterior | 1.3.2 |
| Would a measurement before the sign-off 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 | preposterior | 1.6.1 |
| *Chapter 2. The Language of Capacity 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 over the coming peak? | (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 |

| Planner's question | General form | Mathematical form | Category | Computation | Answered in |
|---|---|---|---|---|---|
| Is it rising, and when does it cross its limit? | (2.21) | $\Pr(\gamma > 0 \mid \mathbf{y})$; $\Pr(t^* \le s \mid \mathbf{y}) = \Pr(f_s(\mathbf{z}) \ge f^* \mid \mathbf{y})$ | 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 |
| 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 built 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 |
| Would the hall stay below 28°C at a load it has not yet carried? | $\Pr(f(\mathbf{z}) \ge f^* \mid \mathbf{y})$ at a scenario node | $\Pr(T^*(Q^*) \le T_{\max} \mid \mathbf{y})$, $T^* = T_c + Q^* z_1$, with $f = -T^*$ | prediction under a threshold | posterior | 4.3.8 |
| *Chapter 5. Thermal Headroom* | | | | | |
| How much load can hall B take before an inlet crosses the limit, and how sure is that number? | $\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$, quantiles | $Q^\dagger$ and $H = Q^\dagger - Q_B$: median and 90 % band | estimation | posterior | 5.3.1 |
| Which rack binds first? | $\Pr(f(\mathbf{z}) \ge f^* \mid \mathbf{y})$ over a family | $\Pr(\arg\min_k Q_k^\dagger = j \mid \mathbf{y})$, $j = \mathrm{B1}, \ldots, \mathrm{B5}$ | prediction under a threshold | posterior | 5.3.2 |
| Can hall B carry its 340 kW design load with 95 % assurance? | $\Pr(f(\mathbf{z}) \ge f^* \mid \mathbf{y}) \ge \kappa$ | $\Pr(Q^\dagger \ge 340\,\mathrm{kW} \mid \mathbf{y})$ against $\kappa = 0.95$ | prediction under a threshold | posterior | 5.3.3 |
| What is the 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 | 5.3.4 |

| Planner's question | General form | Mathematical form | Category | Computation | Answered in |
|---|---|---|---|---|---|
| *Chapter 6. Load Growth Scenarios* | | | | | |
| If the IT load grows 15 % a year, what is the probability the plant is still inside its envelope in year three? | $\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y})$ at a scenario node | $\Pr(Q^\dagger \geq Q_3 \mid \mathbf{y})$, $Q_3 = 456.26\,\mathrm{kW}$ | prediction under a threshold | posterior | 6.3.1 |
| And in each year before and after? | the same, per node | $\Pr(Q^\dagger \geq Q_t \mid \mathbf{y})$, $t = 0, \ldots, 5$ | prediction under a threshold | posterior | 6.3.2 |
| In what year does the hall leave its envelope? | $\Pr(f(\mathbf{z}) \leq f^* \mid \mathbf{y})$, quantiles | $t^\dagger$: median and 90 % band | estimation | posterior | 6.3.3 |
| How much does the answer depend on the growth rate assumed? | the same at stated inputs | $\Pr(Q^\dagger \geq Q_3 \mid \mathbf{y})$ at $\gamma = 12,\ 15,\ 18\,\%/\mathrm{yr}$ | prediction under a threshold | posterior | 6.3.4 |
| *Chapter 7. Design-Day Assessment* | | | | | |
| On the design-day ambient, with every unit running, what is the probability of a violation anywhere? | $\Pr(\max_k f_k(\mathbf{z}) > f^* \mid \mathbf{y})$ | $\Pr(\max_k T_k^* > 27^\circ\mathrm{C} \mid \mathbf{y})$ | prediction under a threshold | posterior | 7.3.2 |
| How warm is each inlet on the design day, and how sure? | moments, $\mathbb{E}[f]$ and $\mathrm{Var}(f)$ | $T_k^*$, its sd, and $\Pr(T_k^* > 27^\circ\mathrm{C} \mid \mathbf{y})$ per rack | estimation | posterior | 7.3.1 |
| Which rack violates, if one does? | $\Pr(A \mid B)$ | $\Pr(T_k^* > 27^\circ\mathrm{C} \mid V,\ \mathbf{y})$ | prediction under a threshold | posterior | 7.3.3 |
| When does the joint answer differ from the worst rack alone? | the joint tail against its bounds | joint, worst marginal, union bound, independence product; the inlets' correlations | prediction under a threshold | posterior | 7.3.4, 7.3.5 |
| *Chapter 8. Load Diversity and Coincident Peaks* | | | | | |
| The racks never all peak at once; what is the probability the coincident peak exceeds the plant's capacity? | $\Pr(y_{\mathrm{new}} > f(\mathbf{z}) \mid \mathbf{y})$, a future value against a derived quantity | $\Pr(S^* > Q^\dagger \mid \mathbf{y})$ | prediction under a threshold | posterior | 8.3.1 |
| How correlated are the racks? | $\Pr(z_i \leq z^* \mid \mathbf{y})$, quantiles | $\rho$: MAP and 90 % band | estimation | posterior | 8.3.2 |
| What is the coincidence factor at the design population? | a derived quantity of the state | 99th-percentile peak over the sum of the racks' 99th percentiles | estimation | posterior | 8.3.3 |
| How much does the covariance move the answer, and how much the capacity band? | conditional probabilities; the spread of a probability | $\Pr(S^* > Q^\dagger \mid \rho, \mathbf{y})$ for $\rho = 0$ to 1; the spread of $\Pr(S^* > Q^\dagger \mid \boldsymbol{\theta}, \mathbf{z})$ with $\rho$ or $Q^\dagger$ known | estimation | posterior | 8.3.4 |
| *Chapter 9. Redundancy Level* | | | | | |
| N+1 or N+2? | $\arg\min_a \mathbb{E}[L(a, \boldsymbol{\theta}) \mid \mathbf{y}]$ | annual shed cost saved, $200\,\mathrm{h} \times \mathrm{R\$}\,1\,500/\mathrm{kWh} \times (E_a - E_{a+1})$, against $c_u$, for N to N+1 and N+1 to N+2 | design | expected consequence | 9.3.3 |

| Planner's question | General form | Mathematical form | Category | Computation | Answered in |
|---|---|---|---|---|---|
| What is each drawing's service availability, and how sure? | $\Pr(f \geq f^* \mid \mathbf{y})$ over an enumeration | $\mathbb{E}[\sum_{\mathbf{s}} P(\mathbf{s})\,\mathbf{1}(\mathrm{Cap}(\mathbf{s}) \geq 1\,000\,\mathrm{kW}) \mid \mathbf{y}]$, with the 90 % band of $U_a$ | prediction under a threshold | posterior | 9.3.2 |
| How much capacity does each carry, all up and in each state? | moments | $H_a$ with all units up, median and band; $\mathrm{Cap}(\mathbf{s})$ by the number of units down | estimation | posterior | 9.3.1 |
| What is the availability band made of? | conditional variance | $\mathrm{sd}(U_{\mathrm{N}+1})$ with $q$ known and with the capacity known | estimation | posterior | 9.3.4 |
| *Chapter 10. Unit Sizing* | | | | | |
| Two 600 kW units or three 400 kW? | $\arg\min_a \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]$ | $\mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]$ for both drawings, and $\Pr(D < 0 \mid \mathbf{y})$ | design | expected consequence | 10.3.4 |
| Which is more available, and at what load? | $\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y})$ over an enumeration | $\mathbb{E}[U(Q^*) \mid \mathbf{y}]$ for $Q^* = 682$ to 1 000 kW; $\Pr(U_{3\times400} < U_{2\times600} \mid \mathbf{y})$ | prediction under a threshold | posterior | 10.3.1 |
| What headroom does each leave? | $\Pr(f(\mathbf{z}) \leq f^* \mid \mathbf{y})$, quantiles | $H_{\mathrm{all}}$ and $H_{-1}$ at 1 000 kW: median and 90 % band | estimation | posterior | 10.3.2 |
| How much does efficiency cost? | $\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$ and quantiles of a derived quantity | $p_E\,(E_{3\times400} - E_{2\times600})$: mean and 90 % band | estimation | posterior | 10.3.3 |
| Is the difference real? | $\Pr(f(\mathbf{z}) > f^* \mid \mathbf{y})$ | $\Pr(D < 0 \mid \mathbf{y})$, with the 90 % band of $D$ | prediction under a threshold | posterior | 10.3.5 |
| *Chapter 11. Expansion Timing* | | | | | |
| When does the next unit have to be ordered so that it is in before the probability of a shortfall passes 5 %? | first crossing of a tail along a scenario chain | $\min\{m : P_m \geq 0.05\} - \ell$ | prediction under a threshold | posterior | 11.3.2 |
| How does the shortfall probability grow along the path? | $\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y})$ at each node | $P_m = \mathbb{E}[\pi_m \mid \mathbf{y}]$, month by month | prediction under a threshold | posterior | 11.3.1 |
| How sure is the date? | quantiles of a derived quantity | $m^\dagger(\mathbf{z}, q)$: median and 90 % band | estimation | posterior | 11.3.3 |
| What sets the date? | moments of a derived quantity | the all-up capacity's band against the plateau of one unit out | estimation | posterior | 11.3.4 |
| *Chapter 12. Oversizing and Stranded Capacity* | | | | | |
| What is the probability that the plant as drawn is more than 30 % oversized at the ten-year load? | $\Pr(f \geq f^* \mid \mathbf{y})$, the upper tail of a ratio | $\Pr(R > 1.30 \mid \mathbf{y})$ | prediction under a threshold | posterior | 12.3.1 |
| And the probability it is undersized? | $\Pr(f < f^* \mid \mathbf{y})$ | $\Pr(R < 1 \mid \mathbf{y})$ | prediction under a threshold | posterior | 12.3.2 |
| What is the answer made of? | $\mathrm{Var}(f \mid \mathbf{y})$, split by input | $\mathrm{Var}(\ln R) = \mathrm{Var}(\ln C) + \mathrm{Var}(\ln Q_{10})$ | estimation | posterior | 12.3.3 |

| Planner's question | General form | Mathematical form | Category | Computation | Answered in |
|---|---|---|---|---|---|
| How much capacity is stranded, and what does it cost? | $\mathbb{E}[L(a,\cdot) \mid \mathbf{y}]$ | $\mathbb{E}[S \mid \mathbf{y}]$, annualized | estimation | expected consequence | 12.3.4 |
| *Chapter 13. Binding Constraints* | | | | | |
| Which element binds first as the load grows: the chiller, the pump, the pipe, or the power feed? | $\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y})$ over a family | $\Pr(e^\dagger(\mathbf{z}) = j \mid \mathbf{y})$, $j$ the chillers, the pipe, the pump, the feed | prediction under a threshold | posterior | 13.3.2 |
| At what load does each bind, and how sure? | $\Pr(f(\mathbf{z}) \leq f^* \mid \mathbf{y})$, quantiles | $Q_e^\dagger$ per element and $Q^\dagger = \min_e Q_e^\dagger$: median and 90 % band | estimation | posterior | 13.3.1 |
| What is the probability that the order flips? | $\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y})$ over a family | $\Pr(\pi(\mathbf{z}) \neq \pi(\hat{\mathbf{z}}) \mid \mathbf{y})$, and each order's probability | prediction under a threshold | posterior | 13.3.3 |
| What decides between the pipe and the chillers? | $\mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y})$, split by entry | the budget of $\mathbf{g}_D^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}_D$, $D = Q^\dagger_{\text{pipe}} - Q^\dagger_{\text{ch}}$ | estimation | posterior | 13.3.4 |
| *Chapter 14. Capacity Importance Measures* | | | | | |
| If I could upgrade one element, which one buys the most headroom? | $\arg\min_a \mathbb{E}[L(a,\mathbf{z}) \mid \mathbf{y}]$ | $\arg\max_e \mathbb{E}[G_e \mid \mathbf{y}]$, and $\mathbb{E}[G_e \mid \mathbf{y}]/p_e$ | design | expected consequence | 14.3.2 |
| What is each element's importance measure? | $\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$ | $I_e = \mathbb{E}[\partial H/\partial r_e \mid \mathbf{y}]$: $\Pr(e_1 = e)$ times the gradient | estimation | posterior | 14.3.1 |
| Why do the gradient and the upgrade disagree? | moments of a minimum | $I_e\,\Delta_e - \mathbb{E}[G_e \mid \mathbf{y}]$ against the distance to the next element | estimation | posterior | 14.3.3 |
| What do two upgrades buy together? | $\arg\min_a \mathbb{E}[L(a,\mathbf{z}) \mid \mathbf{y}]$ | $\mathbb{E}[G_{e,e'} \mid \mathbf{y}]$ for the pipe with the chillers and with the feed | design | expected consequence | 14.3.4 |
| *Chapter 15. Shared Paths and Common-Mode Limits* | | | | | |
| Two halls share one pipe; what is the probability that the shared pipe, not the halls' own units, is what limits them? | $\Pr(\ell = \text{pipe} \mid \text{shortfall},\ \mathbf{y})$ | $\mathbb{E}[P_{L,\text{pipe}}]/\mathbb{E}[P_L]$, at $L = 660$, 759, 820, 860 kW | prediction under a threshold | posterior | 15.3.2 |
| What are the minimal cut sets, today and in a year? | the failing states of the enumeration | sets $S$ with $\mathrm{Cap}(\mathbf{s}_S;\ \hat{\boldsymbol{\theta}}) < L$, orders one and two, at 660 and 759 kW | diagnosis | posterior | 15.3.3 |
| How likely is a shortfall at all? | $\Pr(\mathrm{Cap} < L \mid \mathbf{y})$ | $\mathbb{E}[P_L(\boldsymbol{\theta}) \mid \mathbf{y}]$ by load, with its 90 % band | prediction under a threshold | posterior | 15.3.1 |
| How important is each element? | share of the shortfall from states with it down | $I_e = \mathbb{E}[P_L^{(e)}]/\mathbb{E}[P_L]$ per element, by load | estimation | posterior | 15.3.4 |
| *Chapter 16. As-Built Margin* | | | | | |
| The design says 20 % margin; what does the as-built model say the margin is? | $\Pr(f(\mathbf{z}) \leq f^* \mid \mathbf{y})$, quantiles | $m$: median and 90 % band | estimation | posterior | 16.3.1 |

| Planner's question | General form | Mathematical form | Category | Computation | Answered in |
|---|---|---|---|---|---|
| What is the probability the real margin is under 10 %? | $\Pr(f(\mathbf{z}) < f^* \mid \mathbf{y})$ | $\Pr(m < 0.10 \mid \mathbf{y})$, with the levels 0.15 and 0 | prediction under a threshold | posterior | 16.3.2 |
| What is the margin made of? | $\mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y})$, split by entry | each chiller's $\mathrm{Cap}_i^*$ with its band; the budget of $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$ | estimation | posterior | 16.3.3 |
| What does the drawn economizer add, and when in the year is the margin thinnest? | moments; a scenario chain over months | $m_e$, median and band; $m_k$ for each month | estimation | posterior | 16.3.4 |
| *Chapter 17. Variance Budgets for Design Parameters* | | | | | |
| The conductance of the economizer path is a handbook guess; how much of the headroom band does that one guess carry? | $\mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y})$, split by entry | $g_e$'s share of $\mathrm{Var}(H_c \mid \mathbf{y})$ and of $\mathrm{Var}(H \mid \mathbf{y})$ | estimation | posterior | 17.3.1 |
| What does each entry carry? | $\mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y})$, split by entry | $S_i = \mathrm{Var}(\mathbb{E}[H \mid z_i])/\mathrm{Var}(H)$, all eight entries | estimation | posterior | 17.3.2 |
| What would the band be with the guess pinned down? | $\mathrm{Var}(f(\mathbf{z}) \mid z_i, \mathbf{y})$ | $H$'s band with $g_e$ known; $\mathbb{E}[\mathrm{Var}(H \mid g_e)]$ | estimation | posterior | 17.3.3 |
| How much does a better handbook buy? | $\Pr(f(\mathbf{z}) \le f^* \mid \mathbf{y})$, quantiles | the 90 % bands of $H$ and $H_c$ at handbook factors 1.4, 1.2, 1.1 | estimation | posterior | 17.3.4 |
| *Chapter 18. Measurements Before Commitment* | | | | | |
| Which test or survey on the existing plant would most narrow the build decision? | $\mathbb{E}_{\tilde{y}_a \mid \mathbf{y}}[u]$, $u$ an acceptance indicator | $S_a = \Pr(p_a \notin [0.1,\ 0.9] \mid \mathbf{y})$, ranked, the width beside | design | preposterior | 18.3.2 |
| How much does each narrow the deciding quantity? | $\mathbb{E}_{\tilde{y}_a \mid \mathbf{y}}[u]$, $u$ a spread | width of the 90 % band of $Q^\dagger$ under $\mathcal{N}(\hat{\mathbf{z}}, \boldsymbol{\Sigma}_a)$, and the share removed | estimation | preposterior | 18.3.1 |
| Is narrowing the band the same as settling the decision? | the two measures of one candidate side by side | width removed against $S_a$, by candidate | design | preposterior | 18.3.3 |
| What is each worth per real? | $\mathcal{R}(a) = c_a + \sum_e C_e \Pr(e \mid \mathbf{y}, a)$ | $S_a$ per R$ 10 thousand of $c_a$ | design | expected consequence | 18.3.4 |
| *Chapter 19. The Value of a Pilot* | | | | | |
| Would a one-hall pilot change the decision on the full build? | preposterior of a decision, $\mathbb{E}_{\tilde{y}_a \mid \mathbf{y}}[u]$ | $\Pr(\operatorname{sign} \bar{W}_a \neq \operatorname{sign} \bar{W})$ for the pilot | design | preposterior | 19.3.2 |
| What is the decision now, and how sure? | $\arg\min_a \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]$ | $\bar{W} = \mathbb{E}[W \mid \mathbf{y}]$; $\Pr(W > 0 \mid \mathbf{y})$; $g_e^\dagger$ | design | expected consequence | 19.3.1 |
| Is the pilot worth its cost? | EVSI, $\bar{L}_{\text{now}} - \bar{L}_{\text{after}}$ | $V_a = \mathrm{EVSI}_a/\rho$ against the pilot's net cost | design | expected consequence | 19.3.3 |
| Is there a cheaper way to the same answer? | preposterior, over candidates | flips and $V_a$ for the factory test, both, and a perfect reading (EVPI) | design | preposterior | 19.3.4 |

| Planner's question | General form | Mathematical form | Category | Computation | Answered in |
|---|---|---|---|---|---|
| *Chapter 20. Headroom with the Largest Unit Out* | | | | | |
| With the largest unit out, what is the probability the remaining plant holds the inlets at the design load? | $\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y})$ | $\Pr(Q_L^{\dagger} \geq Q_d \mid \mathbf{y})$, $u^{\star} = \arg\max_u G_u$ out, $T_s^* = 14^{\circ}\mathrm{C}$ | prediction under a threshold | posterior | 20.3.1 |
| Is the joint tail the product of the racks' tails? | joint against marginal tails | $\Pr(\cap_k \{T_k^* \leq 27^{\circ}\mathrm{C}\})$ against $\prod_k \Pr(T_k^* \leq 27^{\circ}\mathrm{C})$ | prediction under a threshold | posterior | 20.3.2 |
| What supply would restore it, and what does that cost the chillers? | $\Pr(f(\mathbf{z}) \leq f^* \mid \mathbf{y})$, quantiles | $T_s^{\dagger} = T_{\max} - \max_k a_k (Q_d + Q_o)$; the chillers' binding load at $T_s^{\dagger}$ | design | posterior | 20.3.3 |
| What does a fourth unit buy? | $\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y})$ under an alternative | the first question with four units, the largest of four out | design | posterior | 20.3.4 |
| *Chapter 21. Cascading Overloads* | | | | | |
| If one unit trips, does the load shift onto the others push a second one past its envelope? | $\Pr(f \geq f^* \mid \mathbf{y})$ along a scenario chain, a conditional probability | $\mathbb{E}[\bar{P} \mid \mathbf{y}]$, and the chiller term alone, at 660, 700, 759 kW on the design day | prediction under a threshold | posterior | 21.3.1 |
| Which first trip matters? | the same, for each first event | $\Pr(E_i \mid \mathbf{y})$ for each $i$; $\Pr(\mathrm{Cap}_{C2} = \max_j \mathrm{Cap}_j \mid \mathbf{y})$ | prediction under a threshold | posterior | 21.3.2 |
| How often a year? | moments, $\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$ | $\mathbb{E}[N_2 \mid \mathbf{y}]$ over the plan year, with its band; the design hours beside | estimation | posterior | 21.3.3 |
| Does the hall side cascade? | $\Pr(f \geq f^* \mid \mathbf{y})$ along the hall-unit chain | $\Pr(\max_x Q_x^* > Q_{\mathrm{sheet}})$ and $\Pr(\max_k T_k^* > 27^{\circ}\mathrm{C})$, a unit out | prediction under a threshold | posterior | 21.3.4 |
| What does Chapter 9's fourth chiller do to the chain? | the same under a design alternative | the first question with four chillers; $\Pr(m_B < 0 \mid \mathbf{y})$ | prediction under a threshold | posterior | 21.3.5 |
| *Chapter 22. Concurrent Maintainability* | | | | | |
| Can any one unit be taken out for service in any month of the year without a violation? | $\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y})$, the minimum over scenarios | $P_{\min} = \min_{u,m} P_{um}$ and its arg min | prediction under a threshold | posterior | 22.3.1 |
| When can each unit come out? | $\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}) \geq \kappa$, over a family | $W_u = \{m : P_{um} \geq 0.95\}$ for each unit | prediction under a threshold | posterior | 22.3.2 |
| What do the named remedies buy? | the same, at a design alternative | $P_{um}$ and $W_u$ with the four remedies built | design | posterior | 22.3.3 |
| What if a second unit fails during the service? | $p(y_{\mathrm{new}} \mid \mathbf{y})$, the predictive | $P_{um}^{(2)}$ with the rates' posterior, $\tau = 72\,\mathrm{h}$ | prediction under a threshold | posterior | 22.3.4 |
| *Chapter 23. Expected Unserved Energy* | | | | | |
| How many kWh of IT load would be shed per year, in expectation, under this design? | $\mathbb{E}[L(\mathbf{z}) \mid \mathbf{y}]$ and its quantiles | $\mathbb{E}[\mathrm{EUE} \mid \mathbf{y}]$, its median and 90 % band | estimation | expected consequence | 23.3.1 |

| Planner's question | General form | Mathematical form | Category | Computation | Answered in |
|---|---|---|---|---|---|
| How many hours is the site short? | the same, counting hours | $\mathbb{E}[\text{LOLE} \mid \mathbf{y}]$, its median and band | estimation | expected consequence | 23.3.2 |
| Is it capacity or availability? | the expectation of each part of a sum | $\mathbb{E}[\text{EUE}_{\text{cap}}]$ against $\mathbb{E}[\text{EUE}_{\text{av}}]$; by limiting element, element down and month | estimation | expected consequence | 23.3.3 |
| What is it worth? | $\mathcal{R} = \sum_e C_e \Pr(e \mid \mathbf{y})$ | $v\,\mathbb{E}[\text{EUE} \mid \mathbf{y}]$, $v = \text{R\$}\,1\,500/\text{kWh}$ | design | expected consequence | 23.3.4 |
| *Chapter 24. Risk–Cost Tradeoffs in Design* | | | | | |
| Where on the cost-versus-expected-shortfall curve does the design sit, and what is the band on that position? | expected consequence over the design variable, eq. (2.23) | $(c_a,\ \mathbb{E}[\text{EUE}(a) \mid \mathbf{y}])$ with its band; the frontier | design | expected consequence | 24.3.1 |
| Which design has the least expected consequence, and how sure? | $\arg\min_a$ $\mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]$ | $a^* = \arg\min_a \mathcal{R}(a)$; $\Pr(a^* \text{ least on a draw})$ | design | expected consequence | 24.3.2 |
| At what shed value does each step pay? | the breakeven of an increment | $v^*_{a\to b}$ on each draw, its median and band | design | expected consequence | 24.3.3 |
| How does the choice move with the shed value? | the minimizer against a stated input | $\arg\min_a \mathcal{R}(a)$ for $v$ from 1 to 3 000 R$/kWh | design | expected consequence | 24.3.4 |
| Where does the curve sit at next year's load? | the same, at a later scenario node | $\mathbb{E}[\text{EUE}(a) \mid \mathbf{y}]$ at 768 to 873 kW | estimation | expected consequence | 24.3.5 |
| Which measurement could change the choice? | expected value of perfect information, per entry, eq. (2.53) | $\min_a \mathcal{R}(a) - \mathbb{E}_{z_k}[\min_a \mathcal{R}(a \mid z_k)]$ | design | preposterior | 24.6.1 |