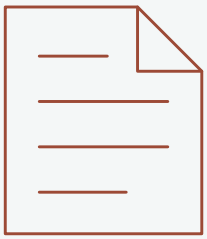
LESSON BOOK



Formulating Probabilistic Engineering Questions

From physical systems and uncertain evidence to decisions, experiments and risk

Draft · 310 pages

Ricardo Marquez, Ph.D.

ricardo-marquez.com

Formulating Probabilistic Engineering Questions

From Physical Systems and Uncertain Evidence to Decisions, Experiments, and Risk

Ricardo Marquez

November 2, 2026 (draft for review)

How to read this book

This book teaches one ability: taking an engineering question stated in ordinary language and writing down the mathematically precise probabilistic problem it asks. The book does not teach how to solve that problem. A reader will not compute the answers here, because the mechanics of solving are covered extensively well in other textbooks, and which technique applies depends on the kind of problem that has been written down. The subject of this book is what comes before the technique.

The place of the subject can be stated with four levels. The first is the engineering intention, what someone is trying to find out or decide. The second is the mathematical question, an expression that says exactly what would answer it. The third is the problem class, the kind of mathematical object that expression turns out to be. The fourth is the solution strategy. Most technical textbooks concentrate on the third and fourth levels. This book concentrates on the first two and teaches enough of the third for the reader to recognize what has been built and to hand it to the right technique.

By the end of the book, a reader should be able to take an unfamiliar engineering question, recognize the intention behind it, identify the uncertain quantities, formulate the relevant events and conditions, express the question mathematically, construct a probabilistic model capable of answering it, and name the problem class. The reader should also be able to recognize when a question is ambiguous, incomplete or unanswerable without further assumptions, and to say what is missing. Such a reader can sit in a meeting with a reliability engineer, a utility operator or a data center consultant, hear an operational question, and write its specification. Software, numerical methods or a specialist may still be needed to evaluate that specification. The reader will know what problem needs solving, what information it requires, and what the answer means.

One procedure runs through every lesson, the seven-step formulation method. Lesson 1 introduces it, and its product is a formulation sheet: one page with the seven steps as fields, the engineering intention at the top and the problem class at the bottom. A blank sheet is in Appendix C.

1. State the engineering question in one unambiguous sentence.
2. Identify the target quantity: a state, parameter, event, future outcome, or decision.
3. Specify what is known: observations, assumptions, operating conditions, and prior information.
4. Define the event or objective mathematically, including thresholds and time horizons.
5. Write the probability or expectation that directly answers the question.
6. Identify the model required to evaluate that expression: variables, physical relations, dependencies, and evidence.
7. Check the formulation: does the expression really answer the original engineering question?

The lessons work on three recurring systems rather than on a new system each time: a two-hall data center cooling plant, which leads; a small electrical network; and a battery string.

Their facts are collected in Appendix D. Returning to a known system with a new question is the teaching device. One physical system supports many questions, and the question is not determined by the system. A single probabilistic model of a physical network can answer many different queries. TERRA is the author's implementation of that kind of model; the lessons concern the queries, and none of them requires it.

The twenty-five lessons share one shape: the engineering question in words; its formulation; the competing formulations, with what each of them answers instead; a table of further questions on one system; the exercises; and what to retain. From Lesson 10 on, a section of worked problems stands before the exercises, each a short statement and its compact run. The exercises are the defining feature of the book. They come in four kinds: formulating a realistic engineering statement, reversing an expression into a sentence, critiquing and repairing a wrong formulation, and rapid one-line questions, because the promise is the ability to hear a question at the pace of speech. From Part III on, a Part closes with a fifth kind, a machine set of ten problems, each solved as a compact run. Some questions are intentionally ambiguous, and the right answer to those states what is missing and what each reading would give. Solutions are in Appendix A. They argue, and they name the common wrong formulations. The rubric in Appendix B scores a formulation against the seven steps and the recurring failure classes.

Contents

How to read this book **ii**

Notation and vocabulary **viii**

I The Language of Engineering Uncertainty **1**

1 What Exactly Is the Engineering Question? **2**
1.1 The engineering question . . . 2
1.2 The formulation . . . 3
1.3 The competing formulations . . . 9
1.4 One system, many questions . . . 12
1.5 The exercises . . . 13
1.6 What to retain . . . 14

2 From Words to Events **15**
2.1 The engineering question . . . 15
2.2 The formulation . . . 15
2.3 The competing formulations . . . 19
2.4 One system, many questions . . . 21
2.5 The exercises . . . 22
2.6 What to retain . . . 23

3 What Is Known, Unknown, and Conditioned Upon? **24**
3.1 The engineering question . . . 24
3.2 The formulation . . . 24
3.3 The competing formulations . . . 29
3.4 One system, many questions . . . 29
3.5 The exercises . . . 30
3.6 What to retain . . . 31

4 Defining Failure, Success, and Acceptable Performance **32**
4.1 The engineering question . . . 32
4.2 The formulation . . . 32
4.3 The competing formulations . . . 39
4.4 One system, many questions . . . 41
4.5 The exercises . . . 42
4.6 What to retain . . . 43

II Translating Questions into Probability 44

5 Threshold and Exceedance Questions 45
5.1 The engineering question . . . 45
5.2 The formulation . . . 45
5.3 The competing formulations . . . 52
5.4 One system, many questions . . . 53
5.5 The exercises . . . 54
5.6 What to retain . . . 55

6 Joint and Conditional Events 57
6.1 The engineering question . . . 57
6.2 The formulation . . . 57
6.3 The competing formulations . . . 64
6.4 One system, many questions . . . 66
6.5 The exercises . . . 66
6.6 What to retain . . . 68

7 Questions About Uncertain Parameters 69
7.1 The engineering question . . . 69
7.2 The formulation . . . 69
7.3 The competing formulations . . . 77
7.4 One system, many questions . . . 78
7.5 The exercises . . . 78
7.6 What to retain . . . 80

8 Questions About Future Performance 81
8.1 The engineering question . . . 81
8.2 The formulation . . . 81
8.3 The competing formulations . . . 89
8.4 One system, many questions . . . 90
8.5 The exercises . . . 91
8.6 What to retain . . . 92

9 Complex Questions with Multiple Conditions 93
9.1 The engineering question . . . 93
9.2 The formulation . . . 93
9.3 The competing formulations . . . 100
9.4 One system, many questions . . . 101
9.5 The exercises . . . 102
9.6 What to retain . . . 104

III Constructing Probabilistic Models 105

10 From Physical Relationships to Uncertain Quantities 106
10.1 The engineering question . . . 106
10.2 The formulation . . . 106
10.3 The competing formulations . . . 113
10.4 One system, many questions . . . 115
10.5 Worked problems . . . 115
10.6 The exercises . . . 118

10.7 What to retain . . . 120

11 From Measurements to Evidence **121**
11.1 The engineering question . . . 121
11.2 The formulation . . . 121
11.3 The competing formulations . . . 126
11.4 One system, many questions . . . 128
11.5 Worked problems . . . 128
11.6 The exercises . . . 131
11.7 What to retain . . . 133

12 From Engineering Knowledge to Priors **134**
12.1 The engineering question . . . 134
12.2 The formulation . . . 134
12.3 The competing formulations . . . 140
12.4 One system, many questions . . . 142
12.5 Worked problems . . . 142
12.6 The exercises . . . 145
12.7 What to retain . . . 147

13 From Component Models to System Questions **148**
13.1 The engineering question . . . 148
13.2 The formulation . . . 148
13.3 The competing formulations . . . 154
13.4 One system, many questions . . . 156
13.5 Worked problems . . . 156
13.6 The exercises . . . 159
13.7 What to retain . . . 161

IV Formulating Decisions and Acceptance Rules **162**

Before Part IV: a page to write now **163**

14 From Probability to Engineering Decisions **164**
14.1 The engineering question . . . 164
14.2 The formulation . . . 164
14.3 The competing formulations . . . 168
14.4 One system, many questions . . . 170
14.5 Worked problems . . . 170
14.6 The exercises . . . 172
14.7 What to retain . . . 174

15 False Acceptance and False Rejection **175**
15.1 The engineering question . . . 175
15.2 The formulation . . . 175
15.3 The competing formulations . . . 180
15.4 One system, many questions . . . 181
15.5 Worked problems . . . 182
15.6 The exercises . . . 184
15.7 What to retain . . . 186

16 Statistical Power and Test Design **187**
16.1 The engineering question . . . 187
16.2 The formulation . . . 187
16.3 The competing formulations . . . 194
16.4 One system, many questions . . . 195
16.5 Worked problems . . . 195
16.6 The exercises . . . 200
16.7 What to retain . . . 203

17 Decisions Under Uncertainty **204**
17.1 The engineering question . . . 204
17.2 The formulation . . . 204
17.3 The competing formulations . . . 210
17.4 One system, many questions . . . 211
17.5 Worked problems . . . 212
17.6 The exercises . . . 215
17.7 What to retain . . . 217

V Formulating Experiments and Information Questions **218**

VI Integrated Case Studies **219**

A Solutions to the exercises **220**

B The scoring rubric **281**

C The blank formulation sheet **286**

D The fact sheets of the recurring systems **288**
D.1 The two-hall cooling plant . . . 288
D.2 The small electrical network . . . 290
D.3 The battery string . . . 292

E How common formulations are evaluated, by problem class **295**

Notation and vocabulary

Every symbol below has one meaning throughout the book. The quantity asked about stands on the left of a conditioning bar and what is taken as known stands on the right. Vectors are bold. The letters A and B as subscripts name the two halls and their cooling units.

Evidence in hand and evidence not yet collected

One distinction in the notation carries much of the book. The letter E is evidence in hand: readings already taken, records already kept and prior information already stated. From Part II on, on the two-hall plant, E is the 40 kW trip test only; a record is written "given the record", and stated rates "given the stated model". It is fixed, and it stands behind the bar of a posterior, as in

$$P(G \geq G^* \mid E),$$

the probability, after what is known has been used, that the opening's conductance reaches the threshold. The symbol $\tilde{E}$ is evidence not yet collected. The readings of a test that has not been run are not known, so $\tilde{E}$ is a random variable, and a probability can be taken over it, as in

$$P(A(\tilde{E}) = 1 \mid G),$$

the probability that a test still to be run would be accepted if the conductance had the value G. A test at a controllable load Q writes its future readings $\tilde{E}_Q$. The first expression is a statement about the plant given a test. The second is a statement about a test given the plant. Lessons 14 to 21 depend on keeping them apart.

The compact run

Every worked formulation is also written as a compact run: one line per step, each line naming the step or the operation applied and its result, with the final expression boxed. The seven steps are the spine and are numbered. Named operations sit inside a step as unnumbered lines: *Intention* under step 1, *Declare* under step 2, *Random here* under step 3, *Expand* under step 5, *Reduce* under step 6, and the problem *Class* under step 7. A competing formulation is written as a run too, with the first wrong line marked $\times$ and its repair beneath it. The box holds an expression, the object of step 5 or its reduction, and never a number or a phrase. The run for

the opening's adequacy reads:

Run, is the opening adequate	
1 Sentence	Given E, one unit out, steady: probability the opening passes 60 kW within 8 K?
Intention	Estimate
2 Target	Parameter G, through an event on it
Declare	G conductance [W/K]; $G^* = Q_{\text{hall}}/\Delta T_{\text{allow}} = 7500\,\text{W/K}$; E the trip test's readings
3 Known	E; one unit out; 60 kW; steady
Random here	G, one unknown value; Q_{hall}, ΔT_{allow} fixed
4 Event	$G \geq G^*$; no horizon
5 Expression	$P(G \geq G^* \mid E)$
6 Model	Heat balances; the opening's law; reading = quantity + offset + noise
7 Check	The opening passes a hall's load within the limit
Class	Posterior exceedance
	$\boxed{P(G \geq G^* \mid E)}$

Two short names stand for parts of the spine: *Parse* is steps 2 to 4, reading a sentence for its random variable, event, condition and rule, and *Translate* is step 5's first line. Four of the operations deserve a definition. *Declare* names every symbol the expression will use, with its meaning and unit, before the expression is written. *Random here* says what is uncertain in the expression, as against what is fixed, set by an intervention do(·), or held behind the bar; a reader who cannot say what is random cannot say what the probability is taken over. *Expand* substitutes a defined quantity's definition into the expression, so that what the expression is about can be seen: a rule $A(\tilde{E})$ that accepts when $P(G \geq G^* \mid \tilde{E}) \geq 0.95$, substituted into $P(A(\tilde{E}) = 1 \mid G)$, shows a probability inside a probability. A run carries an Expand line wherever a rule or a defined event stands in its expression, and it may omit the line only where there is none to substitute. *Reduce* uses the stated model to reach a standard form: a threshold on a parameter, Φ under a stated Gaussian reading, a product under a stated independence. So factoring under a stated independence is a Reduce, and Bayes' rule, which is not a definition, belongs on the Model line. A Gaussian reading writes $\approx$ and says "Gaussian reading". *Settled*, behind a bar, means that the plant has reached its steady temperatures after the change named beside it. A table whose column headers carry no unit holds words, expressions or dimensionless values.

Symbols

Symbol	Unit	Meaning
Probability		
$P(\mathcal{A})$, $P(\mathcal{A} \mid \mathcal{B})$	–	Probability of an event; conditional on what stands behind the bar. Script capitals $\mathcal{A}$, $\mathcal{B}$ are generic events.
Ω	–	The outcomes: every state the system can be in at the time asked about. An event is a set of outcomes.
$\mathcal{A}^c$, $\mathcal{A} \cup \mathcal{B}$, $\mathcal{A} \cap \mathcal{B}$	–	Complement (not $\mathcal{A}$); union ($\mathcal{A}$ or $\mathcal{B}$, or both); intersection ($\mathcal{A}$ and $\mathcal{B}$).
$\mathbb{E}[\,\cdot \mid E\,]$	varies	Expectation given what is known.
θ	varies	A generic parameter: a quantity of the model with one value, unknown.
$p(\theta \mid E)$	varies	A probability density, written in lowercase to keep it apart from P; here the posterior of θ.
Evidence		

Symbol	Unit	Meaning
E	varies	Evidence in hand: readings, records and stated prior information. Fixed. From Part II on, on the two-hall plant, the 40 kW trip test only; a record is "given the record".
$\tilde{E}$, $\tilde{E}_Q$	varies	Evidence not yet collected, a random variable; from a test at load Q.
The objects of a formulation		
$\mathbf{z}$	varies	The state: every unknown of a model at once.
$g(\mathbf{z})$	varies	A limit-state function; $g \geq 0$ is the safe side of a requirement and $g < 0$ its failure domain. A subscript names the limit, as in g_{cap}.
$\hat{\mathbf{z}}$	varies	A point value of the state given E; in this book its most probable value.
$A(\cdot)$, e, $\varnothing$	–	An acceptance rule, a function of evidence with value 1 (accept) or 0, always written with its argument; a placeholder for a test's evidence in a rule's definition; no evidence beyond the priors. On the battery string and the network the argument is the reading itself.
$\mathbb{A}$	–	A set of actions, what can actually be done.
α, $1-\alpha$	–	The posterior risk a rule allows on a reading it accepts; the rule's level.
G_1, G_0	W/K	A stated adequate conductance, a stated inadequate one, behind a bar. Neither is an estimate.
$\tilde{d}$, $d_{\text{crit}}(Q)$	K	The difference of a test not yet run, or a reading on another system with its unit stated; the critical difference, at which a rule's value changes at load Q.
p_{r}, p_{o}	–	The prior inside a rule, fixed when the rule is written; the outer belief about the truth, in an assurance or a prediction. They can differ.
G^*_{50}, G^*_7	W/K	A claim's threshold at another load or limit, as 50 kW over 8 K, or 60 kW over 7 K.
$\mathcal{L}(a,\zeta)$, $\mathcal{L}(a)$	stated	The loss of an action a in a state ζ; an action's cost, in either state.
$\mathcal{L}_{\text{FA}}$, $\mathcal{L}_{\text{FR}}$, $\mathcal{L}_{\text{t}}$	stated	What a false acceptance costs beyond the alternative; what a false rejection costs beyond accepting; a retest's cost.
ξ; κ	–; kW	A share of a load shed; a window's half-width on a meter's offset.
$P_{\text{elec}}(\mathbf{u})$; $\vartheta_{\max}$	kW; °C	A plant's electricity at a setting; a stated limit on a hall's temperature.
H_i	–	A hypothesis, a discrete unknown.
$\text{do}(\mathbf{u})$, $\mathbf{u}$	varies	An intervention that sets the controls $\mathbf{u}$, as against observing their values.
p_{req}	–	The acceptable level of a probability, when a requirement supplies one.
ψ_j	–	A factor of a joint distribution, with j its index.
$s(\cdot)$	–	A structure function: a system's event as a function of its components' states, and s_{IT} for the network's IT load.
$\mathcal{N}(x;\, m,\, \sigma^2)$	–	A normal density at x, with mean m and variance σ^2. Φ is the standard normal distribution function.

Symbol	Unit	Meaning
m_x, σ_x	varies	The center and the spread of a Gaussian reading in the coordinate x.
σ_0	–	A prior's spread in the log; $\log 3$ on the stated priors, "a factor of 3".
ω_i	–	The weight of the i-th component of a mixture, as a drawing's weight.
ε, δ	varies	A reading's error in a measurement relation; a discrepancy, what a relation is known to leave out.
t, τ	h or yr	A time; a time horizon.
T_ℓ	h or yr	The first-passage time: the first time after a stated start at which a rising quantity reaches the level ℓ. A subscript in words names the event instead, as in T_{rt}.
$M(t)$	varies	The running maximum of a stated quantity: the largest value it has taken since the start.
$L(t)$	h	The cumulative hours without service since the start, a path that only rises; $L_{\mathrm{IT}}(t)$ for the network's IT load, and $L_{\mathrm{ex}}(t)$ for the two-hall plant's with the hours of a feeder outage counted as served.
The two-hall plant		
$S(t)$	–	Service indicator: 1 when the cooling service is delivered at time t, 0 otherwise; defined with each hall at its design load.
$\bar{S}(m)$	–	The interval availability of period m, such as a month: the share of its hours in which the service is delivered. A random quantity, not a probability.
$\bar{S}_{\mathrm{ex}}(m)$	–	The same share with the hours of a feeder outage counted as served, for a requirement that excepts them.
G	W/K	Conductance of the wall opening between the halls.
$\hat{G}$	W/K	The point value of G given E.
G^*	W/K	The conductance that carries hall B's full load, with unit B out, within the allowable difference, $Q_{\mathrm{hall}}/\Delta T_{\mathrm{allow}}$; 7500 W/K.
Q_{hall}, Q	kW	IT heat load of one hall in service, 60 kW; trip-test load per hall.
Q_A, Q_B	kW	The load of hall A, or hall B, at the time asked about: a condition when known, an uncertain quantity when not.
ϑ_A, ϑ_B	°C	The temperature of hall A, or hall B. The letter ϑ with a subscript is kept for temperatures, and θ for a parameter.
$\Delta T_{\mathrm{allow}}$	K	Allowable difference of hall B above hall A, 8 K.
ΔT_{full}	K	Difference of hall B above hall A with unit B out, each hall at its full load, the plant settled.
F_A, F_B	–	The events that unit A, or unit B, is unavailable for its own reasons. An outage of the feeder, which makes both units unavailable, is a separate event.
F_{fd}	–	The event that the feeder is out.
D	h	The length of a repair, a random quantity; it sets a random horizon such as $[t_0, t_0 + D]$. D_{tot} is a record's total downtime.
μ_A, μ_{fd}	h	The mean repair of unit A, or of the feeder, a parameter.
q, f	–	The long-run share of time a unit, or the feeder, is down, at the stated rates.

Symbol	Unit	Meaning
c	–	A disturbance window's state, disturbed or calm.
$N_{\mathrm{fd}}(\tau)$, $N_A(\tau)$	–	The number of the feeder's, or unit A's, failures in $[0, \tau]$.
λ_A	1/yr	Unit A's failure rate per year in service, a parameter.
H_{AND}, H_{OR}	–	The parallel structure, in which both units must fail to lose the service; the series structure, in which either unit's loss loses it.
y_A, y_B; b_A, b_B	K	A hall sensor's mean reading over a test; its constant offset. On the network, y with a meter's name is its reading and b_{pdu} the PDU meter's offset, in kW.
k, t_{s}	–, yr	A count of failures in a record, with $\tilde{k}$ one not yet logged and k_{c} a rule's critical count; its exposure, the time in service over which it was counted.
d, σ_d	K	The measured difference of hall B's reading above hall A's in a trip test, and its spread. On the battery string d is a timed reading, in hours.
R, R^*	K/kW	The opening's thermal resistance $1/G$, and the threshold $1/G^*$. $\theta = \log G$ is the logarithm of the conductance.
The electrical network		
$S_{\mathrm{IT}}(t)$	–	1 when the IT load is supplied at time t, 0 otherwise. Kept apart from $S(t)$, the two-hall plant's cooling service.
F_{sup}	–	The event that the UPS's supply is lost: the feeder interrupted, or the incoming breaker open.
F_{ups}	–	The event that the UPS is not carrying the load while its supply is lost: it failed to pick up the load, the loss has outlasted its battery, or the UPS is out for repair with the load on its bypass.
F_{it}, F_{clr}	–	A fault on the IT branch; the breaker above both branches open after a cooling-branch fault not cleared.
D_{sup}	h	The length of a loss of supply, a random quantity. D stays the two-hall plant's repair length.
t_{aut}	min	The autonomy of the UPS battery at full load, 10 min (illustrative).
$N_{\mathrm{ut}}(\tau)$, $N_{\mathrm{IT}}(\tau)$	–	The number of the network's utility-feeder interruptions in $[0, \tau]$; the number of interruptions of the IT load. N_{fd} stays the two-hall plant's feeder.
η	–	The UPS's conversion efficiency, the share of its input it delivers.
P_{load}	kW	The power the UPS delivers to its loads, 1150 kW at full load.
The battery string		
φ	–	The string's state of health: its usable energy as a share of a nameplate string's.
T_{rt}	h	The ride-through time: how long the string carries the 20 kW critical load before reaching its floor.
D_{out}	h	The length of a grid outage at the string, a random quantity.
x_0	–	The state of charge at an outage's start.
γ_i	–	A string's usable energy under hypothesis i as a share of a sound string's: $\gamma_{\mathrm{sound}} = 1$, $\gamma_{\mathrm{weak}} = 0.70$.

Symbol	Unit	Meaning
π_{w}	–	The share of strings with a weak cell, over a list of strings.
μ_{o}, ν_{o}	h, 1/yr	The mean length of an outage; their rate of arrival.
H_{weak}, H_{sound}	–	The hypotheses that one cell of the string is weak, and that none is.
$\vartheta_{\mathrm{pk}}(t)$	K	The temperature of the battery pack.

Bridge to the profession books

The books *Reliability Engineering by Derivation* and *Capacity Planning and Design by Derivation* solve problems of the kind this book writes down. Their notation differs where those books need a symbol for a computation and this book does not. Table 1 sets the two side by side, so that an expression written here can be carried there.

Table 1: The notation of this book against the profession books.

This book	Profession books	Note
E	$\mathbf{y}$	Evidence in hand. A single reading is y in both.
$\tilde{E}$, $\tilde{E}_a$	$\tilde{y}_a$	The outcome of a candidate measurement before it is taken.
a named outcome, such as $S(t)$	y_{new}	A future observation asked about in a predictive probability, which will not be used as evidence.
$\mathbf{z}$	$\mathbf{z}$	The state, with the same meaning.
P	Pr	The same object, printed differently.
g	g, sometimes	Here always a limit state. In the profession books' equipment hall, g is a log-conductance, which this book writes $\log G$.
$A(\cdot)$	—	An acceptance rule has no symbol there. Their A is availability, which this book writes as a probability, $P(S(t) = 1 \mid E)$.
H_i	M_j	A drawing of the plant taken as a discrete unknown.
$\mathrm{do}(\mathbf{u})$	x^*	Their star marks a quantity read at a scenario node, which is an intervention in the sense used here. A star on G^* marks a threshold.

Vocabulary

The book uses the standard term of the field wherever one exists. The terms are defined here briefly and explained where a lesson first needs them.

Availability The probability that a service is being delivered at a given or randomly chosen time. Not a survival probability over a mission time.

Reliability over a horizon The probability that a service is delivered at every time in a stated interval, also called the survivor function or mission reliability.

Aleatory and epistemic uncertainty Variation that stays when everything knowable about the plant is known; lack of knowledge about a quantity that has one value.

Posterior, likelihood, prior What is believed after the evidence; the probability of the evidence given the unknowns; what is believed before it.

Predictive probability The probability of a future outcome, averaged over what remains unknown.

First passage The first time a quantity reaches a stated level.

Limit state The boundary, written as a function, between satisfying a requirement and failing it.

Capacity margin Capacity relative to demand. A statement about capacity, never about availability.

Common-cause failure, minimal cut set A failure that defeats redundant units through something they share; a smallest set of failures sufficient to lose the service.

Observability, identifiability Whether the readings determine a state; whether they determine a parameter.

Nuisance variable An unknown not asked about that must still be carried in the model.

Operating characteristic, power, assurance The probability of acceptance as a function of the true value; its value at a stated truth; its average over what is believed about the truth. The two-hall report uses the name "probability of acceptance" for a posterior, which is not this; Exercise 16.1 takes that up.

Preposterior analysis Reasoning about the posterior that a measurement not yet taken would give.

Expected value of information The expected improvement in a decision from a measurement not yet taken.

Chance constraint A requirement that a stated event hold with at least a stated probability.

Part I

The Language of Engineering Uncertainty

Lesson 1

What Exactly Is the Engineering Question?

1.1 The engineering question

In a design review, an owner asks whether the data center is reliable enough. Everyone in the room understands the sentence, and nobody in the room could yet write down what would answer it. The sentence does not say what counts as a failure, over what period the plant is to be judged, under what operating conditions, or how much probability of failure the owner will accept. Two competent analysts can reply with different numbers, each computed correctly, and neither has to be wrong for the two to disagree, because they have answered different questions.

The plant under review is the one this book returns to most often, and Figure 1.1 shows it. The two large boxes are data halls. Each holds an IT load of 60 kW, which turns into heat, and one cooling unit. The double arrow between the halls is an opening in the wall, and the dashed arrows run from the one electrical feeder that supplies both units. Each unit has the capacity to cool both halls, but it can do so only if the heat of the other hall can reach it through the opening. The drawing therefore claims one spare unit, a claim known in the trade as N+1, and that claim holds only if the opening passes enough heat. The stated facts of the plant are collected in Appendix D.

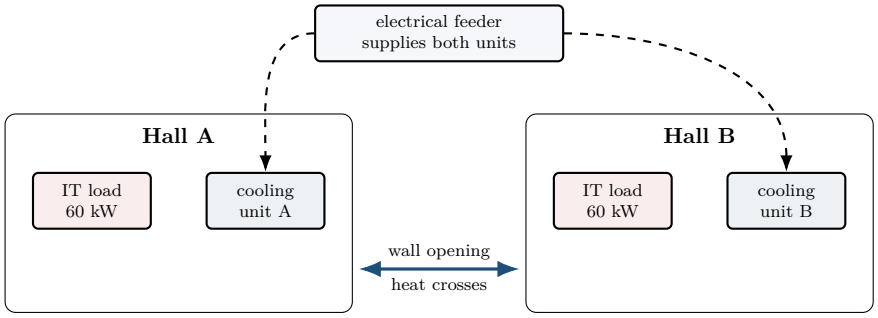


Figure 1.1: The two-hall cooling plant. Each hall holds a 60 kW IT load and one cooling unit, heat crosses the wall opening between the halls, and one electrical feeder supplies both units.

Later in the same review, someone asks how confident the team is that the opening is adequate. This sentence sounds like a narrower version of the first, and it is a different sort of question. It concerns one property of one built object, which has a value today whether or not anyone knows it. The owner's question concerns what the plant will do over time, and it also asks for a verdict.

This lesson takes such sentences apart before any probability is written. It says what kind of

statement each part of a sentence is, what the speaker is trying to do, and what must be supplied before an expression can be written. It then gives a procedure for writing the expression and for reading it back against the sentence it came from. One result deserves stating in advance, because it is unusual in a technical subject: for some sentences the correct outcome of the work is not an expression but a statement of what is missing.

1.2 The formulation

The procedure is built here in the order it is used. A sentence is sorted by kind and then by what the speaker intends, and its uncertain quantities are separated by the kind of uncertainty they carry. The seven-step method follows, and it is then run on the two sentences above. Expressions are read in words throughout. In $P(\mathcal{A} \mid \mathcal{B})$ the vertical bar is read "given": the event $\mathcal{A}$ is what is asked about, and $\mathcal{B}$ is what is taken as known. The letter E stands for the evidence in hand, meaning the readings already taken, the records already kept and the prior information already stated. Lesson 3 treats the bar formally.

1.2.1 Four kinds of statement

A discussion of a plant produces statements of four kinds. A *physical fact* is stated and taken as known. An *uncertain quantity* has a value that is not known. A *decision* is a choice that is still open. A *prediction* says what will happen. The kind of a statement decides where it goes in a formulation, as Table 1.1 shows for eleven sentences from one meeting about the plant. Its first column gives the sentence, its second the kind, and its third the place the statement takes in a formulation.

Table 1.1: Sentences from one meeting about the two-hall plant, sorted by kind.

Sentence	Kind	Place in a formulation
Each hall carries 60 kW of IT load.	Physical fact	A condition, behind the bar
One feeder supplies both cooling units.	Physical fact	The structure of the model
Heat crosses the opening in proportion to the temperature difference.	Physical fact	A relation in the model
Hall B may run at most 8 K above hall A.	Physical fact, stated	The threshold
The opening's conductance lies somewhere near 7500 W/K.	Uncertain quantity	Asked about, or carried in the model
Each hall sensor carries a constant offset, with a prior spread of about half a kelvin.	Uncertain quantity	Carried in the model
The handbook gives each unit 2.5 failures per year.	Physical fact, about the handbook	Prior information
Unit A will fail during the next twelve months.	Prediction	Asked about
Hall B will hold its limit the next time unit B trips.	Prediction	Asked about
The commissioning test will run at 20 kW per hall.	Decision	A choice, then a condition
The handover should be signed rather than a third unit bought.	Decision	The action, not a probability

Most of the rows are straightforward, and four repay a second look. The limit of 8 K is a given of the specification rather than of physics; it is sorted with the facts because it is stated and is not in doubt, and Lesson 4 takes up how a requirement becomes an event. The sentence about heat crossing the opening states the form of a relation as known while leaving its coefficient unknown, so one physical statement contributes both a fact and an uncertain quantity. The handbook sentence is a fact about the handbook. The failure rate of this particular plant

remains an uncertain quantity, and reading the one as the other is a common error. The test load is a decision until the test is run, and after that it is a condition. The other common error in sorting is to treat a decision as though it had a probability of being correct, as in "the probability that a third unit should be bought". A decision has consequences, and probabilities attach to what follows from it.

The four kinds sort statements, and a question is sorted by what it asks about. A question about the conductance is a question about an uncertain quantity, a question about whether unit A will fail is a question about a prediction, and a question about what to buy is a question about a decision. A question about a physical fact is either already answered or a sign that the supposed fact is in doubt.

1.2.2 The engineering intention

The second thing to establish is what the speaker is trying to do. Seven intentions cover the questions in this book: to estimate, predict, compare, decide, optimize, diagnose, or design an experiment. Table 1.2 pairs each one with a sentence about the plant in its second column, and its third column gives the form an answer of that kind usually takes.

Table 1.2: The seven engineering intentions, each with one sentence about the two-hall plant and the usual form of its answer.

Intention	Sentence about the plant	Usual form of the answer
Estimate	How large is the conductance of the opening?	A probability about a quantity that has one value
Predict	Will the service run a year without interruption?	A probability about an outcome still to occur
Compare	Do the readings favor the parallel or the series structure?	A probability for each of several explanations
Decide	Should the N+1 claim be accepted, or a third unit bought?	An action chosen by its consequences
Optimize	Which set point uses the least electricity and keeps the overheating risk under 1 %?	Settings chosen by an objective, under a probability limit
Diagnose	Hall B read hot last night: is the cause the opening, the sensor or unit B?	A probability for each candidate cause of what was seen
Design an experiment	At what load should the trip test run?	A test chosen by what its readings, not yet taken, would do

The intention is named before any mathematics, because the expression follows from it and not from the words of the sentence. "Is the opening adequate?" is an estimate when an engineer asks it out of interest and a decision when an owner asks it with a handover form in hand, and only the second needs a level and a consequence. The intention is therefore the first entry on the formulation sheet.

The third column of Table 1.2 shows that the seven intentions lead to mathematical objects of different kinds, and not merely to different probabilities. Two of them, estimating and predicting, end in a probability of an event. Comparing and diagnosing end in a probability spread over several explanations at once, the structures of the plant in the one case and the candidate causes of a reading in the other. Deciding and optimizing end in a choice, and a probability enters only as an ingredient of the choice. Designing an experiment ends in a choice as well, but it is judged over readings that do not exist yet, which makes it the most unusual object of the seven. A sentence that mixes intentions, such as one that asks how good the plant is and what to do about it, therefore holds more than one question, and each needs its own expression.

1.2.3 Aleatory and epistemic uncertainty

The uncertain statements in Table 1.1 are uncertain in two different ways. The conductance of the built opening has one value, and nobody knows it. This is *epistemic* uncertainty, a lack of knowledge, and measuring the plant reduces it. The month in which unit A next fails varies even when everything knowable about the plant is known. This is *aleatory* uncertainty, a variation that no measurement of the plant removes. A single test separates the two: if a longer or better measurement of this plant, with the plant unchanged, would shrink the uncertainty, then the uncertainty is epistemic. Table 1.3 applies the test to four quantities and gives the reason for each verdict.

Table 1.3: Four quantities of the two-hall plant and the kind of uncertainty each one carries.

Quantity	Kind	Reason
The conductance of the opening	Epistemic	It has one value, which is not known.
The constant offset of a sensor	Epistemic	It has one value; averaging does not reveal it, but a calibration against a reference would.
The noise of a sensor from one minute to the next	Aleatory	The next reading varies however well the sensor is known.
The month in which unit A next fails	Aleatory	It varies even at a known failure rate. The rate itself is epistemic.

The last row shows that both kinds can sit in one sentence, and a single probability can hold both: a probability about the service next year averages over the unknown conductance and over the failures still to occur. The line between the two is a choice made in modeling and not a property of nature (Der Kiureghian and Ditlevsen, *Structural Safety*, 2009). It is worth drawing because it says what a further measurement can buy. A better test narrows what is known about the conductance, while no test of this plant says in which month unit A will fail.

1.2.4 Four levels of a problem

Between the sentence and the number lie four levels, and it helps to know which of them is being worked on. Figure 1.2 stacks them as four boxes, from the engineering intention at the top to the solution strategy at the bottom, with labels at the right marking how this book treats each one. Table 1.4 carries one question down the same ladder. In its second row, G is the conductance of the opening and G^* the conductance that carries one hall's load at the allowable temperature difference. The first row keeps the owner's words, "the cooling system is adequate", and the second reads "adequate" as the opening passing one hall's load. That reading is itself a formulation choice: it makes the question one about capacity, and a statement about capacity is not a statement about availability, a distinction the competing formulations below take up.

Table 1.4: One question carried through the four levels.

Level	Example
1. Engineering intention	How confident is the team that the cooling system is adequate?
2. Mathematical question	$P(G \geq G^* \mid E)$
3. Mathematical problem class	Posterior exceedance with uncertain physical parameters
4. Solution strategy	A Gaussian approximation, numerical integration, sampling, or another method

Most technical textbooks concentrate on the third and fourth levels, and they do so very

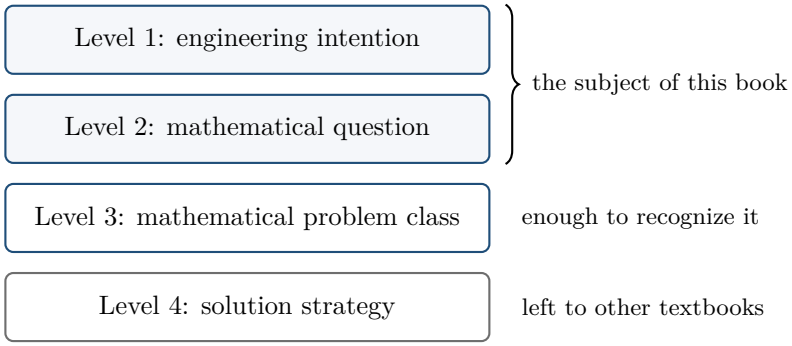


Figure 1.2: The four levels of a problem and how this book treats each one.

well. This book concentrates on the first two and teaches enough of the third for a reader to recognize what has been built and to hand it to the right technique. The fourth level is named here and not used again in the lessons.

1.2.5 The seven-step method

One procedure carries a sentence from the first level to the second. It has seven steps.

1. State the engineering question in one unambiguous sentence.
2. Identify the target quantity: a state, parameter, event, future outcome, or decision.
3. Specify what is known: observations, assumptions, operating conditions, and prior information.
4. Define the event or objective mathematically, including thresholds and time horizons.
5. Write the probability or expectation that directly answers the question.
6. Identify the model required to evaluate that expression: variables, physical relations, dependencies, and evidence.
7. Check the formulation: does the expression really answer the original engineering question?

None of the steps calls for solving anything. Figure 1.3 groups them by what each produces. The first three are carried out in words: the question as one sentence, the thing it asks about, and what is taken as known. The fourth and fifth turn those words into mathematics, an event and then a probability or an expectation of it. The sixth names the model that the expression needs, without building it. The seventh returns to words, and the arrow drawn back to the first box marks that a failed check sends the work back to the sentence, which is where most defects begin.

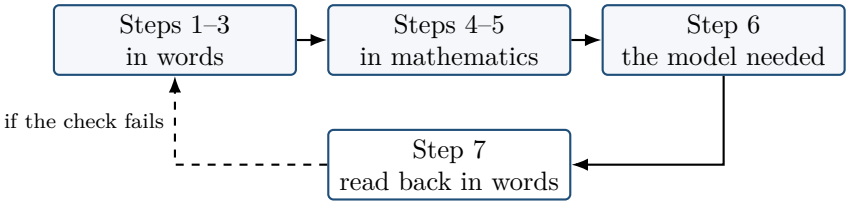


Figure 1.3: The seven steps grouped by what they produce. The dashed arrow returns a formulation that fails its check to the sentence.

The product of the method is a *formulation sheet*, one page with the steps as fields, the intention at the top and the problem class at the bottom; a blank one is in Appendix C. The

last step carries particular weight in this book. Since an expression is not computed here, its correctness cannot be judged from a number, so it is read back in plain words and held against the original sentence. Most errors in formulation survive every step but that one. A sheet may also stop early. When step 1 cannot produce one unambiguous sentence, the sheet records what is missing, and that record is the answer for the time being.

1.2.6 A question that is not ready

The owner's sentence fails at the first step. It cannot be made unambiguous without four things it does not say: what counts as failure, the horizon, the conditions and the level. An analyst can supply two of these provisionally, taking failure to be loss of the cooling service and the conditions to be steady operation at 60 kW per hall. The level cannot be supplied by the analyst, since it belongs to whoever relies on the plant, so it is carried as an open symbol, p_{req}. The horizon has two natural values, each of which yields a different sentence:

> At an hour chosen at random from steady operation, is the service being delivered, and is the probability of that at least the required level?
>
> Over the next twelve months, starting with everything running, is the service delivered at every hour, and is the probability of that at least the required level?

The intention behind both is to decide, on account of the word "enough", and the decision rests on a prediction. The target is a future outcome. To write it, let $S(t) = 1$ mean that the service is delivered at time t, which on this plant requires the feeder to be up and either both units to be up or one unit to be up with $G \geq G^{*}$. The two sentences then become

$$P\big(S(t) = 1 \mid E\big), \qquad t \text{ a random hour}, \tag{1.1}$$

which is the *availability* of the service, and

$$P\big(S(t) = 1 \text{ for every } t \in [0, \tau] \mid \text{all up at } t = 0,\ E\big), \qquad \tau = 1 \text{ yr}, \tag{1.2}$$

which is its *reliability over a horizon*. In both, the evidence E is whatever is in hand: the drawing, the stated rates, any trip test and any outage record. Each probability is then read against p_{req}. The start of the year stands behind the bar in equation (1.2) because it is a condition like any other. A question asked while a unit is already out conditions on that present state instead, and its horizon may itself be random, such as the time until a repair is finished. Table 1.5 sets the two readings side by side, step by step, with each column read downward from the sentence to the problem class.

Reading the sheet's two columns against the owner's sentence decides nothing by itself, and that is the finding. Neither expression is the owner's question until the owner chooses between them. A tenant paid by the hours in service needs the first. A tenant whose workload restarts at great cost needs the second. The analyst's task in the meeting is to bring back both sentences, ask which is meant, and ask what level applies. The sheet's sixth row also carries a lesson of its own. The two readings need different models, so choosing the question has already changed what must be built.

The two readings are not the only ones. Step 5 asks for a probability *or an expectation*, and an owner who thinks of reliability as lost hours is asking for an expectation,

$$\mathbb{E}\Big[\int_0^{\tau} \big(1 - S(t)\big)\,\mathrm{d}t \;\Big|\; E\Big], \qquad \tau = 1 \text{ yr}, \tag{1.3}$$

the expected number of hours in the year during which the service is not delivered. The integral counts the hours out of service along one possible year, and the expectation averages that count

Table 1.5: The seven steps on "Is this data center reliable enough?", for its two natural readings.

Step	Reading at a random hour	Reading over a year
1. Sentence	A random hour of steady operation at full load	The next twelve months at full load, starting with everything up
2. Target	A future outcome: $S(t)$ at one random time	A future outcome: the whole path of S over the year
3. Known	E; 60 kW per hall; steady operation. The level p_{req} is not supplied.	The same, with everything up at $t = 0$. The level p_{req} is not supplied.
4. Event	$S(t) = 1$, with t random	$S(t) = 1$ for every t in $[0, \tau]$, given all up at $t = 0$
5. Expression	Equation (1.1), against p_{req}	Equation (1.2), against p_{req}
6. Model	The structure, which depends on G; each component's fraction of time down; evidence tying the record to the rates and the test to G	All of that, with the process of failure and repair over time, since whether both units are ever down together depends on repairs in progress
7. Read back	The fraction of hours in service. Silent on how many interruptions there are.	A clean year. Silent on how long an interruption lasts.
Problem class	Predictive probability at a random time	Predictive probability over a horizon, a first passage

The two columns as runs:

Run, the owner's sentence read at a random hour

1	**Sentence**	Random hour, steady operation, full load: service delivered, at level p_{req}?
	Intention	Decide, on a prediction
2	**Target**	Future outcome: $S(t)$ at one random time
	Declare	$S(t) = 1$ if the service is delivered at t; E evidence in hand; p_{req} the level
3	**Known**	E; 60 kW per hall; steady operation; p_{req} not supplied
	Random here	$S(t)$, the hour t; G and the rates averaged over
4	**Event**	$S(t) = 1$, with t random
5	**Expression**	Equation (1.1), against p_{req}
6	**Model**	Structure, through G; each component's fraction of time down; E on rates and G
7	**Check**	Fraction of hours in service; silent on the number of interruptions
	Class	Predictive probability at a random time

$$\boxed{P\big(S(t) = 1 \mid E\big) \geq p_{\text{req}}, \quad t \text{ random}}$$

Run, the owner's sentence read over a year

1	**Sentence**	Next twelve months, all up at the start: service at every hour, at level p_{req}?
	Intention	Decide, on a prediction
2	**Target**	Future outcome: the whole path of S over the year
	Declare	$S(t)$, E, p_{req} as above; $\tau = 1$ yr
3	**Known**	E; 60 kW per hall; all up at $t = 0$; p_{req} not supplied
	Random here	The path of S over $[0, \tau]$; G and the rates averaged over
4	**Event**	$S(t) = 1$ for every t in $[0, \tau]$
5	**Expression**	Equation (1.2), against p_{req}
6	**Model**	As above, with the failure and repair process over time
7	**Check**	A clean year; silent on how long an interruption lasts
	Class	Predictive probability over a horizon, a first passage

$$\boxed{P\big(S(t) = 1 \text{ for every } t \in [0, \tau] \mid \text{all up at } t = 0,\ E\big) \geq p_{\text{req}}}$$

over the years the plant might have. Expression (1.3) is not a probability and has no level between zero and one; it is compared with a budget of hours. It carries the same information as the availability of equation (1.1) for a plant in steady operation, expressed in different units, and it says just as little about how many interruptions make up those hours.

The provisional choices made at step 1 deserve the same scrutiny as the horizon. The conditions were taken as steady operation at 60 kW per hall, which is the design load. A plant that runs at half load most of the year needs much less from its opening for most of the year, so an expression conditioned on full load answers a harsher question than the one about the actual year, and the sheet must say which was meant. The level raises a different difficulty. It belongs to whoever relies on the plant: a contract, a tenant's agreement on service levels, or the owner's own appetite for risk. It must also be stated in the same terms as the expression it is compared with. A contract that promises a share of hours supplies a level for equation (1.1) and says nothing about equation (1.2), and a level borrowed from one and applied to the other is a level for a question nobody asked.

What the analyst says in the meeting follows from all this, and it is short. Before any number, two questions go back to the owner. The first is whether the concern is the share of hours the halls are cooled, the chance of a year with no interruption at all, or the hours lost in a year. The second is what value of that quantity would count as enough, and on whose authority. The analyst can add that the candidate quantities differ greatly on this plant, which the next section shows.

Key idea

A sentence is not yet a probability question until it names its event and conditions, its horizon where time enters, and its level where a verdict is asked. Stating which of these is missing, and writing what each reading would give, is a correct result.

1.2.7 A question that is ready

The sentence about the opening does not have this difficulty. Its intention is to estimate, and its target is a parameter of the built plant, the conductance G. Adequacy has a definite meaning here: with one unit out, one hall's load must cross the opening at no more than the allowable difference. That meaning sets the threshold,

$$G^{*} = \frac{Q_{\mathrm{hall}}}{\Delta T_{\mathrm{allow}}} = \frac{60\,\mathrm{kW}}{8\,\mathrm{K}} = 7500\,\mathrm{W/K}, \tag{1.4}$$

and the question becomes

$$P(G \geq G^{*} \mid E). \tag{1.5}$$

The expression has no time horizon, and the omission is deliberate. The conductance is treated as a fixed property of the opening, and a horizon would enter only if that property could drift. Table 1.6 is the filled formulation sheet, with the fields down the left and their entries on the right.

The reading back at step 7 also marks the limits of the expression. It does not say how often the service is lost, whether the claim should be accepted, or how often a test would accept it. Those are other questions about the same plant. The contrast with the previous subsection is plain in the sheet's model row: equation (1.5) needs no failure process at all, while equation (1.2) cannot do without one.

1.3 The competing formulations

Two of the expressions above are both called reliability in ordinary speech, and the size of the difference between them is seldom appreciated. Figure 1.4 is a schematic of one year of

Table 1.6: The formulation sheet filled in for “How confident is the team that the opening is adequate?”.

Field	Entry
Intention	Estimate
1. Sentence	Given the evidence in hand, what is the probability that the opening’s conductance is at least the value that lets one hall’s 60 kW cross at a difference of no more than 8 K, with the other hall’s unit carrying both and the plant steady?
2. Target	A parameter, G, through an event on it
3. Known	E: the trip test’s readings (two hall temperatures, the surviving unit’s duty, two PDU powers), the drawing’s prior on G, the stated spread of the sensor offsets. Conditions: steady, one unit out, 60 kW. Assumptions: heat crosses in proportion to the difference, and each offset is constant over the test.
4. Event	$G \geq G^*$, with G^* from equation (1.4); no horizon
5. Expression	Equation (1.5)
6. Model	Variables: G, the hall temperatures, the heat on the opening, the two loads, the unit’s duty, the two offsets. Relations: a heat balance at each hall and the opening’s law. Readings: each one is its quantity plus an offset and noise.
7. Read back	After these readings, the probability that the opening passes a hall’s load within the limit. That is what “adequate” asked.
Problem class	Posterior exceedance

One line per step:

Run, how confident is the team that the opening is adequate

1	**Sentence**	Given E, one unit out, steady: probability G passes one hall’s 60 kW within 8 K?
	Intention	Estimate
2	**Target**	Parameter G, through an event on it
	Declare	G [W/K]; $Q_{\text{hall}} = 60\,\text{kW}$; $\Delta T_{\text{allow}} = 8\,\text{K}$; $G^* = 7500\,\text{W/K}$; E evidence in hand
3	**Known**	E; steady, one unit out, 60 kW; heat $\propto$ difference, offsets constant (assumptions)
	Random here	G, one unknown value; the offsets; Q_{hall}, ΔT_{allow} fixed
4	**Event**	$G \geq G^*$, equation (1.4); no horizon
5	**Expression**	Equation (1.5)
6	**Model**	Heat balance per hall; the opening’s law; reading = quantity + offset + noise
7	**Check**	After these readings: the opening passes a hall’s load within the limit
	Class	Posterior exceedance

$$\boxed{P(G \geq G^* \mid E)}$$

service, with time running from left to right, the filled bar for the hours in service and the gaps for the interruptions. The marker above the bar picks out one random hour, which is what equation (1.1) asks about, and the bracket below it spans the whole year, which is what equation (1.2) asks about. A random hour is seldom in a gap when the gaps are short, while a whole year contains every gap there is.

Table 1.7 puts numbers on the two formulations for the plant as it stands, one row for each. In both rows, the failure and repair rates are fixed at their stated values and the opening is taken as adequate, so that nothing is uncertain but the failures and repairs themselves, and the whole of the gap between the rows comes from the formulation. The values come from a small calculation with a stated model: each component alternates between up and down with exponentially distributed times to failure and to repair, the components fail independently of one another, and each unit has its own repair. The rates are 2.5 failures per year in service and a 180 h repair for each unit, and 3.69 failures per year in service and a 24 h repair for the feeder. Read instead as failures per calendar year, the same rates give 0.987 and 0.019, which changes nothing below. The values are quoted to show the size of the difference.

The two values describe one plant. It is in service for 98.8 % of its hours and has about a

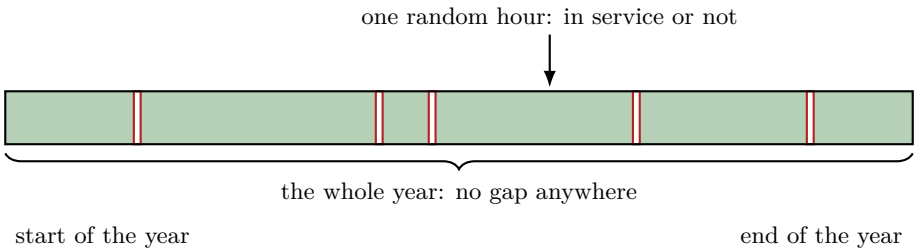


Figure 1.4: Two formulations on one year of service: the state at one random hour and the path over the whole year. The figure is a schematic; the number and the positions of the gaps are illustrative, and their widths are exaggerated so that they can be seen.

Table 1.7: Two formulations of reliability on one plant, with the rates at their stated values and the opening taken as adequate.

Formulation	Expression	Value [–]
Availability at a random hour	$P\big(S(t) = 1 \mid \text{rates}, G \geq G^*\big)$, t random	0.988
No interruption in a year	$P\big(S(t) = 1 \text{ for every } t \in [0, 1 \text{ yr}] \mid \text{all up at } t = 0, \text{ rates}, G \geq G^*\big)$	0.020

2 % chance of an uninterrupted year. It is out of service for about 108 h per year, and a feeder stated to fail 3.69 times per year all but guarantees that some of those hours break the year. A report of "98.8 % reliable" answers the first expression, and an owner who hears it may well believe the second. This is the failure class of this lesson, the *horizon missing*: a sentence or an expression that states no time, so that a statement about one instant can be taken for a statement about a period. Appendix B scores it with the other recurring classes.

The slip, run and repaired:

Run, "98.8 % reliable" offered for a clean year

1	**Sentence**	Next twelve months, all up at the start: service at every hour?
4	**Event** ×	$S(t) = 1$, with t random
	Repair	$S(t) = 1$ for every $t \in [0, 1 \text{ yr}]$
5	**Expression** ×	$P\big(S(t) = 1 \mid \text{rates}, G \geq G^*\big)$, t random: 0.988
	Repair	$P\big(S(t) = 1 \text{ for every } t \in [0, 1 \text{ yr}] \mid \text{all up at } t = 0, \text{ rates}, G \geq G^*\big)$: 0.020
7	**Check**	The wrong form is a share of hours; the sentence asks for a clean year
		$P\big(S(t) = 1 \text{ for every } t \in [0, 1 \text{ yr}] \mid \text{all up at } t = 0, \text{ rates}, G \geq G^*\big)$

The model that step 6 asked for explains why the two values lie so far apart. The availability weighs each interruption by its length, while the year's reliability counts each interruption once, however briefly it lasts. On this plant the feeder dominates both measures: it accounts for about 88 of the 108 hours lost in a year, and with 3.69 failures a year it starts most of the interruptions. Two plants with the same availability can still have very different chances of a clean year, one with rare long outages and the other with frequent short ones. The expected hours lost, equation (1.3), is about 108 h per year on this plant, and it shares the blindness of the availability: it cannot tell those two plants apart either.

The table also isolates one of the two kinds of uncertainty. Fixing the rates and taking the opening as adequate removed everything epistemic from the calculation, so that the 0.020 is entirely aleatory within the stated model. It is the chance of a clean year for a plant whose every property is known and whose components behave as the model says. On the actual plant, with G and the rates unknown, equations (1.1) and (1.2) average over those unknowns as well, and their values move accordingly. One comparison survives every such change, provided both

sides share a start. Under the start of equation (1.2), everything up at $t = 0$, a year with no interruption includes a served hour wherever in $[0, \tau]$ a random hour falls, so the year's probability never exceeds the probability that a random hour of that same year is served. On the stated rates that second probability is 0.988, almost exactly the steady availability, against 0.020. Without a shared start the comparison can fail: one hour after an all-up start, the service is delivered with probability 0.99959, more than the steady value. A repaired report states both quantities and names them: the plant is in service for about 98.8 % of its hours, and the chance that it completes a year with no interruption is about 2 %.

A second competing formulation is close at hand in any review. Equation (1.5) is often available and, after a good test, close to one, so it tends to be offered in reply to the owner's sentence. It answers whether the opening is adequate. The values in Table 1.7 were computed with the opening taken as adequate, which is to say with that probability equal to one, and they remain 0.988 and 0.020. A statement about capacity margin is not a statement about availability, and the two should never be equated. A third reply is no probability at all. "The plant has N+1 cooling" is a fact about the drawing, and it leaves the owner's question where it was.

1.4 One system, many questions

The two-hall plant supports many more questions than the two worked above. Table 1.8 lists six that an engineer might ask of it, with an ordinary-language sentence in the first column, its mathematical form in the second, and the intention in the third. Three symbols in it are taught in later lessons and need only a gloss here. The tilde on $\tilde{E}$ marks readings not yet taken, A is a rule that accepts or rejects a test, and H_{AND} names the parallel structure, in which both units must fail for the service to be lost. The events F_A and F_B are that unit A and unit B are unavailable for their own reasons; an outage of the feeder, which makes both units unavailable at once, is a separate event. The quantity ΔT_{full} is the difference of hall B above hall A with one unit out at full load.

Table 1.8: Six sentences about the two-hall plant, each with its mathematical form and its intention.

Ordinary-language question	Mathematical formulation	Intention
Is the opening adequate?	$P(G \geq G^* \mid E)$	Estimate
Is the opening inadequate?	$P(G < G^* \mid E)$	Estimate
Will a test accept an adequate opening?	$P(A(\tilde{E}) = 1 \mid G)$, for each $G \geq G^*$	Predict
What is the probability of both units being unavailable?	$P(F_A \cap F_B \mid E)$	Predict
Which fault-tree structure is supported?	$P(H_{\mathrm{AND}} \mid E)$	Compare
What is the probability of overheating at full load with one unit out?	$P(\Delta T_{\mathrm{full}} > 8\,\mathrm{K} \mid E)$	Estimate

The table's point lies in what all six sentences share. They concern one physical plant, and that plant did not choose among them. Its expressions draw variously on a parameter of the opening, on readings not yet taken, on the states of two units and on the plant's structure, and a model of the plant would be needed for every one of them. The physical model does not determine the question, so the engineer must formulate the question separately from building the model. A single probabilistic model of a physical network can support all of these queries, as TERRA, the author's implementation of such a model, does; which query is the right one is not something the model can say. The table makes no claim about how its rows relate to one another, and Lesson 9 takes up when two sentences name one quantity.

Two rows repay a closer reading against the earlier subsections. The third is a question about a test rather than about the plant. Its outcome is a prediction about readings not yet taken, so it is filed under that intention, and Lesson 16 writes it in full. The fourth carries no horizon, which makes it a first instance of this lesson's failure class. As written, it means a random hour, and over a year it would be a different expression. Read with the definition above, it also concerns the units' own outages only, so a feeder outage, which leaves both units unavailable, lies outside it; Exercise 1.7 asks what that omission costs.

1.5 The exercises

The exercises are of four kinds: formulate, reverse, critique and repair, and rapid. The solutions are in Appendix A, and the facts of the plant are in Appendix D.

Exercise 1.1 (*formulate*). The owner of the two-hall plant says: "Each hall carries 60 kW and the drawing shows one spare unit. Before the handover is signed, the opening must be shown to carry a hall's load if a unit trips at full load, and it must be decided whether a third unit should be bought instead." Sort each clause by kind. Name the intention of each question. Run the seven steps on the question that can be written today, and state what the other one lacks.

Exercise 1.2 (*formulate*). A cooling system has two chillers supported by a common electrical feeder. The owner wants to know whether commissioning data provide sufficient confidence to accept the redundancy claim, and whether testing at higher load would reduce the risk of accepting an inadequate system. Identify the separate questions and name the intention of each. Sort the statements by kind. Write the expression for the question this lesson reaches, and state in words what the expression of the other must contain.

Exercise 1.3 (*formulate*). Unit B has just tripped at 02:00, and the operator asks: "Will hall B be all right until the repair crew finishes?" Run the seven steps.

Exercise 1.4 (*formulate*). A prospective tenant asks the operator: "Can the cooling survive a failure?" Formulate the question.

Exercise 1.5 (*reverse*). Write one plain sentence for $P(G < G^* \mid E)$.

Exercise 1.6 (*reverse*). Write one plain sentence for $P\big(S(t) = 1 \text{ for every } t \in [0, 2\ \text{yr}] \mid E\big)$.

Exercise 1.7 (*reverse*). Write one plain sentence for $P(F_A \cap F_B \mid E)$, and state what must be added before it can be evaluated.

Exercise 1.8 (*critique and repair*). An owner asks for the chance of getting through next year without an interruption. An analyst computes $P(S(t) = 1 \mid \text{rates}, G \geq G^*) = 0.988$ and reports "a 98.8 % chance of an uninterrupted year". State what the expression answers, and repair it.

Exercise 1.9 (*critique and repair*). Asked whether the service is reliable, an analyst reports $P(G \geq G^* \mid E) = 0.99$ and calls this the service reliability. State what the expression answers, and repair it.

Exercise 1.10 (*critique and repair*). An analyst writes: "The conductance is uncertain because the readings scatter, so this is aleatory uncertainty, and averaging over a longer test will remove it." Name each error, and repair the statement.

Exercise 1.11 (*rapid*). For each sentence, give its kind and, if it is a question, its intention. Where a quantity is uncertain, say which kind of uncertainty it carries.

(a) Each hall has one cooling unit.

(b) What is the conductance of the opening?

(c) Will unit A fail next month?

(d) Should a third unit be added?

(e) Which set point saves the most electricity and keeps hall B within its limit?

(f) Hall B read 31°C last night: is it the sensor or the opening?

(g) Should the trip test run at 40 or 60 kW?

(h) Do the records favor the parallel or the series structure?

(i) Is the hall A sensor's offset more than 0.5 K?

(j) How many times will the feeder fail next year?

1.6 What to retain

- A sentence is not yet a question until it names its event and conditions, its horizon where time enters, and its level where a verdict is asked. Saying what is missing is a correct answer.
- A statement is sorted first, as a physical fact, an uncertain quantity, a decision or a prediction, and the intention is named next. The expression follows from the intention.
- Epistemic uncertainty concerns a quantity with one value, and measurement reduces it. Aleatory uncertainty remains when everything knowable is known.
- The seven steps end by reading the expression back against the original sentence, which is the check available when nothing has been computed.
- Availability at a random hour and no interruption in a year are both called reliability, and on one plant they are 0.988 and 0.020.
- The physical model does not determine the question.

Lesson 2

From Words to Events

2.1 The engineering question

Two sentences about the redundancy of the two-hall plant can be heard in the same meeting. One says that the plant keeps its cooling unless both units are down. The other says that the plant keeps its cooling as long as neither unit is down. They sound like two phrasings of one claim, and they are not. The first is the claim of N+1: the cooling is promised whenever the units are not both down, so one unit may be down. The second promises the cooling only when nothing is down at all, which is all a plant without a spare could promise. One small word separates the two sentences. On this plant, counting the units' own outages only, the chance that the plant is in a state the first sentence does not cover is about forty times smaller than for the second; counting the feeder's outages too, which neither sentence covers, the factor is about eight.

Sentences about a plant are assembled from a short list of such words: not, and, or, neither, at least, at most, unless, only if, throughout, at some time, and given. Each of them except the last is a precise operation on the statements it joins. Once the states the plant can be in have been named, translating a sentence into the event it describes is close to mechanical, and the errors come from reading the small words casually. The exception, given, joins nothing. It says which part of a sentence is taken as known.

This lesson names the states of the plant, translates each of these words into an operation, builds the events of a few sentences from their parts, and reads each event back word by word. Lesson 1 wrote an event whole, as a phrase, and checked it only as a whole at the last step. Building it from parts lets a defect be found at the word that caused it.

2.2 The formulation

The work of this lesson sits in the fourth step of the method, where the event is defined. The other steps run as in Lesson 1 and are condensed here.

2.2.1 Outcomes and events

At a given hour the plant is in one of a small number of states. The feeder is up or out, each unit is up or out for its own reasons, and the opening is adequate or not. Each complete state is an *outcome*, and the set of outcomes is written Ω. An *event* is a set of outcomes, which is the same thing as a statement that is either true or false of the plant at that hour. The units' own outages F_A and F_B are events, and so is the feeder's outage, written F_{fd} and kept separate from the other two. The opening's adequacy $\{G \geq G^*\}$ is an event, and so is $\{S(t) = 1\}$, the delivery of the service.

Table 2.1 lists the states of the three components, one row each, and its last two columns say whether the service is delivered with the opening adequate and with it inadequate. It is the definition of $S(t)$ written out state by state, and it is all this lesson needs to know about the plant.

Table 2.1: The states of the feeder and the two units at one hour, and whether the service is delivered on each side of the opening's threshold.

Feeder out	Unit A's own outage	Unit B's own outage	Served, opening adequate	Served, opening inadequate
no	no	no	yes	yes
no	yes	no	yes	no
no	no	yes	yes	no
no	yes	yes	no	no
yes	no	no	no	no
yes	yes	no	no	no
yes	no	yes	no	no
yes	yes	yes	no	no

The table shows how few rows carry the plant's interest. With the feeder out, nothing else matters. With the feeder up, the opening matters in exactly the two rows where one unit, and only one, has its own outage.

2.2.2 The connectives

Table 2.2 is a dictionary of the words. Its first column gives a word, its second the operation on events that the word performs, and its third an instance on the plant. Script capitals $\mathcal{A}$ and $\mathcal{B}$ stand for any two events.

Table 2.2: The connectives of engineering sentences as operations on events.

Word	Operation	Instance on the plant
not, fails to	complement, $\mathcal{A}^c$	the feeder is not out: F_{fd}^c
and, both	intersection, $\mathcal{A} \cap \mathcal{B}$	both units are out: $F_A \cap F_B$
or, either	union, $\mathcal{A} \cup \mathcal{B}$, inclusive unless stated	either unit is out: $F_A \cup F_B$
exactly one	$(\mathcal{A} \cap \mathcal{B}^c) \cup (\mathcal{A}^c \cap \mathcal{B})$	exactly one unit is out
neither	$\mathcal{A}^c \cap \mathcal{B}^c$	neither unit is out
not both	$(\mathcal{A} \cap \mathcal{B})^c$	the units are not both out
at least one of	union	at least one unit is up: $F_A^c \cup F_B^c$
at least, at most; more than, fewer than	$\geq$, $\leq$; $>$, $<$ on a quantity or a count	the feeder fails at least twice; $G \geq G^*$
$\mathcal{A}$ unless $\mathcal{B}$	$\mathcal{A} \cup \mathcal{B}$	hall B is within its limit unless the feeder is out
$\mathcal{A}$ only if $\mathcal{B}$; if $\mathcal{A}$ then $\mathcal{B}$	$\mathcal{A}^c \cup \mathcal{B}$	unit A is taken out only if unit B is up
throughout, at every time in	intersection over times	the service is delivered throughout the year
at some time in, ever	union over times	an interruption at some time in the year
given, when, with ... out	no operation; the clause is taken as known	with unit B out tonight

Most rows translate as expected, and three deserve a sentence each. "Or" is inclusive: "either unit is out" includes the case of both, and a sentence that means to exclude it says "exactly one" or "but not both". "Unless" is a union, because "$\mathcal{A}$ unless $\mathcal{B}$" requires $\mathcal{A}$ only on the outcomes

where $\mathcal{B}$ fails, and is satisfied wherever $\mathcal{B}$ holds. Applied to the redundancy claim of the opening section, "the cooling is kept unless both units are down" is the event $\{S(t) = 1\} \cup (F_A \cap F_B)$, which demands the cooling on exactly the outcomes $(F_A \cap F_B)^c$: those are the outcomes the claim covers. "Only if" is the reverse of what it sounds like: "unit A is taken out only if unit B is up" forbids one combination, unit A out with unit B also out, and allows every other.

The two rows on "neither" and "not both" are linked by De Morgan's laws,

$$(\mathcal{A} \cup \mathcal{B})^c = \mathcal{A}^c \cap \mathcal{B}^c, \qquad (\mathcal{A} \cap \mathcal{B})^c = \mathcal{A}^c \cup \mathcal{B}^c. \tag{2.1}$$

In words, the first says that "not either" is "neither", and the second says that "not both" is "at least one is not". Equations (2.1) are the source of the most consequential slip of this lesson, which the competing formulations below put a size on.

2.2.3 Thresholds, counts and time

Whether "at least" and "more than" differ depends on what is being compared. The conductance can take any value in a range, and the boundary value G^* itself carries no probability, so $G \geq G^*$ and $G > G^*$ are different events with the same probability. The slip is harmless there, although the translation should still follow the words. A count is different. If $N_{\text{fd}}(\tau)$ is the number of the feeder's failures in $[0, \tau]$, then two failures are at least two and not more than two, so $N_{\text{fd}}(\tau) \geq 2$ and $N_{\text{fd}}(\tau) > 2$ differ by the probability that the count is exactly two.

The words about time are connectives too. The year of Lesson 1, the service delivered at every hour, is an intersection over the hours of the year,

$$\{S(t) = 1 \text{ for every } t \in [0, \tau]\} = \bigcap_{t \in [0,\tau]} \{S(t) = 1\}, \tag{2.2}$$

and the event of equation (1.2) is built from "throughout". By the second of equations (2.1), applied over every hour, the complement of the event in equation (2.2) is a union, an interruption at some time in the year. A third kind of word ranges over something other than hours. In "out for more than a week at a time", the words "at a time" range over the separate outages, and the second worked formulation below turns on that.

2.2.4 The two gates

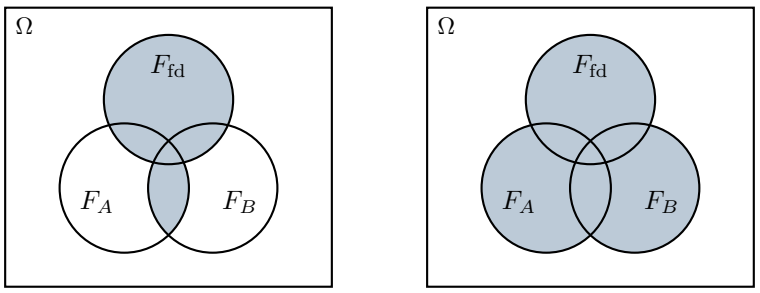
A fault tree combines failure events through gates, and its two basic gates are the intersection and the union: an AND gate passes when all its inputs hold, an OR gate when any of them does. The loss of the service is built from them. With the opening adequate, the service is lost when the feeder is out or both units are out,

$$\{S(t) = 0\} = F_{\text{fd}} \cup (F_A \cap F_B) \qquad \text{on } \{G \geq G^*\}, \tag{2.3}$$

which is the parallel structure Lesson 1 glossed as H_{AND}, after its AND gate. With the opening inadequate, either unit's loss is enough,

$$\{S(t) = 0\} = F_{\text{fd}} \cup F_A \cup F_B \qquad \text{on } \{G < G^*\}, \tag{2.4}$$

the series structure H_{OR}. Figure 2.1 shades the two events on one diagram of the three outages, the feeder's at the top and the units' below it. In the left panel the shading covers the feeder's region and only the overlap of the two units', which is equation (2.3); in the right panel it covers all three regions, which is equation (2.4). The difference between the panels is the two crescents where exactly one unit is out and the feeder is up, the two rows of Table 2.1 in which the opening matters. Here the names H_{AND} and H_{OR} label two loss events only. Treating the structure itself as an unknown to be compared is taught around Lessons 6 and 7.



opening adequate: the AND gate opening inadequate: the OR gate

Figure 2.1: The loss of the service on each side of the opening's threshold, shaded on one diagram of the feeder's outage and the units' own outages. The sizes of the regions are illustrative.

2.2.5 The word given

"Given" and its companions, such as "when", "with unit B out" and "assuming", combine nothing. They mark a clause as known, and a known clause goes behind the bar of the probability rather than into its event. A translation therefore sorts the clauses of a sentence first, into those asked about and those given, and builds only the first kind into the event. What it means for a clause to stand behind the bar, and how that differs from joining it with "and", is the subject of Lesson 3.

2.2.6 The service rebuilt

The first worked formulation rebuilds an event Lesson 1 wrote in words. There the service was delivered when the feeder is up and either both units are up or one unit is up with the opening adequate. Table 2.3 translates the definition clause by clause: the clause in the first column, the word that decides it in the second, and its event in the third.

Table 2.3: Lesson 1's definition of the service, translated clause by clause.

Clause	Word	Event
the feeder is up	not	F_{fd}^c
and	and	$\cap$
both units are up	neither out	$F_A^c \cap F_B^c$
or	or	$\cup$
one unit is up	exactly one out	$(F_A \cap F_B^c) \cup (F_A^c \cap F_B)$
with the opening adequate	and	$\cap \{G \geq G^*\}$

Assembled, the rows give

$$\{S(t) = 1\} = F_{\text{fd}}^c \cap \Big[(F_A^c \cap F_B^c) \cup \Big(((F_A \cap F_B^c) \cup (F_A^c \cap F_B)) \cap \{G \geq G^*\}\Big)\Big]. \tag{2.5}$$

Reading equation (2.5) back clause by clause at step 7 yields two findings. The first concerns "one unit is up", which the table reads as exactly one. Read inclusively, as at least one, it gives the same event, because the case of both units up is already in the first branch. Whether "either" is exclusive can therefore matter or not, depending on the rest of the sentence. The second finding comes from splitting the event on the opening. On $\{G \geq G^*\}$, equation (2.5) reduces to the feeder up and the units not both out, whose complement is equation (2.3); on $\{G < G^*\}$ it reduces to everything up, whose complement is equation (2.4). The two structures are therefore not extra assumptions. They are what one sentence says on each side of G^*. In

the rest of the formulation sheet the target is $S(t)$ at a random hour, the prediction on which the decision about "reliable enough" rests, and the expression is $P(S(t) = 1 \mid E)$.

2.2.7 Counts and durations over a year

The second worked formulation is new. The facilities manager asks for the chance that next year the feeder fails no more than once, and that neither unit is ever out for longer than a week at a time. The words translate in turn. "No more than once" is at most one. "And" is an intersection. "Ever out for longer than a week at a time" is a union over a unit's outages in the year of the event that the outage lasts more than 168 h. "Neither unit" asks for the complement of that union for each unit, and by De Morgan's laws, equations (2.1), the complement of a union is the intersection of the complements: every outage lasts at most 168 h. The event is

$$\{N_{\mathrm{fd}}(\tau) \le 1\} \cap \bigcap_{i \in \{\mathrm{A,B}\}} \{\text{every outage of unit } i \text{ in } [0,\tau] \text{ lasts at most 168 h}\}, \tag{2.6}$$

with $\tau = 1$ yr, and the question asks for its probability given that everything is up at $t = 0$ and given E, the start standing behind the bar. Table 2.4 is the filled formulation sheet, with the fields down the left and their entries on the right; the fourth row is where this lesson's work lies.

The connectives have reached the sixth step as well. The words "at a time" put into the model the whole distribution of the length of a single repair. A sentence about the expected total time out would have needed only the mean length of a repair, while a probability about the total, such as its exceeding a week, needs the total's distribution too. The reading back at step 7 also marks a limit. Equation (2.6) says nothing about the service, since both units can be out together for a day without either outage lasting a week, so a manager who meant "no long loss of cooling" has asked a different question, whose event is built from equation (2.3) rather than from the components one by one.

Key idea

An event is checked one word at a time. Each connective of the sentence should appear in the event as one operation, and each operation in the event should be traceable to one word.

2.3 The competing formulations

Each competing formulation of this lesson is one word of a sentence given the wrong operation. The first changes the meaning of a redundancy claim. The claim that the cooling is kept unless both units are down, $\{S(t) = 1\} \cup (F_A \cap F_B)$, promises the cooling on the outcomes $(F_A \cap F_B)^c$. Written with "neither", it promises the cooling only on $F_A^c \cap F_B^c$. Table 2.5 gives both sets of covered outcomes at a random hour, one row each, with the probability of the set and of its complement, the outcomes the claim leaves uncovered. The values are for the units' own outages, with the stated rates of 2.5 failures per year in service and a 180 h repair for each unit, the two units failing independently and each with its own repair.

For the units' own outages, the uncovered shares differ by a factor of about 40. The feeder's outages are covered by neither sentence, and counting them, with the opening taken as adequate, the uncovered shares become 0.012 and 0.104, a factor of about eight. The first row is what the drawing's N+1 promises, and the second is what a plant with no spare would need. A report that translates the claim with "neither" credits the plant with no redundancy at all, and nothing in its number shows the error until De Morgan's laws are applied to the sentence.

Table 2.4: The formulation sheet filled in for the facilities manager's question.

Field	Entry
Intention	Predict
1. Sentence	Starting with everything up, what is the probability that in the coming year the feeder fails at most once and that every outage of each unit lasts at most 168 h?
2. Target	A future outcome: the feeder's count and the lengths of the units' outages over the year
3. Known	All up at $t = 0$; E: the stated rates and repair times, any outage record. Conditions: steady operation.
4. Event	Equation (2.6): "no more than once" as ≤ 1; "and" as an intersection; "ever ... longer than a week at a time" as a union over outages; "neither" as its complement for each unit, which De Morgan's laws turn into "every outage at most 168 h"
5. Expression	$P(\text{the event of equation (2.6)} \mid \text{all up at } t = 0,\ E)$
6. Model	The feeder's failure process over the year; each unit's failures and the distribution of the length of a single repair
7. Read back	A quiet feeder and no long repair. Silent on whether the service was lost.
Problem class	Predictive probability over a horizon

The same sheet as a run:

Run, the facilities manager's question

1	**Sentence**	All up at the start, next year: feeder fails at most once, every outage at most 168 h?
	Intention	Predict
2	**Target**	Future outcome: the feeder's count; the lengths of the units' outages
	Declare	$N_{\text{fd}}(\tau)$ the feeder's failures in $[0, \tau]$; $\tau = 1$ yr; E rates, repair times, any record
3	**Known**	All up at $t = 0$; E; steady operation
	Random here	$N_{\text{fd}}(\tau)$; the number and length of each unit's outages
4	**Event**	Equation (2.6): at most once; $\cap$; union over outages; its complement per unit
5	**Expression**	$P(\text{the event of equation (2.6)} \mid \text{all up at } t = 0,\ E)$
6	**Model**	The feeder's failure process; each unit's failures; the length distribution of a repair
7	**Check**	A quiet feeder, no long repair; silent on the service
	Class	Predictive probability over a horizon

$$P\Big(\{N_{\text{fd}}(\tau) \leq 1\} \cap \bigcap_{i \in \{\text{A,B}\}} \{\text{every outage of unit } i \text{ lasts at most 168 h}\} \,\Big|\, \text{all up at } t = 0,\ E\Big)$$

Table 2.5: Two translations of one redundancy claim, at a random hour, for the units' own outages on the stated rates.

Translation	Outcomes covered	Covered [–]	Uncovered [–]
Not both units out	$(F_A \cap F_B)^c$	0.998	0.0024
Neither unit out	$F_A^c \cap F_B^c$	0.905	0.095

As a run, with its repair:

Run, "unless both units are down" written with "neither"

1	**Sentence**	The cooling is kept unless both units are down.
4	**Event** ×	Covered $F_A^c \cap F_B^c$, neither unit out; uncovered 0.095
	Repair	Covered $(F_A \cap F_B)^c$, not both out; uncovered 0.0024
7	**Check**	"Unless both" excepts only the joint outage

$$(F_A \cap F_B)^c = F_A^c \cup F_B^c$$

The second competing formulation concerns a count. "The feeder fails at least twice next

year" is $N_{\mathrm{fd}}(\tau) \geq 2$ with τ a calendar year, and with the feeder up at $t = 0$, its stated rate of 3.69 failures per year in service and its 24 h repair, its probability is 0.882. "More than twice" is $N_{\mathrm{fd}}(\tau) > 2$, with probability 0.709. On the conductance the same slip changes nothing, as the previous section showed. The boundary matters exactly where the quantity can sit on it.

The third reads "either" as exclusive. "The service is lost if either unit fails", said of the series structure, read as exactly one unit, drops the case of both units out, which loses the service on either structure. The two readings differ by the probability of both units out with the feeder up, 0.0024 at a random hour on the stated rates, so the slip passes unnoticed in a number while being wrong in kind. The fourth drops "unless". "Hall B stays within its limit unless the feeder is out", written as the limit alone, demands the limit in the hours of a feeder outage too, which is a harsher requirement than the sentence states; written as "within the limit and the feeder up", it makes those hours failures. The translation is the union. Whether an exclusion of this kind belongs in the event at all is a question about requirements, and Lesson 4 takes it up.

These slips form this lesson's failure class, the *connective mistranslated*: a word of the sentence given the wrong operation. Its forms are "neither" for "not both" or the reverse, "either" read as exclusive where the sentence includes both, "more than" for "at least" on a count, "unless" dropped or read as "and not", and a word ranging over outages read as one over total time. It enters at step 4 and is caught at step 7, by reading the event back one word at a time. Appendix B scores it with the other classes.

2.4 One system, many questions

Table 2.6 gives ten sentences about the two-hall plant, with the event of each in the second column and the word that decides it in the third. No probability appears in it, because the sentences differ in their events before anything is computed.

Table 2.6: Ten sentences about the two-hall plant and their events.

Sentence	Event	Deciding word
Both units are out.	$F_A \cap F_B$	both
At least one unit is up.	$F_A^c \cup F_B^c$	at least one
Neither the feeder nor unit A is out.	$F_{\mathrm{fd}}^c \cap F_A^c$	neither
Unit A is out but unit B is not.	$F_A \cap F_B^c$	but not
Exactly one of the three components is out.	$(F_{\mathrm{fd}} \cap F_A^c \cap F_B^c) \cup (F_{\mathrm{fd}}^c \cap F_A \cap F_B^c) \cup (F_{\mathrm{fd}}^c \cap F_A^c \cap F_B)$	exactly
The service is lost, with the opening adequate.	$F_{\mathrm{fd}} \cup (F_A \cap F_B)$	AND gate
The service is lost, with the opening inadequate.	$F_{\mathrm{fd}} \cup F_A \cup F_B$	OR gate
The feeder fails at most once next year.	$N_{\mathrm{fd}}(\tau) \leq 1$	at most
The opening passes at least one hall's load at the limit.	$G \geq G^*$	at least
The service is interrupted at some time next year.	$\bigcup_{t \in [0,\tau]} \{S(t) = 0\}$	at some time

Down the table, the plant and the nouns stay the same, the units, the feeder, the service and the opening, and only the connectives change. Each change of connective is a change of event, and the probabilities of the events can lie as far apart as Table 2.5 showed.

2.5 The exercises

The solutions are in Appendix A, and the facts of the plant in Appendix D.

Exercise 2.1 (*formulate*). Unit B tripped at 01:00, with the feeder and unit A up, and the repair crew will have it back at 06:00. The shift supervisor asks: "What is the chance that hall B keeps its cooling until 06:00, that is, that neither the feeder nor unit A goes down at any time before then?" Translate each connective, write the event and the expression, and check the supervisor's own gloss against the question.

Exercise 2.2 (*formulate*). An insurer's draft clause reads: "The policy pays if the cooling of both halls is lost at once, or if either hall is without cooling for more than 24 consecutive hours, in a policy year." Write each hall's loss of cooling as an event on the outages of the feeder and the units, and then write the event on which the policy pays.

Exercise 2.3 (*formulate*). A tenant's draft contract says: "The operator will have at most one outage next year." Formulate the sentence.

Exercise 2.4 (*reverse*). Write one plain sentence for $P\big(F_{\mathrm{fd}} \cup (F_A \cap F_B) \mid E\big)$, at a random hour.

Exercise 2.5 (*reverse*). Write one plain sentence for $P\big(N_{\mathrm{fd}}(\tau) \le 1 \mid \text{all up at } t = 0,\ E\big)$, with $\tau = 1$ yr, and state how it differs from the same expression with $N_{\mathrm{fd}}(\tau) < 1$.

Exercise 2.6 (*reverse*). Write one plain sentence for $P\big(F_{\mathrm{fd}}^{c} \cap ((F_A \cap F_B^{c}) \cup (F_A^{c} \cap F_B)) \mid E\big)$, at a random hour.

Exercise 2.7 (*critique and repair*). Asked for the probability that at least one unit is available at a random hour, an analyst writes $P(F_A^{c} \cap F_B^{c} \mid E)$. State what the expression answers, and repair it.

Exercise 2.8 (*critique and repair*). A design note translates "hall B stays within its limit unless the feeder is out" as the event $\{\text{hall B within its limit}\} \cap F_{\mathrm{fd}}^{c}$. State what the event describes, and repair it.

Exercise 2.9 (*rapid*). Translate each sentence into an event on the plant. Where a sentence holds a condition rather than an event, say so.

(a) Neither unit is out while the feeder is up.

(b) At most one of the three components is out.

(c) Unit B is out only if unit A is up.

(d) The feeder or unit B is out, but not both.

(e) The feeder fails more than three times next year.

(f) The feeder fails between two and four times next year, inclusive.

(g) The feeder is out at least once next year.

(h) The service holds throughout the repair of unit B.

(i) The cooling is kept unless both units are out.

(j) Unit A is out for more than a week in total next year.

(k) Unit B is out, given that the feeder is up.

2.6 What to retain

- The outcomes come first. An event is a set of them, and every connective of a sentence is an operation on such sets.
- “Neither” is not “not both”. The N+1 claim is a claim of “not both”, and on this plant the outcomes the two translations leave uncovered differ in probability by a factor of about 40 for the units’ own outages, and about eight with the feeder’s outages counted.
- “At least” and “more than” differ on a count and not on a quantity that varies continuously.
- “Either” is inclusive unless the sentence says otherwise, and whether the reading matters depends on the rest of the sentence.
- “Unless” is a union, and “only if” forbids one combination.
- “Throughout” and “at some time” are an intersection and a union over time, and “at a time” ranges over separate outages.
- “Given” builds no event. It moves its clause behind the bar.

Lesson 3

What Is Known, Unknown, and Conditioned Upon?

3.1 The engineering question

The review after a night interruption of the two-hall plant opens with two statements. The electrical engineer says that whenever the feeder goes, the cooling goes. The operations manager concludes that last night the feeder went. Both speak of the same two happenings, the outage of the feeder and the loss of the cooling, but in opposite directions, and the truth of the first says little about the second. The cooling is also lost when both units are down, and nothing in "whenever the feeder goes" says how often that happens by comparison. A third person asks how often the feeder is out and the cooling lost together, which is a third quantity again.

Later in the same review, the crew reports that unit B was found off when they arrived. Someone remarks that if unit B had been switched off deliberately, for a test, the same state of unit B would say something different about unit A. The remark is a good one, and this lesson returns to it at the end of its formulation.

The work of this lesson is to sort the parts of a sentence before writing any probability. Each clause is either what the speaker wants to find out, or what the speaker already knows, or neither. Once the clauses are sorted, the direction of the probability follows, and the direction decides what the model must contain. Lesson 2 built events from the words of a sentence and set the word "given" aside. This lesson decides where every clause goes.

3.2 The formulation

The work of this lesson sits in steps 3 and 5 of the method, what is known and the expression, and through them it reaches step 6. The other steps are condensed.

3.2.1 Three places for a statement

Lesson 1 sorted sentences by kind and gave each a place in a formulation. The places can be made exact, and there are three. A clause that the sentence wants to find out is *asked about*, and it goes to the left of the bar. A clause that is already known goes *behind the bar*: readings already taken, records already kept, stated facts, operating conditions, a present state, a start, and any assumption taken as true. Everything behind the bar is taken as true, and the probability then describes the plant in a world where it is. A quantity that is neither known nor asked about is *carried in the model* and averaged over. The conductance in $P(S(t) = 1 \mid E)$ is of this third kind: the expression neither asks about it nor knows it, and its value still matters.

One question per clause does the sorting: is it known now, is it what is wanted, or is it neither? Table 3.1 applies the question to three sentences about the plant, with the clause

asked about in the second column, the clauses known in the third, and the quantities carried in the fourth.

Table 3.1: Three sentences about the two-hall plant, with each clause in its place.

Sentence	Asked about	Known, behind the bar	Carried in the model
After the trip test, how probable is it that the opening is adequate?	$G \geq G^*$	E: the test's readings, the drawing's prior, the stated sensor spreads	the sensor offsets, the hall temperatures
With unit B out tonight, will hall B hold until the repair is done?	the path of S until unit B returns	unit B out at t_0, tonight's load, E	G; the failures of unit A and the feeder; the repair time
The cooling was lost at 03:00; was the feeder the cause?	F_{fd} at 03:00	$S(03{:}00) = 0$, E	the units' states, G

The second row is Lesson 1's Exercise 1.3, and the table says why each part of its expression stands where it does. The repair time is not known, so it cannot go behind the bar, and it is not what the operator asked about, so it is carried. The third row is the review's question, and the rest of this lesson works on it.

3.2.2 The conditional as a restriction

To condition on a known event is to set aside every outcome in which it does not hold. The conditional probability is then the share, among the outcomes that remain, of those in which the event asked about also holds,

$$P(\mathcal{A} \mid \mathcal{B}) = \frac{P(\mathcal{A} \cap \mathcal{B})}{P(\mathcal{B})}, \qquad P(\mathcal{B}) > 0. \tag{3.1}$$

In equation (3.1), $\mathcal{B}$ is the known event and $\mathcal{A}$ the event asked about. The numerator is the *joint* probability, the share of all outcomes in which both hold, and the denominator restricts attention to $\mathcal{B}$. The equation defines the conditional; it is displayed here to say what the bar means, and no lesson uses it to compute.

Figure 3.1 draws the case of the review, with the opening taken as adequate. In each panel the large rectangle is every outcome at one hour, and the loss of the service on the parallel structure is the union of two regions that do not overlap: the feeder out, and the feeder up with both units out. The two regions partition the loss, since every lost hour has the feeder either out or up. The three panels shade one numerator each and draw its denominator heavy: the feeder's region against everything, which is the joint; the feeder's region against itself, which is the forward conditional, since a feeder outage always loses the service; and the feeder's region against the whole loss, which is the inverted one.

The middle panel is the electrical engineer's statement: the feeder's region lies wholly inside the loss, so the share is one. The right panel is the operations manager's conclusion, and its share is less than one by exactly the region in which the feeder is up and both units are out.

3.2.3 The direction of the bar

The two directions of a conditional are related. Applying equation (3.1) to both and equating the joint probabilities gives Bayes' rule,

$$P(\mathcal{B} \mid \mathcal{A}) = \frac{P(\mathcal{A} \mid \mathcal{B})\, P(\mathcal{B})}{P(\mathcal{A})}. \tag{3.2}$$

Equation (3.2) is displayed for what it says about formulation, and it is not used to compute anything in this lesson. It says that the inverted probability needs two things the forward one

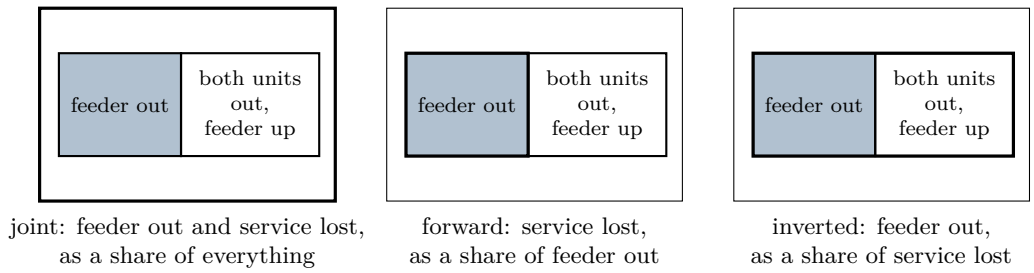


Figure 3.1: One pair of events read three ways, with the opening taken as adequate. The loss of the service is the union of the two inner regions, which do not overlap; the shading is the numerator, and the heavy outline the denominator. The sizes of the regions are illustrative.

does not: how probable $\mathcal{B}$ was before anything was known of $\mathcal{A}$, and how probable $\mathcal{A}$ is overall. On the plant, $P(S(t) = 0 \mid F_{\text{fd}}) = 1$ follows from the structure alone, since both units lose their supply with the feeder, and it needs no rate, reading or record. The inverted $P(F_{\text{fd}} \mid S(t) = 0)$ needs the rates of every cause of a loss, because it asks how the losses divide among them. Turning the bar around has changed what step 6 must supply.

Reporting one direction as the other has a standard name, the *transposed conditional*; in courts it is known as the prosecutor's fallacy. Its most consequential instance in this book concerns evidence. The probability $P(E \mid G)$, for each value $G \geq G^*$, says how expected the readings of a test are if the opening has that adequate conductance. The probability $P(G \geq G^* \mid E)$ says how probable adequacy is after the readings. The second needs what was believed about the conductance before the test, and how the same readings would behave if the opening were inadequate; the first needs neither. A single number $P(E \mid G \geq G^*)$ would already need part of the first ingredient, since it averages over the adequate values in proportion to what was believed about them. Lessons 11 and 12 write those two ingredients, and this lesson names them.

3.2.4 What exactly is known

What stands behind the bar should be exactly what is known, and three cases come up often. The first is a reading. When hall B's sensor reads 31°C, the reading is known and the hall's temperature is not, since the sensor has an offset and noise; Lesson 11 writes the difference. The second is a verdict. Knowing only that a test passed is knowing less than its readings, and a probability conditioned on the verdict answers a different question from one conditioned on the readings. The third is an assumption. Lesson 1 placed the opening's adequacy behind the bar in its Table 1.7, labeled, to isolate the effect of the formulation, and its 0.988 is a statement about a plant whose opening is adequate. Conditioning on an assumption is legitimate. The error is to drop its label and report $P(S(t) = 1 \mid \text{rates}, G \geq G^*)$ as the availability of the plant as it is.

3.2.5 Setting against observing

The crew's report from the review raises one more distinction, which this lesson notes and Lesson 17 develops. Unit B found off and unit B switched off put the same state of unit B behind the bar, and they are different conditions. Finding it off is an observation, and it is evidence about whatever could have taken it down. Switching it off sets its state, and it says nothing about those causes. The book writes the second as do(unit B off), an intervention.

The fact sheet's disturbance windows make the difference concrete. A window of 720 h is disturbed with prior weight 0.1, and in a disturbed window each unit is unavailable with

probability 0.30, against 0.05 in a calm one; within a window of either kind the two units fail independently. Two further assumptions make the comparison clean, and they are stated together: unit B is switched off without regard to the state of the window, and unit A's own failures do not depend on its carrying both halls while unit B is off. Table 3.2 gives unit A's chance of being out under the two conditions on unit B, one row each.

Table 3.2: Unit A's outage when unit B is found out and when it is switched off, on the fact sheet's model of disturbance windows.

Condition on unit B	Expression	Value [–]
Found out	$P(F_A \mid F_B,$ window model)	0.150
Switched off	$P(F_A \mid$ do(unit B off), window model)	0.075

The second row is unit A's unavailability before anything is known of unit B, which on this model is 0.075. Finding unit B out doubles it, because a unit found out is more likely to sit in a disturbed window: its weight rises from 0.1 to 0.4. Without the shared cause, both rows would be 0.05. Figure 3.2 sketches why, with three boxes in each panel, the state of the window above the two units and an arrow from the window to each. In the left panel, unit B found out, a dashed arrow runs back from unit B to the window and down to unit A, which is the path the observation informs. In the right panel, unit B switched off, the arrow into unit B is struck through, so the setting carries no news upward.

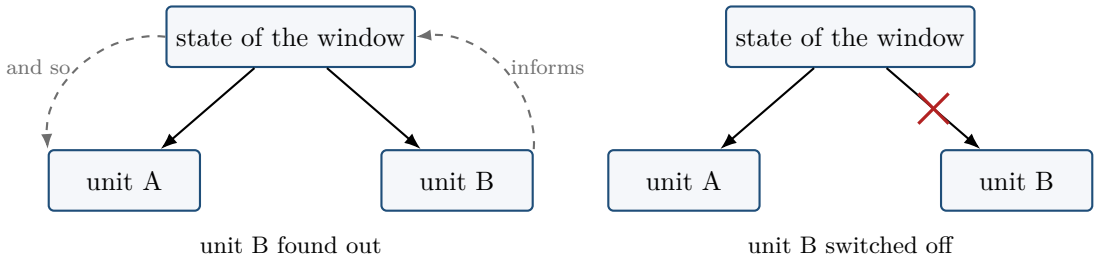


Figure 3.2: Seen against set, as a sketch in words. Solid arrows are the shared cause acting on each unit; the dashed arrows are what an observation of unit B tells about the window and, through it, about unit A. Setting unit B cuts the arrow into it.

The shared cause has a standard name, a common cause, and Lesson 6 teaches the dependence it creates. The failure class of reading a setting as an observation belongs to Lesson 17, where setting a control becomes a decision.

3.2.6 Diagnosing an interruption

The review's question can now be written. The service was lost at t_0, 03:00 last night, and the operations manager wants to know how probable it is that the feeder caused it. Its intention is to diagnose. Table 3.3 is the filled formulation sheet, with steps 3, 5 and 6 in full.

The sheet's fourth row shows that a diagnosis is a probability spread over a list of candidates, and the list must be a partition of what was seen. The feeder's outage and the units' own outages can occur together, so "the feeder out" and "both units out" overlap as events; the partition separates them by asking for the feeder up in every candidate after the first, and only then do the candidates' probabilities sum to one. The list also depends on what is carried. With the opening carried as uncertain, the feeder up with a single unit out is a candidate too. With no log, the stated rates and the opening taken as adequate, the expression has a value, and the competing formulations below use it.

Table 3.3: The formulation sheet filled in for "Was the feeder the cause of last night's loss?".

Field	Entry
Intention	Diagnose
1. Sentence	Given that the service was lost at t_0, and given what is known, what is the probability that the feeder was out at t_0?
2. Target	An event at a past time, F_{fd} at t_0, one of the candidate causes
3. Known	$S(t_0) = 0$; E: the structure, the stated rates, any relay or alarm log. Carried: the units' states and G.
4. Event	F_{fd} at t_0, one member of a partition of the loss: the feeder out; the feeder up with both units out; and, if the opening may be inadequate, the feeder up with exactly one unit out and $G < G^*$.
5. Expression	$P\big(F_{\text{fd}}(t_0) \mid S(t_0) = 0,\ E\big)$
6. Model	The structure; how probable each cause is, from the rates of all of them; with a log in E, how the log behaves under each cause. The forward $P(S(t_0) = 0 \mid F_{\text{fd}}(t_0))$ would have needed the structure only.
7. Read back	The probability that the loss came from the feeder's outage rather than from the other members of the partition. It is not "the feeder always takes the cooling".
Problem class	Posterior probability over candidate causes

In compact form:

Run, was the feeder the cause of last night's loss

1	**Sentence**	Service lost at t_0, given E: probability the feeder was out at t_0?
	Intention	Diagnose
2	**Target**	An event at a past time, F_{fd} at t_0, one candidate cause
	Declare	t_0 03:00 last night; $F_{\text{fd}}(t_0)$ feeder out at t_0; E structure, rates, any log
3	**Known**	$S(t_0) = 0$; E; carried: the units' states, G
	Random here	Which cause acted at t_0
4	**Event**	$F_{\text{fd}}(t_0)$; partition: feeder out; up, both units out; up, one out, $G < G^*$
5	**Expression**	$P\big(F_{\text{fd}}(t_0) \mid S(t_0) = 0,\ E\big)$
6	**Model**	Structure; each cause's probability; with a log, the log under each cause
7	**Check**	The loss came from the feeder; not "the feeder always takes the cooling"
	Class	Posterior probability over candidate causes

$$\boxed{P\big(F_{\text{fd}}(t_0) \mid S(t_0) = 0,\ E\big)}$$

3.2.7 The readings and the opening

A second sentence from commissioning carries both directions at once. The engineer says that if the opening were adequate, a test at 60 kW would read hall B within 8 K of hall A almost every time; that it did read within 8 K; and that the opening is therefore almost certainly adequate. The first clause is a statement about a test given the plant,

$$P\big(\text{a test still to run reads hall B within 8 K of hall A} \mid G\big), \qquad G \geq G^*, \tag{3.3}$$

and the conclusion is a statement about the plant given a test, $P(G \geq G^* \mid E)$, with the readings in E. Step 6 separates the two. Expression (3.3) needs only how the readings behave at a given conductance, through the sensors' offsets and noise. The conclusion needs that, how the readings behave when $G < G^*$, and what was believed about the conductance before the test. The engineer has supplied the first and asserted the second. The first clause is also less certain than it sounds. For a conductance well above G^* a test would pass nearly every time, but near G^* a pass is roughly an even chance, so "almost every time" already assumes that the conductance is not close to the threshold. Lessons 15 and 16 return to the test's side of the bar and to this point.

3.3 The competing formulations

The failure class of this lesson is the *wrong side of the bar*, and it takes three forms. The first, the transposed conditional, has a size on the plant. Table 3.4 gives three expressions on one pair of events, the feeder's outage and the loss of the service at a random hour, one row each, with the expression in the first column, its reading in the second and its value in the third. The values use the stated rates of Lesson 1 (each unit 2.5 failures per year in service with a 180 h repair, the feeder 3.69 with a 24 h repair, all independent, each unit with its own repair) and the opening taken as adequate.

Table 3.4: Three expressions on one pair of events, at a random hour, on the stated rates with the opening taken as adequate.

Expression	Reading	Value [–]
$P(F_{\text{fd}} \cap \{S(t) = 0\} \mid \text{rates}, G \geq G^*)$	the feeder out and the service lost, together	0.010
$P(S(t) = 0 \mid F_{\text{fd}})$	the service lost when the feeder is out, for any G	1
$P(F_{\text{fd}} \mid S(t) = 0,\ \text{rates}, G \geq G^*)$	the feeder out when the service is lost	0.81

An incident report that says "the feeder always takes the cooling, so the feeder caused the loss" reports the second row for the third: 1 for 0.81. The missing 0.19 is the share of losses in which the feeder was up and both units were out together, and that is the very case the N+1 claim concerns.

The inversion and its repair:

Run, "the feeder always takes the cooling, so the feeder caused the loss"

1	**Sentence**	Service lost: was the feeder the cause?
5	**Expression** ×	$P(S(t) = 0 \mid F_{\text{fd}})$: 1
	Repair	$P(F_{\text{fd}} \mid S(t) = 0,\ \text{rates}, G \geq G^*)$: 0.81
7	**Check**	The loss is known, the feeder asked: the loss stands behind the bar

$$\boxed{P(F_{\text{fd}} \mid S(t) = 0,\ \text{rates}, G \geq G^*)}$$

The inversion therefore hides the failures of the redundancy. The same form appears with evidence, when "the readings are very probable if the opening is adequate" is reported as "the opening is very probably adequate", and no number repairs it: the repair is to write $P(G \geq G^* \mid E)$ and name what it needs.

The second form writes a known condition into the event. The first row of Table 3.4, 0.010, reported as "the chance of losing the service when the feeder is out" answers how often the two happen together, which is neither the forward conditional of the second row nor the diagnosis of the third. The third form writes an unknown behind the bar as if it were known: Lesson 1's 0.988, computed with the opening taken as adequate, quoted as the availability of the plant with that condition dropped. A fourth slip, reading a setting as an observation, was shown in Table 3.2 and belongs to Lesson 17.

The class is therefore a statement on the wrong side of the bar, in one of three forms: the asked-about and the known swapped, a known condition written into the event, or an unknown written behind the bar as if known with its label dropped. It enters at step 3 or step 5 and is caught at step 7 by asking of each clause whether it is known now.

3.4 One system, many questions

Table 3.5 gives nine sentences about the two-hall plant, arranged so that neighboring rows use the same events with the bar placed differently. The first column gives the sentence, the second

its expression, and the third its intention.

Table 3.5: Nine sentences about the two-hall plant that differ in what stands behind the bar.

Ordinary-language question	Mathematical formulation	Intention
Is the cooling lost when the feeder is out?	$P(S(t) = 0 \mid F_{\text{fd}})$	Predict
Was the feeder the cause of the lost cooling?	$P(F_{\text{fd}} \mid S(t) = 0,\ E)$	Diagnose
How often are the feeder out and the cooling lost together?	$P(F_{\text{fd}} \cap \{S(t) = 0\} \mid E)$	Predict
Is the opening adequate, after the test?	$P(G \geq G^* \mid E)$	Estimate
Would an adequate opening pass the test?	$P(\text{the test reads within 8 K} \mid G),\ G \geq G^*$	Predict
Is unit A down, now that unit B has been found down?	$P(F_A \mid F_B,\ E)$	Estimate
Is unit A down, now that unit B has been switched off?	$P(F_A \mid \text{do(unit B off)},\ E)$	Estimate
Will hall B hold until the repair, with unit B out?	$P(S(t) = 1 \text{ for every } t \in [t_0, t_0 + D] \mid$ unit B out at $t_0,\ E)$, with D the repair's length	Predict
Was the month disturbed, given that both units failed in it?	$P(\text{disturbed} \mid$ both units failed in the month, $E)$	Diagnose

The plant and its events repeat down the table, and the questions differ only in which statements stand behind the bar. The intention follows what is asked about, not the direction alone. A question about what a cause brings about, at a random time or over a horizon still to come, predicts; a question about which cause lay behind something seen diagnoses; and a question about a state the plant is in now, such as whether unit A is down at this moment, estimates, as Lesson 1 sorted a quantity that has one value.

3.5 The exercises

The solutions are in Appendix A, and the facts of the plant, with its disturbance windows, in Appendix D.

Exercise 3.1 (*formulate*). At 03:00 the fault alarm of unit A and the undervoltage relay of the feeder both fire, and hall B's cooling is lost. The duty engineer asks: "What most probably caused it?" Name the candidate causes, put every clause in its place, and write the expression for each cause.

Exercise 3.2 (*formulate*). The maintenance planner says: "Unit B will be switched off for six hours on Tuesday for service. How likely is the service to hold through it?" Write the expression, and state how it differs from that of Exercise 1.3, in which unit B had tripped.

Exercise 3.3 (*formulate*). The owner asks: "What are the chances the opening is adequate if the test passes?" Formulate the question.

Exercise 3.4 (*reverse*). Write one plain sentence for $P\big(F_{\text{fd}}(t_0) \mid S(t_0) = 0,\ E\big)$.

Exercise 3.5 (*reverse*). Write one plain sentence for $P(S(t) = 0 \mid F_{\text{fd}})$, and state why it needs no evidence.

Exercise 3.6 (*reverse*). Write one plain sentence for $P\big(F_A \mid \text{do(unit B off)},\ E\big)$.

Exercise 3.7 (*critique and repair*). A commissioning engineer writes: “A hall sensor with an offset of more than 0.5 K reads hall B high nine times in ten. The hall B sensor read high, so it very probably has an offset of more than 0.5 K.” State what the argument computes, and repair it.

Exercise 3.8 (*critique and repair*). An analyst writes: “A unit failing at 2.5 per year in service would often log seven failures in three years. Unit A’s record shows seven, so its failure rate is very probably 2.5 per year.” State what the argument computes, and repair it.

Exercise 3.9 (*critique and repair*). An analyst writes $P(F_A \cap F_B \mid E)$ and calls it “the probability that unit A is down when unit B is down”. State what the expression answers, and repair it.

Exercise 3.10 (*rapid*). For each sentence, name what is asked about and what is known, and write the expression.

(a) Unit A has failed twice this year; will it fail again before the year ends?

(b) How often are the feeder and unit A out at the same hour?

(c) Unit A has been found out; is the feeder out too?

(d) If unit A is taken out for maintenance tomorrow, will hall A keep its cooling during the work?

(e) How probable is a reading of 9 K above hall A in a trip test at 60 kW, if the conductance is 8000 W/K?

(f) Taking the opening as adequate, what share of the hours is the service delivered?

(g) How often are both units unavailable at once, for any reason?

(h) Starting with everything up, will the year be free of interruptions?

(i) Given a year free of interruptions, did the feeder never fail?

(j) The feeder failed twice last year; will next year be free of interruptions?

3.6 What to retain

- Every clause has one of three places: asked about, known, or carried in the model. What stands behind the bar is taken as true.
- A conditional probability is a share of the known event; a joint probability is a share of everything. They answer different questions.
- The two directions of a conditional can be 1 and 0.81 on one plant, on its stated rates with the opening taken as adequate. The inverted direction needs how probable each cause was, which the forward one does not.
- What is known is the reading, not the quantity, and a verdict is less than its readings.
- An assumption behind the bar is a condition, and it keeps its label.
- A candidate list for a diagnosis must partition what was seen, or its probabilities do not sum to one.
- A value set is not a value seen: switching unit B off is not finding it off.

Lesson 4

Defining Failure, Success, and Acceptable Performance

4.1 The engineering question

Three documents about the two-hall plant each contain one sentence that matters here. The design specification says that hall B shall not run more than 8 K above hall A. The tenant's contract says that the operator shall cool the halls during at least 99.5 % of the hours in each month, outages of the electrical feeder excepted. The owner's brief says that the plant shall be N+1. The contract is an illustration written for this lesson, since the plant's fact sheet holds no contract, but clauses of its kind are common.

None of the three sentences is a question, and each is silent about something a probability needs. The specification does not say with which unit out, at what load, or for how long. The contract does not say whether 99.5 % is a promise about each month's hours as they fall or about an average over many months, nor whether the excepted hours count as hours of cooling or are left out of the count. The brief does not say what failure the spare unit is meant to cover. None of the three says with what probability it must hold, because each is written as a "shall".

Lesson 1 began with an owner who asked whether the plant was reliable enough, and found that the question lacked four things: its event and conditions, its horizon, and its level. An owner who asks that question is asking whether the plant meets requirements like the three above, and the four missing things are the four parts of a requirement. The work of the lesson is to take a requirement apart into those parts, to say which of them the document supplies and where the others must come from, and to write the probability statement the requirement implies.

4.2 The formulation

The work of this lesson sits in step 4 of the method, the event with its thresholds and horizon, and in the level against which the expression of step 5 is read. The other steps run as before.

4.2.1 The four parts of a requirement

A requirement has four parts. The *event* is what should hold. The *conditions* are the circumstances under which it is asked: operating conditions, and any exclusions. The *horizon* is the time over which it is judged. The *level* is the probability its event must reach. A requirement's verdict therefore always has one shape,

$$P(\text{the event, over the horizon} \mid \text{the conditions, } E) \geq p_{\text{req}}, \tag{4.1}$$

in which the event is built from the words of the requirement as in Lesson 2, the conditions stand behind the bar as in Lesson 3, the horizon is stated as in Lesson 1, and the level p_{req} must be supplied by someone. Three later lessons take up the shape of equation (4.1): Lesson 5 writes its event as a limit state, Lesson 14 turns its verdict into a decision, and Lesson 17 gives it its name in optimization, a chance constraint.

A rule decides when a part counts as missing. The conditions count as stated when the requirement names the operating state in which its event is judged (which components are out, the load, and whether the plant has settled) or says that it holds in every state; its exclusions belong to the conditions, and so does the start of any horizon, which stands behind the bar. The event counts as stated when every threshold it uses is stated somewhere, in the requirement or in the plant's facts. A part supplied by the analyst rather than by the document is a choice, and the sheet labels it as one.

Table 4.1 takes the three requirements of the meeting apart, one row each, with the requirement as written in the first column and its four parts in the others. Where the document says nothing, the entry says so.

Table 4.1: Three requirements on the two-hall plant taken apart into their four parts.

Requirement as written	Event	Conditions	Horizon	Level
Hall B shall not run more than 8 K above hall A.	hall B's difference at most 8 K	not stated: which unit is out, what load, settled or not	not stated	not stated
Cooling during at least 99.5 % of the hours in each month, feeder outages excepted.	the month's served share at least 0.995	the exception is stated; its form is not	each month	not stated
The plant shall be N+1.	not stated: what failure the spare must cover	one unit out, presumably at design load	not stated	not stated

The last column is empty in every row, and that is typical. The second row also holds a number, 0.995, that looks like a level and is not one, for a reason the subsection on the performance level gives.

4.2.2 Failure at the level of a function

A unit that is out is a *component failure*. A hall that loses its cooling is a *functional failure*: the plant no longer performs a function it is required to perform, which is the standard definition of a failure applied to the function rather than to a part. Both halls losing their cooling at once is a total loss. Connectives build the hall events from the outages. Hall B loses its cooling when the feeder is out, or when unit B is out and either unit A is out or the opening cannot pass hall B's load to unit A,

$$\{\text{hall B loses its cooling}\} = F_{\text{fd}} \cup \left(F_B \cap \left(F_A \cup \{G < G^*\}\right)\right), \tag{4.2}$$

and hall A's event is the same with A and B exchanged. "Some hall" is the union of the two hall events. "Both halls" is their intersection, which reduces to $F_{\text{fd}} \cup (F_A \cap F_B)$. Lesson 1's service event $S(t) = 1$ means that no hall loses its cooling, so Lesson 1 defined failure at the level of the hall, and did so without calling it a choice.

The level at which failure is defined decides which unknowns matter. Table 4.2 gives the probability of each definition at a random hour, one row each, with the opening taken as adequate in the second column and as inadequate in the third. The values use the stated rates

Table 4.2: Two definitions of failure, at a random hour on the stated rates.

Definition of failure	Opening adequate [–]	Opening inadequate [–]
Some hall loses its cooling	0.012	0.104
Both halls lose their cooling	0.012	0.012

(each unit 2.5 failures per year in service with a 180 h repair, the feeder 3.69 with a 24 h repair), with the components independent and each unit repaired on its own.

Under the first definition the opening matters: an inadequate opening multiplies the probability of failure by about eight. Under the second it does not enter at all, since the event $F_{\text{fd}} \cup (F_A \cap F_B)$ contains no G. A requirement written with failure as total loss therefore cannot test the N+1 claim: a plant whose opening is inadequate meets it exactly as well as a plant whose opening is adequate.

4.2.3 States between success and failure

At one hour the plant is in one of four bands of outcomes, drawn in Figure 4.1 as a strip divided from left to right: full service with the spare available, full service on one unit with the redundancy spent, one hall lost, and both halls lost. The three brackets below the strip mark three definitions of failure, each a different set of bands.

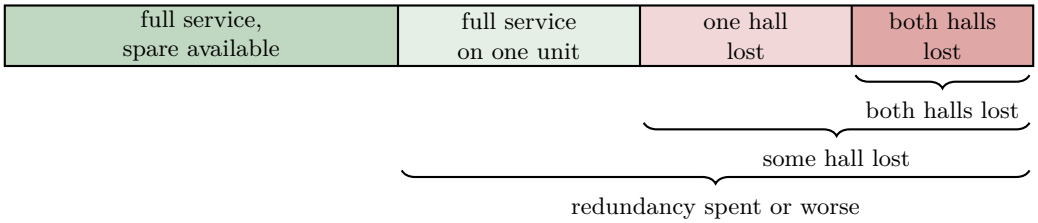


Figure 4.1: The outcomes at one hour in four bands, and three definitions of failure as sets of bands. The widths of the bands are illustrative.

Table 4.3 places the plant's states in the bands, with the state of the feeder and the units in the first column and the band on each side of the opening's threshold in the other two. It is Lesson 2's Table 2.1 with its last two columns refined from "served or not" into the four bands.

Table 4.3: The band of each state of the plant at one hour, on each side of the opening's threshold.

Feeder and units	Opening adequate	Opening inadequate
Feeder up, both units up	full service, spare available	full service, spare available
Feeder up, exactly one unit out	full service on one unit	one hall lost
Feeder up, both units out	both halls lost	both halls lost
Feeder out	both halls lost	both halls lost

The second row carries the table's point. With the opening adequate, a unit's outage moves the plant into the second band, where the service continues on the spare; with the opening inadequate, the same outage moves it straight into the third, and the second band is empty. The degraded state of full service on one unit exists only because the opening works, which is the N+1 claim seen from the side of the outcomes.

A requirement that names success as the first band and failure as the last two leaves the second band in neither set, so its success is no longer the complement of its failure. A state of

this kind, with performance or redundancy reduced short of failure, is a *degraded state*, and a requirement can be written about it directly; Exercise 4.2 is one. Each requirement must say where the middle bands go, and a translation that does not ask has made the choice silently.

4.2.4 Conditions and exclusions

Operating conditions, such as the design load, which unit is out and whether the plant has settled, go behind the bar, as Lesson 3 placed them. An exclusion, such as "outages of the feeder excepted", has two translations, and they answer different questions. In the first, the excepted hours count as served, which is the translation of "unless" as a union:

$$P\big(\{S(t) = 1\} \cup F_{\mathrm{fd}} \mid E\big), \tag{4.3}$$

the probability that at a random hour the operator is not in breach. In the second, the excepted hours are removed, and the exclusion stands behind the bar:

$$P\big(S(t) = 1 \mid F_{\mathrm{fd}}^{c},\ E\big), \tag{4.4}$$

the probability that the service is delivered at a random hour while the feeder is up. On the stated rates, with the opening taken as adequate, the two are 0.99764 and 0.99761, almost equal, and their meanings still differ: equation (4.3) is about the operator's obligation and equation (4.4) about the plant's performance when its supply is present. Their difference is the probability that the feeder is out times the probability of a loss while it is up, here 2.4×10^{-5}, and that holds whatever the dependence between the feeder and the units. The two nearly agree because both factors are small on this plant, and they part wherever either is large. Which translation a contract means is a question for its author, or for its definition of an hour of service. An exclusion can be an "unless" in the event or a condition behind the bar, and the translation must say which of the two it is.

4.2.5 A performance level inside the event

The contract's "at least 99.5 % of the hours in each month" is a statement about each month's hours as they fall. The share of the hours of month m in which the service is delivered is

$$\bar{S}(m) = \frac{1}{730\ \mathrm{h}} \int_{\text{month } m} S(t)\,\mathrm{d}t, \tag{4.5}$$

with 730 h a twelfth of a year. Equation (4.5) defines the *interval availability* of the month. It is not a probability but a random quantity: one value for each month as it happens, unknown in advance. The requirement's event is $\bar{S}(m) \geq 0.995$, which allows 3.65 h without service in a month. Its probability is not the availability. The availability of Lesson 1, the probability of service at a random hour, is the mean of $\bar{S}(m)$, while the requirement asks how often $\bar{S}(m)$ clears 0.995.

As written, equation (4.5) counts an hour of feeder outage as an hour without service. The contract excepts those hours, and the share it means counts them as served, which replaces the integrand by the larger of $S(t)$ and the indicator of a feeder outage,

$$\bar{S}_{\mathrm{ex}}(m) = \frac{1}{730\ \mathrm{h}} \int_{\text{month } m} \max\big(S(t),\ \mathbf{1}\{F_{\mathrm{fd}} \text{ at } t\}\big)\,\mathrm{d}t, \tag{4.6}$$

the interval availability with the exclusion kept in the event. The other translation of the exclusion would instead divide by the hours with the feeder up.

The clause therefore carries two levels of different kinds. The *performance level*, 0.995, sits inside the event and is set by the contract. The *probability level*, p_{req}, sits outside it, and the

contract does not set it: it belongs to whoever bears the cost of a breach. “In each month” over a contract year is “throughout” applied to months, the intersection of twelve such events.

Figure 4.2 shows why a mean and a share can disagree, with one bar per month for the hours lost in three months and a horizontal line at the allowance of 3.65 h. Two bars lie under the line and one lies far over it, and the dashed line marks the mean of the three. Because a single long loss lifts the mean, the mean can lie over the allowance while most months lie under it.

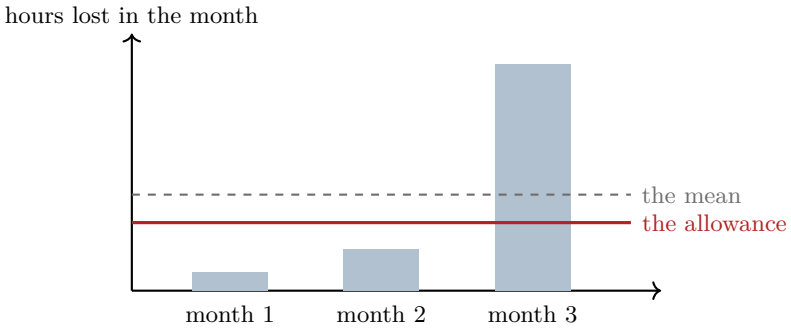


Figure 4.2: Hours lost in three months against the allowance of a monthly requirement, with their mean. The heights are illustrative.

4.2.6 The level

The level comes from someone, and the document that states the event seldom states it. Table 4.4 lists the four sources of requirements met in this lesson, one row each, with what each usually supplies: the event, a performance level inside it, and a probability level outside it.

Table 4.4: What each source of a requirement on the two-hall plant usually supplies.

Source	Event	Performance level	Probability level
Design specification	yes: hall B within 8 K of hall A	the 8 K itself	seldom: a “shall”
Contract	yes: the halls cooled	yes: 99.5 % of each month’s hours	seldom
Owner’s brief	rarely: “N+1” names a configuration	no	sometimes, as an appetite for risk
Insurer	yes: the covered loss	yes: more than 24 consecutive hours	no: the insurer prices the event

The table’s last column is the sparse one. A specification states an event and expects it to hold; a contract sets a performance level and leaves the probability of meeting it to the operator; the owner’s appetite for risk is the most common source of a probability level, and an insurer replaces the level by a price. A “shall” is not a requirement that the probability be one. Read that way, no plant could meet it, so the reading is useless. The event comes from the document and the level from elsewhere, or the requirement is read as a design condition, checked in the design model at its stated conditions. A level must be stated for the expression it is compared with: a level written for a share of hours says nothing about a year without interruption.

4.2.7 The stated limit as a requirement

The limit “hall B may run at most 8 K above hall A” is a stated physical fact of the plant. As a requirement its event is clear, and its conditions and horizon are not, which yields two readings. Table 4.5 runs the seven steps for both, each column read downward from the sentence to the problem class.

The expression of the design reading is the complement of the sixth row of Lesson 1’s Table 1.8. This lesson says nothing of how that row relates to the others; Lesson 9 takes up when two sentences name one quantity. The specification as written supports both readings, and the two need different models at step 6, as the two readings of “reliable enough” did. Which one the specification means is its author’s choice, and the analyst’s task is to bring both back and ask.

4.2.8 The contract clause

The contract clause can now be written in full. Table 4.6 is its filled formulation sheet, with the exclusion translated as in equation (4.3), the excepted hours counting as served, and the twelve months of a contract year.

The sheet’s sixth row shows the requirement reaching into the model. A requirement on the mean would have needed only the long-run fractions of time down, while a requirement on each month’s share needs the process that produces the losses.

4.2.9 The N+1 claim as a requirement

The owner’s brief is the third requirement of the meeting, and the shortest. “The plant shall be N+1” names a configuration, one spare unit, and a configuration is a fact about the drawing. Read as a requirement, it must name the failure the spare is meant to cover, and the four parts follow from that choice. The failure is a hall losing its cooling, equation (4.2), since the spare exists to keep both halls cooled. The conditions are one unit out, either one, with the other unit and the feeder up, the halls at their design load and the plant settled. The feeder stands behind the bar as up because the spare cannot cover it: an outage of the feeder takes both units, and a requirement that left it in the event would test the feeder as well as the spare. The horizon is none, since the requirement concerns a settled state. The level is the owner’s, and the brief does not state it. Table 4.7 is the filled formulation sheet.

Taken apart, the brief’s two words become the question Lesson 1 asked about the opening, equation (1.5). The route matters more than the destination. The equality in the sheet’s fifth row holds only because failure was defined at the hall, by the first row of Table 4.2, and only because the feeder was placed behind the bar. With failure defined as total loss, the opening would have dropped out and the claim could not have failed. With the feeder left in the event, the probability could not exceed 0.990, the share of time the feeder is up on the stated rates, so the claim would fail any level above that for a reason no second unit can change. The same two words support a requirement that tests the spare, one that tests nothing, and one that tests the feeder under the spare’s name.

Key idea

A requirement is complete when it states its event, its conditions, its horizon and its level. A document usually supplies two or three of them, and the probability statement it implies cannot be written until the rest are found or asked for.

Table 4.5: The seven steps on the stated limit, for its design reading and its operating reading.

Step	Design reading	Operating reading
Intention	Decide, resting on an estimate	Decide, resting on a prediction
1. Sentence	With unit B out and each hall at its 60 kW design load, once the plant has settled, is hall B within 8 K of hall A, with at least the required probability?	At every hour of the next year, whatever is out and whatever the load, is hall B within 8 K of hall A, with at least the required probability?
2. Target	A state of the plant under a stated condition	A path over a year
3. Known	Unit B out; 60 kW per hall; settled; E. The level p_{req} is not supplied.	All up at $t = 0$; E; the loads as they will be. The level is not supplied.
4. Event	$\Delta T_{\text{full}} \leq \Delta T_{\text{allow}}$; no horizon, the state being settled	hall B's difference at most 8 K for every t in $[0, \tau]$, $\tau = 1$ yr. It is not defined while the feeder or both units are out; here those hours count as failures
5. Expression	$P(\Delta T_{\text{full}} \leq \Delta T_{\text{allow}} \mid$ unit B out, Q_{hall} per hall, settled, $E)$, against p_{req}	P(hall B's difference $\leq 8\,\text{K}$ for every $t \in [0, \tau] \mid$ all up at $t = 0$, $E)$, against p_{req}
6. Model	The heat balance of each hall with one unit out, the opening's law with G, the loads, E	All of that at every load the year brings, with the failure and repair process
7. Read back	Whether the plant copes with the design case. Silent on how often it occurs.	Whether the limit holds through a year of operation. Silent on how far a breach goes.
Problem class	Posterior exceedance	Predictive probability over a horizon

Both readings as runs:

Run, the stated limit, design reading

1	**Sentence**	Unit B out, 60 kW per hall, settled: hall B within 8 K of hall A, at level p_{req}?
	Intention	Decide, on an estimate
2	**Target**	A state under a stated condition
	Declare	ΔT_{full} hall B above hall A, unit B out, settled [K]; $\Delta T_{\text{allow}} = 8\,\text{K}$; $Q_{\text{hall}} = 60\,\text{kW}$
3	**Known**	Unit B out; 60 kW per hall; settled; E; p_{req} not supplied
	Random here	ΔT_{full}, through G
4	**Event**	$\Delta T_{\text{full}} \leq \Delta T_{\text{allow}}$; no horizon
5	**Expression**	$P(\Delta T_{\text{full}} \leq \Delta T_{\text{allow}} \mid$ unit B out, Q_{hall} per hall, settled, $E)$, against p_{req}
6	**Model**	Heat balance with one unit out; the opening's law with G; the loads; E
7	**Check**	Copes with the design case; silent on how often it occurs
	Class	Posterior exceedance
		$\boxed{P(\Delta T_{\text{full}} \leq \Delta T_{\text{allow}} \mid \text{unit B out, } Q_{\text{hall}} \text{ per hall, settled, } E) \geq p_{\text{req}}}$

Run, the stated limit, operating reading

1	**Sentence**	Every hour of next year, whatever is out, any load: hall B within 8 K, at level p_{req}?
	Intention	Decide, on a prediction
2	**Target**	A path over a year
	Declare	Hall B's difference above hall A at t [K]; $\tau = 1$ yr
3	**Known**	All up at $t = 0$; E; the loads as they will be; p_{req} not supplied
	Random here	Which unit is out and when; the loads; G
4	**Event**	Difference at most 8 K for every t in $[0, \tau]$; feeder or both units out: failures
5	**Expression**	P(hall B's difference $\leq 8\,\text{K}$ for every $t \in [0, \tau] \mid$ all up at $t = 0$, $E)$
6	**Model**	The same balances at every load, with the failure and repair process
7	**Check**	The limit holds through a year; silent on how far a breach goes
	Class	Predictive probability over a horizon
		$\boxed{P(\text{hall B's difference} \leq 8\,\text{K for every } t \in [0, \tau] \mid \text{all up at } t = 0,\ E) \geq p_{\text{req}}}$

Table 4.6: The formulation sheet filled in for the contract clause, with feeder outages counted as served.

Field	Entry
Intention	Decide, resting on a prediction
1. Sentence	Over the coming contract year, starting from the plant in steady operation at $t = 0$, is every month's share of hours with the halls cooled, counting the hours of a feeder outage as cooled, at least 0.995, with at least the required probability?
2. Target	A future outcome: the twelve months' served shares
3. Known	The start: the plant in steady operation at $t = 0$, each component up or out with its long-run share of time, since the contract year begins wherever the plant happens to be; E: the stated rates, the structure, any record. The exclusion is in the event. The level p_{req} is not supplied by the contract.
4. Event	$\bar{S}_{\text{ex}}(m) \geq 0.995$ for every $m = 1, \ldots, 12$, with $\bar{S}_{\text{ex}}(m)$ from equation (4.6)
5. Expression	$P\big(\bar{S}_{\text{ex}}(m) \geq 0.995$ for every $m = 1, \ldots, 12 \mid$ steady operation at $t = 0,\ E\big)$, against p_{req}
6. Model	The failure and repair process over the year, with the distribution of each repair's length and not only its mean: the allowance of 3.65 h a month is far below a unit's mean repair of 180 h, so whether a month clears it depends mostly on whether a loss begins in it at all
7. Read back	The probability that no month of the year breaches the clause. Silent on the size of a breach and on the total hours lost.
Problem class	Predictive probability of an interval availability over a horizon

Tersely:

Run, the contract clause

1	**Sentence**	Contract year, steady at $t = 0$: every month's share ≥ 0.995, feeder hours served?
	Intention	Decide, on a prediction
2	**Target**	Future outcome: the twelve months' served shares
	Declare	$\bar{S}_{\text{ex}}(m)$ month m's served share, equation (4.6); $m = 1, \ldots, 12$; p_{req} not supplied
3	**Known**	Steady operation at $t = 0$; E rates, structure, any record
	Random here	The twelve shares; the exclusion is in the event
4	**Event**	$\bar{S}_{\text{ex}}(m) \geq 0.995$ for every $m = 1, \ldots, 12$
5	**Expression**	$P\big(\bar{S}_{\text{ex}}(m) \geq 0.995$ for every $m = 1, \ldots, 12 \mid$ steady operation at $t = 0,\ E\big)$
6	**Model**	Failure and repair process; each repair's length distribution, not only its mean
7	**Check**	No month breaches; silent on a breach's size and on the total hours lost
	Class	Predictive probability of an interval availability over a horizon

$$P\big(\bar{S}_{\text{ex}}(m) \geq 0.995 \text{ for every } m = 1, \ldots, 12 \mid \text{steady operation at } t = 0,\ E\big) \geq p_{\text{req}}$$

4.3 The competing formulations

The clause of the previous section has been read in several ways in practice, and the readings disagree in size. Table 4.8 gives five formulations of it, one row each, with the formulation in the first column, the treatment of the exclusion in the second and the value in the third. All five use the stated rates with the opening taken as adequate. The first two are closed forms at a random hour. The last three come from a simulation of the stated model started in steady operation; the two monthly values are quoted to two figures and the yearly one to one figure, since a year started with everything up instead differs from it in the third figure.

The one clause supports three verdicts. Read as the availability with the exclusion dropped, 0.988 against 0.995, the plant is in breach. With the exclusion kept, 0.998, it is not. Read in the contract's own words, a month meets the clause with probability 0.98, and a year in which no month is in breach has probability about 0.8, so about one contract year in five contains a breach month. The first comparison commits two errors at once: it drops the exclusion, and it

Table 4.7: The formulation sheet filled in for “The plant shall be N+1”.

Field	Entry
Intention	Decide, resting on an estimate
1. Sentence	With either unit out, the other unit and the feeder up, each hall at its 60 kW design load and the plant settled, does every hall keep its cooling, with at least the required probability?
2. Target	A state of the plant under a stated condition, and through it the conductance G
3. Known	One unit out; the other unit and the feeder up; 60 kW per hall; settled; E. The level p_{req} is not supplied.
4. Event	No hall loses its cooling: the complement of the union of the two hall events of equation (4.2). Under the conditions, by Table 4.3, the event holds exactly when $G \geq G^*$.
5. Expression	P(no hall loses its cooling | one unit out, the other unit and the feeder up, design load, settled, E), which equals $P(G \geq G^* \mid E)$, against p_{req}
6. Model	The heat balance of each hall, the opening’s law, the readings with their offsets and noise
7. Read back	Whether the spare can do its job once a unit is out. Silent on the feeder, which the claim cannot cover, and on how often a unit is out.
Problem class	Posterior exceedance

Its run:

Run, the plant shall be N+1

1	**Sentence**	Either unit out, other and feeder up, 60 kW, settled: every hall cooled, at p_{req}?
	Intention	Decide, on an estimate
2	**Target**	A state under a stated condition; through it G
	Declare	Hall events, equation (4.2); G [W/K]; $G^* = 7500\,\text{W/K}$; p_{req} not supplied
3	**Known**	One unit out; the other unit and the feeder up; 60 kW per hall; settled; E
	Random here	G; the rest behind the bar
4	**Event**	No hall loses its cooling; here, exactly $G \geq G^*$
5	**Expression**	P(no hall loses its cooling | the conditions of step 3)
6	**Model**	Heat balance per hall; the opening’s law; readings with offsets and noise
	Reduce	$P(G \geq G^* \mid E)$, by Table 4.3
7	**Check**	The spare can do its job; silent on the feeder and on how often a unit is out
	Class	Posterior exceedance

$$\boxed{P(G \geq G^* \mid E) \geq p_{\text{req}}}$$

Table 4.8: Five formulations of one contract clause on the stated rates, with the opening taken as adequate. The clause is illustrative.

Formulation	Exclusion	Value [–]
Availability, read against 0.995	dropped	0.988
Availability, read against 0.995	kept, as served	0.998
A given month meets 0.995	dropped	0.75
A given month meets 0.995	kept, as served	0.98
Every month of a year meets 0.995	kept, as served	0.8

compares a mean with a performance level that was set for each month’s share as it falls. The verdict on the contract also needs a probability level, which the contract does not state.

Marked and repaired:

Run, the clause read as the availability, exclusion dropped

1	**Sentence**	Every month's share at least 0.995, feeder hours counted as cooled?
4	**Event** ×	$S(t) = 1$ at a random hour, feeder outages as losses
	Repair	$\bar{S}_{\text{ex}}(m) \geq 0.995$ for every $m = 1, \ldots, 12$, feeder outages as served
5	**Expression** ×	$P(S(t) = 1 \mid \text{rates}, G \geq G^*)$, t random, against 0.995: 0.988, a breach
	Repair	$P\big(\bar{S}_{\text{ex}}(m) \geq 0.995 \text{ for every } m \mid \text{steady operation at } t = 0, \text{ rates}, G \geq G^*\big)$: ≈ 0.8
7	**Check**	The clause is on each month's share, exclusion kept; a mean is not a month

$$P\big(\bar{S}_{\text{ex}}(m) \geq 0.995 \text{ for every } m = 1, \ldots, 12 \mid \text{steady operation at } t = 0,\ E\big) \geq p_{\text{req}}$$

A second competing formulation chooses the definition of failure badly. An assessment of the N+1 claim that defines failure as total loss obtains 0.012 whether the opening is adequate or not, as Table 4.2 showed, and so it cannot fail. The repair defines failure as a hall losing its cooling. A third reads a "shall" as a probability level of one, so that any expression short of one fails; the repair says that the level is not stated and asks for it. A fourth places an exclusion without saying where. On this plant equations (4.3) and (4.4) nearly agree, but only because the feeder is seldom out and a loss while it is up is rare, and a translation that relies on the agreement has made an assumption about the plant that another plant need not share.

These errors are this lesson's failure class, the *requirement without its conditions or its level*: a requirement translated with one of its parts missing or misplaced. Its forms are the operating conditions or the exclusions dropped; an exclusion placed in the event or behind the bar without saying which; failure defined at the wrong level of function, a component or the whole plant where the requirement concerns a hall; a performance level read as a probability level, or a mean compared with a level set for each period's share; and a verdict drawn against a level nobody stated, or against a level stated for a different expression. A missing horizon remains Lesson 1's class. This one enters at step 1 or step 4 and is caught at step 7, by naming for each of the four parts the document that supplied it.

4.4 One system, many questions

Table 4.9 takes six requirements on the two-hall plant apart in the same four columns as Table 4.1, now with the parts filled in where a reading has been chosen.

Table 4.9: Six requirements on the two-hall plant and their four parts.

Requirement	Event	Conditions	Horizon	Level
Hall B within 8 K of hall A, design reading	$\Delta T_{\text{full}} \leq \Delta T_{\text{allow}}$	unit B out, 60 kW per hall, settled	none	not stated
N+1	no hall loses its cooling	one unit out, the other unit and the feeder up, design load	none	the owner's
The contract clause	$\bar{S}_{\text{ex}}(m) \geq 0.995$	steady operation at $t = 0$	each month of the year	not stated
Ride through a unit trip	the service held until the repair ends	a unit out at t_0, the rest up	the random repair time	the owner's
At most one interruption a year	at most one interruption of S	all up at $t = 0$	one year	not stated: the insurer prices it
Hall B within 8 K in the trip test	the test's reading of the difference at most 8 K	a test at load Q	the test	none: a rule

Every requirement fits the same four columns, and almost every one leaves the last column to someone other than its author. Two rows deserve a note. The ride-through row is Lesson 1's Exercise 1.3 restated as a requirement. The last row is not about the plant at all: its event concerns a reading of a test, which differs from the plant's state by the sensors' offsets and noise. It is an acceptance criterion, and Lesson 14 writes the rule that applies it.

4.5 The exercises

The solutions are in Appendix A, and the facts of the plant in Appendix D.

Exercise 4.1 (*formulate*). An insurer's clause reads: "Cover applies if any hall is without cooling for more than 24 consecutive hours in a policy year, other than during planned maintenance." Take the clause apart into its four parts, write its event, and write the expression the insurer needs.

Exercise 4.2 (*formulate*). The operations manager says: "The plant is acceptable if it runs on one unit for no more than 400 hours a year and neither hall is ever without cooling for more than 4 hours at a time." Take the statement apart, and write the expression for each requirement it holds.

Exercise 4.3 (*formulate*). A prospective tenant says: "The cooling must never fail." Formulate the requirement.

Exercise 4.4 (*formulate*). The commissioning plan says: "The plant passes commissioning if a trip test at 60 kW settles hall B no more than 8 K above hall A." State what the event of this requirement is about, and how it differs from the design reading of the specification in Table 4.5.

Exercise 4.5 (*reverse*). Write one plain sentence for

$$P(\bar{S}_{\mathrm{ex}}(m) \geq 0.995 \mid \text{steady operation at the start of month } m,\ E),$$

and state which threshold in it came from the contract.

Exercise 4.6 (*reverse*). Write one plain sentence for

$$P(\Delta T_{\mathrm{full}} \leq \Delta T_{\mathrm{allow}} \mid \text{unit B out},\ Q_{\mathrm{hall}} \text{ per hall, settled},\ E) \geq p_{\mathrm{req}},$$

and name the document that supplies each part.

Exercise 4.7 (*reverse*). Write one plain sentence for $P(F_{\mathrm{fd}} \cup (F_A \cap F_B)$ at some $t \in [0, \tau] \mid$ all up at $t = 0,\ E)$, with $\tau = 1$ yr, and state which definition of failure it uses.

Exercise 4.8 (*critique and repair*). An auditor writes: "The plant's availability is 0.988, below the contract's 99.5 %, so the operator is in breach." State what the comparison answers, and repair it.

Exercise 4.9 (*critique and repair*). An assessment of the N+1 claim defines failure as the loss of all cooling, finds a probability of 0.012 at a random hour, and concludes that the opening is adequate. State what the assessment answers, and repair it.

Exercise 4.10 (*critique and repair*). A reviewer writes: "The specification says shall, so the required probability is one. The probability that hall B stays within 8 K is below one, so the plant fails the specification." State the error, and repair it.

Exercise 4.11 (*rapid*). For each requirement, name the parts it does not state, by the rule of this lesson: event, conditions, horizon, level.

(a) Hall B shall stay within 8 K of hall A.

(b) The service shall be available 99.9 %.

(c) No more than two interruptions a year.

(d) With unit A out at design load, hall A shall stay within its limit once settled.

(e) The plant shall survive an outage of the feeder.

(f) The cooling shall ride through a unit trip.

(g) Unit A shall not fail more than twice a year.

(h) At least 99 % confidence that the opening is adequate.

(i) The halls shall be kept cool except during planned maintenance.

(j) At most 4 hours without cooling in any month, with 95 % probability.

4.6 What to retain

- A requirement has four parts: its event, its conditions, its horizon and its level. A document usually supplies two or three of them.
- Failure is defined at a level: a component, a hall, or the plant. The choice decides which unknowns matter, and on this plant a definition by total loss makes the opening irrelevant.
- When a requirement names a state between success and failure, success is no longer the complement of failure.
- An exclusion goes in the event or behind the bar, and the two translations answer different questions.
- A performance level inside the event is not a probability level. On this plant, with feeder outages excepted, the availability is 0.998 while about 2 % of months breach a monthly clause of 99.5 %; with the exception dropped, the availability is 0.988 while about 75 % of months meet the clause.
- A “shall” states an event, not a probability of one.

Part II

Translating Questions into Probability

Lesson 5

Threshold and Exceedance Questions

5.1 The engineering question

Both halls of the two-hall plant are running at about 55 kW on a summer night when cooling unit B trips. The repair crew cannot come before morning, and the operator asks whether hall B will stay within its limit until then. The commissioning engineer, reached by telephone, answers from the report of the trip test run a month earlier. At the best estimate of the opening that the test gave, a load of 55 kW would settle hall B about one and a half kelvin below its limit, so there is margin to spare. A colleague on the same call asks a different thing: given what the trip test showed, how probable is it that hall B goes over the limit tonight?

The two answers sound like one answer expressed in two units, and they are not. The first is a distance from a boundary, measured at one chosen value of a quantity nobody knows exactly. The second is a probability that the plant is on the wrong side of the boundary, and it depends on how much is still unknown about that quantity as well as on where its best estimate lies. On this plant and this test the two disagree sharply: the margin sounds comfortable, while the probability is not small. The size of the disagreement is given in the competing formulations below.

Most engineering limits have the same form. They compare something the plant does, such as a temperature rise, the duty of a unit or the heat carried through an opening, with something it must not exceed, and the thing the plant does depends both on quantities that are not known and on conditions that change from night to night. A limit of this kind can be written once, as a function of those quantities and conditions that is positive on the safe side. Three different questions can then be asked of it. One asks how far from the boundary the plant sits at a chosen point. One asks how probable it is that the plant is on the wrong side. One asks under which conditions the plant stays on the right side at all.

This lesson writes the limit as a function, says which of its arguments are unknown and which are conditions of the question, and then asks the question that was asked: a margin, a probability or a set of conditions.

5.2 The formulation

The work of this lesson sits in step 4 of the method, where the event is now written as a function of the model's unknowns, and in step 6, because a limit cannot be written as a function without stating what the limited quantity depends on. The other steps run as in Part I.

The lessons of Part II use one body of evidence for the two-hall plant, so that the numbers of one lesson can be set against those of another. The evidence E is a single trip test at 40 kW per hall, with unit B taken out and the plant read after it settled, together with the drawing's prior on the opening's conductance and the stated spreads of the sensors in Appendix D. Its

measured difference of hall B above hall A was 4.78 K. The test comes from the plant's case study, where it is one registered draw on a plant whose candidate drawings lie close to the threshold. Every probability about the opening in this part of the book is read from one model: the case study's model of that test, with the drawing's prior on the conductance and the sensors' offsets and noise, kept to the conductance alone. On it, after the test, the probability that the opening is adequate, $P(G \geq G^{*} \mid E)$, is 0.82. The text calls this evidence "E, the 40 kW trip test" wherever it stands behind a bar.

5.2.1 The limit-state function

A *limit state* is the boundary between satisfying a requirement and failing it. It is written as a function of the state and of the operating conditions, positive on the safe side and negative where the requirement fails, and the set of values on which it is negative is the requirement's *failure domain* (Ditlevsen and Madsen, *Structural Reliability Methods*, 1996). For hall B the limit is the stated one: hall B may not run more than 8 K above hall A. With unit B out and the plant settled, all of hall B's heat must cross the opening to unit A, and by the opening's law the heat it carries is its conductance times the temperature difference across it. Hall B's settled rise above hall A is therefore its load divided by the conductance. The limit state is the allowable difference minus that rise,

$$g(G;\ Q_B) = \Delta T_{\text{allow}} - \frac{Q_B}{G}, \tag{5.1}$$

in kelvin. Its first argument, the conductance G, is uncertain and is the quantity the trip test informs. Its second, the load Q_B that hall B carries at the time asked about, is a condition of the question: the operator tonight knows it, a planner a year ahead does not. The semicolon in equation (5.1) separates the unknown from the condition. The function is positive when hall B settles inside its limit, zero on the boundary, and negative when the opening cannot pass hall B's load within 8 K.

Writing equation (5.1) has already done part of step 6. The function could not be written without stating that the rise is set by the opening's law, that unit B is out, that all of hall B's heat goes through the opening, and that the plant has settled. Each of those statements is now visible in the formulation, where a plain inequality on "hall B's difference", as Lesson 4 wrote it in Table 4.5, kept them out of sight.

The same boundary has more than one function. The expression $G\,\Delta T_{\text{allow}} - Q_B$, in kilowatts, is positive exactly where equation (5.1) is positive, because the conductance is positive, so the two define the same failure domain. The choice of function changes the units and the size of a margin, defined below, and it never changes the event. A limit state is a description of a set; any function that is negative on that set and only there describes it.

5.2.2 Three answers from one limit state

Three different questions can be asked of one limit state, and Table 5.1 lists them, one row each, with the expression in the second column and the kind of object it is in the third. The point value $\hat{G}$ in the first row is the conductance at the most probable state of the case study's model of the test, 8428 W/K for the 40 kW trip test, as the case study reports. It lies near the middle of what the test leaves known about the conductance, though it is not exactly the most probable conductance taken alone.

The *safety margin* is the value of the limit-state function at a stated point. It is a distance from the boundary measured in the function's units, and it is a property of the point chosen as well as of the plant. The *exceedance probability*, often called the probability of failure, is the probability that the unknowns lie in the failure domain given what is known. It depends on where the boundary lies relative to the whole spread of what is known, and not only relative to

Table 5.1: Three questions asked of one limit state, with unit B out and the plant settled.

Question	Expression	Kind of object
The margin at a point	$g(\hat{G};\ Q_B)$	A number in the units of g, here kelvin
The exceedance probability	$P\big(g(G;\ Q_B) < 0 \mid$ unit B out, Q_B, settled, $E\big)$	A probability between zero and one
The envelope at level p	the set of Q_B with $P(g(G;\ Q_B) < 0 \mid$ unit B out, Q_B, settled, $E) \le p$	A set of operating conditions, here a range of loads

its most probable value. The *operating envelope* is the set of operating conditions in which the limit state holds; since its boundary depends on unknowns, the useful envelope is the set of conditions in which the exceedance probability stays below a stated level p.

Figure 5.1 shows the first two answers on one picture. The horizontal axis is hall B's settled rise above hall A at a stated load, and the curve is the weight that what is known of the conductance puts on each rise. The vertical line is the limit of 8 K, and the shaded area beyond it is the exceedance probability. The marker is the rise at $\hat{G}$, and the double arrow from it to the limit is the margin. The two are read from different features of the picture: the margin from one point on the axis, the probability from an area that depends on the whole curve.

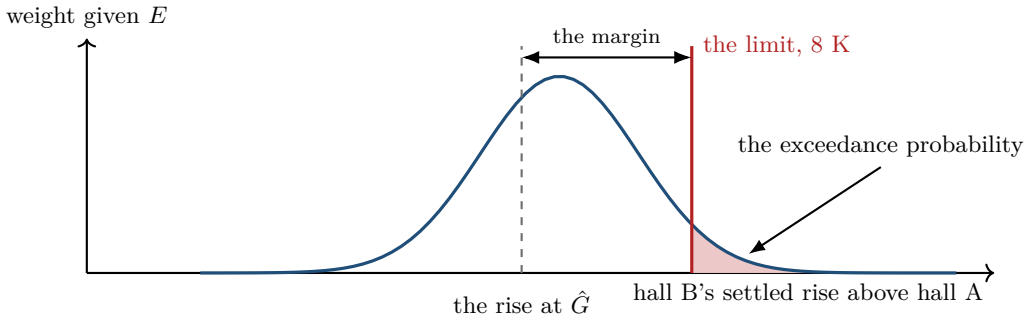


Figure 5.1: One limit state and two of its answers, at one load with unit B out. The margin is read at the point value; the exceedance probability is the shaded weight beyond the limit. The curve's shape is illustrative.

The picture makes plain why the margin is not a probability, and why it should never be compared with a probability level such as p_{req}. A margin is a number in kelvin, or in kilowatts if the other function is chosen, and it can be large while the shaded area is not small, when the curve is wide, or small while the shaded area is tiny, when the curve is narrow. Lesson 1's rule that a capacity margin is never equated with an availability gains a companion here: a safety margin is never equated with an exceedance probability either.

5.2.3 Several limits on one requirement

A requirement can carry more than one limit. With unit B out, unit A must carry the heat of both halls within its capacity of 130 kW, and hall B's heat must cross the opening within the stated difference. The first limit is a second limit state,

$$g_{\text{cap}}(Q_A,\ Q_B) = 130\,\text{kW} - (Q_A + Q_B), \tag{5.2}$$

which on the stated values involves no unknown at all. The requirement fails if either limit fails, so its failure domain is the union of the two, $\{g_{\text{cap}} < 0\} \cup \{g < 0\}$. The union is written on

the two failure domains and not on one combined function, because the two functions are in different units. A requirement whose failure domain is the union of its limit states' domains is a series system of limit states, and the word "or", translated as a union, now applies to regions in the space of the unknowns and conditions rather than to states of components.

At the design load of 60 kW in each hall, equation (5.2) gives a margin of 10 kW, certain on the stated values. The two limit states are in different units, kilowatts and kelvin, so the requirement has no single margin. It has one margin per limit state, and one probability, the probability of the union.

The plant has more limits than these two, and a formulation must say which of them a question concerns. Table 5.2 lists four, one row each, with the limit-state function in the second column, its uncertain arguments in the third and its conditions in the fourth. The last column says whether the plant's stated facts supply everything the function needs.

Table 5.2: Four limits of the two-hall plant written as limit-state functions. Each function carries the units of its threshold.

Limit	Limit-state function	Uncertain	Conditions	Stated in full
Hall B above hall A, unit B out	$\Delta T_{\text{allow}} - Q_B/G$, in K	G	unit B out, Q_B, settled	yes
Hall A above hall B, unit A out	$\Delta T'_{\text{allow}} - Q_A/G$, in K	G	unit A out, Q_A, settled	no: the threshold is not stated
The surviving unit's capacity	$130\,\text{kW} - (Q_A + Q_B)$, in kW	none	one unit out, both loads	yes
Hall B's own temperature, unit B out	$30^\circ\text{C} - \vartheta_A - Q_B/G$, in K	G, and ϑ_A unless held	unit B out, Q_B, settled	no: the 30°C is not stated (illustrative)

The second row is a limit the plant's documents do not state: the fact sheet gives the allowable difference of hall B above hall A and nothing for the reverse, so a question about unit A tripping cannot be written until someone supplies it. The fourth row shows a limit that becomes the first only under a condition. A limit of 30°C on hall B, a value no document of the plant states and which is used here only as an illustration, equals the stated 8 K difference when hall A sits at its set point of 22°C, and with unit B out hall A is held there only if unit A carries both halls, which is the third row. A limit on an absolute temperature therefore brings the capacity limit with it. Thresholds can be uncertain too: the capacity of 130 kW is stated here, and on a plant where it is not known, the third row's function would gain an unknown of its own.

5.2.4 Conditions as arguments

Lesson 4 placed a requirement's operating conditions behind the bar. A limit state shows why they belong there. The load Q_B, which unit is out, and whether the plant has settled are arguments of g, and a condition that is known fixes an argument: the expression then concerns the plant at that value. A condition that is not known, such as the load hall B will carry at some moment still to come, cannot be fixed behind the bar. It stays among the unknowns, and the event becomes one over both the conductance and the load. This lesson treats only questions in which the load is known; Lesson 9 takes up the case in which it is not.

The same reasoning separates two kinds of number that both look like conditions. The allowable difference of 8 K is a stated threshold, part of the requirement. The conductance at which hall B would just reach that threshold at a stated load, $Q_B/\Delta T_{\text{allow}}$, is a consequence of the threshold and the load; it is the point at which the limit state crosses zero, and it moves

when the load moves.

5.2.5 The envelope as a set

Figure 5.2 draws the operating envelope with unit B out in the plane of the two halls' loads, hall A's on the horizontal axis and hall B's on the vertical. The sloping line is the capacity limit of equation (5.2), fixed and certain. The three horizontal lines are the opening's limit of equation (5.1), which depends only on hall B's load, drawn three times: at the point value $\hat{G}$, and at the loads whose exceedance probability given E is 0.05 and 0.01. The envelope at each level is the region below both its horizontal line and the sloping one.

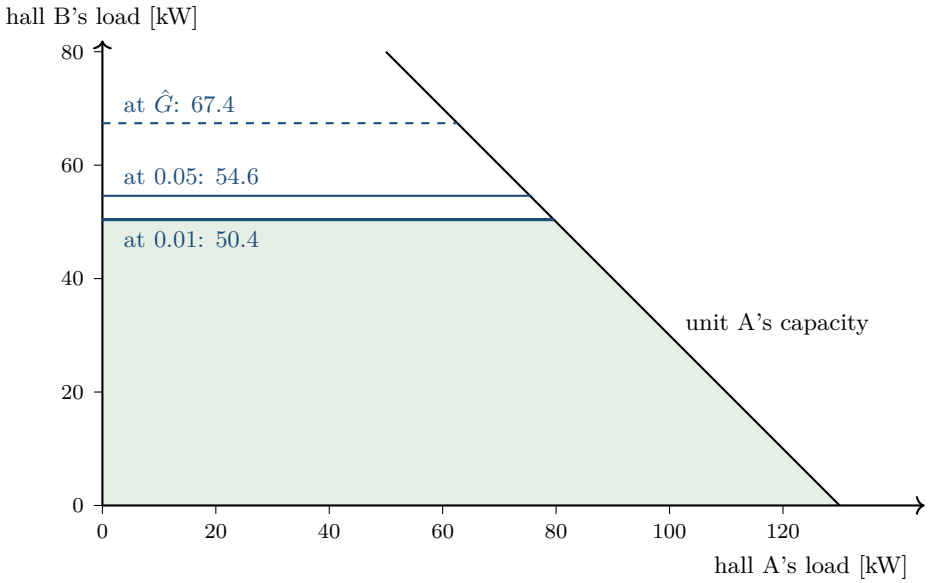


Figure 5.2: The operating envelope with unit B out: the capacity limit and the opening's limit drawn at the point value and at two levels of exceedance probability given E, the 40 kW trip test. The shaded region is the envelope at a level of 0.01.

Read from the figure, the envelope drawn at the point value reaches 67.4 kW for hall B, the envelope held at an exceedance probability of 0.05 reaches 54.6 kW, and the envelope at 0.01 reaches 50.4 kW. An envelope drawn at a point value describes a plant whose conductance is known to equal that value. On the actual plant the load at which the margin at $\hat{G}$ reaches zero is a load at which the exceedance probability is close to one half, since $\hat{G}$ sits near the middle of what is known.

5.2.6 Tonight's load with unit B out

The night shift's question can now be written. Table 5.3 is its formulation sheet, with steps 4 and 6 in full and the others condensed.

The expression carries no horizon, and the reason is in the sheet's third and sixth rows: the state is taken as settled and the conditions as held fixed until morning. Its seventh row names what that assumption hides. A second failure overnight would turn the question into one about a path, which needs the dependence of Lesson 6 and the forms of time of Lesson 8. The sheet's first row also deserves a remark. The intention is to predict because the operator asks about tonight's outcome; the same expression asked by an engineer out of interest in the opening would be an estimate. The expression is the same either way, and the intention decides what is done with it.

Table 5.3: The formulation sheet filled in for "Will hall B stay within its limit tonight with unit B out?".

Field	Entry
Intention	Predict, since tonight's outcome is still to occur; it rests on what is known of G
1. Sentence	With unit B out from now until the repair and both halls at 55 kW, once the plant has settled, what is the probability, given the 40 kW trip test, that hall B runs more than 8 K above hall A?
2. Target	An event on a settled state under a stated condition, and through it the parameter G
3. Known	Unit B out; $Q_A = Q_B = 55\,\mathrm{kW}$, behind the bar; E, the 40 kW trip test, with the drawing's prior on G and the sensors' stated spreads. Assumption: the loads hold through the night.
4. Event	$g(G;\ Q_B) < 0$ from equation (5.1), with $Q_B = 55\,\mathrm{kW}$. The capacity limit gives $g_{\mathrm{cap}}(55, 55) = 20\,\mathrm{kW} > 0$ on the stated values, so the union reduces to the opening's failure domain.
5. Expression	$P\big(\Delta T_{\mathrm{allow}} - Q_B/G < 0 \mid \text{unit B out},\ Q_B = 55\,\mathrm{kW},\ \text{settled},\ E\big)$
6. Model	The opening's law and the heat balance with unit B out; what is known of G after the test, which needs how the test's readings behave at each G (the sensors' offsets and noise) and the drawing's prior; the capacity limit, stated. No failure process: the night is taken as one settled state with unit B out throughout.
7. Read back	The probability that hall B is over its limit tonight once the plant has settled, if the loads stay at 55 kW. Silent on unit A or the feeder failing during the night, and on the minutes before the plant settles.
Problem class	Posterior exceedance

As a run:

Run, will hall B stay within its limit tonight with unit B out

1	**Sentence**	Unit B out, $Q_A = Q_B = 55\,\mathrm{kW}$, settled, given E: hall B over 8 K above hall A?
	Intention	Predict; rests on G
2	**Target**	An event on a settled state; through it G
	Declare	$g(G;\ Q_B)$ [K], equation (5.1); g_{cap} capacity limit [kW]; E the 40 kW trip test
3	**Known**	Unit B out; $Q_A = Q_B = 55\,\mathrm{kW}$; settled; E; loads held overnight (assumption)
	Random here	G only
4	**Event**	$g(G;\ 55\,\mathrm{kW}) < 0$; $g_{\mathrm{cap}}(55, 55) = 20\,\mathrm{kW} > 0$, so the opening's domain alone
5	**Expression**	$P\big(g(G;\ Q_B) < 0 \mid \text{unit B out},\ Q_B = 55\,\mathrm{kW},\ \text{settled},\ E\big)$
	Expand	$P\big(\Delta T_{\mathrm{allow}} - Q_B/G < 0 \mid \text{unit B out},\ Q_B = 55\,\mathrm{kW},\ \text{settled},\ E\big)$
6	**Model**	The opening's law; heat balance, unit B out; G after the test; no failure process
	Reduce	$P(G < 55\,\mathrm{kW}/8\,\mathrm{K} \mid E) = P(G < 6875\,\mathrm{W/K} \mid E)$
7	**Check**	Tonight, once settled; silent on a second failure
	Class	Posterior exceedance

$$\boxed{P(G < 6875\,\mathrm{W/K} \mid E)}$$

5.2.7 The operating envelope as a question

The operations manager asks a question of a different kind: up to what load can hall B be run if unit B is out? The intention is to decide, since the answer sets an operating limit, and the target is not an event but a set of conditions, the loads at which hall B may be run. In the terms of step 2 it is a decision expressed as a set. For each candidate load the event and the expression are those of Table 5.3 with Q_B in place of 55 kW, and the envelope is

$$\{Q_B : P\big(g(G;\ Q_B) < 0 \mid \text{unit B out},\ Q_B,\ \text{settled},\ E\big) \le p\}, \tag{5.3}$$

with the level p not supplied by the manager. Table 5.4 is the formulation sheet, with the rows that differ from the night shift's sheet written out.

Table 5.4: The formulation sheet filled in for "Up to what load can hall B be run with unit B out?".

Field	Entry
Intention	Decide: the answer sets an operating limit
1. Sentence	With unit B out and the plant settled, for which loads of hall B is the probability of exceeding hall B's limit, given the 40 kW trip test, at most a stated level p?
2. Target	A decision expressed as a set of operating conditions, the loads of hall B, and its upper end
3. Known	Unit B out; settled; E, the 40 kW trip test. The level p is not supplied by the manager. Hall A's load enters only through the capacity limit, which the stated values satisfy at every load a hall can carry.
4. Event	For each candidate Q_B, the failure domain $g(G;\ Q_B) < 0$
5. Expression	The set of equation (5.3)
6. Model	As in Table 5.3, evaluated at every candidate load
7. Read back	A range of loads valid at the chosen level, for a settled state with unit A healthy. Silent on how long the plant may be run there
Problem class	Posterior exceedance, as a function of an operating condition

Compactly:

Run, up to what load can hall B be run with unit B out

1	**Sentence**	Unit B out, settled, given E: for which Q_B is exceeding the limit at most p?
	Intention	Decide: sets an operating limit
2	**Target**	A decision as a set of loads of hall B, and its upper end
	Declare	Q_B candidate load [kW]; $g(G;\ Q_B)$ as above; p the level, not supplied
3	**Known**	Unit B out; settled; E; hall A's load only through the capacity limit
	Random here	G; each candidate Q_B is set
4	**Event**	For each Q_B: $g(G;\ Q_B) < 0$
5	**Expression**	The set of equation (5.3)
6	**Model**	As in Table 5.3, at every candidate load
7	**Check**	Loads valid at level p, settled, unit A healthy; silent on how long
	Class	Posterior exceedance, as a function of an operating condition
		$\boxed{\{Q_B : P\big(g(G;\ Q_B) < 0 \mid \text{unit B out},\ Q_B,\ \text{settled},\ E\big) \le p\}}$

Read back, equation (5.3) is a range of loads valid at the chosen level and for a settled state. The manager's sentence gave no level, and the analyst must ask for one; the envelope at the point value answers a different question, about a plant whose conductance is known. Choosing the largest load in the set is an extreme of a function of one condition, which Lesson 17 names as an optimization.

5.2.8 The margin's own question

None of this makes the margin a wrong quantity. It is the right answer to a question of its own kind, and engineering practice asks that question often. A design check reads a limit state at stated values chosen for the purpose, such as a conservative conductance taken from the drawing's tolerance and the design load, and asks whether the function is positive there. Lesson 4 called this reading a requirement as a design condition, checked in the design model at its stated conditions. Its verdict is the sign of the margin at those values, and no probability is asked, because the conservatism was put into the values chosen rather than into a level.

The two readings differ in what they hold fixed. The design check fixes the arguments at stated values and reads a number. The exceedance question leaves the uncertain arguments uncertain, weighs them by what is known, and reads a probability. A margin evaluated at the point value $\hat{G}$ belongs to neither. It is not a design check, because $\hat{G}$ is not a conservative value

chosen for the purpose; on this test it lies near the middle of what is known. And it is not an exceedance probability. The failure class below is this hybrid: a margin at the point value offered where a probability was asked.

5.2.9 A limit question that is not ready

At the handover meeting the owner asks whether the plant is within its limits. The sentence names no limit, no condition and no time, and it asks for a verdict without a level, so it fails at the first step. Four readings are natural, and each gives a different expression.

In the most literal reading, the plant now, with both units running, the opening carries almost no heat and none of the limits of Table 5.2 is active: each of them concerns a plant with one unit out. The question about the present state is an estimate, and its honest answer is that the limits are not the plant's present concern, since they speak of a condition the plant is not in. A second reading is the design case of Lesson 4's Table 4.5, one unit out at the design load, with the surviving unit's capacity added as a union; it needs the probability of the union of the first and third rows at that load, given E, against a level the owner supplies. A third reading is tonight's condition, the expression of Table 5.3 at tonight's loads. A fourth reading asks whether every limit holds at every hour of the coming year, whatever is out and whatever the load. That is a question about a path, since which unit is out and what the load is change over the year, and its expression needs the failure and repair process that Lesson 8 writes.

The analyst's reply in the meeting is short. It names the four readings, asks which condition the owner has in mind, and asks for the level. It also says that a margin will not answer any of them: the owner asked for a verdict, and a verdict on a limit with an uncertain argument is a probability compared with a level.

Key idea

One limit state answers three different questions. Its value at a point is a margin, in the limit's units; the weight of its failure domain is an exceedance probability; the conditions in which that weight stays below a level form an envelope. Only the probability is compared with a probability level.

5.3 The competing formulations

The most common competing formulation reads the margin as the verdict. Table 5.5 sets the two answers side by side at two night loads, one row each, with unit B out and the plant settled and with E the 40 kW trip test. The second column is the margin at $\hat{G} = 8428$ W/K and the third the exceedance probability. The values come from a small calculation with the case study's stated model of the test.

Table 5.5: The margin at the point value against the exceedance probability, at two loads of hall B, with unit B out, the plant settled and E the 40 kW trip test.

Hall B's load [kW]	Margin at $\hat{G}$ [K]	$P(g < 0 \mid \text{unit B out},\ Q_B,\ \text{settled},\ E)$ [–]
50	2.07	0.008
55	1.47	0.056

Between the two rows the margin falls by less than a third, while the exceedance probability rises sevenfold. At 55 kW, the night shift's load, a margin of about one and a half kelvin sits beside a probability of exceeding the limit of about one in eighteen. The margin cannot show this, because it is read at one point and says nothing about how much of what is known lies

across the boundary; the probability depends on the margin and on the spread of the unknowns together. Two plants with the same margin can therefore have very different exceedance probabilities, and a plant with the larger margin can have the larger probability. This is the failure class of this lesson, the *point value read against the limit*.

With the wrong lines marked:

Run, the margin at the point value read as the verdict

1	**Sentence**	Unit B out, both halls at 55 kW, settled: hall B within its limit tonight?
4	**Event** ×	The sign of $g(\hat{G};\ 55\,\mathrm{kW})$, with $\hat{G} = 8428\,\mathrm{W/K}$
	Repair	$g(G;\ 55\,\mathrm{kW}) < 0$, with G uncertain
5	**Expression** ×	$g(\hat{G};\ 55\,\mathrm{kW}) = 1.47\,\mathrm{K} > 0$: within the limit
	Repair	$P\big(g(G;\ Q_B) < 0 \mid \text{unit B out},\ Q_B = 55\,\mathrm{kW},\ \text{settled},\ E\big)$: 0.056
7	**Check**	A number in kelvin is a design check; the sentence asks how probable
		$\boxed{P\big(g(G;\ Q_B) < 0 \mid \text{unit B out},\ Q_B = 55\,\mathrm{kW},\ \text{settled},\ E\big)}$

The reason the probability moves so much faster than the margin lies in the event itself. At 50 kW hall B exceeds its limit only if the conductance is below $50\,\mathrm{kW}/8\,\mathrm{K} = 6250$ W/K, and at 55 kW only if it is below 6875 W/K. Both thresholds lie below the point value of 8428 W/K, in the lower tail of what the trip test left uncertain, and a tail's weight grows quickly as its boundary moves toward the middle. The margin, measured at the point value, sees none of this: it falls by the same amount for every added kilowatt. A margin is a statement about one plant, the one whose conductance is exactly $\hat{G}$; the exceedance probability is a statement about every plant the evidence still allows, each weighed by how probable it is.

A second competing formulation draws the envelope at the point value. "Hall B can be run to 67 kW with unit B out" is the envelope of Figure 5.2 at $\hat{G}$. Held at an exceedance probability of 0.05, the same envelope ends at 54.6 kW, about 13 kW lower, and at 0.01 it ends at 50.4 kW. The first number describes a plant whose conductance is known to be 8428 W/K, which this plant is not.

A third checks one limit state of a requirement that has two. Reporting that unit A's capacity margin is 10 kW at the design load, and concluding that the plant meets its requirement with unit B out, answers the capacity limit and leaves the opening's limit unexamined. The requirement's failure domain is the union of the two, and its probability is the probability of the union. A fourth turns a margin into a probability by nothing at all, as in "a margin of 1.5 K, so 98 % safe". A number in kelvin has become a percentage without any step that could make it one.

These are the forms of this lesson's failure class, the *point value read against the limit*: a requirement whose arguments are uncertain, judged at one point value of them, where the probability of exceedance was asked. Its forms are a margin at the point value reported as the verdict; a margin reported as a probability, an exceedance probability or an availability; an envelope drawn at the point value as though the unknowns were known; and one limit state checked where the requirement holds several. The last form sits close to a class that Lesson 13 will own, a component's event substituted for the system's; the difference is that here the parts are limits of one requirement, and there they are components of a system. The class enters at step 4, when the event is written at a point, or at step 5, when a margin is written where a probability was asked, and it is caught at step 7 by asking whether the expression is a number in the limit's units or a probability.

5.4 One system, many questions

Table 5.6 gives nine questions about the limits of the two-hall plant, with the limit state or expression each one concerns in the second column and the kind of answer in the third. It holds no values, because the point is that one limit, written once, answers questions of different kinds

before anything is computed. Lesson 1's Table 1.8 and Lesson 4's Table 4.9 asked the same plant other questions.

Table 5.6: Nine questions about the limits of the two-hall plant.

Ordinary-language question	Limit state or expression	Kind of answer
How far is hall B from its limit tonight, at the point value?	$g(\hat{G};\ Q_B)$	A margin, in kelvin
How probable is it that hall B exceeds its limit tonight?	$P(g(G;\ Q_B) < 0 \mid$ unit B out, Q_B, settled, $E)$	A probability
Up to what load can hall B run with unit B out, at 1 % risk?	The set of Q_B with that probability at most 0.01	A set of loads
Can unit A carry both halls at their design load?	$g_{\text{cap}}(60, 60) = 10\,\text{kW}$	A margin on stated values
Does either limit fail tonight with unit B out?	$P(\{g_{\text{cap}} < 0\} \cup \{g < 0\} \mid$ unit B out, Q_A, Q_B, settled, $E)$	A probability of a union
Would a trip test at 40 kW keep hall B within its limit?	$P(40\,\text{kW}/G \le \Delta T_{\text{allow}} \mid E)$	A probability
What conductance would leave a margin of 2 K at 55 kW?	The G with $g(G;\ 55\,\text{kW}) = 2\,\text{K}$	A value, no probability
How much would widening the opening by a fifth add to the margin at 55 kW?	$g(1.2\,\hat{G};\ 55\,\text{kW}) - g(\hat{G};\ 55\,\text{kW})$	A difference of margins
Is the opening adequate?	$P(G \ge G^* \mid E)$, with the limit state $G - G^*$	A probability

One limit state stands behind most of the rows, and it yields margins, probabilities, a set of loads, a single value of the conductance and a difference of margins. The kind of answer has to be chosen before any number is computed, because each kind answers a different question. One kind needs a name: a *margin on stated values* is a margin whose arguments are all stated, as the capacity limit's are, so that it carries no uncertainty on the stated model and its probability of being negative is zero or one. Two rows deserve a note. The sixth asks whether a test still to be run could be run safely; it is a question about a test's feasibility, which Lesson 21 takes up, and it is named here only because its expression has the same form as the others. The last row takes Lesson 1's first question and writes its event as a limit state on the conductance alone, $G - G^* \ge 0$.

5.5 The exercises

The solutions are in Appendix A, and the facts of the plant in Appendix D. Where an exercise needs evidence on the opening, it is E, the 40 kW trip test.

Exercise 5.1 (*formulate*). The facilities engineer says: "Unit A goes out for its service tomorrow, with hall A at 50 kW, so unit B will have to carry both halls, and hall A's heat will have to come through the opening. I want the chance that either of those goes wrong." Write the two limit states, the failure domain and the expression, and state what the sentence leaves out.

Exercise 5.2 (*formulate*). An operator on the night shift asks: "With unit B out, what is the highest load hall B may carry so that the chance of going over its limit, once the plant has settled, is at most 2 %?" Formulate the question.

Exercise 5.3 (*formulate*). The tenant of hall B asks: "If unit B trips tonight, how probable is

it that my hall stays at or below 30°C?" Write the limit state on hall B's temperature, the expression, and every condition it needs.

Exercise 5.4 (*formulate*). The insurer's engineer asks: "What is the plant's safety margin?" Formulate the question.

Exercise 5.5 (*reverse*). Write one plain sentence for $P\big(G\,\Delta T_{\text{allow}} - Q_B < 0 \mid \text{unit B out},\ Q_B = 45\,\text{kW},\ \text{settled},\ E\big)$.

Exercise 5.6 (*reverse*). Write one plain sentence for $g(\hat{G};\ Q_B)$ with $Q_B = 50\,\text{kW}$, and state what it is not.

Exercise 5.7 (*reverse*). Write one plain sentence for the set

$$\{Q_B : P(g(G;\ Q_B) < 0 \mid \text{unit B out},\ Q_B,\ \text{settled},\ E) \le 0.05\}.$$

Exercise 5.8 (*critique and repair*). A shift report says: "Hall B is at 52 kW with unit B out. The margin at the point value is 1.8 K, which is 23 % of the 8 K limit, so the risk of exceeding the limit is under a quarter." State what the report computes, and repair it.

Exercise 5.9 (*critique and repair*). A design note says: "With unit B out, unit A's capacity margin is 10 kW at the design load, so the plant meets its requirement with unit B out." State what the note checks, and repair it.

Exercise 5.10 (*critique and repair*). A commissioning summary says: "At the point value of the conductance, hall B can carry 67 kW with unit B out. That figure will be published as the limit for operations." State what the figure is, and repair the published limit.

Exercise 5.11 (*rapid*). For each sentence, say whether it asks for a margin, an exceedance probability, an envelope or a margin on stated values, and write the limit state it concerns.

(a) How many kilowatts below its capacity is unit A when both halls are at 58 kW?

(b) How likely is hall B to be over its limit at 48 kW with unit B out?

(c) For which loads of hall B, with unit B out, is the chance of going over the limit below one in twenty?

(d) By how much does hall B clear its limit at 40 kW if the conductance is 9000 W/K?

(e) Is the chance that either limit fails, with both halls at 52 kW and unit B out, above 10 %?

(f) What is hall B's settled rise at the point value with unit B out at 50 kW?

(g) What load of hall B brings the margin at the point value to zero?

(h) Unit A trips with hall A at 50 kW; can hall A stay within its limit?

(i) How likely is unit A's capacity to be exceeded if the two halls together draw 125 kW?

(j) Would doubling the conductance double the margin at a given load?

5.6 What to retain

- A limit is a function of the unknowns and the conditions, positive on the safe side. Writing it names what the model must supply.

- A margin, an exceedance probability and an envelope answer three different questions about one limit state. A margin is in the limit's units, and only the probability is compared with a probability level.
- On one plant and one trip test, with unit B out and hall B at 55 kW, a margin of 1.5 K sits beside an exceedance probability of 0.06.
- A requirement with several limits fails on the union of their failure domains. It has a margin for each limit and one probability.
- A known condition fixes an argument of the limit state and goes behind the bar; an unknown one stays among the unknowns.
- An envelope drawn at the point value is the envelope of a plant whose unknowns are known.

Lesson 6

Joint and Conditional Events

6.1 The engineering question

A reliability review of the two-hall plant reaches its redundancy. One engineer reads from the outage records: each cooling unit is unavailable about one hour in twenty, the two units are redundant, and so both are unavailable together about one hour in four hundred. A second engineer points to the drawing. One feeder supplies both units, and when it goes, both go. A third recalls the disturbance windows in the record: in a bad month both units fail far more often than in a calm one. All three are describing the same two units, and each is describing a different reason why both at once is more probable than the product of the two.

The same morning, in the electrical room of a facility nearby, a review hears a similar argument about the supply of its IT racks. The utility's feeder is out about two hours a year, and the UPS is under repair about two and a half hours a year, so the chance that both are out together is negligible. But the UPS is called upon only when the feeder is out, and its battery carries the full load for ten minutes. Its readiness is not a separate matter from the feeder's state. It is a question that arises only when the feeder has already failed.

Redundancy is worth exactly what the probability of its parts failing together is worth, and "together" is where every dependence of a system lives. Redundant parts can share a part, so that one failure is the failure of both. A cause can act on both at once. One part can be called into service by the other's failure and then fail itself. A failure can spread from one part to the next. Each of these makes a joint failure more probable than the product of the parts' probabilities, and none of them shows in the parts' records taken one at a time.

This lesson writes the joint failure of redundant parts, names what ties them, and puts that tie into the model, where it can be seen and checked.

6.2 The formulation

The work of this lesson sits in step 6, where the owner's method asks for the model's dependencies and Part I wrote "independent", and in step 4, because a shared part changes the event itself. The other steps run as before. This lesson also brings in the second recurring system, the small electrical network.

6.2.1 The joint by the chain rule

Lesson 3 defined the conditional probability, equation (3.1), as the probability of the joint divided by that of the condition. Read the other way, it says how to build a joint from a conditional. The probability that two events both hold is the probability of one of them times

the probability of the other once the first is known:

$$P(\mathcal{A} \cap \mathcal{B}) = P(\mathcal{A} \mid \mathcal{B})\, P(\mathcal{B}) = P(\mathcal{B} \mid \mathcal{A})\, P(\mathcal{A}). \tag{6.1}$$

Equation (6.1) is the multiplication rule, called here the chain rule. Independence is its special case, the case in which knowing that $\mathcal{B}$ holds does not change the probability of $\mathcal{A}$:

$$P(\mathcal{A} \mid \mathcal{B}) = P(\mathcal{A}), \qquad \text{so that} \qquad P(\mathcal{A} \cap \mathcal{B}) = P(\mathcal{A})\, P(\mathcal{B}). \tag{6.2}$$

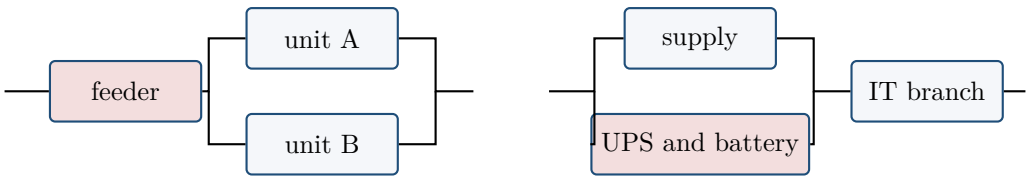
Equation (6.2) is a property of a model, not of the words "two separate units". Two units in separate rooms, with separate records, can still share a supply, a cause or a load, and whether they do is a statement step 6 must make. Every number of Part I took the plant's components as independent, and Lesson 2 said so once; this lesson asks where that assumption fails.

The two directions of equation (6.1) are equal, and they are not equally easy to supply. The useful direction is the one in which the conditional is what the physics or the record provides. For a part that is called into service when another fails, the probability that it fails given the call is a property of that part, which a test or a data sheet can supply; the probability of the call given that the part has failed has no physical source at all. Choosing the direction of the chain is part of the formulation, and the second worked formulation below turns on it.

6.2.2 Redundancy and its cut sets

A *minimal cut set* is a smallest set of failures that loses the service (Vesely and others, *Fault Tree Handbook*, 1981). The service $S(t)$ and the two structures of the two-hall plant are defined, here and in the rest of the book, with each hall at its design load, so that "the opening adequate" means $G \geq G^*$ for hall B's full load with unit B out. With the opening adequate, the two-hall plant's minimal cut sets are the feeder alone and the two units together. Lesson 2's loss event on the parallel structure, equation (2.3), is their union. The list makes one point plainly: a part shared by a redundancy is a cut set of order one, and the redundant pair protects only against failures that are not shared.

Figure 6.1 draws the cut sets of both recurring systems met so far as reliability block diagrams, in which the service flows from left to right and is lost when every path is broken. On the left, the two-hall plant: the feeder in series with the pair of units in parallel. On the right, the network's IT load: a branch in series with a pair of sources in parallel, the supply and the UPS's battery. The shared element of each is shaded.



the two-hall plant, opening adequate — the network's IT load

Figure 6.1: The two recurring systems met so far as reliability block diagrams. On the left the shaded feeder is a cut set by itself; on the right the shaded UPS is called on only when the supply has failed, so its failure is conditional on the other path's.

The right panel needs a word of caution, and it is the subject of the rest of the lesson. A block diagram draws the supply and the UPS as two parallel paths, which invites the reader to treat them as independent and multiply their unavailabilities. The two paths are not alike. The supply is the normal path; the UPS's battery carries the load only while the supply is lost, and only for ten minutes. Its failure is a failure given a demand, and the diagram cannot show that.

6.2.3 Four kinds of dependence

What ties two failures together comes in four kinds, and Table 6.1 lists them, one row each. The second column says what ties the failures, the third gives an instance on the two recurring systems, and the fourth says how the joint is written once the tie is in the model.

Table 6.1: Four kinds of dependence between failures, with an instance of each and the form of the joint.

Kind	What ties the failures	Instance	How the joint is written
A shared component	One part whose failure is the failure of both	The feeder supplying both cooling units	The shared part's event in the union: $F_{\text{fd}} \cup (F_A \cap F_B)$
A shared cause	A condition that raises both failures	The disturbance windows	By conditioning on the cause: the units independent within a window, the joint summed over the window's states
A demand	One part called into service by the other's failure	The UPS, called on when its supply is lost	By the chain rule in the demand's direction: $P(F_{\text{sup}})\,P(F_{\text{ups}} \mid F_{\text{sup}})$
A propagation	A failure that causes another	A branch fault that its breaker fails to clear; the surviving unit carrying both halls	By the conditional of the second failure given the first

The first row returns to Lesson 1's fourth question, the probability of both units being unavailable, and to Exercise 1.7. The events F_A and F_B are the units' own outages, and the feeder's outage F_{fd} makes both units unavailable at once. A unit is unavailable for any reason when it is out for its own reasons or the feeder is out, so both are unavailable on $(F_A \cup F_{\text{fd}}) \cap (F_B \cup F_{\text{fd}})$, which is $F_{\text{fd}} \cup (F_A \cap F_B)$. The shared component appears in the event itself, which is why this kind of dependence enters at step 4.

The second row is the disturbance windows of the plant's record, which Lesson 3 used to separate seeing from setting. A window of 720 h is disturbed with prior weight 0.1; in a disturbed window each unit is unavailable with probability 0.30, in a calm one with 0.05, and within a window of either kind the two units fail independently. The joint is built by summing over the window's state, the cause that ties them:

$$P(F_A \cap F_B) = \sum_{c \in \{\text{disturbed},\ \text{calm}\}} P(F_A \mid c)\,P(F_B \mid c)\,P(c). \tag{6.3}$$

Equation (6.3) is the law of total probability, a probability written as a sum over a partition of what is unknown, here the window's two states, applied to the joint with the units independent given each state. The units are *conditionally independent* given the window and dependent without it: conditional independence is not independence. Knowing that unit B is out makes a disturbed window more probable, which raises the chance that unit A is out too. On the window model, $P(F_A \mid F_B) = 0.150$, against $P(F_A) = 0.075$ before anything is known of unit B. The shared cause appears in the model, not in the event, which is why this kind enters at step 6.

Table 6.2 shows the two terms of equation (6.3) on the window model, one row for each state of the window, with the state's weight, each unit's probability of being out in that state, both units' probability given the state, and the term itself. The stated values supply the first three columns, and the fourth is their product.

The disturbed windows hold one tenth of the weight and four fifths of the joint. That is what a shared cause does: it concentrates the joint failures in the states where both parts are weak, while each part's own record, averaged over the states, shows only a modest unavailability of 0.075.

Table 6.2: The two terms of the law of total probability for both units out, on the window model.

Window's state c	Weight $P(c)$ [–]	Each unit out [–]	Both out given c [–]	Term [–]
calm	0.9	0.05	0.0025	0.00225
disturbed	0.1	0.30	0.0900	0.00900

The third row concerns a part held in reserve. A standby part fails in two ways: on demand, when it is called and does not start, and during its mission, when it starts and then stops before the demand ends. For the UPS, the first is a failure to pick up the load and the second is a battery that runs out before the feeder returns. Both are defined only given the demand, so the joint must be written in that direction, $P(F_{\text{sup}})\,P(F_{\text{ups}} \mid F_{\text{sup}})$, with the second factor read from the UPS's probability of failing to pick up and from how often a loss of supply outlasts the battery.

The fourth row is a failure that causes another. A fault on the cooling branch that its breaker fails to clear opens the breaker above both branches, and the IT racks lose their supply. The probability that the IT branch loses its supply given a fault on the cooling branch is the probability that the cooling branch's breaker fails to clear, a conditional supplied by the protection. With unit B out, unit A carries both halls and works harder, and if its failure rate rises with its duty, the probability of its failure during unit B's repair is conditional on unit B being out. Exercise 6.2 writes that case; the fact sheet states no such rate.

6.2.4 Conditional events in a redundancy

The redundancy's own question is a conditional one: given that one part has failed, how probable is the failure of the other? For the two cooling units it is $P(F_A \mid F_B)$, given the stated model, 0.150 on the window model against 0.075 for unit A alone. For the network it is $P(F_{\text{ups}} \mid F_{\text{sup}})$. When an engineer asks how good the backup is, the question is almost always this conditional, and the joint is the answer to a different question: how often both fail together, at a random hour. Lesson 3 named the error of writing the joint for the conditional as the second form of its class, and this lesson does not teach it again. It does add one observation. The conditional of a redundancy is where its dependence lives, so a formulation that writes the conditional as the second part's own unavailability has assumed independence without saying so.

Table 6.3 sets the conditional of each recurring redundancy beside the unconditional probability of the backup part failing, one row for each system. The third column is the redundancy's conditional, the fourth the backup's own probability at a random hour. The first row uses the window model; the second the network's illustrative inputs, with the conditional read at a feeder interruption.

Table 6.3: The conditional of a redundancy against its backup's own probability of failing, at a random hour.

Redundancy	Conditional	Conditional [–]	Backup's own [–]
The two cooling units	$P(F_A \mid F_B)$, window model	0.150	0.075
The supply and the UPS	$P(F_{\text{ups}} \mid$ a feeder interruption), illustrative	0.85	2.7×10^{-4}

The first row doubles; the second differs by more than three orders of magnitude. The backup's own probability in the second row is the share of time the UPS is under repair, which

says how often it is unavailable when nothing calls on it. The conditional says how often it fails when it is called, and that is the question the redundancy exists to answer.

6.2.5 The small electrical network

The second recurring system arrives here, and its fact sheet is in Appendix D. Figure D.2 draws it as a single line, labeled in words. The utility feeder supplies the facility through the incoming breaker, whose relay opens it on a fault and can open it spuriously. The UPS, with a battery and a static bypass, takes the supply and feeds the PDU, which feeds two branches on their own breakers, the IT racks of 900 kW and the cooling plant of 250 kW. Five meters read the chain. The chain, its loads, its efficiency of 0.96, its meters and three explanations of a meter mismatch come from a case study of a facility's power chain. That study is a steady snapshot with no time history, so the rates of the feeder's interruptions, the protection's behavior, the battery's autonomy of 10 min and the UPS's probability of failing to pick up the load are this book's illustrative inputs, and every number below that rests on them is labeled so. The network's feeder is not the two-hall plant's feeder; the two systems are separate.

6.2.6 Both cooling units unavailable

Lesson 1's fourth question returns here with the feeder and the windows in it. Table 6.4 is its formulation sheet, with steps 4 and 6 in full.

The sheet's sixth row is where this lesson's work shows. Part I would have written "each component's alternation of failure and repair, independent". Here the row names three ties and says which of them the model keeps: the feeder as a part shared by both units, the window as a cause shared by both, and the feeder's rise in a disturbed window, kept out and said to be. A reader of the sheet can see exactly which dependence the number will carry. The competing formulations below size the first two ties separately, each on its own model, and the sheet's expression, which carries both, is written here and not evaluated.

6.2.7 The IT load without supply

The facility's question about its IT racks is "How probable is it that the IT racks are without power at a random hour?". Table 6.5 is its formulation sheet.

The sheet's fifth row chooses the direction of the chain. $P(F_{\text{ups}} \mid F_{\text{sup}})$ is supplied by the UPS's per-demand probability, its battery and the distribution of the losses' lengths. The other direction, the probability that the supply is lost given that the UPS is not carrying, has no source; the UPS's "not carrying" is not even defined while the supply is present. Its sixth row also shows why the distribution of a loss's length enters the model and not only its mean: whether a loss outlasts ten minutes depends on how long losses are, one by one. With exponential lengths, the share of a loss's time that lies beyond the autonomy is a closed expression of the autonomy and the mean length.

6.2.8 Independence as a stated assumption

Every model makes some independence assumptions, since no model ties everything to everything. The rule this lesson adds is that each one is written in step 6, beside the dependences the model keeps, as the sheet of Table 6.4 does. Three questions find the ties a formulation must consider. Is there a part that both failures need, and is it in the event? Is there a condition that raises both, and is it in the model, with the parts independent only given it? Is one failure called into play, or made more likely, by the other? A tie that survives the three questions and is still left out is an assumption, and it is labeled as one; a number computed under it is labeled with it, as every number of this book is labeled with its stated model.

Table 6.4: The formulation sheet filled in for "What is the probability of both units being unavailable?".

Field	Entry
Intention	Predict: an outcome at a random hour
1. Sentence	At an hour chosen at random from steady operation, what is the probability that neither cooling unit can run, for any reason, given the stated rates and the record of disturbance windows?
2. Target	An event at a random hour
3. Known	Steady operation; the stated model: the feeder's stated rates, the structure (one feeder for both units), and the window model with its prior weight, which gives each unit's own probability of being out (0.05 in a calm window, 0.30 in a disturbed one, 0.075 on average) in place of the unit's stated rates
4. Event	$(F_A \cup F_{\text{fd}}) \cap (F_B \cup F_{\text{fd}}) = F_{\text{fd}} \cup (F_A \cap F_B)$, with F_A, F_B the units' own outages
5. Expression	$P\big(F_{\text{fd}} \cup (F_A \cap F_B)\big)$ given the stated model, with $F_A \cap F_B$ summed over the window's state as in equation (6.3)
6. Model	The feeder's alternation of failure and repair; the feeder's outage as a shared component; the window's state as a shared cause, the units independent given it; the feeder's outages independent of the units' own, stated as an assumption. The fact sheet's windows also raise the feeder's unavailability, to 0.08, and that further tie is named here and left out.
7. Read back	The probability that the plant has no cooling unit available at a random hour, whatever the reason. Silent on the opening, which matters only when one unit is left.
Problem class	Predictive probability at a random time, with dependence

Stripped to one line a step:

Run, what is the probability of both units being unavailable

1	**Sentence**	Random hour, steady, stated model: neither unit can run, for any reason?
	Intention	Predict
2	**Target**	An event at a random hour
	Declare	F_A, F_B units' own outages; F_{fd} feeder's outage; c window's state, disturbed or calm
3	**Known**	Steady; feeder's rates, one feeder, windows (a unit out 0.05 calm, 0.30 disturbed)
	Random here	The three components' states at the hour; the window's state
4	**Event**	$(F_A \cup F_{\text{fd}}) \cap (F_B \cup F_{\text{fd}}) = F_{\text{fd}} \cup (F_A \cap F_B)$
5	**Expression**	$P\big(F_{\text{fd}} \cup (F_A \cap F_B)\big)$, given the stated model
6	**Model**	Shared: feeder, window; feeder independent of units (assumption); rise left out
	Reduce	$P(F_A \cap F_B) = \sum_c P(F_A \mid c)\, P(F_B \mid c)\, P(c)$, equation (6.3)
7	**Check**	No unit available at a random hour, for any reason; silent on the opening
	Class	Predictive probability at a random time, with dependence

$$\boxed{P\big(F_{\text{fd}} \cup (F_A \cap F_B)\big), \text{ given the stated model}}$$

Table 6.6 applies the three questions to five pairs of parts of the two recurring systems, one pair a row. The second column names any part the two share, the third any cause that acts on both, the fourth any demand or spread from one to the other, and the last says whether the stated model keeps the tie.

Two of the five rows hold a tie the stated model leaves out, and both are said to be. The second row shows that the same record can tie different pairs to different degrees: the windows raise the units' unavailability sixfold and the feeder's eightfold, and a formulation that keeps the first and drops the second has made a choice that must be written.

The answers are not always known. The fact sheet of the network states no tie between its two branch breakers, and nothing says whether a breaker that once failed to clear makes the other more likely to fail, through a shared age or a shared maintenance crew. A formulation

Table 6.5: The formulation sheet filled in for "How probable is it that the IT racks are without power at a random hour?".

Field	Entry
Intention	Predict
1. Sentence	At an hour chosen at random from steady operation, what is the probability that the IT load is without supply because its supply was lost and the UPS did not cover it, given the network's structure and its illustrative rates?
2. Target	An event at a random hour, $S_{\mathrm{IT}}(t) = 0$, from one cause
3. Known	Steady operation; the structure (the UPS between the supply and the PDU, its battery and its bypass); the fact sheet's rates, per-demand probabilities and autonomy, labeled illustrative
4. Event	$F_{\mathrm{sup}} \cap F_{\mathrm{ups}}$: the UPS's supply lost, and the UPS not carrying the load while it is lost, because it failed to pick up the load, the loss has lasted longer than t_{aut}, or the UPS is out for repair with the load on its bypass. Two further cut sets, a fault on the IT branch and a fault on the cooling branch that its breaker fails to clear, are named and set aside.
5. Expression	$P(F_{\mathrm{sup}} \cap F_{\mathrm{ups}}) = P(F_{\mathrm{sup}})\, P(F_{\mathrm{ups}} \mid F_{\mathrm{sup}})$
6. Model	The supply's losses, their rates and the distribution of their lengths; the UPS's probability of failing to pick up per demand; its autonomy; the UPS's own repairs, which matter only when they coincide with a loss of supply
7. Read back	The share of hours the IT load is without supply because its supply was lost and the UPS did not cover it. Silent on how many separate interruptions those hours make, which Lesson 8 takes up.
Problem class	Predictive probability at a random time, with a demand

The sheet's run:

Run, how probable is it that the IT racks are without power at a random hour

1	**Sentence**	Random hour, steady, illustrative rates: IT load without supply, loss not covered?
	Intention	Predict
2	**Target**	An event at a random hour, $S_{\mathrm{IT}}(t) = 0$, from one cause
	Declare	$S_{\mathrm{IT}}(t)$ IT load supplied; F_{sup} supply lost; F_{ups} UPS not carrying; t_{aut} [min]
3	**Known**	Steady; the structure; rates, demand probabilities, autonomy (illustrative)
	Random here	The supply's state; each loss's length; the UPS's pickup and repairs
4	**Event**	$F_{\mathrm{sup}} \cap F_{\mathrm{ups}}$; two branch cut sets named and set aside
5	**Expression**	$P(F_{\mathrm{sup}} \cap F_{\mathrm{ups}}) = P(F_{\mathrm{sup}})\, P(F_{\mathrm{ups}} \mid F_{\mathrm{sup}})$
6	**Model**	Losses' rates and lengths; pickup per demand; autonomy; the UPS's repairs
7	**Check**	Share of hours without supply from an uncovered loss; silent on how many breaks
	Class	Predictive probability at a random time, with a demand

$$\boxed{P(F_{\mathrm{sup}})\, P(F_{\mathrm{ups}} \mid F_{\mathrm{sup}})}$$

cannot invent that tie and must not deny it. It states the assumption it makes and says what a record of the breakers' tests would decide.

6.2.9 A redundancy claim that is not ready

At the facility's handover, the electrical contractor states that the IT load has a redundant supply, and the owner asks how much that is worth. The sentence names a redundancy without saying against what failure, over what time, or at what level, so it fails at the first step. Three readings come naturally, and each needs a different dependence in its model.

Read at a random hour, the question is the second worked formulation's, $P(F_{\mathrm{sup}} \cap F_{\mathrm{ups}})$, whose model must carry the demand. Read as "the IT load rides through the next utility interruption", it is a conditional on the demand, $P(F_{\mathrm{ups}}$ does not occur $\mid$ an interruption begins), a probability about one event, in which the distribution of the interruption's length is the whole of the

Table 6.6: The three questions applied to five pairs of parts.

Pair	Shared part	Shared cause	Demand or spread	Kept by the stated model
Unit A and unit B	the feeder	the disturbance windows	load sharing during a repair	the feeder and the windows; load sharing not stated
The feeder and a unit	none	the windows raise both	none	no: the feeder's rise in a disturbed window is left out
The supply and the UPS	none	none stated	a demand: the UPS is called when the supply is lost	yes, through the per-demand probability and the autonomy
The IT and cooling branches	the UPS and the PDU	none stated	a spread through protection	yes
The two branch breakers	none	age or maintenance, possibly	none	no: not stated, an assumption either way

question. Read as "the IT load is protected against every failure in the chain", it is not a claim of redundancy at all: the PDU and each branch lie in series with both sources, as Figure 6.1 shows for the IT branch, and a fault there, or a cooling-branch fault its breaker does not clear, defeats the redundancy entirely. The most literal reading, that two sources exist, is a fact about the drawing, as "the plant has N+1" was in Lesson 1, and it leaves the owner's question where it was.

The analyst's reply names the three readings and asks which the owner means, and it says one thing more: the worth of this redundancy is decided by the demand's length beside the battery's autonomy, which no count of sources can show.

Key idea

The probability that redundant parts fail together is a joint, not a product. Write it by the chain rule in the direction the physics supplies, find the shared part, the shared cause, the demand and the spread, and state every independence the model keeps.

6.3 The competing formulations

The competing formulations of this lesson all write a joint as the product of its parts' own probabilities. Table 6.7 gives three such products beside the same joints with the dependence kept, one row for each of the first three kinds of Table 6.1, all at a random hour from steady operation. The first row uses the stated rates of Lesson 1's model (each unit 2.5 failures per year in service with a 180 h repair, the feeder 3.69 with a 24 h repair), the second the window model, and the third the network's illustrative inputs. The values come from a small calculation with each stated model.

Figure 6.2 draws the same three pairs on a logarithmic scale, one line per row of the table, with the product as an open circle and the joint with its dependence as a filled one. The length of each segment is the logarithm of the ratio, and a step of the axis is a factor of ten.

The first row is the rubric's instance of the class. The product is the probability that both units are down for their own reasons, the units taken as independent; the joint with the feeder counted is five times larger, because the feeder, which takes both units at once, accounts for most of the hours in which neither unit can run. The second row keeps every part and drops

Table 6.7: Three joints written as products of their parts' own probabilities and with the dependence kept, at a random hour. The third row rests on illustrative inputs.

Joint	As a product [–]	Dependence kept [–]	Ratio [–]
Both cooling units unavailable: the feeder shared; stated rates	0.0024	0.012	5
Both units out for their own reasons: the window model	0.0056	0.011	2
IT load without supply from a loss of supply: illustrative	6.4×10^{-8}	1.9×10^{-4}	about 3000

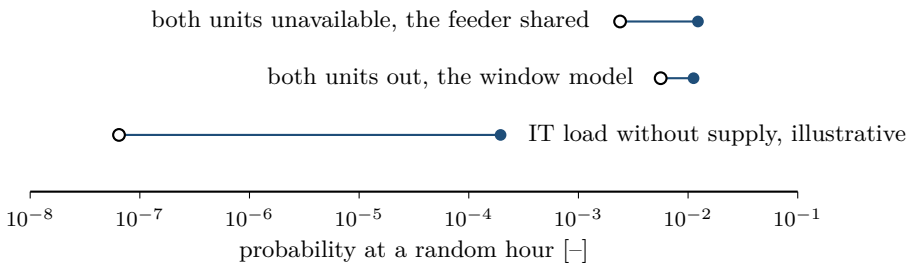


Figure 6.2: The three joints of Table 6.7 as products (open circles) and with their dependence kept (filled circles), on a logarithmic axis.

only the tie between them. Within a window the units are independent, and the product of their marginal probabilities is still wrong by a factor of two, because the marginals average over windows of both kinds while the joint is dominated by the disturbed ones.

The product, marked and repaired:

Run, both units unavailable written as a product

1	**Sentence**	Random hour: neither unit can run, for any reason?
4	**Event** ×	$F_A \cap F_B$, own outages only
	Repair	$F_{\text{fd}} \cup (F_A \cap F_B)$, the shared feeder in the event
5	**Expression** ×	$P(F_A)\,P(F_B)$, stated rates: 0.0024
	Repair	$P\big(F_{\text{fd}} \cup (F_A \cap F_B)\big)$, stated rates: 0.012
7	**Check**	"Any reason" includes the part both units need

$$\boxed{P\big(F_{\text{fd}} \cup (F_A \cap F_B)\big),\ \text{given the stated model}}$$

The third row multiplies the supply's unavailability by the UPS's unavailability for repair, as if the UPS's readiness had nothing to do with its supply. With the demand kept, the UPS fails to cover a feeder interruption 85 % of the time on these inputs, because a battery of ten minutes meets interruptions that last an hour on average; the IT load is then without supply about 1.7 h a year. The product puts the same quantity some three thousand times lower. The values of the third row rest on illustrative inputs, and the size of the ratio is the point rather than its digits: a reserve part's failure, taken at its unconditional probability, can be wrong by orders of magnitude when its mission is short beside its demand.

A fourth competing formulation asserts independence from separateness: two units with separate records are taken as independent. Separate records show each part's own behavior and nothing about the tie between them; the window model's units have separate records and are dependent.

These are the forms of the recurring class this lesson sharpens, *dependence ignored*: a joint event evaluated as if its parts were independent when the model ties them. Its four forms are a shared component left out of the cut sets; a shared cause's tie dropped, with parts that are

independent only given the cause; a failure conditional on a demand taken at its unconditional probability; and a propagating failure treated as two separate events. It enters at step 4, where a shared part should appear in the event, or at step 6, where a tie should appear in the model, and it is caught at step 7 by asking, for each pair of parts, what could fail both.

6.4 One system, many questions

Table 6.8 gives nine questions about the electrical network, each with its expression in the second column and, in the third, the kind of dependence it involves. It holds no values. The point is that a network of few parts yields joints of every kind.

Table 6.8: Nine questions about the electrical network and the dependence each involves.

Ordinary-language question	Mathematical formulation	Dependence
How probable is it that the UPS fails to carry the next interruption?	$P(F_{\text{ups}} \mid \text{an interruption begins})$	A demand
Does a fault on the cooling branch take the IT racks with it?	$P(\text{IT branch without supply} \mid$ a cooling-branch fault$)$	A propagation, through protection
Are the IT and cooling branches without supply together at a random hour?	$P(\text{both branches without supply})$	A shared component, the UPS and the PDU
Is the IT load without supply at a random hour?	$P(F_{\text{sup}} \cap F_{\text{ups}}) +$ $P(\text{IT branch faulted}) +$ $P(\text{uncleared cooling-branch fault})$	A demand, and two separate cut sets
Is the feeder interrupted while the UPS is under repair?	$P(F_{\text{sup}} \cap \{\text{UPS under repair}\})$	None on the stated model
With the UPS on its bypass for repair, does an interruption reach the IT load?	$P(S_{\text{IT}}(t) = 0 \mid$ UPS under repair, $F_{\text{sup}}) = 1$	A condition behind the bar, not a dependence
Does a spurious trip of the incoming relay cost the IT load anything?	$P(F_{\text{ups}} \mid \text{a spurious trip})$	A demand, short
If one branch breaker has failed to clear before, is the other more likely to?	$P(\text{second fails to clear} \mid$ first failed to clear$)$	A shared cause if one exists; not stated
Do the two meters on either side of the UPS read high together?	$P(\text{both offsets positive})$	A shared cause in calibration; not stated

Each pair of the network's parts is tied by a different mechanism, or by none, and step 6 must say which, or the joint is wrong. The fourth row adds three cut sets because they are disjoint enough on the stated model; a branch fault during a loss of supply is counted twice, and its probability is negligible beside any term, which the sheet would say. The last two rows ask about dependences that the fact sheet does not state. A formulation of either must name an assumption one way or the other, and say what record would decide it.

6.5 The exercises

The solutions are in Appendix A, and the facts of both systems in Appendix D.

Exercise 6.1 (*formulate*). The protection engineer asks: "If the cooling branch faults next year, how likely is it that the IT racks go down with it, and how likely is it that this happens at all next year?" Write the two expressions, and state which of them is a conditional and which a joint.

Exercise 6.2 (*formulate*). Unit B has failed, and its repair takes 180 h on average. The vendor of unit A says that while unit A carries both halls its failure rate is twice its normal rate; take that figure as illustrative. Write the probability that hall B keeps its cooling until unit B returns, and name each dependence in it.

Exercise 6.3 (*formulate*). The minutes of a design review record the reliability lead's remark that the two cooling units are independent, and the action list carries an item that reads "Confirm the independence of the cooling units." Write what the action item asks the team to establish.

Exercise 6.4 (*formulate*). The hall operator says: "For the IT room to stay up through a utility interruption, both its racks and its cooling plant must keep their power. Write the chance that both keep it through the next interruption." Write the expression, and state how the two events are related.

Exercise 6.5 (*reverse*). Write one plain sentence for $P(F_A \mid F_B,\ \text{window model})$, and state why it differs from $P(F_A)$.

Exercise 6.6 (*reverse*). Write one plain sentence for $P(F_{\text{sup}})\,P(F_{\text{ups}} \mid F_{\text{sup}})$, and name the event it is the probability of.

Exercise 6.7 (*reverse*). Write one plain sentence for

$$\sum_{c\,\in\,\{\text{disturbed},\ \text{calm}\}} P(F_A \mid c)\,P(F_B \mid c)\,P(c).$$

Exercise 6.8 (*critique and repair*). A protection study says: "Each load branch faults 0.1 times a year and is repaired in 4 hours, so the probability that both branches are without power in the same hour is the square of $0.1 \times 4/8760$, about 2×10^{-9}." State what the study computes, and repair it.

Exercise 6.9 (*critique and repair*). Unit A has just failed. The duty engineer says: "Unit B fails independently of unit A, so its chance of failing during unit A's repair is the same as over any other 180 hours." State what the statement assumes, and repair it.

Exercise 6.10 (*critique and repair*). An electrical engineer writes: "The incoming relay's spurious trips are reclosed in five minutes on average, well inside the UPS battery's ten minutes, so they cost the IT load nothing." State what the argument assumes, and repair it.

Exercise 6.11 (*rapid*). For each pair of events, say whether the stated model makes them independent, and if not, name what ties them.

(a) Unit A out for its own reasons; the feeder out; at a random hour.

(b) The IT branch and the cooling branch without supply, during an interruption the UPS fails to cover.

(c) The UPS failing to pick up the load at two successive interruptions.

(d) Unit A's and unit B's own outages, within a disturbed window.

(e) The UPS's supply lost; the UPS under repair.

(f) Hall A and hall B without cooling at the same hour, with the opening adequate.

(g) A fault on the IT branch; the cooling branch without supply.

(h) The PDU meter reading high; the UPS's efficiency below nominal.

(i) Two interruptions of the feeder in the same month.

(j) Unit B failing during unit A's repair; unit A's repair being long.

6.6 What to retain

- The probability that redundant parts fail together is the whole value of the redundancy, and it is a joint, not a product.
- The chain rule writes the joint as a conditional times a marginal. Choose the direction in which the physics supplies the conditional.
- A shared component is a cut set of order one; a shared cause makes parts independent only given the cause; a standby part fails given the demand; a failure can spread.
- Independence is an assumption of step 6, written where it is used.
- On the two-hall plant's stated rates, both units unavailable is 0.012 with the feeder counted and 0.0024 for their own outages alone; on the network's illustrative inputs, the IT load is without supply about three thousand times more often than the product of two unavailabilities says.

Lesson 7

Questions About Uncertain Parameters

7.1 The engineering question

At the handover of the two-hall plant, the commissioning engineer reports that after the trip test at 40 kW the team is 82 % confident that the opening is adequate. The owner writes down "82 % chance the cooling works". The two sentences use one number for two different things. The first is about a property of the built opening, which has one value whether or not anyone knows it. The second is about what the plant will do, hour by hour, which depends on that property and also on failures that have not happened yet. On this plant the two are different numbers, and the second is the larger.

The same week, in a facility's battery room, a battery string carries a critical load of 20 kW whenever the grid fails. The site's specification asks for two hours at that load. The facilities manager asks two questions in one breath: whether the string has the capacity for its two hours, and whether it will carry the load through the next grid outage. The first is about the string, the second about an event, and with the string's health uncertain and the outages of uncertain length, the two get answers that point in opposite directions. The string very probably does not meet its specification, and it very probably carries the next outage. The sizes are given in the competing formulations below.

Both meetings mix a question about a quantity that has one value with a question about an outcome that depends on it. Lesson 1 sorted the first kind under estimating and the second under predicting, and it separated the uncertainty of a quantity nobody knows from the variation that remains when everything knowable is known. It did not say how the probability of an outcome is built from what is known of the quantity, and so it left the reader without a way to tell the two numbers apart once both have been computed.

This lesson decides whether a question is about a quantity with one value or about an outcome that depends on it, writes each in its own form, and says what the outcome's probability averages over.

7.2 The formulation

The work of this lesson sits in step 2, where the target is named as a parameter or an outcome, and in step 5, where an outcome's expression carries the parameter as an average rather than as a condition. Step 6 gains three words. This lesson also brings in the third recurring system, the battery string.

As in Lesson 5, the evidence on the two-hall plant is E, the 40 kW trip test: one test at 40 kW per hall, its measured difference of 4.78 K, the drawing's prior on the conductance and

the sensors' stated spreads. After it, $P(G \geq G^* \mid E) = 0.82$.

7.2.1 Parameters and outcomes

A *parameter* is a quantity of the model with one value, unknown. Its uncertainty is epistemic, and a question about it asks to estimate, as Lesson 1 sorted the conductance. An *outcome* is still to occur. It depends on the parameters and on what has not happened yet, so it carries both kinds of uncertainty, and a question about it asks to predict. Table 7.1 pairs five parameters of the three recurring systems with an outcome each one bears on, one pair a row, with the system in the first column, the parameter in the second and the outcome in the third.

Table 7.1: Five parameters of the recurring systems and an outcome each one bears on.

System	Parameter, with one value	Outcome, still to occur
Two-hall plant	The opening's conductance G	The cooling delivered at a random hour
Two-hall plant	Unit A's failure rate λ_A	Unit A failing in the coming year
Two-hall plant	Which drawing of the opening is right	The number of interruptions in the coming year
Electrical network	The UPS's efficiency	The utility energy drawn next month
Battery string	The state of health φ, and whether a cell is weak	The next grid outage carried to its end

One test separates the columns. If the parameter were known exactly, would the outcome still be uncertain? For every row the answer is yes: a unit with a known rate may or may not fail this year, and a string of known health may or may not meet an outage longer than it can carry. When the answer is yes, the question about the parameter and the question about the outcome are different questions, and their probabilities are different numbers.

7.2.2 Prior, likelihood and posterior

Three words name the parts of what is known of a parameter. What is believed about it before the evidence is its *prior*, written as a probability density $p(\theta)$ for a parameter θ that varies continuously; the lowercase letter keeps a density apart from a probability. The probability of the evidence given the parameter, read as a function of the parameter, is the *likelihood*, $P(E \mid \theta)$. What is believed after the evidence is the *posterior*, $p(\theta \mid E)$. Lesson 3's Bayes' rule, equation (3.2), is the relation between them: the posterior is proportional to the likelihood times the prior. This lesson uses the three words as names for parts of step 6 and builds none of them. Lesson 11 writes a likelihood from a sensor and a record, and Lesson 12 writes priors from drawings, handbooks and judgment.

For the opening, the prior is the drawing's, centered on 7500 W/K with a spread of a factor of three; the likelihood is how probable the 40 kW test's measured difference would be at each conductance, given the sensors' offsets and noise; the posterior is what the test leaves known, from which $P(G \geq G^* \mid E) = 0.82$ is read. That number is a statement about the parameter.

7.2.3 The predictive probability

The probability of an outcome takes the outcome's probability at each value of the parameter and weighs it by how probable that value is, given the evidence:

$$P(\mathcal{A} \mid E) = \int P(\mathcal{A} \mid \theta,\ E)\, p(\theta \mid E)\, \mathrm{d}\theta. \tag{7.1}$$

Equation (7.1) is the *predictive probability.* It is Lesson 6's law of total probability, equation (6.3), with the parameter as the unknown it sums over. In most formulations the outcome depends on the evidence only through the parameter: once the parameter is known, the readings already taken say nothing more about the outcome, and $P(\mathcal{A} \mid \theta,\ E)$ reduces to $P(\mathcal{A} \mid \theta)$. That reduction is an assumption of step 6, and the rest of this lesson makes it and says so; a sensor whose offset also bears on the outcome would break it. The two factors under the integral carry the two kinds of uncertainty: $P(\mathcal{A} \mid \theta)$ holds what still varies at a known parameter, the aleatory part, and $p(\theta \mid E)$ holds what is not known of the parameter, the epistemic part. The predictive averages the first over the second. The expression is written here and not evaluated.

A different expression is often offered in its place. The *plug-in* fixes the parameter at a point value and reads the outcome's probability there,

$$P(\mathcal{A} \mid \theta = \hat{\theta}), \tag{7.2}$$

which is the probability of the outcome for a system whose parameter is known to equal $\hat{\theta}$. Figure 7.1 shows the two on one axis of the parameter. The upper curve is the posterior weight; the lower curve is the outcome's probability at each value of the parameter, rising as the parameter improves. The plug-in is the lower curve read at one point. The predictive is the lower curve averaged under the upper one, and it differs from the plug-in whenever the lower curve bends where the upper one has weight.

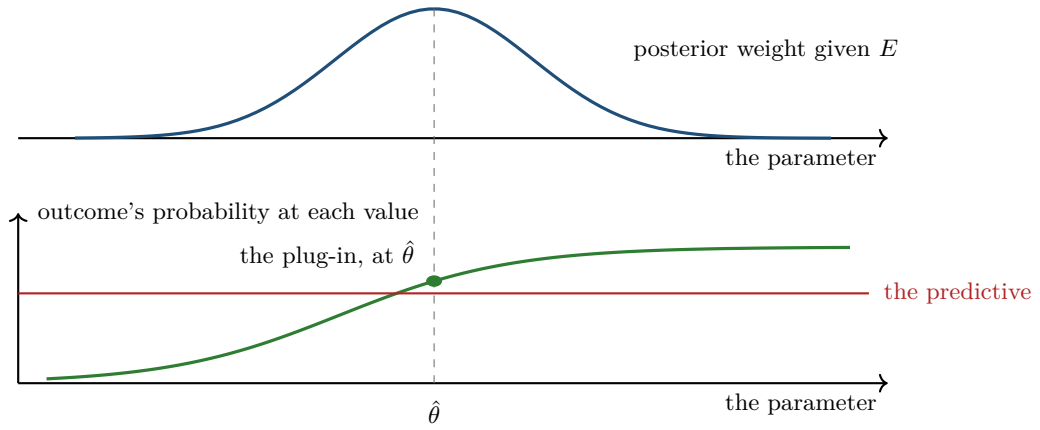


Figure 7.1: The predictive probability and the plug-in. The predictive averages the lower curve under the posterior weight of the upper; the plug-in reads the lower curve at one point. The shapes are illustrative.

With the predictive written, the difference between the two kinds of question can be seen in every field of the formulation sheet, and Table 7.2 sets the two kinds of sheet side by side, one field a row. The second column holds what a sheet about a parameter contains, the third what a sheet about an outcome that depends on it contains.

The sixth row carries the table's point. An outcome's sheet needs everything a parameter's sheet needs and one thing more: the process by which the parameter, together with chance, produces the outcome. For the cooling service that is the failure and repair of the units and the feeder; for the battery string, the length of the next outage. A sheet about the outcome that names no such process has, in fact, been written about the parameter.

7.2.4 A hypothesis as a discrete unknown

Some unknowns take one of a few values. The two-hall plant's case study offers two candidate drawings of the opening for the plant as built, a near pair whose heat-flow capacities of 54 and

Table 7.2: The fields of a formulation sheet for a parameter and for an outcome that depends on it.

Field	A parameter	An outcome that depends on it
Intention	Estimate	Predict
2. Target	A quantity with one value, or an event on it	An event still to occur
3. Known	The evidence on the parameter	The same evidence, and a start or a present state if time enters
4. Event	A set of values of the parameter, such as $G \geq G^*$	A set of future states, such as $S(t) = 1$
5. Expression	$P(\theta \in \text{the set} \mid E)$	$\int P(\mathcal{A} \mid \theta)\, p(\theta \mid E)\, \mathrm{d}\theta$
6. Model	The prior and the likelihood of the parameter	All of that, and the process that turns the parameter into the outcome
Problem class	Posterior exceedance	Predictive probability

66 kW correspond to conductances of 6750 and 8250 W/K. The first lies below the threshold G^*, which puts the plant on the series structure H_{OR}; the second lies above it, which puts the plant on the parallel structure H_{AND}. The two hypotheses are the two values of one discrete unknown, each with a prior weight, and Lesson 1's fifth question, "Which fault-tree structure is supported?", can now be written: $P(H_{\text{AND}} \mid E)$ is a posterior over two values. The predictive of an outcome is then a sum rather than an integral,

$$P\big(S(t) = 1 \mid E\big) = \sum_{i \,\in\, \{\text{AND},\, \text{OR}\}} P\big(S(t) = 1 \mid H_i,\ \text{rates}\big)\, P(H_i \mid E), \tag{7.3}$$

which averages the outcome's probability under each structure by the probability of that structure. The standard name for such an average over hypotheses is model averaging.

The same question about the opening can thus be asked on two models of the one unknown. With the conductance continuous and the drawing's prior, the 40 kW trip test gives $P(G \geq G^* \mid E) = 0.82$. With the conductance one of the two drawings, each with prior weight one half, the same test gives $P(H_{\text{AND}} \mid E) = 0.784$. Both answer "Is the opening adequate?", and they differ because the models differ: the continuous one allows every conductance and weighs them by the drawing's prior, while the discrete one allows only the two values the drawings name. Choosing between them is a decision of step 6, and the sheet must say which was made. Lesson 19 treats competing structural hypotheses in full, including whether a list of them is complete; this lesson says only that a hypothesis is a parameter with a few values, and that its predictive is a sum.

Table 7.3 sets the two models of the one unknown side by side, one feature a row. The second column is the continuous model, the third the model of the two drawings.

Table 7.3: Two models of the opening's one unknown, with E the 40 kW trip test.

Feature	G continuous	G one of two drawings
Values allowed	Every positive conductance	6750 or 8250 W/K
Prior	The drawing's, centered on 7500 W/K, a factor of three either way	Weight one half on each drawing
The question	$P(G \geq G^* \mid E)$	$P(H_{\text{AND}} \mid E)$
Value	0.82	0.784
What it assumes	Nothing about which values are possible beyond the prior's spread	That one of the two drawings is right

The fifth row is the one a reader of either number needs. The discrete model's answer is

conditional on its list of hypotheses being complete, and a conductance between or beyond the two drawings has no place in it. The continuous model has no such list, and pays for it by spreading its prior over values no drawing proposes. Neither number is the true probability that the opening is adequate; each is that probability on its model, and the sheet names the model.

7.2.5 A function of parameters

Some parameters are functions of others. The availability of the cooling service at the stated structure is a function of the units' and the feeder's failure and repair rates, each a parameter with one value; so the availability itself has one value, unknown when the rates are. The plant's case study gives what is known of it after a three-year outage record, on its own priors for the rates and with the parallel structure taken as true. Table 7.4 lists it, one quantity a row: the value at the stated rates, three points of its posterior spread, and its posterior mean.

Table 7.4: The availability of the cooling service on the parallel structure after a three-year outage record, as the case study reports it.

Quantity	Value [–]
At the rounded unavailabilities 0.05 and 0.01, a plug-in	0.9875
Posterior 5 % point	0.9770
Posterior median	0.9874
Posterior 95 % point	0.9930
Posterior mean	0.9865

The table describes two different things. Its middle three rows describe what is known about the plant: its availability, a parameter, lies between 0.977 and 0.993 with probability 0.9. Its last row answers a question about an outcome. The probability that a random hour is served, given the record, is the predictive of equation (7.1) with the availability's posterior under the integral, which is the availability's posterior mean, 0.9865. The first row, 0.9875, is a plug-in, the case study's value at the rounded unavailabilities of 0.05 for a unit and 0.01 for the feeder; at the stated rates themselves it is 0.988. A band on a predictive probability is therefore a band on the parameter in disguise: an outcome at a random hour is served or not, and its probability given the evidence is one number.

7.2.6 The battery string

The third recurring system arrives here, and its fact sheet is in Appendix D. Figure D.3 draws it in words: the grid feed, marked as lost; the string as a tank between its state of charge at the outage's start and its reserve floor; the 20 kW critical load; and the waste heat sent to the room. Its source is a case study of one 75 kWh pack, 80 % full at the start of an outage with a floor at 20 %, discharging at an efficiency of 0.90; it carries the load for 2.025 h. The source is deterministic. What an owner of a string in service does not know is this book's, labeled illustrative: the state of health φ of a sound string, equally likely anywhere between 0.85 and 1.00; a weak cell, with prior probability 0.10, which leaves 0.70 of the usable energy; grid outages of exponential length with a mean of one hour; and a site specification of two hours, set just under the nameplate's 2.025 h so that this lesson can show the contrast it shows. The ride-through T_{rt} is $\varphi \times 2.025$ h for a sound string and 0.70 of that with a weak cell.

7.2.7 The opening and the hour

The handover meeting's two sentences are written side by side in Table 7.5, one column for each, read downward from the sentence to the problem class.

Table 7.5: The seven steps on "Is the opening adequate?" and on "Is the cooling delivered?", with E the 40 kW trip test.

Step	The opening adequate	A random hour served
Intention	Estimate	Predict
1. Sentence	Given the 40 kW trip test, what is the probability that the opening's conductance is at least G^*?	Given the 40 kW trip test and the stated rates, what is the probability that the cooling is delivered at an hour chosen at random from steady operation?
2. Target	A parameter, G	An outcome, $S(t)$ at a random t
3. Known	E: the test, the drawing's prior, the sensors' spreads	The same E; the stated rates, taken as known; steady operation
4. Event	$G \geq G^*$	$S(t) = 1$, which on the plant's structure depends on G
5. Expression	$P(G \geq G^* \mid E)$	$\int P(S(t) = 1 \mid G,\ \text{rates})\, p(G \mid E)\, \mathrm{d}G$
6. Model	The opening's law, the readings' model, the prior on G	All of that, and the failure and repair process on each structure
7. Read back	Whether the built opening passes a hall's load	How often the cooling is on, averaged over what is known of the opening
Problem class	Posterior exceedance	Predictive probability at a random time

Each column as a run:

Run, is the opening adequate

1	**Sentence**	Given E: probability $G \geq G^*$?
	Intention	Estimate
2	**Target**	A parameter, G
	Declare	G [W/K]; $G^* = 7500\,\mathrm{W/K}$; E the 40 kW trip test
3	**Known**	E: the test, the drawing's prior, the sensors' spreads
	Random here	G, one unknown value
4	**Event**	$G \geq G^*$
5	**Expression**	$P(G \geq G^* \mid E)$
6	**Model**	The opening's law; the readings' model; the prior on G
7	**Check**	The built opening passes a hall's load
	Class	Posterior exceedance

$$\boxed{P(G \geq G^* \mid E)}$$

Run, is the cooling delivered

1	**Sentence**	Given E and the stated rates: probability the cooling is on at a random hour?
	Intention	Predict
2	**Target**	An outcome, $S(t)$ at a random t
	Declare	$S(t)$ service at t; $p(G \mid E)$ density of G; the rates taken as known (assumption)
3	**Known**	The same E; the stated rates; steady operation
	Random here	The components' states at the hour; G, averaged over
4	**Event**	$S(t) = 1$, which depends on G through the structure
5	**Expression**	$\int P(S(t) = 1 \mid G,\ \text{rates})\, p(G \mid E)\, \mathrm{d}G$
6	**Model**	As above, with failure and repair on each structure
	Reduce	Two-term sum, equation (7.3), weight $P(G \geq G^* \mid E)$ on the parallel structure
7	**Check**	How often the cooling is on, averaged over the opening
	Class	Predictive probability at a random time

$$\boxed{\int P(S(t) = 1 \mid G,\ \text{rates})\, p(G \mid E)\, \mathrm{d}G}$$

The fifth row of the sheet's second column is an integral over the conductance, and it takes a simple form here. At a known conductance, the probability of service at a random hour takes only two values, one for each structure, because the service and the two structures are defined with each hall at its design load, so that the conductance matters only through which side of G^* it lies on. The integral is therefore the two-term sum of equation (7.3), with $P(G \geq G^* \mid E)$ as the weight of the parallel structure. That is the sense in which the outcome's probability is built from the parameter's: the parameter's probability enters it as a weight and not as the answer.

7.2.8 The string and the next outage

The facilities manager's two questions about the battery string are written side by side in Table 7.6. The evidence here is the fact sheet's prior information only: the nameplate, the stated range of health, the prior on a weak cell, and the outages' mean length. No test of the string has been made.

The predictive in the sheet's second column averages over three unknowns of two kinds. The state of health is epistemic and varies continuously; the weak cell is epistemic and takes one of two values; the outage's length is aleatory. The expression in its fifth row is therefore an integral over φ, a sum over the two hypotheses, and an average over D_{out}, all at once. It is written and not evaluated.

7.2.9 An outcome decided by its parameters

The test of the first subsection has a second answer. If knowing the parameter exactly would leave nothing to chance, the outcome's probability is a parameter's probability under another name. A planned switch-off of the grid supply for 90 minutes is a setting, written do(grid off for 1.5 h) by Lesson 3's rule; its length is fixed. Whether the string carries the load through it depends only on whether $T_{\text{rt}} \geq 1.5$ h, which is decided once φ and the hypothesis are known. Its probability is therefore $P(T_{\text{rt}} \geq 1.5 \text{ h} \mid E)$, a question about the parameters. The distinction between parameter and outcome matters exactly where something is still left to chance after the parameters are fixed, and the formulation must say whether it is.

7.2.10 A parameter that changes

A parameter has one value only over the time the formulation covers. The opening's conductance is fixed while the wall stands, and Lesson 1 said a horizon would enter only if it could drift. A unit's failure rate, the UPS's efficiency and the string's state of health can all change: units age, conversion losses grow, cells fade. When a question reaches far enough ahead that its parameter may change, the parameter becomes part of the outcome. "What will the string's ride-through be in three years?" asks about the state of health at a future time, an outcome, and its model needs the process of fading, which the fact sheet does not state. The formulation must then say either that the parameter is taken as fixed over the horizon, an assumption to be labeled, or what process changes it. A formulation that holds a parameter fixed across a long horizon without saying so has made the choice silently.

7.2.11 A question about a unit that is not ready

At a maintenance review the reliability engineer asks how sure the team is about unit A. The sentence names a unit and no quantity, so it fails at the first step: "sure about" can be about a parameter or an outcome, and the sentence does not say which. Four readings come naturally. Read as a question about the unit's failure rate, it is a parameter question, $P(\lambda_A \in \text{a stated range} \mid \text{the record})$, which needs the range. Read as a question about the

Table 7.6: The seven steps on "Does the string meet its two-hour specification?" and on "Will it carry the next grid outage?".

Step	The specification	The next outage
Intention	Estimate, for a decision on the specification	Predict
1. Sentence	Given what is known of the string, what is the probability that it can carry the critical load for at least two hours from the stated start?	Given what is known of the string and of the grid's outages, what is the probability that the next outage ends before the string reaches its floor?
2. Target	Parameters: φ and the weak cell, through T_{rt}	An outcome: the next outage carried
3. Known	The nameplate, the stated start and floor, the range of φ, the prior on H_{weak}	The same, and the outages' length distribution
4. Event	$T_{\mathrm{rt}} \geq 2$ h, with $T_{\mathrm{rt}} = \varphi \times 2.025$ h, or 0.70 of that with a weak cell	$D_{\mathrm{out}} \leq T_{\mathrm{rt}}$
5. Expression	$P(T_{\mathrm{rt}} \geq 2 \text{ h} \mid E)$	$P(D_{\mathrm{out}} \leq T_{\mathrm{rt}} \mid E)$, averaged over φ, the weak cell and D_{out}
6. Model	The energy balance of the string; the prior on φ; the weak-cell hypothesis	All of that, and the distribution of an outage's length
7. Read back	Whether the string can do what the specification asks	How often an outage like the next one is carried to its end
Problem class	Posterior exceedance	Predictive probability

The two columns, run by run:

Run, does the string meet its two-hour specification

1	**Sentence**	Given what is known of the string: probability it carries the load at least 2 h?
	Intention	Estimate, for a decision
2	**Target**	Parameters φ and the weak cell, through T_{rt}
	Declare	φ state of health [–]; H_{weak}, H_{sound}; T_{rt} ride-through [h]; E the fact sheet's prior
3	**Known**	The nameplate; start and floor; the range of φ; the prior on H_{weak}
	Random here	φ and the hypothesis, each one unknown value
4	**Event**	$T_{\mathrm{rt}} \geq 2$ h; $T_{\mathrm{rt}} = \varphi \times 2.025$ h, or 0.70 of that with a weak cell
5	**Expression**	$P(T_{\mathrm{rt}} \geq 2 \text{ h} \mid E)$
6	**Model**	The string's energy balance; the prior on φ; the weak-cell hypothesis
7	**Check**	The string can do what the specification asks
	Class	Posterior exceedance

$$\boxed{P(T_{\mathrm{rt}} \geq 2 \text{ h} \mid E)}$$

Run, will the string carry the next grid outage

1	**Sentence**	Given the string and the outages: probability the next outage ends before the floor?
	Intention	Predict
2	**Target**	An outcome: the next outage carried
	Declare	D_{out} the next outage's length [h]; T_{rt}, φ, the hypotheses as above
3	**Known**	The same; the outages' length distribution
	Random here	D_{out}, aleatory; φ and the weak cell, epistemic, averaged over
4	**Event**	$D_{\mathrm{out}} \leq T_{\mathrm{rt}}$
5	**Expression**	$P(D_{\mathrm{out}} \leq T_{\mathrm{rt}} \mid E)$, averaged over φ, the weak cell and D_{out}
6	**Model**	As above, with the distribution of an outage's length
7	**Check**	How often an outage like the next is carried to its end
	Class	Predictive probability

$$\boxed{P(D_{\mathrm{out}} \leq T_{\mathrm{rt}} \mid E)}$$

unit's next year, it is an outcome, the predictive probability that unit A fails at least once in the coming year from a named start, averaged over the rate. Read as a question about the unit now, whether it is up at this moment, it is an estimate of a present state, as Lesson 1 sorted such questions. Read most literally, as a question about the team's confidence, it asks for no probability about the unit at all, but for how wide the posterior of its rate is, which is a description of what is known rather than a probability of an event. The analyst's reply names the four readings, asks which one is meant, and says that the first two can differ greatly on a unit whose rate is well known but whose year is still to come.

Key idea

A parameter has one value; an outcome is still to occur and depends on it. The probability of the outcome is the average, over the posterior of the parameter, of its probability at each value. A parameter's probability enters that average as a weight, never as the answer.

7.3 The competing formulations

The headline competing formulation reports a parameter's probability for an outcome's. Table 7.7 gives two pairs, one row each: on the two-hall plant, with E the 40 kW trip test and the rates at their stated values (each unit 2.5 failures per year in service with a 180 h repair, the feeder 3.69 with a 24 h repair, all independent); and on the battery string, on the fact sheet's illustrative inputs. The values come from small calculations with the stated models.

Table 7.7: A parameter question and an outcome question on the same evidence, for two systems. The second row rests on illustrative inputs.

Parameter question	Value [–]	Outcome question	Value [–]
The opening adequate, $P(G \geq G^* \mid E)$	0.82	A random hour served, $P(S(t) = 1 \mid E)$	0.971
The string's ride-through at least two hours, $P(T_{\mathrm{rt}} \geq 2\ \mathrm{h} \mid E)$	0.074	The next outage carried, $P(D_{\mathrm{out}} \leq T_{\mathrm{rt}} \mid E)$	0.834

On the plant, the parameter's probability is lower than the outcome's. An inadequate opening loses the service only while one unit is out, so even on the series structure the cooling is on most of the time: at the stated rates the availability is 0.988 on the parallel structure and 0.896 on the series one, and the predictive weighs them by 0.82 and 0.18. On the string the gap is far wider and in the same direction, for a different reason. A string that falls short of two hours still carries most outages, because most outages are short. An owner told "82 %" for the cooling has been told the probability of a property of the opening; a manager told "7 %" for the next outage has been told the probability of meeting a specification. Either error can be made in either direction, and the size of the error depends on how far the outcome is from the parameter, which no rule of thumb supplies.

A second competing formulation offers the plug-in for the predictive. The availability at the rounded unavailabilities, 0.9875, quoted as the probability that a random hour is served, where the three-year record leaves the rates uncertain and the predictive is 0.9865. On the string, the next outage carried by a sound string of middling health, 0.925, with the weak cell ignored, is 0.846, against the predictive 0.834. Both differences are small here, and the text says so. The error is in kind, and its size depends on how uncertain the parameter is and on how far the outcome's probability bends across that uncertainty. A plug-in that ignores a hypothesis of small prior weight but large effect, such as the weak cell, can be wrong by much more on another system.

The swap, line by line:

Run, "82 %" told for the cooling

1	**Sentence**	Given E and the stated rates: cooling delivered at a random hour?
2	**Target** ×	A parameter, G
	Repair	An outcome, $S(t)$ at a random t
5	**Expression** ×	$P(G \geq G^* \mid E)$: 0.82
	Repair	$P(S(t) = 1 \mid E)$, averaged over G: 0.971
7	**Check**	Knowing G exactly leaves the hour's service uncertain

$$\boxed{\int P(S(t) = 1 \mid G,\ \text{rates})\, p(G \mid E)\, \mathrm{d}G}$$

A third reads a band on a parameter as a band on the outcome's probability, as in "the probability that a random hour is served is between 0.977 and 0.993". The band in Table 7.4 describes what is known of the availability, a parameter; the probability that a random hour is served is one number. A fourth reads a hypothesis's probability as the probability of the outcome it favors: "$P(H_{\text{AND}} \mid E) = 0.784$, so the plant keeps its cooling 78 % of the time". The predictive of equation (7.3) uses that probability as a weight.

These are the forms of the recurring class this lesson sharpens, *parameter against outcome*: a probability about a quantity with one value reported as the probability of an outcome that depends on it, or the reverse. Its four forms are a parameter's posterior probability reported for an outcome's predictive probability; a plug-in at a point value reported as the predictive; a band on a parameter, or on a function of parameters such as the availability, read as a band on the outcome's probability; and a hypothesis's probability read as the probability of the outcome it favors. The class enters at step 2, where the target is named, or at step 5, where the parameter appears as a condition instead of an average, and it is caught at step 7 by asking whether knowing the parameter exactly would leave the answer uncertain.

7.4 One system, many questions

Table 7.8 gives nine questions about the battery string, each with its expression in the second column and, in the third, whether it concerns a parameter or an outcome. It holds no values.

The string's few parameters stand behind every outcome in the table, and every outcome's probability averages over them. Two rows deserve a note. The fifth is an outcome whose probability equals a parameter question, for the reason given above. The seventh reads one outcome's probability as a long-run share of a year's outages, which holds only if the string's health stays fixed across the year: since every outage of the year meets the same string, its outcomes are tied by a shared unknown, Lesson 6's shared cause, and a year in which the string turns out weak is a year in which many outages go uncarried together.

7.5 The exercises

The solutions are in Appendix A, and the facts of the three systems in Appendix D. Where an exercise needs evidence on the opening, it is E, the 40 kW trip test.

Exercise 7.1 (*formulate*). The reliability engineer says: "Unit A's record shows seven failures in three years of service. I want to know whether its failure rate is below the handbook's 2.5 a year, and the chance that it fails more than twice next year." Write both expressions, name the target of each, and state what the second averages over.

Exercise 7.2 (*formulate*). The battery vendor suspects that one cell of the string is weak.

Table 7.8: Nine questions about the battery string, sorted as parameter or outcome.

Ordinary-language question	Mathematical formulation	Parameter or outcome
Is the string's state of health above 0.9?	$P(\varphi > 0.9 \mid E)$	A parameter
Does the string have a weak cell?	$P(H_{\text{weak}} \mid E)$	A parameter, discrete
Is the ride-through at least 1.5 h?	$P(T_{\text{rt}} \geq 1.5 \text{ h} \mid E)$	A parameter, a function of two
Will the string carry the next grid outage?	$P(D_{\text{out}} \leq T_{\text{rt}} \mid E)$	An outcome
Will a planned 90-minute switch-off of the grid supply be carried?	$P(T_{\text{rt}} \geq 1.5 \text{ h} \mid \text{do(grid off for 1.5 h)}, E)$	An outcome, settled by the parameters
How many outages next year will the string fail to carry?	The distribution of that count given steady operation at $t = 0$ and E	An outcome
What share of next year's outages will it carry?	$P(D_{\text{out}} \leq T_{\text{rt}} \mid E)$, read as a share	An outcome, read as a frequency
How long can the string carry 20 kW, given what is known?	The posterior of T_{rt} given E	A parameter, as a whole distribution
How warm will the pack be at the end of the next outage?	The predictive of the pack's temperature at the outage's end	An outcome

The site manager asks whether it is, and whether it matters for next week's planned 90-minute switch-off of the grid supply for maintenance. Write the expressions.

Exercise 7.3 (*formulate*). An insurer's renewal form for the critical load has a line that reads "Reliability of the backup battery string: _____". Write what should go on the line.

Exercise 7.4 (*formulate*). On the electrical network the meters do not add up. The owner asks: "Is the UPS degraded, and if it is, how much more will next month's utility energy be?" Write the expressions, and state what the second averages over.

Exercise 7.5 (*reverse*). Write one plain sentence for

$$\int P\big(N_A(\tau) \geq 1 \mid \lambda_A\big)\, p(\lambda_A \mid \text{the record})\, \mathrm{d}\lambda_A,$$

with τ one calendar year and unit A up at $t = 0$.

Exercise 7.6 (*reverse*). Write one plain sentence for $\sum_i P\big(S(t) = 1 \mid H_i, \text{ rates}\big)\, P(H_i \mid E)$, with t a random hour.

Exercise 7.7 (*reverse*). Write one plain sentence for $P(\varphi \geq 0.95 \mid E)$, and name one outcome whose probability depends on it.

Exercise 7.8 (*critique and repair*). A capacity test gives the string's state of health as 0.93 with a standard deviation of 0.03. The engineer writes: "So the probability of carrying the next outage is between 0.84 and 0.86." State what the range describes, and repair the statement.

Exercise 7.9 (*critique and repair*). On an hour when the network is healthy, its meters favor the nominal explanation with probability 0.99. The report concludes: "So the UPS will run at 96 % efficiency next month with probability 0.99." State what the 0.99 is, and repair the conclusion.

Exercise 7.10 (*critique and repair*). After the plant's record, the posterior mean of unit B's failure rate is about 2.4 a year in service. An analyst writes: "The probability that unit B fails at least once next year is $1 - e^{-2.4} = 0.91$." State what the analyst computed, and repair it.

Exercise 7.11 (*rapid*). For each question, say whether it is about a parameter or an outcome, and write its expression.

(a) Is unit B's mean repair time longer than a week?

(b) Will unit B's next repair take longer than a week?

(c) Is the UPS's efficiency below 0.94?

(d) Will the utility input exceed 1250 kW at the next reading?

(e) Which drawing of the opening is right?

(f) How many feeder interruptions will the network have next year?

(g) Is the PDU meter's offset larger than 10 kW?

(h) Will the next utility interruption be covered by the UPS?

(i) Is the opening's conductance above 9000 W/K?

(j) Will hall B be within its limit the next time unit B trips, with hall B at 50 kW?

7.6 What to retain

- A parameter has one value; an outcome is still to occur and depends on it. Ask whether knowing the parameter exactly would leave the answer uncertain.
- The probability of an outcome averages its probability at each value of the parameter over the posterior: a predictive, not a plug-in.
- A hypothesis is a discrete unknown. $P(H_{\text{AND}} \mid E)$ is a posterior over two values, and its predictive is a sum over them.
- On the 40 kW trip test, the opening is adequate with probability 0.82 and a random hour is served with probability 0.97 at the stated rates; on the battery string's illustrative inputs, the two-hour specification is met with probability 0.07 and the next outage carried with 0.83.
- A band belongs to a parameter. A predictive probability is one number.

Lesson 8

Questions About Future Performance

8.1 The engineering question

At ten past two in the morning the utility feeder that supplies a facility's electrical network is interrupted, and the UPS picks up the IT load from its battery. Two people call the duty engineer within a minute. The data-center manager asks whether the IT racks will have power at twenty to three, when a batch job must start. The customer's representative asks whether the racks will keep their power, without a break, until twenty to three, because a break of any length restarts every machine. The two questions sound alike, and they are not. The first is about one moment; the second is about every moment until then. The battery lasts ten minutes, and the feeder may come back at any time, so the racks can lose their power and have it again by twenty to three. On this network the two probabilities differ by more than a factor of two, as the competing formulations below show.

Two more questions are heard that night. In the battery room of another facility, the battery string carrying a critical load will reach its floor at some time if the grid stays out, and the moment at which it does is itself the uncertain quantity. At the two-hall plant, the month's hours without cooling climb toward the allowance of an illustrative service contract, and whether they will have reached it by the end of the month is a question about a total that only grows.

Every question about future performance refers to time in one of a small number of ways. It can ask about a moment chosen at random, about a stated moment, about every moment of an interval, or about whether something has happened by a stated moment, which is the same as asking when it first happens. Lesson 1 met two of these and found that they can differ enormously. This lesson names all four, writes each from its start, and says when two of them coincide.

8.2 The formulation

The work of this lesson sits in step 4, in the quantifier over time that the event carries, and in step 3, where the start or the present state stands behind the bar. Every horizon names its start, and "throughout" and "at some time" are an intersection and a union over times. The worked system is the electrical network, whose fact sheet is in Appendix D; every number about it rests on the sheet's illustrative inputs, and the text says so with each.

8.2.1 Four forms of a question about time

Table 8.1 gives the four forms on one quantity, the supply of the network's IT load, $S_{\mathrm{IT}}(t)$, which is one while the load is supplied. Each row gives the form, a question in that form, its expression, and the start or present state it needs.

Table 8.1: Four forms of a question about time, on the supply of the network's IT load.

Form	Question	Expression	Start or present state
At a random time	Is the IT load supplied at an hour chosen at random?	$P(S_{\mathrm{IT}}(t) = 1)$, t random	steady operation
At a stated time	Is it supplied at t_1?	$P(S_{\mathrm{IT}}(t_1) = 1 \mid \text{the state at } t_0,\ E)$	a named state at t_0
Throughout	Is it supplied at every time from t_0 to t_1?	$P(S_{\mathrm{IT}}(t) = 1 \text{ for every } t \in [t_0, t_1] \mid$ the state at $t_0,\ E)$	a named state at t_0
By a time	Has it been lost by t_1?	$P(T \le t_1 \mid \text{the state at } t_0,\ E)$, with T the first time after t_0 at which $S_{\mathrm{IT}} = 0$	a named state at t_0

Lesson 1 met the first form and the third, a random hour against a year without interruption. The last two rows are one event read two ways: the load has been lost by t_1 exactly when it was not supplied throughout $[t_0, t_1]$, so "by a time" is the complement of "throughout". Two forms are therefore new beside Lesson 1's, the stated time and the time by which something has happened, and the second of them is the complement of a form Lesson 1 met. What is new is the object in the last row, the time T at which something first happens.

The four forms also ask different things of the model at step 6, and Table 8.2 says what each needs, one form a row. The second column names what the model must supply, the third whether the start matters.

Table 8.2: What each form of a question about time needs from the model.

Form	What the model must supply	Does the start matter
At a random time	Each component's long-run share of time up and down, and the structure	No, by construction: steady operation
At a stated time	How the state moves from the start to the stated time: rates of failure and repair, or of return	Yes
Throughout	All of that, and every path's history: whether the quantity left its set at any moment of the interval	Yes
By a time	The same as throughout, read as the distribution of the first-passage time	Yes

The second column grows down the table. The random-time form needs only shares of time, which is why it is the easiest to compute and the one most often offered; the path forms need the whole process. Lesson 1 found the same thing when its two readings of "reliable enough" needed different models, and this table generalizes it: the form of time a question takes decides how much of the system's behavior the model must contain.

8.2.2 The first-passage time

The first-passage time is the first time after a stated start at which a quantity that rises reaches a stated level,

$$T_\ell = \inf\{\, t \ge t_0 : X(t) \ge \ell \,\}, \tag{8.1}$$

where X is the rising quantity, ℓ the level and t_0 the start. For the IT load, X can be the number of times the load has been lost since t_0 and ℓ one. The first-passage time is a random quantity, and its distribution, $P(T_\ell \le \tau \mid \text{the start},\ E)$, answers every question of the form "by τ" at once, for every horizon τ. Three standard objects are read from it. The probability of

reaching the level by τ is the distribution at τ. The probability of not reaching it within τ is its complement, the survivor function, the reliability over a horizon. The expected first-passage time, for a failure the mean time to failure, is one number about the whole distribution.

That last number deserves a caution in advance. A mean time is not a horizon. A unit whose mean time to failure is longer than a season can still be likely to fail within the season, because the distribution of a first-passage time is spread, and a comparison of a horizon with a mean compares a question about one horizon with a summary of all of them.

8.2.3 The rising path

The book's form for a question about crossing a level is to write it on a path that only rises: a rising path that has reached the level. A count of events rises. The cumulative hours without service since a start,

$$L(t) = \int_{t_0}^{t} \big(1 - S(s)\big)\, \mathrm{d}s, \tag{8.2}$$

rises. The energy drawn from a battery during an outage rises. And any quantity whatever can be turned into a rising path by taking its running maximum,

$$M(t) = \max_{s \in [t_0,\, t]} X(s), \tag{8.3}$$

the largest value it has taken since the start. For a path that only rises, "at or above the level at τ" and "has reached the level by τ" are the same event, because a rising path that has reached a level stays there: $L(\tau) \geq 3.65$ h is the event $T_{3.65\,\mathrm{h}} \leq \tau$, since equation (8.1) counts the level as reached when the path is at it. For a quantity that can fall again, they are not. Hall B's temperature, or the battery pack's, can rise past a limit and fall back, so its value at τ says nothing about whether the limit was crossed, and the crossing is written on its running maximum, $M(\tau) \geq \ell$.

Figure 8.1 draws the two cases side by side, with time on each horizontal axis and a dashed level across each panel. On the left, a rising staircase, a count or a cumulative total, crosses the level once and stays above it. On the right, a path rises past the level and falls back below it before the horizon, marked by the vertical line; its running maximum, drawn as the dashed staircase above it, rises with it and then stays where it was.

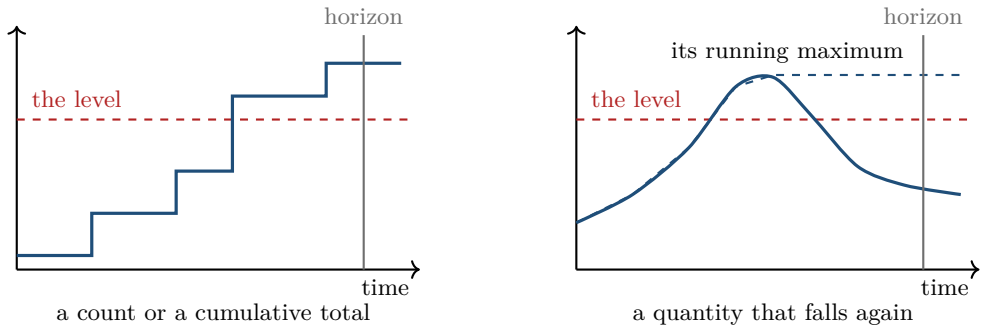


Figure 8.1: A rising path and a path that recovers, each against a level. The rising path is above the level at the horizon exactly when it has reached it; the recovering path is below the level at the horizon though it crossed it, and only its running maximum records the crossing. The shapes are illustrative.

The rule this figure supports is the lesson's: a question with "by", "ever" or "at some time" is written as a rising path that has reached its level, and a question about a quantity that can recover is written on that quantity's running maximum, its count of crossings or its cumulative

total, never on its value at the horizon. The rule applies to Lesson 4's illustrative contract clause. Its event $\bar{S}_{\mathrm{ex}}(m) \geq 0.995$ says that the month's cumulative hours without service, with the hours of a feeder outage counted as served, are at most 3.65 h at the month's end; since that total only rises, the event is the same, up to an event of probability zero, as the first passage to 3.65 h falling after the end of the month.

8.2.4 Starts and present states

The start of a horizon stands behind the bar. The forms at a stated time and throughout make the start quantitative, because the same question asked from different starts has different answers. Table 8.3 asks one question, whether the IT load is supplied, from two starts: steady operation, with nothing known of the present; and the moment a feeder interruption begins, with the UPS up, asked about five and thirty minutes later, with the interruption's later course unknown. The values rest on the fact sheet's illustrative inputs.

Table 8.3: One question, whether the IT load is supplied, from two starts, on the network's illustrative inputs.

Start, behind the bar	$P(S_{\mathrm{IT}} = 1)$ [–]
Steady operation, a random hour	0.99976
A feeder interruption begins, the UPS up; asked about five minutes later	0.98
The same start; asked about thirty minutes later	0.39

The first row answers a question asked when nothing is known of the present; the others answer questions asked at the start of an event, which condition on the event's present state then. A question asked thirty minutes into an interruption that is known to be still on is a third start, and on these inputs its answer is zero, since the battery's ten minutes of autonomy are spent. Quoting the first for either of the others gives the right answer to a question nobody asked, and the table shows how far it can be from the one asked. A question about the present state itself, such as whether the IT load is supplied now, is an estimate and not a prediction; it is answered from what is known now, not from a model of what will happen.

8.2.5 Counts over a horizon

A count over a horizon is a rising path, so every question about a count is a first-passage question. "At least one interruption by τ" is the first passage of the count to one, $T_1 \leq \tau$. "At most one interruption in the year" is the event that the count has not reached two by the year's end, $T_2 > \tau$, the second interruption's time falling after the horizon. "At least k failures" is the k-th failure's time at or before the horizon. Written this way, a count's event and a time's event are one, and either may be used; the choice is made by which is easier to state for the question asked.

Every count also carries a time base: a count over a calendar year from a named start differs from a count over a year of a component's own service when the component spends time out. On the network's illustrative inputs, from steady operation at the start of a calendar year, the IT load's interruptions arrive at 1.88 a year, and the probability of at most one interruption in the year is 0.44, against 0.15 for none. The two numbers are two values of the same distribution, the second interruption's time and the first's, read at the same horizon.

8.2.6 Horizons that are themselves random

Some horizons are not fixed. Lesson 1's operator asked whether hall B would be all right until the repair crew finished, and the horizon was the repair's length D, a random quantity, and the event is an intersection over $[t_0, t_0 + D]$. A random horizon adds one more unknown to step 6: the distribution of its length, and whether that length is tied to the path. The battery's race of the next subsection is the same kind of question, with the outage's length as the horizon, and so is "the UPS's battery lasts the whole interruption", whose horizon is the interruption's length. When the horizon is random, "throughout" is still the complement of "by", but both are read at a horizon that must itself be averaged over. A fixed horizon of the same mean answers a different question: "the next 180 hours" and "until unit B's repair ends" have the same expected length and different probabilities, because the repair's length is spread: averaged over a spread of lengths, the chance of surviving a repair at a constant failure rate is higher than the chance of surviving the mean length, since short repairs gain more than long ones lose. At a constant rate the direction holds for any spread, and the gap grows with it; a model in which the failure rate rises during a long repair can move it the other way.

8.2.7 Two first passages, and the parameters

Some questions compare two first passages. The battery string carries an outage to its end if the grid returns before the string reaches its floor, which is the event $D_{\text{out}} \le T_{\text{rt}}$: a race between the outage's end and the string's first passage to its floor. It is an outcome and a race. The ride-through T_{rt} is itself a first-passage time, the first time at which the energy drawn from the string, a rising path, reaches the usable energy. With the string's health uncertain, the race's probability is also an average over the parameters, Lesson 7's predictive:

$$P(D_{\text{out}} \le T_{\text{rt}} \mid E) = \sum_{i \,\in\, \{\text{sound},\, \text{weak}\}} P(H_i \mid E) \int P(D_{\text{out}} \le T_{\text{rt}} \mid \varphi,\ H_i)\, p(\varphi \mid H_i,\ E)\, \mathrm{d}\varphi. \quad (8.4)$$

Equation (8.4) is written and not evaluated. It shows the forms of this part of the book meeting in one expression: a race of two first passages, inside an average over a continuous parameter, inside a sum over a discrete one.

8.2.8 The interruption in the night

The control room's two questions, and a third about when the load is first lost, are written side by side in Table 8.4, one column each, with the fields that differ written out.

Figure 8.2 shows why the sheet's first two columns part. Time runs from the interruption at the left, with the battery's ten minutes marked and two vertical lines at five and thirty minutes. Each of the three bars is one way the night can go, shaded where the IT load is supplied: the feeder returns within the battery's ten minutes; the feeder returns after the battery has run out; the UPS fails to pick up the load and the feeder returns later still.

Read along the line at thirty minutes, two of the three paths are supplied, so they count for "at 02:40". Read along the whole interval, only the first path is supplied throughout, so only it counts for "throughout until 02:40". The second path is the one the two forms disagree about: the load was lost and has come back. That path exists because the network's quantity recovers when the feeder returns. The figure also shows why the second column of the sheet stops depending on the horizon once it passes the battery's ten minutes: after that, the only paths still supplied throughout are those on which the feeder came back within the ten minutes, and later moments add none.

Table 8.4: The seven steps on three questions asked after the feeder is interrupted at 02:10, on the network's illustrative inputs.

Step	At 02:40	Throughout until 02:40	When first lost
Intention	Predict	Predict	Predict
1. Sentence	With the feeder interrupted at $t_0 = 02{:}10$ and the UPS up, what is the probability that the IT load is supplied at 02:40?	With the same start, what is the probability that it is supplied at every moment from 02:10 to 02:40?	With the same start, what is the distribution of the first time after 02:10 at which the IT load is without supply?
2. Target	A state at a stated time	A path over an interval	A random time
3. Known	The feeder out at t_0 and the UPS up, a present state; the structure; the illustrative inputs	The same	The same
4. Event	$S_{\mathrm{IT}}(t_0 + 30\ \mathrm{min}) = 1$	$S_{\mathrm{IT}}(t) = 1$ for every $t \in [t_0, t_0 + 30\ \mathrm{min}]$	$T \le t_0 + s$, for each s
5. Expression	$P(S_{\mathrm{IT}}(t_0 + 30\ \mathrm{min}) = 1 \mid F_{\mathrm{sup}}$ at t_0, UPS up)	$P(S_{\mathrm{IT}}(t) = 1$ for every $t \in [t_0, t_0 + 30\ \mathrm{min}] \mid F_{\mathrm{sup}}$ at t_0, UPS up)	$P(T \le t_0 + s \mid F_{\mathrm{sup}}$ at t_0, UPS up), a function of s
6. Model	The UPS's pickup per demand; its autonomy; the distribution of the interruption's length; the feeder's return restores the supply	The same; a return after the battery runs out does not undo the loss	The same
7. Read back	Whether the power is on when the job starts	Whether the machines ran without a break	When the first break comes
Problem class	Predictive probability at a stated time	Predictive probability over a horizon	The distribution of a first-passage time

One start, three forms, column by column:

Run, the three questions after the interruption at 02:10

1	**Sentence**	F_{sup} at $t_0 = 02{:}10$, UPS up: (a) on at 02:40? (b) on throughout? (c) first lost?
	Intention	Predict
2	**Target**	(a) a state at a stated time; (b) a path over an interval; (c) a random time
	Declare	$S_{\mathrm{IT}}(t)$ supplied; F_{sup} supply lost; $t_{\mathrm{aut}} = 10$ min (illustrative); T first loss; s since t_0
3	**Known**	F_{sup} at t_0, UPS up, a present state; the structure; illustrative inputs
	Random here	The UPS's pickup; the interruption's length; through them S_{IT} and T
4	**Event**	(a) $S_{\mathrm{IT}}(t_0 + 30\ \mathrm{min}) = 1$; (b) $S_{\mathrm{IT}}(t) = 1$ for every $t \in [t_0, t_0 + 30\ \mathrm{min}]$; (c) $T \le t_0 + s$
5	**Expression**	The event's probability, given F_{sup} at t_0 and the UPS up, in each form
6	**Model**	Pickup; autonomy; interruption's length; (b) a late return does not undo the loss
7	**Check**	(a) power on when the job starts; (b) no break; (c) when the first break comes
	Class	(a) at a stated time; (b) over a horizon; (c) distribution of a first-passage time

(a) $P(S_{\mathrm{IT}}(t_0 + 30\ \mathrm{min}) = 1 \mid F_{\mathrm{sup}}$ at t_0, UPS up)
(b) $P(S_{\mathrm{IT}}(t) = 1$ for every $t \in [t_0, t_0 + 30\ \mathrm{min}] \mid F_{\mathrm{sup}}$ at t_0, UPS up)
(c) $P(T \le t_0 + s \mid F_{\mathrm{sup}}$ at t_0, UPS up), $\quad s \ge 0$

8.2.9 The battery string through the next day

The site manager of the battery room asks: if the grid fails at 14:00, will the pack stay below its limit of 318 K, and will the critical load be carried until the grid returns? The pack's temperature $\vartheta_{\mathrm{pk}}(t)$ rises while the string discharges and falls once the grid returns and the string rests. Three events about it are therefore different: below the limit at the outage's end,

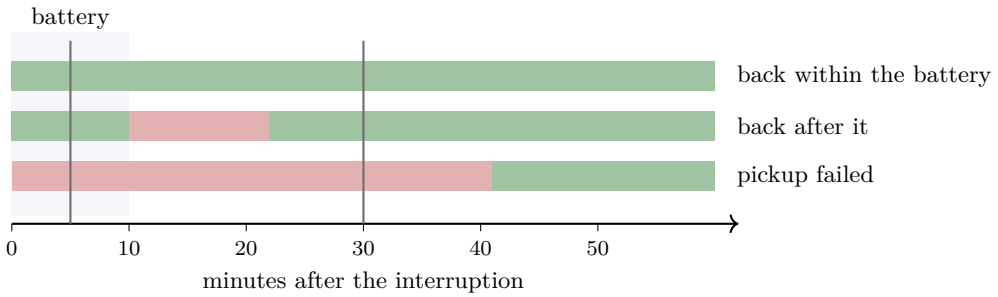


Figure 8.2: Three ways the night can go after the feeder is interrupted: green where the IT load is supplied, red where it is not. At thirty minutes the first two paths are supplied, but only the first was supplied throughout. The times are illustrative.

below the limit at 14:00 the next day, and below the limit at every moment in between. The third is written on the running maximum,

$$P\big(M(t_0 + 24\ \mathrm{h}) < 318\,\mathrm{K} \mid \text{the grid lost at } t_0,\ E\big), \qquad M(t) = \max_{s \in [t_0,\, t]} \vartheta_{\mathrm{pk}}(s). \tag{8.5}$$

During the discharge both the state of charge and the pack's temperature are monotone on the stated model: the charge only falls and the temperature only rises. Within the discharge the value at its end and the crossing therefore coincide; across the whole day they do not, because the temperature falls again once the discharge stops, at the outage's end or at the string's floor, whichever comes first. The load's question is the race of equation (8.4) with its start behind the bar. The source's simulation gives a peak of 303.8 K in an outage of 1.5 h and 307.2 K in one of 3 h; the second run was let discharge past the string's floor, and a string that stops at its floor would peak near 305 K. Either way, on these inputs the energy, not the temperature, is the limit that binds; the formulation of equation (8.5) does not depend on that.

Table 8.5 writes the site manager's two questions as sheets, one column each: the pack below its limit at every moment of the day after the grid fails, and the load carried until the grid returns. The fields that both share, the start and the evidence, are written once in the third row.

The sheet's sixth row shows that the two questions share a model even though their events differ: the earlier of the outage's end and the string's floor ends the heating in the first, and the outage's length ends the race in the second. A formulation that answered them with two separate models, one with the outage's length and one without, would let a long outage carry the load in one and fail to heat the pack in the other.

8.2.10 A question about time that is not ready

Before the summer, the owner of the two-hall plant asks whether the cooling will hold through the summer. The sentence names a period and asks about it without saying how the period is quantified, so it fails at the first step, and each of the four forms gives a reading. Read at a random time, it is the share of summer hours in which the cooling is delivered, an availability over the season. Read throughout, it is the probability that the cooling is delivered at every hour of the summer, from a named start such as steady operation on its first day: a clean summer. Read by a time, it is the probability that the summer's cumulative hours without cooling stay below some allowance by its end, which needs an allowance the owner has not given. Read as a first-passage time, it is the date of the first interruption, which the owner may want if a maintenance window must be planned around it. On the stated rates the first two readings differ as they do over a year, so the choice is not a matter of wording. The analyst's

Table 8.5: The seven steps on the battery room's two questions after the grid fails at t_0.

Step	The pack below its limit	The load carried
Intention	Predict	Predict
1. Sentence	With the grid lost at t_0, what is the probability that the pack stays below 318 K at every moment of the next 24 hours?	With the grid lost at t_0, what is the probability that the string carries the load until the grid returns?
2. Target	A path over a fixed horizon	A race between two first passages
3. Known	The grid lost at t_0, a present state; the string's state of charge at t_0; the room's temperature; E: the fact sheet's stated values and its illustrative inputs	
4. Event	$M(t_0 + 24\text{ h}) < 318\,\text{K}$, on the pack's temperature	$D_{\text{out}} \le T_{\text{rt}}$
5. Expression	Equation (8.5)	Equation (8.4), with the start behind the bar
6. Model	The pack's heat balance; the heating ends at the earlier of the outage's end and the string's floor	The energy balance of the string, its parameters, and the outage's length
7. Read back	Whether the pack ever reached its limit in the day. Silent on how close it came	Whether the load ever lost its supply. Silent on how much margin was left at the grid's return
Problem class	Predictive probability over a horizon, on a running maximum	Predictive probability of a race

And as runs:

Run, the pack below its limit

1	**Sentence**	Grid lost at t_0: probability the pack stays below 318 K throughout the next 24 h?
	Intention	Predict
2	**Target**	A path over a fixed horizon
	Declare	$\vartheta_{\text{pk}}(t)$ pack temperature [K]; $M(t)$ running maximum [K]; E fact sheet (illustrative)
3	**Known**	Grid lost at t_0; state of charge at t_0; room temperature; E
	Random here	The outage's length, φ, the weak cell, through ϑ_{pk}
4	**Event**	$M(t_0 + 24\text{ h}) < 318\,\text{K}$
5	**Expression**	Equation (8.5)
6	**Model**	Heat balance; heating ends at the earlier of the outage's end and the string's floor
7	**Check**	The pack ever at its limit in the day; silent on how close it came
	Class	Predictive probability over a horizon, on a running maximum

$$\boxed{P\big(M(t_0 + 24\text{ h}) < 318\,\text{K} \mid \text{the grid lost at } t_0,\ E\big)}$$

Run, the load carried until the grid returns

1	**Sentence**	Grid lost at t_0: probability the string carries the load until the grid returns?
	Intention	Predict
2	**Target**	A race between two first passages
	Declare	D_{out} the outage's length [h]; T_{rt} the ride-through [h], itself a first passage
3	**Known**	The same
	Random here	D_{out}; φ and the weak cell, averaged over
4	**Event**	$D_{\text{out}} \le T_{\text{rt}}$
5	**Expression**	Equation (8.4), with the start behind the bar
6	**Model**	The string's energy balance and parameters; the outage's length
7	**Check**	The load ever lost its supply; silent on the margin at the grid's return
	Class	Predictive probability of a race

$$\boxed{P(D_{\text{out}} \le T_{\text{rt}} \mid \text{the grid lost at } t_0,\ E)}$$

reply names the four readings, asks which one matters to the owner, and asks for the allowance if the third is meant.

Key idea

A question about time takes one of four forms: at a random time, at a stated time, throughout, or by a time. Each needs its start or present state behind the bar. Write a crossing on a path that only rises; on such a path "at" and "by" coincide, and on a path that recovers they do not.

8.3 The competing formulations

The main competing formulation answers a question about a path with an expression about a moment. Table 8.6 gives the three forms after the interruption at 02:10, at two times, one form a row, on the network's illustrative inputs: the feeder's interruptions exponential in length with a mean of one hour, the UPS failing to pick up the load with probability 0.02 per demand, and a battery of ten minutes. The values come from a small calculation with that model.

Table 8.6: Three forms of one question after the feeder is interrupted at t_0 with the UPS up, at two times, on illustrative inputs.

Expression	5 min after [–]	30 min after [–]
$P(S_{\mathrm{IT}}(t_0+s)=1 \mid F_{\mathrm{sup}} \text{ at } t_0,\ \text{UPS up})$, at a stated time	0.982	0.393
$P(S_{\mathrm{IT}}(t)=1 \text{ throughout } [t_0,t_0+s] \mid F_{\mathrm{sup}} \text{ at } t_0,\ \text{UPS up})$	0.980	0.150
$P(T \le t_0+s \mid F_{\mathrm{sup}} \text{ at } t_0,\ \text{UPS up})$, by a time	0.020	0.850

At five minutes the forms nearly agree, because the battery is still carrying the load. At thirty they part by a factor of 2.6. A manager told "39 %" in answer to "will the machines run without a break until 02:40" has been given the moment for the path. The second and third rows sum to one, because they are complementary events; the first row belongs to neither. This is the failure class of this lesson, in its first form.

The moment for the path, repaired:

Run, "39 %" told for a run without a break

1	**Sentence**	Feeder out at 02:10, UPS up: machines run without a break until 02:40?
4	**Event** ×	$S_{\mathrm{IT}}(t_0+30\ \mathrm{min})=1$, the state at the horizon
	Repair	$S_{\mathrm{IT}}(t)=1$ for every $t \in [t_0, t_0+30\ \mathrm{min}]$
5	**Expression** ×	$P(S_{\mathrm{IT}}(t_0+30\ \mathrm{min})=1 \mid F_{\mathrm{sup}} \text{ at } t_0,\ \text{UPS up})$: 0.393
	Repair	$P(S_{\mathrm{IT}}(t)=1 \text{ for every } t \in [t_0, t_0+30\ \mathrm{min}] \mid F_{\mathrm{sup}} \text{ at } t_0,\ \text{UPS up})$: 0.150
7	**Check**	A loss in the middle, later repaired, breaks it
		$P(S_{\mathrm{IT}}(t)=1 \text{ for every } t \in [t_0, t_0+30\ \mathrm{min}] \mid F_{\mathrm{sup}} \text{ at } t_0,\ \text{UPS up})$

The year shows the same gap on a longer scale. From steady operation, the IT load is supplied at a random hour with probability 0.99976, and the probability of a calendar year with no interruption of it is 0.15, with 1.88 interruptions expected a year; the IT load is without supply about 2.1 h a year, from losses of supply and from faults on its branch together, on the illustrative inputs. This is Lesson 1's contrast of 0.988 against 0.020, met again on a second system. It is the reason the class of the missing horizon recurs: a question with no stated time is usually read as the random-hour form, which is the most flattering.

A second competing formulation writes a crossing on the value at the horizon of a quantity that falls again. "Hall B was at 27°C at 06:00, so it never exceeded 30°C during the night" reads a running maximum from one sample of the path. A third compares a horizon with a mean time: "the mean time to failure is longer than the season, so the unit will last the season".

Exercise 8.8 sizes that error on the plant's stated rate. A fourth writes "throughout" where "at" was asked, the reverse of the first form: the batch job needs power at 02:40 and does not care whether the racks restarted at 02:25, so the first row of Table 8.6, not the second, is its answer.

One related error belongs to an earlier lesson. Quoting the random-hour value, 0.99976, at 02:15 in answer to "will there be power at 02:40" drops the present state from behind the bar. That is Lesson 1's rule on present states, not a form of this lesson's class.

These are the forms of this lesson's class, the *"at a time" read as "by a time"*, with the first passage missed: a question about one moment answered with an expression about a path, or the reverse. Its four forms are the state at the horizon for "by the horizon" or "throughout"; "throughout" for "at the horizon"; a crossing written on the value at the horizon of a quantity that can fall again, instead of on its running maximum, count or cumulative total; and a first-passage time's distribution summarized by its mean and compared with a horizon. It enters at step 4, at the time quantifier, and is caught at step 7 by asking whether a break in the middle of the interval would change the answer the speaker wants. A missing horizon remains Lesson 1's class.

8.4 One system, many questions

The two-hall plant returns as the table's system. Table 8.7 gives nine questions about it, each about time, with its form in the second column and its expression in the third.

Table 8.7: Nine questions about time on the two-hall plant.

Ordinary-language question	Form	Expression
What share of hours is the cooling on?	At a random time	$P(S(t) = 1 \mid E)$
Unit B tripped at 02:00; is the cooling on at 06:00?	At a stated time	$P(S(06{:}00) = 1 \mid \text{unit B out at } 02{:}00,\ E)$
Does the cooling hold until unit B's repair ends?	Throughout a random horizon	$P(S(t) = 1 \text{ for every } t \in [t_0, t_0 + D] \mid$ unit B out at $t_0,\ E)$
Will the plant have an interruption before the year ends?	By a time	$P(T \le \tau \mid \text{all up at } t = 0,\ E)$, T the first interruption
When will the month's hours without cooling reach the illustrative contract's allowance?	A first-passage time	The distribution of $T_{3.65\,\mathrm{h}}$ on $L_{\mathrm{ex}}(t)$ from the month's start, feeder outages counted as served
Will the month breach the clause?	By a time, on a rising path	$P(L_{\mathrm{ex}}(730\ \mathrm{h}) \ge 3.65\ \mathrm{h}) = P(T_{3.65\,\mathrm{h}} \le 730\ \mathrm{h})$
Will unit A fail during the summer?	By a time	$P(N_A(\tau) \ge 1 \mid$ unit A up at the summer's start), given the stated model
Did hall B exceed its limit at some time last night?	By a time, on a past path	The running maximum of the recorded difference; an estimate if the readings carry error
How long, on average, until the plant's next interruption?	Mean of a first-passage time	$\mathbb{E}[T \mid \text{steady operation at } t_0,\ E]$

The plant, its service and its events are the same in every row, and the form of time decides the expression. Two rows are one quantity read two ways, the fifth and the sixth, since the month's lost hours only rise; the eighth asks about a path already run, so it is an estimate of what happened rather than a prediction.

8.5 The exercises

The solutions are in Appendix A, and the facts of the three systems in Appendix D.

Exercise 8.1 (*formulate*). Unit B tripped at 02:00, and the repair crew will arrive at 06:00 and restore it at once. The operator asks for the probability that hall B has its cooling at 06:00, just before the crew restores the unit, and for the probability that it has its cooling throughout until then. Write both, and state whether they differ on this plant.

Exercise 8.2 (*formulate*). The grid fails at 14:00. The site manager asks: "Will the critical load have power at 15:30, and will the string reach its floor before the grid comes back?" Formulate the question.

Exercise 8.3 (*formulate*). A customer asks the network's operator: "How long will the UPS last?" Formulate the question.

Exercise 8.4 (*formulate*). An illustrative service clause reads: "The IT load shall be without power for no more than 30 minutes in total in any calendar year." Write the event on a rising path, the equivalent event on a first-passage time, and the expression.

Exercise 8.5 (*reverse*). Write one plain sentence for $P(N_A(\tau) \geq 1 \mid \text{unit A up at } t_0)$, given the stated model, with τ three calendar months, and write the same event with a first-passage time.

Exercise 8.6 (*reverse*). Write one plain sentence for

$$P\Big(\max_{t \in [t_0,\, t_0 + 6\ \mathrm{h}]} (\vartheta_B(t) - \vartheta_A(t)) > 8\,\mathrm{K} \mid \text{unit B out at } t_0,\ E\Big).$$

Exercise 8.7 (*reverse*). Write one plain sentence for $P(S_{\mathrm{IT}}(t_0 + 1\ \mathrm{h}) = 1 \mid F_{\mathrm{sup}} \text{ at } t_0,\ \text{UPS up})$.

Exercise 8.8 (*critique and repair*). A planner writes: "A cooling unit fails on average once every 4.8 months of service. The summer is three months, so the unit will almost certainly get through the summer without failing." State the error, and repair it.

Exercise 8.9 (*critique and repair*). After a grid outage of three hours, a log shows the battery pack at 301 K at 14:00 the following day, 24 hours after the outage began. The report says: "The pack stayed below its 318 K limit throughout." State what the log supports, and repair the conclusion.

Exercise 8.10 (*critique and repair*). During last night's interruption, which began at 02:10, the operator checked the IT racks at 02:15 and at 02:45 and found them powered both times. The shift report says: "The UPS carried the load throughout." State what the checks establish, and repair the conclusion.

Exercise 8.11 (*rapid*). Name the form of each question (at a random time, at a stated time, throughout, by a time, a first-passage time, or a present state), and write its expression.

(a) Will the IT racks have power at noon tomorrow?

(b) Will the string still be above its floor when the grid returns?

(c) What share of the year does hall B run on one unit?

(d) Will the feeder be interrupted before the month ends?

(e) Is the IT load powered right now?

(f) When will unit A next fail?

(g) Will the IT load's hours without supply this year have reached half an hour by the end of June?

(h) Will the UPS's battery last the whole of the next interruption?

(i) Did the battery pack reach 310 K at any moment during yesterday's outage?

(j) On average, how long until the IT load's first interruption next year?

8.6 What to retain

- A question about time takes one of four forms: at a random time, at a stated time, throughout, by a time. Each needs its start or present state behind the bar.
- The first-passage time is a random quantity. Its distribution answers every "by a time" question at once, and its mean answers none of them.
- Write a crossing on a path that only rises: a count, a cumulative total, a running maximum. On such a path "at" and "by" coincide; on a path that recovers they do not.
- Thirty minutes after a feeder interruption begins, on the network's illustrative inputs, the IT load has power with probability 0.39 and has had it without a break with probability 0.15.
- A question about the present is an estimate; a question about a moment still to come is a prediction.

Lesson 9

Complex Questions with Multiple Conditions

9.1 The engineering question

At the handover of a facility that holds all three of the book's systems, the owner's engineer reads a memorandum aloud. It was written for this lesson and is illustrative, but memoranda of its kind are common:

> Before the facility is accepted, three assurances are needed. The opening must carry a hall's full load if a unit trips, and hall B must not overheat at full load with a unit out. If unit B trips during the summer peak, hall B must not overheat either. The UPS must carry the IT load through a utility interruption, and the control-room battery must give its two hours. All of this is wanted with at least 95 % confidence over the first year.

The meeting hears three assurances, a confidence and a year, and nods. A careful reading finds more, and less. The first two sentences about hall B are one question asked twice in two vocabularies, one about the opening and one about a temperature. The third sounds like the second and is a different question, because the summer peak names a season and does not fix the load hall B will carry when the unit trips. "Through a utility interruption" does not say which interruption, the next one or every one of the first year. The confidence and the year are stated once, at the end, and they cannot attach to every clause in the same way: the opening's capacity does not change over a year, while the number of interruptions does. And "all of this" may mean each assurance separately or all of them together, which are different numbers.

Every earlier lesson took one question at a time, already separated from its neighbors. Real questions arrive in paragraphs, in which clauses overlap, repeat one another, share conditions and leave some unstated. Turning such a paragraph into mathematics is not a matter of writing one expression for it. It is a matter of finding how many questions it holds and how they relate.

This lesson takes a paragraph apart into its clauses, writes each clause's expression, sets the expressions against one another, and only then decides what the paragraph asks.

9.2 The formulation

The work of this lesson sits in step 1, which now produces several sentences from one paragraph, and in step 7, which now reads each expression back against the others as well as against its own clause. It uses every skill of this part of the book: limit states from Lesson 5 to see when two events are one, joints with their dependence from Lesson 6 for the conjunction, the parameter against the outcome from Lesson 7, and the forms of time from Lesson 8 to attach each clause's

horizon. The evidence on the two-hall plant is E, the 40 kW trip test, as throughout this part. The three systems in the memorandum's facility are separate systems with separate fact sheets, and nothing ties one to another; Lesson 25 couples them.

9.2.1 Taking a paragraph apart

A paragraph is taken apart in six moves.

1. List the clauses, one assertion or question each.
2. Sort each by kind and name its intention.
3. Attach to each its event, conditions, horizon and level, and mark the conditions the paragraph states once for several clauses.
4. Write each clause's expression.
5. Compare the expressions in pairs: one quantity, nested, overlapping, or unrelated.
6. Decide whether the paragraph asks for each expression separately or for their conjunction, and write the conjunction with its dependence.

The first four moves are Part I and this part applied clause by clause. The last two are new, and they are where a paragraph's mathematics is decided. Figure 9.1 draws the result of the six moves on the memorandum: its five clauses in a column on the left, the expressions they become on the right, and labeled lines between the expressions for the relations found at the fifth move.

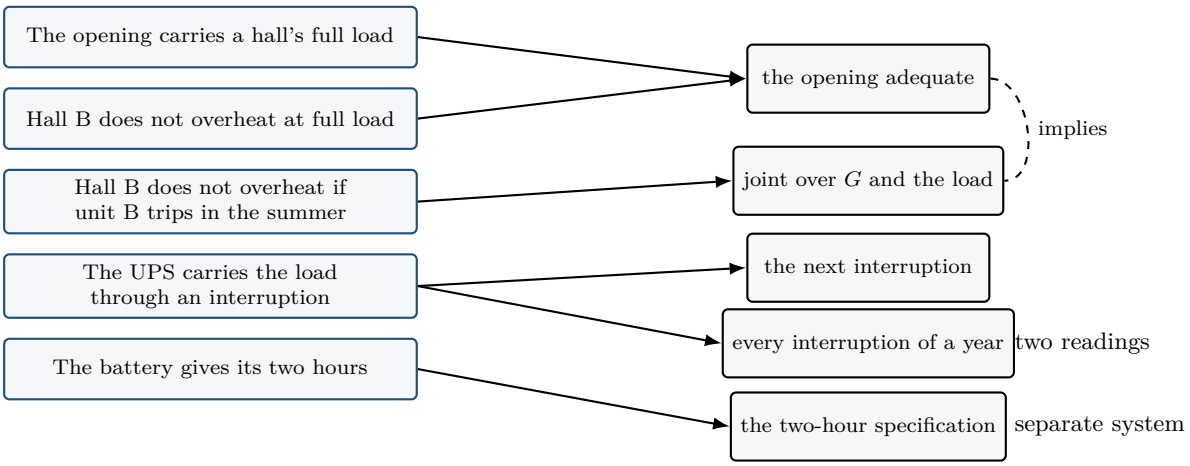


Figure 9.1: The memorandum's five clauses and the expressions they become. Two clauses become one expression, one clause becomes two readings, and the dashed line marks an implication between expressions.

The figure is not a one-to-one map, and that is its point. Two arrows meet at the first expression, two leave the fourth clause, and one line joins two expressions on the right. A formulation that wrote one expression per clause would have five, two of them the same quantity and one of them a choice made silently.

9.2.2 Two sentences, one quantity

Lesson 1's table held six sentences about the two-hall plant, and it made no claim about how its rows relate. Its second row asked whether the opening is inadequate and its sixth the probability of overheating at full load. Lesson 4 noted that the design reading of the 8 K limit is the

complement of the sixth row, and Lesson 5 wrote the limit state for any load. Here the relation is stated. The memorandum's first clause says "a hall" and "a unit"; the fact sheet states a limit for hall B above hall A only, so the clause is read for hall B with unit B out, and its other half, hall A with unit A out, cannot be written until a limit for hall A is stated.

With unit B out and the plant settled, hall B's rise above hall A is its load divided by the opening's conductance, by Lesson 5's limit state, equation (5.1). With the load fixed at the full 60 kW, the rise exceeds 8 K exactly when the conductance is below 60 kW divided by 8 K, which is the threshold G^* of equation (1.4). The two events are the same set of outcomes, so

$$P(\Delta T_{\text{full}} > 8\,\text{K} \mid \text{unit B out},\ Q_{\text{hall}},\ \text{settled},\ E) = P(G < G^* \mid E), \tag{9.1}$$

and on E, the 40 kW trip test, both sides are 0.18. The number was computed twice, once as an event on the conductance and once as an event on the settled rise, each on its own grid, and the two agree, as equation (9.1) says they must.

The equality holds because of five conditions, and dropping any one of them can part the two sides: the load fixed at the full 60 kW; unit B out, so that all of hall B's heat crosses the opening; the plant settled, so that the rise is the steady one; the opening's law, which makes the rise the load over the conductance; and no drift, so that the conductance in service is the one the trip test measured. The recognition rule the example illustrates is this: *two expressions name one quantity when their events, written on the same unknowns, are the same set under the conditions behind both bars.* Writing both events on the same unknowns is the test. Here the second sentence's event, written on the conductance, is the first sentence's event, and nothing else is needed to see it.

The equality as a run:

Run, hall B overheating at full load with unit B out

1	**Sentence**	Unit B out, 60 kW, settled, given E: probability hall B is over 8 K above hall A?
	Intention	Estimate
2	**Target**	An event on a settled state; through it G
	Declare	ΔT_{full} [K]; $Q_{\text{hall}} = 60\,\text{kW}$; G [W/K]; $G^* = 7500\,\text{W/K}$; E the 40 kW trip test
3	**Known**	Unit B out; Q_{hall}; settled; E
	Random here	G only
4	**Event**	$\Delta T_{\text{full}} > 8\,\text{K}$
5	**Expression**	$P(\Delta T_{\text{full}} > 8\,\text{K} \mid \text{unit B out},\ Q_{\text{hall}},\ \text{settled},\ E)$
	Expand	$\Delta T_{\text{full}} = Q_{\text{hall}}/G$: $P(Q_{\text{hall}}/G > 8\,\text{K} \mid E)$
6	**Model**	The opening's law; the readings' model; the prior on G; no drift (assumption)
	Reduce	$Q_{\text{hall}}/G > 8\,\text{K}$ exactly when $G < G^*$
7	**Check**	The complement of "the opening adequate": one quantity
	Class	Posterior exceedance

$$\boxed{P(G < G^* \mid E)}$$

9.2.3 A third sentence that looks the same

The memorandum's third clause is "If unit B trips during the summer peak, hall B must not overheat." The word "if" marks a condition, and the trip goes behind the bar. The summer peak names a period, not a load. At the moment of the trip hall B carries whatever it carries then, Q_B, which is uncertain, so the event $\{Q_B/G > \Delta T_{\text{allow}}\}$ is a joint event over the conductance and the load, and its probability averages the second sentence's probability over the load:

$$P(Q_B/G > \Delta T_{\text{allow}} \mid \text{unit B out},\ \text{settled},\ E) = \int P(G < Q_B/\Delta T_{\text{allow}} \mid E)\, p(Q_B)\, \mathrm{d}Q_B, \tag{9.2}$$

with $p(Q_B)$ the distribution of hall B's load at a summer trip, taken as independent of the conductance. This book has no record of the plant's summer loads, so the distribution is

illustrative: equally likely anywhere between 45 and 60 kW. On it, equation (9.2) is 0.047, against 0.18 at the full load.

The average is easier to read with the inner probability shown at a few loads. Table 9.1 gives it at four loads across the band, one a row, each the probability that hall B exceeds its limit at that load with unit B out, given the 40 kW trip test.

Table 9.1: The probability that hall B exceeds its limit at a stated load with unit B out, given the 40 kW trip test.

Hall B's load [kW]	$P(G < Q_B/\Delta T_{\text{allow}} \mid E)$ [–]
45	0.0003
50	0.008
55	0.056
60	0.18

The inner probability rises steeply across the band, from almost nothing at its low end to the full-load value at its top, so the average lies far below the top. A band reaching higher, or a load distribution weighted toward its top, would bring the third sentence's value closer to the second's; a band reaching lower would take it toward zero. The third sentence's answer depends on a distribution the memorandum does not state, and the analyst must ask for it or label it, as this lesson does.

The two sentences share every word a listener notices, hall B, overheat, a unit trips, and they differ in one condition that the second states and the third does not. The third also has a further reading, "unit B trips during the summer and hall B overheats", a joint over unit B's failure in the season and the overheating, which needs unit B's failure process over the summer, Lesson 8's "by a time". The memorandum's "if" does not ask for it, and the analyst says so.

Figure 9.2 draws all three sentences on one picture, with hall B's load on the horizontal axis and the conductance on the vertical. The sloping line through the origin is the limit, $G = Q_B/\Delta T_{\text{allow}}$; hall B overheats below it. The vertical line at 60 kW is the second sentence's condition, and the part of it below the slope is its failure set; the horizontal line at G^* meets the slope exactly there, which is the first sentence's threshold. The shaded band from 45 to 60 kW is the third sentence's range of loads, and its failure set is the part of the band below the slope.

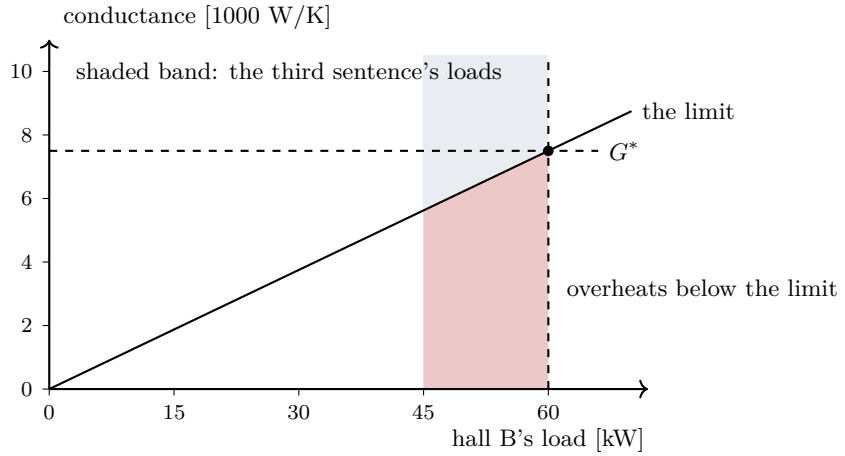


Figure 9.2: Three sentences on the plane of hall B's load and the conductance. At the full 60 kW the limit and the threshold G^* meet, so the first two sentences are one event; across the band of summer loads the failure set (red) is a region over both unknowns. The band is illustrative.

9.2.4 Four relations between expressions

The fifth move compares expressions in pairs, and a pair stands in one of four relations. Table 9.2 lists them, one a row, with what each means for the events, an instance from the memorandum or the plant, and what it means for a set of answers reported together.

Table 9.2: Four relations between the events of two expressions.

Relation	The events	Instance	For a set of answers
One quantity	The same set under the conditions of both	The opening adequate; hall B within 8 K at full load with unit B out	Report once; never multiply, never count as two confirmations
Nested	One contained in the other	The opening adequate implies hall B within its limit at any load below full	The joint is the smaller event; its probability is the smaller one
Overlapping	Each can hold without the other, and both can hold	Unit A failing next year; the service interrupted next year	The joint needs the model of both; neither implies it
Unrelated	On separate systems or independent unknowns	The opening adequate; the battery giving its two hours	The joint is a product, an assumption to state

Two of the rows are traps for a set of answers. The first: a quantity reported twice looks like two pieces of good news, or two requirements met, when it is one. The fourth: "unrelated" is the relation a formulation reaches for by default, and it is the one to state as an assumption. Across the memorandum's systems it holds by construction in this part of the book; within one system it seldom does.

9.2.5 Conditions and their scope

A condition stated once can govern every clause, only the last, or some of them. "With at least 95 % confidence over the first year" holds two conditions, a probability level and a horizon. Table 9.3 reads them clause by clause, one clause a row: its kind and intention, its expression, the horizon the paragraph gives it, and the level.

The fourth column carries this subsection's point. The year attaches naturally to the UPS clause, read as every interruption, and to nothing else: the opening's adequacy and the battery's specification are questions about parameters, and a parameter has no horizon unless it can change within it. Attaching the year to them adds nothing to their numbers and suggests to a reader that a path was considered. The level of the last column is the memorandum's, and it is ambiguous between each clause and the conjunction of all of them, which is the next subsection's question.

9.2.6 The conjunction

"All of this with 95 % confidence", read as the conjunction, asks for one joint over every clause,

$$P\Big(\bigcap_{k} \mathcal{A}_k \mid E\Big) \geq 0.95, \tag{9.3}$$

against each $P(\mathcal{A}_k \mid E) \geq 0.95$ separately, with $\mathcal{A}_k$ the event of the k-th expression. Writing the joint needs each pair's relation from the fifth move. The first two clauses are one event, so they enter equation (9.3) once. The third is implied by the first, since an opening that carries the full load carries any smaller one, so their joint is the first. The plant's clauses, the network's and the battery's are on separate systems, unrelated by construction, so their joint is a product, an assumption stated as such. The joint is never larger than the smallest of its members. On

Table 9.3: The memorandum's clauses with the scope of its shared conditions.

Clause	Kind; intention	Expression	Horizon	Level
The opening carries a hall's full load if a unit trips	Parameter; decide on an estimate	$P(G \geq G^* \mid E)$, for hall B with unit B out; hall A's limit is not stated	None: a property of the opening, unless it drifts	0.95
Hall B does not overheat at full load with a unit out	Parameter; decide on an estimate	$P(\Delta T_{\text{full}} \leq 8\,\text{K} \mid$ unit B out, Q_{hall}, settled, $E)$, the same event	None	0.95
Hall B does not overheat if unit B trips in the summer	An outcome on a condition; decide on a prediction	$P(Q_B/G \leq \Delta T_{\text{allow}} \mid$ unit B out, settled, $E)$, Q_B uncertain	The summer, through the load's distribution	0.95
The UPS carries the IT load through an interruption	Outcome; decide on a prediction	The next: $P(F_{\text{ups}}$ does not occur $\mid$ an interruption begins); every one of the year: the probability of no interruption of the IT load from a loss of supply in $[0, \tau]$, from steady operation at $t = 0$	The first year, if every interruption is meant	0.95
The battery gives its two hours	Parameter; decide on an estimate	$P(T_{\text{rt}} \geq 2\,\text{h} \mid E)$	The first year only if the string's health can fall within it, which its fact sheet does not state	0.95

these inputs every clause but the summer one, at 0.953, falls short of 0.95 when read separately, so the conjunction does too, and its value adds nothing the reader needs; the text leaves it unevaluated.

9.2.7 The memorandum

Table 9.4 is the set of formulation sheets for the memorandum, condensed to one row per expression after the comparison, with its value on this part's evidence and the relation found. The plant's values use E, the 40 kW trip test; the network's and the battery's rest on their fact sheets' illustrative inputs; the summer load is illustrative.

Five clauses became six expressions, two of which are one quantity, and one clause needed two readings. Read separately against the memorandum's 95 %, every clause but the summer one falls short on these inputs, so the conjunction does too. What the analyst brings back to the owner's engineer is the table, the two choices it rests on (which interruption the UPS clause means, and whether the level is for each clause or for all of them), and the one equality found, which removes a clause the engineer believed to be a separate assurance.

9.2.8 A short paragraph with nested conditions

The facilities manager says: "If unit B trips while the feeder is up, and if hall B is above 50 kW at the time, hall B must stay within its limit until the repair is finished." The sentence holds three conditions, two of which narrow the third. The trip and the feeder's state at the trip go behind the bar, as a present state. "Above 50 kW" is a condition on the uncertain load, so the load's distribution is restricted to the loads above 50 kW, up to the full 60 kW. "Until the repair is finished" is a horizon, Lesson 8's "throughout" with the random length D of the repair, so

Table 9.4: The memorandum's expressions after the comparison, with their values and relations.

Expression	Value [–]	Relation to the others
$P(G \geq G^* \mid E)$, hall B with unit B out	0.82	The same event as the next
$P(\Delta T_{\text{full}} \leq 8\,\text{K} \mid \text{unit B out},\ Q_{\text{hall}},\ \text{settled},\ E)$	0.82	The same event as the previous
$P(Q_B/G \leq \Delta T_{\text{allow}} \mid \text{unit B out},\ \text{settled},\ E)$, summer load illustrative	0.953	Implied by the first; joint over G and the load
The next utility interruption covered by the UPS, illustrative	0.150	Separate system
No interruption of the IT load from a loss of supply in the first year, illustrative	0.169	The reading "every interruption"; separate system
$P(T_{\text{rt}} \geq 2\ \text{h} \mid E)$, illustrative	0.074	Separate system

The whole memorandum as one run:

Run, the memorandum

1	**Sentence**	(a) opening; (b) hall B, full load; (c) hall B, summer trip; (d) UPS; (e) battery, 2 h
	Intention	Decide; estimates for (a), (b), (e); predictions for (c), (d)
2	**Target**	Six expressions on three systems, and each pair's relation
	Declare	G, G^*; ΔT_{full} [K]; ΔT_{allow}; Q_B summer load [kW] (illustrative); F_{ups}; T_{rt} [h]; τ first year; $\mathcal{A}_k$ event k
3	**Known**	E, 40 kW trip test; fact sheets; (b), (c) unit B out, settled; level 0.95, scope unstated
	Random here	G; Q_B in (c); interruptions in (d); φ, the weak cell in (e)
4	**Event**	(a) $G \geq G^*$, hall B; (b) $\Delta T_{\text{full}} \leq 8\,\text{K}$; (c) $Q_B/G \leq \Delta T_{\text{allow}}$; (d) no F_{ups}, next or all of $[0, \tau]$; (e) $T_{\text{rt}} \geq 2$ h
5	**Expression**	Each event's probability, Table 9.4
	Expand	(b), with $\Delta T_{\text{full}} = Q_{\text{hall}}/G$, is (a)
6	**Model**	Opening, load; interruptions, UPS; energy; systems independent (assumption)
	Reduce	(a), (b) one event; (c) implied by (a); across systems a product
7	**Check**	Six expressions from five clauses; 95 % for each or for all?
	Class	Posterior exceedances and predictive probabilities

$$P(\mathcal{A}_k \mid E) \geq 0.95 \text{ for each } k, \quad \text{or} \quad P\Big(\bigcap_k \mathcal{A}_k \mid E\Big) \geq 0.95$$

the event is a path, and on the plant's long repairs it brings in unit A's and the feeder's failures during the repair, as in Exercise 6.2. The expression is

$$P(Q_B/G \leq \Delta T_{\text{allow}} \text{ and } S(t) = 1 \text{ for every } t \in [t_0, t_0 + D] \mid \text{unit B out at } t_0,\ \text{feeder up at } t_0,\ Q_B > 50\,\text{kW},\ \text{settled},\ E). \quad (9.4)$$

Read back, equation (9.4) says what each condition removed: the loads at or below 50 kW, the trips that came with a feeder outage, and nothing about the time after the repair. It also says what the manager may have meant and did not: "while the feeder is up" holds at the trip and does not keep the feeder up during the repair, so the event's path includes the feeder's possible failure. If the manager meant the feeder up throughout, that is a fourth condition, behind the bar, and it answers a different question about a world in which the feeder holds.

The sentence as a run:

Run, the facilities manager's nested conditions	
1 Sentence	Unit B trips, feeder up, hall B above 50 kW: within its limit until the repair ends?
Intention	Predict
2 Target	A path over a random horizon
Declare	t_0 the trip; Q_B load at the trip [kW]; D the repair's length [h]; $S(t)$ service at t
3 Known	Unit B out at t_0; feeder up at t_0; $Q_B > 50\,\text{kW}$; settled; E
Random here	G; Q_B in $(50, 60]$ kW; D; unit A's and the feeder's failures
4 Event	$\{Q_B/G \le \Delta T_{\text{allow}}\} \cap \{S(t) = 1 \text{ for every } t \in [t_0, t_0 + D]\}$
5 Expression	Equation (9.4)
6 Model	The opening's law; load above 50 kW; repair's length; unit A's and feeder's failures
7 Check	Feeder up at the trip, not held up through the repair
Class	Predictive probability over a random horizon
	$P\big(Q_B/G \le \Delta T_{\text{allow}} \text{ and } S(t) = 1 \text{ for every } t \in [t_0, t_0 + D]$ $\quad \mid \text{unit B out at } t_0,\ \text{feeder up at } t_0,\ Q_B > 50\,\text{kW},\ \text{settled},\ E\big)$

9.2.9 A paragraph that is not ready

Before the summer, the owner says: "Make sure the cooling and its supply are good enough for the summer, and tell me what to fix first." The paragraph holds three clauses and none of them is ready. "The cooling is good enough" names a system and no event, horizon or level, the failure of "reliable enough". "Its supply is good enough" names a second system with the same gaps, and "its" is itself ambiguous, since in this part of the book the two-hall plant has its own feeder and the network is a separate facility. "What to fix first" is not a probability question at all: it asks to decide among actions, and a decision needs the consequences of each action and a way to weigh them, which Lesson 14 writes.

The six moves still apply, and they produce a useful reply. The first move lists the three clauses. The second sorts them: two predictions resting on estimates, and one decision. The third finds every part missing from the first two and notes that the summer supplies a horizon to both. The fourth writes, for each of the first two, the expressions its natural readings would give: the cooling delivered at every hour of the summer, or at a random summer hour, or hall B within its limit whenever unit B trips at a summer load; and the IT load's supply through the next interruption, or through all of the summer's. The fifth notes that the cooling readings share the opening's conductance and so are tied, while the supply readings are on a separate system. The sixth asks whether the owner wants each assurance or both at once, and sets the decision aside until its consequences are stated. The analyst returns with that list and three questions, which is a correct result for a paragraph that is not ready.

Key idea

A paragraph holds several questions. Write each clause's expression with its own conditions, horizon and level, then compare the expressions: one quantity, nested, overlapping or unrelated. Only then decide whether the paragraph asks for each or for all.

9.3 The competing formulations

The first competing formulation counts one quantity as two. The memorandum's first two clauses, each with probability 0.82, are reported as two independent assurances, so that "both hold" is 0.82×0.82. Table 9.5 gives the two readings, one a row, on E, the 40 kW trip test.

Multiplying makes the plant look worse than it is, by 0.15. The same duplicate can also make it look better: a list of assurances that counts the two clauses as two confirmations reports

Table 9.5: The memorandum's first two clauses, read as one event and as two independent events, on the 40 kW trip test.

Reading	Expression	Value [–]
One event, recognized	$P(G \geq G^* \mid E)$	0.82
Two events, multiplied	$P(G \geq G^* \mid E) \times P(\Delta T_{\text{full}} \leq 8\,\text{K} \mid$ unit B out, Q_{hall}, settled, $E)$	0.67

two findings where there is one. Either way the error is in step 1, where the paragraph was split into clauses, and in step 7, where the expressions were not read back against one another. The duplicate, marked:

Run, one quantity counted as two

1	**Sentence**	Opening carries hall B's full load, and hall B does not overheat at full load?
5	**Expression** ×	$P(G \geq G^* \mid E) \times P(\Delta T_{\text{full}} \leq 8\,\text{K} \mid$ unit B out, Q_{hall}, settled, $E)$: 0.67
	Repair	$P(G \geq G^* \mid E)$: 0.82
	Expand	$\Delta T_{\text{full}} = Q_{\text{hall}}/G$, so $\{\Delta T_{\text{full}} \leq 8\,\text{K}\} = \{G \geq G^*\}$
7	**Check**	On the same unknowns, one set; report it once

$$\boxed{P(G \geq G^* \mid E)}$$

The second answers the look-alike with the duplicate's expression. The full-load value, 0.18, reported as the probability that hall B overheats if unit B trips this summer, where the summer's illustrative load gives 0.047; or the reverse, the summer's 0.953 quoted as the design verdict on the opening. The two probabilities of overheating differ by a factor of almost four, and each is the right answer to the other sentence.

The third folds several questions into one expression. A single number for the whole memorandum, the joint of every clause, is reported where the owner's engineer wanted to know which assurance fails; or one clause's expression is answered for the paragraph, as when the opening's 0.82 is reported as "the facility's confidence". The repair is the set of expressions and their relations, with the conjunction added only if it was asked for. The fourth gives a condition the wrong scope: "over the first year" attached to the opening's adequacy, which yields the same number under a misleading label, since the conductance has one value; or "95 % confidence" read as a level for the conjunction where the paragraph meant each clause, or the reverse.

These are the forms of this lesson's failure class, *several questions folded into one expression*: a paragraph's questions written as fewer or more expressions than it holds, or with their relations misread. Its four forms are several questions folded into one expression, or one expression answered for all; one quantity counted as two, a duplicate multiplied or listed as independent assurance; a look-alike answered with another clause's expression; and a condition stated once applied to clauses it does not govern, or withheld from those it does. It enters at step 1, when the paragraph is split, and is caught at step 7 by reading each expression back against the others: are two of them one event, does one imply another, and does a condition in one belong in another?

9.4 One system, many questions

Table 9.6 gives eight pairs of sentences about the two-hall plant, the system to which the owner's sixth row belongs, with the relation between the events of each pair in the third column. It holds no values; the point is the relation.

Sentences about one plant relate in a small number of ways, and each way changes what a set of answers means together. The first, fourth and fifth pairs look different and are one event;

Table 9.6: Eight pairs of sentences about the two-hall plant and the relation between their events.

First sentence	Second sentence	Relation
Both units are unavailable at a random hour.	Both halls lose their cooling at a random hour, with the opening adequate.	One event: $F_{\text{fd}} \cup (F_A \cap F_B)$
The plant has the parallel structure.	The opening is adequate.	One event, $G \geq G^*$; on the two-drawing model also one event, with a different probability (Lesson 7)
Hall B is over its limit at 50 kW with unit B out.	Hall B is over its limit at 55 kW with unit B out.	Nested: the first implies the second
The month's hours without cooling exceed 3.65 h, feeder outages excepted.	The month's served share is below 0.995, feeder outages excepted.	One event, on a month of 730 h
Unit A fails at least once next year.	Unit A's first failure after the year's start comes before its end.	One event (Lesson 8)
The service is interrupted next year.	The feeder fails next year.	Nested: the second implies the first, which can happen without it
Hall B overheats if unit B trips at full load.	Hall B overheats if unit B trips on a night at 50 kW.	Different conditions on one limit state; the second's failure set lies inside the first's
The opening is adequate.	The plant keeps its cooling at a random hour.	Different: a parameter and an outcome (Lesson 7)

the third and seventh look alike and are nested, and the sixth is nested too; the last looks like a restatement and is a different kind of question.

9.5 The exercises

The solutions are in Appendix A, and the facts of the three systems in Appendix D. These are the hardest exercises of this part of the book: three of the four formulate exercises are paragraphs.

Exercise 9.1 (*formulate*). The site's specification for the battery room reads: "The string shall give two hours at the 20 kW critical load without the pack exceeding 318 K, and shall carry every grid outage of its first year. Each requirement is to be met with 90 % confidence." Take the paragraph apart, write each expression, compare them, and state whether the 90 % applies to each requirement or to all of them.

Exercise 9.2 (*formulate*). The network's service brief says: "The IT load shall lose power no more than once a year, and no loss of power shall last more than ten minutes; the UPS shall convert at an efficiency of at least 0.95 at full load, which is to say that its losses at full load shall not exceed 60 kW." Take the brief apart, write each expression, and find the clauses that name one quantity.

Exercise 9.3 (*formulate*). The owner's engineer says: "The cooling and the power need to be reliable at the same time." Formulate the statement.

Exercise 9.4 (*formulate*). The maintenance planner says: "Unit A goes out for service on Tuesday for six hours. If the feeder stays up and hall A is at 50 kW, hall A must stay within its

limit; and if unit B also fails, the halls must be shut down in order." Take the statement apart, place every condition, and state which clause is not a probability question.

Exercise 9.5 (*reverse*). Write one plain sentence for $P(\{G \geq G^{*}\} \cap \{60\,\mathrm{kW}/G \leq 8\,\mathrm{K}\} \mid E)$, and state what it equals.

Exercise 9.6 (*reverse*). Write one plain sentence for

$$P(Q_B/G > \Delta T_{\mathrm{allow}} \mid \text{unit B out},\ \text{settled},\ E),$$

with Q_B uncertain, and state how it differs from

$$P(\Delta T_{\mathrm{full}} > \Delta T_{\mathrm{allow}} \mid \text{unit B out},\ Q_{\mathrm{hall}},\ \text{settled},\ E).$$

Exercise 9.7 (*reverse*). Write one plain sentence for

$$P(\{T_{\mathrm{rt}} \geq 2\ \mathrm{h}\} \cap \{\text{every grid outage in } [0,\tau] \text{ ends before the floor}\} \mid \text{steady operation at } t=0,\ E),$$

with τ one year.

Exercise 9.8 (*critique and repair*). The site manager asks whether the battery string meets its two-hour specification and whether it will carry the next grid outage. The engineer's reply is one line: "The battery room is 83 % reliable." State what the figure answers, and repair the reply.

Exercise 9.9 (*critique and repair*). A handover note on the battery room reads: "Over the first year, with 95 % confidence, the string gives its two hours: $P(T_{\mathrm{rt}} \geq 2\ \mathrm{h}$ over the first year $\mid E) = 0.074$." State what is wrong with the expression, and repair it.

Exercise 9.10 (*critique and repair*). The owner asks two things: whether the opening is adequate, and whether the next utility interruption will reach the IT racks. The analyst replies with one number: "The probability that the opening is adequate and the next interruption is covered is $0.82 \times 0.15 = 0.12$." State what the number answers, and repair the reply.

Exercise 9.11 (*rapid*). For each pair, say whether the two sentences name one quantity, nested events, overlapping events or different quantities, and why.

(a) "Unit A is down for its own reasons"; "unit A's outage log shows it down".

(b) "The string reaches its floor before the grid returns"; "the next outage is not carried".

(c) "The feeder fails at least twice next year"; "the feeder fails more than once next year".

(d) "Hall B is over its limit at 52 kW with unit B out"; "the opening cannot pass 52 kW within 8 K".

(e) "The IT load loses power in the next interruption"; "the UPS fails to pick up the load in the next interruption".

(f) "The UPS is degraded"; "the utility meter reads 66 kW above its expected value".

(g) "Hall B is within its limit at full load with unit B out"; "hall A is within its limit at full load with unit A out".

(h) "The IT load is never without power next year"; "the feeder is never interrupted next year".

(i) "The cooling holds until unit B's repair ends"; "the cooling holds for the next 180 hours".

(j) "Both halls lose their cooling"; "the service is lost", with the opening adequate.

9.6 What to retain

- A paragraph holds several questions. List the clauses, write each expression with its own conditions, horizon and level, and only then compare them.
- Two sentences name one quantity when their events, written on the same unknowns, are one set under the conditions behind both bars. With unit B out, the full load fixed at 60 kW and the plant settled, $P(\Delta T_{\text{full}} > 8\,\text{K} \mid \text{unit B out},\ Q_{\text{hall}},\ \text{settled},\ E) = P(G < G^{*} \mid E)$, which is 0.18 on the 40 kW trip test.
- A sentence that looks the same can carry an uncertain condition the other fixes. With the summer's load uncertain, illustrative, the same limit gives 0.047, a joint over the conductance and the load.
- A duplicate multiplied turns 0.82 into 0.67.
- A condition stated once governs some clauses and not others; a parameter has no horizon unless it can change within it.
- "All of this with 95 % confidence" is either a level for each clause or one for the conjunction, and the paragraph must say which.

Part III

Constructing Probabilistic Models

Lesson 10

From Physical Relationships to Uncertain Quantities

10.1 The engineering question

The commissioning engineer has one trip test at 40 kW per hall and one reading from it, and two analysts run that reading through one stated model. One reports that the opening is adequate with probability 0.77, the other with probability 0.82. Neither has made an arithmetic error, and they used the same reading and the same prior. Asked what they did differently, the first says that he treated the logarithm of the conductance as bell-shaped, and the second that she treated its reciprocal as bell-shaped. The manager asks which number goes in the report, and why a choice of variable should change a probability at all.

The question is wider than one report. Before anything about a plant can be computed, someone has to say what in the model is a fixed condition, what is a property of the plant that nobody knows, what is a reading, and what the physical law leaves out. A model described in words hides those choices, and each of them ends up behind the bar of the answer whether or not it is said.

This lesson sorts a model's symbols into their roles, writes the relation between a reading and a parameter, and names the quantity that carries the bell shape.

10.2 The formulation

The work sits in step 6, where the model is now written as relations, and in step 3, where a condition, a parameter and a stated value are told apart by name. On the two-hall plant the evidence is E, the 40 kW trip test, whose measured difference of hall B above hall A is $d = 4.7828\,\mathrm{K}$.

10.2.1 Five roles

Each symbol of a model has one of five roles. A *condition* is known or set, and it stands behind the bar. A *parameter* has one unknown value. A *variable* follows from conditions and parameters through a relation, or it varies from occasion to occasion with a stated distribution, as tonight's load or the next outage's length does. An *observation* is a reading, in hand or still to come. A *discrepancy* is what a relation is known to leave out. A number taken as given, such as a limit, is marked *stated* and has no role in the uncertainty. Table 10.1 sorts the symbols of two systems, one role a row, with the test for each role in the second column.

A role belongs to a symbol in a question and not to the symbol for good. A load is a condition at the design point, a variable on a summer night, and a decision when a test is being chosen.

Table 10.1: The five roles, and stated values, on two systems.

Role	Test	Two-hall plant	Battery string
Condition	Known or set; goes behind the bar	Q; unit B out; settled	The 20 kW load; the start and the floor
Parameter	One unknown value; nothing has changed it	G; the sensors' offsets	φ; the weak cell
Variable	Follows from others through a relation, or varies from occasion to occasion with a stated distribution	ϑ_A, ϑ_B; the heat through the opening; tonight's load	The energy drawn; ϑ_{pk}; the next outage's length
Observation	A reading, in hand or to come	The hall readings; d	A timed reading of a discharge
Discrepancy	What a relation is known to leave out	δ: heat through the walls	δ: rate effects the energy balance omits
Stated	A number taken as given	$\Delta T_{\mathrm{allow}}$; G^{*}	The nameplate's 2.025 h

10.2.2 The physical relation and the measurement relation

Two relations are written and kept apart. The *physical relation* holds among variables and parameters. With unit B out and the plant settled, hall B's heat crosses the opening, and

$$Q_B = G\,(\vartheta_B - \vartheta_A) + \delta, \tag{10.1}$$

where δ is the heat the opening's law leaves out. The *measurement relation* ties a reading to the model's quantities: a reading is its quantity plus an error. Figure 10.1 sets the two relations of the trip test in a chain, the first between the middle boxes and the second to the right of them. The next lesson derives the error of d from the sensors, and here its spread is taken as stated, $\sigma_d = 0.711\,\mathrm{K}$. With $\delta = 0$ as a labeled assumption, the two relations combine into

$$d = \frac{Q}{G} + \varepsilon, \qquad \varepsilon \text{ normal, mean 0, spread } \sigma_d. \tag{10.2}$$

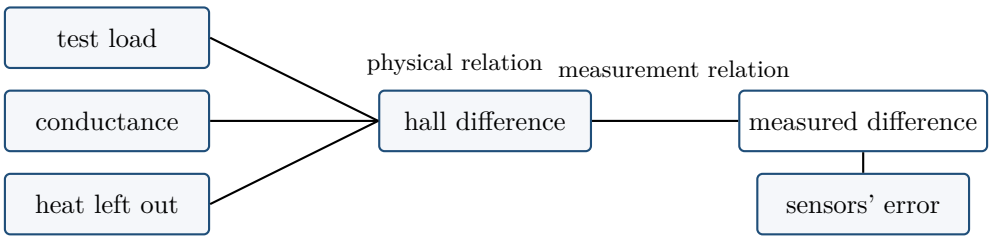


Figure 10.1: The two relations of the trip test. A line joins the quantities a relation ties. The physical relation gives the hall difference from the load, the conductance and the heat left out; the measurement relation gives the measured difference from it and the sensors' error.

A discrepancy is named, given a role, and either carried or labeled as dropped. It is not estimated here. Whether a test can separate δ from G is a question about what a measurement can tell, and Lesson 18 asks it.

10.2.3 Three relations on one opening

One model can be written in several coordinates. Table 10.2 writes equation (10.2) in three, one a row: in the conductance, in its logarithm $\theta = \log G$, and in the opening's thermal resistance

$R = 1/G$. The last two columns say whether the relation is linear in the coordinate and how the adequacy threshold reads in it.

Table 10.2: One measurement relation in three coordinates, with $R^* = 1/G^* = 0.1333$ K/kW.

Coordinate	Measurement relation	Linear in it	Adequacy event
G	$d = Q/G + \varepsilon$	no	$G \geq G^*$
$\theta = \log G$	$d = Qe^{-\theta} + \varepsilon$	no	$\theta \geq \log G^*$
$R = 1/G$	$d = QR + \varepsilon$	yes	$R \leq R^*$

The three events are one event, so changing coordinate changes neither the plant nor the answer, provided the prior goes with it. The stated prior is $\log G$ normal about $\log G^*$ with spread $\sigma_0 = \log 3$. Since $\log(R/R^*) = -\log(G/G^*)$, the prior in R is

$$p(R) \;\propto\; \frac{1}{R}\,\exp\Big[-\frac{(\log R - \log R^*)^2}{2\sigma_0^2}\Big], \tag{10.3}$$

where the factor $1/R$ is the Jacobian of the change of variable. This is the same prior, written elsewhere; a normal on R would be a different model. The posterior in R is the likelihood $\mathcal{N}(d;\, QR,\, \sigma_d^2)$ times this prior. Figure 10.2 draws that one posterior in the three coordinates as solid curves, with the side of the threshold that passes shaded. The shaded probability is the same in the three panels, and the curve has a different shape in each.

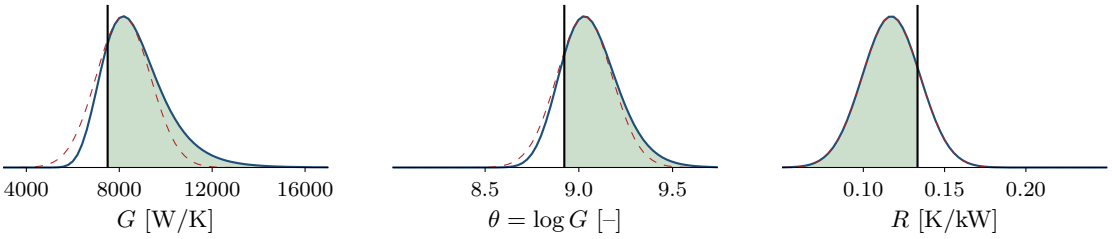


Figure 10.2: One posterior on E, the 40 kW trip test, in three coordinates, each density scaled to a peak of one. Solid, the posterior; shaded, the side of the threshold that passes, 0.818 in every panel; dashed, the normal curve of the Gaussian reading in that coordinate; the vertical line, the threshold.

The three coordinates give three formulations of one question, and they differ line by line. Table 10.3 holds them as one run, lettered by coordinate, with the expression, its model line and its reduction in each.

The threshold's direction turns over in the third coordinate, since a larger conductance is a smaller resistance. The modes differ too: 8174 W/K in the density of G, 8347 W/K in that of θ and 8535 W/K in that of R. They are three points of three densities, and none of them is the conductance.

10.2.4 The quantity that carries the Gaussian

A posterior is often reported through a *Gaussian reading*: its density in one coordinate is read as normal about the density's mode, with the spread its curvature gives, and the tail at the threshold is a value of Φ, the standard normal distribution function. In R,

$$P(R \leq R^* \mid E) \;\approx\; \Phi\Big(\frac{R^* - m_R}{\sigma_R}\Big), \tag{10.4}$$

Table 10.3: One question in three coordinates, as one lettered run, on E.

Run, the opening adequate, in G, in θ and in R		
1	**Sentence**	Given E, unit B out, settled: probability the opening is adequate?
	Intention	Estimate
2	**Target**	(a) G; (b) $\theta = \log G$; (c) $R = 1/G$: one parameter, three coordinates
	Declare	$d = 4.7828\,\mathrm{K}$; $Q = 40\,\mathrm{kW}$; $\sigma_d = 0.711\,\mathrm{K}$; $\sigma_0 = \log 3$; $G^* = 7500\,\mathrm{W/K}$; $R^* = 1/G^*$
3	**Known**	E; unit B out; settled; $\delta = 0$ (assumption)
	Random here	The parameter alone, in whichever coordinate
4	**Event**	(a) $G \geq G^*$; (b) $\theta \geq \log G^*$; (c) $R \leq R^*$: one event
5	**Expression**	(a) $P(G \geq G^* \mid E)$; (b) $P(\theta \geq \log G^* \mid E)$; (c) $P(R \leq R^* \mid E)$
6	**Model**	(a) $d = Q/G + \varepsilon$; (b) $d = Qe^{-\theta} + \varepsilon$; (c) $d = QR + \varepsilon$
	Reduce	(a) $\Phi((m_G - G^*)/\sigma_G)$; (b) $\Phi((m_\theta - \log G^*)/\sigma_\theta)$; (c) $\Phi((R^* - m_R)/\sigma_R)$
7	**Check**	Exact in all three: 0.818. Readings: (a) 0.719; (b) 0.767; (c) 0.818
	Class	Posterior exceedance
		$\boxed{P(G \geq G^* \mid E) = P(\theta \geq \log G^* \mid E) = P(R \leq R^* \mid E)}$

with m_R the mode and σ_R the spread. In numerical work this is the Laplace approximation; how the mode and curvature are found is Appendix E's concern. The exact probability does not depend on the coordinate, and a Gaussian reading does, which the dashed curves of Figure 10.2 show. The coordinate is therefore part of the formulation. Table 10.4 sizes this on three tests, one a row: the exact value in the second column and the reading in each coordinate in the next three. The last two columns are the checklist's numbers below: how far the reading's point d/Q lies from the prior's bound at $R = 0$, counted in reading spreads σ_d/Q, which comes to d/σ_d, and the slope of the prior's log density at that point, times the reading's spread. The values come from a small calculation on the stated model.

Table 10.4: The Gaussian reading in each coordinate against the exact probability that the opening is adequate, with the reading's distance from the prior's bound. The second and third rows are tests other than E, the 5 kW test an illustration.

Test	Exact [–]	In G [–]	In $\log G$ [–]	In R [–]	Bound [spreads]	Straightness [–]
E, 40 kW, $d = 4.7828\,\mathrm{K}$	0.818	0.719	0.767	0.818	6.7	0.135
20 kW, $d = 2\,\mathrm{K}$, the reading a truth of 80 kW would give	0.882	0.704	0.788	0.888	2.8	0.271
5 kW, $d = 0.5\,\mathrm{K}$, the same truth	0.679	0.255	0.563	0.955	0.70	1.08

The first row gives the rule. The observation relation is linear in R with a fixed spread, so the likelihood is exactly Gaussian there, and multiplying it by a prior whose log density is straight only moves its center. The third row gives the rule's condition. The reading lies within a spread of the prior's bound, the prior shapes the posterior, and the reading in R is wrong by 0.276. On that test it loses to the reading in $\log G$, wrong by 0.116; the reading in G, wrong by 0.424, is the worst of the three on every row. Lesson 16 returns to this reduction with a reading not yet taken, where the acceptance probability in R takes the form of equation (10.4), under the same condition.

Principle 10.1. *Pick the coordinate in which the observation relation is linear with a fixed spread, and transform the prior rather than replace it. The Gaussian reading there is good when the prior has one width and one hump, the reading's point lies inside the prior's range and at least 3 reading spreads from its nearest bound, and the prior's log density is nearly straight across the reading's width at the point. Name the coordinate of any Gaussian reading; give the exact value where one is available.*

The rule is a checklist of four steps, taken in order.

1. **Coordinate.** Find the coordinate in which the observation relation is linear with a fixed spread, and write the reading's point and its spread there. If there is none, stop.
2. **Prior's shape.** The screen covers a prior of one width and one hump: normal in the quantity, log-normal with a spread of at most a factor of 3, or uniform between bounds. Any other prior stops the procedure: a mixture of explanations or of widths, a heavy-tailed or skewed prior, two humps. A prior normal in the reading's coordinate, with no bound of any kind, passes steps 3 and 4 at once, exactly.
3. **Range and bounds.** The reading's point must lie inside the prior's range and at least 3 reading spreads from its nearest bound. A physical bound counts even where the sheet's prior does not state it, such as $\eta \leq 1$, and a normal prior with a bound takes this step and the next. In R the bound is $R = 0$ and the distance is the point over its spread, which for one reading is d/σ_d.
4. **Straightness.** The slope of the prior's log density at the point, times the reading's spread, is at most 3/4. In R, on the stated prior, that is $(\sigma_d/d)\,\left|1 + \log\big((d/Q)/R^*\big)/\sigma_0^2\right|$.

If all four pass, use the reading and name its coordinate. If any fails, write the exact expression and name the step that failed.

This is a screening test and not a guarantee, and its cuts are round numbers. It was tested on one family, the log-normal prior in R, for absolute errors at moderate probabilities, on 855 illustrative tests at loads from 0.25 to 60 kW and at 3 to 14 spreads from the bound. On the stated prior 782 pass, with a largest error of about 0.04, and across prior widths of a factor of 1.1 to 3 it is about 0.05. Wider priors reach 0.17 on the weakest tests, whose posterior has a second hump near $R = 0$, so step 2 stops at a factor of 3. Far in the tail the absolute error stays small and the ratio does not: an illustrative 4 kW test at 3 spreads reads 36 times too high at an exact 0.001. Step 4 catches what a distance cannot: an illustrative 0.5 kW test reading 3 K lies 4.2 spreads from the bound, its slope is 0.99, and its reading gives 0.94 for an exact 0.056. Where the posterior has two humps, no mode carries it, as close to an illustrative 2 kW test at 2.1 K, with humps at 0.094 and 0.55 K/kW. Many refusals are of readings that would have been fine, and a refusal costs only writing the exact expression. Table 10.5 runs the checklist on the three tests, lettered by test.

10.2.5 The opening after the trip test

The first worked formulation asks whether the opening is adequate, in the coordinate the rule picks. Table 10.6 is its sheet, with steps 3 and 6 in full.

The exact value is 0.818 and the reading in R gives 0.818. The sheet's sixth row carries the lesson: it names the relation, the prior's coordinate and the checklist's numbers, and an expression quoted without them would not say what 0.818 rests on.

10.2.6 The string after a timed discharge

The second worked formulation is on the battery string, where the relation is already linear and the prior is what a normal reading would lose. Suppose a timed discharge at the critical load, a

Table 10.5: The rule run on the three tests of Table 10.4.

Run, which quantity carries the Gaussian, on three tests	
1 Sentence	(a) 40 kW; (b) 20 kW; (c) 5 kW: may adequacy be quoted by a Gaussian reading?
Intention	Estimate, with a choice of reading
2 Target	A parameter, R, and the reading that may carry it
Declare	d/Q the reading's point; σ_d/Q its spread in R; $R = 0$ the prior's bound
3 Known	Each test's Q and d; σ_d; $\delta = 0$ (assumption); the prior on $\log G$
Random here	R only; the coordinate is chosen, not random
4 Event	$R \le R^*$
5 Expression	$P(R \le R^* \mid d)$, for each test
6 Model	$d = QR + \varepsilon$; one hump; prior transformed; spreads: (a) 6.7; (b) 2.8; (c) 0.70
Reduce	(a) All four pass: Gaussian reading in R; (b), (c) step 3 fails: exact
7 Check	(a) 0.818 for 0.818; (b) exact 0.882, reading 0.888; (c) exact 0.679, reading 0.955
Class	Posterior exceedance
	$\boxed{\text{(a) } P(R \le R^* \mid d) \approx \Phi\big((R^* - m_R)/\sigma_R\big); \quad \text{(b), (c) } P(R \le R^* \mid d)}$

Table 10.6: The formulation sheet filled in for the opening after the 40 kW trip test, in R.

Field	Entry
Intention	Estimate
1. Sentence	With unit B out and the plant settled, given the 40 kW trip test, what is the probability that the opening's thermal resistance is at most R^*, the value that passes hall B's 60 kW within 8 K?
2. Target	A parameter, R, through an event on it
3. Known	E: $d = 4.7828\,\text{K}$ at $Q = 40\,\text{kW}$. Conditions: unit B out, settled. Stated: σ_d, and the prior on $\log G$. Assumption: $\delta = 0$.
4. Event	$R \le R^*$, the same event as $G \ge G^*$
5. Expression	$P(R \le R^* \mid E)$
6. Model	The measurement relation $d = QR + \varepsilon$, linear in R; the prior of equation (10.3), transformed. If a reading is used, the coordinate is named and the checklist's numbers are stated: 6.7 spreads, straightness 0.135.
7. Read back	After this test, the probability that the opening passes hall B's full load within the limit. Silent on whether the service is delivered at any hour.
Problem class	Posterior exceedance

One line per step, with the reading's reduction:

Run, the opening after the 40 kW trip test	
1 Sentence	Given E, unit B out, settled: probability the opening's resistance is at most R^*?
Intention	Estimate
2 Target	A parameter, R, through an event on it
Declare	d [K]; $Q = 40\,\text{kW}$; R [K/kW]; $R^* = 0.1333$ K/kW; $\sigma_d = 0.711\,\text{K}$; $\sigma_0 = \log 3$
3 Known	E; unit B out; settled; $\delta = 0$ (assumption)
Random here	R only; d is in hand
4 Event	$R \le R^*$
5 Expression	$P(R \le R^* \mid E)$
6 Model	$p(R \mid E) \propto \mathcal{N}(d;\, QR,\, \sigma_d^2)\, R^{-1}\, \mathcal{N}(\log R;\, \log R^*,\, \sigma_0^2)$, linear in R
Reduce	Gaussian reading in R: $\Phi\big((R^* - m_R)/\sigma_R\big)$, $m_R = 0.1172$, $\sigma_R = 0.0178$ K/kW
7 Check	The event $G \ge G^*$; the coordinate named; 6.7 spreads, straightness 0.135
Class	Posterior exceedance
	$\boxed{P(R \le R^* \mid E) \approx \Phi\big((R^* - m_R)/\sigma_R\big)}$

test the source study does not have, reads $d = 2.00$ h with a spread of 0.03 h, both illustrative. For a sound string the relation is $d = 2.025 \text{ h} \times \varphi + \delta + \varepsilon$, and the fact sheet's prior on φ is uniform from 0.85 to 1.00. A normal reading in φ puts 0.20 of its weight above 1, where no string can be. Table 10.7 is the sheet.

Table 10.7: The formulation sheet filled in for the string after a timed discharge; the reading and its spread are illustrative.

Field	Entry
Intention	Estimate, for a decision on the specification
1. Sentence	A timed discharge at 20 kW reads 2.00 h with a spread of 0.03 h: what is the probability that the string gives at least its two hours?
2. Target	A parameter, φ, through T_{rt}
3. Known	The reading; 20 kW; the stated start and floor. Assumptions: a sound string; $\delta = 0$.
4. Event	$T_{\text{rt}} \geq 2$ h, that is $\varphi \geq 2/2.025 = 0.9877$
5. Expression	$P(\varphi \geq 0.9877 \mid d, \text{ sound})$
6. Model	The relation, linear in φ; the uniform prior with its bounds, carried. The reading lies 0.83 spreads from the bound at $\varphi = 1$, so no Gaussian reading.
7. Read back	What the test says of the string's two hours. Silent on the weak cell and on any outage.
Problem class	Posterior exceedance

The same sheet as a run:

Run, the string after a timed discharge

1	**Sentence**	Timed discharge at 20 kW reads 2.00 h: probability the string gives its 2 h?
	Intention	Estimate, for a decision
2	**Target**	A parameter, φ, through T_{rt}
	Declare	d [h]; φ [–]; $T_{\text{rt}} = 2.025 \text{ h} \times \varphi$; $m_\varphi = d/2.025$; $\sigma_\varphi = 0.03/2.025$ (illustrative)
3	**Known**	d; 20 kW; sound string (assumption); $\delta = 0$ (assumption)
	Random here	φ only
4	**Event**	$\varphi \geq 0.9877$
5	**Expression**	$P(\varphi \geq 0.9877 \mid d, \text{ sound})$
6	**Model**	$p(\varphi \mid d) \propto \mathcal{N}(d;\, 2.025\,\varphi,\, 0.03^2)\ \mathbf{1}[0.85 \leq \varphi \leq 1]$; 0.83 spreads from 1
	Reduce	$\big(\Phi(a_1) - \Phi(a_{0.9877})\big)/\big(\Phi(a_1) - \Phi(a_{0.85})\big)$, with $a_x = (x - m_\varphi)/\sigma_\varphi$
7	**Check**	The bound carried; silent on the weak cell and on any outage
	Class	Posterior exceedance

$$\boxed{P(\varphi \geq 0.9877 \mid d, \text{ sound})}$$

With the bound carried the value is 0.373, and with it ignored the normal reading gives 0.50. The relation is linear here, and the checklist still stops at step 3, because the reading lies too near the bound.

10.2.7 The opening with its discrepancy carried

The two sheets above drop δ with a label. Carrying it changes more than one line, which is the reason to write it once. With δ in equation (10.1), the difference once the plant has settled is $(Q_B - \delta)R$, so the discrepancy enters the measurement relation and the event as well. The relation is no longer linear in (R, δ). It is linear in $(R, \delta R)$, where the prior is no longer a product, and a run written on R alone would hide that. The spread of δ and the assumption that it is the same at the test load and at full load are both the analyst's, and both are labeled. One reading cannot separate δ from R, which is the question Lesson 18 asks; here it is only

written.

Run, the opening, the heat through the walls carried	
1 Sentence	Given E, with the heat through the walls unknown: hall B within 8 K at 60 kW?
Intention	Estimate
2 Target	Parameters R and δ, through an event on both
Declare	R [K/kW]; δ heat left out [kW]; $Q_{\text{hall}} = 60\,\text{kW}$; $Q = 40\,\text{kW}$; $p(\delta)$ (illustrative)
3 Known	E; unit B out; settled; δ the same at both loads (assumption)
Random here	R and δ; δ asked about by no one, and carried
4 Event	$(Q_{\text{hall}} - \delta)R \le \Delta T_{\text{allow}}$
5 Expression	$P\big((Q_{\text{hall}} - \delta)R \le \Delta T_{\text{allow}} \mid E\big)$
6 Model	$p(R, \delta \mid E) \propto \mathcal{N}(d;\, (Q - \delta)R,\, \sigma_d^2)\, p(R)\, p(\delta)$, bilinear
7 Check	With $\delta = 0$ it returns the first run; the threshold on R moves with δ
Class	Posterior exceedance
	$\boxed{P\big((Q_{\text{hall}} - \delta)R \le \Delta T_{\text{allow}} \mid E\big)}$

10.2.8 A model request that is not ready

At the handover the owner says: “Build a model of the opening.” The sentence names a system and no question, and it fails at step 1, since the roles cannot be assigned without one. For a question about the design case the load is a condition. For a question about tonight it is an uncertain quantity. For a question about which test to run it is a decision. Each reading gives another list of roles and another model, and a list of symbols with no question has no way of saying which should stay. The analyst’s reply asks which question the model must answer.

Key idea

Every symbol of a model has a role, and a physical relation and a measurement relation are written separately. One model has several coordinates. An exact probability does not depend on the coordinate, a Gaussian reading does, and so the coordinate is part of the formulation.

10.3 The competing formulations

Each competing formulation here leaves off something a number rests on. The first reports a Gaussian reading as if it were the probability, which is what the first analyst of the opening did.

Run, the coordinate dropped	
1 Sentence	Given E: probability the opening is adequate?
5 Expression	$P(G \ge G^* \mid E)$
6 Reduce ×	0.77, reported as the probability, no coordinate and no reading named
Repair	0.818 exact; or “0.77, a Gaussian reading in $\log G$”
7 Check	The same model and reading give 0.72 in G; the number must carry its coordinate
	$\boxed{P(G \ge G^* \mid E)}$

The second moves to R and takes a convenient prior there. A normal prior on R, with the lognormal’s center and its spread at that center, is not the stated prior. The model has changed, and the exact value with it.

Run, the prior replaced	
3 Known	E; the stated prior on $\log G$
6 Model ×	R normal, 0.133 ± 0.146 K/kW, cut at 0: exact 0.779 on E, 0.819 at 20 kW
Repair	The prior of equation (10.3): exact 0.818 on E, 0.882 at 20 kW
7 Check	Another prior is another model, however close its center
	$\boxed{p(R) \propto R^{-1}\,\mathcal{N}(\log R;\, \log R^*,\, \sigma_0^2)}$

The third keeps the prior and loses the Jacobian. The mode stays where it was in $\log G$, at $G = 8347$ W/K, when the density in R has its mode at 8535 W/K.

Run, the Jacobian dropped

5	**Expression**	$P(R \le R^* \mid E)$
6	**Reduce** ×	Reading in R, the density without its $1/R$: 0.779 on E, 0.820 at 20 kW
	Repair	The density with its factor: 0.818 on E, 0.888 at 20 kW
7	**Check**	A density changes with its coordinate; a probability of one event does not

$$\boxed{P(R \le R^* \mid E) \approx \Phi\big((R^* - m_R)/\sigma_R\big)}$$

The fourth applies the rule and forgets its condition: the reading in R on the weak test of Table 10.4, used because the relation is linear.

Run, the rule without its condition

1	**Sentence**	5 kW test, $d = 0.5$ K: probability the opening is adequate?
6	**Reduce** ×	Gaussian reading in R, "the relation is linear": 0.955
	Repair	Exact: 0.679; the reading is 0.70 spreads from the bound, so step 3 fails
7	**Check**	How far is the reading from the prior's bound? Under one spread

$$\boxed{P(G \ge G^* \mid d)}$$

The fifth states a physical relation as a fact. The sheet of Table 10.6 rests on $\delta = 0$, on the plant settled and on the stated spread. Quoted with none of them, 0.82 reads as a fact about the opening. No number sizes this one: how much δ matters is unknown until it is modeled, and the sheet must say so.

Run, the relation stated as fact

1	**Sentence**	Given E: is the opening adequate?
3	**Known** ×	E only; nothing else behind the bar
	Repair	E; unit B out; settled; $\delta = 0$ (assumption); the stated σ_d and prior
7	**Check**	Would another δ change the number? Then it is on the sheet

$$\boxed{P(G \ge G^* \mid E), \text{ given the stated model}}$$

Three smaller slips belong beside these. A change of coordinate can turn a threshold over and leave the inequality as it was.

Run, the threshold not turned over

1	**Sentence**	Given E: is the opening adequate, in R?
4	**Event** ×	$R \ge R^*$, copied from $G \ge G^*$: 0.182
	Repair	$R \le R^*$, since $R = 1/G$: 0.818
7	**Check**	A larger resistance passes less heat; the event must read that way

$$\boxed{P(R \le R^* \mid E)}$$

A spread belongs to its coordinate too. The reading in θ on E has a spread of 0.147, which is a share of the conductance and not a number of watts per kelvin.

Run, a spread read in the wrong coordinate

	Declare	$\sigma_\theta = 0.147$ [–], the spread of $\log G$
6	**Model** ×	"G is 8347 W/K, plus or minus 0.147"
	Repair	$\log G$ has spread 0.147: about 15 % of G, not symmetric in G
7	**Check**	Each number carries the unit of its own coordinate

$$\boxed{\mathcal{N}(\theta;\, m_\theta,\, \sigma_\theta^2)}$$

And a discrepancy can be absorbed into the noise, which treats a heat that is there on every

test as an error that differs from test to test.

Run, a discrepancy folded into the noise

3 Known	E; heat through the walls not zero (a stated concern)
6 Model ×	$d = Q/G + \varepsilon$, with σ_d widened "to cover it"
Repair	$Q_B = G(\vartheta_B - \vartheta_A) + \delta$, with δ carried; or $\delta = 0$, labeled
7 Check	δ enters the event, $(Q_{\text{hall}} - \delta)R \le \Delta T_{\text{allow}}$; widened noise leaves the event as it was

$$\boxed{d = (Q - \delta)/G + \varepsilon}$$

These are the forms of the recurring class this lesson sharpens, *model conditioning dropped*: an expression stated without the structure, relation or reading it rests on, so that an assumption is read as a fact. Its five forms are a relation or a discrepancy assumption left off; the coordinate of a Gaussian reading left off; a prior replaced in a change of variable and reported as the same model; the Jacobian dropped; and a rule of thumb applied without its condition. The class enters at step 3 or step 6, and it is caught at step 7 by asking what, if changed, would change the number and is not written.

10.4 One system, many questions

Table 10.8 gives ten phrases that come up about the battery string, with the role each has in a question about the string's two hours, and the relation where a reading enters.

Table 10.8: Ten phrases about the battery string and the role of each, in a question about its two hours.

Phrase	Role, or relation
The string's state of health	Parameter φ, bounded by its prior
The critical load, 20 kW	Condition, stated
The pack's temperature during the discharge	Variable, through its heat balance
A timed discharge reading	Observation, $d = 2.025\ \text{h} \times \varphi + \delta + \varepsilon$, linear in φ
The state-of-charge reading at an outage's start	Observation, about the start and not about φ
The heat the pack's model leaves out	Discrepancy, in the heat balance
The discharge efficiency, 0.90	Stated; a parameter if the source is doubted
Whether one cell is weak	A discrete parameter
An energy-meter reading of a discharge	Observation, linear in φ through the stated usable energy
The next outage's length	Variable, with a stated distribution; not a parameter

One string holds all ten, and each role decides what steps 3 and 6 must say. The last row is there because a length that varies from outage to outage has no one value to estimate.

10.5 Worked problems

Eight problems on the three systems follow, each with its run.

Problem 10.1. The prior is the near pair's two drawings, 6750 or 8250 W/K, each within 2 %, and a 40 kW test reads 5.33 K (illustrative). May adequacy be quoted by a Gaussian reading in R?

Run, a prior with two humps

1	**Sentence**	Two-drawing prior, 40 kW test reads 5.33 K: a Gaussian reading of adequacy in R?
	Intention	Estimate, with a choice of reading
2	**Target**	A parameter, R, and the reading that may carry it
	Declare	Point 0.133, spread 0.018 K/kW; drawings 0.121, 0.148 K/kW; $\omega_i = \frac{1}{2}$ (choice)
3	**Known**	d; unit B out; settled; $\delta = 0$, and the as-built opening one of the two (assumptions)
	Random here	R, about one of two values
4	**Event**	$R \le R^*$, which holds on the 8250 W/K drawing and fails on the other
5	**Expression**	$P(R \le R^* \mid d)$
6	**Model**	$d = QR + \varepsilon$, linear; two humps, so step 2 stops the checklist
	Reduce	The 8250 W/K drawing's share, with $\ell_i = \int \mathcal{N}(d;\, QR,\, \sigma_d^2)\, p_i(R)\, \mathrm{d}R$
7	**Check**	0.531 on these weights, 0.911 at 0.9 and 0.1; the sheet names step 2
	Class	Posterior exceedance

$$\boxed{P(R \le R^* \mid d) = \omega_{8250}\ell_{8250} / (\omega_{8250}\ell_{8250} + \omega_{6750}\ell_{6750})}$$

Problem 10.2. On the meter-offset explanation the PDU meter reads 1192 kW and the UPS output meter 1150 kW (illustrative). What is the probability that the PDU offset exceeds 30 kW?

Run, the offset, read from a difference of two meters

1	**Sentence**	Offset explanation, $y_{\mathrm{pdu}} = 1192$, $y_{\mathrm{ups}} = 1150\,\mathrm{kW}$: probability $b_{\mathrm{pdu}} > 30\,\mathrm{kW}$?
	Intention	Estimate
2	**Target**	A parameter, b_{pdu}, through an event on it
	Declare	b_{pdu} [kW], $0 \pm 45\,\mathrm{kW}$, normal (choice); m_b, σ_b its posterior center and spread
3	**Known**	The explanation; noises 2 and 2 kW; the load's prior flat (choice)
	Random here	b_{pdu}; the load cancels
4	**Event**	$b_{\mathrm{pdu}} > 30\,\mathrm{kW}$
5	**Expression**	$P(b_{\mathrm{pdu}} > 30\,\mathrm{kW} \mid y_{\mathrm{pdu}}, y_{\mathrm{ups}})$
6	**Model**	$y_{\mathrm{pdu}} - y_{\mathrm{ups}} = b_{\mathrm{pdu}} + \varepsilon_3 - \varepsilon_2$, spread 2.83 kW; a normal prior: step 2 passes it
	Reduce	$1 - \Phi\big((30 - m_b)/\sigma_b\big)$, a normal times a normal
7	**Check**	Φ is exact on one explanation; the explanations together are a mixture: step 2 stops
	Class	Posterior exceedance

$$\boxed{P(b_{\mathrm{pdu}} > 30\,\mathrm{kW} \mid y_{\mathrm{pdu}}, y_{\mathrm{ups}}) = 1 - \Phi\big((30 - m_b)/\sigma_b\big)}$$

Problem 10.3. The owner of the battery room says: "Put the pack's temperature on a sounder footing with a Gaussian reading."

Run, a reading asked for with nothing read

1	**Sentence**	"Put the pack's temperature on a sounder footing with a Gaussian reading."
	Intention	Not stated: no reading, no event, no horizon
	Declare	R_{p} pack-to-room resistance [K/kW]; y a heater test's reading [K]; M peak [K]
3	**Known**	No temperature test is stated; the fact sheet holds simulated peaks only
4	**Event**	(a) peak under 318 K, 1.5 h outage; (b) the same, next outage; (c) none, heat leak
5	**Expression**	(a) a probability at a stated length; (b) one over D_{out}; (c) a density, after a test
6	**Model**	(a), (b) the pack's heat balance; (c) a relation for a test that has not been run
7	**Check**	A Gaussian reading needs a reading; the analyst asks which question is meant
	Class	Not ready

$$\boxed{\text{(a) } P(M(t_0 + 1.5\ \mathrm{h}) < 318\,\mathrm{K});\ \text{(b) } P(M(t_0 + D_{\mathrm{out}}) < 318\,\mathrm{K});\ \text{(c) } p(R_{\mathrm{p}} \mid y)}$$

Problem 10.4. A timed discharge is started from a state of charge of 0.60 and reads 1.30 h, spread 0.03 h (illustrative). What is the probability that the string gives its 2 h from the usual 0.80?

Run, a test from another start

1	**Sentence**	Discharge from 0.60 reads 1.30 h: probability of 2 h from 0.80?
	Intention	Estimate
2	**Target**	A parameter, φ; the test's start and the event's start differ
	Declare	d [h]; $x_0 = 0.60$ in the test; $m_\varphi = d/1.35$; $\sigma_\varphi = 0.03/1.35$; $a_x = (x - m_\varphi)/\sigma_\varphi$
3	**Known**	d; 20 kW; sound, $\delta = 0$, energy in proportion to $x_0 - 0.2$ (assumptions)
	Random here	φ only
4	**Event**	2.025 h $\times\, \varphi \geq 2$ h, that is $\varphi \geq 0.9877$
5	**Expression**	$P(\varphi \geq 0.9877 \mid d,\ \text{sound})$
6	**Model**	$d = 2.025\ \text{h} \times \varphi \times 0.4/0.6 + \varepsilon = 1.35\ \text{h} \times \varphi + \varepsilon$; 1.67 spreads from 1.35 h
	Reduce	Step 3 fails: $\big(\Phi(a_1) - \Phi(a_{0.9877})\big)/\big(\Phi(a_1) - \Phi(a_{0.85})\big)$
7	**Check**	The bound in hours moved with the test's start; the event did not
	Class	Posterior exceedance

$$\boxed{P(\varphi \geq 0.9877 \mid d,\ \text{sound})}$$

Problem 10.5. A trip test is being chosen. If the opening is as the 66 kW drawing says, what test load keeps the expected reading 3 spreads from the prior's bound?

Run, the test load as the decision

1	**Sentence**	If G is the drawn 8.25 kW/K: which Q puts the reading 3 spreads from $R = 0$?
	Intention	Design an experiment
2	**Target**	A decision, Q; here the load is neither condition nor variable
	Declare	Q [kW]; $G_\text{d} = 8.25$ kW/K; $\sigma_d = 0.711$ K; expected reading Q/G_d [K]
3	**Known**	σ_d; the drawn value, taken as if true (assumption)
	Random here	Nothing: a condition on the design
4	**Event**	None: a set of test loads
5	**Expression**	$\{Q : Q/(G_\text{d}\,\sigma_d) \geq 3\}$
6	**Model**	The relation at its expected reading; the distance is d/σ_d
	Reduce	$Q \geq 3\,\sigma_d\,G_\text{d} = 17.6$ kW
7	**Check**	The degenerate form, at an expected reading; Lesson 16 writes the full one
	Class	Preposterior design

$$\boxed{\{Q : Q \geq 3\,\sigma_d\,G_\text{d}\}}$$

Problem 10.6. Suppose the heat through the walls is proportional to the hall difference, with an unknown conductance G_w. Given E, does hall B hold 8 K at 60 kW?

Run, a discrepancy that scales with the difference

1	**Sentence**	Wall heat proportional to the difference, given E: hall B within 8 K at 60 kW?
	Intention	Estimate
2	**Target**	An event on the total path, not on the opening alone
	Declare	G_w [W/K]; $R_\text{t} = 1/(G + G_\text{w})$ [K/kW]; R^*; $Q = 40$ kW
3	**Known**	E; unit B out; settled; the wall heat's form (assumption)
	Random here	R_t; G and G_w apart are not asked about
4	**Event**	60 kW $\times\, R_\text{t} \leq 8$ K, that is $R_\text{t} \leq R^*$
5	**Expression**	$P(R_\text{t} \leq R^* \mid E)$
6	**Model**	$Q_B = (G + G_\text{w})(\vartheta_B - \vartheta_A)$; $d = QR_\text{t} + \varepsilon$; a prior on $\log(G + G_\text{w})$ (choice)
7	**Check**	No separation needed for this event; the opening's own adequacy needs G alone
	Class	Posterior exceedance

$$\boxed{P(R_\text{t} \leq R^* \mid E)}$$

Problem 10.7. Unit A's duty meter, with a noise of 2 % of reading, reads the total load as 81.6 kW. In which coordinate does the total carry a Gaussian reading?

Run, noise that scales with the quantity

1	**Sentence**	Duty meter reads 81.6 kW, noise 2 % of reading: which coordinate for the total?
	Intention	Estimate, with a choice of reading
2	**Target**	A parameter, the total load Q_t [kW], and its coordinate
	Declare	y [kW]; Q_t [kW]; ε of spread 0.02 [–]; $\log Q_\mathrm{t}$ normal (choice)
3	**Known**	y; unit B out; settled; the unit's duty equals the total (assumption)
	Random here	Q_t only
4	**Event**	None: a density
5	**Expression**	$p(\log Q_\mathrm{t} \mid y)$
6	**Model**	$y = Q_\mathrm{t}(1+\varepsilon)$: in Q_t the spread grows; $\log y \approx \log Q_\mathrm{t} + \varepsilon$, spread fixed: step 1
	Reduce	Gaussian reading in $\log Q_\mathrm{t}$: no bound in that coordinate, normal prior
7	**Check**	The log relation is approximate; at 2 % the coordinate barely matters
	Class	Posterior of a parameter

$$p(\log Q_\mathrm{t} \mid y) \approx \text{const} \times \mathcal{N}(\log y;\, \log Q_\mathrm{t},\, 0.02^2)\, p(\log Q_\mathrm{t})$$

Up to about 10 % noise the log relation passes step 1, with an error of at most about 0.05, the screen's own size, for a reading within two prior spreads of the prior's center; beyond that step 1 fails, and the error reaches 0.16 at 30 %.

Problem 10.8. On the network, the load and the efficiency are both unknown, read by the UPS output meter and the utility meter. May η be quoted by a Gaussian reading?

Run, no coordinate that passes the first step

1	**Sentence**	Load and efficiency both unknown, two meters: a Gaussian reading of η?
	Intention	Estimate, with a choice of reading
2	**Target**	A parameter, η; the load P_load is nuisance
	Declare	y_ut, y_ups [kW]; P_load [kW]; $P_\mathrm{in} = P_\mathrm{load}/\eta$ the input [kW]
3	**Known**	A steady snapshot; noises 3 and 2 kW; the shape of $p(\eta)$ (choice)
	Random here	η; P_load, integrated out
4	**Event**	None: a density
5	**Expression**	$p(\eta \mid y_\mathrm{ut}, y_\mathrm{ups})$
6	**Model**	$y_\mathrm{ut} = P_\mathrm{in} + \varepsilon_1$, $y_\mathrm{ups} = P_\mathrm{load} + \varepsilon_2$: linear in $(P_\mathrm{load}, P_\mathrm{in})$, no product prior
7	**Check**	Step 1 fails, so the exact expression stands; the reading may still be close
	Class	Posterior of a parameter

$$p(\eta \mid y_\mathrm{ut}, y_\mathrm{ups}) \propto \int \mathcal{N}(y_\mathrm{ut};\, P_\mathrm{load}/\eta,\, 3^2)\, \mathcal{N}(y_\mathrm{ups};\, P_\mathrm{load},\, 2^2)\, p(P_\mathrm{load})\, \mathrm{d}P_\mathrm{load}\; p(\eta)$$

10.6 The exercises

The solutions are in Appendix A, and the facts of the three systems in Appendix D. On the two-hall plant, E is the 40 kW trip test.

Exercise 10.1 (*formulate*). A commissioning note says: "The hall temperatures are logged during the trip test. The heat through the opening follows the drawing's conductance, and any leakage through the walls is small. The test load is held at 40 kW. We want the conductance." Sort every quantity in the note by role, write the physical and measurement relations, and label each assumption.

Exercise 10.2 (*formulate*). A change request on the analysis reads: "Re-parameterize the opening model so that the numbers are better." Formulate the request.

Exercise 10.3 (*formulate*). Run this one: given E, unit B out, settled, what is the probability that hall B can carry 50 kW within 8 K, read in R?

Exercise 10.4 (*formulate*). Run this one: a timed discharge at 20 kW reads 1.80 h with an illustrative spread of 0.03 h. What is the probability that the string gives at least 1.75 h?

Exercise 10.5 (*formulate*). On the network, the utility meter reads the UPS's input while the load the UPS delivers is known. Write the measurement relation for the efficiency η, say in which coordinate it is linear, and write what must be done with the fact sheet's prior on η under the degraded explanation before a Gaussian reading in that coordinate.

Exercise 10.6 (*formulate*). The maintenance planner writes a unit's total downtime in a record as its number of repairs times the mean repair μ. Write the relation in μ and in the repair rate $1/\mu$, say in which it is linear, and state what must go with a prior on $\log \mu$ if the coordinate changes.

Exercise 10.7 (*formulate*). The owner asks: "What is the thermal resistance of the opening, given the trip test?" Formulate the question.

Exercise 10.8 (*formulate*). A ticket reads: "Report the opening's adequacy from the trip test, in the right coordinate." Formulate the ticket.

Exercise 10.9 (*formulate*). A review note on the network says: "The utility meter and the UPS output meter are read in the same hour. The efficiency is taken as 0.96 from the nameplate. We want to know how much of the input is lost, and the PDU meter's reading looks high." Sort every quantity by role, write the relation for the loss, and label each assumption.

Exercise 10.10 (*formulate*). The string's discharge efficiency of 0.90 is doubted. Write the relation of a timed discharge at 20 kW with the efficiency as a parameter beside φ, say in which coordinate it is linear, and say what one reading can determine.

Exercise 10.11 (*reverse*). Write one plain sentence for

$$P(60\,\mathrm{kW} \times R > 8\,\mathrm{K} \mid \text{unit B out},\ Q_{\mathrm{hall}},\ \text{settled},\ E).$$

Exercise 10.12 (*reverse*). A 1 kW test heater is held in the battery pack until the pack is steady, and y is a reading of its rise above the room, with R_{p} the pack-to-room thermal resistance in K/kW. Write one plain sentence for

$$p(R_{\mathrm{p}} \mid y) \propto \mathcal{N}(y;\, 1\,\mathrm{kW} \times R_{\mathrm{p}},\, (1\,\mathrm{K})^2)\ R_{\mathrm{p}}^{-1}\ \mathcal{N}(\log R_{\mathrm{p}};\, \log 16.7,\, \sigma_0^2),$$

and state what the factor R_{p}^{-1} does.

Exercise 10.13 (*reverse*). Write one plain sentence for $y = 2.025\ \mathrm{h} \times \varphi \times 0.70^{\,v} + \varepsilon$, with $v = 1$ if a cell is weak and $v = 0$ if not, naming each symbol's role and saying in which unknown the relation is linear.

Exercise 10.14 (*reverse*). Write one plain sentence for $\mathcal{N}(y_{\mathrm{it}} + y_{\mathrm{cl}};\, P_{\mathrm{load}},\, (3\,\mathrm{kW})^2 + (3\,\mathrm{kW})^2)$, read as a function of P_{load}, and name what it assumes of the two meters.

Exercise 10.15 (*critique and repair*). A report says: "The Gaussian reading in R has mode 0.1172 K/kW and spread 0.0178 K/kW. Inverting, the conductance is 8535 W/K, plus or minus 1/0.0178, about 56 W/K, so it clears 7500 W/K beyond doubt." State what is wrong, and repair it.

Exercise 10.16 (*critique and repair*). A colleague moves the fact sheet's prior on the efficiency under the degraded explanation, 0.94 ± 0.04, to $1/\eta$, and writes it "normal, center 1/0.94, spread 0.04, since the shift is small". State the error, and repair it.

Exercise 10.17 (*critique and repair*). A timed discharge reads 1.95 h with an illustrative spread of 0.03 h. A note says: "That is $(1.95 - 1.72)/0.03 = 7.7$ spreads from the prior's bound, so the Gaussian reading in φ passes, and the string gives its 2 h with probability $\Phi((1.95 - 2)/0.03) = 0.048$." State the error, and repair it.

Exercise 10.18 (*critique and repair*). A report says: "The reading in $\log G$ gives 0.767 and the reading in R gives 0.818, so we report their average, 0.79, as the best estimate that the opening is adequate." State what is wrong, and repair it.

Exercise 10.19 (*critique and repair*). A commissioning note says: "With the drawing's 8250 W/K and the 4.7828 K measured, the opening passes 39.5 kW of the 40 kW, so the heat through the walls is 0.5 kW, and we carry that value." State the errors, and repair it.

Exercise 10.20 (*rapid*). Name each quantity's role in a question about hall B tonight, or say that it has none.

(a) Unit A's 130 kW capacity, taken from its nameplate.

(b) Hall A's set point, held at 22 °C by the operator.

(c) The test load of a trip test run last month.

(d) The duty of unit A tonight.

(e) Whether unit B is out tonight, before anyone has looked.

(f) Tomorrow's outage length at the string.

(g) A reading of the outside air temperature taken tonight.

(h) The coordinate of a Gaussian reading.

(i) Whether the drawing or the as-built opening is the one in service.

(j) "The drawing's 8250 W/K".

10.7 What to retain

- Every symbol has a role in a question: condition, parameter, variable, observation, discrepancy, or stated.
- The physical relation and the measurement relation are written separately, and a discrepancy is carried or labeled as dropped.
- One model has several coordinates. A change of variable carries its Jacobian, and the probability of one event does not move.
- A Gaussian reading depends on its coordinate. Pick the coordinate in which the observation relation is linear, transform the prior, and run the checklist: a fixed spread, a prior of one width and one hump, 3 spreads inside the prior's range, straightness at most 3/4. If any fails, write the exact expression and name the step.
- On one reading: 0.818 exact, 0.767 in $\log G$, 0.818 in R. On a weak test: 0.679 exact, 0.955 in R.

Lesson 11

From Measurements to Evidence

11.1 The engineering question

The test engineer quotes the sensors' data sheet at the handover. Each hall sensor has a noise of 0.3 K, thirty samples were averaged, and so the difference between the halls is known to 0.08 K. On that spread the opening clears its threshold beyond any doubt. The commissioning lead objects that the sensors were never calibrated against each other, that each could be off by half a kelvin for the whole test, and that with that in the sum the same test gives about four chances in five. The reliability engineer, down the table, has seven failures of unit A in a three-year log and writes "2.33 per year" on the board as if it were the rate.

Each of them has read a number off an instrument or a log and called it the quantity. A reading is not the quantity. It is the quantity plus what the instrument adds, and what a reading says about the plant depends on writing that sum. A log is the same: a count says something about a rate only together with the exposure it was counted over.

This lesson writes each reading as a sum, names what does not average out, and writes a record as a count with its exposure.

11.2 The formulation

The work sits in step 3, where a reading is told from the quantity it reads, and in step 6, where the evidence is tied to the unknowns. Lesson 7 defined the likelihood as the probability of the evidence given the parameters, read as a function of them, and this lesson writes one from an instrument's statements. On the two-hall plant E stays the 40 kW trip test; the plant's outage record is another piece of evidence, and it is written "given the record".

11.2.1 A reading as a sum

A reading is a quantity plus an offset plus noise. The *quantity* is what the model calls for. The *noise* differs from sample to sample and averages down with the number of samples. The *offset* stays the same over a test and does not average down at all. A hall sensor's mean over a test is written

$$y_i = \vartheta_i + b_i + \bar{\varepsilon}_i, \qquad i \in \{A, B\}, \tag{11.1}$$

with b_i the sensor's offset and $\bar{\varepsilon}_i$ the mean of its noise over the samples. Instruments rarely state the sum in these terms, and Table 11.1 takes seven statements of the kind a data sheet or a test plan makes, one a row, with the term each supplies and what it leaves open.

Table 11.1: What an instrument's statement supplies to a reading's sum.

Statement	Term it supplies	Left open
"Noise 0.3 K per sample"	Noise; averages as one over the root of the number of samples	Independence from sample to sample
"Mean of 30 samples"	The number of samples	Which noise it divides
"2 % of reading"	Noise proportional to the quantity	Per sample or per test
"Calibrated at the factory"	An offset's spread, if a certificate states one	The spread, and the time since
"Accuracy ±0.5 K"	Not determined: an offset, a noise or a tolerance band	Which of them, and what the 0.5 is a width of
"Resolution 1 kW"	Noise, of a rounded kind	Whether it is small beside the rest
No statement	Nothing; a labeled assumption	The whole term

11.2.2 Nuisance variables

An offset is an unknown nobody is asking about, and it stays in the model all the same: a *nuisance variable*. It has two treatments, and the sheet says which. It may be *carried*, as a parameter with a prior, in a model of all the unknowns. It may be *integrated out*,

$$p(d \mid G) = \int p(d \mid G, b)\, p(b)\, \mathrm{d}b, \tag{11.2}$$

which for normal terms leaves a normal with the variances added. The two treatments ask the same question of the same model; the report's whole model carries its offsets, and the first worked formulation below integrates them. Dropping a nuisance variable is a third thing. It is an assumption, that the variable is zero, and it is labeled as one. What counts as nuisance depends on the question: the PDU meter's offset is nuisance for a question about the UPS's efficiency, and it is the target of a question about the meter.

11.2.3 The trip test's reading

The first worked formulation derives the spread that Lesson 10 took as stated. With unit B out and the plant settled, the measured difference is

$$d = y_B - y_A = \frac{Q}{G} + (b_B - b_A) + (\bar{\varepsilon}_B - \bar{\varepsilon}_A). \tag{11.3}$$

The fact sheet gives each offset a prior spread of 0.5 K and each sensor a noise of 0.3 K per sample over 30 samples. Integrating out the two offsets leaves d normal about Q/G, with variance

$$\sigma_d^2 = 2\,(0.5\,\mathrm{K})^2 + \frac{2\,(0.3\,\mathrm{K})^2}{30} = (0.7113\,\mathrm{K})^2, \tag{11.4}$$

which is the report's floor on a measured difference. Table 11.2 is the sheet, and Figure 11.1 shows what the sum is made of.

The sheet's seventh row is the point of writing a likelihood at all. The expression says nothing yet about what the conductance probably is, and the prior of the next lesson is what turns it into that. In Figure 11.1 the three bars are the noise of the difference, the offsets and both together: the offsets are nearly the whole of σ_d.

Table 11.2: The formulation sheet filled in for the likelihood of the trip test.

Field	Entry
Intention	Estimate: the likelihood serves the posterior on G
1. Sentence	With unit B out, the plant settled and a test load of 40 kW, how probable is the measured difference, for each value of the conductance?
2. Target	The likelihood $p(d \mid G)$, a function of G
3. Known	The test load; unit B out; settled. Stated: the sensors' noise, the number of samples, and each offset's prior spread. Assumptions: $\delta = 0$; the two offsets independent; the noise independent from sample to sample.
4. Event	None: a density of a reading
5. Expression	Equation (11.2), with $b = b_B - b_A$
6. Model	Equation (11.3); a normal prior on each offset; normal noise
7. Read back	How well each conductance explains the reading. It is not a probability of the conductance, and it stays the same whatever is believed about G.
Problem class	Likelihood, nuisance integrated out

As a run, ending in the normal it reduces to:

Run, the likelihood of the trip test

1	**Sentence**	Unit B out, settled, $Q = 40\,\mathrm{kW}$: how probable is d, for each G?
	Intention	Estimate: the likelihood serves the posterior
2	**Target**	The likelihood $p(d \mid G)$, a function of G
	Declare	y_A, y_B means [K]; $b = b_B - b_A$, the offsets [K]; noise 0.3 K per sample; $n = 30$
3	**Known**	Q; unit B out; settled; offsets' spread 0.5 K; $\delta = 0$ (assumption)
	Random here	d, through the offsets and the noise; G held at each value
4	**Event**	None: a density of a reading
5	**Expression**	$p(d \mid G) = \int p(d \mid G, b)\, p(b)\, \mathrm{d}b$
	Expand	$d = y_B - y_A$, with each reading as quantity + offset + noise
6	**Model**	Offsets normal, independent (assumption); noise independent across samples
	Reduce	$\mathcal{N}(d;\, Q/G,\, \sigma_d^2)$, $\sigma_d^2 = 2(0.5)^2 + 2(0.3)^2/30$
7	**Check**	A function of G, not a probability of G
	Class	Likelihood, nuisance integrated out

$$\boxed{p(d \mid G) = \mathcal{N}(d;\, Q/G,\, (0.7113\,\mathrm{K})^2)}$$

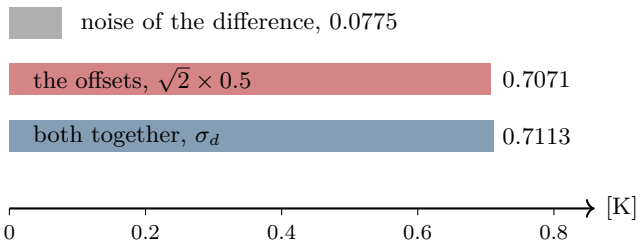


Figure 11.1: The spread of the measured difference and what it is made of. The noise is a tenth of the total, so more samples buy almost nothing.

11.2.4 The five meters

The second worked formulation is on the electrical network, whose meters carry a different difficulty: a nuisance under one explanation is the question under another. Under each explanation H_i of the fact sheet, the five readings are the sums

$$y_{\mathrm{ut}} = \frac{P_{\mathrm{load}}}{\eta} + \varepsilon_1, \quad y_{\mathrm{ups}} = P_{\mathrm{load}} + \varepsilon_2, \quad y_{\mathrm{pdu}} = P_{\mathrm{load}} + b_{\mathrm{pdu}} + \varepsilon_3, \quad y_{\mathrm{it}} = P_{\mathrm{it}} + \varepsilon_4, \quad y_{\mathrm{cl}} = P_{\mathrm{cl}} + \varepsilon_5, \tag{11.5}$$

with $P_{\text{load}} = P_{\text{it}} + P_{\text{cl}}$ and the meters' noises as stated, 3, 2, 2, 3 and 3 kW. The efficiency and the offset have stated spreads that differ by explanation, so the likelihood of one explanation integrates over them. The scenario truths of the fact sheet, 0.91 and +42 kW, appear nowhere in the expression, since a likelihood is written from what each explanation assumes. Table 11.3 is the sheet.

Table 11.3: The formulation sheet filled in for the five meters under one explanation.

Field	Entry
Intention	Diagnose: the likelihood serves the comparison of the explanations
1. Sentence	How probable are the five readings under each explanation of a mismatch?
2. Target	The likelihood $p(\mathbf{y} \mid H_i)$, one number per explanation
3. Known	A steady snapshot; the meters' noises; each explanation's stated spreads on η and on b_{pdu}. Assumption: the noises are independent.
4. Event	None: a density of five readings
5. Expression	$\int p(\mathbf{y} \mid \eta, b_{\text{pdu}}, P_{\text{it}}, P_{\text{cl}})\, p(\eta, b_{\text{pdu}} \mid H_i)\, p(P_{\text{it}}, P_{\text{cl}})\, \mathrm{d}\eta\, \mathrm{d}b_{\text{pdu}}\, \mathrm{d}P_{\text{it}}\, \mathrm{d}P_{\text{cl}}$
6. Model	Equations (11.5); the fact sheet's spreads under each explanation; a prior on the loads (a choice)
7. Read back	One number for each explanation. A comparison also needs the explanations' priors, which Lesson 19 takes up.
Problem class	Likelihood of an explanation, nuisance integrated out

The same, in a run:

Run, the meters under one explanation

1	**Sentence**	Five readings, explanation H_i: how probable are they?
	Intention	Diagnose: the likelihood serves the comparison
2	**Target**	The likelihood $p(\mathbf{y} \mid H_i)$, one number per explanation
	Declare	$\mathbf{y}$ the five readings [kW]; η [–]; b_{pdu} [kW]; $P = (P_{\text{it}}, P_{\text{cl}})$ [kW]; H_i
3	**Known**	Noises 3, 2, 2, 3, 3 kW; a steady snapshot; independent noises (assumption)
	Random here	$\mathbf{y}$, through η, b_{pdu}, the loads and the noises; H_i held
4	**Event**	None: a density of five readings
5	**Expression**	$\int p(\mathbf{y} \mid \eta, b_{\text{pdu}}, P)\, p(\eta, b_{\text{pdu}} \mid H_i)\, p(P)\, \mathrm{d}\eta\, \mathrm{d}b_{\text{pdu}}\, \mathrm{d}P$
6	**Model**	The five sums; η, b_{pdu} as the fact sheet states under each H_i; $p(P)$ (choice)
7	**Check**	One number per explanation; a comparison also needs their priors
	Class	Likelihood of an explanation, nuisance integrated out

$$\boxed{p(\mathbf{y} \mid H_i)}$$

11.2.5 A plant's own record

A plant's own outage record is evidence, and it has a likelihood like any other reading. Unit A's three-year record holds 7 failures over 26280 h in service, with 1260 h down over 7 repairs. A count needs its *exposure*, the time at risk over which it was counted, on a stated clock, and here the clock is hours in service. With the count Poisson about the rate times the exposure, the record's likelihood for the rate is

$$P(k \mid \lambda_A) = \frac{(\lambda_A t_{\text{s}})^k}{k!}\, e^{-\lambda_A t_{\text{s}}}, \qquad k = 7, \quad t_{\text{s}} = 3 \text{ yr in service.} \tag{11.6}$$

A record of an outcome is a reading too: an outage of length D_{out} carried to its end says that $T_{\text{rt}} \geq D_{\text{out}}$, with the likelihood $\mathbf{1}[T_{\text{rt}} \geq D_{\text{out}}]$. Table 11.4 writes one question on the count. It is "given the record" throughout, since on this plant E is the trip test, which says nothing about a failure rate. The stated prior on $\log \lambda_A$, centered at 2.5 per year in service with a spread of a factor of 3, is used here as stated; where it comes from is Lesson 12's question.

Table 11.4: The formulation sheet filled in for unit A's rate, given its record, with a run for the rate and one for the repairs.

Field	Entry
Intention	Estimate
1. Sentence	Given unit A's three-year record and the stated prior, what is the probability that its failure rate is below 2.5 per year in service?
2. Target	A parameter, λ_A, through an event on it
3. Known	The record, as a count with its exposure. Assumption: a constant rate over the record.
4. Event	$\lambda_A < 2.5$ per year in service
5. Expression	$P(\lambda_A < 2.5 \mid \text{the record})$
6. Model	Equation (11.6); the stated prior
7. Read back	A property of unit A, on service time. Silent on next year's count and on the trip test.
Problem class	Posterior exceedance

Then the same record twice, once per question:

Run, unit A's rate, given its record

1	**Sentence**	Given unit A's record and the stated prior: probability its rate is below 2.5?
	Intention	Estimate
2	**Target**	A parameter, λ_A, per year in service
	Declare	$k = 7$ failures; $t_s = 3$ yr in service; λ_A [1/yr]; prior: $\log \lambda_A$, center 2.5, a factor of 3
3	**Known**	The record, a count with its exposure; a constant rate (assumption)
	Random here	λ_A, one unknown value; k is in hand
4	**Event**	$\lambda_A < 2.5$
5	**Expression**	$P(\lambda_A < 2.5 \mid \text{the record})$
6	**Model**	$p(\lambda_A \mid \text{the record}) \propto P(k = 7 \mid \lambda_A)\ p(\lambda_A)$, a Poisson count on service time
7	**Check**	A property of unit A; not next year's count; not given E
	Class	Posterior exceedance

$$\boxed{P(\lambda_A < 2.5 \mid \text{the record})}$$

Run, unit A's mean repair, given its record

1	**Sentence**	Given 7 repairs totaling 1260 h: what is known of unit A's mean repair?
	Intention	Estimate
2	**Target**	A parameter, μ_A
	Declare	μ_A mean repair [h]; 7 repairs; 1260 h total; prior: $\log \mu_A$, center 180 h, a factor of 3
3	**Known**	The record's repairs and their total; each repair finished (assumption)
	Random here	μ_A, one unknown value; the lengths are in hand
4	**Event**	None: a density
5	**Expression**	$p(\mu_A \mid \text{the record})$
6	**Model**	Exponential lengths (illustrative): $p(\text{the lengths} \mid \mu_A) = \mu_A^{-7} e^{-1260/\mu_A}$
7	**Check**	A repair still open at the record's end would be censored, and changes this line
	Class	Posterior of a parameter

$$\boxed{p(\mu_A \mid \text{the record}) \propto \mu_A^{-7} e^{-1260/\mu_A}\ p(\mu_A)}$$

A count alone gives a center and not a rate, and Figure 11.2 sets the point k/t_s beside what the record leaves known. The three bars are unit A's posterior bands, 5 % to 95 %, for records of one, three and ten years in service.

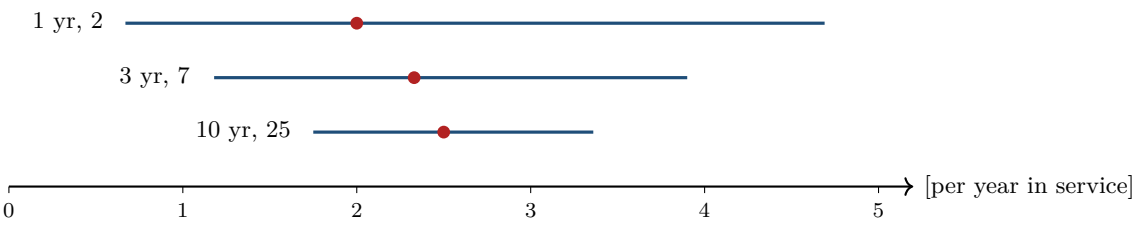


Figure 11.2: What a record leaves known of unit A's rate: the posterior band from 5 % to 95 % on the stated prior, for records of three lengths, each labeled with its exposure and its count, and the point k/t_{s} marked.

11.2.6 A sensor statement that is not ready

The test plan says: "The sensor is accurate to ±0.5 K." The sentence does not say which term of the sum it supplies, and it fails at step 1. Read as a constant offset with a spread of 0.5 K, the report's reading, it gives $\sigma_d = 0.711\,\mathrm{K}$. Read as a noise of 0.5 K per sample, it gives about 0.13 K over 30 samples. Read as a hard band of ±0.5 K, it gives a bounded prior on each offset, uniform in form, and about 0.42 K on the difference. Read as a 95 % band, it gives about 0.37 K on the difference. The readings part by a factor of five, and the analyst asks which is meant and asks for the calibration record.

Key idea

A reading is the quantity plus an offset plus noise. Noise averages and an offset does not. A nuisance variable is carried or integrated out, and the sheet says which. A likelihood is a function of the unknown, not a probability of it, and a record is a count with its exposure, on a stated clock.

11.3 The competing formulations

Each competing formulation below reads a number off an instrument or a log and leaves out what the instrument or the log adds. The first is the shortest route from a reading to a verdict.

Run, the reading taken for the quantity

1	**Sentence**	Given E: is the opening adequate?
5	**Expression** ×	$G = Q/d = 40\,\mathrm{kW}/4.7828\,\mathrm{K} = 8363\,\mathrm{W/K}$, above G^*: adequate
	Repair	$P(G \geq G^* \mid E)$: 0.818
7	**Check**	A second test would read otherwise; the first number leaves no room for that

$$\boxed{P(G \geq G^* \mid E)}$$

The second keeps the noise, which the data sheet states, and drops the offsets, which it does not. This is the test engineer's sentence, and it has the largest size in the lesson.

Run, the noise kept, the offsets dropped

3	**Known**	E; noise 0.3 K per sample, 30 samples
6	**Model** ×	d normal about Q/G, spread 0.0775 K: $P(G \geq G^* \mid d) = 1.000$
	Repair	Offsets integrated out, spread 0.7113 K: 0.818
7	**Check**	Would a sensor swap move the reading? Then the offsets belong in the sum

$$\boxed{P(G \geq G^* \mid E)}$$

The third accepts the offset's existence and reads it as noise, so that more samples seem to

help.

Run, an offset read as noise

1	**Sentence**	"Take 3000 samples and the spread will fall tenfold."
6	**Model** ×	σ_d scaled by one over the root of the number of samples
	Repair	Only the noise scales: 0.7113 K at 30 samples, 0.7071 K at 3000
7	**Check**	An offset does not average; a calibration or a swap changes the model instead

$$\boxed{\sigma_d^2 = 2(0.5\,\mathrm{K})^2 + 2(0.3\,\mathrm{K})^2/n}$$

The fourth reads a likelihood as a statement about the parameter, Lesson 3's transposition in this lesson's setting.

Run, a likelihood read as a posterior

1	**Sentence**	"The reading is most likely at 8363 W/K, so the opening is probably that."
5	**Expression** ×	The peak of $p(d \mid G)$, read as what G probably is
	Repair	$p(G \mid E) \propto p(d \mid G)\, p(G)$; median 8534 W/K
7	**Check**	Which side of the bar is the conductance on?

$$\boxed{p(G \mid E)}$$

The fifth reads a count as a rate. Figure 11.2 showed what a band says that the point does not.

Run, a count read as a rate

1	**Sentence**	Given unit A's record: what is its failure rate?
5	**Expression** ×	$\lambda_A = k/t_\mathrm{s} = 2.33$, a value, no likelihood
	Repair	$p(\lambda_A \mid \text{the record})$; band 1.18 to 3.90 per year in service
7	**Check**	Seven failures in three years could come from a rate of 1.5 or of 3.5

$$\boxed{p(\lambda_A \mid \text{the record}) \propto P(k \mid \lambda_A)\, p(\lambda_A)}$$

The sixth counts over the wrong clock. The book's rates are per year in service, and a calendar exposure holds the hours of repair as well.

Run, the exposure on the wrong clock

3	**Known**	$k = 7$; 26280 h in service; 1260 h down
6	**Model** ×	Exposure 3.14 calendar years: 2.23 per year, quoted as a rate in service
	Repair	Exposure 3 years in service: 2.33 per year in service
7	**Check**	Which clock does the rate use, and does the repair's time count?

$$\boxed{P(k \mid \lambda_A) \text{ with } t_\mathrm{s} = 3 \text{ yr in service}}$$

The seventh changes the unit of exposure altogether. A failure to pick up the load is counted over demands, and an hours-based rate has no meaning for it.

Run, a demand count on a time clock

1	**Sentence**	Illustrative log, 1 failure in 40 demands in 5 years: probability of failing on demand?
6	**Model** ×	Poisson, rate 0.2 per year in service: a rate of pick-up failures in time
	Repair	Binomial: 1 of 40 demands; the exposure is the demands
7	**Check**	Would a quiet year with no demand count as a success? It would not

$$\boxed{p(p_\mathrm{d} \mid \text{the log}) \propto p_\mathrm{d}(1 - p_\mathrm{d})^{39}\, p(p_\mathrm{d})}$$

The eighth writes a record under the wrong evidence, a slip the notation invites on this

plant.

Run, the record written under E

1	**Sentence**	Given unit A's record: is its rate below 2.5?
5	**Expression** ×	$P(\lambda_A < 2.5 \mid E)$
	Repair	$P(\lambda_A < 2.5 \mid \text{the record})$
7	**Check**	E is the 40 kW trip test, and it says nothing of a failure rate

$$\boxed{P(\lambda_A < 2.5 \mid \text{the record})}$$

These are the forms of this lesson's failure class, the *reading taken for the quantity*: a reading or a log entry used as if it were the model's quantity, with what the instrument or the record adds left out. Its four forms are the reading for the quantity; the noise kept and a nuisance variable dropped; an offset read as noise that averages; and a count read as a rate, its exposure dropped or on the wrong clock. The class enters at step 5 or step 6, and it is caught at step 7 by asking whether another reading tomorrow would have to agree with the number. A likelihood read as a posterior is the wrong side of the bar in a new setting.

11.4 One system, many questions

Table 11.5 gives ten things the electrical network's meters and logs supply, with the relation or likelihood each enters and what is nuisance in it. It holds no values.

Table 11.5: Ten readings and logs on the network, with their relation and what is nuisance in each.

Reading or log	Relation or likelihood	Nuisance
The utility meter, for the efficiency	$y_{\text{ut}} = P_{\text{load}}/\eta + \varepsilon_1$	P_{load}, unless read
The PDU meter, for its own calibration	$y_{\text{pdu}} = P_{\text{load}} + b_{\text{pdu}} + \varepsilon_3$	P_{load}; the offset is the target
The PDU meter, for the load	The same	The offset
IT plus cooling, against the UPS meter	Two readings of one load	Each meter's error
A meter replaced mid-year	Two offsets, one a meter	Each offset
A one-year log of feeder interruptions	A count with exposure, on a stated clock	None
A log with an unknown start	A count with its exposure unknown	The start
Pick-up failures among demands	A count over demands	None
A battery test log	A reading of autonomy at its load	The test load, if not recorded
"The meter says 1198 kW"	A reading, not the input	Its noise, and any offset

The second and third rows are one meter and two questions, and the nuisance changes with the question. The last row is the failure class at its shortest.

11.5 Worked problems

Seven problems on the three systems follow, each with its run.

Problem 11.1. The string's battery management system displays "state of health 92 %" (illustrative). Write the likelihood for φ.

Run, a display that may not be a reading

1	**Sentence**	"State of health 92 %": the likelihood for φ?
	Intention	Not stated: how the number is made is unstated
3	**Known**	Missing: whether anything was measured; its error; when it was last updated
5	**Expression**	(a) noise about φ; (b) an offset as well; (c) a vendor's formula: no relation to φ
6	**Model**	(a) $0.92 = \varphi + \varepsilon$; (b) $0.92 = \varphi + b + \varepsilon$; (c) a prior on φ, by Lesson 12's rule
7	**Check**	Did this string's energy enter the number? The analyst asks how it is made
	Class	Not ready

$$\boxed{\text{(a), (b) } p(0.92 \mid \varphi); \quad \text{(c) } p(\varphi)}$$

Problem 11.2. The 40 kW trip test is repeated with the two hall sensors exchanged (illustrative). Write the likelihood of the two measured differences for G.

Run, the sensors exchanged

1	**Sentence**	Two tests at 40 kW, the sensors exchanged: how probable, for each G?
	Intention	Estimate
2	**Target**	The likelihood $p(d_1, d_2 \mid G)$
	Declare	d_1, d_2 [K]; $b = b_B - b_A$ [K]; each test's noise 0.0775 K
3	**Known**	Q; unit B out; settled; $\delta = 0$; the offsets travel with the sensors (assumptions)
	Random here	The two readings; b, integrated out
4	**Event**	None; a density
5	**Expression**	$\int \mathcal{N}(d_1;\, Q/G + b,\, 0.0775^2)\, \mathcal{N}(d_2;\, Q/G - b,\, 0.0775^2)\, p(b)\, \mathrm{d}b$
6	**Model**	The offset changes sign with the exchange; the noise differs from test to test
	Reduce	The mean of the two, normal about Q/G, spread 0.0548 K; their difference reads b
7	**Check**	0.055 K for 0.711 K: a change of model, not more samples
	Class	Likelihood, nuisance integrated out

$$\boxed{p(d_1, d_2 \mid G) \propto \mathcal{N}((d_1 + d_2)/2;\, Q/G,\, 0.0548^2)}$$

Problem 11.3. The battery room's log holds two grid outages (illustrative): one of 0.4 h, and one still running at 1 h when the log ended. Write the likelihood for the mean outage length.

Run, a length still open

1	**Sentence**	One outage of 0.4 h, one still running at 1 h: how probable, for each mean length?
	Intention	Estimate
2	**Target**	The likelihood in μ_o, the mean outage length [h]
	Declare	μ_o [h]; 0.4 h a finished length; 1 h a lower bound on the second
3	**Known**	The log; exponential (illustrative); its end unrelated to the outage (assumption)
	Random here	The lengths; μ_o held at each value
4	**Event**	None; a density and a probability
5	**Expression**	$p(0.4 \text{ h} \mid \mu_\mathrm{o})\ P(D_\mathrm{out} > 1 \text{ h} \mid \mu_\mathrm{o})$
6	**Model**	$p(x \mid \mu_\mathrm{o}) = e^{-x/\mu_\mathrm{o}}/\mu_\mathrm{o}$; $P(D_\mathrm{out} > x \mid \mu_\mathrm{o}) = e^{-x/\mu_\mathrm{o}}$
7	**Check**	The open outage adds exposure and no count; dropping it would favor short lengths
	Class	Likelihood

$$\boxed{\mu_\mathrm{o}^{-1}\, e^{-1.4 \text{ h}/\mu_\mathrm{o}}}$$

Problem 11.4. A UPS battery test at half load, 575 kW, reads 22 min with a spread of 1 min (illustrative). Write the likelihood for the autonomy at the full 1150 kW.

Run, a reading of another quantity

1	**Sentence**	A half-load test reads 22 min: how probable, for each full-load autonomy?
	Intention	Estimate
2	**Target**	The likelihood $p(y \mid t_{\text{aut}})$, a function of t_{aut}
	Declare	y [min]; t_{aut} at 1150 kW [min]; 1 min the reading's spread (illustrative)
3	**Known**	The test load; the energy the same at both loads (assumption)
	Random here	y; t_{aut} held at each value
4	**Event**	None; a density
5	**Expression**	$p(y \mid t_{\text{aut}})$
6	**Model**	$y = 2\,t_{\text{aut}} + \varepsilon$: the assumption is what ties the two loads
7	**Check**	The reading is not of the autonomy; a real battery gives less at a higher rate
	Class	Likelihood

$$\boxed{p(y \mid t_{\text{aut}}) = \mathcal{N}(y;\, 2\,t_{\text{aut}},\, (1\ \text{min})^2)}$$

Problem 11.5. Unit A's three-year record holds 7 failures in service and 7 repairs totaling 1260 h. Given it, what is the probability that the unit's long-run share of time down exceeds 0.05?

Run, one record, two parameters

1	**Sentence**	Given unit A's record: probability its share of time down exceeds 0.05?
	Intention	Estimate
2	**Target**	A derived parameter, q_A, of two others
	Declare	λ_A [1/yr]; μ_A [h]; $q_A = \lambda_A\mu_A/(8760\ \text{h} + \lambda_A\mu_A)$; priors on $\log\lambda_A$, $\log\mu_A$
3	**Known**	The count with its exposure; the repairs' total; the fact sheet's priors
	Random here	λ_A and μ_A
4	**Event**	$q_A > 0.05$
5	**Expression**	$P(q_A > 0.05 \mid \text{the record})$
6	**Model**	$P(k = 7 \mid \lambda_A)\,p(\lambda_A)\ \mu_A^{-7} e^{-1260/\mu_A}\,p(\mu_A)$; exponential lengths (illustrative)
7	**Check**	Two factors of one record; the two parameters stay independent on these priors
	Class	Posterior exceedance

$$\boxed{P(q_A > 0.05 \mid \text{the record})}$$

Problem 11.6. Suppose the trip test's load is not set but read from the PDU meter, with a noise of 1 %. Write the likelihood of the measured difference for G.

Run, a condition that becomes a nuisance

1	**Sentence**	Load read at 1 %, not set: how probable is d, for each G?
	Intention	Estimate
2	**Target**	The likelihood $p(d, y_{\text{Q}} \mid G)$
	Declare	y_{Q} the PDU reading [kW]; Q the true load [kW]; $\sigma_d = 0.711$ K; $p(Q)$ (choice)
3	**Known**	Unit B out; settled; $\delta = 0$; the load steady over the test (assumptions)
	Random here	d, y_{Q}; Q, integrated out
4	**Event**	None; a density
5	**Expression**	$\int \mathcal{N}(d;\, Q/G,\, \sigma_d^2)\,\mathcal{N}(y_{\text{Q}};\, Q,\, (0.01\,Q)^2)\,p(Q)\,\mathrm{d}Q$
6	**Model**	The two relations; the load's error reaches d through Q/G
7	**Check**	1 % of 4.78 K is 0.048 K beside 0.711 K: droppable, with a label
	Class	Likelihood, nuisance integrated out

$$\boxed{p(d, y_{\text{Q}} \mid G) = \int \mathcal{N}(d;\, Q/G,\, \sigma_d^2)\,\mathcal{N}(y_{\text{Q}};\, Q,\, (0.01\,Q)^2)\,p(Q)\,\mathrm{d}Q}$$

Problem 11.7. The PDU meter reads 1150 kW in one snapshot and 1180 kW in a later one (illustrative). Write the likelihood for the rise in load if (a) the meter is the same, and (b) it was replaced between them.

Run, an offset that cancels, or does not		
1	**Sentence**	A rise on the PDU meter: (a) one meter; (b) replaced between: the likelihood?
	Intention	Estimate
2	**Target**	The likelihood in the rise, ΔP [kW]
	Declare	y_1, y_2 [kW]; ε_a, ε_b the noises, 2 kW; b_1, b_2 offsets [kW], spread 1.5 kW each
3	**Known**	Two snapshots; the offsets' stated spread; (a) one offset; (b) two (assumption)
	Random here	$y_2 - y_1$; the offsets, integrated out
4	**Event**	None; a density
5	**Expression**	$p(y_2 - y_1 \mid \Delta P)$
6	**Model**	(a) $y_2 - y_1 = \Delta P + \varepsilon_\mathrm{b} - \varepsilon_\mathrm{a}$; (b) the same, plus $b_2 - b_1$
	Reduce	Normal about ΔP: (a) spread 2.83 kW; (b) spread 3.54 kW
7	**Check**	Whether an offset is nuisance at all depends on the question
	Class	Likelihood, nuisance integrated out
		(a) $\mathcal{N}(y_2 - y_1;\, \Delta P,\, 2.83^2)$; (b) $\mathcal{N}(y_2 - y_1;\, \Delta P,\, 3.54^2)$

11.6 The exercises

The solutions are in Appendix A, and the facts of the three systems in Appendix D. On the two-hall plant, E is the 40 kW trip test and an outage record is written "given the record".

Exercise 11.1 (*formulate*). The commissioning specification reads: "Each hall has one calibrated PT100, logged once a minute for 30 minutes. The surviving unit's duty is read from its own meter, within 2 %. PDU power is metered within 1 %." Write each reading as a sum, name the nuisance variables, write the likelihood of one test given G and the loads, and say what the specification leaves unsaid.

Exercise 11.2 (*formulate*). Run this one: unit A's duty meter reads 81.6 kW with a noise of 2 %, and the two halls' loads are unknown. Write the likelihood for their total.

Exercise 11.3 (*formulate*). The test plan says: "Use the floor for the spread of the measured difference." Formulate the line.

Exercise 11.4 (*formulate*). On the network, a reader of the PDU meter wants the UPS's efficiency. Write the likelihood for the efficiency under the meter-offset explanation, and say what is nuisance in it.

Exercise 11.5 (*formulate*). The string's state of charge at an outage's start is read with an error of spread 0.03. Write the event "the string carries the next outage" and its probability given that reading, and say what the reading is a reading of.

Exercise 11.6 (*formulate*). The network's commissioning sheet has the line: "UPS output meter within 0.5 %, verified at the last shutdown." Write the meter's term in the likelihood of y_ups.

Exercise 11.7 (*formulate*). The planner asks: "What does the feeder's one-year count, 4 failures, say about its repair time?" Formulate the question.

Exercise 11.8 (*formulate*). Run this one: the UPS output meter and the two branch meters are read in one snapshot. Write the likelihood for the total load from all three.

Exercise 11.9 (*formulate*). The planner puts the two units' three-year records together: 14 failures in 6 unit-years in service. Write the likelihood, say what the pooling assumes, and write what replaces it if the assumption is dropped.

Exercise 11.10 (*formulate*). The battery room's log (illustrative) says: "Four grid outages in two years, of 0.5, 0.9, 1.2 and 1.6 h, all carried." Say what the log supplies about the outages and what it supplies about the string, and write the likelihood of each.

Exercise 11.11 (*reverse*). Write one plain sentence for $\mathcal{N}(d;\, 60\,\mathrm{kW}/G,\, \sigma_d^2)$ read as a function of G, and another for the same expression read as a function of d.

Exercise 11.12 (*reverse*). Write one plain sentence for $(\lambda_{\mathrm{fd}})^4 e^{-\lambda_{\mathrm{fd}}}/4!$, with the exposure and its clock named.

Exercise 11.13 (*reverse*). Write one plain sentence for

$$\int \mathcal{N}(y;\, P + b,\, (2\,\mathrm{kW})^2)\, \mathcal{N}(b;\, 0,\, (45\,\mathrm{kW})^2)\, \mathrm{d}b.$$

Exercise 11.14 (*reverse*). Write one plain sentence for $(1 - p_{\mathrm{d}})^6\, p(p_{\mathrm{d}})$, with p_{d} the UPS's probability of failing to pick up on demand, and say what its exposure is.

Exercise 11.15 (*critique and repair*). A network review says: "We average 60 one-minute PDU readings, noise 2 kW each, so the hourly mean has a spread of 0.26 kW, and that rules out the meter-offset explanation." State what is wrong, and repair it.

Exercise 11.16 (*critique and repair*). A branch review says: "The IT branch's log shows no faults in its first 8 months, so the fault rate is zero and that cut set can be dropped from the load's event." State the errors, and repair it.

Exercise 11.17 (*critique and repair*). A note on the network says: "The UPS and PDU meters read the same current transformer, so with two meters the offset averages down by the root of two." State the errors, and repair it.

Exercise 11.18 (*critique and repair*). A test plan says: "We calibrated the two hall sensors against each other in one bath yesterday, so the offsets' difference now has a spread of 0.05 K, and σ_d is 0.05 K." State what is wrong, and repair it.

Exercise 11.19 (*critique and repair*). A network review says: "The UPS output and PDU meters agree to within 1 kW, so the PDU offset is at most 1 kW and the load is known to 1 kW." State the errors, and repair it.

Exercise 11.20 (*rapid*). Name the term each statement supplies: noise, offset, a count, an exposure, or none. Say what else is needed where something is.

(a) "Class A PT100, tolerance ±0.15 °C at 0 °C."

(b) "Repeated readings at a fixed load scatter by 0.1 K."

(c) "Fifty-three operating days since the last overhaul."

(d) "The sensor has drifted 0.2 K since its last calibration."

(e) "Two failures in the log."

(f) "The load bank's controller holds the load to ±0.5 kW."

(g) "The load bank was set to 40 kW."

(h) “The meter reads 4 % high at every load.”

(i) “The log starts at commissioning.”

(j) “The reading was logged once a minute.”

11.7 What to retain

- A reading is the quantity plus an offset plus noise. Noise averages and an offset does not.
- A nuisance variable is carried or integrated out, and the sheet says which. Dropping one is a labeled assumption.
- A likelihood is a function of the unknown, not a probability of it.
- A plant’s own record is evidence: a count with its exposure, on a stated clock. On the two-hall plant it is written “given the record”.
- On one test: 1.000 with the offsets dropped and 0.818 with them in. On one log: 2.33 as a point, 1.18 to 3.90 as what the record leaves known.

Lesson 12

From Engineering Knowledge to Priors

12.1 The engineering question

At a design review the drawing says the opening passes 66 kW. The specification says the plant shall hold 8 K. The handbook says a unit like this one fails 2.5 times a year in service. The manufacturer's fleet list shows forty units, of which this plant's two are members. The sister plant, of the same design, has its own log. And the commissioning lead says of the opening, "I would be surprised if it were off by more than a factor of two." The analyst has to put all of this into the model, or decide what stays out.

Each of those is knowledge, and none of them is yet a probability. One says something about this very plant, and another about plants like it. One is a requirement and not knowledge at all. And one, the fleet list, may already hold the plant's own record, which someone is about to use a second time.

This lesson writes what each source asserts as a distribution, says what its spread means, and applies a short rule to any record, to decide whether it is evidence, a prior or neither.

12.2 The formulation

The work sits in step 3, in the words "prior information" that each lesson so far has used in a clause. On the two-hall plant E stays the 40 kW trip test, and the plant's log is "the record".

12.2.1 What a prior says

A prior has five parts, and a sheet states them all: the *quantity* it is about, its *family*, its *center*, its *spread*, and the *coordinate* those are written in. A source seldom supplies them all. The analyst supplies the rest, and each part so supplied is labeled as a choice. Table 12.1 takes six sources, one a row, with what each asserts, the prior it can support, and what is left for the analyst.

12.2.2 What a spread means

"A spread of a factor of 3" is not one thing. The report takes it as one standard deviation of $\log G$, so that the standard deviation is $\log 3$. Read as a 90 % interval, the same words give a standard deviation of $\log 3/1.645$. "Within 2 %" may be a standard deviation, an interval or a hard band. In this book, a stated factor in a sheet, such as "a factor of 3", is one standard deviation of the log of that factor unless the sheet says otherwise, and any other reading is

Table 12.1: Six sources of knowledge about the plant, and the prior each can support. The tolerance of 2 % and the spread of a factor of 3 are the report's.

Source	It asserts	A prior it can support	The analyst supplies
The near pair's two drawings	54 or 66 kW at 8 K, tolerance 2 %	Two lognormals on G	The weights, and whether the list is complete
The report's unrestricted prior	Center G^*, a factor of 3	$\log G$ normal	Nothing; it is stated
The handbook's 2.5 per year	A center	$\log \lambda_A$ normal, center 2.5	The spread and the family
A fleet record of other units	A center and a between-unit spread, if published	A prior on this unit's rate	How this unit may differ
A nameplate with a range	A bound	φ uniform on the range	The form within the bound
A person's judgment	A center and a rough width	$\log G$ normal	What kind of width; who said it

written out. Figure 12.1 draws the two readings of a factor of 3 on one axis, with the threshold marked.

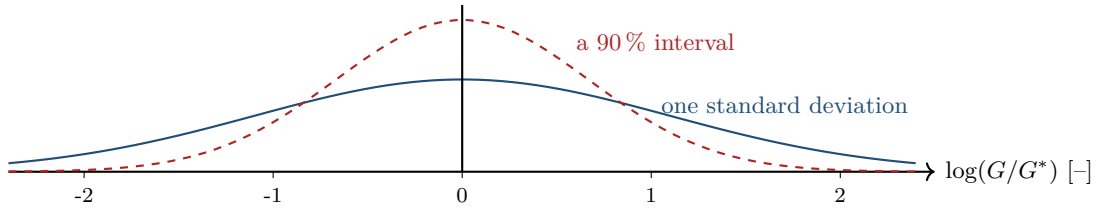


Figure 12.1: The prior on the log of the conductance under two readings of "a spread of a factor of 3": one standard deviation of $\log 3$, solid; a 90 % interval, dashed. The vertical line is the threshold.

How much the words matter depends on the evidence. Table 12.2 gives the probability that the opening is adequate under four priors, one a row, on E and on an illustrative weak test of 5 kW with $d = 0.5$ K. The values come from a small calculation on the stated relation, and the range of the flat prior is illustrative.

Table 12.2: The probability that the opening is adequate, under four priors, on a strong and on a weak test.

Prior	On E [–]	On 5 kW [–]
$\log G$ normal, center G^*, a factor of 3 as one standard deviation	0.818	0.679
The same, a factor of 3 as a 90 % interval	0.813	0.604
$\log G$ normal, center 8250 W/K, a factor of 3 as one standard deviation	0.821	0.700
Flat on G, from 1000 to 40000 W/K	0.857	0.896

On E the first three rows lie within 0.01, because the test leaves the prior little to do. On the weak test they span 0.10, and the prior is much of the answer. The last row is flat, and it is not neutral.

12.2.3 No prior is flat in every coordinate

A prior flat on G is not flat on $\log G$ or on R, since a change of coordinate carries its change-of-variable factor. A sheet that says "uninformative" has therefore chosen a coordinate without saying which. The flat prior of Table 12.2 puts most of its weight on large conductances, because a range in G from 7500 to 40000 W/K is far wider than one from 1000 to 7500 W/K. On the weak test that alone moves the answer from 0.679 to 0.896, and the top of the range moves it again: 0.778, 0.896 and 0.948 at 20000, 40000 and 80000 W/K, with E staying at 0.857. A flat prior is a choice like any other: the sheet names its coordinate and its range, and it says why.

12.2.4 Evidence or prior

Part of a sheet's step 3 is to place each statement that reaches it. The rule is four questions asked in order, and then one check. Figure 12.2 draws them.

1. **Is it a requirement?** A "shall", a limit or a target says what the plant must do, not what it is. It goes to step 4.
2. **Is it a statement about an uncertain quantity in the model, or the same quantity on another item, in hand?** If not, it is *none of these*, and the sheet says so: a design basis, which is a condition of the question; a record not yet obtained; a record with no relation to anything in the model, such as an oil analysis or a thermal image with no relation written.
3. **Did this very item produce it, across a gap the model can carry, with no stated change bearing on the quantity that a relation does not carry, and can its likelihood be written?** Then it is *evidence*. Three parts of the question need saying.
   - The gap between a record and the question counts unless the model holds the quantity constant across it, or a relation carries it. The sheet states which. Last year's trip test is evidence because the model holds G constant.
   - A change is one bearing on the quantity. A repainted hall is not one; exchanged sensors change the offsets and not G; an overhaul of a unit changes its rate. A change that a relation carries leaves a record as evidence: a commissioning test, then a modification stated as $G \times 0.8$, with its noise.
   - A likelihood can be written when there is a relation in the model, or one the sheet adds with its noise stated, or a count with its exposure stated or bounded on both sides. An unknown exposure carried as a bounded nuisance is still evidence, and so is a count known as a range, such as 2 or 3 failures. A count shared by this item and others is evidence on their sum. A one-sided exposure, such as at least 2 years, is not bounded: ask. A record whose item, exposure or noise is wholly unstated is not ready: ask.
4. **Otherwise** it is a *prior* on this item's quantity, with a spread for how this item differs from the source.

The check applies to evidence and to a prior: **has any part of it already entered the model?** A handbook built from a fleet, a fleet that contains this unit, a prior set after looking at the record. A recollection of a logged record is that record again. If so, the overlap is removed or the source is used once, and the sheet says which. Where the overlap's size is unknown, keep the record and do not build the prior from the source; the sheet says so, and asks.

Table 12.3 applies the rule to eight records of the design review and the commissioning file, one a row, with the outcome of each and what the sheet must add.

The first row is why the plant's own record is evidence. The seventh row is the easy one to get wrong, since a drawing feels like a fact about the plant. It is a fact about the document.

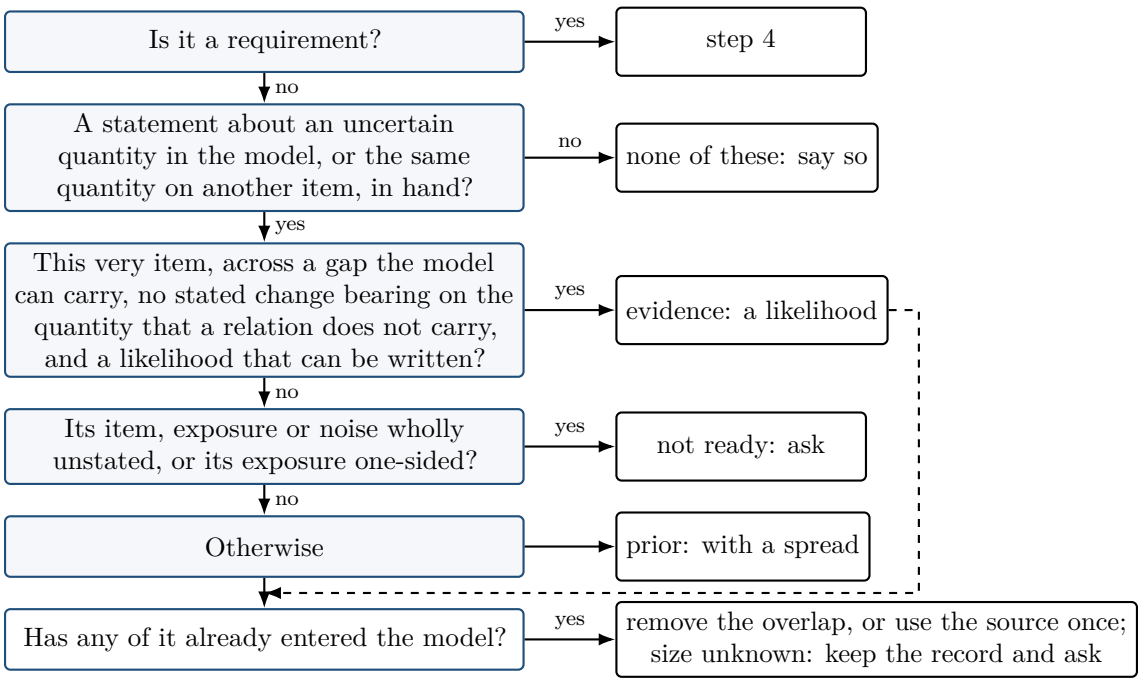


Figure 12.2: The rule for placing a statement. A requirement and a statement with no place in the model leave early, and a record with something wholly unstated goes back with a question; evidence and a prior both pass through the check.

Table 12.3: The rule applied to eight records. The fleet list, the sister plant and last year's test are illustrations.

Record	Outcome	The sheet must add
The plant's own three-year log	Evidence	Its likelihood, a count with its exposure
The sister plant's log	Prior	A spread for how plants of one design differ
The handbook's 2.5 per year	Prior	A spread and a family; the handbook's source, if known
The fleet list, this plant's units in it	Prior, overlap removed	The fleet less these units; a population model is named and left to the literature
Last year's trip test on this opening	Evidence	The model holds G constant across the year, as the sheet states; a likelihood beside this year's; are the sensors the same? Then one offset
The specification's 8 K	Requirement	Nothing here; it is the stated limit, at step 4
The drawing's 66 kW	Prior	A spread: a drawing is a document, and the as-built opening has the value
The commissioning lead's "a factor of two"	Prior	What kind of width; who said it

12.2.5 Unit A's rate from the handbook

The first worked formulation turns a handbook's number into a prior on one unit's rate. Table 12.4 is the sheet, with its run, and a second run writes the near pair's drawings as a prior on the opening.

Put through unit A's three-year record, the prior on the rate gives a band from 1.18 to 3.90 per year in service. On a one-year record holding two failures, its spread moves the band: from

Table 12.4: The formulation sheet filled in for a prior on unit A's rate, with a run for it and one for the opening.

Field	Entry
Intention	Estimate: a prior serves the posterior
1. Sentence	The handbook gives 2.5 failures per year in service for a unit of this design: what prior on unit A's rate does it support?
2. Target	A distribution on the parameter λ_A
3. Known	The handbook's value. By the rule, it is a document, so a prior. Assumption: the unit is of the handbook's design. Choice: the spread and the family.
4. Event	None: a density
5. Expression	$p(\log \lambda_A) = \mathcal{N}(\log \lambda_A;\, \log 2.5,\, (\log 3)^2)$
6. Model	A lognormal family; a spread that stands for how a unit may differ from the handbook, judged and not derived
7. Read back	What is believed of the rate before this unit's own record is used. Silent on the record, which stays out of it.
Problem class	Prior

The handbook's rate, then the drawings:

Run, a prior on unit A's rate, from the handbook

1	**Sentence**	The handbook gives 2.5 per year in service: what prior on λ_A does it support?
	Intention	Estimate: a prior serves the posterior
2	**Target**	A distribution on λ_A
	Declare	λ_A [1/yr]; center 2.5 (handbook); spread $\sigma_0 = \log 3$ (choice); coordinate $\log \lambda_A$
3	**Known**	The handbook's value; a unit of that design (assumption); no plant record used
	Random here	λ_A; the center and the spread are set
4	**Event**	None: a density
5	**Expression**	$p(\log \lambda_A) = \mathcal{N}(\log \lambda_A;\, \log 2.5,\, (\log 3)^2)$
6	**Model**	Lognormal (choice); the spread judged, not derived (assumption)
7	**Check**	A document, so a prior; the record stays out of it
	Class	Prior

$$\boxed{p(\log \lambda_A) = \mathcal{N}(\log \lambda_A;\, \log 2.5,\, (\log 3)^2)}$$

Run, a prior on the opening, from the near pair's two drawings

1	**Sentence**	The drawings give 54 or 66 kW at 8 K, tolerance 2 %: what prior on G?
	Intention	Compare: the weights serve the two structures
2	**Target**	A distribution on G, over two hypotheses
	Declare	H_{AND}: 8250 W/K; H_{OR}: 6750 W/K; 0.02 the deviation of $\log G$ (assumption); weights $\frac{1}{2}$ (choice)
3	**Known**	The two drawings; the as-built opening is one of them (assumption)
	Random here	Which drawing holds; G within its tolerance
4	**Event**	None: a density
5	**Expression**	$\frac{1}{2}\mathcal{N}(\log G;\, \log 8250,\, 0.02^2) + \frac{1}{2}\mathcal{N}(\log G;\, \log 6750,\, 0.02^2)$
6	**Model**	Two lognormals; the unrestricted prior is the reading without the assumption
7	**Check**	A document, so a prior; a list of two says nothing of a third value
	Class	Prior, over hypotheses

$$\boxed{p(\log G) = \sum_i \tfrac{1}{2}\,\mathcal{N}(\log G;\, \log G_i,\, 0.02^2)}$$

0.67 to 4.69 at a factor of 3, and from 1.31 to 4.04 at a factor of 1.5. The less the record holds, the more of the answer the prior is.

12.2.6 The string's priors

The second worked formulation is on the battery string, where a nameplate range and a vendor's rate are the sources, both illustrative. The sheet's product is a probability before any test of this string.

Table 12.5: The formulation sheet filled in for the string's priors and what they imply for its two hours; the inputs are illustrative.

Field	Entry
Intention	Estimate, for a decision on the specification
1. Sentence	From the nameplate range and the vendor's rate of weak cells, and with no test of this string, what is the probability that it gives at least two hours?
2. Target	Priors on φ and on H_{weak}, and an event on them
3. Known	The fact sheet's statements, each a prior by the rule. Choice: uniform on the range. Assumption: φ independent of the weak cell.
4. Event	$T_{\text{rt}} \geq 2$ h, that is $\varphi \geq 0.9877$ and a sound string, since a weak cell cannot reach it
5. Expression	$P(T_{\text{rt}} \geq 2 \text{ h})$ under the prior
6. Model	$T_{\text{rt}} = \varphi \times 2.025$ h, and 0.70 of that with a weak cell; the two priors
7. Read back	A probability before any test. A fleet range would change it.
Problem class	Prior probability of an event

In a run, down to its reduction:

Run, the string's priors

1	**Sentence**	From the nameplate range and the weak-cell rate: probability of at least 2 h?
	Intention	Estimate, for a decision
2	**Target**	Priors on φ and H_{weak}, and an event on them
	Declare	φ on 0.85 to 1.00 (illustrative); $P(H_{\text{weak}}) = 0.10$ (illustrative); $T_{\text{rt}} = 2.025 \text{ h} \times \varphi$
3	**Known**	The fact sheet's statements; no test of this string
	Random here	φ and the hypothesis, each one unknown value
4	**Event**	$\varphi \geq 0.9877$, and sound
5	**Expression**	$P(T_{\text{rt}} \geq 2 \text{ h})$
6	**Model**	Uniform on the range (choice); φ independent of the weak cell (assumption)
	Reduce	$P(\text{sound})\ P(\varphi \geq 0.9877 \mid \text{sound}) = 0.90 \times (1 - 0.9877)/0.15$
7	**Check**	0.074, before any test; a fleet range changes it
	Class	Prior probability of an event

$$\boxed{P(T_{\text{rt}} \geq 2 \text{ h}) = P(\text{sound})\, P(\varphi \geq 0.9877 \mid \text{sound})}$$

Table 12.5 gives 0.074. With an illustrative fleet range of 0.90 to 1.00, the same run gives 0.111. The fleet's range is a record of other strings, a prior by the rule, and it is written as what it is: the analyst's reading of a record, and not a fact about this string.

12.2.7 Three records through the rule

The rule can be run like any other formulation. Three cases that need an argument are these, lettered.

Run, evidence or prior, three records

1	**Sentence**	(a) The sister plant's log; (b) a fleet list with this unit; (c) last year's test here
	Intention	Estimate: each serves the opening or the rate
2	**Target**	Each record's place at step 3
	Declare	λ_A [1/yr]; G [W/K]; (a) plant 2's log; (b) a list of 40 units; (c) d_0, b [K]
3	**Known**	(a) other items; (b) other items and this one; (c) this item, no stated change
	Random here	λ_A in (a), (b); G, b in (c)
4	**Event**	None: each is a factor of the joint
5	**Expression**	(a) $p(\lambda_A)$; (b) $p(\lambda_A)$, this unit taken out; (c) a likelihood in G, below
6	**Model**	(a) a spread for how plants differ; (b) the same; (c) the same sensors, so one b (ask)
7	**Check**	Already entered? (b) yes, so removed; (c) the offset, if the sensors are the same
	Class	(a), (b) Prior; (c) Likelihood

$$\text{(a), (b) } p(\lambda_A); \quad \text{(c) } \int p(d \mid G, b)\, p(d_0 \mid G, b)\, p(b)\, \mathrm{d}b$$

Case (c) is the lesson's failure class in a small form. Two tests with the same sensors share their offset, and a product of two likelihoods, each with the offset integrated out, would use that offset twice: two tests at one load leave a spread of 0.709 K on their mean difference, not 0.503 K. So the sheet asks which sensors were used.

12.2.8 A request that is not ready

The owner says: "Use whatever we know about the opening." The sentence names no source and no quantity, so step 3 cannot be written. Read as the drawing, it asks for a prior. Read as the trip test, it asks for evidence. Read as both, it asks for a posterior with the drawing behind the likelihood. Read as the commissioning lead's opinion, it asks for a prior with a name on it. The analyst lists the sources, places each by the rule, and asks whether the list is complete.

Key idea

A prior has a quantity, a family, a center, a spread and a coordinate. A source seldom supplies them all, and what the analyst supplies is labeled. A requirement goes to step 4. A statement about no uncertain quantity in the model or the same quantity on another item, or not in hand, is none of these. A record of this very item, unchanged, with a likelihood that can be written, is evidence. Anything else is a prior, with a spread for how this item differs. Then ask whether any of it has already entered the model.

12.3 The competing formulations

The headline error of this lesson counts one record twice. A prior is set to what the plant's own log already says, and the log is then used as evidence. Each use looks defensible, and together they narrow a band with no new failure observed.

Run, one record counted twice

1	**Sentence**	Given unit A's record: what is known of its rate?
3	**Known** ×	A prior set to what the record says, and the record as evidence: band 1.45 to 3.43
	Repair	The record once, on the handbook's prior: band 1.18 to 3.90
7	**Check**	Narrower by a factor of 1.37, no new failure; where else did this record go in?

$$p(\lambda_A \mid \text{the record}) \propto P(k \mid \lambda_A)\, p(\lambda_A), \quad p(\lambda_A) \text{ not from the record}$$

Figure 12.3 draws the two bands on one axis, with the record's point for reference. A center alone taken from the log would leave the band almost where it was, 1.17 to 3.88; it is the spread that narrows it. The second form hides the same error inside a source.

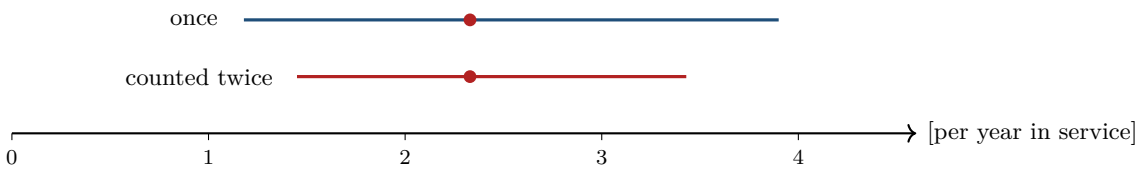


Figure 12.3: Unit A's rate after its record, the band from 5 % to 95 %: the record used once, and counted twice. The point is the record's count over its exposure.

Run, a fleet list that contains the unit

1	**Sentence**	Given unit A's record and the fleet's rate: what is known of its rate?
3	**Known** ×	The fleet's rate as the prior, unit A's failures in it, and unit A's record as evidence
	Repair	The fleet less unit A as the prior; the record as evidence
7	**Check**	Has any of it already entered? Here, yes

$$\boxed{p(\lambda_A \mid \text{the record}) \propto P(k \mid \lambda_A)\, p_{\text{fleet less A}}(\lambda_A)}$$

The third takes a document for the fact it describes. It is filed under model conditioning dropped, since the assumption is what vanishes.

Run, a document read as the fact

1	**Sentence**	Is the opening adequate?
5	**Expression** ×	"The drawing says 66 kW, so $G = 8250\,\text{W/K}$": $P(G \geq G^*) = 1$
	Repair	The drawing as a prior on the as-built G; then $P(G \geq G^* \mid E)$
7	**Check**	Would a trip test be pointless? Then the drawing is being read as the opening

$$\boxed{P(G \geq G^* \mid E)}$$

The fourth leaves a spread in words, and Table 12.2 gave its size.

Run, a spread with no meaning stated

3	**Known**	"$\log G$ spread a factor of 3"
6	**Model** ×	Standard deviation or interval not said: 0.818 or 0.813 on E; 0.679 or 0.604 at 5 kW
	Repair	One standard deviation of $\log 3$, written as such
7	**Check**	Could a reader rebuild the prior from the sheet?

$$\boxed{p(\log G) = \mathcal{N}(\log G;\, \log G^*,\, (\log 3)^2)}$$

The fifth writes "uninformative" and names no coordinate.

Run, a flat prior with no coordinate

3	**Known**	"No prior knowledge assumed"
6	**Model** ×	Flat on G, 1000 to 40000 W/K: 0.857 on E, 0.896 at 5 kW
	Repair	Name the coordinate and the range, and say why; the stated prior gives 0.818, 0.679
7	**Check**	Flat on G favors large conductances; it is a choice

$$\boxed{p(G) = \mathbf{1}[G \text{ in the stated range}] \,/\, \text{its width}}$$

The sixth makes a requirement into a prior, and the rule's first question stops it. It belongs

to the class of a requirement without its parts, written at the wrong step.

Run, a requirement written as a prior

1	**Sentence**	Will hall B hold 8 K at full load with unit B out?
3	**Known** ×	"The plant shall hold 8 K, so ΔT_{full} has a prior centered below 8 K"
	Repair	No prior from a "shall"; the 8 K is the limit in step 4
7	**Check**	Would the sentence still be true of an inadequate opening?

$$\boxed{P(\Delta T_{\text{full}} \le 8\,\text{K} \mid \text{unit B out},\ Q_{\text{hall}},\ \text{settled},\ E)}$$

The first two runs are forms of this lesson's class, *one record counted twice*: the same information used as prior and as evidence, or twice as either. Its four forms are a record in the prior and in the likelihood; a source that contains the item; a prior tuned after seeing the evidence; and two documents quoting one source, read as two. The class enters at step 3, and it is caught at step 7 by asking of each record where else it went in.

12.4 One system, many questions

Table 12.6 gives ten more records and statements about the two-hall plant, with the rule's outcome for each and what the sheet must add. It holds no values.

Table 12.6: Ten records about the two-hall plant, through the rule.

Record or statement	Outcome	The sheet must add
A repair invoice for unit B	Evidence	A likelihood on one repair's length; the check: is it already in the log's total?
A vendor's mean time to failure	Prior	A center only: a spread and a family
An as-built survey of the opening	Evidence	A relation from what it measures to G, with its noise stated; without one, none of these
A simulation study of the opening	Prior	Not a reading of this opening: a spread for the study's error
The design basis, 60 kW	None of these	A condition of the question
Last month's feeder log from the utility	Evidence	A stated clock; the check: is that month in the plant's record?
A neighboring substation's log	Prior	A spread for how the feeders differ
An operator's "about two trips a year"	Prior, or nothing	A recollection of the log is the log again, so keep the log and drop the sentence; with no log, a prior, with who said it
The plant's log since unit B's overhaul	Evidence	On the overhauled unit only
A fleet average published after this plant joined	Prior, overlap removed	This plant taken out; if its share is unknown, keep the record, build no prior from the average, and ask

The fifth row leaves at the rule's second question: a condition is about nothing uncertain in the model, and it stands behind the bar. The eighth is the check at work, since a person's memory of a log is not a second source.

12.5 Worked problems

Seven problems on the three systems follow, each with its run.

Problem 12.1. The relay's maker says (illustrative): "fewer than 1 spurious trip per year in service, at 95 %." Write the prior on the rate.

Run, a one-sided statement

1	**Sentence**	"Fewer than 1 per year, at 95 %": what prior on the spurious-trip rate?
	Intention	Estimate
2	**Target**	A distribution on λ_{tr}
	Declare	λ_{tr} [1/yr]; center 0.5, the fact sheet's (choice); 1 the 95th percentile (the maker's)
3	**Known**	A document about this make; this relay of that make (assumption)
	Random here	λ_{tr}
4	**Event**	None; a density
5	**Expression**	$\mathcal{N}(\log \lambda_{\text{tr}};\, \log 0.5,\, (\log 2/1.645)^2)$
6	**Model**	Lognormal (choice); one percentile cannot set both the center and the spread
7	**Check**	A document, so a prior; the center is the analyst's, and it fixes the spread, 0.42
	Class	Prior

$$\boxed{p(\log \lambda_{\text{tr}}) = \mathcal{N}(\log \lambda_{\text{tr}};\, \log 0.5,\, (\log 2/1.645)^2)}$$

Problem 12.2. Two handbooks, from surveys that share no plants, give a unit's rate as 2.5 and 3.2 failures per year in service (illustrative). Write the prior on unit A's rate.

Run, two sources that do not overlap

1	**Sentence**	Two surveys, 2.5 and 3.2 per year in service: what prior on λ_A?
	Intention	Estimate
2	**Target**	A distribution on λ_A
	Declare	λ_A [1/yr]; centers 2.5, 3.2; $\sigma_0 = \log 3$ (choice); weights $\frac{1}{2}$ (choice)
3	**Known**	Two documents; no plant in both (stated); no survey holds this unit (assumption)
	Random here	λ_A
4	**Event**	None; a density
5	**Expression**	$\frac{1}{2}\mathcal{N}(\log \lambda_A;\, \log 2.5,\, \sigma_0^2) + \frac{1}{2}\mathcal{N}(\log \lambda_A;\, \log 3.2,\, \sigma_0^2)$
6	**Model**	A mixture: neither survey is this unit, so agreement does not shrink the spread
7	**Check**	Already entered? Not here; a shared plant or source would be asked about
	Class	Prior, over hypotheses

$$\boxed{p(\log \lambda_A) = \tfrac{1}{2}\mathcal{N}(\log \lambda_A;\, \log 2.5,\, \sigma_0^2) + \tfrac{1}{2}\mathcal{N}(\log \lambda_A;\, \log 3.2,\, \sigma_0^2)}$$

Problem 12.3. A vendor tested 20 strings and found 2 with a weak cell (illustrative), which leaves the weak-cell share π_{w} uncertain. What is the prior probability that this string has one?

Run, a prior over a share

1	**Sentence**	2 weak strings in 20, the share uncertain: prior probability this string is weak?
	Intention	Estimate
2	**Target**	One probability, $P(H_{\text{weak}})$, through the share
	Declare	π_{w} the share [–]; flat prior on it (choice); this string not among the 20 (assumption)
3	**Known**	Other strings, so a prior on this one, with no spread beyond the share's
	Random here	π_{w}; whether this string is weak
4	**Event**	H_{weak}
5	**Expression**	$P(H_{\text{weak}})$
6	**Model**	$p(\pi_{\text{w}} \mid \text{the list}) \propto \pi_{\text{w}}^2(1-\pi_{\text{w}})^{18}$, a count over strings
	Reduce	$\int \pi_{\text{w}}\, p(\pi_{\text{w}} \mid \text{the list})\, \mathrm{d}\pi_{\text{w}} = 3/22$, against the count's 0.10
7	**Check**	One share over strings, not a population model; one string uses only its mean
	Class	Prior probability of an event

$$\boxed{P(H_{\text{weak}}) = \int \pi_{\text{w}}\, p(\pi_{\text{w}} \mid \text{the list})\, \mathrm{d}\pi_{\text{w}}}$$

Problem 12.4. Unit B's record holds 4 failures in its first two years in service, then an overhaul, then 3 failures in the next year (illustrative split). What is known of its rate now?

Run, one log, split by a stated change

1	**Sentence**	4 failures in 2 yr, an overhaul, 3 failures in 1 yr: what is known of λ_B now?
	Intention	Estimate
2	**Target**	A parameter, λ_B since the overhaul
	Declare	$k = 3$, $t_s = 1$ yr; $k_0 = 4$, 2 yr before; σ_o a log spread for the overhaul (choice); p_o
3	**Known**	After: evidence. Before, a stated change: a prior, from the fact sheet's old rate
	Random here	λ_B
4	**Event**	None; a density
5	**Expression**	$p(\lambda_B \mid k = 3)$
6	**Model**	A Poisson count after; p_o: what the 4 in 2 yr left of the old rate, widened by σ_o
7	**Check**	Each failure used once; a single count of 7 in 3 yr would be a rate across the change
	Class	Posterior of a parameter

$$\boxed{p(\lambda_B \mid k = 3) \propto P(k = 3 \mid \lambda_B)\, p_o(\lambda_B)}$$

Problem 12.5. With no source for the UPS's efficiency, an analyst says "somewhere from 0.90 to 1.00." How much does the coordinate of a flat prior matter for $P(\eta \geq 0.95)$?

Run, a flat prior over a narrow range

1	**Sentence**	"0.90 to 1.00": $P(\eta \geq 0.95)$, flat on (a) η; (b) $1/\eta$?
	Intention	Estimate
2	**Target**	One event under two priors
	Declare	η [–]; the range is a person's statement, so a prior; each shape a choice
3	**Known**	The range; no reading
	Random here	η
4	**Event**	$\eta \geq 0.95$
5	**Expression**	$P(\eta \geq 0.95)$
6	**Model**	(a) $p(\eta)$ flat; (b) $p(1/\eta)$ flat, so $p(\eta) \propto \eta^{-2}$
7	**Check**	0.500 and 0.474: little changes over a narrow range; the coordinate is still named
	Class	Prior probability of an event

$$\boxed{\text{(a)} \int_{0.95}^{1} \frac{\mathrm{d}\eta}{0.10}; \quad \text{(b)} \int_{0.95}^{1} \frac{\eta^{-2}\,\mathrm{d}\eta}{1/0.90 - 1}}$$

Problem 12.6. A reviewer says: "Use a conservative prior on the conductance."

Run, a prior asked for by its effect

1	**Sentence**	"Use a conservative prior on the conductance."
	Intention	Not stated: conservative for which decision?
3	**Known**	Missing: the event and the side of it that is safe; which part of the prior may move
5	**Expression**	(a) a lower center; (b) a wider spread; (c) a weight on the series structure
6	**Model**	(a), (b) $\log G$ normal; (c) a mixture on the near pair's drawings
7	**Check**	At 5 kW the wider prior raises adequacy, 0.604 to 0.679: wide is not cautious
	Class	Not ready

$$\boxed{\text{(a) } \mathcal{N}(\log G;\, \log G_{\text{low}},\, \sigma_0^2);\ \text{(b) } \mathcal{N}(\log G;\, \log G^*,\, \sigma_{\text{wide}}^2);\ \text{(c) } \sum_i \alpha_i\, p_i(G)}$$

Problem 12.7. The utility says that outages at the site "last about an hour." Is that a prior on the next outage's length?

Run, a statement about a mean		
1	**Sentence**	Outages "last about an hour": a prior on the next outage's length?
	Intention	Estimate
2	**Target**	A parameter, the mean length μ_o; not D_{out} itself
	Declare	μ_o [h]; center 1 h (the utility's); σ_u a log spread, asked for; D_{out} [h]
3	**Known**	A document, about this site's outages; no count behind it, so no likelihood
	Random here	μ_o, one value; D_{out} varies from outage to outage
4	**Event**	None; a density
5	**Expression**	$\mathcal{N}(\log \mu_o; \log 1 \text{ h}, \sigma_u^2)$
6	**Model**	Lognormal (choice); given μ_o, D_{out} exponential (illustrative), with no prior
7	**Check**	With μ_o known the next length still varies; the prior is on the mean
	Class	Prior

$$\boxed{p(\log \mu_o) = \mathcal{N}(\log \mu_o; \log 1 \text{ h}, \sigma_u^2)}$$

12.6 The exercises

The solutions are in Appendix A, and the facts of the three systems in Appendix D.

Exercise 12.1 (*formulate*). The design basis memo says: "Units are as the manufacturer's data, mean time between failures 3500 h in service, within a factor of two. The opening was built to the drawing, 66 kW at 8 K, tolerance 2 %. Unit A failed twice in its first year here. The cooling shall hold 8 K with one unit out." Apply the rule to each clause, write each prior or likelihood, and say what each stated width means or ask.

Exercise 12.2 (*formulate*). A design note says: "Put the stated prior on the feeder's rate." Formulate the note.

Exercise 12.3 (*formulate*). Run this one: the commissioning lead says the opening passes 66 kW at the 8 K limit, give or take a factor of two, and is 90 % sure of that. Write the prior on G.

Exercise 12.4 (*formulate*). A ticket reads: "Use the sister plant's history for our failure rates." Formulate the ticket.

Exercise 12.5 (*formulate*). The battery vendor reports that 1 string in 10 had a weak cell, in a fleet of 200 that includes this site's. Write the prior on H_{weak}, and say what to do about this site.

Exercise 12.6 (*formulate*). On the network, the UPS's 0.2 failures per year in service is a labeled illustrative input. Write what would make it a prior, and what a two-year log of this UPS would be.

Exercise 12.7 (*formulate*). The string's acceptance test, run four years ago, read 2.00 h. The owner asks: "Is that evidence about today's state of health?" Formulate the question.

Exercise 12.8 (*formulate*). Run this one: the pack's data sheet gives a thermal mass of 3.0×10^6 J/K, "±10 %". Write the prior on the thermal mass.

Exercise 12.9 (*formulate*). A site survey gives the opening's area as 2.1 m^2. Apply the rule, and say what would change the outcome.

Exercise 12.10 (*formulate*). A commissioning file holds three discharge tests: one by the vendor at the factory on a string of the same design; one on this string on site, before a cell was replaced; and one on this string since the replacement. Apply the rule to each, and write what each contributes to what is known of φ today.

Exercise 12.11 (*reverse*). Write one plain sentence for $\log \mu_{\mathrm{fd}}$ normal, center $\log 24$ h, spread $\log 2$, with μ_{fd} the feeder's mean repair.

Exercise 12.12 (*reverse*). Write one plain sentence for $p(\varphi) = \mathbf{1}[0.90 \le \varphi \le 1.00]/0.10$, and state what it does to a test reading that points below 0.90.

Exercise 12.13 (*reverse*). Write one plain sentence for $\int P(k = 7 \mid \lambda_A)\, p(\lambda_A)\, \mathrm{d}\lambda_A$, and say what a very small value of it would ask of the prior.

Exercise 12.14 (*reverse*). With ζ the string's usable energy as a share of a nameplate string's, write one plain sentence for

$$p(\zeta) = 0.90\, \frac{\mathbf{1}[0.85 \le \zeta \le 1]}{0.15} + 0.10\, \frac{\mathbf{1}[0.595 \le \zeta \le 0.70]}{0.105},$$

and say what each term stands for.

Exercise 12.15 (*critique and repair*). A report says: "The fact sheet's range left the string's two hours in doubt. The test reading looked strong, so we narrowed the prior on the state of health to 0.95 to 1.00, and on that prior the string passes." State the error, and repair it.

Exercise 12.16 (*critique and repair*). A network note says: "With no prior knowledge assumed, we took the PDU meter's offset as uniform within ±45 kW, as the fact sheet gives under the offset explanation." State what the sentence hides, and repair it.

Exercise 12.17 (*critique and repair*). A network memo says: "The UPS's failure rate is 0.2 from the manufacturer's fleet and 0.2 from the site's maintenance contract, which quotes the manufacturer, so two sources agree, and its availability at a random hour is therefore known." State the errors, and repair it.

Exercise 12.18 (*critique and repair*). A note says: "The prior on unit A's rate is the average over the fleet list of units now in service, 2.1 per year, and the list has forty units, so the spread is small." State the errors, and repair it.

Exercise 12.19 (*critique and repair*). A report says: "The as-built opening has the drawing's tolerance of 2 %, so the prior on $\log G$ has a standard deviation of 0.02 about $\log 8250$; the 40 kW test barely moved it, which confirms the drawing." State the errors, and repair it.

Exercise 12.20 (*rapid*). Place each by the rule: evidence, prior, requirement, overlap to remove, or none of these.

(a) A brochure's "99.9 % availability" for this model of unit.

(b) This unit's oil analysis from last month.

(c) "The load shall not exceed 60 kW per hall."

(d) Three trip tests averaged in one report, which is then used beside the raw tests.

(e) A consultant's estimate of this unit's rate, made after reading its log.

(f) A thermal image of the opening taken today.

(g) A retired engineer's "these feeders never last".

(h) The commissioning test of a similar opening in another building.

(i) Next year's record of this unit.

(j) A forecast of next summer's outside temperatures.

12.7 What to retain

- A prior has a quantity, a family, a center, a spread and a coordinate. What the analyst supplies is labeled.
- A spread's meaning is stated, and "flat" names a coordinate.
- The rule: a requirement goes to step 4; no uncertain quantity in the model or the same quantity on another item, or not in hand, is none of these; this very item's record, unchanged, with a likelihood, is evidence; anything else is a prior with a spread. Then ask whether any of it has already entered the model.
- A record counted twice narrows a band with no new information: 1.18 to 3.90 becomes 1.45 to 3.43.
- A prior matters as much as the evidence leaves room for: within 0.01 on the trip test, 0.10 on a weak test, among the normal priors.

Lesson 13

From Component Models to System Questions

13.1 The engineering question

The reliability review has three spreadsheets on the table. One gives each cooling unit's availability, 95 %. One gives the opening's adequacy from the trip test, 82 %. One gives the feeder's availability, 99 %. The owner asks for the plant's availability, and someone reaches for a calculator: is it the 95 %, or the product of the three, or something else? The test engineer then wants to know whether one more test would change the feeder's number. The answer should be plain from the plant, and it takes an hour of argument, because nothing on the table says what depends on what.

No spreadsheet is wrong. What is missing is the plant: which parts feed which, which quantities are computed from others, and which readings speak to which unknowns. A system question is not a component's question at a larger scale. It asks about an event that the components produce together.

This lesson writes the whole model as a product of factors, draws it to show what is tied to what, and writes the system's event through the structure. The drawing is a statement about dependence, and nothing here computes anything with it.

13.2 The formulation

The work sits in step 6, where Lessons 10 to 12's parts are assembled, and in step 4, where a system event is written through the structure. On the two-hall plant E is the 40 kW trip test.

13.2.1 The joint as a product of factors

A model of a system is one joint distribution over its unknowns and its readings, written as a product of *factors*, $\prod_j \psi_j$, each saying one thing and naming only the symbols it ties. On the two-hall plant at a random hour, with the rates at their stated values and the windows left out,

$$p(G, d, F_A, F_B, F_{\mathrm{fd}}, S) = p(G)\; p(d \mid G)\; P(F_A)\; P(F_B)\; P(F_{\mathrm{fd}})\; \mathbf{1}\big[S = s(F_A, F_B, F_{\mathrm{fd}}, G)\big]. \quad (13.1)$$

Each factor is of a kind the reader has met, and Table 13.1 lists them, one a row, with its kind, the quantities it names and the lesson's job it does. A symbol that appears in two factors is what makes two parts dependent, so dependence can be read from the product: Lesson 6's shared part and shared cause are written here as factors.

Table 13.1: The factors of equation (13.1).

Factor	Kind	Quantities	What it says
$p(G)$	Prior	G	What is believed of the opening before the test
$p(d \mid G)$	Reading	d, G	The trip test's likelihood
$P(F_A)$, $P(F_B)$	Component	F_A; F_B	Each unit's share of time out, at the stated rates
$P(F_{\text{fd}})$	Component	F_{fd}	The feeder's share of time out, at the stated rates
$\mathbf{1}[S = s(\cdot)]$	Structure	S, F_A, F_B, F_{fd}, G	The service as a function of its parts

13.2.2 The factor graph as a dependency diagram

The product is drawn as a *factor graph*: a circle for each quantity, shaded if it is observed; a square for each factor, joined by a line to the quantities it names; and a dark square for a deterministic factor. Figure 13.1 draws equation (13.1). It has three uses, which are this lesson's rule.

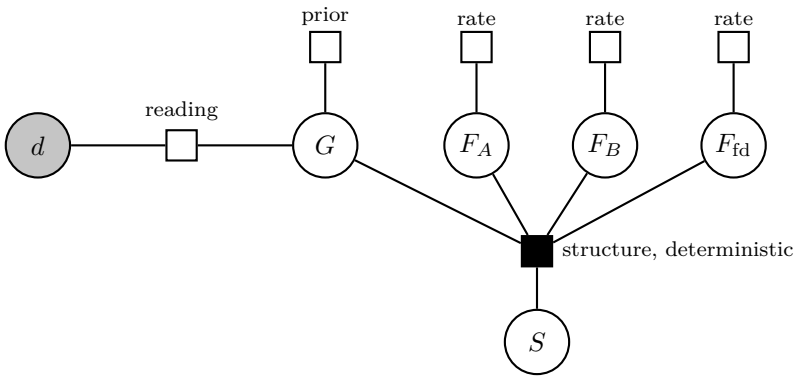


Figure 13.1: The two-hall plant at a random hour as a factor graph. Circles are quantities, the shaded one observed. Squares are factors, the dark one deterministic. A line joins a factor to a quantity it names. The drawing shows what is tied to what, and it has no direction.

1. **What is tied to what.** Two quantities can be dependent only if a path joins them that no known quantity cuts. A reading on a factor does not cut the path through that factor: two unknowns in one reading, such as R and δ in Lesson 10, become tied when it is read.
2. **What is observed.** The shaded circles are the readings in hand, and they stand behind the bar of a query. A query can put more behind its bar, a state known at the hour or an outcome read, and for that query those count as shaded too.
3. **What a query needs.** A factor's output is the quantity left of its bar in the product: S in the structure, d in the reading, G in the prior. It is read from the product, since the figure has no direction. A factor falls away when its output is neither asked about, nor read, nor an input of a factor that stays; repeat. Then keep what is joined to the asked-about quantity by a path no known quantity cuts. The rule assumes each quantity is the output of one factor, and that a known quantity is known exactly. When a reading falls away and the asked-about quantity would change what it reads, ask why. If it is already in the model as another reading, it stays out. If a stated or known value cuts

it, release the value or say it is unused. If it belongs to another hour or item, it is not a reading of this query. Otherwise the rule has found a missing factor: the sheet says so and adds the relation.

On Figure 13.1, $P(G \geq G^* \mid E)$ needs the prior and the reading, and nothing else. The structure's output is neither asked about nor read, so that factor falls away; the states' factors then feed nothing that stays, so they fall away too. So a feeder known to be out leaves the opening at 0.818: its factor is read and stays, but no path joins it to G. Read the service, and the structure stays, and it can tie its inputs: given $S = 0$ with the feeder up, the probability that the opening is adequate is 0.101, and given $S = 0$ with the feeder out it is 0.818, since the feeder alone explains the loss. That last is an argument from a value, and the rule does not make it: pruning by a value is outside the rule, because the diagram cannot show it. The figure is not a flowchart. No message passes along a line, no order of computation is implied, and how a query is evaluated is the subject of Appendix E and of nothing here. It is a picture of the product.

The first use has a consequence that surprises on a first meeting. The conductance and the offsets' difference have independent priors, and the trip test's reading names both, with the offsets carried rather than integrated out as in Lesson 11. Once it is read they are tied.

Run, two unknowns tied by one reading

1	**Sentence**	Given the 40 kW reading: are the conductance and the offsets' difference tied?
	Intention	Estimate
2	**Target**	Dependence between two parameters, after a reading
	Declare	G [W/K]; $b = b_B - b_A$ [K]; d [K]; $Q = 40\,\mathrm{kW}$
3	**Known**	d; unit B out; settled; independent priors on G and b (assumption)
	Random here	G, b
4	**Event**	None: a joint density
5	**Expression**	$p(G, b \mid d)$
6	**Model**	$p(G)\,p(b)\,p(d \mid G, b)$; the reading's factor names both
7	**Check**	Independent before the reading; after it, a larger b goes with a larger G
	Class	Joint posterior

$$\boxed{p(G, b \mid d) \propto p(G)\,p(b)\,p(d \mid G, b)}$$

13.2.3 A derived quantity

A derived quantity is a function of others: S of the states and G, ΔT_{full} of G and Q_{hall}, T_{rt} of φ and the weak cell. It hangs from a deterministic factor and has no prior of its own. A prior beside the kept factor would quietly reweight its parent: an illustrative prior on ΔT_{full}, normal about 7 K with spread 2 K, moves the prior probability that $G \geq G^*$ from 0.50 to 0.77. Dropping the factor instead cuts the tie the function states. A *structure function* is the derived quantity that gives a system's event from its components. On this plant, with each hall at its design load,

$$\{S = 1\} = s(F_A, F_B, F_{\mathrm{fd}}, G) = \begin{cases} F_{\mathrm{fd}}^c \cap (F_A \cap F_B)^c & \text{on } \{G \geq G^*\}, \\ F_{\mathrm{fd}}^c \cap F_A^c \cap F_B^c & \text{on } \{G < G^*\}. \end{cases} \tag{13.2}$$

Writing an event on S and then substituting equation (13.2) is a run's Expand, and a system question is where that line carries the most.

13.2.4 The plant's cooling

The first worked formulation asks for the plant's availability from its parts. The same diagram serves every query of Table 13.5: one model of a physical network supports many queries, which

Table 13.2: The formulation sheet filled in for the plant's cooling at a random hour, from its parts.

Field	Entry
Intention	Predict
1. Sentence	Given the trip test and the stated rates, at an hour chosen at random from steady operation, does no hall lose its cooling?
2. Target	A system event, $S(t) = 1$
3. Known	E; the stated rates; steady operation; each hall at its design load. Assumptions: the components fail independently; the windows are left out.
4. Event	$S(t) = 1$, through equation (13.2)
5. Expression	$P(S(t) = 1 \mid E)$, given the stated model
6. Model	Equation (13.1); every factor of Figure 13.1
7. Read back	The plant's event at a random hour. Silent on the path over a year, and on the windows, left out.
Problem class	Predictive probability at a random time, with dependence

From the diagram's factors, in a run:

Run, the plant's availability, from its parts

1	**Sentence**	Given E and the stated rates: at a random hour, does no hall lose its cooling?
	Intention	Predict
2	**Target**	A system event, $S(t) = 1$
	Declare	F_A, F_B, F_{fd}; G, G^*; s the structure; down: $q = 0.0489$ a unit, $f = 0.0100$ the feeder
3	**Known**	E; stated rates; steady; design load; windows out (assumption)
	Random here	The three states at the hour; G
4	**Event**	$S(t) = 1$
5	**Expression**	$P(S(t) = 1 \mid E)$, given the stated model
	Expand	s: on $G \geq G^*$, feeder up, units not both out; on $G < G^*$, all up
6	**Model**	$p(G)\, p(d \mid G)\, P(F_A)\, P(F_B)\, P(F_{\mathrm{fd}})\, \mathbf{1}[S = s(\cdot)]$
	Reduce	$(1-f)(1-q^2)\, P(G \geq G^* \mid E) + (1-f)(1-q)^2\, P(G < G^* \mid E)$
7	**Check**	$0.98763 \times 0.818 + 0.89561 \times 0.182 = 0.971$; silent on the path over a year
	Class	Predictive probability at a random time, with dependence
		$\boxed{P(S(t) = 1 \mid E),\ \text{given the stated model}}$

is the design of TERRA, the author's implementation of this kind of model. Table 13.2 is the sheet.

The value 0.971 comes out of the diagram's own factors, and the reduction shows why: the conductance enters only through which structure holds. The sheet's sixth row uses every factor of the figure, where the opening's adequacy used two.

13.2.5 The IT load

The second worked formulation is on the electrical network, which has more kinds of factor: the supply's losses, the UPS's pick-up on demand, the battery's autonomy, the two branches, their breakers, and five meters. Figure 13.2 draws it, with the failure part on the left and the meters on the right, each outlined by the query it serves. On the fact sheet's stated model the two parts share no factor, and the third use applies at once. For the probability that the IT load is without supply the meters fall away; for the UPS's efficiency the failure part falls away. One network holds two queries and two sub-diagrams. Table 13.3 lists the left part's factors, and Table 13.4 is the sheet for the first query.

The sheet's sixth row is where the third use does its work: whether or not the meters are read, no path joins them to the load's event. Now ask at the hour it was read: is the IT load without supply at the hour its meter read near zero? The load's state would change what the

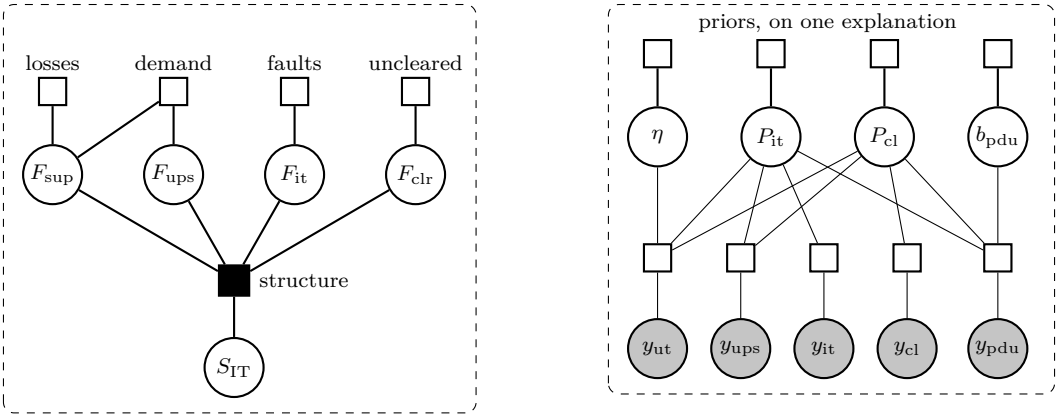


Figure 13.2: The network as a factor graph, at a random hour on the left and for a snapshot of its meters on the right. Left: the supply's losses, the UPS's demand, the IT branch's faults, the uncleared cooling-branch faults, and the structure giving the IT load's supply. Right: the efficiency, the two loads and the PDU offset, with one factor for each of the five meters' readings, shaded as observed. A dashed outline marks a query's sub-diagram.

Table 13.3: The factors of the network's left part, in Figure 13.2, on its illustrative inputs.

Factor	Kind	What it says
$P(F_{\text{sup}})$	Component	The supply's share of time lost, from its rates and lengths
$P(F_{\text{ups}} \mid F_{\text{sup}})$	Demand	The UPS not carrying when called: pick-up, autonomy, repair
$P(F_{\text{it}})$	Component	A fault on the IT branch, for its 4 h repair
$P(F_{\text{clr}})$	Demand	A cooling-branch fault not cleared: its rate, times 0.01, for the 1 h reset of the breaker above
$\mathbf{1}[S_{\text{IT}} = s_{\text{IT}}(\cdot)]$	Structure	The load's supply, through its three cut sets

meter reads, and the reading falls away all the same. It is not already in the model as another reading, no value cuts it, and it belongs to this hour, so the rule has found a missing factor. The relation missing is $P_{\text{it}} = S_{\text{IT}}(t)\, P_{\text{it,on}}$, with $P_{\text{it,on}}$ the load drawn when supplied, which ties what the meter reads to the load's supply; the fact sheet does not state it, so the sheet says so and adds it.

13.2.6 The string's race

On the battery string the same three uses hold, and the derived quantities stand out. Figure 13.3 draws the next outage: the state of health and the weak cell each with a prior, the ride-through hanging from a deterministic factor, the outage's length with its distribution, and the event hanging from a second deterministic factor.

13.2.7 A diagram request that is not ready

The operations director says: "Draw me the factor graph of the data center." A diagram with no question has no stated boundary, so the sentence fails at step 1. Read as the plant's cooling at a random hour, it is Figure 13.1. Read as one test's inference, it is the prior, the reading and the sensors' offsets, and the failures are absent. Read as the cooling and the power together, it

Table 13.4: The formulation sheet filled in for the IT load's supply at a random hour, from its parts.

Field	Entry
Intention	Predict
1. Sentence	At a random hour, on the illustrative rates: is the IT load without supply?
2. Target	A system event, $S_{\mathrm{IT}}(t) = 0$
3. Known	Steady; the structure; the fact sheet's rates, demands, autonomy (illustrative).
4. Event	$(F_{\mathrm{sup}} \cap F_{\mathrm{ups}}) \cup F_{\mathrm{it}} \cup F_{\mathrm{clr}}$
5. Expression	$P(S_{\mathrm{IT}}(t) = 0)$, given the stated model
6. Model	The left part of Figure 13.2. The meters' factors fall away.
7. Read back	The load's event at a random hour. Silent on how many separate breaks.
Problem class	Predictive probability at a random time, with a demand

And its run, the structure expanded:

Run, the IT load's supply, from its parts

1	**Sentence**	Random hour, steady, illustrative rates: is the IT load without supply, any cause?
	Intention	Predict
2	**Target**	A system event, $S_{\mathrm{IT}}(t) = 0$
	Declare	F_{sup}, F_{ups}; F_{it} an IT branch fault; F_{clr} a cooling fault not cleared; s_{IT} the structure
3	**Known**	Steady; the structure; rates, demand probabilities, autonomy (illustrative)
	Random here	The states at the hour; each loss's length; the pick-up; the clearing
4	**Event**	$S_{\mathrm{IT}}(t) = 0$
5	**Expression**	$P(S_{\mathrm{IT}}(t) = 0)$, given the stated model
	Expand	s_{IT}: $(F_{\mathrm{sup}} \cap F_{\mathrm{ups}}) \cup F_{\mathrm{it}} \cup F_{\mathrm{clr}}$
6	**Model**	Factors for each loss, the demand, the branches, the breakers; meters not needed
	Reduce	Sum of the three cut sets' probabilities, disjoint enough here (assumption)
7	**Check**	The load's event; meters silent here; silent on how many breaks
	Class	Predictive probability at a random time, with a demand

$$\boxed{P(S_{\mathrm{IT}}(t) = 0) \approx P(F_{\mathrm{sup}})\, P(F_{\mathrm{ups}} \mid F_{\mathrm{sup}}) + P(F_{\mathrm{it}}) + P(F_{\mathrm{clr}})}$$

The other query holds the same network and uses none of those factors:

Run, the UPS's efficiency, from the same network

1	**Sentence**	From a snapshot of the five meters, on one stated explanation: what is η?
	Intention	Estimate
2	**Target**	A parameter, η
	Declare	η [–]; P_{it}, P_{cl} the loads [kW]; b_{pdu} [kW]; $\mathbf{y}$ the five readings [kW]
3	**Known**	A steady snapshot; $\mathbf{y}$; the meters' noises; the explanation's stated spreads
	Random here	η; the loads and the offset, integrated out
4	**Event**	None: a density
5	**Expression**	$p(\eta \mid \mathbf{y})$, on that explanation
6	**Model**	A prior on each unknown; the five readings' factors; every failure factor falls away
7	**Check**	No rate, pick-up or breaker in it; silent on the load's supply
	Class	Posterior of a parameter

$$\boxed{p(\eta \mid \mathbf{y})}$$

needs a coupling that the fact sheets do not state, and Lesson 25 writes it. The analyst asks what the diagram is for.

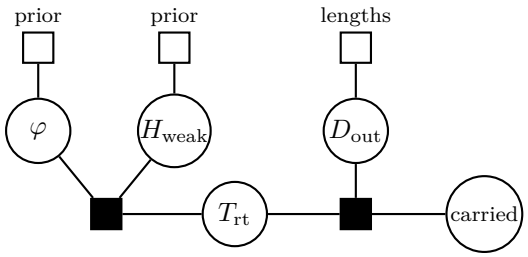


Figure 13.3: The string's next outage as a factor graph. The ride-through and the event hang from deterministic factors and have no priors of their own. Nothing is observed.

The run writes it, with the ride-through's definition in the Expand line:

Run, the next outage carried, from the string's parts

1	**Sentence**	On the fact sheet's priors, with no test: is the next outage carried to its end?
	Intention	Predict
2	**Target**	A derived event, on D_{out} and T_{rt}
	Declare	φ; H_{weak} the weak cell; T_{rt} derived [h]; D_{out} the outage's length [h]; all illustrative
3	**Known**	The fact sheet; no test of this string
	Random here	φ, the hypothesis, D_{out}; T_{rt} follows
4	**Event**	$D_{\text{out}} \le T_{\text{rt}}$
5	**Expression**	$P(D_{\text{out}} \le T_{\text{rt}})$
	Expand	$T_{\text{rt}} = 2.025 \text{ h} \times \varphi$, and 0.70 of that with a weak cell
6	**Model**	Two priors, a length distribution, two deterministic factors
7	**Check**	T_{rt} has no prior of its own
	Class	Predictive probability of a race

$$\boxed{P(D_{\text{out}} \le T_{\text{rt}})}$$

Key idea

A system's model is one joint, written as a product of factors. The factor graph draws that product: it shows what is tied to what, what is observed and which factors a query needs, and it computes nothing. A derived quantity hangs from a deterministic factor and has no prior of its own. One model answers many queries, and the question chooses the sub-diagram.

13.3 The competing formulations

Each competing formulation answers a question about what the parts produce together with an expression about the parts. The first takes one component's number for the plant's.

Run, a component's event for the system's

1	**Sentence**	Given E and the stated rates: is the plant's cooling delivered at a random hour?
4	**Event** ×	Unit A up: 0.951
	Repair	$S(t) = 1$, through the structure: 0.971
7	**Check**	Would the expression notice the feeder or the opening? It would not

$$\boxed{P(S(t) = 1 \mid E)}$$

The two numbers are close only by the accident of these rates, and the error is in kind. A unit's availability says nothing of the feeder or of the opening. The second multiplies whatever is on the table, and one of its three numbers is a statement about a parameter, which is Lesson 7's

class as well.

Run, the spreadsheets multiplied

1	**Sentence**	"What is the plant's availability?"
5	**Expression** ×	$0.95 \times 0.82 \times 0.99 = 0.77$
	Repair	The structure, with $P(G \geq G^* \mid E)$ as a weight: 0.971
7	**Check**	A parameter's probability is not an availability, and two units do not multiply

$$\boxed{P(S(t) = 1 \mid E)}$$

The third is subtler, since each factor in it is a right answer to something. The two hall events are multiplied, as if the halls were independent, and they share the feeder and the opening.

Run, the hall events multiplied

1	**Sentence**	Given E and the stated rates, at a random hour: do both halls keep their cooling?
5	**Expression** ×	$\big(1 - P(\text{hall B loses its cooling})\big)^2 = (1 - 0.0207)^2 = 0.959$
	Repair	$P(S(t) = 1 \mid E) = 0.971$
7	**Check**	The feeder and G are on both halls' paths in the diagram

$$\boxed{P(S(t) = 1 \mid E)}$$

The fourth gives a derived quantity its own prior. The opening's difference and its conductance then stand as two unknowns, and the tie is cut.

Run, a derived quantity given its own prior

3	**Known**	E; a prior on G
6	**Model** ×	A second prior on ΔT_{full}, beside the prior on G, the two independent
	Repair	A deterministic factor, $\Delta T_{\text{full}} = Q_{\text{hall}}/G$, and no prior on ΔT_{full}
7	**Check**	Can ΔT_{full} and G disagree? On the model they cannot

$$\boxed{\mathbf{1}[\Delta T_{\text{full}} = Q_{\text{hall}}/G]}$$

The fifth takes a correct answer to a smaller query. The sub-diagram of the opening's adequacy is the prior and the reading, and its 0.82 is reported as if the rest of the figure were not there.

Run, a sub-diagram's answer for the whole

1	**Sentence**	Given E and the stated rates: is the plant's cooling delivered at a random hour?
5	**Expression** ×	$P(G \geq G^* \mid E) = 0.82$, reported as the plant's
	Repair	$P(S(t) = 1 \mid E)$, on every factor of the figure: 0.971
7	**Check**	Which factors did it use? The prior and the reading; no unit and no feeder

$$\boxed{P(S(t) = 1 \mid E)}$$

The sixth reads the diagram as a procedure.

Run, a factor graph read as an algorithm

1	**Sentence**	"Information flows along the lines, so the feeder's record updates the opening."
6	**Model** ×	Lines read as a flow with a direction and an order
	Repair	Lines say what a factor names. F_{fd} and G meet only at the structure, not read
7	**Check**	Any factor joining them but the structure? No. Its output read? No: it falls away

$$\boxed{p(G \mid E,\ \text{the feeder's record}) = p(G \mid E)}$$

The first five are the forms of this lesson's failure class, *a component event substituted for the system event*: a question about what the parts produce together answered with an expression about a part. Its four forms are a part's event for the system's; parts' probabilities combined as a product, with the structure unwritten; a derived quantity given independent standing, its tie cut; and a sub-diagram's answer reported for the whole. It enters at step 4, 5 or 6, and it is caught at step 7 by asking which part's failure the expression would not notice. The class is close to two others. Lesson 5's fourth form checks one limit of a requirement where several

hold, and here the parts are components of a system. Lesson 6's class drops a tie between parts, and here the structure itself is missing. The sixth near-miss is outside the class, an error about the drawing rather than about the event, and it stands here because the drawing invites it.

13.4 One system, many questions

Table 13.5 holds Figure 13.1 still and asks ten queries of it, some with one more factor added, each with the factors it needs and what falls away. It holds no values.

Table 13.5: Ten queries on the two-hall plant, with the factors each needs and those that fall away.

Query	Needs	Falls away
The opening adequate	Prior on G; the reading	Every state and the structure
A unit's rate, given the record	Its prior; its count factor	G, the reading, the feeder
Service at a random hour	All of Figure 13.1	Nothing
Service through a year	All of it, with the process over time	Nothing; the factors grow
Both units out, with the windows	The units' states; a window factor	G, the reading, the feeder, the structure
Which structure	Prior on the hypotheses; the reading	The states
Overheating at a stated load	Prior on G; the reading; a deterministic rise	The states, for unit B out as given
A test's reading not yet taken	Its factor; what is known of G	The states
The feeder's rate	Its prior; its count factor	Everything else
Hall B's rise with unit B out at 50 kW	As the overheating query	The states

The plant did not change between the rows, and the question chose the sub-diagram. The fourth row keeps every factor and adds time to them, and the eighth reaches ahead to Part IV. Three of its rows are written out below, lettered.

Run, three rows of the table

1	**Sentence**	(a) Is $G \geq G^*$, given E? (b) Is λ_A above 2.5, given the record? (c) A clean year?
	Intention	(a), (b) Estimate; (c) Predict
2	**Target**	(a), (b) A parameter; (c) A system event over a horizon
	Declare	G [W/K]; λ_A [1/yr]; $\tau = 1$ yr; the rest as in Table 13.2
3	**Known**	(a) E; (b) the unit's record; (c) E, the stated rates, a start with all up
	Random here	(a) G; (b) λ_A; (c) the states' paths and G
4	**Event**	(a) $G \geq G^*$; (b) $\lambda_A > 2.5$; (c) $S(t) = 1$ for all $t \leq \tau$
5	**Expression**	(a) $P(G \geq G^* \mid E)$; (b) $P(\lambda_A > 2.5 \mid \text{the record})$; (c) below
6	**Model**	(a) Prior, reading; (b) prior, count; (c) all of Figure 13.1, with the process over time
7	**Check**	(a) No failure process; (b) no G; (c) nothing falls away
	Class	(a), (b) Posterior exceedance; (c) Predictive probability over a horizon

$$\boxed{\text{(c) } P(S(t) = 1 \text{ for all } t \leq \tau \mid \text{all up at } t = 0,\ E)}$$

13.5 Worked problems

Eight problems on the three systems follow, each with its run.

Problem 13.1. The stated rates are replaced by each component's three-year record. Is the service delivered at a random hour, given E and the records?

Run, the rates uncertain

1	**Sentence**	Given E and the three records, a random hour: does no hall lose its cooling?
	Intention	Predict
2	**Target**	A system event, $S(t) = 1$
	Declare	$\boldsymbol{\rho}$ every component's rate and mean repair; $\mathcal{R}$ the three records
3	**Known**	E; $\mathcal{R}$; the fact sheet's priors; steady; windows out (assumption)
	Random here	The states at the hour; G; $\boldsymbol{\rho}$
4	**Event**	$S(t) = 1$
5	**Expression**	$P(S(t) = 1 \mid E, \mathcal{R})$
6	**Model**	State factors gain parents, $P(F_i \mid \boldsymbol{\rho})$; each record adds count and repair factors
	Reduce	$\int P(S(t) = 1 \mid E, \boldsymbol{\rho})\, p(\boldsymbol{\rho} \mid \mathcal{R})\, \mathrm{d}\boldsymbol{\rho}$
7	**Check**	The same event; stated rates are the case with $\boldsymbol{\rho}$ fixed
	Class	Predictive probability at a random time, with dependence

$$\boxed{\int P(S(t) = 1 \mid E, \boldsymbol{\rho})\, p(\boldsymbol{\rho} \mid \mathcal{R})\, \mathrm{d}\boldsymbol{\rho}}$$

Problem 13.2. Suppose the two units are taken to have one unknown rate, with their repair time known. Are both units out at a random hour, given their records?

Run, a path through a shared parameter

1	**Sentence**	One unknown rate for both units, given the records: both out at a random hour?
	Intention	Predict
2	**Target**	A joint event of two components
	Declare	λ the one rate [1/yr]; $q(\lambda) = \lambda\mu/(8760\ \mathrm{h} + \lambda\mu)$; $\mu = 180$ h; $k = 14$, 6 unit-years
3	**Known**	The records; one rate (assumption); the states independent given it (assumption)
	Random here	The two states at the hour; λ
4	**Event**	$F_A \cap F_B$
5	**Expression**	$P(F_A \cap F_B \mid \text{the records})$
6	**Model**	λ a circle joined to both state factors and to the count: an unread path
	Reduce	$\int q(\lambda)^2\, p(\lambda \mid \text{the records})\, \mathrm{d}\lambda$
7	**Check**	Not the square of $\int q\, p$: the rate is shared; only a known rate cuts the path
	Class	Predictive probability at a random time, with dependence

$$\boxed{\int q(\lambda)^2\, p(\lambda \mid \text{the records})\, \mathrm{d}\lambda}$$

Problem 13.3. The IT load is found without supply at a randomly chosen inspection. How probable is it that a fault on the IT branch is why?

Run, the structure read, on the network

1	**Sentence**	IT load found without supply, a random inspection: is an IT branch fault why?
	Intention	Diagnose
2	**Target**	One cut set, given the load's event
	Declare	S_{IT}; F_{sup}, F_{ups}; F_{it}; F_{clr}
3	**Known**	$S_{\mathrm{IT}}(t) = 0$; steady; the rates and demand probabilities (illustrative)
	Random here	Which cut set holds
4	**Event**	F_{it}
5	**Expression**	$P(F_{\mathrm{it}} \mid S_{\mathrm{IT}}(t) = 0)$
6	**Model**	The structure's output is read: it stays and can tie its inputs; no path to the meters
	Reduce	$P(F_{\mathrm{it}}) \big/ \big(P(F_{\mathrm{sup}})P(F_{\mathrm{ups}} \mid F_{\mathrm{sup}}) + P(F_{\mathrm{it}}) + P(F_{\mathrm{clr}})\big)$, cut sets disjoint
7	**Check**	0.19 on these inputs: an event's share, not a statement about the branch's rate
	Class	Posterior probability over candidate causes

$$\boxed{P(F_{\mathrm{it}} \mid S_{\mathrm{IT}}(t) = 0) \approx \frac{P(F_{\mathrm{it}})}{P(F_{\mathrm{sup}})P(F_{\mathrm{ups}} \mid F_{\mathrm{sup}}) + P(F_{\mathrm{it}}) + P(F_{\mathrm{clr}})}}$$

Problem 13.4. A grid outage of 1.6 h was not carried to its end (illustrative). Does the string have a weak cell?

Run, a race read

1	**Sentence**	An outage of 1.6 h not carried: is a cell weak?
	Intention	Diagnose
2	**Target**	A hypothesis, through a read event
	Declare	$D_{\text{out}} = 1.6$ h; T_{rt} derived [h]; sound T_{rt} from 1.72 h; weak T_{rt} from 1.20 to 1.42 h
3	**Known**	The outage and its outcome; the fact sheet's priors; start 0.80 (assumption)
	Random here	The hypothesis; φ
4	**Event**	H_{weak}
5	**Expression**	$P(H_{\text{weak}} \mid D_{\text{out}} = 1.6 \text{ h},\ D_{\text{out}} > T_{\text{rt}})$
6	**Model**	The event's output is read: it stays; T_{rt}'s factor feeds it, with φ and the cell
	Reduce	A value argument: a sound string cannot fail at 1.6 h, so a cell is weak
7	**Check**	A probability of one on one model: is the range right, and was the start 0.80?
	Class	Posterior probability over candidate causes

$$\boxed{P(H_{\text{weak}} \mid D_{\text{out}},\ D_{\text{out}} > T_{\text{rt}}) \propto P(D_{\text{out}} > T_{\text{rt}} \mid D_{\text{out}}, H_{\text{weak}})\, P(H_{\text{weak}})}$$

Problem 13.5. The UPS's supply is known to be lost and the UPS is known to be carrying the load. Is the IT load without supply at this hour?

Run, known states cut the paths

1	**Sentence**	Supply lost, UPS known to be carrying, illustrative rates: IT load without supply?
	Intention	Estimate: a present state
2	**Target**	The load's event, two of its inputs known
	Declare	F_{sup} known; F^c_{ups} known; F_{it}; F_{clr}
3	**Known**	F_{sup}; F^c_{ups}; branches' rates (illustrative), independent of the supply (assumption)
	Random here	The branch states at the hour
4	**Event**	$S_{\text{IT}}(t) = 0$
5	**Expression**	$P(S_{\text{IT}}(t) = 0 \mid F_{\text{sup}}, F^c_{\text{ups}})$
	Expand	s_{IT}, the first cut set empty: $F_{\text{it}} \cup F_{\text{clr}}$
6	**Model**	The known states replace the losses' and the demand's factors; the branches' remain
7	**Check**	A known quantity cuts the path; the supply's rate does not enter
	Class	Prior probability of an event

$$\boxed{P(S_{\text{IT}}(t) = 0 \mid F_{\text{sup}}, F^c_{\text{ups}}) = P(F_{\text{it}} \cup F_{\text{clr}})}$$

Problem 13.6. In the next outage, is the load carried to its end and the pack kept under 318 K?

Run, two events on one parent

1	**Sentence**	Next outage, no test of this string: carried to its end, and the pack under 318 K?
	Intention	Predict
2	**Target**	Two derived events at once
	Declare	D_{out} [h]; T_{rt} derived [h]; M the pack's running maximum [K], derived; φ; H_{weak}
3	**Known**	The fact sheet; a stated start; a string at its floor stops heating (assumption)
	Random here	D_{out}; φ; the hypothesis
4	**Event**	$\{D_{\text{out}} \le T_{\text{rt}}\} \cap \{M(t_0 + D_{\text{out}}) < 318\,\text{K}\}$
5	**Expression**	$P\big(D_{\text{out}} \le T_{\text{rt}},\ M(t_0 + D_{\text{out}}) < 318\,\text{K}\big)$
6	**Model**	Two deterministic factors on D_{out} and φ, so not a product of the two probabilities
	Reduce	A value argument: carried means under 2.03 h, and the pack is then at most 305.0 K
7	**Check**	The second event adds nothing on these inputs; a hotter room would change that
	Class	Predictive probability of a race

$$\boxed{P\big(D_{\text{out}} \le T_{\text{rt}},\ M(t_0 + D_{\text{out}}) < 318\,\text{K}\big) = P(D_{\text{out}} \le T_{\text{rt}})}$$

Problem 13.7. A manager asks: "Draw the factors for hall A's own cooling."

Run, a structure with a limit missing	
1 Sentence	"Draw the factors for hall A's own cooling."
Intention	Not stated: the fact sheet states no limit for hall A above hall B
3 Known	Missing: hall A's limit, and a load at which to ask
4 Event	(a) hall B's 8 K mirrored (assumption); (b) its own limit, asked for; (c) no limit
5 Expression	(a) $P(F_{\text{fd}} \cup (F_A \cap (F_B \cup \{G < G^*\})))$; (b) its own threshold; (c) $P(F_{\text{fd}} \cup (F_A \cap F_B))$
6 Model	(a), (b) the plant's factors, a hall's structure; (c) the feeder's and the units' factors
7 Check	Equation (13.2) already mirrors the limit onto hall A, as (a) does; the analyst asks
Class	Not ready
	(a), (b) $P(F_{\text{fd}} \cup (F_A \cap (F_B \cup \{G < G_{\text{A}}\})))$; (c) $P(F_{\text{fd}} \cup (F_A \cap F_B))$

Problem 13.8. One more trip test is proposed. Taking its reading as if in hand, which of (a) the opening's adequacy, (b) the feeder's share of time down and (c) the service could it move?

Run, what a new reading reaches	
1 Sentence	A new reading d_{n}: does it move (a) the opening; (b) the feeder; (c) the service?
Intention	Estimate: a question about the diagram
2 Target	Three sub-diagrams
Declare	d_{n} [K]; a new reading's factor on G; S; F_{fd}
3 Known	E; d_{n} as if in hand; the stated rates
Random here	G; the states
4 Event	(a) $G \geq G^*$; (b) F_{fd}; (c) $S(t) = 1$
5 Expression	Below, one for each
6 Model	(a) joined to the new factor; (b) structure unread: it falls; (c) output asked: it stays
7 Check	(a), (c) can move; (b) cannot. Over readings not yet taken is Part IV's
Class	Factors of a query
	(a) $P(G \geq G^* \mid E, d_{\text{n}})$; (b) $P(F_{\text{fd}})$; (c) $P(S(t) = 1 \mid E, d_{\text{n}})$

13.6 The exercises

The solutions are in Appendix A, and the facts of the three systems in Appendix D. On the two-hall plant, E is the 40 kW trip test. The last exercise closes Part III.

Exercise 13.1 (*formulate*). The facility's reliability memo says: "Unit A and unit B share one feeder. The opening's conductance has been tested. Each unit logs its outages. The UPS rides through short feeder interruptions, and the string carries the control room. Tell me what we need to answer 'will the halls and the racks both keep running this hour?' " Write the factors, draw the diagram, and name what cannot be written yet.

Exercise 13.2 (*formulate*). Run this one: given the stated rates, at a random hour, is the feeder up and at least one unit up?

Exercise 13.3 (*formulate*). An engineer sketches four boxes on a whiteboard, one each for unit A, unit B, the feeder and "reliability", with arrows from the first three to the fourth, and asks that it be filled in. Formulate what is asked.

Exercise 13.4 (*formulate*). Run this one: at a random hour, on the illustrative rates, is the IT load without supply because the cooling branch's breaker failed to clear a fault?

Exercise 13.5 (*formulate*). Write the product of factors for the string with a timed discharge reading included. Mark what $P(T_{\text{rt}} \geq 2 \text{ h} \mid d)$ needs and what it does not.

Exercise 13.6 (*formulate*). Add the disturbance windows to Figure 13.1, and say which queries of the lesson's last table change.

Exercise 13.7 (*formulate*). The trip test also logs the surviving unit's duty and the two PDU powers, and the plant keeps its outage logs. With those added to the diagram and the loads among the unknowns, the owner asks: "Which of these bear on the opening's adequacy?" Formulate the question.

Exercise 13.8 (*formulate*). A review note says: "Put the static bypass into the IT load's diagram." Formulate the note.

Exercise 13.9 (*formulate*). The pack's thermal mass is doubted and given a prior. Write the product of factors for the pack staying under 318 K in an outage of stated length, and say what falls away by the rule, and what a value argument removes, at a length of 1 h and at 3 h.

Exercise 13.10 (*formulate*). A maintenance manager writes: "Both units are serviced by one contractor, and when the contractor is short of staff every repair runs long." Add what the note says to the plant's factors, and say which queries of the lesson's table it changes.

Exercise 13.11 (*reverse*). Write one plain sentence for this factor, and say what kind of factor it is:

$$\mathbf{1}\big[\text{hall B loses its cooling} \iff F_{\text{fd}} \cup (F_B \cap (F_A \cup \{G < G^{*}\}))\big].$$

Exercise 13.12 (*reverse*). Write one plain sentence for $\sum_i P(H_i)\,P(D_{\text{out}} \le T_{\text{rt}} \mid H_i)$, and say which quantities it has summed or integrated out.

Exercise 13.13 (*reverse*). Write one plain sentence for $p(G)\,\mathcal{N}(d;\, 60\,\text{kW}/G,\, \sigma_d^2)$, normalized over G, and say which of the plant's factors it leaves out.

Exercise 13.14 (*reverse*). Write one plain sentence for $\sum_c P(F_A \mid c)\,P(F_B \mid c)\,P(c)\;P(F_{\text{fd}})$ read as a factor of a joint over F_A, F_B, F_{fd}, and say what it ties and what it leaves untied.

Exercise 13.15 (*critique and repair*). A network summary says: "The IT load is lost when the supply is lost and the UPS does not carry it, so its unavailability is $P(F_{\text{sup}})\,P(F_{\text{ups}} \mid F_{\text{sup}})$." State what is wrong, and repair it.

Exercise 13.16 (*critique and repair*). A note after an incident says: "The cooling service was lost at 03:00 with the feeder and unit A both up. The structure is a deterministic box, so that event says nothing about the opening, which stays at 0.82." State the error, and repair it.

Exercise 13.17 (*critique and repair*). A network review says: "Since the five meters agree, the UPS is healthy, so the IT load cannot be lost in a utility interruption." State the errors, and repair it.

Exercise 13.18 (*critique and repair*). A review says: "The opening is adequate with probability 0.82, so the plant runs in parallel for 82 % of its hours and in series for the rest." State what is wrong, and repair it.

Exercise 13.19 (*critique and repair*). A note on the string says: "The ride-through hangs from a deterministic factor, so it has no uncertainty, and $P(T_{\text{rt}} \ge 2\ \text{h})$ is either 0 or 1." State the error, and repair it.

Exercise 13.20 (*rapid*). For each query on the string or the network, name one factor it needs and one that falls away.

(a) The string's state of health, after a timed discharge.

(b) The pack's peak temperature in an outage of a stated length.

(c) The next outage carried.

(d) The UPS's efficiency, from the five meters.

(e) The PDU meter's offset.

(f) The IT branch faulted, given the record of its faults.

(g) The load's supply in a feeder interruption, the UPS known up.

(h) Whether the weak cell is present, with no test.

(i) The IT load lost, given a fault on the cooling branch.

(j) The number of interruptions of the IT load in a year.

Exercise 13.21 (*machine set*). Ten problems to close Part III. Run each one like a machine.

(1) Hall B at 45 kW, unit B out, settled, given E: probability over its limit?

(2) Tonight, unit B out, hall B's load normal about 50 kW, spread 5 kW (illustrative), settled, given E: over its limit?

(3) The reading of a 60 kW trip test, as a likelihood in R.

(4) A timed discharge reads 1.30 h, spread 0.03 h (illustrative): is a cell weak?

(5) The PDU meter's certificate says "offset within ±5 kW at 95 %": the prior on the offset.

(6) Last year's trip test, the opening enlarged since: evidence or prior on today's G?

(7) The IT load without supply, through its first cut set.

(8) The next outage carried, after a timed discharge reading d, the weak cell not ruled out.

(9) Hall B keeps its cooling at a random hour, given E and the stated rates.

(10) Which factors of Figure 13.1 does $P(F_A \mid S = 0,\ E)$ need?

13.7 What to retain

- A system's model is one joint, written as a product of factors, each naming the quantities it ties.
- The factor graph draws that product. It shows what is tied to what, what is observed, and which factors a query needs. It computes nothing.
- A derived quantity hangs from a deterministic factor and has no prior of its own; Expand writes its definition into the event.
- One model, many queries: the question chooses the sub-diagram.
- A part's number is not the system's: 0.951 for a unit, 0.971 for the service, and 0.959 if the halls are multiplied.

Part IV

Formulating Decisions and Acceptance Rules

Before Part IV: a page to write now

A report on the two-hall plant has a subsection titled “The probability of acceptance and the acceptance margin”. Two passages from it follow.

> The N+1 claim is signed off once, on one test, and then relied on. The question this arm answers is what rule turns one trip test into an acceptance of the claim, that the opening can carry one hall’s load within the allowable rise, and how often that rule accepts a true claim and rejects a false one across the scatter of tests the plant could actually run. It matters because the test is run light, at a fraction of the hall’s load, and the claim is about full load: a rule that works on the design case and fails on real readings accepts plants that will overheat, and one that is too strict refuses plants that are fine.

> The power … lags the design case: at the 70 kW cell, where the full-load design case crosses 0.95, 89 of 200 realized tests accept; … A design that is expected to accept is not one that accepts on most days.

Write what each item below asks for, as far as it can be written with Lessons 1 to 13: in symbols where possible, and in words where not. For each, mark what is random, what is fixed and what stands behind the bar. Where nothing can be written yet, say what is missing. The page is not scored here. Keep the answer. Lesson 23 returns to this page and prints its solution there, and Part IV is what will have changed in between.

(a) The claim.

(b) The rule that “turns one trip test into an acceptance”.

(c) How often the rule “accepts a true claim”.

(d) How often it “rejects a false one”.

(e) “The test is run light … and the claim is about full load”: which load goes where.

(f) “The full-load design case crosses 0.95”.

(g) “89 of 200 realized tests accept”.

(h) The last sentence, as two expressions.

Lesson 14

From Probability to Engineering Decisions

14.1 The engineering question

The commissioning engineer reports that the opening is adequate with probability 0.82. The manager asks the only question that matters to her: do we sign? The engineer says that a probability is not a yes or a no. The manager asks who decides, and how. The owner's contract says 95 %, which the number does not reach. A colleague points out that the same reading would pass at 80 %, and that it also passes if it is simply held against what the limit implies for this load. Three people now hold three answers from one test, and none of them has made an arithmetic error.

What is missing is the decision itself. Nobody has written down what can be done if the claim is not accepted, what goes wrong under each choice, or the rule that turns a test into a choice. A probability says something about the plant. A signature does something to it.

This lesson writes those parts, and it writes the rule as a function of a test's evidence, so that what the rule did with this test and what it will do with the next one are different things.

14.2 The formulation

A decision is the first target in this book that is not a quantity. The work sits at step 2, where the target becomes an action or a rule's value, and at step 5, where Expand now substitutes a rule. On the two-hall plant E is the 40 kW trip test.

14.2.1 The parts of a decision

A decision has four parts. The *actions* are a set $\mathbb{A}$ of what can actually be done. On the opening they are to accept the claim, to modify the opening, to retest, and to shed hall B's load when a unit is lost. "Do not accept" is not an action until one of the others is named. The *states* are the claim's event and its complement, $G \geq G^*$ and $G < G^*$. The *consequences* are what each action leads to in each state, and Table 14.1 sets them out in words; Lesson 17 puts numbers in them.

The fourth part is the *threshold*, and a decision has two of them, of different kinds. Lesson 4 named them. The performance level stands inside the event: G^* comes from the hall's load and the stated limit. The probability level stands outside it: 0.95 is the owner's choice, and it carries his attitude to the consequences. Here it is written $1 - \alpha$, with α the posterior risk the owner will carry on a reading he accepts. Both are stated, and neither can be read from a test.

Table 14.1: The actions on the opening against its two states, with the consequences in words.

Action	$G \geq G^*$	$G < G^*$
Accept the claim	The plant goes into service as drawn	A hall overheats when a unit is lost in service
Modify the opening	Money spent on a sound opening	The opening made adequate, at that cost
Retest	A crew and a load bank, and a later decision	The same, and a later decision
Shed hall B's load on a unit's loss	Load lost that could have stayed	Load lost, and no overheating

14.2.2 The rule as a function

A *decision rule* takes a test's evidence and returns an action. An *acceptance rule* is a decision rule on a stated claim, with two values. For the N+1 claim,

$$A(e) = \mathbf{1}\big[P(G \geq G^* \mid e) \geq 1 - \alpha\big], \tag{14.1}$$

where e is a placeholder, like x in a function, for whatever evidence a test supplies, and the value 1 means accept. The rule can be written before any test is run, and it should be. Three objects now need to be kept apart, and Table 14.2 sets them side by side.

Table 14.2: A posterior, a rule's value and a rule's value on a test not yet run, on the opening at a level of 0.95.

Object	What it is	On the opening
$P(G \geq G^* \mid E)$	A number between 0 and 1, a statement about G	0.818
$A(E)$	The rule's value on evidence in hand, 0 or 1	0
$A(\tilde{E})$	The rule's value on a test not yet run, a random variable until the test is read	No value yet

With E in hand, $A(E)$ is known, and nothing is left to take a probability over. Any probability over a rule's outcome is therefore over $\tilde{E}$, and from here on the book writes it so. Figure 14.1 draws the two paths through one rule: from evidence in hand to one action, and from a test not yet run to a value that is not yet known.

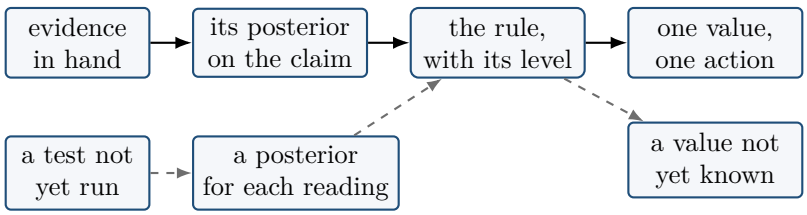


Figure 14.1: One rule on two paths. Solid, evidence in hand: the rule returns one value and an action follows. Dashed, a test not yet run: for each reading it could give the rule has a value, and which one holds is not known until the test is read.

14.2.3 One test and several rules

The rule is part of the formulation, and a different rule is a different formulation. Table 14.3 holds E still and changes the rule. The first three rows are equation (14.1) at three levels. The

last three are rules with no probability in them: the posterior's mode against G^*, the reading turned into a conductance, and the reading held against what the threshold implies at this load.

Table 14.3: Six rules on one reading, $d = 4.78$ K at $Q = 40$ kW.

Rule, as a function of the evidence	On E	Value [–]
$P(G \geq G^* \mid e) \geq 0.80$	$0.818 \geq 0.80$	1
$P(G \geq G^* \mid e) \geq 0.95$	$0.818 < 0.95$	0
$P(G \geq G^* \mid e) \geq 0.99$	$0.818 < 0.99$	0
The posterior's mode in $\log G$, against G^*	$8347\,\mathrm{W/K} \geq 7500\,\mathrm{W/K}$	1
The reading as a conductance, against G^*	$Q/d = 8363\,\mathrm{W/K} \geq 7500\,\mathrm{W/K}$	1
The reading against the threshold's difference	$d = 4.78\,\mathrm{K} \leq Q/G^* = 5.33\,\mathrm{K}$	1

Four rules accept and two do not, on one reading and with no error anywhere. The last three take no account of how far the reading could have scattered. Whether that matters cannot be decided from one reading: it takes a probability over readings, which is Lesson 16's. Here it is enough that the sentence "the test passed" names no rule, and so it is not ready.

14.2.4 The N+1 acceptance on the trip test

The first worked formulation takes the owner's level and the plant's test. Table 14.4 is the sheet, and the step to watch is the fifth, where the expression is a value and the rule is written out in its Expand line.

The rule's value is 0. The sheet's seventh row says what that leaves open: the opening is adequate with probability 0.818, and one of the other actions must now be chosen.

14.2.5 Two rules on one string

On the battery string the rule's inner probability is itself a choice, since it can be a parameter's exceedance or a predictive. Two rules follow, on the fact sheet's priors with no test: keep the string if it meets its two-hour specification, and keep it if the next outage is carried. The sentence "keep or replace" had not separated them. Both levels are illustrative.

Run, keep the string, on two rules

1	**Sentence**	No test, level 0.80: keep the string if (a) it gives its 2 h; (b) the next outage is carried?
	Intention	Decide
2	**Target**	The value of a rule, with a different inner probability in each reading
	Declare	T_{rt} [h]; D_{out} the next outage's length [h]; $A(\cdot)$; $\varnothing$ no test; $\mathbb{A}$ keep, replace; level 0.80 (illustrative)
3	**Known**	The fact sheet's priors; no test of this string
	Random here	φ, the weak cell, and in (b) D_{out}, all inside the rule
4	**Event**	(a) $T_{\mathrm{rt}} \geq 2$ h; (b) $D_{\mathrm{out}} \leq T_{\mathrm{rt}}$
5	**Expression**	$A(\varnothing)$, the rule on no evidence beyond the priors
	Expand	(a) $\mathbf{1}[P(T_{\mathrm{rt}} \geq 2\text{ h}) \geq 0.80]$; (b) $\mathbf{1}[P(D_{\mathrm{out}} \leq T_{\mathrm{rt}}) \geq 0.80]$
6	**Model**	The string's relation and priors; in (b) the outage's exponential length
	Reduce	(a) $\mathbf{1}[0.074 \geq 0.80]$; (b) $\mathbf{1}[0.834 \geq 0.80]$
7	**Check**	Opposite values on one string, because they are two claims
	Class	Decision rule

(a) $\mathbf{1}[P(T_{\mathrm{rt}} \geq 2\text{ h}) \geq 0.80]$; (b) $\mathbf{1}[P(D_{\mathrm{out}} \leq T_{\mathrm{rt}}) \geq 0.80]$

Table 14.4: The formulation sheet filled in for the N+1 acceptance on E, at a level of 0.95.

Field	Entry
Intention	Decide
1. Sentence	Given the 40 kW trip test, at the owner's level of 0.95: is the N+1 claim accepted?
2. Target	The value of a rule on evidence in hand
3. Known	E; unit B out; settled; the stated model; the level, the owner's. The other actions: modify, retest, shed.
4. Event	The claim $G \geq G^*$, inside the rule
5. Expression	$A(E)$, with equation (14.1) at $\alpha = 0.05$
6. Model	The posterior on G of Lessons 10 to 12
7. Read back	What the rule does with this reading. Silent on which action follows a 0, and on how often such a rule accepts.
Problem class	Decision rule

Its run, with the rule in the Expand line:

Run, the N+1 acceptance on the trip test

1	**Sentence**	Given E, at a level of 0.95: is the N+1 claim accepted?
	Intention	Decide
2	**Target**	The value of a rule on evidence in hand
	Declare	$A(\cdot)$ a function of a test's evidence; $G^* = 7500\,\mathrm{W/K}$; $1-\alpha = 0.95$; $\mathbb{A}$ accept, modify, retest, shed
3	**Known**	E; unit B out; settled; the stated model; the level (stated)
	Random here	Nothing in $A(E)$; G only inside the rule's posterior
4	**Event**	The claim $G \geq G^*$, inside the rule
5	**Expression**	$A(E)$
	Expand	$\mathbf{1}[P(G \geq G^* \mid E) \geq 0.95]$
6	**Model**	The posterior of Lessons 10 to 12
	Reduce	$\mathbf{1}[0.818 \geq 0.95]$
7	**Check**	A value, not a probability; 0 says the rule did not accept, and P stays 0.818
	Class	Decision rule

$$\boxed{A(E) = \mathbf{1}[P(G \geq G^* \mid E) \geq 0.95]}$$

14.2.6 A rule on an explanation

On the network the state is an explanation and the action a dispatch. The PDU meter is read against the UPS output meter, so that $y = y_{\mathrm{pdu}} - y_{\mathrm{ups}}$, and a technician is sent if an offset in the PDU meter is more likely than not. The level 0.5 is illustrative, and so are the two readings

the run tries.

Run, a technician sent on an explanation

1	**Sentence**	Given the meters' difference y: send a technician if an offset is more likely than not?
	Intention	Decide
2	**Target**	A rule's value, with a hypothesis as its state
	Declare	y [kW]; H_{off} the offset explanation; $\mathbb{A}$ send, do not send; level 0.5 (illustrative)
3	**Known**	y; the fact sheet's weights and spreads; one fault at a time
	Random here	The explanation, inside the rule
4	**Event**	H_{off}
5	**Expression**	$A(y)$
	Expand	$\mathbf{1}[P(H_{\text{off}} \mid y) \geq 0.5]$
6	**Model**	Bayes' rule on the two spreads of y, 3.2 kW and 45 kW
	Reduce	$y = 8\,\text{kW}$: $\mathbf{1}[0.077 \geq 0.5]$; $y = 12\,\text{kW}$: $\mathbf{1}[0.802 \geq 0.5]$
7	**Check**	The same form as a rule on a quantity; 1 sends a person, and proves no offset
	Class	Decision rule

$$\boxed{A(y) = \mathbf{1}[P(H_{\text{off}} \mid y) \geq 0.5]}$$

14.2.7 A decision request that is not ready

The project manager says: "Should we sign off?" Which claim is meant is not stated: the opening, or the whole of the redundancy statement with the feeder in it. Nor is the test, the level or what happens otherwise. Read as the opening at the owner's level, it is the run above. Read as the service at a random hour, the inner probability is $P(S(t) = 1 \mid E)$, and a different level is owed. Read as a question about the contract, it is not a formulation at all, and it goes to whoever holds the contract. The analyst asks.

Key idea

A decision has actions, states, consequences and thresholds. A rule is a function of the evidence, written before the test. On evidence in hand its value is known, and it says what the rule did and not how probable the claim is. Before the test its value is random.

14.3 The competing formulations

Each competing formulation takes a statement about the plant for a statement about a rule, or the reverse. The first reads a rule's 1 as the claim.

Run, accept read as true

1	**Sentence**	"At a level of 0.80 the rule gave 1, so the opening is adequate."
5	**Expression** ×	$A(E) = 1$, reported as $P(G \geq G^* \mid E) = 1$
	Repair	$A(E) = 1$ at that level, and $P(G \geq G^* \mid E) = 0.818$
7	**Check**	Would the sentence change if the level changed? It should not, and the 1 would

$$\boxed{P(G \geq G^* \mid E)}$$

The second is its mirror.

Run, not accepted read as false

1	**Sentence**	"The rule at 0.95 gave 0, so the opening is inadequate."
5	**Expression** ×	$A(E) = 0$, reported as $P(G < G^* \mid E) = 1$
	Repair	$P(G < G^* \mid E) = 0.182$; the opening is more likely adequate than not
7	**Check**	A 0 says the evidence did not reach the level

$$\boxed{P(G < G^* \mid E)}$$

The third reads a level as what was delivered, which is an error about every accepted reading but one.

Run, a level read as the probability delivered

1	**Sentence**	"Accepted at 95 %, so the claim holds with probability 0.95."
5	**Expression** ×	$P(G \geq G^* \mid e) = 0.95$, for an accepted e
	Repair	$P(G \geq G^* \mid e) \geq 0.95$, and its value on that reading is stated
7	**Check**	The level bounds the probability from below, and only the reading gives its value

$$\boxed{P(G \geq G^* \mid e)}$$

The fourth has no rule in it.

Run, a posterior read as a decision

1	**Sentence**	"The probability is 0.82, which is high, so we accept."
2	**Target** ×	A probability, taken for the action
	Repair	An action, by a stated rule: $A(E)$ at a level stated before the reading
7	**Check**	"High" against what? No level, no alternative action

$$\boxed{A(E) = \mathbf{1}[P(G \geq G^* \mid E) \geq 1 - \alpha]}$$

The fifth has a level, chosen too late. A level of 0.80 accepts E, and the 0.95 does not.

Run, the level set after the reading

1	**Sentence**	"The test gave 0.82, so an 80 % criterion is reasonable for this plant."
3	**Known** ×	The level, chosen from the reading it is meant to judge
	Repair	The level stated before the test, from the consequences
7	**Check**	A level that moves with the reading always accepts: $A(\cdot)$ is no longer a function of e alone

$$\boxed{A(e) = \mathbf{1}[P(G \geq G^* \mid e) \geq 1 - \alpha],\ \alpha \text{ stated first}}$$

The sixth reports one rule under the name of another, which matters because the two can part.

Run, a point-value rule reported as the posterior rule

1	**Sentence**	"Accepted at 95 % confidence: the mode in $\log G$ is 8347 W/K, against 7500 W/K."
5	**Expression** ×	$\mathbf{1}[8347\,\mathrm{W/K} \geq G^*]$, a rule on the mode, named as the rule at 0.95
	Repair	$\mathbf{1}[P(G \geq G^* \mid E) \geq 0.95] = 0$, against the point rule's 1
7	**Check**	Where is the 0.95 in the expression? Nowhere

$$\boxed{\mathbf{1}[P(G \geq G^* \mid E) \geq 0.95]}$$

The seventh leaves the action set empty.

Run, one action called "do not accept"

1	**Sentence**	"If the rule gives 0, we do not accept."
2	**Target** ×	An action set of accept and its negation
	Repair	$\mathbb{A}$ = accept, modify, retest, shed, with a rule or a person to choose among the last three
7	**Check**	What happens on Monday if the rule gives 0? The sentence does not say

$$\boxed{A(\cdot) \text{ into } \mathbb{A} = \{\text{accept, modify, retest, shed}\}}$$

The first four are the forms of this lesson's failure class, *rule against posterior*: what a rule did with this evidence reported as how probable the claim is, or the reverse. Its forms are accept read as true; not accepted read as false; a level read as the probability delivered; and a posterior read as a decision with no rule. It enters at step 2 or step 5, and it is caught at step 7 by asking whether the sentence's verb is about G or about an action. The last three are errors in building the rule. They stand here because they arrive with it.

14.4 One system, many questions

Table 14.5 holds the battery string still and sorts ten sentences by the part of a decision each one is. It holds no values, and the last row is an error that sounds like a part.

Table 14.5: Ten sentences on the battery string, by the part of a decision each one is.

Sentence	Part
Keep the string, or replace it	The action set
The string's state of health is above 0.90	A state, an event on φ
The string gives its 2 h	A state, an event on T_{rt}
Replace it if the next outage is less than 90 % likely to be carried	A rule, with a predictive inside
The 90 %	A threshold on a probability, stated
The 2 h	A threshold on a quantity, stated
A failed ride-through loses the critical load	A consequence
The timed discharge read 2.00 h	Evidence in hand, the rule's argument
Next year's discharge, not yet run	A test not yet run; the rule's value is random
"It passed, so the string is fine"	A rule's value read as a state

14.5 Worked problems

Seven problems on the three systems follow, each with its run.

Problem 14.1. The contract says: accept at 0.95, modify the opening below 0.50, and retest between them. What does the rule do on E?

Run, three outcomes

1	**Sentence**	Given E: accept at 0.95, modify below 0.50, retest between?
	Intention	Decide
2	**Target**	A rule with three values
	Declare	$A(\cdot)$ into $\mathbb{A}$ = accept, retest, modify; levels 0.95, 0.50 (illustrative)
3	**Known**	E; unit B out; settled; the stated model; the two levels
	Random here	Nothing in the value; G inside the rule's posterior
4	**Event**	The claim $G \geq G^*$, inside the rule
5	**Expression**	$A(E)$
	Expand	By where $P(G \geq G^* \mid E)$ falls: at least 0.95; 0.50 to 0.95; below 0.50
6	**Model**	The posterior
	Reduce	0.818 falls between: retest
7	**Check**	Two levels and three values; the band is what an acceptance clause usually means
	Class	Decision rule

$$A(e) = \begin{cases} \text{accept} & P(G \geq G^* \mid e) \geq 0.95 \\ \text{retest} & 0.50 \leq P(G \geq G^* \mid e) < 0.95 \\ \text{modify} & P(G \geq G^* \mid e) < 0.50 \end{cases}$$

Problem 14.2. A 60 kW test reads exactly 8 K. A rule on the reading says accept if $d \le Q/G^*$. What does it do, and what is missing from it?

Run, a rule with no probability in it

1	**Sentence**	A 60 kW test reads 8 K: accept if the reading is at most Q/G^*?
	Intention	Decide
2	**Target**	A rule's value on a stated reading
	Declare	$d = 8\,\mathrm{K}$; $Q = 60\,\mathrm{kW}$; $Q/G^* = 8\,\mathrm{K}$; $A(\cdot)$
3	**Known**	The reading; the rule. Nothing about the offsets or a prior enters it
	Random here	Nothing
4	**Event**	$d \le Q/G^*$, on the reading
5	**Expression**	$A(d)$
	Expand	$\mathbf{1}[d \le Q/G^*]$
6	**Model**	None is used
	Reduce	$\mathbf{1}[8 \le 8]$
7	**Check**	Value 1, and $P(G \ge G^* \mid d)$ on that reading is 0.535
	Class	Decision rule

$$\boxed{A(d) = \mathbf{1}[d \le Q/G^*]}$$

Problem 14.3. Suppose unit A logged one failure in its first year in service (illustrative; the fact sheet's one-year record logs two). Accept that its rate is below 3.5 per year if that has probability at least 0.90, given the record, on the stated prior.

Run, a rule on a count

1	**Sentence**	One failure in one year in service: accept $\lambda_A < 3.5$ at a level of 0.90?
	Intention	Decide
2	**Target**	A rule's value on a count
	Declare	k a count; $t_\mathrm{s} = 1$ yr; λ_A [1/yr]; level 0.90 and limit 3.5 (illustrative)
3	**Known**	The record; the stated prior on $\log \lambda_A$
	Random here	λ_A, inside the rule
4	**Event**	$\lambda_A < 3.5$
5	**Expression**	$A(k)$
	Expand	$\mathbf{1}[P(\lambda_A < 3.5 \mid k) \ge 0.90]$
6	**Model**	Poisson count at $\lambda_A t_\mathrm{s}$; the prior
	Reduce	$k = 1$: $\mathbf{1}[0.948 \ge 0.90]$; the fact sheet's $k = 2$ gives $\mathbf{1}[0.853 \ge 0.90]$
7	**Check**	On this record the rule is a statement about k, and given the record, never "given E"
	Class	Decision rule

$$\boxed{A(k) = \mathbf{1}[P(\lambda_A < 3.5 \mid k) \ge 0.90]}$$

Problem 14.4. The procedure says the acceptance applies "only to a test run to steady temperatures". A trip test was cut short at ten minutes. What is the rule's value?

Run, a rule with a precondition

1	**Sentence**	A trip test cut short, the rule defined only on a test run to steady temperatures: its value?
	Intention	Decide
2	**Target**	A rule's value, where the rule has a domain
	Declare	$A(\cdot)$ defined on steady tests; the cut-short test's readings, not E
3	**Known**	Not steady; the stated model holds only after steady temperatures (stated)
	Random here	Nothing
4	**Event**	The claim, inside a rule that does not apply
5	**Expression**	No value: the rule's argument is outside its domain
6	**Model**	The measurement relation needs a steady plant, so no likelihood is stated
7	**Check**	Neither 1 nor 0; "the rule said no" would be false; the action is to rerun
	Class	Decision rule

$$\boxed{A(e),\ e \in \{\text{the evidence of steady tests}\}}$$

Problem 14.5. The string's owner asks: "Is 90 % good enough?"

Run, a level asked about with nothing to attach it to	
1 Sentence	"Is 90 % good enough?"
Intention	Not stated
3 Known	Missing: of which event; against which consequences; for which action
4 Event	(a) $T_{\text{rt}} \geq 2$ h; (b) $D_{\text{out}} \leq T_{\text{rt}}$; (c) a level on either, as a loss ratio
5 Expression	(a), (b) a probability, compared with 0.90; (c) a question about the level, Lesson 17's
7 Check	(a) 0.074, (b) 0.834, each below 0.90, for different reasons; the analyst asks
Class	Not ready
	$\boxed{\text{(a) } P(T_{\text{rt}} \geq 2 \text{ h}) \text{ against } 0.90; \quad \text{(b) } P(D_{\text{out}} \leq T_{\text{rt}}) \text{ against } 0.90}$

Problem 14.6. The limit is tightened from 8 K to 7 K, and the level stays at 0.95. What does the rule do on E?

Run, the performance level moved	
1 Sentence	Given E, the limit at 7 K, level 0.95: is the claim accepted?
Intention	Decide
2 Target	A rule's value, with a new claim inside it
Declare	$G_7^* = 60\,\text{kW}/7\,\text{K} = 8571\,\text{W/K}$, the claim's threshold at the new limit; $A(\cdot)$
3 Known	E; unit B out; settled; the new limit (stated); the same level
Random here	G, inside the rule
4 Event	$G \geq G_7^*$, a different event on the same plant
5 Expression	$A(E)$
Expand	$\mathbf{1}[P(G \geq G_7^* \mid E) \geq 0.95]$
6 Model	The same posterior
Reduce	$\mathbf{1}[0.489 \geq 0.95]$
7 Check	The test did not change; the claim did: 0.818 becomes 0.489
Class	Decision rule
	$\boxed{A(E) = \mathbf{1}[P(G \geq G_7^* \mid E) \geq 0.95]}$

Problem 14.7. The meter crew has three work orders, one per explanation of the network's meters: none, service the UPS, recalibrate the PDU meter. The foreman issues the one the most probable explanation calls for.

Run, a rule on three explanations	
1 Sentence	Given a snapshot of the meters: issue the work order of the most probable explanation?
Intention	Decide
2 Target	A rule from evidence to one of three actions
Declare	$\mathbf{y}$ the five readings [kW]; H_i, i = nominal, ups, off; $\mathbb{A}$ the three orders
3 Known	$\mathbf{y}$; the fact sheet's weights and spreads; one fault at a time
Random here	The explanation, inside the rule
4 Event	None singled out: the three posteriors together
5 Expression	$A(\mathbf{y})$
Expand	The order for $\arg\max_i P(H_i \mid \mathbf{y})$
6 Model	The meters' factors under each explanation; Bayes' rule
7 Check	No level at all: a 0.40 explanation can win
Class	Decision rule
	$\boxed{A(\mathbf{y}) = \text{the order for } \arg\max_i P(H_i \mid \mathbf{y})}$

14.6 The exercises

The solutions are in Appendix A, and the facts of the three systems in Appendix D. On the two-hall plant, E is the 40 kW trip test.

Exercise 14.1 (*formulate*). A commissioning procedure says: “After the test the engineer reports the probability that the opening is adequate. If it reaches 0.9, the plant is handed over. If not, the contractor either enlarges the opening or repeats the test at a higher load, as the contractor chooses. Handing over a plant that cannot carry a hall costs the owner the hall’s equipment. Enlarging a sound opening costs the contractor the work.” Write the action set, the states, the consequences in words, the thresholds and the rule, and say what the rule does not decide.

Exercise 14.2 (*formulate*). The owner writes: “Sign off the N+1 claim on the trip test if the probability that the opening is adequate reaches 0.80.” Formulate the request.

Exercise 14.3 (*formulate*). Run this one: given E, at a level of 0.80, is the claim that hall B carries 50 kW within 8 K with unit B out accepted?

Exercise 14.4 (*formulate*). Write, as a function of the five meters’ readings, a rule that sends the UPS’s service crew when the degraded explanation has probability at least 0.5. State its action set and what it leaves undecided.

Exercise 14.5 (*formulate*). Run this one: a timed discharge reads 2.00 h with a spread of 0.03 h (illustrative), on a string taken as sound. At a level of 0.90, is the string kept, if it is kept when it gives its 2 h?

Exercise 14.6 (*formulate*). The plant’s reliability lead proposes replacing a unit if, given its three-year record, its rate is more likely than not above 3.5 per year. Write the rule, say what its value 1 means here, and say which clock the record must be on.

Exercise 14.7 (*formulate*). For the string, write a rule into keep, shed a share of the critical load in the next outage, and replace, with two levels on the probability that the next outage is carried. Say which of its parts are illustrative.

Exercise 14.8 (*formulate*). The facility manager writes: “Sign off the string.” Formulate the request.

Exercise 14.9 (*formulate*). An acceptance clause reads: “The plant is accepted if the expected rise of hall B at full load, unit B out, is under 8 K.” Write it as a rule, give its value on E, and hold it against the rule at 0.95.

Exercise 14.10 (*formulate*). A utility agreement says the feeder is accepted over a three-year record if the logged failures do not exceed three years at the stated rate. Write it as a rule, and give its value on the fact sheet’s record.

Exercise 14.11 (*reverse*). Write one plain sentence for this rule, and say what its argument is:

$$A(d) = \mathbf{1}\big[P(D_{\text{out}} \le T_{\text{rt}} \mid d) \ge 0.9\big].$$

Exercise 14.12 (*reverse*). Write one plain sentence for $\mathbf{1}[P(G < G^{*} \mid e) \le 0.05]$, and say how it stands to the rule at 0.95.

Exercise 14.13 (*reverse*). Write one plain sentence for $\mathbf{1}[P(H_{\text{ups}} \mid \mathbf{y}) \ge 0.2]$, and say in words which mistake a level this low treats as the worse.

Exercise 14.14 (*reverse*). Write one plain sentence for $\mathbf{1}[P(N_{\text{fd}}(1\text{ yr}) \le 3 \mid \text{the record}) \ge 0.8]$, and say what kind of probability is inside it.

Exercise 14.15 (*critique and repair*). A handover note says: "The test passed the acceptance rule, so the opening is adequate." State what is wrong, and repair it.

Exercise 14.16 (*critique and repair*). A design review says: "The posterior is 0.82. That is well above even odds, so the claim stands." State what is wrong, and repair it.

Exercise 14.17 (*critique and repair*). A note says: "We signed at 95 %, so there is exactly a 5 % chance we are wrong about this plant." State what is wrong, and repair it.

Exercise 14.18 (*critique and repair*). A commissioning record says: "The first test was close, so the rule was reapplied at a lower level, with last year's test added to the evidence, and it passed." State the errors, and repair it.

Exercise 14.19 (*critique and repair*). A status report reads: "$A(E) = 0$ at 0.95, so the probability that the opening is adequate is reported as zero until it is retested." State the error, and repair it.

Exercise 14.20 (*rapid*). Name each item: a number about the plant, a rule's value, a rule, a threshold on a quantity, a threshold on a probability, an action, or a consequence.

(a) 62.5 kW as the load to be carried.

(b) "The rule said yes."

(c) Open a second duct.

(d) A level of 0.99.

(e) Hall B above its limit for six hours.

(f) "Accept if the difference is under 5.3 K."

(g) The posterior's mode, 8347 W/K.

(h) Shed the cooling branch.

(i) "Reject if two of three tests reject."

(j) Next quarter's test's outcome, before the test.

14.7 What to retain

- A decision has actions, states, consequences and thresholds, and its two thresholds differ in kind: one inside the event, one on its probability.
- A rule is a function of the evidence, stated before the test. $A(E)$ is a value; $A(\tilde{E})$ is random until the test is read.
- Accepted is not true, not accepted is not false, and a level is not the probability delivered.
- One reading, several rules: 0.818 passes at 0.80, fails at 0.95, and a rule with no probability in it passes without asking.

Lesson 15

False Acceptance and False Rejection

15.1 The engineering question

The report states its rule at 95 %. The manager reads two lines in it. At a 70 kW opening, tested at full load, the confidence reaches 95 %. At that same opening, only 89 of 200 simulated tests pass. She asks which line is wrong. The owner, reading the contract beside her, asks another question: we accept at 95 %, so does 5 % of the bad plants get through? The engineer is sure he has heard that said, and is sure it is wrong, and cannot say why.

Nothing in the report is wrong. Three different things go by the name of a false acceptance, and each holds something different fixed: the plant, the test's outcome, or the reading in hand. Asked of one rule and one test, the same words have honest answers that differ by a factor of more than fifty.

This lesson tells the three apart, by asking of each sentence what is random and what is held. It says what the rule's 5 % is a bound on, and it leaves to the next lesson why each rate has the value it has.

15.2 The formulation

The work sits at step 3, where *Random here* now decides the answer, and at step 5, where the expression has to agree with that line. On the two-hall plant E is the 40 kW trip test, the rule is equation (14.1) at a level of 0.95, and a test means fresh sensors unless a run says otherwise.

15.2.1 The four cells

A test and a truth give four cells: accepted or not, against adequate or not. A *false acceptance* is the cell accepted and inadequate, and a *false rejection* is the cell not accepted and adequate. Table 15.1 fills them in for a 40 kW test, with a plant's conductance drawn from the report's prior and the test's readings then drawn at that conductance.

Table 15.1: Truth against outcome for a 40 kW test at a level of 0.95, under the prior on $\log G$. Each entry is a joint probability.

	Accepted [–]	Not accepted [–]	Row [–]
$G < G^*$	0.0013	0.4987	0.50
$G \geq G^*$	0.4135	0.0865	0.50

The table can be read in two directions, and Figure 15.1 marks them. Along a row it says what a test does to a plant of one kind. Down a column it says what an outcome says about a

plant. Off the table altogether stands a third thing: one reading in hand, which is not a cell at all. The whole table exists only under a prior over plants, a point the second reading will need.

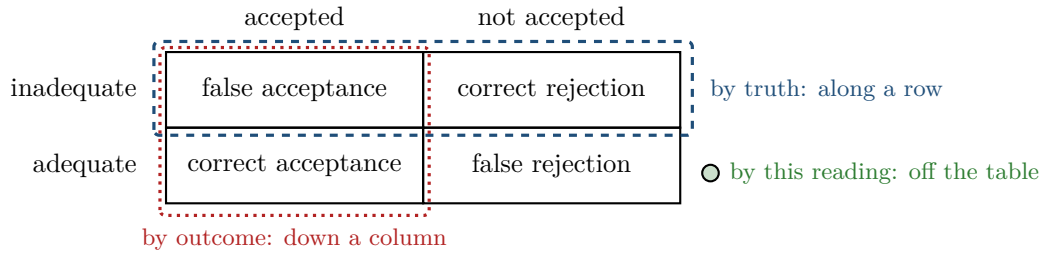


Figure 15.1: Three ways of reading a false acceptance. The dashed outline, along a row: a plant's kind is held and the outcome varies. The dotted outline, down a column: the outcome is held and the plant's kind varies. The filled point, one reading in hand: nothing about other tests enters.

15.2.2 Three conditionings

By truth. At a stated inadequate plant, how often does the rule accept? With G_0 a stated conductance below G^*,

$$P(A(\tilde{E}_Q) = 1 \mid G = G_0) = P\big(P(G \geq G^* \mid \tilde{E}_Q) \geq 0.95 \,\big|\, G = G_0\big). \tag{15.1}$$

The right side is the Expand of the left. Random here: the readings, and nothing else. Fixed: G. This is a *sampling rate*, a property of the test and the rule at that truth, and of no plant in particular. Written against G_1, a stated adequate conductance, the same expression is how often a sound plant passes, and 1 less it is a false rejection. Taken as a function of the truth it is the *operating characteristic*. Averaged over a class of plants, by the prior restricted to that class, it is $P(A(\tilde{E}_Q) = 1 \mid G < G^*)$, which is what "5 % of the bad plants" literally says, and it needs the prior's shape below G^*.

By outcome. Of the plants the rule accepts, what share are inadequate? The expression is $P(G < G^* \mid A(\tilde{E}_Q) = 1)$. Random here: the conductance and the readings, with only the outcome held. It needs a prior over plants, as a population, and on another population it has another value. Bayes' rule goes on the Model line, and it reduces to a ratio of integrals over that prior,

$$P(G < G^* \mid A(\tilde{E}_Q) = 1) = \frac{\int_{G<G^*} P(A(\tilde{E}_Q) = 1 \mid G)\, p(G)\, \mathrm{d}G}{\int P(A(\tilde{E}_Q) = 1 \mid G)\, p(G)\, \mathrm{d}G}. \tag{15.2}$$

By this reading. Given what this test read, how probable is it that the opening is inadequate? The expression is $P(G < G^* \mid E)$. Random here: G. Fixed: the readings, all of them. This is the *posterior risk*. Table 15.2 sets them side by side, with the class average in the first reading's second row, and it is the lesson's center.

Jargon: consumer's and producer's risk, predictive value

Reliability demonstration testing calls the two sampling rates at stated truths the *consumer's risk* (a false acceptance at a stated inadequate value) and the *producer's risk* (a false rejection at a stated adequate value). In acceptance sampling those are written β and α, and this book's α is neither: it is a posterior risk. Diagnostics calls the probability of the state given the test's outcome a *predictive value*: the 0.0032 below is the complement of a predictive value of 0.9968. JCGM 106 calls the third conditioning the specific risk and the joint cell the global risk.

Table 15.2: Three conditionings of a false acceptance.

Reading	Expression	Random	Held	A prior over G
By truth	$P(A(\tilde{E}_Q) = 1 \mid G = G_0)$	The readings	G	Only inside the rule
By truth, over a class	$P(A(\tilde{E}_Q) = 1 \mid G < G^*)$	The readings, G within the class	The class	Inside the rule, and below G^*
By outcome	$P(G < G^* \mid A(\tilde{E}_Q) = 1)$	G, the readings	The outcome	Inside the rule, and as the population
By this reading	$P(G < G^* \mid E)$	G	The readings	Yes

15.2.3 What the rule's level bounds

A rule at $1 - \alpha$ keeps $P(G < G^* \mid e) \leq \alpha$ on every reading it accepts. That is the whole of what the level says directly: it is a bound on the third conditioning, on accepted readings. Where the rule's own prior is the population, the second conditioning is the average of those risks over the accepted readings, so it too is at most α, but only as that average. It says nothing about the first. The report puts it in these words, which with its bar mean a bound.

> α is the risk of a false acceptance the owner will carry

That sentence means the third conditioning, and it is the one a reader will mishear. Table 15.3 takes one rule and one test and gives the six honest answers to "what is the chance of a false acceptance?".

Table 15.3: Six answers to one sentence, for a 40 kW test and the rule at 0.95.

The question asked	Conditioning	Value [–]
On the reading at which the rule just accepts	This reading	0.05
On E itself, which the rule did not accept	This reading	0.182
Among the plants accepted, under the prior	Outcome	0.0032
Averaged over the prior's inadequate plants	Truth, over a class	0.0026
At an opening of 54 kW	Truth	0.010
At an opening just below G^*	Truth	0.065

The last row exceeds the level's 0.05. A plant at exactly G^* is adequate, so 0.065 there is the limit of the false acceptance just below it. A rule on a posterior does not hold that rate to α: at 60, 40 and 20 kW it is 0.060, 0.065 and 0.078, and the report's counts at G^* are 17, 13 and 15 of 200. Each of the six numbers is correct, and they differ by a factor of 70.

15.2.4 The named case

The manager's two lines are one plant and one rule. Table 15.4 reads them from the report's data for each of its test loads, and its two value columns are kept apart because they are two questions. The first is the design case: the posterior $P(G \geq G^* \mid d = Q/G_1)$ at the reading the truth predicts, with no probability over readings. The second is a count of 200 simulated tests, which estimates the sampling rate $P(A(\tilde{E}_Q) = 1 \mid G = G_1)$.

One line holds the reading and lets G vary, and the other holds G and lets the reading vary. Neither is wrong, and neither is an estimate of the other. That the count sits near half is stated here and not explained. Lesson 16 shows why.

Table 15.4: The design case against the count at the cell where the design case first reaches 0.95, on the report's data. Each row is a different plant.

Test load [kW]	Cell, the true capacity [kW]	Design case there [–]	Tests accepted [of 200]
60	70	0.956	89
40	80	0.978	131
20	100	0.958	103

15.2.5 A false acceptance at a stated truth

The first worked formulation takes the near pair's smaller drawing as the truth and the plant's usual test. The sheet is Table 15.5, and its third row is the one to read.

Table 15.5: The formulation sheet filled in for a false acceptance at a 54 kW opening, from a 40 kW test.

Field	Entry
Intention	Predict
1. Sentence	The opening's true capacity is 54 kW. How often does a 40 kW test accept it at 0.95?
2. Target	An outcome of a test not yet run, at a fixed truth
3. Known	$G = G_0$, stated and not evidence; $Q = 40\,\text{kW}$; unit B out; settled; fresh sensors on each test (assumption). No evidence in hand.
4. Event	$A(\tilde{E}_{40}) = 1$
5. Expression	$P(A(\tilde{E}_{40}) = 1 \mid G = G_0)$
6. Model	The readings' relation at G_0; the rule's own prior, inside the rule
7. Read back	How often this rule passes a plant of this size. Silent on any plant actually in hand, and it is not α.
Problem class	Operating characteristic at a stated truth

Its run, with the rule substituted:

Run, a false acceptance at a stated truth

1	**Sentence**	The opening's true capacity is 54 kW: how often does a 40 kW test accept it at 0.95?
	Intention	Predict
2	**Target**	An outcome of a test not yet run, at a fixed truth
	Declare	$\tilde{E}_{40}$ readings not yet taken; $A(\cdot)$ at 0.95; $G_0 = 54\,\text{kW}/8\,\text{K} = 6750\,\text{W/K}$; $G^* = 7500\,\text{W/K}$
3	**Known**	$G = G_0$ (stated); 40 kW; unit B out; settled; fresh sensors (assumption)
	Random here	The readings only; G is held
4	**Event**	$A(\tilde{E}_{40}) = 1$
5	**Expression**	$P(A(\tilde{E}_{40}) = 1 \mid G = G_0)$
	Expand	$P(P(G \geq G^* \mid \tilde{E}_{40}) \geq 0.95 \mid G = G_0)$; inner G integrated, outer held
6	**Model**	The readings' relation at G_0; the rule's own prior
7	**Check**	About no plant in hand, and not α: 0.010, which the next lesson derives
	Class	Operating characteristic at a stated truth

$$\boxed{P(A(\tilde{E}_{40}) = 1 \mid G = G_0)}$$

15.2.6 The share of accepted plants that are inadequate

The second holds the outcome, and so a prior enters the formulation as a statement about a population.

Run, the share of accepted plants that are inadequate		
1	**Sentence**	Plants drawn from the prior, each given a 40 kW test: what share of those accepted are inadequate?
	Intention	Predict
2	**Target**	A state, given a test's outcome
	Declare	$\tilde{E}_{40}$; $A(\cdot)$ at 0.95; $p(G)$ the prior, read as a population (assumption)
3	**Known**	The outcome alone; the population; fresh sensors (assumption)
	Random here	G and the readings
4	**Event**	$G < G^*$
5	**Expression**	$P(G < G^* \mid A(\tilde{E}_{40}) = 1)$
	Expand	The rule's condition behind the bar: $P(G \geq G^* \mid \tilde{E}_{40}) \geq 0.95$
6	**Model**	Bayes' rule on the population
	Reduce	Equation (15.2), at $Q = 40\,\text{kW}$
7	**Check**	A fleet's share: 0.0032, and 0.0097, 0.0018 at 20, 60 kW; silent on a plant with its reading in hand
	Class	Predictive value of a test's outcome
		$\boxed{P(G < G^* \mid A(\tilde{E}_{40}) = 1)}$

15.2.7 An alarm on the network

The third is the same pair on a hypothesis, and its size surprises. An alarm flags when the PDU and UPS output meters differ by at least 6 kW, a limit that is illustrative, and y is their difference as in Lesson 14.

Run, an alarm, two ways		
1	**Sentence**	An alarm at 6 kW: (a) with no offset, how often does it flag? (b) Of flags, how many have no offset?
	Intention	Predict
2	**Target**	(a) A sampling rate; (b) a state, given the outcome
	Declare	$\tilde{d}$ the difference not yet read [kW]; flag the event $\lvert\tilde{d}\rvert \geq 6\,\text{kW}$; H_{off}; limit (illustrative)
3	**Known**	The fact sheet's spreads; (a) no offset explanation, for a meter drawn from its $0 \pm 1.5\,\text{kW}$; (b) its weights 0.05 and 0.95
	Random here	(a) The reading and the small offset; (b) the explanation and the reading
4	**Event**	(a) A flag; (b) no offset
5	**Expression**	(a) $P(\text{flag} \mid H^c_{\text{off}})$; (b) $P(H^c_{\text{off}} \mid \text{flag})$
	Expand	flag: $\lvert\tilde{d}\rvert \geq 6\,\text{kW}$, in each
6	**Model**	$\tilde{d}$ normal about 0, spread 3.2 kW, or 45 kW under H_{off}; (b) Bayes' rule
7	**Check**	(a) 0.061, across such meters, not over the hours of one; (b) 0.564: more than half of alarms find no offset
	Class	(a) Operating characteristic; (b) predictive value of a test's outcome
		$\boxed{\text{(a) } P(\text{flag} \mid H^c_{\text{off}}); \quad \text{(b) } P(H^c_{\text{off}} \mid \text{flag})}$

15.2.8 A sentence that is not ready

A vendor's slide says: "Our false acceptance is below 1 %." At which truth, or across which population, or on which reading? Read by truth, on this plant's rule, it holds at 54 kW from a 40 kW test, at 0.0095, and fails just below G^*, at 0.065. Read by outcome, it holds, at 0.0032, on this prior and on no other. Read by this reading, it fails on E, at 0.182, and the slide has no reading. The analyst asks which one is meant.

Key idea

False acceptance has three conditionings: by truth, by outcome, and by this reading. Ask what is random and what is held. The level bounds the third directly, on accepted readings, and the second only as its average. A sampling rate belongs to a test at a truth, and a posterior risk to a reading, and neither can stand in for the other.

15.3 The competing formulations

Each competing formulation takes a statement conditioned one way and uses it as if it were conditioned another. The first is the owner's sentence.

Run, the level read as a sampling rate

1	**Sentence**	"We accept at 95 %, so 5 % of the bad plants get through."
5	**Expression** ×	$P(A(\tilde{E}_Q) = 1 \mid G < G^*) = \alpha$
	Repair	The class average on the prior: 0.0026 from a 40 kW test; or a stated truth: 0.0095 at 54 kW
7	**Check**	The level is a bound on $P(G < G^* \mid e)$, and neither a truth nor a class is held in it

$$\boxed{P(A(\tilde{E}_Q) = 1 \mid G < G^*)}$$

The second runs the other way, from a rate to a plant.

Run, a sampling rate read as a posterior risk

1	**Sentence**	"Only 1 % of such plants pass, so this accepted plant is adequate with probability 0.99."
5	**Expression** ×	$P(G \geq G^* \mid e) = 1 - P(A(\tilde{E}_{40}) = 1 \mid G = G_0)$
	Repair	$P(G \geq G^* \mid e)$ from this reading: 0.95 on one that just passes
7	**Check**	The sentence transposes the bar, and it drops the reading in hand

$$\boxed{P(G \geq G^* \mid e)}$$

The third reads a posterior as how often.

Run, a posterior read as a frequency

1	**Sentence**	"At a 70 kW opening the full-load design case is 0.956, so the test passes 96 % of the time."
5	**Expression** ×	$P(A(\tilde{E}_{60}) = 1 \mid G = G_1) = P(G \geq G^* \mid d = 60\,\mathrm{kW}/G_1)$
	Repair	Two expressions: 0.956 for the right side, about half for the left, with 89 of 200 by count
7	**Check**	The right side has no probability over readings

$$\boxed{P(A(\tilde{E}_{60}) = 1 \mid G = G_1)}$$

The fourth reads a count as a statement about the plant, whose capacity the count had stated.

Run, a count read as a probability of adequacy

1	**Sentence**	"89 of 200 passed, so the opening is only 45 % likely to be adequate."
5	**Expression** ×	$P(G \geq G^* \mid \cdot) = 0.445$
	Repair	A count estimating $P(A(\tilde{E}_{60}) = 1 \mid G = G_1)$, with G_1 stated at 8750 W/K
7	**Check**	The simulated plant is adequate by construction; what is low is how often the test says so

$$\boxed{P(A(\tilde{E}_{60}) = 1 \mid G = 8750\,\mathrm{W/K})}$$

The fifth averages where a truth was stated. The number is the class average, correct for

the whole table's row and wrong for the opening asked about.

Run, a population average for a stated truth

1	**Sentence**	At an opening just below G^*: how often does a 40 kW test accept?
3	**Known** ×	$G < G^*$ over the prior, for a truth that was stated: 0.0026
	Repair	G_0 stated, just below G^*: $P(A(\tilde{E}_{40}) = 1 \mid G = G_0)$, 0.065
7	**Check**	The prior puts most of its inadequate plants far below G^*, which are easy to refuse

$$\boxed{P(A(\tilde{E}_{40}) = 1 \mid G = G_0),\ G_0 \text{ just below } G^*}$$

The sixth holds the outcome with the reading in hand, and it throws the reading away.

Run, the outcome for the reading

1	**Sentence**	"The trip test did not accept, so the opening is adequate with probability 0.148."
3	**Known** ×	$A(\tilde{E}_{40}) = 0$ alone, with the readings of E in hand
	Repair	$P(G \geq G^* \mid E) = 0.818$: the reading says more than its outcome
7	**Check**	0.148 averages over every reading that fails; this one failed narrowly

$$\boxed{P(G \geq G^* \mid E)}$$

These are the forms of this lesson's failure class, *posterior risk read as a sampling error rate*: a statement conditioned one way used as if it were conditioned another. Its four forms are the level read as a sampling rate; a sampling rate read as a posterior risk; a posterior value read as a frequency, or a count as a probability of adequacy; and a population average or an outcome used where a truth or a reading is stated. It enters at step 3 or step 5, and it is exposed by writing *Random here* honestly, since the expression and that line then disagree. It sharpens Lesson 3's transposed conditional, whose events were stated: here one side of the bar is a test's outcome over readings not yet taken.

15.4 One system, many questions

Table 15.6 holds the network's alarm still and sorts ten statements about it by conditioning. The rates are those of the run above and, for the stated offset, its model.

Table 15.6: Ten statements about one alarm on the network, by conditioning, with what is random in each.

Statement	Conditioning	Random
With no offset explanation, it trips on 6 % of such meters	Truth, over the small offset	The reading, the offset
With an offset, it trips 89 % of the time	Truth, over the offset's prior	The reading, the offset
With an offset of 8 kW, it trips 76 % of the time	Truth, stated	The reading
Of alarms raised, 56 % have no offset	Outcome	Explanation, reading
Of quiet hours, how many hide an offset	Outcome, the complement	Explanation, reading
Tonight reads 12 kW: an offset is 80 % likely	This reading	The explanation
Tonight the alarm tripped, and y is not given	Outcome only	Explanation
The limit of 6 kW	Stated	Nothing
Whether the limit should move	A decision: Lesson 17	
"The alarm is 97 % accurate"	Not ready: on which row	

15.5 Worked problems

Seven problems on the three systems follow, each with its run.

Problem 15.1. Of the prior's adequate plants, how often does a test at the report's rule refuse, by test load? How does that stand to one 80 kW opening?

Run, a false rejection, over the class and at a truth

1	**Sentence**	Tests at 20, 40, 60 kW, rule at 0.95: how often refused, (a) over adequate plants; (b) at 80 kW?
	Intention	Predict
2	**Target**	An outcome of a test not yet run: (a) over a class; (b) at a fixed truth
	Declare	(b) $G_1 = 80\,\text{kW}/8\,\text{K} = 10\,000\,\text{W/K}$; $\tilde{E}_Q$; $A(\cdot)$ at 0.95; (a) the prior above G^*
3	**Known**	(a) $G \geq G^*$; (b) $G = G_1$ (stated); Q; unit B out; settled; fresh sensors (assumption)
	Random here	(a) The readings and G in its class; (b) the readings only
4	**Event**	$A(\tilde{E}_Q) = 0$
5	**Expression**	(a) $P(A(\tilde{E}_Q) = 0 \mid G \geq G^*)$; (b) $P(A(\tilde{E}_Q) = 0 \mid G = G_1)$
	Expand	(b) $P(P(G \geq G^* \mid \tilde{E}_Q) < 0.95 \mid G = G_1)$; (a) the same, averaged
6	**Model**	The readings' relation; the rule's prior; (a) the prior's shape above G^*
7	**Check**	(a) 0.369, 0.173, 0.113; (b) 0.684, 0.358, 0.104: the class holds plants far above their need
	Class	Operating characteristic, (a) over a class and (b) at a stated truth

$$\boxed{\text{(a) } P(A(\tilde{E}_Q) = 0 \mid G \geq G^*); \quad \text{(b) } P(A(\tilde{E}_Q) = 0 \mid G = 10\,000\,\text{W/K})}$$

Problem 15.2. Unit A is accepted over one year in service when its count is at most one, which is the rule of Problem 14.3. If its true rate is 3.5 per year, how often does that happen?

Run, a rate above its level, on a count

1	**Sentence**	Accept at $k \leq 1$ in one year in service, true rate 3.5: how often accepted?
	Intention	Predict
2	**Target**	An outcome of a record not yet kept, at a fixed truth
	Declare	$\tilde{k}$ a count not yet logged; λ_A [1/yr]; $t_\text{s} = 1$ yr; rule at 0.90 (illustrative)
3	**Known**	$\lambda_A = 3.5$ (stated, not evidence); the exposure; a Poisson count (assumption)
	Random here	The count
4	**Event**	$A(\tilde{k}) = 1$
5	**Expression**	$P(A(\tilde{k}) = 1 \mid \lambda_A = 3.5)$
	Expand	$P(\tilde{k} \leq 1 \mid \lambda_A = 3.5)$
6	**Model**	Poisson at $\lambda_A t_\text{s}$
	Reduce	$e^{-3.5}(1 + 3.5)$
7	**Check**	0.136 at the claim's edge, above the rule's $\alpha = 0.10$: the same finding as at G^*
	Class	Operating characteristic at a stated truth

$$\boxed{P(\tilde{k} \leq 1 \mid \lambda_A = 3.5) = e^{-3.5}(1 + 3.5)}$$

Problem 15.3. The contractor proposes to accept if either a 20 kW test or a 40 kW test accepts, run in turn with fresh sensors. How often does a 54 kW opening pass?

Run, either of two tests

1	**Sentence**	A 54 kW opening, accept if the 20 kW or the 40 kW test accepts: how often accepted?
	Intention	Predict
2	**Target**	An outcome of a two-test rule, at a fixed truth
	Declare	$\tilde{E}_{20}$, $\tilde{E}_{40}$; $G_0 = 6750\,\mathrm{W/K}$; each rule at 0.95
3	**Known**	$G = G_0$ (stated); fresh sensors on each test (assumption)
	Random here	Two sets of readings, independent given G_0
4	**Event**	$\{A(\tilde{E}_{20}) = 1\} \cup \{A(\tilde{E}_{40}) = 1\}$
5	**Expression**	$P(A(\tilde{E}_{20}) = 1 \text{ or } A(\tilde{E}_{40}) = 1 \mid G = G_0)$
	Expand	Each $A(\tilde{E}_Q) = 1$ is $P(G \geq G^* \mid \tilde{E}_Q) \geq 0.95$
6	**Model**	Independent given G_0 (assumption)
	Reduce	$1 - (1 - 0.033)(1 - 0.010)$
7	**Check**	0.043: the rate grows with each try, and the level printed on the rule does not change
	Class	Operating characteristic at a stated truth

$$\boxed{1 - \prod_{Q} \big(1 - P(A(\tilde{E}_Q) = 1 \mid G = G_0)\big)}$$

Problem 15.4. The sales engineer says: "The test is 95 % reliable."

Run, a percentage with nothing behind its bar

1	**Sentence**	"The test is 95 % reliable."
	Intention	Not stated
3	**Known**	Missing: of what; at which truth; over which population; on which reading
4	**Event**	(a) A sound plant passes; (b) an accepted plant is sound; (c) the rule's level
5	**Expression**	(a) by truth; (b) by outcome; (c) a statement about the rule, with no expression
7	**Check**	(a) 0.203 at 66 kW from a 40 kW test; (b) 0.9968 on the prior; (c) 0.95. The analyst asks
	Class	Not ready

$$\boxed{\text{(a) } P(A(\tilde{E}_Q) = 1 \mid G = G_1); \quad \text{(b) } P(G \geq G^* \mid A(\tilde{E}_Q) = 1)}$$

Problem 15.5. The level is raised from 0.95 to 0.99 at a 40 kW test. What happens to a false acceptance at 54 kW and to a pass at 66 kW?

Run, a stricter level

1	**Sentence**	40 kW test, level 0.99 against 0.95: acceptance at (a) 54 kW; (b) 66 kW?
	Intention	Compare
2	**Target**	Two rates under two rules
	Declare	$G_0 = 6750\,\mathrm{W/K}$; $G_1 = 8250\,\mathrm{W/K}$; $A_{p_{\mathrm{req}}}(\cdot)$ the rule at level $p_{\mathrm{req}} = 0.95$, 0.99
3	**Known**	Each truth (stated); 40 kW; fresh sensors (assumption)
	Random here	The readings only
4	**Event**	$A_{p_{\mathrm{req}}}(\tilde{E}_{40}) = 1$, at each level
5	**Expression**	(a) $P(A_{p_{\mathrm{req}}}(\tilde{E}_{40}) = 1 \mid G = G_0)$; (b) the same at G_1
	Expand	$P(P(G \geq G^* \mid \tilde{E}_{40}) \geq p_{\mathrm{req}} \mid G = \cdot)$
6	**Model**	The readings' relation; the rule's prior
7	**Check**	(a) 0.0095 to 0.0012; (b) 0.203 to 0.065: both move, so the level trades one for the other
	Class	Operating characteristic at a stated truth

$$\boxed{P(A_{p_{\mathrm{req}}}(\tilde{E}_{40}) = 1 \mid G = G_0) \text{ and } P(A_{p_{\mathrm{req}}}(\tilde{E}_{40}) = 1 \mid G = G_1)}$$

Problem 15.6. Forty plants drawn from the prior have been accepted from a 40 kW test. How many are expected to be inadequate?

Run, a fleet	
1 Sentence	Forty plants accepted from a 40 kW test, drawn from the prior: how many inadequate, on average?
Intention	Predict
2 Target	An expected count, given an outcome
Declare	$n = 40$; $p(G)$ the population; $A(\cdot)$ at 0.95
3 Known	Each outcome; the population (assumption); independent plants, fresh sensors (assumptions)
Random here	Each plant's G
4 Event	Each plant's $G < G^*$
5 Expression	$\mathbb{E}$[inadequate | all accepted]
Expand	Each plant's $A(\tilde{E}_{40}) = 1$ is $P(G \geq G^* \mid \tilde{E}_{40}) \geq 0.95$
6 Model	By outcome, plant by plant
Reduce	$40 \times P(G < G^* \mid A(\tilde{E}_{40}) = 1) = 40 \times 0.0032$
7 Check	0.13; a plant whose reading is on the desk uses its reading, and not this share
Class	Predictive value of a test's outcome
	$\boxed{n\,P(G < G^* \mid A(\tilde{E}_{40}) = 1)}$

Problem 15.7. A timed-discharge rule keeps a string at a level of 0.90 if it gives its 2 h. How often is a sound new string, or one at 0.99 of nameplate, refused? Both are adequate.

Run, sound strings refused	
1 Sentence	Timed-discharge rule at 0.90, $\varphi = 1$ and $\varphi = 0.99$: how often is the string refused?
Intention	Predict
2 Target	An outcome of a test not yet run, at a fixed truth
Declare	$\tilde{d}$ the timed reading not yet taken [h]; spread 0.03 h (illustrative); $T_{\text{rt}} = 2.025\ \text{h} \times \varphi$
3 Known	φ (stated); sound string (assumption)
Random here	The reading only
4 Event	$A(\tilde{d}) = 0$
5 Expression	$P(A(\tilde{d}) = 0 \mid \varphi)$
Expand	$P(P(T_{\text{rt}} \geq 2\ \text{h} \mid \tilde{d}) < 0.90 \mid \varphi)$
6 Model	The timed relation; the prior's bound at $\varphi = 1$, inside the rule
7 Check	0.978 and 0.996: the rule almost never passes an adequate string, and Lesson 16 shows why
Class	Operating characteristic at a stated truth
	$\boxed{P(A(\tilde{d}) = 0 \mid \varphi)}$

15.6 The exercises

The solutions are in Appendix A, and the facts of the three systems in Appendix D. On the two-hall plant, E is the 40 kW trip test, and "the rule" is the one at 0.95 unless an exercise says otherwise.

Exercise 15.1 (*formulate*). A vendor's data sheet says: "95 % acceptance confidence. Under 1 % false pass. Field-proven on 200 sites, of which 4 % were later reworked." Write each clause as an expression, marking what is random and what is held, or say what is missing.

Exercise 15.2 (*formulate*). Run this one: at a 90 kW opening, how often does a 40 kW test refuse?

Exercise 15.3 (*formulate*). The fact sheet's weights on the explanations are replaced by 0.5 for an offset and 0.5 for none. Which of the alarm's two rates moves and which does not? Write both, and say why.

Exercise 15.4 (*formulate*). Run this one: of plants drawn from the prior and accepted from a 20 kW test by the rule at 0.99, what share are inadequate?

Exercise 15.5 (*formulate*). A timed-discharge rule is run on one string every year, and also on a list of strings once each. Write the rate each one estimates, and say which needs a share of weak strings.

Exercise 15.6 (*formulate*). The commissioning agent writes: "At a true capacity of 55 kW, what is the false-pass rate of a 60 kW trip test at the report's rule?" Formulate the question.

Exercise 15.7 (*formulate*). A utility's account manager says: "We almost never reject a good feeder." Write the expression for a stated rule on the record's count, with its truth stated, and say what "good" has to be given.

Exercise 15.8 (*formulate*). Of 1000 plants commissioned (illustrative), 2 were inadequate and accepted, 98 inadequate and not accepted, 700 adequate and accepted, 200 adequate and not accepted. Write all four conditional rates with what is held in each, and say which of them a level would bound if the rule were a posterior one.

Exercise 15.9 (*formulate*). The commissioning agent writes: "What is the false-pass rate of the trip test?" Formulate the question.

Exercise 15.10 (*formulate*). The PDU meter is accepted at each calibration if it agrees with the UPS output meter within 6 kW. Write the false rejection at a stated offset of zero, and say what is random in it.

Exercise 15.11 (*reverse*). Write one plain sentence for $P(A(\tilde{E}_{40}) = 0 \mid G = 12\,000\,\mathrm{W/K})$.

Exercise 15.12 (*reverse*). Write one plain sentence for $P(G \geq G^* \mid A(\tilde{E}_{60}) = 0)$, and say what it needs that $P(A(\tilde{E}_{60}) = 0 \mid G = G_1)$ does not.

Exercise 15.13 (*reverse*). Write one plain sentence for $P(G < 6250\,\mathrm{W/K} \mid E)$, and say why no rule appears in it.

Exercise 15.14 (*reverse*). Write one plain sentence for $P(P(G \geq G^* \mid \tilde{E}_{20}) \geq 0.99 \mid G = 8000\,\mathrm{W/K})$, saying which G is which.

Exercise 15.15 (*critique and repair*). A proposal says: "With $\alpha = 0.05$ the test has a 95 % chance of catching an inadequate plant." State what is wrong, and repair it.

Exercise 15.16 (*critique and repair*). A meter report says: "The alarm trips on only 6 % of healthy meters, so tonight's trip means this meter has an offset with probability 0.94." State what is wrong, and repair it.

Exercise 15.17 (*critique and repair*). A string report says: "At nameplate the posterior at the expected reading is 0.595, so about 60 % of timed discharges on a new string will pass the rule at 0.90." State what is wrong, and repair it.

Exercise 15.18 (*critique and repair*). A reviewer writes: "We ran 200 tests on a 60 kW opening at full load and 17 passed, so the rule's true level is 0.915, not 0.95." State the errors, and repair it.

Exercise 15.19 (*critique and repair*). A report on the string fleet says: "Half the strings in our program are weak, so the prior is right, and the fleet's share of 0.3 % inadequate among accepted strings applies to tonight's string, which read low." State the errors, and repair it.

Exercise 15.20 (*rapid*). For each, say by truth, by outcome, by this reading, or not ready.

(a) “A string at 0.90 of nameplate fails this test nine times in ten.”

(b) “Three of the last ten rejected units turned out to be fine.”

(c) “Given today’s reading, the efficiency is below 0.93 with probability 0.2.”

(d) “The feeder test catches bad feeders.”

(e) “Of breakers passed by the relay test, 2 % later fail to clear.”

(f) “This relay would pass one time in fifty if it were out of calibration by 5 %.”

(g) “The audit is 99 % sure.”

(h) “This morning’s reading leaves a 3 % chance the breaker is out of calibration.”

(i) “Sound strings fail this test almost always.”

(j) “Given that it failed and nothing else, the unit is probably fine.”

15.7 What to retain

- “False acceptance” has three conditionings: by truth, by outcome, by this reading. Ask what is random and what is held.
- The level bounds the posterior risk on accepted readings, directly, and the share of accepted plants only as its average. It does not bound a sampling rate, which here reaches 0.065 just below G^{*}.
- 0.956 and 89 of 200 are two questions about one plant.
- The share of accepted plants that are inadequate needs a prior over plants, as a population, and it is silent on a plant with its reading in hand.

Lesson 16

Statistical Power and Test Design

16.1 The engineering question

The drawing says 66 kW. The commissioning test can be run at 40 kW, perhaps at 60. The owner is about to book the crew and the load banks, and he asks: if this plant is exactly what its drawing says, how often will the test pass? The engineer can already say what any one reading would imply. She cannot say, before a reading exists, how often the rule will say yes. The answer, at 40 kW, is about one time in five.

The question feels harder than the ones before it, and the sentence that states it in a report is harder still to read. It holds two probabilities, one inside the other, and they are over different things. A reader who has no trouble writing the rule can lose the thread here, in the place where the rule's own outcome becomes the thing that is uncertain.

There is a path through it that is mechanical. This lesson gives the path, four moves in a fixed order, and it ends at a standard normal tail.

16.2 The formulation

The work sits at step 5's Expand, which now produces a nested probability, and at step 6's Reduce, which takes it to Φ. On the two-hall plant, E is the 40 kW trip test, and a test means fresh sensors unless a run says otherwise.

16.2.1 Parse, Translate, Expand, Reduce

The path has four moves, each on a step of the sheet. *Parse* (steps 2 to 4) reads the sentence for four things, and Table 16.1 does it for the owner's sentence.

Table 16.1: Parsing "how often will a test at load Q accept, at a true conductance G_1?".

Parse	The item	In the sentence
Random variable	The test's evidence $\tilde{E}_Q$, here its difference $\tilde{d}$	"a test"
Event	$A(\tilde{E}_Q) = 1$	"accept"
Condition	$G = G_1$	"at a true conductance"
Rule	$A(e) = \mathbf{1}[P(G \geq G^* \mid e) \geq 0.95]$	defined elsewhere

Translate (step 5) writes the event behind the condition:

$$P(A(\tilde{E}_Q) = 1 \mid G = G_1). \tag{16.1}$$

Expand (step 5) substitutes the rule:

$$P\big(P(G \geq G^* \mid \tilde{E}_Q) \geq 0.95 \,\big|\, G = G_1\big). \tag{16.2}$$

This is a probability inside a probability, and it is worth stopping on. Figure 16.1 marks its two parts. The inner probability is over G, for each reading the test could give, since the rule does not know the truth. The outer one is over the readings, with G held at G_1. There are two G's in the expression and they do different work, which is why the outer one is written G_1. *Reduce* (step 6) then takes it to a standard form, by two facts, each from a stated model.

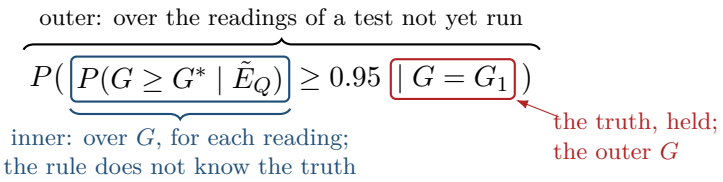


Figure 16.1: The nested expression. The inner probability, boxed in blue, is over G and is computed for each possible reading. The outer one, marked above, is over the readings. The condition, boxed in red, holds the truth while the readings vary.

16.2.2 The inner as a threshold on the reading

With one reading and one unknown, the larger $\tilde{d}$ is, the smaller $P(G \geq G^* \mid \tilde{d})$ is. So the event inside is a threshold on the reading: the rule accepts exactly when $\tilde{d} \leq d_{\text{crit}}(Q)$, where the *critical difference* is defined by

$$P(G \geq G^* \mid d_{\text{crit}}(Q)) = 0.95, \tag{16.3}$$

one number per test load. Table 16.2 gives it in three ways. The first column is exact. The second uses a Gaussian reading in R, the coordinate in which Lesson 10 found the relation linear:

$$P(G \geq G^* \mid \tilde{d}) \approx \Phi\left(\frac{R^* - m_R(\tilde{d})}{\sigma_R}\right), \quad \text{Gaussian reading in } R, \tag{16.4}$$

which may be used only where Principle 10.1's four steps pass. The table's last column runs them at the critical reading. Step 1 passes in R, with the fixed spread σ_d, and step 2 passes on the stated prior, log-normal with a factor of 3. Steps 3 and 4 depend on the load.

Table 16.2: The critical difference at 0.95. The last column runs Lesson 10's checklist at the exact critical reading: the distance from $R = 0$ in reading spreads, and the slope.

Q [kW]	Exact [K]	Reading in R [K]	Flat in R, $QR^* - 1.645\sigma_d$ [K]	Checklist
20	1.659	1.710	1.497	Step 3 refuses: 2.33 spreads, under 3; slope 0.260
40	4.259	4.260	4.163	Passes: 5.99 spreads, slope 0.136
60	6.894	6.894	6.830	Passes: 9.69 spreads, slope 0.090

At 40 and 60 kW the reading agrees with the exact value within 0.002 K. At 20 kW the checklist refuses at step 3, the book writes the exact expression, and the 1.710 shows what was

avoided. The procedure works by saying no. Run at 40 kW, the inner reads as follows.

Run, the acceptance condition, in R		
1	**Sentence**	A 40 kW test, the rule at 0.95: for which readings does it accept, written in R?
	Intention	Decide: which readings the rule accepts, before the outer probability
2	**Target**	The set of readings the rule accepts
	Declare	$\tilde{d}$ [K]; $R = 1/G$, $R^* = 0.1333$ K/kW; $m_R(\tilde{d})$, σ_R the reading's center and spread [K/kW]
3	**Known**	The stated prior; $\sigma_d = 0.711$ K; steps 1 to 4 of Principle 10.1 pass here (5.99 spreads, slope 0.136)
	Random here	G, inside the rule; $\tilde{d}$ read as a variable
4	**Event**	$P(G \geq G^* \mid \tilde{d}) \geq 0.95$
5	**Expression**	$\{\tilde{d} : P(R \leq R^* \mid \tilde{d}) \geq 0.95\}$
6	**Model**	$\tilde{d} = QR + \varepsilon$, linear in R; the prior transformed
	Reduce	$\Phi((R^* - m_R(\tilde{d}))/\sigma_R) \geq 0.95$, Gaussian reading in R: $\tilde{d} \leq 4.26$ K
7	**Check**	Exact, 4.259 K; a threshold on the reading, one per load
	Class	Decision rule
		$\boxed{\{\tilde{d} : (R^* - m_R(\tilde{d}))/\sigma_R \geq 1.645\} = \{\tilde{d} \leq d_{\text{crit}}(40\,\text{kW})\}}$

16.2.3 The reading over tests

The measurement relation already says it, with G held at G_1: $\tilde{d}$ is normal about Q/G_1 with spread σ_d. The set point cancels in the difference, and the sensors' offsets and the averaged noise add, as Lesson 11 found. So

$$P(A(\tilde{E}_Q) = 1 \mid G = G_1) = P(\tilde{d} \leq d_{\text{crit}}(Q) \mid G_1) = \Phi\left(\frac{d_{\text{crit}}(Q) - Q/G_1}{\sigma_d}\right). \qquad (16.5)$$

Given $d_{\text{crit}}(Q)$, this Φ is exact under the stated model: the prior enters only through $d_{\text{crit}}(Q)$, and a critical difference found by the reading in R carries that reading's error in. It is not a Gaussian reading, and the outer step asks nothing of the checklist. The only Gaussian reading in this lesson is equation (16.4), in the inner, in R. Keeping the two Φ's apart is the lesson's main caution: one is a statement about G for a reading, and the other about a reading for a G.

16.2.4 How often a test accepts

Table 16.3 is the sheet for the owner's sentence, and its run is the lesson's model.

For the 66 kW drawing, $G_1 = 8250$ W/K, the run gives 0.141 at 20 kW, 0.203 at 40 kW and 0.297 at 60 kW.

16.2.5 The operating characteristic

Equation (16.5) as a function of the truth is the operating characteristic, plotted in Figure 16.2 for the three test loads, with the report's counts of 200 as crosses. At a stated adequate truth its value is the test's *power*. Below G^* it is Lesson 15's false acceptance. It is one curve and one expression.

The reading's density shows where each value comes from. Figure 16.3 draws a 40 kW test's difference at two truths, 54 and 66 kW, with the critical difference marked. The area below that line is the rate at each truth: 0.010, barely visible, and 0.203. A larger opening moves the whole bell to the left, and the rule's line stays where the rule put it.

One more form can be read by eye. On a prior flat in R, $d_{\text{crit}}(Q) = QR^* - 1.645\,\sigma_d$, the fourth column of Table 16.2, and the power becomes

$$\Phi\left(\frac{Q(R^* - 1/G_1)}{\sigma_d} - 1.645\right), \quad \text{on a prior flat in } R. \qquad (16.6)$$

Table 16.3: The formulation sheet filled in for how often a test accepts at a stated truth.

Field	Entry
Intention	Predict
1. Sentence	A test at load Q, the rule at 0.95, a true conductance G_1: how often will the test accept?
2. Target	An outcome of a test not yet run, at a fixed truth
3. Known	$G = G_1$, stated and not evidence; Q; unit B out; settled. Assumptions: $\delta = 0$; fresh sensors on each test. No evidence in hand.
4. Event	$A(\tilde{E}_Q) = 1$
5. Expression	Equation (16.1), expanded as equation (16.2)
6. Model	$\tilde{d}$ given G_1, normal about Q/G_1, spread σ_d; the prior, in the inner
7. Read back	How often the rule accepts this plant. Silent on any plant in hand, and it assumes the drawing.
Problem class	Operating characteristic at a stated truth

Its run, with Parse at steps 2 to 4, Translate and Expand at step 5, and Reduce at step 6:

Run, how often a test accepts at a stated truth

1	**Sentence**	Test at Q, rule at 0.95, true conductance G_1: how often will the test accept?
	Intention	Predict
2	**Target**	An outcome of a test not yet run, at a fixed truth
	Declare	$\tilde{E}_Q$ readings not yet taken; $\tilde{d}$ their difference [K]; $A(\cdot)$ the rule; G_1 [W/K]; $\sigma_d = 0.711\,\mathrm{K}$; $d_{\mathrm{crit}}(Q)$ [K]
3	**Known**	$G = G_1$ (stated, not evidence); Q; unit B out; settled; $\delta = 0$, fresh sensors (assumptions)
	Random here	$\tilde{d}$ only; G is fixed; no evidence in hand
4	**Event**	$A(\tilde{E}_Q) = 1$
5	**Expression**	$P(A(\tilde{E}_Q) = 1 \mid G = G_1)$
	Expand	$P(P(G \geq G^* \mid \tilde{E}_Q) \geq 0.95 \mid G = G_1)$; inner G integrated, outer held
6	**Model**	$\tilde{d}$ given G_1 normal about Q/G_1, spread σ_d; the prior, in the inner
	Reduce	Inner falls with $\tilde{d}$: $\{\tilde{d} \leq d_{\mathrm{crit}}(Q)\}$, $P(G \geq G^* \mid d_{\mathrm{crit}}(Q)) = 0.95$; then Φ
7	**Check**	No reading in hand; a function of G_1 and Q; 60 kW, 8250 W/K: 0.30
	Class	Operating characteristic at a stated truth

$$\boxed{\Phi\big((d_{\mathrm{crit}}(Q) - Q/G_1)/\sigma_d\big)}$$

Two things can be read off it. The load multiplies the margin, so a light test is a weak test. And on that form, at $G_1 = G^*$, the rate is exactly 0.05, which is where a level and a sampling rate coincide, and it is why they are so often run together. On the stated prior they part: 0.060, 0.065 and 0.078 at 60, 40 and 20 kW. Equation (16.6) gives 0.267 at 60 kW and 66 kW, against 0.297.

16.2.6 Power, assurance and three predictions

The drawing is not known to be right. *Assurance* averages the power over what is believed about the truth, with p_{o} that outer belief,

$$P(A(\tilde{E}_Q) = 1) = \int P(P(G \geq G^* \mid \tilde{E}_Q) \geq 0.95 \,|\, G = G_1)\, p_{\mathrm{o}}(G_1)\, \mathrm{d}G_1. \tag{16.7}$$

The outer truth has not gone away: the Expand shows it being integrated, by p_{o}, where the power holds it at one value. Two things must now be stated. The first is which evidence the inner uses. The rule may read the new test alone, $P(G \geq G^* \mid \tilde{E}_Q)$, or pool it with what is in hand, $P(G \geq G^* \mid E, \tilde{E}_Q)$. An engineer's default is to pool, and a contract's rule applied to one test reads that test alone. Both are formulations and they are different rules: for a 60 kW test

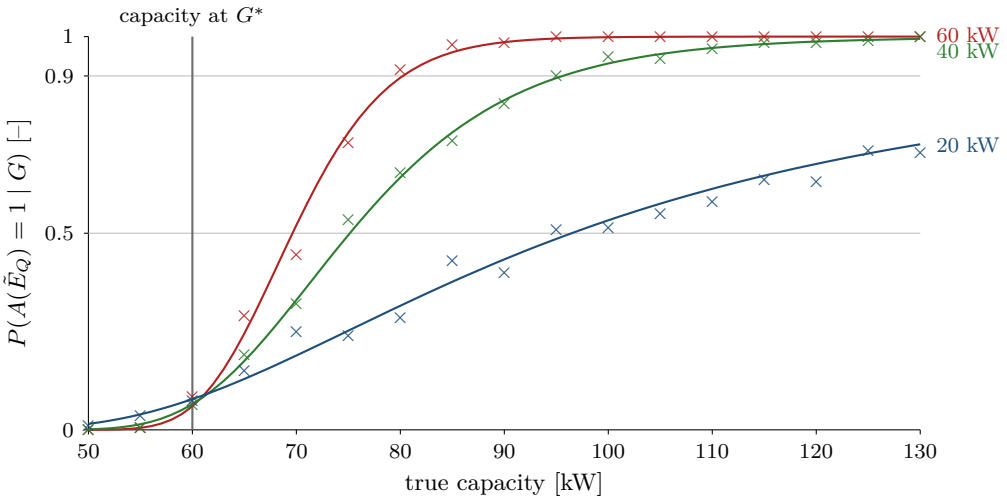


Figure 16.2: The operating characteristic of the rule at 0.95, against the true capacity, for three test loads. Curves, equation (16.5) at the exact critical difference. Crosses, the report's counts of 200 simulated tests. The gray vertical marks the capacity at G^*, and the light horizontals 0.5 and 0.9.

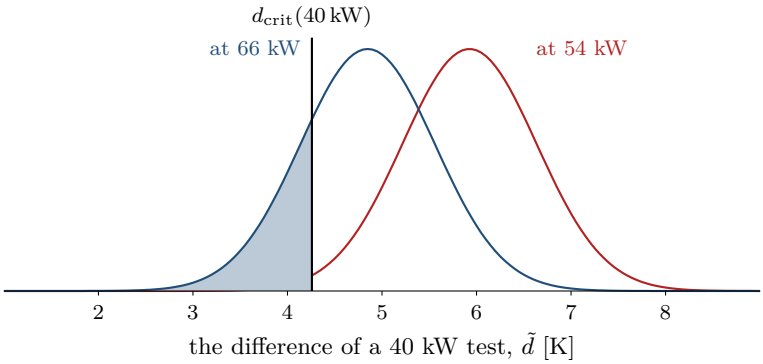


Figure 16.3: The difference of a 40 kW test at two truths, each density scaled to a peak of one, with the critical difference at 4.26 K. The shaded area below it, as a share of the area under its curve, is the probability that the test accepts.

on fresh sensors, given E, the rule alone has a critical difference of 6.894 K and accepts with probability 0.457, and the rule pooled has 7.025 K and 0.498. The second is the outer belief. It is separate from the rule's prior p_{r}, which is fixed when the rule is written. With nothing in hand, p_{o} may be the stated prior, and then the two coincide. Given E, $p_{\mathrm{o}}(G) = p(G \mid E)$, and the rule's prior has not moved. A sheet states both. In this lesson the rule reads the new test alone and p_{r} is the stated prior, unless a run says otherwise.

With that on the table, three questions about tests and plants stand apart, and they have names. Table 16.4 sets them side by side.

Assurance is the second with no evidence behind the bar. The next run writes it, with the

Table 16.4: Three questions about one plant and its tests.

Name	Expression	Random	Fixed	Asks
Posterior inference	$P(G \geq G^* \mid E)$	G	The readings taken	Is the claim supported by this test?
Posterior prediction	$P(A(\tilde{E}_Q) = 1 \mid E)$, the inner alone or pooled, as stated	The next readings; G, by $p(G \mid E)$	The readings taken	Given this test, will the next accept?
Sampling-distribution prediction	$P(A(\tilde{E}_Q) = 1 \mid G = G_1)$	The readings not yet taken	The truth	At this truth, how often would a test accept?

outer integral in its Expand line, and its seventh step holds the neighbors beside it.

Run, assurance of a 60 kW test

1	**Sentence**	Before any test, on the stated prior: how likely is a 60 kW test to accept?
	Intention	Predict
2	**Target**	An outcome of a test not yet run, with the truth unknown
	Declare	$\tilde{E}_{60}$; $\tilde{d}$ [K]; $p_\mathrm{r} = p_\mathrm{o}$ log-normal about G^*, a factor of 3; $d_\mathrm{crit}(60\,\mathrm{kW}) = 6.89\,\mathrm{K}$
3	**Known**	The prior; 60 kW; unit B out; settled; fresh sensors (assumption). No evidence in hand
	Random here	$\tilde{d}$ and G
4	**Event**	$A(\tilde{E}_{60}) = 1$
5	**Expression**	$P(A(\tilde{E}_{60}) = 1)$
	Expand	$\int P(P(G \geq G^* \mid \tilde{E}_{60}) \geq 0.95 \mid G = G_1)\, p_\mathrm{o}(G_1)\, \mathrm{d}G_1$
6	**Model**	The outer Φ at each G_1
	Reduce	$\int \Phi((6.89\,\mathrm{K} - 60\,\mathrm{kW}/G_1)/\sigma_d)\, p_\mathrm{o}(G_1)\, \mathrm{d}G_1$
7	**Check**	0.444; at the drawing alone, 0.297; on the near pair at equal weights (choice), 0.150
	Class	Assurance

$$\boxed{\int \Phi\big((d_\mathrm{crit}(60\,\mathrm{kW}) - 60\,\mathrm{kW}/G_1)/\sigma_d\big)\, p_\mathrm{o}(G_1)\, \mathrm{d}G_1}$$

16.2.7 Reading the crossings

A report's crossing sentences are of two kinds, and Table 16.5 holds them in separate columns. A design-case crossing is where the posterior at the expected reading first reaches 0.95 on the report's grid; the report calls that cell the acceptance margin. A power crossing is where equation (16.5) first reaches 0.9.

Table 16.5: The two kinds of crossing, on the report's grid of capacities. Design-case columns on the left, power columns on the right, the closed form at the exact critical difference.

Test load [kW]	Design case reaches 0.95, cell [kW]	Its value [–]	Count there [of 200]	Power reaches 0.9, by count, cell [kW]	Power reaches 0.9, closed form, cell [kW]
60	70	0.956	89	80 (183 of 200)	85 (0.896 at 80)
40	80	0.978	131	95 (180 of 200)	100 (0.895 at 95)
20	100	0.958	103	None (141 of 200 at 130)	None (0.726 at 130)

The count near half at a crossing has one cause. The design case puts the reading at Q/G_1. Where the design case is exactly 0.95, $Q/G_1 = d_\mathrm{crit}(Q)$ by equation (16.3), so the argument of equation (16.5) is zero and the power is $\Phi(0) = 1/2$. At 60 kW and 70 kW, $Q/G_1 = 6.857\,\mathrm{K}$

against $d_{\text{crit}} = 6.894\,\text{K}$: $\Phi(0.05)$, or 0.521, and the count is 89 of 200.[1] The rest of the column is the grid: the first cell past 0.95 is past it by some margin. At 40 kW, 75 kW has 0.9497 and 107 of 200, and 80 kW has 0.978 and 131.

16.2.8 The owner's paragraph, fully written

Lesson 1 stopped short of this statement. On the two-hall plant, whose two units are its two chillers, it can now be written whole.

> A cooling system has two chillers supported by a common electrical feeder. The owner wants to know whether commissioning data provide sufficient confidence to accept the redundancy claim, and whether testing at higher load would reduce the risk of accepting an inadequate system.

Run, the owner's paragraph

1	**Sentence**	(a) Do the data suffice to accept? Would a higher load reduce the risk, (b) by truth, (c) by outcome?
	Intention	(a) Decide; (b), (c) design an experiment, here as a comparison
2	**Target**	(a) A rule's value on evidence in hand; (b), (c) two rates across test loads
	Declare	Claim $G \geq G^*$, one unit out, feeder up; E; level 0.95 (assumed: "sufficient" hides it); $G_0 = 6750\,\text{W/K}$
3	**Known**	(a) E; (b), (c) the rule, and no reading of any new test; (c) the prior as a population
	Random here	(a) G; (b) $\tilde{d}$ only; (c) G and $\tilde{d}$
4	**Event**	(a) $G \geq G^*$; (b) accepted, at G_0; (c) inadequate, among those accepted
5	**Expression**	(a) $A(E)$; (b) $P(A(\tilde{E}_Q) = 1 \mid G = G_0)$; (c) $P(G < G^* \mid A(\tilde{E}_Q) = 1)$
	Expand	(a) $\mathbf{1}[P(G \geq G^* \mid E) \geq 0.95]$; (b) $P(P(G \geq G^* \mid \tilde{E}_Q) \geq 0.95 \mid G = G_0)$
6	**Model**	(a) The posterior; (b) the reading at G_0; (c) Bayes' rule on the population
	Reduce	(a) $\mathbf{1}[0.818 \geq 0.95]$; (b) $\Phi((d_{\text{crit}}(Q) - Q/G_0)/\sigma_d)$; (c) the ratio of Lesson 15
7	**Check**	(a) 0; (b) 0.033, 0.010, 0.003 at 20, 40, 60 kW; (c) 0.0097, 0.0032, 0.0018. The feeder: in no event
	Class	(a) Decision rule; (b) operating characteristic; (c) predictive value of a test's outcome

(a) $A(E)$; (b) $P(A(\tilde{E}_Q) = 1 \mid G = G_0)$; (c) $P(G < G^* \mid A(\tilde{E}_Q) = 1)$

The answer to the second question is yes, on either reading of "risk", and the paragraph did not separate them. Three more things are owed. The same load also raises the chance that a sound plant passes, from 0.141 to 0.297 at 66 kW. Which load to choose, at what cost of testing, is Lesson 21's. And no test load bears on the feeder: an outage of the feeder defeats both chillers, and the claim concerns the loss of one. A planned test is not evidence, so nothing in the run conditions on a test that has not happened.

16.2.9 A request that is not ready

The owner's representative asks: "Is the test good enough?" Good at which truth, at what rate, and against what cost of refusing a sound plant? Read as power at the drawing, it is 0.297 from a 60 kW test. Read as false acceptance at a stated truth, it is 0.003 at 54 kW. Read as the rule itself, it is a question about the rule's level, Lesson 17's. Not one of those is "good" until someone says what it is good for.

[1] The count is 2.1 binomial spreads below 0.521, and 2.7 below 0.540 on the count's own rule. The report scores its counts on a coarser grid of $\log G$ than its closed form, with the threshold's cell counted whole, and that rule's critical differences are 1.671, 4.282 and 6.929 K. On those the closed form crosses 0.9 at the count's cells, 0.9043 at 80 kW and 0.9006 at 95 kW, which is why the table's last two columns differ; scatter is not the cause. Cell by cell, the report's counts on that rule show what scatter gives, a sum of squared residuals of 36.9 over 41 cells, so the 89 is one low draw. The report's prose gives 142 of 200 at 130 kW from a 20 kW test, where its data give 142 at 125 kW and 141 at 130 kW.

Key idea

Parse, Translate, Expand, Reduce. The nested expression has two G's and two probabilities: the inner over G, for each reading, and the outer over readings, at a held truth. The inner gives a threshold on the reading, and the outer gives $\Phi((d_{\text{crit}}(Q) - Q/G_1)/\sigma_d)$. Where the design case is 0.95, that is one half.

16.3 The competing formulations

Each competing formulation loses track of what is random. The first writes the rule on evidence already in hand.

Run, evidence in hand where a test not yet run belongs

1	**Sentence**	At a true conductance G_1, how often would the trip test accept?
5	**Expression** ×	$P(A(E) = 1 \mid G = G_1)$, with the readings in hand
	Repair	$P(A(\tilde{E}_{40}) = 1 \mid G = G_1)$: over readings not yet taken
7	**Check**	With E in hand, $A(E)$ is 0 whatever G_1 is: the expression answers nothing

$$\boxed{P(A(\tilde{E}_{40}) = 1 \mid G = G_1)}$$

The second replaces the rule with the physics. It is a clean formula, and a different test.

Run, a rule replaced by a threshold from the physics

1	**Sentence**	At 66 kW, how often does a 40 kW test accept at the rule at 0.95?
6	**Model** ×	It accepts when $\tilde{d} \le Q/G^* = 5.33\,\text{K}$: $\Phi((5.33 - 4.85)/0.711) = 0.752$
	Repair	d_{crit} from the posterior rule, 4.26 K: 0.203
7	**Check**	At 54 kW the first accepts 0.202 and the rule 0.010: the threshold holds no prior and no floor

$$\boxed{\Phi\big((d_{\text{crit}}(40\,\text{kW}) - 40\,\text{kW}/G_1)/\sigma_d\big)}$$

The third puts the expected reading into the rule, which is the report's design case, and calls it how often.

Run, the design case as power

1	**Sentence**	A 70 kW opening, a 60 kW test: how often will it pass?
6	**Model** ×	$\tilde{d}$ set to Q/G_1: $\Phi((R^* - m_R(Q/G_1))/\sigma_R)$, about 0.95
	Repair	$\tilde{d}$ random about Q/G_1: $\Phi((d_{\text{crit}} - Q/G_1)/\sigma_d)$, about one half
7	**Check**	The wrong line uses the inner's Φ for the outer's, and no probability over readings remains

$$\boxed{\Phi\big((d_{\text{crit}}(60\,\text{kW}) - 60\,\text{kW}/G_1)/\sigma_d\big)}$$

The fourth keeps the form and drops a nuisance. It is Lesson 11's class with a size on power.

Run, the noise's spread for the floor

1	**Sentence**	A 70 kW opening, a 60 kW test: how often will it pass?
6	**Model** ×	Spread of $\tilde{d}$ from the averaged noise alone, 0.077 K: 0.682
	Repair	$\sigma_d = 0.711\,\text{K}$, with the offsets: 0.521
7	**Check**	Over a test with fresh sensors the offsets are random, and they are most of the spread

$$\boxed{\Phi\big((d_{\text{crit}}(60\,\text{kW}) - 60\,\text{kW}/G_1)/\sigma_d\big)}$$

The fifth states a truth nobody knows.

Run, power at one truth as how likely the test is to pass

1	**Sentence**	"How likely is a 60 kW test to pass?", on a drawing nobody has verified
3	**Known** ×	$G = 8250\,\mathrm{W/K}$, as a fact: 0.297
	Repair	State it as an assumption and say power; or carry the prior: assurance, 0.444
7	**Check**	Is the drawing evidence that G has this value? It is a prior (Lesson 12)

$$\boxed{\int P(A(\tilde{E}_{60}) = 1 \mid G = G_1)\, p_{\mathrm{o}}(G_1)\, \mathrm{d}G_1}$$

The sixth writes one G for two.

Run, the inner and the outer confounded

1	**Sentence**	At a stated G_1, how often does a test accept?
5	**Expression** ×	Expand as $P(P(G_1 \geq G^* \mid \tilde{E}_Q) \geq 0.95 \mid G = G_1)$
	Repair	The inner over G, unknown to the rule; the outer holds $G = G_1$
7	**Check**	With G_1 in the inner nothing is random in it, so the expression is 1 for every adequate G_1

$$\boxed{P(P(G \geq G^* \mid \tilde{E}_Q) \geq 0.95 \mid G = G_1)}$$

These are the forms of this lesson's failure class, *realized against prospective evidence*: a test's outcome treated as if its readings were in hand, or the reverse. Its five forms are evidence in hand where a test not yet run belongs; the expected reading plugged in; a rule replaced by a threshold that no longer holds the posterior; one truth used where the truth is unknown; and the inner and outer G confounded. It enters at step 3, 5 or 6, and it is caught at step 3 by asking what is left to be random. The fourth near-miss is Lesson 11's class meeting this one.

16.4 One system, many questions

Table 16.6 holds the two-hall plant still and names ten sentences about a test. It holds no values.

Table 16.6: Ten sentences about a test on the two-hall plant, each named.

Sentence	Which
Given the trip test, is the opening adequate?	Posterior inference
If it is what the drawing says, how often will a 60 kW test pass?	Sampling-distribution prediction
Given the trip test, will a 60 kW test pass?	Posterior prediction
What the 60 kW test would read if the drawing holds	A design-case reading, no probability
The rule's posterior at that reading	The design case
How often a test would pass, the opening not known	Assurance
How often an undersized opening would pass	False acceptance, by truth
Did the test pass?	A rule's value on E
At what capacity the test passes nine times in ten	A crossing of power
At what capacity the expected reading reaches 0.95	A crossing of the design case

16.5 Worked problems

Eight problems on the three systems follow, each with its run.

Problem 16.1. A timed-discharge rule keeps the string at a level of 0.90 if it gives its 2 h. The spread is 0.03 h (illustrative), and the string is taken as sound. Which readings keep it, and can a string at nameplate expect one?

Run, a critical reading above nameplate

1	**Sentence**	Rule at 0.90 on $T_{\mathrm{rt}} \geq 2$ h, a timed reading: which readings keep the string?
	Intention	Decide: which readings the rule accepts
2	**Target**	The set of readings the rule accepts
	Declare	$\tilde{d}$ the timed reading [h]; 0.03 h (illustrative); $T_{\mathrm{rt}} = 2.025$ h $\times\, \varphi$; φ uniform on 0.85 to 1, in the rule
3	**Known**	Sound string (assumption); the prior's bound at $\varphi = 1$
	Random here	φ, inside the rule; $\tilde{d}$ read as a variable
4	**Event**	$P(T_{\mathrm{rt}} \geq 2 \text{ h} \mid \tilde{d}) \geq 0.90$
5	**Expression**	$\{\tilde{d} : P(T_{\mathrm{rt}} \geq 2 \text{ h} \mid \tilde{d}) \geq 0.90\}$
6	**Model**	$\tilde{d} = 2.025$ h $\times\, \varphi + \varepsilon$; here the inner rises with the reading, to 0.90 at 2.085 h
	Reduce	$\{\tilde{d} \geq 2.085 \text{ h}\}$, a set above nameplate
7	**Check**	Nameplate centers on 2.025 h, two spreads under it: the bound is why Lesson 15's sound strings fail
	Class	Decision rule

$$\boxed{\{\tilde{d} : P(T_{\mathrm{rt}} \geq 2 \text{ h} \mid \tilde{d}) \geq 0.90\} = \{\tilde{d} \geq 2.085 \text{ h}\}}$$

Problem 16.2. The UPS is accepted if it picks up the load in n demand tests in a row. If it fails to pick up with probability p_{d} per demand, how often is it accepted?

Run, an exact tail

1	**Sentence**	Accept after n pick-ups in a row, a true failure share p_{d}: how often accepted?
	Intention	Predict
2	**Target**	An outcome of a test not yet run, at a fixed truth
	Declare	$\tilde{k}$ failures in n demands; p_{d} per demand; rule $\mathbf{1}[\tilde{k} = 0]$
3	**Known**	p_{d} (stated); demands independent (assumption); illustrative
	Random here	$\tilde{k}$
4	**Event**	$\tilde{k} = 0$
5	**Expression**	$P(A(\tilde{k}) = 1 \mid p_{\mathrm{d}})$
	Expand	$P(\tilde{k} = 0 \mid p_{\mathrm{d}})$
6	**Model**	Binomial
	Reduce	$(1 - p_{\mathrm{d}})^n$, a count's tail: no Φ
7	**Check**	$p_{\mathrm{d}} = 0.02$: 0.545 at $n = 30$, 0.133 at $n = 100$: more tests, and a good unit passes less often
	Class	Operating characteristic at a stated truth

$$\boxed{(1 - p_{\mathrm{d}})^n}$$

Problem 16.3. The rule of Problem 14.3, at 0.90 on a rate below 3.5, is applied to three years in service in place of one. How do its critical count and its rate of acceptance change, at true rates of 2.5 and 3.5?

Run, a longer record

1	**Sentence**	One year in service against three, the rule at 0.90: the critical count, and how often accepted?
	Intention	Compare
2	**Target**	An operating characteristic, across exposures
	Declare	$\tilde{k}$ a count not yet logged; λ a true rate [1/yr]; $t_\mathrm{s} = 1$, 3 yr; $k_\mathrm{c}(t_\mathrm{s})$ the critical count
3	**Known**	$\lambda_A = \lambda$ (stated); a Poisson count (assumption); the stated prior, inside the rule
	Random here	$\tilde{k}$
4	**Event**	$A(\tilde{k}) = 1$
5	**Expression**	$P(A(\tilde{k}) = 1 \mid \lambda_A = \lambda)$
	Expand	$P(P(\lambda_A < 3.5 \mid \tilde{k}) \geq 0.90 \mid \lambda_A = \lambda)$
6	**Model**	The inner falls with $\tilde{k}$: $k_\mathrm{c} = 1$ at one year, 7 at three (0.902, and 0.832 at 8)
	Reduce	$\sum_{k \leq k_\mathrm{c}(t_\mathrm{s})} e^{-\lambda t_\mathrm{s}} (\lambda t_\mathrm{s})^k / k!$, a count's tail
7	**Check**	At 2.5: 0.287 to 0.525. At 3.5, the claim's edge: 0.136 to 0.179. The count is discrete, so both rise
	Class	Operating characteristic at a stated truth

$$\boxed{\sum_{k \leq k_\mathrm{c}(t_\mathrm{s})} e^{-\lambda t_\mathrm{s}} (\lambda t_\mathrm{s})^k / k!}$$

Problem 16.4. The owner asks for "power 0.9 at a 10 % margin".

Run, a target with its parts missing

1	**Sentence**	"Power 0.9 at a 10 % margin."
	Intention	Not stated: a requirement on a test, or a claim about one
3	**Known**	Missing: margin over what; at which test load; under which rule
4	**Event**	(a) 66 kW, full-load test; (b) 66 kW, 40 kW test; (c) a margin on the conductance's estimate
5	**Expression**	(a), (b) $P(A(\tilde{E}_Q) = 1 \mid G = 8250\,\mathrm{W/K}) \geq 0.9$; (c) no such form
7	**Check**	(a) 0.297, and 0.9 takes about a 34 % margin (80.2 kW); (b) 0.203. The analyst asks
	Class	Not ready

$$\boxed{\text{(a), (b) } \Phi\big((d_\mathrm{crit}(Q) - Q/8250\,\mathrm{W/K})/\sigma_d\big) \text{ against } 0.9}$$

Problem 16.5. One test is run at 20 kW and then at 40 kW on the same sensors, with one posterior from both readings. How often does a true 66 kW opening pass?

Run, two readings, no single critical difference

1	**Sentence**	Same sensors, 20 and 40 kW, one posterior, true 66 kW: how often accepted?
	Intention	Predict
2	**Target**	An outcome of a two-reading test, at a fixed truth
	Declare	$\tilde{d}_1$, $\tilde{d}_2$ [K]; b the offsets' difference [K]; $\mathcal{C}$ the set the rule accepts; $G_1 = 8250\,\mathrm{W/K}$
3	**Known**	G_1 (stated); one sensor pair at both loads (assumption)
	Random here	$\tilde{d}_1$, $\tilde{d}_2$, tied by the shared b
4	**Event**	$(\tilde{d}_1, \tilde{d}_2) \in \mathcal{C}$
5	**Expression**	$P(P(G \geq G^* \mid \tilde{d}_1, \tilde{d}_2) \geq 0.95 \mid G = G_1)$
6	**Model**	$\tilde{d}_i = Q_i/G_1 + b + \varepsilon_i$; the inner has no one threshold, and its level set is a curve
	Reduce	Stops at the integral of a bivariate normal over $\mathcal{C}$
7	**Check**	About 0.75, against 0.141 and 0.203 singly: b cancels in $\tilde{d}_2 - \tilde{d}_1$. No single d_crit, so no one Φ
	Class	Operating characteristic at a stated truth

$$\boxed{P\big((\tilde{d}_1, \tilde{d}_2) \in \mathcal{C} \mid G = G_1\big)}$$

Problem 16.6. Given the 40 kW trip test, how likely is a 60 kW test on fresh sensors to accept, the rule applied to the new reading alone?

Run, the next test, given this one

1	**Sentence**	Given E: will a 60 kW test on fresh sensors accept, the rule on its reading alone?
	Intention	Predict
2	**Target**	An outcome of a test not yet run, with the truth known only through E
	Declare	$\tilde{E}_{60}$; $\tilde{d}$ [K]; $p(G \mid E)$; $d_{\text{crit}}(60\,\text{kW})$
3	**Known**	E; fresh sensors (assumption), so the two tests are independent given G
	Random here	$\tilde{d}$ and G; the readings of E fixed
4	**Event**	$A(\tilde{E}_{60}) = 1$
5	**Expression**	$P(A(\tilde{E}_{60}) = 1 \mid E)$
	Expand	$\int P(P(G \geq G^* \mid \tilde{E}_{60}) \geq 0.95 \mid G = G_1)\, p(G_1 \mid E)\, \mathrm{d}G_1$; the inner on $\tilde{E}_{60}$ alone
6	**Model**	The outer Φ at each G_1; the posterior from E as the outer belief; the rule's prior unmoved
	Reduce	$\int \Phi((d_{\text{crit}}(60\,\text{kW}) - 60\,\text{kW}/G_1)/\sigma_d)\, p(G_1 \mid E)\, \mathrm{d}G_1$
7	**Check**	0.457, against the rule's 0 on E and 0.818 for the claim; pooled with E it would be 0.498
	Class	Predictive probability

$$\boxed{\int \Phi\big((d_{\text{crit}}(60\,\text{kW}) - 60\,\text{kW}/G_1)/\sigma_d\big)\, p(G_1 \mid E)\, \mathrm{d}G_1}$$

Problem 16.7. A 70 kW opening is tested at 60 kW five times on the same pair of sensors, and accepted if any test accepts. How often is it accepted? The pair's offsets stay as they are, and only the noise is fresh.

Run, the same sensors

1	**Sentence**	70 kW, 60 kW test, five repeats on one sensor pair, any accept: how often accepted?
	Intention	Predict
2	**Target**	An outcome of repeated tests, at a fixed truth
	Declare	b the offsets' difference [K], spread $\sqrt{2} \times 0.5\,\text{K}$; noise alone 0.077 K
3	**Known**	$G_1 = 8750\,\text{W/K}$ (stated); b the same on all five (assumption)
	Random here	b, once; the noise, five times
4	**Event**	At least one acceptance in five
5	**Expression**	$P(\bigcup_i \{A(\tilde{E}_i) = 1\} \mid G = G_1)$
	Expand	Each $A(\tilde{E}_i) = 1$ is $\tilde{d}_i \leq d_{\text{crit}}$, by the rule's threshold
6	**Model**	Given b, independent repeats, each $\Phi((d_{\text{crit}} - Q/G_1 - b)/0.077)$
	Reduce	$\int \big(1 - (1 - \Phi(\cdot))^5\big)\, p(b)\, \mathrm{d}b$
7	**Check**	0.571, not the 0.975 of fresh sensors: a repeat repeats the offset, and most pairs are nearly decided
	Class	Operating characteristic at a stated truth

$$\boxed{\int \Big(1 - \big(1 - \Phi\big((d_{\text{crit}} - Q/G_1 - b)/0.077\big)\big)^5\Big)\, p(b)\, \mathrm{d}b}$$

Problem 16.8. Under the offset explanation, the PDU meter is accepted if its offset is within $\pm\kappa$ with probability at least 0.95, from one difference against the UPS output meter. Write how often that happens at a stated offset, for $\kappa = 5$ and $\kappa = 8$ kW (illustrative).

Run, a rule that may never accept, and a window

1	**Sentence**	Accept the meter if $P(\lvert b_{\text{pdu}}\rvert \le \kappa \mid \tilde{d}) \ge 0.95$: how often, at a stated offset?
	Intention	Predict
2	**Target**	An outcome of a reading not yet taken, at a fixed truth
	Declare	$\tilde{d}$ the difference [kW], spread 2.83 kW; prior 0 ± 45 kW; posterior spread 2.82 kW; κ (illustrative)
3	**Known**	b_{pdu} (stated); the offset explanation (assumption)
	Random here	$\tilde{d}$ only
4	**Event**	$A(\tilde{d}) = 1$
5	**Expression**	$P(P(\lvert b_{\text{pdu}}\rvert \le \kappa \mid \tilde{d}) \ge 0.95 \mid b_{\text{pdu}})$
6	**Model**	$\tilde{d}$ normal about b_{pdu}; the inner peaks at $\tilde{d} = 0$, at 0.923 for $\kappa = 5$, 0.995 for $\kappa = 8$
	Reduce	$\kappa = 5$: never. $\kappa = 8$: $\lvert\tilde{d}\rvert \le 3.37$ kW, so two Φ's
7	**Check**	At $\kappa = 8$: 0.766 at offset 0, 0.051 at offset 8 kW; at $\kappa = 5$ no reading can accept
	Class	Operating characteristic at a stated truth

$$\boxed{\kappa = 8\,\text{kW}\text{:}\ \Phi\big((3.37 - b_{\text{pdu}})/2.83\big) - \Phi\big((-3.37 - b_{\text{pdu}})/2.83\big)}$$

16.6 The exercises

The solutions are in Appendix A, and the facts of the three systems in Appendix D. On the two-hall plant, E is the 40 kW trip test, "the rule" is the one at 0.95 unless an exercise says otherwise, and a test means fresh sensors.

Exercise 16.1 (*formulate*). A report on the two-hall plant has a subsection titled "The probability of acceptance and the acceptance margin". Its prose follows, with its equations taken out; an ellipsis marks each omission.

Passage A.

> The question this arm answers is what rule turns one trip test into an acceptance of the claim, that the opening can carry one hall's load within the allowable rise, and how often that rule accepts a true claim and rejects a false one across the scatter of tests the plant could actually run.

Passage B.

> ... the readings ... are accepted as supporting it when the posterior puts enough probability there, ... where the bar reads "given": the probability is computed from the readings of the test, not before them, and α is the risk of a false acceptance the owner will carry. The power of the test at a truth ... is the probability, over the readings that truth generates, that the rule accepts ...

Passage C.

> ... the probability of acceptance falls as the observed hall difference ... rises, so the rule accepts exactly the tests whose difference lies below a critical value, one number per test load ... over the tests that could happen, the observed difference is normal about the design-case value $Q/G(C)$ with the floor ... as its spread ...

Passage D.

> The power ... lags the design case: at the 70 kW cell, where the full-load design case crosses 0.95, 89 of 200 realized tests accept; the power reaches 0.9 in the 80 kW cell from a full-load test, the 95 kW cell from a 40 kW test, and nowhere on the grid from a 20 kW test ...

From the prose alone:

(a) Write the claim as an event, with G^* from the hall's load and the allowable rise.

(b) Write the rule as a function of the evidence, at $\alpha = 0.05$.

(c) Parse the sentence on power in passage B into the random variable, the event, the condition and the rule. Translate it, and Expand it, saying which G is which.

(d) From passage C's first clause, write the critical difference's definition, and the inner in R where Lesson 10's checklist allows it. Run the checklist at 20 kW, where $d_{\text{crit}} = 1.66\,\text{K}$.

(e) From passage C's second clause, Reduce the power to a Φ form, stating what is assumed.

(f) The report writes the truth as a capacity C, with $G(C) = C/\Delta T_{\text{allow}}$. Write the power as a function of C, and evaluate it at $Q = 60\,\text{kW}$, $C = 70\,\text{kW}$ by hand, from $d_{\text{crit}} = 6.89\,\text{K}$.

(g) Sort passage D's crossings into those of the design case and those of the power, and write each kind's expression. Say what (f) has to do with "89 of 200".

(h) The report calls one quantity "the probability of acceptance". Which expression is it? Is it how often a test accepts? And which of Lesson 15's three conditionings is "the risk of a false acceptance"?

Exercise 16.2 (*formulate*). The note below is constructed for this exercise. It is illustrative, and no such report exists. It concerns the battery string, taken as sound, with its state of health uniform between 0.85 and 1.00.

> A. The string is signed off once a year on one half-depth timed discharge, from 0.80 to 0.50 state of charge at the critical load, which takes 1.0125 h at nameplate. The claim is that a full discharge would give at least 1.90 h. The question is what rule turns one such test into an acceptance of the claim, and how often that rule passes a string that meets the claim and one that does not, across the scatter of the tests that could be timed.
>
> B. A test is accepted when the probability of the claim, given its timed reading, is at least 0.95. The power at a stated state of health is the probability, over the readings such a string gives, that the rule accepts.
>
> C. The probability of the claim rises as the timed reading rises, so the rule accepts exactly the tests that read at or above a critical reading, 1.007 h. Over the tests that could happen, the reading is normal about its design value, the half-depth time at that state of health, with a spread of 0.03 h.
>
> D. At a state of health of 1.00, the first cell of a grid of 0.01 at which the design case reaches 0.95, 104 of 200 simulated tests accept. The power reaches 0.9 nowhere between 0.85 and 1.00. At 0.93, 3 of 200 accept.

From the note alone:

(a) Write the claim as an event on φ.

(b) Write the rule as a function of the reading.

(c) Parse the sentence on power in B. Translate it, and Expand it, saying which φ is which.

(d) Write the critical reading's definition, and run Lesson 10's checklist there in φ. Write the inner in the form the checklist allows.

(e) Reduce the power to a Φ form, stating what is assumed, and say what differs from the form on the opening, and why.

(f) Evaluate it by hand at $\varphi = 1.00$ and $\varphi = 0.93$, and hold the values against the counts.

(g) Sort D's sentences into those of the design case and those of the power, write each kind's expression, and say why the power cannot reach 0.9.

(h) Which of D's sentences is a false acceptance, and by which conditioning?

Exercise 16.3 (*formulate*). Run this one: a 60 kW test, the rule at 0.99, whose critical difference is 6.41 K, a true conductance of 10 000 W/K. How often does it accept?

Exercise 16.4 (*formulate*). The commissioning plan has a checkbox: "Report the test's power at the 66 kW drawing, at the report's rule, for a 60 kW test." Formulate what the box asks.

Exercise 16.5 (*formulate*). Write the assurance of the timed-discharge rule at 0.90 on a string taken as sound, over the fact sheet's prior on its state of health, and say what its size says about the rule.

Exercise 16.6 (*formulate*). Run this one: how often does a 40 kW test falsely accept a 6000 W/K opening?

Exercise 16.7 (*formulate*). Under the offset explanation, the PDU meter is accepted if $P(b_{\mathrm{pdu}} \le 5\,\mathrm{kW} \mid \tilde{d}) \ge 0.95$ (illustrative), from one difference against the UPS output meter. Reduce how often it is accepted at a stated offset to a Φ form.

Exercise 16.8 (*formulate*). The commissioning plan has a checkbox: "Confirm the test has adequate power." Formulate what the box asks.

Exercise 16.9 (*formulate*). The feeder's record is to be kept for three years in service, and a count of at most 11 accepts. Write the operating characteristic as a function of the feeder's rate, and evaluate its form at the stated rate.

Exercise 16.10 (*formulate*). The sensors are recalibrated, so that each offset's spread is 0.25 K in place of 0.5 K, and σ_d becomes about 0.36 K. Write the power at the 66 kW drawing from a 60 kW test, and say which parts of the expression the change reaches.

Exercise 16.11 (*formulate*). The test engineer asks: "Given what the 40 kW test showed, how likely is a second 40 kW test to pass, on its own reading alone, if the sensors stay where they are?" Formulate the question.

Exercise 16.12 (*reverse*). Write one plain sentence for $P\big(\Phi((R^{*} - m_R(\tilde{d}))/\sigma_R) \geq 0.95 \,\big|\, G = G_1\big)$, and say when the expression may be used.

Exercise 16.13 (*reverse*). Write one plain sentence for $\Phi((4.26\,\mathrm{K} - 40\,\mathrm{kW} \times R_1)/0.711\,\mathrm{K})$, where R_1 is a stated resistance, and check its units.

Exercise 16.14 (*reverse*). Write one plain sentence for $\sum_{k=0}^{2} e^{-5\lambda}(5\lambda)^k/k!$ read as a rule's operating characteristic, and name the rule.

Exercise 16.15 (*critique and repair*). A handover note says: "The plant passed a test of power 0.9, so the opening is adequate with probability 0.9." State what is wrong, and repair it.

Exercise 16.16 (*critique and repair*). A site procedure says: "Accept if the measured difference is at most 8 K, at whatever load the test can reach." State what is wrong, and repair it, with its size at a 20 kW test on a 30 kW opening.

Exercise 16.17 (*critique and repair*). A proposal says: "The test has power 0.99 at 130 kW, so it is a good test." State what is wrong, and repair it.

Exercise 16.18 (*critique and repair*). A plan says: "We simulated 60 kW tests at the posterior's mode in $\log G$ and report 0.34 as the chance that the next test passes." State the error, and repair it.

Exercise 16.19 (*critique and repair*). A plan says: "For assurance we centered the prior on the drawing with a 3 % spread, since the drawing is reliable, and the result, 0.30, shows the test is adequate." State the errors, and repair it.

Exercise 16.20 (*rapid*). For each expression, say what is random and what is held.

(a) $P(R > 0.15\ \mathrm{K/kW} \mid E)$.

(b) $P(A(\tilde{E}_{20}) = 0 \mid G = 9000\,\mathrm{W/K})$.

(c) $P(A(\tilde{E}_{20}) = 0 \mid E)$.

(d) $A(E)$.

(e) $P(\tilde{d} > 7\,\mathrm{K} \mid G = 8000\,\mathrm{W/K})$, at $Q = 60\,\mathrm{kW}$.

(f) $P(\tilde{k} \leq 3 \mid \lambda_{\mathrm{fd}} = 3.69)$.

(g) $P(\varphi \geq 0.95 \mid d)$.

(h) $P(A(\tilde{d}) = 1)$ on the string, with no test.

(i) $d_{\mathrm{crit}}(40\,\mathrm{kW})$.

(j) $P(G \geq 6000\,\mathrm{W/K} \mid d = 5\,\mathrm{K})$, at $Q = 40\,\mathrm{kW}$.

16.7 What to retain

- Parse, Translate, Expand, Reduce: random variable, event, condition, rule; $P(A(\tilde{E}_Q) = 1 \mid G = G_1)$; the nested form; $\Phi((d_{\text{crit}}(Q) - Q/G_1)/\sigma_d)$.
- Two G's and two probabilities. The inner is over G, for each reading. The outer is over the readings, at a held truth.
- The outer Φ is exact under the stated model. A Gaussian reading stands only in the inner, in R, where the checklist passes.
- Power is at a stated truth, and assurance averages over belief. Posterior inference, posterior prediction and sampling-distribution prediction are three questions.
- Where the design case is 0.95, the power is one half. Design-case crossings and power crossings are different columns.

Lesson 17

Decisions Under Uncertainty

17.1 The engineering question

The plant's operator has just read the acceptance report, and she has three questions, which to her are one. What is the probability that hall B overheats if its cooling unit is lost tonight? What set point makes that least likely? And what is the cheapest set point that keeps the risk under one percent? The engineer begins answering the third as a variation on the first. The three share one plant, one test and nearly the same words, and they are three different kinds of problem. The first number she reaches for, from last year's operating log, answers none of them.

Her manager has a question of another kind. The owner's contract says 95 %, and he asks where that came from and whether it is right for this plant. Nobody remembers. It was in the last contract.

This lesson puts numbers on the consequences of Lesson 14, chooses among actions by what they are expected to cost, and makes the change from observing a setting to choosing one explicit in the notation.

17.2 The formulation

The work sits at step 2, where the target is an action or a setting, and at step 3, where a setting goes behind the bar as $\mathrm{do}(\mathbf{u})$. On the two-hall plant, E is the 40 kW trip test. This lesson's numbers on losses, limits and electricity are illustrative, and each is labeled where it is used.

17.2.1 Expected loss

A *loss* $\mathcal{L}(a, \zeta)$ states what an action a costs in a state ζ. Table 17.1 puts numbers in Lesson 14's table, in stated units, and its numbers are illustrative. The action to take is the one with the least expected loss, weighted by what is known of the state:

$$a^{*} = \arg\min_{a \in \mathbb{A}} \mathbb{E}[\mathcal{L}(a, G) \mid E] = \arg\min_{a \in \mathbb{A}} \sum_{\zeta} \mathcal{L}(a, \zeta)\, P(\zeta \mid E). \tag{17.1}$$

This is the first use in the book of a posterior as an input and not as an answer.

Modify is chosen, at 10. Had the posterior been 0.95, accepting would cost 9.5, and at 0.99, 1.9. Modifying and shedding cost the same in either state on this table, so only the top row depends on what is believed about the opening.

17.2.2 What a level says

Between accepting and one alternative, equation (17.1) gives a rule. Let $\mathcal{L}_{\mathrm{FA}}$ be what a false acceptance costs beyond the alternative, and $\mathcal{L}_{\mathrm{FR}}$ what a false rejection costs beyond

Table 17.1: An illustrative loss table on the opening, in stated units, with the expected loss given E, on which $P(G \geq G^* \mid E) = 0.818$.

Action	$G \geq G^*$ [units]	$G < G^*$ [units]	Expected, given E [units]
Accept the claim	0	190	34.6
Modify the opening	10	10	10.0
Shed hall B's load on a unit's loss	25	25	25.0

accepting. With $p = P(G \geq G^* \mid e)$, accepting has the smaller expected loss when its excess in the inadequate state is outweighed by its saving in the adequate one, $(1 - p)\,\mathcal{L}_{\mathrm{FA}} \leq p\,\mathcal{L}_{\mathrm{FR}}$. Collecting p gives $\mathcal{L}_{\mathrm{FA}} \leq p\,(\mathcal{L}_{\mathrm{FA}} + \mathcal{L}_{\mathrm{FR}})$, so

$$P(G \geq G^* \mid e) \geq \frac{\mathcal{L}_{\mathrm{FA}}}{\mathcal{L}_{\mathrm{FA}} + \mathcal{L}_{\mathrm{FR}}}. \tag{17.2}$$

So a level is a statement about losses: 0.80 is a ratio $\mathcal{L}_{\mathrm{FA}}/\mathcal{L}_{\mathrm{FR}}$ of 4, 0.95 of 19, and 0.99 of 99. Table 17.1's ratio is 180 to 10, which gives 0.947. The owner's 0.95 was stated in Lesson 14 as his, and it can now be read: a false acceptance costs about nineteen times a false rejection. If that is not true of this plant, the level is wrong for it. This is the first of two ways to write *risk tolerance*, a tolerance on a loss. The other, on a probability, comes next.

17.2.3 The ladder of three questions

The operator's three sentences stand on hall B at 60 kW with unit B out. Hall A's set point is the control, so that $\vartheta_B(\mathbf{u}) = \mathbf{u} + 60\,\mathrm{kW} \times R$, with a limit $\vartheta_{\max}$ of 30°C, which is 22°C plus the stated 8 K, and a lowest set point of 18°C, both illustrative. That relation holds only with unit B out, and every box in this section carries E, unit B out, the stated load and the plant settled, whether or not it prints them. Setting $\mathbf{u}$ is assumed to leave R unchanged and to hold hall A at $\mathbf{u}$, which is what makes the relation survive $\mathrm{do}(\mathbf{u})$. Plant electricity is $P_{\mathrm{elec}}(\mathbf{u}) = 120\,\mathrm{kW}$, both halls at 60 kW, over a coefficient of performance of 4.0 at 22°C, rising by 0.15 per kelvin, also illustrative. Table 17.2 writes the three, each a different problem class.

Table 17.2: Three neighboring sentences, three problem classes, on hall B at 60 kW with unit B out, on the illustrative limit and electricity.

Sentence	Expression	Class	Answer
What is the probability of overheating?	$P(\vartheta_B(\mathbf{u}) > \vartheta_{\max} \mid E, \mathrm{do}(\mathbf{u} = 22°\mathrm{C}))$	Probability estimation	0.182
What settings minimize it?	$\arg\min_{\mathbf{u}} P(\vartheta_B(\mathbf{u}) > \vartheta_{\max} \mid E, \mathrm{do}(\mathbf{u}))$	Risk optimization	$\mathbf{u} = 18°\mathrm{C}$, at the bound; risk under 0.0001; 35.29 kW
What settings minimize electricity, with the risk under 1 %?	$\arg\min_{\mathbf{u}} \mathbb{E}[P_{\mathrm{elec}}(\mathbf{u}) \mid E]$ subject to $P(\vartheta_B(\mathbf{u}) > \vartheta_{\max} \mid E, \mathrm{do}(\mathbf{u})) \leq 0.01$	Chance-constrained optimization	$\mathbf{u} = 20.48°\mathrm{C}$; 31.81 kW

Figure 17.1 plots the risk and the electricity against the set point, with the three answers marked. The risk falls as the set point falls, and the electricity rises, and the sentences differ in which of those they care about. The second question has no cost in it, and here the risk is monotone, so the bound decides. That is a finding about the sentence and this relation: where a risk rises both ways, too little of a setting and too much, the optimum is interior.

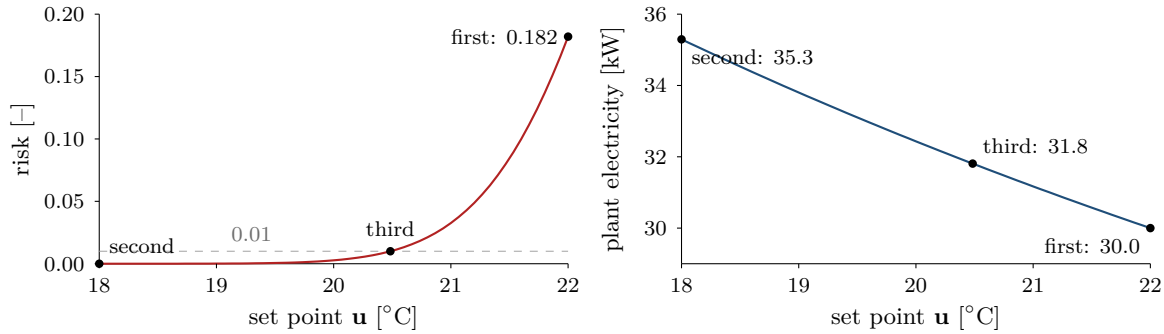


Figure 17.1: The ladder on hall B, on the illustrative limit and electricity. Left, the risk given E against the set point, with the 0.01 line dashed. Right, plant electricity against the set point. The dots mark the answers to the three sentences of Table 17.2.

The three runs that follow change lines 1, 2 and 5, with what those lines bring with them, and nothing else about the plant. The first is an estimate at a setting.

Run, the first rung, the probability of overheating

1	**Sentence**	Given E, hall B at 60 kW, unit B out, set point 22°C: is hall B over 30°C?
	Intention	Estimate
2	**Target**	A parameter's event, at a set control
	Declare	$\mathbf{u}$ hall A's set point [°C]; $\vartheta_B(\mathbf{u}) = \mathbf{u} + 60\,\text{kW} \times R$; $\vartheta_{\max} = 30°\text{C}$ (illustrative)
3	**Known**	E; do($\mathbf{u} = 22°\text{C}$); unit B out; 60 kW; settled; R unchanged by $\mathbf{u}$, hall A held at $\mathbf{u}$ (assumptions)
	Random here	R; $\mathbf{u}$ is set
4	**Event**	$\vartheta_B(\mathbf{u}) > \vartheta_{\max}$
5	**Expression**	$P(\vartheta_B(\mathbf{u}) > \vartheta_{\max} \mid E, \text{do}(\mathbf{u} = 22°\text{C}))$
6	**Model**	The opening's relation; the posterior on R
	Reduce	$P(R > 8\,\text{K}/60\,\text{kW} \mid E) = P(G < G^* \mid E)$
7	**Check**	0.182; the claim's complement, with the limit written as a temperature
	Class	Probability estimation; here a posterior exceedance

$$\boxed{P(\vartheta_B(\mathbf{u}) > \vartheta_{\max} \mid E, \text{do}(\mathbf{u} = 22°\text{C}))}$$

The second asks for a setting and states no cost.

Run, the second rung, the setting that minimizes the risk

1	**Sentence**	Given E, hall B at 60 kW, unit B out: which set point minimizes the risk over 30°C?
	Intention	Optimize
2	**Target**	A setting
	Declare	As above; $\mathbf{u} \geq 18°\text{C}$ (illustrative)
3	**Known**	E; unit B out; 60 kW; settled; $\mathbf{u}$ chosen, so do($\mathbf{u}$) for each candidate
	Random here	R; $\mathbf{u}$ is the decision
4	**Event**	$\vartheta_B(\mathbf{u}) > \vartheta_{\max}$; objective its probability
5	**Expression**	$\arg\min_{\mathbf{u}} P(\vartheta_B(\mathbf{u}) > \vartheta_{\max} \mid E, \text{do}(\mathbf{u}))$
6	**Model**	As above
	Reduce	The risk rises with $\mathbf{u}$, so the lowest allowed $\mathbf{u}$
7	**Check**	18°C, under 0.0001, 35.29 kW: the bound decides, since the sentence prices nothing
	Class	Risk optimization

$$\boxed{\arg\min_{\mathbf{u}} P(\vartheta_B(\mathbf{u}) > \vartheta_{\max} \mid E, \text{do}(\mathbf{u}))}$$

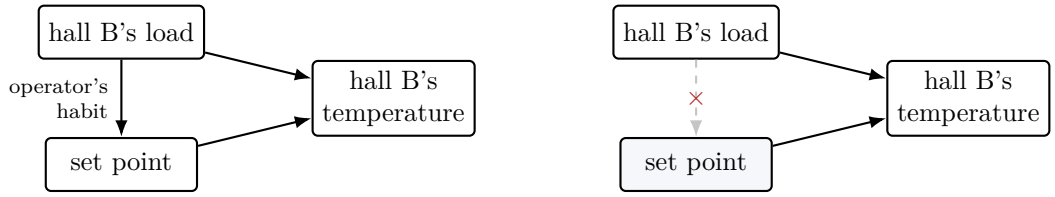
The third has the operator's objective and a tolerance on a probability, which is a *chance*

constraint.

Run, the third rung, a chance-constrained set point		
1	**Sentence**	Given E, hall B at 60 kW, unit B out: least electricity, with the risk over 30°C at most 1 %?
	Intention	Optimize
2	**Target**	A setting, under a constraint on a probability
	Declare	As above; $P_{\text{elec}}(\mathbf{u}) = 120\,\text{kW}/(4.0 + 0.15(\mathbf{u} - 22°\text{C}))$ (illustrative); tolerance 0.01
3	**Known**	E; unit B out; 60 kW; settled; $\mathbf{u}$ chosen; P_{elec} taken as known (assumption)
	Random here	R; nothing in P_{elec} here, so the expectation is trivial, and the sheet says so
4	**Event**	$\vartheta_B(\mathbf{u}) > \vartheta_{\max}$, in the constraint; objective P_{elec}
5	**Expression**	$\arg\min_{\mathbf{u}} \mathbb{E}[P_{\text{elec}}(\mathbf{u}) \mid E]$ s.t. $P(\vartheta_B(\mathbf{u}) > \vartheta_{\max} \mid E, \text{do}(\mathbf{u})) \leq 0.01$
6	**Model**	P_{elec} falls with $\mathbf{u}$ and the risk rises, so the constraint binds; Principle 10.1 passes on E
	Reduce	$\Phi((R_{\mathbf{u}} - m_R)/\sigma_R) = 0.99$, Gaussian reading in R, with $R_{\mathbf{u}} = (30°\text{C} - \mathbf{u})/60\,\text{kW}$
7	**Check**	Exact $\mathbf{u} = 20.48°$C, the reading 20.49°C; 31.81 kW, 3.49 kW under the second rung
	Class	Chance-constrained optimization
		$\arg\min_{\mathbf{u}} \mathbb{E}[P_{\text{elec}}(\mathbf{u}) \mid E]$ s.t. $P(\vartheta_B(\mathbf{u}) > \vartheta_{\max} \mid E, \text{do}(\mathbf{u})) \leq 0.01$

17.2.4 Setting against observing

A set point in a log is not a setting chosen at random. Suppose an operator raises it only on light nights: this is an illustrative log, constructed for this lesson, in which 24°C appears only with hall B's load between 30 and 50 kW. Figure 17.2 draws the difference. It is not Lesson 13's factor graph: its arrows mean "sets". In the log, hall B's load sets the operator's choice and also sets the temperature, so the set point carries news about the load. Under do($\mathbf{u}$), the arrow into the set point is cut, and the setting carries no news of anything. By Lesson 12's rule the log is this plant's own record, so its temperature readings could be evidence on R. They are not used here: only its pairing of set points with loads is, as a stated distribution of the load.



in the log: the setting is found — under do($\mathbf{u}$): the setting is set

Figure 17.2: A set point found and a set point set. Left, in the illustrative log, the operator's habit ties the set point to the load, so seeing the set point says something about the load. Right, under do($\mathbf{u}$), that arrow is cut, and the load is whatever it is.

With the load independent of R and E (assumption), a setting's risk is taken over the plant's own marginal load $p(Q_B)$, and not over the loads that came with it in the log:

$$P(\vartheta_B > \vartheta_{\max} \mid E, \text{do}(\mathbf{u})) = \int P(\mathbf{u} + Q'R > \vartheta_{\max} \mid E)\, p(Q')\, \text{d}Q'. \tag{17.3}$$

Table 17.3 sizes it, on the 40 kW test's posterior. Every row is a value of the model on E, and none is a frequency read from the log. With unit B out, the physics is the same in all of them. The first two rows agree: given the load, found and set agree, provided the load is the only arrow into the set point. What moves the number is the load. The log has nothing to say about 24°C at 60 kW. In general an observed value and a set one agree when nothing that moves the setting also moves the outcome, once what is stated has been conditioned on, or when the setting was assigned at random, and a log alone cannot say which holds.

Table 17.3: The 24°C set point with unit B out, found in the illustrative log and set, over three illustrative distributions of hall B's load, on the illustrative 30°C limit.

Expression	Hall B's load	Value [–]
$P(\vartheta_B > \vartheta_{\max} \mid E, \mathbf{u} = 24°\text{C}$ found in the log)	Uniform on 30 to 50 kW	0.103
$P(\vartheta_B > \vartheta_{\max} \mid E, \text{do}(\mathbf{u} = 24°\text{C}))$, over the log's loads	Uniform on 30 to 50 kW	0.103
$P(\vartheta_B > \vartheta_{\max} \mid E, \text{do}(\mathbf{u} = 24°\text{C}))$, over the plant's marginal	Uniform on 30 to 60 kW	0.289
$P(\vartheta_B > \vartheta_{\max} \mid E, \text{do}(\mathbf{u} = 24°\text{C}), Q_B = 60\,\text{kW})$	60 kW	0.833

The rule from Lesson 3 stands: a decision goes behind the bar as a setting, written do(·), and a planned action is not an observation. "We will test at 60 kW" adds no evidence until a reading exists, and "we accepted" adds none about G. An action whose worth depends on what it will reveal is priced differently. After the reading the best action on it would be taken, so its expected loss is the average, over the readings not yet taken, of the least expected loss on each: a minimum inside an expectation. That expression is written once in the worked problems below, and Lesson 20 takes it up.

17.2.5 Accepting or modifying

The first full sheet takes Table 17.1 to its action. Table 17.4 is the sheet.

17.2.6 Bypassing the UPS

The same distinction arises on the network, on a setting that is either on or off. To repair the UPS, the load is put on the static bypass for 12 h, chosen for a quiet Sunday. Its loss-of-load exposure is the supply's, at the illustrative rates of 2.0 interruptions and 0.5 spurious trips per year.

Run, the UPS bypassed by choice

1	**Sentence**	The load put on the bypass for a 12 h repair: is it lost at least once in those hours?
	Intention	Predict
2	**Target**	An event over a horizon, under a set state
	Declare	do(bypass); $\tau = 12$ h; rates 2.0, 0.5 per year (illustrative); no UPS between the supply and the load
3	**Known**	do(bypass), chosen; supply up at the start; Poisson losses (assumption)
	Random here	The supply's losses in the horizon
4	**Event**	At least one loss of supply in $[0, \tau]$
5	**Expression**	$P(N_{\text{ut}}(\tau) + \text{trips} \geq 1 \mid \text{do(bypass)})$
6	**Model**	Two Poisson streams, independent of the setting (assumption)
	Reduce	$1 - e^{-2.5 \times 12/8760}$
7	**Check**	0.0034. Found on the bypass in a log, the UPS got there by failing, and the plant's state comes with it
	Class	Predictive probability

$$\boxed{1 - e^{-(2.0+0.5)\tau},\ \tau \text{ in years}}$$

17.2.7 A request that is not ready

The director says: "Run it as efficiently as possible." Efficient in what, under which limit, at what risk, and over which settings? Read as least electricity with no limit, the answer is the highest set point the units allow, and the hall overheats. Read as the third rung, the tolerance is missing, and it is the owner's to give. Read as cost per unit of IT load served, the objective

Table 17.4: The formulation sheet filled in for the choice among actions on the opening, on the illustrative loss table.

Field	Entry
Intention	Decide
1. Sentence	Given the 40 kW trip test and the loss table: accept the claim, modify the opening, or shed?
2. Target	An action
3. Known	E; the loss table (illustrative); the stated model. The actions are chosen, not drawn.
4. Event	The two states, weighting each action; no horizon, since the table's units carry it
5. Expression	Equation (17.1)
6. Model	The posterior of Lessons 10 to 12
7. Read back	The action with the least expected loss on this table. It depends on the table as much as on the posterior.
Problem class	Expected-loss minimization

Its run, with the expectation expanded:

Run, accept, modify or shed

1	**Sentence**	Given E and the loss table: accept the claim, modify the opening, or shed?
	Intention	Decide
2	**Target**	An action
	Declare	$\mathbb{A}$ accept, modify, shed; $\mathcal{L}(a, \zeta)$ [units] (illustrative); ζ the states $G \geq G^*$, $G < G^*$
3	**Known**	E; the loss table; the stated model
	Random here	G only; the actions are chosen
4	**Event**	The states, weighting each action
5	**Expression**	$\arg\min_{a \in \mathbb{A}} \mathbb{E}[\mathcal{L}(a, G) \mid E]$
	Expand	$\mathcal{L}(a, G \geq G^*)\, P(G \geq G^* \mid E) + \mathcal{L}(a, G < G^*)\, P(G < G^* \mid E)$
6	**Model**	The posterior
	Reduce	Accept 190×0.182; modify 10; shed 25
7	**Check**	Modify, at 10; the rule at 0.95 agrees here, and need not on another table
	Class	Expected-loss minimization

$$\boxed{\arg\min_{a \in \mathbb{A}} \mathbb{E}[\mathcal{L}(a, G) \mid E]}$$

is another one altogether, with the load as a setting. Three readings, three problem classes, and the analyst asks.

Key idea

A decision is an argmin, and the posterior is one of its two inputs. A level states a ratio of losses. Three neighboring sentences are three problem classes: probability estimation, risk optimization, chance-constrained optimization. A setting is written do($\mathbf{u}$), a value in a log is not one, and a plan or a decision is not evidence.

17.3 The competing formulations

Each competing formulation answers the wrong rung, or treats a setting as a reading. The first stops one rung too low.

Run, a minimum risk with no cost

1	**Sentence**	"Least electricity, with the risk under 1 %."
5	**Expression** ×	$\arg\min_{\mathbf{u}} P(\vartheta_B(\mathbf{u}) > \vartheta_{\max} \mid E, \mathrm{do}(\mathbf{u}))$
	Repair	The third rung: electricity as the objective, the risk in a constraint
7	**Check**	35.29 kW against 31.81 kW, on the illustrative inputs, for a risk nobody asked for

$$\boxed{\arg\min_{\mathbf{u}} \mathbb{E}[P_{\mathrm{elec}}(\mathbf{u}) \mid E] \quad \text{s.t.} \quad P(\vartheta_B(\mathbf{u}) > \vartheta_{\max} \mid E, \mathrm{do}(\mathbf{u})) \le 0.01}$$

The second is Lesson 7's plug-in, on a loss.

Run, a plug-in expected loss

1	**Sentence**	Given E and the illustrative loss table: accept or modify?
6	**Model** ×	G at the posterior's mode in $\log G$, 8347 W/K, adequate: accept costs 0
	Repair	The posterior as weights: accept costs 34.6
7	**Check**	The point value hides the state with the large loss, which is what the decision turns on

$$\boxed{\arg\min_{a \in \mathbb{A}} \mathbb{E}[\mathcal{L}(a, G) \mid E]}$$

The third is the log's number used for the choice.

Run, found for set

1	**Sentence**	"The illustrative log shows 24°C is fine: the model gives 0.10 there." Can it be set on a full-load night?
3	**Known** ×	$\mathbf{u} = 24$°C found, with the load that came with it, for $\mathbf{u}$ chosen at 60 kW
	Repair	$\mathrm{do}(\mathbf{u} = 24\text{°C})$, $Q_B = 60$ kW: 0.833
7	**Check**	Who or what set this value? In the log, an operator who was looking at the load

$$\boxed{P(\vartheta_B > \vartheta_{\max} \mid E, \mathrm{do}(\mathbf{u} = 24\text{°C}), Q_B = 60\,\mathrm{kW})}$$

The fourth writes a constraint with no probability in it. It feels conservative, and it sets the risk at even odds.

Run, a constraint at a point value

1	**Sentence**	Least electricity, with the risk over the illustrative 30°C limit at most 1 %.
5	**Expression** ×	Constraint $\mathbf{u} + 60\,\mathrm{kW} \times m_R \le 30$°C: $\mathbf{u} \le 22.97$°C
	Repair	$P(\vartheta_B(\mathbf{u}) > \vartheta_{\max} \mid E, \mathrm{do}(\mathbf{u})) \le 0.01$: $\mathbf{u} \le 20.48$°C
7	**Check**	At 22.97°C the risk is 0.500: the constraint holds at the center, by construction

$$\boxed{P(\vartheta_B(\mathbf{u}) > \vartheta_{\max} \mid E, \mathrm{do}(\mathbf{u})) \le 0.01}$$

The fifth reads a plan as evidence.

Run, a plan read as evidence

1	**Sentence**	"We have scheduled a 60 kW test, so the probability the opening is adequate has gone up."
3	**Known** ×	The plan, behind the bar of $P(G \ge G^* \mid \cdot)$
	Repair	$P(G \ge G^* \mid E)$ stays 0.818 until a reading exists; a plan is $\tilde{E}_{60}$, over which a probability can be taken
7	**Check**	What would the plan have to be, to move G? Nothing about a schedule can

$$\boxed{P(G \ge G^* \mid E)}$$

The sixth reads a decision as evidence.

Run, a decision read as evidence

1	**Sentence**	"We accepted the plant, so we now take $G \geq G^*$ in the service model."
3	**Known** ×	The acceptance, behind the bar as if it observed G
	Repair	$P(S(t) = 1 \mid E)$, with G still uncertain; the signature changes no reading
7	**Check**	0.818 stays 0.818 after a signature, and the service's number should keep the mixture

$$\boxed{P(S(t) = 1 \mid E)}$$

The seventh has a tolerance with no event to attach it to.

Run, a level with no horizon

1	**Sentence**	"Keep the risk under 1 %.", with no statement of per what
4	**Event** ×	Overheating, with no conditions and no horizon
	Repair	Stated: at the moment unit B is lost, hall B at 60 kW; or over a year, which is a different event
7	**Check**	Per loss of a unit, and per year, differ by the rate of losses

$$\boxed{P(\vartheta_B(\mathbf{u}) > \vartheta_{\max} \mid E, \mathrm{do}(\mathbf{u}), \text{unit B out}, Q_B = 60\,\mathrm{kW}) \leq 0.01}$$

The third, fifth and sixth are forms of this lesson's failure class, *intervention read as observation*: a setting or an action treated as a reading, or a value in a record treated as what follows from choosing it. Its four forms are a found setting for a set one; a planned action read as evidence; a decision updating belief; and a relation assumed to survive do(·) when only the log has been seen, as in the eighth problem below. It enters at step 3 or step 5, usually the first, and it is caught by asking who or what set this value. The first and seventh are errors of rung and of event, and the second and fourth are Lesson 7's plug-in, on a loss and in a constraint.

17.4 One system, many questions

Table 17.5 holds the network still and gives ten sentences, each with its target and whether a do(·) is needed. It holds no values.

Table 17.5: Ten sentences on the network, with the target of each and whether a setting is behind its bar.

Sentence	Target	A do(·)
The probability that the load is lost in tonight's storm	A probability	No
The same, with the UPS bypassed for repair	A probability	Yes
The same, when the log shows it bypassed	A probability	No, and another number
Replace the UPS, or recalibrate the meter	An action, by expected loss	The action is chosen
Lowest-cost dispatch with the risk under 1 %	A setting, chance constraint	Yes
The bypass schedule that minimizes the chance of loss	A setting, risk optimization	Yes
"A bypass is planned on Sunday": the effect on the efficiency's estimate	None: a plan is not evidence	Not on that estimate
A flag raised, and the crew then sent	A rule's value, then a setting	Yes, for the crew
The efficiency if the load falls by itself	A prediction	No
The efficiency if the load is shed by choice	A prediction	Yes

17.5 Worked problems

Eight problems on the three systems follow, each with its run.

Problem 17.1. A fourth row is added to the loss table: retest at 60 kW, at a cost $\mathcal{L}_{\mathrm{t}}$ (illustrative), on fresh sensors, then choose among the others on what the two tests show together. Write its expected loss.

Run, a retest as an action

1	**Sentence**	Retest at 60 kW for $\mathcal{L}_{\mathrm{t}}$, then choose: what is its expected loss, given E?
	Intention	Decide
2	**Target**	One action's expected loss, where the action reveals something
	Declare	$\tilde{E}_{60}$; $\mathcal{L}_{\mathrm{t}}$ [units] (illustrative); $\mathbb{A}$ accept, modify, shed; $\mathcal{L}(a, \zeta)$ (illustrative)
3	**Known**	E; the loss table; fresh sensors (assumption); the retest is a setting, its readings not evidence yet
	Random here	$\tilde{E}_{60}$ and G
4	**Event**	The states, on each possible reading
5	**Expression**	$\mathcal{L}_{\mathrm{t}} + \mathbb{E}[\min_a \mathbb{E}[\mathcal{L}(a, G) \mid E, \tilde{E}_{60}] \mid E]$
6	**Model**	The reading's relation; the posterior, for each reading; left there
7	**Check**	A min inside an expectation; compared with 10 for modify; Lesson 20 takes it up
	Class	Expected-loss minimization

$$\boxed{\mathcal{L}_{\mathrm{t}} + \mathbb{E}\big[\min_{a \in \mathbb{A}} \mathbb{E}[\mathcal{L}(a, G) \mid E, \tilde{E}_{60}] \,\big|\, E\big]}$$

Problem 17.2. On the string, a share ξ of the critical load can be shed from the start of an outage. On a sound string, what is the smallest share that carries the next outage to its end with probability at least 0.95?

Run, a shed share as a control

1	**Sentence**	A sound string, next outage: the smallest shed share with the outage carried at 0.95?
	Intention	Optimize
2	**Target**	A setting, under a chance constraint
	Declare	ξ the share shed [–]; $T_{\mathrm{rt}}(\xi) = 2.025$ h $\times\, \varphi/(1-\xi)$; D_{out} [h]; level 0.95 (illustrative)
3	**Known**	$\mathrm{do}(\xi)$, chosen; sound string (assumption); the fact sheet's φ and D_{out}
	Random here	φ and D_{out}; ξ is set
4	**Event**	$D_{\mathrm{out}} \le T_{\mathrm{rt}}(\xi)$
5	**Expression**	$\min \xi$ s.t. $P(D_{\mathrm{out}} \le T_{\mathrm{rt}}(\xi) \mid \mathrm{do}(\xi)) \ge 0.95$
6	**Model**	The string's relation, with a smaller load lengthening it; exponential lengths
	Reduce	$\int_{0.85}^{1} (1 - e^{-2.025\varphi/(1-\xi)})\, \mathrm{d}\varphi / 0.15 \ge 0.95$, which binds
7	**Check**	$\xi = 0.377$; 0.846, 0.917, 0.976 at $\xi = 0$, 0.25, 0.5
	Class	Chance-constrained optimization

$$\boxed{\min \xi \;\; \text{s.t.} \;\; P(D_{\mathrm{out}} \le T_{\mathrm{rt}}(\xi) \mid \mathrm{do}(\xi)) \ge 0.95}$$

Problem 17.3. Lesson 14 sent a technician when an offset was more likely than not. If a visit costs 1 and an offset left in place costs 9 (illustrative), what level should the rule have had?

Run, a level from a loss ratio

1	**Sentence**	A visit costs 1, an offset left in place 9: at what level should the technician be sent?
	Intention	Decide
2	**Target**	A rule's level, from its losses
	Declare	$\mathbb{A}$ send, do not send; $\mathcal{L}(\text{send}, \cdot) = 1$; $\mathcal{L}(\text{do not send}, H_{\text{off}}) = 9$ (illustrative)
3	**Known**	The loss table; the rule's y as its argument
	Random here	The explanation, given y
4	**Event**	H_{off}
5	**Expression**	Send when $\mathbb{E}[\mathcal{L}(\text{send}, \cdot) \mid y] \leq \mathbb{E}[\mathcal{L}(\text{do not send}, \cdot) \mid y]$
6	**Model**	Equation (17.2), with the roles of the two actions as stated
	Reduce	$1 \leq 9\, P(H_{\text{off}} \mid y)$, so $P(H_{\text{off}} \mid y) \geq 1/9$
7	**Check**	0.111, not 0.5. A level of 0.5 says a left offset costs two visits, which nobody chose
	Class	Decision rule

$$\boxed{A(y) = \mathbf{1}[P(H_{\text{off}} \mid y) \geq 1/9]}$$

Problem 17.4. The contract says: "Keep overheating risk acceptable."

Run, a tolerance with no number

1	**Sentence**	"Keep overheating risk acceptable."
	Intention	Not stated
3	**Known**	Missing: which hall; at what load; per what; at what tolerance; over which settings
4	**Event**	(a) on a unit's loss at full load; (b) over a year; (c) a loss-based tolerance
5	**Expression**	(a), (b) a chance constraint; (c) an expected loss
7	**Check**	(a) is the third rung, once a number is given; (b) needs the rates; (c) needs the losses
	Class	Not ready

$$\boxed{\text{(a), (b) } P(\cdot) \leq \text{a stated tolerance}; \quad \text{(c) } \arg\min_{\mathbf{u}} \mathbb{E}[\mathcal{L}(\mathbf{u}, R) \mid E]}$$

Problem 17.5. Two controls are available on hall B's loss: hall A's set point, and a share ξ of hall B's load shed. Write the settings that keep the risk at 1 %.

Run, two controls

1	**Sentence**	Set point and shed share together: which settings keep the risk over 30°C at 1 %?
	Intention	Optimize
2	**Target**	A set of settings
	Declare	$\mathbf{u} = (\mathbf{u}_1, \xi)$; $\vartheta_B(\mathbf{u}) = \mathbf{u}_1 + (1-\xi)\, 60\,\text{kW} \times R$; $R_{0.99} = 0.1586$ K/kW; limit (illustrative)
3	**Known**	E; unit B out; settled; do($\mathbf{u}$)
	Random here	R
4	**Event**	$\vartheta_B(\mathbf{u}) > \vartheta_{\max}$
5	**Expression**	$\{\mathbf{u} : P(\vartheta_B(\mathbf{u}) > \vartheta_{\max} \mid E, \text{do}(\mathbf{u})) = 0.01\}$
6	**Model**	The risk rises with each of the two, so its level set is a curve
	Reduce	$\mathbf{u}_1 + (1-\xi)\, 60\,\text{kW} \times R_{0.99} = 30°\text{C}$, with $R_{0.99}$ the posterior's 0.99 point
7	**Check**	One constraint, two controls: no single answer until an objective is stated
	Class	Chance-constrained optimization, its boundary; the objective is still owed

$$\boxed{\mathbf{u}_1 + (1-\xi)\, 60\,\text{kW} \times R_{0.99} = \vartheta_{\max}}$$

Problem 17.6. The plant adopts a schedule: hall A's set point is chosen from hall B's load at each hour, $\mathbf{u} = \pi(Q_B)$. Write the risk under that policy.

Run, a policy as a setting

1	**Sentence**	A schedule $\mathbf{u} = \pi(Q_B)$: what is the risk over 30°C on a unit's loss?
	Intention	Predict
2	**Target**	An event, under a set rule for the setting
	Declare	$\pi(\cdot)$ the schedule; Q_B [kW], with Q' its values and a stated distribution; $\vartheta_B = \pi(Q_B) + Q_B R$
3	**Known**	E; $\mathrm{do}(\mathbf{u} = \pi(Q_B))$; unit B out; settled
	Random here	R and Q_B; the arrow from load to setting is now designed
4	**Event**	$\vartheta_B > \vartheta_{\max}$
5	**Expression**	$P(\vartheta_B > \vartheta_{\max} \mid E, \mathrm{do}(\mathbf{u} = \pi(Q_B)))$
6	**Model**	Total probability over Q_B; the load independent of R and E (assumption)
	Reduce	$\int P(\pi(Q') + Q'R > \vartheta_{\max} \mid E)\, p(Q')\, \mathrm{d}Q'$
7	**Check**	A policy is a setting. Same arrow as the log's, and here it is known and chosen
	Class	Probability estimation

$$\boxed{\int P(\pi(Q') + Q'R > \vartheta_{\max} \mid E)\, p(Q')\, \mathrm{d}Q'}$$

Problem 17.7. An overheating is priced at 10 units of electricity (illustrative) and added to the objective, in place of the 1 % constraint, with $18°\mathrm{C} \le \mathbf{u} \le 22°\mathrm{C}$. Which set point follows, and what does it do to the risk?

Run, two ways of writing a tolerance

1	**Sentence**	Least electricity plus 10 per overheating, against the 1 % constraint: which set point?
	Intention	Optimize
2	**Target**	A setting, by expected loss
	Declare	$\mathcal{L}(\mathbf{u}, R) = P_{\mathrm{elec}}(\mathbf{u}) + 10 \times \mathbf{1}[\vartheta_B(\mathbf{u}) > \vartheta_{\max}]$ [kW equivalents] (illustrative); $18°\mathrm{C} \le \mathbf{u} \le 22°\mathrm{C}$
3	**Known**	E; unit B out; 60 kW; settled; $\mathrm{do}(\mathbf{u})$
	Random here	R
4	**Event**	$\vartheta_B(\mathbf{u}) > \vartheta_{\max}$, in the loss
5	**Expression**	$\arg\min_{18 \le \mathbf{u} \le 22} \mathbb{E}[\mathcal{L}(\mathbf{u}, R) \mid E]$
6	**Model**	The posterior on R; the range matters, since above it overheating is certain and the cost falls again
	Reduce	$\arg\min_{18 \le \mathbf{u} \le 22} \big(P_{\mathrm{elec}}(\mathbf{u}) + 10\, P(\vartheta_B(\mathbf{u}) > \vartheta_{\max} \mid E, \mathrm{do}(\mathbf{u}))\big)$
7	**Check**	21.35°C, with risk 0.065; at a price of 100, 20.20°C and 0.005. Each tolerance is told something else
	Class	Expected-loss minimization

$$\boxed{\arg\min_{18 \le \mathbf{u} \le 22} \big(P_{\mathrm{elec}}(\mathbf{u}) + 10\, P(\vartheta_B(\mathbf{u}) > \vartheta_{\max} \mid E, \mathrm{do}(\mathbf{u}))\big)}$$

Problem 17.8. A log shows the UPS's efficiency measured only at light load, which is when maintenance tests are run. The planner asks what the efficiency would be if the load were reduced by choice. When do the two agree?

Run, a relation read from a log		
1	**Sentence**	Efficiency found at light load in a log, against the efficiency if the load is set light: when do they agree?
	Intention	Compare
2	**Target**	Two expressions, and the condition that makes them equal
	Declare	η; P_{load} [kW], with Q' a stated value; v whatever else came with a light hour (a maintenance state, a mode)
3	**Known**	The log; whether v affects η is not stated
	Random here	η, and v where it varies
4	**Event**	None: two distributions
5	**Expression**	$p(\eta \mid P_{\text{load}} = Q')$ against $p(\eta \mid \text{do}(P_{\text{load}} = Q'))$
6	**Model**	They agree if nothing that moves the load also moves the efficiency, or if the load was assigned at random
7	**Check**	The log alone cannot say; a stated v can be conditioned on, and an unstated one cannot
	Class	Posterior of a parameter, under a setting

$$\boxed{p(\eta \mid P_{\text{load}} = Q') \;\; \text{against} \;\; p(\eta \mid \text{do}(P_{\text{load}} = Q'))}$$

17.6 The exercises

The solutions are in Appendix A, and the facts of the three systems in Appendix D. On the two-hall plant, E is the 40 kW trip test. The limit of 30°C, the coefficient of performance, every loss table and the log are illustrative. The last exercise closes Part IV.

Exercise 17.1 (*formulate*). A plant manager writes: "We have three options for the opening: leave it, enlarge it for 12 units, or fit a dedicated bypass duct for 30, which rules out overheating. A hall overheating costs 150. I want no more than a one-in-a-hundred chance of that, and I want to spend as little as possible." Separate the objective, the constraint, the states and the actions, and write the problem given E.

Exercise 17.2 (*formulate*). The energy manager's ticket reads: "Find the best set point." Formulate the ticket.

Exercise 17.3 (*formulate*). Run this one: at a set point of 23°C, with a limit of 31°C, given E, unit B out and hall B at 60 kW, is hall B over its limit?

Exercise 17.4 (*formulate*). The limit is lowered to 28°C. Rewrite the third rung, and say what moves and what does not.

Exercise 17.5 (*formulate*). Run this one: on the fact sheet's priors, with no test, keep the string, shed a quarter of the load in the next outage, or replace it? An outage not carried costs 100, shedding a quarter costs 5, and replacing costs 20, for a new string with $\varphi = 1$ (illustrative).

Exercise 17.6 (*formulate*). In the lesson's loss table, the false acceptance's cost is doubled to 380. Write the level it implies, and say whether the action on E changes.

Exercise 17.7 (*formulate*). Write, by expected loss over the three explanations of the network's meters, the choice among keeping the UPS, servicing it and recalibrating the PDU meter, given the five readings. State the table's shape, and say why a single level cannot replace it.

Exercise 17.8 (*formulate*). The energy manager's ticket reads: "Find the set point that minimizes plant electricity with hall B's overheating risk at most 5 %, on the trip test, on a unit's loss at 60 kW." Formulate the ticket.

Exercise 17.9 (*formulate*). A trip test takes unit B out by choice. Last year's log also shows unit B found tripped on three nights. Write what stands behind the bar in each, and say which one supports the N+1 claim as it stands.

Exercise 17.10 (*formulate*). Lesson 14's rule at 0.95 and this lesson's argmin agree on E. Write the condition on a two-action loss table under which the rule and the argmin agree on every reading.

Exercise 17.11 (*reverse*). Write one plain sentence for $\arg\min_a \sum_i \mathcal{L}(a, H_i)\, P(H_i \mid \mathbf{y})$, with a an action on the network.

Exercise 17.12 (*reverse*). Write one plain sentence for $P(\vartheta_A(\mathbf{u}) > \vartheta_{\max} \mid E, \mathrm{do}(\mathbf{u}), Q_A = 45\,\mathrm{kW}) \le 0.05$, naming each symbol's role.

Exercise 17.13 (*reverse*). Write one plain sentence for $P(\vartheta_B > \vartheta_{\max} \mid \mathbf{u} = 23°\mathrm{C},\ \text{the log})$, and say what question it does not answer.

Exercise 17.14 (*reverse*). Write one plain sentence for $\min \xi$ subject to $P(D_{\mathrm{out}} \le T_{\mathrm{rt}}(\xi) \mid d, \mathrm{do}(\xi)) \ge 0.99$, and name its class.

Exercise 17.15 (*critique and repair*). A note says: "The log shows the plant is safe at 24°C. It never overheated there." State what is wrong, and repair it.

Exercise 17.16 (*critique and repair*). A shift note says: "The PDU meter's calibration is booked for next week, so the offset explanation can be dropped from tonight's diagnosis." State what is wrong, and repair it.

Exercise 17.17 (*critique and repair*). A design note says: "We chose the set point so that hall B stays under 31°C at the mean resistance, so the risk is low." State what is wrong, and repair it.

Exercise 17.18 (*critique and repair*). A review says: "The best action at 0.82 is to modify the opening, so the probability that the opening is adequate must be low." State the error, and repair it.

Exercise 17.19 (*critique and repair*). A proposal says: "Minimizing expected loss with overheating priced in guarantees at most a 1 % risk of overheating." State the error, and repair it.

Exercise 17.20 (*rapid*). For each, say whether the value is observed or set, and name the problem class where one applies.

(a) The relay was found open.

(b) The relay is opened for a test.

(c) Least load shed with the string carrying at 0.9.

(d) The set point in last June's log.

(e) The set point chosen for tonight.

(f) The lowest risk the plant can reach.

(g) Replace or repair, on stated costs.

(h) The load drops because a hall is empty.

(i) The load is shed by the operator.

(j) A test planned for next week.

Exercise 17.21 (*machine set*). Ten problems to close Part IV. Run each one like a machine.

(1) The rule at 0.99 for the claim that the opening carries 50 kW within 8 K: write it as a function, and its value on E.

(2) How often does a 40 kW test falsely accept a 58 kW opening, at the report's rule?

(3) How often does a 20 kW test accept a 110 kW opening, at the report's rule?

(4) Before any test, how likely is a 20 kW test to accept, on the stated prior?

(5) Given E, how likely is a 40 kW test on fresh sensors to accept, the rule on its reading alone?

(6) Of plants accepted from a 40 kW test by the rule at 0.99, under the prior, what share are inadequate?

(7) The timed-discharge rule at 0.90: how often is a string at 0.95 of nameplate kept?

(8) Given $y = 8\,\text{kW}$, recalibrate the PDU meter for 5, or do nothing, at 40 if an offset is there (illustrative)?

(9) With unit A out and hall B's set point at 22°C, is hall A at 50 kW over 30°C (illustrative), given E?

(10) At a set point of 22°C, the largest load on hall B with the risk over 30°C at most 1 %, on a unit's loss, given E?

17.7 What to retain

- A decision is an argmin, and the posterior is only one of its two inputs: on an illustrative table, 34.6 against 10, where a point value says 0.
- A level states a ratio of losses: 0.95 is about 19 to 1. Risk tolerance can be written on a loss or on a probability, and the two do not say the same thing.
- Three neighboring sentences are three problem classes: probability estimation, risk optimization, chance-constrained optimization.
- do($\mathbf{u}$) is the setting, and a value in a log comes with what set it. On the illustrative log, 24°C has 0.103 on the log's loads, 0.289 on the plant's own, and 0.833 at full load: the bar and the load both matter.
- Decisions and plans are not evidence.

Part V

Formulating Experiments and Information Questions

Part VI

Integrated Case Studies

Appendix A

Solutions to the exercises

The solutions show the reasoning, not only the expression, and they name the common wrong formulations and the question each of those does answer. A formulation that differs from the one given here is not necessarily wrong. The test is step 7 of the method: read the expression back in words and hold it against the original sentence. Appendix B gives the scoring.

Lesson 1

Exercise 1.1. The statement has six clauses, and each has a kind. “Each hall carries 60 kW” is a physical fact. “The drawing shows one spare unit” is also a physical fact, but a fact about the drawing; whether the spare can do its job is an uncertain quantity. “Before the handover is signed” refers to a decision, the signing. “The opening must be shown to carry a hall’s load” is a question about an uncertain quantity, the conductance. “If a unit trips at full load” is a condition, and it goes behind the bar. “Whether a third unit should be bought instead” is a decision.

Two questions follow. The first is an estimate, and it can be written today: the seven steps give the sheet of Table 1.6, with the expression $P(G \geq G^* \mid E)$. That estimate feeds a decision, the handover, and as a decision it lacks its level: a first form is “sign if $P(G \geq G^* \mid E) \geq p_{\text{req}}$”, with p_{req} not supplied. The second question, whether to buy a third unit, is a decision that lacks both a level and its consequences: the cost of the unit, and the cost of the outages it would prevent. Neither can be written until those are supplied. The common wrong formulation is one expression for the whole statement. The common wrong sort treats a decision as an event with a probability, as in “the probability that a third unit should be bought”, which has no meaning until consequences are stated.

Exercise 1.2. The two chillers and the common feeder are physical facts. The statement holds two questions. The first has the intention to decide, and it rests on an estimate: $P(\mathcal{A} \mid E)$, where $\mathcal{A}$ is the event that the redundancy claim holds, compared with a level. Both parts need supplying. The claim must be defined, for example as one chiller carrying the design load within its limit, and the word “sufficient” is where the level is hiding. The common feeder bounds the claim, because a feeder outage defeats both chillers, so the claim can only concern the loss of one chiller.

The second question has the intention to design an experiment. The phrase “the risk of accepting an inadequate system” has two readings that the sentence does not separate: the probability that an accepted system is inadequate, and the probability that an inadequate system is accepted. Lesson 15 is about that difference. Under the second reading, the expression is a probability over readings not yet taken, one for each candidate test load: it contains an acceptance rule, a stated inadequate system behind the bar, and the future readings $\tilde{E}_Q$ on

the left, and the question compares it across loads. Lessons 16 and 21 write it in full. The common wrong formulation runs both questions together as a single probability that the system is adequate.

Exercise 1.3. The intention is to predict. The sentence: given that unit B went out at time t_0 and given what is known, what is the probability that the service is delivered at every time until unit B returns? The target is a future path. What is known includes that unit B is out, tonight's load, the rates, and the readings taken since 02:00, which belong in E because they are evidence in hand about tonight's temperatures. With D the repair time, itself random, the event is $S(t) = 1$ for every t in $[t_0, t_0 + D]$, with the limit read at tonight's load rather than the full load. The expression is

$$P\big(S(t) = 1 \text{ for every } t \in [t_0, t_0 + D] \mid \text{unit B out at } t_0,\ E\big).$$

It conditions on the present state, unit B out, rather than on an all-up start, and its horizon is random. The model needs G, the failure processes of unit A and of the feeder, the distribution of the repair time and tonight's load. Read back: the probability that nothing else goes wrong and the opening copes until the repair is finished, which is what the operator asked. The problem class is a predictive probability over a random horizon. One common wrong formulation is $P(G \geq G^* \mid E)$, which has no horizon and leaves out unit A and the feeder failing during the repair. Another uses a fixed horizon, which drops the randomness of the repair time.

Exercise 1.4. This question is intentionally ambiguous, and it is not ready. All four items are missing. The failure is not named: one unit, the feeder, or anything at all. The horizon is not named: an instant, the length of a repair, or the contract year. The conditions are not named: full load or the tenant's actual load. The level is not named: what probability would count as "can survive". Four readings give four expressions.

(a) The loss of one unit, in steady operation at full load, asked as a matter of capacity: an estimate, $P(G \geq G^* \mid E)$. This says that the remaining unit can carry both halls once settled; surviving the loss through the whole of its repair is the path probability of reading (b).

(b) The most literal reading: the cooling is kept until the failed unit is back. This is the expression of Exercise 1.3, $P\big(S(t) = 1 \text{ for every } t \in [t_0, t_0 + D] \mid \text{the unit out at } t_0,\ E\big)$.

(c) The loss of the feeder. Both units go with it, so $P(S(t) = 1 \mid \text{feeder out at } t,\ E) = 0$ on the stated structure. No evidence is needed, and the honest answer is no.

(d) Any failure over the contract year: the prediction of equation (1.2).

Each reading is then compared with a level that the tenant must supply. The right answer asks the tenant which reading is meant and says what each gives; by the rubric of Appendix B, the expression for each reading is enough, and a full sheet for each is not required. The common wrong response picks reading (a) because it is the one in which the plant does well.

Exercise 1.5. Given what is known, this is the probability that the opening's conductance falls short of what carries one hall's load at the allowable difference. It is not "the probability that the plant fails", since it is a statement about the opening alone.

Exercise 1.6. Given what is known, this is the probability that the service runs for two years without a single interruption. It is not "the service is available for two years", which could be taken for a fraction of the hours. The expression states no start, and a complete one would put the start behind the bar, as equation (1.2) does.

Exercise 1.7. Given what is known, this is the probability that both units are unavailable for their own reasons. A time is missing, and a random hour is the usual reading. It is not the

probability that both units are unavailable, because a feeder outage takes both units out at once and lies outside the event as defined; reading it as the plain sentence is the failure class *dependence ignored*. At a random hour, on the stated rates, the expression is 0.0024, while both units are unavailable for any reason with probability 0.012, about five times as much. Nor is it the probability of losing the service, since on the series structure one unit is enough to lose it.

Exercise 1.8. The expression answers what fraction of hours the plant is in service. The failure class is the horizon missing. The repair is equation (1.2) with the same conditions, $P(S(t) = 1 \text{ for every } t \in [0, 1 \text{ yr}] \mid \text{all up at } t = 0, \text{ rates}, G \geq G^*)$, which on these rates is 0.020. The gap is not due to the feeder alone. For the unit pair by itself, with the feeder left out and the same conditions (the units' own outages, the opening taken as adequate, both units up at $t = 0$ for the year), the two values are 0.998 and 0.803.

Exercise 1.9. The expression answers whether the opening is adequate, a question about a parameter of the plant, where the owner asked about the service, an outcome. The failure class is *parameter against outcome*. A statement about capacity is not an availability, and with the opening certainly adequate the service is still lost whenever the feeder or both units are out. The repair chooses a reading of the question, as in Table 1.5, and writes equation (1.1) or equation (1.2), in which the conductance is one unknown among several.

Exercise 1.10. There are two errors. The conductance has one value, so the uncertainty about it is epistemic. The scatter of the readings is aleatory, and it is one of the reasons the conductance is not known. The second error concerns averaging. A longer test reduces the effect of the noise on the mean of the readings, but it does nothing to the sensors' constant offsets, which are also epistemic and leave a floor; a calibration against a reference would address those. The repaired statement: the conductance is epistemic, and a longer test reduces only the noise's share of what is unknown about it.

Exercise 1.11. The kind of each sentence comes first, using the four kinds, with a question sorted by what it asks about. The intention follows where the sentence is a question, and then the kind of uncertainty where a quantity is uncertain.

(a) A physical fact.

(b) A question about an uncertain quantity; estimate; epistemic.

(c) A question about a prediction; predict; aleatory.

(d) A question about a decision; decide.

(e) A question about a decision, the settings; optimize. The overheating it guards against is uncertain in both ways, the conductance epistemic and the failures and loads aleatory. The sentence lacks the probability level: "keeps hall B within its limit" must say with what probability.

(f) A question about an uncertain quantity, the cause of the reading, which has one value; diagnose; epistemic.

(g) A question about a decision, the test; design an experiment. The test's readings, not yet taken, are aleatory given the plant, and the test is chosen to reduce the epistemic uncertainty about the conductance.

(h) A question about an uncertain quantity, the plant's structure; compare; epistemic.

(i) A question about an uncertain quantity; estimate; epistemic.

(j) A question about a prediction; predict; aleatory, at a rate that is itself epistemic.

Lesson 2

Exercise 2.1. The intention is to predict. The present state is not part of the event: unit B out after its trip, with the feeder and unit A up at t_0, is known, and it goes behind the bar. Unit B's trip was found, not chosen, so it is an observation; a planned outage would be a setting, which Lesson 3 distinguishes. The supervisor's gloss joins two outages with "neither ... nor", which is the intersection of their complements, and ranges over time with "at any time before then", which is an intersection over $[t_0, t_1]$, with t_0 now and t_1 06:00. The event as glossed is therefore $F_{\mathrm{fd}}(t)^c \cap F_A(t)^c$ for every t in $[t_0, t_1]$, and the expression is

$$P\big(F_{\mathrm{fd}}(t)^c \cap F_A(t)^c \text{ for every } t \in [t_0, t_1] \mid \text{unit B out, the feeder and unit A up at } t_0,\ E\big).$$

Reading it back against the question at step 7 shows that the gloss is narrower than the question. With unit B out, hall B is cooled only through the opening, so "hall B keeps its cooling" also needs the opening to pass the hall's load at tonight's load. The repaired event adds $\cap \{G \geq G^*\}$, with the threshold read at the load actually carried, and the model then carries G as well as the failure processes of the feeder and unit A. The horizon is taken as fixed at 06:00, as the crew promised; if the repair may run late, it is random, as in Exercise 1.3. Two wrong translations are common. Reading "neither ... nor" as "not both" allows the feeder to fail as long as unit A does not, which is the connective mistranslated. Reading "until 06:00" as the single hour at 06:00 asks about one instant where the sentence asked about every hour before it.

Exercise 2.2. Each hall's loss is built first. Hall B loses its cooling when the feeder is out, or when unit B is out and either unit A is out or the opening cannot pass hall B's load to unit A:

$$F_{\mathrm{fd}} \cup \big(F_B \cap (F_A \cup \{G < G^*\})\big).$$

The fact sheet states the limit only for hall B above hall A, with unit B out; hall A's loss is taken to be the same event with A and B exchanged, the same limit holding the other way, and the solution says so. The loss of both halls at once is the intersection of the two, which reduces to $F_{\mathrm{fd}} \cup (F_A \cap F_B)$, and "at once" at some time in the policy year is a union over the hours of the year. The second branch, "either hall is without cooling for more than 24 consecutive hours", is a union over the two halls of the event that some interval longer than 24 h lies wholly inside that hall's loss. The policy pays on the union of the two branches over the policy year. The expression is the probability of that union given the plant's state at the start of the policy year, all up if nothing else is known, and E. "Either" is inclusive here, since a loss of both halls for a day is plainly covered, and "more than" on a duration that varies continuously is harmless. The model needs the length of each loss, not only its occurrence. Two translations are tempting and wrong: both halls lost at once written as $F_A \cap F_B$ alone, which forgets that the feeder takes both units with it; and "more than 24 consecutive hours" read as 24 h in total over the year, a word over separate losses read as one over total time.

Exercise 2.3. This question is intentionally ambiguous, and it is not ready. Its connective is clear, at most one, ≤ 1, and its noun is not. The missing items are the following. The event: an outage of what, the tenant's hall, the service to either hall, the feeder, or any component; and what counts as one outage, since two component outages that overlap may be one interruption or two. The horizon: "next year" needs its start, which goes behind the bar. The level: the sentence is a promise, so the probability with which it must hold is not given, and Lesson 4 takes up promises. The conditions are steady operation, and the sentence names no exclusion. Four readings give four expressions, each over a calendar year $[0, \tau]$ and given all up at $t = 0$ and E. Write $N_{\mathrm{T}}(\tau)$ for the number of interruptions of the tenant's hall and $N_S(\tau)$ for the number of interruptions of the service of Lesson 1.

(a) The tenant's hall: $N_{\mathrm{T}}(\tau) \le 1$.

(b) The service of Lesson 1: $N_S(\tau) \le 1$.

(c) The feeder: $N_{\mathrm{fd}}(\tau) \le 1$, whose probability over a calendar year, with the feeder up at $t = 0$ and on its stated rate and repair time, is 0.118.

(d) Any component: $N_{\mathrm{fd}}(\tau) + N_A(\tau) + N_B(\tau) \le 1$.

The solution uses reading (a), since the tenant's interest is the tenant's own hall, and says so to the tenant. The point of the exercise is that the ambiguity sits in the noun: translating the connective perfectly does not remove it. The common wrong response is reading (c), the one that can be written at once from the fact sheet.

Exercise 2.4. Given what is known, this is the probability that at a random hour the feeder is out or both units are out for their own reasons. It is the loss of the service when the opening is adequate, equation (2.3). It is not the probability that both units are out, which leaves the feeder out, nor the loss of the service on the series structure, which also counts one unit out.

Exercise 2.5. Starting with everything up, and given what is known, this is the probability that the feeder fails at most once in the coming calendar year. With $N_{\mathrm{fd}}(\tau) < 1$ it becomes the probability of no failure at all. With the feeder up at $t = 0$, its stated rate and its 24 h repair, the two are 0.118 and 0.025: on a count, the strict and the loose inequality differ by the probability of the boundary value, here exactly one failure.

Exercise 2.6. Given what is known, this is the probability that at a random hour the feeder is up and exactly one unit is out, which is the plant running on its spare. It is not the probability of losing the service, which this event avoids whenever the opening is adequate, and it is not "at least one unit out", which would include both.

Exercise 2.7. The expression answers whether both units are available, which is "neither unit is out". The failure class is the connective mistranslated: "at least one" was written as "both". By the second of equations (2.1), the repair is

$$P(F_A^c \cup F_B^c \mid E) = P\big((F_A \cap F_B)^c \mid E\big).$$

On the stated rates, for the units' own outages at a random hour, the two are 0.905 and 0.998, so the error overstates the probability of failing the requirement, no unit available: 0.095 against 0.0024, about 40 times too large.

Exercise 2.8. The event describes hall B within its limit while the feeder is up, so every hour of a feeder outage is counted as a failure, where the sentence excused those hours. The failure class is the connective mistranslated: "unless" was read as "and not". The repair is the union,

$$\{\text{hall B within its limit}\} \cup F_{\mathrm{fd}},$$

which holds on every outcome where hall B is within its limit and on every outcome where the feeder is out. Whether the excused hours belong in the event as here, or should be removed from the count by standing behind the bar, is a separate choice about requirements, made in Lesson 4.

Exercise 2.9. Each item is an event on the plant at one hour unless the item names a period, in which case the event is over $[0, \tau]$, a calendar year.

(a) $F_A^c \cap F_B^c \cap F_{\mathrm{fd}}^c$; "while" is an "and" here, since nothing is given.

(b) The complement of "at least two out": no component out, or exactly one out. Writing it as "not all three out" is the connective mistranslated.

(c) $F_B^c \cup F_A^c$, which is $(F_A \cap F_B)^c$: "only if" forbids the one combination of both out.

(d) $(F_{\text{fd}} \cap F_B^c) \cup (F_{\text{fd}}^c \cap F_B)$; "but not both" makes the "or" exclusive.

(e) $N_{\text{fd}}(\tau) > 3$, that is $N_{\text{fd}}(\tau) \geq 4$ on a count.

(f) $\{N_{\text{fd}}(\tau) \geq 2\} \cap \{N_{\text{fd}}(\tau) \leq 4\}$; "between" is an intersection of two bounds, and "inclusive" fixes the boundaries.

(g) $N_{\text{fd}}(\tau) \geq 1$, or equally $\bigcup_{t\in[0,\tau]} F_{\text{fd}}(t)$ for the feeder up at $t = 0$: a union over time and a count give the same event.

(h) $\bigcap_{t\in[t_0,t_0+D]}\{S(t) = 1\}$, with t_0 the start of the repair and D its length. The item holds a condition: "the repair of unit B" means unit B is out at t_0, which is known and stands behind the bar, not in the event.

(i) $\{S(t) = 1\} \cup (F_A \cap F_B)$; "unless" is a union, and the claim covers the outcomes $(F_A \cap F_B)^c$.

(j) The total of unit A's outage lengths in $[0, \tau]$ exceeds 168 h. This is a total, not each outage, and it differs from the quantifier "at a time".

(k) Not an event to be built. F_B is asked about and F_{fd}^c is given, so it stands behind the bar: $P(F_B \mid F_{\text{fd}}^c)$. Lesson 3 treats the bar.

Lesson 3

Exercise 3.1. The intention is to diagnose. The candidates come first, from the structure. Hall B loses its cooling when the feeder is out, or when unit B is out and either unit A is out or the opening is inadequate (the hall event of Exercise 2.2). These causes can occur together, since a unit can be out for its own reasons while the feeder is out, so they are made into a partition of hall B's loss at t_0: the feeder out; the feeder up with both units out; the feeder up with unit B out, unit A up and the opening inadequate. The relay points to the first and unit A's alarm to the second. Unit A out alone does not take hall B's cooling, so the alarm cannot be the whole story.

Every clause then takes its place. Known, behind the bar: the alarm fired, the relay fired, hall B's cooling was lost at t_0, and E. The alarms are readings, not states. That the relay fired is known; that the feeder was out is not, since a relay can fire on a sag without an outage. Asked about: which cause held. Carried: the states of the three components and G. For each candidate c the expression is

$$P(c \text{ at } t_0 \mid \text{alarm of unit A fired, relay fired, hall B lost at } t_0,\ E),$$

and the three sum to one, because they partition the loss. The model needs the structure, how probable each cause is from the rates, and how each alarm behaves under each cause, which Lesson 11 writes. Two wrong formulations are common. Writing $P(\text{relay fired} \mid F_{\text{fd}})$ is the transposed conditional. Writing F_{fd} behind the bar because the relay fired puts an unknown there as if it were known. Both are the wrong side of the bar.

Exercise 3.2. The intention is to predict, and the expression is

$$P(S(t) = 1 \text{ for every } t \in [t_0, t_0 + 6\text{ h}] \mid \text{do(unit B off at } t_0),\ \text{the feeder and unit A as found at } t_0,\ \text{Tuesday's load},\ E).$$

The start is the plant as found on Tuesday at t_0, normally the feeder and unit A up, and Tuesday's load stands behind the bar with it, since the opening's threshold is read at the load carried. It differs from Exercise 1.3 in two ways. There unit B was found out, which is an

observation. On the fact sheet's model of disturbance windows, a unit found out is evidence of a disturbed window, which raises the chance that unit A or the feeder fails during the repair. Here unit B is set off, so the weight of a disturbed window stays where it was. The horizon also differs in kind: there it was the random length of a repair, and here it is a fixed six hours. The common wrong formulation copies Exercise 1.3's expression with "unit B out at t_0" behind the bar, which reads the planned outage as a found one.

Exercise 3.3. This question is intentionally ambiguous, and it is not ready. The event asked about, $G \geq G^*$, is clear. What is missing is what stands behind the bar: whether the test has been run, which test it is (its load), and what will be known about it, the readings or only the verdict. There is no horizon, since the conductance is fixed, and no level is asked. Four readings give four expressions. The tilde on $\tilde{E}$ marks readings not yet taken, as in Lesson 1, and A is the rule that passes or fails the test.

(a) The test has been run and its readings are in hand: $P(G \geq G^* \mid E)$, with the readings in E.

(b) The test has been run and only its verdict is reported: $P(G \geq G^* \mid \text{the test passed},\ E_0)$, with E_0 the evidence before the test. This is a weaker condition than (a).

(c) The test has not been run: the probability that the opening is adequate in the cases where the test would pass, $P(G \geq G^* \mid A(\tilde{E}) = 1,\ E_0)$.

(d) The transposed reading the owner may mean, how often an adequate opening passes: $P(A(\tilde{E}) = 1 \mid G)$ for $G \geq G^*$.

The solution uses (a) if the test has been run and (c) if it has not, and says so. Lessons 14 to 16 write the rule and the last two expressions in full. The common wrong response answers (d), the test's side of the bar, when the owner asked about the opening.

Exercise 3.4. Given that the service was lost at t_0, and given what is known, this is the probability that the feeder was out at that time. It is not the probability that a feeder outage loses the service, which is the forward direction.

Exercise 3.5. This is the probability that the service is lost while the feeder is out. On the stated structure both units lose their supply with the feeder, so the value is one. No rate, reading or record enters, because the structure alone decides it.

Exercise 3.6. With unit B switched off deliberately, and given what is known, this is the probability that unit A is out for its own reasons. It is not "given that unit B has failed": switching unit B off carries no news about a shared cause, so on the window model the value is unit A's unconditional 0.075, not the 0.150 that follows a unit found out.

Exercise 3.7. The argument computes the probability of a high reading given a large offset, nine in ten, and reports it as the probability of a large offset given a high reading. The failure class is the wrong side of the bar, in its first form, the transposed conditional. The repair writes

$$P(\text{offset of the hall B sensor} > 0.5\,\text{K} \mid \text{hall B read high},\ E),$$

whose model needs two things the engineer did not give: how probable a large offset was before the reading, which the fact sheet's prior spread of 0.5 K on an offset supplies, and how often a sensor with a small offset reads high, which happens whenever hall B really is warm. A high reading that is common under both says little about the offset. The repaired expression can be written and its model named; it cannot be evaluated from the engineer's sentence.

Exercise 3.8. The argument computes how probable a record of seven failures in three years is if the rate is 2.5 per year, and reports it as how probable the rate 2.5 is given the record. The failure class is the wrong side of the bar, the transposed conditional. The repair writes

$$P(\text{unit A's failure rate near 2.5 per year} \mid \text{seven failures in three years of service},\ E),$$

whose model needs what was believed about the rate before the record, such as the handbook's value, and how probable seven failures are at other rates; a record that many rates would produce does not single out one of them. The record stands behind the bar as evidence, which Lesson 11 writes.

Exercise 3.9. The expression answers how probable it is that both units are down together. The failure class is the wrong side of the bar, in its second form: the known condition, unit B down, was written into the event. The repair is $P(F_A \mid F_B,\ E)$. The numbers depend on the model, and the solution names it. On the fact sheet's model of disturbance windows, the repaired expression is 0.150 and the analyst's joint is 0.011, more than ten times smaller; the joint is the conditional multiplied by the probability that unit B is down. On the stated rates without disturbance windows, the same joint is 0.0024, at a random hour and for the units' own outages, so the model changes the joint too. Neither joint is the probability that both units are unavailable for any reason, which also counts the feeder's outage.

Exercise 3.10. For each item, the clause asked about goes on the left and the clause known behind the bar.

(a) Asked: a further failure this year; known: two failures so far, which is evidence. $P(N_A(\tau) - N_A(t) \geq 1 \mid N_A(t) = 2,\ E)$, with t now and τ the year's end.

(b) Asked: both together; nothing about either is known. $P(F_{\text{fd}} \cap F_A \mid E)$, a joint at a random hour.

(c) Asked: the feeder now; known: unit A found out. $P(F_{\text{fd}} \mid F_A,\ E)$; a present state, so the intention is to estimate.

(d) Asked: hall A's cooling over the work; known: unit A set out, an intervention, and the plant as found. $P(\text{hall A keeps its cooling for every } t \text{ in the work} \mid \text{do(unit A out at } t_0),\ E)$.

(e) Asked: the reading; known: the conductance and the load. $P(\text{the test reads 9 K} \mid G = 8000\,\text{W/K},\ Q = 60\,\text{kW})$, a statement about a test given the plant; for a reading that varies continuously, a small interval around 9 K.

(f) Asked: the service; known: an assumption, labeled. $P(S(t) = 1 \mid G \geq G^*,\ E)$, t a random hour. The assumption keeps its label in any report.

(g) Asked: both units unavailable for any reason; nothing known. $P\big(F_{\text{fd}} \cup (F_A \cap F_B) \mid E\big)$ at a random hour, since the feeder's outage takes both units. Writing $P(F_A \cap F_B \mid E)$ alone is dependence ignored.

(h) Asked: the year; known: the start. $P(S(t) = 1 \text{ for every } t \in [0, \tau] \mid \text{all up at } t = 0,\ E)$.

(i) Asked: the feeder's count; known: a year free of interruptions. $P(N_{\text{fd}}(\tau) = 0 \mid \text{no interruption in } [0, \tau])$, which is one on the stated structure, for any value of G, since every feeder outage interrupts the service. The forward question, a year free of interruptions given no feeder failure, is the unit pair's survival over the year: 0.803 on the stated rates, for the units' own outages, with the opening taken as adequate and both units up at $t = 0$.

(j) Asked: next year; known: last year's record, which is evidence and goes behind the bar. $P(\text{no interruption next year} \mid \text{the record},\ E)$.

Lesson 4

Exercise 4.1. The clause takes apart as follows. The event is a loss of cooling at some hall lasting more than 24 h: the union over the two halls, and over the policy year, of an interval longer than 24 h lying wholly inside that hall's loss, with each hall's loss built as in equation (4.2). The hall events are Exercise 2.2's. The conditions are the start, the plant's state at the beginning of the policy year (all up at $t = 0$ if nothing else is known), and the exclusion "other than during planned maintenance". The exclusion has two translations, and for an event over a year they differ more than at a random hour. Counted as cooled, the planned hours join the event as a union, so a loss that runs into a maintenance period is cut at its edge. Removed, the planned periods are taken out of the year, and a loss must last more than 24 h within the hours that remain; that is a condition on the path, and it needs the maintenance schedule in the model. The clause does not say which, and the solution asks the insurer. The horizon is the policy year. There is no level, because the insurer does not require the event to hold or fail: it prices it. The intention is to decide on a premium, resting on a prediction, and the expression is the probability of the event given the start and E, or the expected number of such losses in the year if a claim can recur. Two wrong formulations are common: 24 h in total over the year, where the clause asks about consecutive hours; and planned maintenance left in the event, which counts the insurer's own exclusion as a claim. Both drop or misplace a part of the requirement.

Exercise 4.2. The statement holds two requirements. The first is about a degraded state, and "runs on one unit" has two readings: the hours with exactly one unit out and the service delivered, the second band of Figure 4.1; or the hours with the feeder up and exactly one unit out, whether or not the service survives, the second and third bands together. With the opening adequate the two coincide, and the solution uses the first and says so. The requirement is that those hours number at most 400 in the year, and it is ambiguous again between each year's hours as they fall, an event on a realized total that needs a probability level, and an average over years, an expectation compared with 400 h directly:

$$P(\text{hours on one unit in } [0,\tau] \le 400 \mid \text{all up at } t = 0,\ E) \ge p_{\text{req}},$$

or

$$\mathbb{E}[\text{hours on one unit in } [0,\tau] \mid \text{all up at } t = 0,\ E] \le 400.$$

The solution writes both and asks which is meant. The second requirement is about functional failure: for each hall and every loss of its cooling in the year, the length of the loss is at most 4 h, Lesson 2's word "at a time" ranging over losses. Its expression is the probability of that event given all up at $t = 0$ and E, against p_{req}. The conditions are steady operation after that start, the horizon is one year, and neither requirement states a probability level. The common wrong formulations are a single expression for both requirements, and the hours on one unit read as hours without cooling, which moves the degraded band into the failure set.

Exercise 4.3. This requirement is intentionally ambiguous, and it is not ready. The event is missing: which failure, the tenant's hall, any hall, both halls, or any component. The horizon is missing: the lease term, with its start behind the bar. The conditions are missing: the load, and whether planned maintenance or a feeder outage is excused. The level is missing, because "never" read as a probability of one can be met by no plant. Three readings give three expressions, each over the lease term $[0,\tau]$ given all up at $t = 0$ and E.

(a) The tenant's hall never loses its cooling: the probability that the event of equation (4.2), for that hall, holds at no t in $[0,\tau]$.

(b) No total loss: the probability that $F_{\text{fd}} \cup (F_A \cap F_B)$ holds at no t in $[0,\tau]$.

(c) No component ever fails: $P(N_{\mathrm{fd}}(\tau) + N_A(\tau) + N_B(\tau) = 0 \mid \text{all up at } t = 0,\ E)$, which is small for any plant and is not what a tenant means.

Each is compared with a level the tenant must supply. The solution uses reading (a), since the tenant's interest is the tenant's own hall, and says so. The honest reply in the meeting is that no plant meets "never" literally, followed by three questions: which failure, over what term, and with what probability. The common wrong response is reading (b), the definition by which the plant looks best, as in Exercise 1.4.

Exercise 4.4. The event of the commissioning requirement is about a reading of a test, the measured difference of hall B above hall A, which is the true difference plus the sensors' offsets and noise. It is an acceptance criterion: an event on readings that decides a verdict. The design reading of Table 4.5 has its event on the plant's state, $\Delta T_{\mathrm{full}} \le \Delta T_{\mathrm{allow}}$, under the same condition. The two events are linked by the model of the readings, and they are not the same event: a plant within the limit can read over it, and a plant over the limit can read within it. The solution writes both events, names the gap between them, and stops; Lesson 14 writes the rule and its outcome as a function of the evidence. The common wrong formulation treats "passes commissioning" as "meets the specification".

Exercise 4.5. Starting from the plant in steady operation at the beginning of month m, and given what is known, this is the probability that the halls are cooled during at least 99.5 % of the month's hours, an hour of feeder outage counting as cooled. The threshold 0.995 is the contract's performance level. The probability itself is still to be compared with a probability level that the contract does not give.

Exercise 4.6. With unit B out, each hall at its design load and the plant settled, the probability, given what is known, that hall B is within the stated limit of hall A is required to be at least p_{req}. The event comes from the specification. The conditions do not: the specification states none, and these are the analyst's choice for the design reading, labeled as such. The level comes from the owner, and E is the commissioning evidence and the plant's stated facts.

Exercise 4.7. Starting with everything up, and given what is known, this is the probability of a total loss of cooling, both halls at once, at some time in the coming year. It uses the definition of failure as total loss, so the opening does not enter it, and it cannot say whether the N+1 claim holds.

Exercise 4.8. The comparison sets the mean share of hours in service, with the exclusion dropped, against a performance level that the contract set for each month's share with the exclusion kept. The failure class is the requirement without its conditions or its level, in two of its forms at once: the exclusion dropped, and a mean compared with a level set for each period. If the auditor speaks of months already past, breach is a matter of record: each month's share is known from the log and no probability is needed, only the clause's own event, $\bar{S}_{\mathrm{ex}}(m) \ge 0.995$, checked month by month. For the coming year the repair writes

$$P\big(\bar{S}_{\mathrm{ex}}(m) \ge 0.995 \text{ for every } m = 1, \dots, 12 \mid \text{steady operation at } t = 0,\ E\big),$$

which on the stated rates, with the opening taken as adequate, is about 0.8, while a single month meets the clause with probability 0.98, as Table 4.8 gives. A breach in some month of the year therefore has probability about 0.2, and whether that is acceptable needs a probability level the contract does not state.

Exercise 4.9. The assessment answers how often both halls lose their cooling at once. That event, $F_{\mathrm{fd}} \cup (F_A \cap F_B)$, does not involve the opening, so its value at a random hour on the stated

rates is 0.012 whether the opening is adequate or not, and it says nothing about adequacy. The failure class is the requirement without its conditions or its level, in the form of failure defined at the wrong level of function: the whole plant, where the N+1 claim concerns a hall. The repair defines failure as a hall losing its cooling, equation (4.2). Under the condition of the claim, one unit out with the other unit and the feeder up, that event holds exactly when the opening is inadequate, so its conditional probability is $P(G < G^* \mid E)$ and the opening enters fully. Without that condition, at a random hour, "some hall loses its cooling" is 0.104 with the opening inadequate against 0.012 with it adequate.

Exercise 4.10. The specification states an event and no level, and reading "shall" as $p_{\text{req}} = 1$ makes every plant fail, since no uncertain quantity is certain. The failure class is the requirement without its conditions or its level: a verdict drawn against a level nobody stated. The repair has two forms. Either the owner supplies p_{req}, and the design reading's expression of Table 4.5 is compared with it, or the requirement is read as a design condition, checked in the design model at the stated conditions.

Exercise 4.11. By the lesson's rule, the conditions count as stated when the operating state is named or the requirement says it holds in every state, with exclusions and any start among the conditions; the event counts as stated when every threshold it uses is stated somewhere. Each item names the missing parts.

(a) Conditions (which unit is out, what load, settled or not), horizon, level.

(b) The event in part (which definition of failure, and whether 99.9 % is a share as it falls or a probability), conditions (the operating state, exclusions), horizon, level.

(c) The event (an interruption of what), conditions (the start of the year, the operating state, exclusions), level. The horizon, a year, is given.

(d) The event's threshold: no limit for hall A is stated anywhere, since the fact sheet states hall B's limit above hall A only. And the level. The conditions are stated and no horizon is needed.

(e) The event in part ("survive" must be defined, presumably the service held), the horizon (through the outage, whose length is random), and the level. The condition, the feeder out, is stated. On the stated structure both units are out with the feeder, so the probability is zero and the requirement fails any positive level; the honest answer says so.

(f) Conditions (the load, which unit, the state of the rest at the trip), horizon (until the repair ends), level.

(g) Conditions (the start, the operating state), level. The event is a count of a component's failures, not a function, and the horizon is a year.

(h) None: the event $G \geq G^*$ is stated with its threshold, its condition is the evidence E, it needs no horizon, and the level is 0.99.

(i) The event (cool within what limit), conditions (the operating state and the form of the exclusion), horizon, level.

(j) Conditions (which hall, the operating state, exclusions, the start) and the horizon's term (the months of which period). The level is given, but not whether 95 % applies to each month separately or to all the months of the term jointly, and the two differ greatly.

Lesson 5

Exercise 5.1. A planned service is a setting, not an observation, so the condition is do(unit A off) (Lesson 3). Two limits apply. Unit B must carry both halls within its ca-

pacity, $g_{\mathrm{cap}} = 130\,\mathrm{kW} - (Q_A + Q_B)$, and hall A's heat must cross the opening the other way, which gives a limit state of the same form as equation (5.1) with the roles of the halls exchanged, $g' = \Delta T'_{\mathrm{allow}} - Q_A/G$. Here $\Delta T'_{\mathrm{allow}}$ is a limit on hall A above hall B. The fact sheet states no such limit: its 8 K is for hall B above hall A only. The failure domain is the union of the two, and the expression is

$$P\big(\{g_{\mathrm{cap}} < 0\} \cup \{g' < 0\} \mid \mathrm{do}(\text{unit A off}),\ Q_A = 50\,\mathrm{kW},\ Q_B,\ \text{settled},\ E\big).$$

The sentence leaves out hall B's load, without which neither limit can be evaluated, and hall A's limit, which must be asked for rather than borrowed from hall B. Two wrong formulations are common: copying hall B's 8 K to hall A without saying so, which is a threshold nobody stated, and checking only the opening, which drops one limit state of two.

Exercise 5.2. The intention is to decide, since the answer is an operating limit, and the target is a set of loads and its upper end. For one candidate load the expression is

$$P\big(\Delta T_{\mathrm{allow}} - Q_B/G < 0 \mid \text{unit B out},\ Q_B,\ \text{settled},\ E\big),$$

and the envelope is the set of loads for which it is at most 0.02, equation (5.3) with $p = 0.02$. The rule the operator wants is the largest load in that set. The model is what is known of G given the trip test and the opening's law. Read back, the rule holds for a settled state with unit A healthy through the night and hall A's load within unit A's capacity; the solution states those conditions with the rule. The common wrong formulation is the load at which the margin at $\hat{G}$ is zero, which is the envelope drawn at the point value and has an exceedance probability near one half, not 0.02.

Exercise 5.3. Hall B's settled temperature is hall A's temperature plus the rise, so a limit on hall B's own temperature is $g = 30^{\circ}\mathrm{C} - \vartheta_A - Q_B/G$, a difference in kelvin; the tenant's 30°C is the tenant's, not a stated value of the plant. It becomes the stated 8 K difference only under the condition that hall A is held at its set point of 22°C, and with unit B out that needs unit A to carry both halls, so the capacity limit joins as a union:

$$P\big(\{g_{\mathrm{cap}} < 0\} \cup \{g < 0\} \mid \text{a trip of unit B},\ \text{unit A and the feeder up},\ Q_A,\ Q_B,\ \text{settled},\ E\big).$$

The trip has not happened, so it is a condition of the question, not a found event, and so are the states of unit A and the feeder, which the solution takes as up and says so. Hall A's set point is taken as held at 22°C, an assumption the solution labels; the fact sheet gives the set point a prior spread of 2 K, and carrying that spread makes ϑ_A an unknown of the limit state. The tenant did not give tonight's loads, and the solution asks for them. Two wrong formulations are common: the rise alone, which silently assumes that hall A is held, and an absolute limit read as the difference limit with no condition on hall A.

Exercise 5.4. This question is intentionally ambiguous, and it is not ready. Missing are which limit is meant (the opening's, the surviving unit's capacity, or a hall's own temperature), under which condition (which unit is out and at what load), at which point (the point value, or a design value), and whether a margin or a probability is wanted. Four readings give four answers.

(a) The opening's margin at a stated load and the point value, $g(\hat{G};\ Q_B)$, a number in kelvin.

(b) The surviving unit's capacity margin at the design load, $g_{\mathrm{cap}}(60, 60) = 10\,\mathrm{kW}$, stated and certain.

(c) The probability of exceeding the opening's limit at a stated load, $P(g(G;\ Q_B) < 0 \mid \text{unit B out},\ Q_B,\ \text{settled},\ E)$, which is what an insurer can price.

(d) The most literal reading, "the" margin of the plant, the smallest over all limits and conditions, which is not a single number at all, because the margins are in different units.

The solution uses reading (c) for an insurer, who prices a probability, and says so, asking for the load at which it is to be read. The common wrong response is to quote reading (b), because it is stated and certain, which answers a limit the insurer did not ask about.

Exercise 5.5. Given the 40 kW trip test, this is the probability that, with unit B out and hall B at 45 kW, the opening cannot pass hall B's load within the allowable difference once the plant has settled. The limit state is written in kilowatts, and the event is the same as with the form in kelvin of equation (5.1).

Exercise 5.6. It is how far, in kelvin, hall B's settled rise would lie below its limit with unit B out and hall B at 50 kW, if the conductance were exactly its point value. It is not a probability, and it is not a statement that hall B will stay within its limit.

Exercise 5.7. It is the set of hall B's loads for which, with unit B out and the plant settled, the probability of exceeding hall B's limit is at most 5 % given the trip test: the operating envelope at a level of 0.05.

Exercise 5.8. The report computes a margin, divides it by the limit, and calls the ratio a risk. A ratio of kelvin to kelvin has no unit, but nothing has made it a probability. The failure class is the point value read against the limit, in the form of a margin reported as a probability. The repair is the exceedance probability,

$$P(\Delta T_{\text{allow}} - Q_B/G < 0 \mid \text{unit B out},\ Q_B = 52\,\text{kW},\ \text{settled},\ E),$$

whose model is what is known of G after the trip test. The margin enters it only as the distance at one point.

Exercise 5.9. The note checks one of the requirement's two limit states, the capacity, which is stated and certain, and draws the requirement's verdict from it. The failure class is the point value read against the limit, in the form of one limit state checked where the requirement holds several. The repair writes the union $\{g_{\text{cap}} < 0\} \cup \{g < 0\}$ at the load in question, with its probability given E. On the stated values the capacity term adds nothing to that probability at the design load, so the requirement's probability is the opening's, which the note never examined.

Exercise 5.10. The figure is the envelope drawn at the point value: the load at which the margin at $\hat{G}$ is zero. At that load hall B is on the boundary if the conductance equals its point value, and the probability of exceeding the limit is close to one half. The failure class is the point value read against the limit, in the form of an envelope drawn at the point value. The repair is to choose a level p with the owner and to publish the largest load whose exceedance probability given E is at most p, with the conditions it holds under: on this test, 54.6 kW at a level of 0.05 and 50.4 kW at 0.01, with unit B out and the plant settled.

Exercise 5.11. Each item names the kind of answer and the limit state.

(a) A margin on stated values: $g_{\text{cap}}(58, 58) = 14\,\text{kW}$.

(b) An exceedance probability of the opening's limit state at $Q_B = 48\,\text{kW}$, with unit B out and the plant settled.

(c) An envelope, at a level of 0.05; on the 40 kW trip test it reaches 54.6 kW.

(d) A margin at a stated conductance: $8\,\text{K} - 40\,\text{kW}/(9000\,\text{W/K}) = 3.56\,\text{K}$, with no probability involved, since the conductance is given.

(e) The exceedance probability of the union of the capacity and opening limit states, compared with a level of 0.1.

(f) Not a question about a limit: a point value of the rise, 5.93 K, from which the margin follows.

(g) The boundary of the envelope drawn at the point value.

(h) Not ready: no limit for hall A above hall B is stated, so the limit state cannot be written.

(i) A margin on stated values, $g_{\text{cap}} = 5\,\text{kW}$: the probability that the capacity is exceeded is zero on the stated values.

(j) No. The margin is $\Delta T_{\text{allow}} - Q_B/G$, which is not proportional to G; doubling the conductance halves the rise, and the margin grows by half the rise.

Lesson 6

Exercise 6.1. The sentence holds two questions. The first is a conditional: given a fault on the cooling branch, the probability that the IT racks lose their supply,

$$P(\text{IT branch without supply} \mid \text{a cooling-branch fault}),$$

which on the stated structure is the probability that the cooling branch's breaker fails to clear its fault, 0.01 per demand on the network's illustrative inputs. It needs no rate, because the fault is given. The second is a joint over the coming year: the probability that at least one fault on the cooling branch occurs and its breaker fails to clear it,

$$P(\text{an uncleared cooling-branch fault in } [0,\tau] \mid \text{steady operation at } t = 0,\ E),$$

with $\tau = 1$ yr and "an" meaning at least one, with the count on calendar time from the named start. On the stated model such faults arrive at a constant rate, the branch's fault rate times the probability of failing to clear. Two wrong formulations are common. The joint reported for the first question answers how often the event happens, not how likely it is given the fault, Lesson 3's second form. And the two branches' own fault probabilities multiplied ignore the propagation through the protection, which is this lesson's class.

Exercise 6.2. The intention is to predict from a present state. Unit B was found out at t_0, which is an observation, and stands behind the bar with E, the 40 kW trip test, and the stated model of failures and windows. The event is that hall B keeps its cooling at every time until unit B returns: the feeder up, unit A up, and the opening passing hall B's load within the limit, for every t in $[t_0, t_0 + D]$, with D unit B's repair length, a random horizon. At the design load that last condition is $G \geq G^*$; at a lower load it is $Q_B/G \leq \Delta T_{\text{allow}}$, and the load then belongs behind the bar. The expression is

$$P\big(S(t) = 1 \text{ for every } t \in [t_0, t_0 + D] \mid \text{unit B out at } t_0,\ \text{design load},\ E\big).$$

Three dependences are in it. Unit A's failure rate is conditional on unit B's state, five per year during the repair instead of 2.5, on the vendor's illustrative figure: a propagation through load sharing. The window's state is a shared cause, and a unit found out raises the weight of a disturbed window, which raises unit A's and the feeder's rates during the repair (Lesson 3). The feeder is a shared component. The model needs unit A's failure process at its raised rate, the feeder's, the distribution of the repair's length, the window model and what is known of G. It is Exercise 1.3 with its dependences written. The common wrong formulation keeps Exercise 1.3's model with unit A at its normal rate, which ignores the load sharing.

Exercise 6.3. The action item is intentionally ambiguous, and it is not ready: "independence" names no events. Missing are which events (the units' own outages, or their unavailability for any reason), under which condition (given the window's state, or with it unknown), and at what time (a random hour, or over the other's repair). Four readings give four answers.

(a) Their own outages given the window's state: independent by the stated model, $P(F_A \cap F_B \mid c) = P(F_A \mid c)\,P(F_B \mid c)$ for each state c.

(b) Their own outages with the window unknown: dependent, since $P(F_A \mid F_B) = 0.150$ against $P(F_A) = 0.075$.

(c) Their unavailability for any reason: dependent through the feeder, a shared component.

(d) Unit A's failure during unit B's repair: dependent if load sharing raises unit A's rate, as in Exercise 6.2.

The most literal reading, whether the events F_A and F_B are independent, is (b), and its answer is no. The solution uses (b) and says so. The common wrong response answers (a), because the stated model says "independent within a window", and so reports conditional independence as independence.

Exercise 6.4. During an interruption both branches are fed by the UPS, so each keeps its power exactly when the UPS carries the load throughout. The two events are one event, perfectly dependent, and their joint is the probability of that one event:

$$P\big(F_{\text{ups}} \text{ does not occur during the interruption} \mid \text{an interruption begins}\big) = (1 - 0.02) \times P(D_{\text{sup}} \le t_{\text{aut}}),$$

which is 0.15 on the illustrative inputs, since an interruption of mean length one hour ends within ten minutes in about one case in six and the UPS picks up the load with probability 0.98. The product of the two branches' probabilities would square it, to 0.023, counting one event twice. The class is dependence ignored, in the form of a shared component left out: the UPS is shared by both branches.

Exercise 6.5. Given that unit B is found out for its own reasons, this is the probability that unit A is out for its own reasons too, on the model of disturbance windows. It is larger than $P(F_A)$, 0.150 against 0.075, because finding unit B out makes a disturbed window more probable, and a disturbed window raises unit A's chance as well.

Exercise 6.6. It is the probability that the UPS's supply is lost and that, while it is, the UPS is not carrying the load: the probability that the IT load is without supply because of a loss of supply, at a random hour. The second factor is the UPS's failure given the demand, not its unavailability.

Exercise 6.7. It is the probability that both units are out for their own reasons at a random hour, with the two independent within a window of either kind and the window's state unknown: the joint averaged over the shared cause.

Exercise 6.8. The study computes the probability that two independent branch faults overlap, and treats the two branches as if nothing tied them. The failure class is dependence ignored, in the form of a shared component left out: both branches hang from the UPS and its supply. The repair adds two cut sets. The dominant one is the shared supply, $F_{\text{sup}} \cap F_{\text{ups}}$: the supply lost and the UPS not covering it, which takes both branches at once, about 1.9×10^{-4} at a random hour on the illustrative inputs, roughly a hundred thousand times the study's value. The second is the propagation through protection: a fault on either branch that its breaker fails to clear opens the breaker above both branches until it is reset, at a rate of $2 \times 0.1 \times 0.01$ per year, each lasting about one hour, so about 2.3×10^{-7} at a random hour. The study's term remains, and it is negligible beside both.

Exercise 6.9. The statement assumes that the units' own failures are independent without conditioning on anything. On the stated model they are independent only given the window's

state. Unit A's failure raises the weight of a disturbed window, and with it unit B's chance of failing during the repair. The failure class is dependence ignored, in the form of a shared cause's tie dropped. The repair writes

$$P\big(\text{unit B fails in } [t_0, t_0 + D] \mid \text{unit A failed at } t_0\big),$$

given the stated model of failures and windows, with the window's state carried in the model and its weight updated by unit A's failure. If unit A's repair also loads unit B, that is a second tie, as in Exercise 6.2.

Exercise 6.10. The argument assumes that a trip of average length is the only kind of trip, and that the UPS always picks up the load. A spurious trip is a demand on the UPS, and the UPS's failure is conditional on it in two ways. It fails to pick up the load with probability 0.02 per demand, and it fails during its mission when the trip outlasts the battery, which for the fact sheet's exponential reclose times of mean five minutes (illustrative) happens with probability e^{-2}, about 0.135. The failure class is dependence ignored, in the form of a failure conditional on a demand taken at its unconditional probability, here zero; the argument also compares a mean with a horizon, a form Lesson 8 names. The repair writes

$$P(F_{\text{ups}} \mid \text{a spurious trip}) = 0.02 + 0.98 \times P(D_{\text{sup}} > t_{\text{aut}} \mid \text{a spurious trip}),$$

which is 0.15 on the illustrative inputs: about one trip in seven interrupts the IT load.

Exercise 6.11. Each item names the tie, or its absence, on the stated model.

(a) Independent on the stated model, the windows aside.

(b) The same event: both branches are fed by the UPS, a shared component.

(c) Independent on the stated model: the probability applies per demand.

(d) Independent given the window, which is the condition stated.

(e) Independent processes on the stated model; only their overlap matters, because it leaves the bypass with no source.

(f) One event, not two: with the opening adequate, both halls lose their cooling exactly on $F_{\text{fd}} \cup (F_A \cap F_B)$.

(g) Tied by the protection, through the IT branch's breaker failing to clear.

(h) Dependent: on the stated model only one fault acts at a time, so the two explanations of one mismatch are mutually exclusive, and excluding one another is the strongest dependence there is. Lesson 24 compares them.

(i) Independent counts in disjoint stretches of time, on the stated model of a constant rate.

(j) Tied by exposure: a longer repair leaves more time for unit B to fail, so the two are dependent even at constant rates. Load sharing, if it raises unit B's rate, is a second tie.

Lesson 7

Exercise 7.1. The sentence holds two questions. The first is about a parameter, unit A's failure rate, and its intention is to estimate:

$$P\big(\lambda_A < 2.5 \text{ per year in service} \mid \text{the record}\big),$$

conditioned on the record, with the handbook's value as prior information; on the two-hall plant E is kept for the 40 kW trip test. Whether the handbook is a prior and the record evidence,

and how the record's likelihood is written, are Lessons 12 and 11; here they are named. The second is about an outcome, and its intention is to predict:

$$P\big(N_A(\tau) \geq 3 \mid \text{unit A up at } t = 0,\ \text{the record}\big) = \int P\big(N_A(\tau) \geq 3 \mid \lambda_A\big)\, p(\lambda_A \mid \text{the record})\, \mathrm{d}\lambda_A,$$

with τ one calendar year and "more than twice" written as at least three, since the count takes whole values (Lesson 2). It averages the count's probability at each rate over what the record leaves known of the rate. The time base is stated: the rate per year in service, the count over the calendar year from the named start. Two wrong formulations are common: the count's probability at the record's own rate of seven in three years, which is a plug-in, and the parameter question reported as the answer to the second.

Exercise 7.2. The sentence holds two questions. Whether a cell is weak is a question about a discrete parameter, $P(H_{\text{weak}} \mid E)$, with E the vendor's evidence, which the statement does not give; the intention is to diagnose, since a cause is sought for the vendor's suspicion. The second is about the planned switch-off, which is a setting, do(grid off for 1.5 h), of fixed length. The string carries the load through it when its ride-through is at least 1.5 h:

$$\begin{aligned} P\big(T_{\text{rt}} \geq 1.5 \text{ h} \mid \text{do(grid off for 1.5 h)},\ E\big) &= P(H_{\text{sound}} \mid E)\, P(\varphi \times 2.025 \text{ h} \geq 1.5 \text{ h} \mid H_{\text{sound}},\ E) \\ &\quad + P(H_{\text{weak}} \mid E)\, P(0.70\,\varphi \times 2.025 \text{ h} \geq 1.5 \text{ h} \mid H_{\text{weak}},\ E). \end{aligned}$$

On the fact sheet's range of health the first conditional is one and the second zero, since a sound string lasts at least 1.72 h and a string with a weak cell at most 1.42 h, so the answer is $P(H_{\text{sound}} \mid E)$: the outcome is decided by the parameters alone, and the weak cell decides it. On the prior alone this is 0.90. Two wrong formulations are common: the planned switch-off written as a found outage, and the weak cell's probability read as the probability of failing the switch-off without asking what the string's health does under each hypothesis.

Exercise 7.3. The line is intentionally ambiguous, and it is not ready to fill: the word reliability names no quantity. Missing are whether a property of the string or an outcome is meant; if an outcome, over what horizon and from what start; and the level. Four readings give four answers.

(a) The specification's parameter question, $P(T_{\text{rt}} \geq 2 \text{ h} \mid E)$, 0.074 on the illustrative inputs.

(b) The next outage, $P(D_{\text{out}} \leq T_{\text{rt}} \mid E)$, 0.834.

(c) Every outage of the coming year carried, from steady operation at $t = 0$: a predictive over the year's count of outages and their lengths, in which the string's health and its weak cell are shared by every outage, so that the outages' outcomes are tied by a shared unknown (Lesson 6) and the probability is not a power of one outage's.

(d) The most literal reading, the string's own reliability as a component, its probability of failing as equipment, which the fact sheet does not state at all.

Each is compared with a level the insurer must supply. The solution uses reading (c), since an insurer pays for every outage the string does not carry, and says so. The common wrong response gives (a) or (b) alone, whichever looks better.

Exercise 7.4. The sentence holds two questions. Whether the UPS is degraded is a question about a discrete unknown, the network's three explanations of the mismatch (nominal operation, a degraded UPS, an offset in the PDU meter), with stated prior weights of 0.90, 0.05 and 0.05: $P(H_{\text{ups}} \mid E)$, with E the five readings. The extra energy is an expectation over the coming month,

$$\mathbb{E}\Big[\int_{\text{month}} \Big(\frac{P_{\text{load}}}{\eta} - \frac{P_{\text{load}}}{0.96}\Big)\mathrm{d}t \;\Big|\; E\Big],$$

with η the UPS's efficiency and P_{load} the load it carries. The fact sheet (Table D.2) states what each explanation assumes of η: under the nominal explanation and under the meter offset it is 0.96 ± 0.004, held near nominal; under a degraded UPS it is 0.94 ± 0.04, uncertain. The expectation therefore averages over the explanations and, within each, over η: the sum over i of the expected extra energy given H_i and E, times $P(H_i \mid E)$. The owner's "if it is" asks for the conditional, the expected extra energy given H_{ups} and E, a different expression. The solution writes both and says that the conditional is what the owner said and the unconditional what a budget needs. The common wrong formulation is the extra energy at the efficiency of the case study's degraded scenario, 0.91, which is the truth that scenario's readings were made from and not a value the explanation assumes: a plug-in under one hypothesis.

Exercise 7.5. Given what the evidence says about unit A's failure rate, it is the probability that unit A fails at least once in the coming year, averaged over what is known of the rate. It is not "the probability that the rate is high".

Exercise 7.6. It is the probability that the cooling is delivered at a random hour, given the rates and given what the evidence says about which drawing of the opening is right, averaged over the two drawings. It is not $P(H_{\text{AND}} \mid E)$.

Exercise 7.7. Given what is known, it is the probability that the string's state of health is at least 0.95 of a nameplate string's. One outcome that depends on it is the next grid outage being carried to its end, whose probability averages over the state of health.

Exercise 7.8. The range is the plug-in probability of carrying the next outage evaluated across the spread of the state of health: a band on the parameter, carried through a function of it. The failure class is parameter against outcome, in the form of a band on a parameter read as a band on the outcome's probability. The repair is one number,

$$P(D_{\text{out}} \le T_{\text{rt}} \mid E) = \int P(D_{\text{out}} \le T_{\text{rt}} \mid \varphi)\, p(\varphi \mid E)\, \mathrm{d}\varphi,$$

with the weak-cell hypothesis kept in the model unless the test rules it out. The spread of φ remains a statement about the string.

Exercise 7.9. The 0.99 is the posterior probability of a hypothesis, $P(H_{\text{nominal}} \mid E)$, on one hour's readings. The failure class is parameter against outcome, in the form of a hypothesis's probability read as the probability of the outcome it favors. The repair separates two questions. If the efficiency is held fixed over the month, "the efficiency is near 0.96" has a probability that adds to $P(H_{\text{nominal}} \mid E)$ the weight of the other explanations under which the efficiency is also near 0.96, as it is under the meter offset. If the efficiency may drift during the month, the question is a prediction and needs a model of the drift, which the fact sheet does not state; the report must say so.

Exercise 7.10. The analyst computed a plug-in: the probability of at least one failure for a unit whose rate is known to equal the posterior mean. The failure class is parameter against outcome, in the form of a plug-in reported as the predictive. The repair writes

$$\int \left(1 - e^{-\lambda_B \times 1\ \text{yr}}\right) p(\lambda_B \mid \text{the record})\, \mathrm{d}\lambda_B,$$

with unit B up at $t = 0$ and the coming calendar year as the horizon. From an up start, the time to the first failure runs on calendar time, so a rate per year in service applies to the whole year with no correction. The posterior here is given the plant's record, its own evidence on the rate. The averaged value is smaller than the plug-in, because $1 - e^{-\lambda_B}$ bends downward as the rate grows, and the solution says so without computing it.

Exercise 7.11. Each item names the target and its expression.

(a) A parameter: the probability, given the plant's repair record, that unit B's mean repair time exceeds 168 h.

(b) An outcome: the probability that unit B's next repair lasts more than 168 h, averaged over what is known of the mean.

(c) A parameter: $P(\eta < 0.94 \mid E)$.

(d) An outcome: a predictive over the efficiency, the load and the meter's noise. A reading not yet taken is evidence not yet collected, the $\tilde{E}$ of Lesson 1's table.

(e) A parameter, discrete: $P(H_{\mathrm{AND}} \mid E)$ against $P(H_{\mathrm{OR}} \mid E)$.

(f) An outcome: the predictive distribution of the number of interruptions over the coming calendar year from steady operation at $t = 0$.

(g) A parameter under one explanation: the probability that the offset exceeds 10 kW given the meter-offset explanation, under which the fact sheet puts the offset at 0 ± 45 kW, weighed by that explanation's probability; under the other two the offset is 0 ± 1.5 kW and 10 kW is out of reach.

(h) An outcome: $P(F_{\mathrm{ups}}$ does not occur $\mid$ an interruption begins, $E)$.

(i) A parameter: $P(G > 9000\,\mathrm{W/K} \mid E)$.

(j) An outcome on a condition, settled by the parameter once the plant has settled:

$$P\big(Q_B/G \le \Delta T_{\mathrm{allow}} \mid \text{unit B out},\ Q_B = 50\,\mathrm{kW},\ E\big).$$

Lesson 8

Exercise 8.1. Unit B's trip is a found event, so "unit B out at 02:00" stands behind the bar as a present state, and the crew's arrival fixes the horizon at 06:00. The two expressions are

$$P\big(S(06{:}00) = 1 \mid \text{unit B out at } 02{:}00,\ E\big)$$

for the moment just before the crew restores the unit and

$$P\big(S(t) = 1 \text{ for every } t \in [02{:}00,\ 06{:}00] \mid \text{unit B out at } 02{:}00,\ E\big)$$

for the path. They differ by the paths on which the cooling is lost and regained within the four hours: unit A or the feeder failing and returning by 06:00. With a unit's mean repair of 180 h and the feeder's of 24 h on the stated rates, such a return within four hours is rare, so on this plant the two nearly coincide. The path recovers too slowly for the moment and the path to part, and the solution says so. The cooling is defined with hall B at its design load, and only at that load does it come to $G \ge G^{*}$ with unit B out; both expressions then carry that condition and therefore what is known of G (Lesson 7). The common wrong formulation answers either question with the steady availability, which drops the present state.

Exercise 8.2. The question is ready: the sentence holds two questions on the battery string, and the second is a race. The grid's loss at $t_0 = 14{:}00$ is a present state behind the bar. The load has power at 15:30 if the grid is back by then or the string has not yet reached its floor:

$$P\big(D_{\mathrm{out}} \le 1.5\ \mathrm{h} \text{ or } T_{\mathrm{rt}} > 1.5\ \mathrm{h} \mid \text{the grid lost at } t_0,\ E\big).$$

The string reaches its floor before the grid returns on the event $T_{\mathrm{rt}} < D_{\mathrm{out}}$, a race between two first passages,

$$P(T_{\mathrm{rt}} < D_{\mathrm{out}} \mid \text{the grid lost at } t_0,\ E),$$

the complement of Lesson 7's next outage carried. Both average over the state of health and the weak-cell hypothesis. The common wrong formulation reads "power at 15:30" as "power throughout until 15:30", which is $P(T_{\mathrm{rt}} > \min(D_{\mathrm{out}}, 1.5\ \mathrm{h}) \mid \ldots)$. The two differ by the paths on which the string reaches its floor and the grid returns afterwards but before 15:30, $T_{\mathrm{rt}} < D_{\mathrm{out}} \leq 1.5$ h: the load is lost for a while and carried again at 15:30. On the fact sheet's inputs that needs a weak cell, since a sound string lasts at least 1.72 h, and the solution says so: the forms part only where the path can recover within the horizon.

Exercise 8.3. This question is intentionally ambiguous, and it is not ready. Missing are what the UPS must last against (an interruption, its own next failure, or its service life), from what start, and at what load. Four readings give four answers.

(a) The battery's autonomy at full load, $t_{\mathrm{aut}} = 10$ min, a stated parameter with no probability.

(b) The time to the UPS's next failure from steady operation, a first-passage time with distribution $P(T_{\mathrm{UPS}} \leq \tau \mid \text{steady operation at } t_0,\ E)$; its mean is five years on the illustrative rate of 0.2 failures a year.

(c) The time until the IT load is next lost, the first passage of the count of its interruptions to one, at 1.88 interruptions a year on the illustrative inputs.

(d) The most literal reading, the UPS's service life, which the fact sheet does not state at all.

The solution uses reading (c) for a customer, whose interest is the load, and says so. The common wrong response is reading (a), the one number on the data sheet.

Exercise 8.4. The rising path is the year's cumulative hours without supply of the IT load, $L_{\mathrm{IT}}(t) = \int_{t_0}^{t} (1 - S_{\mathrm{IT}}(s))\,\mathrm{d}s$, from the start of the calendar year. The clause's event is $L_{\mathrm{IT}}(\tau) \leq 0.5$ h at the year's end, which is the same, up to an event of probability zero, as $T_{0.5\,\mathrm{h}} > \tau$, since the total only rises and T counts the level as reached when the total is at it. The expression is

$$P\big(L_{\mathrm{IT}}(\tau) \leq 0.5\ \mathrm{h} \mid \text{steady operation at the year's start},\ E\big),$$

with τ the calendar year. The level is missing. The model needs the distribution of each interruption's length and not only their rate, because the total is a sum of lengths. Two wrong formulations are common: the expected hours without supply, 2.1 a year on the illustrative inputs, compared with 0.5 h, which compares a mean with a level set for each year's total (Lesson 4's class); and "at most one interruption", which counts breaks instead of minutes.

Exercise 8.5. Starting with unit A up, it is the probability that unit A fails at least once in the next three calendar months. The same event is $T_A \leq t_0 + 3$ months, with T_A the time of unit A's first failure after t_0.

Exercise 8.6. With unit B out from t_0, it is the probability that at some moment in the next six hours hall B runs more than 8 K above hall A, including the hours before the plant settles. It is not the probability that hall B is more than 8 K above hall A after six hours.

Exercise 8.7. With the feeder interrupted at t_0 and the UPS up, it is the probability that the IT load has power one hour later, whether or not it lost its power in between.

Exercise 8.8. The planner compares a horizon with the mean of the time to failure. The failure class is "at a time" read as "by a time", in the form of a first-passage time's distribution summarized by its mean and compared with a horizon. The repair writes

$$P\big(N_A(\tau) = 0 \mid \text{unit A up at the summer's start}\big),$$

given the stated model, with τ the three calendar months of the summer. From an up start the time to the first failure runs on calendar time, so on the stated rate of 2.5 failures per year in service, taken as fixed, it is $e^{-2.5/4}$, about 0.535: close to an even chance, not almost certain.

Exercise 8.9. The log supports the pack's temperature at one moment, after the string has rested and cooled. The failure class is "at a time" read as "by a time", in the form of a crossing written on the value at the horizon of a quantity that falls again. The repair is the running maximum over the 24 hours,

$$M(t_0 + 24\ \text{h}) = \max_{s \in [t_0,\, t_0 + 24\ \text{h}]} \vartheta_{\text{pk}}(s) < 318\,\text{K},$$

which needs the temperature through the outage, when it rose; the log at 14:00 says nothing about the peak. The source's simulation puts the peak of a three-hour outage at 307.2 K, and the solution adds that the record of this outage, not the model, must show it.

Exercise 8.10. The checks establish the state at two moments, not the path between them. The battery lasts ten minutes, so the racks may have lost their power after 02:20 and regained it when the feeder returned before 02:45. The failure class is "at a time" read as "by a time", in the form of the state at two times reported for "throughout". The repair is a question about a past path,

$$P\big(S_{\text{IT}}(t) = 0 \text{ at some } t \in [\text{02:15},\ \text{02:45}] \mid \text{powered at 02:15 and at 02:45},\ E\big),$$

which on the illustrative inputs is about 0.65: given both checks, a loss between them is more likely than not. Its intention is to estimate, since the path has already run. The UPS's event log, if it records transfers to and from the battery, would decide it, and the report should cite the log rather than the checks.

Exercise 8.11. Each item names the form and the expression.

(a) At a stated time, from the named state now, t_0: $P(S_{\text{IT}}(\text{noon}) = 1 \mid \text{the state at } t_0,\ E)$; with noon a day away the answer is close to the steady value, which the solution may use and say so.

(b) A race: $P(T_{\text{rt}} > D_{\text{out}} \mid \text{the grid lost at } t_0,\ E)$.

(c) At a random time, read as a share: the probability, at a random hour, that the feeder is up and exactly one unit is out; or the share of the year's hours in that state, a random quantity like Lesson 4's interval availability.

(d) By a time: $P(N_{\text{ut}}(\tau) \geq 1 \mid \text{the state at } t_0,\ E)$, with t_0 now and τ the rest of the month, the count on calendar time.

(e) A present state, an estimate: $P(S_{\text{IT}}(\text{now}) = 1 \mid E)$, answered from what is known now.

(f) A first-passage time: the distribution of T_A from the present state.

(g) By a time, on a rising path: the probability that L_{IT} at the end of June is at least 0.5 h, from steady operation at the year's start, which is the first passage of L_{IT} to half an hour falling before the end of June.

(h) Throughout the interruption: the probability that the UPS picks up the load and the interruption ends within t_{aut}, given that an interruption begins.

(i) By a time, on a past path: the running maximum of the pack's recorded temperature over the outage, against 310 K; an estimate if the readings carry error, and a value read from the record if they do not.

(j) The mean of a first-passage time: $\mathbb{E}[T \mid \text{steady operation at the year's start},\ E]$, with T the first interruption.

Lesson 9

Exercise 9.1. The paragraph holds three clauses, and the comparison finds two of them tied. The first is a parameter question, two hours at 20 kW,

$$P(T_{\mathrm{rt}} \geq 2 \text{ h} \mid E),$$

0.074 on the fact sheet's illustrative inputs. The second, "without the pack exceeding 318 K", is a second limit state on the same two-hour discharge, written on the running maximum of the pack's temperature over the discharge. With the first, it forms one requirement with two limit states, which fails on the union of their failure domains (Lesson 5); on the stated model the energy limit binds first and the pack peaks far below 318 K, so the union's probability is the first's. The third is about every outage of the first year, from steady operation at $t = 0$:

$$P(\text{every grid outage in } [0, \tau] \text{ ends before the string's floor} \mid \text{steady operation at } t = 0,\ E),$$

a predictive over the year's count of outages and their lengths. The string's health and its weak cell are shared by every outage, so the outages' outcomes are tied through those parameters (a shared cause), and the probability is not a power of one outage's. The paragraph says "each", so the 90 % is a level for each requirement separately. Two wrong formulations are common: the third clause read as the next outage, which drops the horizon; and the first two read as independent requirements and multiplied, when one limit state almost always decides both.

Exercise 9.2. The brief holds three clauses, which become four expressions. "No more than once a year" is a count from steady operation at the start of the calendar year:

$$P(N_{\mathrm{IT}}(\tau) \leq 1 \mid \text{steady operation at } t = 0,\ E),$$

0.44 on the illustrative inputs. "No loss of power shall last more than ten minutes" quantifies over the year's losses and needs each loss's length. The efficiency clause is a parameter question, $P(\eta \geq 0.95 \mid E)$. Its gloss, losses at full load at most 60 kW, is a second expression: at the full 1150 kW the UPS's loss is $1150\,\mathrm{kW} \times (1/\eta - 1)$, which is at most 60 kW exactly when $\eta \geq 1150/1210$, about 0.9504. No two clauses name one quantity: the two efficiency clauses are nested events on one parameter, with nearly the same threshold, and the loss clause is the stricter, and the joint of the two is the loss clause. Even the nesting holds only because the load is fixed at full; the brief's "which is to say" asserts an equality that is not one. Two wrong formulations are common: the two efficiency clauses treated as two requirements and multiplied, which counts the stricter clause's event twice; and the first two clauses folded into "at most one short loss a year", which drops one of them.

Exercise 9.3. This statement is intentionally ambiguous, and it is not ready. Missing are which events are meant (the two-hall plant's cooling service, the network's IT supply, or both halls' cooling), the horizon, the level, and what "at the same time" adds. In this part of the book the two-hall plant and the network are separate systems with separate feeders, and the solution says so. Four readings give four expressions.

(a) The most literal: both delivered at a random hour, $P(S(t) = 1 \text{ and } S_{\mathrm{IT}}(t) = 1 \mid E)$, a product if the systems are independent on their stated models: 0.971×0.99976 on this part's inputs.

(b) Both delivered at every hour of the first year, a joint over two paths from steady operation at $t = 0$.

(c) Each reliable separately at its own level: two expressions, not one.

(d) One conditional on the other, the cooling given that the power is lost, which on separate systems equals the unconditional.

The solution uses reading (b) for a handover and says so; Lesson 25 couples the two systems and returns to the statement. The common wrong response is one expression on one system.

Exercise 9.4. Unit A's service is a setting, do(unit A off at t_0 for 6 h) (Lesson 3). "If the feeder stays up" is a condition on a future path, and it has two places. Read as a condition, it goes behind the bar as the feeder up throughout $[t_0, t_0 + 6\ \text{h}]$, and the expression answers a question about a world in which the feeder holds. Read as part of the requirement, it enters the event as an exclusion, "hall A within its limit, or the feeder out", an "unless". The solution writes both and asks which is meant. "Hall A at 50 kW" is a condition on the load. "Within its limit" needs a limit for hall A above hall B, which the fact sheet does not state, as Exercise 5.1 found; and, as there, unit B must carry both halls within its 130 kW capacity, so the expression also needs hall B's load, which the statement leaves out. The last clause is a procedure, not a probability question: it says what to do on the event that unit B fails during the service, whose probability,

$$P\big(N_B(t_0 + 6\ \text{h}) - N_B(t_0) \geq 1 \mid \text{do(unit A off at } t_0),\ E\big),$$

is what the planner needs in order to plan for it. Two wrong formulations are common: the feeder's condition dropped, and the shutdown clause written as a probability of shutting down in order.

Exercise 9.5. Given the 40 kW trip test, it is the probability that the opening is adequate and that hall B, carrying its full load with unit B out, once the plant has settled, is within 8 K of hall A. The two events are one event, so the probability equals $P(G \geq G^* \mid E)$, 0.82.

Exercise 9.6. It is the probability that hall B goes over its limit, once settled, the next time unit B trips, at whatever load hall B carries then. The second fixes the load at the full 60 kW; the first averages over the load, so it is smaller whenever the load can be below full.

Exercise 9.7. Starting from steady operation, it is the probability that the string meets its two-hour specification and also carries every grid outage of the coming year. It is a conjunction of a parameter clause and an outcome clause that share the string's health, so its probability is not in general the product of theirs.

Exercise 9.8. The figure is the probability of carrying the next outage, $P(D_{\text{out}} \leq T_{\text{rt}} \mid E) = 0.834$ on the illustrative inputs, reported for both questions. The failure class is several questions folded into one expression, in the form of one expression answered for all. The repair gives two expressions and two answers: $P(T_{\text{rt}} \geq 2\ \text{h} \mid E) = 0.074$ for the specification, and 0.834 for the next outage, side by side with the reason they differ, which is that most outages are short. The word "reliable" names neither, and the reply should not use it.

Exercise 9.9. A horizon is attached to a parameter clause, and the note's own number refutes its claim: 0.074 is far below the 95 % it asserts. On the fact sheet the string's health has one value and does not change within a year, so "over the first year" adds nothing to the number and suggests to a reader that a path was considered. The failure class is several questions folded into one expression, in the form of a condition with the wrong scope. The repair is either $P(T_{\text{rt}} \geq 2\ \text{h} \mid E)$ at commissioning, with the horizon dropped and the note saying why; or, if the author meant that the string must still give two hours at the year's end, a model of the string's ageing, which the fact sheet does not state and which the note must ask for.

Exercise 9.10. The number answers whether both events hold, on two separate systems taken as independent, which the owner did not ask. The failure class is several questions folded into

one expression. The repair gives two expressions and two answers: $P(G \geq G^{*} \mid E) = 0.82$ for the opening on the 40 kW trip test, and, for the racks, the probability that the next interruption reaches them, $1 - 0.15 = 0.85$, a prediction on the network's illustrative inputs. It also notes that the owner's second question is the complement of "covered", which the analyst's number had silently turned around.

Exercise 9.11. Each item names the relation and its reason.

(a) Different: a state and a reading of it (Lesson 3's "the known item is the reading").

(b) One event: the string reaches its floor before the grid returns exactly when the outage is not carried.

(c) One event on a count: at least two is more than one.

(d) One event, at that fixed load: with unit B out, the rise is 52 kW over the conductance.

(e) Nested: a failure to pick up the load implies a loss of power, which can also come from a battery that runs out.

(f) Different: an explanation and a reading that is evidence for it.

(g) Different: the stated limit covers hall B above hall A only, so the second has no stated threshold.

(h) Overlapping: neither implies the other, since branch faults interrupt the load with the feeder up and the UPS covers short interruptions.

(i) Different: a random horizon and a fixed one with the same mean.

(j) One event: on the parallel structure both are $F_{\mathrm{fd}} \cup (F_A \cap F_B)$.

Lesson 10

Exercise 10.1. The test load is a condition. The hall temperatures ϑ_A, ϑ_B are variables, and the logged readings are observations. The conductance is a parameter, and "the drawing's conductance" is not its value but information about it, a prior that Lesson 12 writes. Wall leakage is a discrepancy. The note says it is small, not zero, so it is either carried with a stated spread or dropped with a label. The relations are $Q_B = G(\vartheta_B - \vartheta_A) + \delta$ with unit B out and the plant settled, which the note leaves unsaid and the sheet must ask for, and reading = temperature + error. The common wrong formulation treats the conductance as stated and asks what the test confirms, which has no uncertain quantity left to ask about.

Exercise 10.2. This request is intentionally ambiguous, and it is not ready. "Better" has no referent: better against the exact value, against a report's number, or against the cost of a run. It names no test, and it does not say whether the model may change. Three readings give three different pieces of work. A new coordinate with the same model and the prior transformed leaves every exact probability unchanged and changes only a Gaussian reading, which is what the word usually means. A new prior is a new model, and its answers change for a reason the request should state. A new relation, such as carrying δ, is a new model too. The solution uses the first reading and says so, and it asks which reading the requester wants checked and on which test, since the rule's checklist can fail on a weak one.

Exercise 10.3.

Run, hall B at 50 kW, read in R

1	**Sentence**	Given E, unit B out, settled: probability hall B carries 50 kW within 8 K?
	Intention	Estimate
2	**Target**	A parameter, R, through an event on it
	Declare	$d = 4.7828\,\mathrm{K}$; $Q = 40\,\mathrm{kW}$; R [K/kW]; $Q_B = 50\,\mathrm{kW}$; $\Delta T_{\mathrm{allow}} = 8\,\mathrm{K}$; m_R, σ_R
3	**Known**	E; unit B out; settled; $\delta = 0$ (assumption)
	Random here	R only
4	**Event**	$Q_B R \le \Delta T_{\mathrm{allow}}$, that is $R \le 0.16$ K/kW
5	**Expression**	$P(R \le 0.16 \mid E)$
6	**Model**	$d = QR + \varepsilon$; the prior transformed; all four steps pass, as on E
	Reduce	$\Phi\big((0.16 - m_R)/\sigma_R\big)$, Gaussian reading in R
7	**Check**	Below the design load; one event on the opening
	Class	Posterior exceedance

$$\boxed{P(R \le 0.16 \mid E) \approx \Phi\big((0.16 - m_R)/\sigma_R\big)}$$

Exercise 10.4.

Run, the string at 1.75 h

1	**Sentence**	Timed discharge reads 1.80 h: probability the string gives at least 1.75 h?
	Intention	Estimate
2	**Target**	A parameter, φ, through T_{rt}
	Declare	$d = 1.80$ h; spread 0.03 h (illustrative); φ [–]; $T_{\mathrm{rt}} = 2.025\ \mathrm{h} \times \varphi$
3	**Known**	d; 20 kW; sound string (assumption); $\delta = 0$ (assumption)
	Random here	φ only
4	**Event**	$\varphi \ge 1.75/2.025 = 0.8642$
5	**Expression**	$P(\varphi \ge 0.8642 \mid d,\ \text{sound})$
6	**Model**	$d = 2.025\ \mathrm{h} \times \varphi + \varepsilon$; uniform on 0.85 to 1.00; 2.6 spreads: step 3 fails
	Reduce	The normal about 1.80/2.025, spread 0.03/2.025, cut at the bounds
7	**Check**	Here the lower bound is the near one; silent on the weak cell
	Class	Posterior exceedance

$$\boxed{P(\varphi \ge 0.8642 \mid d,\ \text{sound})}$$

The exact value is 0.956 and the uncut normal gives 0.952, a small difference here; the checklist writes the exact expression all the same.

Exercise 10.5. With the load known, the meter reads $y = P_{\mathrm{load}}/\eta + \varepsilon$. This is nonlinear in η and linear in $1/\eta$, so the rule picks $1/\eta$. Under the degraded explanation the fact sheet puts η at 0.94 ± 0.04, which states no shape, so the analyst supplies one and labels it. That prior is transformed by its change of variable, with the factor η^2, and it is not written as a normal on $1/\eta$ with center $1/0.94$ and the same spread. The Gaussian reading then runs the checklist. A physical bound counts at step 3 though the sheet's prior does not state it: $\eta \le 1$ is $1/\eta \ge 1$. The meter's 3 kW on a load of 1150 kW is a spread of 0.0026 in $1/\eta$, and a reading near 1200 kW lies about 17 spreads from that bound. The change in the prior across a spread is 0.022, so all four steps pass, and the reading may be used, in that coordinate. The common wrong formulation keeps η as the coordinate because it is the one in the fact sheet, which is a choice the sheet should then state.

Exercise 10.6. The total is $D_{\mathrm{tot}} = k\mu$ for k repairs at a mean of μ, linear in μ. In the repair rate $\nu_{\mathrm{r}} = 1/\mu$ it is $D_{\mathrm{tot}} = k/\nu_{\mathrm{r}}$, nonlinear, so the rule picks μ for this relation. A prior on $\log \mu$ moves to μ with the factor $1/\mu$, and to ν_{r} with the factor $1/\nu_{\mathrm{r}}$. The relation as written treats the mean as if the total held no scatter; a record's total is itself a sum of lengths, whose likelihood Lesson 11 writes.

Exercise 10.7.

Run, the opening's thermal resistance	
1 Sentence	Given E: what is known of the opening's resistance?
Intention	Estimate
2 Target	A parameter, R, its whole distribution
Declare	R [K/kW]; d [K]; σ_d [K]; $\sigma_0 = \log 3$
3 Known	E; unit B out; settled; $\delta = 0$ (assumption)
Random here	R only
4 Event	None: a density, not an event
5 Expression	$p(R \mid E)$
6 Model	$d = QR + \varepsilon$; the prior transformed
7 Check	The question is ready; a summary of it names its coordinate
Class	Posterior of a parameter

$$\boxed{p(R \mid E) \propto \mathcal{N}(d;\, QR,\, \sigma_d^2)\, R^{-1}\, \mathcal{N}(\log R;\, \log R^*,\, \sigma_0^2)}$$

Exercise 10.8. This ticket is ready, whatever it may look like. "The right coordinate" is what Principle 10.1 picks and "the trip test" is E, so nothing is left to ask.

Run, the ticket, read with the lesson's rule	
1 Sentence	Given E, unit B out, settled: is the opening adequate, in the rule's coordinate?
Intention	Estimate
2 Target	A parameter, R, through an event on it
Declare	R [K/kW]; R^*; $d = 4.7828$ K; $Q = 40$ kW; m_R, σ_R
3 Known	E; unit B out; settled; $\delta = 0$ (assumption)
Random here	R only
4 Event	$R \le R^*$
5 Expression	$P(R \le R^* \mid E)$
6 Model	$d = QR + \varepsilon$; the prior transformed; 6.7 spreads, slope 0.135: all four pass
Reduce	$\Phi\big((R^* - m_R)/\sigma_R\big)$, Gaussian reading in R, $m_R = 0.1172$, $\sigma_R = 0.0178$ K/kW
7 Check	"Right" is the checklist's outcome, stated with its numbers; 0.818, and exactly 0.818
Class	Posterior exceedance

$$\boxed{P(R \le R^* \mid E) \approx \Phi\big((R^* - m_R)/\sigma_R\big)}$$

Exercise 10.9. The two meter readings are observations. The input and the output are variables, each following from the load and the efficiency, or the load a parameter of the snapshot, as the question takes it. The efficiency is a parameter; "0.96 from the nameplate" is information about it unless the note means it as stated, and the note must say which, since a stated efficiency leaves nothing to ask about the loss. The loss is a variable, $P_{\text{load}}(1/\eta - 1)$, with a relation that is linear in $1/\eta$ and in P_{load} separately and not in the two together. "The PDU meter's reading looks high" is a remark about another observation, and on the fact sheet it would raise the offset explanation; it has no role in the loss until that is asked. Assumptions to label: a steady hour, both meters read over the same interval, and no offset on the two meters. The common wrong formulation subtracts the two readings and calls the result the loss, which reads two readings as their quantities.

Exercise 10.10. With η_{d} the discharge efficiency, the string delivers 45 kWh $\times\, \varphi \times \eta_{\text{d}}$ to the load, so a timed reading at 20 kW is $y = 2.25$ h $\times\, \varphi\, \eta_{\text{d}} + \varepsilon$. It is linear in the product $\kappa = \varphi\, \eta_{\text{d}}$ and in neither factor when the other is unknown. One reading determines the product and leaves the two tied, so a prior on each stays in the answer, and whether a test can separate them is Lesson 18's question. The ride-through at the same load depends only on the product, so a question about the two hours can be written on κ; a question about the waste heat cannot, since that depends on η_{d} alone. The common wrong formulation keeps 0.90 as stated and gives all of the shortfall to φ.

Exercise 10.11. Given the 40 kW trip test, with unit B out, hall B at its full load and the plant settled, it is the probability that hall B runs more than 8 K above hall A. It is Lesson 9's quantity, written on R.

Exercise 10.12. It is what is known of the pack-to-room thermal resistance after one reading of the steady rise under the test heater, up to a constant, on a lognormal prior about 16.7 K/kW, the fact sheet's 60 W/K, with a log spread σ_0. The relation is linear in R_{p}, and the factor R_{p}^{-1} carries the prior from $\log R_{\mathrm{p}}$ to R_{p}. Without it the density would be another model's.

Exercise 10.13. A timed reading of the string is its nameplate ride-through times its state of health, taken down to 0.70 if a cell is weak, plus an error. Here y is an observation, φ a parameter, v a discrete parameter, 2.025 h and 0.70 stated, and ε the reading's error. The relation is linear in φ for each value of v, and not in v. With v unknown the prior on the product has two humps, so step 2 stops the checklist and the exact expression is written; with v stated, the prior is uniform and the checklist runs.

Exercise 10.14. It says how well each total load explains the sum of the IT and cooling meters' readings, with the two noises' variances added. It assumes that the two meters' errors are independent, and that neither has an offset, which the fact sheet's explanations do not put on these two meters.

Exercise 10.15. A spread does not invert as a value does: 1/0.0178 is not a spread of G, and its units are not even watts per kelvin. Near the mode a spread of 0.0178 K/kW in R is about 1300 W/K in G, and the distance to 7500 W/K is less than one such spread. The 8535 W/K is the reciprocal of the mode of the density in R, not a mode in G. The failure class is model conditioning dropped, in the form of a reading's coordinate left off. The repair stays in the reading's coordinate and turns the threshold over: $P(R \le R^{*} \mid E) \approx \Phi((R^{*} - m_R)/\sigma_R) = 0.818$, a Gaussian reading in R, and nowhere near beyond doubt.

Exercise 10.16. A normal on $1/\eta$ is a replaced prior and not a transformed one, whatever its center. The failure class is model conditioning dropped, in the form of a prior replaced in a change of variable and reported as the same model. Its spread is off as well: near 0.94 a spread of 0.04 in η is about 0.045 in $1/\eta$. That the shift is small is a claim to be shown. The fact sheet's 0.94 ± 0.04 also states no shape, and a normal on η would put 0.067 of its weight above 1. The repair chooses a shape for η and labels it, cuts it at 1, and writes the density in $1/\eta$ with its factor η^{2}; if a normal is used for convenience, it says so as an approximation and sizes it.

Exercise 10.17. The note counts from the wrong bound. The distance is to the prior's nearest bound, and for a reading of 1.95 h that is the upper one, 2.025 h, which is 2.5 spreads away. That fails step 3, so the checklist writes the exact expression and names the step. The failure class is model conditioning dropped, in the form of a rule of thumb without its condition. The repair is the normal about 1.95/2.025 cut at both bounds,

$$P(\varphi \ge 0.9877 \mid d,\ \text{sound}) = \frac{\Phi(a_1) - \Phi(a_{0.9877})}{\Phi(a_1) - \Phi(a_{0.85})}, \qquad a_x = \frac{x - 1.95/2.025}{0.03/2.025},$$

which is 0.042, against the note's 0.048. The size is small here, and the error is in which bound was read.

Exercise 10.18. An average of two readings is not a probability of anything. The two numbers are two approximations to one exact value, and the exact value is available: 0.818. The failure class is model conditioning dropped, in the form of a reading's coordinate left off, since 0.79 belongs to no coordinate at all. Averaging is no method: on the illustrative 5 kW test the

average, 0.759, happens to beat both readings, since they err in opposite directions, and on E it is worse than the reading in R, with no way to tell which case is at hand. The repair reports 0.818, or names one reading and its coordinate.

Exercise 10.19. There are two errors. The note reads one measured difference as the hall difference, so it takes the reading for its quantity and leaves out the 0.71 K spread on it; at that spread, the same subtraction could give several kilowatts of either sign. It also takes the drawing's conductance as the opening's, and with G unknown one reading cannot separate δ from it. The failure class is model conditioning dropped, in the form of a relation or a discrepancy assumption left off. The repair carries both unknowns, with $p(G, \delta \mid E) \propto \mathcal{N}(d;\, (Q - \delta)/G,\, \sigma_d^2)\, p(G)\, p(\delta)$, and it says that one reading leaves the two tied.

Exercise 10.20. Each item names the role, in that question.

(a) Stated.

(b) A condition: known, and set.

(c) A condition of that test, and part of the evidence; no role in tonight's load.

(d) A variable: it follows from the loads through the heat balance.

(e) A variable if no one has looked, with the unit's share of time out as its distribution; a condition once known.

(f) None: it belongs to another system.

(g) An observation, if the model has a relation for it; otherwise it has no role.

(h) None in the plant: a choice in the formulation, which the sheet states.

(i) A discrete parameter: one unknown fact, not something that varies from night to night.

(j) Stated, if the sheet takes it as given; otherwise a document about the parameter G, not its value.

Lesson 11

Exercise 11.1. The hall readings are $y_i = \vartheta_i + b_i + \bar{\varepsilon}_i$; the duty reading is the unit's duty times $1 + \varepsilon$, with ε of spread 0.02; each PDU reading is the hall's load times $1 + \varepsilon$, with spread 0.01. The offsets b_A and b_B are nuisance variables, and so are the loads if the question is about G alone. Hall A's temperature enters the difference only through the set point, and it is nuisance too. The likelihood of one test is the product of the densities of its five readings, the two halls, the duty and the two PDUs, with the temperatures written through the opening's law, and with the offsets integrated out or carried, as the sheet says. Written on the difference, it is four densities. The specification leaves three things unsaid, and the solution asks for each: "calibrated" names no spread for the offset; "within" does not say whether 2 % and 1 % are a noise, an offset or a band; and no noise per sample is given for the sensors. The common wrong formulation takes "calibrated" to mean that the offset is zero, which drops a nuisance variable with no label.

Exercise 11.2.

Run, the duty meter's likelihood	
1 Sentence	Duty meter reads 81.6 kW, noise 2 %: how probable is that, for each total load?
Intention	Estimate: the likelihood serves the posterior
2 Target	The likelihood $p(y \mid Q_\mathrm{t})$, a function of Q_t
Declare	y the duty reading [kW]; $Q_\mathrm{t} = Q_A + Q_B$ [kW]; noise proportional, 0.02
3 Known	y; unit B out; settled; the unit's duty equals the total (assumption)
Random here	y, through the noise; Q_t held at each value
4 Event	None: a density of a reading
5 Expression	$p(y \mid Q_\mathrm{t})$
6 Model	$y = Q_\mathrm{t}(1+\varepsilon)$, ε normal, spread 0.02
7 Check	A function of the total; it does not divide it between the halls
Class	Likelihood

$$\boxed{p(y \mid Q_\mathrm{t}) = \mathcal{N}(y;\, Q_\mathrm{t},\, (0.02\,Q_\mathrm{t})^2)}$$

Exercise 11.3. This line is ready. "The floor" is a term of the fact sheet: the 0.71 K that the sensors' offsets and noise put on the spread of a measured hall difference, which this lesson derives.

Run, the floor in the likelihood	
1 Sentence	With the fact sheet's floor as the spread of d: how probable is d, for each G?
Intention	Estimate
2 Target	The likelihood $p(d \mid G)$
Declare	d [K]; Q [kW]; $\sigma_d = 0.7113\,\mathrm{K}$, the floor
3 Known	Q; unit B out; settled; $\delta = 0$ (assumption)
Random here	d; G held at each value
4 Event	None; a density
5 Expression	$p(d \mid G)$
6 Model	Equation (11.3), the offsets integrated out, which is what the floor is
7 Check	A floor: no more samples go below it, and a calibration would change the model
Class	Likelihood, nuisance integrated out

$$\boxed{p(d \mid G) = \mathcal{N}(d;\, Q/G,\, (0.7113\,\mathrm{K})^2)}$$

Exercise 11.4. The relations are the sums of the lesson's second worked formulation. The explanation holds the efficiency at 0.96 ± 0.004 and the PDU offset at 0 ± 45 kW, and a question about η makes b_pdu and the loads the nuisance. The likelihood for the efficiency is

$$p(\mathbf{y} \mid \eta, H_\mathrm{off}) = \int p(\mathbf{y} \mid \eta, b_\mathrm{pdu}, P)\, p(b_\mathrm{pdu} \mid H_\mathrm{off})\, p(P)\, \mathrm{d}b_\mathrm{pdu}\, \mathrm{d}P.$$

The PDU meter's reading alone bears on η only through P_load, and with an offset of spread 45 kW it says little of the load; the UPS output meter says more. The common wrong formulation reads η off the ratio of the PDU and utility readings, which takes the offset as zero, under the one explanation that says it is not.

Exercise 11.5. The reading is of the start, not of the string's health. Written $y_\mathrm{soc} = x_0 + \varepsilon$ with spread 0.03, it is an observation about the state of charge x_0 at the outage's start. The usable energy above the floor, and so the ride-through, depends on x_0: $T_\mathrm{rt} = \varphi \times 2.025\ \mathrm{h} \times (x_0 - 0.2)/0.6$ for a sound string. The event is $D_\mathrm{out} \le T_\mathrm{rt}$, and x_0 is a nuisance variable with the reading as its evidence:

$$P(D_\mathrm{out} \le T_\mathrm{rt} \mid y_\mathrm{soc}) = \int P(D_\mathrm{out} \le T_\mathrm{rt} \mid \varphi, x_0)\, p(\varphi)\, p(x_0 \mid y_\mathrm{soc})\, \mathrm{d}\varphi\, \mathrm{d}x_0.$$

The density of x_0 given the reading needs a prior on x_0, which the fact sheet does not state, so the analyst supplies one and labels it. The weak cell is weighed in as in Lesson 7. The common wrong formulation reads a state of charge of 0.80 as a stated value and drops the term.

Exercise 11.6. This line is intentionally ambiguous, and it cannot enter the likelihood as it stands. It does not say which term of the sum it supplies, whether 0.5 % is the width of a standard deviation or of a hard band, or what "verified" did to the meter. Three readings give three terms, at the full 1150 kW. As a noise of spread 5.75 kW per reading, it replaces the fact sheet's 2 kW and averages. As a constant offset of spread 5.75 kW, it does not average, and it ties every reading of that meter. As a hard band, it gives a uniform offset of spread about 3.3 kW. "Verified at the last shutdown" may also mean that an offset was measured then, which leaves only the drift since. The analyst asks for the manufacturer's definition, the verification record and its date, and in the meantime carries the fact sheet's 2 kW with the line's reading labeled as open. The common wrong response treats 0.5 % as noise because noise is the term that averages down.

Exercise 11.7.

Run, what a count says of a repair time

1	**Sentence**	Given the feeder's one-year count, 4 failures: what is known of its mean repair?
	Intention	Estimate
2	**Target**	A parameter, μ_{fd}
	Declare	$k = 4$; $t_{\text{s}} = 1$ yr in service; μ_{fd} [h]; prior: $\log \mu_{\text{fd}}$, center 24 h, a factor of 3
3	**Known**	The count with its exposure, in service; no repair length logged
	Random here	μ_{fd}
4	**Event**	None; a density
5	**Expression**	$p(\mu_{\text{fd}} \mid k = 4)$
6	**Model**	A Poisson count in the rate alone: on the service clock μ_{fd} does not enter
7	**Check**	The answer is the prior. On a calendar clock the exposure would hold repairs
	Class	Posterior of a parameter

$$\boxed{p(\mu_{\text{fd}} \mid k = 4) = p(\mu_{\text{fd}})}$$

Exercise 11.8.

Run, the total load, from three meters

1	**Sentence**	UPS output and two branch meters: how probable, for each total load?
	Intention	Estimate: the likelihood serves the posterior
2	**Target**	The likelihood in P_{load}, the two branches' split being nuisance
	Declare	y_{ups}, y_{it}, y_{cl} [kW]; noises 2, 3, 3 kW; $P_{\text{load}} = P_{\text{it}} + P_{\text{cl}}$ [kW]
3	**Known**	Steady snapshot; no offsets here; independent noises; flat on the split (assumptions)
	Random here	The readings; the split, integrated out
4	**Event**	None: a density of three readings
5	**Expression**	$p(y_{\text{ups}}, y_{\text{it}} + y_{\text{cl}} \mid P_{\text{load}})$
6	**Model**	$y_{\text{ups}} = P_{\text{load}} + \varepsilon_2$; $y_{\text{it}} + y_{\text{cl}} = P_{\text{load}} + \varepsilon_4 + \varepsilon_5$, spread $\sqrt{18}$ kW
7	**Check**	Two independent looks at one load; the PDU and utility meters would add terms
	Class	Likelihood, nuisance integrated out

$$\boxed{\mathcal{N}(y_{\text{ups}};\, P_{\text{load}},\, (2\,\text{kW})^2)\,\mathcal{N}(y_{\text{it}} + y_{\text{cl}};\, P_{\text{load}},\, 18\ \text{kW}^2)}$$

Exercise 11.9. Pooled, the likelihood is $P(14 \mid \lambda) = (6\lambda)^{14} e^{-6\lambda}/14!$, a count with the summed exposure. It assumes that the two units have one rate, which their common design makes plausible and which is an assumption all the same. Dropped, the likelihood is a product, $P(k_A \mid \lambda_A)\, P(k_B \mid \lambda_B)$, with a prior on each rate, and the sheet needs the two counts separately, which the pooled line has thrown away. The pooling also changes the question, since a band on one common rate is narrower than a band on either unit's rate. The common wrong formulation pools the counts to tighten the band and then quotes the result for unit A.

Exercise 11.10. About the outages, the log supplies four lengths and an exposure of two years, on the calendar clock, since an outage does not wait for service. The count's likelihood is Poisson

in the arrival rate, $(2\nu_{\mathrm{o}})^4 e^{-2\nu_{\mathrm{o}}}/4!$, and the lengths' is $\mu_{\mathrm{o}}^{-4} e^{-4.2\ \mathrm{h}/\mu_{\mathrm{o}}}$ if they are exponential, which is illustrative. About the string, each "carried" is a reading of an event, and the likelihood of the ride-through is $\mathbf{1}[T_{\mathrm{rt}} \geq 1.6\ \mathrm{h}]$, the longest outage being the one that binds. That says T_{rt} is at least 1.6 h, so it rules out a weak cell on the fact sheet's ranges, whose highest value is 1.42 h. It leaves φ as it was, and removing the hypothesis moves the probability of the string's 2 h from 0.074 to 0.082. The log also leaves unsaid what the state of charge was at each start, which moves T_{rt}. The common wrong formulation averages the four lengths and calls that the ride-through.

Exercise 11.11. As a function of G, it says how well each conductance explains one reading from a 60 kW test: a likelihood, not a probability of G. As a function of d, it says how the reading of such a test would scatter at a given conductance: a sampling density, which is what Lesson 16 will ask about a reading not yet taken. One formula, two questions.

Exercise 11.12. It is the probability of four feeder failures in one year in service if the feeder's rate is λ_{fd} per year in service. The exposure is one year on the service clock. Read as a function of the rate, it is the likelihood of a one-year record holding four failures.

Exercise 11.13. It is the density of a meter's reading when the true power is P, the noise has a spread of 2 kW, and the meter's offset is unknown with a spread of 45 kW. With the offset integrated out it is normal about P, with the two variances added, so the noise hardly matters.

Exercise 11.14. It is what is known of the pick-up's failure probability after six demands that all succeeded, up to a constant, on a stated prior. The exposure is the six demands, not any span of time, and a quiet year adds nothing to it. A clean run leaves a range of values, so it is no reason to set p_{d} to zero.

Exercise 11.15. The sentence divides the whole spread by the root of the number of readings, and only the noise scales that way. The failure class is reading taken for the quantity, in the form of an offset read as noise. Under the meter-offset explanation the reading is $y_{\mathrm{pdu}} = P_{\mathrm{load}} + b_{\mathrm{pdu}} + \varepsilon$, with an offset of spread 45 kW that sixty readings share, so the hourly mean keeps a spread of about 45 kW, and the 0.26 kW is its noise alone. Nothing here rules the explanation out, since the offset is the very thing it allows. The repair writes the sum, and it compares the PDU reading with what carries no such offset, the UPS output and the sum of the two branches, under each explanation.

Exercise 11.16. There are two errors. A count of zero is read as a rate of zero, which is the class's count read as a rate: a zero has a likelihood too, $P(k = 0 \mid \lambda_{\mathrm{it}}) = e^{-\lambda_{\mathrm{it}} t_{\mathrm{s}}}$ with $t_{\mathrm{s}} = 2/3$ yr if the clock is service, and that leaves a range of rates, wide over so short an exposure. The review then removes a cut set from the structure because its probability looks small, which changes the event: the load's event still holds F_{it}, whatever its rate. The repair writes $p(\lambda_{\mathrm{it}} \mid \text{the log}) \propto e^{-\lambda_{\mathrm{it}} t_{\mathrm{s}}}\, p(\lambda_{\mathrm{it}})$, on a stated prior, and keeps the cut set with that rate behind it. The sheet asks whether the eight months are in service or on the calendar.

Exercise 11.17. There are two errors. A shared current transformer gives the two readings one offset, so the two meters are tied by a part they share, which is Lesson 6's class of dependence ignored. An offset also does not average at all, which is this lesson's class in the form of an offset read as noise. Averaging two readings with one shared offset leaves that offset whole. The repair writes both readings with a common term, $y_{\mathrm{ups}} = P_{\mathrm{load}} + b + \varepsilon_2$ and $y_{\mathrm{pdu}} = P_{\mathrm{load}} + b + \varepsilon_3$, which are not two independent looks at P_{load}. The fact sheet does not state that the meters share a transformer, so the note's premise is itself an assumption to label.

Exercise 11.18. The sentence replaces the whole spread by the part that has shrunk. The sum still holds the noise, so $\sigma_d^2 = (0.05\,\mathrm{K})^2 + 2(0.3\,\mathrm{K})^2/30$ and σ_d is 0.092 K, with the noise now the larger part. The failure class is reading taken for the quantity, in the form of a term of the sum dropped, here the noise where the lesson's first error dropped the offsets. The plan should also state how the bath's own error reaches the difference and how far the offsets may drift between yesterday and the test. The repair writes the sum and labels the drift.

Exercise 11.19. There are three errors. A difference of two meters reads the difference of their offsets, so two meters that share a transformer, or that drift together, can agree and both be off. A 1 kW agreement is inside the noise of that difference, whose spread is 2.83 kW, so it does not bound anything to 1 kW. And the load's spread is not the agreement: on one meter it carries that meter's noise and offset. The failure class is reading taken for the quantity. The repair writes $y_{\mathrm{pdu}} - y_{\mathrm{ups}} = b_{\mathrm{pdu}} + \varepsilon_3 - \varepsilon_2$, with any shared term named, and it writes the load's likelihood from all the meters under each explanation.

Exercise 11.20. Each item names the term.

(a) An offset, as a tolerance band: its form within the band, and the range it applies over, are needed.

(b) Noise.

(c) An exposure, on the service clock; a count is needed.

(d) An offset, which changes with time: a drift, and the times of the test and the calibration are needed.

(e) A count: its exposure and clock are needed.

(f) Not determined: an error on a condition, a noise or an offset, and the ± may be a band. Ask, and label the load as uncertain or as set.

(g) None: a condition, unless the setting's own error is in question.

(h) An offset proportional to the quantity.

(i) The exposure's start: its end is needed.

(j) None by itself: it gives the number of samples when the length of the test is known.

Lesson 12

Exercise 12.1. Four clauses, four outcomes. "Mean time between failures 3500 h in service, within a factor of two" is a document about units of this make, so a prior on the rate: a center of $8760/3500 = 2.5$ per year in service, in the coordinate $\log \lambda_A$, with "a factor of two" read by the book's default as one standard deviation of $\log 2$, labeled, and asked about, since a memo is not a sheet. The drawing is a document, so a prior on the as-built G, centered at 8250 W/K, which is 66 kW over 8 K, with the tolerance as its spread as a labeled choice and asked about in the same way, since "tolerance 2 %" may be a standard deviation or a band and says what the drawing permits, and under the assumption that the as-built opening is the drawn one. "Unit A failed twice in its first year here" is this very item in the period covered, so evidence: $k = 2$, $t_{\mathrm{s}} = 1$ yr in service, if the clock is service, which the memo should state. "The cooling shall hold 8 K with one unit out" is a requirement, which goes to step 4. The check finds one question to ask: whether the manufacturer's figure includes this plant. The common wrong formulation reads the mean time as the unit's rate, a document read as the fact.

Exercise 12.2. This note is ready. "The stated prior" is in the fact sheet: on the log of the feeder's rate, centered at 3.69 per year in service, with a spread of a factor of 3, which this lesson reads as one standard deviation.

Run, the stated prior on the feeder's rate

1	**Sentence**	The fact sheet's prior on the feeder's rate: write it.
	Intention	Estimate
2	**Target**	A distribution on λ_{fd}
	Declare	λ_{fd} [1/yr]; center 3.69; $\sigma_0 = \log 3$ (stated); coordinate $\log \lambda_{\text{fd}}$
3	**Known**	The fact sheet; this feeder of that kind (assumption)
	Random here	λ_{fd}
4	**Event**	None; a density
5	**Expression**	$\mathcal{N}(\log \lambda_{\text{fd}};\, \log 3.69,\, (\log 3)^2)$
6	**Model**	Lognormal, as stated
7	**Check**	A document, so a prior; the feeder's log stays out of it
	Class	Prior

$$\boxed{p(\log \lambda_{\text{fd}}) = \mathcal{N}(\log \lambda_{\text{fd}};\, \log 3.69,\, (\log 3)^2)}$$

Exercise 12.3.

Run, the lead's judgment as a prior

1	**Sentence**	66 kW give or take a factor of two, 90 % sure: what prior on G?
	Intention	Estimate: a prior serves the posterior
2	**Target**	A distribution on G
	Declare	G [W/K]; center 66 kW/8 K; a factor of 2 as a 90 % interval; coordinate $\log G$
3	**Known**	The lead's statement; who said it, and when, on the sheet
	Random here	G
4	**Event**	None: a density
5	**Expression**	$\mathcal{N}(\log G;\, \log 8250,\, (\log 2/1.645)^2)$
6	**Model**	Lognormal (choice); the interval symmetric in the log (assumption)
7	**Check**	A person, so a prior; is the 66 kW the drawing's, has the lead seen the test? Ask
	Class	Prior

$$\boxed{p(\log G) = \mathcal{N}(\log G;\, \log 8250,\, (\log 2/1.645)^2)}$$

Exercise 12.4. This ticket is intentionally ambiguous, and it is not ready. It does not say which rates, whether the sister plant is of the same design and duty, or how the history is to be used. Three readings give three models. The rule's reading: the sister plant's log is a prior on this plant's rate, with a spread for how the plants differ. A second reading replaces this plant's record with the sister plant's, which throws evidence away and answers a question about another plant. A third pools the two logs as if one plant, which takes the plants as identical; with a spread for their differences, that is a population model, which is named here and left to the literature. The solution uses the first and says so, and it asks for the sister plant's design and whether its log already feeds any handbook in use.

Exercise 12.5. The fleet is other items, so its 1 in 10 is a prior on this string. By the share's mean, with a flat prior on the share labeled as a choice, $P(H_{\text{weak}})$ is $21/202 = 0.104$, and the bare 0.10 is close to it on a list this long. The check then asks whether this site's string entered it. It did, so the site is taken out: with its status unknown, the list is 20 or 19 in 199, which gives 0.104 or 0.100, and if the string has been tested, that test is evidence and is used only once. The difference is small here, and the sheet still says so, since the form of the error is what recurs. The common wrong formulation uses 0.10 and any test of this string together without asking.

Exercise 12.6. It needs a spread and a coordinate, say $\log \lambda_{\text{ups}}$ normal about $\log 0.2$ with a stated spread, and a source, since the fact sheet marks the value as the book's illustrative input

and not a record. A two-year log of this UPS is this very item in the period the model covers, so it is evidence: a count with its exposure in service. The common wrong formulation replaces the 0.2 with the log's count over two years, a count read as a rate.

Exercise 12.7. The answer turns on what the sheet says of the four years, and either labeled reading scores. The fact sheet states no ageing, so it states neither.

Run, the acceptance test, on two readings of the gap		
1	**Sentence**	The string's 2.00 h test of four years ago: evidence about today's φ?
	Intention	Estimate
2	**Target**	Its place at step 3, under (a) φ constant; (b) φ ageing
	Declare	$d_0 = 2.00$ h; φ today's [–]; σ_t the test's spread, asked for; σ_a ageing's spread (choice)
3	**Known**	This very item, no stated change; the gap: (a) held constant; (b) not carried
	Random here	φ
4	**Event**	None; a factor of the joint
5	**Expression**	(a) $p(d_0 \mid \varphi)$; (b) $p(\varphi)$, centered near $d_0/2.025$
6	**Model**	(a) $d_0 = 2.025\ \mathrm{h} \times \varphi + \varepsilon$; (b) the same center, the spread σ_a, cut at the range
7	**Check**	(a) evidence; (b) a prior; the sheet states which, and uses the test once
	Class	(a) Likelihood; (b) Prior
		$\boxed{\text{(a) } \mathcal{N}(d_0;\, 2.025\,\varphi,\, \sigma_\mathrm{t}^2);\ \text{(b) } \mathcal{N}(\varphi;\, d_0/2.025,\, \sigma_\mathrm{a}^2)\, \mathbf{1}[0.85 \le \varphi \le 1]}$

Exercise 12.8.

Run, a prior on the pack's thermal mass		
1	**Sentence**	Data sheet: 3.0×10^6 J/K, "±10 %": what prior on the thermal mass?
	Intention	Estimate: a prior serves the posterior
2	**Target**	A distribution on C_p
	Declare	C_p [J/K]; center 3.0×10^6; 10 % read as a hard band (choice, asked about)
3	**Known**	A document about this pack's design; this pack of that design (assumption)
	Random here	C_p
4	**Event**	None: a density
5	**Expression**	$\mathbf{1}[2.7 \times 10^6 \le C_\mathrm{p} \le 3.3 \times 10^6]/(0.6 \times 10^6)$
6	**Model**	Uniform within the band (choice); any labeled reading of the 10 % scores
7	**Check**	A document, so a prior; the fact sheet states the value, and this sheet doubts it
	Class	Prior
		$\boxed{p(C_\mathrm{p}) = \mathbf{1}[2.7 \times 10^6 \le C_\mathrm{p} \le 3.3 \times 10^6]/(0.6 \times 10^6)}$

Exercise 12.9. It is in hand, and it comes from this very item. The model has no relation from an area to a conductance, so at question 2 it is a record with no relation to anything in the model, and the outcome is none of these; the sheet says so. A stated relation changes that: if G is written as a coefficient times the area, with the coefficient's spread stated, the area becomes a reading of the opening with its own error, and it is evidence. The common wrong formulation multiplies by a handbook's coefficient and treats the product as the conductance, which reads a document as the fact and drops the coefficient's spread.

Exercise 12.10. The factory test is of another string, so it is a prior on this one, with a spread for how strings of one design differ, and it asks whether the vendor's figures already hold this string. The on-site test before the replacement is of this string, in hand, with a stated physical change since, so it fails question 3 and is a prior too, moved and widened for what the new cell changed. The test since the replacement is this very item, unchanged, with a timed reading whose relation can be written, so it is evidence. The check then asks about overlap: if the first two tests shaped one prior, each is used once, and neither is also a source of the other. The common wrong formulation treats all three as readings of today's φ and multiplies their likelihoods.

Exercise 12.11. It is a prior that puts the feeder's mean repair near 24 h, with a factor of two as one standard deviation in the log.

Exercise 12.12. It is a uniform prior on the state of health from 0.90 to 1.00. A reading that points below 0.90 cannot move the posterior below the bound, whatever it says; a prior that gives a value no weight cannot be argued out of that by any evidence, which is a reason to widen it or to carry a hypothesis, such as the weak cell.

Exercise 12.13. It is how probable a count of seven was, on the prior alone, with the rate averaged over: the record's probability before the record is used. A very small value asks whether the prior was wrong for this unit or the record is odd, and it is a question to put to the sheet and not a reason to tune the prior after the fact.

Exercise 12.14. It is the prior on the string's usable energy, before any test: with probability 0.90 the string is sound and the share is uniform from 0.85 to 1.00, and with probability 0.10 it has a weak cell and the share is 0.70 of that, uniform from 0.595 to 0.70. The first term stands for a sound string and the second for a weak one, and the gap between them is why one reading can tell the two apart.

Exercise 12.15. The report sets a prior after seeing the reading it will then be updated with. The failure class is one record counted twice, in the form of a prior tuned after the evidence: the reading narrowed the range, and it then counts again in the likelihood. Nothing has been learned about the string beyond the one reading. The repair states the prior as it stood before the test, the fact sheet's 0.85 to 1.00, and updates it once. The range may be narrowed only on information other than this reading, such as a fleet record, and the sheet then says which.

Exercise 12.16. The sentence hides a choice of shape and a reading of a spread. On the fact sheet $\pm$ is a spread, so 0 ± 45 kW is one: read as a normal, its standard deviation is 45 kW; a uniform within ± 45 kW has a standard deviation of about 26 kW, so the two are different priors, and the second is narrower than the source. "No prior knowledge" is also false, since the bound and the width came from the fact sheet. The class is model conditioning dropped, in the form of a relation or an assumption left off. The repair states the shape as a choice and writes its width: normal with a standard deviation of 45 kW, or, if the uniform is kept for a reason, within ± 78 kW so that the spread matches, and it says which.

Exercise 12.17. There are two errors. Two documents that quote one source are one source, the class's fourth form: there is one 0.2, and its spread is not narrowed by a second citation. The memo then passes from a rate to "availability therefore known", which reads confidence about a parameter as knowledge of an outcome, Lesson 7's parameter against outcome; and an availability needs the repair time as well as the failure rate. The repair writes one prior on the rate, with a spread, cites the manufacturer, and states the predictive at a random hour with what it needs.

Exercise 12.18. There are three errors. A list of units now in service leaves out the ones retired, which are likely to be the ones that failed most, so the average is biased low; the question for the sheet is what the list omits. A fleet list holding this plant's two units has its overlap to remove, which is the lesson's class in the form of a source that contains the item. And the spread of a fleet's average says how well the average is known, not how far this unit may lie from it, which is what a prior's spread is for. The repair asks for the retired units, takes this plant's out, and writes a spread for how a unit differs from the fleet; a population model is named and left to the literature.

Exercise 12.19. There are two errors. A drawing's tolerance says what the drawing permits, and it is not what is believed of the as-built opening. Written as the prior with no label, as here, it reads a document as the fact, and the spread that belongs here is for how the as-built opening may differ from the drawing. And the claim reverses the logic: a prior that tight would move little whatever the test read, so a test that barely moves it confirms nothing. A prior that gives other values no weight cannot be argued out of its center. The repair uses a labeled spread for the as-built opening, or carries the drawings beside a wider hypothesis, and it then asks what the reading says on that prior. The second is a mixture, so Lesson 10's step 2 stops the checklist and the exact expression is written.

Exercise 12.20. Each item names the outcome and the question that decides it.

(a) Prior: a document about a model, and on an availability, which must be tied to a rate and a repair time.

(b) None of these: no relation from oil to failure is in the model. With one written, it would be evidence.

(c) Requirement.

(d) Overlap to remove: the report and the raw tests are one set of readings.

(e) Prior; the overlap's size is unknown, since it was made after reading the log: keep the log, build no prior from the estimate, and ask.

(f) None of these: no relation from an image to G is in the model. With one written, with its noise, it would be evidence.

(g) Prior: a person's judgment, with who said it.

(h) Prior: another item, with a spread for how the openings differ.

(i) None of these: not in hand, a reading not yet taken.

(j) None of these: the model has no quantity for the outside temperature. With a relation to the load added, ask again.

Lesson 13

Exercise 13.1. Three diagrams, none holding both the halls and the racks. The plant's factors are those of the lesson's first figure, with each unit's count factor added for its log. The network's and the string's are those of its second and third. "This hour" is a stated time, so its start must be named. "The halls and the racks both" is a conjunction of two systems, and the fact sheets state no tie between them: separate feeders, in this part of the book. Under independence, labeled as an assumption, the joint is a product. Whether that assumption holds is what cannot be written yet, since a real facility may share a supply, and Lesson 25 writes the coupling. The common wrong formulation multiplies the two availabilities and does not say that it assumed anything.

Exercise 13.2.

Run, the feeder up and at least one unit up	
1 Sentence	Stated rates, a random hour: is the feeder up and at least one unit up?
Intention	Predict
2 Target	An event on three states; not the service
Declare	F_A, F_B, F_{fd}; q a unit down; f the feeder down
3 Known	Steady; stated rates; components independent (assumption)
Random here	The three states at the hour
4 Event	$F_{\text{fd}}^c \cap (F_A \cap F_B)^c$
5 Expression	$P\big(F_{\text{fd}}^c \cap (F_A \cap F_B)^c\big)$, given the stated model
6 Model	Three component factors; no G, no structure
Reduce	$(1-f)(1-q^2)$
7 Check	The parallel structure's availability; silent on the opening, so not the service
Class	Predictive probability at a random time
	$\boxed{(1-f)(1-q^2)}$

Exercise 13.3. This sketch is intentionally ambiguous, and it is not ready. Readings split by horizon and level score as well as these, split by structure. "Reliability" is not a quantity, so the fourth box names no event; its horizon and level are missing as well. The structure is unstated, the opening is absent, and the arrows suggest a direction that a dependency does not have. Three readings give three structure functions. As the parallel structure, the event is $F_{\text{fd}}^c \cap (F_A \cap F_B)^c$. As the series structure, it is $F_{\text{fd}}^c \cap F_A^c \cap F_B^c$. As the service, the opening decides between the two, and a fourth box without G cannot be filled in at all. The solution uses the third, with the horizon asked for. The common wrong response multiplies three availabilities into the fourth box.

Exercise 13.4.

Run, the uncleared fault's cut set	
1 Sentence	Random hour, illustrative rates: IT load lost to an uncleared cooling fault?
Intention	Predict
2 Target	One cut set of the load's event
Declare	F_{clr} the breaker above open after an uncleared fault; λ_{cool} cooling faults [1/h]
3 Known	Steady; the structure; fault rate, clearing probability, 1 h reset (illustrative)
Random here	Whether the hour falls in such a reset; the clearing on demand
4 Event	F_{clr}
5 Expression	$P(F_{\text{clr}})$
6 Model	Rate, demand and reset: the 1 h sets the length, not the branch's 4 h repair
Reduce	$\lambda_{\text{cool}} \times 0.01 \times 1$ h
7 Check	One term of the load's event, not the load's event
Class	Predictive probability at a random time, with a demand
	$\boxed{P(F_{\text{clr}}) = \lambda_{\text{cool}} \times 0.01 \times 1 \text{ h}}$

Exercise 13.5. The product is $p(\varphi)\,P(H_{\text{weak}})\,\mathbf{1}[T_{\text{rt}} = \cdot]\,p(d \mid \varphi, H_{\text{weak}})\,p(D_{\text{out}})\,\mathbf{1}[\text{carried} = \cdot]$. The query on T_{rt} given d needs the two priors, the reading and the ride-through's deterministic factor, whose output is asked about. The event's factor has an output that is neither asked about, nor read, nor an input of a factor that stays, so it falls away, and then the outage's length, an input of nothing that stays, falls away with it. The common wrong formulation writes a prior on T_{rt} beside the one on φ.

Exercise 13.6. A circle for the window's state c, with its prior square, and a line from c to each unit's factor, which now reads $P(F_A \mid c)$; a line to the feeder's factor too, if its rise in a disturbed window is kept. The units are now tied by a path through c that nothing known

cuts. Of the table's queries, service at a random hour and service through a year change. The both-units-out row already asks for a window factor, and it now finds one on the diagram. The opening's adequacy, which structure, the overheating and test queries do not change, since the structure's output is unread, it falls away, and the new part feeds nothing that stays. A unit's or the feeder's rate given its record changes only if the record's counts are modeled by window.

Exercise 13.7.

Run, which readings bear on the opening	
1 Sentence	With duty and PDU readings added, which readings bear on $G \geq G^{*}$?
Intention	Estimate: a question about the diagram
2 Target	The factors $P(G \geq G^{*} \mid \cdot)$ needs
Declare	d the difference [K]; duty and PDU readings [kW]; the outage logs; the loads [kW]
3 Known	The stated model's diagram; which readings exist
Random here	G; the offsets and the loads, carried
4 Event	$G \geq G^{*}$
5 Expression	$P(G \geq G^{*} \mid$ the readings that reach it)
6 Model	Hall readings through d; duty, PDU through the loads; logs only at the structure
7 Check	The structure's output is not read, so it falls away; no path then joins the logs to G
Class	Factors of a query
	$\boxed{P(G \geq G^{*} \mid d,\ \text{duty},\ \text{PDU readings})}$

Exercise 13.8. This note is ready. The fact sheet says what the static bypass does: if the UPS itself fails, it feeds the load directly from the supply. So it adds no new circle, and the place to put it is the definition of F_{ups}, in the demand factor.

Run, the bypass, in the demand factor	
1 Sentence	Where does the static bypass enter the IT load's factors?
Intention	Estimate: a question about the diagram
2 Target	A factor of the diagram
Declare	F_{sup}; F_{ups} with its three causes; t_{aut}; a UPS's repair with the load on its bypass
3 Known	The fact sheet's protection and its bypass; the bypass does not fail (assumption)
Random here	Whether the UPS is out for repair when the supply is lost
4 Event	F_{ups} given F_{sup}: no pick-up, or the battery outlasted, or the UPS on its bypass
5 Expression	$P(F_{\text{ups}} \mid F_{\text{sup}})$
6 Model	The bypass is the third cause: it keeps the load up until the supply is lost
7 Check	A UPS's repair alone loses nothing, so no cut set is added
Class	Factors of a query
	$\boxed{P(F_{\text{ups}} \mid F_{\text{sup}})}$

Exercise 13.9. With D_{out} at the stated length, the product is $p(C_{\text{p}})\,p(\varphi)\,P(H_i)\,\mathbf{1}[T_{\text{rt}} = \cdot]\,\mathbf{1}[M = \cdot]$, where the running maximum follows from the pack's heat balance, the thermal mass, and how long the string heats, which is the smaller of D_{out} and T_{rt}. By the rule nothing falls away at either length, since T_{rt} is an input of the factor asked about and φ and the cell are its inputs. A value argument then does more, and the diagram cannot show it. At 1 h no string has reached its floor, since the lowest ride-through is 1.20 h, so T_{rt} does not enter the maximum on any value it can take, and φ and the hypothesis can be dropped. At 3 h every string has reached its floor, so they cannot, under the assumption that a string at its floor stops heating. The common wrong formulation keeps φ at 1 h without seeing that it can go, or drops it at 3 h.

Exercise 13.10. The note adds a quantity that both units' repairs depend on: a circle for the contractor's state, with a prior asked for, and a line to each unit's repair factor, which now reads $P(F_i \mid$ the contractor's state). The two units are then tied by a path that nothing known cuts, as a shared cause ties them. Of the table's queries, both units out, service at a random

hour and service through a year change. A unit's rate does not, since the note is about repairs and not about failures, and the opening's adequacy does not, since the structure's output is unread. The note gives no value, so the sheet asks how often the contractor is short and how far repairs run over; a record of repair lengths by month could say. The common wrong formulation raises each unit's mean repair and keeps the units independent.

Exercise 13.11. Hall B loses its cooling when the feeder is out, or when unit B is out and either unit A is out or the opening cannot pass its load. It is a deterministic factor, so hall B's event has no prior of its own.

Exercise 13.12. It is the probability that the next outage is carried to its end, weighed over whether a cell is weak. Inside each term the outage's length and the state of health have been integrated out, and the sum takes out the hypothesis, so nothing but the event is left.

Exercise 13.13. It is what is known of the conductance after one 60 kW test. Every failure factor is left out: the structure's output is neither asked about nor read, so it falls away, and the states' factors then feed nothing that stays.

Exercise 13.14. It is the joint of the two units' and the feeder's states at a random hour, on a model in which a window's state moves both units together and the feeder is separate. It ties the two units through c, and it leaves the feeder untied from either, since $P(F_{\mathrm{fd}})$ has no c in it; that is the assumption to label, given that the fact sheet raises the feeder too in a disturbed window.

Exercise 13.15. The summary reports one cut set for the load's event. The failure class is a component event substituted for the system event, in the form of a sub-diagram's answer for the whole. The structure has three cut sets, and the expression would not notice a fault on the IT branch, which by itself takes the load and is about a fifth of the load's unavailability on the illustrative inputs, or an uncleared cooling fault, which is far smaller and still there. The repair writes the structure: $P(S_{\mathrm{IT}}(t) = 0) \approx P(F_{\mathrm{sup}})\,P(F_{\mathrm{ups}} \mid F_{\mathrm{sup}}) + P(F_{\mathrm{it}}) + P(F_{\mathrm{clr}})$.

Exercise 13.16. The note treats a read output as if it were unread. A factor whose output is read stays, and it can tie its inputs, deterministic or not. On the structure, with the opening adequate and the feeder and unit A both up, the service cannot be lost, so on the stated model this event gives $P(G \geq G^{*} \mid S = 0,\ F_{\mathrm{fd}}^{c},\ F_{A}^{c},\ E) = 0$, not 0.82. The failure class is model conditioning dropped, since the observation has been left off what the number rests on. The repair writes the query with the event behind the bar. It also asks whether the record is right and whether the structure is complete, because a probability of zero on one model is a reason to check that model.

Exercise 13.17. There are two errors. The meters bear on the efficiency and the offset, and on the diagram they share no factor with the pick-up or the autonomy, so a sub-diagram's answer has been reported for the whole: this lesson's class. And a statement about a state has been read as a certainty about an outcome, Lesson 7's class, since a healthy UPS with ten minutes of battery still does not carry an interruption of an hour. The repair is the load's event in an interruption: $P(F_{\mathrm{ups}} \mid F_{\mathrm{sup}})$, with the pick-up and the autonomy, on the fact sheet's illustrative inputs, with the meters' evidence entering only if the efficiency bears on the autonomy, which the fact sheet does not state.

Exercise 13.18. The sentence reads a probability about one fixed value as a share of time. The opening has one conductance, so the plant has one structure at every hour of the year, and it is the analyst who does not know which. The failure class is parameter against outcome,

Lesson 7's, in the form of a parameter's probability read as an outcome's. The arithmetic of the service happens to survive at a random hour, since 0.971 averages over the same weights. It does not survive over a year: the interval availability is one structure's or the other's for the whole year, and its distribution has two humps. The repair writes the weights as what is known of G, and it writes a year's availability on each structure.

Exercise 13.19. A deterministic factor says that T_{rt} follows from its parents, not that its parents are known. The ride-through is uncertain because φ and the cell are, and $P(T_{\mathrm{rt}} \geq 2\ \mathrm{h})$ is 0.074 on the fact sheet's priors. The error reads a statement about a relation as a statement about a value. The failure class is model conditioning dropped, since the parents' uncertainty has been left off what the number rests on. The repair writes $P(T_{\mathrm{rt}} \geq 2\ \mathrm{h}) = \sum_i P(H_i) \int \mathbf{1}[2.025\ \mathrm{h} \times \gamma_i\, \varphi \geq 2\ \mathrm{h}]\, p(\varphi)\, \mathrm{d}\varphi$, with $\gamma_{\mathrm{weak}} = 0.70$ and $\gamma_{\mathrm{sound}} = 1$.

Exercise 13.20. Each item names one factor needed and one that falls away; any valid factor scores.

(a) The reading's factor; the outage's length.

(b) The pack's heat balance; the weak cell's prior, by a value argument, for lengths under 1.20 h, when no string has reached its floor.

(c) The prior on φ; nothing on its own diagram falls away.

(d) The five readings' factors; the pick-up.

(e) The PDU reading's factor, with the others for the loads; the failure factors.

(f) The branch's count factor; the meters.

(g) The IT branch's factor; the supply's rate, since the interruption is given.

(h) The weak cell's prior; everything else.

(i) The clearing's demand factor; the meters.

(j) The rate factors of the load's losses; the meters.

Exercise 13.21. Each problem is a run.

Run, (1) hall B at 45 kW		
1	**Sentence**	Unit B out, $Q_B = 45\,\mathrm{kW}$, settled, given E: over its limit?
	Intention	Estimate
2	**Target**	An event at a stated load; through it G
	Declare	$g(G;\ Q_B) = \Delta T_{\mathrm{allow}} - Q_B/G$ [K]; E the 40 kW trip test
3	**Known**	Unit B out; $Q_B = 45\,\mathrm{kW}$; settled; E
	Random here	G only
4	**Event**	$g(G;\ 45\,\mathrm{kW}) < 0$
5	**Expression**	$P(g(G;\ 45\,\mathrm{kW}) < 0 \mid \text{unit B out},\ \text{settled},\ E)$
6	**Model**	The opening's law; what is known of G after the test
	Reduce	$P(G < 5625\,\mathrm{W/K} \mid E)$
7	**Check**	A stated load, no occasion: Estimate; $R > 0.1778$ K/kW is as good
	Class	Posterior exceedance
		$\boxed{P(G < 5625\,\mathrm{W/K} \mid E)}$

Run, (2) hall B tonight

1	**Sentence**	Tonight, unit B out, load unknown, settled, given E: hall B over its limit?
	Intention	Predict: an occasion to come
2	**Target**	An outcome, through Q_B and G
	Declare	Q_B normal about 50 kW, spread 5 kW (illustrative); G [W/K]; $\Delta T_{\text{allow}} = 8\,\text{K}$
3	**Known**	Unit B out; settled; E; load independent of the opening (assumption)
	Random here	Q_B, varying from night to night; G, one unknown value
4	**Event**	$Q_B/G > \Delta T_{\text{allow}}$
5	**Expression**	$P(Q_B/G > \Delta T_{\text{allow}} \mid \text{unit B out, settled, } E)$
6	**Model**	The opening's law; $p(Q_B)$; $p(G \mid E)$
	Reduce	$\int P(Q_B > \Delta T_{\text{allow}} G)\, p(G \mid E)\, \mathrm{d}G$
7	**Check**	The load is averaged over, not set at its center
	Class	Predictive probability

$$\boxed{\int P(Q_B > \Delta T_{\text{allow}} G)\, p(G \mid E)\, \mathrm{d}G}$$

Run, (3) a 60 kW test's likelihood

1	**Sentence**	A 60 kW test's reading: how probable, for each R?
	Intention	Estimate: the likelihood serves the posterior
2	**Target**	$p(d \mid R)$, a function of R
	Declare	d [K]; $Q = 60\,\text{kW}$; R [K/kW]; $\sigma_d = 0.711\,\text{K}$
3	**Known**	Unit B out; settled; $\delta = 0$ (assumption); offsets integrated out
	Random here	d; R held at each value
4	**Event**	None: a density of a reading
5	**Expression**	$p(d \mid R)$
6	**Model**	$d = 60\,\text{kW} \times R + \varepsilon$
7	**Check**	A function of R, linear in its mean
	Class	Likelihood, nuisance integrated out

$$\boxed{p(d \mid R) = \mathcal{N}(d;\, 60\,\text{kW} \times R,\, \sigma_d^2)}$$

Run, (4) a weak cell, from one timed reading

1	**Sentence**	A timed discharge reads 1.30 h, spread 0.03 h: is a cell weak?
	Intention	Diagnose
2	**Target**	A hypothesis, H_{weak}
	Declare	$d = 1.30$ h; φ [–]; $\gamma_{\text{weak}} = 0.70$, $\gamma_{\text{sound}} = 1$; $P(H_{\text{weak}}) = 0.10$ (illustrative)
3	**Known**	d; 20 kW; the fact sheet's priors; start 0.80; $\delta = 0$ (assumptions)
	Random here	Which hypothesis; φ, integrated out under each
4	**Event**	H_{weak}
5	**Expression**	$P(H_{\text{weak}} \mid d)$
6	**Model**	$p(d \mid H_i) = \int \mathcal{N}(d;\, 2.025\,\gamma_i\,\varphi,\, 0.03^2)\, p(\varphi)\, \mathrm{d}\varphi$; the prior on the hypotheses
7	**Check**	1.30 h lies in the weak range, 1.20 to 1.42 h, and far below the sound range
	Class	Posterior probability over candidate causes

$$\boxed{P(H_{\text{weak}} \mid d) \propto p(d \mid H_{\text{weak}})\, P(H_{\text{weak}})}$$

Run, (5) a certificate's prior

1	**Sentence**	"Offset within ±5 kW at 95 %": what prior on the PDU meter's offset?
	Intention	Estimate: a prior serves the posterior
2	**Target**	A distribution on b_{pdu}
	Declare	b_{pdu} [kW]; center 0; 5 kW a 95 % half-width, so a standard deviation of 5/1.96 kW
3	**Known**	A document about this meter; no drift since the certificate (assumption)
	Random here	b_{pdu}
4	**Event**	None: a density
5	**Expression**	$\mathcal{N}(b_{\text{pdu}};\, 0,\, (5\,\text{kW}/1.96)^2)$
6	**Model**	Normal (choice); the certificate states a width and no shape
7	**Check**	A statement about the meter, with no reading to write a likelihood from, so a prior
	Class	Prior

$$\boxed{p(b_{\text{pdu}}) = \mathcal{N}(b_{\text{pdu}};\, 0,\, (5\,\text{kW}/1.96)^2)}$$

Run, (6) a test before a stated change

1	**Sentence**	Last year's trip test, opening enlarged since: evidence or prior on today's G?
	Intention	Estimate: its place serves the posterior
2	**Target**	Its place at step 3
	Declare	m_0, σ_{r} what d_0 left of $\log G$; a_{e}, σ_{e} the enlargement in $\log G$, asked for
3	**Known**	In the model, in hand; this opening, with a stated physical change since
	Random here	G
4	**Event**	None: a factor of the joint
5	**Expression**	$p(\log G)$, what last year's test left known, moved and widened by the enlargement
6	**Model**	Lognormal (choice); the enlargement independent of what the test left (assumption)
7	**Check**	Fails the third question on the change, so a prior; not a likelihood beside this year's
	Class	Prior

$$\boxed{p(\log G) = \mathcal{N}(\log G;\, m_0 + a_{\text{e}},\, \sigma_{\text{r}}^2 + \sigma_{\text{e}}^2)}$$

Run, (7) the first cut set

1	**Sentence**	Random hour, illustrative rates: supply lost and the UPS not carrying?
	Intention	Predict
2	**Target**	One cut set of the load's event
	Declare	F_{sup}; F_{ups}; $t_{\text{aut}} = 10$ min (illustrative)
3	**Known**	Steady; the structure; rates, demand probability, autonomy (illustrative)
	Random here	The supply's state; a loss's length; the pick-up
4	**Event**	$F_{\text{sup}} \cap F_{\text{ups}}$
5	**Expression**	$P(F_{\text{sup}} \cap F_{\text{ups}})$
6	**Model**	Losses' rates and lengths; pick-up; autonomy; repairs
	Reduce	$P(F_{\text{sup}})\, P(F_{\text{ups}} \mid F_{\text{sup}})$
7	**Check**	Not independent: F_{ups} is a failure on demand
	Class	Predictive probability at a random time, with a demand

$$\boxed{P(F_{\text{sup}} \cap F_{\text{ups}}) = P(F_{\text{sup}})\, P(F_{\text{ups}} \mid F_{\text{sup}})}$$

Run, (8) the next outage, after a timed discharge

1	**Sentence**	After a timed reading d, the weak cell not ruled out: is the next outage carried?
	Intention	Predict
2	**Target**	An outcome; a predictive over φ and the hypotheses
	Declare	d [h]; φ [–]; $T_{\mathrm{rt}} = 2.025\ \mathrm{h} \times \gamma_i \times \varphi$; $\gamma_{\mathrm{weak}} = 0.70$, $\gamma_{\mathrm{sound}} = 1$; D_{out} [h]
3	**Known**	d; the outages' length distribution (illustrative)
	Random here	D_{out}; φ and the hypothesis, averaged over
4	**Event**	$D_{\mathrm{out}} \le T_{\mathrm{rt}}$
5	**Expression**	$P(D_{\mathrm{out}} \le T_{\mathrm{rt}} \mid d)$
6	**Model**	$d = 2.025\ \mathrm{h} \times \gamma_i \times \varphi + \varepsilon$; the bound carried; the weak cell's prior
	Reduce	$\sum_i P(H_i \mid d) \int P(D_{\mathrm{out}} \le 2.025\ \mathrm{h} \times \gamma_i\, \varphi)\, p(\varphi \mid d, H_i)\, \mathrm{d}\varphi$
7	**Check**	A predictive, not a plug-in at the reading
	Class	Predictive probability of a race

$$\boxed{\sum_i P(H_i \mid d) \int P(D_{\mathrm{out}} \le 2.025\ \mathrm{h} \times \gamma_i\, \varphi)\, p(\varphi \mid d, H_i)\, \mathrm{d}\varphi}$$

Run, (9) hall B keeps its cooling

1	**Sentence**	Given E and the stated rates, a random hour: does hall B keep its cooling?
	Intention	Predict
2	**Target**	One hall's event; part of the service
	Declare	F_A, F_B, F_{fd}; q a unit down; f the feeder down; G, G^*
3	**Known**	E; stated rates; steady; design load; windows out (assumption)
	Random here	The states at the hour; G
4	**Event**	Not $F_{\mathrm{fd}} \cup \big(F_B \cap (F_A \cup \{G < G^*\})\big)$
5	**Expression**	$1 - P\big(F_{\mathrm{fd}} \cup (F_B \cap (F_A \cup \{G < G^*\})) \mid E\big)$, given the stated model
6	**Model**	The states' factors, the prior and the reading; a hall's structure, not the plant's
	Reduce	$1 - \big(f + (1-f)\, q\, (q + (1-q)\, P(G < G^* \mid E))\big)$
7	**Check**	0.979, above the plant's 0.971: on the mirrored limit hall A can lose it alone
	Class	Predictive probability at a random time, with dependence

$$\boxed{1 - \big(f + (1-f)\, q\, (q + (1-q)\, P(G < G^* \mid E))\big)}$$

Run, (10) what a read service needs

1	**Sentence**	Which factors of the plant's diagram does $P(F_A \mid S = 0,\ E)$ need?
	Intention	Estimate: a question about the diagram
2	**Target**	A sub-diagram
	Declare	$p(G)$; $p(d \mid G)$; the three states' factors; the structure
3	**Known**	d observed; $S = 0$ read
	Random here	F_A, F_B, F_{fd}, G
4	**Event**	F_A
5	**Expression**	$P(F_A \mid S = 0,\ E)$
6	**Model**	The structure's output is read: it stays; its inputs' factors feed it; G joins d
7	**Check**	Unread, $P(F_A \mid E)$ would need only unit A's factor
	Class	Factors of a query

$$\boxed{p(G)\, p(d \mid G)\, P(F_A)\, P(F_B)\, P(F_{\mathrm{fd}})\, \mathbf{1}[S = s(\cdot)]}$$

Lesson 14

Exercise 14.1. Actions: hand over; enlarge the opening; repeat the test at a higher load. States: $G \ge G^*$ and $G < G^*$, with "adequate" taken as the N+1 claim, which the procedure should say. Consequences, from its words: a handover on an inadequate opening costs the owner the hall's equipment, and enlarging a sound opening costs the contractor the work; the others are not stated. Thresholds: G^* inside the event, 0.9 outside it. The rule is

$A(e) = \mathbf{1}[P(G \geq G^* \mid e) \geq 0.9]$, where 1 is the handover. It does not decide between the other two actions: “as the contractor chooses” is a person, not a function of the evidence, and the contractor bears only one of the two costs. The common wrong formulation stops at the 0.9 and calls the rule complete.

Exercise 14.2. This one is ready. It names the claim, the evidence and the level, and only the other action is left open, which the value on E does not need.

Run, the N+1 claim at the owner’s 0.80		
1	**Sentence**	Given E: accept the N+1 claim if its probability reaches 0.80?
	Intention	Decide
2	**Target**	A rule’s value on evidence in hand
	Declare	$G^* = 7500\,\mathrm{W/K}$; $A(\cdot)$; level 0.80, stated
3	**Known**	E; unit B out; settled; the stated model
	Random here	G, inside the rule
4	**Event**	$G \geq G^*$
5	**Expression**	$A(E)$
	Expand	$\mathbf{1}[P(G \geq G^* \mid E) \geq 0.80]$
6	**Model**	The posterior
	Reduce	$\mathbf{1}[0.818 \geq 0.80]$
7	**Check**	Value 1; the claim’s probability is still 0.818
	Class	Decision rule
		$\boxed{A(E) = \mathbf{1}[P(G \geq G^* \mid E) \geq 0.80]}$

Exercise 14.3.

Run, a 50 kW claim at a level of 0.80		
1	**Sentence**	Given E, level 0.80: is the claim that hall B carries 50 kW within 8 K accepted?
	Intention	Decide
2	**Target**	The value of a rule on evidence in hand
	Declare	$G^*_{50} = 50\,\mathrm{kW}/8\,\mathrm{K} = 6250\,\mathrm{W/K}$, the claim’s threshold; $A(\cdot)$; level 0.80
3	**Known**	E; unit B out; settled; the stated model
	Random here	G, inside the rule
4	**Event**	$G \geq G^*_{50}$
5	**Expression**	$A(E)$
	Expand	$\mathbf{1}[P(G \geq G^*_{50} \mid E) \geq 0.80]$
6	**Model**	The posterior
	Reduce	$\mathbf{1}[0.992 \geq 0.80]$
7	**Check**	A claim about 50 kW, and silent on 60 kW
	Class	Decision rule
		$\boxed{A(E) = \mathbf{1}[P(G \geq G^*_{50} \mid E) \geq 0.80]}$

Exercise 14.4. $A(\mathbf{y}) = \mathbf{1}[P(H_{\text{ups}} \mid \mathbf{y}) \geq 0.5]$, where 1 sends the crew and 0 does not. The action set is send, do not send. It leaves undecided what happens when no explanation reaches 0.5, and whether the PDU meter’s explanation has an action of its own. The posterior is taken over the three explanations with their weights 0.90, 0.05 and 0.05, so a rule on the degraded explanation still uses the other two. The common wrong formulation holds the efficiency’s reading against a limit, which is a rule on a point value.

Exercise 14.5.

Run, the string kept on a timed reading		
1	**Sentence**	A timed reading of 2.00 h, a sound string, level 0.90: kept, if it gives its 2 h?
	Intention	Decide
2	**Target**	The value of a rule on a reading in hand
	Declare	$d = 2.00$ h; spread 0.03 h (illustrative); $T_{\mathrm{rt}} = 2.025$ h $\times\, \varphi$; level 0.90
3	**Known**	d; sound string (assumption); φ uniform on 0.85 to 1.00
	Random here	φ, inside the rule
4	**Event**	$T_{\mathrm{rt}} \geq 2$ h
5	**Expression**	$A(d)$
	Expand	$\mathbf{1}[P(T_{\mathrm{rt}} \geq 2 \text{ h} \mid d) \geq 0.90]$
6	**Model**	$d = 2.025$ h $\times\, \varphi + \varepsilon$; the prior with its bound
	Reduce	$\mathbf{1}[0.373 \geq 0.90]$
7	**Check**	A reading of exactly 2 hours does not keep it: the rule asks for more than the reading
	Class	Decision rule

$$\boxed{A(d) = \mathbf{1}[P(T_{\mathrm{rt}} \geq 2 \text{ h} \mid d) \geq 0.90]}$$

Exercise 14.6. $A(k) = \mathbf{1}[P(\lambda_A > 3.5 \mid \text{the record}) > 0.5]$, where here the value 1 means replace, so 1 is the costly action and not an acceptance. With $k = 7$ over $t_{\mathrm{s}} = 3$ yr in service and the stated prior, the record is a count with its exposure, on the clock of time in service. On a calendar clock the same count over three years of operation is a count read on the wrong exposure, which is Lesson 11's class. The rule is on the rate, a parameter, and it says nothing about next year's count.

Exercise 14.7. With $p = P(D_{\mathrm{out}} \leq T_{\mathrm{rt}} \mid \cdot)$: keep when $p \geq p_1$, shed a share when $p_2 \leq p < p_1$, replace when $p < p_2$. Shedding changes the event, since a smaller load gives a longer T_{rt}, so p here must be the probability with no shedding, and the rule has to say so. Both levels and the share are illustrative, as is the outage's distribution; the critical load is stated. The common wrong formulation writes one level and treats shedding as if it left the event alone.

Exercise 14.8. This request is intentionally ambiguous, and it is not ready. It names no claim, no evidence, no level and no alternative. Three readings. If the claim is the specification, on the fact sheet with no test, the expression is $\mathbf{1}[P(T_{\mathrm{rt}} \geq 2 \text{ h}) \geq p]$, whose inner value is 0.074. If it is the next outage, it is $\mathbf{1}[P(D_{\mathrm{out}} \leq T_{\mathrm{rt}}) \geq p]$, whose inner value is 0.834. If it is a timed discharge the manager has in mind, the argument is that reading, and it has to be supplied. In every reading p is missing, and so is what happens otherwise. The analyst asks for the claim, the evidence, the level and the other action.

Exercise 14.9. $A(e) = \mathbf{1}[\mathbb{E}[Q_{\mathrm{hall}}/G \mid e] < 8\,\mathrm{K}]$. Its inner quantity is an expectation, and no probability level appears. On E the expected rise is 7.03 K, so its value is 1, while the rule at 0.95 gives 0. The two rules part because a mean below the limit is compatible with a good deal of probability above it: here 0.182. The clause is a clean rule, but it carries no statement about how much risk it accepts, so it cannot be read as "95 %" or as any level.

Exercise 14.10. $A(k) = \mathbf{1}[k \leq 3 \text{ yr} \times 3.69/\text{yr}] = \mathbf{1}[k \leq 11.07]$, given the record. The fact sheet's three-year record has $k = 11$, so its value is 1. The rule is on the count, with a threshold from the stated rate, and no probability is in it, so a count of 12, which is well within what a rate of 3.69 produces, would refuse a feeder running exactly at its stated rate. The common wrong formulation calls the rule a test of the rate; it asks only whether this count fell at or below its mean.

Exercise 14.11. On the reading of a timed discharge, keep the string if the next outage is at least 90 % likely to be carried to its end. Its argument is the reading d, and the rule has

a predictive inside it, averaged over what the reading leaves unknown and over the outage's length.

Exercise 14.12. Accept when the probability of the opening being inadequate, given the evidence, is at most 0.05. It is the same rule as the one at 0.95, with the claim's complement in place of the claim: the two have the same value on every reading.

Exercise 14.13. Send the crew when the degraded explanation has probability at least 0.2, given the five readings. A level this low treats a degraded UPS left in service as the worse mistake and a needless visit as the lesser one. Lesson 17 writes how much worse, as a ratio of losses.

Exercise 14.14. Given the feeder's record, accept if the feeder is at least 80 % likely to fail at most three times next year. Inside is a predictive probability of a count over a horizon, so the rate has been averaged over, and the rule is not on the rate.

Exercise 14.15. Failure class: rule against posterior, accept read as true. The note answers "what did the rule do with this test?" and reports it as a fact about the plant. Repair: name the rule and its level, say that it accepted, and give $P(G \geq G^* \mid e)$ on that reading.

Exercise 14.16. Failure class: rule against posterior, a posterior read as a decision. There is no rule and no level, and "even odds" is a level chosen after the reading. The 0.82 answers how probable the claim is. Repair: state the level from the consequences, and write $A(E)$; at the owner's 0.95, its value is 0.

Exercise 14.17. Failure class: rule against posterior, a level read as the probability delivered. The rule says the risk is at most 5 % on an accepted reading, and it equals 5 % only on the reading where the rule just accepts. On a stronger reading it is smaller. Repair: give $P(G < G^* \mid e)$ on the reading in hand. Lesson 15 takes up what else this sentence is mistaken for.

Exercise 14.18. Two errors, from two classes. The level was set after the reading, so the rule is no longer a function of the evidence with a stated level; this lesson's near-miss. And a further test was added because the first failed, with no statement of how last year's test bears on the plant today, which is Lesson 12's question: it may be evidence or a prior, it may already be in the prior, and the sheet must say so. The expression reported answers whether some level and some evidence can be found that pass. Repair: the stated level; both tests in one posterior, if last year's passes that lesson's rule, which is a rule on pooled evidence and is stated as one (Lesson 16); a value of the rule on that.

Exercise 14.19. Failure class: rule against posterior, not accepted read as false. The rule's value has been written in place of the posterior. The expression reported answers what the rule did, and the probability is still 0.818; a retest would change it only when its readings are in hand. Repair: report $A(E) = 0$ and $P(G \geq G^* \mid E) = 0.818$ side by side.

Exercise 14.20. (a) A threshold on a quantity, once a limit makes it a load to be carried. (b) A rule's value. (c) An action. (d) A threshold on a probability. (e) A consequence. (f) A rule, with no probability in it. (g) A number about the plant, given E. (h) An action, on the network. (i) A rule, with three tests as its argument. (j) A rule's value on a test not yet run: random until the test is read, so it is not yet a number.

Lesson 15

Exercise 15.1. "95 % acceptance confidence" is a level: $P(\text{claim} \mid e) \geq 0.95$ on an accepted reading, a bound by this reading, if the rule is a posterior one, which the sheet does not say. "Under 1 % false pass" is a rate with no truth: by truth it is $P(A(\tilde{E}) = 1 \mid \text{a stated truth})$, or by outcome $P(\text{inadequate} \mid A(\tilde{E}) = 1)$ on some population, and the sheet says neither. "4 % of 200 sites later reworked" is a count, so an estimate by outcome at best, if reworked means inadequate and if all 200 were accepted: the rework may have other causes, and the 200 are the vendor's population and not this plant's. The common wrong formulation reads all three as one number stated three ways.

Exercise 15.2.

Run, a false rejection at 90 kW

1	**Sentence**	A 90 kW opening, a 40 kW test, rule at 0.95: how often refused?
	Intention	Predict
2	**Target**	An outcome of a test not yet run, at a fixed truth
	Declare	$G_1 = 90\,\text{kW}/8\,\text{K} = 11\,250\,\text{W/K}$; $\tilde{E}_{40}$; $A(\cdot)$ at 0.95
3	**Known**	$G = G_1$ (stated); unit B out; settled; fresh sensors (assumption)
	Random here	The readings only
4	**Event**	$A(\tilde{E}_{40}) = 0$
5	**Expression**	$P(A(\tilde{E}_{40}) = 0 \mid G = G_1)$
	Expand	$P(P(G \geq G^* \mid \tilde{E}_{40}) < 0.95 \mid G = G_1)$
6	**Model**	The readings' relation at G_1; the rule's prior
7	**Check**	0.161; the report's count is 34 of 200 refused
	Class	Operating characteristic at a stated truth

$$\boxed{P(A(\tilde{E}_{40}) = 0 \mid G = 11\,250\,\text{W/K})}$$

Exercise 15.3. $P(\text{flag} \mid H^c_{\text{off}})$ does not move, at 0.061, since it is by truth and no weight is in it. $P(H^c_{\text{off}} \mid \text{flag})$ moves from 0.564 to 0.064, since it is by outcome, and the weights are its population. The common wrong formulation calls the second a property of the alarm.

Exercise 15.4.

Run, the share inadequate at a stricter level

1	**Sentence**	Plants from the prior, 20 kW test, rule at 0.99: what share of those accepted are inadequate?
	Intention	Predict
2	**Target**	A state, given a test's outcome
	Declare	$\tilde{E}_{20}$; $A_{0.99}(\cdot)$; $p(G)$ the prior, read as a population (assumption)
3	**Known**	The outcome alone; the population; fresh sensors (assumption)
	Random here	G and the readings
4	**Event**	$G < G^*$
5	**Expression**	$P(G < G^* \mid A_{0.99}(\tilde{E}_{20}) = 1)$
	Expand	The rule's condition behind the bar: $P(G \geq G^* \mid \tilde{E}_{20}) \geq 0.99$
6	**Model**	Bayes' rule on the population
	Reduce	The ratio of equation (15.2), with this rule
7	**Check**	0.0023; under 0.01, since with this prior as the population it averages risks each at most 0.01
	Class	Predictive value of a test's outcome

$$\boxed{P(G < G^* \mid A_{0.99}(\tilde{E}_{20}) = 1)}$$

Exercise 15.5. On one string, year by year, with fresh errors and no ageing: $P(A(\tilde{d}) = 1 \mid \varphi)$ at this string's φ, a rate by truth, estimated by the share of its years that pass. No share of

weak strings enters. On a list of strings once each: $P(A(\tilde{d}) = 1)$ averaged over the list, and any "of those accepted" statement is by outcome, which needs the list's share of weak strings π_{w} and the spread of φ across it. Only the second needs a population. The first must state whether the string ages, since otherwise the years do not repeat one truth.

Exercise 15.6. This one is ready: it names the truth, the test load and the rule.

Run, a false pass at 55 kW

1	**Sentence**	True capacity 55 kW, a 60 kW test, the report's rule: how often does it pass?
	Intention	Predict
2	**Target**	An outcome of a test not yet run, at a fixed truth
	Declare	$G_0 = 55\,\mathrm{kW}/8\,\mathrm{K} = 6875\,\mathrm{W/K}$; $\tilde{E}_{60}$; $A(\cdot)$ at 0.95
3	**Known**	$G = G_0$ (stated); unit B out; settled; fresh sensors (assumption)
	Random here	The readings only
4	**Event**	$A(\tilde{E}_{60}) = 1$
5	**Expression**	$P(A(\tilde{E}_{60}) = 1 \mid G = G_0)$
	Expand	$P(P(G \geq G^* \mid \tilde{E}_{60}) \geq 0.95 \mid G = G_0)$
6	**Model**	The readings' relation at G_0; the rule's prior
7	**Check**	0.005; the report's count is 1 of 200
	Class	Operating characteristic at a stated truth

$$\boxed{P(A(\tilde{E}_{60}) = 1 \mid G = 6875\,\mathrm{W/K})}$$

Exercise 15.7. $P(A(\tilde{k}) = 0 \mid \lambda_{\mathrm{fd}} = \lambda_1)$, a false rejection by truth, where $\tilde{k}$ is a count not yet logged, on a stated exposure. "Good" has to be given a value λ_1, and the manager's sentence has none; at the stated rate of 3.69, a rule of $\tilde{k} \leq 11$ over three years in service refuses about four records in ten (0.43). A sentence with no rate and no rule behind it cannot be written.

Exercise 15.8. By truth: $P(\text{accepted} \mid \text{inadequate}) = 2/100 = 0.02$, holding the plant's kind; $P(\text{not accepted} \mid \text{adequate}) = 200/900 = 0.222$. By outcome: $P(\text{inadequate} \mid \text{accepted}) = 2/702 = 0.0028$, holding the outcome; $P(\text{adequate} \mid \text{not accepted}) = 200/298 = 0.671$. All four are on this population. A level bounds the posterior risk on each accepted reading directly, and these four average over readings. Only the third can be read against it, and only as that average, which falls at or under the level when the rule's prior is this population.

Exercise 15.9. This question is intentionally ambiguous, and it is not ready. It does not say at which truth, over which population or at which load and level. By truth, at a stated G_0 and load: $P(A(\tilde{E}_Q) = 1 \mid G = G_0)$, which on the report's rule from 40 kW is 0.0095 at 54 kW and 0.065 just below G^*. Averaged over the prior's inadequate plants: 0.0026. By outcome: $P(G < G^* \mid A(\tilde{E}_Q) = 1)$, 0.0032 on the prior. By a reading in hand: $P(G < G^* \mid E)$, 0.182, though "rate" does not suggest it. The most literal reading is the first, and it needs a truth. The agent is asked for the truth or the population, the load and the rule.

Exercise 15.10. With the offset stated at zero, the rule refuses when $|\tilde{d}| \geq 6\,\mathrm{kW}$, where $\tilde{d}$ is the difference of the two meters at the next calibration. The false rejection is $P(|\tilde{d}| \geq 6\,\mathrm{kW} \mid b_{\mathrm{pdu}} = 0)$. Random: the two meters' noises, so that $\tilde{d}$ is normal about 0 with spread 2.83 kW. Reduced under that stated model it is $2(1 - \Phi(6/2.83))$, which is 0.034. It differs from the alarm's 0.061, which averages over the offset's 0 ± 1.5 kW under the other explanations: a stated truth against an average over it.

Exercise 15.11. If the opening's conductance is 12 000 W/K, how often a 40 kW test would fail to accept it, over the readings that plant could give.

Exercise 15.12. Of the plants a 60 kW test does not accept, the share that are adequate. It needs a prior over plants as a population, which the second expression does not need except inside the rule.

Exercise 15.13. Given the 40 kW trip test, how probable it is that the opening cannot carry 50 kW within the limit. No rule appears because the question is about the plant, on evidence in hand, and nothing is being accepted or refused.

Exercise 15.14. If the opening's conductance is 8000 W/K, how often a 20 kW test gives readings on which the rule at 0.99 accepts. The inner G is what the rule does not know, integrated over its posterior for each possible reading. The outer one is the truth, held at 8000 W/K while the readings vary.

Exercise 15.15. Failure class: posterior risk read as a sampling error rate, the level read as a sampling rate. Two errors in one sentence. The 0.05 is a bound on the posterior risk on accepted readings, and it is not a rate at any truth. And "catching an inadequate plant" is $P(A(\tilde{E}_Q) = 0 \mid G = G_0)$, which has no one value: 0.990 at 54 kW and 0.935 just below G^* from a 40 kW test. Repair: state a truth, and give that rate.

Exercise 15.16. Failure class: posterior risk read as a sampling error rate, a sampling rate read as a posterior risk. The report has transposed the bar: 0.061 is $P(\text{flag} \mid H_{\text{off}}^c)$, and it is read as $1 - P(H_{\text{off}} \mid \text{flag})$. With the fact sheet's weights a flag alone gives an offset 0.436, and tonight's reading can say more. Repair: $P(H_{\text{off}} \mid y)$ from tonight's y, which at 8 kW is 0.077 and at 12 kW is 0.802.

Exercise 15.17. Failure class: posterior risk read as a sampling error rate, a posterior value read as a frequency. The 0.595 is $P(T_{\text{rt}} \geq 2 \text{ h} \mid d = 2.025 \text{ h})$, at one reading, and no probability over readings is in it; it is under the rule's level, so that reading itself does not pass. The report's question is $P(A(\tilde{d}) = 1 \mid \varphi = 1)$, which is 0.022. Repair: quote that, and name which is which.

Exercise 15.18. Two errors. Failure class: posterior risk read as a sampling error rate, the level read as a sampling rate. The 17 of 200 is a rate at one truth, at G^*, the limit of the false acceptance just below it, and nothing makes it equal α; here it is about 0.06 in closed form and 0.085 by this count. The second error takes 1 less a rate at G^* and calls it a level. That quantity is how often this plant is refused. Repair: report 17 of 200 at that truth, and leave the level as what the rule states.

Exercise 15.19. Two classes combined. A population average for a plant with its reading in hand is this lesson's class, the outcome used where a reading is stated: tonight's string read low, and that reading goes behind the bar, not the fleet's share. Inside it sits Lesson 12's question: "half are weak" is a statement about the list, which could set a prior on this string, and a prior built from that list is not shown to be "right" by the list; if tonight's string is on the list, it is counted twice. Repair: a prior from the list with this string taken out, and $P(\text{the claim} \mid \text{tonight's reading})$. The expression reported answers what share of a fleet with those proportions would be inadequate among those accepted.

Exercise 15.20. (a) By truth. (b) By outcome. (c) By this reading. (d) Not ready: no truth and no rate. (e) By outcome. (f) By truth. (g) Not ready: of what, given what. (h) By this reading. (i) By truth, if "sound" is given a value. (j) By outcome: with the reading in hand it would be by this reading, and "nothing else" is what makes it the outcome.

Lesson 16

Exercise 16.1. (a) $G \geq G^*$, with $G^* = Q_{\text{hall}}/\Delta T_{\text{allow}} = 60\,\text{kW}/8\,\text{K} = 7500\,\text{W/K}$; unit B out, settled, behind the bar. It is one event in $\log G$, the report's coordinate, and in R, as $R \leq R^*$.

(b) $A(e) = \mathbf{1}[P(G \geq G^* \mid e) \geq 0.95]$. "From the readings of the test, not before them" is what makes e the argument.

(c) The random variable is $\tilde{E}_Q$, "the readings that truth generates". The event is $A(\tilde{E}_Q) = 1$, "that the rule accepts". The condition is $G = G_1$, "at a truth". The rule is (b). Translate: $P(A(\tilde{E}_Q) = 1 \mid G = G_1)$. Expand: $P(P(G \geq G^* \mid \tilde{E}_Q) \geq 0.95 \mid G = G_1)$. The inner G is integrated by the posterior, for each possible reading. The outer one is held at G_1. The inner is over G, and the outer over the readings.

(d) The inner falls with $\tilde{d}$, with one reading and one unknown, so the event inside is $\{\tilde{d} \leq d_{\text{crit}}(Q)\}$, with $P(G \geq G^* \mid d_{\text{crit}}(Q)) = 0.95$. In R, $\Phi((R^* - m_R(\tilde{d}))/\sigma_R) \geq 0.95$ under Principle 10.1. At 20 kW, $d_{\text{crit}} = 1.66\,\text{K}$ puts the point 2.33 reading spreads from $R = 0$, under the 3 required: step 3 refuses, and the exact d_{crit} is written there. The reading would have given 1.71 K.

(e) Assumed: the measurement relation $\tilde{d} = Q/G + \varepsilon$, normal, with spread σ_d from the floor; $\delta = 0$; unit B out; settled; fresh sensors on each test.

Run, the report's power	
1 Sentence	At a truth, over the readings that truth generates: how often does the rule accept?
Intention	Predict
2 Target	An outcome of a test not yet run, at a fixed truth
Declare	$\tilde{E}_Q$; $\tilde{d}$ [K]; $A(\cdot)$ of (b); G_1 [W/K]; Q [kW]; σ_d [K]; $d_{\text{crit}}(Q)$ [K]
3 Known	$G = G_1$, stated; Q; unit B out; settled; $\delta = 0$, fresh sensors (assumptions)
Random here	$\tilde{d}$ only; no evidence in hand
4 Event	$A(\tilde{E}_Q) = 1$
5 Expression	$P(A(\tilde{E}_Q) = 1 \mid G = G_1)$
Expand	$P(P(G \geq G^* \mid \tilde{E}_Q) \geq 0.95 \mid G = G_1)$
6 Model	$\tilde{d}$ given G_1 normal about Q/G_1, spread σ_d
Reduce	$\{\tilde{d} \leq d_{\text{crit}}(Q)\}$; then Φ, exact under this model
7 Check	A function of Q and G_1, of no reading; it rises with G_1
Class	Operating characteristic at a stated truth
	$\boxed{\Phi\big((d_{\text{crit}}(Q) - Q/G_1)/\sigma_d\big)}$

(f) $G(C) = C/\Delta T_{\text{allow}}$, so $Q/G(C) = Q\Delta T_{\text{allow}}/C$, and the power at C is $\Phi((d_{\text{crit}}(Q) - Q\Delta T_{\text{allow}}/C)/\sigma_d)$. At $Q = 60\,\text{kW}$, $C = 70\,\text{kW}$: $60 \times 8/70 = 6.857\,\text{K}$, so $\Phi((6.89 - 6.857)/0.711) = \Phi(0.05)$, about one half.

(g) Design case: "the full-load design case crosses 0.95", at the 70 kW cell. Its expression is $P(G \geq G^* \mid d = Q/G(C))$, one number for each C, with no probability over readings. Power: "the power reaches 0.9", at 80 kW, 95 kW and nowhere. Its expression is (e). Part (f) is the design-case crossing seen from the other side: where the design case is 0.95, $Q/G(C)$ is about d_{crit}, the argument is near 0, and the power near one half. "89 of 200" is a count on that same event, 0.445, against about 0.52.

(h) The report's "probability of acceptance" is the posterior, $P(G \geq G^* \mid e)$, on a test's evidence e, and it is not how often a test accepts: the report's name for that is the power. The name is the trap, since in standard use, and in this book's vocabulary, the probability of acceptance is the operating characteristic. "The risk of a false acceptance" is α, a bound on the posterior risk on accepted readings: Lesson 15's third conditioning, by this reading. It is not $P(A(\tilde{E}_Q) = 1 \mid G = G_0)$, which at a 54 kW opening from a 40 kW test is 0.010.

The wrong answers this exercise catches are a rule with no function; G where G_1 is held; 0.956 for how often; the inner's Φ for the outer's; noise alone for σ_d; and passage D's crossings taken as one kind.

Exercise 16.2. (a) $\varphi \geq \varphi^*$, with $\varphi^* = 1.90\,\text{h}/2.025\,\text{h} = 0.9383$, the same event as $T_{\text{rt}} \geq 1.90\,\text{h}$ on a sound string.

(b) $A(\tilde{d}) = \mathbf{1}[P(\varphi \geq \varphi^* \mid \tilde{d}) \geq 0.95]$, a function of the reading.

(c) The random variable is the reading not yet taken, $\tilde{d}$. The event is $A(\tilde{d}) = 1$. The condition is $\varphi = \varphi_1$, a stated state of health. The rule is (b). Translate: $P(A(\tilde{d}) = 1 \mid \varphi = \varphi_1)$. Expand: $P(P(\varphi \geq \varphi^* \mid \tilde{d}) \geq 0.95 \mid \varphi = \varphi_1)$. The inner φ is integrated by the posterior, for each possible reading, and the outer one is held at φ_1.

(d) The inner rises with $\tilde{d}$, so the event inside is $\{\tilde{d} \geq d_{\text{crit}}\}$, with $P(\varphi \geq \varphi^* \mid d_{\text{crit}}) = 0.95$ and $d_{\text{crit}} = 1.007$ h. The checklist in φ: step 1 passes, since $\tilde{d} = 1.0125 \text{ h} \times \varphi + \varepsilon$ is linear with a fixed spread, $0.03/1.0125 = 0.0296$ in φ. Step 2 passes, on a uniform prior. Step 3 refuses: the point $d_{\text{crit}}/1.0125 \text{ h} = 0.9946$ lies 0.18 reading spreads from the bound at 1, under the 3 required. So the exact expression is written, the normal cut at the prior's bounds, with $m = \tilde{d}/1.0125$ h and $s = 0.0296$:

$$P(\varphi \geq \varphi^* \mid \tilde{d}) = \frac{\Phi((1-m)/s) - \Phi((\varphi^* - m)/s)}{\Phi((1-m)/s) - \Phi((0.85-m)/s)}.$$

(e) Assumed: the stated relation, normal, with a spread of 0.03 h; a sound string; the timing error drawn afresh on each test.

Run, how often the half-depth test accepts	
1 Sentence	A half-depth timed test, the rule at 0.95, a stated state of health: how often does it accept?
Intention	Predict
2 Target	An outcome of a test not yet run, at a fixed truth
Declare	$\tilde{d}$ the timed reading [h]; φ_1 stated; $\varphi^* = 0.9383$; $d_{\text{crit}} = 1.007$ h; spread 0.03 h (illustrative)
3 Known	$\varphi = \varphi_1$, stated; sound string; fresh timing error (assumptions); no reading in hand
Random here	$\tilde{d}$ only
4 Event	$A(\tilde{d}) = 1$
5 Expression	$P(A(\tilde{d}) = 1 \mid \varphi = \varphi_1)$
Expand	$P(P(\varphi \geq \varphi^* \mid \tilde{d}) \geq 0.95 \mid \varphi = \varphi_1)$; inner φ integrated, outer held
6 Model	$\tilde{d}$ given φ_1 normal about $1.0125 \text{ h} \times \varphi_1$, spread 0.03 h
Reduce	The inner rises with $\tilde{d}$: $\{\tilde{d} \geq d_{\text{crit}}\}$; then an upper tail, $1 - \Phi(\cdot) = \Phi(-\cdot)$
7 Check	Rises with φ_1; the argument is design value less critical reading
Class	Operating characteristic at a stated truth
	$\boxed{\Phi\big((1.0125 \text{ h} \times \varphi_1 - d_{\text{crit}})/0.03 \text{ h}\big)}$

On the opening a larger conductance gives a smaller difference, so the rule accepts below its critical value and the argument is critical less design. Here a healthier string gives a longer time, so the rule accepts above, and the argument is design less critical. The sign follows from which way the inner moves with the reading, and substituting into the opening's formula gets it wrong.

(f) At 1.00: $(1.0125 - 1.007)/0.03 = 0.18$, so $\Phi(0.18) = 0.57$, against 104 of 200, or 0.52. At 0.93: $(0.9416 - 1.007)/0.03 = -2.18$, so 0.015, against 3 of 200, or the same.

(g) Design case: "the first cell ... at which the design case reaches 0.95". Its expression is $P(\varphi \geq \varphi^* \mid d = 1.0125 \text{ h} \times \varphi_1)$, a posterior at the expected reading, with no probability over readings. Power: "reaches 0.9 nowhere", "3 of 200" and "104 of 200", all from (e). Where the design case is exactly 0.95 the design value equals d_{crit}, so the argument is 0 and the power one half; the cell at 1.00 is just past that point. The power reaches 0.9 when $1.0125 \text{ h} \times \varphi_1 \geq 1.007 + 1.28 \times 0.03$ h, at $\varphi_1 = 1.03$, which the bound rules out.

(h) "At 0.93, 3 of 200 accept": 0.93 is under φ^*, so it is a false acceptance by truth, a sampling rate at a stated inadequate string, 0.015. It is not the rule's α, and no share of accepted strings has been stated.

Exercise 16.3.

Run, a stricter rule at 10 000 W/K		
1	**Sentence**	A 60 kW test, the rule at 0.99, true conductance 10 000 W/K: how often does it accept?
	Intention	Predict
2	**Target**	An outcome of a test not yet run, at a fixed truth
	Declare	$\tilde{E}_{60}$; $\tilde{d}$ [K]; $G_1 = 10\,000\,\mathrm{W/K}$; d_{crit} at $0.99 = 6.41$ K; $\sigma_d = 0.711\,\mathrm{K}$
3	**Known**	G_1 (stated); 60 kW; unit B out; settled; fresh sensors (assumption)
	Random here	$\tilde{d}$ only
4	**Event**	$A_{0.99}(\tilde{E}_{60}) = 1$
5	**Expression**	$P(A_{0.99}(\tilde{E}_{60}) = 1 \mid G = G_1)$
	Expand	$P(P(G \geq G^* \mid \tilde{E}_{60}) \geq 0.99 \mid G = G_1)$
6	**Model**	$\tilde{d}$ normal about $60\,\mathrm{kW}/G_1 = 6\,\mathrm{K}$, spread σ_d
	Reduce	$\Phi((6.41 - 6)/0.711)$
7	**Check**	0.717; at 0.95 the same plant passes 0.896 of the time
	Class	Operating characteristic at a stated truth
		$\boxed{\Phi\big((d_{\mathrm{crit}} - 60\,\mathrm{kW}/G_1)/\sigma_d\big),\ d_{\mathrm{crit}} \text{ at } 0.99}$

Exercise 16.4. This one is ready.

Run, the test's power at the drawing		
1	**Sentence**	The 66 kW drawing, the report's rule, a 60 kW test: what is the power?
	Intention	Predict
2	**Target**	An outcome of a test not yet run, at a fixed truth
	Declare	$\tilde{E}_{60}$; $\tilde{d}$ [K]; $G_1 = 8250\,\mathrm{W/K}$; $d_{\mathrm{crit}}(60\,\mathrm{kW}) = 6.89\,\mathrm{K}$; $\sigma_d = 0.711\,\mathrm{K}$
3	**Known**	G_1 as an assumption, the drawing; unit B out; settled; fresh sensors (assumption)
	Random here	$\tilde{d}$ only
4	**Event**	$A(\tilde{E}_{60}) = 1$
5	**Expression**	$P(A(\tilde{E}_{60}) = 1 \mid G = G_1)$
	Expand	$P(P(G \geq G^* \mid \tilde{E}_{60}) \geq 0.95 \mid G = G_1)$
6	**Model**	$\tilde{d}$ normal about 7.27 K, spread σ_d
	Reduce	$\Phi((6.89 - 7.27)/0.711)$
7	**Check**	0.297; it assumes the drawing
	Class	Operating characteristic at a stated truth
		$\boxed{\Phi\big((d_{\mathrm{crit}}(60\,\mathrm{kW}) - 60\,\mathrm{kW}/8250\,\mathrm{W/K})/\sigma_d\big)}$

Exercise 16.5. $P(A(\tilde{d}) = 1) = \int_{0.85}^{1} (1 - \Phi((2.085 - 2.025\varphi)/0.03))\,\frac{\mathrm{d}\varphi}{0.15}$, for a string taken as sound. Its value is 0.0008. The rule almost never accepts any string that the prior allows. That says the rule is built badly for this test, since a reading of more than nameplate is needed before a bounded prior can give 0.90 to the top 8 % of its range. A sound string is not at fault: the rule asks more of the reading than its spread can give. With the weak cell included the value is smaller still, by the factor 0.9.

Exercise 16.6.

Run, a false acceptance at 6000 W/K		
1	**Sentence**	A 40 kW test, the rule at 0.95, true conductance 6000 W/K: how often does it accept?
	Intention	Predict
2	**Target**	An outcome of a test not yet run, at a fixed truth
	Declare	$\tilde{E}_{40}$; $\tilde{d}$ [K]; $G_0 = 6000\,\mathrm{W/K}$; $d_{\mathrm{crit}}(40\,\mathrm{kW}) = 4.26\,\mathrm{K}$; $\sigma_d = 0.711\,\mathrm{K}$
3	**Known**	G_0 (stated); 40 kW; unit B out; settled; fresh sensors (assumption)
	Random here	$\tilde{d}$ only
4	**Event**	$A(\tilde{E}_{40}) = 1$
5	**Expression**	$P(A(\tilde{E}_{40}) = 1 \mid G = G_0)$
	Expand	$P(P(G \geq G^* \mid \tilde{E}_{40}) \geq 0.95 \mid G = G_0)$
6	**Model**	$\tilde{d}$ normal about $40\,\mathrm{kW}/G_0 = 6.67\,\mathrm{K}$, spread σ_d
	Reduce	$\Phi((4.26 - 6.67)/0.711)$
7	**Check**	0.0004; a truth this far below G^* almost never passes
	Class	Operating characteristic at a stated truth
		$\boxed{\Phi\big((d_{\mathrm{crit}}(40\,\mathrm{kW}) - 40\,\mathrm{kW}/G_0)/\sigma_d\big)}$

Exercise 16.7. The inner is normal with center $0.9961\,\tilde{d}$ and spread 2.823 kW, under the offset explanation, and it falls with $\tilde{d}$. So the rule is $\{\tilde{d} \leq (5 - 1.645 \times 2.823)/0.9961\} = \{\tilde{d} \leq 0.358\,\mathrm{kW}\}$, and the outer is $\Phi((0.358 - b_{\mathrm{pdu}})/2.83)$: 0.550 at offset 0, and 0.003 at offset 8 kW. Here the inner's normal is exact, since the prior is normal in the reading's coordinate with no bound, which passes steps 3 and 4 of Principle 10.1 at once. A meter with no offset is accepted only a little over half the time, because the limit is close to the reading's own spread.

Exercise 16.8. This box is intentionally ambiguous, and it is not ready. "Power" names no truth, and "adequate" no level. Three readings. At the drawing, with a stated level: $\Phi((d_{\mathrm{crit}}(Q) - Q/8250\,\mathrm{W/K})/\sigma_d) \geq 0.9$, which fails at every allowed load, 0.297 at most. At a stated larger capacity: the same at that truth, which holds at 80 kW from a full-load test, by count. Over belief, which would be assurance: $\int \Phi(\cdot)\, p_{\mathrm{o}}(G_1)\, \mathrm{d}G_1$, 0.444 at 60 kW. The plan's author is asked for the truth, the level and the test load, and whether a false rejection is what "adequate" guards against.

Exercise 16.9. $P(\tilde{k} \leq 11 \mid \lambda_{\mathrm{fd}}) = \sum_{k=0}^{11} e^{-3\lambda_{\mathrm{fd}}} (3\lambda_{\mathrm{fd}})^k / k!$, a count's tail, with $\tilde{k}$ not yet logged. Exposure 3 yr in service, a Poisson count (assumption). No Φ is needed, since the count's tail is the standard form. At the stated 3.69 it is 0.571, so a feeder exactly at its stated rate fails 43 % of the time, as Exercise 15.7 found.

Exercise 16.10. The form is the same: $\Phi((d'_{\mathrm{crit}}(60\,\mathrm{kW}) - 60\,\mathrm{kW}/8250\,\mathrm{W/K})/\sigma'_d)$, with $\sigma'_d = 0.362\,\mathrm{K}$. The change reaches two places. The denominator is one. The other is the rule itself, since the inner's posterior uses the same spread, so $d'_{\mathrm{crit}}(60\,\mathrm{kW}) = 7.42\,\mathrm{K}$ in place of 6.89 K. The power is 0.659, against 0.297, at no change in load: at this plant better sensors do more than the allowed load can. A formulation that updates only the denominator gives the wrong d_{crit} and the wrong answer.

Exercise 16.11. This one is ready, and it is not the fresh-sensors question. The same pair at

the same load repeats the offset and Q/G, so what changes between the tests is the noise alone.

Run, a second 40 kW test on the same sensors		
1	**Sentence**	Given E, the sensors left in place: will a second 40 kW test pass, on its reading alone?
	Intention	Predict
2	**Target**	An outcome of a test not yet run, given a reading on the same sensors
	Declare	$d = 4.78\,\mathrm{K}$ read; $\tilde{d}$ [K]; b the offsets' difference; noise of each test 0.077 K
3	**Known**	E; same load; same offsets, G unchanged (assumptions); the rule on the new reading alone
	Random here	The two tests' noises; G and b cancel in the difference
4	**Event**	$\tilde{d} \le d_{\mathrm{crit}}(40\,\mathrm{kW})$
5	**Expression**	$P(A(\tilde{E}_{40}) = 1 \mid E)$, same sensors
	Expand	$P(P(G \ge G^* \mid \tilde{E}_{40}) \ge 0.95 \mid E)$
6	**Model**	$\tilde{d} - d$ normal about 0, spread 0.110 K; the prior on $Q/G + b$ flat across 0.077 K (assumption)
	Reduce	$\{\tilde{d} \le d_{\mathrm{crit}}(40\,\mathrm{kW})\}$; then $\Phi((4.26 - 4.78)/0.110)$
7	**Check**	Below 0.001, against 0.335 on fresh sensors: a retest on the same pair measures the same offset
	Class	Predictive probability
		$\boxed{\Phi\big((d_{\mathrm{crit}}(40\,\mathrm{kW}) - d)/(\sqrt{2} \times 0.077\,\mathrm{K})\big)}$

Exercise 16.12. If the opening's conductance is G_1, how often a test's reading is one on which the Gaussian reading in R of the claim's probability reaches 0.95. It may be used only at loads where Principle 10.1's four steps pass at the critical reading. At 20 kW step 3 refuses, and the exact inner is used.

Exercise 16.13. At a stated resistance R_1, how often a 40 kW test accepts. In kelvin throughout: 40 kW times R_1 in K/kW is K, and the difference of two kelvin values over 0.711 K is dimensionless.

Exercise 16.14. How often a record over five years in service, with a true rate λ, shows at most two failures; the rule is $\mathbf{1}[\tilde{k} \le 2]$ on such a record. It is a Poisson tail, and no Φ is in it.

Exercise 16.15. Failure class: posterior risk read as a sampling error rate, a sampling rate read as a posterior risk, with Lesson 14's rule against posterior beside it. The note answers "how often does this test accept such a plant?" and reports it as how probable this plant's adequacy is. Power is at a stated truth, and this plant's truth is not stated. Repair: $P(G \ge G^* \mid e)$ on the reading in hand, at least the rule's level if the rule was a posterior one.

Exercise 16.16. Failure class: realized against prospective evidence, a rule replaced by a threshold that no longer holds the posterior; and the limit of 8 K belongs to full load. The rule is not a function of Q, so a light test passes nearly anything. At 20 kW on a 30 kW opening, which is 3750 W/K and half of what is needed, $Q/G = 5.33\,\mathrm{K}$, and $\Phi((8 - 5.33)/0.711)$ is above 0.999. Repair: the rule at a stated level, with $d_{\mathrm{crit}}(Q)$ for each load. The procedure's rule answers whether hall B ran within the limit during the test, not whether it would at full load.

Exercise 16.17. Failure class: realized against prospective evidence, one truth used where the truth is unknown, and here not the truth that decides. At 130 kW from a 40 kW test the power is 0.994, and no claim is near it. Where the decision lies, near 60 to 70 kW, the same test gives 0.065 to 0.330. Repair: report the operating characteristic near the decision, and state which false rejections matter.

Exercise 16.18. Failure class: realized against prospective evidence, one truth used where the truth is unknown. A point value has been used for the integral, as in Lesson 7's plug-in.

The expression reported answers how often a test passes if the opening is at the posterior's mode in $\log G$, 8347 W/K, which is 0.34 from $\Phi((6.89 - 60\,\mathrm{kW}/8347\,\mathrm{W/K})/0.711)$. Repair: $\int \Phi((d_{\mathrm{crit}} - 60\,\mathrm{kW}/G_1)/\sigma_d)\, p(G_1 \mid E)\, \mathrm{d}G_1$, which is 0.457, higher because Φ is far from straight here, with the rule on the new reading alone, as stated.

Exercise 16.19. Two classes. The outer belief has been chosen to agree with the drawing and so to nearly reproduce the power at the drawing: with the rule's prior held at the stated one while only the outer belief narrows, an assurance on a 3 % spread is the 0.297 again, and calling it assurance is one truth used where the truth is unknown. Narrowing the rule's prior too would change the rule, and the 0.30 would not follow. And the outer belief states what is believed about G: the fact sheet's 2 % tolerance is a tolerance on the drawing's value, conditional on the drawing being the plant's, and the test exists because that is doubted, which is Lesson 12's question. Also 0.30 does not show the test adequate on any reading. Repair: the stated prior, 0.444, beside the power at the drawing, 0.297, each named, and a stated level to read them against.

Exercise 16.20. (a) Random R; held, the readings of E. (b) Random, the readings of a 20 kW test; held, G. (c) Random, those readings and G; held, the readings of E. (d) Nothing random; the value is known. (e) Random $\tilde{d}$; held, G and Q. (f) Random, the count; held, the rate. (g) Random φ; held, the timed reading. (h) Random, the reading and φ; nothing held. (i) Nothing random: a stated number per load and rule. (j) Random G; held, a stated reading and load, with no test behind them.

Lesson 17

Exercise 17.1. Actions: leave, enlarge (12), fit the duct (30). States: $G \geq G^*$ and $G < G^*$ for the opening as it stands. The objective is the action's cost $\mathcal{L}(a)$, and the constraint is the probability of overheating at most 0.01:

$$\arg\min_{a \in \mathbb{A}} \mathcal{L}(a) \quad \text{s.t.} \quad P(\text{overheat} \mid E, \mathrm{do}(a)) \leq 0.01.$$

Leaving it has a risk of 0.182, which fails the constraint. The note does not say that enlarging makes the opening adequate, or by how much, and the expression needs that: if it does, enlarging costs 12 and passes; if not, the duct at 30 passes. The note also prices an overheating at 150, which is the other way of writing the tolerance: expected loss in place of the limit gives leave 27.3, enlarge 12, duct 30, and the same action follows here. The sheet says which it uses. The common wrong formulation treats "as little as possible" as the whole objective and drops the constraint.

Exercise 17.2. This ticket is intentionally ambiguous, and it is not ready: best against what, under which limit, at which load, on what evidence. Three readings. Least risk: the second rung, $\arg\min_{\mathbf{u}} P(\vartheta_B(\mathbf{u}) > \vartheta_{\max} \mid E, \mathrm{do}(\mathbf{u}))$, which the bound decides at 18°C. Least electricity with a tolerance: the third rung, with the tolerance asked for. Least electricity and an overheating priced: $\arg\min_{\mathbf{u}} \mathbb{E}[\mathcal{L}(\mathbf{u}, R) \mid E]$, with the price asked for. Least electricity with nothing else is the most literal reading, and it gives the highest set point the units allow. The manager is asked for the objective, the tolerance or the price, and the condition under which it is asked.

Exercise 17.3.

Run, the limit and the set point moved together		
1	**Sentence**	Given E, unit B out, 60 kW, set point 23°C, limit 31°C: is hall B over its limit?
	Intention	Estimate
2	**Target**	A parameter's event, at a set control
	Declare	$\mathbf{u} = 23°\text{C}$; $\vartheta_{\max} = 31°\text{C}$ (illustrative); $\vartheta_B(\mathbf{u}) = \mathbf{u} + 60\,\text{kW} \times R$
3	**Known**	E; do($\mathbf{u}$); unit B out; 60 kW; settled
	Random here	R
4	**Event**	$\vartheta_B(\mathbf{u}) > \vartheta_{\max}$
5	**Expression**	$P(\vartheta_B(\mathbf{u}) > \vartheta_{\max} \mid E, \text{do}(\mathbf{u} = 23°\text{C}))$
6	**Model**	The opening's relation; the posterior
	Reduce	$P(60\,\text{kW} \times R > 8\,\text{K} \mid E) = P(G < G^* \mid E)$
7	**Check**	0.182: the rise is still 8 K, so this is the first rung's event
	Class	Probability estimation; here a posterior exceedance

$$\boxed{P(G < G^* \mid E)}$$

Exercise 17.4. $\arg\min_{\mathbf{u}} \mathbb{E}[P_{\text{elec}}(\mathbf{u}) \mid E]$ s.t. $P(\mathbf{u} + 60\,\text{kW} \times R > 28°\text{C} \mid E, \text{do}(\mathbf{u})) \leq 0.01$. The event's threshold moves, and with it the set point, by the same 2 K, to 18.48°C, which is just above the illustrative bound of 18°C; the electricity rises to 34.56 kW. The 0.99 point of R does not move, since nothing about the opening changed.

Exercise 17.5.

Run, keep, shed or replace the string		
1	**Sentence**	No test: keep the string, shed a quarter in the next outage, or replace it, on the stated losses?
	Intention	Decide
2	**Target**	An action
	Declare	$\mathbb{A}$ keep, shed, replace; $\mathcal{L}$: not carried 100, shed 5, replace 20 (illustrative); $T_{\text{rt}}(\xi)$; D_{out}
3	**Known**	The fact sheet's priors, with the weak cell; no test; a new string at $\varphi = 1$ (stated)
	Random here	φ, the weak cell, D_{out}; the action is chosen
4	**Event**	$D_{\text{out}} \leq T_{\text{rt}}(\xi)$, for each action's ξ and string
5	**Expression**	$\arg\min_a \mathbb{E}[\mathcal{L}(a, \cdot)]$
	Expand	$\mathcal{L}(a) + 100\,(1 - P(D_{\text{out}} \leq T_{\text{rt}} \mid \text{do}(a)))$
6	**Model**	The string's relation; exponential lengths
	Reduce	Keep 16.6; shed 14.2; replace 33.2
7	**Check**	Shed; replacing buys little here, since a new string also misses long outages (0.868)
	Class	Expected-loss minimization

$$\boxed{\arg\min_{a \in \mathbb{A}} \big(\mathcal{L}(a) + 100\,(1 - P(D_{\text{out}} \leq T_{\text{rt}} \mid \text{do}(a)))\big)}$$

Exercise 17.6. Equation (17.2) with $\mathcal{L}_{\text{FA}} = 380 - 10 = 370$ and $\mathcal{L}_{\text{FR}} = 10$ gives $370/380 = 0.974$. The action on E does not change: accepting costs 69.2 and modifying 10, as before. Only a posterior between 0.947 and 0.974 would change the choice.

Exercise 17.7. A table of three actions by three explanations, $\mathcal{L}(a, H_i)$, and the choice

$$\arg\min_a \sum_i \mathcal{L}(a, H_i)\, P(H_i \mid \mathbf{y}).$$

A single level cannot replace it, since equation (17.2) is for two actions and two states. Each pair of actions gives a boundary, and those boundaries partition the three posteriors. The common wrong formulation acts on the most probable explanation, as Problem 14.7 did, which ignores what each mistake costs.

Exercise 17.8. This one is ready, and its form is the third rung. The limit and the electricity's model are the two things it leaves to the fact sheet and to the lesson's illustrative values, which the sheet labels.

Run, a chance-constrained set point at 5 %

1	**Sentence**	Given E, a unit's loss at 60 kW: least electricity, with hall B's risk at most 5 %?
	Intention	Optimize
2	**Target**	A setting, under a constraint on a probability
	Declare	$\mathbf{u}$; $\vartheta_B(\mathbf{u}) = \mathbf{u} + 60\,\text{kW} \times R$; $\vartheta_{\max} = 30°\text{C}$; $P_{\text{elec}}(\mathbf{u})$ (illustrative); tolerance 0.05
3	**Known**	E; unit B out; 60 kW; settled; $\mathbf{u}$ chosen; P_{elec} taken as known (assumption)
	Random here	R; the expectation on P_{elec} is trivial here
4	**Event**	$\vartheta_B(\mathbf{u}) > \vartheta_{\max}$, in the constraint; objective P_{elec}
5	**Expression**	$\arg\min_{\mathbf{u}} \mathbb{E}[P_{\text{elec}}(\mathbf{u}) \mid E]$ s.t. $P(\vartheta_B(\mathbf{u}) > \vartheta_{\max} \mid E, \text{do}(\mathbf{u})) \le 0.05$
6	**Model**	P_{elec} falls with $\mathbf{u}$ and the risk rises, so the constraint binds
	Reduce	$\mathbf{u} = 30°\text{C} - 60\,\text{kW} \times R_{0.95}$, with $R_{0.95}$ the posterior's 0.95 point
7	**Check**	21.21°C, 30.91 kW: a looser tolerance saves 0.90 kW against the 1 % rung; the tolerance is the owner's
	Class	Chance-constrained optimization

$$\boxed{\arg\min_{\mathbf{u}} \mathbb{E}[P_{\text{elec}}(\mathbf{u}) \mid E] \quad \text{s.t.} \quad P(\vartheta_B(\mathbf{u}) > \vartheta_{\max} \mid E, \text{do}(\mathbf{u})) \le 0.05}$$

Exercise 17.9. The trip test: do(unit B out), a setting, with the plant at a stated load and the readings as evidence on G. The log's nights: F_B found, an observation, with whatever came with it, such as a disturbance or a high load. The test supports the claim as it stands, since the claim is a statement about the opening with unit B out and the test's setting is known not to have been chosen by what else was happening. The log's nights could be evidence too, if their load and temperatures were read and the reasons for each trip stated; as three nights it is a count of trips, and no reading of the opening.

Exercise 17.10. On two actions, with the level $1 - \alpha$ equal to $\mathcal{L}_{\text{FA}}/(\mathcal{L}_{\text{FA}} + \mathcal{L}_{\text{FR}})$, for every reading e. The rule is then the argmin. With a third action, or with another ratio, the two can part: Table 17.1 has the ratio 18, so a posterior between 0.947 and 0.95 would be accepted by the argmin and refused by the rule.

Exercise 17.11. Given the five readings, the action on the network with the least expected loss, averaged over the three explanations by their posterior weights.

Exercise 17.12. With hall A's load at 45 kW and a setting chosen, the probability that hall A exceeds its limit, given the trip test, is at most 0.05. Here $\mathbf{u}$ is a setting, chosen; $\vartheta_{\max}$ is a stated limit; E is evidence; Q_A is a known condition; and the 0.05 is a tolerance. The opening's relation is what remains random.

Exercise 17.13. Among the log's hours with the set point at 23°C, how often hall B was over the limit. It does not answer what happens if 23°C is chosen on a given night, unless the log's hours are known to be like that night.

Exercise 17.14. The smallest share of the load to shed, given a timed reading, so that the next outage is carried with probability at least 0.99. Its class is chance-constrained optimization.

Exercise 17.15. Failure class: intervention read as observation, a found setting for a set one. The note answers how the hours at 24°C went, with the loads those hours had. Set on a full-load night, on this lesson's illustrative limit, it is 0.833. "Never overheated" is also a count over light hours, which says little about a heavy one. Repair: do($\mathbf{u} = 24°\text{C}$) with the load stated.

Exercise 17.16. Failure class: intervention read as observation, a planned action read as evidence. A booking moves no probability about tonight's meter. The note answers what the diagnosis would be after a calibration that has not happened. Repair: keep H_{off} with its weight and tonight's reading, $P(H_{\text{off}} \mid \mathbf{y})$. Once the calibration has been done and read, it is evidence, or a stated setting of the offset, and the sheet says which.

Exercise 17.17. Failure class: parameter against outcome, a plug-in at a point value, inside a constraint. The note's constraint holds at the center of R, so about half the posterior lies over it: on this limit, $\mathbf{u} \leq 23.97°\text{C}$ with a risk of one half. It answers whether the temperature at the average opening is within the limit. Repair: the chance constraint with a stated tolerance; at 1 % on a 31°C limit, $\mathbf{u} \leq 21.48°\text{C}$.

Exercise 17.18. Failure class: intervention read as observation, a decision updating belief. The action depends on the loss table as much as on the posterior. On the lesson's table the argmin modifies at 0.818, because one state is costly, and the probability is still 0.818. The expression reported would answer what a modification says about the opening, and it says nothing. Repair: report the action and the posterior side by side, with the table.

Exercise 17.19. None of the listed classes applies: two tolerances are taken as one, which is the lesson's point on risk tolerance. The sentence answers what expected loss would do if it held a limit. An expected loss trades the risk against the cost at the stated price, and the risk that results is whatever the trade gives: 0.065 at a price of 10, and 0.005 at 100, on the illustrative inputs. It guarantees no level. Repair: if 1 % is required, state it as a constraint.

Exercise 17.20. (a) Observed. (b) Set. (c) Set; chance-constrained optimization. (d) Observed. (e) Set. (f) Set, and risk optimization. (g) Set, since an action is chosen; expected-loss minimization. (h) Observed. (i) Set. (j) Neither yet: a plan, which is not evidence, and its readings are not yet taken.

Exercise 17.21.

Run, (1) a rule at 0.99 on the 50 kW claim

1	**Sentence**	The rule at 0.99 for the claim that the opening carries 50 kW within 8 K: as a function, and on E?
	Intention	Decide
2	**Target**	A rule, and its value on evidence in hand
	Declare	$G^*_{50} = 6250\,\text{W/K}$, the claim's threshold; $A(\cdot)$; $1 - \alpha = 0.99$
3	**Known**	E; unit B out; settled; the stated model
	Random here	G, inside the rule
4	**Event**	$G \geq G^*_{50}$
5	**Expression**	$A(E)$
	Expand	$\mathbf{1}[P(G \geq G^*_{50} \mid E) \geq 0.99]$
6	**Model**	The posterior
	Reduce	$\mathbf{1}[0.992 \geq 0.99]$
7	**Check**	Value 1; a claim about 50 kW, and silent on 60 kW
	Class	Decision rule

$$\boxed{A(e) = \mathbf{1}[P(G \geq G^*_{50} \mid e) \geq 0.99]}$$

Run, (2) a false acceptance at 58 kW

1	**Sentence**	A 40 kW test, the rule at 0.95, a 58 kW opening: how often accepted?
	Intention	Predict
2	**Target**	An outcome of a test not yet run, at a fixed truth
	Declare	$G_0 = 7250\,\mathrm{W/K}$; $\tilde{d}$ [K]; $d_{\mathrm{crit}}(40\,\mathrm{kW}) = 4.26\,\mathrm{K}$; $\sigma_d = 0.711\,\mathrm{K}$
3	**Known**	G_0 (stated); unit B out; settled; fresh sensors (assumption)
	Random here	$\tilde{d}$ only
4	**Event**	$A(\tilde{E}_{40}) = 1$
5	**Expression**	$P(A(\tilde{E}_{40}) = 1 \mid G = G_0)$
	Expand	$P(P(G \geq G^* \mid \tilde{E}_{40}) \geq 0.95 \mid G = G_0)$
6	**Model**	$\tilde{d}$ normal about 5.52 K, spread σ_d
	Reduce	$\Phi((4.26 - 5.52)/0.711)$
7	**Check**	0.038; a sampling rate, and not the rule's α
	Class	Operating characteristic at a stated truth

$$\boxed{\Phi\big((d_{\mathrm{crit}}(40\,\mathrm{kW}) - 40\,\mathrm{kW}/G_0)/\sigma_d\big)}$$

Run, (3) power at 110 kW from a light test

1	**Sentence**	A 20 kW test, the rule at 0.95, a 110 kW opening: how often accepted?
	Intention	Predict
2	**Target**	An outcome of a test not yet run, at a fixed truth
	Declare	$G_1 = 13\,750\,\mathrm{W/K}$; $\tilde{d}$ [K]; $d_{\mathrm{crit}}(20\,\mathrm{kW}) = 1.659\,\mathrm{K}$, exact (step 3 refuses the reading in R)
3	**Known**	G_1 (stated); unit B out; settled; fresh sensors (assumption)
	Random here	$\tilde{d}$ only
4	**Event**	$A(\tilde{E}_{20}) = 1$
5	**Expression**	$P(A(\tilde{E}_{20}) = 1 \mid G = G_1)$
	Expand	$P(P(G \geq G^* \mid \tilde{E}_{20}) \geq 0.95 \mid G = G_1)$
6	**Model**	$\tilde{d}$ normal about 1.45 K, spread σ_d
	Reduce	$\Phi((1.659 - 1.45)/0.711)$
7	**Check**	0.613, at nearly twice the need: a light test is a weak test
	Class	Operating characteristic at a stated truth

$$\boxed{\Phi\big((d_{\mathrm{crit}}(20\,\mathrm{kW}) - 20\,\mathrm{kW}/G_1)/\sigma_d\big)}$$

Run, (4) assurance of a 20 kW test

1	**Sentence**	Before any test, on the stated prior: how likely is a 20 kW test to accept?
	Intention	Predict
2	**Target**	An outcome of a test not yet run, with the truth unknown
	Declare	$\tilde{E}_{20}$; $\tilde{d}$ [K]; $p(G)$ the stated prior; $d_{\mathrm{crit}}(20\,\mathrm{kW})$
3	**Known**	The prior; unit B out; settled; fresh sensors (assumption). No evidence in hand
	Random here	$\tilde{d}$ and G
4	**Event**	$A(\tilde{E}_{20}) = 1$
5	**Expression**	$P(A(\tilde{E}_{20}) = 1)$
	Expand	$\int P(P(G \geq G^* \mid \tilde{E}_{20}) \geq 0.95 \mid G = G_1)\, p_{\mathrm{o}}(G_1)\, \mathrm{d}G_1$
6	**Model**	The outer Φ at each G_1; p_{o} the stated prior
	Reduce	$\int \Phi((d_{\mathrm{crit}}(20\,\mathrm{kW}) - 20\,\mathrm{kW}/G_1)/\sigma_d)\, p_{\mathrm{o}}(G_1)\, \mathrm{d}G_1$
7	**Check**	0.319; a statement about the test before any plant is known
	Class	Assurance

$$\boxed{\int \Phi\big((d_{\mathrm{crit}}(20\,\mathrm{kW}) - 20\,\mathrm{kW}/G_1)/\sigma_d\big)\, p_{\mathrm{o}}(G_1)\, \mathrm{d}G_1}$$

Run, (5) a 40 kW test, given this one

1	**Sentence**	Given E: will a 40 kW test on fresh sensors accept, the rule on its reading alone?
	Intention	Predict
2	**Target**	An outcome of a test not yet run, with the truth known only through E
	Declare	$\tilde{E}_{40}$; $\tilde{d}$ [K]; $p(G \mid E)$; $d_{\text{crit}}(40\,\text{kW})$
3	**Known**	E; fresh sensors (assumption), so the tests are independent given G
	Random here	$\tilde{d}$ and G
4	**Event**	$A(\tilde{E}_{40}) = 1$
5	**Expression**	$P(A(\tilde{E}_{40}) = 1 \mid E)$
	Expand	$\int P(P(G \geq G^* \mid \tilde{E}_{40}) \geq 0.95 \mid G = G_1)\, p(G_1 \mid E)\, \mathrm{d}G_1$; the inner alone
6	**Model**	The outer Φ at each G_1; the posterior from E as the outer belief
	Reduce	$\int \Phi((d_{\text{crit}}(40\,\text{kW}) - 40\,\text{kW}/G_1)/\sigma_d)\, p(G_1 \mid E)\, \mathrm{d}G_1$
7	**Check**	0.335; the rule's value on E was 0, and that does not fix the next
	Class	Predictive probability

$$\boxed{\int \Phi\big((d_{\text{crit}}(40\,\text{kW}) - 40\,\text{kW}/G_1)/\sigma_d\big)\, p(G_1 \mid E)\, \mathrm{d}G_1}$$

Run, (6) the share inadequate at 0.99

1	**Sentence**	Plants from the prior, a 40 kW test, the rule at 0.99: what share of those accepted are inadequate?
	Intention	Predict
2	**Target**	A state, given a test's outcome
	Declare	$\tilde{E}_{40}$; $A_{0.99}(\cdot)$; $p(G)$ as a population (assumption)
3	**Known**	The outcome alone; the population; fresh sensors (assumption)
	Random here	G and the readings
4	**Event**	$G < G^*$
5	**Expression**	$P(G < G^* \mid A_{0.99}(\tilde{E}_{40}) = 1)$
	Expand	The rule's condition behind the bar: $P(G \geq G^* \mid \tilde{E}_{40}) \geq 0.99$
6	**Model**	Bayes' rule on the population
	Reduce	The ratio of equation (15.2), with this rule
7	**Check**	0.0006; a fleet's share, silent on a plant in hand
	Class	Predictive value of a test's outcome

$$\boxed{P(G < G^* \mid A_{0.99}(\tilde{E}_{40}) = 1)}$$

Run, (7) the string's rule at 0.95 of nameplate

1	**Sentence**	The timed-discharge rule at 0.90, $\varphi = 0.95$: how often is the string kept?
	Intention	Predict
2	**Target**	An outcome of a test not yet run, at a fixed truth
	Declare	$\tilde{d}$ [h]; spread 0.03 h (illustrative); critical reading 2.085 h, at or above
3	**Known**	φ (stated); sound string (assumption)
	Random here	$\tilde{d}$ only
4	**Event**	$A(\tilde{d}) = 1$
5	**Expression**	$P(A(\tilde{d}) = 1 \mid \varphi = 0.95)$
	Expand	$P(P(T_{\text{rt}} \geq 2\ \text{h} \mid \tilde{d}) \geq 0.90 \mid \varphi = 0.95)$
6	**Model**	$\tilde{d}$ normal about 2.025 h $\times$ 0.95; the rule's prior
	Reduce	$1 - \Phi((2.085 - 1.924)/0.03)$
7	**Check**	Below 0.001: such a string does not give 2 h, and the rule is right to refuse it
	Class	Operating characteristic at a stated truth

$$\boxed{1 - \Phi\big((2.085\ \text{h} - 2.025\ \text{h} \times 0.95)/0.03\ \text{h}\big)}$$

Run, (8) recalibrate or do nothing

1	**Sentence**	Given $y = 8\,\text{kW}$: recalibrate for 5, or do nothing, at 40 if an offset is there?
	Intention	Decide
2	**Target**	An action
	Declare	$\mathbb{A}$ recalibrate, nothing; $\mathcal{L}$ 5; 40 on H_{off}, 0 otherwise (illustrative); y [kW]
3	**Known**	y; the fact sheet's weights and spreads; the loss table
	Random here	The explanation
4	**Event**	H_{off}
5	**Expression**	$\arg\min_a \mathbb{E}[\mathcal{L}(a, \cdot) \mid y]$
6	**Model**	Bayes' rule on the two spreads of y
	Reduce	Nothing: $40 \times 0.077 = 3.1$; recalibrate: 5
7	**Check**	Do nothing; the level here is 1/8, and 0.077 is under it
	Class	Expected-loss minimization

$$\boxed{\arg\min_{a \in \mathbb{A}} \mathbb{E}[\mathcal{L}(a, \cdot) \mid y]}$$

Run, (9) hall A at 50 kW, with unit A out

1	**Sentence**	Unit A out, hall B's set point 22°C, hall A at 50 kW, given E: is hall A over 30°C?
	Intention	Estimate
2	**Target**	A parameter's event, at a set control
	Declare	$\mathbf{u}$ hall B's set point; $\vartheta_A(\mathbf{u}) = \mathbf{u} + 50\,\text{kW} \times R$; $\vartheta_{\max} = 30°\text{C}$ (illustrative)
3	**Known**	E; do($\mathbf{u} = 22°\text{C}$); unit A out; 50 kW; settled; the limit mirrored onto hall A (assumption)
	Random here	R
4	**Event**	$\vartheta_A(\mathbf{u}) > \vartheta_{\max}$
5	**Expression**	$P(\vartheta_A(\mathbf{u}) > \vartheta_{\max} \mid E, \text{do}(\mathbf{u} = 22°\text{C}))$
6	**Model**	The opening's relation, the heat now running from A to B; the same R
	Reduce	$P(50\,\text{kW} \times R > 8\,\text{K} \mid E)$
7	**Check**	0.008; the fact sheet states no limit on A above B, so the assumption is the sheet's
	Class	Probability estimation; here a posterior exceedance

$$\boxed{P(50\,\text{kW} \times R > 8\,\text{K} \mid E)}$$

Run, (10) the largest load at 1 %

1	**Sentence**	Set point 22°C, a unit's loss, given E: the largest load on hall B with the risk at most 1 %?
	Intention	Optimize
2	**Target**	A setting, here a load, under a chance constraint
	Declare	Q_B set [kW]; $\vartheta_B = 22°\text{C} + Q_B R$; $\vartheta_{\max} = 30°\text{C}$ (illustrative); $R_{0.99} = 0.1586$ K/kW
3	**Known**	E; do(Q_B); unit B out; settled
	Random here	R; the load is set
4	**Event**	$Q_B R > 8\,\text{K}$
5	**Expression**	$\max Q_B$ s.t. $P(Q_B R > 8\,\text{K} \mid E, \text{do}(Q_B)) \leq 0.01$
6	**Model**	The risk rises with the load, so the constraint binds
	Reduce	$Q_B = 8\,\text{K}/R_{0.99}$
7	**Check**	50.4 kW: the same 0.99 point as the third rung, with the load and not the set point as the control
	Class	Chance-constrained optimization

$$\boxed{\max Q_B \quad \text{s.t.} \quad P(Q_B R > 8\,\text{K} \mid E, \text{do}(Q_B)) \leq 0.01}$$

Appendix B

The scoring rubric

A formulation cannot be checked by computing it, so it is checked by this rubric, whether the reader scores their own work or a second reader scores it. The rubric crosses the seven steps of the method with the failure classes that recur across the book.

Scoring a formulation

Each of the seven steps scores 2 when it is complete and correct, 1 when it is present but incomplete, and 0 when it is missing or wrong, for a total out of 14. The intention at the top of the sheet and the problem class at the bottom are marked right or wrong and are reported beside the total. Two steps gate the result: a formulation is accepted only when step 5 and step 7 both score 2, because an expression that does not answer the question is not rescued by a good model, and step 7 is the only check the reader has without a computation. A question that is not ready has a different correct answer. It scores full marks when step 1 names each missing item and the answer gives the expression for each reading; the other steps are not required for every reading. Every point lost is labeled with a failure class where one applies.

The seven-step score applies to the *formulate* exercises that ask for a full formulation. A formulate exercise that asks only for some parts, such as an event built from its words, the four parts of a requirement, or what an event is about, is scored on those parts alone, 2 for each part complete and correct, 1 when incomplete and 0 when missing or wrong; the gate on steps 5 and 7 applies only when an expression is asked for. The other kinds are scored more simply. A *reverse* exercise scores 2 when the sentence says what the expression asks about, what it takes as known and its horizon, 1 when one of these is wrong or missing, and 0 when the sentence reads the expression as a different question, such as a parameter read as an outcome. A *critique and repair* exercise scores 1 for naming the question the wrong expression does answer, 1 for naming its failure class, and 2 for a repair that passes step 7, out of 4. A *rapid* set scores one point for each item whose every part is right. A *machine set* is scored problem by problem, out of 20: 2 when the run's box answers the sentence and its lines hold, 1 when the box is right and a line is wrong or missing, and 0 when the box answers another question.

The recurring failure classes

Table B.1 defines the eight recurring classes, each with an instance on the two-hall plant of Appendix D and the question that the wrong expression does answer. Each lesson owns at least one failure class, some of them particular to that lesson, and adds its own to this list when it is written.

Table B.1: The eight recurring failure classes, with an instance of each on the two-hall plant.

Failure class	Instance	What the wrong expression answers instead
Horizon missing	A probability of "no interruption" written with no interval	The state at one time, not the path over a period
Wrong side of the bar	Three forms (Lesson 3). The transposed conditional: $P(E \mid G)$, for $G \geq G^*$, reported for $P(G \geq G^* \mid E)$. A known condition written into the event: $P(F_A \cap F_B \mid E)$ for $P(F_A \mid F_B, E)$. An assumption behind the bar reported as a fact	How expected the readings are if the opening is adequate; how often the two outages happen together; the question for a plant on which the assumption holds
Parameter against outcome	Four forms (Lesson 7). A parameter for an outcome: $P(G \geq G^* \mid E)$, 0.82, reported as the chance of service, which is 0.971. A plug-in at a point value for the predictive. A band on a parameter read as a band on the outcome's probability. A hypothesis's probability read as the probability of the outcome it favors	Whether the opening is adequate, and nothing about whether a unit will fail; the outcome for a plant whose parameter is known; what is known of the parameter; the weight of one structure
Rule against posterior	Four forms (Lesson 14). Accept read as true: $A(E) = 1$ reported as "the claim is true". Not accepted read as false. A level read as the probability delivered. A posterior read as a decision with no rule	What a rule did with this evidence, not how probable the claim is; how probable the claim is, with no rule to act on it
Realized against prospective evidence	Five forms (Lesson 16). $P(A(E) = 1 \mid G)$ written with E in hand. The expected reading plugged in: the design case read as power. A rule replaced by a threshold that no longer holds the posterior. One truth used where the truth is unknown. The inner and outer G confounded	Nothing, in the first: with the readings fixed, the rule's value is already known, and the expression should be written over $\tilde{E}$. A different question in the others: a posterior at one reading; another rule's curve; a statement for a plant at that truth
Dependence ignored	Four forms (Lesson 6). A shared component left out: $P(F_A \cap F_B)$, given the stated model, reported as the chance that both units are unavailable, with the feeder's outage, which takes both at once, left out. A shared cause's tie dropped: the product of the marginals on the window model. A demand-conditioned failure at its unconditional probability: the UPS's unavailability for repair used for its failure to carry an interruption. A propagation as two separate events	The chance that both units are down for their own reasons only; the joint of parts independent within a window as if they were independent across windows; how often a supply loss and a UPS repair overlap; how often two faults overlap

Table B.1: The eight recurring failure classes, with an instance of each on the two-hall plant.

Failure class	Instance	What the wrong expression answers instead
Intervention read as observation	Four forms (Lesson 17). A probability given a set point found in the log, used for the set point chosen. A planned action read as evidence. A decision updating belief. A relation assumed to survive do(·) when only the log has been seen	What usually happens when that value is seen, not what follows from setting it; the belief before the plan or the decision
Model conditioning dropped	Five forms (Lesson 10): a relation or discrepancy assumption left off; a Gaussian reading's coordinate left off (0.77 for 0.818); a prior replaced in a change of variable; the Jacobian dropped; a rule of thumb without its condition	The same question with its assumption hidden, as if it were a fact; a different model or reading, reported as the same

The classes owned by single lessons

Table B.2 adds the classes that belong to one lesson each, in the order of the lessons, with the step at which each enters and an instance on the two-hall plant. Each is caught, like the recurring ones, by reading the expression back at step 7.

Table B.2: The failure classes owned by single lessons, with an instance of each on the two-hall plant.

Lesson	Failure class	Enters	Instance	What the wrong expression answers instead
2	Connective mistranslated	Step 4	$F_A^c \cap F_B^c$ written for "the units are not both out"	Whether neither unit is out, a plant with no spare in use
4	Requirement without its conditions or its level, including failure defined at the wrong level of function	Step 1 or 4	The availability, 0.988, compared with a clause of 99.5 % of each month's hours with feeder outages excepted	The mean share of hours with the exclusion dropped, against no stated probability level
5	Point value read against the limit: a margin as the verdict or as a probability; an envelope at the point value; one limit state for several	Step 4 or 5	A margin of 1.47 K at the point value reported as "safe", with hall B at 55 kW and unit B out	How far the plant is from the boundary if its conductance equals the point value; the exceedance probability is 0.056
8	"At a time" read as "by a time": the state at the horizon for a path, or the reverse; a crossing on the value at the horizon; a mean time for a horizon	Step 4	"Hall B was at 27°C at 06:00, so it never exceeded 30°C in the night"	The state at one moment, not whether a quantity that falls again ever crossed its limit, which its running maximum records

Table B.2: The failure classes owned by single lessons, with an instance of each on the two-hall plant.

Lesson	Failure class	Enters	Instance	What the wrong expression answers instead
9	Several questions folded into one expression: folding; one quantity counted as two; a look-alike; a condition's scope	Step 1	The opening adequate and hall B within 8 K at full load, each 0.82, multiplied to 0.67	Two questions where the paragraph held one; the event is 0.82
11	Reading taken for the quantity: the reading for the quantity; noise kept and a nuisance variable dropped; an offset read as noise; a count read as a rate	Step 5 or 6	σ_d taken as 0.08 K, with the sensors' offsets dropped	What the test would say with sensors known to have no offset; 1.000 for 0.818
12	One record counted twice: in the prior and the likelihood; a source containing the item; a prior tuned after the evidence; two documents quoting one source	Step 3	A prior set to what unit A's log already says, and the log then used as evidence	The band of a record twice as long that nobody kept: 1.45 to 3.43 for 1.18 to 3.90
13	A component event substituted for the system event: a part's event for the system's; parts' probabilities combined as a product; a derived quantity given independent standing; a sub-diagram's answer for the whole	Step 4, 5 or 6	A unit's availability given as the plant's	What one part does, or a product of parts: 0.951 or 0.959 for 0.971
15	Posterior risk read as a sampling error rate: the level read as a sampling rate; a sampling rate read as a posterior risk; a posterior value read as a frequency, or a count as a probability of adequacy; a population average or an outcome where a truth or a reading is stated	Step 3 or 5	"We accept at 95 %, so 5 % of inadequate plants get through"	How much posterior risk the rule allows on an accepted reading; the rate is 0.010 at 54 kW and 0.065 at G^{*}, from a 40 kW test

Where each class enters

Table B.3 marks the steps at which each class usually enters a formulation. Reading the expression back at step 7 is where each of them should be caught, so that column is marked for every class. A class found at step 7 but entered earlier is repaired at the step where it entered.

Table B.3: The seven steps crossed with the recurring failure classes. A filled circle marks a step at which the class usually enters, and an open circle the step at which it is caught. The columns are step numbers.

Failure class	Step 1	Step 2	Step 3	Step 4	Step 5	Step 6	Step 7
Horizon missing	●			●			○
Wrong side of the bar			●		●		○
Parameter against outcome		●			●		○
Rule against posterior		●			●		○
Realized against prospective evidence			●		●	●	○
Dependence ignored				●		●	○
Intervention read as observation			●		●		○
Model conditioning dropped			●			●	○

Appendix C

The blank formulation sheet

Table C.1 is the blank sheet, with the seven steps of the method as fields, the engineering intention at the top and the problem class at the bottom. It is meant to be copied and filled in by hand, one sheet per question.

Table C.1: The formulation sheet. The fields hold words and expressions.

The question as heard.
Engineering intention (circle one): estimate predict compare decide optimize diagnose design an experiment
Step 1. State the engineering question in one unambiguous sentence.
Step 2. Identify the target quantity: a state, parameter, event, future outcome, or decision.
Step 3. Specify what is known: observations, assumptions, operating conditions, and prior information.
Step 4. Define the event or objective mathematically, including thresholds and time horizons.
Step 5. Write the probability or expectation that directly answers the question.
Step 6. Identify the model required to evaluate that expression: variables, physical relations, dependencies, and evidence.
Step 7. Check the formulation: does the expression really answer the original engineering question?
Problem class, for example: posterior exceedance; predictive probability; expected-loss minimization; chance-constrained optimization; preposterior design.

Appendix D

The fact sheets of the recurring systems

Each recurring system is stated here in full, so that a lesson can use it without any other source. A fact sheet records what is stated about a system. It does not say what to ask of it. The two-hall plant, the small electrical network and the battery string are given below. The systems are separate: the network's feeder is not the two-hall plant's, and the battery string is not the network's UPS.

D.1 The two-hall cooling plant

The system. Figure D.1 draws the plant: two boxes for the halls, each holding its IT load and its cooling unit; a double arrow for the opening in the wall between them; and dashed arrows from the one electrical feeder to the two units. Two data halls, A and B, each carry an IT load of 60 kW in service. Each hall has one cooling unit with a capacity of 130 kW, enough to carry both halls. Heat crosses the wall opening in proportion to the temperature difference between the halls, with a conductance G that nobody has measured. One electrical feeder supplies both units, and when it is out, both units are out.

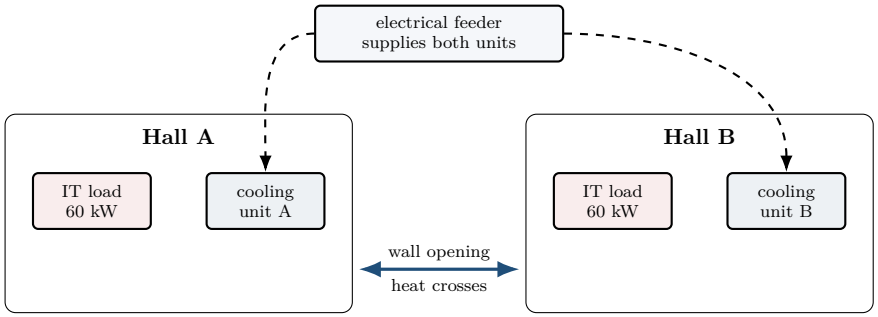


Figure D.1: The two-hall plant. Each hall holds a 60 kW IT load and one cooling unit; heat crosses the wall opening between the halls with conductance G; one electrical feeder supplies both units.

The limit. With one unit out, the other hall's heat must cross the opening. The stated limit is that hall B may not run more than 8 K above hall A, whose set point is 22°C. The source derives the 8 K as the span from that set point to the top of the ASHRAE allowable range for IT equipment; this book treats it as the stated limit. No limit is stated for hall A above hall B. The conductance that passes hall B's full load, with unit B out and the plant settled, at that

difference is

$$G^{*} = \frac{Q_{\text{hall}}}{\Delta T_{\text{allow}}} = \frac{60\,\text{kW}}{8\,\text{K}} = 7500\,\text{W/K}.$$

The two structures. The plant has one of two success structures, depending on the opening. If the opening passes a hall's load, the units form a parallel pair: the service is lost only when both units are down or when the feeder is down. This is the structure called H_{AND}, from the AND gate of its fault tree. If the opening cannot, the units form a series pair: the service is lost when either unit is down or the feeder is down. This is H_{OR}. The drawing claims the first, which is the N+1 claim. The candidate drawings carry heat-flow capacities of 30 and 80 kW (a far pair) or 54 and 66 kW (a near pair), which are conductances of 3750 and 10 000 W/K, or 6750 and 8250 W/K, at the allowable difference.

The trip test. A test takes unit B out, so that hall B's heat must cross the opening to unit A, and reads the plant after it settles. It runs at 20, 40 or 60 kW per hall. It logs each hall's temperature once a minute for 30 minutes, the duty of the surviving unit, and the two PDU powers. Each temperature sensor has noise of 0.3 K per sample and a constant offset that averaging does not remove, with a prior spread of 0.5 K. Together they put a floor of 0.71 K on the spread of a measured hall difference. Unit A's duty meter has a noise of 2 % and the PDU meters 1 %.

The records. The plant keeps an outage record for each component, a count of failures over an exposure with each repair logged. The stated rates give each unit a failure rate of 2.5 per year in service and a mean repair time of 180 h, and the feeder 3.69 per year in service and 24 h; a rate per year in service counts failures per 8760 h of the component's own up time. These are unavailabilities of about 0.05 and 0.01. The plant may also have a record of disturbance windows, 720 h each, in which a shared cause raises the unavailabilities to 0.30 per unit and 0.08 for the feeder, with a prior weight of 0.1 on a window being disturbed. In a calm window the unavailabilities are the rounded values above, 0.05 per unit and 0.01 for the feeder; Lessons 3 and 6 use the calm and disturbed values as given here, and say so where they do.

Table D.1 collects the stated values. Everything in the table is stated, and the unknowns of the plant are the conductance, the sensor offsets, the rates of this particular plant, and its structure. The lessons add a few things that the source does not state, each labeled where it is used: exponential repair lengths; a 5 kW test; a fleet list of forty units, a sister plant and a trip test of last year, all illustrative; a prior on the hall loads, which is the analyst's choice; and, in Lesson 17, an absolute limit of 30°C and a lowest set point of 18°C, a coefficient of performance and a plant electricity, loss tables, and a constructed operating log, all illustrative.

Table D.1: The stated values of the two-hall plant.

Quantity	Unit	Value
IT heat load per hall in service, Q_{hall}	kW	60
Cooling-unit capacity	kW	130
Stated limit on hall B above hall A, ΔT_{allow}	K	8
Set point	°C	22
Conductance threshold, G^*	W/K	7500
Trip-test loads per hall, Q	kW	20, 40, 60
Temperature sensor noise per sample	K	0.3
Samples per test	–	30
Prior spread of a sensor's offset	K	0.5
Prior spread of the set point	K	2
Floor on the spread of a measured difference	K	0.71
Duty meter, PDU meter noise	–	2 %, 1 %
Failure rate of a unit, of the feeder, per year in service	1/yr	2.5, 3.69
Mean repair time of a unit, of the feeder	h	180, 24
Unavailability of a unit, of the feeder	–	about 0.05, about 0.01
Prior on $\log G$: center	W/K	7500
Prior on $\log G$: spread	–	a factor of 3
A drawing's tolerance on its conductance	–	2 %
Prior on the log rate of a unit, of the feeder: center	1/yr	2.5, 3.69
Prior on the log mean repair of a unit, of the feeder: center	h	180, 24
The four priors above: spread	–	a factor of 3
Record lengths	unit-yr	1, 3, 10
Failures logged, unit	–	2, 7, 25
Failures logged, feeder	–	4, 11, 37
Downtime in the 3-year record, each unit, the feeder	h	1260, 264
Disturbance window	h	720

What is constructed. The plant comes from the case study *Reliability of a cooling service from trip tests, outage logs and disturbance windows, by probabilistic inference on a two-hall data center*. Its readings and records are synthetic, generated from stated models with stated noise. No real two-hall record exists. Its thermal physics is steady and linear, read after settling.

D.2 The small electrical network

The system. Figure D.2 draws the network as a single line from left to right. The utility feeder supplies the facility through the incoming breaker, whose protective relay opens it on a fault and can also open it spuriously. The UPS, a double-conversion unit with a battery and a static bypass (the bypass and the battery are this book's additions, illustrative), takes that supply and feeds the PDU, which feeds two branches, each on its own breaker: the IT racks, 900 kW, and the cooling plant, 250 kW. While the UPS's supply is lost, the UPS carries the load from its battery for its autonomy; if the UPS itself fails, the static bypass feeds the load directly from the supply. The circles marked m are the five meters: the utility input, the UPS output, the PDU, the IT load and the cooling load.

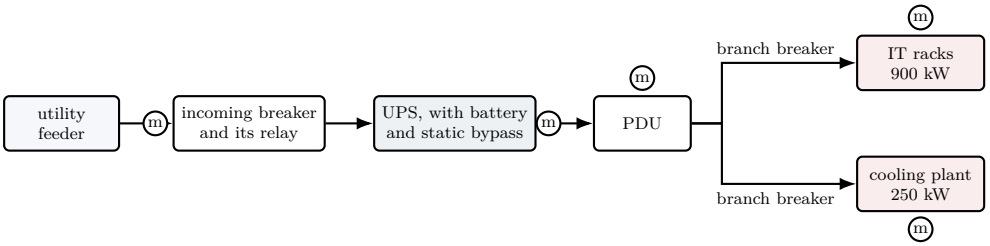


Figure D.2: The small electrical network: the feeder, its protection, the UPS and the two load branches. The circles marked m are the five meters.

Protection. The protection's layout is this book's, illustrative, since the source has none. A fault on a branch is cleared by that branch's breaker, so that the other branch stays supplied. If the branch breaker fails to clear its fault, the breaker above both branches opens, and both are lost until it is reset. The incoming relay's spurious trips are losses of supply like the feeder's interruptions, short ones, which the UPS must cover; their times to reclose are exponential, like the interruptions' lengths.

The meters and their explanations. At nominal efficiency the UPS turns 1197.9 kW from the utility into the 1150 kW its loads draw. When the meters do not add up, three explanations are stated, one fault at a time: nominal operation with meter noise, a degraded UPS efficiency, and an offset in the PDU meter, with prior weights 0.90, 0.05 and 0.05. Each explanation leaves its own quantity uncertain. Under a degraded UPS the efficiency η is about 0.94 ± 0.04, and under the other two it is 0.96 ± 0.004. Under the meter-offset explanation the PDU meter's offset is about 0 ± 45 kW, and under the other two it is 0 ± 1.5 kW. The source's three scenarios set a degraded efficiency of 0.91 and an offset of +42 kW as their truths; those are the values its readings were generated from, not values an explanation assumes.

Table D.2 collects the values. Its last column says where each comes from: stated by the source study, derived from stated values, or illustrative, supplied by this book because the source does not state it.

Table D.2: The values of the small electrical network and their source.

Quantity	Unit	Value	Source
IT load; cooling plant load	kW	900; 250	stated
UPS efficiency, nominal	–	0.96	stated
Utility input at nominal efficiency	kW	1197.9	derived
Meter noise: utility, UPS output, PDU, IT, cooling	kW	3, 2, 2, 3, 3	stated
Prior weights: nominal, UPS degraded, PDU meter offset	–	0.90, 0.05, 0.05	stated
UPS efficiency under a degraded UPS; under the other explanations	–	0.94 ± 0.04; 0.96 ± 0.004	stated
PDU offset under the offset explanation; under the others	kW	0 ± 45; 0 ± 1.5	stated
Scenario truths: a degraded efficiency; a PDU meter's offset	–; kW	0.91; +42	stated, as truths
Feeder interruptions; mean length, exponential	1/yr; h	2.0; 1	illustrative
Spurious trips; mean time to reclose, exponential	1/yr; min	0.5; 5	illustrative
Faults per branch; mean repair	1/yr; h	0.1; 4	illustrative
A branch breaker fails to clear its fault	per demand	0.01	illustrative
Reset of the breaker above both branches	h	1	illustrative
UPS fails to pick up the load	per demand	0.02	illustrative
UPS battery autonomy at the full 1150 kW, t_{aut}	min	10	illustrative
UPS failures; mean repair, the load on the bypass	1/yr; h	0.2; 12	illustrative

Rates are per year in service, as for the two-hall plant, and every component's failures are independent of the others' unless a lesson states a tie.

What is constructed. The network comes from the case study *Power chain loss localization*, a steady snapshot of a facility's power chain in which three explanations of a meter mismatch are ranked from five readings. Its chain, loads, efficiency, meters and explanations are stated there, with synthetic readings. It has no time history, so every rate, the protection's layout, the static bypass, the battery and its autonomy, and the per-demand probabilities are this book's illustrative inputs. The sheet states no fault upstream of the UPS, no record, and no tie among the parts beyond those above.

D.3 The battery string

The system. When the grid fails, a single series string of cells, the battery string, carries a steady critical load of 20 kW by itself. Figure D.3 draws it: the grid feed, marked as lost; the string as a tank, drawn between its state of charge at the outage's start and its reserve floor, with the usable energy between them; the load; and the waste heat that discharging sends through the pack to the room. The string may be discharged from its state of charge at the outage's start, 0.80 of its capacity, down to the floor, 0.20, which leaves 45 kWh of a 75 kWh nameplate. Discharging at an efficiency of 0.90 draws 22.2 kW from the string to deliver 20 kW, so 40.5 kWh reach the load, and the string carries it for 2.025 h before reaching its floor. The 2.2 kW lost in the discharge heats the pack, which leaks the heat to a room held at 300 K; the pack's limit is 318 K.

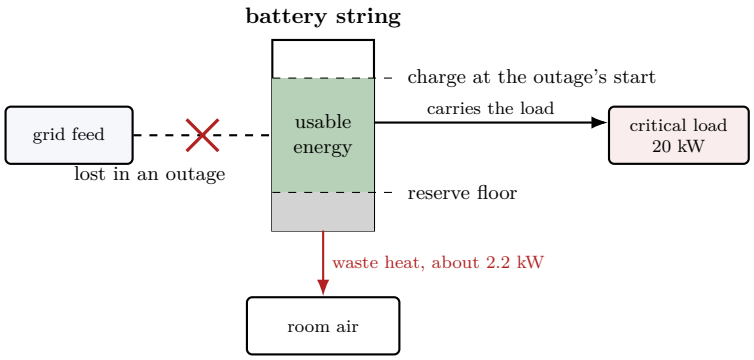


Figure D.3: The battery string carrying its critical load through a grid outage, with the usable energy between the state of charge at the outage's start and the reserve floor.

What is uncertain. The source states the string as new and certain. This book adds what an owner of a string in service does not know. The string's state of health, its usable energy as a share of a nameplate string's, is uncertain, and for a sound string it is taken as equally likely anywhere between 0.85 and 1.00. One cell may be weak, with prior probability 0.10, and a string with a weak cell keeps 0.70 of a sound string's usable energy. Grid outages arrive about 1.5 times a year with lengths spread exponentially about a mean of one hour. The state of charge at an outage's start is read with an error of standard deviation 0.03. A site specification asks for two hours at the critical load, a value set just under the nameplate's 2.025 h so that the lessons can show a specification usually missed beside outages usually carried. All of these are illustrative.

Table D.3 collects the values with their source.

Table D.3: The values of the battery string and their source.

Quantity	Unit	Value	Source
Nameplate capacity	kWh	75	stated
State of charge at the outage's start; reserve floor	–	0.80; 0.20	stated
Discharge efficiency	–	0.90	stated
Critical load	kW	20	stated
Usable energy above the floor; deliverable after the loss	kWh	45.0; 40.5	stated
Ride-through at nameplate	h	2.025	stated
Waste heat while carrying the load	kW	2.22	stated
Pack temperature limit; room temperature	K	318; 300	stated
Pack thermal mass; pack-to-room heat leak	J/K; W/K	3.0×10^6; 60	stated
Peak pack temperature in an outage of 1.5 h; of 3 h	K	303.8; 307.2	stated
State of health of a sound string, φ, uniform	–	0.85 to 1.00	illustrative
A weak cell: prior probability; usable energy left	–	0.10; 0.70	illustrative
Grid outages; mean length, exponential	1/yr; h	1.5; 1	illustrative
Error of the reading of the state of charge	–	0.03	illustrative
Required ride-through at the critical load	h	2	illustrative

With these inputs the ride-through is $T_{\mathrm{rt}} = \varphi \times 2.025$ h for a sound string, between 1.72 and 2.03 h, and 0.70 of that with a weak cell, between 1.20 and 1.42 h.

What is constructed. The string comes from the case study *How long can this battery carry the load?*, a deterministic simulation of one pack through outages of 1.5 and 3 h, checked against a one-line energy balance. It has no sensors and no uncertainty. The book calls the pack a string, a single series string of cells, and supplies everything uncertain on the sheet as illustrative inputs. The sheet states no test of the string, no failure of the string as a component, and no ageing over time.

Appendix E

How common formulations are evaluated, by problem class

The problem class at the bottom of a formulation sheet is the hand-off from this book to the techniques of other textbooks. This appendix is organized by that class. For each class it will say in a few lines how the expression is commonly evaluated and where the standard treatment is found. It is the only place in the book where solution techniques are named in any detail, and it derives none of them. Table E.1 lists the classes the lessons have named so far. The appendix is completed with the last lesson.

Table E.1: The problem classes and the form of each.

Problem class	Form
Posterior exceedance	$P(g(\mathbf{z}) \geq 0 \mid E)$ for a parameter or a state
Predictive probability	The probability of a future outcome, averaged over the posterior of what is unknown
Probability estimation	The probability of a stated event at a stated setting, $P(\text{event} \mid E, \text{do}(\mathbf{u}))$; on a parameter's event, a posterior exceedance
Decision rule	$A(e)$, a function from evidence to an action, and its value on evidence in hand
Operating characteristic at a stated truth	$P(A(\tilde{E}_Q) = 1 \mid G = G_1)$, a probability over readings not yet taken, with the truth held
Assurance	$\int P(A(\tilde{E}_Q) = 1 \mid G)\, p(G)\, \mathrm{d}G$, the operating characteristic averaged over what is believed about the truth
Predictive value of a test's outcome	$P(G < G^* \mid A(\tilde{E}_Q) = 1)$, a state given an outcome, under a stated population
Expected-loss minimization	The action minimizing the expected consequence given E
Risk optimization	$\arg\min_{\mathbf{u}} P(\text{a stated event} \mid E, \text{do}(\mathbf{u}))$, with no cost in the objective; a bound usually decides it
Chance-constrained optimization	An objective optimized over $\mathbf{u}$, subject to a stated event holding with at least a stated probability
Preposterior design	An objective over tests, taken over $\tilde{E}$ before the readings exist
Posterior of a parameter	$p(\theta \mid E)$, the whole density
Joint posterior	$p(\theta_1, \theta_2 \mid E)$, kept whole because the parameters are tied
Likelihood	$p(E \mid \theta)$ as a function of θ, with any nuisance integrated out
Prior	A distribution on a quantity, stated before the evidence
Prior probability of an event	$P(\mathcal{A})$ under the prior, with no evidence
Posterior probability over candidate causes	$P(H_i \mid E)$, a weight on each explanation
Factors of a query	The set of factors a query needs, found on the diagram; nothing is evaluated
Not ready	No evaluation; the hand-off goes back to the requester, with what is missing