

Control Engineering by Derivation

Deriving the control law, and the move under a limit and over a horizon

Draft · 266 pages

Ricardo Marquez, Ph.D.

ricardo-marquez.com

Control Engineering by Derivation

Volume II

The law, the limits and the horizon

Ricardo Marquez

November 2, 2026 (work in progress: chapters are added as they are written)

How to read this book

This book teaches one skill: to take a control engineer's question, write it as a probability or an expectation over the unknowns of a physical model, derive what that expression requires, compute it, check it, and act on it. Volume I holds Parts I to III, Chapters 1 to 11: the method, the closed loop on the graph, the control problem every later chapter works from, and the belief the controller acts on. This second volume holds Part IV, which derives the control law from that problem, and Part V, which adds limits and a horizon. Chapter numbers run continuously across the volumes; a reference to a chapter numbered 11 or below points to Volume I, whose results this volume builds on and does not repeat.

Every chapter follows the progression fixed in Volume I's preface: step 0, the plant graph; step 1, the question as a probability or expectation over the state; step 2, the graph extended for the question and the records attached as rows; step 3, the posterior of the state, derived and computed once; step 4, the answers, each read from that posterior down to its number; step 5, evaluation and checks against the full calculation and the constructed truth; step 6, the report with its band; step 7, the next measurement, ranked by the width it removes. Each chapter also names which restriction of the problem of Chapter 7 it works under, and what that gives up.

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

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

Contents

How to read this book **ii**

Notation and vocabulary **vi**

IV Deriving the Control Law **1**

12 Certainty equivalence and the linear-quadratic law **2**
12.1 A vendor's gain beside a posterior 2
12.2 The posterior of the state (step 3) 5
12.3 The answers (step 4) 5
12.4 The checks (step 5) 18
12.5 The report (step 6) 19
12.6 The next measurement (step 7) 21
12.7 What to retain 22
12.8 Exercises 23

13 Integral action **25**
13.1 A flat trend and a sheet with no integrator 25
13.2 The posterior of the state (step 3) 28
13.3 The answers (step 4) 29
13.4 The checks (step 5) 41
13.5 The report (step 6) 42
13.6 The next measurement (step 7) 43
13.7 What to retain 45
13.8 Exercises 45

14 Tuning as an objective **47**
14.1 Two tunings and no objective 47
14.2 The posterior of the state (step 3) 50
14.3 The answers (step 4) 51
14.4 The checks (step 5) 62
14.5 The report (step 6) 63
14.6 The next measurement (step 7) 64
14.7 What to retain 67
14.8 Exercises 68

15 Saturation and anti-windup **70**
15.1 The loop at its stop 70
15.2 The posterior of the state (step 3) 74
15.3 The answers (step 4) 76

15.4 The checks (step 5) . . . 90
15.5 The report (step 6) . . . 90
15.6 The next measurement (step 7) . . . 92
15.7 What to retain . . . 94
15.8 Exercises . . . 95

V Limits and the Horizon **97**

16 Chance-constrained set points **98**
16.1 A set point the sheet would raise . . . 98
16.2 The posterior of the state (step 3) . . . 102
16.3 The answers (step 4) . . . 103
16.4 The checks (step 5) . . . 115
16.5 The report (step 6) . . . 116
16.6 The next measurement (step 7) . . . 117
16.7 What to retain . . . 120
16.8 Exercises . . . 120

17 Expected-cost set points **122**
17.1 A probability on the sheet and no consequence . . . 122
17.2 The posterior of the state (step 3) . . . 124
17.3 The answers (step 4) . . . 125
17.4 The checks (step 5) . . . 141
17.5 The report (step 6) . . . 141
17.6 The next measurement (step 7) . . . 143
17.7 What to retain . . . 145
17.8 Exercises . . . 145

18 MPC as a restriction **147**
18.1 A plan on the commissioning sheet . . . 147
18.2 The posterior of the state (step 3) . . . 152
18.3 The answers (step 4) . . . 153
18.4 The checks (step 5) . . . 164
18.5 The report (step 6) . . . 164
18.6 The next measurement (step 7) . . . 167
18.7 What to retain . . . 170
18.8 Exercises . . . 170

19 Forecast uncertainty and tightening **172**
19.1 A flat trend and a tube on the sheet . . . 172
19.2 The posterior of the state (step 3) . . . 176
19.3 The answers (step 4) . . . 178
19.4 The checks (step 5) . . . 192
19.5 The report (step 6) . . . 194
19.6 The next measurement (step 7) . . . 195
19.7 What to retain . . . 197
19.8 Exercises . . . 198

20 Recourse and the receding horizon **199**
20.1 The plan restated . . . 199
20.2 The posterior of the state (step 3) . . . 203
20.3 The answers (step 4) . . . 204
20.4 The checks (step 5) . . . 216
20.5 The report (step 6) . . . 217
20.6 The next measurement (step 7) . . . 219
20.7 What to retain . . . 221
20.8 Exercises . . . 222

21 Thermal and electrical storage **223**
21.1 Two stores and a column on the sheet . . . 223
21.2 The posterior of the state (step 3) . . . 227
21.3 The answers (step 4) . . . 228
21.4 The checks (step 5) . . . 241
21.5 The report (step 6) . . . 242
21.6 The next measurement (step 7) . . . 244
21.7 What to retain . . . 246
21.8 Exercises . . . 246

A Atlas of control questions **248**

Notation and vocabulary

Every symbol below has the same meaning in every chapter, except where a chapter's table says otherwise. The unknown always stands on the left of a conditioning bar and the evidence on the right. Vectors and matrices are bold. Symbols of the larger plants are added here as the chapters that introduce them land.

Symbol	Unit	Meaning
The plant graph and its state		
$\mathbf{z}$, z_i	varies	State vector of the plant graph: every free quantity of one telemetry window at once; its i-th entry.
$\mathbf{z}_k$, k	varies, –	The state at controller step k; the step index, each step 30 s on the closed-loop hall.
$\mathbf{y}$, y	varies	The evidence: all readings and prior information treated as known; a single reading, subscripted by what it reads (y_{T_s}, y_{Q_A}).
$r_k(\mathbf{z})$	varies	Row k of the model: a physical relation that must hold ($r_k = 0$), a reading compared with the state, or a prior.
w_k, σ_k	varies	Weight (precision) of row k and the standard deviation it corresponds to, $w_k = 1/\sigma_k^2$. A row may equivalently be written normalized, r_k/σ_k, with unit weight.
σ_p	varies	Tolerance of a physics row; an exact constraint is the limit $\sigma_p \to 0$.
The equipment hall of Part I		
T, T_c	°C	Temperature of the hall air; the supply air held by the cooling unit's discharge controller, pinned.
Q, G	kW, kW/K	Heat load of the hall; effective conductance of its air path to the supply air.
g, G_{ref}	–, kW/K	Log-conductance $\ln(G/G_{\text{ref}})$, the entry of the state that carries the conductance; its reference value.
y_T, σ_T	°C, K	The hall sensor's reading and the standard deviation of its error.
The closed-loop hall of Part II		
r, e	°C, K	The set point of the discharge loop; the control error $e = y_c - r$ on the discharge sensor's reading.
c, u	–	The valve command of the PI controller; the valve position, from 0 to 1, between its stops.
K_p, K_i	1/K, 1/(K s)	Proportional and integral gains.
K_Q	K/kW	Process gain of the hall to its load, $\partial T/\partial Q = 1/G$.

Symbol	Unit	Meaning
W, g_w, T_w	kW/K, –, °C	Coil conductance to chilled water with the valve fully open; its log-derate; the chilled-water temperature, pinned.
C_h, C_c, τ_a	kJ/K, kJ/K, s	Thermal capacity of the hall air; of the coil and supply air; the actuator's time constant.
$T_{\max}$	°C	A limit on a temperature.
f^*, G^*, T^*	varies	A requirement level on a state entry or a derived quantity; a starred plant quantity is also a scenario value, the quantity at a set point not yet applied.
b_j	varies	Bias of reading j, an entry freed when a sensor is suspected.
Probability		
$\Pr(A)$, $\Pr(A \mid B)$	–	Probability of event A; conditional on event B.
$p(\cdot)$	varies	A probability density; $p(\mathbf{z} \mid \mathbf{y})$ is the posterior of the state, $p(\mathbf{y} \mid \mathbf{z})$ the likelihood, $p(\mathbf{z})$ the prior.
$\mathbb{E}$, Var, Cov	varies	Expectation, variance, covariance.
$\mathcal{N}(\mu, \sigma^2)$	–	Normal distribution with mean μ and variance σ^2; Φ its standard cumulative distribution function, φ its density; $\Phi_2(u_1, u_2; \varrho)$ that of two standard normals of correlation ϱ.
χ^2_ν	–	Chi-square distribution with ν degrees of freedom.
Inference objects		
$\mathcal{L}(\mathbf{z})$	–	The objective, minus the log posterior up to a constant: $\mathcal{L} = \frac{1}{2} r(\mathbf{z})^\mathsf{T} r(\mathbf{z})$ for normalized Gaussian rows.
$\hat{\mathbf{z}}$	varies	The maximum a posteriori (MAP) state, the minimizer of $\mathcal{L}$.
$\mathbf{J}$	varies	Jacobian of the normalized rows with respect to the state, evaluated at $\hat{\mathbf{z}}$.
$\mathbf{\Lambda}$	varies	Information matrix (Gauss–Newton Hessian of $\mathcal{L}$), $\mathbf{J}^\mathsf{T}\mathbf{J}$ for normalized rows, $\mathbf{J}^\mathsf{T}\mathbf{W}\mathbf{J}$ with $\mathbf{W} = \operatorname{diag}(w_k)$ for unnormalized ones.
$\mathbf{\Sigma}$, Σ_{ij}	varies	Posterior covariance of the state under the Laplace approximation, $\mathbf{\Lambda}^{-1}$; its entries. The posterior spread of entry i is written $\sqrt{\Sigma_{ii}}$; σ_k always names a row's standard deviation.
$f(\mathbf{z})$, $\mathbf{g}$	varies	A derived quantity of the state and its gradient $\partial f / \partial \mathbf{z}$ at $\hat{\mathbf{z}}$.
$\mathbf{A}$, $\mathbf{A}_{\text{free}}$, $\mathbf{A}_{\text{dep}}$	varies	Jacobian of the physics rows on the full state; its blocks on the free and the dependent entries (Chapter 4).
$\mathbf{N}$, $\mathbf{J}_{\text{rows}}$, $\mathbf{J}_{\text{red}}$	varies	Null-space basis of the physics rows, $[\mathbf{I};\ -\mathbf{A}_{\text{dep}}^{-1}\mathbf{A}_{\text{free}}]$; the Jacobian of the evidence rows (readings and priors); the reduced Jacobian $\mathbf{J}_{\text{rows}}\mathbf{N}$ of the reduced state.
f_L, f_U	varies	The ends of the credible interval of a derived quantity.
y_{new}, $\tilde{y}_a$	varies	A future reading (the predictive) and the outcome of a candidate reading a before it is taken (the preposterior).
e_k	–	Normalized residual of reading k: its residual divided by the standard deviation of that residual.

Symbol	Unit	Meaning
M, M_j, B_{AB}	–	A drawing (model) of the plant; the members of a family of candidates; the Bayes factor $p(\mathbf{y} \mid M_A)/p(\mathbf{y} \mid M_B)$.
$p(\mathbf{y} \mid M)$	–	Evidence (marginal likelihood) of model M.
$\mathbf{J}_a$, $\mathbf{\Lambda}_a$, $\mathbf{\Sigma}_a$	varies	The Jacobian rows of a candidate reading, its information $\mathbf{J}_a^\mathsf{T}\mathbf{J}_a$, and the preposterior covariance $(\mathbf{\Sigma}^{-1} + \mathbf{\Lambda}_a)^{-1}$.
γ, t^*	varies/h, h	Slope of an entry's trend across windows; the time at which a rising trend reaches a limit from below (as in the companion book).
Decisions		
α, κ	–	The stated risk of a chance constraint, $\Pr(f \le f^* \mid \mathbf{y}) \ge 1 - \alpha$; the stated confidence $\kappa = 1 - \alpha$.
$U(a)$, $L(a, \cdot)$	varies	Preposterior value of a candidate reading a; the loss of an action a.

Vocabulary

Plant graph The physical system as nodes (places where a conserved quantity balances), edges (paths along which it flows) and sources, with the relations that must hold among them. Built before any question is asked.

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

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

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

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

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

Set point, process gain The value a loop is asked to hold; the change of the controlled quantity per unit change of an input.

Chance constraint A limit required to hold with a stated probability, $\Pr(f \le f^* \mid \mathbf{y}) \ge 1 - \alpha$.

Back-off The distance a set point stands inside its limit because the quantity is uncertain: $\Phi^{-1}(1 - \alpha)$ times the quantity's spread, for a normal posterior.

Scenario node The plant re-solved with a set point or an actuator pinned to a value not yet applied, and the covariance carried through.

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

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

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

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

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

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

Part IV

Deriving the Control Law

Chapter 12

Certainty Equivalence and the Linear-Quadratic Law

12.1 A vendor's gain beside a posterior

The commissioning sheet holds the vendor's PI, $K_p = 0.10$ per kelvin and $K_i = 0.002$ per kelvin second, and now it holds a posterior as well: a spread of 0.0485 on g_w and of 0.19 K on the sensor's offset. One colleague says that a controller designed from a model needs that band in its gain. Another says that the textbook result removes it. The trend at 30 kW shows the reading at 18.00°C.

The engineer would conclude from the trend that the loop is doing its job and that the argument can wait. The sheet has no line for either claim, and no reason for its 0.10. The question of this chapter is

when the band changes the move at all, what the law is when it does not, and how that law does against the PI at a load step.

Jargon: linear-quadratic regulator, Riccati recursion, cost-to-go weight, LQG, certainty equivalence, estimation term, duality, regret, feedforward

The *linear-quadratic regulator* is the law that minimizes a quadratic cost on linear rows with the state known. The *Riccati recursion* is the backward pass that computes its gains, and the *cost-to-go weight* $\mathbf{S}_k$ is the matrix it carries: what a state at step k will cost from there on. *LQG* is the same law fed by the filter's mean. *Certainty equivalence* is the use of that mean as if it were the state when choosing the move (Chapter 7's word); here it is derived, under stated hypotheses. The *estimation term* is the part of the expected cost that comes from the law acting on an estimate and not on the state. *Duality* of control and estimation is the algebraic identity between the control recursion and the filter's covariance recursion under a map of symbols. *Regret* is the excess expected cost of a law computed at the wrong plant, over that plant's own law. *Feedforward* is a move computed from a metered disturbance before the controlled quantity responds.

Key idea

Stack the moves and the states of the future as entries, with the cost as rows, and eliminate them from the last one back, as Chapter 9 eliminated the past from the first one forward. The recursion that results is the law. Under linear rows, a quadratic cost, additive noise independent of the reading errors and of the start, a known gain, no binding limit and one controller that remembers all it has read and done, the band adds terms to the expected cost and none to the move. That is a result about this restriction, and it says nothing about a disturbance the rows do not hold.

12.1.1 The restriction

The problem is equation (7.1), the policy form of the expected consequence (2.23). This chapter works in the row of Chapter 7's map, Table 7.6, labeled "known gain, quadratic, no limit": a known linear plant, a quadratic ℓ, and the last line of the problem removed. **Two couplings of Section 7.1.3 are kept: 1, the move is an entry of the graph, and 2, the move produces the readings. Two fall: 3, the move does not change the coming Σ, and 4, the band does not reach the move.** That the last two fall is the result of Sections 12.3.3 and 12.3.5. It is not an assumption, and nothing below sets them to zero by hand. Chapter 7 showed this on one node and two moves; this chapter derives it for the closed-loop hall and for any horizon.

12.1.2 The sampled hall (step 0)

The plant is the closed-loop hall at a pinned point, read open loop so that the move is free: Chapter 6's fit, $g = 0.01833$, $g_w = -0.16764$, $b = 0.05708$, at 30 kW and $r = 18°\mathrm{C}$. That point is a choice: $T = 22.8521°\mathrm{C}$, $T_c = 17.9429°\mathrm{C}$, $u = 0.5623$. The state $\mathbf{x}_k = (\delta T, \delta T_c, \delta u)^\mathsf{T}$ holds the deviations of the hall air, the coil's discharge air and the valve at controller step k, 30 s apart. The move v_k is the valve command's deviation, held over the step and made after the reading of step k, as in Chapter 7. Discretizing the linearized rows with the command held gives

$$\mathbf{x}_{k+1} = \mathbf{A}\mathbf{x}_k + \mathbf{B}v_k + \mathbf{w}_k, \qquad \mathbf{y}_k = \mathbf{C}\mathbf{x}_k + \boldsymbol{\varepsilon}_k, \tag{12.1}$$

$$\mathbf{A} = \begin{bmatrix} 0.96034 & 0.02699 & -0.0744 \\ 0.32393 & 0.26008 & -2.17187 \\ 0 & 0 & 0.60653 \end{bmatrix}, \qquad \mathbf{B} = \begin{bmatrix} -0.01443 \\ -0.74494 \\ 0.39347 \end{bmatrix}, \tag{12.2}$$

with $\mathbf{C} = \mathbf{I}$ when all three readings are used and $\mathbf{C} = (0, 1, 0)$ with the discharge reading alone. The eigenvalues of $\mathbf{A}$ are 0.97261, 0.60653 and 0.24781, time constants of 1080.2, 60.0 and 21.5 s. The steady gain of the valve on T_c by this pair is $-7.01178\,\mathrm{K}$ per unit. The pair agrees with Chapter 8's map at the same point (largest difference 0.0), and with the plant's Runge–Kutta at a kick of 10^{-4} (largest difference $1.7 \times 10^{-8}\,\mathrm{K}$).

The choices, named here. The cost uses Chapter 7's spreads, 0.5 K on T_c and 0.2 on the command, so $q = 4$ and $\rho = 25$, and it runs over $N = 40$ moves (20 min):

$$J = \sum_{k<N} \left(\mathbf{x}_k^\mathsf{T}\mathbf{Q}\mathbf{x}_k + \rho\, v_k^2\right) + \mathbf{x}_N^\mathsf{T}\mathbf{Q}\mathbf{x}_N, \qquad \mathbf{Q} = \mathrm{diag}(0, q, 0). \tag{12.3}$$

Only the ratio ρ/q sets the gain: multiplying both by ten moves it by 6.7×10^{-16}. The cost buys a law with no tuning step, and it fails as a description of what the plant's owner wants wherever a load asks the valve to stay away from the pinned point, because it charges the distance from that point. Section 12.3.6 shows what this choice does. The plant noise has spreads $(0.02\,\mathrm{K}, 0.02\,\mathrm{K}, 0.002)$: Chapter 9's value on the air, the same on the coil, and a stated small one on the valve. The readings have $(0.50\,\mathrm{K}, 0.10\,\mathrm{K}, 0.010)$, the plant's sensors, and the filter reads T_c as $y_c - \hat{b}$. **The law reads $v_k = \mathbf{K}_k\mathbf{x}_k$, with the sign of Chapter 7: warm air opens the valve.**

12.1.3 The control engineer's questions (step 1)

Table 12.1 lists the seven questions. The first column is the question in the engineer's words, the second the general form, the third its mathematical form on this hall, the fourth its category, the fifth the computation and the last the subsection. The second question has no form of its own: it is algebra.

Table 12.1: The questions asked of the law on the sampled hall.

Control engineer's question	General form	Mathematical form	Category	Computation	Answered in
What is the law with no band and no limit?	expected consequence over moves, (2.23)	$\mathbf{K}_k$ from the stacked rows	design	expected consequence	12.3.1
Is the recursion Chapter 9's in disguise?	–	$\mathbf{A} \to \mathbf{A}^\mathsf{T}$, $\mathbf{B} \to \mathbf{C}^\mathsf{T}$, $\mathbf{Q} \to \mathbf{W}$, $\mathbf{R} \to \mathbf{V}$	–	none: algebra	12.3.2
When does the band change the move?	expected consequence at a belief, (2.23)	v_0 and $\mathbb{E}J$ against $\mathbf{\Sigma}$	design	expected consequence	12.3.3, 12.3.4
What does the band cost, if not the move?	moments of a derived quantity, (2.10)	$\sum_k \operatorname{tr}(\mathbf{K}_k^\mathsf{T} H_k \mathbf{K}_k \mathbf{\Sigma}_{k\|k})$	estimation of a derived quantity	posterior	12.3.5
How does the law do against the PI?	prediction at a scenario node, closed loop, (2.10)	load and set-point steps, noise	prediction	posterior	12.3.6
What does fit error cost in closed loop?	moments of a derived quantity, (2.10)	regret over $p(\mathbf{z} \mid \mathbf{y})$	estimation of a derived quantity	posterior	12.3.7
Which measurement most reduces the expected cost?	next measurement, (2.23), (4.73)	the law's cost term after each candidate	design of measurement	preposterior	12.6.1

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

The rows on the past. These are Chapter 9's, now with three entries per step: a prior row on $\mathbf{x}_0$, physics rows $(\mathbf{x}_{k+1} - \mathbf{A}\mathbf{x}_k - \mathbf{B}v_k)$ scaled by the plant noise, and reading rows scaled by the sensors. Each is linear, and its errors are taken as independent and normal. The move sits in a residual and has no line in the Jacobian, which is the zero Chapter 7 called load-bearing in equation (7.6). That choice buys a covariance known before any reading. It fails for a valve whose effect depends on the state, and Section 12.3.4 tests what a band on the gain does.

The rows on the future. Following equation (7.17), the moves and future states are entries, $\mathbf{z} = (v_0, \mathbf{x}_1^\mathsf{T}, v_1, \mathbf{x}_2^\mathsf{T}, \dots, v_{N-1}, \mathbf{x}_N^\mathsf{T})^\mathsf{T}$, from a given start $\mathbf{x}_0$. The rows are the cost's spreads and the physics at a tolerance σ_p:

$$\mathbf{r}(\mathbf{z}) = \begin{bmatrix} \sqrt{\rho}\, v_k \\ (\mathbf{x}_{k+1} - \mathbf{A}\mathbf{x}_k - \mathbf{B}v_k)/\sigma_p \\ \mathbf{Q}^{1/2}\mathbf{x}_{k+1} \end{bmatrix}, \qquad \mathbf{J} = \begin{bmatrix} * & 0 & 0 & 0 & \cdots \\ * & * & 0 & 0 & \cdots \\ 0 & * & 0 & 0 & \cdots \\ 0 & 0 & * & 0 & \cdots \\ 0 & * & * & * & \cdots \\ \vdots & & & & \ddots \end{bmatrix}, \tag{12.4}$$

with $\mathcal{L} = \frac{1}{2}\mathbf{r}^\mathsf{T}\mathbf{r}$, which on the physics is half the cost (12.3) less the fixed term of $\mathbf{x}_0$. The Jacobian's pattern, by block, in the columns $v_0, \mathbf{x}_1, v_1, \mathbf{x}_2$, is a property of the graph: a move row touches its move, a physics row touches a move and its two neighboring states, and a cost row touches its state. The information matrix of these normalized rows is $\mathbf{\Lambda} = \mathbf{J}^\mathsf{T}\mathbf{J}$; with unnormalized rows it is $\mathbf{J}^\mathsf{T}\mathbf{W}\mathbf{J}$ with $\mathbf{W}$ a weight matrix, and from here on the letter $\mathbf{W}$ means the plant noise's covariance. The information is block tridiagonal: at $N = 2$ the stack has 8 entries and, at $\sigma_p = 0.01\,\mathrm{K}$, 46 of the 64 elements are nonzero. **The rows on the future are not fit like the rows on the past.** Equation (7.2) averages over the future noise, so these rows hold no noise as a free entry; at the end of Section 12.3.1 the tolerance σ_p goes to zero.

The law's rivals. The vendor's PI reads y_c alone and moves the valve by K_p times the error plus its integral. It keeps a state of its own, with no row for the hall air or the valve position.

12.2 The posterior of the state (step 3)

Every answer below uses one of two beliefs: the belief on the plant's entries, from which $\mathbf{A}$ and $\mathbf{B}$ are read, and the belief on the sampled state, which the law acts on.

12.2.1 The belief on the plant

This is Chapter 6's posterior, cited and not refit, on its reduced state: $g = 0.01833$ with a spread of 0.04210, $g_w = -0.16764$ with 0.04846, and $b = 0.05708$ with 0.18986 K, with $\mathrm{corr}(g, b) = -0.9160$ and $\mathrm{corr}(g_w, b) = 0.9924$. These correlations are the ridge Chapter 6 found. The mean gives equation (12.2). The spreads are used in Sections 12.3.4, 12.3.7 and 12.6.1, by that chapter's reweighted draws: 3000 draws, seed 1203, effective size 2005. Chapter 8's priors on the dynamic entries are drawn independently beside them, seed 1204.

12.2.2 The belief on the sampled state

On the rows of the past the belief is Chapter 9's recursion. Eliminating the older state by equation (9.4) is the predict and adding the reading's row is the update; equations (9.14) and (9.15) with matrices in place of numbers read

$$\begin{gathered}\hat{\mathbf{x}}_{k|k-1} = \mathbf{A}\hat{\mathbf{x}}_{k-1|k-1} + \mathbf{B}v_{k-1}, \qquad \mathbf{P}_{k|k-1} = \mathbf{A}\boldsymbol{\Sigma}_{k-1|k-1}\mathbf{A}^{\mathsf{T}} + \mathbf{W}, \\ \boldsymbol{\Upsilon}_k = \mathbf{P}_{k|k-1}\mathbf{C}^{\mathsf{T}}\left(\mathbf{C}\mathbf{P}_{k|k-1}\mathbf{C}^{\mathsf{T}} + \mathbf{V}\right)^{-1}, \qquad \boldsymbol{\Sigma}_{k|k} = (\mathbf{I} - \boldsymbol{\Upsilon}_k\mathbf{C})\mathbf{P}_{k|k-1}, \\ \hat{\mathbf{x}}_{k|k} = \hat{\mathbf{x}}_{k|k-1} + \boldsymbol{\Upsilon}_k\left(\mathbf{y}_k - \mathbf{C}\hat{\mathbf{x}}_{k|k-1}\right),\end{gathered} \tag{12.5}$$

with $\mathbf{W}$ and $\mathbf{V}$ the covariances of the plant noise and the readings. The move enters the mean and not the covariance. With all three readings the steady covariances are

$$\mathbf{P} = \begin{bmatrix} 3.5007 & 1.3395 & 0.0004 \\ 1.3395 & 1.0444 & -0.0085 \\ 0.0004 & -0.0085 & 0.0061 \end{bmatrix} \times 10^{-3}, \quad \boldsymbol{\Sigma}_\infty = \begin{bmatrix} 3.2943 & 1.1970 & 0.0013 \\ 1.1970 & 0.9393 & -0.0072 \\ 0.0013 & -0.0072 & 0.0058 \end{bmatrix} \times 10^{-3}, \tag{12.6}$$

in units of K and of valve, so that the steady spreads are 0.0574 K, 0.0306 K and 0.0024; with y_c alone they are 0.0598 K, 0.0315 K and 0.0025. Neither matrix contains a move or a reading.

A belief to test with. A state at the start of a horizon needs a belief. The chapter uses $\hat{\mathbf{x}}_0 = (0, 1, 0)^{\mathsf{T}}$ and a base band with spreads $(0.3535\,\mathrm{K}, 0.0707\,\mathrm{K}, 0.0071)$, a prior of one reading's spread updated by one reading, scaled by a factor s, so that $\boldsymbol{\Sigma}_{0|0}$ is diagonal and its spreads are s times those. This belief is a stated test case, wide on purpose, and it is not a state the loop visits: the coil forgets a kick of T_c in 21.5 s.

12.3 The answers (step 4)

Each answer reads the rows of Section 12.1.4 and the beliefs of Section 12.2. Each route is named where it is used: the recursion; the stack, which is least squares on all rows at once; the condensed solve, which is least squares on the moves alone; the closed form of the expected cost; the Lyapunov route, the stationary covariance of the closed loop; fresh linear records, with seeds and counts; the plant, nonlinear and noiseless; and the draws.

12.3.1 The law from the stacked rows

Control engineer's question. What is the law with no band and no limit?

General form. The expected consequence (2.23), with the sequence of moves as the action and the plant known. With no noise left in the future, it is the minimum of the rows' objective,

$$\min_{v_{0:N-1},\,\mathbf{x}_{1:N}} \ \mathcal{L}(\mathbf{z}) = \tfrac{1}{2}\mathbf{r}(\mathbf{z})^{\mathsf{T}}\mathbf{r}(\mathbf{z}). \tag{12.7}$$

Specialized form. Equation (12.7) with the rows (12.4), which is equation (2.23) with $L = J$ of equation (12.3) on the known plant. The modeling choice entering here is a known $(\mathbf{A}, \mathbf{B})$: it is what makes these rows linear in the entries, and Section 12.3.4 tests a band on it.

The solution is Chapter 9's elimination run from the other end, and Figure 12.1 draws it. Time runs from left to right. The green boxes are states, the blue box is the move, the black square is the physics row between neighbors and the open boxes labeled with a weight are what the cost leaves on an entry. The top row is the state of the stack before the two strokes, with the last entry shaded; the middle row is after the first stroke, and the bottom row after the second.

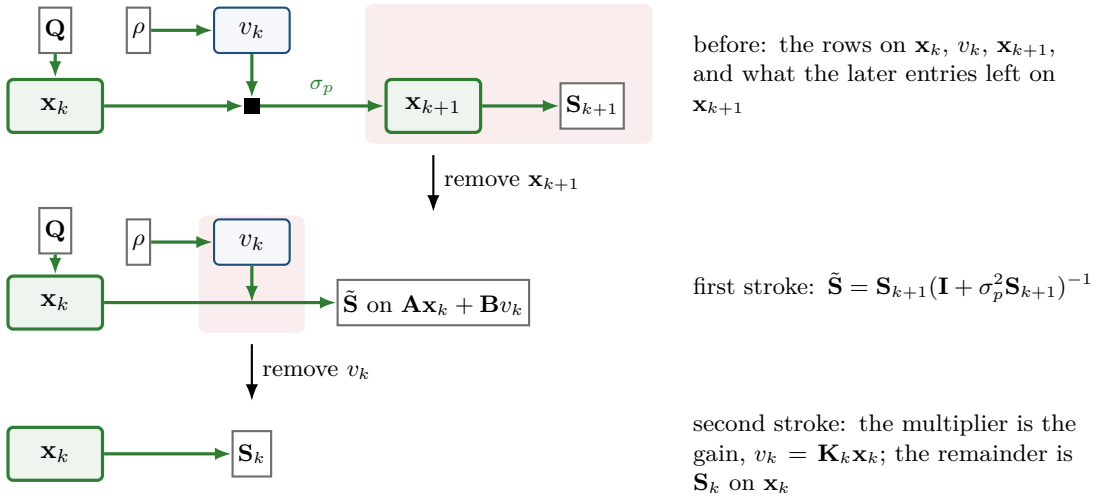


Figure 12.1: The elimination of the stack from the last entry back, equations (12.8) and (12.9). Green: states. Blue: the move. Black square: the physics row at tolerance σ_p. Open boxes: the weights the cost rows and the later entries leave. Shaded: what the next stroke removes. The bottom row has the form of the top row's right end, one step earlier.

The first stroke. Suppose everything after step $k+1$ has been eliminated and has left the weight $\mathbf{S}_{k+1}$ on $\mathbf{x}_{k+1}$, starting from $\mathbf{S}_N = \mathbf{Q}$. Then $\mathbf{x}_{k+1}$ appears in the physics row and that weight only. With $\mathbf{M} = [\mathbf{A} \ \ \mathbf{B}]$ the coupling to the remaining entries $(\mathbf{x}_k, v_k)$, its block of the information is $\mathbf{I}/\sigma_p^2 + \mathbf{S}_{k+1}$ and its off-diagonal block is $-\mathbf{M}/\sigma_p^2$. Equation (9.4) with $x = \mathbf{x}_{k+1}$, the same step that gave Chapter 9's predict, leaves

$$\frac{\mathbf{M}^{\mathsf{T}}\mathbf{M}}{\sigma_p^2} - \frac{\mathbf{M}^{\mathsf{T}}}{\sigma_p^2}\Big(\frac{\mathbf{I}}{\sigma_p^2} + \mathbf{S}_{k+1}\Big)^{-1}\frac{\mathbf{M}}{\sigma_p^2} = \mathbf{M}^{\mathsf{T}}\tilde{\mathbf{S}}\mathbf{M}, \qquad \tilde{\mathbf{S}} = \mathbf{S}_{k+1}\big(\mathbf{I} + \sigma_p^2\mathbf{S}_{k+1}\big)^{-1}. \tag{12.8}$$

This is Chapter 7's shrunk weight, equation (7.18), as a matrix, and $\sigma_p \to 0$ gives $\mathbf{S}_{k+1}$.

The second stroke. The entries $(\mathbf{x}_k, v_k)$ now hold the information $\left[\begin{smallmatrix} \mathbf{Q}+\mathbf{A}^\mathsf{T}\tilde{\mathbf{S}}\mathbf{A} & \mathbf{A}^\mathsf{T}\tilde{\mathbf{S}}\mathbf{B} \\ \mathbf{B}^\mathsf{T}\tilde{\mathbf{S}}\mathbf{A} & \rho+\mathbf{B}^\mathsf{T}\tilde{\mathbf{S}}\mathbf{B} \end{smallmatrix}\right]$. Equation (9.4) with $x = v_k$ divides by the move's own element and leaves the Schur complement on $\mathbf{x}_k$. The multiplier in that division is what the conditional mean of the removed entry reads, and on these rows that mean is the move:

$$H_k = \rho + \mathbf{B}^\mathsf{T}\tilde{\mathbf{S}}\mathbf{B}, \qquad \mathbf{K}_k = -H_k^{-1}\mathbf{B}^\mathsf{T}\tilde{\mathbf{S}}\mathbf{A}, \qquad \mathbf{S}_k = \mathbf{Q} + \mathbf{A}^\mathsf{T}\tilde{\mathbf{S}}\mathbf{A} + \mathbf{A}^\mathsf{T}\tilde{\mathbf{S}}\mathbf{B}\mathbf{K}_k. \tag{12.9}$$

Equations (12.8) and (12.9) are the Riccati recursion, derived and not stated. With $\sigma_p \to 0$ it is the hard law, whose form Chapter 7 reached by derivatives in equation (7.15). With a move of several entries, ρ becomes a matrix $\mathbf{R}$ and H_k a matrix, by the same steps.

Building blocks. $\mathbf{A}$ and $\mathbf{B}$ from equation (12.2), read at the mean of the belief of Section 12.2.1; $\mathbf{Q}$ and ρ from equation (12.3); the two strokes.

The plant's numbers. On Chapter 7's example the recursion returns that chapter's by-hand values, $K_1 = 0.062560$, $K_0 = 0.063023$, $S_1 = 4.12821$. On the hall at one move, hard, $\mathbf{B}^\mathsf{T}\mathbf{Q}\mathbf{B} = 2.21975$, so $H = 27.21975$; with $\mathbf{B}^\mathsf{T}\mathbf{Q}\mathbf{A} = (-0.96525, -0.77497, 6.47166)$ this gives $\mathbf{K} = (0.03546, 0.02847, -0.23776)$, and the remainder is

$$\mathbf{S}_0 = \begin{bmatrix} 0.3855 & 0.30951 & -2.58468 \\ 0.30951 & 4.2485 & -2.07517 \\ -2.58468 & -2.07517 & 17.3294 \end{bmatrix}, \tag{12.10}$$

the weight the next stroke back starts from. Table 12.2 gives the gains by the number of moves left, and Figure 12.2 draws them, with each panel's horizontal axis the moves left on a logarithmic scale and the dotted lines the steady values.

Table 12.2: The gains by the number of moves left, the hard recursion, $\rho/q = 6.25$. "Steady" solves the algebraic equation.

Moves left [–]	1	2	3	5	10	40	80	steady
K_T [1/K]	0.035461	0.079716	0.100159	0.110914	0.114221	0.118018	0.118281	0.11829
K_c [1/K]	0.028471	0.034208	0.034079	0.033990	0.034078	0.034201	0.034209	0.03421
K_u [–]	−0.237756	−0.442459	−0.498050	−0.510781	−0.511758	−0.512621	−0.512681	−0.51268

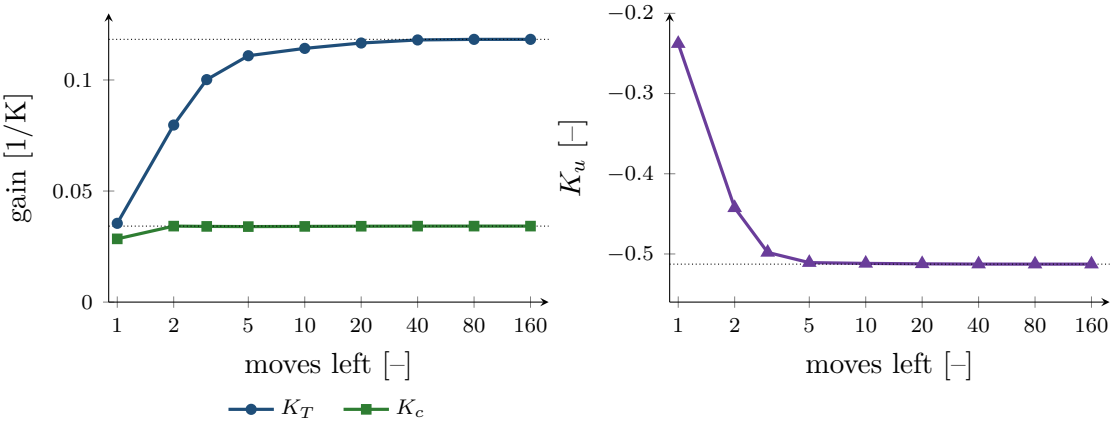


Figure 12.2: The gains of Table 12.2, with 20 and 160 moves left added. Left: the gains on the hall air and the discharge air. Right: the gain on the valve's position. Dotted: the steady values.

The gains settle backward from the end of the horizon: within 10, 1 and 0.1 % of steady at 3, 8 and 33 moves left. The algebraic equation and the recursion at 400 moves differ by 3.3×10^{-16}. The steady weight, which Section 12.3.2 uses, is

$$\mathbf{S} = \begin{bmatrix} 3.02667 & 0.40804 & -6.01739 \\ 0.40804 & 4.25697 & -2.32452 \\ -6.01739 & -2.32452 & 28.72346 \end{bmatrix}. \tag{12.11}$$

What the law is. It is not a PI. Its gain on T_c is about a third of the vendor's K_p. Per kelvin its largest gain is on the hall air, and it subtracts the valve's position. The coil forgets a kick of T_c in 21.5 s, so what persists is the air. The moduli of the closed loop's eigenvalues are 0.95572 and a pair at 0.33913 for the law, and 0.94412, a pair at 0.78412 and 0.27803 for the vendor's PI. A start at $(0, 1, 0)$ gives $v_0 = 0.03420$. Across ρ/q from 0.25 to 400, K_c runs from 0.19041 to 0.00084 and the largest modulus stays between 0.951 and 0.972: the cost moves the fast pair and not the hall.

The stack against the recursion. Table 12.3 solves the stack whole at $N = 40$ from $(0, 1, 0)$, for each tolerance. The first column is σ_p, the second the largest gap in the moves from the recursion at that same σ_p, the third the stack's first move, the fourth the largest gap from the hard law and the last the cost-to-go.

Table 12.3: The stack against the recursion, $N = 40$, start $(0, 1, 0)$. Stated tolerance for the second column: 10^{-7}.

σ_p [K]	Against recursion, same σ_p [–]	v_0, stack [–]	Against hard law [–]	Cost-to-go [–]
1.0000	3.70×10^{-17}	0.006186	2.80×10^{-2}	4.05323
0.1000	1.42×10^{-16}	0.031873	2.33×10^{-3}	4.24598
0.0200	1.93×10^{-16}	0.034085	1.16×10^{-4}	4.25640
0.0100	1.58×10^{-15}	0.034172	2.93×10^{-5}	4.25677
0.0010	1.07×10^{-14}	0.034201	2.94×10^{-7}	4.25690
0.0001	4.81×10^{-14}	0.034201	2.94×10^{-9}	4.25690

At every σ_p the stack and the recursion agree to rounding, so the two strokes are the stack's solution at that σ_p by algebra. Against the hard law the gap goes as σ_p^2. Condensed least squares, with the physics hard, meets the hard recursion to 3.3×10^{-17}. **A soft physics row on the future is Chapter 7's optimism**: at $\sigma_p = 0.10$ K the first move is 0.9319 of the hard law's, because the stack has planned on noise that does part of the valve's work. The stack is the law only with the rows hard and no noise left free. It is not the general solution of the problem (7.1), and Section 12.3.4 leaves it in two directions.

12.3.2 The recursion and the filter

Control engineer's question. Is the recursion Chapter 9's in disguise?

Equations (12.8)–(12.9) with $\sigma_p \to 0$ and $\tilde{\mathbf{S}} = \mathbf{S}_{k+1}$ give $\mathbf{S}_k = \mathbf{Q} + \mathbf{A}^\mathsf{T}\mathbf{S}_{k+1}\mathbf{A} - \mathbf{A}^\mathsf{T}\mathbf{S}_{k+1}\mathbf{B}(\mathbf{R} + \mathbf{B}^\mathsf{T}\mathbf{S}_{k+1}\mathbf{B})^{-1}\mathbf{B}^\mathsf{T}\mathbf{S}_{k+1}\mathbf{A}$. The filter's rows (12.5) give, on the predicted covariance,

$$\mathbf{P}_{j+1} = \mathbf{W} + \mathbf{A}\mathbf{P}_j\mathbf{A}^\mathsf{T} - \mathbf{A}\mathbf{P}_j\mathbf{C}^\mathsf{T}\left(\mathbf{V} + \mathbf{C}\mathbf{P}_j\mathbf{C}^\mathsf{T}\right)^{-1}\mathbf{C}\mathbf{P}_j\mathbf{A}^\mathsf{T}. \tag{12.12}$$

Under the map $\mathbf{A} \to \mathbf{A}^\mathsf{T}$, $\mathbf{B} \to \mathbf{C}^\mathsf{T}$, $\mathbf{Q} \to \mathbf{W}$, $\mathbf{R} \to \mathbf{V}$, the first line becomes equation (12.12) by algebra, with $\mathbf{S}_{N-j}$ in the role of $\mathbf{P}_j$, and the control gain's transpose with its sign reversed becomes the predictor gain $\mathbf{A}\boldsymbol{\Upsilon}_j$. On the hall, $N = 40$ and with the filter's first predicted variance $\mathbf{P}_0$ as the terminal weight, the largest gaps are $4.9 \times 10^{-17}\,\mathrm{K}^2$ on the covariance and

6.7×10^{-16} on the gain, against a stated tolerance of 10^{-10}. With $\mathbf{A}$ changing halfway, from 30 to 45 kW, the control recursion, run backward, needs the sequence read from its far end, and then the gap is $4.2 \times 10^{-17}\,\mathrm{K}^2$; with the same index it is $0.0053\,\mathrm{K}^2$.

What the identity is, and what it is not. It is an algebraic identity between two Riccati equations, which is why the two elimination routes look alike. It is not one problem. The map reverses the index, and it is no time-reversed plant, since $\mathbf{A}^\mathsf{T}$ is an adjoint and not $\mathbf{A}^{-1}$. On the hall's own numbers $\mathbf{S}$ of equation (12.11) and $\mathbf{P}$ of equation (12.6) are unrelated matrices from unrelated inputs. Doubling $\mathbf{W}$ leaves $\mathbf{K}$ unchanged (gap 0.0) and moves $\boldsymbol{\Upsilon}$ by 0.057. Doubling ρ leaves $\boldsymbol{\Upsilon}$ unchanged (gap 0.0) and moves K_c by -0.01403. One recursion takes readings forward. The other has none, and is solved backward before any reading exists.

12.3.3 Certainty equivalence

Control engineer's question. When does the band change the move?

General form. The expected consequence (2.23) of the move at a belief, with later moves by the policy, which is one step of the recursion (7.3):

$$V_k(\hat{\mathbf{x}}_{k|k}, \boldsymbol{\Sigma}_{k|k}) = \min_{v_k} \; \mathbb{E}\big[\mathbf{x}_k^\mathsf{T}\mathbf{Q}\mathbf{x}_k + \rho v_k^2 + V_{k+1} \mid \hat{\mathbf{x}}_{k|k}, \boldsymbol{\Sigma}_{k|k}\big]. \tag{12.13}$$

Specialized form. Equation (12.13), which is (2.23) with the action v_k, on the rows (12.1) with the hard physics, $\mathbf{W}$ and $\mathbf{V}$ now the noise that the future averages, and the belief (12.5). Two choices enter: the noise is additive, of mean zero, independent from step to step, and independent of state and move; and the plant noise, the reading errors and the start are independent of each other, which makes $\boldsymbol{\Sigma}_{k|k}$ the covariance of the estimate's error. Normal errors would add that $\hat{\mathbf{x}}_{k|k}$ is the conditional mean, and the theorem below keeps that clause apart.

The one fact the derivation needs is the mean of a quadratic:

$$\mathbb{E}\big[\mathbf{x}^\mathsf{T}\mathbf{S}\mathbf{x} \mid \hat{\mathbf{x}}, \boldsymbol{\Sigma}\big] = \hat{\mathbf{x}}^\mathsf{T}\mathbf{S}\hat{\mathbf{x}} + \operatorname{tr}(\mathbf{S}\boldsymbol{\Sigma}). \tag{12.14}$$

The last term has no move in it, since $\boldsymbol{\Sigma}$ does not depend on v. Now induct. Suppose $V_{k+1} = \hat{\mathbf{x}}_{k+1}^\mathsf{T}\mathbf{S}_{k+1}\hat{\mathbf{x}}_{k+1} + \operatorname{tr}(\mathbf{S}_{k+1}\boldsymbol{\Sigma}_{k+1|k+1}) + c_{k+1}$, with c_{k+1} free of the moves; at $k+1 = N$ this is equation (12.14) with $c_N = 0$. Seen from step k, the next mean $\hat{\mathbf{x}}_{k+1|k+1}$ is a quantity with mean $\mathbf{m} = \mathbf{A}\hat{\mathbf{x}}_{k|k} + \mathbf{B}v_k$ and covariance $\mathbf{P}_{k+1|k} - \boldsymbol{\Sigma}_{k+1|k+1}$, the part of the prediction that the reading will resolve, as in equation (7.9). Averaging over the coming reading, the two terms in $\boldsymbol{\Sigma}_{k+1|k+1}$ cancel:

$$\mathbb{E}_{\mathbf{y}_{k+1}}[V_{k+1}] = \mathbf{m}^\mathsf{T}\mathbf{S}_{k+1}\mathbf{m} + \operatorname{tr}(\mathbf{S}_{k+1}\mathbf{A}\boldsymbol{\Sigma}_{k|k}\mathbf{A}^\mathsf{T}) + \operatorname{tr}(\mathbf{S}_{k+1}\mathbf{W}) + c_{k+1}. \tag{12.15}$$

The move is only in $\rho v_k^2 + \mathbf{m}^\mathsf{T}\mathbf{S}_{k+1}\mathbf{m}$, whose minimum is the second stroke with $\sigma_p = 0$: $v_k = \mathbf{K}_k\hat{\mathbf{x}}_{k|k}$, the gain of equation (12.9) acting on the estimate. Since $\mathbf{Q} + \mathbf{A}^\mathsf{T}\mathbf{S}_{k+1}\mathbf{A} = \mathbf{S}_k + \mathbf{K}_k^\mathsf{T}H_k\mathbf{K}_k$ by that equation, the remaining terms collect into the same form at step k, which closes the induction:

$$\boxed{\mathbb{E}J = \hat{\mathbf{x}}_0^\mathsf{T}\mathbf{S}_0\hat{\mathbf{x}}_0 + \operatorname{tr}(\mathbf{S}_0\boldsymbol{\Sigma}_{0|0}) + \sum_{k<N} \operatorname{tr}(\mathbf{S}_{k+1}\mathbf{W}) + \sum_{k<N} \operatorname{tr}\big(\mathbf{K}_k^\mathsf{T}H_k\mathbf{K}_k\boldsymbol{\Sigma}_{k|k}\big)} \tag{12.16}$$

The theorem, with its hypotheses. The move depends on the belief only through $\hat{\mathbf{x}}$, and $\mathbf{\Sigma}$ adds terms that no move can reach, when: (1) the rows are linear; (2) the cost is quadratic; (3) the noise is additive, of mean zero, independent from step to step, and does not depend on state or move; (4) the gain is known; (5) no limit binds; (6) the plant noise, the reading errors and $\mathbf{x}_0$ are independent of each other; and (7) one controller makes every move and keeps every reading and every move it has made. Hypothesis 7 is what lets the induction average over the coming reading as one that the next move will hold, with all that came before. It is not a formality: a distributed controller with partial recall can fail the result, even on linear rows, a quadratic cost and normal errors, because a move can then be worth making for what it tells another controller. These hypotheses are sufficient, and none is shown necessary: a limit that does not bind leaves the move alone, as the next subsection shows.

What normal errors add. They are not on the list, because the algebra needs second moments only. Equation (12.18) below, with the filter's rows (12.5), gives the terms of equation (12.16) for any errors of these variances. Normality is needed only for the filter's optimality: it makes the filter's mean the conditional mean, so that no other estimator does better. Without it the result is about this filter, and a nonlinear estimator may beat it. On 40 000 records at $s = 1$ with all readings, whose estimation term is 0.18710, Laplace errors of the same variances give a gap of 0.18730 (se 0.00098, seed 9002) and two-point errors 0.18577 (se 0.00045, seed 9003), which is -3.0 se. A second seed, 9013, gives 0.18604, -2.4 se, and twelve more, 9101 to 9112, average 0.18712 (se 0.00013), $+0.2$ se. Carrying the second moments through the rows, with no law for the errors, returns the term to 1.7×10^{-15}, so the two low seeds are the tail of the draws.

Building blocks. $\mathbf{S}_k$, $\mathbf{K}_k$, H_k from equations (12.8)–(12.9); $\mathbf{\Sigma}_{k|k}$ and $\mathbf{W}$ from equation (12.5); the test belief of Section 12.2.2.

The plant's numbers. Table 12.4 takes $\hat{\mathbf{x}}_0 = (0, 1, 0)^\mathsf{T}$, $N = 40$, and scales the band by s. Each row is one s. The second column minimizes the closed form over v_0, with the law afterward; the third reads the same minimizer from fresh records, by the parabola through three first moves on the same draws, 20 000 records per row, seeds $1201 + 100s$. The next four columns are the terms of equation (12.16), the eighth their total and the last the total from the records. Figure 12.3 draws them. The upper left panel has v_0 against s, with the closed form as a line and the records as points with bars of one se. The lower panel stacks the four terms at each s. The upper right panel, which the next subsection uses, has the move with a limit.

Table 12.4: The first move and the expected cost against the band, $\hat{\mathbf{x}}_0 = (0, 1, 0)^\mathsf{T}$, $N = 40$. Records: 20 000 per row, se in parentheses.

s [–]	v_0, closed [–]	v_0, records [–]	At the mean [–]	tr($\mathbf{S}_0\mathbf{\Sigma}$) [–]	Noise [–]	Estimation [–]	Total [–]	Total, records [–]
0.0001	0.034201	0.03419 (0.00003)	4.25690	0.00000	0.10917	0.05889	4.42496	4.42524 (0.00065)
0.25	0.034201	0.03428 (0.00008)	4.25690	0.02454	0.10917	0.08377	4.47437	4.47236 (0.00144)
1	0.034201	0.03449 (0.00030)	4.25690	0.39267	0.10917	0.18710	4.94583	4.95555 (0.00666)
2	0.034201	0.03322 (0.00059)	4.25690	1.57067	0.10917	0.37985	6.31659	6.28463 (0.01932)
4	0.034201	0.03332 (0.00119)	4.25690	6.28269	0.10917	1.09780	11.74655	11.84192 (0.07328)

The move stays at 0.034201 **while the cost goes from** 4.42496 **to** 11.74655**.** The closed-form minimizer differs from the law by at most 2.3×10^{-15}. The largest record gap on v_0 is 1.7 se, at $s = 2$. The total by records stands within 1.7 se of the closed form in every row. The closed form of the expected cost at v_0 differs from the total at the law by 8.9×10^{-16} at $s = 1$.

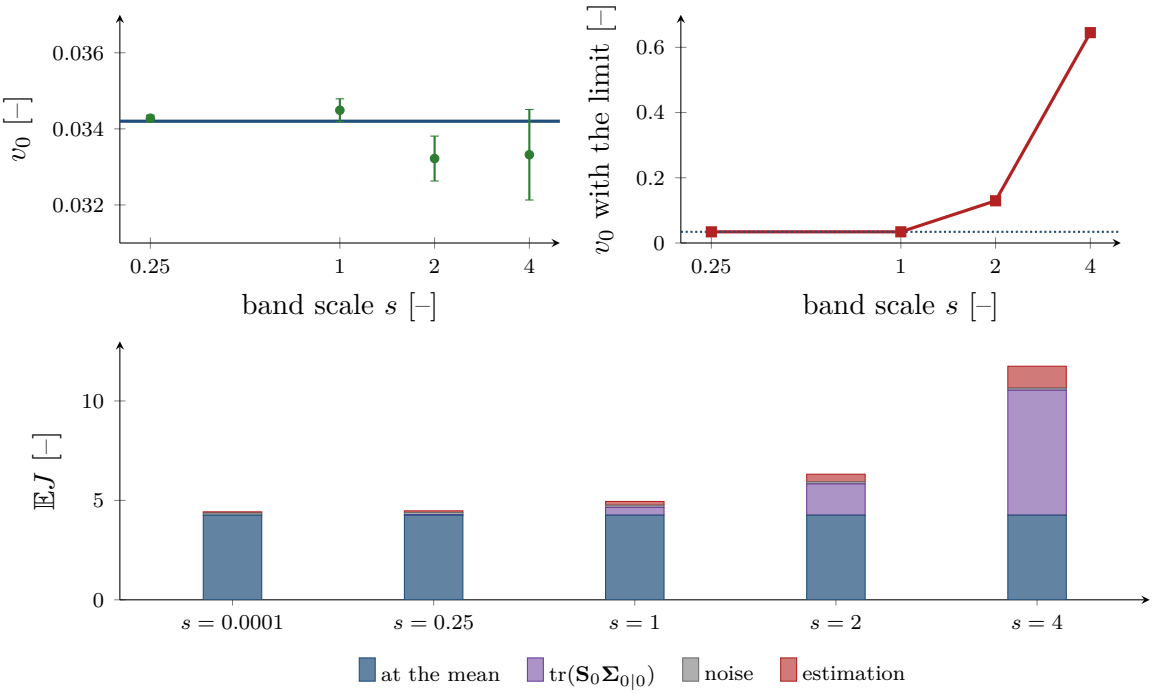


Figure 12.3: Certainty equivalence on the hall. Upper left: the first move against the band's scale, closed form (line) and records (points, one se); the axis is narrow, 0.031 to 0.037. Lower: the terms of equation (12.16) at each s, stacked. Upper right: the first move with the limit of Section 12.3.4 (red) against the law's (dotted); at $s = 4$ its command is above the stop.

12.3.4 Where each hypothesis fails

The second part of the answer to the same question is what happens outside the hypotheses. Two are tested on the hall; a third, a disturbance the rows do not hold, is in Section 12.3.6.

A gain with a band. Hypothesis 4 fails when $\mathbf{A}$ and $\mathbf{B}$ are uncertain. Over one move with the steady $\mathbf{S}$ as the cost-to-go, averaging the same quadratic over the plant's band gives the cautious move

$$v = -\left(\rho + \mathbb{E}[\mathbf{B}^{\mathsf{T}}\mathbf{S}\mathbf{B}]\right)^{-1}\mathbb{E}[\mathbf{B}^{\mathsf{T}}\mathbf{S}\mathbf{A}]\,\hat{\mathbf{x}}, \tag{12.17}$$

where $\mathbb{E}[\mathbf{B}^{\mathsf{T}}\mathbf{S}\mathbf{B}]$ holds a term in the covariance of $\mathbf{B}$ and $\mathbb{E}[\mathbf{B}^{\mathsf{T}}\mathbf{S}\mathbf{A}]$ one in the covariance of $\mathbf{B}$ with $\mathbf{A}$, so the band reaches the move through both. This is coupling 4 by the second route of Section 7.1.3. Table 12.5 evaluates it over the belief of Section 12.2.1 at $\hat{\mathbf{x}} = (0, 1, 0)^{\mathsf{T}}$, with Chapter 8's dynamic priors at each spread in the log. The first column is that spread and the second the resulting relative spread of B_{T_c}. Then come three moves: the certainty-equivalent move at the fit, the same at the belief's mean $\mathbb{E}[\mathbf{A}, \mathbf{B}]$, and the cautious one. Then the cautious move's change against each of the first two, and the excess cost of each of them as a share of the one-move cost.

The change against the fit mixes two things, because the certainty-equivalent move here is taken at the fit, and the fit is not the mean of the draws, while Chapter 7's rows take the prior's own mean. With the fit's band alone, the $-0.23\,\%$ is entirely that mean shift, $-0.24\,\%$, and the two expectations add 0.00. Against $\mathbb{E}[\mathbf{A}, \mathbf{B}]$ the change is $+0.00$, -1.74, -6.45 and $-15.28\,\%$, and that chapter's scalar one-move formula, equation (7.25), at the same relative spreads gives -0.00, -1.77, -6.51 and $-15.31\,\%$. So the sizes match. The nearest comparison in Section 7.3.4 is its cautious move, $15.87\,\%$ below certainty equivalence at $s/\hat{\theta} = 0.50$; its $15.64\,\%$ belongs to the exact route over two moves. The split on the hall differs. At Chapter 8's prior, where

Table 12.5: A gain with a band: one move, the steady $\mathbf{S}$, 3000 draws, effective size 2005, seeds 1203 and 1204. These are this example's values.

sd log dyn. [–]	sd B_{T_c} / B_{T_c} [–]	v, at the fit [–]	v, at $\mathbb{E}[\mathbf{A},\mathbf{B}]$ [–]	v, cautious [–]	Against the fit [%]	Against $\mathbb{E}[\mathbf{A},\mathbf{B}]$ [%]	Excess, fit's [%]	Excess, $\mathbb{E}[\mathbf{A},\mathbf{B}]$'s [%]
0.00	0.0068	0.03421	0.03413	0.03413	−0.23	+0.00	0.000	0.000
0.15	0.1530	0.03421	0.03438	0.03378	−1.25	−1.74	0.002	0.004
0.30	0.3010	0.03421	0.03529	0.03301	−3.51	−6.45	0.016	0.058
0.50	0.4850	0.03421	0.03740	0.03169	−7.37	−15.28	0.058	0.299

B_{T_c} has a spread of 15.3 %, the mean shift is +0.50 %, $\mathbb{E}[\mathbf{B}^\mathsf{T}\mathbf{S}\mathbf{B}]$ alone gives −0.40 % and the numerator $\mathbb{E}[\mathbf{B}^\mathsf{T}\mathbf{S}\mathbf{A}]$ alone −1.35 %: the same dynamic entries move $\mathbf{A}$ and $\mathbf{B}$ together, and the covariance of $\mathbf{B}$ on its own is the minor part. The excess cost of ignoring the band stays low because a cost is flat at its minimum: at most 0.299 % for the move at the mean and 0.058 % for the move at the fit. These sizes are this example's, on one move.

A limit. Hypothesis 5 fails when a limit binds. Take δT_c at the next step at most +0.55 K with probability 0.95, a stated limit on one step, with $\mathbf{S}_1$ unchanged, so the move is the larger of the law's and the limit's. The law's next mean of δT_c is 0.23460 K. Table 12.6 gives the spread of that quantity at each s, the move the limit needs and the move made.

Table 12.6: A limit on the next step: $\delta T_c \le +0.55$ K at 0.95. The last column is the expected cost over the law's at that s.

s [–]	sd next T_c [K]	v needed [–]	Move [–]	Against law [%]	Added cost [–]
0.25	0.03544	−0.31095	0.03420	+0.00	0.000000
1.00	0.11870	−0.12709	0.03420	+0.00	0.000000
2.00	0.23487	0.12940	0.12940	+278.36	0.301340
4.00	0.46845	0.64517	0.64517	+1786.40	12.410992

The upper right panel of Figure 12.3 draws the fourth column. The band leaves the valve alone until the limit binds, between $s = 1$ and $s = 2$, and moves it from there on. That is Chapter 6's back-off, equation (6.3), inside the law. At $s = 1$ the limit is present and the move is unchanged, which is the sense in which the hypotheses are sufficient and not necessary. At $s = 4$ the move 0.64517 means a command of 1.2075, above the stop, so the linear rows do not hold there and that row is their arithmetic and not something the valve can do; at $s = 2$ the command is 0.6917. How to carry a stop in the law is Chapter 15's question.

12.3.5 What the band costs

Control engineer's question. What does the band cost, if not the move?

General form. Moments of a derived quantity, equation (2.10), with the derived quantity the cost of the record.

Specialized form. Equation (2.10) with $f(\mathbf{z}) = J$ of equation (12.3), and the policy free. Telescoping $\mathbf{x}_{k+1}^\mathsf{T}\mathbf{S}_{k+1}\mathbf{x}_{k+1} - \mathbf{x}_k^\mathsf{T}\mathbf{S}_k\mathbf{x}_k$ through the strokes, for any moves and on any record,

$$J = \mathbf{x}_0^\mathsf{T}\mathbf{S}_0\mathbf{x}_0 + \sum_{k<N}\Big[H_k\big(v_k - \mathbf{K}_k\mathbf{x}_k\big)^2 + \mathbf{w}_k^\mathsf{T}\mathbf{S}_{k+1}\mathbf{w}_k + 2(\mathbf{A}\mathbf{x}_k + \mathbf{B}v_k)^\mathsf{T}\mathbf{S}_{k+1}\mathbf{w}_k\Big]. \qquad (12.18)$$

The last term has mean zero when $\mathbf{w}_k$ is independent of what came before and of mean zero. So in expectation the moves pay only through $H_k(v_k - \mathbf{K}_k\mathbf{x}_k)^2$. For the law fed by the filter, $v_k - \mathbf{K}_k\mathbf{x}_k = -\mathbf{K}_k(\mathbf{x}_k - \hat{\mathbf{x}}_{k|k})$, whose second moment is the estimation term of equation (12.16). The gap between the law fed by the filter and the law fed by the truth is therefore that term and nothing else.

Building blocks. $\mathbf{K}_k$, H_k from equation (12.9), $\mathbf{\Sigma}_{k|k}$ from equation (12.5), on the belief of Section 12.2.2 at $s = 1$.

The plant's numbers. Table 12.7 runs the two laws on the same draws, with a stated tolerance of 3 se. Each row is one set of readings. The columns are the records and seed, the paired gap with its se, the closed-form estimation term and their difference in se.

Table 12.7: Separation in simulation, $s = 1$, $N = 40$. The gap is filter-fed minus truth-fed, on the same draws.

Readings	Records, seed	Gap (se) [–]	Estimation term [–]	Difference [se]
All	40 000, 1210	0.18691 (0.00070)	0.18710	−0.27
y_c only	40 000, 1211	0.20872 (0.00079)	0.20890	−0.23
All, second seed	40 000, 1220	0.18745 (0.00070)	0.18710	+0.50

The band's price on this example is 0.18710 **of** 4.94583 **in estimation**, with $\mathrm{tr}(\mathbf{S}_0\mathbf{\Sigma}_{0|0}) = 0.39267$ beside it, which no estimator could avoid once the start is uncertain. The totals agree too: filter-fed 4.94116 (se 0.00466) against the closed 4.94583, and truth-fed 4.75426 (se 0.00429) against 4.75873.

Coupling 3 on the page. On the same draws, the estimate's error at step 10 under the law and under no moves at all differs by at most 8.9×10^{-16} K on any of 5000 records, seed 1212. The sample variance of T_c's error is $0.001211\,\mathrm{K}^2$ against $\mathbf{\Sigma}_{10|10}$'s $0.001213\,\mathrm{K}^2$. Coupling 4 is equation (12.9): neither $\mathbf{W}$, $\mathbf{V}$ nor $\mathbf{\Sigma}$ enters it. Stationary and per step, with all readings, $\mathrm{tr}(\mathbf{SW}) = 0.003028$ and the estimation term $\mathrm{tr}(\mathbf{K}^\mathsf{T} H \mathbf{K} \mathbf{\Sigma}_\infty) = 0.001945$, and with y_c alone that term is 0.002111. The Lyapunov route returns the same sums, 0.004973 and 0.005139.

12.3.6 The law against the PI

Control engineer's question. How does the law do against the PI?

General form. Moments, equation (2.10), of the controlled air and the cost at a closed-loop scenario node: the set point is pinned, the valve is free and determined by the controller's row, and at a stop it is pinned again.

Specialized form. Equation (2.10) with f the stationary cost per step and the spread of T_c under noise, or f the path of T_c at a step. A comparison is fair only on a stated basis, so this subsection uses three: the same cost, with each loop's own best tuning; equal valve travel; and the event. Noise comparisons are linear and stationary, by the Lyapunov route. Steps are on the plant at the fit, noiseless.

Building blocks. $\mathbf{K}$ and $\mathbf{\Upsilon}$ at steady from equations (12.9) and (12.5), $\mathbf{P}$ from equation (12.6); the PI's row with its integrator as a fourth state.

Noise, the model's own case. Table 12.8 gives each loop's stationary cost per step, its spread on T_c and its valve travel per hour, with white noise only. The PI is a causal feedback on y_c, so by equation (12.18) and Section 7.4.2 it cannot beat the law on y_c on this cost and these rows, and it does not. The retuned row is the best PI on a grid of 80 values of K_p from 0.005 to 2 and 50 ratios K_i/K_p from 0.0005 to 0.2.

Table 12.8: The loops under the design's white noise, stationary, by the Lyapunov route on the linear plant.

Loop	J per step [–]	sd T_c [K]	Travel [/h]
Law + filter, all readings	0.00497	0.03345	0.23
Law + filter, y_c only (the PI's information)	0.00514	0.03420	0.23
PI 0.10, 0.002	0.01431	0.04632	0.48
PI retuned on this cost, $K_p = 0.0388$, $K_i = 0.000074$	0.00621	0.03708	0.27

Retuning decides fairness. At the vendor PI's noise travel, the law on y_c with $\rho = 0.1455$ has a spread on T_c of 0.03317 K against 0.04632 K. At the law's travel, the PI on the line $K_i/K_p = 0.02$, with $K_p = 0.0190$, has 0.03725 K against 0.03420 K, and $J = 0.00682$ against 0.00514. But **the retuned PI is no loop to run**: its load-step peak is 0.8611 K, and it is still 0.2080 K off after 100 min. A cost with white noise only does not ask for an integrator.

With the record's load wander. Table 12.9 adds a load wander with a step-to-step correlation of 0.95, which is not in the law's rows. Each row is a spread of that wander. The information is not alike across the columns: the law can read all the readings, and the PI reads y_c alone. The columns are the cost per step of the law unchanged on all readings, of the same law on y_c, which is the PI's information, of the vendor's PI, of a PI retuned at that wander, and of the law on all readings with its filter's **W** retuned to hold the wander; the gain **K** stays where it was.

Table 12.9: Cost per step with a load wander outside the law's rows, spread on T_c in K in parentheses. The record's own wander is 0.5 kW.

sd load [kW]	Law, all readings [–]	Law, y_c only [–]	PI 0.10, 0.002 [–]	PI retuned at wander [–]	Law, all, filter's **W** retuned [–]
0.5	0.00692 (0.0391)	0.00723 (0.0403)	0.01560 (0.0466)	0.00880 (0.0398)	0.00675 (0.0366)
1.0	0.01275 (0.0524)	0.01350 (0.0545)	0.01946 (0.0476)	0.01447 (0.0441)	0.01126 (0.0428)
2.0	0.03610 (0.0874)	0.03857 (0.0916)	0.03489 (0.0514)	0.03276 (0.0516)	0.02666 (0.0570)
4.0	0.12947 (0.1650)	0.13885 (0.1734)	0.09663 (0.0641)	0.09645 (0.0625)	0.08129 (0.0874)

The unchanged law loses to the vendor's PI on this cost from 2 kW on either set of readings. On the PI's own information it beats the retuned PI to 1 kW, 0.01350 against 0.01447, and loses from 2 kW. The retuned filter holds its lead with readings the PI does not have. On the spread of T_c the PI is better from 2 kW even against the retuned filter (0.0514 against 0.0570 K). The theorem is about the rows it was given.

The steps. Table 12.10 runs the two events on the plant at the fit, noiseless: a load step from 30 to 45 kW, and a set-point step from 18 to 17°C at 30 kW, with the law's target taken from the fit's steady state. The cost column is J over 40 steps, charged to the new valve, 0.8435 at the load step and 0.7534 at the set-point step. Figures 12.4 and 12.5 draw T_c's error against the event's target, with the horizontal axis in controller steps of 30 s.

Table 12.10: The steps on the plant at the fit, noiseless. Load: peak and final error of T_c. Set point: settling to 0.1 K. Commands above 1 are clipped by the plant: "computed" charges the command and "received" the clipped command the plant acts on. A row with a command above 1 is outside hypothesis 5, and the row at $\rho = 0.146$ most of all.

Loop	Load peak [K]	Load end [K]	Load J [–]	Set point: settle [s]	Set point J, computed [–]	Set point J, received [–]
Law, fed the truth	0.2977	+0.2977	34.15	390	11.68	11.64
Law + filter, all readings	**0.7271**	**+0.7271**	93.99	390	11.72	11.64
Law + filter, **W** with 0.5 kW wander	0.5787	+0.5787	72.43	390	11.70	11.64
Law at the PI's travel, $\rho = 0.146$, y_c; outside the theorem	0.5711	+0.5711	76.70	90 (command 1.981)	60.21	13.76
Law + filter, metered load fed forward	0.2205	0.0000	10.95	390	11.72	11.64
PI 0.10, 0.002	0.1977	0.0000	24.93	300	20.43	20.43
PI 0.41, 0.0082	0.0860	0.0000	17.51	120 (command 1.099)	17.31	15.84
PI at the law's travel, $K_p = 0.019$	0.5065	−0.0013	70.67	1230	52.24	52.24

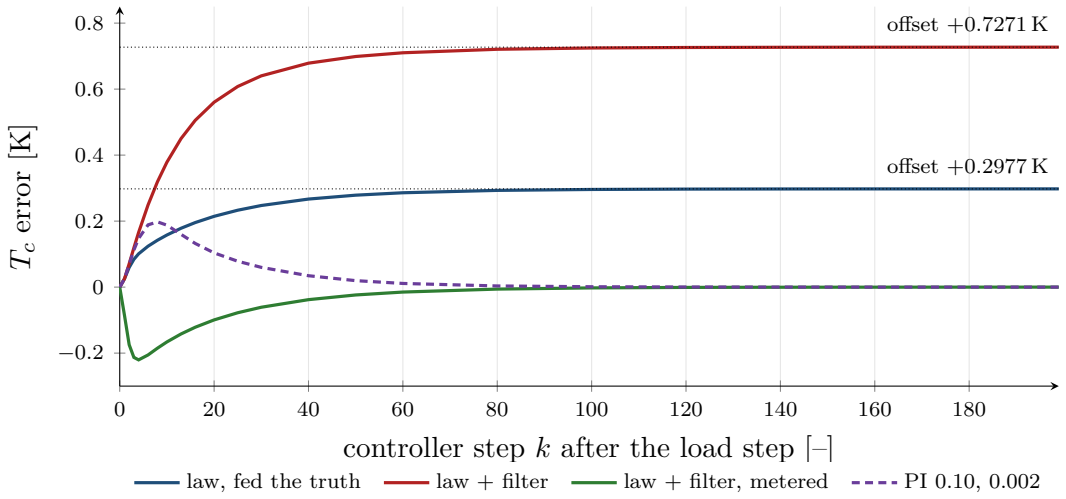


Figure 12.4: The load step from 30 to 45 kW on the plant at the fit, noiseless. Each curve is T_c's error against $r - \hat{b}$. Dotted: the two offsets the law leaves.

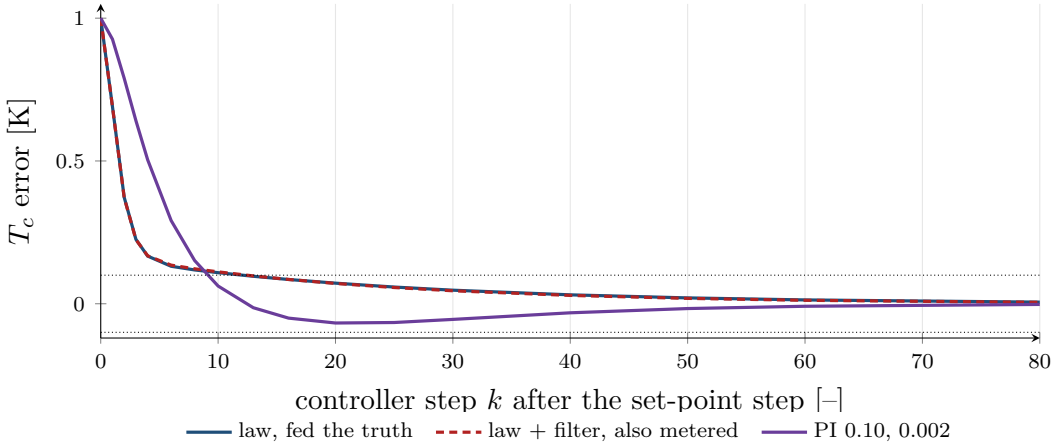


Figure 12.5: The set-point step from 18 to 17°C at 30 kW, same plant, the first 80 of 200 steps. Each curve is T_c's error against the new target. Dotted: the settling band of ±0.1 K. The metered row coincides with the law and filter, because the load does not change.

The set point: the law wins on this cost, by its target and its state feedback together. The law and filter settle later than PI 0.10, at 390 against 300 s, and cost less, 11.72 against 20.43. A step-sized start alone does not do it: PI 0.10 started at the same new valve, 0.7534, costs 18.97, and PI 0.41 costs 18.16, or 15.00 as received. Figure 12.5 shows what the law adds: the target, from the fit's steady state, puts the deviation on all three states at once, the feedback drives the valve past its new value, and the error falls in about four steps and then creeps. Its command touches the stop in a small way, 1.007 with the filter and 1.004 fed the truth, and that command is clipped. Charging the unclipped command misleads where it goes further: the 60.21 of the row at $\rho = 0.146$ is 13.76 on the command the plant receives, below PI 0.10 at 20.43, and PI 0.41 falls from 17.31 to 15.84. Those runs are outside hypothesis 5, so they say what the plant does and nothing about the theorem. The law holds r only at the fit, as Section 12.3.7 shows.

The load: the law leaves an offset; the PI does not. The rows hold no disturbance entry and the law has no integral action, so a load away from the pinned point leaves an error: the plant's +0.2977 K with the law fed the truth, and, by the linear route, +0.3665 K with the full state fed back against +1.9715 K with the valve left open loop. Figure 12.4 shows the two offsets as level lines. The PI's curve returns to zero. Building that entry and that action is the work of Chapter 13.

A result the design did not ask for: separation does not protect this case. The filter raises the offset from 0.2977 to 0.7271 K. Its rows hold no load, so its belief on T falls behind: on the linear route, +2.0052 K against the true +3.2746 K, with a T_c offset of +0.8200 K. That belief is 0.1501 of the 0.1670 by which the valve falls short at steady, 90 %. The theorem's noise has mean zero, which keeps the last term of equation (12.18) out of the expectation and makes the estimate's error unbiased. A step does not have mean zero, so the error has a mean. What it pays is the same square, biased; the trace no longer predicts it. On the linear plant over 40 steps, the gap in the pinned cost from the truth-fed law to the filter-fed one is 28.11, which that equation at the steady $\mathbf{S}$ and H splits into 22.69 from $H(v - \mathbf{K}\mathbf{x})^2$, +12.19 from the cross term and −6.77 at the terminal weight, so 81 % is the square. At steady that square is 0.927 per step, against $\operatorname{tr}(\mathbf{K}^\mathsf{T} H \mathbf{K} \boldsymbol{\Sigma}_\infty) = 0.001945$. The theorem is intact, and a hypothesis is what failed.

A second result: the offset is partly the cost's doing. The cost (12.3) charges distance from the pinned valve. Holding the valve the load needs would cost $\rho\,\delta v^2 = 1.9763$ per step,

while the law's own offset costs $q\,e^2 = 0.5373$ and the valve it holds costs 1.3098. Of the linear offset of 0.36652 K, 0.22235 K is the cost's, since it remains with the load known and the cost unchanged, and 0.14416 K is the missing entry's. A target, with the cost charged to the valve re-solved at the load, gives 0. The metered row is that target, and not only a feedforward, because the law re-solves its pinned point at the metered load. On the cost as it stands, over 40 steps on the linear plant, the law fed the truth pays 27.44, PI 0.10 pays 32.36, PI 0.41 pays 33.27, and the moves solved knowing the step pay 25.41, which is the bound, with nothing below it. So the truth-fed law beats both PIs on its own cost, and the PIs' zero offset is what that cost does not ask for. The loop that can be run does not: the law and filter pay 55.56 (47.43 on the plant, against 31.18 and 33.01), and stay worse at 2000 steps, 6630 against 3898. The cost column of Table 12.10, charged to the new valve, puts the PIs ahead of the truth-fed law as well, 24.93 against 34.15. Chapter 13 owes the entry, the target and the change of cost together. The metered row of Figure 12.4 works here because the meter is taken noiseless and the fit is the plant; that row also belongs to Chapter 13.

12.3.7 What fit error costs

Control engineer's question. What does fit error cost in closed loop?

General form. Moments of a derived quantity, equation (2.10), over the plant's posterior.

Specialized form. Equation (2.10) with $f(\mathbf{z})$ the regret on the plant $\mathbf{z}$,

$$f(\mathbf{z}) = \frac{J_{\text{stat}}\big(\text{law and filter at the fit, on } \mathbf{z}\big)}{J_{\text{stat}}\big(\text{the law and filter of } \mathbf{z}, \text{ on } \mathbf{z}\big)} - 1, \tag{12.19}$$

stationary at 30 kW, each by the Lyapunov route, and averaged by the draws of Section 12.2.1. Chapter 8's sensor lag is in no band here, which limits what this can say.

Building blocks. The draws; $\mathbf{K}$ and $\boldsymbol{\Upsilon}$ at the fit; each drawn plant's $(\mathbf{A}, \mathbf{B})$ and its own steady law.

The plant's numbers. Table 12.11 gives one row per band, and adds the stated truth as its last row, which the fit does not see. The columns are the share of plants the law holds stable, the regret's mean, median, 95 % quantile and largest value, and, for scale, the median excess of the vendor's PI over each plant's own law. Figure 12.6 draws the regret on the left, on a logarithmic axis, by band, with the truth's as a dotted marker.

Table 12.11: Regret of the law and filter at the fit. Draws: 3000, effective size 2005. The percentages are this example's.

Band	Stable [%]	Mean (se) [%]	Median [%]	95 % [%]	Largest [%]	PI over law [%]
Fit only	100.00	0.0387 (0.0015)	0.0195	0.1429	2.9974	187.577
+ dynamic, sd 0.15 (Chapter 8's)	100.00	0.2595 (0.0075)	0.1449	0.9205	3.1570	188.312
+ dynamic, sd 0.30	100.00	0.9020 (0.0263)	0.4530	3.2209	9.0169	189.532
Stated truth, one plant	stable	0.0779	–	–	–	–

Fit error barely matters to this law on this cost, since a regret is flat at the minimum and the law sits far below any instability. It is no claim of robustness in general. Chapter 8's sensor lag is in no band, and white noise is the only disturbance here.

The stated truth. The record's truth is $g = -0.03390$, $g_w = -0.10$, $b = 0.30\,\mathrm{K}$. The entries of $\mathbf{A}$ move by at most 0.01842 against the fit, and B_{T_c} is -0.74392 against -0.74494. The gain at the truth is $(0.11062, 0.03307, -0.50566)$ against $(0.11829, 0.03421, -0.51268)$ at the fit. On the truth's plant the stationary cost per step is 0.004683 at the fit's law and 0.004679 at its own, a regret of 0.0779 %. A kick, noiseless, gives 0.0015 %. The truth's dynamic entries are the priors' centers, so these rows test nothing about the priors, the caution of Chapter 8.

The offset passes through. On the truth at 30 kW, T_c settles at 17.6916°C under the law and 17.7000°C under the PI, with $r = 18$°C and $r - b = 17.70$°C. Across the posterior, the error of T_c against r is $-0.0538\,\mathrm{K}$ (sd 0.1796) under the law and $-0.0541\,\mathrm{K}$ (sd 0.1790) under the PI. **The law does as the PI does, with no integrator, because the fit read Q/u:** the valve at the fit's point is the truth's valve. At 45 kW with the point held at 30 kW the error is $+0.6735\,\mathrm{K}$ (sd 0.2005), and with the metered point $-0.0536\,\mathrm{K}$ (sd 0.1835); on the truth at 45 kW, the law gives 18.3878°C, the metered law 17.6851°C and the PI 17.7000°C. Off the fit's ridge, with the plant at the fit and a further offset δb on the reading, the law moves T_c by $-0.467\,\delta b$ with all readings and $-0.556\,\delta b$ with y_c alone (at $\delta b = 0.1\,\mathrm{K}$; at $0.3\,\mathrm{K}$, -0.462 and -0.549), against the PI's $-1.000\,\delta b$. That plant could not have produced this window, so the fraction is reported and not read as the law's advantage. At the set-point step on the truth, the reading settles at 17.0316°C under the law and 17.0000°C under the PI.

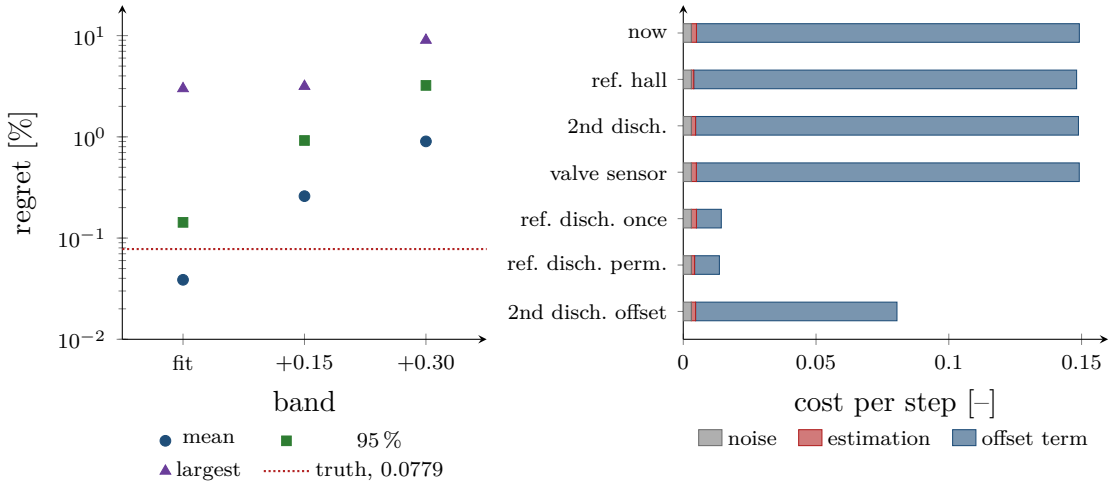


Figure 12.6: Left: the regret of Table 12.11 by band, on a logarithmic axis; the dotted marker is the stated truth. Right: the cost per step of Table 12.15 for each candidate, stacked: the noise term, the estimation term and the offset term.

12.4 The checks (step 5)

The answers have set their routes side by side. This section collects them and takes σ_p to its limit.

12.4.1 The limit of the physics row

On the recursion, the steady K_c at 200 moves is 0.015398, 0.031874, 0.034093, 0.034180, 0.034209 and 0.034210 at $\sigma_p = 0.5$, 0.1, 0.02, 0.01, 0.001 and $0\,\mathrm{K}$. The limit is the hard law, and the gap at $0.01\,\mathrm{K}$ is the size that Table 12.3 shows on the moves.

12.4.2 The routes against each other

Table 12.12 has one row per check, naming what is compared, the values and the result. The first four hold by algebra, at rounding. The rest are agreements with a stated size, and the restriction checks are inequalities that must hold if the computation follows the derivation.

Table 12.12: The checks of the chapter.

Check	Values	Result
Stack against recursion, same σ_p	at most 4.81×10^{-14}; tolerance 10^{-7}	holds, by algebra
Duality, covariance and gain	$4.9 \times 10^{-17}\,\mathrm{K}^2$, 6.7×10^{-16}; tolerance 10^{-10}	holds, by algebra
Closed-form minimizer against the law	at most 2.3×10^{-15}	holds, by algebra
Steady equation's residual in Joseph form	7.1×10^{-14}	holds, by algebra
Chapter 7's by-hand gains	0.062560, 0.063023, 4.12821 returned	agrees
Lyapunov against the closed form, stationary cost	0.004973 both; candidates' gaps at most 1.2×10^{-17}	agrees
Separation's gap against the term, two seeds	-0.27 se, $+0.50$ se; tolerance 3 se	within
The gap with other errors of the same variances	Laplace $+0.2$ se; two-point -3.0 se, -2.4 se, twelve more seeds $+0.2$ se; by second moments 1.7×10^{-15}	agrees
Records against the stationary cost, seed 1205, 2000 records, 400 steps	0.00496 (se 0.00002) against 0.00497, -0.43 se; reference hall sensor 0.00391 (se 0.00001) against 0.00393, -1.98 se	within 3 se
Restriction, noise	best PI 0.00621 $\geq$ law on y_c 0.00514 $\geq$ law, all 0.00497	holds
Restriction, load step on its own cost	nothing below the bound 25.41	holds
Output on a second run	compared byte by byte	identical

12.5 The report (step 6)

The band does not move the valve on this plant class, by algebra and to 10^{-15} on the hall; it costs 0.18710 of 4.94583 in estimation on this example. The derived law is not PI-shaped. On this cost it beats the PI under noise on either basis of comparison and at the set-point step, and on this example it leaves 0.7271 K under a load step, an offset the filter makes worse, and there the law with its filter loses to both PIs on its own cost, 55.56 against 32.36 and 33.27. The surprise is that last one: the theorem was not violated, since a step is outside its hypotheses, and the estimate still cost the loop more than twice the offset of the law fed the truth, because the rows hold no load. It is the same square, biased: 0.927 per step, where the trace predicts 0.001945.

12.5.1 The answers and their bands

Table 12.13 gives the answers in the order of Table 12.1. Each row is a question; the second column is the answer, the third its band or basis and the last the record's stated truth where one exists.

12.5.2 The entry in the commissioning log

The entry gives the law's gains, $(0.11829, 0.03421, -0.51268)$, with the cost that produced them, $q = 4$ and $\rho = 25$ charged on the valve's distance from the pinned point, and the statement that this cost is a choice. It gives the pinned point, 30 kW and $T_c = 17.9429$°C at $u = 0.5623$, one operating point with one unscheduled gain. It states that the law holds no disturbance entry,

Table 12.13: The report on the law. Percentages and costs are this example's, on the stated cost.

Control engineer's question	Answer	Band or basis	Stated truth
What is the law with no band and no limit?	$\mathbf{K} = (0.11829, 0.03421, -0.51268)$ at steady, by the recursion	$\rho/q = 6.25$, a choice; fit's point	(0.11062, 0.03307, −0.50566)
Is the recursion Chapter 9's in disguise?	An identity of equations, not one problem: $\mathbf{S}$ and $\mathbf{P}$ are unrelated on the hall	$4.9 \times 10^{-17}\,\mathrm{K}^2$; tolerance 10^{-10}	–
When does the band change the move?	Not under the seven hypotheses: 0.034201 at every s. With a gain's band, −1.74 % against the move at the belief's mean; with a binding limit, 0.12940 at $s = 2$	2.3×10^{-15}; records within 1.7 se	–
What does the band cost, if not the move?	0.18710 in estimation and 0.39267 at the start, of 4.94583	gap 0.18691 (se 0.00070), seed 1210	–
How does the law do against the PI?	Noise: 0.00514 on y_c against 0.00621 retuned, 0.01431 vendor. Set point: 11.72 against 20.43. Load: 0.7271 K offset against 0.0000, and 55.56 against 32.36 on the law's own cost	same cost, retuned; event named; equal travel in the text	–
What does fit error cost in closed loop?	Regret, mean 0.2595 % (se 0.0075); offset error −0.0538 K (sd 0.1796) against r	95 % at 0.9205 %; 100.00 % stable	0.0779 %; 17.6916°C
Which measurement most reduces the expected cost?	The law's term: reference hall sensor. The whole cost: one reference discharge reading	Table 12.15	–

so an offset is expected at any load away from that point: 0.7271 K with the filter at 45 kW, on the plant at the fit, noiseless. It states that the offset of the discharge sensor passes through, −0.0538 K (sd 0.1796) against r, under the law as under the PI. The law is not offered as a replacement for the vendor's PI on this record: the load step goes the other way, on the law's own cost as well, and the comparison at the set point turns on this cost.

12.5.3 Common mistakes

Table 12.14 lists five mistakes this chapter's plant makes easy. Each row names the mistake, its form on this hall and the correction.

Table 12.14: Five mistakes in reading a derived law.

Mistake	On this hall	Correction
Reading separation as "the band never moves the valve"	0.034201 at every s; but −1.74 % with a gain's band and 0.12940 with a limit that binds	check the seven hypotheses; they are sufficient, not necessary
Calling a law better or worse than a PI without the basis	0.00514 against 0.01431 on this cost; against the retuned PI 0.00621; the retuned PI peaks at 0.8611 K	state the cost, the basis and the event
Taking separation to protect a load step	0.2977 K fed the truth, 0.7271 K with the filter	a step does not have mean zero; put the disturbance in the rows
Reading duality as one problem	doubling $\mathbf{W}$ leaves $\mathbf{K}$ alone; doubling ρ leaves $\boldsymbol{\Upsilon}$ alone	it is an identity between two equations, with unrelated inputs; the map reverses the index, and $\mathbf{A}^{\mathsf{T}}$ is an adjoint, not $\mathbf{A}^{-1}$
Planning with a soft physics row on the future	at $\sigma_p = 0.10$ K the first move is 0.9319 of the law's	hold the future's rows hard and average their noise

12.6 The next measurement (step 7)

The rows are linear, so what each candidate would do to the law's cost is known before it is read, and the gain $\mathbf{K}$ is untouched by it.

12.6.1 The candidates

Control engineer's question. Which measurement most reduces the expected cost?

General form. The expected consequence (2.23) after a candidate's row, with the belief taken from the preposterior (4.73) before the reading exists.

Specialized form. Equation (2.23) with L the stationary cost per step and the candidate's row added to $\mathbf{C}$ and $\mathbf{V}$ of equation (12.5), which by equation (4.73) gives its steady $\mathbf{\Sigma}_\infty$. Per step, by equation (12.16),

$$J_{\text{stat}} = \operatorname{tr}(\mathbf{SW}) + \operatorname{tr}\big(\mathbf{K}^{\mathsf{T}} H \mathbf{K} \mathbf{\Sigma}_\infty\big). \tag{12.20}$$

The first term is 0.003028, which no reading reaches, and the second is 0.001945. Outside the law's model stands **the offset term**: with the set point corrected by $\hat{b}$, the error of T_c against r costs $q\,\mathrm{Var}(b) = 4.00 \times 0.03605 = 0.144193$ per step, whatever the move. That is the Laplace variance of Chapter 6, which applies to Table 12.15 and its closed-form update; the reweighted draws behind the spread of 0.1790 K in Section 12.3.7 give 0.1281, on which one reference reading removes 89.29 % in place of 90.40 %, and the ranking is the same. It is a term in $\mathbf{\Sigma}$ with no move in it, by the same theorem, and it is about thirty times the rest.

Building blocks. $\mathbf{S}$, $\mathbf{K}$, H of the steady law; $\mathbf{\Sigma}_\infty$ from each candidate; $\mathrm{Var}(b)$ from the belief of Section 12.2.1. A discharge reference reading of variance σ_a^2 takes that variance to $\mathrm{Var}(b) - \mathrm{Var}(b)^2/(\mathrm{Var}(b) + \sigma_a^2)$.

The plant's numbers. Table 12.15 has one row per candidate. The columns are the estimation term after it, the share of that term removed, the law's sum J, the offset term, and the share of J plus the offset term removed. The right panel of Figure 12.6 draws the same three terms, stacked.

Table 12.15: The candidates, on the law's cost at the steady filter with all readings. Percentages are this example's.

Candidate	Estim. [–]	Removed of it [%]	J [–]	Offset term [–]	Removed of J + offset [%]
Reference hall sensor 0.1 K, permanent	0.000905	53.46	0.003934	0.144193	0.70
Second discharge sensor 0.10 K, permanent	0.001607	17.36	0.004636	0.144193	0.23
Valve position sensor 0.002	0.001909	1.82	0.004938	0.144193	0.02
Reference discharge sensor 0.05 K, one reading	0.001945	0.00	0.004973	0.009351	90.40
Reference discharge sensor 0.05 K, permanent	0.001229	36.82	0.004257	0.009351	90.88
Second discharge sensor, its offset sd 0.2 K, permanent	0.001607	17.36	0.004636	0.075843	46.05

12.6.2 The recommendation

The ranking turns on the term. The law's own term wants the reference hall sensor, which removes 53.46 % of it. The whole expected cost wants one reference discharge reading, which removes 90.40 % on this example and nothing from the law's term. As a permanent instrument, the second discharge sensor with its own offset of 0.2 K removes 46.05 %. So the recommendation is the reference reading first. A load meter has no row here, because the law's state has no load entry, and its value is not a term in $\mathbf{\Sigma}$: that row is Chapter 13's. A test of C_c or τ_a is bounded by the regret of Table 12.11, with a mean of 0.2595 % at Chapter 8's prior.

12.6.3 What the analysis showed

Two colleagues gave two readings of the same band, and on this restriction the second one is right: the band prices the move's outcome without moving it. The gain is computed from the plant's mean and the cost, and the band only says what the result will cost. That settles less than the argument suggested. The law it derives has no integrator and no load row, and the first step the plant takes off its pinned point exposes that, by more with the filter than without it. The sheet's reading of 18.00 meanwhile says nothing about any of it. The engineer takes back three things: a derived law carries the rows and the cost it was derived from, so its offset is read there and not in the theorem; the PI's zero offset is a property this cost did not ask for, yet the loop that can be run, law and filter together, still loses to the PI on that same cost; and the largest term in the expected cost has no move in it, so the first thing worth buying is a measurement.

12.6.4 Limits of the model

The plant. This is one operating point and one unscheduled gain on a nonlinear plant. Steps are noiseless and at the fit; noise comparisons are linear and stationary. There is white noise and no load entry, and Chapter 8's sensor lag is in no regret.

The stated choices. The noises, the cost and the start $(0, 1, 0)$ are stated, and the start is not a state the loop visits, since the coil forgets it in 21.5 s. The ratio ρ/q moves every comparison with the PI. Equal travel and equal cost are two fair bases, and the line $K_i/K_p = 0.02$ is a choice at equal travel.

The stop. The stop is clipped by the plant and is not in the law. Table 12.10 gives J on the command computed and on the command received: 60.21 and 13.76 at the command 1.981, 17.31 and 15.84 at 1.099, and at most 1.007 in the law's other set-point rows. Those runs are outside hypothesis 5. Chapters 15 and 18 and Chapter 22 keep what falls here.

The constructed record. Percentages are this example's, and the record is constructed from a stated truth, so the agreements test the method on its own model. They do not test the model against a real hall. Adding the load to the rows, with integral action and a changed cost, is the work of Chapter 13.

12.7 What to retain

Definition 12.1 (The linear-quadratic law and certainty equivalence)**.** Eliminating the stacked rows of a known linear plant with a quadratic cost, from the last entry back, gives the Riccati recursion, equations (12.8) and (12.9), and its gain is the law. Under linear rows, a quadratic cost, additive noise of mean zero and independent of state and move, a known gain, no binding

limit, plant noise, reading errors and start independent of each other, and one controller with full recall, the move is that gain on the filter's mean, and the band adds the terms of equation (12.16) and leaves the move alone. The algebra needs second moments only; normal errors add that the filter is the best estimator. These are sufficient conditions for a restriction of the problem (7.1) that keeps couplings 1 and 2.

Three conclusions carry beyond the hall, as statements of form and not of size. The recursion is the same elimination as the filter's, so a law is a fit of rows; the identity with the filter is one of equations and not of problems. A theorem's hypotheses are about the rows given to it, so a disturbance the rows do not hold is outside it however good the estimator. And a comparison with a PI means nothing without its cost, its basis and its event. The chapter's questions and their forms are collected with the other chapters' in Appendix A.

Table 12.16 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 12.16: The terms of the chapter and where each is defined.

Term	Expression	Units	Meaning	Defined in
Sampled plant	eq. (12.1)	mixed	deviations at a pinned point	12.1.2
Cost	eq. (12.3)	–	the stated choice	12.1.2
Rows of the plan	eq. (12.4)	–	the stack	12.1.4
Filter, matrix form	eq. (12.5)	mixed	the belief on the state	12.2.2
Shrunk weight	$\tilde{\mathbf{S}}$, eq. (12.8)	–	the physics row eliminated	12.3.1
Riccati recursion, gain	H_k, $\mathbf{K}_k$, $\mathbf{S}_k$, eq. (12.9)	mixed	the move eliminated	12.3.1
Duality	eq. (12.12)	–	an identity between two equations	12.3.2
Expected cost	eq. (12.16)	–	the terms in the band	12.3.3
Cautious move	eq. (12.17)	–	the band on the gain reaches the move	12.3.4
Completed square	eq. (12.18)	–	what any move pays against the law	12.3.5
Regret	eq. (12.19)	–	excess over the plant's own law	12.3.7
Stationary cost	eq. (12.20)	–	the candidate's term	12.6.1

12.8 Exercises

Exercise 12.1. *Formulation.* A level loop has a tank with level L_k, a metered inflow F_k and a valve move v_k on the outflow. Write the stacked rows of its plan and name its couplings. (With $\mathbf{x}_k$ the deviations of the level and the valve, the rows are $\sqrt{\rho}\, v_k$, $(\mathbf{x}_{k+1} - \mathbf{A}\mathbf{x}_k - \mathbf{B}v_k - \mathbf{E}\, y_{F,k})/\sigma_p$ and $\mathbf{Q}^{1/2}\mathbf{x}_{k+1}$, the metered inflow being a known input. Couplings 1 and 2 are kept, since the move is an entry and produces the level readings. Couplings 3 and 4 fall under the same seven hypotheses.)

Exercise 12.2. *Derivation.* Derive the two strokes of the elimination. (Removing $\mathbf{x}_{k+1}$ by equation (9.4) leaves $\mathbf{M}^\mathsf{T}[\mathbf{I}/\sigma_p^2 - (\mathbf{I}/\sigma_p^2 + \mathbf{S})^{-1}/\sigma_p^4]\mathbf{M}$, and the bracket is $(\mathbf{I} + \sigma_p^2\mathbf{S})^{-1}\mathbf{S} = \tilde{\mathbf{S}}$. Removing v_k from the block with element $H_k = \rho + \mathbf{B}^\mathsf{T}\tilde{\mathbf{S}}\mathbf{B}$ gives the multiplier $-H_k^{-1}\mathbf{B}^\mathsf{T}\tilde{\mathbf{S}}\mathbf{A} = \mathbf{K}_k$ and the remainder $\mathbf{Q} + \mathbf{A}^\mathsf{T}\tilde{\mathbf{S}}\mathbf{A} + \mathbf{A}^\mathsf{T}\tilde{\mathbf{S}}\mathbf{B}\mathbf{K}_k$.)

Exercise 12.3. *Derivation.* Show why $\mathrm{tr}(\mathbf{S}\boldsymbol{\Sigma})$ has no move in it. (Equation (12.14) separates the mean from the covariance. The filter's rows (12.5) give $\boldsymbol{\Sigma}_{k|k}$ with no v and no $\mathbf{y}$, since the move sits in a residual and has no line in the Jacobian. The derivative of the term in v_k is zero.)

Exercise 12.4. *Derivation.* Show why an unknown gain breaks it, from $\mathbb{E}[\mathbf{B}^\mathsf{T}\mathbf{S}\mathbf{B}]$. (With $\bar{\mathbf{B}}$ the mean and $\mathbf{C}_B$ the covariance of $\mathbf{B}$, $\mathbb{E}[\mathbf{B}^\mathsf{T}\mathbf{S}\mathbf{B}] = \bar{\mathbf{B}}^\mathsf{T}\mathbf{S}\bar{\mathbf{B}} + \operatorname{tr}(\mathbf{S}\mathbf{C}_B)$ multiplies v^2, so the term in the band now holds the move. The minimizer is equation (12.17), whose denominator grows with the band. Its numerator $\mathbb{E}[\mathbf{B}^\mathsf{T}\mathbf{S}\mathbf{A}]$ moves too when an uncertain entry reaches both $\mathbf{A}$ and $\mathbf{B}$, and on the hall that is the larger part.)

Exercise 12.5. *Application.* Find the steady K_c at $\rho/q = 25$. (That is $q = 4$ and $\rho = 100$. Run equations (12.8) and (12.9) with $\sigma_p = 0$ from $\mathbf{S}_N = \mathbf{Q}$ on the pair (12.2) until they settle: $\mathbf{K} = (0.05318, 0.01136, -0.18704)$, so $K_c = 0.01136$ per K.)

Exercise 12.6. *Application.* Find the estimation term at $s = 2$ from the rest of the row. (By equation (12.16) and Table 12.4, the total is 6.31659, the term at the mean 4.25690, $\operatorname{tr}(\mathbf{S}_0\boldsymbol{\Sigma}_{0|0}) = 1.57067$ and the noise term 0.10917, so the estimation term is $6.31659 - 4.25690 - 1.57067 - 0.10917 = 0.37985$.)

Chapter 13

Integral Action

13.1 A flat trend and a sheet with no integrator

The trend after the 45 kW step is reassuring. The discharge reading has come back to its set point of 18°C, with a mean error of +0.0172 K over steps 140 to 159. The PI row on the commissioning sheet has $K_i = 0.002$ per kelvin second, a value given by its vendor with no reason attached. A second sheet offers a model-based controller. Its one line on the subject reads "a cost on the move gives integral action, so no offset", and its block list shows no integrator.

The engineer would conclude two things: that a flat reading means the air is held, and that an integrator is a line in a cost, to be bought with the cost. Neither sheet tests either conclusion. The question of this chapter is

where the integrator comes from, and what a flat reading does not say.

Jargon: integrator, disturbance entry, disturbance model, steady target, effort cost, move cost, velocity form, feedforward, controller's own stability

An *integrator* is a pole of the controller at $z = 1$: a part of the controller's state that adds up the reading's error. A *disturbance entry* is an entry d added to the state for something constant that nothing reads. It is a *disturbance model* once the place it enters is chosen: its column in the physics rows and its coefficient in the reading rows. The *steady target* is the state and valve position at which the plant rests under a unit of d with the reading at its set point. An *effort cost* is charged on the valve's own position and a *move cost* on its change between steps. The *velocity form* writes the plant in differences between steps. *Feedforward* moves the valve from a reading of the disturbance itself. A controller has its *own stability* when its poles lie on or inside the unit circle, so that with the loop opened it does not grow by itself; an integrator still ramps on a constant error.

Key idea

Integral action needs two things together: the entry must be estimated, and the law must steer it to a target that holds the reading. Then the controller has a pole at $z = 1$, by algebra, for any gains. A cost on the move does not make that pole. An offset on the sensor has the same columns and a law with no pole, and through a loop that does hold the reading it passes to the air with the opposite sign.

13.1.1 The restriction this chapter works under

This chapter lies on the row "Known gain, quadratic, no limit" of Table 7.6, a restriction of the problem (7.1) in plant class and objective. The plant's gain stays pinned at the fit, Chapters 22 to 25 drop that, and no command reaches a stop, Chapter 15 takes that up. Couplings 1 and 2

of Section 7.1.3 are kept: the move is an entry of the rows, and the move makes the readings. Couplings 3 and 4 go by result, which Section 13.3.8 checks on this plant. The linear-quadratic law's gain comes from the Riccati recursion of Chapter 12 and the filter from Chapter 9; neither is re-derived.

13.1.2 The closed-loop hall about the fit (step 0)

The plant is the closed-loop hall of Chapter 9, with entries $\mathbf{z}_k = (T_k, T_{c,k}, u_k)^\mathsf{T}$, the hall air, the discharge air and the valve, at 30 s steps. For design it is linear about Chapter 6's fit at 30 kW, $Q = 30.0581$ kW, $g = 0.01833$, $g_w = -0.16764$, $b = 0.05708$ K. Its operating point is $T = 22.8616$°C, $T_c = 17.9429$°C, $u = 0.5634$, and all entries below are deviations from it. For every record it is nonlinear at the stated truth. One step of the linear plant, with c_k the command, is

$$\mathbf{z}_{k+1} = \mathbf{A}\mathbf{z}_k + \mathbf{B}c_k, \qquad \mathbf{A} = \begin{bmatrix} 0.960333 & 0.026979 & -0.074372 \\ 0.323747 & 0.259699 & -2.1705 \\ 0 & 0 & 0.606531 \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} -0.014425 \\ -0.744639 \\ 0.393469 \end{bmatrix}. \tag{13.1}$$

Its poles are 0.97258, 0.60653 and 0.24745, the first the hall's slow mode with the valve held, 36.5 steps, and its zeros from the command to T_c are 0.95035 and −0.53666. The steady reading moves −6.9982 K per unit of valve and +0.13118 K/kW of load, so the 15 kW step, open loop, would move it +1.9676 K.

The scenario node. Every what-if here is closed loop. The set point is pinned at $r = 18$°C, the valve is free and determined by the controller's row, and no command on any record reaches a stop (share 0.0000). The records carry the load step from 30 to 45 kW at step 80, over steps 0 to 159, with no set-point move. Controllers start bumpless, a stated choice: the controller's state is set so that its output holds the valve where it is.

Figure 13.1 draws the plant. The green boxes are the free nodes of Figure 9.1's kind, joined by solid physical edges, the dashed arrow the discharge reading, and the blue block the controller. The purple boxes are the same entry d drawn at each of the four places this chapter tries: on the command, on the coil, on the hall and on the reading. The note at the bottom names the direction $\mathbf{v}$ that the next section derives.

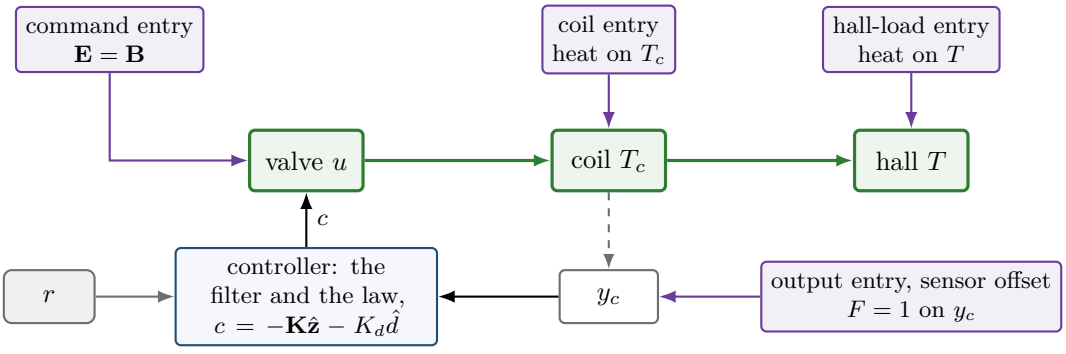


Figure 13.1: The chain with the disturbance entry d at each of its places. Green: free nodes. Gray: the pinned set point. Dashed: the reading y_c, the only reading the controller takes in most of the chapter. Purple: the four places of d, of which any one controller's model holds one. The bottom line is the direction the reading cannot see.

13.1.3 The control engineer's questions (step 1)

The engineer asks eight questions. Table 13.1 lists them. The first column is the question in the engineer's words, the second the general form of its category, the third its mathematical form on this hall, the fourth its category, the fifth the computation and the last the subsection.

Table 13.1: The questions asked of the integrator on the closed-loop hall.

Control engineer's question	General form	Mathematical form	Category	Computation	Answered in
Where is the integrator?	expected consequence, the law, (2.23)	a pole of $\mathbf{A}_c$ at $z = 1$	design	expected consequence	13.3.1
What does the cost on the move do?	the same, by cost, (2.23)	steady y against the cost row	design	expected consequence	13.3.2
Which disturbance does the loop assume?	moments, then the law, (2.10)	$\mathbf{E}$, F of the entry; load-step response	design	posterior	13.3.3
How many entries can one reading carry?	moments, (2.10)	rank of the steady rows	estimation	posterior	13.3.4
Does the meter remove the need?	expected consequence, (2.23)	feedforward against the entry, two events	design	expected consequence	13.3.5
What can it not remove?	threshold at a scenario node's entry, (2.12)	$T_c \to r - b$; $\mathbf{\Sigma}$ of b	estimation	posterior	13.3.6
Is the derived controller better than the PI?	expected consequence, fresh records, (2.23)	travel, scatter, mean square, margins	design	expected consequence	13.3.7
Does the band of d enter the law?	expected consequence, (2.23)	$\mathbf{K}, K_d$ against $\mathbf{\Sigma}$	design	expected consequence	13.3.8

13.1.4 The evidence and the target as rows (step 2)

The entry. A disturbance model adds one entry, with a column $\mathbf{E}$ in the physics and a coefficient F on the reading,

$$\mathbf{z}_{k+1} = \mathbf{A}\mathbf{z}_k + \mathbf{B}c_k + \mathbf{E}d_k, \qquad d_{k+1} = d_k + w_{d,k}, \qquad y_k = \mathbf{C}\mathbf{z}_k + Fd_k + \varepsilon_k, \tag{13.2}$$

with y_k the discharge reading less r and $\mathbf{C} = (0, 1, 0)$. Four models are tried. The hall load has $\mathbf{E} = (0.008156, 0.001664, 0)^{\mathsf{T}}$ per kW, the coil heat $\mathbf{E} = (0.001664, 0.054642, 0)^{\mathsf{T}}$ per kW, the command $\mathbf{E} = \mathbf{B}$, and in each $F = 0$; the output entry has $\mathbf{E} = \mathbf{0}$ and $F = 1$. The constant is carried by a walk in the filter. The walk buys a filter that keeps listening to the reading and it fails for a disturbance that moves faster than the walk says, which Section 13.3.3 measures.

The rows. With $\mathbf{z}_{a,k} = (\mathbf{z}_k^{\mathsf{T}}, d_k)^{\mathsf{T}}$ the stack of Chapter 9 holds, for each step, the rows

$$\mathbf{r} = \begin{bmatrix} \boldsymbol{\sigma}_x^{-1}(\mathbf{z}_{k+1} - \mathbf{A}\mathbf{z}_k - \mathbf{B}c_k - \mathbf{E}d_k) \\ (d_{k+1} - d_k)/\sigma_d \\ (y_k - \mathbf{C}\mathbf{z}_k - Fd_k)/\sigma_c \end{bmatrix}, \qquad \mathbf{J} = \begin{bmatrix} -\boldsymbol{\sigma}_x^{-1}\mathbf{A} & -\boldsymbol{\sigma}_x^{-1}\mathbf{E} & \boldsymbol{\sigma}_x^{-1} & \mathbf{0} \\ \mathbf{0} & -1/\sigma_d & \mathbf{0} & 1/\sigma_d \\ -\mathbf{C}/\sigma_c & -F/\sigma_c & \mathbf{0} & 0 \end{bmatrix}, \tag{13.3}$$

in the columns $\mathbf{z}_k$, d_k, $\mathbf{z}_{k+1}$, d_{k+1}. The command sits in the residual and has no line in the Jacobian, the zero of equation (7.6) again. The information matrix of these normalized rows is $\mathbf{J}^{\mathsf{T}}\mathbf{J}$; with unnormalized rows it is $\mathbf{J}^{\mathsf{T}}\mathbf{W}\mathbf{J}$. The entry's weight lies in the physics rows or in a reading row, since no entry here has both $\mathbf{E} \neq \mathbf{0}$ and $F \neq 0$, and that column is what this chapter turns on.

The steady target. A constant d has a resting place with the reading at its set point. It is the pair $(\mathbf{z}_s, u_s)$, per unit of d, that solves

$$(\mathbf{A}-\mathbf{I})\mathbf{z}_s + \mathbf{B}u_s + \mathbf{E} = \mathbf{0}, \qquad \mathbf{C}\mathbf{z}_s + F = 0, \tag{13.4}$$

four equations in four unknowns. The law is obtained by changing coordinates. With $\tilde{\mathbf{z}}_k = \mathbf{z}_k - \mathbf{z}_s d$ and $\tilde{c}_k = c_k - u_s d$, for constant d,

$$\tilde{\mathbf{z}}_{k+1} = \mathbf{A}\mathbf{z}_k + \mathbf{B}c_k + \mathbf{E}d - \mathbf{z}_s d = \mathbf{A}\tilde{\mathbf{z}}_k + \mathbf{B}\tilde{c}_k + \big[(\mathbf{A}-\mathbf{I})\mathbf{z}_s + \mathbf{B}u_s + \mathbf{E}\big]d = \mathbf{A}\tilde{\mathbf{z}}_k + \mathbf{B}\tilde{c}_k, \tag{13.5}$$

and $y = \mathbf{C}\tilde{\mathbf{z}} + (\mathbf{C}\mathbf{z}_s + F)d = \mathbf{C}\tilde{\mathbf{z}}$. In these coordinates d has gone from the rows. The cost about the target is $\sum_k q\, y_k^2 + \rho\, \tilde{c}_k^2$, whose law by Chapter 12 is $\tilde{c} = -\mathbf{K}\tilde{\mathbf{z}}$ with the gain of $(\mathbf{A}, \mathbf{B}, q\mathbf{C}^{\mathsf{T}}\mathbf{C}, \rho)$. Returning to the plant's coordinates and using the fit's estimates in place of the entries, which is the row's restriction,

$$c_k = -\mathbf{K}\hat{\mathbf{z}}_k - K_d\hat{d}_k, \qquad K_d = -(\mathbf{K}\mathbf{z}_s + u_s). \tag{13.6}$$

The stated choices at step 0. The costs are $q = 4$ and $\rho = 25$, Chapter 7's spreads of 0.5 K and 0.2. They buy gains of a known origin, and a different pair changes the sizes throughout; Chapter 14 takes up the choice. The filter's plant noise has spreads 0.001, 0.001 and 0.0001 per step, and the walk is 0.30 kW per step, both Chapter 10's. The other entries' walks are matched to it by steady effect on the reading, which buys a like comparison of the models. The linear design at one operating point buys constant gains, and it fails away from the fit by what every record, run at the truth, then shows.

13.2 The posterior of the state (step 3)

Every answer reads one posterior, the fit's belief on the state with its entry, in the steady form its recursion settles to. The controller is that belief with the law applied.

13.2.1 The belief on the state and its entry

The rows (13.3) with the hall-load entry have $\mathbf{A}_a$ and $\mathbf{C}_a$ as below, with $\mathbf{B}_a = (\mathbf{B}^{\mathsf{T}}, 0)^{\mathsf{T}}$. The recursion (9.14)–(9.15), in matrix form, reaches a steady gain $\boldsymbol{\Upsilon}$, the symbol of Chapter 9, and a steady covariance,

$$\mathbf{A}_a = \begin{bmatrix} \mathbf{A} & \mathbf{E} \\ \mathbf{0} & 1 \end{bmatrix}, \quad \mathbf{C}_a = \begin{bmatrix} \mathbf{C} & F \end{bmatrix} = (0,1,0,0), \quad \boldsymbol{\Upsilon} = \begin{bmatrix} 0.270837 \\ 0.115463 \\ -0.000002 \\ 2.821494 \end{bmatrix},$$

$$\boldsymbol{\Sigma} = \begin{bmatrix} 6.3772\times10^{-3} & 2.7084\times10^{-3} & 1.3006\times10^{-8} & 6.9199\times10^{-2} \\ 2.7084\times10^{-3} & 1.1546\times10^{-3} & -1.8622\times10^{-8} & 2.8215\times10^{-2} \\ 1.3006\times10^{-8} & -1.8622\times10^{-8} & 1.5820\times10^{-8} & 1.5097\times10^{-7} \\ 6.9199\times10^{-2} & 2.8215\times10^{-2} & 1.5097\times10^{-7} & 1.3025 \end{bmatrix}, \tag{13.7}$$

on (T, T_c, u, d), in K, K, [–], kW. The spreads are 0.0799 K, 0.0340 K, 0.000126 and 1.141 kW, with correlations of d with T and T_c of 0.759 and 0.728. The sensor's spread is 0.10 K. A correlation of T with T_c of 0.998 says the reading ties the two nearly into one number. The last entry of $\boldsymbol{\Upsilon}$ moves the estimated load 2.821 kW per kelvin of innovation.

13.2.2 The controller the belief and the law make

With $\boldsymbol{\xi}_k$ the predicted belief, $\hat{\mathbf{z}}_{a,k} = (\mathbf{I} - \boldsymbol{\Upsilon}\mathbf{C}_a)\boldsymbol{\xi}_k + \boldsymbol{\Upsilon} y_k$, and $\mathbf{K}_a = (\mathbf{K}, K_d)$, the law (13.6) and one predict give the controller from reading to command,

$$\begin{aligned} \boldsymbol{\xi}_{k+1} &= \mathbf{A}_c\boldsymbol{\xi}_k + \mathbf{B}_c y_k, \quad \mathbf{A}_c = (\mathbf{A}_a - \mathbf{B}_a\mathbf{K}_a)(\mathbf{I} - \boldsymbol{\Upsilon}\mathbf{C}_a), \quad \mathbf{B}_c = (\mathbf{A}_a - \mathbf{B}_a\mathbf{K}_a)\boldsymbol{\Upsilon}, \\ c_k &= \mathbf{C}_c\boldsymbol{\xi}_k + D_c y_k, \quad \mathbf{C}_c = -\mathbf{K}_a(\mathbf{I} - \boldsymbol{\Upsilon}\mathbf{C}_a), \qquad\qquad D_c = -\mathbf{K}_a\boldsymbol{\Upsilon}. \end{aligned} \tag{13.8}$$

At $\rho = 25$ and the hall-load entry, $\mathbf{K} = (-0.11817, -0.034141, 0.512083)$ and equation (13.4) gives $\mathbf{z}_s = (0.163639, 0, 0.018744)^{\mathsf{T}}$ and $u_s = 0.01874$ per kW, so that $K_d = -0.00901$. Numerically,

$$\mathbf{A}_c = \begin{bmatrix} 0.958628 & -0.258848 & -0.066986 & 0.008026 \\ 0.235753 & 0.157597 & -1.789183 & -0.005042 \\ 0.046496 & -0.010707 & 0.405042 & 0.003543 \\ 0 & -2.821494 & 0 & 1 \end{bmatrix}, \qquad \mathbf{B}_c = \begin{bmatrix} 0.285334 \\ 0.07668 \\ 0.024141 \\ 2.821494 \end{bmatrix}, \tag{13.9}$$

with $\mathbf{C}_c = (0.11817, -0.027216, -0.512083, 0.009006)$ and $D_c = 0.061357$ per kelvin. Nothing in the display is labeled an integrator.

13.3 The answers (step 4)

Each answer reads the belief and the controller of Section 13.2. Where a record is run it is nonlinear at the truth, and the route is named.

13.3.1 The pole at one

Control engineer's question. Where is the integrator?

General form. The law is an action in the expected consequence (2.23), and the problem (7.1) is restricted to its row. The question is a property of the policy that results.

Specialized form. Equation (2.23) with the action the law (13.6), under the cost about the target. Take the unit-entry direction with the target's state,

$$\mathbf{v} = \begin{bmatrix} \mathbf{z}_s \\ 1 \end{bmatrix}. \tag{13.10}$$

Two facts follow from equation (13.4) and the definition of K_d. The second row of the target gives

$$\mathbf{C}_a\mathbf{v} = \mathbf{C}\mathbf{z}_s + F = 0, \tag{13.11}$$

so $\mathbf{v}$ is a direction the reading cannot see, and $(\mathbf{I} - \boldsymbol{\Upsilon}\mathbf{C}_a)\mathbf{v} = \mathbf{v}$ whatever $\boldsymbol{\Upsilon}$ is. Next, $\mathbf{K}_a\mathbf{v} = \mathbf{K}\mathbf{z}_s + K_d = -u_s$, so

$$(\mathbf{A}_a - \mathbf{B}_a\mathbf{K}_a)\mathbf{v} = \begin{bmatrix} \mathbf{A}\mathbf{z}_s + \mathbf{E} + \mathbf{B}u_s \\ 1 \end{bmatrix} = \begin{bmatrix} \mathbf{z}_s \\ 1 \end{bmatrix} = \mathbf{v}, \tag{13.12}$$

by the first row of the target. Therefore, with the definition in equation (13.8),

$$\boxed{\mathbf{A}_c\mathbf{v} = \mathbf{v}}, \tag{13.13}$$

exactly, by algebra, for any $\mathbf{K}$ and any filter gain, under two hypotheses: the target (13.4) exists, and $\mathbf{C}_a\mathbf{v} = \mathbf{0}$ on every reading row. The first is not automatic. Take the reading $T - T_c$

in place of y_c: the target's matrix has rank 3 of 4 for any entry, because at rest the hall ties $T - T_c$ to the load whatever the valve does. The load entry's right side is then inconsistent, with a residual of 0.008, and the coil entry's is consistent with no unique answer. With the hypotheses met, the controller has a pole at $z = 1$ because the entry is estimated and the law holds the target, and no cost has been mentioned in this derivation. Equation (13.13) holds with no Riccati gain and no Kalman gain in it.

It is an integrator in the plain sense. Let $\mathbf{w}^\mathsf{T}$ be the left vector of $\mathbf{A}_c$ at $z = 1$, with $\mathbf{w}^\mathsf{T}\mathbf{v} = 1$. Then $I_k = \mathbf{w}^\mathsf{T}\boldsymbol{\xi}_k$ obeys

$$I_{k+1} = \mathbf{w}^\mathsf{T}\mathbf{A}_c\boldsymbol{\xi}_k + \mathbf{w}^\mathsf{T}\mathbf{B}_c y_k = I_k + (\mathbf{w}^\mathsf{T}\mathbf{B}_c)\, y_k, \tag{13.14}$$

a running sum of the reading's error. Its share of the command is $\mathbf{C}_c\mathbf{v}\, I_k$, and $\mathbf{C}_c\mathbf{v} = -\mathbf{K}_a\mathbf{v} = u_s$. So the integral gain is

$$K_i\, \Delta t = u_s\, (\mathbf{w}^\mathsf{T}\mathbf{B}_c). \tag{13.15}$$

The target gives the valve per unit of entry, and the filter gives the rate at which the reading's error is taken as entry. So a pole is not yet integral action: that takes $u_s \neq 0$, which, with the state observable at $z = 1$, is the same as full rank of the steady rows (13.18). With $\mathbf{E} = (\mathbf{A} - \mathbf{I})\mathbf{e}_1$ the target is state alone, $u_s = 0$: rank 3 of 4, $\mathbf{C}_c\mathbf{v} = 0$, the pole has no share of the command, the closed loop keeps a pole at 1, and no steady filter is stable. With a stable loop, one reading and $\mathbf{w}^\mathsf{T}\mathbf{B}_c \neq 0$, a steady I_k requires $y = 0$. With several reading rows it requires only $\mathbf{w}^\mathsf{T}\mathbf{B}_c\mathbf{y} = 0$, a weighted sum, which Section 13.3.4 takes up.

Building blocks. $\mathbf{A}_c$, $\mathbf{B}_c$ and $\mathbf{C}_c$ of equations (13.8) and (13.9); $\boldsymbol{\Upsilon}$ of equation (13.7); the target (13.4).

The plant's numbers. For the hall-load entry $\mathbf{v} = (0.163639, 0, 0.018744, 1)^\mathsf{T}$ per kW. The first row of equation (13.9) gives 0.163639 back, and the second entry of $\mathbf{v}$ is the zero of equation (13.11). The computed residuals $\|\mathbf{A}_c\mathbf{v} - \mathbf{v}\|$ are at most 2.8×10^{-17} for the four models, and they stay there over 200 random pairs $(\boldsymbol{\Upsilon}, \mathbf{K})$, seed 1313, against a stated tolerance of 10^{-9}. With $\mathbf{w}^\mathsf{T}\mathbf{B}_c = 3.792475$ per kelvin and $u_s = 0.018744$, equation (13.15) over 30 s gives $K_i = 0.002370$ per kelvin second, beside the vendor's 0.002. Setting K_d 10 % off its target, with the same filter, moves the nearest pole to a gap of 0.01266 from one: the integrator is gone.

The scalar check. On Chapter 7's one-node example with an entry, $x_{k+1} = ax_k - \theta u_k + d$, where $a = 0.36788$ and $\theta = 4.48780$, read x without noise. One step of the rows gives $\hat{d}_k = x_k - ax_{k-1} + \theta u_{k-1}$. The target is $u_s = d/\theta$, so the law is $u_k = Kx_k + \hat{d}_k/\theta$, with $K = 0.063026$. Subtracting u_{k-1},

$$u_k - u_{k-1} = \Big(K + \frac{1}{\theta}\Big)x_k - \frac{a}{\theta}\, x_{k-1} = \frac{a}{\theta}(x_k - x_{k-1}) + \Big(K + \frac{1-a}{\theta}\Big)x_k, \tag{13.16}$$

exactly a PI, with $K_p = a/\theta = 0.08197$ and $K_i\, \Delta t = 0.20388$, by algebra. The matrices at a walk of 0.05 and a sensor's spread of 0.10 K give controller poles 1.00000 and 0.04514.

The velocity form. Differencing the same example removes d from the rows: with the state $(x_k - x_{k-1}, x_k)$ and a move cost, the law is $\Delta u = +0.067489\, \Delta x + 0.168139\, x$, with a pole at one. The same numbers come by another route. The move's law on (x, u_{k-1}), with gains K_1 and K_2, fed with $\hat{d}$ read from one step as above, gives $\Delta u = K_2(a/\theta)\Delta x + (K_2(1-a)/\theta - K_1)x$, with gaps of 1.4×10^{-17} and 5.6×10^{-17}. So **on this scalar example differencing is the estimator of d with its gain at one**. On the hall, with three states, it is three entries, one per state, read with gain identity: the move's law through them equals the differenced state's law, largest gap 1.9×10^{-14} over 20 test vectors, seed 1313. The integrator comes with the differencing, and the cost on the move only sets the two gains.

13.3.2 The cost on the move

Control engineer's question. What does the cost on the move do?

General form. The expected consequence (2.23) at three cost rows, read at the steady state of a constant disturbance.

Specialized form. Equation (2.23) with the action the steady valve position and ℓ one of $qy^2 + \rho c^2$, $qy^2 + \rho(c - u_s d)^2$ and $qy^2 + \rho(c - c_{\text{prev}})^2$. Under the last two, the steady state with $y = 0$ costs nothing, so the target (13.4) is the resting place. Under the first it is not. With g_u and g_d the steady reading per unit of valve and of d,

$$\min_u \; q(g_u u + g_d)^2 + \rho u^2 : \qquad u = -\frac{q g_u g_d}{\rho + q g_u^2}, \qquad y = \frac{\rho}{\rho + q g_u^2}\, g_d. \tag{13.17}$$

The share $\rho/(\rho + q g_u^2)$ is left on the reading, whether d is known or estimated, and the controller has no pole at one.

Building blocks. $g_u = -6.9982\,\text{K}$ and $g_d = 1.9676\,\text{K}$ for the step; the controller (13.8) under each law.

The plant's numbers. Table 13.2 gives the steady reading error for the 15 kW step on the fit's linear plant. Each row is a cost; the columns are the law with d known, with d estimated, the same cost with no entry and a filter, and whether the controller has its pole at one.

Table 13.2: Steady reading error for the 15 kW step, the fit's linear plant, $\rho = 25$; open loop +1.9676 K.

Cost row	d known [K]	d estimated [K]	No entry [K]	Controller's pole at one
Effort ρc^2	+0.22268	+0.22268	+1.94808	no
About the target $\rho(c - u_s d)^2$	0	0	+1.94808	with the entry
Move $\rho(\Delta c)^2$	0	0	+1.94475	with the entry

The table reads by rows. **The sheet's line is false as stated.** With no entry and the state read exactly, a move cost at $\rho = 25$ gives a controller's pole of 0.44654 and leaves +0.09234 K. With the entry a move cost leaves no offset, and neither does the cost about the target, so the move's cost is not needed even for that. With no entry and the filter it leaves +1.94475 K. Effort's share is 0.11317 here, and 0.00508, 0.33795 and 0.83619 at $\rho = 1$, 100 and 1000. Charged on the valve's whole position it also fights the steady input: holding 30 kW at the fit's 0.5634 leaves +0.44623 K, with the valve short by 0.06376.

The no-entry column is mostly the filter's doing, so it is split. With the state read without noise and the move cost on $(\mathbf{z}, c_{\text{prev}})$, the error is +0.27482, +0.09234, +0.04282 and +0.02570 K at $\rho = 1000$, 25, 1 and 0.01, with the controller's pole at 0.80449, 0.44654, 0.13332 and 0.00424, never at one. The differenced state, read the same way, holds the reading at zero, with the closed loop's radius at about 0.950. At the truth, on the nonlinear plant with no noise, target and move leave 0.00000 K with the entry, effort +0.20728 K, and the laws with no entry +1.78734 and +1.79158 K.

Against Chapter 12. That chapter reports this law leaving +0.2977 K fed the truth and +0.7271 K with its filter, and +1.948 K here looks like a contradiction. The law is the same; the filter is not. Those two are on the nonlinear plant at the fit, linearized at $u = 0.5623$, gain −7.0118 K, and that filter reads y_c, y_T and y_u with plant noise of spreads 0.020, 0.020 and 0.002. The filter here reads y_c alone with spreads 0.001, 0.001 and 0.0001, at $u = 0.5634$, gain −6.9982 K, with the same ρ and step. Run at that chapter's point, this script returns its +0.2977 and +0.7271 K and, linear, its +0.3665 and +0.8200 K. On the linear plant with the state exact, the law at this chapter's point leaves +0.3667 K. From +0.8200 K, moving to this operating point gives +0.8197, reading y_c alone +0.8678 and this plant noise +1.9481 K: the point accounts for 0.0003 K, the readings for 0.048 and the plant noise for 1.080. At the truth the two filters move T_c by +0.6967 and +1.8115 K at this point, and the reading's errors are +0.6853 and +1.7918 K. A real difference remains, and it is the filter's stated noise, a stated choice in each chapter: a filter told that its model is nearly exact hardly moves on the reading. The law's own share is the +0.3667 K, and the rest is the filter failing to carry a load it has no entry for.

13.3.3 Where the entry goes

Control engineer's question. Which disturbance does the loop assume?

General form. Moments (2.10) of the entry under its rows, then the law; every model has a pole at one by equation (13.13), so the model chooses what else the controller is.

Specialized form. Equation (2.10) with $f(\mathbf{z}) = d$ under the model's $(\mathbf{E}, F)$ of equation (13.2). The same $\mathbf{K}$ serves all four, since it is computed with no reference to d. What differs is K_d through equation (13.4) and $\boldsymbol{\Upsilon}$ through the columns. The assumption entering here is that one constant carries everything the model does not hold; it fails when the event is at a place the entry does not name, which the T_w row tests.

Building blocks. $\mathbf{K}$ and the four targets; the controller (13.8); the given PI as a reference.

The plant's numbers. Table 13.3 sets the four controllers beside the given PI at $\rho = 25$. The columns give K_d, the controller's poles and zeros, the nearest PI by least squares on the relative error up to the crossover, the largest relative error up to the crossover and up to three times it, and the noiseless load step's peak and steps to stay within 0.02 K at the truth. Figure 13.2 draws the poles and zeros of the four controllers against the plant's poles. Each panel is the z-plane near the positive real axis, with crosses for the controller's poles, open circles for its zeros and gray triangles for the plant's poles.

In frequency, at this ρ and walk, a PI is a fair reduction for the coil and command entries, with 15.4 and 15.1 % relative error up to the crossover. In the figure those two controllers have a slow pole and a slow zero nearly together. The fairness stops there. On the records (seeds 13000 to 13999, count 1000) the nearest PI travels ×1.499 and ×1.483 of its controller, 0.5357/h (se 0.0020) against 0.3574 and 0.5176/h (se 0.0020) against 0.3490, with mean square up 0.00150 K^2 (se 0.00002) and 0.00141 K^2 (se 0.00001). Nor is the ranking general: the hall-load entry at walks of 0.03 and 0.10 kW has an error of 0.060 and 0.063, closer than the coil's 0.154, and the coil's own is 0.333 at $\rho = 316$. At $\rho = 25$ and walk 0.30 kW the hall-load controller has its slow pole at 0.97171 with the zero off at 0.93908, a second slow mode that no PI has. Its nearest PI has a negative K_p, a gain margin of 1.86 and a phase margin of 30.5°, and on the records (seeds 13000 to 13999) it travels 0.4872/h (se 0.0026) against the controller's 0.2536 (se 0.0011), with the mean square equal: paired difference −0.00018 K^2 (se 0.00013). That is no reduction. **The output entry's zeros are the plant's poles**: it cancels the hall's

Table 13.3: The four entries on one reading, $\rho = 25$, the walk matched at 0.30 kW, and the given PI.

Entry	K_d [–]	Controller's poles	Zeros	Nearest PI: K_p [/K], K_i [/(K s)]	Rel. error to crossover, ×3	Peak [K]	Steps to 0.02 K
Given PI	–	1	–	0.10, 0.002	–	+0.1868	52
Hall load	−0.00901	1, 0.97171, 0.27478 ± 0.15633i	0.93908, 0.60654, 0.24749, 0	−0.05368, 0.001831	0.610, 1.358	+0.2406	68
Coil heat	−0.02834	1, 0.95084, 0.25035 ± 0.17202i	0.95956, 0.60654, 0.24510, 0	0.13374, 0.001246	0.154, 0.330	+0.2139	74
Command	+1.00000	1, 0.95093, 0.25735 ± 0.16861i	0.95936, 0.58913, 0.24797, 0	0.12982, 0.001289	0.151, 0.330	+0.2124	72
Output	−0.36838	1, 0.92917, 0.21694 ± 0.18123i	0.97259, 0.60653, 0.24746, 0	0.26846, 0.000763	0.568, 1.067	+0.1969	130

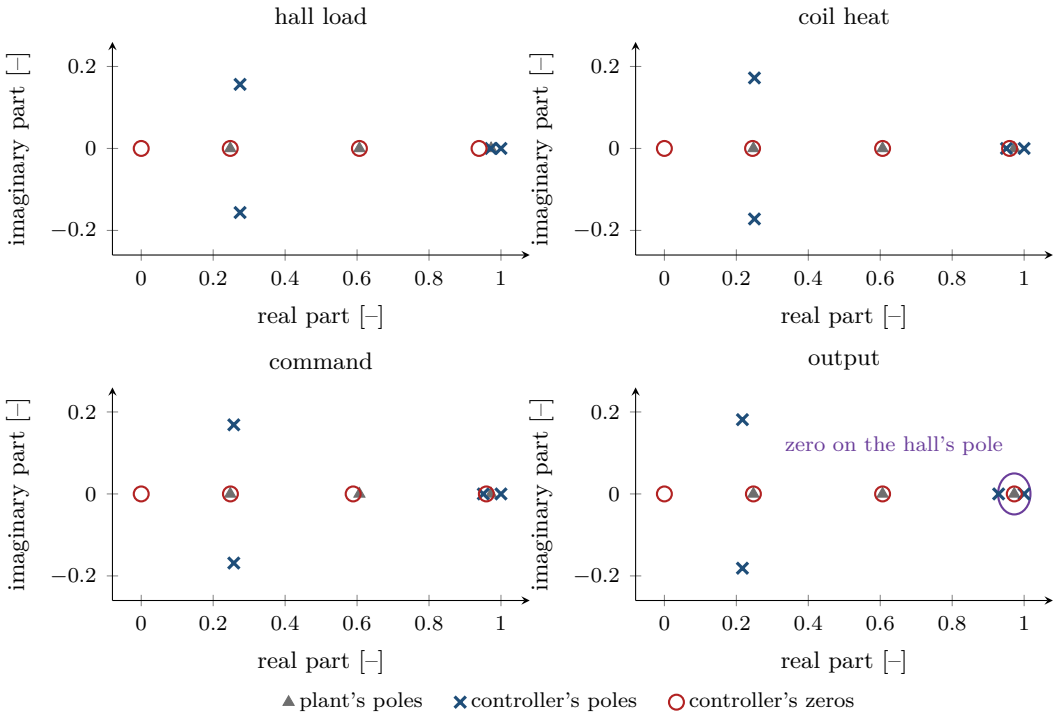


Figure 13.2: Poles and zeros of the four controllers of Table 13.3 against the plant's. Every controller has its pole at $z = 1$. In the output panel the purple ring marks the zero at 0.97259 on the plant's 0.97258.

mode, so the load's tail stays, and it takes 130 steps. Read noiselessly at step 399, its $\hat{d}$ has moved +1.95904 K, and the load and coil entries' have moved 14.93931 kW each.

The same travel. The rows above are at one ρ, and their travel differs. Figure 13.3 repeats the comparison with each model's ρ and walk selected to the same predicted travel budget, the selection Section 13.3.7 states. The horizontal axis is the controller step, the vertical axis the mean T_c over 1000 fresh records, seeds 13000 to 13999, about its mean over steps 40 to 79. The left panel is the load step at step 80 and the right panel a step of +0.5 K in T_w at step 80 with the load held.

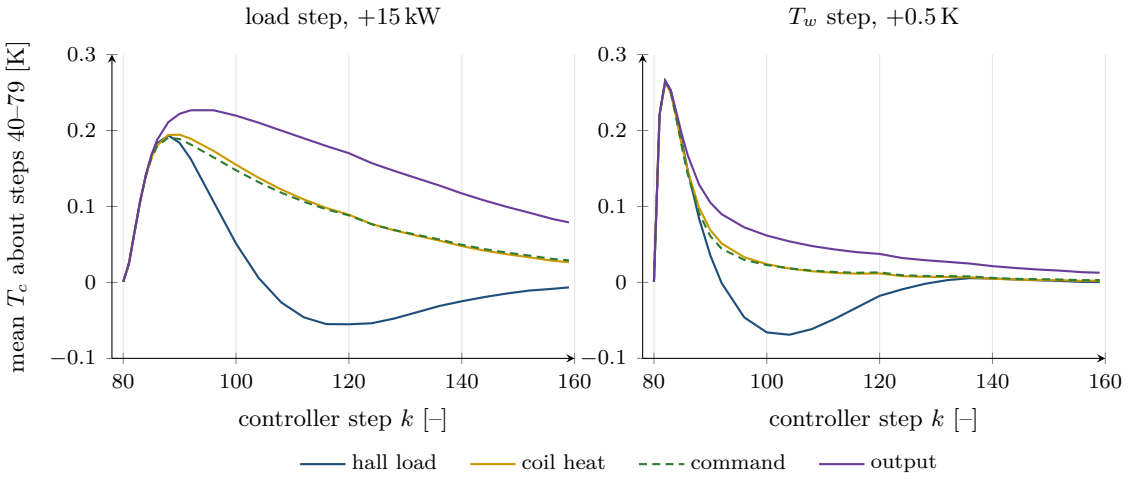


Figure 13.3: Mean T_c after two events, by entry, each selected to the given PI's predicted travel. Means over 1000 records, seeds 13000 to 13999.

On the load step the hall-load controller returns fastest and undershoots, while the output entry's tail is the longest, and the coil and command lines lie together. On the T_w step the peaks lie between +0.2642 and +0.2656 K and the tails change places. Mean square over steps 80 to 119, load step: hall load 0.01216, command 0.02058, coil 0.02179, output 0.03887 K^2. On the T_w step the order changes: command 0.00937, coil 0.00956, hall load 0.01133, output 0.01246 K^2. **Choosing the model is choosing the integrator.** For this hall's load step, unmetered, the hall-load entry, which ranks third on the T_w step; with the meter read, a coil entry for what the meter misses, Section 13.3.5. The hall-load model has a limit of its own: at $\rho = 25$ the controller's own pole crosses the circle between walks of 0.65 and 0.70 kW, at 1.00247 at 0.70 kW and 1.02396 at 1 kW, where K_i's residue is −0.007550, and at a walk of 0.30 kW it stands at 1.01069 for $\rho = 316$. That is the controller's own stability, and the selection below keeps it as a rule.

13.3.4 How many entries one reading carries

Control engineer's question. How many entries can one reading carry?

General form. Moments (2.10), read at the steady rows: the entries' constants can be separated only if the rows at $z = 1$ have full column rank.

Specialized form. Equation (2.10) with $f(\mathbf{z})$ the pair of constants, under the steady rows

$$\begin{bmatrix} \mathbf{A} - \mathbf{I} & \mathbf{E} \\ \mathbf{C} & \mathbf{F} \end{bmatrix} \tag{13.18}$$

with one column of **E**, **F** per entry and one row of **C**, **F** per reading. The modeling choice to note is that G and the coil stay pinned in these rows. Without that, a hall reading would not separate the load from the offset, which is Chapter 9's false pin.

Building blocks. The columns of equation (13.2), and the recursion of Section 13.2.1 run from a wide prior.

The plant's numbers. Table 13.4 lists the rank for each set of rows and entries, with the smallest singular value after scaling the columns to unit norm.

Table 13.4: Rank of the steady rows (13.18). Full rank with the hall reading rests on G and the coil pinned.

Readings	Entries	Rank [–]	Smallest singular value [–]
y_c	load	4 of 4	1.27×10^{-1}
y_c	offset	4 of 4	2.82×10^{-2}
y_c	load, offset	4 of 5	1.33×10^{-1}
y_c	load, coil	4 of 5	1.46×10^{-1}
y_c, y_T	load, offset	5 of 5	7.59×10^{-2}
y_c, y_T	load, coil	5 of 5	1.26×10^{-1}
y_c, y_T	load, offset, coil	5 of 6	1.33×10^{-1}
y_c, y_Q	metered load, coil	5 of 5	9.34×10^{-2}

One reading holds one constant. On y_c with the command held, the load's steady column is $+0.13118\,\mathrm{K/kW}$ and the offset's is 1, one number each, so any offset equals a load of $7.623\,\mathrm{kW}$ per kelvin. The recursion on both, with a prior on d of $5\,\mathrm{kW}$, on b of $0.20\,\mathrm{K}$, and a wide prior on the state, gives a spread of b of 0.2000, then $0.1732\,\mathrm{K}$ at $n = 40$, 160, 1000 and 10000 readings, and a spread of d of 5.0000, 1.3874, 1.3230, 1.3203 and $1.3200\,\mathrm{kW}$, with the correlation going to -1.0000. The floor is the prior's and the pinned coil's. The first steps' transient does differ, with a column cosine of 0.8249 over 10 steps and 0.9352 over 120, and only a known step time could use that.

One reading more. With several rows a steady I_k requires only $\mathbf{w}^{\mathsf{T}}\mathbf{B}_c\mathbf{y} = 0$, so the test has two parts: the pole, and $\mathbf{w}^{\mathsf{T}}\mathbf{B}_c$ with its weight on y_c alone. For the pole, equation (13.11) must hold on *every* row of $\mathbf{C}_a$. Add the hall reading y_T with one load entry. The direction $\mathbf{v}$ has a T entry of $+0.163639\,\mathrm{K}$ per kW, so the hall row gives $\mathbf{C}_a\mathbf{v} = (0, 0.163639)^{\mathsf{T}} \neq \mathbf{0}$. The derivation of the pole fails at its first line. The computed nearest pole lies at a gap of 0.0402, and the reading's error at the truth, noiseless, is $-0.01421\,\mathrm{K}$ after the load step and $+0.05774\,\mathrm{K}$ after T_w steps $+0.5\,\mathrm{K}$. A second entry on the hall row restores it: with an offset f on y_T, the direction $(\mathbf{z}_s^{\mathsf{T}}, 1, -0.163639)^{\mathsf{T}}$ satisfies equation (13.11) on both rows and the pole returns, gap 3.33×10^{-16}, input weights $(1, 0)$, errors -0.00008 and $+0.00025\,\mathrm{K}$ at step 399, a transient, and 0 at step 2999. So does a coil entry beside the load entry: weights $(1, 0)$, error 0. An extra entry on the discharge row with b out of the target does not, gap 0.0672. With the valve position y_u as the extra reading the gap is 9.51×10^{-5} and the error is below the digits printed. **The pole is not enough, and the count is not the mechanism.** The coil entry alone on y_c, y_T has $\mathbf{C}_a\mathbf{v} = 0$ on both rows, since its target leaves T where it is. The pole stays, gap 8.9×10^{-16}, and yet the input weights are $(1, 0.0036)$: the running sum balances the reading against the hall row's constant error, and the reading's error is $-0.00942\,\mathrm{K}$. The command entry alone on y_c, y_T, y_u also keeps its pole, with weights $(-0.0369, -0.0005, 1)$, nearly all on the valve's row, and an error of $+1.35146\,\mathrm{K}$. One entry per reading is the generic way to meet the test, and

the test is what must be checked. The warning is fair for this entry on these rows and is no argument against sensors. The sizes are this example's.

13.3.5 The meter against the entry

Control engineer's question. Does the meter remove the need?

General form. The expected consequence (2.23) over routes, on fresh records under two events.

Specialized form. Equation (2.23) with the actions the routes of Table 13.5 and L read as travel, peak, mean square and offset. The meter's reading y_Q enters either as feedforward, a gain on a filtered reading, or as a row of the filter. The steady gain from the plant is $K_{\text{ff}} = 1/(W(T_c - T_w)) = 0.01874$ per kW at the fit, 0.01867 at the truth. The first-order filter is set at the hall's own lag with the coil held, $C_h/G = 19.6$ steps, gain 0.0497 per step, because the heat reaches the coil only through the hall. This buys a valve that moves with the heat, and it fails if the lag is wrong, as the unfiltered row shows.

Building blocks. The steady columns; the controller (13.8) with rows y_c, y_Q, whose meter row has its own entry, or none; the given PI.

The plant's numbers. Table 13.5 gives 1000 records, seeds 13000 to 13999. Each row is a route. The columns are the travel on steps 40 to 79, then the load step's peak, mean square over steps 80 to 119 and reading offset over steps 140 to 159, then the T_w step's peak and offset. Standard errors lie at or below 0.0038/h on travel, 0.0015 K on a peak, 0.00019 K^2 on a mean square and 0.0008 K on an offset. Figure 13.4 draws five of the routes, with the same axes as Figure 13.3.

Table 13.5: The meter against the entry, 1000 records. Bold: the offset a route with no integrator leaves.

Route	Travel [/h]	Load peak [K]	Mean square [K^2]	Reading offset [K]	T_w peak [K]	T_w offset [K]
Given PI	0.4078	+0.1867	0.01397	+0.0074	+0.2635	−0.0010
PI + meter, first-order at the hall's lag	0.4113	+0.0413	0.00167	+0.0004	+0.2638	−0.0004
PI + meter, unfiltered	0.9469	−0.4954	0.03853	+0.0088	+0.2635	−0.0007
P + meter, hall's lag, no integrator	0.4373	+0.0431	0.00068	−0.0120	+0.3017	**+0.2906**
Meter alone	0.0518	+0.0520	0.00089	−0.0166	+0.4748	**+0.4538**
Filter, meter + coil entry, $\rho = 25$	0.3986	−0.1517	0.01193	−0.0116	+0.2654	+0.0012
Filter, meter + coil entry, $\rho = 1$	0.8613	−0.0082	0.00129	+0.0008	+0.2481	−0.0001
Filter, meter, no entry	0.1721	−0.1592	0.01075	−0.0241	+0.3452	**+0.3346**

The meter removes the load step from the integrator's work, not the integrator. At the hall's lag the mean square falls to 0.119 of the given PI's, at 1.008 of its travel, and on the T_w step the right panel shows it adds nothing. Unfiltered, the valve opens ahead of the heat and the peak changes sign, −0.4954 K. Without the integrator the load step is even tidier, but the T_w step stays on the reading: +0.2906 K with the P row, +0.4538 K with the meter alone, +0.3346 K with the filter and no entry, whose nearest pole is at a gap of 4.86×10^{-2}. A meter reading high by 1 kW opens the valve by 0.01874, so its bias on the reading is negative. With no feedback it is −0.1312 K by the fit's columns and −0.13384 K with the meter alone at the truth. The P row brings it to −0.08249 K, and a loop whose integrator takes y_c alone brings it to 0. The ×0.119 assumes a meter that is right. On the same records, the mean square is

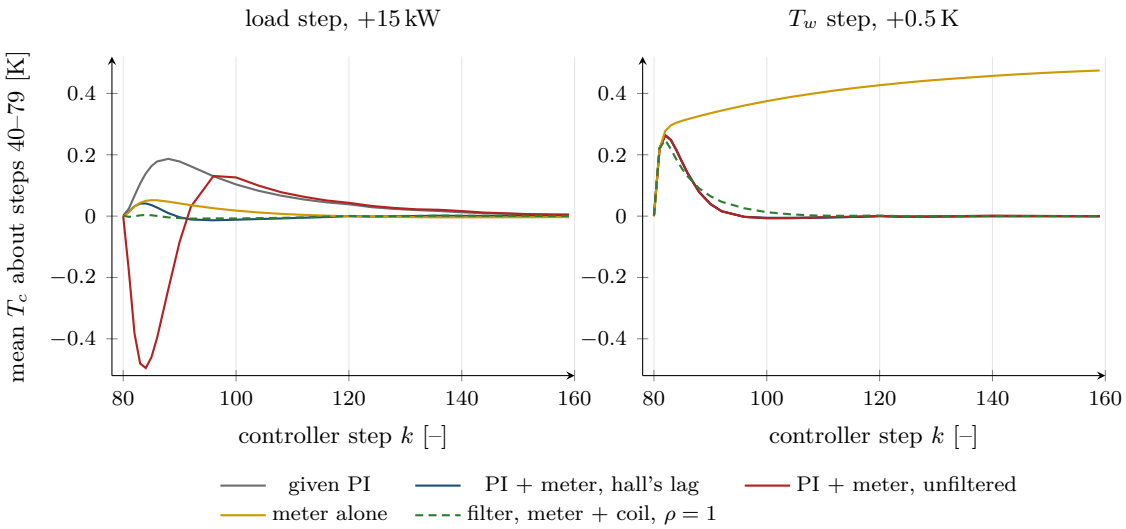


Figure 13.4: The meter on two events. Means over 1000 records, seeds 13000 to 13999. On the right the given PI and its two meter routes lie on one line.

×0.127 and ×0.130 of the given PI's with the meter's gain 10 % low and high, and ×0.285 and ×0.252 with the filter's lag halved and doubled. A false 15 kW meter step, with the load held, peaks at −0.1862 K (se 0.0012), with a mean square of 0.01457 K^2 (se 0.00016) against 0.00169 for the given PI on the same records with no event.

13.3.6 What the integrator cannot remove

Control engineer's question. What can it not remove?

General form. A threshold on an entry at a scenario node, equation (2.12), whose moments are (2.10).

Specialized form. Equation (2.10) with $f(\mathbf{z}) = T_c$ at the closed-loop node where r is pinned. The reading's row with its offset is $y_{c,k} = T_{c,k} + b + \varepsilon_k$. Equation (13.14), with the input weight on this row alone, requires the mean of $y_c - r$ to be zero at steady state, so

$$\mathbb{E}[T_c] = r - b, \qquad \mathrm{Var}(T_c) = \Sigma_{bb}, \tag{13.19}$$

by algebra, for any controller with a pole at one whose input weight $\mathbf{w}^\mathsf{T}\mathbf{B}_c$ lies on this reading alone, whatever its model. **The offset passes through with the opposite sign.** No integrator can remove it, because the integrator's input is the sum $T_c + b$: the two have one coefficient each in one row, and the sum is what the running sum drives to r. Removing it takes a row that reads b apart from that sum. The target must then move by $\hat{b}$, which is what the routes below do.

Same columns, different law. The output entry and the offset entry have the same columns, $\mathbf{E} = \mathbf{0}$ and $F = 1$, hence the same filter. The target is what differs. The output entry's law asks the reading to come to r, with $\mathbf{v} = (-1, -1, 0.142893, 1)^\mathsf{T}$ and $\mathbf{C}_a\mathbf{v} = 0$. The offset entry's law asks the air to come to r, with $\mathbf{z}_s = \mathbf{0}$, $u_s = 0$, $K_d = 0$, so $\mathbf{v} = (0, 0, 0, 1)^\mathsf{T}$ and $\mathbf{C}_a\mathbf{v} = 1$. Equation (13.11) fails and the offset's law has no pole at one. The integrator is the entry and the target together.

Building blocks. The rows (13.2) with $F = 1$; the target; Chapter 6's Σ_{bb} and the rows of Chapter 10 for the routes that separate b.

The plant's numbers. With no noise, at 30 kW, an offset appearing at step 40 leaves the reading's error below 3×10^{-7} K and $T_c - r$ at +0.20000, −0.30000 and −0.50000 K for $b = -0.20$, +0.30 and +0.50 K, on the given PI, the load entry and the output entry alike. The same-columns pair has a gap in $\boldsymbol{\Upsilon}$ of 0.0, K_d of −0.36838 against 0, and nearest poles at 0.0 and 0.0443. The offset's law leaves +1.8166 K of reading after the load step, $T_c = 19.51660$°C.

Figure 13.5 shows both sides. On the left, the horizontal axis is the controller step and the vertical axis a deviation from r, with the onset of $b = +0.30$ K at step 40, no noise: the solid lines are the reading and the dashed lines the true T_c, each for the given PI and the load entry, and the dotted line is the truth's $r - b$. On the right, each route's estimate of b is shown as its mean with a bar of one standard deviation across 300 records, with the truth dotted. Table 13.6 gives the routes that separate the offset, on 300 fresh records, seeds 13000 to 13299, measured on steps 220 to 239, with the target moved by $\hat{b}$. A row is a route on a loop; the columns are the air's error with its se and RMS, the reading's error, the hall air, and $\hat{b}$. The reference is one reading of spread 0.05 K at step 79, noise seed 131, on Chapter 6's fit. The second load is Chapter 10's rows on steps 40 to 79 and 140 to 159.

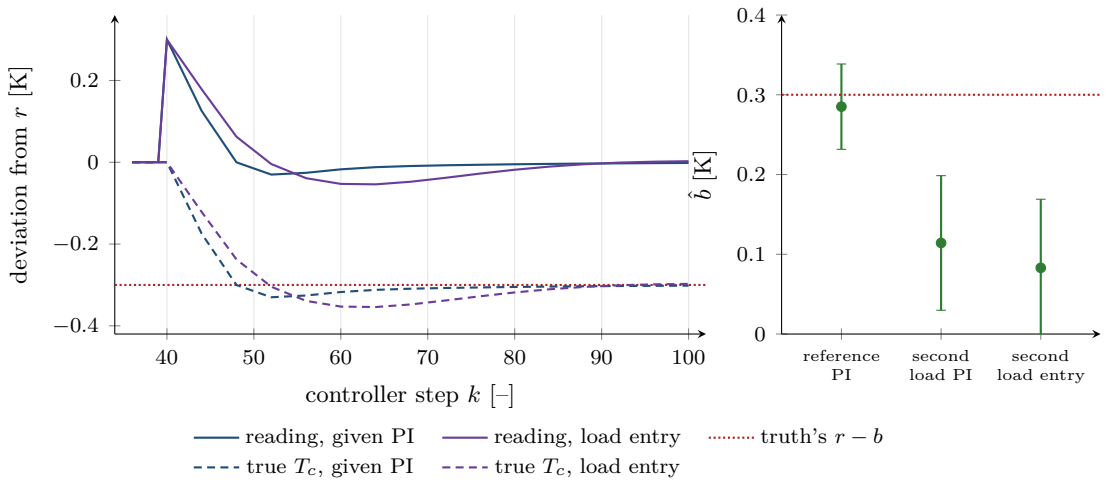


Figure 13.5: The offset through its onset and across records. Left: $b = +0.30$ K from step 40, no noise. The reading returns to r and the air goes to $r - b$. Right: $\hat{b}$ by route, mean with a bar of one sd across 300 records, seeds 13000 to 13299; dotted, the truth.

Table 13.6: The routes that separate the offset, 300 records. The truth is $b = +0.30$ K, $r - b = 17.70$°C.

Route	Loop	$T_c - r$ (se) [K]	RMS [K]	Reading $-r$ [K]	Hall air [°C]	$\hat{b}$ mean, sd across, stated sd [K]
None	PI	−0.3007 (0.0013)	0.3015	+0.0007	25.454	–
None	load entry	−0.3008 (0.0013)	0.3017	+0.0006	25.452	–
One reference reading 0.05 K	PI	−0.0156 (0.0034)	0.0604	+0.2858	25.739	0.2851, 0.0535, 0.0484
Second load	PI	−0.1860 (0.0051)	0.2059	+0.1153	25.564	0.1142, 0.0844, 0.1645
Second load	load entry	−0.2184 (0.0051)	0.2357	+0.0830	25.534	0.0830, 0.0860, 0.1641
Oracle	PI	−0.0007 (0.0013)	0.0231	+0.3006	25.754	0.30

The left panel is the chapter's second finding. Both readings return to r and stay; both airs

settle at the dotted line, 0.30 K below where the reading says. The table shows what removing it means. The second load's shrink toward the prior leaves 0.619 and 0.726 of the offset, with the truth within two stated standard deviations on 0.967 and 0.897 of the records, so its band covers the truth while its mean stays short of it. The reference leaves 0.052. Removing the offset moves the reading off r by $\hat{b}$ and raises the hall air by what is removed, from 25.454 toward 25.754°C. Whether to do that is a matter for the set-point layer, Chapter 16. A $\hat{d}$ that drifts as a fault is the concern of Chapter 26.

13.3.7 The derived controller against the PI

Control engineer's question. Is the derived controller better than the PI?

General form. The expected consequence (2.23), with the actions the two families and the consequence read on fresh records.

Specialized form. Equation (2.23) with a a member of either family. The comparison is fair by construction. Selection is at the fit's linear plant with no use of the truth, and both families have two knobs, (K_p, K_i) against (ρ, walk). Within a budget of predicted travel, ×0.25, ×0.5, ×1 and ×2 of the given PI's 0.4137/h, each takes the least predicted mean square subject to Pr(stable at 20 kW) ≥ 0.95 on Chapter 8's priors (sd log 0.15), a phase margin of at least 30° at the fit, and, as a stated rule, the controller's own poles on or inside the circle. The stability row never binds; the pole rule binds at ×0.5 only, 0.03568 against 0.03540 K^2. A third family holds the PI to the filter's margins at the same budget.

Building blocks. The controller (13.8); the plant (13.1) for selection; the records for the numbers, paired on seeds 13000 to 13999.

The plant's numbers. Figure 13.6 draws the mean square over the load step against travel on the records, on logarithmic axes, for the free PI, the PI held to the filter's margins and the filter, with each point's gain and phase margin at 30 kW as its label and the given PI as a single point. Table 13.7 gives the budget ×1 in full. Each row is a route; the columns are travel with its se, the spread of T_c, the load peak, the mean square with its se, the T_w step's mean square, the overshoot of a noiseless +0.5 K set-point step, the margins and the factor on the plant's gain to the boundary, at its 5 % point at 20 kW.

Table 13.7: Budget ×1 on 1000 records, paired. Selected filter: $\rho = 31.6$, walk 0.487 kW. PI free: $K_p = 0.0619$, $K_i = 0.002472$. PI at the filter's margins: $K_p = 0.0807$, $K_i = 0.002033$.

Route	Travel (se) [/h]	sd T_c [K]	Peak [K]	Mean square (se) [K^2]	T_w mean square [K^2]	Set-point overshoot [K]	GM [–], PM	Delay margin [s]	Factor, 5 % [–]
Given PI	0.4078 (0.0017)	0.0322	+0.1867	0.01397 (0.00016)	0.00889	+0.0744	14.09, 59.8°	144	8.71
PI ×1 free	0.3740 (0.0024)	0.0395	+0.1890	0.01130 (0.00014)	0.00962	+0.1753	5.03, 45.6°	95	3.07
PI ×1 at the filter's margins	0.3532 (0.0017)	0.0330	+0.1937	0.01405 (0.00016)	0.00936	+0.0995	11.64, 55.7°	134	6.41
Filter ×1	0.3593 (0.0016)	0.0334	+0.1935	0.01216 (0.00015)	0.01133	+0.1468	11.53, 50.4°	145	7.77

No route dominates, and that is the result, at each stated basis. Against the free PI at ×1, the filter has a mean square ×1.076 on the load step, with ×0.846 of the scatter and a gain margin of 11.53 against 5.03. Against the PI held to its margins, it has ×0.866 on the load step and ×1.210 on T_w, an event its model does not name, with more set-point overshoot, +0.1468 against +0.0995 K. These points come from coarse grids that leave part of each budget unspent, so their ratios confound the route with the shortfall. Filter over free PI,

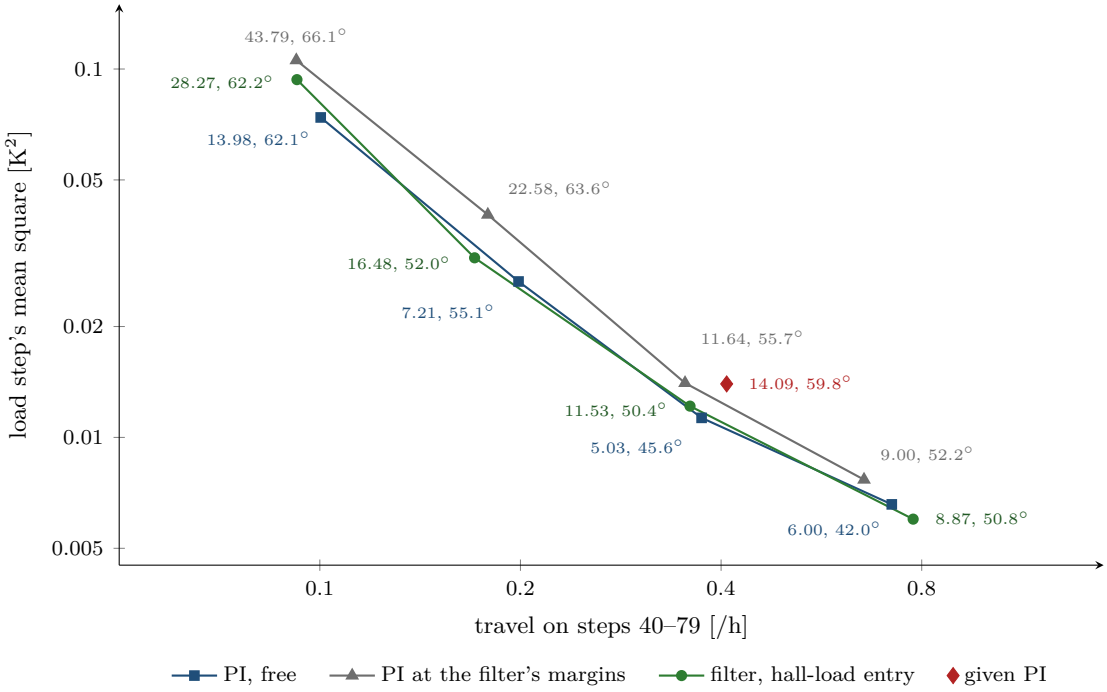


Figure 13.6: Mean square against travel on 1000 records, seeds 13000 to 13999, at the four budgets. Labels: gain margin, phase margin at the fit, 30 kW.

at the four budgets, they run: mean square ×1.266, ×1.159, ×1.076, ×0.912, at travel ×0.921, ×0.859, ×0.961, ×1.077; the ×0.859 is not like with like, nor is the ×1.185 of travel against the PI at the filter's margins at ×2. Fine grids spend the budgets (PI 161 × 81, filter 61 × 97, same ranges and screen; seeds 13000 to 13999, count 1000, paired). Filter over free PI, travel is then ×0.985, ×0.962, ×1.003, ×1.005 and mean square ×1.141 (difference +0.01036, se 0.00007), ×0.978 (−0.00057, se 0.00007), ×1.024 (+0.00025, se 0.00008), ×0.928 (−0.00045, se 0.00005), so the sign at ×0.5 reverses and "no winner" stands more strongly. Scatter is ×0.907, ×0.911, ×0.878, ×0.853 and the T_w step's mean square ×1.073, ×1.170, ×1.173, ×1.071. Over the PI at the filter's margins, mean square is ×0.988, ×0.837, ×0.857, ×0.808 at travel ×0.988, ×0.961, ×1.021, ×1.066, and T_w ×1.000, ×1.170, ×1.220, ×1.161. **Retuning the PI alone gets more than the filter, by giving up margin, and held to the margins it gets nothing**: at ×1 against the given gains the mean square changes −0.00267 (se 0.00004) for the free PI, with its gain margin down from 14.09 to 5.03, −0.00181 (se 0.00008) for the filter, and +0.00008 (se 0.00001) for the PI held to the filter's margins, at scatter +0.0073 and +0.0012 K for the first two. Designing the filter at the truth's plant changes little, 0.01245 against 0.01216 K^2. The derived controller's claim is what it says, its assumed disturbance. The comparison is on one load step and one T_w step.

13.3.8 The band of the entry and the law

Control engineer's question. Does the band of d enter the law?

General form. The expected consequence (2.23), on the row of Table 7.6.

Specialized form. Equation (2.23) with L the cost in the target's coordinates, $J = q\,\mathrm{Var}(y) + \rho\,\mathrm{Var}(c - u_s d)$, by a Lyapunov equation with d a walk, and the action the law scaled by a factor. Chapter 12 holds the result in general for this row; it is checked here, stated by the walk.

Coupling 3 is gone because the command has no line in the Jacobian (13.3). Coupling 4 is gone if $\mathbf{K}$ and K_d do not move with $\mathbf{\Sigma}$.

Building blocks. $\mathbf{\Sigma}$ of equation (13.7), moved by the walk; $\mathbf{K}$, K_d of equation (13.6).

The plant's numbers. Across walks of 0.03, 0.30 and 3 kW, the gap in $\mathbf{K}$ and K_d is 0.0, against a tolerance of 10^{-9}. With the law scaled by 0.8, 0.9, 1, 1.1 and 1.2, the expected cost has its least value at a factor of 1 at each walk: 0.0015230, 0.0175993 and 0.6789750. $\mathbf{\Sigma}$ of d moves $\mathbf{\Upsilon}$, so the controller's numbers do move: K_i's residue is 0.000174 at 0.03 kW and 0.002370 at 0.30 kW. The band reaches those numbers through the filter, never through the law.

13.4 The checks (step 5)

The answers have set the routes side by side. This section tests the machinery against earlier chapters and its own limits, and sets the answers against the record's stated truth.

13.4.1 The routes against each other

Table 13.8 collects the checks, one to a row, with what is compared, the values and the result. Algebra is what makes the first two hold; the others are agreements to a stated size.

Table 13.8: The checks of the chapter.

Check	Values	Result
$\mathbf{A}_c\mathbf{v} = \mathbf{v}$, four models and 200 random pairs, seed 1313	largest 2.8×10^{-17}; tolerance 10^{-9}	holds, by algebra
Scalar law through the entry against the differenced state	gaps 1.4×10^{-17}, 5.6×10^{-17}	holds, by algebra
Stacked first move, horizon 800, against $-K_d$	effort gap 1.5×10^{-16}, move 1.3×10^{-17}; tolerance 10^{-6} of K_d	agrees
This plant map against the series' closed-loop map	largest gap 0.0; tolerance 10^{-9}	agrees
Batch records against the record, seed 51, steps 0–159	T_c 3.6×10^{-15} K, u 4.4×10^{-16}	agrees
Chapter 10's raw route, seeds 10000–10999	0.4112, 0.0323, 0.1863, 0.01400, each equal to the digits printed there	agrees
Chapter 8's K_{95} at 20 kW, 20 000 draws (effective 12 420)	0.8709 against 0.8693	reported, at fewer draws than the 200 000 there
Seed 51, two loads	$\hat{b}$ 0.1075, band 0.1646 by Gauss–Newton, 0.1658 by Newton in Chapter 10	$\hat{b}$ gap 0.0000
Steady filter against the recursion, 2000 steps	1.7×10^{-14}	agrees
One period from the truth's steady state	moves it by less than 10^{-9}	agrees

13.4.2 The limits of the entry

As the walk goes to zero the pole stays and its gain vanishes: $K_i = 5.812 \times 10^{-4}$, 5.998×10^{-5}, 6.201×10^{-6} and 6.230×10^{-7} at walks of 10^{-1} to 10^{-4} kW, with the nearest pole at 10^{-15} or closer. Equation (13.15) says why: $\mathbf{w}^{\mathsf{T}}\mathbf{B}_c \to 0$ with the target unchanged. As $\sigma \to 0$ on the state, the limit of a physics row, K_i converges at the walk of 0.30 kW: 0.002370, 0.002371, 0.002371 at $\times 1$, $\times 0.1$ and $\times 0.01$.

13.4.3 The answers against the stated truth

The truth has $g = -0.03390$, $g_w = -0.10$ and $b = +0.30\,\mathrm{K}$. Noiseless, $\hat{d}$ moves 14.939 kW for the step of 15 kW, because it is the valve's true move of 0.2800 read through the fit's 0.018744 per kW; the truth's value is 0.018668. On 1000 fresh records, $\hat{d}$'s move less the true load's move is +0.177 kW (se 0.022) with a spread of 0.694 kW. On seed 51 it moved 15.227 kW against 14.108 kW, one draw. On that record the mean T_c over steps 140 to 159 is 17.6724°C with the filter and 17.6807°C with the PI, against the truth's $r - b = 17.70$°C.

13.5 The report (step 6)

The integrator is a disturbance entry together with a target that holds the reading against it, by algebra. A cost on the move does not make it: the cost about the target leaves the same zero offset. And the reading it zeroes is not the air, which here is 0.30 K below it. These sizes are this example's.

13.5.1 The answers and their bands

Table 13.9 gives the answers in the order of Table 13.1. Each row is a question, with the answer, its band or se, and the record's stated truth where one applies.

Table 13.9: The report on the integrator. Records are seeds 13000 to 13999 unless stated.

Control engineer's question	Answer	Band	Stated truth
Where is the integrator?	A pole at $z = 1$ from the entry and the target, for any gains; $K_i = 0.002370$ at $\rho = 25$	residual 2.8×10^{-17}; tolerance 10^{-9}	–
What does the cost on the move do?	Sets gains; no offset with the entry, as about the target. No entry: +1.94475 K. Effort: +0.22268 K	linear plant; at the truth +0.20728 K	–
Which disturbance does the loop assume?	The chosen entry; at one travel budget, load step: load 0.01216, command 0.02058, coil 0.02179, output 0.03887 K^2; T_w step: 0.01133, 0.00937, 0.00956, 0.01246 K^2	se at or below 0.00025 K^2 on the load step	load step 15 kW
How many entries can one reading carry?	One. A further reading must pass two tests, the pole and $\mathbf{w}^{\mathsf{T}}\mathbf{B}_c$ on y_c alone: load on y_T, gap 0.0402; coil on y_T, pole kept, error −0.00942 K; load with hall offset, weights $(1, 0)$, error 0	rank 4 of 5; sd b floor 0.1732 K	–
Does the meter remove the need?	No. Mean square ×0.119 on the load step, nothing on T_w; without an integrator +0.2906 K on T_w	se 0.0003 K on that offset	–
What can it not remove?	The sensor offset: $T_c - r = -0.3007$ K under a flat reading	se 0.0013 K, 300 records; prior band sd b 0.1899 K	$b = +0.30$ K; $r - b = 17.70$°C
Is the derived controller better than the PI?	No winner. On fine grids, at near-equal travel, mean square ×1.141, ×0.978, ×1.024, ×0.928 of the free PI over the four budgets, ×0.988, ×0.837, ×0.857, ×0.808 of the PI at the filter's margins; at ×1 scatter ×0.878, ×1.026; T_w ×1.173, ×1.220	paired se at or below 0.00009 K^2 on the mean-square differences	–
Does the band of d enter the law?	No: gap 0.0 across walks, least cost at the factor 1. It moves the controller through $\boldsymbol{\Upsilon}$	tolerance 10^{-9}	–

13.5.2 The entry in the commissioning log

The entry states the controller's assumed disturbance, with its column and its target, and where its integrator's input weight falls, on the reading held or on others. It sets a band on T_c beside the flat trend: the reading's error near zero, with T_c at $r - b$ and a spread for b of 0.1899 K from Chapter 6. It records that the meter goes in with its lag and that the integrator stays. For the vendor's offer, it records that the sheet's statement on the move cost is false as stated, and asks for the entry and the target in its place. The verdict on the two controllers is by the table, at each stated basis, with no choice recorded here.

13.5.3 Common mistakes

Table 13.10 lists five mistakes this plant makes easy. Each row names the mistake, its form on this hall and the correction.

Table 13.10: Five mistakes about integral action.

Mistake	On this hall	Correction
Reading a flat trend as air held	mean error +0.0172 K with T_c at 17.6807°C on seed 51	the running sum zeroes the reading; state T_c with b's band
Crediting a cost with the integrator	the move cost with no entry leaves +1.94475 K; about the target, with the entry, 0	ask for the entry and the target
Adding a sensor to the filter and checking only the count or the pole	y_T on the load entry: gap 0.0402, error −0.01421 K; on the coil entry: pole kept, error −0.00942 K	test the pole and $\mathbf{w}^{\mathsf{T}}\mathbf{B}_c$ on y_c alone; one entry per reading is the generic way
Entering the offset as if it were the output entry, or the reverse	same columns; the offset's law has no pole and leaves +1.8166 K	the law, not the filter, says which it is
Taking the meter for a replacement	T_w step +0.4538 K with the meter alone	the meter does the load; the integrator stays

13.6 The next measurement (step 7)

Two questions call for two instruments, since one concerns what the integrator removes and the other what it cannot. Nothing is refit: the law stays, one entry per added reading, on the linear-normal preposterior (4.73).

13.6.1 The candidates

Table 13.11 lists the candidates, the permanent instruments among them. Each row is a candidate. The columns are the spread of the first entry in the steady filter, the linear load step's peak and mean square, the predicted travel, and the spread of the true T_c from Chapter 6's $\mathbf{\Sigma}$ and the stated rows.

13.6.2 The recommendation

What the integrator removes wants the meter, and that is already there: it is the row to install, with its lag, before any other. What it cannot remove wants the discharge reference: one reading takes the spread on the true T_c from 0.1899 to 0.0484 K, and no row in the top five moves it at all. A better meter worsens the peak, −0.2044 against −0.1296 K, because the law at $\rho = 25$ acts ahead of the heat, which is a matter of tuning, Chapter 14.

Table 13.11: The candidates at $\rho = 25$, walk 0.30 kW, each added reading with its own entry. Linear at the fit.

Candidate	sd first entry [kW]	Peak [K]	Mean square [K^2]	Travel [/h]	sd true T_c [K]
y_c alone	1.141	+0.2641	0.02610	0.2575	0.1899
Meter 1 kW (permanent, present)	0.508	−0.1296	0.00784	0.1342	0.1899
Meter 0.2 kW	0.173	−0.2044	0.01481	0.0827	0.1899
Hall 0.5 K, own offset	1.113	+0.2494	0.02195	0.2510	0.1899
Reference hall 0.1 K (permanent), own offset	0.937	+0.1577	0.00645	0.2092	0.1899
One reference discharge reading 0.05 K	–	–	–	–	**0.0484**
Second load, 20, 40, 80 steps	–	–	–	–	0.1646, 0.1589, 0.1549

13.6.3 What the analysis showed

The sheet said a cost on the move gives integral action, and the trend looked as though the air was held. The analysis found the integrator elsewhere: in an estimated entry and a target that holds the reading, by algebra and with any gains, so that the move's cost only shares out what the differencing or the entry has already built. The target is the part easily missed, since the same columns with a different target have no pole. A filter with one more reading can lose the pole, or keep it with its input weight off the reading held, and either way leave an offset on that reading. The derived controller did not win its comparison; what it offers is a named disturbance, and on this load step a retuned PI gains more mean square than it does by giving up gain margin, and nothing when held to its margins. The engineer takes back three questions to put to any offered controller: which disturbance it assumes, where its integrator's input weight falls, and what it says about the air beneath a flat reading.

13.6.4 Limits of the model

The design. This is a linear design at one operating point, a constant d carried through a walk, and $\hat{d}$ read through the fit's coil. The plant's gain stays pinned at the fit.

The stop. No stop is reached, and the filter's input is the unclipped command. Chapter 15 takes that up.

The stated choices. q, ρ, the walks and the 30° and pole rules are stated choices. The budgets' comparison is on one load step and one T_w step.

What the priors do not hold. A sensor lag and a low load are in no prior. The factor to the boundary at the fit, 20 kW, is 3.61 for the PI ×1 and 8.56 for the filter ×1; with a 5 s lag 3.00 and 6.61; at 5 kW, 1.85 and 4.78.

The constructed record. Fresh records share the truth, so they test the method on its own model. The record is constructed from a stated truth. All percentages and ratios here are this example's. What sets ρ against an objective is the next chapter's question.

13.7 What to retain

Definition 13.1 (Integral action by derivation)**.** With an entry d in the rows (13.2) and the law (13.6) steered to the target (13.4), the direction $\mathbf{v} = (\mathbf{z}_s^{\mathsf{T}}, 1)^{\mathsf{T}}$ satisfies $\mathbf{A}_c\mathbf{v} = \mathbf{v}$, by algebra, for any $\mathbf{K}$ and any filter gain, provided the target exists and $\mathbf{C}_a\mathbf{v} = \mathbf{0}$ on every reading row. That is a pole of the controller at $z = 1$, whose gain is $K_i\,\Delta t = u_s\,\mathbf{w}^{\mathsf{T}}\mathbf{B}_c$; it is integral action only with $u_s \neq 0$, which is full rank of the steady rows, and it holds the reading only with $\mathbf{w}^{\mathsf{T}}\mathbf{B}_c$ on that row alone.

Three statements carry beyond the hall, as statements of form and not of size. The integrator is the entry and the target together, so neither the entry alone nor a cost makes it. An added reading must keep the pole, by equation (13.11), and keep the integrator's input on the reading held; an entry per reading is the generic way to do both, not the mechanism. And what enters the reading's row with the air's coefficient goes to the air with the opposite sign through any controller whose integrator takes that row alone, by equation (13.19).

Table 13.12 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 13.12: The terms of the chapter and where each is defined.

Term	Expression	Units	Meaning	Defined in
Disturbance entry	$\mathbf{E}$, F, eq. (13.2)	mixed	a constant with its columns	13.1.4
Steady target	$\mathbf{z}_s$, u_s, eq. (13.4)	per unit	rest with the reading held	13.1.4
The law	K_d, eq. (13.6)	–	the regulator in the target's coordinates	13.1.4
Controller	$\mathbf{A}_c$, eq. (13.8)	mixed	belief and law, reading to command	13.2.2
The direction	$\mathbf{v}$, eq. (13.10)	per unit	what the reading cannot see	13.3.1
The pole	eq. (13.13)	–	the integrator, for any gains	13.3.1
Integral gain	eq. (13.15)	/(K s)	target's valve times filter's rate	13.3.1
Velocity form	eq. (13.16)	–	one entry per state, gain identity; on the scalar, $\hat{d}$ with gain one	13.3.1
Effort's share	eq. (13.17)	–	the offset an effort cost leaves	13.3.2
Steady rows	eq. (13.18)	–	the rank of what a reading holds	13.3.4
Offset's law	eq. (13.19)	K	$T_c \to r - b$	13.3.6
Controller's own stability	poles on or inside the circle	–	no growth by itself with the loop opened; an integrator still ramps on a constant error	13.3.7

13.8 Exercises

Exercise 13.1. *Formulation.* A tank's level x is read each step, its inflow c is commanded and its outflow d is not metered. Write the entry and the target, and say what the output-entry controller cancels. (With $\beta = \Delta t/A$: $x_{k+1} = x_k + \beta c_k - \beta d_k$, so $\mathbf{A} = 1$, $\mathbf{B} = \beta$, $\mathbf{E} = -\beta$, $F = 0$. The target (13.4) gives $u_s = 1$, $x_s = 0$: the inflow rises by the outflow and the level stays. The output entry has $\mathbf{E} = 0$, $F = 1$, with $x_s = -1$, $u_s = 0$, and there the question has no steady answer. Its steady rows have rank 1 of 2, so $\mathbf{C}_c\mathbf{v} = u_s = 0$: the controller's pole at 1 has no share of the command and cancels within the controller, the closed loop keeps a pole at 1, and no steady filter is stable, the recursion's spread on the level still at 0.836 after 200 000 steps.

With $\beta = 0.1$, $K = 1$, a walk of 0.01 and a sensor's spread of 0.1, an outflow of 1 leaves a level error of -1.94629 under the output entry and 0 under the outflow entry. The pattern of Table 13.3 does not carry over.)

Exercise 13.2. *Derivation.* Show $\mathbf{A}_c\mathbf{v} = \mathbf{v}$. (From the target (13.4): $\mathbf{C}_a\mathbf{v} = \mathbf{C}\mathbf{z}_s + F = 0$, so $(\mathbf{I} - \boldsymbol{\Upsilon}\mathbf{C}_a)\mathbf{v} = \mathbf{v}$. From $K_d = -(\mathbf{K}\mathbf{z}_s + u_s)$: $\mathbf{K}_a\mathbf{v} = -u_s$. Then $(\mathbf{A}_a - \mathbf{B}_a\mathbf{K}_a)\mathbf{v} = (\mathbf{A}\mathbf{z}_s + \mathbf{E} + \mathbf{B}u_s, 1)^\mathsf{T} = \mathbf{v}$.)

Exercise 13.3. *Derivation.* Find the scalar PI gains. (With $\hat{d}_k = x_k - ax_{k-1} + \theta u_{k-1}$ and $u_k = Kx_k + \hat{d}_k/\theta$, subtracting u_{k-1} gives equation (13.16), so $K_p = a/\theta$ and $K_i\,\Delta t = K + (1 - a)/\theta$. With $a = 0.36788$, $\theta = 4.48780$ and $K = 0.063026$: $K_p = 0.08197$ and $K_i\,\Delta t = 0.063026 + 0.14085 = 0.20388$.)

Exercise 13.4. *Derivation.* Say why two readings on one entry lose the pole. (Equation (13.11) must hold on each row of $\mathbf{C}_a$. One entry gives one free number in $\mathbf{v}$, fixed by the target at the controlled row, so the second row has $\mathbf{C}_a\mathbf{v} \neq 0$ in general; on the hall y_T gives $+0.163639$ per kW, and then $(\mathbf{I} - \boldsymbol{\Upsilon}\mathbf{C}_a)\mathbf{v} \neq \mathbf{v}$. A second entry with $F = 1$ on that row adds a free number that cancels it. The pole alone is not the end: the coil entry has $\mathbf{C}_a\mathbf{v} = 0$ on both rows, but input weights $(1, 0.0036)$ and an error of $-0.00942\,\mathrm{K}$.)

Exercise 13.5. *Application.* Find effort's offset for the 15 kW step. (By equation (13.17) with $g_u = -6.9982\,\mathrm{K}$, $q = 4$ and $\rho = 25$: $qg_u^2 = 195.901$, so the share is $25/220.901 = 0.11317$. The open-loop effect is $+1.9676\,\mathrm{K}$, and the offset is $0.11317 \times 1.9676 = +0.22268\,\mathrm{K}$.)

Chapter 14

Tuning as an Objective

14.1 Two tunings and no objective

The commissioning sheet holds two tunings for the discharge loop. The plant runs $K_p = 0.10$ per kelvin and $K_i = 0.002$ per kelvin second. The sheet of Chapter 8 writes $K_p = 0.41$ and $K_i = 0.0082$. One noisy record from that chapter sets them side by side at 20 kW. After a 10 kW load step the peak error is 0.141 K against 0.061 K. The spread of T_c is 0.042 K against 0.100 K, and the valve travels 0.41 against 2.19 per hour. One tuning is faster and one is quieter.

The sheet calls one of them "well tuned", and an engineer reading only the peak would take 0.061 K and agree. Chapter 8 declined to weigh the two columns, since a threshold on stability has no price on travel. A weighing needs an objective, and neither tuning arrived with one. The question of this chapter is

what either tuning minimizes, and what would have to be true for it to be the right one.

Jargon: LQ law, inverse optimality, return difference, spectral factor, static output feedback, fixed-structure gains, event rate, gain scheduling, expected value of perfect information

An *LQ law* (linear-quadratic) is the linear row on the whole state that minimizes a sum of quadratic stage costs on a linear plant; Chapter 12 derives it. *Inverse optimality* is the question asked backward: given a row, is there a cost it is the LQ law of. The *return difference* is $1 + L$, with L the loop's transfer function opened at the actuator. A *spectral factor* is a polynomial whose squared modulus on the unit circle equals a given nonnegative function. *Static output feedback* is a row that acts on what is read and not on the whole state. *Fixed-structure gains* are the best values of such a row's few numbers under a stated objective, here the PI's two. The *event rate* is the stated number of load events per hour. *Gain scheduling* is a row whose gains are chosen from a known operating condition. The *expected value of perfect information* is the most that knowing the state exactly could save.

Key idea

"Optimal" is not a property of a tuning. It needs a route, a disturbance and a form. On the hall's whole state, at this 30 s step and with no sensor lag, a derived test finds no quadratic cost for any stabilizing PI on the grid; at slower sampling some pass, so whether a PI has an LQ objective depends on the sampling. On the two-entry model the PI sees, some tunings have one, and the plant's reset is too fast to be among them. As the best of its own two gains, a PI can be read under a stated objective and a stated disturbance. Each reading answers a different question.

14.1.1 Three readings of an optimal tuning

Chapter 7 said that a controller is a partial answer to one problem, and the PI row of Chapter 5 was its first line: a law with no objective behind it. Chapter 12 solved the forward problem, from a cost to a row. Asked backward, from a row to a cost, the question splits, because "is this row optimal" leaves two things unsaid: on which state, and over which rows.

Figure 14.1 sets out the three routes this chapter names. Each column is a route. The first box says what the route treats as the state and the second what it asks; the third, in red, says what the route hides. The FULL route takes the whole augmented state of the hall and every linear row on it. The REDUCED route takes a two-entry model fitted to the step response, on which the PI is already a full-state row. The BEST-PI route keeps the whole plant and restricts the row to the PI's own two gains.

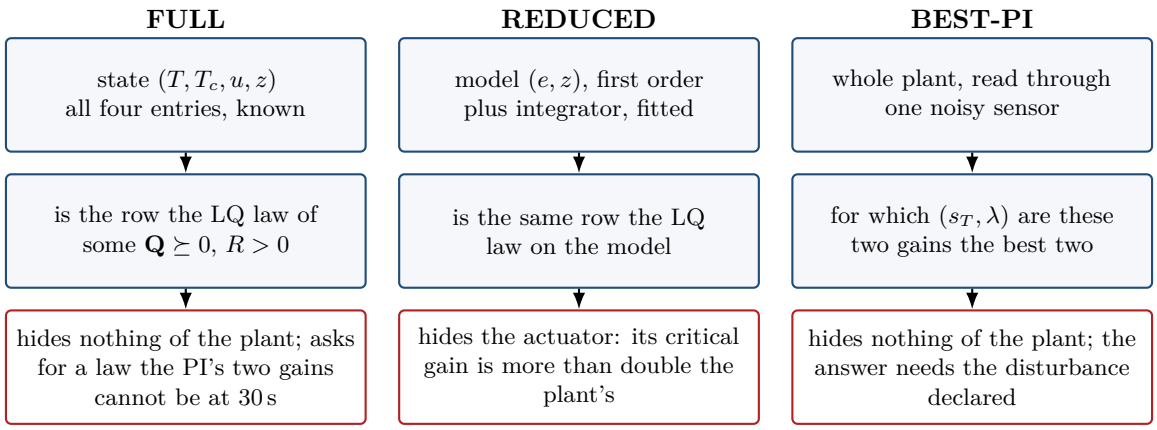


Figure 14.1: The three routes of the chapter. Blue-outlined boxes: what each route treats as the state and what it asks of a tuning. Red-outlined boxes: what the route hides. A verdict on one route says nothing on another.

14.1.2 The hall, the restriction and the objective (step 0)

The plant is the closed-loop hall of Chapter 8 at a pinned load: the map (8.6), 30 s apart, with the offset of Chapter 6 free and the priors of Chapter 8 on the capacities.

The restriction. This chapter works under the row "Known gain, quadratic, no limit" of Table 7.6, asked backward. The couplings kept are 1 and 2 of Section 7.1.3. Then, from Section 14.3.3, it works under that row's fixed-form variant: the row is held to the PI's two gains and the band is averaged over them, which brings coupling 4 back through a gain it multiplies. Coupling 3 stays given up, and the stop is Chapter 15's.

The scenario node. Every what-if but one is closed loop: the set point is pinned at $r = 18°\mathrm{C}$, the load at Q^*, and the valve's deviation stays free inside its stops, determined by the controller's row. The exception is the valve bump of step 7, which is open loop and pins the actuator.

The map on the differenced state. With an integrator entry $z_{k+1} = z_k + \Delta t\, e_k$ and $\Delta t = 30\,\mathrm{s}$, the map is one pair $(\mathbf{A}, \mathbf{B})$ on $\mathbf{x} = (T, T_c, u, z)^\mathsf{T}$ with the command's change as the input, and the PI is one row $\mathbf{F}$,

$$\mathbf{x}_{k+1} = \mathbf{A}\mathbf{x}_k + \mathbf{B}v_k, \qquad v_k = \mathbf{F}\mathbf{x}_k, \qquad \mathbf{F} = (0,\ K_p,\ 0,\ K_i). \tag{14.1}$$

This pair is the one-step pair of equation (8.10) with the row of z added, not the continuous $\mathbf{A}$ of equation (8.5). The same three symbols serve two readings. On deviations, v_k is the command's deviation and z_k the integral of the error. On differences, take $\mathbf{x}$ to hold each physical entry's change over one step, and set $z_k = \Delta t\, e_{k-1}$. The PI differenced is then

$$\Delta c_k = K_p(e_k - e_{k-1}) + K_i \Delta t\, e_{k-1}, \tag{14.2}$$

with the same $\mathbf{A}$, $\mathbf{B}$ and $\mathbf{F}$. A noiseless 10 kW step over 200 steps shows a largest difference in c of 4.8×10^{-17} between the two forms. This matters for the cost. A weight on z is a temperature spread, a weight on v is a spread of valve moves, and a load that wanders without bound still has stationary differences.

The modeling choices. Each of the following is stated, not read from the record.

- Disturbances: the record's load wander, autoregressive with coefficient 0.95 and spread 0.5 kW; sensor noise of 0.10 K; and 10 kW load events at a stated rate λ, at baseline one per two hours, 0.5 per hour, the record's count. These buy a stationary cost, and they fail if the load is otherwise; a row at another rate tests it.

- Objective: per step,

$$J = \frac{\mathbb{E}[e^2]}{s_T^2} + \frac{\mathbb{E}[(\Delta c)^2]}{s_m^2}, \qquad s_m = 0.2 \text{ per step.} \tag{14.3}$$

  The two spreads say what size of error and what size of valve move are worth the same. Only their ratio can move the gains.

- The baseline $s_T = 0.25$ K **was taken from the inverse map of Section 14.3.3, not chosen independently.** The agreement between the plant's tuning and the baseline's best gains is therefore not evidence for either. Chapter 7's 0.5 K, equation (7.5), is a row beside it in Table 14.7.

14.1.3 The control engineer's questions (step 1)

Table 14.1 lists the seven questions. The first column is the question in the engineer's words, the second the general form, the third its mathematical form on this hall, the fourth the category, the fifth the computation and the last the subsection. The first three ask a tuning for its objective, on three routes. The next three run forward, from an objective to a tuning. The last is the next measurement.

14.1.4 The evidence and the disturbances as rows (step 2)

The evidence is that of Section 8.1.4, unchanged, with the sparsity pattern (8.7): the readings act on Q, g, g_w and b, and the last three columns, the capacities, have a line from their own priors and from nothing else. The discharge reading stays out for the reason given there. The tuning adds no row on the past. What it adds is on the future, which is averaged. By the rule of Section 7.1.2, the disturbances of step 0 are not entries to fit but quantities to average over.

The disturbances' covariance. The BEST-PI route needs the stationary variances of (14.3). On the differenced closed loop, with augmented state $\boldsymbol{\xi}$ holding the plant's changes, the last error, the last noise and the wander's level, the one-step matrix $\mathbf{M}(K)$ gives a stationary covariance

$$\mathbf{S} = \mathbf{M}\mathbf{S}\mathbf{M}^{\mathsf{T}} + \sigma_n^2 \mathbf{g}_1 \mathbf{g}_1^{\mathsf{T}} + \sigma_w^2 \mathbf{g}_2 \mathbf{g}_2^{\mathsf{T}} + \frac{\lambda \Delta t}{3600} (10\,\text{kW})^2\, \mathbf{g}_3 \mathbf{g}_3^{\mathsf{T}}, \tag{14.4}$$

Table 14.1: The questions asked of a tuning. $\mathbf{y}$ is Chapter 8's: Chapter 6's three readings and three priors, with three priors on the dynamic entries added.

Control engineer's question	General form	Mathematical form	Category	Computation	Answered in
Is this PI an LQ law on the hall's state?	inverse of the expected consequence, (14.8)	$\gamma \geq 1$, $\gamma\lvert\chi_c\rvert^2 - \lvert\chi\rvert^2 \geq 0$	design	– (algebra at the fit), then posterior	14.3.1
And on the model the PI sees?	the same on a reduction, (14.8)	on (e, z), $\mathbf{Q}$ diagonal	design	– (algebra at the fit)	14.3.2
What is it the best PI for?	restricted expected consequence, inverted, (14.8)	$\nabla_K[qV_e + \rho V_m] = 0$ for (q, λ)	design	expected consequence	14.3.3
What does the ratio buy?	the forward map, (2.23)	$s_T/s_m \mapsto K_p$, decay, margin	design	expected consequence	14.3.4
Does the band move the gains?	expected consequence over the posterior, (2.23)	$\min_K \mathbb{E}[J(K, \mathbf{z}) \mid \mathbf{y}]$	design	expected consequence	14.3.5
Does a schedule do better than the band?	the same, by scenario node, (2.23)	$\min_{K(Q^*)}$ against $\min_K$	design	expected consequence	14.3.6
Which test most improves the tuned cost?	next measurement, (2.16)	$\mathbb{E}_{\mathbf{y}}[\min_K \mathbb{E}(J \mid \mathbf{y})]$	design of measurement	preposterior	14.6

with $\sigma_n = 0.10\,\mathrm{K}$ and $\sigma_w = 0.1561\,\mathrm{kW}$ the wander's innovation. Here $\mathbf{g}_1$ carries the sensor noise through K_p, while $\mathbf{g}_2$ and $\mathbf{g}_3$ carry load through the load column of the one-step map. A load event is counted as a white increment with the variance of a rate-λ arrival of 10 kW steps: a device for the cost, not a forecast of the load. The right side is linear in λ, so

$$V_e(K, \lambda) = V_{e0}(K) + \lambda V_{e1}(K), \qquad V_m(K, \lambda) = V_{m0}(K) + \lambda V_{m1}(K), \tag{14.5}$$

where $V_e = \mathbb{E}[e^2]$ is of the true T_c and $V_m = \mathbb{E}[(\Delta c)^2]$. With $q = 1/s_T^2$ and $\rho = 1/s_m^2$, the objective is $J = qV_e + \rho V_m$. The result (14.5) is what makes the inverse question of Section 14.3.3 solvable.

14.2 The posterior of the state (step 3)

Every answer reads one posterior, that of Chapter 8. Each route reads it at the fit first, where the algebra is, and then across the draws, where the band enters.

14.2.1 The rows and the covariance

The state is equation (8.8), the rows are those of Section 8.2.1, $\mathcal{L} = \frac{1}{2}\mathbf{r}^\mathsf{T}\mathbf{r}$, and they run on the reduced state of that chapter. The information matrix of the normalized rows is $\mathbf{J}^\mathsf{T}\mathbf{J}$; with unnormalized rows it is $\mathbf{J}^\mathsf{T}\mathbf{W}\mathbf{J}$. The MAP state has $\hat{g} = 0.01833$, $\hat{g}_w = -0.16764$, $\hat{b} = 0.05708\,\mathrm{K}$, and the block of the covariance on those three is

$$\mathbf{\Sigma}_{(g,g_w,b)} = \begin{bmatrix} 0.001772 & -0.001827 & -0.007321 \\ -0.001827 & 0.002348 & 0.00913 \\ -0.007321 & 0.00913 & 0.036048 \end{bmatrix}, \tag{14.6}$$

with spreads 0.0421, 0.04846 and 0.18986 K. The capacities carry their priors, log-sd 0.15, independent of this block, since no row reads them.

The draws are those of the draws route of Section 8.2.3: a t_5 proposal with seed 14, $N = 4000$, effective size 2403, with the dynamic entries drawn independently at seed 15. The posterior mean of $(g, g_w, b, \ell_c, \ell_a, \ell_h)$ is (0.0207, −0.1669, 0.0541, 0.0014, −0.0022, 0.0054). Standard errors below come from the weights and are given in parentheses.

14.2.2 The map at the MAP state

At the fit and 20 kW, equation (14.1) reads

$$\mathbf{A} = \begin{bmatrix} 0.96105 & 0.0299 & -0.07994 & 0 \\ 0.35879 & 0.33435 & -2.42739 & 0 \\ 0 & 0 & 0.60653 & 0 \\ 0 & 30 & 0 & 1 \end{bmatrix}, \qquad \mathbf{B} = \begin{bmatrix} -0.01523 \\ -0.80008 \\ 0.39347 \\ 0 \end{bmatrix}. \tag{14.7}$$

The open-loop eigenvalues are 1, 0.97773, 0.60653 and 0.31767, and $\det \mathbf{A} = 0.18839$. This map's radius at $K_p = 0.41$ agrees with the plant's own to 4.4×10^{-16}. The pair $(\mathbf{A}, \mathbf{B})$ is controllable: the smallest singular value of the controllability matrix, with z in units of 300 K s, is 1.08×10^{-3}.

Table 14.2 gives what each tuning does on this map. Each row is a tuning, with its gains, the stationary spreads of the error and of the valve move from equation (14.4) at 0.5 events per hour, and J at the baseline. These are stationary figures on the linear map, unlike the single noisy record of the opening scene. They are the numbers an inverse question has to explain.

Table 14.2: What each tuning does at the fit, 20 kW, $\lambda = 0.5$ per hour, sensor noise 0.10 K.

Tuning	K_p [1/K]	K_i [1/(K s)]	sd e [K]	sd Δc [–]	J, baseline [–]
Plant's	0.10	0.0020	0.0542	0.0116	0.05029
Sheet's	0.41	0.0082	0.1018	0.0580	0.24991
0.41, K_i held	0.41	0.0020	0.0664	0.0648	0.17547
Doubled	0.20	0.0040	0.0644	0.0243	0.08102
Halved	0.05	0.0010	0.0710	0.0059	0.08154
Slow reset	0.10	0.0005	0.1090	0.0138	0.19485
Rated limit	0.80	0.0160	0.2367	0.1942	1.83977

14.3 The answers (step 4)

Each answer starts from $\hat{\mathbf{z}}$ and the map (14.7) built on it, then carries the band where the question needs it. All the inverse questions take one general form.

14.3.1 The FULL route, an LQ law on the hall's state

Control engineer's question. Is this PI an LQ law on the hall's state?

General form. The forward problem is equation (2.23), $a^* = \arg\min_a \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]$. The inverse question fixes the action and a class $\mathcal{L}$ of losses, and asks for the set

$$\{\, L \in \mathcal{L} \;:\; a^{\text{given}} \in \arg\min_a \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}] \,\}. \tag{14.8}$$

Nothing of the hall appears here. The class has to be stated, since an empty or a universal answer can come from the class alone.

Specialized form. Equation (14.8) with a^{given} the row $\mathbf{F}$ of (14.1), on the state $\mathbf{x}$, and $\mathcal{L}$ the sums of stage costs $\mathbf{x}^{\mathsf{T}}\mathbf{Q}\mathbf{x} + Rv^2$, with no cross term, counted at the samples. The hypotheses are: a single input; $(\mathbf{A}, \mathbf{B})$ controllable; the loop stable; $\mathbf{Q} \succeq 0$; $R > 0$. With a known plant and this class the law does not depend on the noise, by Chapter 12, so the expectation goes and the test is algebra at the fit. The law has $\mathbf{F} = -\mathbf{B}^{\mathsf{T}}\mathbf{P}\mathbf{A}/d$ with $d = R + \mathbf{B}^{\mathsf{T}}\mathbf{P}\mathbf{B}$, and $\mathbf{P}$ solves $\mathbf{P} = \mathbf{Q} + \mathbf{A}^{\mathsf{T}}\mathbf{P}\mathbf{A} - d\,\mathbf{F}^{\mathsf{T}}\mathbf{F}$.

Take z on the unit circle and write $\mathbf{H} = (\mathsf{z}\mathbf{I} - \mathbf{A})^{-1}\mathbf{B}$. Expanding the left side gives, for any symmetric $\mathbf{P}$,

$$\mathbf{P} - \mathbf{A}^{\mathsf{T}}\mathbf{P}\mathbf{A} = (\bar{\mathsf{z}}\mathbf{I} - \mathbf{A}^{\mathsf{T}})\mathbf{P}(\mathsf{z}\mathbf{I} - \mathbf{A}) + (\bar{\mathsf{z}}\mathbf{I} - \mathbf{A}^{\mathsf{T}})\mathbf{P}\mathbf{A} + \mathbf{A}^{\mathsf{T}}\mathbf{P}(\mathsf{z}\mathbf{I} - \mathbf{A}). \tag{14.9}$$

Multiply by $\mathbf{H}^*$ on the left and $\mathbf{H}$ on the right, and use $\mathbf{P} - \mathbf{A}^{\mathsf{T}}\mathbf{P}\mathbf{A} = \mathbf{Q} - d\,\mathbf{F}^{\mathsf{T}}\mathbf{F}$ and $\mathbf{B}^{\mathsf{T}}\mathbf{P}\mathbf{A} = -d\,\mathbf{F}$:

$$\mathbf{H}^*\mathbf{Q}\mathbf{H} - d|\mathbf{F}\mathbf{H}|^2 = \mathbf{B}^{\mathsf{T}}\mathbf{P}\mathbf{B} - d\,\mathbf{F}\mathbf{H} - d\,(\mathbf{F}\mathbf{H})^*. \tag{14.10}$$

Adding R to both sides and collecting gives the return-difference identity, with $L = -\mathbf{F}\mathbf{H}$,

$$d\,|1 + L|^2 = R + \mathbf{H}^*\mathbf{Q}\mathbf{H}. \tag{14.11}$$

For a single input $1 + L = \chi_c/\chi$, where $\chi(\mathsf{z}) = \det(\mathsf{z}\mathbf{I} - \mathbf{A})$ and $\chi_c(\mathsf{z}) = \det(\mathsf{z}\mathbf{I} - \mathbf{A} - \mathbf{B}\mathbf{F})$. Multiply by $|\chi|^2$, and write $\mathbf{Q} = \mathbf{C}^{\mathsf{T}}\mathbf{C}$:

$$d\,|\chi_c|^2 = R\,|\chi|^2 + \|\mathbf{C}\,\mathrm{adj}(\mathsf{z}\mathbf{I} - \mathbf{A})\mathbf{B}\|^2. \tag{14.12}$$

The last term has lower degree, because $\mathrm{adj}(\mathsf{z}\mathbf{I} - \mathbf{A})\mathbf{B}$ has degree at most $n - 1$ with n the size of the state. That pins the constant. A polynomial p times $p(1/\mathsf{z})$ has top coefficient $p_n p_0$, and the constant term of a characteristic polynomial is $(-1)^n$ times a determinant, so comparing the coefficients of z^n gives $d\det(\mathbf{A} + \mathbf{B}\mathbf{F}) = R\det\mathbf{A}$. Hence

$$\gamma \equiv \frac{\det\mathbf{A}}{\det(\mathbf{A} + \mathbf{B}\mathbf{F})} = \frac{d}{R}, \qquad \Phi(\omega) \equiv \gamma\,|\chi_c(\mathrm{e}^{\mathrm{i}\omega})|^2 - |\chi(\mathrm{e}^{\mathrm{i}\omega})|^2. \tag{14.13}$$

Since $\mathbf{P} \succeq 0$ gives $d \geq R$, and the last term of (14.12) is nonnegative, a cost requires

$$\gamma \geq 1 \quad \text{and} \quad \Phi(\omega) \geq 0 \text{ for all } \omega, \tag{14.14}$$

with no knowledge of $\mathbf{Q}$. The first condition is a quick check and not a second requirement. No eigenvalue of $\mathbf{A}$ or of the stable $\mathbf{A} + \mathbf{B}\mathbf{F}$ lies outside the circle, so by Jensen's formula the means of $\ln|\chi|^2$ and $\ln|\chi_c|^2$ over the circle vanish, and the mean of $\ln(\gamma|\chi_c|^2/|\chi|^2)$ is $\ln\gamma$; at 0.10/0.002 both read -0.09001. Then $\Phi \geq 0$ makes that ratio at least one everywhere, which implies $\gamma \geq 1$. The test is also sufficient. If (14.14) holds, Φ is a nonnegative trigonometric polynomial of degree at most $n - 1$, because γ cancelled the top term. It has a spectral factor N of that degree, and controllability makes $N = \mathbf{c}\,\mathrm{adj}(\mathsf{z}\mathbf{I} - \mathbf{A})\mathbf{B}$ for some row $\mathbf{c}$. The cost $\mathbf{Q} = R\,\mathbf{c}^{\mathsf{T}}\mathbf{c}$ with that R then has (14.12), and χ_c, being stable, is its unique stable factor, so $\mathbf{F}$ is that cost's LQ law. The integrator's mode is seen, because $\Phi(0) = \gamma|\chi_c(1)|^2 > 0$. The test does not depend on R except as a scale, so a row is either an LQ law for some weight or for none.

Building blocks. The map (14.7) at $\hat{\mathbf{z}}$. The determinant gives γ a closed form: $1/\gamma = 1 + \mathbf{F}\mathbf{A}^{-1}\mathbf{B}$, and the row of z in $\mathbf{A}$ makes the last entry of $\mathbf{A}^{-1}\mathbf{B}$ equal to $-\Delta t$ times its T_c entry, so

$$\frac{1}{\gamma} = 1 + (K_p - K_i\Delta t)\,x_c, \qquad x_c = [\mathbf{A}^{-1}\mathbf{B}]_{T_c} = +2.35454. \tag{14.15}$$

At 0.10/0.002 this gives 0.913925, as the determinants do. The sign of x_c comes from the actuator's lag. Positive x_c means $\gamma \geq 1$ needs $K_i/K_p \geq 1/\Delta t = 0.03333$, a reset time at or under one sampling period. That closed form explains γ and not Φ.

The plant's numbers. Table 14.3 applies the test to the named tunings. Each row is a tuning, with γ, the smallest Φ over the circle, which falls at $\omega = \pi$ for each PI, and the verdict. Figure 14.2 draws what the test sees, in the form of (14.11) divided by R: $\gamma|1 + L|^2$ against the frequency in radians per step on logarithmic axes, for the plant's tuning, the sheet's and a true LQ row, with the line at one that (14.14) requires it to stay above.

Table 14.3: The FULL route, fit, 20 kW. The last row is the LQ law of a stated cost, a check that the test passes what it should.

Tuning	γ [–]	smallest Φ [–]	Verdict
0.10 / 0.002	0.9139	−7.420	none: $\gamma < 1$
0.41 / 0.0082	0.7214	−23.94	none: $\gamma < 1$
0.41 / 0.002	0.5482	−36.08	none: $\gamma < 1$
0.20 / 0.004	0.8415	−13.65	none: $\gamma < 1$
0.05 / 0.001	0.9550	−3.879	none: $\gamma < 1$
0.10 / 0.0005	0.8332	−13.36	none: $\gamma < 1$
0.80 / 0.016	0.5703	−36.78	none: $\gamma < 1$
LQ row, **Q** on e and z	4.017863	+0.002532	exists; $d/R = 4.017863$

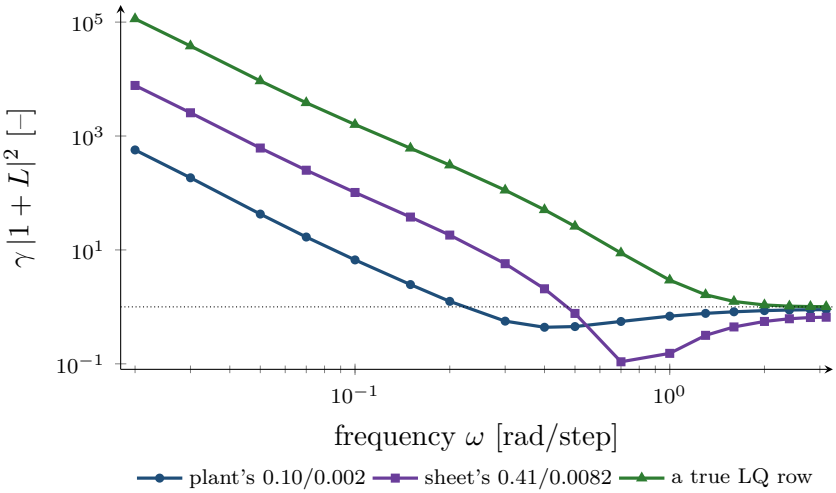


Figure 14.2: What the test (14.14) sees on the circle, FULL route, fit, 20 kW, from equations (14.11) and (14.13): $\gamma|\chi_c|^2/|\chi|^2$. Dotted: the value one, which an LQ law stays at or above. The steep rise at low frequency is the integrator.

The LQ row stays above the line at every frequency, reaching 1.00772 at $\omega = \pi$. The two PI rows fall below it above about 0.22 and 0.47 radians per step and stay below, with the sheet's lowest drawn point at 0.10859, and they end at 0.89416 and 0.65855. Here $\gamma < 1$ alone settles them, but it does not settle the question. At this 30 s step, on a grid of 81 by 81 tunings, K_p from 0.002 to 2 and K_i/K_p from 0.0005 to 0.5, 4920 **are stable,** 1155 **of those have** $\gamma \geq 1$**, and none has a cost**. The test $\Phi \geq 0$ fails on every one of the 1155.

How far that reaches has to be said. "None of 4920" is a grid at the fit, 20 kW and one sampling period; a tuning between its points is not covered by it, and another period is not either, as the continuous limit below shows. The derived test (14.14) is what generalizes: it decides any tuning on any plant that meets its hypotheses, and the closed form (14.15) excludes in advance every tuning whose reset is longer than one period on this plant at this period. Beyond the grid, the two named tunings fail across the conditions tried: γ runs from 0.9244 to 0.8916 for the plant's and 0.7490 to 0.6673 for the sheet's between 5 and 45 kW, at the truth

they are 0.9136 and 0.7205, and the share of the posterior's weight with a cost, on the first 1000 draws, is 0.0000 for each. Those two are named tunings and those are finite checks.

A second route, without frequencies. The law's condition $(R+\mathbf{B}^{\mathsf{T}}\mathbf{P}\mathbf{B})\mathbf{F}+\mathbf{B}^{\mathsf{T}}\mathbf{P}\mathbf{A}=0$ is linear in $\mathbf{P}$, and so is $\mathbf{Q}(\mathbf{P})=\mathbf{P}-(\mathbf{A}+\mathbf{B}\mathbf{F})^{\mathsf{T}}\mathbf{P}(\mathbf{A}+\mathbf{B}\mathbf{F})-R\mathbf{F}^{\mathsf{T}}\mathbf{F}$. A cost exists if and only if some $\mathbf{P}$ meeting the first has $\lambda_{\min}\mathbf{Q}\geq 0$, a convex search. Cutting planes on that search, in coordinates (0.5 K, 0.5 K, 0.2, 300 K s), give upper bounds of -0.0123, -0.0057 and -0.1506 on $\lambda_{\min}\mathbf{Q}$ for 0.10/0.002, 0.41/0.0082 and 0.41/0.002. No feasible cost exists on any of them, by a route that uses no frequency. The true LQ row gives an upper bound of $+0.00528$ and a best found of $+0.00477$.

What a diagonal cost would be. With $R=25$, the matrix $\mathbf{P}$ has ten entries, and the four gain conditions and six off-diagonal entries of $\mathbf{Q}(\mathbf{P})$ make ten linear equations. So a diagonal cost is unique if it exists. At 0.10/0.002 it has $q_T=-7.29$, $q_e=+0.733$, $q_u=-44.4$ and $q_z=+9.14\times10^{-5}$: two rewards, so it is not a cost. The last entry follows the identity at low frequency. At $\mathsf{z}=1$, $\chi=0$, and $\mathrm{adj}(\mathbf{I}-\mathbf{A})\mathbf{B}$ lies along z alone, the null vector of $\mathbf{I}-\mathbf{A}$, with $\chi_c(1)=-K_i$ times its entry. So (14.12) leaves $dK_i^2=q_z$, that is

$$q_z=R\gamma K_i^2. \tag{14.16}$$

The continuous limit. The test counts cost at the samples, so the sampling period might be thought to matter, and it does. In continuous time the law is $\mathbf{F}=-\mathbf{B}^{\mathsf{T}}\mathbf{P}/R$ with $\mathbf{P}\succeq 0$, so $\mathbf{F}\mathbf{B}=-\mathbf{B}^{\mathsf{T}}\mathbf{P}\mathbf{B}/R$. The PI has no gain on the valve's state and $\mathbf{B}$ enters only there, so $\mathbf{F}\mathbf{B}=0$. That forces $\mathbf{B}^{\mathsf{T}}\mathbf{P}\mathbf{B}=0$, hence $\mathbf{P}\mathbf{B}=0$, hence $\mathbf{F}=0$. So no nonzero PI on T_c with an actuator state is a continuous LQ law. That result is only relative degree: it covers any row with no actuator gain, and it says nothing about a sampled row, which the test (14.14) decides period by period.

Table 14.4 gives the sampled verdicts. Each row is a period, with x_c, γ at 0.10/0.002, and the count of tunings that pass on a log grid of K_p from 0.0005 to 3 and K_i/K_p from 0.0001 to 1: with no lag on grids of 45 by 45 and 70 by 70, and with the 5 s sensor lag taken up below on the first. Below 30 s, at 0.10/0.002, γ approaches one from below and nothing passes without the lag. Passes start between 45 and 48 s. At 60 s, 0.07/0.002 passes, a reset of 35 s and so under one period, with $\gamma=8.386$ and smallest $\Phi=+0.0084$, and a round trip by the spectral factor returns exactly (0, 0.07, 0, 0.002). The named tunings still fail at 60 s: the plant's has $\gamma=1.5439$ and smallest $\Phi=-0.1864$, the sheet's $\gamma=-2.2500$. **So on this full state, at this** 30 s **step and with no lag, no PI on the grid passes; at slower sampling some do; and whether a PI has an LQ objective depends on the sampling.**

Table 14.4: Tunings that pass the test (14.14) by controller period, fit, 20 kW. A dash marks a grid not run.

Period [s]	x_c [–]	γ at 0.10/0.002 [–]	No lag, 45^2	No lag, 70^2	5 s lag, 45^2
1	+0.00151	0.9999	0	–	88
5	+0.04057	0.9964	0	–	0
15	+0.43983	0.9701	0	–	18
30	+2.35454	0.9139	0	0	469
45	+7.19346	0.9329	0	0	160
48	+8.71595	0.9663	3	5	88
60	+17.61518	1.5439	54	134	2
120	+290.59281	−0.0252	64	157	0

Why the cross term is forbidden. The stage cost $(v - \mathbf{F}\mathbf{x})^2$ is semidefinite and has value zero on the stable law $v = \mathbf{F}\mathbf{x}$. So every stabilizing row is optimal for it, and with a free cross term allowed the question (14.8) has an answer for every tuning and says nothing. The test forbids it for that reason, and a claim of optimality that allows a free one is empty. That applies only to a free one. A continuous cost counted between the samples under the hold has, once sampled, a restricted cross term, and it is a class of its own. Cutting planes on that class at 30 s found none feasible on any of 60 stable tunings, with upper bounds on $\lambda_{\min}$ of -0.0159, -0.0334 and -0.4341 at 0.10/0.002, 0.41/0.0082 and 0.41/0.002.

The verdict is fragile. With a first-order 5 s discharge sensor, a five-entry state, $x_{\text{read}} = -36.34876$ flips the sign in (14.15). Low-gain cells then pass: at 0.02/0.0004, $\gamma = 1.4100$ with smallest $\Phi = +7.679 \times 10^{-7}$, and the round trip by the spectral factor returns 0.02 and 0.0004. The plant's and the sheet's still fail. The verdict reads the plant's relative degree, and five seconds of sensor nobody modeled changes it, so a designer cannot lean on it.

14.3.2 The REDUCED route, an LQ law on the model the PI sees

Control engineer's question. And on the model the PI sees?

General form. The inverse of equation (2.23), equation (14.8), with the plant replaced by a stated reduction.

Specialized form. Equation (14.8), the inverse of (2.23), with a^{given} the row (K_p, K_i) on a two-entry state (x, z),

$$x_{k+1} = a\,x_k - \theta\,v_k, \qquad z_{k+1} = z_k + \Delta t\,x_k, \tag{14.17}$$

where x is the discharge error's change on the differenced reading. A first-order model buys a state on which the PI is already full-state feedback, and it fails by having no actuator. The test (14.14) applies with the same hypotheses and $n = 2$. Two gains now meet two diagonal weights, $\mathbf{Q} = \text{diag}(q_d, q_z)$, so the diagonal cost is one linear solve: $\mathbf{P}$ has three entries, against two gain conditions and one off-diagonal entry of $\mathbf{Q}(\mathbf{P})$. The two weights read in plain terms as $q_d = 1/s_d^2$, a spread on the error's rate in K per step, and $q_z = 1/(\Delta t\,s_T)^2$.

A diagonal cost is one form of cost, and $\Phi \geq 0$ is the test; the diagonal cost is read where $q_d \geq 0$. A guide to where that holds comes from the model's continuous version, with $\dot{x} = -x/\tau + g\,v$ in magnitude, where $d = R$ and $|\chi_c|^2 - |\chi|^2$ is $(g^2/R)(q_d\omega^2 + q_z)$, with $\chi = s(s + 1/\tau)$ and $\chi_c = s^2 + (1/\tau + gK_p)s + gK_i$. Comparing the coefficients of ω^2, $q_d \geq 0$ means

$$\frac{K_i}{K_p} \leq \frac{1}{\tau} + \frac{gK_p}{2}. \tag{14.18}$$

So on this model a cost on the error's rate that is not a reward needs a reset that is not too fast, by a number the model sets.

Building blocks. The map's step response from (14.7), least squares over 20 steps (10 minutes, the window of Section 8.3.4): $a = 0.81542$, $\theta = 1.18459$, a time constant of 147 s, a steady gain of 6.4176 K per unit against the plant's 10.5177, an rms misfit of 0.149 K per unit, and $g = 0.04365$ K per unit second. Windows of 10 and 40 steps give 0.80640 and 0.87225 for a, with misfits of 0.1662 and 0.3563. **The mismatch has a size: the model's critical gain at ratio** 0.02 **is** 2.1893 **against the plant's** 1.0129**.** The plant's verdict below does not rest on that gap. Models with K_c set to 1.0129 and τ from 50 to 300 s all keep the smallest Φ negative at 0.10/0.002, from -0.3252 to -1.135, so K_c off by two does not drive it. The sheet's fails likewise from 60 s up, and passes at 50 s with $+0.1947$.

The plant's numbers. Table 14.5 gives the route's verdict. Each row is a tuning, with the diagonal weight on the rate at $R = 25$, the temperature spread that q_z means, the smallest Φ and the verdict. The temperature spread is (14.16) again: $s_T = s_m/(\Delta t\, K_i\sqrt{\gamma})$, with $\gamma = 1.0617$ at the plant's.

Table 14.5: The REDUCED route, on the model (14.17). The rate spread s_d exists only where $q_d > 0$.

Tuning	q_d at $R = 25$ [–]	s_T [K]	smallest Φ [–]	Verdict
0.10 / 0.002	−1.966	3.235	−0.4361	none: needs a reward on the error's rate
0.41 / 0.0082	−8.299	0.710	−1.752	none
0.41 / 0.002	+8.797	2.337	+0.01028	exists, $s_d = 0.337$ K/step
0.10 / 0.0005	+0.428	12.48	+0.0003602	exists, $s_d = 1.529$ K/step

Some tunings do have a cost here, so the first route's emptiness was not a property of PI as such. The largest ratio with a cost, by a scan of 400, is 0.00754, 0.01116 and 0.01472 at $K_p = 0.10$, 0.41 and 0.80, against the plant's 0.02; the guide (14.18) reads 0.00898, 0.01575 and 0.02426, and it overstates the sampled limit. The plant's reset is 50 s against the model's 147 s. **The reset is too fast for any LQ cost on this model.** At 0.41/0.002 the round trip returns 0.41 and 0.002 by both the diagonal cost and the rank-one spectral factor.

Figure 14.3 sets the three routes' verdicts on one plane. Each panel is a grid with rows of K_p and columns of K_i/K_p. A cell's mark is the route's verdict at that tuning, at the fit and 20 kW: + a cost or objective exists, − none, × unstable. In the third panel, + means an objective with at most 10 events per hour, a stated plausibility line, and a one above it. The plant's and the sheet's cells are outlined. The first two panels are read here and the third in the next subsection.

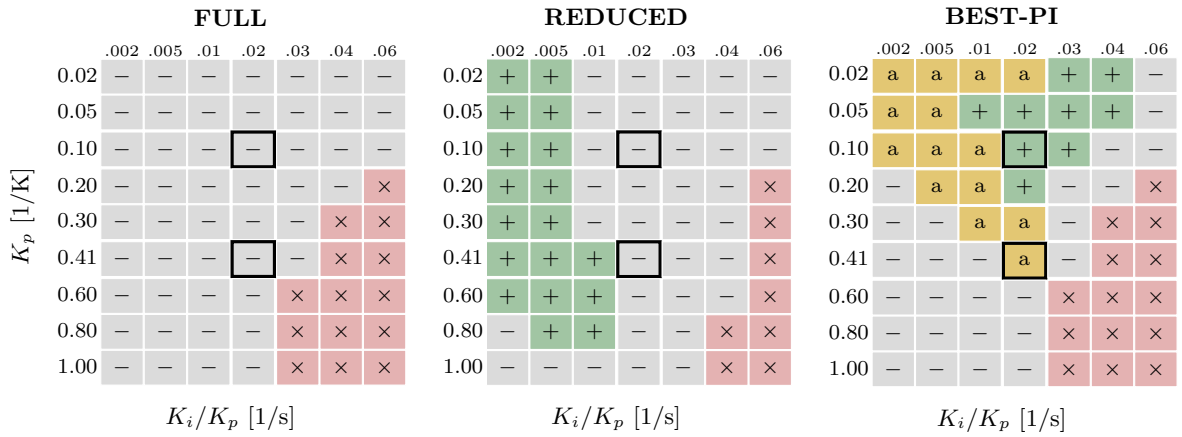


Figure 14.3: The three routes' verdicts on the gain plane, fit, 20 kW. Green +: a cost exists (in the third panel, an objective with at most 10 events per hour). Gold a: an objective only above 10 events per hour; one event per step is 120 per hour. Gray −: none. Red ×: unstable. Outlined: the plant's 0.10/0.002 and the sheet's 0.41/0.0082, both at ratio 0.02. The FULL panel is one coarse view of the 81-by-81 grid of the text.

14.3.3 The BEST-PI route, the best of its own form

Control engineer's question. What is it the best PI for?

General form. The inverse of equation (2.23), equation (14.8), with the actions restricted to the row's form and the class $\mathcal{L}$ a stated objective with free numbers.

Specialized form. Equation (14.8), the inverse of (2.23), with the action the two gains $K = (K_p, K_i)$ on the whole plant, read through one noisy sensor, and $\mathcal{L}$ the objective (14.3) with $s_m = 0.2$ held and $q = 1/s_T^2$ and the event rate λ free. This is static output feedback, so the answer depends on the disturbances, as the LQ law's does not: the fixed form is where the noise reaches the gains. The stated disturbance is the assumption that enters here, and the rows of noise and rate below test it. Stationarity in the log of the gains is

$$\nabla_K \left[q\, V_e(K, \lambda) + \rho\, V_m(K, \lambda) \right] = 0. \tag{14.19}$$

These are two equations in (q, λ). So the two gradients must be parallel. By (14.5) each gradient is linear in λ, so the condition is a quadratic in λ,

$$\det\left[\nabla V_{e0} + \lambda \nabla V_{e1}, \quad \nabla V_{m0} + \lambda \nabla V_{m1} \right] = 0, \tag{14.20}$$

and with $\mathbf{e}$ and $\mathbf{m}$ the two columns at a root, $q = -\rho\, \mathbf{e}^{\mathsf{T}}\mathbf{m} / \mathbf{e}^{\mathsf{T}}\mathbf{e}$ follows directly. An objective exists if and only if a root has $\lambda \geq 0$, $q > 0$ and a positive-definite Hessian. Where two roots qualify, the lower rate is reported, by rule.

Building blocks. **S** of (14.4), at the map (14.7); gradients and Hessian by differences in the log of the gains, step 10^{-3}.

The plant's numbers. Table 14.6 gives the implied objective. Each row is a tuning, with the implied spread and rate and the verdict. The third panel of Figure 14.3 shows the same on the plane.

Table 14.6: The BEST-PI route, fit, 20 kW, sensor noise 0.10 K, $s_m = 0.2$. The rate of one event per step is 120 per hour.

Tuning	s_T [K]	Events per hour [1/h]	Verdict
0.10 / 0.002	0.253	0.530	objective; second root 8.582 K, 277.141 per hour
0.41 / 0.0082	0.430	**229.789**	arithmetic only: an rms load change of 13.84 kW per step on a 20 kW load
0.20 / 0.004	0.247	5.805	objective
0.05 / 0.001	0.300	0.055	objective
0.41 / 0.002	–	–	none: the root has $q < 0$
0.80 / 0.016	–	–	none: no root with $\lambda \geq 0$

The round trip, the forward map at each implied objective, returns the gains to the digits printed: 0.1000 and 0.00200 at the plant's, 0.4100 and 0.00820 at the sheet's. **Two gains meet two free numbers, so existence is a low bar; what can fail is sign, curvature and plausibility.** That 0.530 sits near the record's 0.5 is a coincidence of constructed values, not a finding, and the objective's baseline was picked here to begin with. The line at 10 events per hour in the figure is a stated choice and so is the reading of 120 per hour as a physical bound, which depends on the 10 kW size, since only the rate times the size squared enters. The size-free reading is the rms load change per step: 13.84 kW at the sheet's rate and 0.66 kW at the plant's, against a 20 kW load.

The rate moves with sensor noise as the noise squared, since with the wander small the two enter almost only as λ/σ_n^2. At sensor noise of 0.02, 0.05, 0.10 and 0.20 K, the plant's tuning

implies 0.009, 0.123, 0.530 and 2.160 events per hour, while its s_T stays near a quarter of a kelvin. The sheet's implies 9.228 per hour at 0.02 K and 229.789 at 0.10 K. **So the sheet's tuning fits a quiet sensor and a busy load, not this hall as stated.**

On the first 1000 draws, the plant's tuning has an objective on all of the weight, with s_T from 0.192 to 0.316 K and 0.425 to 0.669 events per hour between the 5 % and 95 % points. The sheet's has one on 0.5410 of the weight, at 75.257 to 252.007 per hour.

The restriction checks. A restriction cannot beat what it restricts, and the costs at the baseline keep that order. The state known gives 0.00028; the best PI with no sensor noise 0.00073; LQG, a filter on the reading, 0.04682, by two routes that agree to those digits; the best PI with the noise 0.05028. The PI's price over the filter is 7.4 %, and 14.9 % at $s_T = 0.10$ K, 2 per hour. These are this example's.

14.3.4 The designer's table, from a ratio to gains

Control engineer's question. What does the ratio buy?

General form. The expected consequence (2.23), with the action the row and the objective stated.

Specialized form. Equation (2.23) at the fit with the loss (14.3) on three routes. The reduced LQ law on (14.17) with $\mathbf{Q} = \mathrm{diag}(0, 1/(\Delta t\, s_T)^2)$ is then played on the plant, like with like, since what matters is what it does there. The full LQ law on (14.1) has $\mathbf{Q} = \mathrm{diag}(0, 0, 0, 1/(\Delta t\, s_T)^2)$ and reads the whole state, so its row is not a PI. The best PI minimizes J over the two gains. Both LQ laws come from the discrete Riccati solver of Chapter 12.

Building blocks. The pair (14.7), the model (14.17), and V_e, V_m of (14.4). The margin is the factor on both gains together at which the loop first stops settling; the pair is the modulus of Section 8.2.4.

The plant's numbers. Table 14.7 is the table for a designer, at 0.5 events per hour, the fit and 20 kW. Each row is a spread s_T at $s_m = 0.2$; the columns give the reduced law's K_p with its margin and cost on the plant, the full law's margin, and the best PI's gains, pair, margin and cost. Figure 14.4 draws the gain plane on logarithmic axes: the best-PI curves as s_T falls from 1 to 0.05 K, one curve for each event rate; the reduced LQ curve; the plant's stability boundary; and the two tunings.

Table 14.7: From the ratio to the gains. Costs are J of (14.3) on the plant at $\lambda = 0.5$ per hour, like with like. The 0.25 K row is the baseline, picked from the inverse map; the 0.50 K row is Chapter 7's spread.

s_T	Reduced LQ, on the plant			Full LQ	Best PI			
[K]	K_p [1/K]	margin [–]	cost [–]	margin [–]	K_p / K_i	pair [–]	margin [–]	cost [–]
2.00	0.2870	5.46	0.0407	6.09	0.0256 / 0.00140	0.8610	5.02	0.00139
1.00	0.4448	3.40	0.1141	3.98	0.0393 / 0.00167	0.8597	5.06	0.00413
0.50	0.6573	2.19	0.3497	2.76	0.0645 / 0.00185	0.8395	7.23	0.01404
0.25	0.9181	1.46	1.5043	2.05	0.0990 / 0.00197	0.7982	10.34	0.05028
0.10	1.2638	unstable		1.53	0.1356 / 0.00205	0.7306	9.90	0.29004
0.05	1.4325	unstable		1.31	0.1468 / 0.00207	0.6965	9.51	1.13756

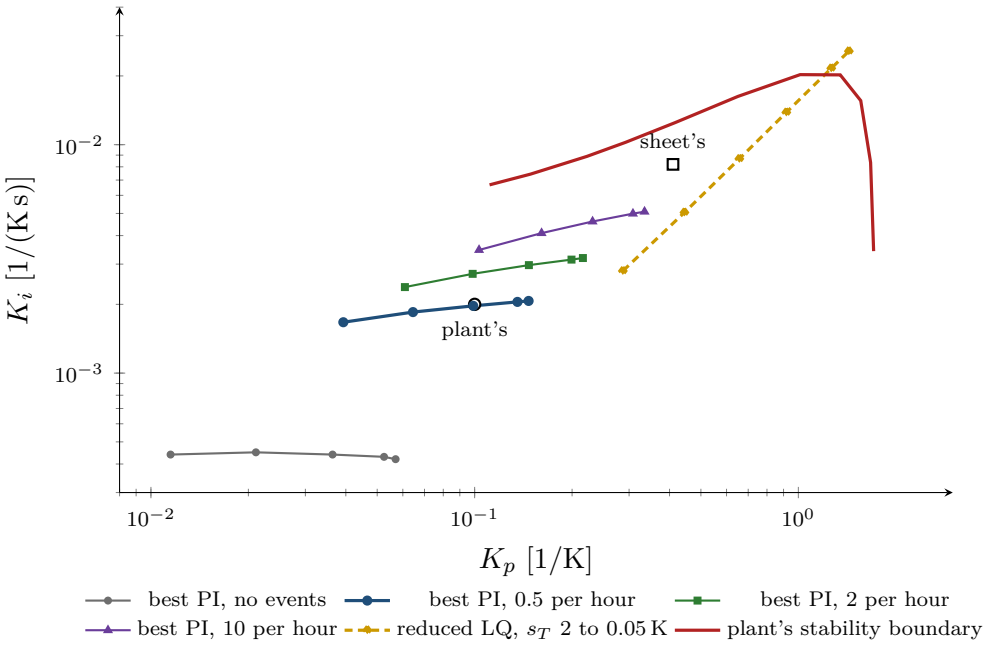


Figure 14.4: The gain plane, fit, 20 kW, sensor noise 0.10 K, $s_m = 0.2$. Along each curve s_T falls left to right, on the best-PI curves from 1 to 0.05 K. Red: where the loop stops settling, with the ratio K_i/K_p held along each ray from the origin; tunings above and to the right of it are unstable. The sheet's tuning lies on no best-PI curve of this figure.

Each column of the table is on one plant, and the differences are large. The best PI's gain does not follow the reduced law's. At the baseline the reduced law's 0.9181 has a margin of 1.46 and costs 1.5043 on the plant, about 30 times the best PI's. That factor is the sensor's noise and not the model's mismatch: certainty equivalence ignores the sensor, so the law's gain is set with no noise in view. With no noise the reduced law costs 0.000735 against the best PI's 0.000726, a factor of 1.014. Only the instability at 0.10 K is mismatch, since the model's critical gain is more than double the plant's. **The inverse map is a tool only on the right route.** The full LQ law's margins can be compared with the sampled guarantee: the identity (14.11) with $\mathbf{Q} \succeq 0$ gives $|1 + L|^2 \geq R/d$, so L stays out of the disc of radius $\sqrt{R/d}$ about -1, and a factor s on the gains cannot bring $sL = -1$ until $s \geq 1/(1 - \sqrt{R/d})$. At $s_T = 0.25$ K, $d/R = 3.8208$ guarantees 2.047, and the realized margin is 2.05.

Temperature weight alone tops out. At $s_T = 0.001$ K the best PI's K_p is 0.1514: past that, at this noise and rate, more gain widens the temperature's own spread, because the reading's noise feeds the error it is meant to remove. In the figure the best-PI curves bunch at their right ends, while the event rate moves them across the plane. The stated rate does as much to the tuning as the stated ratio.

14.3.5 The band and the fixed gains

Control engineer's question. Does the band move the gains?

General form. The expected consequence (2.23) over the posterior, for a fixed row.

Specialized form. Equation (2.23) with the action K and the loss $J(K, \mathbf{z})$,

$$\min_K \ \mathbb{E}\big[J(K, \mathbf{z}) \mid \mathbf{y}\big] \approx \min_K \sum_i w_i \ \min\big(J(K, \mathbf{z}_i),\ J_{\text{cap}}\big), \tag{14.21}$$

with $\mathbf{z}_i$ and w_i the draws of Section 14.2.1. Here is the fixed-form variant of the restriction: $\boldsymbol{\Sigma}$ enters through a gain it multiplies, coupling 4. An LQ law's gain with a known plant would not feel it, and a fixed row on an uncertain plant can. The linear cost diverges at the stability boundary while the plant's stops bound it, Section 8.3.6, so any draw is capped at the simulated cycle at $1.02\,K_c$ at the truth, $J_{\text{cap}} = 3.383$ at the baseline. That is the choice entering here. Near the optimum no draw reaches the cap, the largest J at the band's gains being 0.073, so the cap rows test nothing there; at 0.90/0.018 the cap holds 0.6409 of the weight.

Building blocks. The posterior's draws; $\mathbf{S}$ of (14.4) on each draw's map (8.6); a paired standard error on the excess, the difference of the fit's K and the band's on the same draws.

The plant's numbers. Table 14.8 compares the two at three objectives. Each row is an objective, with the gains best at the fit, the gains best with the band, the band's move in each gain, the excess expected cost of using the fit's gains with its paired standard error, and the probability of stability. Figure 14.5 draws the expected cost across the gains: in the left panel along the line $K_i = 0.02\,K_p$, with the fit's cost and the band's, and in the right panel across K_i with K_p held at 0.0999.

Table 14.8: The band and the gains. N = 4000, effective size 2403, cost capped at 3.383 at the baseline, draws of $(g, g_w, b, \ell_c, \ell_a, \ell_h)$. Sensor noise 0.10 K except where stated.

Objective	K at the fit	K with the band	Move in K_p, K_i [%]	Excess (se) [%]	Pr(stable) [–]
Baseline	0.0990 / 0.00197	0.0999 / 0.00197	+0.9, +0.3	+0.002 (0.001)	>0.9999
$s_T = 0.05$ K, 10/h	0.3349 / 0.00509	0.3382 / 0.00511	+1.0, +0.2	+0.005 (0.002)	>0.9999
The same, noise 0.001 K	1.1850 / 0.01598	1.1799 / 0.01563	−0.4, −2.2	+0.341 (0.038)	>0.9999

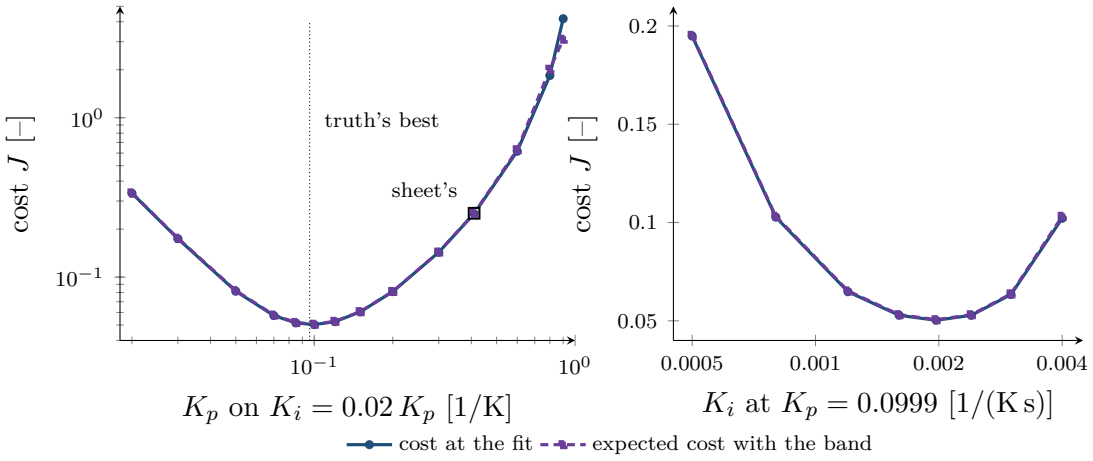


Figure 14.5: The baseline cost across the gains, at the fit and under the band with the capped cost (14.21), draws as in Table 14.8. Left: the line through both named tunings, the sheet's marked; dotted, the gain best at the stated truth, whose dynamic entries sit at the prior's center. Right: across K_i. The two curves nearly coincide around a broad minimum. At $K_p = 0.90$ the band's cost lies below the fit's because the band's is capped at 3.383 and the fit's is not.

The band hardly moves the tuning, and this is stated as small. The move in K_p is about 1 %. $\boldsymbol{\Sigma}$ does enter through the gain it multiplies, but the cost surface is flat near the minimum and the margin there is about 10, far from the boundary where the band would bite. So this is a statement about this plant at this objective. A single prior corner, moved by 1.96

of its own spreads, moves the best K_p to a ratio between 0.880 and 1.142 of the fit's, by C_h; averaged, the corners nearly cancel. With the prior's sd on the dynamic entries at 0.30 the move is +4.0 % and the excess +0.047 % (se 0.010), and with an extremely quiet sensor at $s_T = 0.05$ K and 10 per hour the gains go high enough that the excess is +0.341 %. A plant nearer its boundary would feel more.

Under the band, the sheet's tuning costs 4.969 times the best at the baseline and holds the envelope of the pair of Section 8.3.4, 0.8609, with probability 0.5243; the plant's costs 1.000 times the best and holds it with probability 0.9878. That baseline was picked to favor the plant. At the sheet's own implied objective, at the fit, the plant's tuning costs 6.818 times the best. Retuning to each of 200 draws, resampled by weight, saves 0.355 % (se 0.027) of the fixed cost. That is the expected value of perfect information at this objective, and the ceiling on any test of Section 14.6.

14.3.6 The load and a schedule

Control engineer's question. Does a schedule do better than the band?

General form. The expected consequence (2.23) again, with one action for each scenario node against one action for all.

Specialized form. Equation (2.23) with the loss (14.21) averaged over a stated mix of nodes, the six loads 5, 10, 15, 20, 30 and 45 kW equally likely. A fixed tuning has one K; a schedule has $K(Q^*)$ with the load known. The mix is a choice, and it fails if the hall spends its time elsewhere. Each load is pinned in its row, closed loop.

Building blocks. The first 2000 draws, effective size 1175; the map (8.6) at each load; K_c at the fit by load, and K_{95} from the weighted quantile.

The plant's numbers. Table 14.9 compares three choices for two objectives. Each row is an objective; the columns give the best fixed cost on the mix, the excess of gains tuned at 20 kW alone, the schedule's saving and the best K_p at the two ends of the load range. Figure 14.6 draws K_c, K_{95} and the two objectives' best K_p against load on a logarithmic gain axis.

Table 14.9: The load, six loads equally likely, with the band. "Fast" is $s_T = 0.05$ K, 10 events per hour, sensor noise 0.02 K.

Objective	Fixed [–]	Tuned at 20 kW only [%]	Schedule [%]	Best K_p, 5 to 45 kW [1/K]
Baseline	0.05390	+0.46	−1.57	0.1138 to 0.0830
Fast	0.82213	+86.08	−10.68	0.5975 to 0.8413

The direction of K_p in the schedule depends on the objective and on how events scale with load. K_c falls with falling load, from 2.1069 at 45 kW to 0.5116 at 5 kW. With a fixed 10 kW event, the baseline's best K_p rises as the load falls, against K_c, and the fast objective's falls with it, the textbook direction. That baseline result depends on the fixed event. Scaling the event with load, 10 kW at 20 kW, gives at the fit a best K_p of 0.0605 at 5 kW and 0.1360 at 45 kW, the textbook direction. K_i rises with load at both objectives, 0.00170 to 0.00226 at the baseline and 0.00741 to 0.01519 at the fast one, so only K_p diverges. No cause is claimed here: the height of the gain does not separate the two directions, since at the fit a sensor of 0.02 K at the baseline's spread and rate gives 0.2617 at 5 kW falling to 0.2233 at 45 kW, while $s_T = 0.05$ K, 10 per hour gives 0.3110 rising to 0.3602.

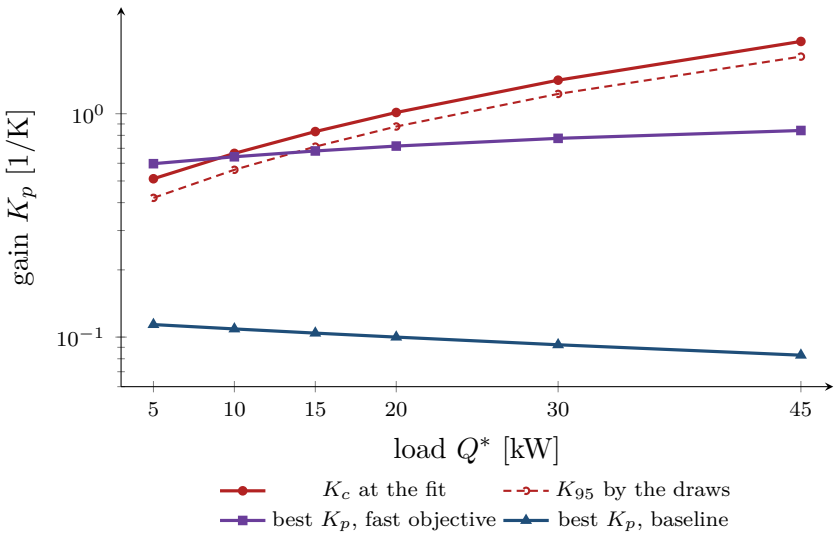


Figure 14.6: Gains against load with the band. K_c and K_{95} are on the line of ratio 0.02; each best gain has its own K_i, so it is not compared with them on that line, and the schedule's probability of stability stays above 0.9999 at every load. With the event held at 10 kW, the two objectives' best gains run in opposite directions.

A schedule on known load is worth 1.57 % and 10.68 % here. Perfect information about the plant at the same objectives is worth 0.355 % and 1.539 %, the second by the state route of Section 14.6.1. So in this example knowing the load is worth more than knowing the plant. Gains tuned at 20 kW for the fast objective cost 5.62741 at 5 kW, against 1.17948 for the schedule there, and hold stability there with only 0.9586. Percentages are this example's and rest on the stated mix.

14.4 The checks (step 5)

The answers rest on a stationary linear cost and on routes that can be tested against each other. The tests compare like with like, and a truth row closes them.

14.4.1 The cost against the full calculation

Table 14.10 lists the tests. Each row names a check, what is compared and the result. The nonlinear runs use the stated truth, at K_p with ratio 0.02, 5500 steps, seed 141 in turn.

The linear stationary cost tracks the nonlinear truth to about 2 % in spread. The large event is 2.9 % off in sum e^2, the cost of counting a 10 kW step at a pinned load. A 5 s sensor lag, absent from the band, moves the best PI to 0.1002/0.00191 against 0.0990/0.00197, and ignoring it costs +0.111 %; the plant's tuning then implies 0.277 K and 0.595 per hour, and the sheet's has no objective. So the BEST-PI reading, unlike the FULL verdict, moves little with the sensor at the plant's tuning.

14.4.2 The state route against the draws

The value of a test in Section 14.6 is given by a quadratic about the fit and by reweighted draws. They agree within about two and a half standard errors on every candidate at the baseline. Perfect knowledge is 0.382 % by the first and 0.379 % by the second, against 0.355 % (se 0.027) from the retuned 200 draws, a third route. The reduced state's limit in σ_p is that of Table 8.9

Table 14.10: The checks. Linear values in parentheses are the stationary figures of (14.4) at the truth.

Check	Compared	Result
This map against the plant's	radius at 0.41	difference 4.4×10^{-16}
Differenced form against position	sd e, no events, 0.10/0.002; 0.41/0.0082	0.04264 both; 0.10138 both
Nonlinear truth, $K_p = 0.10$	sd e; sd Δc	0.0433 (0.0424); 0.0112 (0.0114)
Nonlinear truth, $K_p = 0.41$	sd e; sd Δc	0.1002 (0.1012); 0.0575 (0.0580)
A +2 kW event, $K_p = 0.10$, noiseless	sum e^2	0.0099 (0.0100)
A +10 kW event, $K_p = 0.10$, noiseless	sum e^2	0.2431 (0.2501)
Restrictions, baseline	state known; no noise; LQG; best PI	$0.00028 \le 0.00073$; $0.04682 \le 0.05028$
LQG by two routes	closed form; closed-loop covariance	0.04682; 0.04682
Round trip, FULL with the lag	0.02/0.0004 by the spectral factor	0.02, 0.0004
Round trip, FULL at 60 s	0.07/0.002 by the spectral factor	(0, 0.07, 0, 0.002)
Round trip, REDUCED	0.41/0.002 by diagonal cost; by spectral factor	0.41, 0.002 both
Round trip, BEST-PI	forward map at the implied objective	five tunings returned to the digits printed
Dynamic prior narrowed, the limit of the band	prior sd 0.05 against none, baseline	K_p 0.0992 against 0.0990; excess +0.000 %
Second seed for the draws	seeds 16, 17	move +1.0 %; expected cost 0.05066 against 0.05058
Cap times 0.25 and 4	band's K	0.0999, 0.00197, no change; no draw reaches the cap there
Fit against the posterior mean	K from each, under the band	0.0990 against 0.0988; excess of the latter +0.005 %

and is not repeated. Here the corresponding narrowing is of the dynamic priors, in the table's prior row.

14.4.3 The stated truth

At the truth, the best gains at the baseline are 0.0960/0.00193 with cost 0.04868. The band's K costs +0.074 % there. The plant's tuning implies 0.246 K and 0.561 per hour, the sheet's 0.385 K and 206.041 per hour, and the FULL verdict at the truth is none for both. As in Section 8.4.5, the truth's dynamic entries sit at the prior's center by construction, so this row tests the fit's product and path and nothing about the capacities. The dotted line in Figure 14.5 is to be read the same way.

14.5 The report (step 6)

On this full state, at this 30 s step and with no lag, no PI on the grid is an LQ law, by a test that decides any tuning; at slower sampling some are, so the answer depends on the sampling. On the model the PI sees, the plant's reset is too fast for one. As the best of its own form, the plant's tuning fits 0.253 K against a 0.2 move at 0.530 events per hour, a reading whose baseline was picked from this same map; the sheet's fits nothing this hall experiences, since it needs 229.789 events per hour at this sensor, an rms load change of 13.84 kW per step on a 20 kW load. The band moves the gains by about 1 %, on this plant, at these priors. Whether a tuning is optimal has become three questions with three answers.

14.5.1 The answers and their bands

Table 14.11 gives the answers in the order of Table 14.1. Each row is a question, with the answer, its band or reach, and the stated truth.

Table 14.11: The report at 20 kW, sensor noise 0.10 K, unless marked. The truth's dynamic entries sit at the prior's center.

Control engineer's question	Answer	Band or reach	Stated truth
Is this PI an LQ law on the hall's state?	no at 30 s, either tuning: γ 0.9139, 0.7214; Φ −7.420, −23.94	the test decides any tuning under its hypotheses; 0 of 4920 is a grid at one period; passes start between 45 and 48 s; share 0.0000 on 1000 draws	none, γ 0.9136, 0.7205
And on the model the PI sees?	no: q_d −1.966, −8.299; yes at 0.41/0.002, q_d +8.797	a model with misfit 0.149 K/unit and K_c 2.1893 against 1.0129	–
What is it the best PI for?	plant's: 0.253 K, 0.530/h; sheet's: 0.430 K, 229.789/h	plant's 0.192–0.316 K, 0.425–0.669/h; sheet's on 0.5410 of the weight, 75.257–252.007/h	0.246 K, 0.561/h; 0.385 K, 206.041/h
What does the ratio buy?	at 0.25 K, 0.5/h: 0.0990/0.00197, margin 10.34, pair 0.7982, cost 0.05028	the table's rows; reduced law at 0.10 K unstable on the plant	best 0.0960/0.00193
Does the band move the gains?	+0.9 %, +0.3 %; excess +0.002 % (se 0.001)	at most +0.341 % across the rows; +0.047 % at prior sd 0.30	band's K +0.074 %
Does a schedule do better than the band?	1.57 % at the baseline, 10.68 % fast, against perfect information's 0.355 %, 1.539 %; direction of K_p depends on the objective and on event scaling	the stated mix of six loads	–
Which test most improves the tuned cost?	set-point +0.5 K, permanent 5 s logger: 0.265 % (draws 0.265, se 0.020), payback 40.5 days	ceiling 0.355 % (se 0.027); white noise at 5 s assumed	–

14.5.2 The entry in the commissioning log

The entry records the tuning, $K_p = 0.10$, $K_i = 0.002$, and the objective it is declared to serve: $s_T = 0.25$ K, $s_m = 0.2$ per step, 0.5 events per hour of 10 kW, sensor noise 0.10 K. It names the route, BEST-PI, and says the baseline spread was picked from the inverse map. It gives the margin on both gains, 10.13 for the logged tuning, and Pr(stable) > 0.9999 under the band with the stated priors. It gives the load range, 20 kW for these figures; between 5 and 45 kW the best K_p is 0.1138 to 0.0830 with a fixed 10 kW event, and a fixed K_p tuned at 20 kW costs +0.46 % on the stated mix. It says no LQ claim is made on either route at this 30 s period. A tuning without a stated disturbance declares no objective, and the entry says so of any tuning that lacks one.

14.5.3 Common mistakes

Table 14.12 lists six mistakes this chapter's plant makes easy. Each row names the mistake, its form on this hall and the correction.

14.6 The next measurement (step 7)

The question is now which test most improves the tuned cost, not which narrows a band. The two rank tests differently, because the value counts only what moves the gains. The record's

Table 14.12: Six mistakes in reading a tuning as optimal.

Mistake	On this hall	Correction
Saying a PI "is optimal" without a route	none on FULL at 30 s, 0.41/0.002 on REDUCED, 0.10/0.002 on BEST-PI	name the state, the sampling period, the class of costs and the form
Taking a grid for a theorem	0 of 4920, at the fit, 20 kW and 30 s; 0.07/0.002 passes at 60 s	the derived test generalizes; say how far each reaches
Allowing a free cross term	$(v - \mathbf{F}\mathbf{x})^2$ makes every stabilizing row optimal	forbid it, and state that the test does
Reading 0.530 as a finding about the record	it sits near 0.5 by a coincidence of constructed values, and moves as noise squared	treat the rate as declared, with its sensor noise
Using the reduced inverse map on the plant	reduced law at 0.10 K unstable, from the mismatch; at 0.25 K, about 30 times the best PI's cost, from the noise it ignores	score any law on the plant with its sensor, like with like
Recommending a tuning without its disturbance	the sheet's fits at 9.228/h with a 0.02 K sensor and 229.789/h with this one	declare s_T, s_m, the rate and the noise with the gains

load step leaves the spreads of $\ln C_c$ and $\ln \tau_a$ at 0.1494 and 0.0936, against 0.0738 and 0.0235 for the set-point step with the 5 s logger, and it is worth more, 0.294 % against 0.265 %.

14.6.1 The preposterior of the tuned cost

A candidate's rows give the preposterior of equation (6.23) with the information of Section 8.6.1, over 40 steps, replacing $\mathbf{\Sigma}$ by $\mathbf{P}_{\text{post}}$ on the six entries. The value of a test is by (2.16) the saving in $\mathbb{E}_{\mathbf{y}}[\min_K \mathbb{E}(J \mid \mathbf{y})]$, read by two routes. The state route takes the best gains as locally linear in the entries, $\ln K^*(\boldsymbol{\theta}) = \ln K^* + \mathbf{S}_K(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}})$, and the excess cost as a quadratic with Hessian $\mathbf{H}$ in the log of the gains, so

$$\text{value} = \tfrac{1}{2}\,\mathrm{tr}[\mathbf{H}\mathbf{S}_K(\mathbf{P}_{\text{prior}} - \mathbf{P}_{\text{post}})\mathbf{S}_K^{\mathsf{T}}]. \tag{14.22}$$

With $\mathbf{P}_{\text{post}} = 0$ this is the expected value of perfect information. The draws route reweights the pool by a normal likelihood with that information, at 300 outer truths, seed 142, on an 11-by-11 grid of gains with factors from 0.7 to 1.4; small inner samples bias it upward, and it serves only as a check. The baseline's sensitivities are, for $\mathrm{d}\ln K_p$ on $(g, g_w, b, \ell_c, \ell_a, \ell_h)$, $(0.359, -0.488, 0.083, 0.147, 0.336, -0.443)$. So at this objective they lean on g_w and C_h about three times as much as on C_c, and the ranking of Section 8.6.2 does not carry over.

14.6.2 The candidates

Table 14.13 lists the candidates. Each row is one; the columns give the value at the baseline by the state route and by the draws, the value by the state route at the fast objective ($s_T = 0.05$ K, 10 per hour, 0.02 K), the one-time cost of running the test in units of J per step, and the days of operation at this value that pay it back. Perfect knowledge is 0.382 % and 0.379 % at the baseline by the two routes.

14.6.3 The recommendation

No test is worth much at this objective. Every candidate saves under a third of one percent, against a ceiling of 0.355 % (se 0.027). The first-ranked row is the record's own load step, 0.294 %, a test already taken; the fit of the dynamic entries is not done in this book's plan up to here, and that step test is what would do it. Of the new tests, the set-point step of +0.5 K with the permanent 5 s logger pays back soonest, in 40.5 days, and is the one to run if

Table 14.13: The candidates, as a percentage of the tuned expected cost now, 0.050606 at the baseline, 0.651379 at the fast objective. The valve bump is open loop and pins the actuator; the rest are closed loop.

Candidate	State [%]	Draws (se) [%]	Fast, state [%]	One-time cost [–]	Payback [days]
Reference discharge sensor, 0.05 K	0.020	0.013 (0.003)	0.015	–	–
Set-point +0.5 K, $K_p = 0.41$, 30 s log	0.193	0.189 (0.015)	1.028	15.63	55.5
Valve bump +0.20, open, 5 s logger	0.282	0.300 (0.024)	1.164	513.25	1250.2
Set-point +0.5 K, permanent 5 s logger	**0.265**	**0.265 (0.020)**	**1.305**	**15.63**	**40.5**
Record's load step, 30 to 45 kW	0.294	0.283 (0.026)	0.801	–	–

any is; at the fast objective it is first by the state route, 1.305 % of a ceiling of 1.539 %, with a payback of 13.6 days. Its value rests on white noise at 5 s, the amendment of Section 8.6.2. A reference discharge sensor teaches almost nothing, 0.020 %.

Those figures are the value of information. Changing the plant is a different thing, and the two are not ranked on one scale here: a test has a one-time cost and a payback, and a sensor has no price in this book. The loop's own discharge sensor made quieter teaches nothing and changes what the tuning reads. At 0.05 K the expected cost falls by 56.81 % retuned to 0.1498/0.00299 and by 46.45 % with the gains left alone; at 0.02 K by 84.63 % and 59.46 %.

The logger's white-noise assumption permits one more reading: the mean of its six samples in a period, with noise of 0.0408 K. With that mean's delay modeled, at the fit, it saves 60.7 % retuned to 0.1709/0.00301 and 48.7 % unretuned, against 56.9 % and 46.7 % for the 0.05 K sensor at the fit. It rests on the same assumption as the logger's test row. The recommendation at this objective is therefore that no test earns its place on the gains it would move. If a quieter reading is had, by a sensor at its price or by that mean if the noise is white, retune to it.

14.6.4 What the analysis showed

The sheet called one tuning well tuned on the strength of a lower peak; the analysis finds that the word needs a route, a disturbance and a form, and that a tuning has an objective only relative to all three. Read as a law on the hall's whole state, at this 30 s step and with no lag, neither tuning has a cost and no PI on the grid does; at 60 s some do, so that claim should not be made of a PI, or denied of one, without its sampling period. Read on the model it sees, the plant's reset is too fast. Read as the best of its own two gains, the plant's tuning makes sense of a quarter-kelvin spread and a load event every two hours, and that reading was chosen by looking; the sheet's tuning makes sense of a busy load and a quiet sensor, which this hall as stated does not have. The band, which the book has followed through every chapter, barely moves the gains here, because the cost is flat where the tuning sits; a known load is worth more than perfect knowledge of the plant, and a quieter reading, retuned, lowers the cost much further, though it has no price here. What would improve this loop's cost is a better reading of the discharge air, and a test that fits its dynamics would not.

14.6.5 Limits of the model

The objective. The objective, the event rate and the sensor noise are stated. Baseline s_T was picked from the inverse map, so the baseline's closeness to the plant's tuning is not evidence. Percentages are this example's.

The cost. It is counted at the samples, with no cross term, and between samples it is uncounted; the class counted between them was tried only at 30 s on 60 tunings. It is a stationary linear cost; the cap stands in for the cycle of Section 8.3.6, and the stop itself is Chapter 15's.

The events. They are small-signal at a pinned load. A 10 kW event is +2.9 % off linear in sum e^2, and a random walk of load is an accounting device, not a forecast.

The verdicts. FULL verdicts are on the grid and at the stated period; the closed form (14.15) explains γ, not Φ. A second BEST-PI root exists for most tunings, and the lower-rate root is reported by rule.

Structure and the sampler. The 5 s lag and rate limits are not in the band, and the sampler stays ideal. K_c at periods of 30, 15 and 5 s is 1.0129, 1.8201 and 4.8734, and with the lag 0.7081, 1.0759 and 1.7284.

The constructed record. The record is constructed from a stated truth, so the truth's agreement tests the method on its own model and not the model on a real hall. Chapter 15 takes up the stop.

14.7 What to retain

Definition 14.1 (The test for an LQ law)**.** For a single-input pair $(\mathbf{A}, \mathbf{B})$, controllable, with $\mathbf{A} + \mathbf{BF}$ stable and $\det \mathbf{A} \neq 0$, the row $\mathbf{F}$ is the LQ law of some $\mathbf{Q} \succeq 0$, $R > 0$, with no cross term and cost at the samples, if and only if $\gamma = \det \mathbf{A} / \det(\mathbf{A} + \mathbf{BF}) \geq 1$ and $\gamma |\chi_c|^2 - |\chi|^2 \geq 0$ on the unit circle.

Three conclusions on this hall go with the method. A tuning question is three questions, on the state, on a model and on the form, and an answer to one is not an answer to another; the first answer also changes with the sampling period. A fixed form depends on the disturbances, so a tuning declares its disturbance or declares no objective. And the band reaches a fixed gain through a gain it multiplies, by an amount this plant's flat cost makes small and another plant's need not. The atlas of questions is Appendix A.

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

Table 14.14: The terms of the chapter and where each is defined. Units for the closed-loop hall.

Term	Expression	Units	Meaning	Defined in
Objective	eq. (14.3)	–	spreads on error and move	14.1.2
Differenced PI	eq. (14.2)	–	same **A**, **B**, **F** on changes	14.1.2
Stationary covariance	eq. (14.4)	mixed	linear in the event rate	14.1.4
Inverse question	eq. (14.8)	–	a class of costs for a given action	14.3.1
Return-difference identity	eq. (14.11)	–	the law on the circle	14.3.1
γ	eq. (14.13), (14.15)	–	the constant the degree pins	14.3.1
Φ	eq. (14.13)	–	nonnegative for an LQ law	14.3.1
Reduced model	eq. (14.17)	mixed	what the PI sees	14.3.2
Implied objective	eq. (14.19), (14.20)	K, 1/h	(s_T, λ) of a tuning	14.3.3
Expected cost with the band	eq. (14.21)	–	fixed gains over the draws	14.3.5
Value of a test	eq. (14.22)	%	the saving in the tuned cost	14.6.1

14.8 Exercises

Exercise 14.1. *Formulation.* A sampled PI holds a tank's level by a pump's speed, with a pump lag. Write the inverse question on the two routes that need no disturbance, and say what the third route must declare. (On the state of level, pump and integrator: for which $\mathbf{Q} \succeq 0$, $R > 0$, without a cross term, is the row the law, by (14.14). On a first-order model of level with an integrator, the same with $n = 2$, the model stated with its misfit. As the best two gains, equation (14.19), the route must declare the inflow's disturbance, its event rate and the sensor's noise.)

Exercise 14.2. *Derivation.* Obtain γ from the identity's leading coefficient. (Equation (14.12) is $d\,\chi_c(\mathsf{z})\chi_c(1/\mathsf{z}) = R\,\chi(\mathsf{z})\chi(1/\mathsf{z}) + N(\mathsf{z})N^{\mathsf{T}}(1/\mathsf{z})$ with N of degree at most $n-1$. The coefficient of z^n on each product is the leading coefficient, one, times the constant term $(-1)^n \det(\cdot)$, and N contributes nothing. So $d\det(\mathbf{A}+\mathbf{BF}) = R\det\mathbf{A}$ and $d/R = \gamma$. On the true LQ row this reads 4.017863 both ways.)

Exercise 14.3. *Derivation.* Say why $\mathbf{FB} = 0$ rules out a continuous LQ law. (The law is $\mathbf{F} = -\mathbf{B}^{\mathsf{T}}\mathbf{P}/R$, so $\mathbf{FB} = -\mathbf{B}^{\mathsf{T}}\mathbf{PB}/R$. A PI on T_c has no gain on the actuator's state, the only state $\mathbf{B}$ enters, so $\mathbf{FB} = 0$. With $\mathbf{P} \succeq 0$, $\mathbf{B}^{\mathsf{T}}\mathbf{PB} = 0$ forces $\mathbf{PB} = 0$, hence $\mathbf{F} = 0$. That is relative degree and says nothing of a sampled row: γ at 0.10/0.002 rises only to 0.9999 at 1 s, and 0.07/0.002 passes at 60 s.)

Exercise 14.4. *Derivation.* Say why $(v - \mathbf{Fx})^2$ empties the question. (It is a semidefinite stage cost, $\mathbf{Q} = \mathbf{F}^{\mathsf{T}}\mathbf{F}$, cross term $-\mathbf{F}^{\mathsf{T}}$, $R = 1$, whose value is zero on $v = \mathbf{Fx}$. No policy can do better than zero, so every stabilizing row is optimal for its own such cost, and (14.8) with a free cross term in $\mathcal{L}$ holds for every tuning.)

Exercise 14.5. *Application.* Find the reduced diagonal cost at 0.41/0.002. (On (14.17) with $a = 0.81542$, $\theta = 1.18459$ and $R = 25$, $\mathbf{P}$ has three entries. The two gain conditions and the off-diagonal of $\mathbf{P} - \mathbf{A}_c^{\mathsf{T}}\mathbf{P}\mathbf{A}_c - R\mathbf{F}^{\mathsf{T}}\mathbf{F} = \mathrm{diag}(q_d, q_z)$ are three linear equations; the diagonal then gives $q_z = 2.034 \times 10^{-4}$, which is $R\gamma K_i^2$ with $\gamma = 2.0344$, and $q_d = 8.797$, so $s_d = 0.337\,\mathrm{K}$ per step and $s_T = 2.337\,\mathrm{K}$.)

Exercise 14.6. *Application.* Find the implied event rate of the plant's tuning. (From (14.4) at 0.10/0.002, $\lambda = 0$ and $\lambda = 1$ give the columns of (14.20) by differences in the log of the gains. The quadratic has two nonnegative roots, $\lambda = 277.141$ with $s_T = 8.582\,\mathrm{K}$ and $\lambda = 0.530$ with $s_T = 0.253\,\mathrm{K}$, both with positive Hessian. By the rule the lower rate is reported: 0.530 per hour.)

Chapter 15

Saturation and Anti-Windup

15.1 The loop at its stop

The commissioning engineer has the trend open after the set point went to 17°C. The valve position reads 1.00, the discharge reading stands at 17.46°C, the loop's error is +0.46 K and is not closing, and the command sits at 1.03. On the desk lies Chapter 6's table, which prints $\Pr(T^{*} \le 27.5^{\circ}\mathrm{C}) = 0.1492$ at 60 kW and calls the node void. The load will fall back within the hour, and the vendor's sheet offers an anti-windup option.

The engineer would conclude that the loop is simply short of valve for the moment, that it will come back when the load does, and that 0.1492 is at least a rough figure for the hall air at the higher load. Nothing on the trend tests any of that. The row that defined every earlier node is the one the valve has just stopped obeying. The question of this chapter is

what the plant and the law are at the stop, what the stuck hour costs when the load falls, and what the reading says that it could not say before.

Jargon: stop, regime, margin at the stop, kink, windup, anti-windup, back-calculation, tracking time, conditional integration, read-valve row, active set, recovery, stop test, leverage of the stop

The *stop* is the bound of the actuator. A *regime* is the set of rows that holds on one side of it. The *margin at the stop* m is the distance in supply air between what the set point asks and what the fully open coil can give; the loop is at the stop when $m < 0$. The *kink* is the corner the plant's answer has at $m = 0$. *Windup* is the growth of a controller's internal state while the actuator cannot follow it, and *anti-windup* is any treatment that stops it. *Back-calculation* feeds the difference between the limited and the computed command into the integrator, at a rate set by the *tracking time* T_t. *Conditional integration* holds the integrator while the command is at a stop and the error pushes further. The *read-valve row* feeds the same integrator from the valve's reading in place of a stored number for the stop. An *active set* is the list of moves of a plan that sit on a bound. *Recovery* is the time from the load's fall until the supply air stays in a stated band. A *stop test* is a move made to put the valve on its stop on purpose, and the *leverage of the stop* is $1 - u^{*}$, the valve's room at the held condition it is compared with.

Key idea

The stop swaps one row of the graph, so the plant has two regimes and the regime is an event on the state. The law's row changes with the plan's active set, not with the valve's position, and it can change before the valve gets there. Any controller can be rewritten, by algebra alone, with a row that reads its own command. Windup is that row's command not being what the plant received; the true input with a gain large enough stops it, and the true input with no gain changes nothing. At the stop the discharge reading stops being the loop's own and becomes a row on the coil and the sensor together.

15.1.1 The boundary of the map

Table 7.6 carries this chapter as the row "the stop kept": the plant class of Chapters 12 to 14, known gain, quadratic cost, with the stop added, which ends linearity. Couplings 1, 2 and 4 of Section 7.1.3 remain in the problem (7.1); coupling 4 enters through the two limits of the window. Coupling 3 is given up: no policy here chooses a move for what it will teach, and the stop test of the last section is priced as a measurement, not chosen by a law. The law's study restricts further still and says so where it is done (Section 15.3.4). The disturbance entry of Chapter 13 and the spreads of Chapter 14 are taken as derived, as is the Riccati recursion of Chapter 12. The law studied here is not that chapter's: it charges moves and is derived in Section 15.3.4. Those chapters had one fixed row for the whole plant. This chapter is where that stops being true.

15.1.2 The closed-loop hall at its stop (step 0)

The plant is the closed-loop hall of Chapter 6, on its reduced state $\mathbf{z} = (Q,\ g,\ g_w,\ b)^{\mathsf{T}}$ of equation (6.5). Every what-if is a scenario node, closed loop while the valve is free. At the stop the actuator is pinned again. The controller's row of equation (6.4), $T_c + b = r$, leaves the graph and a pin takes its place,

$$\underbrace{r_{\mathrm{l}} = [T_c + b - r]/\sigma_p}_{\text{holding}} \qquad \text{becomes} \qquad \underbrace{r_{\mathrm{s}} = [u - 1]/(0.01\,\sigma_p)}_{\text{at the stop}}, \tag{15.1}$$

while the hall and coil rows stay. The count of entries does not change: one row out and one row in, so sparsity is unchanged and the regime is the only new thing. The coil's row now solves for the supply air and not for the valve. At a node with load Q^* and set point r, with W and G the coil's and path's conductances, the two regimes give

$$\begin{aligned} \text{holding:}\quad & T_c = r - b, \qquad T_{\mathrm{hold}} = r - b + \frac{Q^*}{G}, \qquad u = \frac{Q^*}{W(r - b - T_w)}, \\ \text{at the stop:}\quad & T_c = T_w + \frac{Q^*}{W}, \quad T_{\mathrm{sat}} = T_w + \frac{Q^*}{W} + \frac{Q^*}{G}, \quad u = 1. \end{aligned} \tag{15.2}$$

The first line is equation (6.7). Where it asks for $u > 1$ the node is what Section 6.1.2 called void, and the second line is what the plant does there. The two are joined by one function of the state, the margin

$$\begin{gathered} m(\mathbf{z}; r) = r - b - T_w - \frac{Q^*}{W} = T_{\mathrm{hold}} - T_{\mathrm{sat}}, \\ T^* = \max\bigl(T_{\mathrm{hold}},\ T_{\mathrm{sat}}\bigr), \\ u^* = \min\Bigl(1,\ \frac{Q^*}{W(r - b - T_w)}\Bigr) \quad \text{for } r - b - T_w > 0, \quad u^* = 1 \text{ otherwise.} \end{gathered} \tag{15.3}$$

The clause matters only because the fraction turns negative at $r - b - T_w \le 0$, where no valve holds the set point and $m < 0$ anyway. The loop is at the stop exactly when $m < 0$. Which

regime holds is therefore not a mode the engineer chooses but an event on $\mathbf{z}$, with a probability of its own.

Figure 15.1 draws both graphs. Each panel runs from the pinned chilled water on the left through the coil and valve to the supply air, the hall air and the load, with the controller's row and the discharge sensor above. Green is free, gray is pinned and blue is fixed by the rows; dashed arrows are readings. The two panels differ in three places: the valve goes from free to pinned, the controller's row goes out of the graph, and the discharge reading goes from a line with no entry to a row on the coil, the load and the offset.

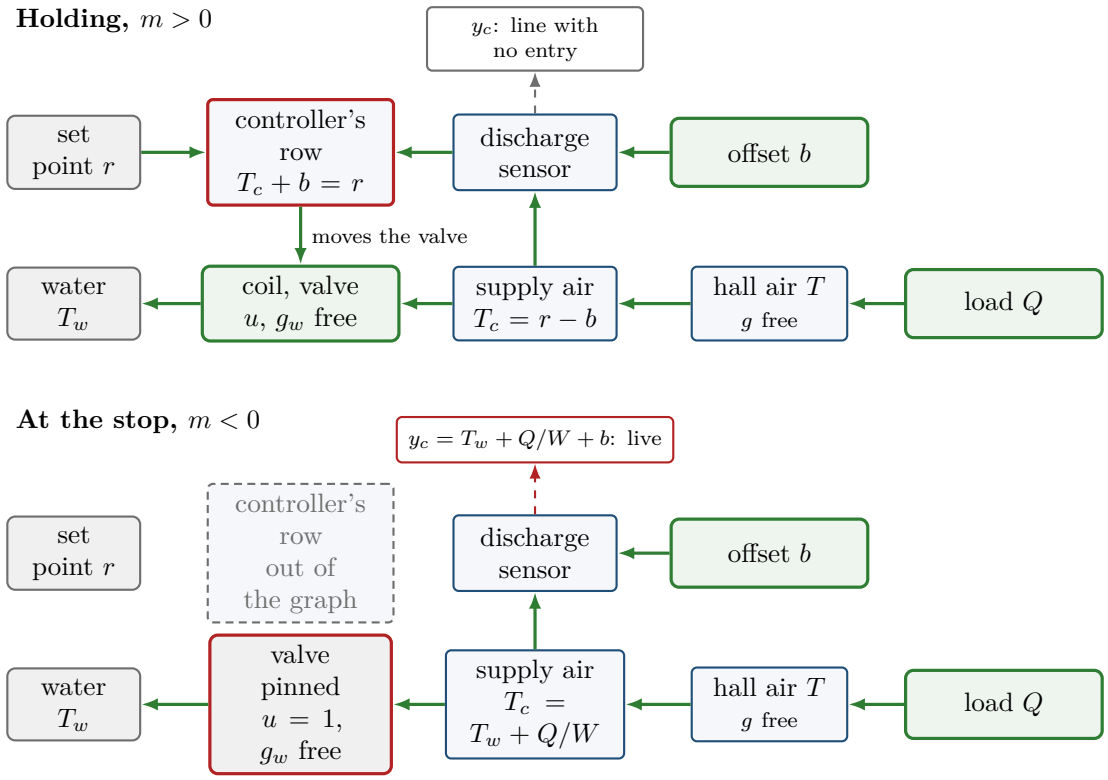


Figure 15.1: The two graphs of one hall, equation (15.1). Above: the set point is pinned, the valve is free, and the controller's row (red) consumes the discharge reading. Below: the valve is pinned, the controller's row has left, the supply air floats, and the discharge reading is a row on the coil, the load and the offset. The hall, valve and meter readings and the priors are rows not drawn.

15.1.3 The control engineer's questions (step 1)

Table 15.1 lists the seven questions. The first column is the question in the engineer's words, the second the general form, the third its mathematical form on this hall, the fourth its category, the fifth the computation and the last the subsection. The first three read the steady plant across the kink; the next three are about the law and its memory; the last asks what the stop reveals.

15.1.4 The evidence as rows (step 2)

The belief at a node is Chapter 6's fit of the twenty minutes at about 30 kW, its six rows unchanged. What is new is which readings carry entries in each regime, and the choices the two regimes rest on.

Table 15.1: The questions asked of the hall at its stop. T^*, u^* and m follow equation (15.3); $\mathbf{y}$ is Chapter 6's evidence unless a row says otherwise.

Control engineer's question	General form	Mathematical form	Category	Computation	Answered in
Where is the stop, and with what probability is the loop at it?	threshold, (2.12)	$\Pr(m < 0 \mid \mathbf{y})$	prediction under a threshold	posterior	15.3.1
What is the hall air across the kink?	moments of a mixture, (2.10)	$p(T^* \mid \mathbf{y})$, two regimes	prediction under a threshold	posterior	15.3.2
Does Chapter 6's window survive?	joint chance constraint, (6.20)	$\Pr(T^* \le L,\ u^* < 1 \mid \mathbf{y})$	prediction under a threshold	posterior	15.3.3
What is the law at the stop?	expected consequence, the stop kept, (15.13)	bounded least squares against the clipped law	design	expected consequence	15.3.4
Why does the integrator wind, and what stops it?	a rewriting of the controller, (15.16)	the row on c or on $\mathrm{sat}(c)$, with its gain	design	algebra	15.3.5
What does each treatment cost on recovery?	expected consequence of fixed laws, (2.23)	recovery, fall below $r - b$	design	expected consequence	15.3.6
What does the stop reveal?	moments, (2.10)	$p(b, g_w \mid$ rows with y_c live$)$	estimation	posterior	15.3.7

Which row acts on which entry. Chapter 5's zero row, the discharge line of equation (5.8), holds only while the controller's row does. From equation (15.2), with $*$ for an entry that can be nonzero and columns Q, g, g_w, b,

$$\mathbf{J}_{\text{hold}} = \begin{bmatrix} * & * & 0 & * \\ * & 0 & * & * \\ 0 & 0 & 0 & 0 \\ * & 0 & 0 & 0 \end{bmatrix} \begin{matrix} y_T \\ y_u \\ y_c \\ y_Q \end{matrix}, \qquad \mathbf{J}_{\text{stop}} = \begin{bmatrix} * & * & * & 0 \\ 0 & 0 & 0 & 0 \\ * & 0 & * & * \\ * & 0 & 0 & 0 \end{bmatrix} \begin{matrix} y_T \\ y_u \\ y_c \\ y_Q \end{matrix}. \tag{15.4}$$

Three facts are read here. The zero row moves: at the stop the discharge line opens and the valve's line closes, since $u = 1$ is a pin and its reading only confirms the regime. The hall reading loses the offset's column at the stop and gains the coil's, so the floating air does not depend on the sensor. And no entry is added: the stop reveals by changing the columns, not the state.

The rows of a stop. A window at the stop is compared with a held window, each at its own load, so the state of that fit is $(g,\ g_w,\ b,\ Q_1,\ Q_2)^{\mathsf{T}}$. Section 15.2.2 writes the nine rows. The pattern in those columns is equation (15.4) applied to each window, with the priors of Chapter 6 on g, g_w and b.

The modeling choices. Five enter here, each tested later.

- *Steady regimes at a node.* Equation (15.2) is a steady statement. It buys a closed form for every node of Sections 15.3.1 to 15.3.3, and it says nothing of how the loop arrives; the simulations of Section 15.3.6 carry the time.
- *The water pinned.* T_w is held at 14°C. This is what lets the stop's row name the offset, and Section 15.3.7 releases it, since at the stop the discharge line reads T_w with the same coefficient as b.

- *Coil duty linear in the valve.* $W u\,(T_c - T_w)$ with W one constant is what makes a fully open valve comparable with a partly open one. A real valve's characteristic would break the comparison, and the constructed record cannot test it.
- *The load pinned.* The node pins Q^*. Section 15.3.2 unpins it once as a stated what-if; the forecast belongs to Chapter 19.
- *Stated costs and tunings.* The cost spreads are Chapter 7's, $q = 4$ and $\rho = 25$, with ρ at 1 and 400 beside them. Back-calculation has $T_t = T_i$. The read-valve row's gain and the estimator's walk follow stated rules, given in Section 15.3.6.

15.2 The posterior of the state (step 3)

Two posteriors serve the answers. The first is the belief every node reads, and it is not new. The second adds rows the stop has made live.

15.2.1 The belief at a node

The pair is that of Section 6.2.2, with the curvature term kept: on (Q, g, g_w, b),

$$\hat{\mathbf{z}} = \begin{bmatrix} 30.05809 \\ 0.01833 \\ -0.16764 \\ 0.05708 \end{bmatrix}, \qquad \boldsymbol{\Sigma} = \begin{bmatrix} 0.02496 & 0.00088 & 0.00076 & -0.00027 \\ 0.00088 & 0.00177 & -0.00183 & -0.00732 \\ 0.00076 & -0.00183 & 0.00235 & 0.00913 \\ -0.00027 & -0.00732 & 0.00913 & 0.03605 \end{bmatrix}, \tag{15.5}$$

in kW, K and mixed units, with $\mathrm{corr}(g_w, b) = 0.99242$ and spreads 0.15798, 0.0421, 0.04846 and 0.18986. At the mean, $\hat{W} = 13.5305\,\mathrm{kW/K}$ and $\hat{G} = 6.1110\,\mathrm{kW/K}$; the stated truth has 14.4774 and 5.8000. The exact posterior is carried by Chapter 6's reweighted draws: a t_5 proposal about $\hat{\mathbf{z}}$, seed 6, $N = 2 \times 10^6$, effective size 1 209 705. That correlation will decide the first three answers, because b and g_w are the two entries in the margin.

15.2.2 The rows of a window at the stop

A stop test holds 30 kW throughout, with $r = 18°\mathrm{C}$ to step 159 and $r = 16°\mathrm{C}$ from step 160. The held window is steps 120 to 159 and the stop window steps 300 to 339, so the hall has had 140 steps to settle. This is a stated move that the record does not hold; its records are fresh ones from the stated truth. With window means over the forty steps, each row divided by its sensor's spread over those steps, the residual vector is

$$\mathbf{r}(\mathbf{z}) = \begin{bmatrix} (r_1 - b + Q_1/G - y_{T1})/\sigma_T \\ (Q_1/(W(r_1 - b - T_w)) - y_{u1})/\sigma_u \\ (Q_1 - y_{Q1})/\sigma_Q \\ (T_w + Q_2/W + Q_2/G - y_{T2})/\sigma_T \\ (T_w + Q_2/W + b - y_{c2})/\sigma_c \\ (Q_2 - y_{Q2})/\sigma_Q \\ (g - 0.22314)/0.30 \\ g_w/0.15 \\ b/0.20 \end{bmatrix}, \qquad \mathcal{L} = \tfrac{1}{2}\,\mathbf{r}^{\mathsf{T}}\mathbf{r}, \tag{15.6}$$

on $\mathbf{z} = (g,\ g_w,\ b,\ Q_1,\ Q_2)^{\mathsf{T}}$, with $r_1 = 18°\mathrm{C}$. Each iteration solves the normal equations $(\mathbf{J}^{\mathsf{T}}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^{\mathsf{T}}\mathbf{r}$; the information matrix of these normalized rows is $\mathbf{J}^{\mathsf{T}}\mathbf{J}$, and with unnormalized rows it is $\mathbf{J}^{\mathsf{T}}\mathbf{W}\mathbf{J}$. Table 15.2 runs them on the first fresh record of the test, whose window means are $y_{T1} = 22.8272°\mathrm{C}$, $y_{u1} = 0.5659$, $y_{Q1} = 29.6825\,\mathrm{kW}$, $y_{T2} = 21.3567°\mathrm{C}$, $y_{c2} = 16.3837°\mathrm{C}$

Table 15.2: Gauss–Newton iterations of the rows (15.6), first fresh record of the stop test, seed 15500. Iterations 4 and 5 print as iteration 6 in these columns but for the gradient.

k [–]	g [–]	g_w [–]	b [K]	Q_1 [kW]	Q_2 [kW]	$\mathcal{L}$ [–]	$\lVert\mathbf{J}^{\mathsf{T}}\mathbf{r}\rVert$ [–]
0	0.22314	0.00000	0.00000	29.6825	30.3500	2779.6899	2.40×10^{4}
1	−0.03283	−0.18807	0.12906	29.6768	30.3648	38.7742	3.39×10^{3}
2	−0.01058	−0.15998	0.15612	29.6679	30.3722	1.2267	4.35×10^{1}
3	−0.01026	−0.15976	0.15584	29.6687	30.3714	1.2204	7.32×10^{-3}
6	−0.01026	−0.15976	0.15584	29.6687	30.3714	1.2204	6.35×10^{-9}

and $y_{Q2} = 30.3500\,\text{kW}$. Each row of the table is an iteration; the columns give the five entries, the objective and the norm of its gradient.

At the MAP state the Jacobian and the Gauss–Newton covariance are

$$\mathbf{J} = \begin{bmatrix} -63.1924 & 0 & -12.6491 & 2.1299 & 0 \\ 0 & -357.9239 & 93.1085 & 12.0640 & 0 \\ 0 & 0 & 0 & 6.3246 & 0 \\ -64.6891 & -28.1701 & 0 & 0 & 3.0575 \\ 0 & -140.8505 & 63.2456 & 0 & 4.6376 \\ 0 & 0 & 0 & 0 & 6.3246 \\ 3.3333 & 0 & 0 & 0 & 0 \\ 0 & 6.6667 & 0 & 0 & 0 \\ 0 & 0 & 5 & 0 & 0 \end{bmatrix}, \tag{15.7}$$

$$\mathbf{\Sigma} = \begin{bmatrix} 0.000228 & -0.000152 & -0.000512 & -0.000347 & 0.001166 \\ -0.000152 & 0.000315 & 0.000914 & 0.001776 & -0.000921 \\ -0.000512 & 0.000914 & 0.002870 & 0.003837 & -0.003777 \\ -0.000347 & 0.001776 & 0.003837 & 0.023222 & 0.001675 \\ 0.001166 & -0.000921 & -0.003777 & 0.001675 & 0.023387 \end{bmatrix}. \tag{15.8}$$

The new line is the fifth of equation (15.7): it is the line that was all zeros while the loop held. In the column of b its entry 63.2456 is second among the reading rows, after the valve's 93.1085. The valve's line at the held window has the ratio 93.1085/357.9239 between b and g_w, and the new line has 63.2456/140.8505. Two lines at different ratios separate two entries, and this is the whole of what the stop adds. Here $\hat{g}_w = -0.15976$ with spread 0.01774, $\hat{b} = 0.15584$ with spread 0.05357, and $\text{corr}(g_w, b)$ falls from 0.99242 to 0.9610. These values are one record's draw; Section 15.3.7 takes the test over two hundred records, and the first record stands low on b.

15.2.3 The stacked rows of the record

The record has a stop of its own, unsettled, at steps 162 onward: 76 of the 200 steps from 40 to 239 have a command at or above one, and the true valve is at one on 68. Section 10.3.6 stacked those steps in time, with the commands known. Two fits on the same stacked rows are needed here: all rows kept, and the discharge rows at the stop's steps dropped. Gauss–Newton from Chapter 6's fit, eight iterates, seed 51, gives

$$\text{kept: } \hat{g}_w = -0.08807\ (0.02475), \quad \hat{b} = 0.34532\ (0.08539);$$
$$\text{dropped: } \hat{g}_w = -0.16217\ (0.03082), \quad \hat{b} = 0.03821\ (0.12429),$$

with $\text{corr}(g_w, b)$ of 0.99251 and 0.99252. The last step of each is below 10^{-6}.

15.3 The answers (step 4)

Each node's answer reads the pair (15.5) by a named route, and the draws of Section 15.2.1 stand beside it. The law and the recovery answers use fixed laws scored on the plant, and the last answer reads the posteriors of Sections 15.2.2 and 15.2.3.

15.3.1 The kink and the event of the stop

Control engineer's question. Where is the stop, and with what probability is the loop at it?

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

Specialized form. Equation (2.12) with $f = -m$ of equation (15.3), $f^* = 0$ and the Laplace normal of step 3. The margin's gradient on (Q, g, g_w, b), the load being pinned at the node, is

$$\mathbf{g}_m = \mathbf{g}_{T_{\text{hold}}} - \mathbf{g}_{T_{\text{sat}}} = \left(0,\ 0,\ Q^*/W,\ -1\right)^{\mathsf{T}}, \qquad \Pr(m < 0 \mid \mathbf{y}) \approx \Phi\left(\frac{-m(\hat{\mathbf{z}}; r)}{\sqrt{\mathbf{g}_m^{\mathsf{T}} \boldsymbol{\Sigma}\, \mathbf{g}_m}}\right). \tag{15.9}$$

Chapter 6 read the same event as a normal on u; equation (15.9) is a normal on m, which is what the plant's corner is in. Both are set beside the direct route: the reweighted draws pushed through m itself.

Building blocks. $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$ of equation (15.5); $\hat{W}$; the gradient in equation (15.9); Φ.

The plant's numbers. Differentiating equation (15.2) gives the two regimes' slopes. Holding, $dT^*/dr = 1$ and $dT^*/dQ = 1/G = 0.1636\,\text{K/kW}$ at the fit's mean. At the stop, $dT^*/dr = 0$ and $dT^*/dQ = 1/G + 1/W = 0.2375\,\text{K/kW}$; the truth's are 0.1724 and 0.2415. At $r = 18°\text{C}$ the kink is at 53.350 kW on the mean and 53.566 kW in the truth; at 45 kW it is at $r = 17.3829°\text{C}$ and 17.4083°C. Figure 15.2 draws this. The left column has the set point on its horizontal axis at 45 kW, the right the load at $r = 18°\text{C}$. The upper panels give the hall air and the supply air, the lower panels the valve; solid curves are the fit's mean and dotted curves the stated truth.

The left column shows why there is a floor: below the corner the set point does nothing at all. The right shows a worse fact for a load request: past the corner each further kilowatt costs more temperature, 0.2375 against 0.1636 K/kW, because the supply air now floats as well.

Table 15.3 gives the event. Each row is a node; the columns give the margin's mean and spread, the probability by the normal on u, by the normal on m and direct with its standard error, and the truth's regime. The spread of m is 0.0359 K at 45 kW and 0.0248 K at 54.99 kW, far below the 0.1899 K of the offset alone. That is the valve reading at work. By equation (15.9) the margin reads b and g_w in a combination, and $\text{corr}(g_w, b) = 0.99242$ runs along it. So the record does not know the offset, and it knows the stop's position well. The three routes agree to about 0.01: at $r = 17.38°\text{C}$, 0.0029 K below the mean's kink at 45 kW, they give 0.5321, 0.5322 and 0.5231 (se 0.0005), the last differing from the first two by more than sampling. At the kink itself, $r = 17.3829°\text{C}$, the normal on m gives 0.5000 by construction and the direct route 0.4911 (se 0.0005), with seed 15 at 0.4905.

15.3.2 The hall air across the kink

Control engineer's question. What is the hall air across the kink?

General form. Moments and a tail of a derived quantity under the posterior, equation (2.10), with the threshold form (2.12) for the tail.

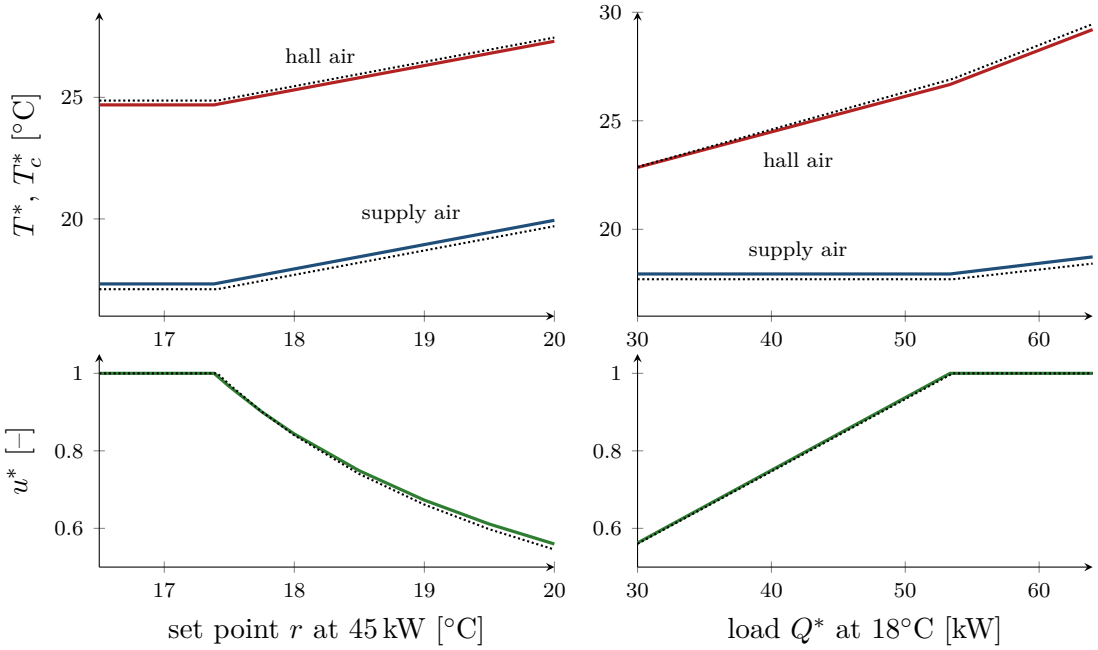


Figure 15.2: The kink of equation (15.3), closed loop until the valve pins. Solid: the fit's mean. Dotted: the stated truth. Left, against the set point at 45 kW: the hall air stops falling at the corner. Right, against the load at 18°C: the hall air's slope steepens at the corner.

Table 15.3: The event of the stop, $\Pr(m < 0 \mid \mathbf{y})$, by three routes. Draws with seed 6, effective size 1 209 705.

Q^* [kW]	r [°C]	m, mean [K]	sd m [K]	Normal on u [–] / on m [–]	Direct (se) [–]	Truth
45	17.32	−0.0629	0.0359	0.9466 / 0.9602	0.9545 (0.0003)	stop
45	17.38	−0.0029	0.0359	0.5321 / 0.5322	0.5231 (0.0005)	stop
45	17.40	+0.0171	0.0359	0.3130 / 0.3169	0.3076 (0.0004)	stop
45	17.44	+0.0571	0.0359	0.0430 / 0.0558	0.0495 (0.0003)	hold
54.99	18.10	−0.0212	0.0248	0.8056 / 0.8037	0.8058 (0.0004)	hold
54.99	18.19	+0.0688	0.0248	0.0038 / 0.0028	0.0041 (0.0000)	hold
60	18.50	+0.0085	0.0353	0.4055 / 0.4049	0.4098 (0.0004)	hold

Specialized form. Equation (2.10) with $f = T^* = \max(T_{\text{hold}}, T_{\text{sat}})$ of equation (15.3), read by four named routes.

- *hold*: Chapter 6's state route on the holding regime alone, a normal with gradient $\mathbf{g}_{T_{\text{hold}}}$.
- *cdiff*: the delta method (4.53) through the kinked function, with its gradient by a central difference of step h; a single normal.
- *two-regime*: the Laplace normal with both regimes linear, so that T_{hold} and T_{sat} are jointly normal. With means μ_1, μ_2, spreads s_1, s_2, $\theta = \text{sd}(T_{\text{hold}} - T_{\text{sat}}) = \text{sd}(m)$ and $a = (\mu_1 - \mu_2)/\theta$, the standard moments of a maximum of two normals are

$$\begin{aligned} \mathbb{E}[T^*] &= \mu_1\Phi(a) + \mu_2\Phi(-a) + \theta\varphi(a), \\ \mathbb{E}[T^{*2}] &= (\mu_1^2 + s_1^2)\Phi(a) + (\mu_2^2 + s_2^2)\Phi(-a) + (\mu_1 + \mu_2)\theta\varphi(a), \end{aligned} \tag{15.10}$$

  and $\Pr(T^* \le L) = \Pr(T_{\text{hold}} \le L,\ T_{\text{sat}} \le L)$, the integral (6.22) with these two quantities in place of T^* and u^*.
- *direct*: the reweighted draws through the kinked function. This is the reference for the other three.

The two-regime route assumes normality of the state, not of T^*; it is an approximation like the others, and it is tested below where the state is not close to normal.

Building blocks. $\hat{\mathbf{z}}$, $\boldsymbol{\Sigma}$; $\mathbf{g}_{T_{\text{hold}}} = (0,\ -Q^*/G,\ 0,\ -1)^{\mathsf{T}}$ and $\mathbf{g}_{T_{\text{sat}}} = (0,\ -Q^*/G,\ -Q^*/W,\ 0)^{\mathsf{T}}$; equation (15.10); the limit $L = 27.5$°C.

The plant's numbers. Table 15.4 gives the routes at the nodes that test them. Each row is a route at a node; the columns give the mean, spread, 95th percentile and probability of the limit. Figure 15.3 draws the density at the mean's kink: the horizontal axis is the hall air and the vertical axis its density, the markers the draws in bins, the solid curve the two-regime route, the dashed curve hold and the dotted curve cdiff. The left panel is Chapter 6's posterior at 45 kW, the right the what-if with the valve's rows dropped at 54 kW.

Table 15.4: The hall air by route across the kink, $L = 27.5$°C. Direct: seed 6, effective size 1 209 705; a printed 1.0000 is above 0.9999. Cdiff at $h = 10^{-6}$. Truth's T^*: 24.867, 24.896 and 27.370°C at the three nodes.

Node	Route	Mean T^* [°C]	sd [K]	q95 [°C]	$\Pr(T^* \le L)$ [–]
45 kW, kink, r 17.383°C	direct	24.6975	0.1677	24.9802	1.0000
	two-regime	24.7039	0.1681	24.9804	1.0000
	hold	24.6896	0.1559	24.9461	1.0000
	cdiff	24.6896	0.1677	24.9655	1.0000
45 kW, window's edge, r 17.438°C	direct	24.7395	0.1573	24.9998	1.0000
	hold	24.7445	0.1559	25.0010	1.0000
54.99 kW, best r 18.189°C	direct	27.1218	0.2186	27.4818	0.9580 (se 0.0003)
	two-regime	27.1307	0.2186	27.4903	0.9544
	hold	27.1307	0.2186	27.4902	0.9544
	cdiff	27.1307	0.2186	27.4902	0.9544

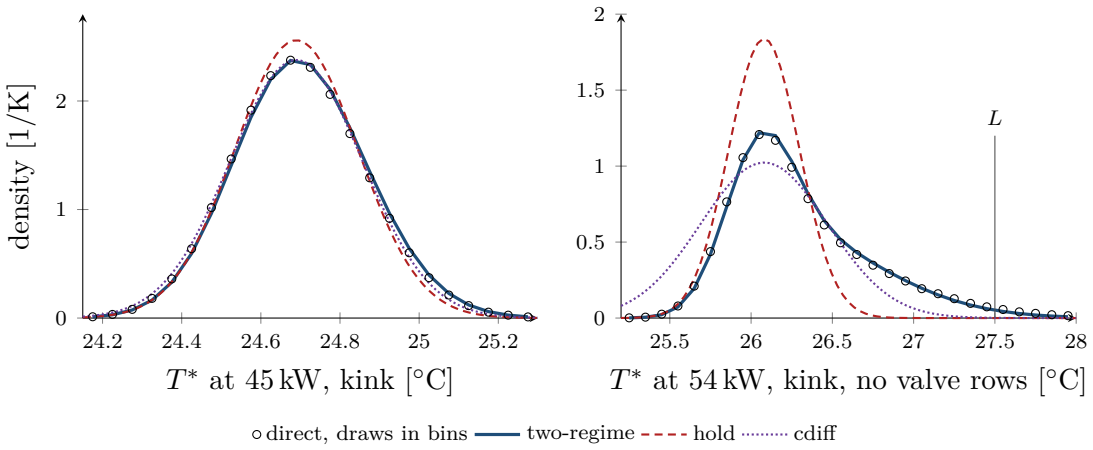


Figure 15.3: The density of the hall air at the mean's kink, $m = 0$ at $\hat{\mathbf{z}}$. Dashed and dotted are normals with the means and spreads printed for the routes: left 24.6896 with 0.1559 and 0.1677 K, right 26.080 with 0.217 and 0.390 K. Left: Chapter 6's posterior, one hump, no visible corner. Right: the valve's rows dropped, a stated what-if with 500 000 draws, seed 15, effective size 429 274; the regimes separate and the tail is the second regime's.

The surprise. A single-normal band might be expected to mislead across the kink. **On Chapter 6's posterior, at or above the mean's corner, it barely does.** There the largest gap at a 95th percentile is 0.0341 K, hold against direct at the kink at 45 kW, against a back-off of 0.2565 K there and 0.3595 K at 54.99 kW. Further into that tail at the same corner the gap is 0.0517 K at the 99th percentile and 0.0700 K at the 99.9th, 14.5 % of its back-off, and the largest gap between the two distribution functions is 0.0286; every probability of the limit at 45 kW is 1.0000, so that node tests the band's width and not a decision. Below the corner hold has no standing, since it reads a regime the plant has left: its gap is 0.1666 K at $r = 17.25$°C, 45 kW, and 0.4652 K at $r = 18$°C, 60 kW, and the void row of Table 15.5 is the same gap read as a probability. With the limit near the corner, 55.3 kW and 0.05 K below it, hold gives 0.9696 and direct 0.9532 (se 0.0003). At that load's kink, two-regime is the worse of the two normals: 0.9452 against hold's 0.9504, direct 0.9494 (se 0.0003). At the window's edge the gap is 0.0011 K. At closure's set point the direct probability stands 0.0035 above the two-regime one, and two-regime and hold agree to 0.0000, so that gap is Chapter 6's non-normality and none of it is the kink. The left panel of Figure 15.3 shows the reason. The corner lifts the mean of a maximum by at most $\theta\,\varphi(0)$ in equation (15.10), and θ is the spread of the margin: 0.0143, 0.0099 and 0.0141 K at 45, 54.99 and 60 kW. Since $\operatorname{corr}(g_w, b) = 0.99242$ the two regimes read nearly one direction of the state, so the two normals nearly coincide where they meet.

Where the central difference does lie. This is not a licence for the central difference. Its answer depends on which side of the corner it stands. One millikelvin of set point either side of the kink at 45 kW gives $\operatorname{sd}(T^*)$ of 0.1559 against 0.1805 K and $\operatorname{sd}(u^*)$ of 0.0108 against 0.0000. At the kink exactly, its gradient $(0,\ -7.364,\ -1.663,\ -0.5)^{\mathsf{T}}$ is half of each regime's in the columns where they differ, an average of two rows no plant has. Its 0.1677 K sits beside direct for a reason that can be stated. At $m = 0$ equation (15.10) gives $\operatorname{Var}(T^*) = (s_1^2 + s_2^2)/2 - \theta^2/(2\pi)$, and the averaged gradient gives $(s_1^2 + s_2^2)/2 - \theta^2/4$, so

$$\operatorname{Var}(T^*) = \operatorname{Var}_{\text{cdiff}} + \Big(1 - \frac{2}{\pi}\Big)\frac{\theta^2}{4} \qquad \text{at } m = 0. \tag{15.11}$$

The formula gives 0.1681, 0.3443 and 0.4230 K at the three kinks of this subsection, from cdiff's 0.1677, 0.3143 and 0.3899 K with θ of 0.0359, 0.4662 and 0.5438 K, and two-regime prints the

same three values. So the match is small θ at exactly $m = 0$, and it goes as θ grows or the node moves off the corner. On the valve there is no such match: cdiff gives mean 1.0000 with spread 0.0054 at the kink, and direct 0.9958 with 0.0060.

What would have made it larger. The band of the margin is what matters, and two stated what-ifs widen it.

- *The valve's rows dropped.* Hall and load meter alone at the same window, g_w on its prior, 54 kW at the kink. Now $\theta = 0.5438$ K. The 95th percentile is 27.200°C direct and 26.721°C by cdiff; hold has mean 26.080°C and spread 0.217 K. The probability of the limit is 0.9798 direct, 0.9872 two-regime, 0.9999 cdiff and 1.0000 hold. The mixture is real here, as the right panel shows, and two-regime is not exact either.
- *The load not pinned.* $Q^* \sim \mathcal{N}(53.350,\ s^2)$, independent of the belief, at $r = 18$°C, with $s = 2$ kW; the forecast itself is Chapter 19's. The probability is 0.9446 direct (se 0.0002), 0.9660 cdiff and 0.9835 hold, **on opposite sides of** 0.95. The 95th percentiles are 27.5269, 27.4184 and 27.3109°C, with skew +0.370. At $s = 0.5$ kW the three are 0.9996, 0.9998 and 0.9999, and at 5 kW they are 0.7559, 0.7901 and 0.8363.

The valve reading narrows the kink, and a load band opens it. Here the single normal does mislead.

15.3.3 The window across the kink

Control engineer's question. Does Chapter 6's window survive?

General form. The joint chance constraint (6.20), the threshold form (2.12) with the indicator of an intersection.

Specialized form. Equation (6.20) with $f_1 = T^*$ and $f_2 = u^*$ of the kinked plant. On $m > 0$ equation (15.3) gives $T^* = T_{\text{hold}}$ and $u^* = u_{\text{hold}} < 1$, and $u_{\text{hold}} > 1$ is $m < 0$ wherever $r - b - T_w > 0$. So, the event $m = 0$ having no probability,

$$\{T_{\text{hold}} \le L,\ u_{\text{hold}} \le 1\} = \{T^* \le L,\ u^* < 1\}. \tag{15.12}$$

The joint that Chapter 6 computed on the holding regime is therefore already the joint of the kinked plant. The hall air alone is another matter: $\Pr(T^* \le L)$ now has the second regime in it, and below the floor it stops at $\Pr(T_{\text{sat}} \le L)$. That tail is equation (2.12) on T_{sat}, whose closure is the load at which it reaches 0.95 with the valve released from protection.

Building blocks. The identity (15.12); the draws; the means and spreads of T_{sat} from $\mathbf{g}_{T_{\text{sat}}}$ and $\mathbf{\Sigma}$.

The plant's numbers. Chapter 6's joint treatment survives by identity, not by approximation: the two sides of equation (15.12) on the draws, with $u_{\text{hold}} \le 1$ on one side and $u^* < 1$ on the other and no margin in the test, differ by 0.0 over 8 nodes, 2 of them below the floor at the mean. Its normal on u does remain an approximation of the same event. The exact window at 45 kW is 17.440 to 19.942°C, against the state route's 17.438 to 19.937°C. The state route's closure is 54.99 kW. The exact joint's best values at 54.8, 55.0 and 55.2 kW are 0.9702, 0.9526 and 0.9276, on a set-point grid of 0.005 K, unchanged at 0.0025 K, which interpolates to 55.02 kW.

Table 15.5: $\Pr(T^* \le 27.5°\text{C} \mid \mathbf{y})$ past the floor, replacing Chapter 6's void reading.

Node	Chapter 6's void reading [–]	Direct (se) [–]	Truth T^* [°C]
60 kW, $r = 18°$C	0.1492	0.0010 (0.0000)	28.489
56 kW, $r = 18°$C	0.9597	0.8207 (0.0004)	27.523
54.99 kW, $r = 17.5°$C	1.0000	0.9781 (0.0002)	27.279

The void readings are another matter. Table 15.5 replaces them. Each row is a node $u^* > 1$ at the mean; the columns give Chapter 6's reading, the direct probability of the kinked plant with its standard error, and the truth's hall air.

So 0.1492 was not a rough figure for anything: the plant's value is 0.0010. The 56 kW row is the warning, because there a void reading stands above 0.95 for a load whose true hall air breaks the limit. Below the floor T^* stops falling at T_{sat}, so the hall air alone has no lower edge. Its own closure, by the state route on T_{sat}, is 55.30 kW and direct 55.33 kW, with the truth's limit at 55.904 kW, which is also the load at which the truth's window closes. **Protecting the valve costs** 0.30 kW **of load** by the state routes and 0.31 kW direct. At 56 kW the hall air alone holds at 0.8207 and the joint at its best set point at 0.7182, and neither passes. Chapter 6's choice of the valve as protected stands as a choice, and this is its price.

15.3.4 The law at the stop

Control engineer's question. What is the law at the stop?

General form. Equation (7.1) with the stop as a limit that holds with certainty, the expected consequence (2.23) over move sequences,

$$\min_{c_0, c_1, \ldots} \sum_k \ell(\mathbf{z}_k, c_k) \quad \text{subject to} \quad F(\mathbf{z}_{k+1}, \mathbf{z}_k, c_k) = \mathbf{0}, \qquad 0 \le c_k \le 1. \tag{15.13}$$

Specialized form. Equation (15.13) under a further restriction: known plant at the truth's values, known load, the state read and no noise, so that the expectation is one trajectory and everything is restricted but the stop. The cost is $J = \sum_k q\,(T_{c,k} - (r - b))^2 + \rho\,(c_k - c_{k-1})^2$ on the plant, over 120 steps. The linear model and target are the post-step steady point, with u a virtual value above 1 where needed.

This is not the law of Chapter 12. Equation (12.3) charges ρv^2 on the distance from the pinned point, on three states with no memory, and its gains are $(0.11829,\ 0.03421,\ -0.51268)$. A valve move has no target of its own here, so the law used in this chapter is a move-penalized variant, derived as follows. With $\mathbf{x}$ the deviation of $(T,\ T_c,\ u)^{\mathsf{T}}$ from the target and $\bar{c}$ the deviation of the last command, the entry $\mathbf{z} = (\mathbf{x}^{\mathsf{T}},\ \bar{c})^{\mathsf{T}}$ obeys $\mathbf{z}^+ = \mathbf{A}_z\mathbf{z} + \mathbf{B}_z\,(c_k - c_{k-1})$, with $\mathbf{A}_z$ holding $\mathbf{A}$, $\mathbf{B}$ and a last row $(\mathbf{0}^{\mathsf{T}},\ 1)$, and $\mathbf{B}_z = (\mathbf{B}^{\mathsf{T}},\ 1)^{\mathsf{T}}$. The cost is then quadratic in $\mathbf{z}$ and the move, so the Riccati recursion of Chapter 12 applies to these four states and gives

$$c_k = c_{k-1} - \mathbf{K}\,\mathbf{z}_k, \qquad \mathbf{K} = \frac{\mathbf{B}_z^{\mathsf{T}}\mathbf{P}\mathbf{A}_z}{\rho + \mathbf{B}_z^{\mathsf{T}}\mathbf{P}\mathbf{B}_z}. \tag{15.14}$$

Four routes are named.

- *clip*: the move (15.14), clipped to $[0, 1]$, its memory the clipped command.
- *wound*: the same law, its memory the unclipped command.

- *QP(N)*: the receding problem of N moves with the law's Riccati matrix as its terminal cost. The moves are the entries, as in equation (7.17), the rows are linear in them, and bounds make it bounded least squares,

$$\min_{\mathbf{0}\le\mathbf{c}\le\mathbf{1}} \ \tfrac{1}{2}\left\|\mathbf{M}\,\mathbf{c}+\mathbf{M}_x\,\mathbf{x}_0+\mathbf{m}_c\,c_{-1}\right\|^2. \tag{15.15}$$

- *bound*: equation (15.13) itself over the whole horizon on the plant, by bounded Gauss–Newton. It is the value a law's excess is measured against. The plant is bilinear and the solver local, so strictly this is the smallest cost found, and it is checked below every law; a lower optimum cannot be excluded, and then every excess below would be a lower bound on the true one.

For a fixed active set $\mathcal{A}$ of equation (15.15), the free entries solve the normal equations of the columns left, $(\mathbf{M}_{\mathcal{F}}^{\mathsf{T}}\mathbf{M}_{\mathcal{F}})\mathbf{c}_{\mathcal{F}} = -\mathbf{M}_{\mathcal{F}}^{\mathsf{T}}(\cdots)$, so that c_0 is affine in $\mathbf{x}_0$ with gains that depend on $\mathcal{A}$. On the linear model the law is piecewise affine, with a new row for each active set, and that is all the statement covers: the plant is bilinear. The regions follow the plan's active set and not the valve's position. A later move of the plan can sit on the bound while the first is still free, and then the row has changed with the valve below its stop; with the first move on the bound the row is the constant 1. The receding plan's own theory is Chapter 18's and is not developed here.

Building blocks. The move (15.14); the rows of equation (15.15); $q = 4$ and $\rho = 25$; the plant integrated at its fine step for every route.

The plant's numbers. Table 15.6 gives the study. Each row is a scenario; the columns give the bound's cost at $\rho = 25$, the excess over it of clip and of QP(20), the steps at a stop for clip and QP(20), and clip's excess at the other two values of ρ. Figure 15.4 draws two parts of it: the left panels are the command and the supply air on the load step past the stop, by clip and QP(20); the right panel has on its horizontal axis the set-point step e the loop starts from at $r = 17.5$°C, 45 kW, and on its vertical axis the first move of each route.

Table 15.6: The law at the stop, known plant, known load, no noise, 120 steps. Excess over the smallest cost found, in percent of it; these percentages are this example's. The bound lay below every law, by no more than 8.1×10^{-9} the wrong way. Bound's optimality measures 3.5×10^{-11} to 6.4×10^{-6}.

Scenario	J bound [–]	Clip [%]	QP(20) [%]	Stop steps [–]	Clip, ρ 1 [%]	Clip, ρ 400 [%]
r 18 → 17°C, past the stop	107.40080	0.00	0.00	120, 120	0.00	0.06
r 18 → 17.5°C, stop in transit	2.88870	0.00	0.13	18, 18	0.00	0.22
r 18 → 21°C, lower stop	86.43371	0.01	0.01	3, 3	0.02	1.21
Load 30 → 53 kW, no stop touched	0.13823	0.59	0.59	0, 0	0.26	0.75
Load 30 → 60 kW, past the stop	36.59155	0.72	0.54	90, 91	0.71	0.88

Clipping is near-optimal on this plant for this cost and these steps, by computation: the cost on the fast supply air, the step known at once. Its largest excess that can be laid to the stop is 0.88 %, at $\rho = 400$ with the load past the stop. The 1.21 % at $\rho = 400$ on the third row touches no stop, for 0 steps at a stop. The baseline has to be read first: 0.59 % arises on the fourth row, with no stop touched, from a linear law on a bilinear plant, and it is not the stop's. Only the load past the stop has an excess above that baseline, and its absolute size there is 0.26361 of cost, 0.37 of one step's final-offset cost of 0.71113.

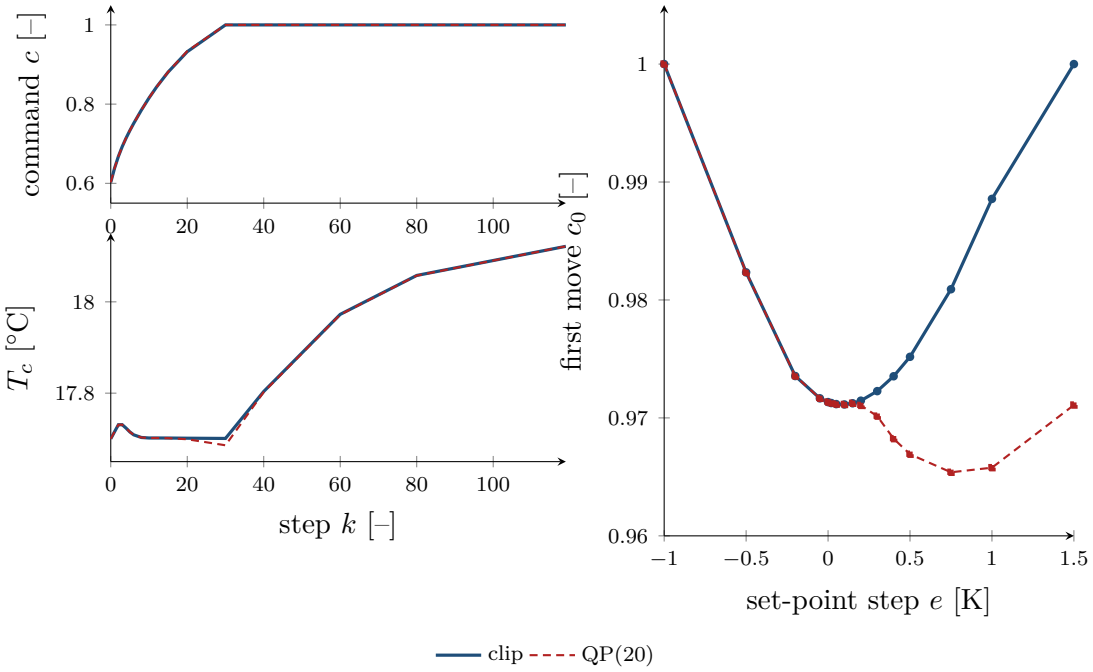


Figure 15.4: Clip against QP(20), $\rho = 25$. Left: the load step 30 → 60 kW at $r = 18$°C, target 17.700°C, straight lines between printed steps. Right: the first move against the step e taken, $r = 17.5$°C, 45 kW; the routes part beyond $e = 0.15$ K, where bounds enter the plan.

Table 15.7 moves the cost and the steps. Each row is one change from Table 15.6, at $\rho = 25$ unless it says otherwise; the columns give the smallest cost found, from one start at the clipped commands, the clipped law's excess over it, and its steps at a stop. With warning, the load steps 30 steps into a 150-step run and the law and the optimum are told the post-step target from the start. The excess runs from 1.53 % to 13.87 %. The hall air is slow, so a charge on it gives the valve something to plan for, and a charge on the supply air does not. Even with only 4 steps at a stop the excess on the hall air is 4.71 %. Near-optimal clipping is therefore a finding about this cost and these steps and not a property of clipping.

Table 15.7: How far clipping's near-optimality reaches. Clipped law against the smallest cost found, known plant, known load, no noise; one start, so each excess is a lower bound on the true one if a lower optimum exists.

Change from Table 15.6	J smallest found [–]	Clip [%]	Stop steps [–]
Hall-air cost, load 30 → 40 kW	60.56559	4.71	4
Hall-air cost, load 30 → 53 kW	349.52257	5.13	82
Hall-air cost, load 30 → 60 kW	665.30913	9.82	116
$\rho = 10^4$, load 30 → 60 kW	102.70252	2.14	79
Load 30 → 53 kW with r 18 → 17.85°C	1.78711	1.53	68
Load 30 → 60 kW, 30 steps' warning	36.52329	13.87	90

The QP can lose to clipping, 0.13 % against 0.00 % on the second row, because its plan runs on the linear model. Nor is it monotone in N: on the second row at $\rho = 400$, QP(5), QP(10) and QP(20) cost 5.70171, 5.70312 and 5.70254. What it does gain is anticipation. On the slow load step it banks cold air before the stop, 17.686°C against 17.701°C at step 30 in the lower left panel, and QP(40) brings the excess to 0.24 %.

The right panel shows the fixed row's end. The two routes agree to $e = 0.15$ K, at which the

plan's stop count is still zero. It then runs 2, 7, 12 and 18 of 20 moves, at $e = 0.20$, 0.30, 0.40 and 0.50 K, and the slope dc_0/de at $e = 0.30$ K is -0.02249 for the QP against $+0.01037$ for clip. They do not differ by a scale: the sign is different. At that e the first move is 0.97015, below the stop, and the row has already changed, because the plan's later moves are on the bound. **The fixed row is gone before the valve reaches its stop, and for this cost it hardly costs.** Wound memory barely matters either, with $J = 36.85516$ on both clip and wound for the load past the stop, because this law's memory leaks: the gain on the last command in the vector $(-0.07759,\ -0.01025,\ 0.25232,\ 0.44234)$ of equation (15.14) is 0.44. That memory exists only because this cost charges moves; the law of Chapter 12 has none to wind. Windup needs an integrator, so it is a fact of the Fit line and not of this law.

15.3.5 Windup as a row

Control engineer's question. Why does the integrator wind, and what stops it?

General form. The Fit line of equation (7.1) takes the move as an input. Any controller with state $\boldsymbol{\xi}$, $\boldsymbol{\xi}^+ = \mathbf{A}\boldsymbol{\xi} + \mathbf{B}e$, $c = \mathbf{C}\boldsymbol{\xi} + \mathbf{D}e$, can be rewritten for any matrix $\mathbf{L}$ by adding $\mathbf{L}(c - \mathbf{C}\boldsymbol{\xi} - \mathbf{D}e) = \mathbf{0}$,

$$\boldsymbol{\xi}^+ = (\mathbf{A} - \mathbf{LC})\,\boldsymbol{\xi} + (\mathbf{B} - \mathbf{LD})\,e + \mathbf{L}\,c. \tag{15.16}$$

This has the shape of an estimator of the controller's own state, driven by the command, and it is only a rewriting: an algebraic identity with no noise in it and no posterior behind it. It is the same controller, for every $\mathbf{L}$, while c is what the plant received. At the stop it is not, and the row with c then reads a valve the plant did not get. Replacing c by $\mathrm{sat}(c)$ makes the row true, but the input alone does not cure anything. The plain PI is $\mathbf{L} = 0$, which reads no c at all, so at $\mathbf{L} = 0$ the true input changes nothing. The cure is the true input with a gain that meets a condition, given below.

Specialized form. Equation (15.16) with the PI row (5.1): $\boldsymbol{\xi} = I$, $\mathbf{A} = \mathbf{C} = 1$, $\mathbf{D} = K_p$, $\mathbf{B} = K_i\Delta t$, and the constant u_0 carried in c. Taking $\mathbf{L} = \beta$ and $\mathrm{sat}(c)$ in the input gives $I^+ = (1-\beta)I + (K_i\Delta t - \beta K_p)e + \beta(\mathrm{sat}(c) - u_0)$. Since $c - u_0 = K_p e + I$, this is

$$I^+ = I + K_i\,\Delta t\,e + \beta\,(\mathrm{sat}(c) - c), \tag{15.17}$$

back-calculation in one line, stable while $|1 - \beta| < 1$. At the stop with a steady error the row settles where $K_i\Delta t\,e = \beta(c - \mathrm{sat}(c))$, so the command stands at $K_i\Delta t\,e/\beta$ past the stop, and **the gain condition** is that β be large against $K_i\Delta t$: only then is that excess small against the stroke. The other treatments are the same row. *Conditional integration* is equation (15.17) with a state-dependent gain, $\beta_k = K_i\Delta t\,e/(c - \mathrm{sat}(c))$ on the steps it holds, which makes $I^+ = I$. The *read-valve row* replaces the stored stop by the valve's reading through its innovation,

$$I^+ = I + K_i\,\Delta t\,e + \beta_{\mathrm{read}}\,\nu, \qquad \nu = y_u - \bar{u}, \qquad \bar{u}^+ = \hat{\bar{u}} + a\,(c - \hat{\bar{u}}), \qquad \hat{\bar{u}} = \bar{u} + \ell\,\nu, \tag{15.18}$$

with $a = 0.3935$ the actuator's lag over one period and $\ell = 0.20$. It needs no number for the stop. Only the estimator of Chapter 13 is an estimator in the full sense, with noise and a filter's gain, and there the same row is the disturbance entry's input: where the model's valve is the unclipped command, $\hat{d}$ integrates the innovation exactly as I does.

Building blocks. The PI row; $K_i\Delta t = 0.0600$ per kelvin per step; $T_i = K_p/K_i = 50.0$ s; $\beta = \Delta t/T_t$ with $T_t = T_i$ stated, so $\beta = 0.60$.

The plant's numbers. Without the row, the integrator grows by $K_i \Delta t\, e$ per step. At the record's +0.460 K this is 0.0276 per step, or 2.21 valve strokes in 80 steps. On the record's own command of 1.0315, conditional integration's gain is $\beta_k = 0.87$. At the stop the read-valve row's innovation settles at $-0.7644\,(c - 1)$, so $\beta_{\text{read}} = 0.10$ acts as back-calculation with $\beta = 0.0764$. At +0.460 K the excess $K_i \Delta t\, e/\beta$ is 0.046 at $\beta = 0.60$, 0.092 at 0.30, 0.276 at 0.10, 0.361 at the read-valve row's 0.0764 and 0.552 at 0.05. Back-calculation at $\beta = 0$ against no row, case A below over its 150 records, differs by 0.0 in the integrator and in T_c. The rewriting is checked in Section 15.4 to 2.2×10^{-16}.

15.3.6 What each treatment costs on recovery

Control engineer's question. What does each treatment cost on recovery?

General form. The expected consequence (2.23) of fixed actions, here fixed laws, with the consequence read as recovery and fall.

Specialized form. Equation (2.23) with the action one of seven routes, on the truth's plant, and the expectation estimated by 150 fresh records, seeds 15000 to 15149, with common noise, so that differences are paired. Recovery is the first step after the load's fall from which T_c stays within 0.10 K of $r - b$ for 40 steps; the fall is the deepest T_c goes below $r - b$. The three cases are as follows. A: the record's stop, the load falling to 30 kW at step 240. B: 60 kW from step 40 to 159, then 30 kW, $r = 18$°C. C: as A with the valve's true top stop at 0.85, not told to any controller. The tunings were stated before recovery was seen. Back-calculation has $T_t = T_i$. The read-valve row takes $\beta_{\text{read}} = 0.10$, the largest on its grid with sd T_c on ordinary running within 1 % of the row's absence, 0.0330 against 0.0328 K. The estimator's walk is 0.010, nearest the PI's load-step peak on the first 50 records, 0.1836 against 0.1911 K, chosen with no reference to recovery.

Building blocks. The rows (15.17) and (15.18); the plant; the seeds and definitions above.

The plant's numbers. Table 15.8 gives the result. Each row is a route; each pair of columns gives a case's recovery as a mean with its standard error, and its mean fall. Figure 15.5 draws case A: the horizontal axis is the controller step, the upper panel the supply air, the lower the valve, each the mean over records, with the target as a dotted line and the fall at step 240 marked. Table 15.9 varies the two gains and Table 15.10 varies the band.

Table 15.8: Recovery from the fall, 150 fresh records, seeds 15000 to 15149. Recovery in steps of 30 s, mean (se), at the 0.10 K band, which Table 15.10 varies; fall in K, mean, se at most 0.0024. "Never": no record recovers by the end of the run, at this band.

	A: the record's stop		B: a load stop		C: top stop 0.85	
Route	Recovery	Fall	Recovery	Fall	Recovery	Fall
PI, none	121.1 (0.6)	0.5995	64.9 (0.8)	1.2090	never, 150 of 150	0.2770
PI, conditional (the plant's)	35.6 (0.9)	0.1246	45.4 (1.0)	0.2698	80.0 (0.5)	0.1486
PI, back-calculation	36.3 (0.8)	0.1278	45.3 (1.0)	0.2732	80.9 (0.5)	0.1507
PI, read valve	41.9 (0.4)	0.2027	45.2 (1.0)	0.4429	46.1 (1.0)	0.0909
Estimator, commanded	116.6 (0.4)	0.5893	72.7 (1.0)	1.1531	never, 150 of 150	0.2770
Estimator, clipped	23.9 (1.1)	0.1030	50.3 (1.2)	0.2315	71.2 (1.0)	0.1276
Estimator, read	25.5 (1.2)	0.1061	52.8 (1.4)	0.2377	38.7 (0.2)	0.0667

The falls come first, because they do not depend on a band. **Windup is large in the fall, and a true row with its gain condition met removes most of it.** In case A the fall goes

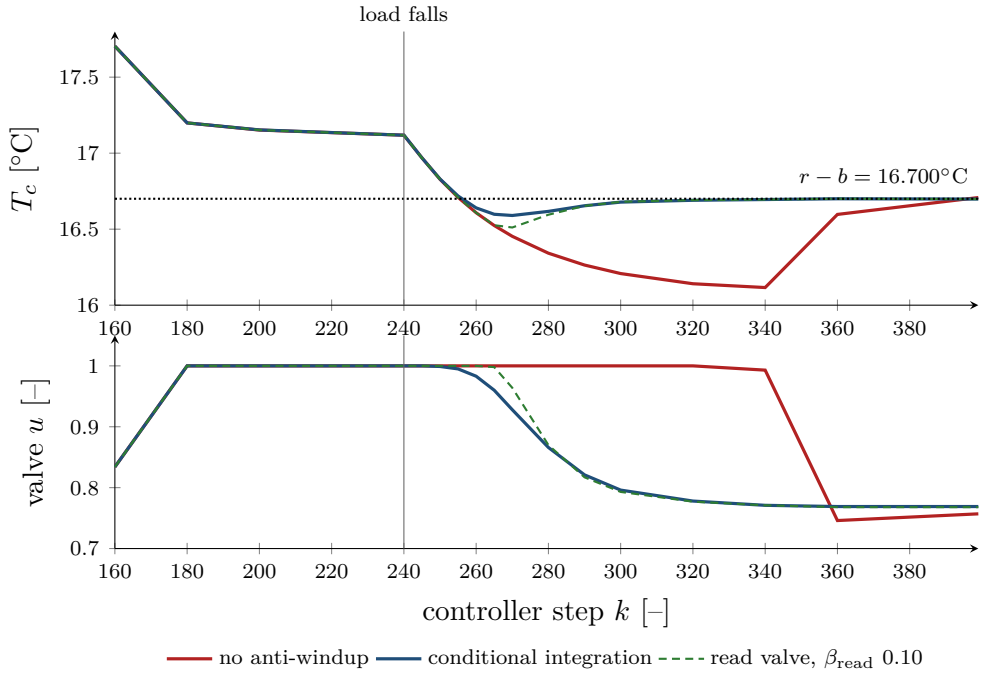


Figure 15.5: The supply air and the valve through the fall, case A, mean over 150 fresh records, straight lines between printed steps. Without anti-windup the valve stays open for about 100 steps after the load has gone, and the air falls past its target.

Table 15.9: The two gains, case A, same records. Sd T_c and travel on ordinary running, steps 40 to 79. Bold: the stated tunings. Beside them, no row has sd 0.0328 K and travel 0.4051 per hour.

Route	β [–]	Recovery, mean (se) [steps]	Fall [K]	sd T_c [K]	Travel [1/h]
Back-calculation	0.05	48.5 (0.3)	0.2916	0.0328	0.4051
	0.10	41.6 (0.4)	0.1909	0.0328	0.4051
	0.30	37.2 (0.7)	0.1331	0.0328	0.4051
	0.60	36.3 (0.8)	0.1278	0.0328	0.4051
	1.00	35.7 (0.9)	0.1255	0.0328	0.4051
Read valve	0.03	61.7 (0.5)	0.4142	0.0328	0.4051
	0.05	50.9 (0.2)	0.3250	0.0328	0.4053
	0.10	41.9 (0.4)	0.2027	0.0330	0.4070
	0.30	35.5 (0.9)	0.1274	0.0350	0.4266
	0.60	36.1 (1.1)	0.1276	0.0414	0.4871

from 0.5995 to 0.1246 K and in case B from 1.2090 to 0.2698 K. The integrator stands at 2.643 at the fall without a row and 0.485 with conditional integration. A true row with a low gain removes less: back-calculation at $\beta = 0.05$ still falls 0.2916 K, and the read valve's integrator stands at 0.788 at the fall. Conditional integration and back-calculation agree within one step at the 0.10 K band, with a paired difference of +0.7 (se 0.2). The estimator with its commanded row winds as the PI does, a fall of 0.5893 K, and with its clipped row it does not, 0.1030 K.

A recovery time is a statement about its band. Table 15.10 gives the three routes at four bands. Each row is a band; each column a case and a route, recovery as a mean with its standard error. Case A's fall of about 0.125 K sits next to the 0.10 K band, so conditional integration's recovery runs 78.0, 35.6, 13.5 and 7.9 steps from band 0.05 to 0.20 K. At 0.10 K the unprotected route takes 121.1 steps against 35.6, a paired difference of +85.5 (se 1.1). At 0.05 K the band is near the loop's own running spread, with about 35 % of records inside it over the last 40 steps on every route, and read valve and conditional integration tie, 77.5 against 78.0.

Table 15.10: Recovery against the band, mean (se) steps, the records and seeds of Table 15.8. Records with no recovery count the 220 steps to the end of the run: at 0.05 K, in case A 30, 4, 4 and in case C 150, 10, 10; at 0.15 K, case C none 66.

	A: the record's stop			C: top stop 0.85		
Band [K]	None	Conditional	Read valve	None	Conditional	Read valve
0.05	157.7 (2.8)	78.0 (3.1)	77.5 (3.1)	220.0 (0.0)	109.4 (3.0)	98.6 (3.4)
0.10	121.1 (0.6)	35.6 (0.9)	41.9 (0.4)	220.0 (0.0)	80.0 (0.5)	46.1 (1.0)
0.15	117.8 (0.4)	13.5 (0.7)	35.2 (0.3)	116.2 (7.5)	48.0 (1.5)	34.3 (0.1)
0.20	116.3 (0.4)	7.9 (0.1)	20.8 (1.0)	30.8 (0.1)	30.8 (0.1)	30.8 (0.1)

Where the stop is known the read valve falls deeper. At its stated gain it falls 0.0781 K deeper than conditional integration in case A and 0.1731 K in case B, and it is 6.3 (se 0.7) steps slower in case A at the 0.10 K band. Table 15.9 says why: at $\beta_{\text{read}} = 0.30$ it matches recovery at 35.5 steps, but at sd T_c of 0.0350 K and travel 5.3 % higher, since the reading's noise reaches the integrator. Back-calculation pays no such price, because its row has no noise in it.

Where the stop is not where the controller thinks it falls less. In case C it falls 0.0909 K against 0.1486 K, and the estimator's read row 0.0667 against 0.1276 K. Its recovery lead depends on the band: 46.1 steps against 80.0 at 0.10 K, a paired difference of −33.8 (se 1.1), 34.3 against 48.0 at 0.15 K, and none at 0.20 K. This is a modest surprise, and its case is an unknown stop, not this plant as constructed. The unprotected route's "never" depends on the band as well, and the stuck valve behind it does not: at 0.20 K all three routes count 30.8 steps, yet none of the unprotected records is inside that band over the last 40 steps. The air passes through for 40 steps on its way down and ends the run 0.2600 K below the target, its valve still at the stop of 0.85.

Like with like. The estimator is a different law, so the comparison with PI needs care. On ordinary running its travel is 0.5755 against 0.4051 per hour, its load-step peak 0.1900 against 0.1987 K. Case A is bimodal for it, with a median of 13 steps against 37 and a tenth percentile of 11 against 13, so its mean of 23.9 does not describe a typical record. Case B reverses the order, +4.9 (se 0.8) steps from conditional integration. There is no one verdict of estimator against PI here, and Chapter 13, which compares the two on ordinary running, reports no winner either. A choice between them would rest on the travel the valve can afford, which neither chapter prices.

15.3.7 What the stop reveals

Control engineer's question. What does the stop reveal?

General form. Moments of an entry under a posterior with new rows, equation (2.10).

Specialized form. Equation (2.10) with $f = b$ or g_w under the posteriors of Sections 15.2.2 and 15.2.3. The reason is the fifth line of equation (15.7). While held the discharge line is zero; at the stop it reads $y_c = T_w + Q/W + b$. As one remark in scalars, with a held window at the same load, $W u_1 (r - b - T_w) = Q = W(y_c - b - T_w)$, so W and Q cancel:

$$b + T_w = \frac{y_c - u_1\, r}{1 - u_1}. \tag{15.19}$$

This shows two things. The denominator is the *leverage*: at a stop at another load Q_2 it is $1 - u^*(Q_2, r)$, so a stop at a load where the valve had room is worth more. And **the row reads** $b + T_w$, the columns of b and T_w in the stop's line being equal. The stop is an identification moment for the sum. **It identifies the offset only with T_w known.** It also rests on Wu linear in the valve, and the constructed record, whose T_w is pinned by construction, can test neither condition.

Building blocks. The posteriors of step 3; the closed form (15.19) as a check; the node forms of Chapter 6 on each posterior for the consequence.

The plant's numbers. Static on the record, equation (15.19) is unusable: steps 120 to 159 give $u_1 = 0.8145$, leverage 0.1855, and steps 200 to 239 give $b = 1.0882$ K against a truth of 0.30 K, since the hall is still settling and the load differs by 0.27 kW. The stacked rows of Section 15.2.3 carry the transient. Table 15.11 sets their consequences beside Chapter 6's. Each row is a posterior; the columns give the offset with its spread, the window at 45 kW, the closure, and T_{sat} at 54.99 kW with its spread. Figure 15.6 draws two views. The left panel is the plane of the offset and the coil, with the one- and two-standard-deviation contours of the kept and dropped fits and the truth as a cross. The right panel has the leverage on its horizontal axis and the spread of b after 40 steps at a stop on its vertical axis, one point per stop load.

Table 15.11: The stop's rows on Chapter 6's questions, state route on each posterior, stacked fits on the record, seed 51: one record's draw.

Posterior	$\hat{b}$ (sd) [K]	Window at 45 kW [°C]	Closure [kW]	T_{sat} at 54.99 kW (sd) [°C]
Chapter 6, 30 kW window	0.05708 (0.18986)	17.438 to 19.937	54.99	27.063 (0.221)
Stacked, the stop's rows kept	0.34532 (0.08539)	17.438 to 20.012	55.51	27.238 (0.059)
Stacked, the stop's rows dropped	0.03821 (0.12429)	17.387 to 20.027	55.90	27.133 (0.067)
Stated truth	0.30	17.408 to 20.041	55.904	27.279

The stop's discharge rows remove 52.1 % of b's variance on the frozen rows of the stack, 0.0859 against 0.1242 K. On this record the MAP moves to 0.34532 K from 0.03821 K, with the truth at -0.53 against $+2.11$ standard deviations on b and -0.48 against $+2.02$ on g_w. That is one record's draw and not a result. Twelve fresh records, seeds 15700 to 15711, give a mean $\hat{b}$ of 0.2624 (se 0.0201) kept against 0.1657 (se 0.0265) dropped, the kept mean sitting beside Chapter 10's 0.2330 over other seeds; both still stand below the truth, and twelve records are thin. The consequence, state route on seed 51, is the second row of Table 15.11: the band on T_{sat} falls from 0.221 to 0.059 K, and the closure moves to 55.51 kW against the truth's 55.904 kW. The dropped fit's 55.90 kW comes from a mean far from the truth's and is no sign of a better fit.

The stop test. The right panel of Figure 15.6 says where to do it: at low load. The stated move of Section 15.2.2, at a hall air of 21.258°C and leverage 0.4399, with T_c drifting -0.0031 K

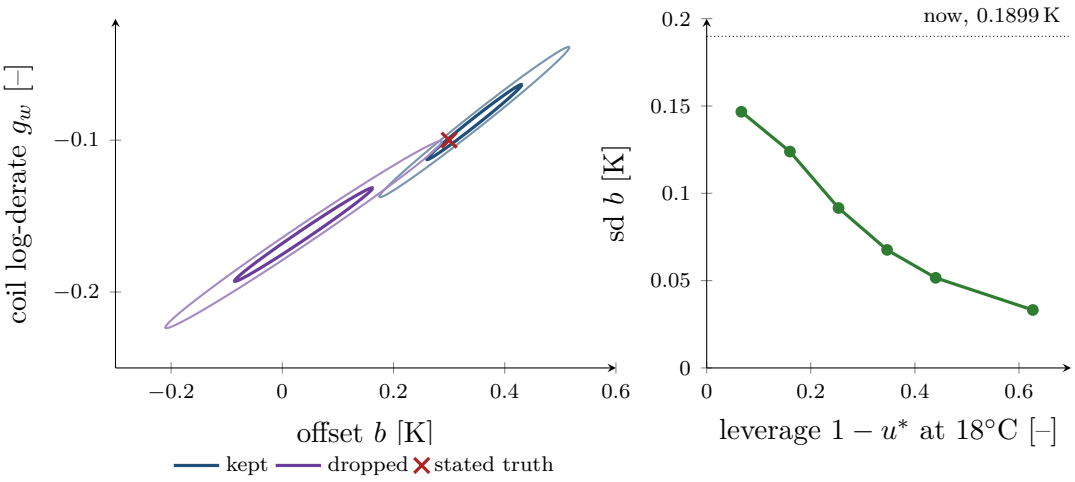


Figure 15.6: What the stop reveals. Left: the Laplace contours of (b, g_w) on the record's stacked rows, seed 51, one and two standard deviations, with the correlations of Section 15.2.3; one record's draw. Right: the spread of b from Chapter 6's rows and 40 steps at a stop, frozen rows at the truth's leverage, at stop loads of 50, 45, 40, 35, 30 and 20 kW from left to right, T_w pinned.

over the window, was run on 200 fresh records, seeds 15500 to 15699. The mean $\hat{b}$ is 0.2850 (se 0.0051). The spread across records is 0.0728 K against a stated 0.0500 K, and the truth lies within two stated spreads in 83.5 % of records on b, 77.0 % on g_w and 90.0 % on g, where a calibrated band gives 95 % (se 1.5). The band is too narrow, because static rows ignore the load's wander through the lagging hall. The closed form (15.19) with the meters' ratio and no priors has mean 0.3005 (se 0.0056) and spread 0.0796 K.

The test's safety is a stated condition of it. The move lasts 180 steps, 90 minutes, from step 160 to step 339. Over the same 200 records the hall air stands at a mean of 21.26°C in the window and reaches at most 23.04°C over the test, safe at 30 kW against the 27.5°C limit. Three things are assumed and not tested: the plant here has no cold-side limit, though the supply air goes down to 15.99°C; T_w is held unmoved with the valve fully open, which a real chilled-water loop need not do; and the test needs a protected integrator. An unprotected PI winds to 5.23 in the test and takes 76.9 (se 0.9) steps to return at the 0.10 K band, against 0.49 and 27.7 (se 0.9) with conditional integration. A stop window at high load has a ceiling of its own, since the hall air floats at the stop: the hall air's closure caps it near 55.30 kW.

The condition. Table 15.12 releases the water, with T_w an entry on a prior added to Chapter 6's window and the test's frozen rows. Each row is a prior spread; the columns give the spread of b before the test and after it, of T_w after it, and of the sum. The sum stays at 0.0512 to 0.0520 K whatever the prior, and the spread of b follows the prior. For $b + T_w$, the stop is an identification moment. For b alone, it is one only as far as T_w is known.

Table 15.12: The stop test with the chilled water on a prior, frozen rows.

Prior sd T_w [K]	sd b before [K]	sd b after [K]	sd T_w after [K]	sd$(b + T_w)$ after [K]
0.01	0.1879	0.0524	0.0100	0.0515
0.05	0.1880	0.0686	0.0486	0.0516
0.10	0.1881	0.0984	0.0899	0.0518
0.20	0.1887	0.1431	0.1432	0.0520
0.50	0.1915	0.1844	0.1895	0.0512

15.4 The checks (step 5)

The tests stand against the routes, the rewriting, the draws, the orders the mathematics fixes and the limit of exact rows.

15.4.1 Routes and identities

Table 15.13 lists them. Each row is a check; the columns give what is compared, the result and its tolerance or standard error. All pass.

Table 15.13: The checks of the chapter.

What is compared	Result	Tolerance
Batch simulator against the record's own, seed 51, 240 steps	gap in T_c and c of 0.0	10^{-9}
Integrator against the plant's, one period	gap 0.0	–
Back-calculation (15.17) against the rewritten form (15.16), 500 draws, seed 15	gap 2.2×10^{-16}	10^{-12}
QP with bounds inactive against the LQ move, $N = 5$ to 40	gap 0.0	10^{-8}
Bound less the best law, over Table 15.6	$+8.1 \times 10^{-9}$	10^{-6}
Identity (15.12) on the draws, $u_{\text{hold}} \le 1$ against $u^* < 1$, 8 nodes, 2 below the floor	gap 0.0	–
Back-calculation at $\beta = 0$ against no row, case A, 150 records	gap 0.0 in I and T_c	–
Draws, seed 15 against seed 6: effective size; Pr(stop) at the kink, 45 kW; joint at closure	1 214 073 against 1 209 705; 0.4905 against 0.4911; 0.9541 against 0.9537	se about 0.0004
Chapter 6's saturated plant at 60 kW	0.0010	its 0.0010
Exact closure, set-point grid 0.005 against 0.0025 K	0.9702, 0.9526, 0.9276 both	printed digits

15.4.2 The limit of exact rows

The rows (15.6) are on the reduced state. Table 15.14 refits the first stop test on the nine-entry state of two windows, with the physical rows at a finite tolerance and the valve's at $0.01\,\sigma_p$. Each column is a tolerance, the last the reduced state. By $\sigma_p = 0.001$ the full state agrees with the reduced state to the fifth decimal in the entries, and to 5×10^{-5} in the spread of b.

Table 15.14: The nine-entry state against the reduced state as the tolerance falls, first fresh stop test. Standard deviations in parentheses.

Entry [unit]	$\sigma_p = 0.100$	$\sigma_p = 0.010$	$\sigma_p = 0.001$	Reduced state
g_w [–]	−0.17311 (0.03912)	−0.16008 (0.01859)	−0.15976 (0.01775)	−0.15976 (0.01774)
b [K]	0.10229 (0.14875)	0.15458 (0.05780)	0.15583 (0.05362)	0.15584 (0.05357)

15.5 The report (step 6)

Chapter 6's window stands by identity: it was the joint of the kinked plant all along. At or above the mean's corner the single normal misses the 95th percentile by at most 0.0341 K, because the valve reading holds the margin narrow; it misleads below the corner, 0.1666 K at 45 kW and 0.4652 K at 60 kW, and with the valve rows dropped or the load unpinned. Past the

floor a void 0.9597 at 56 kW is really 0.8207. With the cost on the fast supply air and the step known at once, clipping's excess over the smallest cost found is at most 1.21 %, 0.88 % where a stop is touched; with the cost on the hall air it is 4.71 to 9.82 %, and with warning of the step 13.87 %. Windup is a fall of 0.5995 K against 0.1246 K, and a true row cures it if its gain is large against $K_i \Delta t$. The stop reads $b + T_w$.

15.5.1 The answers and their bands

Table 15.15 gives the answers in the order of Table 15.1. Each row is a question; the columns give the answer, its band or what it rests on, and the stated truth.

Table 15.15: The report for the hall at its stop. Direct: seed 6; records as stated. Percentages are this example's.

Control engineer's question	Answer	Band or basis	Stated truth
Where is the stop, and with what probability is the loop at it?	At 18°C the kink is at 53.350 kW; at 45 kW, $r = 17.3829$°C. Pr(stop) at 45 kW, $r = 17.38$°C: 0.5321, 0.5322, direct 0.5231; at the kink, normal 0.5000, direct 0.4911	sd m 0.0359 K at 45 kW, 0.0248 K at 54.99 kW; se 0.0005	53.566 kW; 17.4083°C
What is the hall air across the kink?	At the kink, 45 kW: state route (hold) 24.6896°C, sd 0.1559 K; direct 24.6975°C, sd 0.1677 K	q95 gap at the corner 0.0341 K against back-off 0.2565 K, below it 0.1666 K; $\theta\varphi(0)$ 0.0143 K	24.867°C
Does Chapter 6's window survive?	Yes, by identity: 17.438 to 19.937°C, direct 17.440 to 19.942°C. Closure 54.99 kW, direct 55.02 kW. Hall air alone closes at 55.30 kW, direct 55.33 kW; void 0.1492 is 0.0010	valve's price 0.30 kW (0.31 kW direct)	17.408 to 20.041°C; 55.904 kW
What is the law at the stop?	A move-penalized variant, derived here. On the linear model piecewise affine, its row following the plan's active set; constant 1 at the stop. Clip's excess 0.72 %, largest laid to the stop 0.88 %; 0.59 % baseline with no stop touched; 4.71 to 13.87 % on a slow cost or with warning	against the smallest cost found; known plant, load, state; no noise	–
Why does the integrator wind, and what stops it?	The rewritten row reads c, not sat(c); 0.0276 per step, 2.21 strokes in 80 steps. Back-calculation, conditional integration and the read valve are one row, and each needs β large against $K_i \Delta t$; at $\beta = 0$ the true input changes nothing	an identity, gap 2.2×10^{-16}; no noise, no posterior	–
What does each treatment cost on recovery?	Fall, case A: none 0.5995 K; conditional 0.1246; back-calculation 0.1278; read valve 0.2027. Case C: read valve 0.0909 against 0.1486. Recovery, case A, none, conditional, read valve: at the 0.10 K band 121.1, 35.6, 41.9 steps; at 0.20 K 116.3, 7.9, 20.8	fall se at most 0.0024; recovery depends on the band; 150 records; stated tunings	–
What does the stop reveal?	$b + T_w$. Stop test: $\hat{b}$ mean 0.2850 K, sd across 0.0728 K. Stacked, 52.1 % of b's variance	stated band too narrow, 83.5 % within two; T_w pinned	$b = 0.30$ K

15.5.2 The entry in the commissioning log

The entry keeps conditional integration: it is the plant's, its fall matches back-calculation's, 0.1246 against 0.1278 K, and the stop here is known. The vendor's option adds nothing unless the top stop may not be where the controller believes, and a read-valve row is what would answer that, with a shallower fall at a stop of 0.85. Any recovery time it logs names its band. It records that void nodes are not to be read as probabilities, with 0.9597 against 0.8207 at 56 kW as the reason. It states that protecting the valve costs 0.30 kW of load, by choice. And it recommends a stop test at low load, 90 minutes with the integrator protected, to be scheduled once T_w is known and the coil's cold side and the chilled-water loop are cleared for it, since until then the test reads the sum and not the sensor, and the model has no cold-side limit. Each of these rests on a constructed record and on the stated tunings.

15.5.3 Common mistakes

Table 15.16 lists six mistakes this hall makes easy. Each row names the mistake, its form here and the correction.

Table 15.16: Six mistakes at the stop.

Mistake	On this hall	Correction
Reading a void node as a probability	0.9597 at 56 kW, where the plant gives 0.8207	past the floor use T_{sat}, the second regime
Trusting a central difference because it matched	sd 0.1677 K at the kink; 0.1559 and 0.1805 K a millikelvin to either side	the match holds at $m = 0$ with small θ, by equation (15.11); state the two regimes
Assuming the kink always matters, or never	0.0341 K at the corner with the valve rows; 0.1666 K below it; opposite sides of 0.95 with $s = 2$ kW	read θ, the margin's spread, and the side of the corner first
Crediting the stop with a linear law's excess, or calling clipping safe	0.59 % with no stop touched; 9.82 % with the cost on the hall air	measure the baseline off the stop; state the cost and the steps
Calling the read valve better	falls 0.0781 K deeper where the stop is known; 0.0577 K less at a stop of 0.85; its recovery lead gone at a 0.20 K band	state the case, the band and its noise price
Reading the stop as an identification of the offset	sd b 0.0524 to 0.1844 K with the water's prior; the sum 0.0512 to 0.0520 K	the row reads $b + T_w$; pin T_w by measurement

15.6 The next measurement (step 7)

The stop has added two candidates a held loop does not have, one of which is a move. The rows of each are added to Chapter 6's fit by the preposterior (4.73).

15.6.1 The candidates

Table 15.17 ranks them, with frozen rows at Chapter 6's fit, from the present sd b of 0.1899 K, sd T_{sat} at 54.99 kW of 0.2206 K and closure of 54.99 kW. Each row is a candidate; the columns give the spreads and the closure with the means held, and the spread of b again with T_w on a prior of 0.20 K.

On frozen rows the stop test stands near the reference on b, 0.0517 against 0.0484 K, with no instrument, while T_w is pinned. The frozen 0.0517 K is not what the test delivers: across records it measures 0.0728 K against the reference's 0.0484 K, which has no wander to carry, and its stated band covers about 84 %. Released at 0.20 K, a prior that is a stated choice with

Table 15.17: The candidates, frozen rows, means held at Chapter 6's fit.

Candidate	sd b [K]	sd T_{sat} [K]	Closure [kW]	sd b, T_w prior 0.2 K [K]
Stop test, 30 kW, 40 steps (settling 140)	0.0517	0.1149	55.74	0.1431
Record's stop, 45 kW, as if settled	0.1239	0.0938	55.84	0.1436
Held at 45 kW, 40 steps	0.1591	0.1076	55.87	0.1585
Set-point move to 20°C, 40 steps	0.0901	0.1422	55.67	0.1481
Reference discharge reading, 0.05 K	0.0484	0.1649	55.44	0.0483
Permanent second discharge sensor, offset 0.2 K	0.1379	0.1944	55.20	0.1374
Permanent chilled-water sensor, 0.05 K	0.1899	0.2206	54.99	0.1877
Stop test and that sensor	0.0517	0.1149	55.74	0.0676

no source in the record, the test gives 0.1431 K, and only the reference holds. The permanent chilled-water sensor does nothing alone with T_w pinned, as it must, and with the test it restores most of it, 0.0676 K. The ranking depends on the entry: closure wants windows at high load, held or at the stop. The closures hold the means unchanged, so they rank rows and claim no load; on the same footing the hall air's own closure reaches 56.17 kW for the record's stop as if settled, above the truth's 55.904 kW, which is not a claim about the plant. A stop test also moves the plant to teach. That is coupling 3, and choosing such moves is Chapters 22 to 24's business.

15.6.2 The recommendation

For the offset, a reference discharge reading brought once remains the measurement that survives a released T_w, and at 0.0484 K it stands below the 0.0728 K the stop test measures across records. For the band at closure, windows at high load do more, up to the cap near 55.30 kW. The stop test is worth scheduling at low load once a chilled-water reading is trusted, since then it costs no instrument at the discharge and reads the coil as well. The permanent instrument, among these, is the chilled-water sensor, whose value lies in what it lets the stop say.

15.6.3 What the analysis showed

The trend suggested a loop short of valve for a moment, a table figure that was roughly right, and an option to buy; a single-normal band could also be expected to fail at the corner. The analysis finds the table figure meaningless past the floor, the option redundant on this plant unless its stop is not where the controller thinks, and the band sound at or above the corner, because the valve reading has already pinned where the corner is, and unsound below it or once that reading or the load's pin is lost. The law's fixed row ends with the plan's first active bound, before the valve reaches its stop, and for a cost on the fast supply air with the step known at once the plant barely charges for it; a slower cost or warning of the step charges more. The integrator's growth, which does cost, is a rewritten row reading a command the plant did not get, and it stops only with a gain large enough. And the discharge reading, which the loop had consumed since Chapter 5, comes alive at the stop and speaks about a sum of two entries, of which the engineer can only take one from it by measuring the other.

15.6.4 Limits of the model

The water and the valve. T_w is pinned in most of the chapter and Wu is linear in the valve; both carry the reading of the stop, and the record is built to satisfy them.

The law's study. Its plant, load and state are known and it has no noise. Its percentages are this example's and this cost's, and the smallest cost found is a local solver's.

The recovery measure. A recovery time depends on its band, and a route can pass through a wide band for 40 steps and be counted; the falls do not depend on it.

The stop test. No cold-side limit is modeled and T_w is assumed unmoved with the valve fully open. The prior of 0.20 K on T_w is a choice with no source in the record.

Steady nodes. Sections 15.3.1 to 15.3.3 read steady nodes. The stop's static rows are over-confident, 83.5 % within two stated spreads.

The records. Fresh records come from the stated truth; twelve stacked refits are thin; the estimator's linear point is fixed at 30 kW. The what-ifs are stated, not observed.

What belongs elsewhere. The load's forecast is Chapter 19's, the stop in a receding plan Chapter 18's, and a set point chosen after a reading Chapter 16's.

The constructed record. The record is constructed from a stated truth, so its agreement tests the method on its own model and not the model against a real hall.

15.7 What to retain

Definition 15.1 (The stop as a regime)**.** At an actuator's stop the controller's row leaves the graph and a pin takes its place, equation (15.1). The answer is then $\max(T_{\mathrm{hold}}, T_{\mathrm{sat}})$, the regime is the event $m < 0$ on the state, and the first-order effect of the corner is sized by $\mathrm{sd}(m)\,\varphi(0)$.

Three conclusions carry beyond the hall as statements of form. A joint constraint on the holding regime and the stop is already the joint of the kinked plant, by equation (15.12), while a single limit read past the floor is not a probability of anything. Any controller can be rewritten with a row on its own command, equation (15.16), an identity while the command is what the plant received; anti-windup is that row with the true input and a gain large enough, and with no gain the true input changes nothing. And a reading the loop consumes returns as a row when the loop lets go, with whatever else that row carries. Whether clipping suffices, and whether the single normal misleads, were computed here for this plant only, and for clipping for the stated costs and steps only. The chapter's questions join the others in Appendix A.

Table 15.18 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 15.18: The terms of the chapter and where each is defined. Units for the closed-loop hall.

Term	Expression	Units	Meaning	Defined in
The row swapped	eq. (15.1)	–	controller's row out, pin in	15.1.2
Two regimes	eq. (15.2)	°C, –	the plant on each side	15.1.2
Margin at the stop	m, eq. (15.3)	K	$T_{\text{hold}} - T_{\text{sat}}$; negative at the stop	15.1.2
Moved zero row	eq. (15.4)	–	the discharge line opens, the valve's closes	15.1.4
Moments of the maximum	eq. (15.10)	°C	the two-regime route	15.3.2
Cdiff's variance at the kink	eq. (15.11)	K^2	the gap to the maximum's variance	15.3.2
Window's identity	eq. (15.12)	–	the joint is the kinked plant's	15.3.3
Move-penalized law	eq. (15.14)	–	four states, the last command the memory	15.3.4
Bounded plan	eq. (15.15)	–	on the linear model, piecewise affine in the state	15.3.4
Controller rewritten	eq. (15.16)	–	the same law, any $\mathbf{L}$; an identity	15.3.5
Back-calculation	eq. (15.17)	–	the row with sat(c), β large against $K_i\Delta t$	15.3.5
Read-valve row	eq. (15.18)	–	the row with the valve's reading	15.3.5
Recovery	first step in band for 40 steps	steps	cost of the stuck hour, at a stated band	15.3.6
Leverage of the stop	$1 - u^*$, eq. (15.19)	–	the valve's room at the comparison	15.3.7
Stop test	rows (15.6)	–	a move to the stop, reading $b + T_w$	15.2.2

15.8 Exercises

Exercise 15.1. *Formulation.* A level controller moves a pump's speed to hold a tank's level reading at a set point, with a fixed outflow demand and a top speed. Write the two regimes, the margin and the row that swaps. (Holding: the level reading equals the set point and the speed $s = D/k$ meets the demand D, with k the pump's flow at top speed. At the top speed: $s = 1$, the controller's row leaves and the level is free, falling at the rate of $D - k$. The margin is $k - D$, and the regime is the event that it is negative. The level then has no steady value, so this plant has a second regime with no steady node, a difference from the hall.)

Exercise 15.2. *Derivation.* Derive back-calculation from the rewritten controller, and say what happens at $\beta = 0$. (Equation (15.16) with $\mathbf{A} = \mathbf{C} = 1$, $\mathbf{D} = K_p$, $\mathbf{B} = K_i\Delta t$, $\mathbf{L} = \beta$ and sat(c) for c: $I^+ = (1-\beta)I + (K_i\Delta t - \beta K_p)e + \beta(\text{sat}(c) - u_0)$. With $c - u_0 = K_p e + I$, this is $I + K_i\Delta t\, e + \beta(\text{sat}(c) - c)$. Against the direct form on 500 draws the largest gap is 2.2×10^{-16}. At $\beta = 0$ this is the plain PI: the true input enters with no gain and changes nothing.)

Exercise 15.3. *Derivation.* Derive the leverage, and give it at 30 and 45 kW. (At the same load, $W u_1(r - b - T_w) = Q = W(y_c - b - T_w)$; cancelling W, $(1 - u_1)(b + T_w) = y_c - u_1 r$, so the row's sensitivity to the sum is $1 - u_1$, and at a stop at load Q_2 it is $1 - u^*(Q_2, r)$. At the truth and 18°C it is 0.4399 at 30 kW and 0.1599 at 45 kW.)

Exercise 15.4. *Derivation.* Show from equation (15.10) that the corner lifts the mean by $\theta\varphi(0)$ at $m = 0$ at the mean, and give it at 45 kW. (With $\mu_1 = \mu_2 = \mu$, $a = 0$, so $\mathbb{E}[T^*] = \mu(\Phi(0) + \Phi(0)) + \theta\varphi(0) = \mu + \theta\varphi(0)$. With $\theta = \text{sd}(m) = 0.0359$ K this is 0.0143 K.)

Exercise 15.5. *Application.* Find the hall air's own closure with the valve released. (Solve $\Pr(T_{\mathrm{sat}} \le 27.5^{\circ}\mathrm{C} \mid \mathbf{y}) = 0.95$ in the load, with T_{sat} of equation (15.2), its mean rising at $0.2375\,\mathrm{K/kW}$ and its spread $\sqrt{\mathbf{g}_{T_{\mathrm{sat}}}^{\mathsf{T}} \boldsymbol{\Sigma}\, \mathbf{g}_{T_{\mathrm{sat}}}}$ recomputed at each trial, from $27.0627^{\circ}\mathrm{C}$ with $0.2206\,\mathrm{K}$ at $54.99\,\mathrm{kW}$ and $27.3026^{\circ}\mathrm{C}$ with $0.2246\,\mathrm{K}$ at $56\,\mathrm{kW}$. The mean plus 1.6449 spreads meets the limit at $55.30\,\mathrm{kW}$.)

Part V

Limits and the Horizon

Chapter 16

Chance-Constrained Set-Point Selection

16.1 A set point the sheet would raise

The commissioning sheet at 23:00 shows the two-unit hall with its loop on the return air: a loop error of 0.00 K and a return set point of 30.33°C. The coils are entered at their rating and the recirculation fractions at their survey values. The sheet's arithmetic puts rack 3 on its 27°C limit at $r = 32.20$°C for a load of 200 kW, and a warmer return runs warmer chilled water.

The engineer would raise the set point by 1.87 K and ask for a better survey, so as to raise it further with confidence. The sheet has no band, says nothing of the two racks beside rack 3, and cannot say whether a tighter survey raises anything. The operations book asked this question of a pinned supply and, with no chiller-power row, priced only its risk side; here a loop stands between the set point and the supply, and a chiller row is drawn. The question of this chapter is

what set point holds the three inlets at 95 %*, and what a tighter posterior is worth.*

Jargon: chance-constrained set point, binding inlet, fresh and persistent node error, preposterior of the set point, perfect-knowledge limit, coverage over the belief, coverage at one truth, window of set points

A *chance-constrained set point* is the largest input pinned at a scenario node that meets a stated probability on a limit. The *binding inlet* is the rack whose own constraint gives the smallest set point. The *node error* e is the loop reading's error at the node: *fresh* if it is a new draw independent of the fit, *persistent* if it is the offset b_r already in the state. The *preposterior of the set point* is the distribution of the set point a reading not yet taken would produce. The *perfect-knowledge limit* is the set point chosen with the state known; its saving averaged over the belief is the analogue, under the assurance, of the expected value of perfect information. *Coverage over the belief* is how often the limit holds when the truth is drawn from what the belief says. *Coverage at one truth* is the same frequency with the truth held and only the records redrawn. The *window of set points*, from Chapter 6, is here the range between the inlets' ceiling and the supply's floor.

Key idea

The band of the fit enters the move through the limit, as a back-off, and nothing else in this chapter's decision depends on the band. A tighter posterior removes back-off, and it also moves the mean. So it raises the set point in expectation, by an identity, and it can lower it on the reading actually taken. A probability over the belief is not a frequency at one truth. And what the band costs is a price that has to be computed.

16.1.1 The restriction this chapter works under

This chapter lies on the row "One move, no future" of Table 7.6, a restriction of the problem (7.1) in its horizon: the fit is fixed and the decision is a threshold read from it. Coupling 4 of Section 7.1.3, the width of $\mathbf{\Sigma}$ entering the move through the limit, is the chapter's subject, and coupling 1 stands in the plain sense that the set point is an entry of the graph at the node. Coupling 2 is gone because no record follows the move. Coupling 3 is gone because a set point here teaches nothing, and Limits names what that gives up.

16.1.2 The two-unit hall at one node (step 0)

The plant is the two-unit hall of Chapter 11, equations (11.2) and (11.3), on the reduced state of Section 4.2.8,

$$\mathbf{z} = (Q,\ g_1,\ g_2,\ \beta_1,\ \beta_2,\ \beta_3,\ b_r)^{\mathsf{T}},$$

equation (11.4), with $\sum G = G_1 + G_2$ the coils' conductance, $\sum C = 20\,\mathrm{kW/K}$ the air's capacity and $Q_o = 8\,\mathrm{kW}$. One scenario node serves the chapter: the load pinned at $Q^* = 200\,\mathrm{kW}$, so that $Q^* + Q_o = 208\,\mathrm{kW}$ and

$$c = \frac{Q^* + Q_o}{\sum C} = 10.4000\,\mathrm{K}. \tag{16.1}$$

Chapter 11's Table 11.7 gave the setting over load; this chapter takes one decision and asks what stands behind it.

Two graphs, each named at every node. *The loop, closed:* the set point r is pinned and the supply T_s is free. From equation (11.9),

$$T_k^*(\mathbf{z}, e; r) = r - e - (1 - \beta_k)\,c, \qquad T_s^*(\mathbf{z}, e; r) = r - e - \frac{Q^* + Q_o}{\sum G}, \qquad k = 1, 2, 3. \tag{16.2}$$

The hold, open: the supply is pinned, the controller's row is out of the graph and the return floats. From equation (11.8),

$$T_k^*(\mathbf{z}; T_s) = T_s + \frac{Q^* + Q_o}{\sum G} - (1 - \beta_k)\,c. \tag{16.3}$$

Below, x stands for the setting in either graph, r or T_s. It enters each inlet with coefficient one, so a node's inlets are $m_k(\mathbf{z}) + x$ for the mean part m_k.

Two routes for the reading's error at the node. Equation (16.2) holds an error e, and the chapter carries it two ways everywhere,

$$\text{fresh: } e \sim \mathcal{N}(0,\ 0.30^2) \text{ independent of } \mathbf{z}, \qquad \text{persistent: } e = b_r, \tag{16.4}$$

the fresh route being Chapter 11's default and the persistent one its check in Section 11.3.4. The first buys a node whose error the fit says nothing about; it fails if the error stays. The second integrates what persists against its posterior; it fails if the error does not stay. Section 16.3.6 tests each in the other's world.

Figure 16.1 draws the node in both graphs. Each panel runs from the chiller row on the left through the supply, the coils and the return to the three inlets, with the load pinned below. Green is free, gray is pinned and blue is fixed by the rows. The line under each panel names which entries the inlets' gradient reads and which reach only the price.

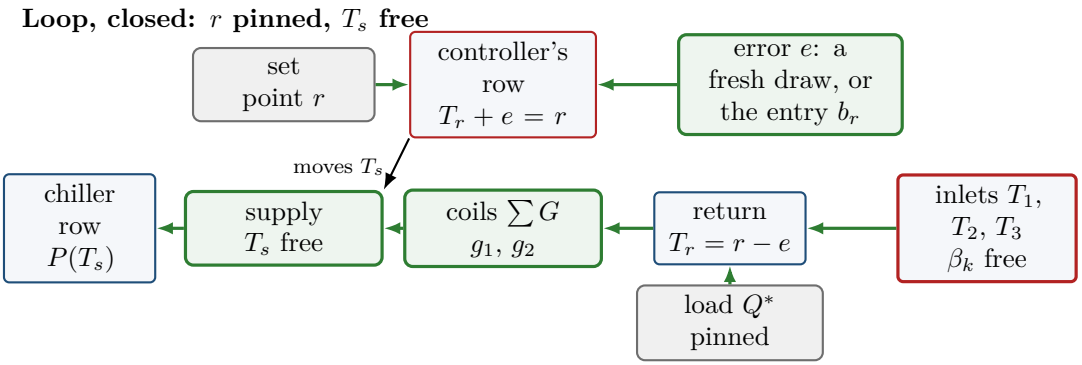


Figure 16.1: The node in its two graphs, equations (16.2) and (16.3). The loop's air balance fixes each inlet from the return, so the coils leave the inlets' band and enter the price through the supply. The hold's inlets float on the coils. The window's five readings and the priors stand behind both as the fit and are not drawn.

16.1.3 The control engineer's questions (step 1)

Table 16.1 lists eight questions. The first column is the question in the engineer's words, the second the general form, the third its mathematical form on this hall, the fourth its category, the fifth the computation and the last the subsection. The first three settle the set point; the next two ask what a tighter band does and what it is worth; the last three ask whether the 0.95 holds, where the window ends, and what to read next.

Table 16.1: The questions asked of the set point at 200 kW. $\mathbf{y}$ is Chapter 11's five live readings and six priors. x is r or T_s.

Control engineer's question	General form	Mathematical form	Category	Computation	Answered in
What set point holds each inlet?	chance constraint inverted, (2.12)	$x_{95,k} = 27 - m_k - \kappa s_k$; loop, hold, two routes	prediction under a threshold	posterior	16.3.1
What does holding all three cost?	joint chance constraint, (2.12)	$\Pr(\max_k T_k^* \le 27^\circ\mathrm{C} \mid \mathbf{y}) \ge 0.95$	prediction under a threshold	posterior	16.3.2
Does the normal hold?	the same, by draws, (2.12)	weighted quantile of the exact posterior	prediction under a threshold	posterior	16.3.3
Does a tighter posterior raise the set point?	preposterior of a threshold, (2.16)	$\mathbb{E}[x_{95}(y)]$, its sd, $\Pr(x_{95}(y) < x_{95})$	sensor placement	preposterior	16.3.4
What is a degree worth?	expected consequence through a price, (2.23)	$\mathbb{E}[P(T_s^*(x))]$ against the largest r at the truth	decision value	expected consequence	16.3.5
Does the 95 % hold on fresh records?	coverage, (16.18)	over the prior; at one truth; with an offset	validation	posterior	16.3.6
When does no set point hold?	window of set points, (2.12)	ceiling against T_s stops, over Q^*	prediction under a threshold	posterior	16.3.7
What next?	preposterior value, (2.16)	saving in \$/h, paired on the draws	sensor placement	preposterior	16.6.1

16.1.4 The evidence as rows (step 2)

The rows of the fit. The evidence is the 23:00 window of Section 11.1.4, unchanged: five live readings and six priors, in the pattern (11.10). The return reading is consumed by the controller's row. The columns of β_2 and β_3 hold only their priors, and this is the fact the chapter turns on.

Which entry reaches which inlet. A node adds no rows to the fit; it reads the state through gradients. Differentiating equations (16.2) and (16.3) in the columns of equation (11.4), with the load pinned, the three inlets' gradients have the patterns

$$\mathbf{G}_{\text{fresh}} = \begin{bmatrix} 0 & 0 & 0 & * & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & * & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & * & 0 \end{bmatrix}, \quad \mathbf{G}_{\text{pers}} = \begin{bmatrix} 0 & 0 & 0 & * & 0 & 0 & * \\ 0 & 0 & 0 & 0 & * & 0 & * \\ 0 & 0 & 0 & 0 & 0 & * & * \end{bmatrix}, \quad \mathbf{G}_{\text{hold}} = \begin{bmatrix} 0 & * & * & * & 0 & 0 & 0 \\ 0 & * & * & 0 & * & 0 & 0 \\ 0 & * & * & 0 & 0 & * & 0 \end{bmatrix}. \tag{16.5}$$

The first column is empty in all three, since the node pins the load. On the fresh route the error stands beside these gradients as a common term in all three inlets. Rack 3's line reaches only β_3, or β_3 with b_r, or β_3 with the two coils, so its band has to be mostly the survey's.

The candidate rows. A later reading adds a line to the Jacobian. An inlet reading at rack k, with the window's own load in the state, is the row

$$r_{T_k} = \left[r - b_r - (1 - \beta_k)(Q + Q_o) / \textstyle\sum C - y_{T_k} \right] / \sigma_{T_k}, \tag{16.6}$$

whose line has entries on Q, β_k and b_r only. A survey is a row on β_k alone, a reference return reading a row on b_r alone, and a reference load meter a row on Q alone. A permanent inlet sensor is twelve windows of equation (16.6) at $0.30/\sqrt{12} = 0.0866\,\mathrm{K}$ if its error is fresh.

The modeling choices. Six enter here.

- *The node at* 200 kW. It buys one decision to examine, and 200 kW is Table 11.7's row and not the fit's $\hat{Q}$. A different load moves everything, which Section 16.3.7 shows for the window.
- *The joint assurance.* 27°C on all three inlets at 0.95 together is the plant's stated policy. No value is placed on a breach, so the 0.95 is an assurance and not an optimum.
- *The price from the plant's clock.* The chiller's coefficient of performance, the ambient of 18.00°C and the tariff of 0.07 \$/kWh are the plant's stated values at 23:00 from the two-hall plant, carried to this hall. They buy a price with no new constants; a different hour or chiller changes it, so 24°C and 0.22 \$/kWh stand beside them.
- *The steady node.* The move is read after the plant has settled, which buys equation (16.2) and ignores the way there; Chapter 15 has the stop in time.
- *Coils independent of the supply.* g_u does not depend on T_s. This buys one fit for every x. The record holds T_s fixed, so it cannot test this, and it would move the price.
- *The candidates' tolerances.* Inlet sensors at 0.30 K, a survey at 0.02, a reference return at 0.05 K and a meter at 0.5 kW. One added row carries an own calibration of 0.2 K on a permanent inlet, a choice borrowed from Chapter 6's discharge sensor.

16.2 The posterior of the state (step 3)

There is one posterior, the fit of Chapter 11, and it is cited and not derived again. What is new is what the node reads from it.

16.2.1 The fit of the window

On the rows (11.11), with normal equations (11.12) and covariance (11.13), the MAP state and its spreads are

$$\hat{\mathbf{z}} = \begin{bmatrix} 194.6273 \\ -0.0735 \\ -0.1340 \\ 0.2204 \\ 0.3500 \\ 0.5000 \\ 0.2802 \end{bmatrix}, \qquad \sqrt{\operatorname{diag} \boldsymbol{\Sigma}} = \begin{bmatrix} 2.7676 \\ 0.0239 \\ 0.0252 \\ 0.0314 \\ 0.0800 \\ 0.0800 \\ 0.1935 \end{bmatrix}, \tag{16.7}$$

in kW, K and dimensionless entries, with $\mathbf{\Sigma}$ the covariance printed in Section 11.2.2,

$$\mathbf{\Sigma} = 10^{-4}\begin{bmatrix} 76590 & 205.6 & 226.1 & 34.05 & 0 & 0 & -2604 \\ 205.6 & 5.719 & -2.394 & 1.638 & 0 & 0 & 10.83 \\ 226.1 & -2.394 & 6.339 & 1.644 & 0 & 0 & 10.09 \\ 34.05 & 1.638 & 1.644 & 9.881 & 0 & 0 & 23.66 \\ 0 & 0 & 0 & 0 & 64.00 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 64.00 & 0 \\ -2604 & 10.83 & 10.09 & 23.66 & 0 & 0 & 374.4 \end{bmatrix}. \tag{16.8}$$

The information matrix of these normalized rows is $\mathbf{J}^{\mathsf{T}}\mathbf{J}$, with the curvature term of equation (11.13) kept; with unnormalized rows it is $\mathbf{J}^{\mathsf{T}}\mathbf{W}\mathbf{J}$. At the fit $\sum G = 12.6261\,\text{kW/K}$. The entries $\hat{\beta}_3$ and $\hat{\beta}_2$ are the survey's own, unread.

16.2.2 The belief at the node

With the gradients of equation (16.5) at $\hat{\mathbf{z}}$, rack 3's lines are

$$\begin{aligned} \mathbf{g}_{\text{fresh}} &= (0,\ 0,\ 0,\ 0,\ 0,\ 10.4000,\ 0)^{\mathsf{T}}, \\ \mathbf{g}_{\text{pers}} &= (0,\ 0,\ 0,\ 0,\ 0,\ 10.4000,\ -1)^{\mathsf{T}}, \\ \mathbf{g}_{\text{hold}} &= (0,\ -8.4864,\ -7.9875,\ 0,\ 0,\ 10.4000,\ 0)^{\mathsf{T}}, \end{aligned} \tag{16.9}$$

and the inlets' covariance at the node is

$$\mathbf{V} = \mathbf{G}\mathbf{\Sigma}\mathbf{G}^{\mathsf{T}} + \sigma_e^2\,\mathbf{1}\mathbf{1}^{\mathsf{T}}, \tag{16.10}$$

with $\sigma_e = 0.30\,\text{K}$ on the fresh route and the second term absent on the other two. Table 16.2 gives the result. Each row is a graph and route; the columns give each inlet's mean at $x = 0$ and spread, then the correlations among the inlets. On the fresh route the two larger correlations come through the common e.

Table 16.2: The inlets at the node by the state route: m_k is the mean at $x = 0$, [*] being K for a loop and °C for the hold; s_k the spread by equation (16.10).

Graph, route	m_1 [*]	s_1 [K]	m_2 [*]	s_2 [K]	m_3 [*]	s_3 [K]	corr(3, 2) [–]	corr(3, 1) [–]	corr(2, 1) [–]
Loop, fresh	−8.1076	0.4437	−6.7600	0.8844	−5.2000	0.8844	0.1151	0.2293	0.2293
Loop, persistent	−8.3878	0.3084	−7.0402	0.8542	−5.4802	0.8542	0.0513	0.0487	0.0487
Hold	8.3663	0.3159	9.7139	0.8610	11.2739	0.8610	0.0663	0.0774	0.0774

16.2.3 The exact posterior by reweighted draws

The reweighted draws of Section 11.4.2 carry the exact posterior beside the Laplace normal: a t_5 proposal on the Laplace normal, seed 11, $N = 10^6$, effective size 779 835. Their β_3 has mean 0.5000 and spread 0.0800; it is normal and unread, so Laplace is exact on it. Their $\sum G$ has mean 12.6258 and spread 0.1703 kW/K, against the delta route's 12.6261 and 0.1700. Outer draws of this posterior, 3000 with seed 1616, carry the preposterior and the price.

16.3 The answers (step 4)

Each answer reads the pair of Section 16.2 at the node, by the route it names. Where a value is one record's draw, or depends on a held truth, the answer says so.

16.3.1 The set point for each inlet

Control engineer's question. What set point holds each inlet?

General form. The threshold form (2.12) as the chance constraint (6.1), inverted for the pinned input through the back-off (6.3).

Specialized form. Equation (2.12) with $f = -T_k^*$ of equation (16.2) or (16.3), $f^* = -27°\mathrm{C}$ and assurance 0.95. The input is additive, so equation (6.3) inverts at once,

$$x_{95,k} = 27 - m_k - \kappa\, s_k, \qquad s_k^2 = \mathbf{g}_k^{\mathsf{T}} \boldsymbol{\Sigma}\, \mathbf{g}_k + [\text{fresh}]\ \sigma_e^2, \qquad \kappa = \Phi^{-1}(0.95) = 1.6449. \quad (16.11)$$

The assumption entering here is that the inlet is normal under the posterior; for the loop it is linear in the unread normal β_3, which is what makes this nearly exact, and Section 16.3.3 tests the rest.

Building blocks. m_k and s_k of Table 16.2, from $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$ through the gradients (16.9); κ.

The plant's numbers. Table 16.3 gives the set points. Each row is a graph and route, then the two references; the columns give rack 3's mean and spread, its back-off, and the setting for each rack. Rack 3 binds on every route. The first and third rows reproduce Table 11.7's 30.745 and 14.310°C, and Φ at each rack-3 value returns 0.950000.

Table 16.3: The setting for each rack at 200 kW, equation (16.11). The sheet is a labelled reference with no band. The truth and the sheet carry zero node error. m_3 is in K for a loop and °C for the hold, marked [*].

Graph, route	m_3 [*]	s_3 [K]	Back-off [K]	x_{95}, rack 3 [°C]	Rack 2 [°C]	Rack 1 [°C]
Loop, fresh (r)	−5.2000	0.8844	1.4548	30.7452	32.3052	34.3777
Loop, persistent (r)	−5.4802	0.8542	1.4050	31.0752	32.6352	34.8805
Hold (T_s)	+11.2739	0.8610	1.4163	14.3098	15.8698	18.1140
Stated truth, largest r; largest T_s				31.4720; 15.3067		
The sheet, r; T_s				32.2000; 17.3429		

The sheet's two settings put the truth's rack 3 at 27.7280 and 29.0362°C, both over the limit. They stand as references. The variances of T_3^* show what the band is made of, in K^2: fresh, β_3 0.6922 and e 0.0900; persistent, β_3 0.6922 and b_r 0.0374; hold, β_3 0.6922 and the coils 0.0492. **The band is the survey's fraction on every graph.** With β_3 known the fresh loop still has a back-off of 0.4935 K, which is $\kappa\sigma_e$.

16.3.2 The joint chance constraint

Control engineer's question. What does holding all three cost?

General form. The threshold form (2.12) with the indicator of an intersection, the joint chance constraint (6.20).

Specialized form. Equation (2.12), as equation (6.20) with three limits, and the largest x that meets it,

$$\Pr\big(\max_k T_k^* \le 27°\mathrm{C} \mid \mathbf{y}\big) = \Phi_3\big(27 - \mathbf{m} - x\mathbf{1};\ \mathbf{V}\big) \ge 0.95, \quad (16.12)$$

with Φ_3 the normal integral under the covariance (16.10), computed by conditioning one inlet at a time on a two-dimensional lattice, and with bisection on x. This reads one probability for one event. No bound on separate probabilities enters, because the normal gives the joint directly.

Building blocks. **m** and **V** of Table 16.2; the lattice.

The plant's numbers. Table 16.4 gives the cost. Each row is a graph and route; the columns give the setting for rack 3 alone, the joint setting, the cost, and the probabilities at rack 3's setting. Figure 16.2 draws the three inlet densities at the fresh joint set point. The horizontal axis is the inlet temperature and the vertical axis the density; the red line is the limit, and the dotted black lines are the stated truth's inlets at that set point with zero node error.

Table 16.4: The joint against rack 3 alone.

Graph, route	Rack 3's x [°C]	Joint x [°C]	Cost [K]	Pr(rack 2) [–]	Pr(all) at rack 3's x [–]
Loop, fresh	30.7452	30.7428	0.0025	0.99967	0.94971
Loop, persistent	31.0752	31.0732	0.0020	0.99974	0.94976
Hold	14.3098	14.3078	0.0021	0.99973	0.94975

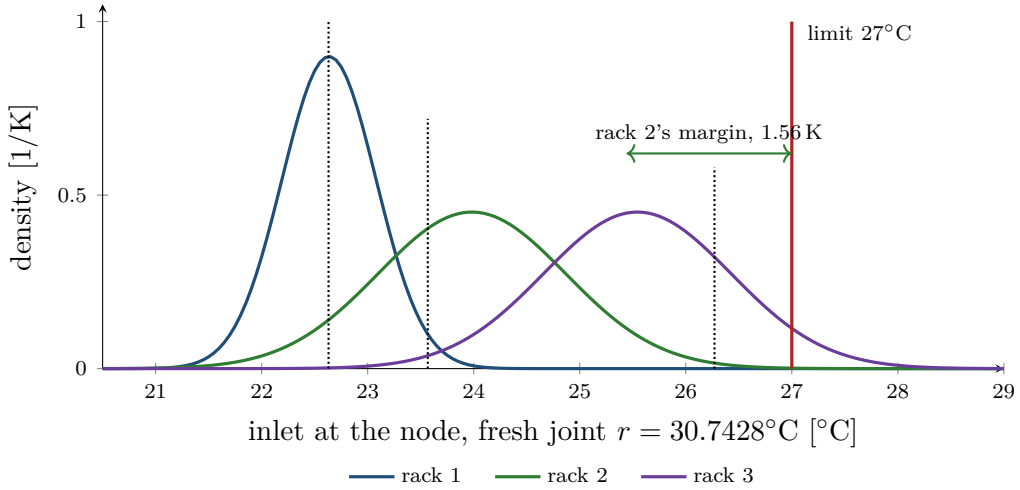


Figure 16.2: The inlets' Laplace normals on the fresh loop at the joint set point, $\mathcal{N}(m_k + r,\ s_k^2)$ with m_k and s_k of Table 16.2. Rack 3's upper 5 % point sits at the limit. Rack 2's margin is the difference of the two racks' settings on the means, drawn from its own upper 5 % point. Dotted: the stated truth's inlets, 22.6308, 23.5668 and 26.2708°C.

The joint costs 0.0025 K **on the fresh loop.** Rack 2's setting stands 1.56 K above rack 3's on the means, so at rack 3's setting rack 2 holds with probability 0.99967 and all three with 0.94971. The joint is nearly free at the window. Section 16.3.4 shows that it stops being free after a reading. The joint set point at other assurances, fresh loop, is 31.0618°C at 0.90 and 30.1418°C at 0.99, about the 30.7428°C at 0.95. The sequential integral against 4×10^6 normal draws is in Section 16.4.

The routes and the plant. The persistent r stands 0.3304 K above the fresh one at the joint. On rack 3 the gap is 0.3299 K, of which 0.2802 K is $\hat{b}_r$, the reading believed high, and 0.0497 K the smaller back-off. Each route's own world takes $\hat{b}_r$ back out: the exact means of the supply are 14.2656 and 14.3109°C with the hold at 14.3078°C, within 0.05 K. That agreement

compares two different worlds. In any one world the two settings differ by the whole 0.3304 K, as Table 16.5 shows. Each row is a world and a setting at this window; the columns give the supply, the power and the probability that all three inlets hold.

Table 16.5: Each setting in each world, over this window's exact posterior, 10^6 reweighted draws.

World	Setting	x [°C]	$\mathbb{E}[T_s]$ [°C]	$\mathbb{E}[P]$ [kW]	Coverage (se) [–]
Persistent	fresh r	30.7428	13.9804	30.6180	0.9794 (0.0001)
Persistent	persistent r	31.0732	14.3109	30.0515	0.9505 (0.0002)
Fresh	fresh r	30.7428	14.2653	30.1300	0.9500 (0.0002)
Fresh	persistent r	31.0732	14.5957	29.5646	0.8975 (0.0003)

The choice of route is a choice of 0.33 K **of supply.** It is worth 0.0397 \$/h in the persistent world and 0.0396 \$/h in the fresh one at 0.07 \$/kWh, which is 23.7 % of the whole back-off and 40.0 % of the best candidate of Table 16.19. Taken in the wrong world, the fresh setting overholds at 0.9794 and the persistent setting falls to 0.8975. The load moves the agreement with the hold too: at the fit the fresh setting's supply less the hold's is −0.1057, −0.0673, −0.0389, −0.0162 and +0.0030 K at 120, 160, 200, 240 and 280 kW.

The band by hand, beside the posterior. Table 16.6 sets the spreads of the fractions and of e by hand, with the means held at the fit. Each row is a choice; the columns give rack 3's setting, the joint setting and rack 3's back-off. This is a closed form beside the posterior, used to see the dependence and for nothing else: the rows hold the means, which a real reading would not.

Table 16.6: The set point against the band, loop, fresh; means held at the fit. The first row is the posterior's.

sd β_3 [–]	sd β_2 [–]	s_e [K]	r_{95}, rack 3 [°C]	Joint r [°C]	Back-off, rack 3 [K]
0.0800	0.0800	0.30	30.7452	30.7428	1.4548
0.0400	0.0800	0.30	31.3564	31.3426	0.8436
0.0200	0.0800	0.30	31.5995	31.5785	0.6005
0.0001	0.0800	0.30	31.7065	31.6830	0.4935
0.1600	0.0800	0.30	29.4188	29.4188	2.7812
0.0200	0.0200	0.05	31.8481	31.8481	0.3519

16.3.3 The normal against the draws

Control engineer's question. Does the normal hold?

General form. The same tail, the threshold form (2.12), read from the exact posterior in place of the Laplace normal.

Specialized form. Equation (2.12) by the reweighted draws $\mathbf{z}_i$, w_i of Section 16.2.3. The probability is the weighted count of the draws with every inlet at or below the limit, and the set point is the weighted 0.95 quantile of the largest x each draw allows. On the fresh route e is drawn with each $\mathbf{z}_i$.

Building blocks. The draws and their effective size; the node's functions (16.2) and (16.3).

The plant's numbers. Table 16.7 sets the two routes side by side. Each row is a graph and route; the columns give the Laplace value and the draws' for rack 3 and the joint, then the probability the draws give at the Laplace joint setting. The gaps of 0.004 to 0.005 K on the persistent route and the hold are 2 to 2.5 standard errors: not sampling but the curved rows through b_r and the coils. The chapter reports the Laplace values, with the draws beside them, and no verdict here turns on the difference.

Table 16.7: The set point by Laplace and by the exact posterior's draws, se about 0.002 K; the last column has se 0.00024.

Graph, route	Rack 3, Laplace [°C]	Rack 3, draws [°C]	Joint, Laplace [°C]	Joint, draws [°C]	Pr(all), draws [–]
Loop, fresh	30.7452	30.7456	30.7428	30.7431	0.95004
Loop, persistent	31.0752	31.0794	31.0732	31.0774	0.95052
Hold	14.3098	14.3049	14.3078	14.3026	0.94938

16.3.4 The claim that a tighter posterior raises the set point

Control engineer's question. Does a tighter posterior raise the set point?

General form. The preposterior (2.16), with u the set point the outcome would produce and $\tilde{y}_a$ the candidate reading.

Specialized form. Equation (2.16) with $u = x_{95}(y)$ of equation (16.11) after the row with line $\mathbf{j}$ and standard deviation σ is added. On linear rows the update of the posterior gives, by the Schur complement of Chapter 3,

$$m'(y) = m + \frac{\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\,\mathbf{j}^\mathsf{T}}{\mathbf{j}\boldsymbol{\Sigma}\mathbf{j}^\mathsf{T} + \sigma^2}\,(y - \hat{y}), \qquad s'^2 = s^2 - \frac{(\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\,\mathbf{j}^\mathsf{T})^2}{\mathbf{j}\boldsymbol{\Sigma}\mathbf{j}^\mathsf{T} + \sigma^2}, \tag{16.13}$$

with $y \mid \mathbf{y} \sim \mathcal{N}(\hat{y},\ \mathbf{j}\boldsymbol{\Sigma}\mathbf{j}^\mathsf{T} + \sigma^2)$. The second relation does not contain y, which is why the band can be known before the reading, equation (4.73). The first has mean m under the predictive: the posterior mean is a martingale. Then $x_{95}(y) = 27 - m'(y) - \kappa s'$ is normal, and with the variance of $m'(y)$ equal to $s^2 - s'^2$ by equation (16.13),

$$\begin{aligned} \mathbb{E}[x_{95}(y) \mid \mathbf{y}] &= x_{95} + \kappa\,(s - s'), \\ \mathrm{sd}[x_{95}(y) \mid \mathbf{y}] &= \sqrt{s^2 - s'^2}, \\ \Pr(x_{95}(y) < x_{95} \mid \mathbf{y}) &= \Phi\!\left(\frac{-\kappa\,(s - s')}{\sqrt{s^2 - s'^2}}\right). \end{aligned} \tag{16.14}$$

The gain in expectation is the back-off removed, and it is never negative. The chance that the set point falls is never zero unless $s' = s$, in which case nothing moves. The rows here are not linear, so equation (16.14) is a closed form beside the refits and not the reported answer: each outer draw is a truth from the exact posterior, its reading is drawn from it, and the full graph is refit, with the five window readings held.

Building blocks. $\boldsymbol{\Sigma}$ of equation (16.8); the gradients (16.9); the candidates' lines from equation (16.6); κ; 3000 outer draws, seed 1616.

Table 16.8: The identity (16.14) against refits of the nonlinear graph, rack 3. In each pair the closed value comes first; 3000 outer draws, standard errors in parentheses.

Reading	Graph, route	s, s' [K]	$\mathbb{E}$ gain [K]	sd [K]	Pr(falls) [–]
Handheld inlet rack 3, 0.3 K	loop, fresh	0.8844, 0.4431	+0.7260, +0.7282 (0.0140)	0.7655, 0.7680	0.1714, 0.1760 (0.0070)
	loop, persistent	0.8542, 0.2980	+0.9149, +0.9160 (0.0146)	0.8005, 0.7987	0.1265, 0.1307 (0.0062)
	hold	0.8610, 0.3305	+0.8726, +0.8764 (0.0147)	0.7951, 0.8042	0.1362, 0.1450 (0.0064)
Permanent inlet, 12 windows	loop, fresh	0.8844, 0.3558	+0.8695, +0.8743 (0.0148)	0.8097, 0.8102	0.1414, 0.1423 (0.0064)
Survey β_3, 0.02	loop, fresh	0.8844, 0.3616	+0.8601, +0.8590 (0.0148)	0.8072, 0.8103	0.1433, 0.1447 (0.0064)
Reference return, 0.05 K	loop, fresh	0.8844, 0.8844	0, 0	0, 0	–, 0
	loop, persistent	0.8542, 0.8334	+0.0342, +0.0377 (0.0035)	0.1873, 0.1899	0.4276, 0.4203 (0.0090)
Load meter, 0.5 kW	loop, persistent	0.8542, 0.8492	+0.0083, +0.0086 (0.0017)	0.0927, 0.0937	0.4643, 0.4543 (0.0091)

The plant's numbers. Table 16.8 sets the closed form against the refits. Each row is a reading and a route; the columns give the two spreads, the gain, its spread and the probability of a fall, each closed and refit. The largest last Gauss–Newton step over all refits is 1.95×10^{-7}. The identity survives the nonlinear graph inside the standard errors.

It does not give the joint. The identity is rack 3's. The joint set point is close to a minimum over racks, and on the same outcomes

$$\mathbb{E}[x_{\text{joint}}(y)] - x_{\text{joint}} = (\mathbb{E}[x_{95,3}(y)] - x_{95,3}) - \mathbb{E}[x_{95,3}(y) - x_{\text{joint}}(y)] + (x_{95,3} - x_{\text{joint}}), \quad (16.15)$$

so the joint's gain is rack 3's, less the joint's mean cost after the reading, plus the cost before it. A reading that moves rack 3's mean and leaves rack 2's makes the cost after large. After the handheld the joint's gain is +0.6404 K against rack 3's +0.7282 K, a cost after of 0.0902 K; after the permanent inlet, +0.7551 against +0.8743 K, a cost after of 0.1217 K. Rack 2 binds in 13.6 % and 19.7 % of the outcomes. Tightening rack 3 hands the decision to a rack no one read. Reading both racks shrinks that cost, to 0.0515 K with permanent inlets at both, and a reading that moves no mean nearly removes it: the in-loop reference of Table 16.19 has a cost after of 0.0016 K, below the cost before of 0.0025 K, so its joint gain of +0.0846 K stands above rack 3's +0.0838 K.

Actual readings. Table 16.9 refits at given values of a handheld reading at rack 3. Before it, T_3 now stands at 24.9841°C with a predictive spread of 0.8810 K. Each row is a reading; the columns give its predictive z, the refit's $\hat{\beta}_3$, rack 3's set point and the joint set point with their changes, and the binding rack. Every row narrows the posterior to about sd β_3 0.0313 and s 0.443 K.

Table 16.9: The set point after one handheld reading at rack 3, 0.3 K, loop, fresh; the rest of the window held. The stated truth's T_3 now is 25.6933°C.

y_{T_3} [°C]	z [–]	$\hat{\beta}_3$ [–]	r_{95} rack 3, change [°C, K]	Joint r, change [°C, K]	Binds
23.2221	−2.000	0.3527	33.0041, +2.2588	32.3010, +1.5582	rack 2
24.9841	0	0.5000	31.4712, +0.7260	31.4540, +0.7113	rack 3
25.6933	+0.805	0.5592	30.8548, +0.1096	30.8531, +0.1103	rack 3
25.8651	+1.000	0.5735	30.7057, −0.0396	30.7047, −0.0380	rack 3
26.5000	+1.721	0.6265	30.1546, −0.5906	30.1546, −0.5882	rack 3

Figure 16.3 draws the whole function. The horizontal axis of both panels is the reading. The upper panel has the set point after that reading, for rack 3 and the joint, with the set point

of no reading as a dashed line. The lower panel has the predictive density of the reading, with the region where the set point falls shaded. The dotted line is the stated truth's T_3.

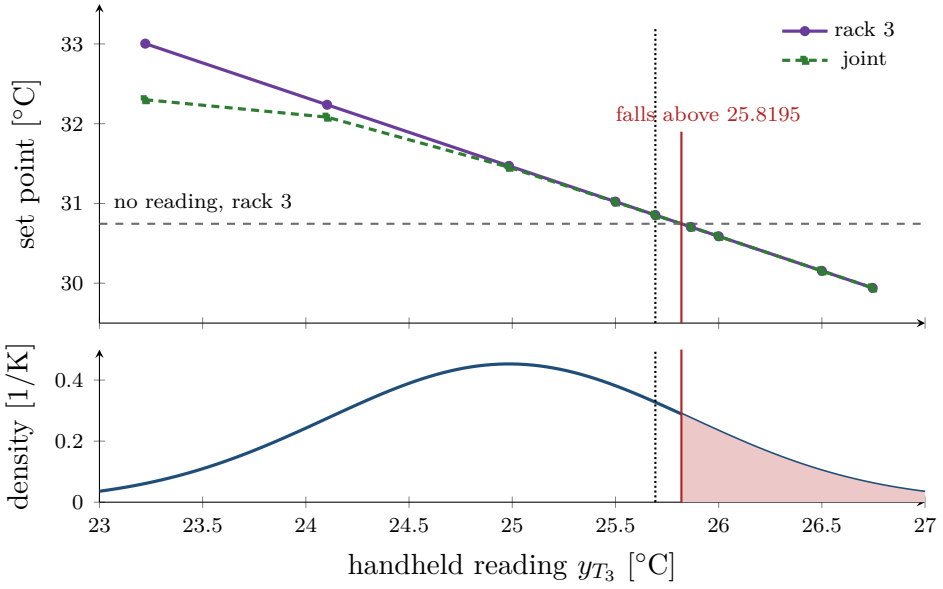


Figure 16.3: The preposterior of the set point, loop, fresh. Upper: refits at nine readings. Lower: the predictive $\mathcal{N}(24.9841,\ 0.8810^2)$. The shaded region narrows the posterior and lowers the set point. Dotted: the stated truth's T_3 now, 25.6933°C. The joint leaves rack 3 at low readings, where rack 2 binds.

Any reading above 25.8195°C**,** +0.948 **predictive standard deviations, narrows Σ and lowers the set point.** The figure is the third line of equation (16.14) seen on the reading's axis. For the joint the crossing is 25.8211°C. That crossing stands 0.126 K from the true T_3, so at the truth a 0.3 K reading lowers rack 3's set point with probability 0.3370, and the joint's with 0.3351, not the belief's 0.1714.

The reference reading. Chapter 10's analogue is a reference reading of the return at 0.05 K, here at the truth's 30.1653°C. On the persistent route it moves $\hat{b}_r$ from 0.2802 (sd 0.1935) to 0.1719 (sd 0.0484) and s from 0.8542 to 0.8334 K, and r_{95} from 31.0752 to 31.0011°C, a change of −0.0741 K: a better sensor and a lower set point. The truth's limit with its offset persisting is 31.6367°C. On the fresh route the same reading moves nothing, 30.7428°C before and after: the error at the node is a new draw, and by equation (16.13) the line on b_r has no covariance with the gradient, whose only entry is on β_3.

16.3.5 The price of a degree

Control engineer's question. What is a degree worth?

General form. The expected consequence (2.23), with the action the set point and the loss the chiller's power.

Specialized form. Equation (2.23) with $a = x$, L the chiller's power at the supply the node asks for, on the chiller row

$$P(T_s) = \frac{Q_{\text{ch}}\,(T_{\text{amb}} + T_{\text{app}} - T_s)}{\eta\,(T_s + 273.15)}, \qquad Q_{\text{ch}} = Q^* + Q_o, \tag{16.16}$$

with $\eta = 0.45$ and $T_{\text{app}} = 15\,\text{K}$ the plant's constants, and the action's value against perfect knowledge

$$\mathcal{R}(x) = \mathbb{E}\big[P(T_s^*(\mathbf{z}, e; x)) \mid \mathbf{y}\big], \qquad \Delta = \mathcal{R}(x) - \mathbb{E}\big[P(T_s^*(\mathbf{z}, 0; x_{\lim}(\mathbf{z}))) \mid \mathbf{y}\big], \tag{16.17}$$

where $x_{\lim}(\mathbf{z}) = \min_k\big[27 + (1 - \beta_k)c\big]$ is the physical limit for that state. This is where the fit's coils arrive: T_s^* is read on the outer draws, so the coil's band enters the price and not the set point. Δ is the analogue of the expected value of perfect information for this decision, and not that quantity itself: the comparator never breaches and the setting breaches 5 % of the time, so the two do not share one objective, and the assurance does part of the work. It is a ceiling on a reading's value as long as the 5 % point of $x_{\lim}$ under any posterior lies below its mean, which holds on these near-normal beliefs. The static sheet is no baseline for it.

Building blocks. The outer draws of the exact posterior; T_s^* of equation (16.2); the joint settings of Section 16.3.2; the price (16.16) at 18.00°C and 0.07 \$/kWh.

The plant's numbers. At $T_s = 14°\text{C}$ the coefficient of performance is 6.8009, and the power 30.5841 kW, or 2.1409 \$/h. The slope is $-1.7162\,\text{kW}$ per kelvin of T_s, which is 0.1201 \$/h per kelvin, and 0.3776 \$/h at 0.22 \$/kWh. At the mean T_s moves one for one with r. Table 16.10 gives the distance to the limit, over the belief and at the stated truth. Each row is a setting compared with its limit; the columns give the supply, the power above the limit, and that in \$/h at the two tariffs.

Table 16.10: The band's price at 200 kW, 18.00°C. Above the rule the difference is paired over the outer draws, se in parentheses; $\mathbb{E}[T_s]$ is that of the reference. Below it the setting runs on the truth's coils, with zero node error or with the truth's offset persisting at the node as named, against the truth's limit; the hold's pinned T_s has neither.

Reference or setting	$\mathbb{E}[T_s]$ or T_s [°C]	Above it [kW]	\$/h at 0.07, 0.22
Belief: fresh loop, all known	15.6785	2.3820 (0.0261)	0.1667, 0.5240
Belief: fresh loop, state known, e remaining	15.1927	1.5561 (0.0241)	0.1089, 0.3423
Belief: hold, all known	15.6785	2.3324 (0.0251)	0.1633, 0.5131
Truth's limit, $r = 31.4720°\text{C}$, 28.3517 kW	15.3067		
Fresh joint r, zero error	14.5774	1.2434	0.0870, 0.2735
Fresh joint r, offset persisting	14.4128	1.5250	0.1067, 0.3355
Persistent joint r, zero error	14.9079	0.6792	0.0475, 0.1494
Persistent joint r, offset persisting	14.7432	0.9602	0.0672, 0.2112
Hold	14.3078	1.7048	0.1193, 0.3750

The whole back-off is worth about **\$0.17/h at this hall at this hour**: 7.91 % of chiller power over the belief and 4.20 % at the truth. Within one world the two loops' settings stand 0.5642 kW apart with zero error and 0.5648 kW apart with the offset persisting. At an ambient of 24°C the fresh setting's distance from the truth's limit is 1.2677 kW. The sheet's r and T_s sit below the limit's power by 1.2350 and 3.4386 kW, and they sit over the limit on the truth; a saving bought that way is not a saving, and the sheet is a labelled reference and never a baseline.

What the loop changes. The operations form on this hall, the open fit with T_s pinned and moved at the node, gives T_s of 14.3108°C with $s = 0.8596\,\text{K}$, against the closed fit's hold at 14.3098°C with 0.8610 K. No number here compares with that book's plant. The loop changes four things at this decision. The coils leave the band and enter the price: the loop's T_s^* has spread 0.2217 K from the coils and 0.3731 K with e, and the curvature of P at that spread costs

0.0008 kW, which is negligible. The return sensor's error enters the band and passes to the air in full. A supply sensor's error is absorbed by the loop, while at the hold it adds 0.0024 K of back-off. And the set point now has a price, which the operations book left uncomputed.

16.3.6 Coverage on fresh records

Control engineer's question. Does the 95 % hold on fresh records?

General form. A set point chosen from the data is a function of the data, and the limit is met or not at the truth. The frequency over truths $\mathbf{z}_0$ from a distribution π and over records $\mathbf{y}$ from each is

$$C(\pi) = \iint \mathbf{1}\big[\max_k T_k^*(\mathbf{z}_0, e;\ x(\mathbf{y})) \le 27^\circ\mathrm{C}\big]\, p(\mathbf{y}, e \mid \mathbf{z}_0)\, \pi(\mathbf{z}_0)\, \mathrm{d}\mathbf{y}\, \mathrm{d}e\, \mathrm{d}\mathbf{z}_0. \tag{16.18}$$

A record is a truth, a window at $r = 30.33^\circ$C, the fit, the joint set point and then the node.

Specialized form. Equation (16.18) with three choices of π: the prior, with Q stated at 200 kW since Q has none; the exact posterior of this window; and a point mass at one truth. Under the prior, with the record's own posterior used to decide, equation (16.18) equals 0.95 for an exact posterior by the definition of a posterior; under a point mass it has no such guarantee.

Building blocks. The joint set point (16.12) refit on each record's $\hat{\mathbf{z}}$ and $\mathbf{\Sigma}$; 20 000 records for the prior, 3000 for the belief, 5000 for each truth; seed 1616.

The plant's numbers over the prior. Table 16.11 gives the result. Each row is a world, the way the error truly behaves, and a route, the way the set point treats it; the columns give the coverage and the set point's mean and spread.

Table 16.11: Coverage of all three inlets with the truth drawn from the prior, 20 000 records.

World, route	Coverage (se) [–]	x mean, sd [°C, K]
fresh, fresh	0.9491 (0.0016)	30.7428, 0
persistent, persistent	0.9507 (0.0015)	30.7871, 0.2272
persistent, fresh	0.9506 (0.0015)	30.7428, 0
fresh, persistent	0.9395 (0.0017)	30.7882, 0.2298
hold, world immaterial	0.9504 (0.0015)	15.9202, 1.5783
fresh, handheld read first	0.9493 (0.0016); rack 3 0.9597	31.3921, 0.6478
fresh, permanent read first	0.9511 (0.0015); rack 3 0.9647	31.5028, 0.6515

A set point chosen after the reading holds 95 %, over the belief it was chosen under. The fresh route holds in both worlds and gives up 0.0443 K on the mean over the prior; the persistent route falls short of 0.95 by about 6 standard errors if the error is in fact fresh. Those figures average over the prior's offsets, and this window's fit reads an offset of 0.2802 K, so at this window Table 16.5 is the sharper statement: 0.3304 K given up, and 0.8975 for the persistent setting in the fresh world. Rack 3's excess over 0.95 in the last two rows is rack 2's share of the joint. Over this window's own belief, 3000 draws, the coverage is 0.9460 with no reading, 0.9503 with the handheld and 0.9517 with the permanent inlet, with standard error 0.004.

The plant's numbers at one truth. Table 16.12 holds the truth, the same on every record, and moves its β_3, the other entries staying at the record's truth. Each row is a truth; the columns give the coverage and set point with no added reading, with the handheld and with the permanent inlet. Figure 16.4 draws the three columns. The horizontal axis is the truth's β_3, the vertical axis the coverage; the left panel is the full range and the right panel the region near 0.95. The red line is 0.95, the gray band the unconditional coverage of the first row of Table 16.11 with two standard errors, and the dotted line the record's truth.

Table 16.12: Coverage at one truth. World and route fresh; other entries at the record's truth. With no reading the setting does not move and the coverage is closed, the tail of the common e; the readings' columns are 5000 records. x in °C, its spread across records in K.

Truth's β_3 [–]	z from survey [–]	No reading: coverage, x	Handheld: coverage (se), x (sd)	Permanent: coverage (se), x (sd)
0.46	−0.50	1.0000, 30.743	0.9896 (0.0014), 31.762 (0.244)	0.9802 (0.0020), 31.932 (0.142)
0.50	0	1.0000, 30.743	0.9692 (0.0024), 31.451 (0.274)	0.9626 (0.0027), 31.592 (0.167)
0.54	+0.50	0.9997, 30.743	0.9470 (0.0032), 31.113 (0.289)	0.9496 (0.0031), 31.215 (0.178)
0.57	+0.875	0.9925, 30.743	0.9282 (0.0037), 30.852 (0.294)	0.9446 (0.0032), 30.923 (0.181)
0.60	+1.25	0.9179, 30.743	0.9092 (0.0041), 30.589 (0.296)	0.9372 (0.0034), 30.628 (0.183)
0.65	+1.875	0.3660, 30.743	0.8770 (0.0046), 30.147 (0.297)	0.9268 (0.0037), 30.136 (0.184)
0.70	+2.50	0.0190, 30.743	0.8344 (0.0053), 29.705 (0.298)	0.9152 (0.0039), 29.644 (0.186)

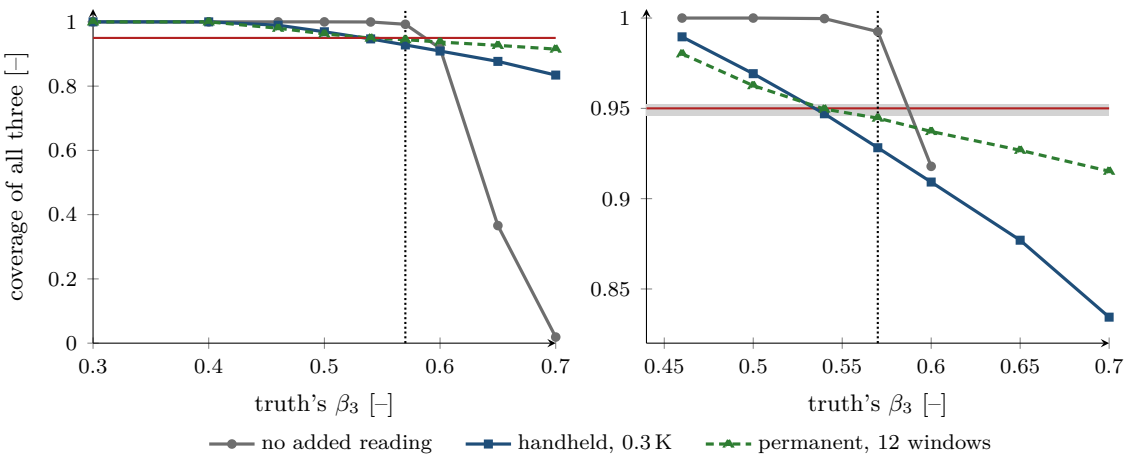


Figure 16.4: Coverage at one truth against the truth's β_3, Table 16.12 with the output's two further rows at 0.30 and 0.40. Without a reading coverage is 1 or nearly 0 according to whether the survey is right. With a reading it compresses toward 0.95 and misses on the far side of the survey. Red: 0.95. Gray band, right: 0.9491 with two standard errors. Dotted: the record's truth, 0.57.

Probability over the belief is not frequency at one truth, and a tighter posterior can lower the second. Without a reading the set point does not move, so a record's coverage is near 1 or near 0 according to whether the survey is right, and the 0.95 is the average over the prior on the survey's error. With a reading, coverage compresses toward 0.95 and misses on the far side of the survey, because the fit shrinks toward it. On the near side rack 2 also caps the setting: at a truth of 0.30 the handheld's x is 32.301°C, against rack 2's 32.3052°C.

The record's truth is one such truth, and it is above the survey. On a second seed at that truth, the handheld lowers coverage from 0.9925 to 0.9278 (se 0.0037) and the permanent inlet to 0.9394 (se 0.0034); the sweep's seed reads 0.9282 and 0.9446, so the permanent's distance from 0.95 is not firm. The two raise the supply by only 0.1131 and 0.1769 K and save 0.1927 kW,

0.0135 \$/h, and 0.3019 kW, 0.0211 \$/h, against the belief's expected 1.0919 and 1.2873 kW. The other routes at this truth with no added reading have coverage 1.0000 at $x = 30.8935$°C (persistent) and 0.9998 at 14.6386°C (hold), whereas the record's own draw gives the hold 14.3078°C.

The sensor's offset. A calibration offset d on the return sensor passes into the plant on top of a fresh node error, $T_r = r - d - e$; the window's reading carries the same d with its own fresh error. This world is neither of the two routes'. Table 16.13 gives the coverage at the truth's β_3 of 0.57. Each row is an offset; the columns give each route's coverage and set point.

Table 16.13: Coverage with a calibration offset on the return sensor and a fresh node error, truth's β_3 0.57. The fresh column is closed, its setting not moving; the persistent column is 5000 records.

d [K]	Fresh: coverage, x [°C]	Persistent: coverage (se), x [°C]
−0.3	0.9238, 30.7428	0.9232 (0.0038), 30.6245
0	0.9925, 30.7428	0.9634 (0.0027), 30.7996
+0.3	0.9997, 30.7428	0.9854 (0.0017), 30.9746

The persistent route reads 58.3 % of the offset, through the discharges and the meter; the fresh route reads none of it. A low offset takes both routes below the 95 % at this truth, and the 0.0006 between them there is well inside the persistent column's standard error of 0.0038. With the truth's β_3 at the survey's 0.50 every entry of this table is 1.0000, so this too is a statement about one truth.

What a reference does about it. A reference return reading at 0.05 K does nothing for the fresh setting, which stays at 30.7428°C. On the persistent route, at $d = -0.3$ K, it moves the coverage from 0.9232 to 0.9368 (se 0.0034) and x from 30.6245 to 30.5347°C, with a spread of 0.2896 K across records, because one reading cannot part the offset from the window's fresh error. With the offset alone that spread is 0.0471 K, x is 30.5353°C at $d = -0.3$ and 31.1196°C at $d = +0.3$ K, and the coverage is 1.0000 in all three rows run with the offset alone.

Removing the offset. Chapter 13 left to this layer whether to remove an offset once found. On the fresh route the setting does not move, so removing it moves the air by d and the choice is one of coverage against price at a stated truth. From $d = +0.3$ K removal gives the coverage back from 0.9997 to 0.9925 and returns 0.0360 \$/h at the fit. From $d = -0.3$ K it buys 0.9238 up to 0.9925 at about that price. The chapter places no value on a breach, so it states the two sides and does not choose; Chapter 17 owns that value.

16.3.7 The window of set points

Control engineer's question. When does no set point hold?

General form. The threshold form (2.12) with the indicator of an intersection, equation (6.20), read as a set of inputs at each load, as in Section 6.3.5.

Specialized form. Equation (2.12) with the supply's lower stop added to the inlets' limit,

$$\Pr\big(\max_k T_k^* \le 27°\text{C},\ 8°\text{C} \le T_s^* \le 20°\text{C} \mid \mathbf{y}\big) \ge 0.95, \tag{16.19}$$

with the closed loop of equation (16.2). Marginally the ceiling is the joint set point (16.12), and the floor is the r at which $\Pr(T_s^* \ge 8°\text{C}) = 0.95$. The all-together value is read by the draws.

Past a stop the actuator is pinned again and the plant is the hold; so the hold's own constraint at $T_s = 8°\mathrm{C}$ stands beside the loop's, an open-loop node.

Building blocks. $\hat{\mathbf{z}}$, $\boldsymbol{\Sigma}$; the supply's spread with e at each load; the draws; the stops of Chapter 11.

The plant's numbers. Table 16.14 gives the window. Each row is a load; the columns give the fresh route's floor, ceiling and width, the persistent width, the all-together edges and the best probability any r reaches, the hold's setting and the truth's width. Figure 16.5 draws it. The horizontal axis is the load and the vertical axis the set point. The solid lines are the marginal floor and ceiling, the dashed ones the all-together edges, and the dotted black lines the truth's. Vertical lines mark the closures.

Table 16.14: The window of return set points against the load. Loop, closed; hold, open. "None": no r reaches 0.95 together.

Q^* [kW]	Floor [°C]	Ceiling [°C]	Width [K]	Persistent width [K]	Together, low, up [°C]	Best Pr [–]	Hold $T_{s,95}$ [°C]	Truth's width [K]
120	18.6799	29.2215	10.5416	10.7709	18.683, 29.210	1.0000	19.1894	11.8041
200	25.0875	30.7428	5.6553	5.8350	25.093, 30.739	1.0000	14.3078	7.3067
240	28.2997	31.4927	3.1930	3.3545	28.307, 31.489	0.9999	11.8669	5.0580
280	31.5161	32.2391	0.7230	0.8692	31.590, 32.229	0.9673	9.4261	2.8092
290	32.3207	32.4253	0.1046	0.2472	none	0.9218	8.8159	2.2471
300	33.1256	32.6113	−0.5142	−0.3749	none	0.8446	8.2057	1.6849

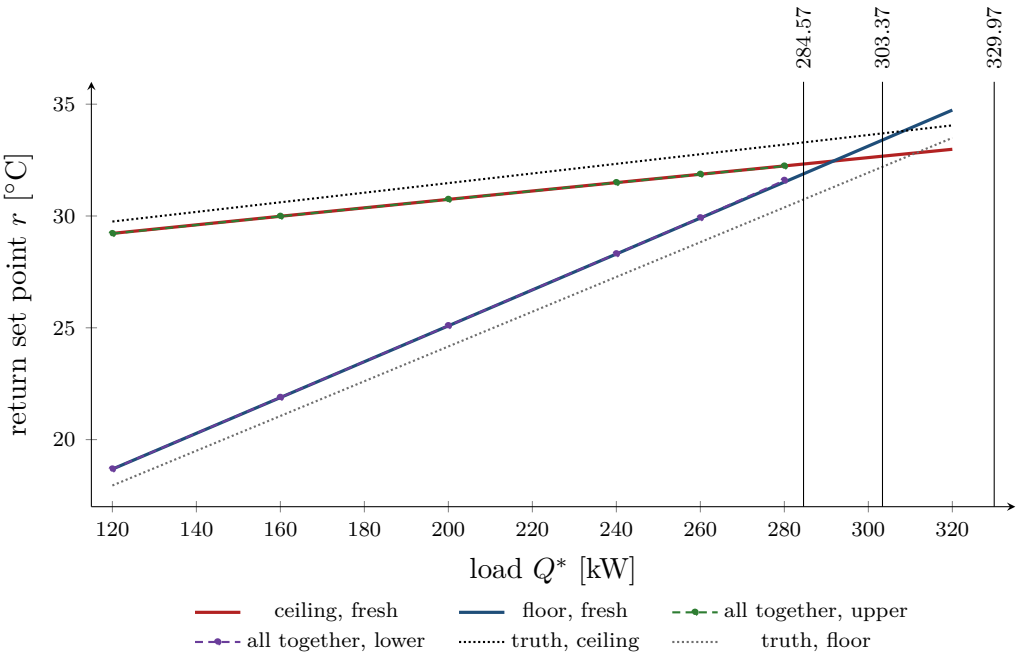


Figure 16.5: The window of set points against the load. Vertical lines, from the left: the all-together closure of the loop, 284.57 kW; the hold at $T_s = 8°\mathrm{C}$ meeting the inlets at 0.95, 303.37 kW; the stated truth's limit, 329.97 kW. The truth's dotted lines are printed to 320 kW.

The window narrows quickly: the floor rises at 0.0803 K/kW, about $1/\sum G$, and the ceiling at 0.0187 K/kW. It **closes at** 284.57 kW on the joint. The marginal edges meet later, at 291.69 kW fresh and 293.97 kW persistent, because near closure the two failures nearly add, since e pushes them opposite ways: marginal edges that meet are a joint probability near 0.90.

The hold at $T_s = 8°\text{C}$ holds 0.95 to 303.37 kW, and the truth's limit is 329.97 kW. So the loop's window closes 18.80 kW before the inlets are lost, because a loop near its stop holds nothing and the floor needs its own back-off; past it the plant is the hold at 8°C. The window does not close at low load: T_s at 20°C sets the ceiling below 113.67 kW, and the width there stays near 12 K less the band. Chapter 15's stop is the same stop in time; only steady state is read here.

16.4 The checks (step 5)

Coverage above checks the refits wherever the setting reads the record. Its first row does not: there x has spread 0, so that row is the normal integral over the prior again and tests only the arithmetic, 0.9491 against the integral's 0.95014. The checks here test the routes, the numerics and the limit of exact rows.

16.4.1 Routes and numerics

Table 16.15 lists them. Each row is a check; the columns give what is compared and the result. All pass.

Table 16.15: The checks of the chapter.

What is compared	Result
Batch fit against Chapter 11's fit	$\mathbf{z}$ gap 2.17×10^{-9}; $\mathbf{\Sigma}$ gap 5.57×10^{-9}; rows gap 1.19×10^{-14}
Newton covariance against Gauss–Newton	r_{95} 30.7452, 30.7452; $T_{s,95}$ 14.3098, 14.3099
Lattice of 144 points against 46 368, joint x	fresh 30.7428 both; persistent 31.0732 both; hold 14.3077, 14.3078
Sequential integral at the joint x against 4×10^6 normal draws, seed 1636	0.95014, 0.94988, 0.94986; se 0.00011
The tail at r_{95} for each route	0.950000 on each
Laplace against reweighted draws	Table 16.7
Identity against refits	Table 16.8

16.4.2 The limit of exact rows

Table 16.16 solves the full state with the node's entries T_c^* and T_3^* added, its physical rows at a finite tolerance, against the reduced state. Each row is a tolerance; the columns give the inlet's value at $x = 0$ and its spread, in each graph with no e. By $\sigma_p = 0.001$ the two agree at the digits printed.

Table 16.16: T_3^* at $x = 0$ on the full state with its physical rows at σ_p, and on the reduced state.

σ_p [row units]	Loop, mean [K]	Loop, sd [K]	Hold, mean [°C]	Hold, sd [K]
0.100	−5.2000	0.8439	11.2682	0.8671
0.010	−5.2000	0.8321	11.2738	0.8611
0.001	−5.2000	0.8320	11.2739	0.8610
Reduced	−5.2000	0.8320	11.2739	0.8610

16.4.3 The fit against the stated truth

The stated truth stands from the fit at +1.94 standard deviations on Q, +0.98 on g_1, +0.56 on g_2, −0.01 on β_1, −0.50 on β_2, +0.875 on β_3 and −0.60 on b_r. The true inlets at the fresh

joint r with zero error are 22.6308, 23.5668 and 26.2708°C, all under the limit, with a gap of 0.7292 K on rack 3. The truth reverses nothing: rack 3 binds there as well.

16.5 The report (step 6)

At 200 kW the joint set point is 30.74°C on the return, against 30.33°C now, or 14.31°C if the supply is held. The band is the survey's β_3, and holding all three inlets costs almost nothing today. On the fresh route a reading at rack 3 raises the joint set point in expectation by 0.6404 K (handheld) to 0.7551 K (permanent) and lowers it in 0.1750 and 0.1423 of outcomes; reading both racks raises it by up to 0.8295 K. The whole back-off is worth about \$0.17/h here at 23:00, and the choice of route about \$0.04/h of that. What the loop changes from a pinned supply is where the coils go, into the price and out of the band.

16.5.1 The answers and their bands

Table 16.17 gives the answers in the order of Table 16.1. Each row is a question; the columns give the answer, its band or what it rests on, and the stated truth.

Table 16.17: The report for the set point at 200 kW. Dollars and percentages are this example's, at 23:00.

Control engineer's question	Answer	Band or basis	Stated truth
What set point holds each inlet?	Rack 3: r 30.7452°C fresh, 31.0752°C persistent; T_s 14.3098°C hold	back-off 1.4548, 1.4050, 1.4163 K; β_3 supplies 0.6922 K^2 on each	31.4720°C; 15.3067°C
What does holding all three cost?	0.0025 K: r 30.7428°C; 31.0732°C; T_s 14.3078°C	rack 2 0.99967 at rack 3's setting	rack 3 binds
Does the normal hold?	Yes at 200 kW: draws 30.7431, 31.0774, 14.3026°C	se 0.002 K; gaps at most 0.0052 K	–
Does a tighter posterior raise the set point?	In expectation: joint +0.6404 K handheld, +0.7551 K permanent. By reading: falls above $y_{T_3} = 25.8195$°C	Pr(falls) 0.1714, refit 0.1760; rack 2 binds in 13.6, 19.7 %	Pr(falls) at the truth 0.3370
What is a degree worth?	0.1201 \$/h per K of T_s; all known 0.1667 \$/h, state known 0.1089 \$/h	2.3820 kW (se 0.0261), 7.91 %; 0.22 \$/kWh multiplies by 3.143	0.0870 \$/h, 4.20 %
Does the 95 % hold on fresh records?	Over the prior 0.9491; route wrong way round 0.9395. At one truth no	se 0.0016, 0.0017; one truth, closed and 5000 records	0.9925; handheld 0.9278
When does no set point hold?	Above 284.57 kW on the joint; hold at 8°C holds to 303.37 kW. Width at 200 kW 5.6553 K	marginal edges 291.69, 293.97 kW	329.97 kW; 7.3067 K
What next?	Inlets at racks 2 and 3, or a survey of both fractions: 0.0989, 0.0974 \$/h on the fresh route	se 0.0016, 0.0017; within se; no instrument cost stated	–

16.5.2 The entry in the commissioning log

The entry sets $r = 30.74$°C at 200 kW on the fresh route, because it holds in both worlds, and it records 31.07°C as what the persistent route would run, with the price of the choice: 0.33 K of supply and 0.0396 \$/h, against 0.8975 if the persistent setting met a fresh error at this window. It names the survey's β_3 as the band. It records in advance what a rack-3 reading would do: above 25.82°C the set point goes down, and that is not a failure of the reading. It states the 0.95 as a probability over the belief, with any calibration offset on the return sensor unseen. It gives the window to 284.57 kW, and the hold at 8°C to 303.37 kW beyond it. The sheet's

32.20°C is noted as a reference with no band, over the limit on the stated truth. Each of these rests on a constructed record.

16.5.3 Common mistakes

Table 16.18 lists six mistakes this hall makes easy. Each row names the mistake, its form here and the correction.

Table 16.18: Six mistakes in choosing a set point.

Mistake	On this hall	Correction
Saying a better reading raises the set point	$y_{T_3} = 26.50$°C narrows s to 0.4438 K and takes r_{95} down 0.5906 K	the gain is an expectation, equation (16.14); state the chance of a fall
Reading rack 3's gain as the joint's	+0.7282 against +0.6404 K	the joint is near a minimum; rank on the joint
Taking 0.95 as a frequency at this hall	0.9925 with no reading, 0.9278 with the handheld, at one truth	it is over the belief; say so in the log
Reading the offset as an error to average	fresh route at $d = -0.3$ K, 0.9238	the loop passes it to the air in full; a reference helps only with the persistent route, 0.9368, or by removing the offset, 0.9925
Comparing with the sheet's setting	1.2350 kW below the limit's power and over the limit	compare with the perfect-knowledge limit
Treating marginal edges as closure	291.69 against 284.57 kW	use the joint of the two failures

16.6 The next measurement (step 7)

A candidate is ranked here by what it would change in the plant's power at the joint set point, a rank of this question and no other.

16.6.1 The candidates by saving

Control engineer's question. What next?

General form. The preposterior (2.16), with u the realized value of the expected consequence (2.23).

Specialized form. Equation (2.16) with $\tilde{y}_a$ the candidate's reading and u the negative of $P(T_s^*)$ of equation (16.16) at the joint set point the refit would choose, on the same outer draw as the truth, compared with the set point of no reading. The preposterior covariance is equation (4.73) and the rest is the refit of Section 16.3.4. The fit's five readings are held, so these values are this window's; the prior-drawn records of Section 16.3.6 give the fleet's view.

Building blocks. The candidates' lines of Section 16.1.4; the 3000 outer draws; 0.07 $/kWh and 18.00°C.

The plant's numbers. Table 16.19 ranks the candidates on the fresh route. Each row is a candidate; the columns give the joint gain, the probability the set point falls and the share of outcomes in which rack 2 binds, then the saving by route. Figure 16.6 draws the savings. The horizontal axis is the saving in $/h and each row of bars is a candidate, one bar per route; the vertical lines are the two perfect-knowledge references of Table 16.10 for the fresh route.

Table 16.19: The candidates, joint set point, saving against no reading paired on 3000 outer draws at 0.07 $/kWh; se on the fresh saving. At 0.22 $/kWh multiply by 3.143. On the persistent route and the hold the reference reading once and in the loop coincide by construction.

Candidate	Gain [K]	Pr(falls) [–]	Rack 2 binds [–]	Fresh (se) [$/h]	Persistent [$/h]	Hold [$/h]
Permanent inlets, racks 2 and 3, 12 windows	+0.8295	0.1457	0.0920	0.0989 (0.0016)	0.1403	0.1280
Survey β_2, β_3, 0.02	+0.8170	0.1467	0.0857	0.0974 (0.0017)	0.1076	0.1049
Permanent inlet rack 3, 12 windows	+0.7551	0.1423	0.1970	0.0901 (0.0014)	0.1234	0.1131
Survey β_3, 0.02	+0.7428	0.1443	0.1963	0.0887 (0.0014)	0.0972	0.0950
Permanent inlet rack 3, own calibration 0.2 K	+0.6935	0.1640	0.1640	0.0828 (0.0014)	0.1065	0.1008
Inlets racks 2 and 3, 0.3 K, once	+0.6926	0.1747	0.0800	0.0826 (0.0016)	0.1050	0.1005
Handheld inlet rack 3, 0.3 K	+0.6404	0.1750	0.1360	0.0764 (0.0014)	0.0958	0.0915
Reference return, permanent, in the loop, 0.05 K	+0.0846	0	0	0.0094 (0.0006)	0.0046	0.0011
Inlet rack 2, 0.3 K	+0.0015	0.0183	0.0013	0.0002 (0.0000)	0.0003	0.0006
Reference return, once	0	0	0	0	0.0046	0.0011
Reference load meter, 0.5 kW	0	0	0	0	0.0011	0.0016

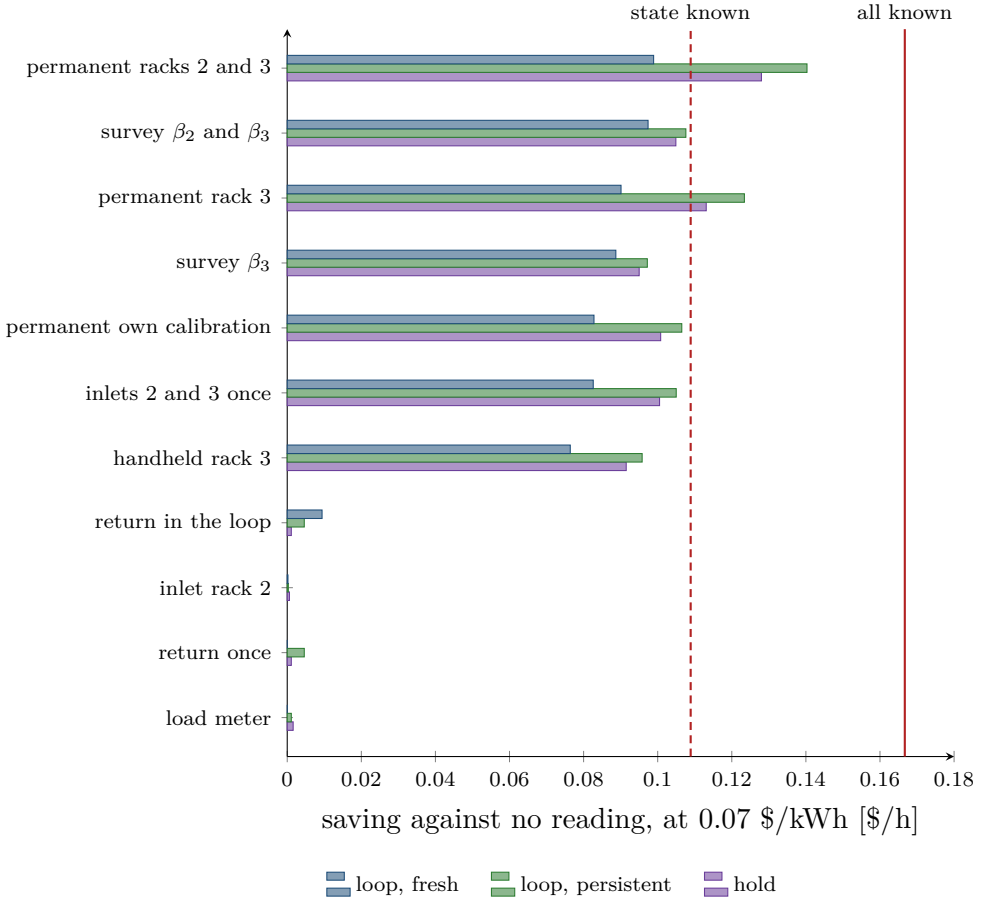


Figure 16.6: The candidates' savings, Table 16.19. Red lines: the fresh loop's perfect-knowledge references, all known 0.1667 $/h and state known 0.1089 $/h; they are fresh-route values and the other routes' bars are read against their own references.

On the fresh route the first two stand within a standard error of each other. A 0.3 K reading at rack 2 alone is worth 0.0002 $/h, nothing, and added to a 0.3 K reading at rack 3 it is worth 0.0062 $/h. **A value is not additive, and the rank is the rank of the question.** A reference return in the loop is worth little on any route, even though a return sensor's error passes in full to the air, since the fit already holds its band in e and rack 3's band is the fraction.

A loop on the inlet, beside the ranking. Moving the loop onto rack 3's inlet itself, a permanent sensor at 12 windows, with an own calibration of 0.2 K, or at 0.3 K, saves 0.1314, 0.1109 and 0.0971 $/h on the fresh comparison, with $\mathbb{E}[r_3]$ of 26.65, 26.47 and 26.35°C and coverage of 0.9473, 0.9490 and 0.9487. It removes the band from the plant's own reading and leaves rack 2 unread, which binds in 36, 29 and 26 % of outcomes. **Its stability is not computed**, so it stands beside the ranking and not in it.

An annual figure. Held all year at this node, which is arithmetic and not a forecast, the best on the fresh route is 867 $/yr at 0.07 $/kWh and 2723 $/yr at 0.22 $/kWh. No instrument cost is stated, so there is no verdict to buy.

16.6.2 The recommendation

For this question, read the fractions and not the return. Between a survey of β_2 and β_3 and permanent inlets at both racks the fresh route does not choose; the persistent route and the hold favor the permanent sensors, 0.1403 against 0.1076 $/h and 0.1280 against 0.1049 $/h, by 30.4 % and 21.9 %. Before either, the log's statement stands: a reading above about 25.82°C at rack 3 lowers the set point, and that outcome is the reading doing its work.

16.6.3 What the analysis showed

The sheet offered a raise and the promise that a better survey would raise it further. The analysis finds the raise over the limit on the stated truth, and the promise true in expectation, by an identity tested on the nonlinear graph, and false reading by reading, as the actual values show. It needs two additions: on the joint, tightening one rack and leaving the next unread hands the decision to that rack; and the expectation is over the belief, so at one truth a tighter posterior may lower both the gain and the coverage. The size of it is modest too. This hall's back-off is worth less than a tenth of its chiller power at this hour, and against a question of a few degrees of a chiller at night, a survey does about what a permanent sensor does on the fresh route, 1.5 % apart; on the persistent route and the hold the permanent sensors lead.

16.6.4 Limits of the model

One move, no future. Coupling 3, the dual effect, is gone: a set point here teaches nothing. Chapter 17 and Chapters 22 to 25 take that up.

The node. The state is steady and one load is pinned, with no forecast, which is Chapter 19's. The fit's window is held and shared by every outer draw.

The price. g_u is independent of T_s, which a record with T_s fixed cannot test, and it would move the price. There is no pump or fan power, one form of coefficient of performance and the two-hall plant's clock and tariff on this hall. The dollars and percentages are this example's.

The sensors. The twelve-window average assumes the inlet sensor's error fresh. The prior on β_3 is a survey's tolerance, tested in Section 11.4.4.

The assurance. No value is placed on a limit breach, so the 0.95 is a stated assurance; the expected cost belongs to Chapter 17.

The records. The truth is shared across the records of a row, and the plant is the model, with no misspecified row. The record is constructed from a stated truth, so its agreement tests the method on its own model and not the model against a real hall.

16.7 What to retain

Definition 16.1 (The chance-constrained set point)**.** At a scenario node the largest input with $\Pr(\max_k f_k \le f^* \mid \mathbf{y}) \ge 1-\alpha$, read through the back-off κs_k for each limit, equation (16.11), or through the joint integral, equation (16.12). The entries in the gradient are the band; on a return loop the coils are not among them.

Definition 16.2 (The preposterior of the set point)**.** On linear rows $x_{95}(y)$ is normal with mean $x_{95} + \kappa(s - s')$, spread $\sqrt{s^2 - s'^2}$ and probability of a fall $\Phi(-\kappa(s-s')/\sqrt{s^2-s'^2})$, equation (16.14).

Three conclusions carry beyond the hall as statements of form. A tighter band raises a threshold set point in expectation and may lower it on any one reading; a joint constraint lowers that gain below the single limit's when the reading tightens one limit and leaves another unread, by equation (16.15), and need not otherwise. A probability over the belief averages over truths, so at one truth a reading can move the frequency either way. And the band's worth is a price on the state's other entries, to be computed against a limit of perfect knowledge and not against a sheet. The sizes, 0.0025 K, \$0.17/h and 284.57 kW, are this hall's. The chapter's questions join the others in Appendix A.

Table 16.20 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 16.20: The terms of the chapter and where each is defined. Units for the two-unit hall.

Term	Expression	Units	Meaning	Defined in
Loop node	eq. (16.2)	°C	r pinned, T_s free	16.1.2
Hold node	eq. (16.3)	°C	T_s pinned, return free	16.1.2
Fresh, persistent error	eq. (16.4)	K	a new draw, or the entry b_r	16.1.2
Inlets' covariance	eq. (16.10)	K^2	gradients on $\mathbf{\Sigma}$, with e	16.2.2
Set point for one inlet	eq. (16.11)	°C	limit less mean less back-off	16.3.1
Binding inlet	the smallest $x_{95,k}$	–	the rack that sets the move	16.3.1
Joint set point	eq. (16.12)	°C	one event on all inlets	16.3.2
Preposterior identity	eq. (16.14)	K, –	gain, spread, chance of a fall	16.3.4
Chiller row	eq. (16.16)	kW	the price of the supply	16.3.5
Perfect-knowledge limit	eq. (16.17)	kW, \$/h	the ceiling on any reading's value	16.3.5
Coverage	eq. (16.18)	–	frequency, over a stated π	16.3.6
Window with stops	eq. (16.19)	°C	the r meeting both ends	16.3.7

16.8 Exercises

Exercise 16.1. *Formulation.* A discharge loop holds a reading at r by moving a coil, with one protected inlet downstream, $T = T_c + a$ and a an unread prior. Write the node and both error routes, and name what leaves the band. (Closed: $T^* = r - e + a$, with e a fresh draw or $e = b$, the state's offset. The coil's derate is no longer in T^*; it enters the valve, as a limit of its own or a price. The band is $\mathrm{Var}(a)$ plus σ_e^2 or Σ_{bb} with its covariance.)

Exercise 16.2. *Derivation.* Derive equation (16.14) from the update of the posterior. (By equation (16.13) $m'(y)$ is linear in y with $\mathbb{E}[m'(y)] = m$, s' is free of y, and the variance of $m'(y)$ is $(\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{j}^\mathsf{T})^2/(\mathbf{j}\boldsymbol{\Sigma}\mathbf{j}^\mathsf{T} + \sigma^2) = s^2 - s'^2$. Substituting in $x_{95}(y) = 27 - m'(y) - \kappa s'$ gives the mean $x_{95} + \kappa(s - s')$ and the spread $\sqrt{s^2 - s'^2}$; the event $x_{95}(y) < x_{95}$ is the normal tail.)

Exercise 16.3. *Derivation.* Say why the handheld's expected joint gain is below rack 3's, what can reverse that, and why a reference return reading is worth nothing on the fresh route. (The joint set point is close to $\min_k x_{95,k}(y)$, and for random quantities $\mathbb{E}[\min_k] \le \min_k \mathbb{E}$, the gap being the outcomes in which another rack binds. By equation (16.15) the cost before the reading is also removed, so a reading that moves no mean, such as the in-loop reference, can leave the joint's gain above rack 3's. On the fresh route the gradient has only a β_3 entry and the reading's line only a b_r entry, and $\boldsymbol{\Sigma}$ has no covariance between them, so $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{j}^\mathsf{T} = 0$ and $s' = s$. The error at the node is a new draw, which no earlier reading can teach.)

Exercise 16.4. *Application.* Find rack 3's set point at 200 kW on the fresh loop. (By equation (16.11), $m_3 = -(1 - 0.5000)(10.4000) = -5.2000$ K, and $s_3^2 = 0.6922 + 0.0900$, so $s_3 = 0.8844$ K and the back-off is 1.4548 K. Then $r_{95} = 27 + 5.2000 - 1.4548 = 30.7452$°C.)

Exercise 16.5. *Application.* Find the probability that the handheld lowers rack 3's set point. (By equation (16.14) with $s = 0.8844$ and $s' = 0.4431$ K: $\kappa(s - s') = 0.7260$ K and $\sqrt{s^2 - s'^2} = 0.7655$ K, a ratio of 0.948, which is the crossing of Figure 16.3 in predictive standard deviations. The normal tail below -0.948 is 0.1714.)

Exercise 16.6. *Application.* Find the fresh window's width at 200 kW. (The ceiling is the joint set point, 30.7428°C. The floor solves $\Pr(T_s^* \ge 8\text{°C} \mid \mathbf{y}) = 0.95$ with T_s^* of equation (16.2), whose mean at r is r less the 16.4738 K implied by 14.2714°C at 30.7452°C and whose spread with e is 0.3731 K: 25.0875°C. The width is $30.7428 - 25.0875 = 5.6553$ K.)

Chapter 17

Expected-Cost Set Points

17.1 A probability on the sheet and no consequence

The commissioning sheet from Chapter 16 is on the table, with the return set point at 30.74°C "at 95 %". The manager asks why 95 and not 99, and the engineer has no answer, because the sheet states a probability and no consequence. A rack alarm has never been priced. The chiller's bill is on the wall, at 0.07 $/kWh.

The engineer would conclude that 95 % is a safe, conventional figure, and that the choice between it and 99 % is a matter of taste. The sheet cannot test that. The question of this chapter is

where cost puts the set point, and what 95 % *claims a breach is worth.*

Jargon: expected-cost set point, excess, price per kelvin-hour, price per violation, critical fractile, gross and net violation price, implied price, duality gap, expected shortfall, tolerance interval

An *expected-cost set point* is the set point that minimizes the expected energy plus an expected price on the limit. The *excess* is the amount by which the hottest inlet stands above its limit, zero otherwise. A *price per kelvin-hour* charges the excess, and a *price per violation* charges the event whatever its depth. The *critical fractile* is the quantile that the ratio of two slopes picks. The *gross violation price* c_v is the price added to the energy, and the *net violation price* $c_n = c_v - c_e$ is the slope of the total past the limit, where c_e is the energy's slope. The *implied price* of a chance constraint is the price that would make its set point optimal; it is a Lagrange multiplier of the limit. A *duality gap* here means a probability that no price per violation reproduces. The *expected shortfall* is the mean depth of the tail below the critical fractile, measured from the band's center. A *tolerance interval* is the range of set points whose expected cost lies within a stated amount of the minimum.

Key idea

A price on the excess turns the limit into a quantile, with the probability of violation at the optimum equal to the energy's slope over the price. A chance constraint then asserts a price, and it can be read back. A price per violation is a different problem: its condition is on a density, it is not convex, and over a range of prices its minimum abandons the limit for the actuator's stop. The band of the posterior both moves the optimum and widens it.

17.1.1 The restriction this chapter works under

This chapter lies on the row "One move, no future" of Table 7.6, with the limit of the problem (7.1) replaced by a price on it. The fit is fixed and one move is chosen. Coupling 4 of Section 7.1.3 is

the chapter: the width of $\mathbf{\Sigma}$ enters the move through the price. Coupling 1 stays, since the move is an entry of the node's rows. Coupling 2 is gone, because no reading follows the move here, and coupling 3 is gone, because a set point teaches nothing at this node. Limits names what those two give up.

17.1.2 The two-unit hall at its node (step 0)

The plant is Chapter 16's: the two-unit hall of Chapter 11 on the reduced state (11.4), $\mathbf{z} = (Q, g_1, g_2, \beta_1, \beta_2, \beta_3, b_r)^\mathsf{T}$. The node pins $Q^* = 200\,\mathrm{kW}$ at 23:00, with Q_o the other heat, so that $c = (Q^* + Q_o)/\sum C = 10.4\,\mathrm{K}$. Ambient is 18°C. The reading's error at the node is the fresh route of equation (11.9), $e \sim \mathcal{N}(0, 0.3^2)$ in K. The joint set point of Chapter 16 is $r_{\mathrm{cc}} = 30.7428°\mathrm{C}$.

The scenario nodes. Each is named wherever it is used.

- *Closed:* r pinned, T_s free, determined by the controller's row,

$$T_r = r - e, \qquad T_s = T_r - \frac{Q^* + Q_o}{\sum G}, \qquad T_k = T_r - (1 - \beta_k)\, c. \tag{17.1}$$

- *At a stop:* T_s at 8 or 20°C, the actuator pinned, open, $T_r = T_s + (Q^* + Q_o)/\sum G$.
- *Open, the hold:* T_s pinned as a chosen value.

Define, as functions of the state,

$$R_k(\mathbf{z}) = 27 + (1 - \beta_k)\, c, \qquad R(\mathbf{z}) = \min_k R_k(\mathbf{z}), \tag{17.2}$$

the set point that puts rack k on the 27°C limit with $e = 0$, and the same for the hottest. On the closed node, $T_{\max} - 27 = r - e - R(\mathbf{z})$.

The chiller. The energy is read through the graph. With $T_{\mathrm{amb}} + T_{\mathrm{app}} = 33°\mathrm{C}$ and $\eta = 0.45$ from the plant's stated chiller, and $L = Q^* + Q_o = 208\,\mathrm{kW}$,

$$P(T_s) = \frac{L}{\mathrm{COP}(T_s)} = L\,\frac{33 - T_s}{\eta\,(T_s + 273.15)}. \tag{17.3}$$

17.1.3 The control engineer's questions (step 1)

The engineer asks nine questions. Table 17.1 lists them. The first column is the question in the engineer's words, the second the general form, the third its mathematical form on this node, the fourth its category, the fifth the computation and the last the subsection.

17.1.4 The evidence as rows and the costs (step 2)

The rows. No row is added. The evidence is the 23:00 window of Chapter 11, five readings and six priors, with the pattern (11.10). Two facts read there carry this chapter: the columns of β_2 and β_3 hold only their priors, and the return reading is consumed by the loop. The band on the hot rack is the survey's. The information matrix of the normalized rows is $\mathbf{J}^\mathsf{T}\mathbf{J}$; with unnormalized rows it is $\mathbf{J}^\mathsf{T}\mathbf{W}\mathbf{J}$.

The three costs. Each is a stated loss on the state, the action being the set point.

$$J_{\mathrm{A}}(r) = E(r) + c_v\,\mathbb{E}\big[(T_{\max} - 27)_+ \mid \mathbf{y}\big], \tag{17.4}$$

$$J_{\mathrm{B}}(r) = E(r) + C_f\,\Pr(T_{\max} > 27 \mid \mathbf{y}), \tag{17.5}$$

$$J_{\mathrm{S}}(r) = E(r) + \sum_k c_k\,\mathbb{E}\big[(T_k - 27)_+ \mid \mathbf{y}\big]. \tag{17.6}$$

Table 17.1: The questions asked of the set point's price. α is $\Pr(T_{\max} > 27 \mid \mathbf{y})$ at the set point.

Control engineer's question	General form	Mathematical form	Category	Computation	Answered in
What does a degree of set point cost?	expected consequence, (2.23)	$E(r)$, $c_e(r) = -E'(r)$	decision value	expected consequence	17.3.1
Where does a price on the excess put it?	the same, minimized, (2.23)	$\Pr(T_{\max} > 27) = c_e(r^*)/c_v$	design	expected consequence	17.3.2
Where does a price per violation put it?	the same, (2.23)	$c_e(r) = C_f\, p(r)$, or a stop	design	expected consequence	17.3.3
Which price does $\alpha = 0.05$ imply?	the inverse, (2.23)	$c_v = c_e/\alpha$, $C_f = c_e/p$	decision value	expected consequence	17.3.4
What does the band do to the optimum?	the same, across $\mathbf{\Sigma}$, (2.23)	location, curvature, interval within a tolerance	design	expected consequence	17.3.5
What is the chance constraint's error if the price is known?	paired, same draws, (2.23)	$J(r_{\mathrm{cc}}) - J(r_{\mathrm{ec}})$; perfect knowledge	decision value	expected consequence	17.3.6
Sum over racks or joint limit?	expected consequence, (2.23)	$\sum_k c_k \Pr(T_k > 27) = c_e$	design	expected consequence	17.3.7
Do the readings rank the same?	preposterior value, (2.16)	$\mathbb{E}[J(r_0) - J(r(y))]$ against the back-off's	sensor placement	preposterior	17.3.8
Does it hold on fresh records?	coverage and realized cost, (2.23)	violations against the α asked	validation	expected consequence	17.3.9

The modeling choices. Each enters here and is tested later.

- *Every consequence price is a stated choice.* None is measured, so every table runs across a range: 0.15 to 1000 \$/K h and 0.1 to 10 000 \$/event. The ranges buy a shape rather than a verdict, and they fail as a recommendation until someone states a real price.
- *An event is one hour in violation at this node.* This makes \$/event equal to \$/h in violation times one hour. It buys comparable units, and a violation with a duration or a recovery would break it.
- *The hottest inlet carries the headline.* It lets $\alpha = 0.05$ meet Chapter 16's joint constraint. Section 17.3.7 tests it against a sum over racks.
- *A consequence linear in the excess, or flat per event.* Neither is a real consequence. Each buys a closed condition, and Limits returns to them.
- *The tolerance on the optimum's width is* 0.005 *\$/h,* with 0.001 and 0.02 beside it. The width depends on it, and the location does not.

17.2 The posterior of the state (step 3)

Every answer reads one posterior, and it is Chapter 11's. Its derivation is cited here and its numbers are displayed, because each cost above is an integral against them.

17.2.1 The fit of the window

The rows (11.11), the normal equations (11.12) and the covariance (11.13) give, in the order of equation (11.4),

$$\hat{\mathbf{z}} = \begin{bmatrix} 194.6273 \\ -0.0735 \\ -0.1340 \\ 0.2204 \\ 0.3500 \\ 0.5000 \\ 0.2802 \end{bmatrix}, \qquad \mathbf{\Sigma} = 10^{-4} \begin{bmatrix} 76590 & 205.6 & 226.1 & 34.05 & 0 & 0 & -2604 \\ 205.6 & 5.719 & -2.394 & 1.638 & 0 & 0 & 10.83 \\ 226.1 & -2.394 & 6.339 & 1.644 & 0 & 0 & 10.09 \\ 34.05 & 1.638 & 1.644 & 9.881 & 0 & 0 & 23.66 \\ 0 & 0 & 0 & 0 & 64.00 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 64.00 & 0 \\ -2604 & 10.83 & 10.09 & 23.66 & 0 & 0 & 374.4 \end{bmatrix}, \tag{17.7}$$

in kW, K and dimensionless entries. At the fit $\sum G = 12.6261\,\text{kW/K}$. The derived quantity that sets the band is rack 3's limit, with gradient $\mathbf{g} = \partial R_3/\partial \mathbf{z} = (0, 0, 0, 0, 0, -c, 0)^{\mathsf{T}}$, so that with the node's error

$$s^2 = \mathbf{g}^{\mathsf{T}} \mathbf{\Sigma} \mathbf{g} + \sigma_e^2 = (10.4)^2 (0.0064) + 0.3^2, \qquad s = 0.8844\,\text{K}. \tag{17.8}$$

17.2.2 The draws of the posterior

The rows are nonlinear in the coils and the answers are nonlinear in r, so the expectations are taken on the posterior itself: Chapter 11's importance draws, a t_5 proposal about $\hat{\mathbf{z}}$, seed 11, $N = 10^6$, effective size 779 835. Each draw carries $(R_k, \sum G)$, and e is drawn beside it. The hottest limit has mean 32.1489°C and spread 0.7812 K, or 0.8368 K with e. The closed forms that use s of equation (17.8) stand beside these and are named where they do.

Every draw's 20°C stop lies past its limit, a share of 1.00000000, with a mean margin of 4.3282 K. So a violation can be reached at the top stop on the whole posterior. That fact returns in Section 17.3.3.

17.3 The answers (step 4)

Each answer is an expected consequence over the draws of Section 17.2.2. The last two questions add outer draws of the same posterior.

17.3.1 What a degree of set point costs

Control engineer's question. What does a degree of set point cost?

General form. The expected consequence (2.23), $\mathcal{R}(a) = \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}]$.

Specialized form. Equation (2.23) with $a = r$, on the closed node (17.1), and L the chiller's bill per hour through equation (17.3):

$$\begin{aligned} E(r) &= 0.07\,\mathbb{E}\big[P(T_s(r, \mathbf{z}, e)) \mid \mathbf{y}\big], \\ c_e(r) = -E'(r) &= 0.07\,\mathbb{E}\Big[\frac{L\,(33 + 273.15)}{\eta\,(T_s + 273.15)^2}\;\mathbf{1}[8 < T_s < 20] \;\Big|\; \mathbf{y}\Big]. \end{aligned} \tag{17.9}$$

The choice that enters is an electrical price with no demand charge and no pump or fan power: it buys a one-line energy, and Limits names what it leaves out.

Building blocks. The draws of Section 17.2.2; the node (17.1); the power (17.3). Since $P'' > 0$ on every draw, E is convex.

The plant's numbers. Table 17.2 gives E and its slope. Each row is a set point, and the last column is E less its tangent at r_{cc}. Figure 17.1 plots the same curve. Its horizontal axis is the set point and its vertical axis the expected bill; the solid curve is $E(r)$, the dashed line the tangent at r_{cc}, and the dotted line the bill at the 20°C stop, 1.4348 \$/h.

Table 17.2: The energy against the set point, closed loop, over the posterior; open at the stop for the share of draws there.

r [°C]	$\mathbb{E}[T_s]$ [°C]	$\mathbb{E}[P]$ [kW]	E [\$/h]	c_e [\$/h per K]	At the stop [–]	E less tangent [\$/h]
26.00	9.5225	38.3909	2.6874	0.12397	0.0000	+0.00954
30.33	13.8525	30.8381	2.1587	0.12026	0.0000	+0.00007
30.7428	14.2653	30.1300	2.1091	0.11991	0.0000	0
32.20	15.7225	27.6463	1.9352	0.11871	0.0000	+0.00088
35.00	18.5225	22.9437	1.6061	0.11643	0.0000	+0.00745
36.00	19.5048	21.3151	1.4921	0.10410	0.1003	+0.01336
37.00	19.9861	20.5205	1.4364	0.00938	0.9187	+0.07766
38.00	20.0000	20.4977	1.4348	0.00000	1.0000	+0.19597

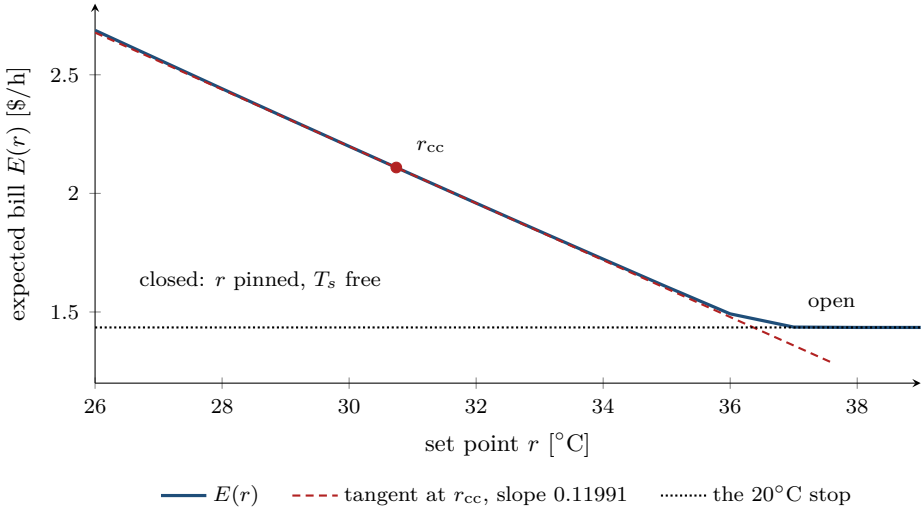


Figure 17.1: The energy of equation (17.9) from Table 17.2's run. Solid: $E(r)$ over the posterior. Dashed: the tangent at r_{cc}. Dotted: the stop, where the actuator is pinned and the loop is open.

The energy is nonlinear, and nearly linear where the decision lives. The slope at r_{cc} is 0.11991 \$/h per K, and over 28 to 33°C it falls 3.42 %, from 0.12223 to 0.11805. The tangent departs by 0.00745 \$/h at 35°C. Central differences match the analytic slope to the digits printed. At 0.22 \$/kWh the slope is 0.37687, and at ambient 24°C it is 0.12226. Past about 36°C the draws reach the stop. The stop's bill is 0.6743 \$/h below r_{cc}'s, and the slope just inside it is 0.11527.

17.3.2 A price on the excess

Control engineer's question. Where does a price on the excess put it?

General form. The action that minimizes the expected consequence (2.23).

Specialized form. Equation (2.23) with the loss of equation (17.4). On the closed node $(T_{\max} - 27)_+ = (r - e - R)_+$, whose derivative in r is $\mathbf{1}[T_{\max} > 27]$ on a draw not at a stop, so while no draw is at a stop

$$J'_{\mathrm{A}}(r) = -c_e(r) + c_v \Pr(T_{\max} > 27 \mid \mathbf{y}). \tag{17.10}$$

Both terms of J_{A} are convex on the live region, the check of Section 17.4.1, and the root is the minimum:

$$\boxed{\Pr(T_{\max} > 27 \mid \mathbf{y})\big|_{r^*} = \frac{c_e(r^*)}{c_v}}. \tag{17.11}$$

This is a critical fractile, with c_v the gross price, added to an energy that keeps falling past the limit. It holds only while no draw is at a stop at r^*, since a draw at the stop moves neither its energy nor its excess.

The same condition in net form. Below the limit the total falls at c_e per kelvin. Past it the total rises at $c_n = c_v - c_e$. Then $-c_e(1-\alpha) + c_n\alpha = 0$ gives

$$\alpha = \frac{c_e}{c_e + c_n} = \frac{c_e}{c_v}. \tag{17.12}$$

The form $c_e/(c_e + c_v)$ therefore holds only if c_v means the net slope. The two statements are one condition, and the chapter uses the gross price and prints c_n beside it. Reading the gross price into the net form gives 0.047584 at 2.4001 \$/K h against the measured 0.049961, and 0.322082 at 0.25 against 0.475106.

The closed form beside the posterior takes rack 3 normal with spread (17.8) and the slope frozen at r_{cc}: $r^* = 32.2 - s\,\Phi^{-1}(1 - c_e/c_v)$.

Building blocks. The draws; c_e of equation (17.9); the excess at each draw from equation (17.2). The minimum is searched on the raw draws and the root (17.11) is solved separately, so that each tests the other.

The plant's numbers. Table 17.3 runs the price. Each row is a price; the columns give the net price, the numerical minimum and the root, the set point with the slope frozen, the probability at the root with its standard error, the slope over the price, and the optimum's energy, excess cost and total. Figure 17.2 shows how the minimum and the tail meet. The upper panel plots J_{A} against r at three prices. The lower panel plots $p(r)$, the density in r of $R + e$, so that the area to the left of any set point is the probability of violation there; the shaded area is the tail at the implied price.

Table 17.3: Form A, on all the draws, closed loop. The price 2.4001 is the one Section 17.3.4 reads from r_{cc}.

c_v [\$/K h]	c_n [\$/K h]	r^*, min [°C]	r^*, root [°C]	Frozen [°C]	α (se) [–]	$c_e(r^*)/c_v$ [–]	E [\$/h]	Excess cost [\$/h]	J^* [\$/h]
0.15	0.0318	32.8272	32.8272	32.8595	0.787949 (0.000465)	0.787948	1.8610	0.1159	1.9769
0.25	0.1312	32.1144	32.1144	32.1243	0.475106 (0.000565)	0.475106	1.9454	0.0795	2.0249
1	0.8804	31.1520	31.1520	31.1536	0.119571 (0.000366)	0.119572	2.0601	0.0515	2.1116
2.4001	2.2802	**30.7428**	**30.7428**	30.7427	0.049961 (0.000240)	0.049961	2.1091	0.0440	2.1531
10	9.8796	30.2102	30.2102	30.2089	0.012035 (0.000110)	0.012036	2.1731	0.0370	2.2101
100	99.8791	29.5199	29.5199	29.5173	0.001209 (0.000028)	0.001209	2.2564	0.0300	2.2864
1000	999.8786	28.9502	28.9502	28.9501	0.000121 (0.000006)	0.000121	2.3254	0.0255	2.3509

The figure shows the three minima lying on the rising side of the density, and the dearer the price the farther into the tail. The curves lie together on the left, where nothing violates and only the energy counts, and part where the tail begins.

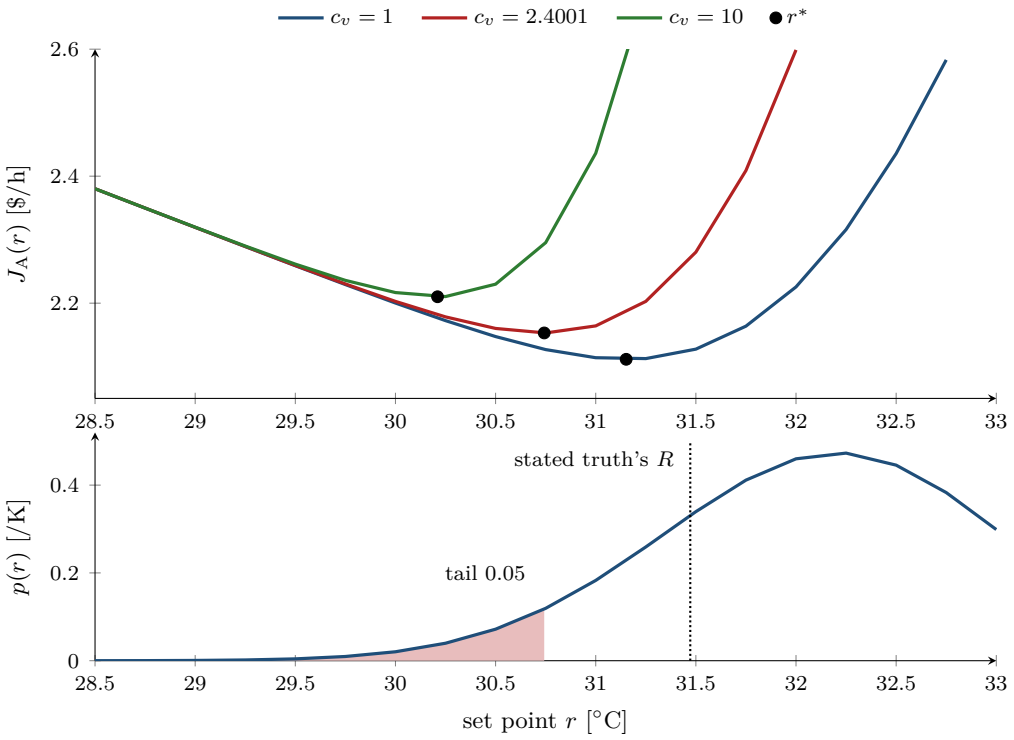


Figure 17.2: The minimum and the tail, closed loop, on all the draws. Upper: J_A of equation (17.4) at three prices, each with its minimum. Lower: the density in r of $R + e$, equation (17.13); the shaded area left of r_{cc} is the probability of violation there. Dotted: rack 3's limit at the stated truth, 31.4720°C, with $e = 0$.

Over the rows the minimum and the root differ by less than 0.00001 K, and α and $c_e(r^*)/c_v$ by 0.000001. **With the slope taken at r^*, the fractile holds on the nonlinear energy to the draws' granularity.** With the slope frozen it is off by up to 0.0323 K at the lowest price and 0.0026 K at 100. The closed form gives 30.7449, 30.2035 and 29.5150°C at the implied price, 10 and 100. The tail resolves to about $\alpha = 10^{-5}$ on these draws, so the last rows are read with their standard errors.

The form's own runaway. Below the slope at the stop, 0.11527 $/K h, a kelvin-hour costs less than the energy it saves on every draw. So J_A falls all the way, and the minimum has every draw at $T_s = 20$°C, open, with $\alpha = 1.00000$: 1.91094 $/h at $c_v = 0.11$ and 1.65125 $/h at 0.05. Equation (17.11) has no meaning there, because its probability counts draws whose excess no longer moves. Just above that slope the minimum is interior, on the stop's shoulder: at 0.116 it stands at 35.5571°C with 0.0067 of the draws at the stop, at 1.93662 $/h against the stop's 1.93691 $/h. At 0.12 and 0.13 it stands at 33.7837 and 33.2424°C.

The other two nodes. At 2.4001 $/K h the closed node on the persistent route, $e = b_r$, gives $r^* = 31.0770$°C with $\alpha = 0.049943$. The open node, the hold, gives $T_s^* = 14.3022$°C with $\alpha = 0.049947$. The fractile (17.11) holds on each, with the node's own slope.

17.3.3 A price per violation

Control engineer's question. Where does a price per violation put it?

General form. The action that minimizes the expected consequence (2.23), whose second line is the certain cost plus a consequence times a probability.

Specialized form. Equation (2.23) with the loss of equation (17.5), $c_a = E(r)$ and $C_e = C_f$. The penalty is differentiated with e integrated by its normal, which removes the step on the draws:

$$\begin{aligned} \Pr(T_{\max} > 27 \mid \mathbf{y}) &= \mathbb{E}\Big[\Phi\Big(\frac{r - R(\mathbf{z})}{\sigma_e}\Big) \;\Big|\; \mathbf{y}\Big], \\ p(r) = \mathbb{E}\Big[\frac{1}{\sigma_e}\varphi\Big(\frac{r - R(\mathbf{z})}{\sigma_e}\Big) \;\Big|\; \mathbf{y}\Big], &\qquad J'_{\mathrm{B}}(r) = -c_e(r) + C_f\, p(r). \end{aligned} \tag{17.13}$$

The condition is on a density, not on a tail. Since c_e is nearly flat and p has one turning point, the equation $c_e(r)/C_f = p(r)$ has no roots or two. The lower root r_{low} is a local minimum and the upper root r_{high} a local maximum. Beyond r_{high} the total falls again, because a violation already certain costs no more, and it falls as far as the 20°C stop, open, where

$$J_{\mathrm{B}}(\text{stop}) = E_{\text{stop}} + C_f \underset{\text{stop}}{\Pr}. \tag{17.14}$$

So J_{B} is not convex, and the stationary point has to be compared with the stop. Two prices mark the range: C_1, where the two roots appear at a tangent, and C_2, where $J_{\mathrm{B}}(r_{\text{low}})$ ties the stop.

Building blocks. The draws; c_e; the density (17.13); $E_{\text{stop}} = 1.4348$ \$/h and $\Pr_{\text{stop}} = 1.000000$ from Section 17.2.2.

The plant's numbers. The density has its mode at 32.2061°C, at 0.47355 per K; the normal beside it gives 0.47674 per K at 0.8368 K. On a 0.02 K grid over six spreads to either side it has one turning point. Table 17.4 runs the price. Each row is a price; the columns give the local minimum, its probability and total, the local maximum, the total at the stop and which one is the minimum. Figure 17.3 shows the two conditions. The upper panel plots J_{B} against r at four prices, with the stop as the flat end of each curve. The lower panel plots $p(r)$ with c_e/C_f as near-horizontal lines, whose crossings are the stationary points.

Table 17.4: Form B, on all the draws, e integrated in the penalty. An event is one hour in violation.

C_f [\$/event]	r_{low} [°C]	α [–]	$J(r_{\text{low}})$ [\$/h]	r_{high} [°C]	J(stop) [\$/h]	The minimum
0.20	none	–	–	none	1.6348	the stop, open
$C_1 = 0.2507$	none	–	–	none	1.6855	the stop, open (tangent)
0.30	31.6951	0.290237	2.0824	32.7122	1.7348	the stop; local minimum loses
0.65	31.0043	0.089123	2.1357	33.3368	2.0848	the stop; local minimum loses
$C_2 = 0.7056$	30.9526	0.079967	2.1404	33.3820	2.1404	tie
1.0188	30.7428	0.050007	2.1601	33.5640	2.4536	interior, closed
5	30.0600	0.007692	2.2296	34.1428	6.4348	interior, closed
100	29.1619	0.000292	2.3290	34.8811	101.4348	interior, closed
10 000	28.1421	0.000002	2.4467	35.6454	10 001.4348	interior, closed

The upper panel shows the local minimum above the stop for $C_f = 0.5$, level with it at C_2 and below it at 1.0188. Every curve turns down past its local maximum and runs to the stop. In the lower panel each dashed line crosses the density twice.

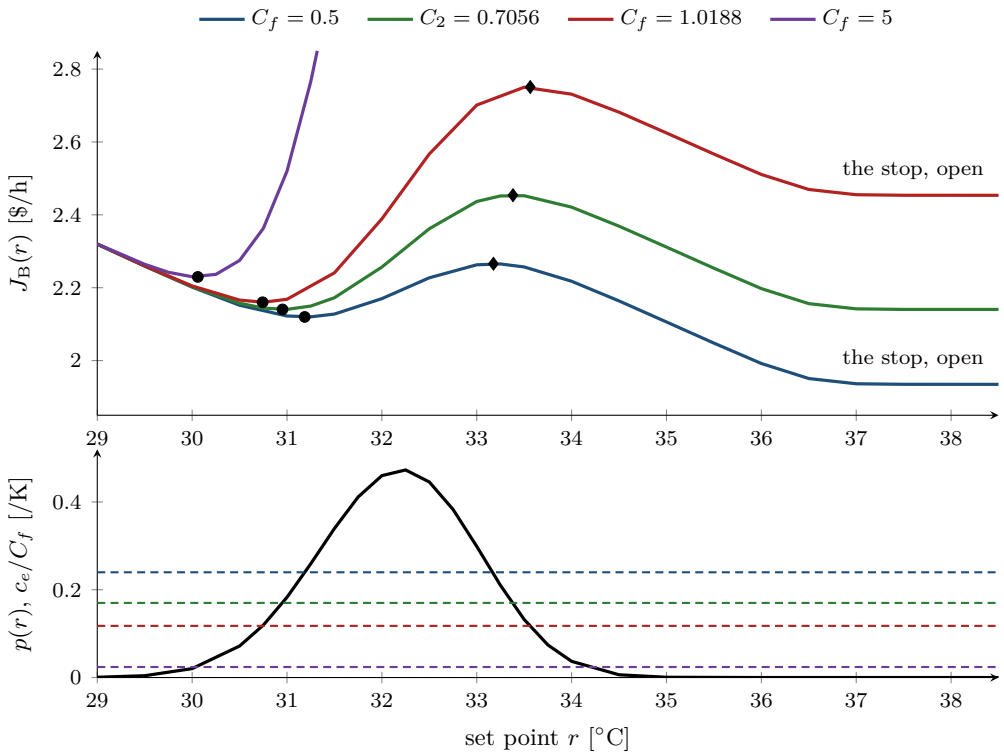


Figure 17.3: The density condition (17.13), on all the draws. Upper: J_B, closed loop until the draws reach the stop, flat at the stop. Circles: local minima. Diamonds: local maxima. The curve at $C_f = 5$ leaves the panel; its stop is at 6.4348 \$/h. Lower: $p(r)$, with dashed lines at c_e/C_f in each curve's color, c_e taken at r_{cc}; crossings are the stationary points.

Below $C_2 = 0.7056$ \$/event the minimum is the stop: $T_s = 20$°C, open, $\mathbb{E}[T_3] = 31.28$°C, violating with probability 1.000000. The penalty has no violation guard in that range. At C_2 the energy saved against r_{low}, 0.6491 \$/h, equals the added penalty $C_2(1-\alpha)$, and below C_2 the saving wins. **No fixed price gives an interior minimum with α above** 0.07997**.** A chance constraint at $\alpha = 0.10$ has a set point, 31.0606°C, and no C_f reproduces it as a global minimum: that is the duality gap of this form. The other boundary is out of reach: by the normal closed form, since the draws do not get there, the 8°C stop needs C_f of order 10^{16} \$/event.

A raw check on 50 000 draws, with a sharp violation and a global grid of 0.05 K, finds the stop at $C_f = 0.5$ and 0.7, and 30.65, 30.00 and 29.25°C at 1.0188, 5 and 100, against the roots 30.7428, 30.0600 and 29.1619°C. The gaps are that subset's noise on a step function. It confirms which basin holds the minimum, not the digits.

17.3.4 The price an α implies

Control engineer's question. Which price does $\alpha = 0.05$ imply?

General form. The expected consequence (2.23) read backward: given the action, the consequence under which it is the minimizer.

Specialized form. Equation (2.23) through its stationary conditions (17.11) and (17.13), solved for the price at the r_α with $\Pr(T_{\max} > 27 \mid \mathbf{y}) = \alpha$:

$$c_v(\alpha) = \frac{c_e(r_\alpha)}{\alpha}, \qquad C_f(\alpha) = \frac{c_e(r_\alpha)}{p(r_\alpha)}. \tag{17.15}$$

Each is a Lagrange multiplier. Minimizing $E(r)$ subject to $\Pr(T_{\max} > 27) \le \alpha$ has the Lagrangian $E + \lambda(\Pr - \alpha)$ and the stationary condition $c_e = \lambda p$, so $\lambda = C_f$. In the same way c_v is the multiplier of a limit on the expected excess at that r. A multiplier reports what the constraint charges at the margin, and the status of the point under form B has still to be checked against the stop.

Building blocks. c_e and p of equations (17.9) and (17.13); the draws' quantile r_α; the stop (17.14). The closed forms beside them are $c_v = c_e/\alpha$ and $C_f = c_e s/\varphi(\Phi^{-1}(1-\alpha))$, with the frozen slope and s of equation (17.8).

The plant's numbers. Table 17.5 reads the price from each α. Each row is a probability; the columns give the set point, the price per kelvin-hour, the mean excess given a violation, the price per violation, and the status under form B. At r_{cc}, $\Pr = 0.04996$ on the draws, which is why that row reads 2.4001 and the 0.05 quantile of the draws, at 30.7431°C, reads 2.3983 and 1.0181. The closed form gives 2.3983 and 1.0283. The draws' own band, two standard errors at r_{cc}, is 2.377 to 2.423 \$/K h and 1.013 to 1.025 \$/event, so the digits past the third are the seed's. At 0.22 \$/kWh the prices are 7.5433 and 3.2019, and at ambient 24°C they are 2.4472 and 1.0388. All of J scales with the tariff at a scaled price.

The size of the price. With no outside price, the hall's own money is the yardstick. The whole hall runs 238.13 kW, or 16.6691 \$/h. **At this node at 23:00, $\alpha = 0.05$ says an hour in breach is worth** 1.02 **\$, which is** 3.7 **minutes of the hall's bill, and a kelvin-hour** 2.40 **\$, which is** 34.29 **chiller-kWh or** 8.6 **minutes of the bill.** Held all year at this node, as stated arithmetic, that is an expected 438 hours in violation and 386 to 446 \$ a year. An expectation is not a typical year, because most of the band is a fixed β_3 that does not redraw

Table 17.5: The price each α implies, closed loop, on all the draws.

α [–]	r_α [°C]	c_v [\$/K h]	$\mathbb{E}$[excess given violation] [K]	C_f [\$/event]	Under that C_f
0.0001	28.9164	1215.22	0.2177	267.99	global minimum
0.001	29.4729	121.07	0.2482	31.749	global minimum
0.01	30.1486	12.041	0.3022	3.9310	global minimum
0.05 (at r_{cc})	**30.7428**	**2.4001**	0.3671	**1.0188**	global minimum
0.08	30.9523	1.4967	0.3971	0.7059	global minimum, the edge
0.10	31.0606	1.1965	0.4149	0.5968	local minimum; the stop wins
0.50	32.1677	0.2375	0.6876	0.2510	local minimum; the stop wins
0.70	32.6044	0.1691	0.8667	0.2798	a local maximum

each hour. With the belief held and e fresh each hour, the chance of more than 1 hour in the year is 0.3438, of more than 438 hours 0.1249, and of more than 4380 hours 0.0400. Half of the expected hours come from the worst 0.0289 of the belief.

The implied C_f is only 1.44 times C_2, the price below which the optimum abandons the limit. That ratio, the 0.07997 and the status "global minimum" all rest on the 20°C stop. With the draws held and the stop alone changed, 18, 19 and 22°C give $C_2 = 0.4458$, 0.5783 and 0.9506 \$/event, with α at C_2 of 0.15065, 0.10432 and 0.05454, and a stop at 22.57°C would tie the 95 % set point itself.

Whether that is sensible depends on what a breach does, which this model does not contain, and it is the owner's to state. What the computation says is narrower: the price is small, it is set by the energy's slope and not by any consequence, and anyone who would pay more than about 1 \$ to avoid an hour over 27°C at this node is not at 95 %.

17.3.5 What the band does to the optimum

Control engineer's question. What does the band do to the optimum?

General form. The minimizer of the expected consequence (2.23), read as the posterior's width varies.

Specialized form. Equation (2.23) with the loss (17.4) at the fixed price 2.4001 \$/K h, on a posterior scaled by f: every draw moved toward the mean, $\mathbf{z} \mapsto \bar{\mathbf{z}} + f(\mathbf{z} - \bar{\mathbf{z}})$, $e \mapsto fe$, means held. The chosen scaling buys a one-parameter family of beliefs; it is a stated what-if and no reading produces it. Three quantities are read: the location $r^*(f)$; the curvature, by differentiating equation (17.10),

$$J''_{\mathrm{A}}(r^*) = c_v\, p(r^*) + E''(r^*); \tag{17.16}$$

and the tolerance interval $\{r : J_{\mathrm{A}}(r) \le J^* + \mathrm{tol}\}$. Beside the posterior stand the closed forms of a normal with a frozen slope, with $\kappa = \Phi^{-1}(1-\alpha)$:

$$r^*(0) - r^*(s) = s\kappa, \qquad J'' = \frac{c_v\, \varphi(\kappa)}{s}, \qquad \text{width} \approx 2\sqrt{\frac{2\,\mathrm{tol}}{J''}} \propto \sqrt{s}. \tag{17.17}$$

As $s \to 0$ the curvature grows without limit and the quadratic estimate goes to zero. The true interval does not, because the cost itself has a kink at the limit, with slopes $-c_e$ below and $c_v - c_e$ above:

$$\mathrm{width}(0) = \frac{\mathrm{tol}}{c_e} + \frac{\mathrm{tol}}{c_v - c_e}. \tag{17.18}$$

Building blocks. All the draws, scaled, so r^* at $f = 1$ is the root of Table 17.3; the density and the slope of the earlier answers; $c_e = 0.11871$ at the no-band set point.

The plant's numbers. Table 17.6 runs f. Each row is a scale; the columns give the spread of $T_{\max}$, the optimum, how far it moved, its total, the curvature, the quadratic width, and the interval's width with its parts below and above r^*. Figure 17.4 plots the two effects against the spread. The left panel has r^* with the tolerance interval's edges, and the right panel has the width with the quadratic estimate and the kink's floor.

Table 17.6: Form A at the fixed price 2.4001 \$/K h across the scaled posterior; tolerance 0.005 \$/h; all the draws.

f [–]	sd $T_{\max}$ [K]	r^* [°C]	Moved [K]	J^* [\$/h]	J'' [\$/h per K^2]	Quadratic width [K]	Width (below, above) [K]
0	0	32.1998	0	1.9354	kink	0	**0.0443** (0.0421, 0.0022)
0.01	0.0088	32.1852	−0.0146	1.9376	27.87	0.0379	0.0591 (0.0458, 0.0133)
0.10	0.0884	32.0539	−0.1459	1.9571	2.790	0.1197	0.1265 (0.0753, 0.0513)
0.25	0.2211	31.8353	−0.3645	1.9897	1.118	0.1892	0.1934 (0.1083, 0.0851)
0.50	0.4409	31.4714	−0.7284	2.0440	0.5606	0.2671	0.2702 (0.1467, 0.1235)
1	0.8368	30.7428	**−1.4570**	2.1531	0.2833	0.3757	**0.3779** (0.2007, 0.1772)
2	1.5467	29.2140	−2.9859	2.3789	0.1489	0.5183	0.5196 (0.2705, 0.2491)
3	2.2188	27.5989	−4.6009	2.6164	0.1028	0.6238	0.6297 (0.3247, 0.3050)

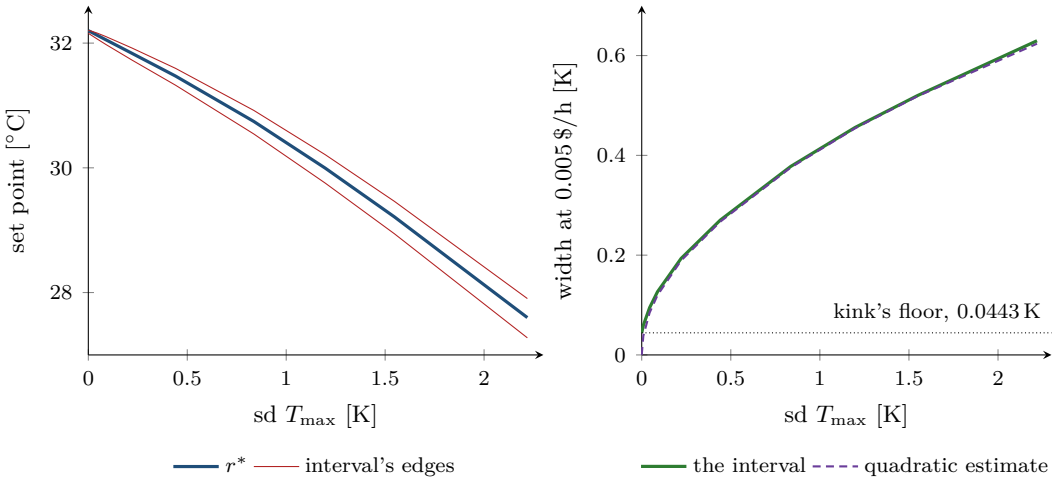


Figure 17.4: The two effects of the band, form A at 2.4001 \$/K h, closed loop. Left: the optimum falls linearly in the spread, with the interval's edges around it. Right: the interval widens as a square root, from the kink's floor (17.18) rather than from zero.

The band moves the optimum, linearly in s, away from the limit: 1.46 K at $f = 1$, against 1.4548 K in closed form. **The band also widens it, as the square root of s at the wide end**, since by equation (17.17) the curvature falls as $1/s$; the closed form gives $J'' = 0.2799$ against 0.2833. It does not start from zero: equation (17.18) gives $0.0421 + 0.0022 = 0.0443$ K. So "the band widens the optimum" is true of a tolerance interval, but at $f = 1$ the move is about four times the width and is what the decision feels. The kink is asymmetric, and the band makes the interval nearly symmetric. At $f = 1$ the tolerances 0.001 and 0.02 \$/h give 0.1682 and 0.7683 K. The price barely changes the kink's width, 0.0478, 0.0443, 0.0426 and 0.0422 K at $c_v = 1$, 2.4001, 10 and 100, and it changes the band's modestly, 0.4203, 0.3779, 0.3348 and 0.2977 K.

Form B across the band. On the first 100 000 draws, at the fixed 1.0188 \$/event, r_{low} reads 32.1960 (a jump above it), 31.3001, 30.7334 and 29.9545°C at $f = 0$, 0.5, 1 and 2, with widths 0.0421, 0.2713, 0.4228 and 0.6659 K. Unlike form A, α rises with the band at a fixed price, 0.021, 0.050 and 0.116 at $f = 0.5$, 1 and 2. Between $f = 2.5$ and 3 the interior minimum loses to the stop, 2.3934 against 2.4534 and then 2.4588 against 2.4524. A wider band at a fixed penalty can remove the optimum altogether.

17.3.6 The chance constraint when the price is known

Control engineer's question. What is the chance constraint's error if the price is known?

General form. The difference of the expected consequence (2.23) at two actions, on the same draws; and the same consequence under perfect knowledge, the physical limit.

Specialized form. Equation (2.23) with L of equation (17.4), closed loop, and the paired difference

$$\Delta(c_v) = J(r_{\text{cc}}) - J(r_{\text{ec}}) = \mathbb{E}\big[L(r_{\text{cc}}, \mathbf{z}, e) - L(r_{\text{ec}}, \mathbf{z}, e) \mid \mathbf{y}\big], \tag{17.19}$$

with r_{ec} the root (17.11) and its standard error taken on the pairs. The reference is perfect knowledge, $\mathbb{E}[\min_r L(r, \mathbf{z}, e) \mid \mathbf{y}]$ with everything known, and the same with $\mathbf{z}$ known and e remaining, integrated by a 41-node Gauss–Hermite rule. The closed form beside it takes the normal with a frozen slope: the cost of the band at the optimum, above the no-band limit, is

$$c_e \, \frac{s \, \varphi(\kappa)}{\alpha}, \tag{17.20}$$

c_e times the expected shortfall $s\varphi(\kappa)/\alpha$; the energy of the back-off alone is $c_e \kappa s$.

Building blocks. The roots of Table 17.3; all 10^6 draws for the gap; 20 000 draws resampled from the posterior, seed 1717, for the reference.

The plant's numbers. Table 17.7 runs the gap. Each row is a price; the columns give the expected-cost set point, its probability and total, the chance-constrained total, and the gap. Figure 17.5 plots the gap. In the left panel both axes are logarithmic, the horizontal axis is the price in each form's own units, and the dotted line marks 0.01 \$/h. The gap is zero at each form's implied price, which a logarithmic axis cannot show. The right panel draws the gap at the stated truth, on a linear vertical axis.

Table 17.7: Expected-cost against chance-constrained, paired on 10^6 draws, closed loop; α at r_{cc} is 0.04996.

c_v [\$/K h]	r_{ec} [°C]	α_{ec} [–]	$J(r_{\text{ec}})$ [\$/h]	$J(r_{\text{cc}})$ [\$/h]	Gap (se) [\$/h]
0.25	32.1144	0.47511	2.0249	2.1137	+0.08882 (0.00012)
1	31.1520	0.11957	2.1116	2.1274	+0.01585 (0.00011)
2.4001	30.7428	0.04996	2.1531	2.1531	0
5	30.4502	0.02403	2.1837	2.2008	+0.01711 (0.00027)
10	30.2102	0.01204	2.2101	2.2925	+0.08242 (0.00081)
100	29.5199	0.00121	2.2864	3.9430	+1.65664 (0.01048)
1000	28.9502	0.00012	2.3509	20.4481	+18.09724 (0.10664)

The regret is mild near the implied price and asymmetric far from it. The set point for 1000 \$/K h stands 1.79 K below r_{cc}. Held there when the true price is 2.40, it costs 0.1723 \$/h more, +0.2163 in energy and −0.0440 in excess. Held at r_{cc} when the true price is

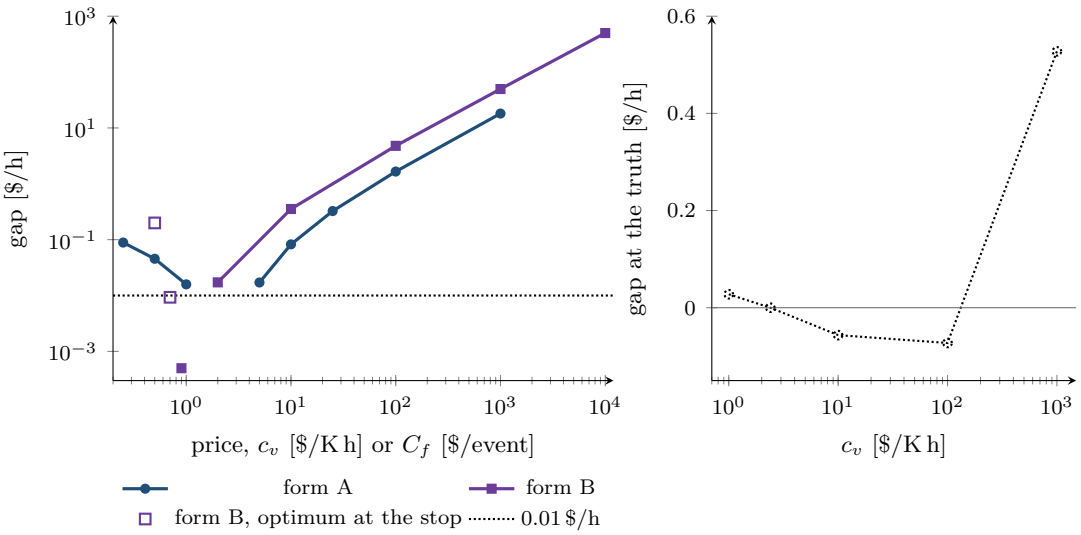


Figure 17.5: The gap $J(r_{\text{cc}}) - J$ at the optimum. Left: over the posterior, on all the draws, form A's gap (17.19) and form B's. Hollow squares are prices where form B's optimum is the stop. Right, dotted: form A at the stated truth, $\beta_3 = 0.57$, with e fresh; constructed, one truth.

1000, the cost is 18.10 $/h. Those are the ends of a 416-fold range, and they do not describe a modest error. With r_{cc} held and the true price a factor below and above the implied one, the regret is 0.0037 and 0.0047 $/h at a factor of 1.5, 0.0103 and 0.0151 at 2, 0.0240 and 0.0434 at 3, and 0.0919 and 0.3082 at 10. The asymmetry appears beyond about a factor of 3, and it comes from c_v multiplying the excess, which makes the gap linear in the price at the high end. On the first 200 000 draws, r_{cc} is within 0.01 $/h of the optimum for c_v from 1.199 to 4.235, a factor of 2.00 down and 1.76 up, and within 0.05 $/h from 0.456 to 7.643. Form B shows the same: +0.0172 at 2, +0.354 at 10, +4.78 at 100 and +49.7 $/h at 1000. It also has +0.0093 and +0.1993 $/h at 0.7 and 0.5, where the optimum is the stop, and that is a gap no one should collect. Here a set point backed off too far costs cents per hour at any price in the range, and one held too high costs dollars only once the true price is tens of times the implied one.

Against perfect knowledge. Table 17.8 sets each against the physical limit. Each row is a price; the columns give the excess of $J(r_{\text{ec}})$ over all known, over the state known with e remaining, and that of $J(r_{\text{cc}})$ over all known, all paired. All known, the total is 1.9412 $/h.

Table 17.8: Above perfect knowledge, 20 000 draws, seed 1717, paired; all known 1.9412 $/h.

c_v [$/K h]	$J(r_{\text{ec}})$ above all known (se) [$/h]	Above state known (se) [$/h]	$J(r_{\text{cc}})$ above all known (se) [$/h]
1	0.1696 (0.0009)	0.1099 (0.0009)	0.1860 (0.0007)
2.4001	**0.2115** (0.0014)	0.1377 (0.0014)	0.2115 (0.0014)
10	0.2698 (0.0026)	0.1767 (0.0026)	0.3501 (0.0063)
100	0.3483 (0.0058)	0.2296 (0.0058)	1.9905 (0.0653)
1000	0.4039 (0.0065)	0.2638 (0.0065)	18.3954 (0.6556)

The closed form (17.20) gives 0.2188 $/h at $\alpha = 0.05$, and the back-off's energy alone gives 0.1744 $/h, a ratio of $\varphi(\kappa)/(\alpha\kappa) = 1.2540$. **Chapter 16's whole back-off, at about** 0.17 **$/h, leaves out the tail it keeps:** 0.21 **$/h at this price.** The cost of the band rises only slowly with price, from 0.17 to 0.40 $/h over three decades.

At the stated truth. The record's truth is constructed: $\beta_3 = 0.57$, $R = 31.4720$°C, with $T_s = 15.3067$°C and 1.9846 \$/h at the limit. With e fresh by quadrature, the best r knowing the state at 2.4001 \$/K h is 30.9779°C, at 2.0585 \$/h, and r_{cc} gives 2.0735 \$/h with Pr = 0.007533. **At this truth the tighter set point loses until the price is high**: $J(r_{cc}) - J(r_{ec})$ is +0.0277 at 1, −0.0564 at 10, −0.0726 at 100 and +0.5277 \$/h at 1000. Rack 3's truth stands 0.73 K above r_{cc} in r. This is one truth.

17.3.7 A sum over racks against the joint limit

Control engineer's question. Sum over racks or joint limit?

General form. The action that minimizes the expected consequence (2.23).

Specialized form. Equation (2.23) with the loss (17.6), closed loop. Differentiating each rack's excess as in equation (17.10),

$$\sum_k c_k \Pr(T_k > 27 \mid \mathbf{y}) = c_e(r^*). \tag{17.21}$$

A joint constraint holds Pr(any) = 0.05 and knows no depth and no rack. The hottest-excess form prices only the worst inlet. The ratio prices are stated: "rack 3 ten times" means $(1, 1, 10)$, so the level of the price moves with the ratio.

Building blocks. Three beliefs: the first 200 000 draws of the window's posterior; the posterior after a handheld reading of rack 3 at 23.2221°C, Laplace normal, 200 000 draws, seed 1718, under which rack 2 binds; and a hand-set belief, not the posterior, with rack 2 given rack 3's mean and spread, independent, as a stated test of when the forms part. The marginals and the slope are read on each.

The plant's numbers. Table 17.9 gives them. Each row is a belief; the columns give the joint set point with the marginals of racks 2 and 3, the uniform price at it in each form, the gap between the hottest and the sum at that price, and the optimum with a rack priced ten times.

Table 17.9: Three racks, closed loop. Prices in \$/K h. The first two rows are posteriors; the third is hand-set.

Belief	Joint r [°C]; rack 2, rack 3 Pr [–]	Hottest, sum [\$/K h]	Gap [K]	Rack 2 ×10: r [°C], Pr(any) [–]	Rack 3 ×10: r [°C], Pr(any) [–]
Window's posterior	30.7428; 0.00034, 0.05017	2.3759, 2.3741	0.0006	31.0938, 0.10763	30.2075, 0.01206
After the handheld reading	32.3010; 0.04957, 0.00056	2.3713, 2.3656	0.0011	31.7675, 0.01190	32.5449, 0.08742
Hand-set, racks 2 and 3 identical	30.4772; 0.02553, 0.02554	2.4024, 2.3523	0.0074	30.1880, 0.02206	30.1748, 0.02114

At a uniform price the sum, the hottest and the joint nearly coincide, to within 0.0074 K even with two racks identical, because near $\alpha = 0.05$ two violations rarely coincide and $\sum_k \Pr(T_k > 27)$ is close to Pr(any): 0.05107 against 0.05001. The window row's implied 2.3759 against 2.4001 is that subset's noise on α, which reads 0.05047 there. **What moves the set point is the price on the binding rack, not the ratio.** The last two columns of the table mix the two. On the window, $(1, 1, 10)$ stands 0.5352 K below the joint's r, and the uniform 10 stands 0.5358 K below it. And $(1, 10, 1)$ stands 0.3510 K above it, where the

uniform 1 stands 0.4029 K above: that rise comes from rack 3's price falling, not from rack 2's rising. With the binding rack held at the uniform 2.3741 \$/K h, a rack 2 ten times dearer pulls r down by 0.0228 K with a gap of +0.0001 \$/h at the joint's r, and one hundred times dearer by 0.1652 K with +0.0061 \$/h. After the reading, with rack 2 binding, the same two rows give 0.0327 K and +0.0002 \$/h, and 0.1596 K and +0.0090 \$/h. So on these two posteriors the forms barely part until the rack that does not bind is about 100 times dearer. Where two racks bind together, in the hand-set row, there is no such rack: $(1, 10, 1)$ stands at 30.1880°C, between the joint's 30.4772°C and the uniform 10's 29.9884°C.

17.3.8 Whether the readings rank the same

Control engineer's question. Do the readings rank the same?

General form. The preposterior expectation (2.16), with the utility the cost saved.

Specialized form. Equation (2.16) with u the realized saving in J_{A} on the outcome's own set point,

$$U(a) = \mathbb{E}_{\tilde{y}_a|\mathbf{y}}\big[\, J(r_0) - J(r(\tilde{y}_a)) \,\big], \qquad \Pr\big(\text{all } T_k \le 27 \mid \mathbf{y}, \tilde{y}_a\big)\big|_{r(\tilde{y}_a)} = 1 - \frac{c_e(r \mid \mathbf{y}, \tilde{y}_a)}{c_v}. \quad (17.22)$$

Each outer row draws a state from the posterior of step 3, draws the reading, and refits. The refit's covariance is the $\boldsymbol{\Sigma}_a$ of equation (4.73) at that row. The set point comes from the row's normal posterior, with $c_e(r \mid \cdot)$ by a Gauss–Hermite rule on the slope, and the cost is realized on the same row. The closed form beside it, normal, frozen slope, one rack, compares this value at α with the back-off's at 0.05, for any reading: $(\varphi(\kappa_\alpha)/\alpha)/\kappa_{0.05}$.

Building blocks. Chapter 16's outer draws, same seed 1616 and order, 3000 rows, with the fit's five readings held; a 55-point lattice for the joint on each row; the largest last step over the refits, 1.95×10^{-7}; closed loop on the fresh route. With no reading on these rows, r_0 is 30.7425°C at the implied price and J is 2.1550 \$/h.

The plant's numbers. Chapter 16's column reproduces, 0.0764 and 0.0989 \$/h. Table 17.10 gives the value of each candidate. The columns give the back-off's value, the value in J at 2.4001 \$/K h with its energy and excess parts, the ratio of the two, and the value at 10 and 100. Figure 17.6 plots the same four columns for the eight candidates that have any value, with the candidates on the horizontal axis.

At the implied price the rank is the same within standard errors, and the readings are worth about 1.2 **times more.** The back-off's value counts the energy and omits the tail that the reading shrinks. The ratios in Table 17.10 carry the sampled e, so they do not resolve this. Table 17.11 repeats the six leading candidates on 10 000 new rows, seed 1747, with e integrated in the realized cost by the 41-node rule. Each row is a candidate, and the columns give the back-off's value, and at each price the value and its ratio to the back-off.

At the implied price the ratios are 1.194 to 1.243, each with a standard error of 0.02 to 0.03, so they agree with the closed form's 1.2540 or lie a little below it. At 10 and 100 the ratio rises, to 1.447 to 1.581 and 1.785 to 2.056, against the closed form's 1.5824 and 2.0162. The readings that include rack 2 keep up with the one-rack closed form, and those of rack 3 alone fall short of it, since rack 2 stays at its prior as the set point backs off. The in-loop reference's 1.543 in Table 17.10 is noise and carries no mechanism. Its posterior does not change, so it can be computed on the draws with no outer rows: a value of 0.0128 (0.0001) against a back-off of 0.0101 \$/h, a ratio of 1.261, and 1.589 and 2.020 at 10 and 100, which are the closed form's ratios within the draws' noise.

Table 17.10: Value of a reading, closed loop, fresh route, 3000 rows, e sampled; standard errors paired against no reading; \$/h. The ratio's standard error is near 0.06 here; Table 17.11 resolves it.

Candidate	Back-off at 0.95 (se) [\$/h]	At 2.4001 (se); energy, excess [\$/h]	Ratio [–]	At 10 (se) [\$/h]	At 100 (se) [\$/h]
Permanent inlets, racks 2 and 3, 12 windows	0.0989 (0.0016)	**0.1278** (0.0047); 0.0988, 0.0290	1.292	0.1618 (0.0084)	0.1911 (0.0138)
Survey β_2, β_3, 0.02	0.0974 (0.0017)	0.1256 (0.0046); 0.0973, 0.0282	1.289	0.1597 (0.0084)	0.1882 (0.0132)
Permanent inlet rack 3, 12 windows	0.0901 (0.0014)	0.1138 (0.0049); 0.0900, 0.0238	1.262	0.1388 (0.0090)	0.1555 (0.0132)
Survey β_3, 0.02	0.0887 (0.0014)	0.1134 (0.0049); 0.0885, 0.0249	1.280	0.1372 (0.0088)	0.1542 (0.0133)
Permanent inlet rack 3, own calibration 0.2 K	0.0828 (0.0014)	0.1054 (0.0048); 0.0826, 0.0228	1.274	0.1299 (0.0090)	0.1394 (0.0149)
Inlets racks 2 and 3, 0.3 K, once	0.0826 (0.0016)	0.1054 (0.0047); 0.0825, 0.0229	1.275	0.1330 (0.0086)	0.1646 (0.0126)
Handheld inlet rack 3, 0.3 K	0.0764 (0.0014)	0.0979 (0.0047); 0.0763, 0.0216	1.281	0.1218 (0.0088)	0.1296 (0.0155)
Reference return, permanent, in the loop, 0.05 K	0.0094 (0.0006)	0.0145 (0.0023); 0.0094, 0.0051	1.543	0.0248 (0.0054)	0.0257 (0.0126)
Inlet rack 2, 0.3 K	0.0002	0.0000 (0.0001)	–	0.0000	0.0000
Reference return, once; reference load meter	0	0	–	0	0

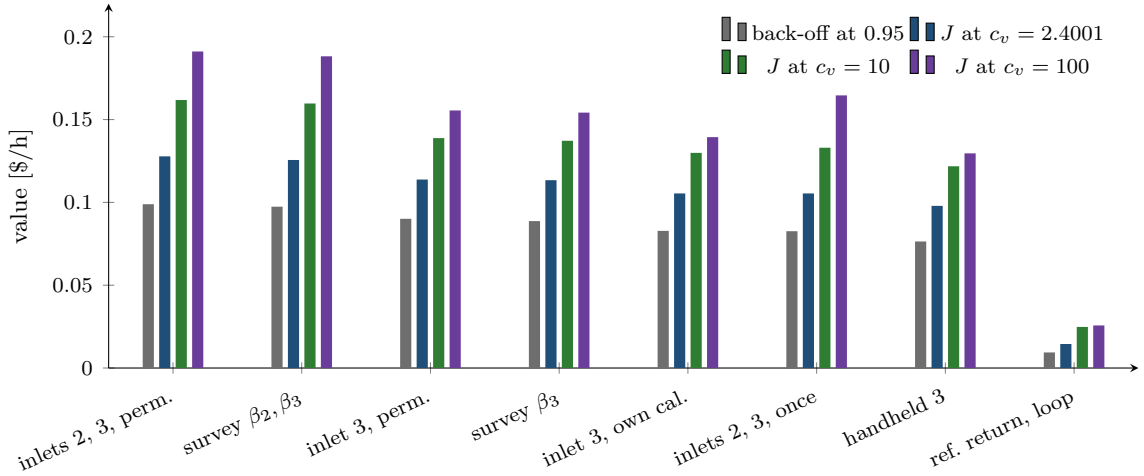


Figure 17.6: Value of each reading from Table 17.10, closed loop, fresh route. Gray: the back-off value at 0.95, Chapter 16's. Colors: value in J_{A} at three stated prices. Standard errors are in the table.

Table 17.11: Value of a reading with e integrated, closed loop, fresh route, 10 000 rows, paired against no reading; \$/h. The back-off's standard error is 0.0008 or 0.0009.

Candidate	Back-off	At 2.4001 (se)	Ratio (se)	At 10 (se)	Ratio (se)	At 100 (se)	Ratio (se)
Permanent inlets, racks 2 and 3	0.0994	0.1231 (0.0019)	1.239 (0.022)	0.1566 (0.0033)	1.576 (0.036)	0.2029 (0.0089)	2.042 (0.092)
Survey β_2, β_3	0.0979	0.1216 (0.0019)	1.243 (0.023)	0.1547 (0.0033)	1.581 (0.036)	0.2004 (0.0089)	2.048 (0.093)
Permanent inlet rack 3	0.0906	0.1083 (0.0021)	1.194 (0.027)	0.1312 (0.0037)	1.447 (0.044)	0.1618 (0.0093)	1.785 (0.105)
Survey β_3	0.0897	0.1076 (0.0021)	1.200 (0.027)	0.1306 (0.0037)	1.457 (0.045)	0.1612 (0.0094)	1.798 (0.107)
Inlets racks 2 and 3, once	0.0825	0.1020 (0.0019)	1.236 (0.026)	0.1300 (0.0033)	1.575 (0.043)	0.1697 (0.0089)	2.056 (0.110)
Handheld inlet rack 3	0.0772	0.0938 (0.0020)	1.215 (0.030)	0.1162 (0.0036)	1.505 (0.050)	0.1464 (0.0092)	1.896 (0.122)

The top two stay on top, and the permanent inlets lead the survey by a small and now resolved margin, +0.0014 (0.0003), +0.0019 (0.0004) and +0.0025 (0.0006) \$/h. The permanent rack-3 inlet and the survey of β_3 differ by +0.0006 (0.0003).

A rack-2 reading gains on a better rack-3 sensor as the price rises. Inlets 2 and 3 at 0.3 K, less the permanent rack-3 inlet, paired, is −0.0063 (0.0010), −0.0012 (0.0019) and +0.0079 (0.0031). The further the set point backs off, the more rack 2 matters. The excess part of the top row falls as the price rises, 0.0238 (0.0021), 0.0184 (0.0034) and 0.0147 (0.0090), and its standard error grows, because violations become rare. The rank correlations of 0.9636, 0.8545 and 0.8545 count the noise among the zero rows, and they are not quoted alone. Two things are not ranked: a loop on the inlet, whose stability is not computed, and form B's value, since the density condition was not solved per row.

17.3.9 Fresh records

Control engineer's question. Does it hold on fresh records?

General form. The expected consequence (2.23), realized over a stated world, set beside the probability the set point asked for.

Specialized form. Equation (2.23) with **z** and **y** drawn: a record is a truth, the closed graph's window at $r = 30.33$°C, five readings and any added one, the fit, the set point from the fractile (17.11), then the closed node with a fresh e. Each record's own posterior is therefore consumed by its set point. Over the prior the truth is drawn per record, and the α asked should be met on average. At one truth it need not be, and the frequency there is not a miscoverage.

Building blocks. Seed 1717 and offsets; the fit and the fractile per record, form A only; $Q = 200$ kW stated; 8000 records over the prior, 5000 and 4000 at the truth, 3000 across truths.

The plant's numbers. Table 17.12 runs the prior. Each row is a price; the columns give the α asked and the violations found, and the gap to r_{cc} on the same records, with no reading and with a handheld reading of rack 3 read first. **The quantile the price asks for holds over the belief.** The handheld column runs low at 1 and 2.4001 by about 2.2 standard errors, on one draw of truths shared across its rows. A fresh draw of 20 000 records reads 0.04945 (se 0.00153) against the 0.04918 asked, and Chapter 16's seed gives 0.0507. No bias is claimed.

Table 17.13 fixes the record's truth, $\beta_3 = 0.57$, as a row. The columns give, with no reading, violations and realized J with the quadrature beside them, and the same with a handheld and a permanent reading read first, with the mean and spread of the set point. The limit with the state known and e fresh is 2.0585 \$/h at 2.4001.

Table 17.12: Fresh records over the prior, 8000 per column, closed loop. Violations with their standard errors.

c_v [\$/K h]	No reading: asked, found (se) [–]	Gap (se) [\$/h]	Handheld first: asked, found (se) [–]	Gap (se) [\$/h]
1	0.11826, 0.11663 (0.00359)	+0.0160 (0.0011)	0.11787, 0.10988 (0.00350)	+0.0096 (0.0006)
2.4001	0.04941, 0.04850 (0.00240)	0	0.04919, 0.04412 (0.00230)	−0.0001
10	0.01190, 0.01100 (0.00117)	+0.0811 (0.0085)	0.01183, 0.01112 (0.00117)	+0.0406 (0.0050)
100	0.00120, 0.00112 (0.00037)	+1.6184 (0.1146)	0.00119, 0.00125 (0.00040)	+0.8752 (0.0683)

Table 17.13: Fresh records at the stated truth, $\beta_3 = 0.57$, closed loop. Quadrature: the no-reading set point at the truth, e integrated.

c_v [\$/K h]	No reading, 5000: violations (se) [–], J (se) [\$/h]	Quadrature [–], [\$/h]	Handheld, 4000: r_{ec} mean (sd) [°C], violations (se) [–], J (se) [\$/h]	Permanent, 4000: r_{ec} mean (sd) [°C], violations (se) [–], J (se) [\$/h]
1	0.15260 (0.00509), 2.0458 (0.0009)	0.143065, 2.0448	31.0585 (0.2927), 0.16275 (0.00584), 2.0689 (0.0014)	31.0862 (0.1818), 0.13400 (0.00539), 2.0559 (0.0011)
2.4001	0.00800 (0.00126), 2.0733 (0.0006)	0.007533, 2.0735	30.8528 (0.2924), 0.07150 (0.00407), 2.0879 (0.0022)	30.9197 (0.1817), 0.05950 (0.00374), 2.0722 (0.0016)
10	0, 2.1361 (0.0005)	0.000013, 2.1355	30.5839 (0.2911), 0.01475 (0.00191), 2.1145 (0.0041)	30.7018 (0.1808), 0.01300 (0.00179), 2.0914 (0.0025)
100	0, 2.2188 (0.0005)	0, 2.2186	30.2387 (0.2854), 0.00150 (0.00061), 2.1504 (0.0087)	30.4204 (0.1760), 0.00050 (0.00035), 2.1156 (0.0034)

At this truth, at 2.4001**, the handheld raises the realized** J, 2.0879 against 2.0733 \$/h, with violations from 0.008 to 0.0715. This is Chapter 16's finding priced, one truth above the survey. **By** 10 **the reading helps**, because the no-reading margin is then paid in energy: 2.1145 against 2.1361 \$/h. Both hold. The permanent's 0.0595 at 2.4001 stands above 0.05 by 2.5 standard errors at this truth, and it is a frequency at one truth, not a miscoverage.

Table 17.14 moves the truth at the implied price. Each row is a truth's β_3; the columns give violations and J with no reading and with the handheld, and the limit with the state known and e fresh.

Table 17.14: Across truths, $c_v = 2.4001$ \$/K h, 3000 records per row, closed loop. The 0.57 row's seed differs from Table 17.13's.

Truth's β_3 [–]	Violations, no reading (se) [–]	Handheld (se) [–]	J, no reading [\$/h]	J, handheld [\$/h]	State known [\$/h]
0.46	0	0.0110 (0.0019)	2.0711	1.9549	1.9224
0.50	0	0.0280 (0.0030)	2.0711	1.9995	1.9718
0.54	0.0007 (0.0005)	0.0477 (0.0039)	2.0713	2.0498	2.0213
0.57	0.0120 (0.0020)	0.0610 (0.0044)	2.0745	2.0890	2.0585
0.60	0.0903 (0.0052)	0.0823 (0.0050)	2.1049	2.1299	2.0959
0.65	0.6333 (0.0088)	0.1160 (0.0058)	2.5115	2.2023	2.1582
0.70	0.9827 (0.0024)	0.1527 (0.0066)	3.5830	2.2809	2.2208

Per record, the largest swings are at the far survey misses: the handheld costs 0.0145 \$/h at the record's truth and saves 1.30 \$/h at $\beta_3 = 0.70$. It also saves 0.0716 \$/h at 0.50, with no survey miss at all. **In expectation most of the value is energy recovered where the rack is at or cooler than surveyed.** The handheld's 0.0938 (0.0020) \$/h of Table 17.11 is 0.0771 of energy and 0.0168 of excess, and it can be split by each outer row's own β_3. Rows below 0.50 are 0.5031 of the mass and contribute +0.0577 \$/h, 0.61 of the value, with their

excess part negative at −0.0095. Rows at 0.65 and above are 0.0298 of the mass and contribute +0.0258 \$/h, 0.28 of the value, and they carry the excess part, at +0.0288.

17.4 The checks (step 5)

The tests stand against Chapter 16's node, the Laplace route, the draws, and the limits of no band and exact rows.

17.4.1 Routes and identities

Table 17.15 lists them. Each row is a check; the columns give what is compared, the result and its tolerance or standard error. All pass. The Laplace route and the draws agree at the implied price, 30.7425 against 30.7428°C, with $\alpha = 0.04996$ on both, and the gap is within the 0.0020 K standard error of r^*.

Table 17.15: The checks of the chapter.

What is compared	Result	Tolerance or se
Node against Chapter 16: r_{cc}; $\mathbb{E}[P]$ on its 3000 rows; handheld and best back-off values	30.7428°C; 30.1061 kW; 0.0764, 0.0989 \$/h	its printed digits
Laplace against draws at the implied price	r 30.7425, 30.7428°C; α 0.04996 both	se r^* 0.0020 K; se J 0.00026 \$/h
Lattices 55, 144, 46 368 at that α	30.7425, 30.7425, 30.7424°C	printed digits
Numerical minimum against the root (17.11), over the prices	below 0.00001 K; α gap 0.000001	the draws
Central difference against the slope (17.9)	equal	printed digits
e integrated against raw draws, Pr at r_{cc}	0.050008, 0.049961	se 0.00024
Density by formula against a count over ±0.02 K	0.11770, 0.11975 per K	about one se
10^6 independent normal draws of the Laplace node, seed 1737	0.04982, asked 0.04996	se 0.00022
No band, $f = 10^{-4}$: r^* against the means' R and the reduced node's	32.1997; 32.1998; 32.2000°C	–
Convexity on the live region, 0.25 K grid, 26 to 35°C: form A, form B	least second difference 5.41×10^{-5}; −0.0226	A convex, B not
Truth row	β_3 0.57 (+0.88 sd), R 31.4720°C, $\sum G$ 12.8670 kW/K, T_s 15.3067°C, 1.9846 \$/h	as Chapter 16

17.4.2 The limit of exact rows

The node's rows are Chapter 16's and no row is added here, so the limit $\sigma_p \to 0$ on the full state is that chapter's table and is cited, not recomputed. The limit that is new here is the band's. As $f \to 0$ the fractile (17.11) leaves r^* at the reduced node's limit less a vanishing margin, 32.1997°C at $f = 10^{-4}$, with $E = 1.9354$ \$/h, $c_e = 0.11871$ and $\alpha = c_e/c_v = 0.04946$. So the no-band optimum sits on the limit and α keeps its value: the band sets where the fractile lies, and the two slopes set which fractile it is.

17.5 The report (step 6)

At this node at 23:00, a price on the excess puts the set point at a quantile with $\alpha = c_e/c_v$. Chapter 16's 95 % is the optimum if a kelvin-hour costs 2.40 \$, or an hour in breach 1.02 \$, 3.7 minutes of the hall's bill. The set point moves little with price, 1.79 K from 2.40 to 1000 \$/K h. Standing that far back at a true 2.40 costs 0.17 \$/h, and staying at r_{cc} at a true 1000 costs

18.10 \$/h, while within a factor of 3 of the implied price the regret stays below 0.05 \$/h either way. The fixed penalty is not a milder version of this: below 0.71 \$/event its minimum is the stop.

17.5.1 The answers and their bands

Table 17.16 gives the answers in the order of Table 17.1. Each row is a question; the columns give the answer, its band or what it rests on, and the stated truth where one applies. Figure 17.5 carries that truth's gap as its dotted line.

Table 17.16: The report for the set point's price. Draws: seed 11; records as stated. Prices are stated choices. Dollars and percentages are this example's.

Control engineer's question	Answer	Band or basis	Stated truth
What does a degree of set point cost?	0.11991 \$/h per K at r_{cc}; falls 3.42 % over 28 to 33°C; stop 1.4348 \$/h	central differences match; 0.22 \$/kWh gives 0.37687	1.9846 \$/h at the limit
Where does a price on the excess put it?	A quantile, $\alpha = c_e(r^*)/c_v$: 31.1520, 30.7428, 30.2102, 29.5199, 28.9502°C at 1, 2.4001, 10, 100, 1000 \$/K h	se r^* 0.0016 to 0.0118 K, rising with price; to the draws; below 0.11527 the stop	30.9779°C at 2.4001, state known
Where does a price per violation put it?	30.7428°C at 1.0188, 30.0600°C at 5 \$/event; the stop below 0.7056	non-convex; no interior α above 0.07997, at the 20°C stop	–
Which price does $\alpha = 0.05$ imply?	2.4001 \$/K h (c_n 2.2802), 1.0188 \$/event	2 se on the draws 2.377 to 2.423, 1.013 to 1.025; closed form 2.3983, 1.0283; 0.22 \$/kWh 7.5433, 3.2019	–
What does the band do to the optimum?	Moves it 1.46 K, widens it to 0.3779 K from a floor of 0.0443 K	tolerance 0.005 \$/h; 0.1682, 0.7683 K at 0.001, 0.02	–
What is the chance constraint's error if the price is known?	+0.08242 at 10, +1.65664 at 100, +18.09724 \$/h at 1000; within 0.01 \$/h for 1.199 to 4.235	paired se; 0.2115 \$/h above all known at 2.4001	+0.0277, −0.0564, −0.0726, +0.5277 at 1, 10, 100, 1000
Sum over racks or joint limit?	Coincide at a uniform price, within 0.0074 K; the binding rack's price moves r, 0.5352 K at 10; the other rack ten times dearer, 0.0228 K	ratios stated; subsets 200 000	–
Do the readings rank the same?	Same rank within se at the implied price; value about 1.2 times; rack 2 gains with price	10 000 rows, e integrated; ratio se 0.02 to 0.03; top two +0.0014 (0.0003)	–
Does it hold on fresh records?	Over the prior 0.04850 (se 0.00240), asked 0.04941; handheld 0.04945 (se 0.00153) on the fresh draw, asked 0.04918	8000 and 20 000 records; form A	at 0.57: 0.00800, J 2.0733 \$/h; handheld 0.07150, 2.0879

17.5.2 The entry in the commissioning log

The entry records the α in use and the price it implies at this node, 0.05 with 2.40 \$/K h or 1.02 \$/event, and names the form of the price. It states that those prices come from the energy's slope and that no consequence has been measured, so whoever owns the racks can compare them with a figure of their own. It records that a fixed penalty below 0.71 \$/event sends the optimum to the stop and so needs a limit beside it. And it records that prices are to be stated by rack if the racks differ, since a joint limit cannot express that, and that the binding rack's price is the one that moves the set point. It makes no change to r_{cc} on this evidence. Each of these rests on a constructed record, one load and one hour.

17.5.3 Common mistakes

Table 17.17 lists six mistakes this node makes easy. Each row names the mistake, its form here and the correction.

Table 17.17: Six mistakes with a price on a limit.

Mistake	On this node	Correction
Writing $c_e/(c_e + c_v)$ with the gross price	0.047584 against 0.049961 at 2.4001; 0.322 against 0.475 at 0.25	c_e/c_v with the gross price; $c_e/(c_e + c_n)$ with the net
Freezing the slope at a distant set point	off by 0.0323 K at 0.15 \$/K h	take c_e at r^*
Using a fixed penalty with no stop in the comparison	$J(r_{\text{low}})$ 2.1357 loses to 2.0848 at 0.65 \$/event	compare with equation (17.14); keep a limit beside it
Treating a form B as implying any α	none above 0.07997	the duality gap; use the constraint or form A
Saying the band only moves the optimum, or only widens it	1.46 K and 0.3779 K at $f = 1$, floor 0.0443 K	state both, with the tolerance
Valuing a reading by its back-off alone	0.0938 against 0.0772 \$/h at the implied price	count the tail that the reading shrinks

17.6 The next measurement (step 7)

The candidates are those valued in Section 17.3.8, by the preposterior (17.22), each refit with $\mathbf{\Sigma}_a$ of equation (4.73). A new row rebuilds the posterior of step 3, and the set point follows from it.

17.6.1 The candidates

Table 17.18 reads what each candidate does to the set point at the implied price. Each row is a candidate; the columns give the mean rise of $r(\tilde{y})$, its spread over outcomes, the chance the set point falls on the reading, and the violation share realized on the rows. Permanent instruments are among them: the inlets on racks 2 and 3, rack 3 with its own calibration, and the in-loop reference. With no reading the violation share on these rows is 0.0533.

Table 17.18: What each candidate does to the set point, form A at 2.4001 \$/K h, closed loop, 3000 rows.

Candidate	$\mathbb{E}[r(\tilde{y})] - r_0$ [K]	sd $r(\tilde{y})$ [K]	Pr(r falls) [–]	Violation share [–]
Permanent inlets, racks 2 and 3, 12 windows	+0.8287	0.7565	0.1457	0.0503
Survey β_2, β_3, 0.02	+0.8162	0.7577	0.1467	0.0477
Permanent inlet rack 3, 12 windows	+0.7538	0.6494	0.1423	0.0480
Survey β_3, 0.02	+0.7416	0.6528	0.1443	0.0457
Permanent inlet rack 3, own calibration 0.2 K	+0.6923	0.6489	0.1640	0.0503
Inlets racks 2 and 3, 0.3 K, once	+0.6917	0.7207	0.1747	0.0527
Handheld inlet rack 3, 0.3 K	+0.6393	0.6444	0.1750	0.0493
Reference return, permanent, in the loop, 0.05 K	+0.0843	0.0002	0.0000	0.0517
Inlet rack 2, 0.3 K	+0.0013	0.0156	0.0183	0.0543

17.6.2 The recommendation

For this question the reading to take is on rack 3, and the survey and the permanent inlets rank together at the top, within 0.0025 \$/h of each other. No instrument cost is stated, so there is no verdict to buy. A price on the excess changes the case for a reading by about a fifth at the implied price and it does not reorder the top; rack 2 matters more only if the owner states a high price. A reading of rack 2 alone stays near nothing until rack 3 is known. The reading raises the set point on most outcomes and lowers it on 0.14 to 0.18 of them, and that is its working as intended.

17.6.3 What the analysis showed

The sheet read as a conventional and safe figure. The analysis finds that it is also a price, a small one, set by the energy's slope and by no consequence at all: on this node, whoever calls 95 % safe is saying that an hour over the limit is worth about what the hall spends in a few minutes. Whether that is right is outside this model. The band moves the optimum and widens it, with the move far the larger and a floor under the width from the kink. A fixed penalty looked like the same idea with a simpler price, and it is a different problem, non-convex, with a range of prices over which the optimum abandons the limit and a probability above which it never returns. A price picks an α, and an α asserts a price. Near the implied price an error in the price costs little either way, and far above it a set point held too high costs far more than one backed off too far, because the price multiplies the excess; so the question that matters is whether the breach's owner has a figure above the implied one.

17.6.4 Limits of the model

One move, no future. Coupling 3 is gone, since a set point teaches nothing here, and coupling 2, since no reading follows it. The node is steady, at one load, with no forecast, which Chapter 19 takes up, and no plan over time, which Chapter 18 does. A stop in time is Chapter 15's.

The consequence. A violation has no duration or recovery, and the hour-long event is stated. A consequence linear in the excess or flat per event is neither form of a real one. No consequence price is measured. A price is a mean, since J is linear in it, so a band on the price does not enter the optimum.

The draws. Rows with α below about 10^{-4} are at the draws' reach and are read by their standard errors. The 3000 outer rows are thin at high prices and the 10 000 still carry a standard error near 0.1 on the ratio at 100, the window's fit is held and shared by the outer draws, and form B's preposterior is not computed. Subsets of 50 000 to 200 000 draws and 8000 prior records are reductions for time and named where used.

The plant. The coils' g_u are independent of T_s, no pump or fan power is drawn, and the plant's constants are its stated ones. The records share a truth at each row, and the plant is the model.

The constructed record. The record is constructed from a stated truth, so its agreement tests the method on its own model and not the model against a real hall. Percentages and dollars are this example's.

17.7 What to retain

Definition 17.1 (The expected-cost set point)**.** With a price c_v added to an energy of slope c_e, the minimizer of $E(r) + c_v\,\mathbb{E}[(T_{\max} - 27)_+ \mid \mathbf{y}]$ puts the probability of violation at $c_e(r^*)/c_v$, equation (17.11). Read backward, a chance constraint at α asserts the price c_e/α, equation (17.15).

Three conclusions carry beyond the hall as statements of form. The width of the posterior enters the move through the price, moving the optimum by its quantile and widening it by a root, down to the kink's floor. A flat price per event has a density condition, equation (17.13), and so needs a comparison with the actuator's stop. And a value of information depends on the cost it is measured in, even where the rank survives. The numbers 2.40, 1.02, 0.08, 1.46 and 18.10 are this node's. The chapter's questions join the others in Appendix A.

Table 17.19 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 17.19: The terms of the chapter and where each is defined. Units for the two-unit hall at its node.

Term	Expression	Units	Meaning	Defined in
Limit in the set point	R, eq. (17.2)	°C	r that puts the hottest on 27°C	17.1.2
Three costs	eqs. (17.4) to (17.6)	\$/h	energy plus a stated price	17.1.4
Energy's slope	c_e, eq. (17.9)	\$/h per K	what a degree saves	17.3.1
Critical fractile	eq. (17.11)	–	$\alpha = c_e(r^*)/c_v$	17.3.2
Gross and net price	c_v; $c_n = c_v - c_e$, eq. (17.12)	\$/K h	added to the energy; slope past the limit	17.3.2
Density condition	eq. (17.13)	/K	$c_e = C_f\,p$; zero or two roots	17.3.3
Implied price	eq. (17.15)	\$/K h, \$/event	the limit's multiplier	17.3.4
Duality gap	$\alpha > 0.07997$ here	–	no C_f reproduces it	17.3.3
Kink's width	eq. (17.18)	K	the interval's floor	17.3.5
Expected shortfall	eq. (17.20)	K, \$/h	the tail's depth, then cost	17.3.6
Sum over racks	eq. (17.21)	–	the condition at unequal prices	17.3.7
Value in cost	eq. (17.22)	\$/h	preposterior saving in J	17.3.8

17.8 Exercises

Exercise 17.1. *Formulation.* A pump holds a tank's level against a demand, and the pump's energy falls as the level set point falls. Write forms A and B for a lower limit on the level, and name the scenario node at the pump's stop. ($J_{\mathrm{A}}(r) = E(r) + c_v\mathbb{E}[(\ell_{\min} - \ell)_+ \mid \mathbf{y}]$ and $J_{\mathrm{B}}(r) = E(r) + C_f\,\Pr(\ell < \ell_{\min} \mid \mathbf{y})$, with r pinned and the speed free on the closed node. At the stop the speed is pinned, the node is open, the controller's row leaves, and form B's total must be compared with its value there.)

Exercise 17.2. *Derivation.* Derive the quantile, with the gross and net forms. (With $T_{\max} - 27 = r - e - R$, $\partial_r(T_{\max} - 27)_+ = \mathbf{1}[T_{\max} > 27]$, so $J'_{\mathrm{A}} = -c_e + c_v\alpha$ and $\alpha = c_e/c_v$. With $c_n = c_v - c_e$, $-c_e(1-\alpha) + c_n\alpha = 0$ gives $\alpha = c_e/(c_e + c_n)$, which is the same condition.)

Exercise 17.3. *Derivation.* Show why form B has no or two stationary points. (The condition is $c_e(r)/C_f = p(r)$. The left side is nearly flat, and p rises to one mode and falls, so a level above the mode misses it and a lower one crosses twice: a local minimum on the rising side and a local maximum on the falling side. The tangent is at $C_1 = 0.2507$ \$/event.)

Exercise 17.4. *Application.* Find the implied c_v at r_{cc}. (By equation (17.15) with the draws' slope 0.11991 and probability 0.04996, $c_v = 2.4001\,\$/\mathrm{K\,h}$.)

Exercise 17.5. *Application.* Say what C_2 means and give it. (The price at which the local minimum ties the stop: $J_{\mathrm{B}}(r_{\mathrm{low}}) = E_{\mathrm{stop}} + C_f \Pr_{\mathrm{stop}}$, at $r_{\mathrm{low}} = 30.9526°\mathrm{C}$ and $2.1404\,\$/\mathrm{h}$. Below it the stop is the minimum. $C_2 = 0.7056\,\$/\mathrm{event}$.)

Exercise 17.6. *Derivation and application.* Show why J'' falls as $1/s$, and give the kink's width at $0.005\,\$/\mathrm{h}$. (For a normal, $p(r^*) = \varphi(\kappa)/s$, so $J'' \approx c_v\varphi(\kappa)/s$ with E'' small. At $s = 0$ the slopes are -0.11871 and $2.4001 - 0.11871$, so the width is $0.005/0.11871 + 0.005/2.2814 = 0.0421 + 0.0022 = 0.0443\,\mathrm{K}$.)

Chapter 18

Model Predictive Control as a Restriction

18.1 A plan on the commissioning sheet

The commissioning sheet offers a supervisory plan, and its one line of method reads “one fit per step, then 95 % limits on every inlet at every step, so 95 % over the hour”. At 20:00 the record’s supervisor reads 14°C on the chiller and 30.5°C on both return set points, with the tank idle. Rack B3 has no inlet meter. The tariff is 0.22 $/kWh now and 0.07 $/kWh after 23:00, and the tank holds about 300 kWh.

The engineer would accept the sheet’s arithmetic. Or, reading the same line with more care, the engineer would notice that twelve windows and ten racks make 120 rows, and would divide the 5 % among them, 0.042 % each. Neither reading says what the plan is. The question of this chapter is

what the plan over the next hour is, what 95 % on each row gives over the hour, and what a fixed plan gives up.

Jargon: model predictive control, plan, fixed plan, copy, per-copy limit, joint limit, Bonferroni’s rule, direct calibration of the joint, certainty-equivalent plan, move limit, hold share, value of stored cooling, net cost, perfect knowledge

Model predictive control sets scenario nodes across time-indexed copies of the graph, with the pins as decisions. A *plan* is the list of those pins over the horizon; a *fixed plan* is one carried out whatever is read after it is made. A *copy* is the graph at one future window. A *per-copy limit* is one chance constraint per rack and window, and the *joint limit* is one chance constraint on all of them together. *Bonferroni’s rule* shares the joint’s α equally among the rows. *Direct calibration of the joint* chooses one tightening multiplier z by draws of the belief and the forecast, so that the joint is met. A *certainty-equivalent plan* has no band in its limits. A *move limit* bounds the change of a set point in a window, and it keeps a plan within its hold rows. The *hold share* is the share of a set-point move a hall completes within a window. The *value of stored cooling* is a stated terminal value of the tank, the water and the hall. *Net cost* is electricity less that value. *Perfect knowledge* is the plan told the draw: a lower bound on cost, not an available plan.

Key idea

One step of a predictive controller is two problems: a fit that conditions rows on the past, and a plan that averages rows on the future. No reading picks a future move. A limit on each row does not hold the row's conjunction, but the reason on this plant is not the number of copies: time is nearly one event, because the entries that matter persist. The loss lies across racks that no reading touches and in the plan's own first moves.

18.1.1 The boundary of the map

Table 7.6 places a model predictive controller on the row of the fixed sequence, and this chapter works on the sequence before it is re-solved. The restriction on the problem (7.1) is a fixed plan chosen at one now: $\mathbf{u}_{\text{now}+1}, \dots, \mathbf{u}_{\text{now}+12}$ are numbers and not a policy. In the couplings of Section 7.1.3, coupling 1 is kept, because the moves are entries of the future copies. Within the plan no future reading exists, so coupling 2 has no effect on the plan's rows, though the past record the fit reads is still closed-loop. Coupling 3 is given up, since the plan does not shape $\mathbf{\Sigma}$. What the restriction gives up beyond the couplings is recourse: a future move may not depend on a future reading. Coupling 4 survives in one form, the band entering only by limits, as back-off. Re-solving, and the realized recourse it restores, are Chapter 20's.

The last line of equation (7.1) is joint over time and stated on the evidence at the moment of deciding. Chapter 7 left that statement to be paid for on a horizon longer than two moves, and this chapter pays it, at one now. Every percentage and dollar below is this example's and not a property of control.

18.1.2 Two halls, one water loop and a tank (step 0)

The plant is drawn at 5-minute windows. Two halls share one chilled-water loop and a tank. Each hall $h \in \{A, B\}$ has a return air $T_{r,h}$ held by a sampled PI through a valve u_h, and the water T_s follows a chiller set point T_{sp} with a lag. With Q_h the rack load, $Q_{o,h}$ the other heat, G_h the hall's coil conductance and q_t the tank's charge, positive to charge,

$$\begin{aligned} C_h\,\dot{T}_{r,h} &= Q_h + Q_{o,h} - Q_{c,h}, & Q_{c,h} &= u_h G_h (T_{r,h} - T_s), \\ C_w\,\dot{T}_s &= Q_{c,A} + Q_{c,B} + q_t - Q_{ch}, & \dot{E} &= \eta_c q_t^{+} - q_t^{-} - \lambda E, \end{aligned} \tag{18.1}$$

and the electricity is $Q_{ch}/\mathrm{COP}(T_s, T_{\text{amb}})$ at the window's price. A rack's inlet mixes the hall's coil air with its return by the recirculation fraction β_k,

$$T_{\text{in},k} = T_{r,h} - (1 - \beta_k)\,\frac{Q_{c,h}}{\sum_u C_u}, \tag{18.2}$$

and each inlet has the limit 27°C. Storage is hall A 4.00, hall B 3.50 and water 11.67 kWh/K against a tank of 600 kWh, and the chiller's capacity is 800 kW. At valve fully open and steady state this is the shift plant of the companion *Operations and Control by Derivation*. What the loops change is that a set point is no longer the air: a move takes time to land, a hall's return sits at $r_h - b_h$ with b_h the offset of its return sensor, and the plan's rows have a hold regime and a stop.

The scenario node. Every what-if is closed loop at the halls: the set points r_A, r_B are pinned, the valves are free and the PI's row is present, and at a stop the actuator is pinned again, as in equation (15.1). The chiller's set point T_{sp} is pinned with its lag. The tank's move q_t is pinned. The plan's start is the end of window 12 (20:00, tariff 0.22 flat over the hour), of window 72 (01:00, 0.07 flat) and, for the question of when moves matter, of window 126 (05:30, the price edge 0.10 to 0.18).

Figure 18.1 draws one step. Time runs from left to right across copies of the graph, and the vertical line is the now. To its left the copies carry reading rows (dashed) and are conditioned. To its right they carry no reading row; the gray boxes are the moves, pinned at numbers, and the copies are averaged over the belief and the forecast. The entries that persist, along the bottom, reach across the line.

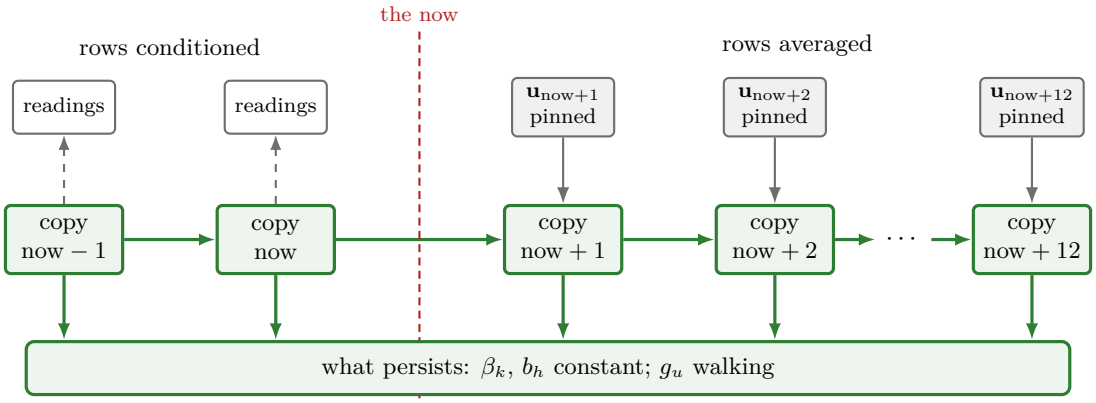


Figure 18.1: One step on time-indexed copies. Left of the now the copies have reading rows (dashed) and are conditioned. Right of it the moves are pins (gray) with no reading row, and the copies are averaged over the belief and the forecast. The entries that persist are shared by all copies, which is why the copies turn out to be nearly one event.

18.1.3 The control engineer's questions (step 1)

Table 18.1 lists the nine questions. The first column is the question in the engineer's words, the second the general form, the third its mathematical form on this plant, the fourth its category, the fifth the computation and the last the subsection. All the plans are scenario nodes across time, which the third column says; the computation is the one that reads them.

18.1.4 The evidence and the plan as rows (step 2)

Two sets of rows are written here, the ones on the past that the fit conditions and the ones on the future that the plan averages. They share the plant and share no line.

Which row acts on which entry. One window's state has 23 entries: T_s, two loads, six log-conductances g_u, ten fractions β_k, two offsets and two valves. With $*$ for an entry that can be nonzero, the columns grouped and the 48 future moves as the last group, the pattern is

$$\mathbf{J} = \begin{bmatrix} * & * & * & 0 & 0 & * & * & 0 \\ * & * & * & * & 0 & 0 & * & 0 \\ 0 & * & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & * & 0 \\ * & 0 & * & * & * & * & 0 & 0 \end{bmatrix} \begin{array}{c} T_s,\ T_r,\ T_d \text{ readings} \\ \text{4 inlet readings} \\ \text{hall and loop meters,} \\ \text{valve readings} \\ \text{priors} \end{array} \tag{18.3}$$

with columns (T_s; Q; g; β_{read}; β_{unread}; b; u; $\mathbf{u}_{\text{future}}$). Two columns matter below. The last is zero in every row, and the fifth, the fractions of racks A2, A4, A5, B1, B3 and B5, has only its priors. The sparsity is a property of the graph: no reading row reaches an unread rack.

The rows of the plan. Copy j of 12 has the return air, water and tank as linear storage rows between windows of length Δ, the belief's means $\hat{G}_h$, $\hat{\beta}_k$, $\hat{b}_h$ and the forecast's means $\hat{Q}_{h,j}$,

Table 18.1: The questions asked of the hour's plan. $\Pr_J$ is the joint of equation (18.13).

Control engineer's question	General form	Mathematical form	Category	Computation	Answered in
Is the plan one fit?	two problems, expected consequence, (2.23)	rank of past rows on the future moves, across copies	design	expected consequence	18.3.1
What does the plan run from?	moments, (2.10)	filter of the fit, $(\hat{\mathbf{z}}_k, \boldsymbol{\Sigma}_k)$	estimation	posterior	18.3.2
How good are its rows?	moments of a prediction, (2.10)	rows and prediction model against the plant, noiseless	design	posterior	18.3.3
What is the plan over the next hour?	expected consequence with limits, (2.23)	linear rows across copies, min $\mathbb{E}$ cost	design	expected consequence	18.3.4
What do 95 % limits per copy give?	joint chance constraint, (2.12)	$\Pr_J$ at $z = 1.645$	prediction under a threshold	posterior	18.3.5
What does holding the joint cost?	the same, by z, (2.12)	Bonferroni against calibrated, across copies	design	expected consequence	18.3.6
What does a fixed plan give up?	expected consequence against a bound, (2.23)	perfect knowledge of the draw; re-planning's first number	design	expected consequence	18.3.7
Do the stop, capacity, mass, tank matter, and when?	what-ifs, (2.23)	pinned moves at scenario nodes, scan by window	design	expected consequence	18.3.8
Does it hold on fresh records?	threshold on fresh records, (2.12)	coverage	prediction under a threshold	posterior	18.3.9

$\hat{p}_j$:

$$\begin{aligned} \tfrac{C_h}{\Delta}(T_{r,h,j} - T_{r,h,j-1}) + Q_{c,h,j} &= \hat{Q}_{h,j} + Q_{o,h}, \\ \tfrac{C_w}{\Delta}(T_{s,j} - T_{s,j-1}) - Q_{c,A,j} - Q_{c,B,j} - q_{t,j} + Q_{ch,j} &= 0, \\ E_j - (1 - \lambda\Delta)E_{j-1} - \Delta(\eta_c q^{+}_{t,j} - q^{-}_{t,j}) &= 0, \\ Q_{c,h,j} &\le u_{\max}\hat{G}_h(T_{r,h,j} - T_{s,j}), \\ T_{r,h,j} - (1 - \hat{\beta}_k)\, Q_{c,h,j} / \textstyle\sum_u C_u &\le 27 - z\, s_{k,j}. \end{aligned} \tag{18.4}$$

The moves are recovered from the path by the hold share a and the water's lag a_w,

$$r_{h,j} = T_{r,h,j-1} + \frac{T_{r,h,j} - T_{r,h,j-1}}{a} + \hat{b}_h, \qquad T_{sp,j} = T_{s,j-1} + \frac{T_{s,j} - T_{s,j-1}}{a_w}, \tag{18.5}$$

and bounded in range and in change per window. The tightening is the state route's spread, with the gradient of equation (18.2) in the hold regime, on (β_k, b_h, Q_h),

$$s^2_{k,j} = \mathbf{g}^{\mathsf{T}}_{k,j}\boldsymbol{\Sigma}_{k,j}\,\mathbf{g}_{k,j}, \qquad \mathbf{g}_{k,j} = \Big(\frac{\hat{Q}_{h,j} + Q_{o,h}}{\sum_u C_u},\ -1,\ -\frac{1 - \hat{\beta}_k}{\sum_u C_u}\Big)^{\mathsf{T}}, \tag{18.6}$$

where $\boldsymbol{\Sigma}_{k,j}$ holds the belief's block on (β_k, b_h) and, for the load, the forecast's band with the fit's load carried by its autoregressive term.

The objective. The plan minimizes, over the moves,

$$\sum_{j=1}^{12} \hat{p}_j\, \Delta\, \frac{Q_{ch,j}}{\mathrm{COP}(T_{s,j}, \hat{T}_{\mathrm{amb},j})} - V(E_{12}) - \kappa\Big[\sum_h C_h(T_{r,h,0} - T_{r,h,12}) + C_w(T_{s,0} - T_{s,12})\Big], \tag{18.7}$$

with a price on any row the first moves cannot reach and a tie-break on moves. V is the value of what the tank holds at the end and κ its slope, which also prices the cooling left in hall and water. The cost's product of Q_{ch} and $1/\mathrm{COP}(T_s)$ is bilinear. It is linearized at the last iterate inside a trust region on T_s and accepted on the true objective, which is sequential linear programming.

The modeling choices. Each buys something and each can fail.

- *Twelve windows.* The outline's hour. It buys a plan a program solves in seconds, and it cannot price its own store, as Section 18.3.4 shows.
- *A walk on g of* 0.005 *per window.* This is the forward band of a coil that may change, a stated form, and a coil with another history would break it.
- *Valve row* $u \le 0.95$. It keeps the plan off the stop, where equation (18.4) stops being the plant. Section 18.3.8 prices it.
- *Move limits of* 0.5 K *per window on r and* T_{sp}*, with r in* 26 *to* 34°C. They keep the plan inside rows that are right at rest and wrong in transit, as Section 18.3.3 measures.
- *Hold share* $a = 0.80$. One number for a transient that Section 18.3.3 prints as 0.87 to 0.92 with an overshoot.
- *Tank valued by a backward recursion over the rest of the shift*, with the shift's end at 0.07 \$/kWh recharged, on means only. It sets every move of the tank and it is the largest choice here. It is a stand-in that Chapter 21 owns.
- *Slack priced, not refused*, at 20 \$/K and 1 \$/kW per window, with a tie-break of 0.002 \$/K and 2×10^{-5} \$/kW. A row out of reach of the first moves is printed and not hidden.

- *Tie-break on travel.* It prices the travel of T_r, T_s and q_t and enters no reported cost. Among plans of nearly equal net it takes the one that moves least, and where a set point is almost free in cost it is what decides that set point, as Section 18.3.6 prints.
- *One multiplier z on every row's spread, target* 0.95 *on inlets alone.* This buys a one-dimensional search, and it gives up any allocation by rack, which Section 18.6.1 prices.

18.2 The posterior of the state (step 3)

The belief is computed once per window and every answer reads it. The future's rows are not in it, which is the first answer.

18.2.1 The rows of one window

On the reduced state of Chapter 6, with the return air a function of the entries by the fit's quasi-steady row, a window k has

$$T_{r,h}(\mathbf{z}_k) = T_s + \frac{Q_h + Q_{o,h}}{u_h G_h}, \qquad G_h = \sum_{u \in h} G_{\mathrm{rated},u}\, \mathrm{e}^{g_u}. \tag{18.8}$$

Its residual vector, each row divided by its sensor's spread, is

$$\mathbf{r}(\mathbf{z}_k) = \begin{bmatrix} (T_s - y_{Ts})/\sigma_{Ts} \\ (T_{r,h}(\mathbf{z}_k) + b_h - y_{Tr,h})/\sigma_{Tr} \\ (T_{d,u}(\mathbf{z}_k) - y_{Td,u})/\sigma_{Td} \\ (Q_h - y_{Q,h})/\sigma_Q \\ (Q_A + Q_B + Q_{o,A} + Q_{o,B} - y_{\mathrm{loop}})/\sigma_{\mathrm{loop}} \\ (T_{\mathrm{in},k}(\mathbf{z}_k) - y_{\mathrm{in},k})/\sigma_{\mathrm{in}}, \quad k \in \{A1, A3, B2, B4\} \\ (u_h - y_{u,h})/\sigma_u \\ \text{priors on } T_s,\ g_u,\ \beta_k,\ b_h \end{bmatrix}, \qquad \mathcal{L} = \tfrac{1}{2}\, \mathbf{r}^{\mathsf{T}} \mathbf{r}, \tag{18.9}$$

which makes 18 reading rows and 19 prior rows. The information matrix of these normalized rows is $\mathbf{J}^{\mathsf{T}}\mathbf{J}$; with unnormalized rows it is $\mathbf{J}^{\mathsf{T}}\mathbf{W}\mathbf{J}$. The filter of Chapter 9 runs window by window: last window's posterior of what persists, (g_u, β_k, b_h), with a walk added on g_u, is the prior row of the next, while T_s, Q_h and u_h are that window's own. Window 1 matches a single fit of that window, with largest mean gap 1.9×10^{-8} and largest spread gap 1.1×10^{-8}.

18.2.2 The posterior at the plan's start

Table 18.2 runs Gauss–Newton at window 12 on the book's record, from last window's posterior and this window's meters. Each row is an iteration, with the entries the block below uses, the objective and the norm of its gradient. The filter takes four.

Table 18.2: Gauss–Newton iterations of the rows (18.9) at window 12, book's record. Six of the 23 entries are shown.

k [–]	T_s [°C]	Q_B [kW]	β_{B4} [–]	β_{B3} [–]	b_B [K]	u_B [–]	$\mathcal{L}$ [–]	$\lVert\nabla\mathcal{L}\rVert$ [–]
0	14.0000	278.7679	0.3957	0.5000	−0.2717	0.8475	12.9669	5.44×10^{2}
1	14.1660	282.0725	0.3971	0.5000	−0.2687	0.8459	5.4524	9.25×10^{-1}
2	14.1658	282.0777	0.3971	0.5000	−0.2688	0.8459	5.4524	3.24×10^{-3}
3	14.1658	282.0778	0.3971	0.5000	−0.2688	0.8459	5.4524	2.01×10^{-5}
4	14.1658	282.0778	0.3971	0.5000	−0.2688	0.8459	5.4524	1.05×10^{-7}

A block of the Jacobian at the last iterate, in the columns $(T_s, Q_B, \beta_{B4}, \beta_{B3}, b_B, u_B)$, is

$$\mathbf{J}_{\text{block}} = \begin{bmatrix} 3.333 & 0.187 & 0 & 0 & 3.333 & -64.711 \\ 3.333 & 0.075 & 0 & 0 & 0 & -64.711 \\ 3.333 & 0.120 & 32.453 & 0 & 0 & -64.711 \\ 0 & 0.250 & 0 & 0 & 0 & 0 \\ 0 & 0.125 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 100.000 \end{bmatrix} \begin{matrix} y_{Tr,B} \\ y_{Td,B2} \\ y_{\text{in},B4} \\ y_{Q,B} \\ y_{\text{loop}} \\ y_{u,B} \end{matrix}. \tag{18.10}$$

The fourth column is zeros. Over all reading rows the largest entry in the column of β_{B3}, and in that of β_{A5}, is 0, and the one row each has is the prior's, $1/0.08 = 12.5$. The return reading alone carries b_B in this block, and the read rack's inlet alone carries its own fraction.

The covariance on $(\beta_{B3}, \beta_{A5}, \beta_{B4}, b_B, b_A, Q_B)$, in mixed units of [–], K and kW, is at windows 12 and 72

$$\mathbf{\Sigma}_{12} = \begin{bmatrix} 0.006400 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0.006400 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0.000103 & 0.000265 & -0.000008 & -0.001354 \\ 0 & 0 & 0.000265 & 0.009964 & -0.000194 & -0.032844 \\ 0 & 0 & -0.000008 & -0.000194 & 0.009971 & 0.006445 \\ 0 & 0 & -0.001354 & -0.032844 & 0.006445 & 10.155714 \end{bmatrix}, \tag{18.11}$$

$$\mathbf{\Sigma}_{72} = \begin{bmatrix} 0.006400 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0.006400 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0.000017 & 0.000048 & -0.000002 & -0.000239 \\ 0 & 0 & 0.000048 & 0.001833 & -0.000039 & -0.006082 \\ 0 & 0 & -0.000002 & -0.000039 & 0.001834 & 0.001205 \\ 0 & 0 & -0.000239 & -0.006082 & 0.001205 & 10.195980 \end{bmatrix}. \tag{18.12}$$

The first two rows and columns do not change in 60 windows. The correlation of β_{B4} with b_B is +0.2611 and +0.2753. The tank's own scalar filter gives 298.84 kWh with spread 1.39 at window 12 and 292.05 with spread 0.59 at window 72, against a reading's 5 kWh.

18.3 The answers (step 4)

Each plan reads the pair $(\hat{\mathbf{z}}_k, \mathbf{\Sigma}_k)$ of Section 18.2.2 and the forecast's rows, and each probability is through the plant's own rows at draws of the pair, with the loops closed by the plant's PI. The draws are 20000, seed 1801, unless stated.

18.3.1 Two problems in one step

Control engineer's question. Is the plan one fit?

General form. An action chosen by its expected consequence, equation (2.23), over a posterior that is an input to it and not the output of it.

Specialized form. Equation (2.23) with a the sequence $\mathbf{u}_{\text{now}+1..\text{now}+12}$ and $p(\mathbf{z} \mid \mathbf{y})$ the belief. Write the rows on the past and the columns of the moves together, $\mathbf{J}_{\text{all}} = [\mathbf{J} \mid \mathbf{0}]$. Since the last block of the pattern (18.3) is zero, the normal equations of the past say nothing about a future move. The restricted problem is

$$\begin{aligned} \min_{\mathbf{u}_{\text{now}+1..\text{now}+12}} \quad & \mathbb{E}\Big[\sum_{j=1}^{12} \ell(\mathbf{z}_{\text{now}+j}, \mathbf{u}_{\text{now}+j}) \;\Big|\; \mathbf{y}_{\text{now}}\Big] \\ \text{subject to} \quad & \Pr_J = \Pr(T_{\text{in},k,j} \le 27^\circ\text{C for all } k, j \mid \mathbf{y}_{\text{now}}) \ge 0.95, \end{aligned} \tag{18.13}$$

the problem (7.1) with its policy line replaced by numbers and its Fit line frozen at the now.

Building blocks. The pattern (18.3), the belief of Section 18.2.2 and the rows (18.4).

The plant's numbers. One window's past rows are 37 on 23 entries and 48 move columns, four moves in twelve copies. The rank is 23 of 71: no reading picks a future move, so "one fit per step" is the wrong words. The same fit gives different plans. With the same mean and $\mathbf{\Sigma}$ entering at $z = 1.645$ against 0, r_B is lower by 1.319 K and T_{sp} by 1.334 K at the last copy; a fit's single answer could not say so. Doubling the cost's weight moves T_{sp} by at most 0.077 K, r by 0.029 K and q_t by 0.0 kW: in the evening the plan sits on limits and not on the price. One step is two problems of different kinds, a filter update of 4 Gauss–Newton iterations on 23 entries and a plan of 6 linear programs on 205 columns. Rows on the past are conditioned; rows on the future are averaged.

18.3.2 The belief a plan runs from

Control engineer's question. What does the plan run from?

General form. Moments of entries under the posterior, equation (2.10).

Specialized form. Equation (2.10) with f an entry of $\mathbf{z}_k$, read as $\hat{z}_i$ and $\sqrt{\Sigma_{ii}}$. A quasi-steady row assumes the hall air at rest. That is what lets one window stand alone in the fit, and it fails in a transit, which the distance below measures.

Building blocks. $\hat{\mathbf{z}}_k$ and $\mathbf{\Sigma}_k$ of Section 18.2.2; the held row $r_h - \hat{b}_h$ for the return, used while the fit's valve is under 0.98.

The plant's numbers. Table 18.3 gives the entries. Each row is an entry, with its stated truth and the mean and spread at the two starts. **Unread entries stay at their priors for ever**: six racks at 0.0800 in every window, while read racks fall to about 0.0041 by window 72. Everything below follows from that. The truth of B3 and A5 lies above the survey, by amounts the record cannot see.

Table 18.3: The belief at the plan's two starts; book's record. Spreads in parentheses.

Entry [unit]	Stated truth	Window 12	Window 72
b_B [K]	−0.25	−0.2688 (0.0998)	−0.1694 (0.0428)
b_A [K]	+0.10	+0.0296 (0.0999)	+0.0334 (0.0428)
β_{B3} [–]	0.57	0.5000 (0.0800)	0.5000 (0.0800)
β_{A5} [–]	0.51	0.4500 (0.0800)	0.4500 (0.0800)
β_{B5} [–]	0.40	0.3500 (0.0800)	0.3500 (0.0800)
β_{A3} [–]	0.40	0.4050 (0.0084)	0.3949 (0.0034)
β_{B4} [–]	0.38	0.3971 (0.0102)	0.3873 (0.0041)

How far from dynamic. At the record's own end-of-window states, the root mean square of $Q + Q_o - Q_c$ is 3.28 and 2.98 kW in halls A and B against a hall meter's spread of 4 kW; as a return-air error, 0.149 and 0.163 K against the reading's 0.3 K. The fit's steady row gives a return 0.186 and 0.173 K from the truth over windows 2 to 144, and the held row $r_h - \hat{b}_h$ gives 0.083 and 0.143 K, which is why the plan starts from the held row.

Calibration, 300 *fresh records at the stated truth, seed* 1803. At windows 12 and 72, b_B lies within two stated spreads on 0.943 and 0.893 of records. The read fractions lie within on 0.923

to 0.960, except β_{A3} at 0.797 at window 72. b_A **is not calibrated**: at window 72 its mean error is −0.1074 K against a stated spread of 0.0428 K (spread across records 0.0437 K), within two on 0.310 of records; at window 12, −0.1092 K and 0.830. This is stated as a finding of the fit. At this truth it makes hall A cooler than believed by that amount, the safe side here, and its cause belongs to a later Part.

18.3.3 The prediction model and the plan's rows

Control engineer's question. How good are its rows?

General form. Moments of a function of the state, equation (2.10), here the gap of a predicted path from the plant's own at the same entries.

Specialized form. Equation (2.10) with f the gap, read at the belief's mean with no noise. Two levels are compared with the plant. The *prediction model* is the plant's own rows (18.1) at a draw of the belief, with the loops closed by the plant's PI and nothing of the truth. The *plan's rows* (18.4) are its hold regime as a linear program. The assumption entering is that a hall moves as a share per window. This buys linearity, and it fails in a transit.

Building blocks. The rows (18.4) and (18.5); the belief's mean entries.

The plant's numbers. The prediction model stays within 0.0002 K of the plant on the inlets for the book's plans, 0.0013 K for a free plan at the edge. For a set-point step in both halls, the share complete at the end of the first window is 0.873 to 0.888 in hall A and 0.907 to 0.923 in hall B, and by the second window it overshoots to 1.057 to 1.093, against the rows' 0.80. The inlet of B3 at the first window stands +0.228 K above its settled value for a 0.5 K step; part of that is the forecast's load rising over the hour, and with the load held flat the overshoot is +0.142 K. Table 18.4 sets the plan's rows against the plant. Each row is a plan; the columns give the largest gap in return air, in inlet over the first three copies and over the last six, and the gap in cost.

Table 18.4: The plan's rows against the plant, noiseless, at the belief's mean entries.

Start [–]	Plan	$\lvert T_r \rvert$ [K]	Inlet, first 3 [K]	Inlet, last 6 [K]	Cost gap [$]
12	book's move limits	0.250	0.276	0.016	−0.0119
12	no move limits	0.249	0.344	0.015	−0.0100
72	book's move limits	0.136	0.136	0.005	−0.0040
72	no move limits	0.146	0.136	0.005	−0.0039
126	book's move limits	0.267	0.161	0.181	+0.0471
126	no move limits	0.825	0.242	0.704	+0.1807

The rows are right at rest and wrong in transit. In the evening and overnight the last six copies are within 0.016 K. At the edge a plan left free moves T_{sp} by up to 3 K a window and leaves its own rows by 0.704 K, so the move limit is what keeps the plan inside its model. All probabilities below are through the prediction model, so the rows' error is priced in z and not ignored.

18.3.4 The plan over the next hour

Control engineer's question. What is the plan over the next hour?

General form. The action minimizing an expected consequence under limits, equation (2.23).

Specialized form. Equation (2.23) restricted by equation (18.13), with the loss the objective (18.7) and the limit carried by $z\,s_{k,j}$ of equation (18.6). The cost is taken on the mean path and the band enters only by the limits, which is coupling 4 in the restricted form; the over-draws columns of the table report the true expectation. At a tank value set by a recursion on means, the plan reads the store's worth as a stated number, and that fails if the rest of the shift differs from its means.

Building blocks. $\hat{\beta}_k$, $\hat{b}_h$, $\hat{G}_h$ and $\mathbf{\Sigma}_k$ of Section 18.2.2; the forecast's means; z from Section 18.3.6; V, concave in the tank's level, and κ, its slope at the start's level, 0.01991 \$/kWh at window 12 and 0.00950 at window 72. At window 12, V is 9.978 \$ at 298.8 kWh and 5.870 \$ at 97.8 kWh. The water's lag is $a_w = 0.9426$, and the plan starts from the fit's T_s, 14.166°C and 13.911°C.

The plant's numbers. Figure 18.2 draws the evening's plans. The horizontal axis of each panel is the copy, one to twelve, and the panels give T_s, r_A, r_B and the tank. The red curve holds no band, the blue curve holds the joint with $z = 2.17$, and the dotted line is the record's supervisor, a labeled reference and not a baseline. The certainty-equivalent plan raises T_{sp} at the move limit to 17.5°C, settling at 17.04°C, with r_A at 33.54 and r_B at 31.70°C. The calibrated plan moves T_{sp} through 14.5, 15.0 and 15.485 and settles at 15.281°C, with r_A from 31.0 to 31.472°C and r_B through 30.0 and 29.829 to 29.961°C. Both discharge the tank at −200 kW throughout, 298.8 to 97.8 kWh. Overnight, at $z = 1.91$, T_{sp} is about 14.43°C, r_A 31.25°C, r_B 30.41°C and the tank idle.

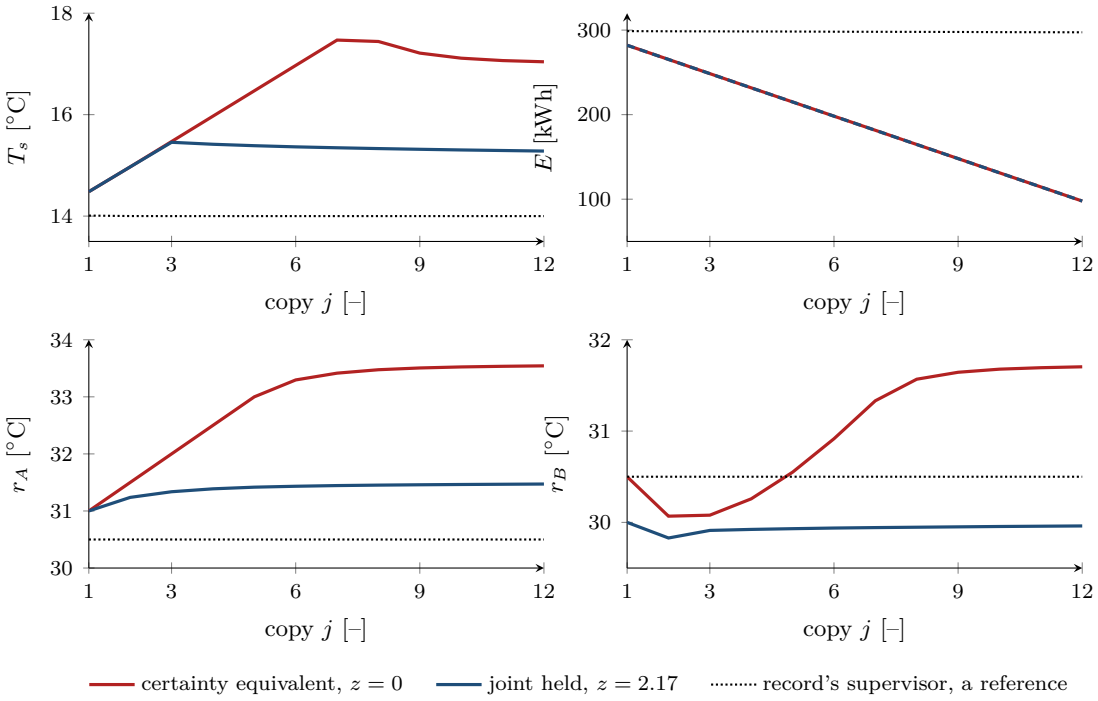


Figure 18.2: The plan from window 12. Top: the water and the tank; the two plans' tank paths coincide, so one is dashed. Bottom: the return set points. Means over draws for the supervisor's T_s and E; the plans' own rows otherwise.

Table 18.5 reports them over draws. Each row is a plan; the columns give the expected electricity, the net cost and the joint. Net means electricity less the value of what is left in tank, water and hall, and it is the only fair column, because the plans leave different amounts behind.

On the calibrated plan's own rows at window 12 it is 14.954 $ of electricity, less V of 5.870 $, less −0.299 $ for the cooling hall and water give up at κ, which is 9.382 $.

Table 18.5: The plans over 20000 draws, seed 1801, standard errors in parentheses.

Start [–]	Plan	Electricity [$]	Net of stored [$]	Joint [–]
12	certainty equivalent	12.845 (0.009)	7.968 (0.009)	0.1615 (0.0026)
12	chance-constrained	14.965 (0.010)	9.394 (0.010)	0.9468 (0.0016)
12	record's supervisor (a reference)	23.917 (0.016)	13.932 (0.016)	0.9082 (0.0020)
72	certainty equivalent	6.167 (0.004)	−0.989 (0.004)	0.3173 (0.0033)
72	chance-constrained	6.109 (0.004)	−0.365 (0.004)	0.9485 (0.0016)
72	record's supervisor (a reference)	6.341 (0.004)	−0.207 (0.004)	0.9534 (0.0015)

The hour's largest lever is priced outside the hour. With the tank pinned to its start, electricity is 21.849 $, against 14.948 $ with no value on what is stored and 14.965 $ at the stated value. Overnight the stated value leaves the tank idle, while no value empties it for 4.250 $. A twelve-window horizon cannot price its own store: the recursion is a stand-in, and Chapter 21 owns it.

18.3.5 What per-copy limits give

Control engineer's question. What do 95 % limits per copy give?

General form. A threshold on the conjunction of events, equation (2.12) with $f = 27°\mathrm{C} - \max_{k,j} T_{\mathrm{in},k,j}$, as in the joint chance constraint (6.20).

Specialized form. Equation (2.12) with $f^* = 0$, read by draws of the belief and the forecast through the prediction model, the plan built with $z = 1.645$ in equation (18.4). The event is over all 120 rows. Three comparisons are read from the same draws: each copy's draws shuffled, which makes copies independent; the racks' events shuffled, which makes racks independent while each stays one event over all its copies; and what persists held at its means.

Building blocks. $\mathbf{\Sigma}_k$; the gradient (18.6), with values below; the delta method (4.53) for the variance budget.

The plant's numbers. At copy 12 from window 12, the gradient of B3's inlet on (β_{B3}, b_B, Q_B) is $(9.9563\,\mathrm{K},\ -1,\ -0.01667\,\mathrm{K/kW})^{\mathsf{T}}$, the load's variance is $9.456\,\mathrm{kW}^2$ and $\mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}\mathbf{g}$ has the parts β 0.6344, b 0.0100 and load $0.0026\,\mathrm{K}^2$, so its spread is 0.8044 K. Table 18.6 gives what the limits give. Each row is a reading, at the two starts.

Figure 18.3 draws the same. The upper left panel has the copy on the horizontal axis and the probability that all racks at that copy are within the limit on the vertical. The upper right has the probability that a rack is within the limit at all copies, for the four racks below one. The lower panel has three bars per start: the actual joint, the joint with racks shuffled and the joint with copies shuffled.

The outline's claim holds, for a different reason than its picture. Per-copy 95 % does not give 95 %: 0.86 and 0.91. But time is nearly one event. β and b persist, with a correlation of 0.99 across the hour, so twelve copies cost little, and were they independent the joint would be 0.2950. **The loss is across racks and in the first moves.** B3 and A5 are unread, independent and both near their limit, because the valve's row pulls hall A up to meet the water hall B permits. Then each row's shortfall comes from the plan's own transient: by copy 0.9709, 0.9428, 0.9086, 0.8682, then about 0.90. With what persists at its means the joint

Table 18.6: Per-copy limits, $z = 1.645$ on all 120 rows. 20000 draws, seed 1801.

Reading	Start 12	Start 72
Joint (se) [–]	0.8579 (0.0025)	0.9146 (0.0020)
B3 over all copies / A5 over all copies [–]	0.9312 / 0.9226	0.9325 / 0.9816
Correlation, B3 copy 1 with copy 12 [–]	0.9914	0.9916
Correlation, B3 with A5, copy 12 [–]	−0.0015	−0.0029
Copies made independent, by shuffling (se) [–]	0.2950 (0.0032)	0.4585 (0.0032)
Racks shuffled, each over all copies [–]	0.8585	0.9150
Spread B3, copy 12: all / what persists only / the rest [K]	0.8066 / 0.8037 / 0.0720	0.8319 / 0.8287 / 0.0741
Rows under 0.95, of 120; smallest [–]	19; 0.9252	11; 0.9354
Union bound on the rows [–]	−0.1803 (empty)	0.2354

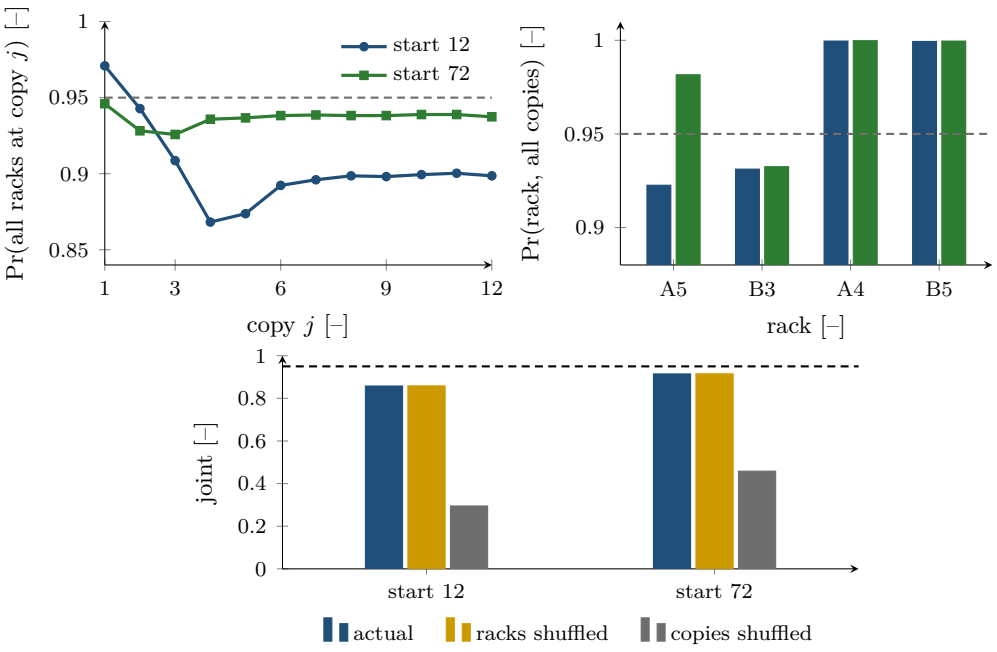


Figure 18.3: What $z = 1.645$ gives. Upper left: all racks by copy. Upper right: each rack over all copies, blue for start 12 and green for start 72; the other six racks read 1.0000. Below: the joint against its two shuffles. The dashed line is 0.95.

is 1.0000, and with only what persists uncertain it is 0.8647 and 0.9187. The union bound on the rows says nothing here, being empty at start 12.

18.3.6 What holding the joint costs

Control engineer's question. What does holding the joint cost?

General form. The threshold of equation (2.12) on the conjunction, held by a choice of tightening, with the consequence of equation (2.23) read for each.

Specialized form. Equation (2.12) on the event of equation (18.13), with one multiplier z on every $s_{k,j}$ in the rows (18.4). Bonferroni's rule takes $z = \Phi^{-1}(1 - 0.05/120) = 3.341$. Direct calibration finds, by bisection on common draws, the smallest z whose plan's joint on those draws reaches 0.95: 5000 draws, seed 18, evaluated on 20000 others, seed 1801. One multiplier assumes every row should give up in proportion to its spread. It buys a one-dimensional search, and so a calibrated plan need not be the cheapest plan that holds the joint.

Building blocks. The belief, the rows (18.4), the spread $s_{k,j}$ of equation (18.6) and the prediction model.

The plant's numbers. Table 18.7 gives the four routes. Each row is a route at a start; the columns give z, the joint, the back-off on B3 at copy 12, the last r_B and T_{sp}, the net cost and the net against the calibrated plan. Figure 18.4 draws net against joint, one panel per start.

Table 18.7: The routes to the joint. Standard errors in parentheses; net over calibrated is paired on the draws.

Start [–]	Plan	z [–]	Joint [–]	Back-off B3 [K]	r_B last [°C]	T_{sp} last [°C]	Net [$]	Over calibrated [$]
12	certainty equivalent	0	0.1615 (0.0026)	0	31.704	17.040	7.968	−1.426 (0.001)
12	per-copy	1.645	0.8579 (0.0025)	1.323	30.385	15.706	8.987	−0.407 (0.000)
12	**directly calibrated**	2.172	0.9468 (0.0016)	1.747	29.961	15.281	9.394	0
12	Bonferroni	3.341	0.9807 (0.0010)	2.688	29.021	14.270	10.487	+1.093 (0.001)
72	certainty equivalent	0	0.3173 (0.0033)	0	31.992	16.016	−0.989	−0.624 (0.000)
72	per-copy	1.645	0.9146 (0.0020)	1.363	30.630	14.650	−0.452	−0.087 (0.000)
72	**directly calibrated**	1.906	0.9485 (0.0016)	1.580	30.413	14.434	−0.365	0
72	Bonferroni	3.341	0.9916 (0.0006)	2.769	29.224	13.243	0.094	+0.460 (0.000)

Against the per-copy plan the joint costs 0.407 $ and 0.424 K of r_B in the evening, and 0.087 $ and 0.217 K overnight. **Bonferroni pays for** 120 **events where there are about two.** It costs a further 1.093 and 0.460 $, and 0.940 and 1.189 K. It also cannot be reached in the first copies within the move limits, with slack of 0.842 and 0.765 K, which is why its joint is 0.9807 and not nearer one. It is safe, not wrong. Out of sample the calibrated joint is 0.9468, in sample 0.9504, up to two standard errors under 0.95, and this is stated. **The plan holds the directly calibrated joint at one now**: a statement on the belief and forecast at that now, about a sequence not revised after it. It is not a guarantee of a plan re-solved each step, which is Chapter 20's.

The overnight numbers are conditional on the tie-break. In the evening hall A's return is pinned by its rows, and removing the tie-break changes nothing. Overnight hall A's set

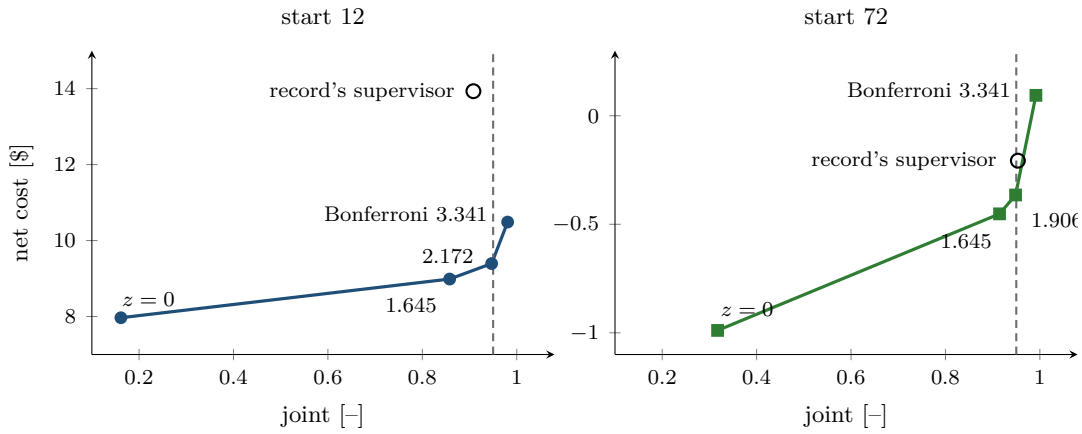


Figure 18.4: Net cost against the joint, by route. The filled marks are the plans of one rule at four values of z. The open circle is the record's supervisor, a reference. The dashed line is 0.95. The axis is linear; the routes crowd near one.

point is almost free in cost, so the tie-break chooses it. With the tie-break removed at the same z, the solver's plan takes r_A up to 31.865°C against 31.333°C, at a paired net of −0.0002 $, and the joint falls from 0.9485 to 0.9268 with A5 over all copies from 0.9895 to 0.9668. Re-calibrated without it, z is 2.055 and net −0.314 $. At $z = 1.645$ the same freedom gives 0.8845 against 0.9146, and A5 0.9495 against 0.9816. So $z = 1.906$, the overnight row of Table 18.6 and the overnight cost of the joint are statements about the plan that moves least among plans of nearly equal net. That is a stated rule of the plan and it must travel with the plan.

18.3.7 What a fixed plan gives up

Control engineer's question. What does a fixed plan give up?

General form. The expected consequence of an action against that of an action told more, equation (2.23).

Specialized form. Equation (2.23) with a the calibrated fixed plan, against the plan re-solved at each draw of the entries, of the exogenous paths or of both, paired over 120 draws, seed 1802. A plan told the entries sits on 27°C with no margin, and the loop's scatter then breaks it (largest inlet mean 27.196°C): its joint is not comparable and it is a lower bound on cost only, the physical limit and no achievable plan. A plan told only the paths keeps z.

Building blocks. The belief, the rule of Section 18.3.6 and the prediction model.

The plant's numbers. Table 18.8 gives the net cost. Each row names what the plan is told; the last column is the fixed plan over that row, paired.

The gap to the limit is not the future: it is the entries no reading touches. Knowing the hour's load and price is worth 0.003 and 0.081 $, and knowing β, b and g is worth nearly the whole gap. The flat tariffs of these two starts do not decide this: at the price edge, window 126, told paths are worth 0.047 $ (se 0.012, 40 draws). The fixed plan's 9.646 and −0.255 $ are the luck of 120 draws, 2.0 and 2.1 standard errors above the 9.394 and −0.365 $ of 20000 draws, with joints of 0.9750 and 0.9667 against 0.9468 and 0.9485; the paired gaps are unaffected. Figure 18.7, in Section 18.6.1, sets these bars beside the measurements that could buy them.

Table 18.8: What a fixed plan gives up, against plans told the draw. Paired, 120 draws, seed 1802; standard errors in parentheses.

Start [–]	Plan told	Net [$]	Fixed over this [$]
12	nothing, the fixed plan	9.646 (0.126)	–
12	everything	8.444 (0.120)	+1.202 (0.042)
12	β, b, g only	8.453 (0.120)	+1.193 (0.042)
12	load, ambient, price only, z kept	9.643 (0.126)	+0.003 (0.005)
72	nothing, the fixed plan	−0.255 (0.052)	–
72	everything	−0.820 (0.055)	+0.565 (0.020)
72	β, b, g only	−0.769 (0.058)	+0.514 (0.021)
72	load, ambient, price only, z kept	−0.336 (0.049)	+0.081 (0.007)

One move, then re-plan. Only the first number belongs here. On 24 fresh records at the stated truth, seed 1806, with the same rule z and 12 windows counted, fixed less re-planned is a net of −0.039 $ (se 0.019) in the evening and −0.069 $ (se 0.007) overnight. **There is no gain at this truth**, slightly the other way, with all 24 inside the limit on both routes. This is consistent with Table 18.8, since a re-plan reads no row on B3 or A5, but it is not decomposed here, it is at one truth, and the count is small. Chapter 20 owns it.

18.3.8 Which moves matter, and when

Control engineer's question. Do the stop, capacity, mass, tank matter, and when?

General form. The expected consequence of a what-if, equation (2.23), read at scenario nodes.

Specialized form. Equation (2.23) with a the plan under one changed line of equation (18.4) and the joint re-held by its own calibration, except in the scan, which holds the rule $z = 2.17$ on 2000 draws of each window's belief. The comparison is of plans each holding the joint. Reading it as a saving against the supervisor assumes the supervisor is only a reference, and its joint is printed beside.

Building blocks. The rule of Section 18.3.6; the rows with $u_{\max}$, capacity, C_h, C_w or the move limit changed.

The stop. Table 18.9 sets the valve's row at 1.00, 0.95 and 0.90. Each row is a plan; the columns give z, the joint, the net, the share of periods at the stop in each hall and the probability a hall touches the stop at all. At 1.00 the plan rides the stop, with 25 to 58 % of periods there against 1 to 6 % at 0.95; that is Chapter 15's kink, with the valve as a protected quantity. The row at 0.95 does not keep the valve from touching: a hall touches in 0.750 to 1.000 of draws, and only the row at 0.90 brings that to 0.005 at the two flat starts. What it buys is time off the stop. It is a choice with a price, 0.25 to 0.54 $ per hour.

Capacity. The chiller reaches 800 kW in the plans started at windows 84 to 114, while they charge: the plan from window 114 charges for six copies at 800 kW and then discharges, with a mean q_t of −60.8 kW. A what-if of 1000 kW at window 96 moves the mean charge from 101.9 to 166.7 kW, with a largest duty of 887.1 kW, and it is worth 0.005 $ per hour on the rows. At the one window priced, it does not matter.

Table 18.9: The valve's row, joint re-held at 0.95. Standard errors of the joint at most 0.0016.

Start [–]	$u_{\max}$ [–]	z [–]	Joint [–]	Net [$]	Share at stop, A / B [–]	Pr(hall ever at stop) [–]
12	1.00	2.250	0.9463	8.858	0.252 / 0.360	1.000
12	0.95	2.172	0.9468	9.394	0.008 / 0.018	0.750
12	0.90	2.094	0.9463	10.091	0.000 / 0.000	0.005
72	1.00	2.023	0.9494	−0.617	0.434 / 0.451	1.000
72	0.95	1.906	0.9485	−0.365	0.015 / 0.017	0.847
72	0.90	1.867	0.9499	−0.044	0.000 / 0.000	0.005
126	1.00	2.102	0.9492	7.341	0.346 / 0.575	1.000
126	0.95	2.039	0.9488	7.786	0.056 / 0.064	1.000
126	0.90	2.398	0.9499	8.503	0.001 / 0.010	0.622

Thermal mass. Each plan holding the joint, the storage rows of hall and water are worth 0.39 % and 0.10 % of the hour's cost at starts 12 and 72. **At the edge the plan does precool**: from window 126, T_s falls to 11.873°C and r_A to 29.571°C before 06:00, worth 0.0886 $ (se 0.0004), 0.86 %, at the book's limits. With the move limits relaxed, net is 7.786, 7.656, 7.412 and 7.410 $. The extra 0.376 $ needs moves of 2 to 3 K a window, where the rows err by 0.704 K; it is stated and not taken.

The scan. Figure 18.5 takes every sixth window as the start. The horizontal axis is the window, the upper panel's vertical axis the net dollars saved against the supervisor as the labeled reference, over that supervisor's electricity for the hour, by three plans: all moves free, set points only, and tank only. The tank-only plan keeps the supervisor's set points, so no valve row or back-off binds it. The lower panel is the tariff. All moves save 15.7 to 21.3 % from 19:30 to 22:30, the tank 10 to 17 % and the set points 6 to 8 %. From 01:00 to 05:00 all moves change the cost by −1.8 to +1.8 % against the reference, at the rule. At the 06:00 price rise moves matter again, 11.0 and 12.0 % saved.

Table 18.10 reads those late windows with every column's joint, because one rule does not hold one joint across the shift. Each row is a start; the columns give the joints, the shares saved at the rule, and all moves with the joint re-held for that window. **From** 02:30 **to** 04:00 **the rule over-tightens**: all moves and set points hold about 0.97 against the supervisor's 0.95, so those set-point costs are not plain losses, and part of "moves barely matter" is the margin; at window 96, +0.0 % at 0.974 becomes +2.5 % at 0.942 with $z = 1.766$. **At windows** 114 **and** 120 **the rule misses**, at 0.943 and 0.940 against the supervisor's 0.956, and there all moves and set points lose on both axes: plain losses. Re-holding the joint makes them larger, −1.6 and −3.2 %. At window 132 the rule and the supervisor both read 0.939.

Table 18.10: The scan's late windows with every joint; 2000 draws, se about 0.005. Saved is net dollars over the supervisor's electricity, positive saves; rule $z = 2.17$. The last three columns re-hold the joint for that window on 2000 other draws, seed 18.

	Joint at the rule [–]				Saved at the rule [%]			All moves, re-held		
Start [–]	Supervisor	Set points	All	Tank	Set points	All	Tank	z [–]	Joint [–]	Saved [%]
90	0.952	0.976	0.975	0.952	−0.7	+0.2	+1.1	1.773	0.942	+2.8
96	0.952	0.974	0.974	0.952	−1.7	+0.0	+2.1	1.766	0.942	+2.5
102	0.952	0.975	0.974	0.951	−2.0	+1.8	+4.2	1.750	0.943	+4.1
108	0.952	0.973	0.972	0.951	−3.8	+0.4	+5.4	1.711	0.945	+3.1
114	0.956	0.949	0.943	0.955	−3.7	−1.1	+2.6	2.312	0.947	−1.6
120	0.956	0.940	0.940	0.956	−5.1	−1.8	+2.7	2.453	0.947	−3.2
126	0.949	0.954	0.953	0.949	−4.0	+11.0	+13.2	2.055	0.947	+11.4
132	0.939	0.939	0.939	0.939	−7.2	+12.0	+17.4	2.445	0.945	+10.7

From window 114 **the tank alone beats all moves**, a nested subset of them, at no worse a joint at 114, 120 and 132; at 126 it saves 13.2 against 11.0 % at 0.949 against 0.953. The tank carries the price rise, and the plan's other moves cost it: its rows have slack out of reach, 0.316 to 0.486 K over those windows, and the valve row forbids the supervisor's stance. This says the plan's rows need the first moves it cannot make there, and not that the supervisor is a baseline: its joint falls under 0.95 at the last window too.

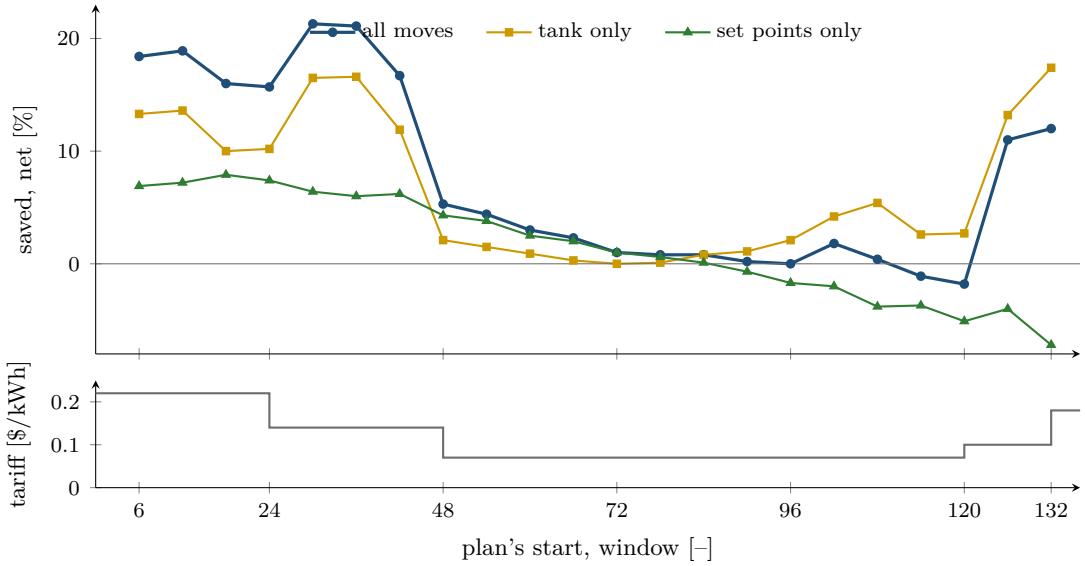


Figure 18.5: Net dollars saved against the record's supervisor as a labeled reference, over its electricity for the hour, every sixth window, rule $z = 2.17$, 2000 draws; the joints differ, Table 18.10. Below: the tariff by window's end. Window 12 is 20:00, window 72 is 01:00 and window 132 is 06:00.

18.3.9 Fresh records

Control engineer's question. Does it hold on fresh records?

General form. The threshold of equation (2.12), compared with how often the event happens when records are drawn.

Specialized form. Equation (2.12) with the event of equation (18.13), each record's own fit and own plan, applied open-loop through the plant's simulation. The two questions differ. At the stated truth the count reports that truth's outcome, one draw of the unread entries. With the truth drawn from the stated priors the count is coverage, and it tests the joint's number.

Building blocks. Each record's belief by the filter of Section 18.2.1; the rule z; the plant.

The plant's numbers. Table 18.11 gives 100 records at the stated truth, seed 1803. Each row is a plan at a start; the columns give the share with all inlets within the limit, the largest inlet's mean and the costs. **The truth is one point**: $\beta_{B3} = 0.57$ sits 0.875 stated spreads above the survey, so the certainty-equivalent plan is over 27°C on every record and every banded plan is under on every record. That is this truth's outcome and not coverage.

Table 18.12 draws the truth: β, b and g from the stated priors, with each record's fit and plan, $n = 200$, seed 1804. Each row is a start; the columns give the realized joint for the two rules.

Table 18.11: Fresh records at the stated truth, $n = 100$, seed 1803. Standard errors in parentheses. A share of 1.000 is not coverage.

Start [–]	Plan	Pr(all ≤ 27) [–]	Mean largest inlet [°C]	Electricity [$]	Net [$]
12	certainty equivalent	0.000	27.891	13.103 (0.108)	8.252 (0.108)
12	per-copy	1.000	26.568	14.702 (0.122)	9.325 (0.121)
12	directly calibrated	1.000	26.175	15.310 (0.127)	9.762 (0.127)
12	Bonferroni	1.000	25.888	16.876 (0.139)	10.937 (0.138)
12	record's supervisor	1.000	26.574	24.204 (0.196)	14.263 (0.196)
72	certainty equivalent	0.000	27.921	6.163 (0.050)	−0.976 (0.050)
72	directly calibrated	1.000	26.317	6.152 (0.050)	−0.330 (0.050)

Table 18.12: Coverage, truth drawn per record. Target 0.95. Standard errors in parentheses, and the distance from 0.95 in standard errors in brackets.

Start [–]	Per-copy [–]	Directly calibrated [–]
12	0.8800 (0.0230) [−3.0]	0.9550 (0.0147) [+0.3]
72	0.8700 (0.0238) [−3.4]	0.9250 (0.0186) [−1.3]

Per-copy misses by 3.0 and 3.4 standard errors. Calibrated, with z taken from the book's record and used as a rule, stands within −1.3 and +0.3. The 0.9250 is not distinguished from 0.95 at this count, nor is it confirmed.

18.4 The checks (step 5)

The tests stand against the plan's independence of the truth, its local optimality and the prediction model.

18.4.1 The plan and its rows

Table 18.13 lists them. Each row is a check, with what is compared and the result.

18.4.2 What the checks do not cover

The rows of Section 18.2.1 are on the reduced state, whose limit of exact physics rows Chapter 6 checked; the limit is not repeated on this plant. The calibration is checked out of sample on the draws and not against an exact posterior: the Laplace normal of the belief is not reweighted here. The four shifts of the T_{sp} path and the three starts test a neighborhood, so the plan is a local optimum checked locally.

18.5 The report (step 6)

The plan is two problems: the fit gives the belief, and a linear program over the future copies gives the moves. Holding the joint needs z of 2.17 and 1.91, the second under the stated tie-break, not 1.645 and not 3.341. What limits the plan is two racks no reading touches. Against the shift plant of the companion book, what the loops add is a set point that takes time to land, so that the rows are right at rest and wrong in transit, and a stop the plan pays to stay off.

Table 18.13: The checks of the chapter.

What is compared	Result
The plan reads no truth: stated truth changed, readings fixed	moves, belief and inlets, gap 0.0
Linearization from $T_s = 10$, 14, 18°C	same net 9.38206 \$, gap 0.0; a local check, not a global proof
T_{sp} path shifted -0.3, -0.1, $+0.1$, $+0.3$ K, the rest free	net higher by 0.30, 0.098, 18, 76 \$; upward shifts pay slack of 0.099 and 0.297 K
$\mathbf{\Sigma} \times 10^{-8}$	largest tightening left 0.1800 K, the forecast's
$z = 0$ with the band zero against the certainty-equivalent plan	gap 0.0
Draws against the state route, $z = 1.645$, copy 12	B3 mean 25.6808 against row 25.6769°C; spread 0.8066 against 0.8044 K; A5 0.9635 against 0.9552 K
Forecast rows, lead 12, 40000 draws, against stated	Q_A spread 3.067 (3.067); halls' correlation 0.0483 (0.0479); ambient spread 0.798 (0.800); price log spread 0.1309 (0.1300)
Prediction model, one Runge–Kutta step per 30 s, against the plant's 10 s	inlet gap 1.8×10^{-4} K; cost gap -1.2×10^{-6} \$
Batch simulation at one stated truth against the plant's, 3 records, 6 windows	readings and inlets, gap 0.0
Filter's window 1 against a single fit	mean gap 1.9×10^{-8}; spread gap 1.1×10^{-8}; tolerance 10^{-6}
Tie-break removed, start 12, where the rows pin r_A	net change 0.0; r_A at the last copy 31.472 both ways
Tie-break removed, start 72, same z	net -0.0002 \$; r_A up to 31.865°C; joint 0.9268 against 0.9485. The overnight z is conditional on it

18.5.1 The answers and their bands

Table 18.14 gives the answers in the order of Table 18.1. Each row is a question; the columns give the answer and its band or basis. The stated truth of the unread entries and the offsets is the table's last row, and Figure 18.6 draws it against the belief at window 12: each panel is the normal of an entry, with the entry on its horizontal axis and the truth a dotted line. The truth is shown to test the method, and no plan read it.

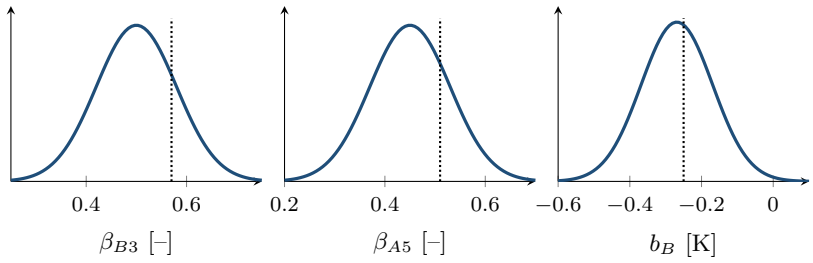


Figure 18.6: The belief at window 12 against the stated truth (dotted). Curves are $\mathcal{N}(\hat{z}_i, \Sigma_{ii})$ with means 0.5000, 0.4500, -0.2688 and spreads 0.0800, 0.0800, 0.0998. The vertical scale is density. No plan read the dotted lines.

18.5.2 The entry in the commissioning log

The entry records the joint limit held at one now by direct calibration, 0.95 on inlets, on stated priors, with z of 2.17 and 1.91 for the two starts on this record, the second for the plan the tie-break chooses. It records the valve row at 0.95, the move limits of 0.5 K per window and the tank's value as stated, a stand-in owned by Chapter 21. It says that the plan is not to be read

Table 18.14: The report for the hour's plan. Percentages and dollars are this example's.

Control engineer's question	Answer	Band or basis
Is the plan one fit?	No: rank 23 of 71. One step is a filter update and a plan of 6 linear programs	the move columns are zero in every past row
What does the plan run from?	A filter of the fit; six racks at 0.0800 for ever; b_B -0.2688 and -0.1694 K	spreads 0.0998, 0.0428 K; b_A not calibrated, within two on 0.310
How good are its rows?	Prediction 1.8×10^{-4} K from the plant; rows within 0.016 K at rest, 0.276 K in transit, 0.704 K with no move limit at the edge	noiseless, belief's mean
What is the plan over the next hour?	T_{sp} to 15.281°C, r_A 31.472, r_B 29.961°C, tank -200 kW; net 9.394 \$. Overnight 14.434, 31.249, 30.413°C, tank idle; net -0.365 \$	se 0.010, 0.004 \$; tank's value stated
What do 95 % limits per copy give?	0.8579 and 0.9146, lost across racks and in transit, not across copies	se 0.0025, 0.0020; copies shuffled 0.2950, 0.4585
What does holding the joint cost?	0.407 and 0.087 \$ over per-copy; Bonferroni a further 1.093 and 0.460 \$	joint 0.9468, 0.9485, at one now; se 0.0016
What does a fixed plan give up?	At most 1.202 and 0.565 \$ against perfect knowledge, nearly all in β, b, g; the hour's load and price 0.003, 0.081 \$; one move then re-plan no gain at this truth	se 0.042, 0.020 \$; a bound, not an available plan; re-planning on 24 records
Do the stop, capacity, mass, tank matter, and when?	Valve row 0.25 to 0.54 \$/h; capacity 0.005 \$/h; mass 0.10 to 0.86 %; all moves 15.7 to 21.3 % of the evening, ±1.8 % from 01:00 to 05:00; from 04:30 the tank alone beats all moves	against the supervisor's electricity, a reference; rule $z = 2.17$, at which all moves' joint runs 0.939 to 0.975
Does it hold on fresh records?	Coverage per-copy 0.87 and 0.88; calibrated 0.925 and 0.955	$n = 200$; se 0.015 to 0.024
Stated truth	$\beta_{B3} = 0.57$, $\beta_{A5} = 0.51$, $b_B = -0.25$ K, $b_A = +0.10$ K	never read by a plan

as a guarantee once it is re-solved. It proposes permanent inlet readings at racks B3 and A5. Each of these rests on a constructed record.

18.5.3 Common mistakes

Table 18.15 lists six mistakes this plant makes easy. Each row names the mistake, its form here and the correction.

Table 18.15: Six mistakes in the hour's plan.

Mistake	On this plant	Correction
"One fit per step"	rank 23 of 71; the same fit gives plans 1.319 K apart in r_B	a fit and a plan: rows conditioned, rows averaged
Per-copy 95 % read as the hour's	0.8579, 0.9146	calibrate the joint, or state what is held
Blaming the count of copies	shuffled 0.2950 against actual 0.8579	the copies are nearly one event; look across racks and at the first moves
Reading Pr = 1.000 on fresh records as coverage	every banded plan under 27°C at this truth	draw the truth: 0.9250 to 0.9550
Comparing plans on electricity alone	12.845 against 14.965 $ leaves different stores	compare net of what is stored, with the joint beside it
Reading one multiplier's zero as the reading's worth	B3 alone −0.004 $ at one z, +0.509 $ at a stated 0.8 K back-off	allocate before ranking; pairs still lead, and savings do not add

18.6 The next measurement (step 7)

The gap to perfect knowledge lies in entries no reading touches, so the candidates are rows on them, priced by what the plan then costs.

18.6.1 The candidates

Each candidate updates $\mathbf{\Sigma}$ by its row through the preposterior (4.73), on frozen rows. The joint is re-calibrated on 4000 draws, seed 1807, with the means held, and the saving is net against no candidate over one hour. Table 18.16 ranks them. Each row is a candidate; the columns give the spreads of the two racks at the two starts and the saving. Figure 18.7 puts the same savings beside the give-up of Table 18.8. The horizontal axis is dollars per hour, with one panel per start. The upper three bars are what a plan told the draw would save, a bound; the lower five are what a candidate saves.

Under one multiplier B3 alone saves nothing in the evening, and that is the multiplier's doing. Once B3 is read its spread is small, and the margin it still needs against the loop's scatter and the plan's transient can be bought only by raising z on every row, from 2.195 to 2.430, which over-protects A5, at 0.9896 against B3's 0.9700. Table 18.17 gives the read rack a stated back-off at every copy and the other racks z times their spread, on the same draws. Each row is an allocation; the columns give z, the joint, each rack over all copies, the net and the saving. Then B3 alone saves +0.384 to +0.509 $ and A5 alone +0.239 to +0.296 $, the larger figures at joints slightly under the target. The water does couple the halls, since 0.509 and 0.296 sum to less than the pair's 1.153, and two readings remain the best candidate, with no gain from allocating them on this coarse grid. A sharper offset or meter buys nothing here. The third decimal is not resolved: re-calibration alone on other draws moves the row of no candidate from 9.394 $ at $z = 2.172$ to 9.445 $ at $z = 2.195$, a change larger than −0.004 and +0.011.

Two readings are not worth the whole gap. The 1.153 $ holds the means, and the 1.202 $ of perfect knowledge is averaged over draws; the bound's own means-held counterpart is

Table 18.16: The candidates, frozen rows, means held, one multiplier; 8000 evaluation draws; saving against none, net, one hour. A permanent reading's row is credited with all the windows since 19:00.

Candidate	sd β_{B3}, β_{A5} [–] start 12 / 72	Start 12 [$]	Start 72 [$]
Permanent inlet readings, B3 and A5, 0.3 K	0.0088, 0.0074 / 0.0034, 0.0029	**+1.153**	**+0.519**
Airflow re-survey, six unread racks, sd 0.03	0.0281, 0.0281	+0.903	+0.388
Permanent inlet reading, B3 alone, 0.3 K	0.0088, 0.0800 / 0.0034, 0.0800	−0.004	+0.299
Reference return reading, 0.05 K, both halls	0.0800, 0.0800	+0.011	+0.000
Hall meters to 1 kW	0.0800, 0.0800	−0.000	+0.000

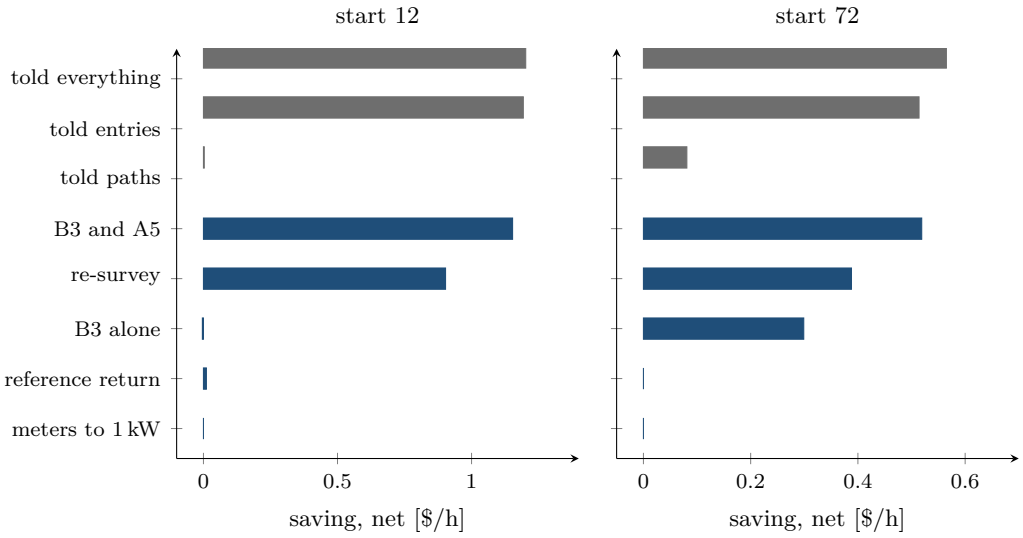


Figure 18.7: The give-up against the candidates. Gray: a plan told the draw (load, ambient and price; β, b, g; everything), a bound from 120 paired draws. Blue: a candidate, means held, one multiplier, 4000 draws calibrating and 8000 evaluating. The two sets use different draws and are means-held against averaged, so they are compared by size only and the text below corrects it.

Table 18.17: Allocation by rack at start 12, means held; 4000 draws calibrating and 8000 evaluating, as Table 18.16; se of the joint about 0.0025.

Read	Back-off [K]	z [–]	Joint [–]	B3, all [–]	A5, all [–]	Net [$]	Saving [$]
none	one z	2.195	0.9489	0.9698	0.9789	9.445	–
B3	one z	2.430	0.9470	0.9700	0.9896	9.449	−0.004
B3	0.6	1.961	0.9513	0.9922	0.9695	9.061	+0.384
B3	0.8	1.789	0.9453	0.9999	0.9521	8.936	+0.509
B3	1.0	1.781	0.9491	1.0000	0.9517	8.943	+0.502
A5	one z	2.336	0.9501	0.9759	0.9900	9.458	−0.013
A5	0.6	1.992	0.9484	0.9606	0.9999	9.206	+0.239
A5	0.8	1.898	0.9455	0.9515	1.0000	9.149	+0.296
B3 and A5	one z	2.383	0.9409	0.9889	0.9820	8.292	+1.153
B3 and A5	0.6	0.008	0.9705	0.9984	0.9965	8.406	+1.039

$9.394 - 7.968 = 1.426$ \$. Its joint is also under the target, at 0.9409. Means held flatters, and a reading moves the means. The preposterior of the pair, with 5×5 outcomes at evenly spaced normal quantiles and 2000 draws each, re-holds the joint in each outcome. Against none at those draws, 9.420 \$, means-held saves +1.154 \$ and the preposterior +0.852 \$, a mean net of 8.568 \$ with a range of 8.125 to 9.127 \$ and a mean joint of 0.9510. That is about seven tenths of the bound, on a small grid. Keeping the means-held z in every outcome would seem to save +0.905 \$, but its mean joint is only 0.8774, as low as 0.6270. For B3 alone at 0.8 K, 8 outcomes give a mean net of 9.101 \$, saving +0.319 \$; at one multiplier the same grid gave 9.487 \$, with a range of 9.128 to 9.973 \$, no gain. The truth's hot racks would have moved the plan cooler, so the saving is not this truth's.

18.6.2 The recommendation

Permanent inlet readings at B3 and A5 together, worth about 0.85 \$ per hour in the evening with the joint re-held, and less than the 1.153 and 0.519 \$ of the means-held table. A single reading is worth having only with its back-off allocated, about 0.32 \$. A one-off airflow re-survey comes second by the means-held table, at 0.903 and 0.388 \$ per hour, with no permanent instrument; its preposterior is not computed. A reference return reading does not earn its place for this question. The readings are the permanent instrument of this chapter, and their value is by the plan's cost and on frozen rows, so the number ranks and does not promise.

18.6.3 What the analysis showed

The sheet said one fit per step and a limit per row made the hour safe, and the careful reading said share the risk among all the rows. Neither survives. The plan is a second problem beside the fit, and limits per row fall short, though not through the count of copies the arithmetic blames: time on this plant is nearly one event, so the shortfall lies across two racks nobody reads and in the plan's own first moves, and sharing among all the rows pays for many events where there are about two. The future was not where the money lay. The hour's load and price are worth little to a plan, and the entries the record cannot see are worth nearly all of what perfect knowledge could give, so the lever is instruments and not a better forecast or, at this truth, a re-solve. Two readings recover about seven tenths of that bound once the means are allowed to move, and one is not worthless, as one multiplier made it look. The largest single choice is the value of the tank, which the hour cannot price by itself.

18.6.4 Limits of the model

The belief. It is a Laplace normal, not reweighted; the fit's rows are quasi-steady; and b_A is not calibrated.

The joint. One multiplier z with no allocation by rack means that calibrated is not the cheapest plan holding the joint, and it can misprice a measurement: B3 alone goes from -0.004 to +0.509 \$ with a stated back-off. The overnight z is conditional on the tie-break. The rule $z = 2.17$ does not hold one joint across the shift. The z from one record is used as a rule on others. The joint is stated at one now and on inlets alone, with no cold-side limit.

The plan. It is a local optimum checked locally. The forecast's bands are taken as stated, with paths by the plant's autoregressive coefficients, and the walk on g is a stated form of the forward band. The tank's value is a stand-in. No pump power or part-load curve is modeled. Sensitivity to the twelve-window horizon is not computed beyond the tank's terminal rows.

The counts. Fresh records are at one truth, coverage rests on 200 and re-planning on 24.

The constructed record. The record is constructed from a stated truth, so its agreement tests the method on its own model and not the model against a real plant. Percentages and dollars are this example's.

18.7 What to retain

Definition 18.1 (A fixed plan on the belief)**.** Model predictive control pins moves across future copies of the graph, minimizes an expected consequence over the belief and the forecast, and holds limits on them, equation (18.13). Its moves are numbers, chosen at one now, and rows on the past have no line on them.

Three conclusions carry beyond this plant as statements of form. One step of the controller is two problems, and the band enters the second only through the first's $\mathbf{\Sigma}$. A limit per row does not hold their conjunction, and how far it falls short depends on what the rows share: entries that persist make copies one event, while independent unread entries add. And a plan's gap to perfect knowledge should be split by what is told, since the split ranks measurements against forecasts. How large each of these is here was computed for this plant only. The chapter's questions join the others in Appendix A. The forecast's tubes are Chapter 19's, the plan re-solved at each step Chapter 20's, and the store's own valuation Chapter 21's.

Table 18.18 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 18.18: The terms of the chapter and where each is defined.

Term	Expression	Units	Meaning	Defined in
Two problems	rank of $[\mathbf{J} \mid \mathbf{0}]$, eq. (18.3)	–	past rows have no line on future moves	18.1.4
Plan's rows	eq. (18.4)	mixed	hold regime, linear in the path	18.1.4
Hold share	a, eq. (18.5)	–	share of a set-point move in a window	18.1.4
Value of stored cooling	V, κ, eq. (18.7)	\$, \$/kWh	stated terminal value, a stand-in	18.1.4
Move limit	0.5 K per window	K	keeps a plan within its hold rows	18.3.3
Fixed plan	eq. (18.13)	–	numbers chosen at one now	18.3.1
Per-copy limit	$z = 1.645$ on each row	–	not the joint	18.3.5
Joint limit	$\Pr_J \geq 0.95$, eq. (18.13)	–	all racks, all copies, one now	18.3.1
Tightening	$z\, s_{k,j}$, eq. (18.6)	K	one multiplier on each row's spread	18.1.4
Bonferroni's rule	$\Phi^{-1}(1 - 0.05/120)$	–	120 events, safe	18.3.6
Direct calibration	bisection on common draws	–	smallest z meeting the joint on draws	18.3.6
Net cost	eq. (18.7)	\$	electricity less what is stored	18.3.4
Perfect knowledge	the plan told the draw	\$	a bound on cost only	18.3.7
Coverage	truth drawn, event counted	–	the joint's test; a count at one truth is not coverage	18.3.9

18.8 Exercises

Exercise 18.1. *Formulation.* One tank of water has a level reading, one pump fills it, and the outflow splits to a metered line and an unmetered one by a fraction ϕ known only from a

survey. Write the two problems for a plan of the pump's speed over the next hour, and say which row no reading touches. (The fit conditions the level, flow and speed rows on the past on the state (L, D, ϕ); if only the total outflow enters the level row, no reading has a line in the column of ϕ and its row is the prior's. The plan chooses the future speeds, numbers that have no line in any of those rows, to minimize expected energy under a joint limit on a quantity that depends on ϕ, with the band entering by back-off. No reading touches the fraction's column.)

Exercise 18.2. *Derivation.* Say why the storage rows are linear, and why the coil's row is a limit in the hold regime. Count the rows at twelve copies and derive Bonferroni's z. (Equation (18.1) is linear in T_r, T_s, E and the duties if each duty is an unknown of the program; the product $u_h G_h (T_r - T_s)$ alone is bilinear. In the hold regime u_h is free, so it is removed, and all the row can require is $u_h \le u_{\max}$: $Q_c \le u_{\max} \hat{G}_h (T_r - T_s)$. Ten racks by twelve copies is 120 rows; $\alpha/120 = 0.05/120$ on each gives $z = \Phi^{-1}(1 - 0.05/120) = 3.341$.)

Exercise 18.3. *Derivation.* From the gradient at copy 12 and the entries of $\mathbf{\Sigma}_{12}$, compute the spread of B3's inlet and the calibrated back-off. (With $\mathbf{g} = (9.9563,\ -1,\ -0.01667)^{\mathsf{T}}$, the parts are $9.9563^2 \times 0.0064 = 0.6344$, 0.0100 from b_B, and $0.01667^2 \times 9.456 = 0.0026\,\mathrm{K}^2$, whose sum has the root 0.8044 K. Times $z = 2.172$ this is the back-off 1.747 K.)

Exercise 18.4. *Application.* From the by-rack figures of start 72, 0.9325 and 0.9816, compute the independent product and set it against the joint. (The product is 0.9153. The joint is 0.9146 and the racks shuffled give 0.9150: two racks independent of each other account for it, whereas copies shuffled give 0.4585.)

Exercise 18.5. *Application.* What does a second reading add in the evening, and why is 1.153 \$ not to be set against 1.202 \$? (From Table 18.17: B3 alone saves −0.004 \$ at one multiplier and +0.509 \$ at a 0.8 K back-off, so the second reading adds +0.644 \$ against the allocated first, on means held. The 1.153 \$ holds the means while the 1.202 \$ is averaged over draws; the means-held counterpart of the bound is $9.394 - 7.968 = 1.426$ \$. With the means allowed to move and the joint re-held the pair saves +0.852 \$.)

Chapter 19

Forecast Uncertainty and Constraint Tightening

19.1 A flat trend and a tube on the sheet

The commissioning sheet draws a tube: the band widens along the horizon, so the limit tightens with lead, and a load and weather forecast sets the back-off. The record at 20:00, window 13, reads hall B's return at 30.36°C and valve B at 0.86. Every read inlet stands under 25°C, and rack B3 has no meter. The loop trend is flat and the published ramp is 25 kW per shift. A belief at the end of window 12 stands at that clock, when window 13 begins, so the tables below label 20:00 as window 12; the trend's tables print a window's own start, 19:55.

The engineer would widen the back-off with lead, buy a better forecast, and read a flat trend as nothing to do until the valve nears its stop. Nothing on the trend or the sheet tests any of the three. The question of this chapter is

how the load and weather bands tighten the plan, and what doing nothing risks before the next decision.

Jargon: lead, tightening schedule, tube, violation by a lead, mass already over, crossing proper, part from below, cost of waiting, trend ramp, stance

The *lead* is the count of windows from now to a future copy. A *tightening schedule* gives the back-off of each rack at each copy, a shape times one multiplier z. A *tube* is the predicted band under an affine feedback policy; the hall loop is one such policy. The *violation by a lead* is the probability that any rack is over its limit at any copy up to that lead, the joint tail. The *mass already over* is the part of a threshold probability held by states with no rising trend, which stand over the level without having reached it from below. The *crossing proper* is the part with a rising trend. The *part from below* is the probability that a draw is under at now and over by the lead. The *cost of waiting* compares planning now with planning after one held window. The *trend ramp* is the ramp estimated from the readings so far, and a *stance* is a set of set points left standing.

Key idea

The stop divides this plant in two. Off it, the hall loop holds the return, so an inlet does not read the chiller's set point, more load cools the worst rack, and the band is an unread entry scaled by the load. At the stop the operations book's rows come back. A forecast's band reaches the inlet only through the part of the load an unread fraction lets through, and no policy on these readings lowers an entry none of them reads. A violation by a lead counts racks already over, so it must be read apart from a crossing.

19.1.1 The restriction this chapter works under

This chapter lies on the row "fixed sequence" of Table 7.6, with one affine variant. The restriction of the problem (7.1) is that moves are fixed now or, in the variant, affine in the filtered hall meter and the valve reading. Of the couplings of Section 7.1.3, couplings 2 and 3 are given up as in Chapter 18. The affine variant takes back a response to the load and the valve. The valve reading touches the offset and g, and neither reading touches β. Coupling 4 enters by limits only. The belief, the planner and the prediction model are Chapter 18's and are not re-derived. The operations book's look-ahead chapter asks the same question on rows with the valve fully open, and the loop changes its answer: Section 19.3.5 sets the two side by side on one belief.

19.1.2 The two-hall plant at its nodes (step 0)

The plant is Chapter 18's two halls, one water loop and a tank, in windows of 5 minutes. Ten racks draw air from two halls, k a rack and h its hall, with $\sum C = 30\,\mathrm{kW/K}$ of air capacity in each hall. An inlet reads the return less the share of the coil's drop that the rack's recirculation fraction β_k does not bring back,

$$T_k = T_{r,h} - (1-\beta_k)\,\frac{Q_{\mathrm{coil},h}}{\sum C}. \tag{19.1}$$

Every what-if is a scenario node and says whether the loop is open or closed. Writing $Q_t = Q_h + Q_o$ for the hall's heat with its other load Q_o, and G_h for its coil's conductance, the steady rows of each node are

$$\begin{aligned}
&\text{closed, off the stop:} && T_{r,h} = r_h - b_h, && T_k = r_h - b_h - (1-\beta_k)\frac{Q_t}{\sum C},\\
&\text{closed, at the stop:} && u_h = 1, && T_k = T_s + Q_t\Big(\frac{1}{G_h} - \frac{1-\beta_k}{\sum C}\Big),\\
&\text{open, valve pinned:} && u_h = u_{\mathrm{now}}, && T_k = T_s + Q_t\Big(\frac{1}{u_{\mathrm{now}}G_h} - \frac{1-\beta_k}{\sum C}\Big).
\end{aligned} \tag{19.2}$$

The closed node pins r_A, r_B, T_{sp} and the tank's move q_t and leaves the valves free, with the valve pinned again at a stop. This is Chapter 18's prediction model. The open node opens the hall loops and pins each valve where this draw's loop has it at now. The feedback node is the closed node with the supervisor's moves affine in what it read one window earlier,

$$r_{h,j} = \bar{r}_{h,j} + K_{r,h}\,\hat{e}_{h,j-1}, \qquad T_{sp,j} = \bar{T}_{sp,j} - K_u(y_{u,j-1} - \bar{u}_{j-1}), \qquad K_{r,h} = \frac{1-\hat{\beta}_{\mathrm{worst},h}}{\sum C}, \tag{19.3}$$

with the move limit kept, $\hat{e}$ the hall meter's deviation through a filter of gain 0.2744, and $K_u = 9\,\mathrm{K}$ per unit of valve, half the steady 18.1. At each copy T_{sp} reads the valve of the hall whose reference is the larger, and $\bar{u}$, the plan's own mean path, is held at its value at copy 12 after the plan's hour, a choice that sets the size of what this policy earns at a stance. Two further nodes appear: fully open, with the valves pinned at 1, the operations book's steady rows; and one unit, closed, with one chiller unit of two, 400 kW. Do-nothing is the closed node at $T_{sp} = 14°\mathrm{C}$, $r = 30.5°\mathrm{C}$, tank idle. It is a premise for a risk and not a recommendation, and the record's supervisor appears only as a labeled reference, never as a baseline for money.

Figure 19.1 draws the three main nodes for one hall. Each row of the figure runs from the water on the left through the coil and valve to the return air and one rack's inlet. Green is free, gray is pinned, and dashed arrows are readings. The red box in each row is the rack's fraction, with no reading row in any of them.

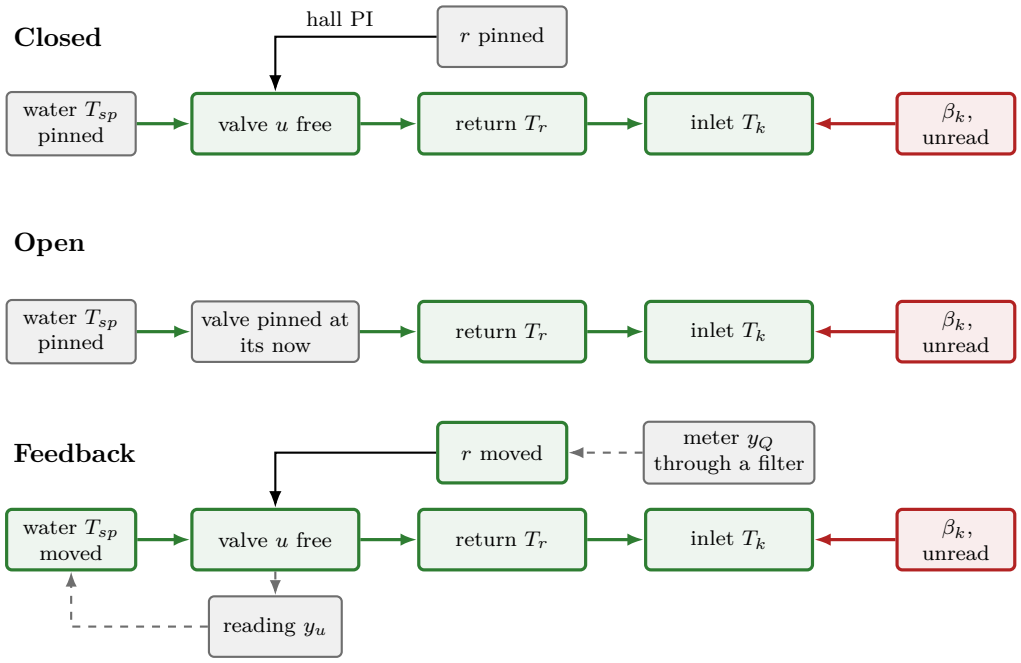


Figure 19.1: The closed, open and feedback nodes for one hall, equations (19.2) and (19.3). Closed: the hall PI moves the valve to hold the pinned r. Open: the valve is pinned at its value at now. Feedback: the closed node, with r moved on the filtered meter and T_{sp} on the valve reading, one window late. No node has a reading of β_k.

19.1.3 The control engineer's questions (step 1)

Table 19.1 lists the eight questions. The first column is the question in the engineer's words, the second the general form, the third its mathematical form on this plant, the fourth its category, the fifth the computation and the last the subsection. Here $T_{k,j}$ is rack k at copy j, and the limit on an inlet is 27°C.

19.1.4 The evidence as rows (step 2)

The rows on the past are Chapter 18's fit, conditioned. The rows on the future are averaged, and they are what this chapter adds to.

The forecast's rows. At the end of window now the load of hall h at lead s follows the published line, with a ramp entry R common to the halls and an autoregressive term ξ_h started at the fit's load less the line at now,

$$Q_h(s) = Q_h^0 + 25\,\frac{\text{now} + s - 1}{143} + (R - 25)\,\frac{s}{143} + \xi_{h,s}, \qquad \xi_{h,s} = 0.8\,\xi_{h,s-1} + \epsilon_{h,s}, \tag{19.4}$$

in kW, with $R \sim \mathcal{N}(25, 8^2)$, $Q_A^0 = 340$, $Q_B^0 = 285$ and ξ of stationary spread 3 kW. The mean runs along the whole published line, and the ramp's band is carried forward from now only, because the level at now is conditioned on the fit. Section 19.3.7 scores what that leaves out. Ambient and price have bands that widen with the root of the lead. The walk on g is Chapter 18's, 0.005 per window.

The rows of the ramp read from the meters. The trend ramp is its own small fit on $\boldsymbol{\theta} = (a_A,\ a_B,\ R)^{\mathsf{T}}$, the two intercepts and the ramp. With $x_t = t/143$ for $t = 0, \ldots, \text{now} - 1$

Table 19.1: The questions asked of the two-hall plant, about what a forecast's band does and what doing nothing risks.

Control engineer's question	General form	Mathematical form	Category	Computation	Answered in
What moves an inlet under the loop?	moments at scenario nodes, (2.10)	$\partial T_k/\partial(T_{sp}, r, Q, \beta, g)$ per node	design	posterior	19.3.1
Which entry makes the band, by lead?	moments, (2.10)	$\mathrm{Var}(\max_k T_{k,j} \mid \mathbf{y})$, groups on one at a time	estimation	posterior	19.3.2
Does the band level off under feedback?	moments at nodes, (2.10)	the band by lead, open, closed, affine	prediction under a threshold	posterior	19.3.3
What does each schedule cost?	expected consequence with limits, (2.23)	$z\,\mathrm{sd}_j$ per copy, joint 0.95; against perfect knowledge	design	expected consequence	19.3.4
What does doing nothing risk before the next decision?	threshold and trend, (2.12), (2.21)	$\Pr(\exists k, j \le s : T_{k,j} > 27 \mid \mathbf{y})$, with $\Pr(\gamma > 0,\ f_s \ge 27 \mid \mathbf{y})$ apart	prediction under a threshold	trend	19.3.5
What does waiting one window cost?	expected consequence, (2.23)	plan now against plan after one held window	design	expected consequence	19.3.6
Trend or published forecast, and when?	trend, (2.21)	$\hat{R}$ from readings against 25 kW; rmse, coverage; the valve's $\Pr(t^* \le s \mid \mathbf{y})$	prediction under a threshold	trend	19.3.7
Where does weather enter?	moments at nodes, (2.10)	ambient band at the rows; one unit; a capacity row	design	posterior	19.3.8

and $\mathbf{V} = 3^2\,[0.8^{|t-t'|}] + 4^2\mathbf{I}$ the covariance of one meter's window errors, whitened by $\mathbf{V}^{-1/2}$, the rows are

$$\mathbf{r}(\boldsymbol{\theta}) = \begin{bmatrix} \mathbf{V}^{-1/2}(a_A\mathbf{1} + R\,\mathbf{x} - \mathbf{y}_{Q_A}) \\ \mathbf{V}^{-1/2}(a_B\mathbf{1} + R\,\mathbf{x} - \mathbf{y}_{Q_B}) \end{bmatrix}, \qquad \mathbf{J} = \begin{bmatrix} \mathbf{V}^{-1/2}\mathbf{1} & \mathbf{0} & \mathbf{V}^{-1/2}\mathbf{x} \\ \mathbf{0} & \mathbf{V}^{-1/2}\mathbf{1} & \mathbf{V}^{-1/2}\mathbf{x} \end{bmatrix}. \tag{19.5}$$

The third column is shared, which is what makes R common to the halls. The information matrix of these normalized rows is $\mathbf{J}^\mathsf{T}\mathbf{J}$; with unnormalized rows it is $\mathbf{J}^\mathsf{T}\mathbf{W}\mathbf{J}$.

Which row acts on which entry. With $*$ for an entry that can be nonzero, the columns $(\beta_k, b_h, g, Q_h, R)$ and the rows of (19.2), the graph gives

$$\mathbf{J}_{\text{closed}} = \begin{bmatrix} * & * & 0 & * & * \end{bmatrix}, \qquad \mathbf{J}_{\text{stop}} = \begin{bmatrix} * & 0 & * & * & * \end{bmatrix}, \qquad \mathbf{J}_{\text{open}} = \begin{bmatrix} * & * & 0 & * & * \end{bmatrix}. \tag{19.6}$$

So β_k has a column at every node, and the offset trades places with g at the stop. On the open node the valve pinned at its now value carries the offset in and cancels the persistent g, leaving only its walk. Section 19.3.1 puts the numbers in.

The modeling choices. Five enter here.

- *Set points held at a stated value, or at a plan's last.* This is what lets the band be read by lead, and it fails as soon as a supervisor moves them. Chapter 20 takes up re-solving.
- *One affine form of feedback.* The gains (19.3) are one stated pair, neither tuned nor optimized. They buy a named policy to compare, and any statement about feedback here is about this form, except the floor, which is general to these readings.
- *The worst inlet as the largest of ten per draw.* This buys a budget of what the limit sees, and it mixes racks across draws.
- *One multiplier per schedule, target joint* 0.95. This keeps the comparison of shapes flat, and a bisection tolerance of 0.03 in z sets what can be resolved.
- *The trend from hall meters only, with a shared ramp.* This buys a fit with one trend entry, and it fails if the halls ramp differently.

19.2 The posterior of the state (step 3)

Two posteriors serve the answers: the belief Chapter 18 built, and the small fit of the ramp. No third fit is needed; the forecast's band stands beside the belief as stated rows.

19.2.1 The belief at now

The belief is Chapter 18's filter, read here at windows 12 and 72, with 23 entries on the reduced state. Table 19.2 gives the entries this chapter reads at each window: the two fractions by mean and spread, and the offset and the load by spread alone, the larger of the two halls. The two fractions stay at their priors for ever, because no row reads them, and the offsets' spreads fall by more than half.

An inlet's mean and variance at lead s on the closed node follow from (19.2), as a derived quantity with its gradient on $\mathbf{z}$ and the forecast's two added terms. On the columns

Table 19.2: The entries of the belief this chapter reads. Mean with $\sqrt{\Sigma_{ii}}$ in parentheses; the last two rows give the larger $\sqrt{\Sigma_{ii}}$ of the two halls.

Entry [unit]	Window 12	Window 72
β_{B3} [–]	0.5000 (0.0800)	0.5000 (0.0800)
β_{A5} [–]	0.4500 (0.0800)	0.4500 (0.0800)
b_h [K]	(0.0999)	(0.0428)
Q_h [kW]	(3.2939)	(3.2956)

$(\beta_k,\ b_h,\ Q_h,\ R)$, with $\phi_s = 0.8^s$,

$$f_s(\mathbf{z}) = r_h - b_h - (1-\beta_k)\frac{Q_t(s)}{\sum C}, \qquad \mathbf{g}_{f_s} = \Big(\frac{Q_t(s)}{\sum C},\ -1,\ -\frac{(1-\hat\beta_k)\phi_s}{\sum C},\ -\frac{(1-\hat\beta_k)\,s/143}{\sum C}\Big)^{\mathsf T},$$
$$\mathrm{Var}(f_s \mid \mathbf{y}) = \mathbf{g}_{f_s}^{\mathsf T}\,\boldsymbol{\Sigma}\,\mathbf{g}_{f_s} + \Big(\frac{1-\hat\beta_k}{\sum C}\Big)^2 3^2(1-\phi_s^2), \tag{19.7}$$

where $\boldsymbol{\Sigma}$ here is the belief's block with 8^2 added in the place of R. For B3 at window 12, lead 1, this gives mean 25.882°C and spread 0.7894 K.

19.2.2 The trend entry of an inlet and of a valve

A crossing needs a trend entry γ. The chapter's quantities have one on each node. Differencing (19.7) along the ramp, an inlet off the stop has, in K per hour, with 143 windows of 5 minutes the ramp's shift of 11.92 hours,

$$\gamma = -\frac{(1-\beta_k)R}{11.92\sum C}, \qquad \mathbf{g}_\gamma = \Big(\frac{\hat R}{11.92\sum C},\ 0,\ 0,\ -\frac{1-\hat\beta_k}{11.92\sum C}\Big)^{\mathsf T}. \tag{19.8}$$

The sign is the fact that matters. Since $1-\beta_k > 0$, a rising ramp makes the inlet fall. For B3 at the mean this is $\hat\gamma = -0.0350$ K/h, spread 0.0125. The valve's demand $u_h(s) = Q_t(s)/(G_h(r_h - b_h - T_s))$ has the opposite sign, with γ the demand's rise per hour, so a rising ramp takes it toward its level at 1. The same gradients read that pair, with the walk on g added to its variance.

19.2.3 The ramp from the meters

The rows (19.5) are linear, so one solve of the normal equations $(\mathbf{J}^{\mathsf T}\mathbf{J})\,\Delta\boldsymbol{\theta} = -\mathbf{J}^{\mathsf T}\mathbf{r}$ gives the fit. The ramp's entry of the covariance comes by Schur complement on the two intercepts. With $s_{11} = \mathbf{1}^{\mathsf T}\mathbf{V}^{-1}\mathbf{1}$, $s_{1x} = \mathbf{1}^{\mathsf T}\mathbf{V}^{-1}\mathbf{x}$ and $s_{xx} = \mathbf{x}^{\mathsf T}\mathbf{V}^{-1}\mathbf{x}$,

$$\Sigma_{RR}^{-1} = 2\Big(s_{xx} - \frac{s_{1x}^2}{s_{11}}\Big), \qquad \Sigma_{RR,\mathrm{both}}^{-1} = \Sigma_{RR}^{-1} + \frac{1}{8^2}, \tag{19.9}$$

the factor 2 being the two halls. "Both" is the normal posterior of the meters on the published forecast. On the book's record, the trend reads $\hat R = 36.17$ ($\sqrt{\Sigma_{RR}} = 5.09$), 32.51 (2.86) and 31.83 (2.46) kW per shift at windows 72, 108 and 120. Both reads 32.95 (4.29), 31.66 (2.70) and 31.24 (2.35). The stated spread of the trend alone is 51.80, 22.14, 13.10, 8.90, 5.09 and 2.15 kW per shift at windows 12, 24, 36, 48, 72 and 132: at window 12 a line through an hour of meters says nothing of a shift's ramp.

19.3 The answers (step 4)

Each answer reads the belief of Section 19.2.1 and the forecast's rows, by a named route. Draws of the belief run through the plant's rows at the node, with the state route of equation (19.7) beside them where one exists.

19.3.1 What moves an inlet

Control engineer's question. What moves an inlet under the loop?

General form. A derived quantity of the state and its gradient, the moments (2.10) read through the delta method (4.53), at each scenario node.

Specialized form. Equation (2.10) with $f = T_{B3}$ of equation (19.2). The choice that enters here is to read steady sensitivities at one copy, which buys one table of signs and fails in transit, where Chapter 18 measured the rows' error. Off the stop,

$$\frac{\partial T_k}{\partial Q} = -\frac{1-\beta_k}{\sum C}, \qquad \frac{\partial T_k}{\partial r} = 1, \qquad \frac{\partial T_k}{\partial T_{sp}} = 0, \qquad \frac{\partial T_k}{\partial \beta_k} = \frac{Q_t}{\sum C}, \tag{19.10}$$

and at the stop $\partial T_k/\partial Q = 1/G_h - (1-\beta_k)/\sum C$, with T_{sp} reaching the inlet one for one and r gone.

Building blocks. $\hat{\mathbf{z}}$ at window 72; the rows (19.2); a scenario node per row, read with B3 at copy 12.

The plant's numbers. Table 19.3 gives the sensitivities. Each row is a node; the columns give the base value of B3, valve B and the five sensitivities.

Table 19.3: What moves B3 at copy 12, window 72, entries at the fit's mean, noiseless, tank idle. $T_{sp} = 14$°C and $r = 30.5$°C unless a row says otherwise. On the open row $\partial/\partial Q_B$ is a load moved after now with the valve held; moving the level at now too re-pins the valve and gives -0.01647.

Node	B3 [°C]	Valve B [–]	$\partial/\partial T_{sp}$ [K/K]	$\partial/\partial r_B$ [K/K]	$\partial/\partial Q_B$ [K/kW]	Per 0.01 β_{B3} [K]	Per 0.01 $\log G_B$ [K]
Closed, off the stop	25.508	0.9204	0.0006	0.9999	−0.01667	0.1033	0.0000
Closed, at the stop, T_{sp} 17°C	27.173	1.0000	0.9970	0.0028	+0.03271	0.1032	−0.1519
Open, valve at its now value	25.539	0.9184	0.9959	0	+0.03697	0.1032	0
Fully open, the operations book's rows	24.177	1.0000	0.9975	0	+0.03286	0.1032	−0.1526

Off the stop, more load cools the worst inlet. The return is held, so more heat widens the drop across the coil: the closed form gives -0.01667 K/kW at the fit's $\beta_{B3} = 0.5000$, and the node gives the same. T_{sp} does not reach an inlet at all; it reaches the valve, at 0.0552 per kelvin on the steady row. At the stop the operations book's signs come back. On the open node a load moved after now warms the inlet, at $+0.03697$ K/kW against the third row's closed form of $+0.03726$, and that is the climb of the open node's mean in Section 19.3.3. The stop is the switch between two plants, and the column of β is the same in both: $Q_t/\sum C = 10.326$ K per unit, so about 0.103 K per 0.01. That column is what the rest of the chapter turns on.

19.3.2 The variance budget by lead

Control engineer's question. Which entry makes the band, by lead?

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

Specialized form. Equation (2.10) with $f = \max_k T_{k,j}$ through the closed node, with one group of the belief or the forecast on and the others at their means, on common draws. The choice is the grouping: β, the offsets b, g with its walk, the load, the weather (ambient and price) and the 30 s sensor. It buys a budget by name, and it does not add exactly, because the largest of ten is not linear.

Building blocks. Draws of the belief and of the rows (19.4); $n = 1000$, seed 1902; the valve's stop share at each lead.

The plant's numbers. Table 19.4 gives the budget at start 12 with the settings held. Figure 19.2 draws it, with lead on the horizontal axis and variance stacked by group on the vertical axis. The left panel has the settings held; the right has the calibrated plan and then its last set points held.

Table 19.4: Variance of the worst inlet by lead, start 12, closed, settings held, $n = 1000$ [K^2]. The share is load and weather over the summed groups.

Lead [–]	All	β	b	g	Load	Weather	Sensor	Share [–]
1 (5 min)	0.5008	0.4867	0.0102	0.0000	0.0026	0.0000	0.0044	0.005
12 (1 h)	0.5137	0.5031	0.0101	0.0001	0.0020	0.0000	0.0044	0.004
72 (6 h)	0.5498	0.5384	0.0101	0.0001	0.0066	0.0000	0.0043	0.012
132 (rest)	0.5869	0.5739	0.0101	0.0002	0.0171	0.0000	0.0046	0.028

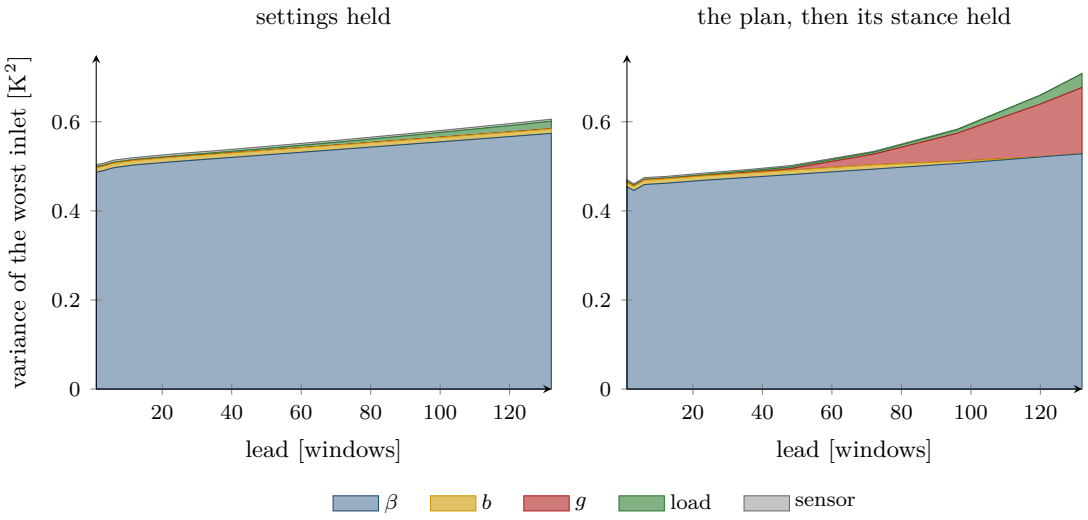


Figure 19.2: The worst inlet's variance stacked by group, start 12, closed, $n = 1000$, each group on with the others at their means. Weather is 0.0000 at every lead and is not drawn. Left: settings held. Right: the plan, then its last set points held, where g grows as the valve reaches its stop.

The sheet's picture is barely visible in the figure. **Load and weather reach** 1 % **of the worst inlet's variance at lead** 65 **(01:20),** 5.4 **hours out, and** 5 % **never within the shift**, the largest share being 0.0283. From window 72 the 1 % comes at lead 70, with largest share 0.0104. At the plan's stance held, it comes at lead 91 from start 12 and never from start 72, with largest share 0.0437 and 0.0078. Weather is 0.0000 at every lead.

What grows is β, from 0.4867 to 0.5739 K^2, and it grows neither by learning nor by forgetting. By (19.10) its gradient is $Q_t / \sum C$, and the mean load rises under a fixed unread entry: the growth is the persistent band scaled by the load. The entry that matters late is g, and only through the stop. At the plan's stance held from start 12, g goes from 0.0000 to 0.1486 K^2 at lead 132, with valves at the stop 0.518 and 0.647 against 0.341 and 0.011 with the settings held. This g band is Chapter 18's stated walk, so its size is that choice's.

In money, plans solved on the rows with tightening from all groups against persistent entries only, at the same z, differ by -0.093, -0.052 and -0.041 \$ for horizons of 12, 36 and 72 windows from start 12, and -0.001, $+0.002$ and $+0.006$ \$ from start 72. The largest B3 back-off moves from 1.831 to 1.827, from 1.832 to 1.834 and from 1.894 to 1.895 K at start 12. The signs mix, inside the planner's resolution, a local optimum at a tolerance of 10^{-3}. These are no resolved cost.

19.3.3 Open, closed and under feedback

Control engineer's question. Does the band level off under feedback?

General form. Moments at scenario nodes, equation (2.10), with the threshold form (2.12) for a violation.

Specialized form. Equation (2.10) with $f = T_{B3,j}$ at each node of (19.2) and (19.3), and equation (2.12) with the indicator of an intersection over racks and copies up to the lead. The state route is equation (19.7) for the closed node. On the open node the persistent g cancels against the valve pinned at its now value, so that route carries the load's move and the walk on g.

Building blocks. Draws of the belief, $n = 2000$, seed 1901; the gains of (19.3); the state route; the rows and the draws at the same settings.

The plant's numbers. Table 19.5 gives the band with the settings held from start 12. Each row is a lead; the columns give B3's spread on the three nodes, the state route on two and the violation by that lead on the three. Figure 19.3 draws B3's spread and mean by lead for all three nodes, with the limit and a dotted marker for the largest B3 of the record.

Table 19.5: The band by lead, start 12, settings held, $n = 2000$. Standard errors of the violation 0.006 to 0.011. The feedback's T_{sp} hit its move limit in 0.0227 of its moves.

Lead [–]	B3 sd [K]			State route [K]		Violation by this lead [–]		
	Open	Closed	Feedback	Open	Closed	Open	Closed	Feedback
1	0.7878	0.7881	0.7881	0.7933	0.7900	0.0695	0.0750	0.0750
12	0.8233	0.7999	0.7988	0.8369	0.8044	0.0980	0.0900	0.0910
36	0.8726	0.8140	0.8146	0.8848	0.8169	0.1815	0.0925	0.0930
72	0.9593	0.8332	0.8311	0.9550	0.8356	0.3010	0.0925	0.0930
132	1.0922	0.8654	0.8580	1.0756	0.8694	0.5050	0.0925	0.0930

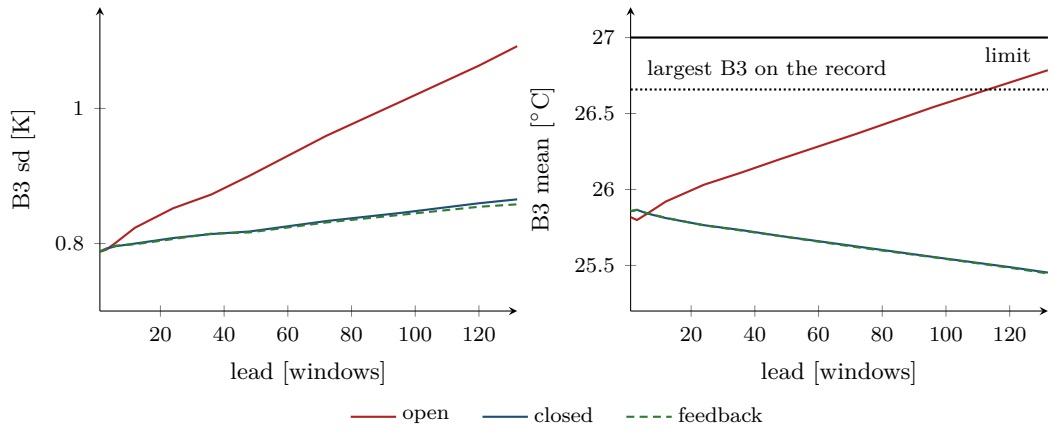


Figure 19.3: B3's band and mean by lead, start 12, settings held, $n = 2000$. Left: open grows +39 % over the shift, closed +10 % and feedback +9 %. Right: open climbs, closed and feedback fall. The dotted line is the stated truth's largest B3, 26.658°C, one draw of noise at one truth.

The sheet's contrast holds, and the feedback that delivers it is the hall loop, not a tube bought on top. Open loop, the band grows 39 % over the shift, and the mean climbs from 25.818 to 26.785°C, so the violation reaches 0.505. Closed, it grows 10 %, and the mean falls from 25.857 to 25.453°C. The hall PI, already there, is the affine policy.

A supervisory feedback adds almost nothing off the stop: 0.8580 against 0.8654 K, and 0.0930 against 0.0925 on the violation. The floor is β alone. With every other entry at its mean, the walk off and no sensor scatter, B3's spread at lead 132 is 0.8477 K closed and 0.8487 K under feedback, against 0.8654 closed with everything on. No affine policy on these readings lowers it, for neither reads β. The persistent entries together are not the floor, since they hold b and g as well and the valve reading does read those: at the plan's stance their spread falls from 0.9278 closed to 0.8857 K under feedback, and β alone stays at 0.8468 and 0.8491. The feedback does cost travel: T_{sp} has a spread of 0.15 to 0.23 K with the valve reading's noise among it. "Levels off" is not exact, since closed still creeps, by the load times β.

It does earn something late at the plan's stance, at the stop. From start 12, with the plan and then its last set points held, valve B at lead 132 is at its stop with probability 0.665 closed and 0.444 under feedback, the violation is 0.0960 (se 0.0066) and 0.0720 (se 0.0058), and T_{sp} settles at 14.968 from 15.3°C. From start 72, lead 72: 0.275 and 0.185; 0.0565 and 0.0545, not distinguished.

19.3.4 The schedules and their cost

Control engineer's question. What does each schedule cost?

General form. The expected consequence (2.23) under limits: the least expected cost whose plan holds the joint chance constraint.

Specialized form. Equation (2.23) with the action a plan solved on rows tightened by $z\,\mathrm{sd}_{k,j}$, where $\mathrm{sd}_{k,j}$ is one of five shapes: Chapter 18's state route, the open node's spread by draws, the closed node's, persistent entries only and the feedback's. The multiplier z is set by bisection so that the joint on 2000 draws is 0.95, and the plan is evaluated on 6000 others. The choice of one multiplier buys a comparison of shapes alone, and it fails to allocate by rack. Perfect knowledge of the same draw, $n = 60$, is a bound on cost only: its joint is not comparable, and it cannot be achieved.

Building blocks. The planner of Chapter 18; the five shapes; seeds 19, 1901 and 1903; the bisection's tolerance of 0.03 in z.

The plant's numbers. Table 19.6 gives the results. Each row is a schedule at a start and a horizon; the columns give z, the back-off at the first and last copy, the joint, net cost, the excess over the state route and the excess over perfect knowledge. Figure 19.4 draws B3's back-off against the copy for each schedule, at start 12 with 12 windows and start 72 with 36.

Table 19.6: Schedules, joint 0.95 on 2000 draws, evaluated on 6000. Net is electricity less the value of what is stored. Excess over perfect knowledge on 60 draws; dashes where not computed.

Start	H	Schedule	z [–]	B3 first, last [K]	Joint (se) [–]	Net [$]	Over state route [$]	Over perfect (se) [$]
12	12	state route	2.125	1.679, 1.709	0.9433 (0.0030)	9.361	0	+1.108 (0.074)
12	12	open loop's growing	2.234	1.742, 1.865	0.9438 (0.0030)	9.487	+0.125	+1.234 (0.075)
12	12	closed, by draws	2.172	1.779, 1.759	0.9438 (0.0030)	9.408	+0.046	+1.155 (0.075)
12	12	persistent only	2.172	1.778, 1.759	0.9438 (0.0030)	9.408	+0.046	+1.154 (0.075)
12	12	feedback's, feedback run	2.172	1.779, 1.762	0.9453 (0.0029)	9.417	+0.055	+1.160 (0.074)
72	12	state route	1.922	1.598, 1.593	0.9558 (0.0027)	−0.358	0	+0.537 (0.038)
72	12	open loop's growing	1.906	1.560, 1.637	0.9557 (0.0027)	−0.353	+0.005	+0.543 (0.038)
72	12	feedback's, feedback run	1.906	1.591, 1.593	0.9538 (0.0027)	−0.359	−0.001	+0.537 (0.038)
72	36	state route	1.938	1.611, 1.627	0.9545 (0.0027)	11.964	0	–
72	36	open loop's growing	1.906	1.570, 1.790	0.9548 (0.0027)	12.027	+0.063	–
72	36	closed, by draws	1.922	1.613, 1.681	0.9550 (0.0027)	11.974	+0.010	–
72	36	feedback's, feedback run	1.938	1.626, 1.700	0.9540 (0.0027)	11.891	−0.073	–

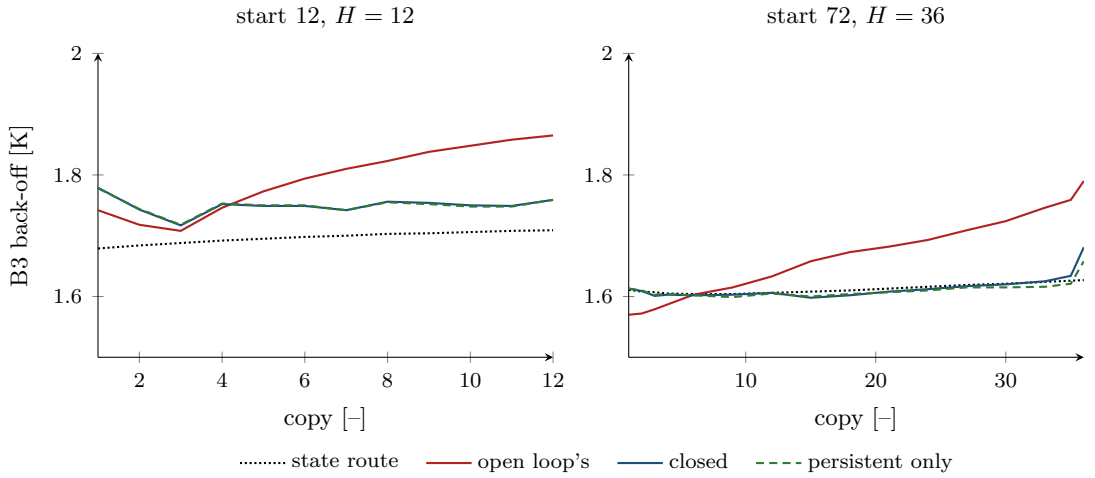


Figure 19.4: B3's back-off by copy, each schedule's z times its shape, Table 19.6. The shapes by draws step up at the last copy on the right. The open loop's shape grows; persistent entries alone trace the closed shape, so the forecast's bands add nothing measurable.

The right back-off is flat from the first copy: for B3 about 1.7 K in the evening and about 1.6 K overnight, with A5's last copy at 1.836 to 2.034 K on the closed shape. The persistent entries set it and the forecast's bands add nothing measurable: persistent only equals closed by draws to 0.001 $ at $H = 12$ and 0.003 $ at $H = 36$. The step at the last copy of the right panel coincides with the stance's own last move: the tank goes idle, T_{sp} falls from 14.203 to 13.703°C, B3's spread goes from 0.8500 to 0.8747 K and valve B's stop share from 0.099 to 0.253. Drawing the tube the open-loop way costs up to 0.125 $ per hour for nothing, 0.005 overnight and 0.063 over three hours.

The gap to the bound is Chapter 18's, and no schedule closes it: 1.11 to 1.23 $ in the evening and 0.54 $ overnight. Differences of about 0.03 in z, about 1 $ per unit of z in the evening, are not resolved, so $+0.046$ and -0.073 are read as about nothing, and only the open-loop row at start 12 stands clear. The joints sit 2.2 se under to 2.1 se over 0.95, out of sample. The z here, from 2000 draws, differs from Chapter 18's 2.172 and 1.906, which the script reproduces on that chapter's route.

19.3.5 Doing nothing

Control engineer's question. What does doing nothing risk before the next decision?

General form. Two forms answer it together. The threshold form (2.12), with the indicator of a union, gives the violation by a lead. The trend form (2.21) splits a threshold probability into the crossing and the mass that never crossed.

Specialized form. Equation (2.12) on the closed node, with f the largest inlet over racks and over copies up to s and $f^* = 27°\mathrm{C}$, gives

$$V(s) = \Pr(\exists\, k,\ j \le s:\ T_{k,j} > 27 \mid \mathbf{y}) = \int \mathbf{1}\Big[\max_{k,\, j \le s} T_{k,j}(\mathbf{z}) > 27\Big]\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}. \tag{19.11}$$

This is the probability of a violation by each lead, and it is not a crossing, because it counts a rack already over at now. The second line of equation (2.21), with f_s of equation (19.7), γ of equation (19.8) and $f^* = 27°\mathrm{C}$, gives the crossing, read under the Laplace posterior by equation (2.22) as

$$\begin{aligned}
\Pr(t^* \le s \mid \mathbf{y}) = \Pr(\gamma > 0,\ f_s \ge 27 \mid \mathbf{y}) &\approx \Phi_2\left(\frac{\hat\gamma}{\sqrt{\mathbf{g}_\gamma^{\mathsf{T}} \boldsymbol{\Sigma}\, \mathbf{g}_\gamma}},\ \frac{\hat f_s - 27}{\sqrt{\mathrm{Var}(f_s \mid \mathbf{y})}};\ \varrho(s)\right), \\
\varrho(s) &= \frac{\mathbf{g}_\gamma^{\mathsf{T}} \boldsymbol{\Sigma}\, \mathbf{g}_{f_s}}{\sqrt{\mathbf{g}_\gamma^{\mathsf{T}} \boldsymbol{\Sigma}\, \mathbf{g}_\gamma \ \mathrm{Var}(f_s \mid \mathbf{y})}}, \\
\Pr(\gamma \le 0,\ f_s \ge 27 \mid \mathbf{y}) &= \Phi\left(\frac{\hat f_s - 27}{\sqrt{\mathrm{Var}(f_s \mid \mathbf{y})}}\right) - \Pr(t^* \le s \mid \mathbf{y}).
\end{aligned} \tag{19.12}$$

The pair is on the holding regime, off the stop. A rising path at or above the level at s has crossed by s, before or after now. The choice to linearize γ, a product of two entries, buys the Φ_2 form. It carries β_k into a sign that β_k cannot change, so it makes the ceiling $\Pr(\gamma > 0 \mid \mathbf{y})$ and the crossing itself too large. Since $1 - \beta_k > 0$, $\gamma > 0$ exactly when $R < 0$, and the crossing with γ's true sign is $\Pr(R < 0,\ f_s \ge 27 \mid \mathbf{y})$, which the table sets beside it, read by 4×10^6 draws with R conditioned on $R < 0$.

Building blocks. $\hat{\mathbf{z}}$ and $\boldsymbol{\Sigma}$ of Section 19.2.1; $\mathbf{g}_{f_s}$ and $\mathbf{g}_\gamma$ of equations (19.7) and (19.8); draws through the plant's rows, $n = 4000$, seed 1904; each draw's ramp, rebuilt from the same stream with a largest load gap of 0.0 kW; the path's least-squares slope, for the crossing along draws.

The plant's numbers. Table 19.7 gives the violation at three windows. Each row is a window and an interval to the next decision; the columns give the violation, the largest single rack-and-copy probability, the largest single rack, the violation with the entries at their means, with the load's paths at their means and the sensor on, and with the sensor off as well, and the valves at their stops at the end. Other windows, over the remainder: 36 0.0645, 60 0.0530, 96 0.0443, 120 0.0437, 132 0.0495.

Table 19.7: The violation by the next decision, closed node, settings held, tank idle, $n = 4000$. The largest rack and copy is B3 in every row.

Window	Clock	Interval	Violation (se) [–]	Largest copy [–]	Largest rack [–]	Entries at means [–]	Load mean, sensor on [–]	Load mean, sensor off [–]	Stop at end, A, B [–]
12	20:00	5 min	0.0695 (0.0040)	0.0683	0.0683	0.0000	0.0703	0.0687	0.000, 0.000
		1 h	0.0875 (0.0045)	0.0707	0.0853	0.0000	0.0845	0.0723	0.000, 0.000
		remainder	0.0915 (0.0046)	0.0707	0.0882	0.0000	0.0863	0.0730	0.326, 0.011
72	01:00	5 min	0.0360 (0.0029)	0.0352	0.0352	0.0000	0.0355	0.0343	0.002, 0.000
		1 h	0.0430 (0.0032)	0.0357	0.0423	0.0000	0.0423	0.0355	0.017, 0.000
		remainder	0.0473 (0.0034)	0.0357	0.0460	0.0000	0.0447	0.0355	0.453, 0.048
108	04:00	5 min	0.0317 (0.0028)	0.0310	0.0310	0.0000	0.0317	0.0310	0.002, 0.066
		1 h	0.0397 (0.0031)	0.0340	0.0390	0.0000	0.0393	0.0315	0.062, 0.078
		remainder	0.0450 (0.0033)	0.0350	0.0437	0.0000	0.0443	0.0370	0.293, 0.309

Table 19.8 splits the threshold probability of B3. Each row is a window and a lead; the first columns give equation (19.12)'s pair, its three probabilities and the crossing with γ's true sign, and the last columns give the draws through the plant's rows, any rack: the violation by that lead, the mass over at now, the part from below, and the crossing along paths. Figure 19.5 draws the violation by lead, the largest row's probability at each lead and the crossing, one panel per window. The horizontal axis is the lead and the vertical axis a probability. In the middle panel the operations book's rows appear dotted, on the same belief.

Table 19.8: The threshold probability split, closed node, settings held. State route: B3, equation (19.12), Laplace. Draws: any rack, $n = 4000$. The crossing along draws is a rack over at s with a positive slope along its whole path.

			State route, B3				Draws, any rack			
Window	Lead	$\varrho(s)$ [–]	Line [–]	Mass [–]	Crossing, Φ_2 [–]	Crossing, true sign [–]	Violation by s [–]	Over at now [–]	Part from below [–]	Crossing, path [–]
12	1	+0.4440	0.0783	0.0771	1.25×10^{-3}	6.99×10^{-5}	0.0695	0.0757	0.0018	0.0000
	12	+0.4553	0.0664	0.0652	1.19×10^{-3}	6.33×10^{-5}	0.0875	0.0757	0.0120	0.0000
	132	+0.5656	0.0360	0.0347	1.23×10^{-3}	6.63×10^{-5}	0.0915	0.0757	0.0160	0.0010
72	1	+0.4468	0.0347	0.0339	8.18×10^{-4}	3.09×10^{-5}	0.0360	0.0333	0.0035	0.0015
	12	+0.4578	0.0358	0.0349	8.65×10^{-4}	3.41×10^{-5}	0.0430	0.0333	0.0097	0.0018
	72	+0.5145	0.0263	0.0254	8.94×10^{-4}	3.54×10^{-5}	0.0473	0.0333	0.0140	0.0020
108	1	+0.4470	0.0280	0.0273	7.25×10^{-4}	2.50×10^{-5}	0.0317	0.0173	0.0145	0.0152
	12	+0.4578	0.0316	0.0308	8.07×10^{-4}	3.02×10^{-5}	0.0397	0.0173	0.0225	0.0165
	36	+0.4808	0.0283	0.0274	8.24×10^{-4}	3.10×10^{-5}	0.0450	0.0173	0.0278	0.0200

Most of the risk is there at once. In the next 5 minutes the violation is 0.0695 of the shift's 0.0915 at 20:00, and 0.0360 of 0.0473 at 01:00. Doing nothing does not build this risk: it is the chance that the rack is over now, on an entry no reading touches. With the entries at their means the violation is 0.0000, every bit of it the belief. With the load's paths at their means and the sensor on it is 0.0863 against 0.0915, and with the sensor off as well 0.0730. By lead 1, 12, 36, 132 at 20:00 it runs 0.0695, 0.0875, 0.0897, 0.0915; the first hour adds 0.0180 and the conditional probability over the last hour is 0.0003.

These are not crossings. At 20:00 the entries' own rows put 0.0757 (se 0.0042) over the limit at now, which is not a crossing but a standing state, and the first window's valve and lag move that to 0.0695 at lead 1. The state route agrees: over a rising ramp B3 has $\hat{\gamma} = -0.0350\,\mathrm{K/h}$, so at lead 132 the threshold line is 0.0360, the mass over with no rising trend is 0.0347, and the crossing by Φ_2 is 1.23×10^{-3}. That number overstates. With γ's true sign the crossing is 6.63×10^{-5}, about 19 times smaller, 25 times at window 72 and 27 at window

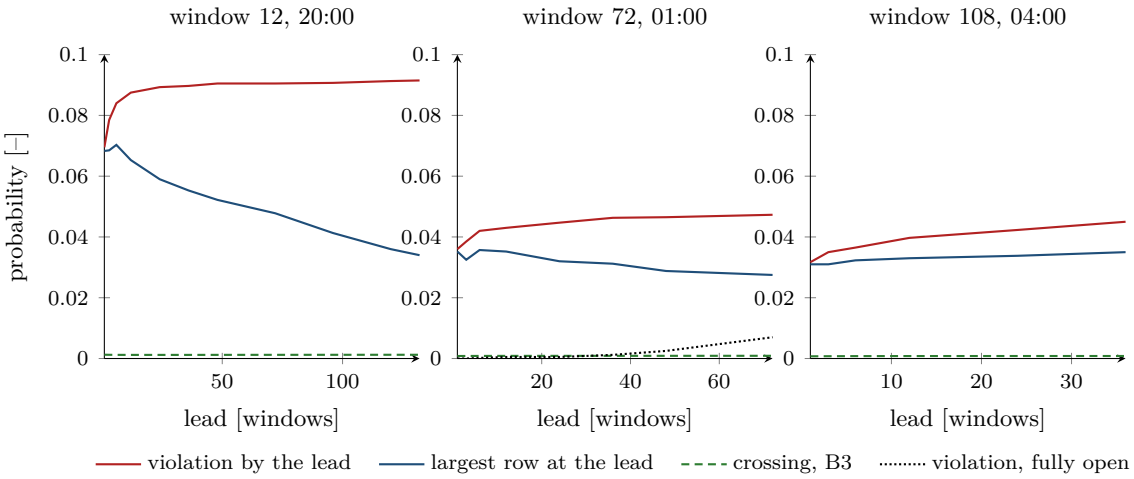


Figure 19.5: Doing nothing, closed node, settings held, $n = 4000$. The violation by a lead rises and flattens; the largest row's probability falls with lead at the first two windows; the crossing of B3 by Φ_2, equation (19.12), stays near zero, and it overstates the true sign's by about twenty times. Dotted: valves fully open, window 72, the same belief, on the operations book's rows.

108, under a ceiling of $\Pr(R < 0) = 8.89 \times 10^{-4}$ and not the 2.59×10^{-3} of the linearized pair. The positive $\varrho(s)$ is the same linearization at short lead: with β_k out of $\mathbf{g}_\gamma$ it is +0.0012 at lead 1 and +0.1416 at lead 132, where the ramp's own term is real. The table keeps Φ_2 because it is the form of equation (2.22), and it is the true sign's value that an inlet's crossing should be read at. Along the draws the trend entry from the draw's own ramp and fraction gives a crossing of 0.0000 on 4000 draws at every lead, for $\Pr(R \le 0)$ is 0.0009 by the published band (0.0013 on the draws).

So on the ramp's term an inlet on this node falls, and that is all γ says, for γ differences the ramp and leaves out the decay of ξ from the fit's load. That decay does move the mean upward over the first hour: the state route rises from 25.365 to 25.428°C over 12 leads at window 108, and from 25.490 to 25.511°C over 6 at window 72. Paths do rise and cross: the crossing along paths is 0.0200 of a line of 0.0360 at 04:00.

The rest is the part from below, 0.0160 at 20:00 and 0.0140 at 01:00, draws under at now and over later. It is mostly the return sensor's scatter acting through the loop, sampled over many copies, which inflates a maximum. At 20:00 the part from below is 0.0110 with the sensor on and the load at its means, 0.0015 with the sensor off and the load on, and 0.0008 with both off; at 01:00, 0.0123, 0.0055 and 0.0030. Of those draws at 20:00 only 0.0625 have the rack rising along its path.

The split at 04:00 is one window's start and not a trend. Table 19.9 prints the split at the windows either side, with one row per window; the columns give the start's return in B3's hall under $r - b$, the mass over at now, the part from below, the violation over the remainder and the share of B3's paths rising over the rest of the shift. The violation holds at 0.0440 to 0.0480 throughout. The split does not: over at now drops to 0.0140 at 04:00 from 0.0290 at 03:55, and it is back to 0.0340 by 04:20. From window 108 on, in most windows, the fit's valve in B3's hall stands at or above Chapter 18's 0.98, so the start's return comes from the filter and not from the loop's row, under $r - b$ at 04:00 and over it by 04:20. Putting the start on the loop's row removes the swing, with the violation unmoved: over at now is then 0.0250 at 04:00, 0.0260 at 04:20, 0.0255 at 04:40 and 0.0235 at 05:00, and the part from below 0.0225, 0.0195, 0.0185 and 0.0210. So the split from 04:00 on is not to be read at one window as printed, in either direction, and the row at window 108 of Table 19.8 is one such window. Read on one start, "already there" weakens slowly and not at 04:00: the part from below goes

from 0.0140 of 0.0473 at 01:00 to about half of the violation at 05:00. The share of B3's paths that rise grows too, 0.179 at 03:00 and 0.716 at 04:40, with valve B at its stop in 0.309 of draws by the end from 04:00.

Table 19.9: The split at neighboring windows, closed node, settings held, $n = 2000$, seed 1980. The return's gap is the start's return in B3's hall under $r - b$, negative when over it, and is zero where the fit's valve there is under 0.98.

Window	Clock	Return's gap [K]	Over at now [–]	From below [–]	Violation [–]	B3 rising [–]
96	03:00	0.000	0.0305	0.0170	0.0475	0.179
100	03:20	0.000	0.0315	0.0145	0.0460	0.194
104	03:40	0.000	0.0320	0.0155	0.0475	0.249
107	03:55	0.000	0.0290	0.0170	0.0460	0.376
108	04:00	+0.217	0.0140	0.0330	0.0470	0.540
109	04:05	+0.147	0.0190	0.0275	0.0465	0.536
110	04:10	+0.099	0.0180	0.0300	0.0480	0.478
112	04:20	−0.084	0.0340	0.0115	0.0455	0.603
116	04:40	−0.108	0.0355	0.0085	0.0440	0.716
120	05:00	−0.165	0.0415	0.0040	0.0445	0.675

The row's probability falls with lead, 0.0683 at lead 1 and 0.0340 at 132, as B3's mean falls from 25.832 to 25.425°C with the load. In the operations book's picture it rises to the horizon. The ratio of the violation over the remainder to the largest single copy is 1.29 at 20:00 and 1.32 at 01:00.

The stop is not where this risk lives in the evening and the night. At 20:00, over the remainder, the violation with a valve at its stop is 0.0937 (1922 draws) and with none 0.0895 (2078); at 01:00, 0.0443 (2890) and 0.0550 (1110). Conditioning on either valve dilutes that test, so remove the stop at equal loop gain, conductance times three with K_p and K_i over three, on $n = 3000$: the violation goes from 0.0957 to 0.0947 at 20:00 and from 0.0500 to 0.0493 at 01:00. At 04:00 it goes from 0.0493 to 0.0397, se about 0.004, so by morning a part of it does live there. Of the part from below, 0.5938 at 20:00, 0.7500 at 01:00 and 0.8468 at 04:00 have a valve at its stop by the end, against 0.4805, 0.7225 and 0.8345 of all draws. Stop shares in the last hours swing with the belief's load at now: at window 120, valve B stands at 0.659 at 5 minutes and 0.136 at 1 hour, a limit on reading them.

What the loop changes against the operations book. On the same belief at window 72, fully open gives a violation over the remainder of 0.0070 (se 0.0013), with B3's mean rising from 24.149 to 24.491°C and its spread from 0.8415 to 0.9537 K. Closed gives 0.0473 (se 0.0034), the mean falling from 25.494 to 25.339°C, and the spread going from 0.8371 to 0.8652 K. The loop closes the valve down to hold the return, so the air is warmer and the tail larger now, and the rising band and rising risk go away. That book's own look-ahead figure is on its own rows and record, and it is not like with like.

19.3.6 The cost of waiting one window

Control engineer's question. What does waiting one window cost?

General form. The expected consequence (2.23) of two actions.

Specialized form. Equation (2.23) with the two actions: plan now, or hold the settings for one window and then plan from the mean state after it. Each is calibrated to the joint 0.95 on 2000 draws, with z capped at 4, and evaluated on 6000 with the same draws. The choice to test one window buys a clean comparison and says nothing about a chain of waits.

Building blocks. The belief at each window; the planner; seeds 1950 and 1951; the difference paired by draw.

The plant's numbers. Table 19.10 gives the comparison. Each row is a window; the columns give z on each route, the joints, the waited route's joint at the z of acting now, the paired cost at that z, and the violation in the first window on each route.

Table 19.10: Plan now against one held window and then the plan, $H = 12$. A z of 4.000 with a joint under 0.95 means the search's cap was reached. Standard errors of the joints about 0.003.

Window	Clock	z now, wait [–]	Joint now, wait [–]	Wait at z now, joint [–]	Wait less now at z now (se) [$]	First window over, now [–]	First window over, wait [–]
12	20:00	2.188, 4.000	0.9530, 0.9242	0.9078	+0.3747 (0.0007)	0.0290	0.0758
36	22:00	2.141, 4.000	0.9517, 0.9440	0.9282	+0.2489 (0.0004)	0.0288	0.0560
72	01:00	1.891, 1.875	0.9475, 0.9460	0.9480	+0.0143 (0.0000)	0.0448	0.0385
108	04:00	1.734, 1.719	0.9472, 0.9463	0.9477	+0.0263 (0.0002)	0.0357	0.0342
126	05:30	2.047, 2.109	0.9545, 0.9500	0.9477	−0.0300 (0.0001)	0.0282	0.0355

In the evening one window spends the joint. The waited window alone is over at 0.0758 against the 0.05 the joint allows, so 0.95 is out of reach at any z. The 4.000 is the search's cap and the +2.08 $ that goes with it is that cap's price, not a meaningful number, so the column to read is at z now: 0.375 $ and 0.045 of joint. At 22:00 the search is also capped, at a joint of 0.9440, 2.0 se under, which is not clearly out of reach. **Overnight, waiting costs 1 to 3 cents and no joint.** At 01:00 the plan's first move is itself warmer than holding, 0.0448 against 0.0385. At the 05:30 edge the sign is slightly in waiting's favor at this one tank value, and z's tolerance does not separate it. The difference is not the forecast, for at 20:00 it is +0.3744 $ with load and weather at their means and +0.3747 $ with the entries frozen. Re-solving both plans with the tank idle leaves +0.1418 $, so 0.2329 $ of it is the tank's window forfeited at the 0.22 $/kWh tariff, the waited route discharging 183.3 kWh against 200.0, and the rest is the window's held set points. The joint's 0.045 does not depend on the tank, and that part is the unread entry against the hold. At 01:00 neither plan moves the tank and the cost is unchanged, +0.0143 $.

A plan's stance left standing. The same question arises after a plan ends. Table 19.11 holds the calibrated plan's last set points with the tank idle, against the record's settings on the same windows. Each row is a start; the columns give the stance, the violation over the next hour and the remainder after the plan's hour, the part from below, and the stops at the end, first for the stance and then for the record's settings.

Table 19.11: After the plan's hour, $n = 3000$, seed 1950. A violation over the windows after the plan; the part from below is under through the plan's hour and over after it. Standard errors 0.003 to 0.005.

		The stance				The record's settings			
Start	T_{sp}, r_A, r_B [°C]	1 h	Rest	From below	Stop A, B	1 h	Rest	From below	Stop A, B
12	15.269, 31.457, 29.948	0.0307	0.0717	0.0337	0.510, 0.648	0.0827	0.0857	0.0027	0.318, 0.009
72	14.446, 31.262, 30.426	0.0413	0.0463	0.0047	0.212, 0.263	0.0427	0.0457	0.0023	0.440, 0.049
108	13.629, 30.575, 30.605	0.0483	0.0500	0.0037	0.033, 0.036	0.0423	0.0467	0.0093	0.283, 0.295

A plan left standing is cooler on B3 and nearer the stop. At 20:00's stance the risk is 0.0307 over the next hour and then more than doubles by morning, and this is the one place in the chapter where the risk really does come from below: 0.0337 of draws are under through the plan's hour and over after it, and of those 0.990 have the rack rising and all have a valve at its stop. On the record's settings the same part is 0.0027.

19.3.7 Trend or published forecast

Control engineer's question. Trend or published forecast, and when?

General form. The trend (2.21): the sign of a trend entry, the crossing time of a level and the forecast.

Specialized form. Equation (2.21) with the trend entry the ramp R, read at $\hat{R}$ of Section 19.2.3, $\hat{R}$ with the prior as both, or the published 25 kW per shift. The third line gives a load forecast at lead 12 and to the shift's end. The second line, on the valve's demand of Section 19.2.2 with $f^* = 1$, gives the valve's crossing by equation (2.22): Φ_2 with $\varrho(s)$, and the never-crossing mass beside it. The modeling choice is that only the hall meters feed the trend, which buys a fit independent of the belief's other rows, and each route is scored on fresh records in three cases, one being the stated truth and two being constructed to test the claim either way.

Building blocks. $\hat{R}$ and Σ_{RR} of equation (19.9); the forecast's rows (19.4); $n = 200$ fresh records per case; for the valve, gradients as in Section 19.2.2.

The plant's numbers. The record's ramp is 30 kW per shift, the forecast's 25 with spread 8. Table 19.12 gives the window from which the trend's error falls below the published forecast's, by case and target; each row is a case. Table 19.13 gives the coverage of the 90 % band at the stated truth, by window. Figure 19.6 draws both, the window on the horizontal axis; the left panel has the ramp's rmse on a log scale for three routes, with the published ramp's error, and the right has the coverage at lead 12.

Table 19.12: The first window from which the trend beats the published forecast, with the rmse of the two beside it [kW per shift for the ramp, kW for the loads]. $n = 200$ records per case.

Ramp case	Ramp	Load, lead 12	Load, shift's end	Both
Stated truth, 30	84 (4.09 against 5.00)	96, 02:55 (4.37 against 4.64)	72, 00:55 (5.65 against 5.81)	beats published to the end from 12 (5.14)
Forecast right, drawn $\mathcal{N}(25, 8^2)$	60 (6.35 against 8.02)	60 (4.73 against 5.05)	48 (8.42 against 8.53)	to the end from 36 (7.51 against 8.53)
What-if 45	36 (12.98 against 20.00)	36 (5.17 against 7.01)	36 (12.01 against 20.12)	at lead 12 from 12 (4.25 against 4.40)

The trend needs about half the shift to say anything against a stated ramp of this band, by the stated spreads of equation (19.9). It wins sooner when the forecast's level error has had time to pile up. The published forecast's coverage decays with the window, from 0.912 to 0.575 at lead 12 at the stated truth, and from 0.887 to 0.510 with the forecast right, so a wrong ramp is only part of it. The band itself leaves out a term its own prior implies. By equation (19.4) it carries the ramp's band forward from now only, and a ramp off its mean also misplaces the line so far, by $(R - 25)(1 - 0.8^s)(\text{now} - 1)/143$. With R on the whole line the coverage holds at 0.902 to 0.895 with the forecast right, and goes from 0.927 to 1.000 at the

Table 19.13: Coverage of the 90 % band, $n = 200$, se about 0.02 to 0.03: at the stated truth, and in the last two columns with the ramp drawn from the forecast's own band.

	Ramp in band			Lead 12				To the end			Lead 12, forecast right
Window	Trend	Both	Pub.	Pub.	Trend	Both	Pub.	Trend	Both	Pub.	Pub., R on whole line
12	0.895	1.000	1.000	0.912	0.720	0.743	1.000	0.882	0.988	0.887	0.902
48	0.880	0.955	1.000	0.848	0.698	0.755	0.950	0.807	0.895	0.743	0.902
72	0.910	0.905	1.000	0.818	0.718	0.755	0.863	0.725	0.797	0.657	0.905
96	0.895	0.890	1.000	0.680	0.740	0.772	0.715	0.710	0.762	0.575	0.932
120	0.910	0.915	1.000	0.585	0.755	0.770	0.562	0.728	0.743	0.517	0.890
132	0.885	0.900	1.000	0.575	0.738	0.748	0.575	0.738	0.748	0.510	0.895

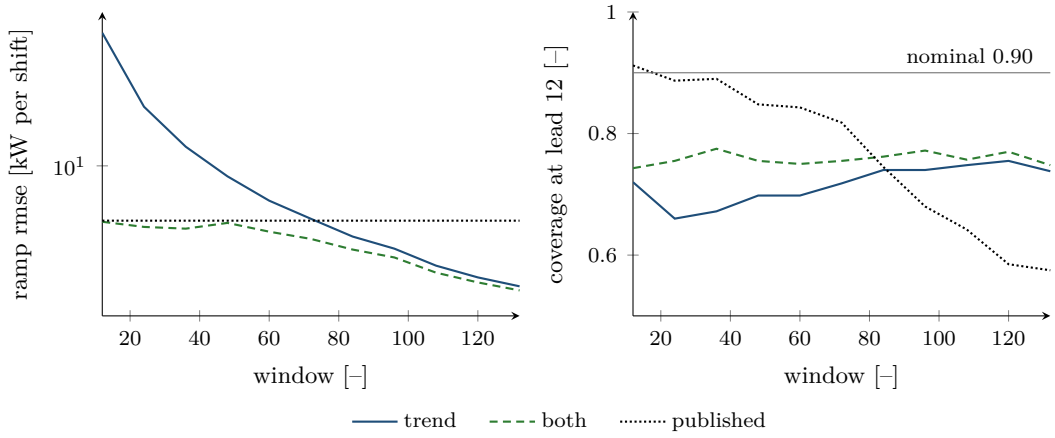


Figure 19.6: The three routes at the stated truth, $n = 200$ records. Left: the ramp's rmse; the dotted line is the published ramp's error, 5 kW per shift, against the stated truth of 30. Right: the coverage of the load's 90 % band at lead 12; the published forecast's decays with the window.

stated truth. Its ramp band's 1.000 is trivial: 5 lies inside 1.645×8 on every record at this truth.

Neither the trend's nor both's load band reaches 0.90 at any window: at lead 12 they hold 0.66 to 0.81 across the three cases, from window 12 on, and only the published route's shortfall is a late one. The trend's ramp meanwhile is calibrated, 0.88 to 0.93 at the stated truth and 0.87 to 0.94 with the ramp drawn. Part of the shortfall is the band's own: it carries the ramp's variance and leaves out the line's level. With the level in it the same forecasts cover 0.76 to 0.86 over the halls, still under nominal at every window and in every case. What remains is a level error of the load forecast, not of the slope. It is stated here as a finding, and its cause is not taken up.

The valve's crossing. Reading the ramp matters for when the valve runs out. Table 19.14 gives the valve's crossing on the book's record. Each row is a window, a route and a hall; the columns give the ramp, the crossing with its mass at the end of the shift, the correlation $\varrho(s)$ there, the draws' stop share beside it, and the crossing's median and early end of its band by inversion. The late end of the band lies beyond the shift in every row.

Table 19.14: The valve's crossing on the book's record, equation (2.22) with $f^* = 1$, read at the shift's end. Draws' stop share at that lead, $n = 4000$. Mass: $\Pr(\gamma \le 0,\ u_s \ge 1 \mid \mathbf{y})$.

Window	Route	$\hat{R}$ (sd)	Hall	Crossing [–]	Mass [–]	$\varrho(s)$ [–]	Draws' stop share [–]	Median; 0.025 point
12	published	25.00 (8.00)	A	0.3260	7.0×10^{-6}	+0.492	0.326	beyond; 01:00
			B	0.0081	1.8×10^{-10}	+0.550	0.011	beyond; beyond
72	published	25.00 (8.00)	A	0.4365	5.4×10^{-5}	+0.377	0.453	beyond; 02:20
			B	0.0418	2.2×10^{-7}	+0.426	0.048	beyond; 06:25
	trend	36.17 (5.09)	A	0.8132	1.0×10^{-13}	+0.252	0.830	03:50; 01:20
			B	0.2815	1.7×10^{-15}	+0.287	0.284	beyond; 04:10
108	published	25.00 (8.00)	A	0.2655	5.2×10^{-5}	+0.263	0.293	beyond; 04:40
			B	0.2946	4.8×10^{-5}	+0.295	0.309	beyond; 04:00
	trend	32.51 (2.86)	A	0.6309	8.2×10^{-31}	+0.116	0.658	06:00; 04:10
			B	0.7086	9.3×10^{-31}	+0.129	0.735	05:35; 04:00

The trend moves the valve's answer the right way, and it still misses at this truth. At window 72 the published route has valve A at 0.4365 by the shift's end with its median beyond the shift, and the trend, reading 36.17 kW per shift, has 0.8132 with its median at 03:50. That read stands 1.2 of its spreads above the record's 30. On the book's record valve A first reaches its stop at 02:55 and valve B at 04:25, while from window 72 the trend gives valve B only 0.2815 by the shift's end, with its median beyond the shift; from window 108 valve B's median is 05:35. The cause is not taken up here. The mass is negligible because a rising ramp is nearly certain, so on this quantity, unlike the inlet, the crossing and the threshold line agree. On the same record the do-nothing answer does not move on the inlets: over the remainder at window 72 the violation is 0.0450, 0.0440 and 0.0440 for published, trend and both, and at window 108 0.0443, 0.0560 and 0.0533, published and trend being 2.1 se apart, the largest gap. **It moves the stop**: valve B at the end stands at 0.045, 0.300 and 0.180 at window 72, and 0.303, 0.726 and 0.674 at window 108. The trend matters to the question of when the valve runs out, and not to whether a rack is over. Short leads at windows 108 and 120 swing with one window's load, so the line and the draws differ there, for instance 0.1585 against 0.066 for valve B at window 108, lead 1.

The forecast on fresh records. Table 19.15 scores the do-nothing violation on fresh records, each with its own fit and forecast, 60 draws per record, settings held. Each row is a truth, a

start and an interval; the columns give the mean predicted, the share observed and the top third by predicted.

Table 19.15: The do-nothing violation on fresh records. Truth drawn: β, b and constant g from the stated priors, the coverage. Stated truth: one point, not coverage.

Truth	n	Window	Interval	Predicted [–]	Observed (se) [–]	Top third: predicted, observed (se) [–]
drawn	150	12	1 h	0.0729	0.0533 (0.0183)	0.1354, 0.0794 (0.0341)
drawn	150	12	remainder	0.1117	0.0867 (0.0230)	0.2175, 0.1525 (0.0468)
drawn	150	72	1 h	0.0621	0.0467 (0.0172)	0.1242, 0.0968 (0.0375)
drawn	150	72	remainder	0.0857	0.0733 (0.0213)	0.1827, 0.1636 (0.0499)
drawn, both	150	72	remainder	0.0896	0.0733 (0.0213)	0.1952, 0.1607 (0.0491)
stated	100	12	remainder	0.0757	0.0000	one truth's outcome
stated	100	72	remainder	0.0543	0.0000	the same

With the truth drawn the shares stand 0.6 to 1.1 se under the prediction, on the conservative side and not distinguished from calibrated at this n. The valve at its stop in the last window, truth drawn, has outcomes 0.320 and 0.173 for A and B against predicted 0.266 and 0.175 at window 72. **At the stated truth the valve's stop share is one truth's outcome**: 0.770 and 0.980 against the belief's 0.279 and 0.036 at window 72, and 0.374 and 0.072 with the ramp read. Records at one truth differ by noise only, so this is not a frequency the belief's share should match and not a coverage statement. The ramp moves hall A's; the rest is this truth's. It does not reach an inlet: no record goes over 27°C, the largest being 26.786°C. The 0.0000 is likewise one truth's outcome.

19.3.8 Where weather enters

Control engineer's question. Where does weather enter?

General form. Moments at scenario nodes, equation (2.10).

Specialized form. Equation (2.10) with f first an inlet and then the hour's cost, with the ambient band on and off on common draws. In this plant ambient enters only the chiller's efficiency, and capacity is the number of units times 400 kW. A what-if row $C_0(1-0.015(T_{\text{amb}}-\bar{T}))$ tests what a capacity row would do. The plant does not have this row, and its slope is a stated number with no source in the record. It buys one answer to a claim that the plant cannot otherwise test.

Building blocks. Draws, $n = 3000$ and $n = 4000$, seed 1910; the one-unit node; the what-if row at a window 120 heat of 697.1 kW at the end.

The plant's numbers. Ambient reaches no inlet: the largest gap is 0.0 K**, on two units and on one.** It reaches cost. The hour's cost has a spread of 1.458 $ at start 12, with ambient 0.480, price 1.371, load 0.107 and entries 0.002; and at start 72, 0.605 $, with 0.234, 0.556, 0.028 and 0.001.

Table 19.16 gives the one-unit node and the what-if row. Each row of the upper part is a start and a set of moves, with the lead at which the violation at a copy first exceeds 0.5, and the violation at three intervals; each row of the lower part is a capacity C_0, with the violation with and without the ambient band.

The sheet's "and chiller capacity" is not in this plant. On one unit, with the tank idle, the violation passes one half by lead 3, 15 minutes, and T_s reaches 28.73°C at 1 hour in the evening.

Table 19.16: Above: one chiller unit of two, closed node, $n = 2000$. Below: the what-if capacity row, which the plant does not have, window 120 to the shift's end, $n = 4000$.

Start	Moves	Lead of $P > 0.5$	5 min [–]	1 h [–]	2 h [–]
12	settings held, tank idle	3	0.2055	1.0000	1.0000
12	settings held, tank at 200 kW	13	0.1070	0.4185	1.0000
12	T_{sp} 8°C, r 26°C, tank at 200 kW	13	0.0060	0.4030	1.0000
72	settings held, tank idle	3	0.1210	1.0000	1.0000
72	settings held, tank at 200 kW	8	0.0540	0.9920	1.0000
72	T_{sp} 8°C, r 26°C, tank at 200 kW	8	0.0165	0.9920	1.0000

C_0 [kW]	Ambient's band	At its mean	Difference	Share binding
800	0.0507	0.0507	+0.0000	0.000
720	0.0507	0.0507	+0.0000	0.003
700	0.0660	0.0507	+0.0152	0.250
690	0.1577	0.0990	+0.0587	0.697
680	0.4313	0.3990	+0.0323	0.941
670	0.7728	0.8273	−0.0545	0.990
660	0.9595	0.9965	−0.0370	0.999
640	1.0000	1.0000	+0.0000	1.000

With the tank at 200 kW that goes to lead 13 in the evening and lead 8 at 01:00. Moving T_{sp} to 8°C and r to 26°C does not change the lead: capacity, not a set point, sets it, though it does lower the first 5 minutes. On the what-if row, ambient matters only inside a margin around the heat, and either way: it moves the joint at 2.9 kW over the heat and at 37.1 kW under it, and not at 22.9 kW over or 57.1 kW under.

19.4 The checks (step 5)

Each check tests one link between the belief and an answer. Table 19.17 lists them; each row gives the check, what it printed and its target.

19.4.1 The checks against the rows and the limits

Reading the $\mathbf{\Sigma}$ rows from the top, the band falls with the scale at lead 1 and stops falling at lead 72. That floor is the stated walk on g, since with it off the spread is 3.6×10^{-15} K. This is the $\sigma_p \to 0$ limit of this chapter's rows: with no belief's band the node returns its noiseless path.

19.4.2 The state route beside the draws

Table 19.5 sets the state route beside the draws. Closed, the gap is at most 0.0052 K in a spread of about 0.8 K; open, at most 0.0184 K, with the sign changing along the shift, because the route's load move is steady. The route of equation (19.7) gives 0.7894 K at lead 1 against the same table's 0.7900, which leaves out the belief's covariance of the load with b and β. On the crossing the two routes differ in kind: the linearized γ carries β into its sign, where the draws do not, which overstates the crossing about twenty times, and Section 19.3.5 reports both.

Table 19.17: The checks. Standard errors in parentheses.

Check	Printed	Target
This chapter's rows against Chapter 18's, no options	inlet gap 0.0, net gap 0.0	10^{-12}
Feedback at zero gains against closed	gap 0.0	10^{-12}
The forecast reads no truth: stated offsets, fractions and log-conductances changed, readings fixed	inlet gap 0.0	0
$\mathbf{\Sigma} \times 10^{-2}$, paths at means, no sensor noise	B3 sd at leads 1, 72: 8.57×10^{-2}, 9.21×10^{-2} K; joint 0.0000	path's largest 25.511°C
$\mathbf{\Sigma} \times 10^{-4}$	8.57×10^{-3}, 3.08×10^{-2} K; joint 0.0000	the same
$\mathbf{\Sigma} \times 10^{-8}$	9.45×10^{-5}, 2.99×10^{-2} K; joint 0.0000	the walk on g
Entries frozen, walk off	sd at lead 72 3.6×10^{-15} K; gap to the noiseless path 0.0	0
Two seeds, $n = 6000$, start 72; start 12	0.0478 (0.0028), 0.0527 (0.0029); 0.0913 (0.0037), 0.0977 (0.0038)	each −1.2 se
First copy, window 72: linearized joint on $(\hat{\mathbf{z}}, \mathbf{\Sigma})$; plant's rows	0.0348 (0.0003); 0.0359 (0.0013)	agree within 1 se
The violation cannot fall with the horizon	0.0365, 0.0410, 0.0435, 0.0465, 0.0485, 0.0485	monotone
Each draw's ramp, rebuilt from the stream, three windows	largest load gap 0.0 kW	0
Ramp estimator on a noiseless 30 kW line, window 96	30.000000 kW per shift	30
Its spread with two windows, with the prior	7.999	the prior's 8
Synthetic records, window 96, $n = 1000$	stated 3.388, observed 3.294; mean error −0.097 (0.104)	agree
Forecast rows, lead 12 against the draws	Q_B sd 3.091 (3.075); ambient 0.789 (0.800); price log sd 0.1289 (0.1300)	agree
Chapter 18's route reproduced	z 2.172, 1.906; last T_{sp} 15.281, 14.434°C	its numbers

19.5 The report (step 6)

The load and weather bands tighten almost nothing here. The right back-off for B3 is about 1.7 K in the evening and 1.6 K overnight, flat, and it is two unread racks. Doing nothing shows a violation of 0.0695 in the next 5 minutes and 0.0915 by morning, at 20:00, and most of that is a state already over, not a crossing: with γ's true sign the inlet's own crossing is 6.63×10^{-5}, because on this node a rising ramp cools it. The rest comes from below and is mostly the return sensor's scatter. The violation holds near 0.045 from 03:00 on, and toward morning the part from below slowly reaches about half of it, with the share of B3's paths rising.

19.5.1 The answers and their bands

Table 19.18 gives the answers in the order of Table 19.1. Each row is a question; the columns give the answer and its band or basis. The last row is the stated truth.

Table 19.18: The report for the two-hall plant. Percentages and money are this example's.

Control engineer's question	Answer	Band or basis
What moves an inlet under the loop?	Off the stop: r, and Q at -0.01667 K/kW; T_{sp} not at all. At the stop: T_{sp}, and Q at $+0.03271$ K/kW	β's column 0.103 K per 0.01 in both
Which entry makes the band, by lead?	β, 0.4867 to 0.5739 K^2; load and weather at most 0.0283 of the variance; weather 0.0000	$n = 1000$; plan's stance held: g to 0.1486 K^2, a stated walk floor, β alone, 0.8477 K; se 0.0058, 0.0066
Does the band level off under feedback?	Open +39 %, closed +10 %, feedback adds almost nothing off the stop (0.8580 against 0.8654 K); at the stance, violation 0.0720 against 0.0960	
What does each schedule cost?	Flat, B3 about 1.7 K evening, 1.6 K overnight; the open loop's growing shape costs +0.125 $ per hour in the evening; persistent only equals closed to 0.001 $	joint 2.2 se under to 2.1 se over; bound's gap 1.11 to 1.23 $, 0.54 $
What does doing nothing risk before the next decision?	Violation at 20:00: 0.0695, 0.0875, 0.0915 by 5 min, 1 h, remainder; at 01:00: 0.0360, 0.0430, 0.0473; at 04:00: 0.0317, 0.0397, 0.0450. Over at now 0.0757, 0.0333, and 0.0173 at 04:00, one window's start. Crossing, B3, true sign, 6.63×10^{-5}, 3.54×10^{-5}, 3.10×10^{-5}	se 0.003 to 0.005; draws' crossing by the ramp's sign 0.0000; Φ_2 overstates about twenty times, 1.23×10^{-3}, 8.94×10^{-4}, 8.24×10^{-4}
What does waiting one window cost?	At z now: +0.375 $ and 0.045 of joint at 20:00; +0.014 $ at 01:00; −0.030 $ at 05:30	joint se 0.003; the +2.08 $ is the search's cap and is not a cost
Trend or published forecast, and when?	Trend beats published on the ramp from window 84 at the stated truth, and on load at lead 12 from window 96. It moves the valve's crossing the right way, and still misses at this truth; not the inlets	$n = 200$; trend's and both's load bands under nominal at every window, the published one late, by a term it leaves out; valve's stop share at the stated truth is one truth's outcome
Where does weather enter?	Cost only: spread 0.480 $ of 1.458 $. Ambient's inlet gap 0.0 K. One unit: over by 15 min, 40 to 65 min with the tank	what-if capacity row matters from about 37 kW under the heat to about 3 kW over it
Stated truth	ramp 30 kW per shift against 25; $\beta_{B3} = 0.57$; $\beta_{A5} = 0.51$; $b_A = +0.10$ K; $b_B = -0.25$ K; largest B3 on the record 26.658°C	one truth, one draw of noise

19.5.2 The entry in the commissioning log

The entry states a back-off flat by copy from the state route, the joint held by direct calibration, and no lead-growing schedule. It counts the hall loops as the feedback and proposes no supervisory feedback off the stop. It proposes a valve-reading feedback on T_{sp} only for a stance left standing for hours. It says that the ramp is to be re-read from the meters from about 01:00, and read for the valve's stop, not the inlets. It uses ambient for cost only. It records the

one-unit loss time as 15 minutes, 40 to 65 minutes with the tank. It marks the forecast's stop share as not trusted at this truth. It logs a do-nothing violation as a probability of being over and not as a time to a limit. And it proposes inlet readings at B3 and A5. Each of these rests on a constructed record and on stated gains.

19.5.3 Common mistakes

Table 19.19 lists six mistakes this plant makes easy. Each row names the mistake, its form here and the correction.

Table 19.19: Six mistakes with a forecast's band.

Mistake	On this plant	Correction
Drawing the tube open loop	+39 % against +10 %; +0.125 $ per hour	read the band on the node that runs: the hall loop is the feedback
Buying a forecast to shrink the back-off	load and weather at most 0.0283 of the variance; ramp band 2 kW saves 0.000 $	read the budget first; here the entry is unread, and elsewhere it may not be
Reading a violation as a crossing	0.0915 by morning, inlet crossing 6.63×10^{-5}	split it: the mass already over, the part from below, the crossing
Reading a flat trend as nothing to do	0.0757 over at now on a flat loop trend	the entry is not in the trend; read it from the belief
Reading 0.0000 on stated-truth records as coverage	predicted 0.0757, observed 0.0000	one truth's outcome; draw the truth
Calling +2.08 $ the cost of waiting	the search's cap at $z = 4$	read at z now: 0.375 $ and 0.045 of joint

19.6 The next measurement (step 7)

The candidates are the permanent instruments that read the entries in the budget, and the better forecast the sheet proposes. Each updates $\mathbf{\Sigma}$ by the preposterior (4.73) with its row, with the means held, which flatters, and with the joint re-calibrated.

19.6.1 The candidates

Table 19.20 ranks them. Each row is a candidate; the columns give the saving, the do-nothing violation to 3 hours and B3's spread at lead 12, each at starts 12 and 72.

Table 19.20: The candidates, joint 0.95 re-calibrated on 2000 draws, evaluated on 3000, seed 1960; do-nothing from the same belief.

Candidate	Saving, start 12, 72 [$]	Do-nothing to 3 h, start 12, 72 [–]	B3 sd lead 12, start 12, 72 [K]
None	0, 0	0.0887, 0.0460	0.8105, 0.8365
Permanent inlet readings, B3 and A5	+1.086, +0.510	0.0002, 0.0002	0.1470, 0.0831
Airflow re-survey, six racks, sd 0.03	+0.871, +0.391	0.0002, 0.0000	0.3046, 0.3013
Reference return reading, 0.05 K	−0.015, 0.000	0.0835, 0.0457	0.8067, 0.8362
Hall meters to 1 kW	+0.025, 0.000	0.0883, 0.0460	0.8103, 0.8366
Weather station	+0.015, +0.007	0.0887, 0.0460	0.8105, 0.8365
Better load forecast, ramp band 2 kW	−0.000, −0.000	0.0890, 0.0460	0.8106, 0.8366

A better forecast is worth nothing here, and a weather station no resolved value. On the same plan, with z the same 2.156 and 1.906, the ambient band against its mean moves

the net by +0.0153 $ (se 0.0088) and +0.0074 $ (se 0.0043). The permanent inlet readings save 1.086 and 0.510 $ per hour, and they take do-nothing's 0.0887 to 0.0002. That 0.0002 assumes the belief's mean stays where it is. The truth's hot rack would move it, as in Chapter 18, so it is not this truth's risk.

19.6.2 The recommendation

The permanent instrument is an inlet reading at each of B3 and A5. It is the one candidate that reaches the entry the budget names. A survey of the airflow comes near it without a permanent sensor, while the entry's value stays where the survey leaves it. Nothing about the forecast, the weather or the meters earns a place until these two entries are read, and the question of whether to learn them by a move belongs to Chapters 22 to 24.

19.6.3 What the analysis showed

The sheet said that the band widens along the horizon, so that the limit tightens with lead, and the flat trend said there was nothing to do. On this plant the first is true of open-loop rows and barely true of the plant: under the hall loop the band grows by a tenth, from the load times an unread entry, with forecasts a small share of it, and the levelling the tube promises is delivered by the loop already installed. Doing nothing's risk at 20:00 is present and not coming, mostly a state already over on an entry no reading touches, and on the ramp's term the rising load cools the inlet, so its crossing is about 10^{-4} of the shift's chance. The risk from below is mostly the return sensor's scatter in the evening and the night. Toward morning the violation stays near 0.045, the part from below slowly reaches about half of it, and the share of rising paths grows. What does cross is the valve. The trend ramp moves its answer the right way against the published ramp, and it still misses at this truth. A stance left standing is the one case where the risk builds with lead through the stop, and it is the only place a supervisory feedback has earned anything.

19.6.4 Limits of the model

This plant's rows. The results follow from the valve being off its stop for most of the shift, a capacity with no ambient row and a band dominated by a survey's 0.08. Change those and the budget changes. Forecast bands are small here, at this plant, and a tube's result does not carry over.

The belief. It is a Laplace normal with Chapter 18's flags unchanged: b_A not calibrated and rows quasi-steady. The walk on g is a stated form and sets the late band. The crossing's pair is linear in a product of entries and holds off the stop, and it overstates the inlet's crossing about twenty times. From window 108 on the start's return often comes from the filter, which moves the split between mass and part from below at each such window, not the violation.

The forecast. Its stop share at the stated truth is one truth's outcome and not a frequency. The trend's and both's load bands fall under nominal at every window and the published one late; part of the shortfall is the band's own, and a level error of the forecast remains, its cause not taken up. The published band conditions the level on the fit and carries the ramp's band from now only. Stop shares at short lead swing with one window's load.

The schedules. One multiplier per schedule, a bisection tolerance of 0.03 in z, 2000 draws to calibrate and a local optimum for the planner, so differences of about 0.03 to 0.1 $ are unresolved. Perfect knowledge rests on 60 draws.

The feedback. The gains are one stated pair, not tuned or optimized, so "adds almost nothing" is of this affine form. The floor's argument, no reading of β, is general to these readings.

The records. Fresh do-nothing records number 150 and 100, with 60 draws each: small, and at the stated truth one point. Do-nothing has no cold-side limit. Waiting was tested for one window, with no chain of waits.

The constructed record. The record is constructed from a stated truth, so its agreement tests the method on its own model and not the model against a real plant. Percentages and money are this example's.

What belongs elsewhere. The result of re-solving the plan is Chapter 20's and the store's value Chapter 21's. Learning from a move is Chapters 22 to 24's, and fault hypotheses Chapter 26's.

19.7 What to retain

Definition 19.1 (Violation, mass and crossing)**.** The violation by a lead is the joint tail (19.11), which counts a state already over. A threshold probability at a lead splits into the crossing, $\Pr(\gamma > 0,\ f_s \geq f^* \mid \mathbf{y})$, and the mass $\Pr(\gamma \leq 0,\ f_s \geq f^* \mid \mathbf{y})$, read as Φ_2 with $\varrho(s)$ by equation (19.12).

Three conclusions carry beyond the plant as statements of form. A band is read on the node that runs: where a loop holds a quantity, that loop is already the feedback a tube asks for, and its sensitivities are not those of the open rows. No policy lowers the band of an entry none of its readings touches, so a budget by entry decides between a better forecast and a sensor before any money is spent. And a quantity's sign of trend decides whether its probability of being over is a crossing at all. The sizes here are this plant's. The chapter's questions join the others in Appendix A.

Table 19.21 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 19.21: The terms of the chapter and where each is defined. Units for the two-hall plant.

Term	Expression	Units	Meaning	Defined in
Node's rows	eq. (19.2)	°C	inlet closed, at the stop, open	19.1.2
Affine policy	eq. (19.3)	°C	r on the meter, T_{sp} on the valve	19.1.2
Trend ramp	rows (19.5), spread (19.9)	kW per shift	ramp from the meters so far	19.2.3
Inlet's trend entry	eq. (19.8)	K/h	negative on a rising ramp	19.2.2
Tube	the band under feedback	K	the hall loop is one	19.3.3
Tightening schedule	$z\,\mathrm{sd}_{k,j}$	K	back-off by copy; here flat	19.3.4
Violation by a lead	eq. (19.11)	–	the joint tail, racks and copies	19.3.5
Crossing and mass	eq. (19.12)	–	rising, or standing over	19.3.5
Part from below	under at now, over by s	–	a path's reading along draws	19.3.5
Cost of waiting	plan now against one held window	$	read at z now	19.3.6
Stance	last set points left standing	°C	the risk that builds through the stop	19.3.6

19.8 Exercises

Exercise 19.1. *Formulation.* Write the three nodes for one hall and one unread rack, and say which entry no policy reaches. (Closed: r and T_{sp} pinned, valve free, $T_k = r - b - (1-\beta_k)Q_t/\sum C$. Open: valve pinned at its now value, $T_k = T_s + Q_t(1/(u_{\text{now}}G) - (1-\beta_k)/\sum C)$. Feedback: closed, with r and T_{sp} moved by (19.3) on the meter and the valve reading. The entry is β_k: no row of the three reads it, so no affine policy on these readings lowers its band.)

Exercise 19.2. *Derivation.* Derive $\partial T_k/\partial Q$ off and at the stop. (Off it, $T_r = r - b$ does not depend on Q, so $\partial T_k/\partial Q = -(1-\beta_k)/\sum C = -0.5/30 = -0.01667\,\text{K/kW}$. At the stop, $T_r = T_s + Q_t/G$ adds $1/G$, so $\partial T_k/\partial Q = 1/G - (1-\beta_k)/\sum C$; the node returns $+0.03271\,\text{K/kW}$.)

Exercise 19.3. *Derivation.* Show that an inlet's threshold probability on a falling trend is nearly all mass, and give the split for B3 at window 12, lead 132. (By equation (19.8) a rising ramp gives $\gamma < 0$, so $\Pr(\gamma > 0 \mid \mathbf{y})$ is small and the first argument of Φ_2 in (19.12) is far below zero. The crossing cannot exceed that ceiling and the rest of the line is mass. With $\hat{\gamma} = -0.0350$, spread 0.0125, $\hat{f}_s = 25.436°\text{C}$ and spread 0.8694 K: line 0.0360, mass 0.0347, crossing 1.23×10^{-3} by Φ_2. Since $\gamma > 0$ exactly when $R < 0$, the true sign's ceiling is 8.89×10^{-4}, and its crossing is 6.63×10^{-5}.)

Exercise 19.4. *Application.* From the 20:00 row of Table 19.7, give the share of the shift's violation that is there in the first 5 minutes. (The violation by the remainder is 0.0915 and by lead 1 0.0695; the share is 0.0695 of 0.0915, or 0.76.)

Exercise 19.5. *Application.* From Table 19.10, say why 4.000 at window 12 is not a cost, and read the cost that is. (The waited window alone is over at 0.0758 against the 0.05 the joint allows, so no z reaches 0.95 and the search stops at its cap. At z now, waiting costs +0.3747 $ and moves the joint from 0.9530 to 0.9078.)

Chapter 20

Recourse and the Receding Horizon

20.1 The plan restated

The commissioning sheet for the plan of Chapter 18 carries one more line: “re-solve every 5 min: feedback corrects the forecast, so it saves money, and the 95 % holds all night”. The trend at 20:05 is open beside it. The return reading of hall B stands at 30.005°C on a set point r_B of 30.0°C, and the re-solved r_A is not the second move the plan had written down five minutes earlier.

The engineer would conclude that re-solving is free insurance, and would shorten the interval if the controller allowed it. Nothing on the sheet or the trend tests that. The plan is thrown away after one move, and the sheet does not say why that is the right part to keep. It does not say what a response is worth where the forecast was already right. And its 95 % was a statement made at one now, about a sequence that the re-solve no longer runs. The question of this chapter is

why only the first move is taken, what a response is worth, and what 95 % *says once the plan is restated.*

Jargon: receding horizon, innovation, realized recourse, anticipated recourse, conditional guarantee, recursive feasibility, decision interval, same-end route, told routes, net

A *receding horizon* is the rule that takes the first move of a plan, reads, and solves again with the horizon moved on; the framework’s name for each solve is scenario nodes across time-indexed copies. An *innovation* is a reading less its prediction, read against its stated band. *Realized recourse* means that a later move depends on a later reading while the plan that chose the present move does not count on it. *Anticipated recourse* is a plan that counts on its own later reaction. A *conditional guarantee* states the limit at each now, on that now’s belief. *Recursive feasibility* asks whether a re-solved problem still has a move. The *decision interval* is the time between re-solves. The *same-end route* re-solves with its horizon shrinking to the fixed plan’s end, which isolates the response from the moved horizon. The *told routes* are given outright what no reading gives; *told everything* is a physical reference and a bound on cost, not something a controller can do. *Net* is the electricity less the value of what is stored at the first plan’s end, the same value for every route.

Key idea

A reading moves the belief on the state and the load, and it cannot move an entry with a zero column in its row. Every move of a plan is chosen under the belief at its solve, and only the first is still current when it is carried out; the later ones are given up because a surprise can make them wrong, not because keeping them costs more. A response therefore acts as a premium, paid in money where nothing unforeseen arrives and paid back in held limits where something does, in proportion to what the plant still has in reserve. Restating the plan replaces a joint guarantee over one sequence by a conditional guarantee at each now, and the frequency over the shift is a different number that has to be measured.

20.1.1 The restriction this chapter works under

This chapter lies on the row "Limit held after every reading" of Table 7.6. The restriction is a fixed sequence re-solved at each window, with the first move taken. Against Chapter 18 it keeps realized recourse, because the next plan reads the next reading. It gives up anticipated recourse, because no plan counts on its own later reaction, and the plan still does not shape $\mathbf{\Sigma}$, which is coupling 3 of Section 7.1.3. Couplings 1, 2 and 4 of the problem (7.1) remain. The amendment of Section 7.3.3, that restating replaces the joint by a conditional guarantee, is made measurable here. A plan that counts on learning is computed in Chapters 22 to 24, not here. The planner and the belief are those of Chapter 18, cited and not re-derived.

20.1.2 The two halls at a re-solve (step 0)

The plant is the one of Chapter 18: two halls on one water loop, a tank, windows of 5 min. The plan's starts are the ends of windows 12 (20:00), 72 (01:00) and 126 (05:30). The state of a window is, on the reduced state of that chapter's fit,

$$\mathbf{z}_k = (T_{s,k},\ Q_{A,k},\ Q_{B,k},\ g_{A1},\dots,g_{B3},\ \beta_{A1},\dots,\beta_{B5},\ b_A,\ b_B,\ u_{A,k},\ u_{B,k})^{\mathsf{T}}, \tag{20.1}$$

twenty-three entries. The tank is not among them: its belief is a separate scalar filter on its own row, with its own mean and spread. The supply water, the two loads and the two valves belong to the window. The six log-conductances, the ten recirculation fractions and the two return-sensor offsets persist. An inlet is $T_{\mathrm{in},i} = T_c + \beta_i(T_r - T_c)$, and four racks, A1, A3, B2 and B4, have a reading.

The scenario node. Every route has the same node: closed loop at the hall, with r_A and r_B pinned, the valves free and the PI's row present; at a stop the actuator is pinned again. T_{sp} is pinned with its lag and q_t is pinned. Every what-if of this chapter says the same, and no route adds a stop of its own.

Truths, kept apart. Two truths drive the records. *Record's truth* is the stated truth on every record, one point. *Prior truth* draws β, b and a constant g from the stated priors, with the meters' rows with noise only, one of the two cases of Chapter 18's coverage. Only a prior truth can test a probability.

Routes. All routes are paired on records and noise from the now on. *Fixed*: one plan at now, held. *Same end*: re-solved each window, first move taken, horizon shrinking to the fixed plan's end. *Receding*: re-solved each window, first move taken, 12 windows ahead each time. *Receding CE* and *fixed CE*: the same two at $z = 0$, the certainty equivalent. *Told load*: the fixed plan told the record's load, ambient and price, with belief and z kept. *Told everything*: the fixed plan told every entry and path of the draw, at $z = 0$.

Figure 20.1 draws one step of the receding route. The copies run left to right with one now between them. Blue copies are conditioned on their readings and gray copies are averaged under the belief. Between the two panels one reading arrives and the conditioned edge moves by one copy. The unread rack's fraction stands above all copies; it joins every inlet row and no reading row.

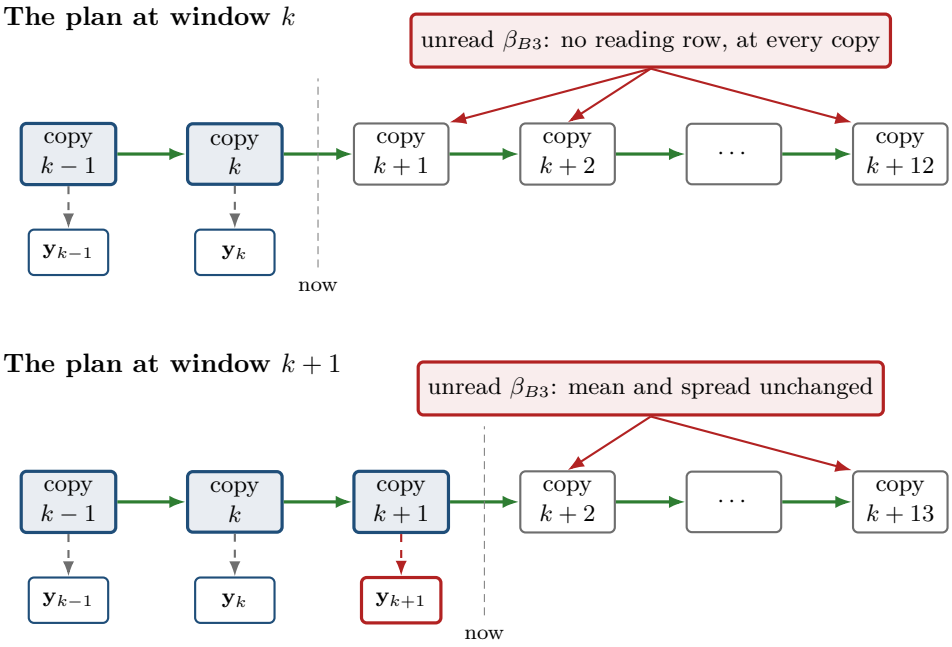


Figure 20.1: One re-solve. Blue: copies conditioned on their readings (the past). Gray: copies averaged under the belief (the future), whose moves are pins with no reading row. Red: the one reading that arrives and the unread entry it does not touch. The first move of the upper plan is carried out; the rest of it is discarded.

20.1.3 The control engineer's questions (step 1)

Table 20.1 lists the seven questions. The first column is the question in the engineer's words, the second the general form, the third its mathematical form on this plant, the fourth its category, the fifth the computation and the last the subsection. Five take a scenario node's reading of the posterior; the last is answered with the candidates of the next measurement.

20.1.4 The evidence and the re-solve as rows (step 2)

The new reading as rows. With $\mathbf{p} = (g, \beta, b)$ the persisting block of equation (20.1) and h the window's reading function, the rows that one reading adds, each divided by its spread, are

$$\mathbf{r}(\mathbf{z}_{k+1}) = \begin{bmatrix} \boldsymbol{\Sigma}_{\mathbf{p},k+1|k}^{-1/2}(\mathbf{p} - \hat{\mathbf{p}}_k) \\ \boldsymbol{\sigma}^{-1}\big(h(\mathbf{z}_{k+1}) - \mathbf{y}_{k+1}\big) \end{bmatrix}, \qquad \mathcal{L} = \tfrac{1}{2}\,\mathbf{r}^{\mathsf{T}}\mathbf{r}, \tag{20.2}$$

with $\boldsymbol{\Sigma}_{\mathbf{p},k+1|k} = \boldsymbol{\Sigma}_{\mathbf{p},k} + \mathrm{diag}(0.005^2 \text{ on } g,\ 0 \text{ on } \beta, b)$, the walk on g of Chapter 18. The information matrix of these normalized rows is $\mathbf{J}^{\mathsf{T}}\mathbf{J}$; with unnormalized rows it is $\mathbf{J}^{\mathsf{T}}\mathbf{W}\mathbf{J}$.

Table 20.1: The questions asked of the re-solved plan. Every what-if is a closed-loop node of Section 20.1.2.

Control engineer's question	General form	Mathematical form	Category	Computation	Answered in
Why take only the first move?	posterior against the planned posterior, (2.10)	innovation against its band; moved entries; plan's second move against the re-solved first	estimation	posterior	20.3.1
What is a response worth?	expected consequence, paired, (2.23)	net, fixed less re-solved, by truth	design	expected consequence	20.3.2
How often must it re-solve?	the same, by interval, (2.23)	5, 15, 60 min, with and without a surprise	design	scenario nodes across time	20.3.3
What does a restated plan hold?	joint chance constraint, restated, (2.12)	conditional at each now against frequency over the shift	prediction under a threshold	posterior	20.3.4
Can the re-solve have no move?	recursive feasibility, (20.13)	least slack at a prohibitive price; the rows' contradiction	design	scenario nodes across time	20.3.5
Where is a response worth more?	what-ifs at scenario nodes, (2.23)	load step, one unit, tank, by node	design	scenario nodes across time	20.3.6
What should be read next?	fresh records with a real reading	saving and value of recourse by instrument	validation	posterior	20.6

Which row acts on which entry. The reading rows have this pattern, with $*$ for an entry that can be nonzero, a property of the graph:

$$\mathbf{J}_{\text{readings}} = \begin{bmatrix} * & 0 & 0 & 0 & 0 & 0 & 0 \\ * & * & * & 0 & 0 & * & * \\ * & * & * & 0 & 0 & 0 & * \\ 0 & * & 0 & 0 & 0 & 0 & 0 \\ * & * & * & * & 0 & 0 & * \\ 0 & 0 & 0 & 0 & 0 & 0 & * \end{bmatrix} \begin{matrix} y_{T_s} \\ y_{T_r} \\ y_{T_d} \\ y_Q \\ y_{T_{\text{in}}} \\ y_u \end{matrix}, \qquad \text{columns } T_s,\ Q,\ g,\ \beta_{\text{read}},\ \beta_{\text{unread}},\ b,\ u. \tag{20.3}$$

The fifth column is zero in every row. That single column decides most of this chapter.

The re-solve. The problem solved at a now is that of Chapter 18, with its priced slack: over the moves $\mathbf{u}_{k+1:k+H}$ of T_{sp}, r_A, r_B and q_t,

$$\begin{aligned} \min_{\mathbf{u},\,\mathbf{s}\geq 0} \quad & \mathbb{E}\big[\text{electricity} - \text{value stored at } k+H \mid \mathbf{y}_{0:k}\big] + \text{price} \cdot \mathbf{s} \\ \text{subject to} \quad & \mu_{i,j}(\mathbf{u}; \hat{\mathbf{z}}_k) + z\ \text{sd}_{i,j}(\mathbf{\Sigma}_k) \leq 27^\circ\text{C} + s_{i,j}, \quad u_{h,j} \leq 0.95 + \text{slack}, \\ & \text{storage, move limits, ranges, capacity, tank bounds, hard,} \end{aligned} \tag{20.4}$$

with i a rack and j a copy. The fixed route solves it once, and the receding route carries out $\mathbf{u}_{k+1}$ and solves it again at $k+1$ with $(\hat{\mathbf{z}}_{k+1}, \mathbf{\Sigma}_{k+1})$.

The modeling choices. Each buys something and can fail.

- *One rule for z.* z is calibrated on the book's record, 2.172, 1.906 and 2.039 at the three starts, the 1.906 under Chapter 18's stated tie-break (2.055 without it), and used as a rule on other records. This buys a like comparison across routes and it fails where another record's belief needs another z; Section 20.3.4 shows it over-delivering by mid-shift.

- *Net as the common column.* Value is taken at the first plan's end for every route, with no price on a violation. This buys a paired difference. It fails once a hall loses hold, where a route can look cheap by overheating, so P and K h are read beside it.
- *A surprise is one step.* Load steps of stated size at lead 2, one unit lost and 250 kWh gone are what-ifs. They buy pairs, and they say nothing of the distribution of real events.
- *Failure fallback.* A re-solve with no solution holds the last feasible plan. It buys a run to the end, and it carries the unit-loss results and the 60 min route under a load step, which is stated where they appear.

20.2 The posterior of the state (step 3)

Each now has one belief, the filter of Chapter 18, and each answer reads it. What is new here is the step from one belief to the next and what that step can and cannot do.

20.2.1 The belief from one window to the next

Each iteration on the rows (20.2) solves the normal equations $(\mathbf{J}^{\mathsf{T}}\mathbf{J})\,\Delta\mathbf{z} = -\mathbf{J}^{\mathsf{T}}\mathbf{r}$, and at the solution

$$\boldsymbol{\Sigma}_{k+1}^{-1} = \mathbf{J}^{\mathsf{T}}\mathbf{J}, \qquad \mathbf{J} = \begin{bmatrix} \boldsymbol{\Sigma}_{\mathbf{p},k+1|k}^{-1/2}\,\mathbf{E}_{\mathbf{p}} \\ \boldsymbol{\sigma}^{-1}\,\partial h/\partial \mathbf{z} \end{bmatrix}, \tag{20.5}$$

with $\mathbf{E}_{\mathbf{p}}$ selecting the persisting entries. Window 1 of the filter is a single fit, and the check against it is Chapter 18's.

Now take an unread entry β_i. Its column in the reading rows is zero by equation (20.3). Its prior is independent of every other entry. So the information matrix has no term joining β_i to anything else, at the first window and, by the same argument, at every later one. The block of equation (20.5) splits, and

$$\hat{\beta}_{i,k+1} = \hat{\beta}_{i,k}, \qquad \Sigma_{\beta_i\beta_i,k+1} = \Sigma_{\beta_i\beta_i,k}, \tag{20.6}$$

exactly, for every reading. This is a row-pattern fact and not a computed result. It is what the plan's back-off stands on, since the inlet's gradient has the entry $T_r - T_c$ in that same column: the column is zero in the reading rows and nonzero in the row of the protected quantity.

20.2.2 The belief at the plan's nows

Table 20.2 gives the belief of the book's record at the first two nows. Each row is an entry; the columns give its mean and spread at windows 12 and 72. These are the diagonal of $\boldsymbol{\Sigma}$ read in its own units. The stated truths belong to the report and are not used here. The tank's belief at the three starts is 298.84, 292.05 and 285.03 kWh.

Along the record the spread of b_B falls from 0.0998 at window 12 to 0.0318 K at window 132, and sd β_{B3} stays 0.0800 at every now, as equation (20.6) requires.

20.2.3 The prediction a plan assumes

The plan at k assumes a posterior at $k+1$. Its reading has the prediction, under the prediction model of Chapter 18 with the belief drawn,

$$\nu_{j,k+1} = y_{j,k+1} - \mathbb{E}\big[h_j(\mathbf{z}_{k+1}) \mid \mathbf{y}_{0:k}\big], \qquad s_j^2 = \mathrm{Var}\big[h_j(\mathbf{z}_{k+1}) \mid \mathbf{y}_{0:k}\big] + \sigma_j^2, \tag{20.7}$$

where ν_j is the innovation of reading j and s_j its stated band. The variance is taken over 400 draws of the belief per record, with the plant's rows closing the hall loops.

Table 20.2: Means and standard deviations of the belief on the book's record (Chapter 18's filter), diagonal of $\mathbf{\Sigma}$.

Entry [unit]	Mean, window 12	sd, window 12	Mean, window 72	sd, window 72
b_A [K]	+0.0296	0.0999	+0.0334	0.0428
b_B [K]	−0.2688	0.0998	−0.1694	0.0428
β_{A5}, unread [–]	0.4500	0.0800	0.4500	0.0800
β_{B3}, unread [–]	0.5000	0.0800	0.5000	0.0800
β_{A3}, read [–]	0.4050	0.0084	0.3949	0.0034
β_{B4}, read [–]	0.3971	0.0102	0.3873	0.0041
Q_A [kW]	335.8467	3.2939	355.5868	3.2956
Q_B [kW]	282.0778	3.1868	301.3652	3.1931

20.3 The answers (step 4)

The first answer reads the belief's movement across one step. The next four read routes run on records through the plant's rows. The sixth reads what-ifs at scenario nodes. Each route's every plan consumes the $(\hat{\mathbf{z}}_k, \mathbf{\Sigma}_k)$ of Section 20.2, through the rows of equation (20.4).

20.3.1 The first move and the innovation

Control engineer's question. Why take only the first move?

General form. Moments of a function of the state under the posterior, equation (2.10), taken before and after one reading.

Specialized form. Equation (2.10) with f an entry of $\mathbf{z}$ and with the evidence $\mathbf{y}_{0:k}$ and then $\mathbf{y}_{0:k+1}$. The reading's effect on a plan's belief is the pair

$$\frac{\hat{z}_{i,k+1} - \hat{z}_{i,k+1|k}}{\sqrt{\Sigma_{ii,k+1|k}}}, \qquad \frac{\sqrt{\Sigma_{ii,k+1}}}{\sqrt{\Sigma_{ii,k+1|k}}}, \tag{20.8}$$

the mean moved in units of the spread before, and the spread after over before. The plan is a function of the belief, $\mathbf{u}^*_{k+1:k+H}(\hat{\mathbf{z}}_k, \mathbf{\Sigma}_k)$, through equation (20.4). Its later moves were chosen under a belief at k, and the belief at $k+1$ replaces it, so the comparison that answers the question is

$$\Delta_k = \mathbf{u}^*_{k+2}(\hat{\mathbf{z}}_{k+1}, \mathbf{\Sigma}_{k+1}) - \mathbf{u}^*_{k+2}(\hat{\mathbf{z}}_k, \mathbf{\Sigma}_k), \tag{20.9}$$

the re-solved first move less the plan's second. The choice entering here is that the innovations are read on fresh records and against the plan's stated band, which lets the band itself be tested; it fails to say anything about readings the plant's rows do not make.

Building blocks. $\hat{\mathbf{z}}_k$, $\mathbf{\Sigma}_k$ of Section 20.2.2; the update (20.5); the innovation and band (20.7); the plan (20.4).

The plant's numbers. On the book's record, one draw, window 13's largest innovations are y_{TinA3} at +1.32 bands and y_{uB} at +1.33, and window 73's is y_{QB}, +12.197 kW, +2.44 bands. Table 20.3 takes this to fresh records, $n = 72$ for each truth and start. Each row is a reading. The first four columns give the root mean square of innovation over band in each cell, and the last the smallest share within two bands over the four. A value near 1 says the band is about the right size.

Table 20.3: Innovations against their stated bands, fresh records, $n = 72$ per cell. R: record's truth. P: prior truth.

Reading	R, start 12 [–]	R, start 72 [–]	P, start 12 [–]	P, start 72 [–]	Smallest share within two bands [–]
y_{T_s}	1.01	0.95	1.04	0.88	0.944
$y_{T_r A}$	0.93	0.99	1.03	0.86	0.917
$y_{T_r B}$	1.15	1.01	0.93	1.06	0.889
y_{QA}	1.29	0.91	1.03	1.04	0.861
y_{QB}	1.02	1.05	1.00	1.02	0.944
$y_{Q\mathrm{loop}}$	0.99	1.02	0.98	1.02	0.931
y_{TinA1}	1.16	1.05	1.17	0.95	0.917
y_{TinA3}	0.97	1.15	1.04	0.97	0.903
y_{TinB2}	1.08	1.04	0.94	0.86	0.944
y_{TinB4}	0.90	1.01	0.90	1.11	0.903
y_E	1.09	0.85	0.92	1.04	0.931
y_{uA}	1.07	1.03	1.02	1.01	0.931
y_{uB}	1.00	1.04	1.04	0.94	0.917
y_P	1.07	0.87	1.10	1.07	0.931

The bands are about right, 0.85 to 1.29, with shares from 0.861 to 1.000 across the cells. There is one exception, stated and not explained. y_{QA} has a mean innovation of +2.398 kW (se 0.736) and +1.854 kW (se 0.509) at the record's truth, against +0.113 and +0.118 kW (se 0.630, 0.632) at prior truth. It is reported as a finding of the band, and its cause belongs to a later Part.

Table 20.4 and Figure 20.2 give the pair (20.8) by class of entry at the record's truth. "Before" is the one-step prediction under the plan's own first move: for a window's own entry its mean and spread over the draws, for what persists the last mean and the spread with the walk. In the table each row is a class and the columns give the two ratios at both starts. In the figure the horizontal axis lists the entries, with the largest of each persisting class, and the vertical axis is the mean moved in units of its spread before. T_s has no bar: its prediction has no spread in the draws, so its row gives the fit less the prediction in K, a gap of 0.953 and 0.902 of the fit's own spread.

Table 20.4: What one reading moves against the one-step prediction, record's truth, $n = 72$. Ranges are over the class.

Entry class	rms(mean moved)/sd before, start 12 [–]	rms(mean moved)/sd before, start 72 [–]	sd after/before, start 12 [–]	sd after/before, start 72 [–]
u_B, u_A	1.120, 1.272	0.977, 1.273	0.5143, 0.5078	0.5304, 0.5217
Q_A, Q_B	1.764, 1.694	1.471, 1.453	1.0214, 1.0278	1.0238, 1.0259
T_s, fit less prediction [K]	0.0826	0.0781	–	–
g (six)	0.409–0.498	0.382–0.485	0.9042–0.9180	0.9019–0.9198
b_A, b_B	0.471, 0.346	0.187, 0.134	0.9649, 0.9649	0.9933, 0.9933
Read β (four)	0.248–0.329	0.105–0.130	0.9606–0.9616	0.9927–0.9930
Unread β (six)	**0.000**	**0.000**	**1.0000**	**1.0000**

The reading moves the state and the load. It never moves the racks that set the limit. The stated inlet band after one update stands at 0.9892 and 0.9995 of what it was before. So what the next plan changes is its starting point, by about one to two predicted spreads on the window's own entries, and its back-off barely changes. The spread ratios on Q

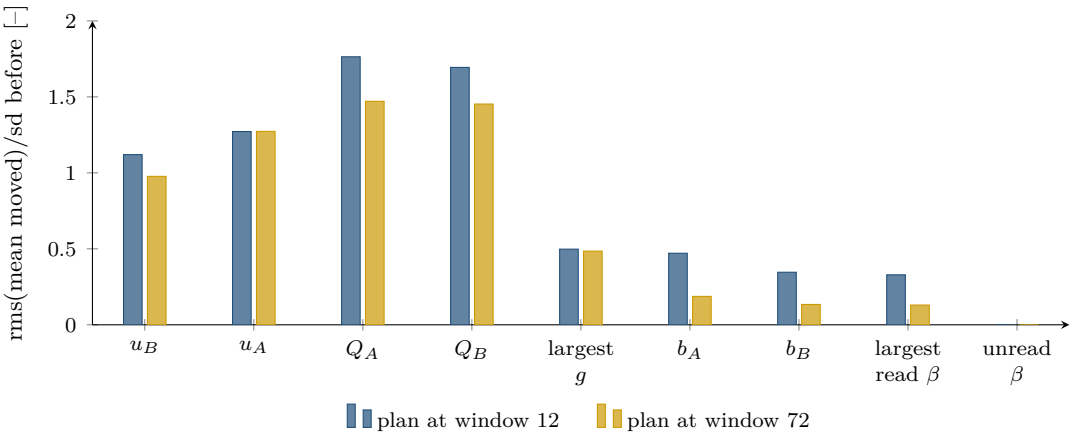


Figure 20.2: The mean moved by one window's readings over its predicted spread, by entry, record's truth, $n = 72$. The last group is exactly zero, by equation (20.6).

sit just above 1 because the fit gives a window's own entries no prior from the last window, so the two spreads are not one conditioned on the other. The change from window k's entry to window $k + 1$'s is larger, up to 8.740 spreads on u_B, and it is not news: it carries the plan's own move and the load's path.

Every move of a plan is chosen under the belief at its solve, and the first is the only one still current when it is carried out. By equation (20.9), re-solved first move less the plan's second, root mean square, at the record's truth: T_{sp} 0.0224 and 0.1234 K, r_A 0.1113 and 0.1139 K, r_B 0.0796 and 0.0304 K. At prior truth: 0.0893 and 0.1613 K, 0.0866 and 0.0831 K, 0.1046 and 0.1146 K, with q_t at 0 and 5.058 kW. The later moves were chosen under a posterior that a reading has replaced. That alone does not condemn them: with no surprise, keeping them is cheaper at equal P, Section 20.3.2. The reason to give them up is the surprise, Section 20.3.6.

One kind of news at a time. On $n = 36$ records the gap is split. With no news and one window more of horizon, it is 0.0000 K on T_{sp} at the record's truth, with r_A at 0 and 0.0068 K; at prior truth it is nonzero on r, up to 0.0835 K, and on q_t, 8.524 kW at start 72. At start 12, record's truth, on r_A, the state read gives 0.0530 K, what persists 0.0995 K and the forecast from the new load 0.0929 K. The three contribute the same order and no one owns the gap. The no-news figures rest on few records: the median is 0.0000 in every cell, the prior-truth r gaps come from 1 or 2 records of 36, and q_t's from 6.

20.3.2 The value of a response

Control engineer's question. What is a response worth?

General form. The expected consequence (2.23) of the routes as actions, paired on records.

Specialized form. Equation (2.23) with a the route, L the net over the 12 windows counted, and the expectation over the truth, the record and the noise. Each route's belief at each now is $(\hat{\mathbf{z}}_k, \mathbf{\Sigma}_k)$ through the plan (20.4). The paired estimate is

$$\Delta_{\text{route}} = \frac{1}{n} \sum_{m=1}^{n} \left[\text{net}_m(\text{fixed}) - \text{net}_m(\text{route}) \right], \tag{20.10}$$

positive when the route is cheaper, with its standard error across records. The choice entering here is to use records at both truths, $n = 60$, and 24 at start 126, which lets the sign be read at the record's truth and over truths drawn; it costs a small sample at the price edge.

Building blocks. The routes of Section 20.1.2; the plan (20.4); the paired difference (20.10); P, the share of records with every inlet at or under 27°C at every window.

The plant's numbers. Table 20.5 gives the routes, and Figure 20.3 draws the four differences. In the table each row is a truth and start, the base columns give the fixed route's net and P, and the rest are the differences (20.10). In the figure the horizontal axis lists the truths and starts and the vertical axis is the difference.

Table 20.5: Fixed less each route, net, 12 windows; $n = 60$, and 24 at start 126. Standard errors in parentheses. Positive means the route is cheaper.

Truth, start	Fixed net [$]	P fixed [–]	Same end [$]	Receding [$]	P receding [–]	Told load [$]	Told everything [$]
Record's 12	9.567 (0.166)	1.0000	−0.044 (0.012)	−0.043 (0.012)	1.0000	−0.026 (0.009)	+0.985 (0.022)
Record's 72	−0.278 (0.061)	1.0000	−0.052 (0.006)	−0.053 (0.006)	1.0000	+0.026 (0.008)	+0.309 (0.009)
Record's 126	8.143 (0.164)	1.0000	−0.136 (0.014)	−0.238 (0.012)	1.0000	−0.064 (0.012)	+0.532 (0.024)
Prior 12	10.421 (0.255)	0.9333	−0.027 (0.012)	−0.027 (0.012)	0.9333	−0.028 (0.010)	+1.482 (0.093)
Prior 72	−0.287 (0.086)	0.8167	−0.034 (0.006)	−0.034 (0.006)	0.8333	+0.034 (0.010)	+0.417 (0.034)
Prior 126	7.335 (0.224)	0.8750	−0.074 (0.013)	−0.129 (0.015)	0.8750	−0.056 (0.011)	+0.767 (0.119)

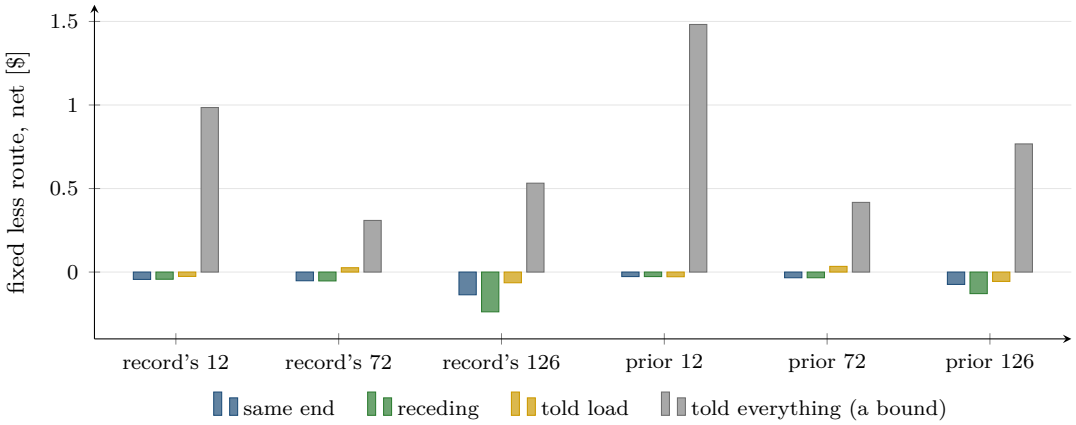


Figure 20.3: Fixed less each route, net, from Table 20.5. Told everything is a bound on cost and not a route a controller can take.

Where nothing surprising happens a response earns nothing in net. Chapter 18's first number was −0.039 $ (se 0.019) on 24 records. Here it is −0.043 $ (se 0.012) on 60, and it holds its sign and shrinks across truths drawn from the prior. Re-solving costs 0.027 to 0.053 $ per hour in this plant at the evening and overnight starts, and buys no measured coverage: P is unchanged but for one record in sixty at prior truth, start 72.

The cap on recourse. Told the hour's load, ambient and price outright, the plan gains at most +0.034 $, and at the evening start it loses. That is the worth of the hour's paths, and it tells no entry of the state. The gap to told everything is not split here: that route also sets $z = 0$ and is told g, b and the read β, which readings do touch. As an order only, on other records, two real readings at unread racks close 0.69 and 0.66 of it at the record's truth and 0.51 and 0.45 at prior truth (Section 20.6). That reference's largest inlets are 27.190 and 27.148°C, so it is a bound on cost only.

Certainty equivalence. Re-solving does not cure it. Receding CE against fixed CE has P at 0.0000 and 0.0000 at the record's truth, and at prior truth 0.1333 against 0.1333 and 0.1833

against 0.1667. Net is 8.096 against 8.083 $ and −0.799 against −0.925 $. Feedback cannot find a band it never reads.

The moved horizon. At start 126, same end and receding part, −0.136 against −0.238 $. The moved horizon sees the 06:00 rise and stores more, 115.6 against 100.0 kWh. The stored value is a stand-in, Chapter 21's, so that part of the difference is a horizon effect and not a response.

Why the sign is negative. The net prices electricity and what is stored, and no row. Fed one channel of news at a time, $n = 20$, the re-solve loses nothing from the state it reads, +0.013 and +0.003 $ at the record's truth. The loss is in the load forecast re-anchored on the fit, −0.045 $ of a total −0.032 $ in the evening and −0.018 of −0.044 $ overnight, and overnight at the record's truth in the persistent entries the fit carries, −0.028 $, whose cause is not taken up here. At prior truth overnight the state itself carries −0.020 $ of −0.047 $. The load runs above the first plan's forecast on 54 of 60 overnight records, by 1.77 kW on average, the ramp being Chapter 19's to read, and record by record the difference tracks that deviation, with correlation −0.56 and −0.57: at the evening start 33 records above lose 0.091 $ and 27 below gain 0.014 $. Each re-solve restores the valve's 0.95 row by cooling the supply, which is why it runs colder, by 0.0434 and 0.1396 K on average over the hour. A fixed plan runs its valve past 0.95 instead, 0.125 and 0.336 unit-windows against 0.080 and 0.158, and pays nothing, since off the stop the net does not see it. Charged at the first plan's own price of that row, 1.155 and 0.528 $ per unit valve per window, the difference goes from −0.044 to +0.008 $ (se 0.007) and from −0.052 to +0.041 $ (se 0.003), at the record's truth; the price is ill-conditioned at prior truth and the test is not read there. On a truth whose loads lie on the forecast's own line the sign is gone: −0.009 (se 0.014) and +0.0005 $ (se 0.006). This is an accounting test and a controlled truth, and it does not say what the row is worth to the plant.

The earlier tests fit this. With nothing to learn, one record with no noise on the forecast's paths, same end differs by −0.0031 and −0.0022 $ at the z rule and by −0.0054 and −0.0425 $ at $z = 0$; the −0.0425 $ goes to +0.0048 $ with the valve priced, the quasi-steady rows in transit breaking the same row. With the same truth and exact readings, $n = 16$, the difference is −0.034 $ (se 0.009) and −0.061 $ (se 0.005), so reading noise is not the cause. With the row at the stop for both routes, $n = 20$, the difference does not close but widens, from −0.036 to −0.106 $ and from −0.057 to −0.103 $ on the same-end route. There is no valve row left to break, and the fixed plan breaks others: it sits at the stop in 0.279 and 0.658 of windows against 0.123 and 0.192, the return stands 0.265 and 0.362 K from its set point against 0.235 and 0.236, and the plan's own inlet runs past its planned row by 2.09 and 2.69 K-windows against 1.88 and 1.32. Priced at that plan's margin, 0.0734 and 0.0319 $ per K per window, the inlet rows close the difference only to −0.091 and −0.060 $ (se 0.016, 0.010). The rest is a hall that has left the rows, which no row of the plan prices, and overnight the routes no longer store alike: the fixed plan spends 0.235 $ more on electricity and credits 0.338 $ more in store. So the valve's row closes the gap where it can be priced, and the stop's case is consistent with this without proving it.

20.3.3 The decision interval

Control engineer's question. How often must it re-solve?

General form. The expected consequence (2.23), with the interval as the action.

Specialized form. Equation (2.23) with a an interval of 5, 15 or 60 min, L the net over 36 windows, and the difference (20.10) taken against 60 min. The plan's own moves are held

between re-solves, 12 windows ahead. The node is read twice, with no surprise and with hall B at +80 kW from lead 2. The choice of one step at one size buys the pair; it says nothing of how often such a step comes.

Building blocks. The plan (20.4) from $(\hat{\mathbf{z}}_k, \mathbf{\Sigma}_k)$ at each re-solve; the difference (20.10); K h over 27, the sum over inlets of time above the limit.

The plant's numbers. Table 20.6 gives both cases, and Figure 20.4 draws them. In the table's upper block each row is a truth and start with no surprise; in its lower block each row is a start and interval at the surprise, with the measures that a net cannot carry. In the figure both horizontal axes are the interval; the left vertical axis, logarithmic, is K h under the surprise, and the right is the net against 60 min.

Table 20.6: The interval, 36 windows. Upper block: no surprise, 60 min less this, net, $n = 24$. Lower block: hall B +80 kW from lead 2, prior truth, $n = 20$.

Truth, start, no surprise	5 min [$]	15 min [$]	P at 5, 15, 60 [–]
Record's 12	−0.143 (0.023)	−0.103 (0.026)	1.0000 all
Record's 72	−0.253 (0.022)	−0.195 (0.017)	1.0000 all
Prior 12	−0.018 (0.027)	−0.012 (0.027)	0.9583 all
Prior 72	−0.226 (0.035)	−0.184 (0.029)	0.8750 all

Start	Interval [min]	P (se) [–]	K h over 27 (se) [K h]	Largest inlet [°C]	Valve at stop [–]	T_s largest [°C]	Tank end [kWh]	Net against 60 min [$]
12	5	0.9000 (0.0671)	0.0232 (0.0227)	25.997	0.113	14.579	0.4	−4.721 (0.582)
12	15	0.9000 (0.0671)	0.0282 (0.0236)	26.108	0.345	14.731	0.4	−3.539 (0.500)
12	60	0.6500 (0.1067)	0.8608 (0.6589)	26.895	0.465	15.297	12.0	–
72	5	0.7000 (0.1025)	0.0810 (0.0605)	26.393	0.172	14.615	223.1	−5.405 (0.405)
72	15	0.6000 (0.1095)	0.1093 (0.0619)	26.686	0.419	14.835	245.2	−4.787 (0.407)
72	60	0.0000 (0.0000)	31.95 (7.2267)	34.182	0.730	21.913	454.9	–

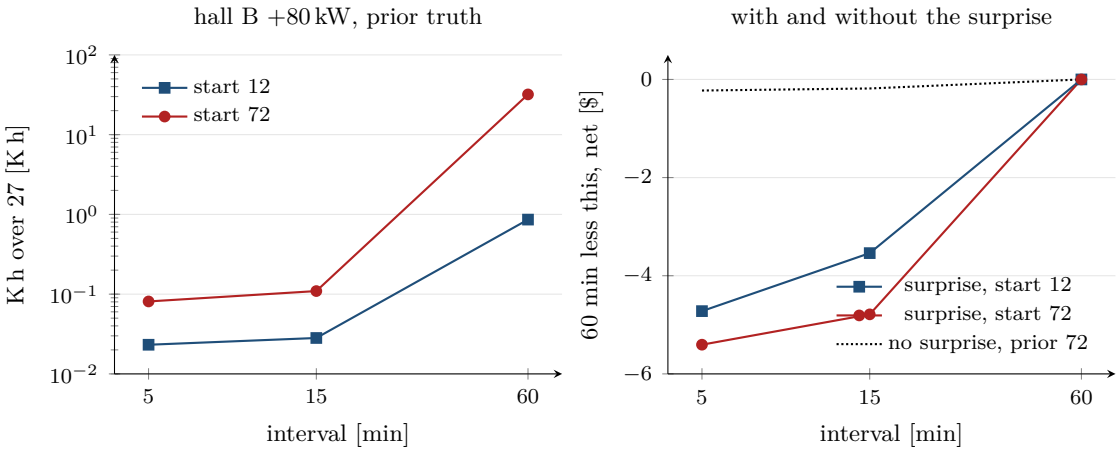


Figure 20.4: The interval. Left: K h over 27 under the surprise, from Table 20.6; at start 72 the 60 min point is mostly carried by a held plan. Right: net against 60 min under the surprise and, dotted, without it. A negative net means the shorter interval spends more.

The interval is not a saving; it is a reaction time, and a late re-solve can find no move. With no surprise the hour is cheapest, or at prior truth, start 12, not separable from the others (−0.018, se 0.027). With one, overnight, the hourly plan ends with the tank charged, 454.9 kWh, while T_s floats to 21.913°C. Most of that is not the wait itself. The hourly re-solves

find no solution on 0.3500 of tries, 0.250 at 60 min and 0.800 at 120 min, and those records run on an expired plan. K h by hour goes 0.782, 5.881 and 25.288, so the first hour is the reaction time's own cost, already over the 5 min route's whole 0.0810. The 16 of 20 records with a held plan average 38.948 K h, and the other four 3.959. Net is the wrong column there: the 60 min route "saves" 5.405 $ by overheating the hall. 15 min gets most of what 5 min gets on P and K h. It does not get it on the valve at the stop, 0.345 against 0.113.

20.3.4 The joint over the shift

Control engineer's question. What does a restated plan hold?

General form. A tail of a function of the state under the posterior, the threshold form (2.12), here the joint chance constraint of equation (6.20).

Specialized form. Equation (2.12) with f the largest inlet over racks and copies and $f^{*} = 27^{\circ}\mathrm{C}$, on the belief at each now. Let A_k be the event that every inlet holds over the next 12 windows under the plan at k. Through the back-off in equation (20.4) the plan states

$$\Pr(A_k \mid \mathbf{y}_{0:k}) \approx 0.95, \tag{20.11}$$

one statement per now, on that now's $(\hat{\mathbf{z}}_k, \mathbf{\Sigma}_k)$. By the total probability rule, equation (7.21), its average over the readings is the unconditional probability of the same event. But the event differs from now to now, and the sequence actually run is none of them. With R_k the event over the next 12 windows along the moves carried out, and $S = \bigcap_k R_k$ the whole shift,

$$R_k \neq A_k, \qquad \Pr(S) = \Pr\Big(\bigcap_{k} R_k\Big), \tag{20.12}$$

and equation (20.11) fixes neither. If the hours were independent, $\Pr(S)$ would be the product of the hours; if they were one event, it would be one hour's. Which is a property of the entries, so it is measured. The assumption entering is that z is held at 2.172, the evening's and the largest of the three, over the shift; a z recalibrated per now is not computed.

Building blocks. $(\hat{\mathbf{z}}_k, \mathbf{\Sigma}_k)$ at each now; equations (20.11) and (20.12); the unread entries' equation (20.6); prior truth, since only a drawn truth tests a frequency.

The plant's numbers. Table 20.7 gives the statements, restated at every window from 12 to the shift's end, prior truth, $n = 72$, 132 windows. Each row is a statement and the second column its value. Figure 20.5 draws the two that matter. On the left the horizontal axis is the hour of the shift and the vertical axis the probability of no window over through that hour: measured, as the product of the hours, and as 0.95 to that power. On the right the bars count the records whose first window over falls in each hour.

The feared repeated-trials effect is absent here, and the reason is not good news. Eleven chances cost 0.9167 against an hour's mean of 0.9495, and not 0.5651, because the violation is mostly one persisting cause on a record: an unread rack with β drawn hot. On this sample's six records that rack's β is +1.89 to +2.77 prior spreads above the survey. Four of the six are over until the shift's last window, one of them on and off, 51 windows in 5 hours. Two are brief, 2 windows and 1, so a record can leave. A hot β does not settle it either: 6 of the 66 records never over have an unread β above +2.172 spreads, and 2 of the six over do not. A second run on other records gives 3 records, 128, 54 and 131 windows, all to the shift's end. In both runs every violation is on an unread rack, in these samples, and by equation (20.6) no reading identifies the cause; re-solving does not remove what persists.

Table 20.7: The joint over the shift, restated each window, z held at 2.172, prior truth, $n = 72$, 132 windows.

Statement	Value
Stated at each now	0.95, conditional on that now's belief
Realized over the next 12 windows, by now: mean, smallest, largest [–]	0.9508, 0.9306, 0.9583
By hour of the shift, no window over [–]	0.931, 0.931, 0.944, 0.944, 0.958 six times, 0.944
Whole shift (se) [–]	**0.9167** (0.0326)
The hours if independent; 0.95 to the eleventh [–]	0.5651; 0.5688
Pr(over next hour | over this hour); | none this hour [–]	0.917 (on 36 hours); 0.0029
First window over, by hour, this sample [records]	5, 1, then 0 nine times
Records over; windows over per such record, mean, largest	6; 59.7, 127
Windows over, each of the six	127, 79, 51, 2, 1, 98
A second run, other records: whole shift (se); records over	0.9583 (0.0235); 3
K h over 27: unread, read racks [K h]	0.0772, 0.0000
Racks ever over [–]	A5 0.042, B3 0.042, others 0.000

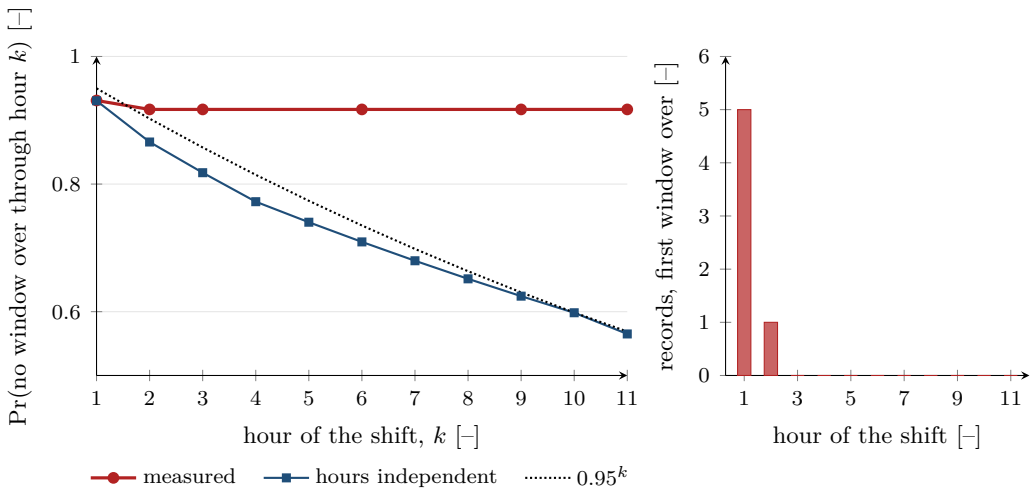


Figure 20.5: The shift, prior truth, $n = 72$. Left: measured against the product of the eleven hours' frequencies, whose last value is 0.5651, and against 0.95^k. Right: the hour in which each of the six violated records first goes over.

The guarantee restated. At each now, conditional on that now's belief, 0.95 for the next hour, under plans not yet made. It is not 0.95 over the shift, and it is not a joint on the sequence actually run. This example's measured shift figure is 0.9167 (se 0.0326), within 1.0 se of 0.95, which does not confirm 0.95; the second run's is 0.9583 (se 0.0235).

The conditional along the book's record. On 2000 draws at each now, equation (20.11) reads 0.9525 at window 12, rises to 0.9810 at window 108, and stands at 0.9615 at window 132, with z held and the spread of b_B falling from 0.0998 to 0.0318 K, and $\mathrm{sd}\,\beta_{B3}$ at 0.0800 throughout. With the persisting entries at their means it is 1.0000 at every now. So the held z over-delivers by mid-shift.

The record's truth. Over $n = 12$ records it gives 1.0000. This is one point, this truth's outcome, and it is not coverage.

What a shift statement would cost. At 15 min, $n = 24$, on the same records: $z = 1.645$ gives 0.7500 (se 0.0884) at 100.115 \$; the rule gives 0.8750 (se 0.0675) at 102.889 \$; $z = 3.0$ gives 0.9583 (se 0.0408) at 107.503 \$. These are read as order only at this n; the rule's 0.8750 lies within 1 se of the 0.9167 above.

20.3.5 Recursive feasibility

Control engineer's question. Can the re-solve have no move?

General form. Whether the set of moves meeting the hard rows and the chance constraint is empty at a now,

$$\mathcal{U}_k = \{\mathbf{u} : \text{hard rows hold},\ \Pr(\text{limits} \mid \mathbf{y}_{0:k}) \geq 1 - \alpha\} \neq \emptyset, \tag{20.13}$$

a property of the belief and the rows, not a cost.

Specialized form. Equation (20.13) on the rows of equation (20.4). Storage, move limits, ranges, capacity and tank bounds hold hard, while inlet and valve rows carry the priced slack of Chapter 18. Two tests are read. Slack left at a prohibitive price, 10 000 \$/K and 1000 \$/kW, says the z rule's rows have no move, and those rows are a conservative restriction of $\mathcal{U}_k$, not $\mathcal{U}_k$ itself: a plan with slack can still meet $\Pr \geq 1 - \alpha$, so that is checked on draws at every event. No solution at all empties $\mathcal{U}_k$ under the hard rows alone. The prohibitive price stands in for an infinite one, by sequential linear programs, so it gives a low slack and not a certified least, and the check is whether it agrees with the stated price's.

Building blocks. The plan (20.4) at $(\hat{\mathbf{z}}_k, \boldsymbol{\Sigma}_k)$; the same plan at the prohibitive price; $P(\text{joint} \mid \text{belief})$ on 4000 draws of the belief.

The plant's numbers. Table 20.8 gives the frequencies on the shift runs of Section 20.3.4. Each row is a measure and the columns are the two truths.

So Chapter 18's slack price is effectively exact on these events, mean 0.2079 against 0.2136 K, the two apart on one event, where the prohibitive-price solve lands 0.2404 K higher. And most events are empty by rule and not by probability: the chance constraint holds on 22 of the 26 after the first plan, smallest over all events 0.7980; the second run has 24 of 27. The first plan's slack comes from the supervisor's start and not from recursion, and the two are kept apart in the table.

Table 20.8: How often the z rule's rows leave a re-solve with slack, shift runs, every window to the end.

Frequency	Prior truth	Record's truth
First plan with inlet slack, of records [–]	0.2222	0.0833
Re-solves after it with inlet slack [–]	0.0028	0.0019
Records with any inlet slack [–]	0.347	0.250
Largest inlet slack [K]	1.161	0.127
No solution [–]	0.0000	0.0000
Every event: slack left at the prohibitive price	42 of 42	4 of 4
Every event: Pr(joint | belief) ≥ 0.95 on 2000 draws	30 of 42	3 of 4
the same, after the first plan	22 of 26	3 of 3

Table 20.9 gives four exhibits on the book's record. Each row is a case, closed loop at the hall, and the columns give what the rows leave at the prohibitive price there and then P conditional on that now's belief under four relaxations. The first case has no what-if at all: record 6 of the record's-truth run, window 89, where the belief holds T_{rB} at 30.535°C against the plan's 30.250°C.

Table 20.9: The relaxations, P conditional on that now's belief, 4000 draws. Standard errors at most 0.0052.

Case	At the prohibitive price	Slack priced [–]	Drop the tail [–]	Hold last plan [–]	No move limit [–]
A reading, window 89	slack 0.0519 K	0.9778	z 2.104: 0.9740	0.9772	0.9778
Load surprise, hall B +80 kW, window 15	slack 0.3300 K, valve 47.381 kW	0.8980	z 1.816: 0.8745	0.9005	0.9375
One unit, evening, told 400 kW	slack 2.8507 K	0.0000	z 0, slack 1.0949 K: 0.0000	0.0628	0.0000
One unit, overnight, window 75	**no solution; slack cannot rescue it**	–	no z down to 0	0.0000, largest inlet 42.118°C	5 windows only, slack 3.7571 K: 0.0000

A case with no move at any price. In the last row T_{sp} in force is 14.129°C, so the move limit caps the row's T_s at 14.797°C, while 676.0 kW against 400 + 200 kW holds T_s no lower than 18.098°C with the halls' heat unstored. The two rows contradict each other. The plan's rows assume a chiller that makes its set point, and a chiller at capacity is outside them.

What guarantee remains in each. With slack priced there is a cost on violation and no probability statement. Dropping the tail leaves a lower level, stated at this now. Holding the last plan leaves the last now's statement, not updated by this reading. No move limit lies outside the plan's rows, by Chapter 18's 0.71 K. The natural case shows that slack against the z rule need not break the joint, 0.9778: it is infeasible by rule, not by probability.

On records. With hall B at +80 kW, record's truth, $n = 16$: fixed has $P = 0.0000$ and 0.5530 K h; slack priced has 1.0000, 0 K h and net 14.724 \$; dropping the tail has the same P at 14.717 \$, with z down to 1.290; holding the plan has 0.0000, 0.4849 K h and 12.375 \$, held on 0.8333 of re-solves. Dropping the tail and pricing the slack differ here by less than a cent, and holding the plan discards the response.

20.3.6 Where a response is worth more

Control engineer's question. Where is a response worth more?

General form. The expected consequence (2.23), read at what-ifs.

Specialized form. Equation (2.23) with a the route and the expectation conditioned on a scenario node: closed loop at the hall, T_{sp} and q_t pinned by the route, and at a stop the valve pinned again. The difference is that of equation (20.10), read beside P and K h. Negative means the response pays money, and what it buys is in the other columns. The unit's status is told to one route alone and otherwise reaches the plan only through T_s in the fit's rows; that choice is what the "told" row prices, and the lost unit here is only a what-if. Isolating a fault is Chapters 27 and 28.

Building blocks. The routes; the difference (20.10); the plan (20.4); the fallback of Section 20.1.4.

The plant's numbers. Table 20.10 gives the nodes, 12 windows, $n = 20$. Each row is a node, truth and start. The columns give P, K h, the net against fixed and the share of periods with a valve at its stop, one value per route in the order named.

Table 20.10: Scenario nodes, closed loop at the hall, $n = 20$. Routes in the order fixed / receding / told or read. Net: fixed less the later routes.

Node	Truth, start	P (se) [–]	K h over 27 [K h]	Net against fixed (se) [$]	Valve at stop [–]
Hall B +40 kW	record's 12	1.000 / 1.000	0 / 0	−1.209 (0.036)	0.431 / 0.027
Hall B +40 kW	record's 72	0.300 (0.102) / 1.000	0.0281 / 0	−0.448 (0.012)	0.458 / 0.071
Hall B +80 kW	record's 12	0.000 / 1.000	0.6008 / 0	−2.381 (0.055)	0.458 / 0.187
Hall B +80 kW	prior 12	0.450 (0.111) / 0.900 (0.067)	0.2231 / 0.0287	−1.879 (0.215)	0.423 / 0.154
Hall B +80 kW	prior 72	0.300 (0.102) / 0.850 (0.080)	1.0676 / 0.0317	−0.738 (0.088)	0.537 / 0.202
One unit lost	record's 72	0 / 0 / 0	**31.9290** / **12.0685** / **1.9556**	−0.627 (0.078), −1.156 (0.017)	0.840 / 0.804 / 0.115
One unit lost	prior 72	0 / 0 / 0	32.7741 / 14.3663 / 5.7358	−0.546 (0.085), −1.019 (0.189)	0.800 / 0.731 / 0.248
One unit lost	record's 12	0 / 0 / 0	0.1264 / 0.1339 / 0.4539	+0.006 (0.009), −0.029 (0.021)	0.560 / 0.537 / 0.004
Tank down 250 kWh	record's 12	1.000 all	0	−0.028 (0.024), −0.092 (0.025)	0.000

A response acts as a premium: a small cost where nothing unforecast arrives, returned in held limits where something does. At prior truth, evening, +80 kW, P goes from 0.450 to 0.900 and K h from 0.2231 to 0.0287, and the receding route spends 1.879 $ more. That is extra spend on cooling the fixed plan does not run, and not a price of P: net is unlike across routes once a hall loses hold.

Recourse is worth what the plant keeps in reserve. The same lost unit is worth K h at night, 19.860 read through T_s and 29.973 told before the loss, and nothing in the evening. Overnight the tank is idle with 290 kWh in reserve, and the response spends it: the largest T_s goes from 28.999 to 18.903°C when told. In the evening the plan is already discharging at

−200 kW. Nothing is left to respond with, and telling the plan makes K h worse, 0.1264 to 0.4539, with no solution on 0.425 of re-solves and the last plan held. A plan that values that reserve is Chapter 21's and is not computed here.

Telling. The unit's status against reading it through T_s goes from 12.0685 to 1.9556 K h, though 0.762 of those re-solves have no solution and hold the last feasible plan. That fallback carries the result, and so does the timing: the "told" route has the status at its first re-solve, one window before the re-solve that sets the move for the loss window. Table 20.11 moves that time, on the same records, unit lost overnight. Each row is a route and the columns give K h and the share of re-solves with no solution at each truth.

Table 20.11: When the plan is told of the lost unit, start 72, $n = 20$, standard errors in parentheses.

Route	K h, record's [K h]	No solution [–]	K h, prior [K h]	No solution [–]
Fixed plan	31.9290 (0.1760)	0.000	32.7741 (3.0911)	0.000
Never told, read through T_s	12.0685 (2.4702)	0.296	14.3663 (2.7074)	0.383
Told before the loss	1.9556 (0.0662)	0.762	5.7358 (2.5759)	0.787
Told at the plan for the loss window	1.8749 (0.0602)	0.754	1.2191 (0.2960)	0.725
Told 5 min late	32.7828 (0.1828)	0.833	32.7765 (2.7515)	0.833
Told 10 min late	28.9367 (0.2321)	0.750	29.8218 (2.7122)	0.750
5 min late, no move limit on failure	6.2428 (0.1795)	0.833	3.8716 (0.5209)	0.833

A status one window late is worse than no wire, and no better than the fixed plan. By then the told rows have no solution, the pre-loss plan is held, and it asks for what the lost unit cannot give, while the untold route keeps solving on its wrong capacity and moving. A loss that falls inside a window reaches the plan no sooner than the next re-solve, so the late row is not a corner case. The last row keeps the same late status and replaces the held plan by a solve without the move limit, shortening the horizon until one exists. That is outside the plan's rows, as in Section 20.3.5, and it recovers most of the value. So the wire's value depends on its latency and on the fallback.

The stop. Figure 20.6 follows the book's record. In the upper panels hall B takes +40 kW from window 14, with window on the horizontal axis: the left gives the valve u_B and the right gives T_{rB} and T_{sp} for the fixed and receding plans. In the lower panels, 250 kWh of the tank are gone at window 14: the left gives the tank, the reading and the belief, and the right the belief's stated spread. Under the fixed plan u_B is 0.9974 at window 14, 0.9993 at window 15, then 1.0000, with T_{rB} at 31.342°C against r_B of 29.961°C at window 24 and the reading at 30.754°C. At the stop the hall's set point stops meaning anything, and the only move left is T_{sp}. The receding plan drops T_{sp} through 15.0, 14.5 and 14.022°C and raises r_B to 30.374°C, coming off the stop: u_B is 0.9572 at window 15 and 0.9771 at window 24, with B3 at 25.796 against 26.522°C. This is the valve as the protected quantity of Chapter 15: a move upstream restores what the integrator cannot.

The tank at its limit. This costs little in net and exposes the filter. With 250 kWh gone at window 14 the tank stands at 15.0 and then −0.2 kWh, its floor in the plant's rows, which stop a discharge at the first step past empty. The readings are 20.3, −3.8 and 6.8 kWh. The belief goes 249.3, 217.7, 189.5 and on to 19.7 kWh at window 24, its stated spread 1.291 falling to 1.000 kWh, while the plan asks −200 kW of an empty tank all hour. The tank's row has no walk, so a reading over 200 kWh away moves the belief only by a small share, on top of the

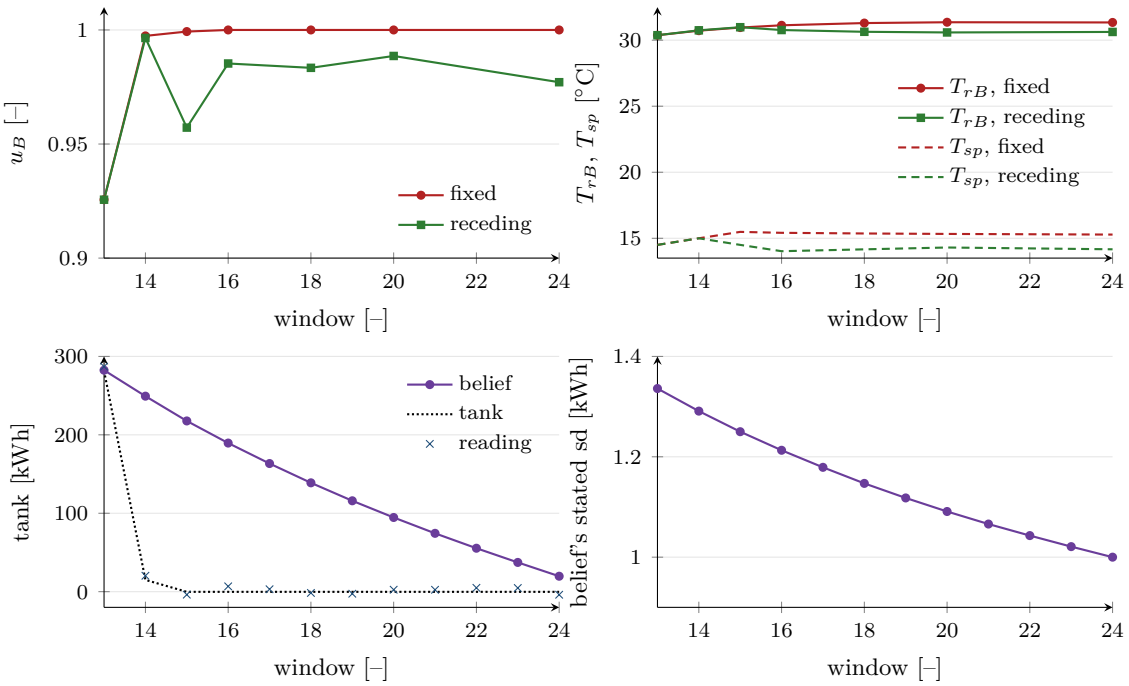


Figure 20.6: The book's record at two what-ifs, start 12. Upper: hall B +40 kW from window 14; solid T_{rB}, dashed T_{sp}. Lower: 250 kWh gone at window 14, receding; the tank's value, the reading and the belief, with the belief's stated spread at right.

asked move's fall: it falls 31.6 kWh in window 15 and 17.6 kWh in window 24. The stated band is smaller than the error by two orders. Taking the reading at face value restores the row, with the tank ending at 2.4 kWh and net at −0.092 $ (se 0.025) in the evening, because nothing else depended on it here. The finding is stated here, and the fix belongs with Chapter 21.

20.4 The checks (step 5)

The routes are tested against the plant's own simulation, the planner against Chapter 18's, and the state route against the closed form of one reading. Table 20.12 lists them.

20.4.1 The routes and the limit of no noise

Each row of Table 20.12 is a check, with its result and what it compares against. The check of $\sigma_p \to 0$ here is the nothing-to-learn row of Section 20.3.2: with no noise, no reading says anything, and the difference falls to −0.0031 and −0.0022 $ at the z rule; it does not fall at $z = 0$, start 72, −0.0425 $, and that is carried into the sign's test.

20.4.2 The closed form of one reading

With β_{B3} alone and a reading row of spread 0.3 K, the update (20.5) on a one-entry state is $1/\sigma^2 = 1/0.08^2 + (T_r - T_c)^2/0.3^2$, with $T_r - T_c$ at 9.860 K, which gives 0.02844. The filter's row also touches the entries β shares with T_r and T_c, and it gives 0.03164. The two are not meant to be identical, and the wider value is the one the plan reads. On the zero column, by contrast, the check is exact: every unread β prints a mean moved of 0.000 and a spread ratio of 1.0000 in every cell.

Table 20.12: The checks. Monte Carlo standard errors appear with every row of the answers.

Check	Result	Against
Stepped rows, 3 records, 24 windows: readings, X, inlets	gap 0.0	the plant's simulation
Filter	entries and $\mathbf{\Sigma}$ gap 0.0	Chapter 18's filter
Cached terminal value	V and slope gap 0.0	Chapter 18's value
Book's plan, window 12, $z = 2.172$	r_B 29.961, T_{sp} 15.281°C	Chapter 18's, the same
Warm re-solve at a 0.25 K region, 44 re-solves	first-move gap 0.0000 K; net gap under 0.00001 $	wide region, no start
A plan reads no truth, 4 records	largest move gap 0.0	truth's entries changed
Interval 12 against 50; run twice	net gap 0.0; gap 0.0	the same routes
Joint at window 12, seed 18	0.9504	Chapter 18's calibration
$\sigma_p \to 0$: nothing to learn, one record	z rule −0.0031, −0.0022 $; $z = 0$ −0.0054, −0.0425 $ same end	noise on the records
Told everything, 4 records, largest inlet	27.172 to 27.250°C	over 27 by construction
The reading's row, B3, one window	sd β 0.08000 to 0.03164	β alone in closed form, 0.02844

20.5 The report (step 6)

Re-solving costs a little where nothing happens, 0.027 to 0.053 $ per hour at the evening and overnight starts, because it holds a row the net does not price, and is worth a great deal where something does: at prior truth, evening, +80 kW, it takes P from 0.450 to 0.900, spending 1.879 $ more. It moves the state and never the racks that set the limit. And 95 % is conditional at each now, while the shift's 0.9167 (se 0.0326) reflects mostly one persisting cause on a record, not many chances.

20.5.1 The answers and their bands

Table 20.13 gives the answers in the order of Table 20.1, with the stated truth as its last row; Figure 20.7 sets the belief beside that truth. In the table each row is a question, and the columns give the answer and its band or basis. In the figure each row is an entry, the bar is the mean with two stated spreads on the book's record at window 72, and the dotted mark is the stated truth, which no plan reads.

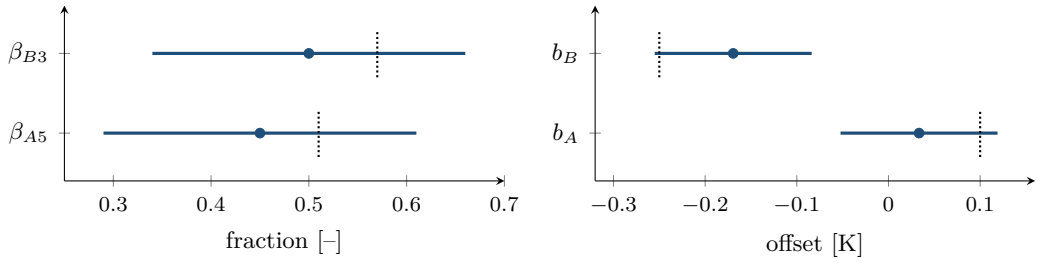


Figure 20.7: The belief at window 72 on the book's record against the stated truth. Bar: mean with two stated spreads, from Table 20.2. Dotted: the stated truth, one point. The unread fractions keep their priors at every now.

20.5.2 The entry in the commissioning log

The entry takes the first move only. It chooses the interval as a reaction time, 15 min by default for a plant of this inertia and shorter if the valve's stop matters, knowing that under the +80 kW

Table 20.13: The report for the re-solved plan. Dollars and percentages are this example's.

Control engineer's question	Answer	Band or basis
Why take only the first move?	The first is the only move still current when carried out; the later ones are given up for the surprise's sake, since with none keeping them is cheaper. T_{sp} re-solved less planned 0.0224, 0.1234 K rms; unread β moved 0.000; inlet band 0.9892, 0.9995 of before	innovation over band 0.85 to 1.29; y_{QA} mean +2.398 kW (se 0.736) a finding of the band
What is a response worth?	A premium. No surprise: −0.043 $ (se 0.012), −0.053 $ (se 0.006) at the record's truth; −0.027 $ (se 0.012), −0.034 $ (se 0.006) at prior truth. The sign is the valve row the fixed plan breaks unpaid: priced, +0.008, +0.041 $	$n = 60$; told load at most +0.034 $; told everything a bound
How often must it re-solve?	A reaction time, and a late re-solve can find no move. No surprise, the hour is cheapest; hall B +80 kW, start 72: K h 0.0810, 0.1093, 31.95 at 5, 15, 60 min, the last with no solution on 0.3500 of re-solves	$n = 20$; se 0.0605, 0.0619, 7.2267
What does a restated plan hold?	0.95 at each now on that now's belief. The shift measures 0.9167, against 0.5651 if the hours were independent	se 0.0326, not confirmation of 0.95; a second run 0.9583; violations on unread racks in both samples
Can the re-solve have no move?	Yes. On one unit, overnight, no solution: 18.098 against 14.797°C. Below 0.95 under every relaxation inside the rows: 0.9005 at most under a load surprise. Slack on the z rule's rows, after the first plan, 0.0028 of re-solves, is mostly not such a case: 22 of 26 hold 0.95	every event read, 2000 draws; 4000 per exhibit
Where is a response worth more?	Where the plant has reserve: one unit overnight 31.9290, 12.0685, 1.9556 K h; in the evening 0.1264, 0.1339, 0.4539 K h. Told one window late, 32.7828	$n = 20$; leans on a held last plan, 0.762 of re-solves
What should be read next?	Inlet readings at B3 and A5: +0.761 $ (se 0.075), +0.188 $ (se 0.017) at prior truth, on the first plan	$n = 20$; Section 20.6
Stated truth	$\beta_{B3} = 0.57$, $\beta_{A5} = 0.51$, $b_A = +0.10$ K, $b_B = -0.25$ K	one point; no plan reads it

step it leaves inlet slack on 0.7250 and 0.7833 of its re-solves against 0.1486 and 0.1736 at 5 min. It states the limit as conditional at each now and never for the shift. It prices slack and logs each use as an event, with the fallback named. It wires the unit's status to the plan, with a latency under one window or a fallback other than the stale plan, since one window late the wire is worth less than none. It asks for a walk on the tank's row. And it proposes inlet readings at B3 and A5. Each of these rests on a constructed record and on the stated choices.

20.5.3 Common mistakes

Table 20.14 lists six mistakes the sheet's line makes easy. Each row names the mistake, its form on this plant and the correction.

Table 20.14: Six mistakes with a re-solved plan.

Mistake	On this plant	Correction
Stating 95 % for the shift	0.95 is stated at each now; the shift measures 0.9167 (se 0.0326)	state the limit conditional at each now, and measure any frequency on drawn truths
Reading $P = 1.0000$ at the record's truth as coverage	every route at start 12	one truth is one point; use prior truth
Expecting feedback to find an unread rack	unread β moved 0.000, spread ratio 1.0000	a zero column reads nothing; add the row
Reading net alone across routes	the 60 min route "saves" 5.405 \$ with T_s at 21.913°C	read P and K h beside net once a hall loses hold
Taking one instrument for half of two	B3 alone −0.665 \$, both +0.685 \$, record's truth, start 12	one reading moves the bind to the other rack
Trusting a stated band that has no walk	tank belief 19.7 kWh, spread 1.000 kWh, tank −0.2 kWh	test innovations against the band; give the row a walk

20.6 The next measurement (step 7)

The candidates here are real readings on each record, entering the filter as real rows, with z recalibrated on the book's record, over 12 windows, $n = 20$. A new row returns to step 3 and rebuilds the belief.

20.6.1 The candidates

Table 20.15 sets the candidates side by side. Each row is a candidate. The columns give z at the two starts, the saving against no instrument's fixed plan at the two truths, and, within the instrument, fixed less receding, which is the value of recourse with it. Each cell lists starts 12 and 72.

The instrument saves on the first plan, not by making recourse valuable: with both readings the value of recourse stays within about ±0.04 \$. Chapter 18's means-held +1.153 and +0.519 \$ were flattered. With real readings they are +0.761 and +0.188 \$ at prior truth, on different records, $n = 20$; the gap between the two statements is partly the z rule and is not separated. B3 alone is worse than nothing at the evening's record's truth, −0.665 \$, as Chapter 18's move of the bind predicted in sign, since its z of 3.000 comes with A5 unread. Its P at start 72, record's truth, is 0.5000 (se 0.1118) fixed and 0.3000 (se 0.1025) receding. That is this truth's outcome on a z from the book's one record, and it is not a coverage statement.

Table 20.15: The candidates, net [$], $n = 20$; each cell starts 12, 72, standard errors in parentheses.

Candidate	z [–]	Saving, record's [$]	Saving, prior [$]	Recourse with it, record's [$]	Recourse with it, prior [$]
None	2.172, 1.906	–	–	−0.053 (0.017), −0.043 (0.008)	−0.016 (0.022), −0.030 (0.008)
Permanent inlet readings, B3 and A5	3.078, 3.109	+**0.685** (0.019), +**0.203** (0.004)	+**0.761** (0.075), +**0.188** (0.017)	+0.029 (0.020), −0.044 (0.009)	+0.004 (0.024), −0.041 (0.009)
Permanent inlet reading, B3 alone	3.000, 2.180	−0.665 (0.038), +0.183 (0.011)	−0.057 (0.171), +0.157 (0.052)	−0.101 (0.027), −0.021 (0.005)	−0.002 (0.021), −0.052 (0.013)
Unit status to the plan	–	K h 12.0685 to 1.9556, one unit lost, overnight; one window late 32.7828	K h 14.3663 to 5.7358; one window late 32.7765	–	–
A walk on the tank's row, or the reading at face value	–	tank end 2.4 against −0.2 kWh; −0.092 (0.025) evening	–	–	–

20.6.2 The recommendation

The permanent instrument is the pair of inlet readings at B3 and A5, recalibrated together. Unit status comes next, as a wire and not a sensor, for the overnight node, and only with the latency or fallback of Table 20.11. The tank's row is a change to the fit and not a measurement. A single inlet reading at B3 is not recommended on the evening's evidence here.

20.6.3 What the analysis showed

The sheet read re-solving as free insurance that also saves money and carries 95 % through the night. The analysis finds a premium: a small steady cost where the plant runs off the forecast's line and a row is held, nothing where the forecast holds, and a large return where something unforecast arrives, limited by what the tank still holds. Only the first move is taken because it is the only one still current when carried out, and because a surprise can make the later ones wrong; with no surprise they were cheaper. The difference between beliefs lies in the state and the load and not in what binds. So the shift's violations are mostly one persisting cause that no reading identifies, and the sheet's sentence has to be rewritten as a conditional statement. The outline's recursive-feasibility case appears here with a real absence of any move, in rows that assume a chiller making its set point. And the first number of Chapter 18 stands: no gain at either truth without a surprise, and the loss is a row the net does not price.

20.6.4 Limits of the model

The sign. The sign is identified as the pricing of a row, by an accounting test at the record's truth and a controlled truth; the price is not read at prior truth, and the evening's forecast channel at the record's truth, larger than on the forecast's own line, is not split here.

Recourse. Anticipated recourse is not computed. The net's stored value is a stand-in, and it carries the horizon effect at start 126.

Unlike columns. Net is unlike across routes once a hall loses hold, so K h and P are read beside it, and no price on violation is chosen.

The rule for z. It comes from the book's record, is used as a rule and is held over the shift.

The sampling. Prior truth at start 72 has fixed P of 0.8167 (se 0.0500), 2.7 se under 0.95, where Chapter 18's was 0.9250 (se 0.0186) on other records; at $z = 2.055$, without that chapter's tie-break, it is 0.8333 (se 0.0481), 2.4 se under. Six violations carry Section 20.3.4. Instrument rows are at $n = 20$. Inlets are sampled at window ends.

The truths and the surprises. Prior truths have a constant g and the meters' rows with noise only; Chapter 18's other meter case is not run here. Surprises are stated steps of one size at one lead, at a window's start. Unit-loss outcomes and the 60 min route under a step lean on a held last plan.

The constructed record. The record is constructed from a stated truth, so its agreement tests the method on its own model and not the model against a real plant. Percentages and dollars are this example's.

20.7 What to retain

Definition 20.1 (The conditional guarantee)**.** A plan restated at each now states the limit on that now's belief, equation (20.11), under plans not yet made. It is a different statement from a joint on the sequence run and from a frequency over the shift, equation (20.12), which depends on how the violations persist and has to be measured on drawn truths.

Three conclusions carry beyond this plant as statements of form. A reading cannot move an entry with a zero column in its row, equation (20.6), however often the plan is restated. The first move is the only part of a plan still current when it is carried out. And recursive feasibility is a question about the rows, equation (20.13), so a re-solve can lose every move through a physical limit the rows do not hold. The dollar and percentage figures are this example's. The chapter's questions join the others in Appendix A, and the next chapter takes up what the reserve is worth.

Table 20.16 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 20.16: The terms of the chapter and where each is defined. Units for the two-hall plant.

Term	Expression	Units	Meaning	Defined in
Receding horizon	eq. (20.4)	–	first move of a re-solved plan	20.1.4
Zero column	eq. (20.3)	–	no reading row on an unread β	20.1.4
Unread entry's update	eq. (20.6)	–	mean and spread unchanged, exactly	20.2.1
Innovation, band	eq. (20.7)	as read	reading less prediction; its stated sd	20.2.3
Gap in the first move	eq. (20.9)	K, kW	re-solved first less plan's second	20.3.1
Paired difference	eq. (20.10)	$	fixed less route, net	20.3.2
Conditional guarantee	eq. (20.11)	–	the limit at each now, on that now's belief	20.3.4
Shift frequency	eq. (20.12)	–	the run's joint, measured	20.3.4
Recursive feasibility	eq. (20.13)	–	the re-solved problem still has a move	20.3.5
Realized recourse	first move re-solved	–	a later move on a later reading, not counted on	20.1.1
Anticipated recourse	not computed	–	a plan counting on its later reaction	20.1.1

20.8 Exercises

Exercise 20.1. *Formulation.* A tank with one pump holds a level, with an unread outflow fraction β and a level reading. Write the re-solved problem, say what is conditional on what, and say which entry no reading touches. (At each now: minimize the expected pump cost over the next moves, subject to the storage rows, the pump's limits and Pr(level $\leq$ top at every copy | $\mathbf{y}_{0:k}$) ≥ 0.95; carry out the first move, read and solve again. Each statement is conditional on that now's belief and on plans not yet made. If β has a zero column in the reading rows, as here, its mean and spread are unchanged at every now, equation (20.6).)

Exercise 20.2. *Derivation.* Show why restating gives a conditional guarantee and not a joint. (Each plan states $\Pr(A_k \mid \mathbf{y}_{0:k}) \approx 0.95$. By equation (7.21) this averages over the readings to a statement about A_k, the event under the plan at k. The moves carried out after k come from later plans, so the run's event R_k differs from A_k, and equation (20.12) has no value fixed by these statements. On this plant, the conditional along the book's record ranges from 0.9525 to 0.9810 with the same z.)

Exercise 20.3. *Derivation.* Derive the rows' contradiction for one unit lost overnight, window 75, with T_s fit at 17.554°C. (Set point in force 14.129°C, move limit 0.5 K: 14.629°C at most. The lag $C_w/W_{\text{ch}} = 42000/400 = 105$ s over 300 s gives a row's T_s capped near $14.629 + (17.554 - 14.629)\text{e}^{-300/105}$. The heat balance with 676.0 kW against $400 + 200$ kW gives T_s of about $17.554 + 76.0 \times 300/42000$ as a floor. No move meets both: 18.098°C against 14.797°C.)

Exercise 20.4. *Application.* From the hours' figures, set the independent product against the measured shift. (The product of 0.931, 0.931, 0.944, 0.944, 0.958 six times and 0.944, carried with the 72 records' unrounded fractions, is 0.5651, against $0.95^{11} = 0.5688$ at the stated level. The measured whole shift is higher because the hours mostly share one cause, an unread rack drawn hot: P(over next hour given over this hour) is 0.917. So 0.5651 against 0.9167.)

Exercise 20.5. *Application.* A candidate instrument has +0.761 $ of saving at prior truth, start 12. Say whether it makes recourse more valuable. (No. Within the instrument, fixed less receding is +0.004 $ (se 0.024) at prior truth, start 12, against −0.016 $ (se 0.022) with none: the saving is on the first plan, and the value of recourse with it is +0.004 $.)

Chapter 21

Thermal and Electrical Storage

21.1 Two stores and a column on the sheet

The commissioning sheet is open at 19:00. The tariff reads 0.22 \$/kWh now and 0.07 \$/kWh after 23:00, and the sheet's line beside it says "a ratio of 3.14, so charge overnight and discharge in the evening; the price forecast is ±10 %, so derate the schedule's saving by 10 % and buy a better price forecast". The chilled-water tank reads 300 kWh, and a column of the sheet says "end the shift at 300". The UPS battery beside it reads 400 kWh, and its line says "reserve 120 kWh, do not touch".

The engineer would accept the arithmetic and the column: the ratio is plainly above the loss, the 10 % looks prudent, and ending where the shift began looks neutral. None of those lines says when to charge. The question of this chapter is

when to charge, what the ±10 % changes, what a kilowatt-hour left at 07:00 is worth, and how much each store should hold back.

Jargon: stance, store, terminal value of a store, certainty equivalence in price, value of recourse, COP arbitrage, round trip, reserve as a chance constraint, guard, ride-through load, wear flip

The *stance* is the hall's chance-constrained hold, a function of the belief alone; it fixes the supply water and with it the price of cooling $\kappa_k = 1/\mathrm{COP}_k$. The *store* is the tank or the battery, with its level as a state. The *terminal value of a store* is the value function at the horizon, $V(E)$, here in a cyclic form over an idle day. *Certainty equivalence in price* says that a fixed schedule on a linear cost depends on the price only through its mean. The *value of recourse* (the value of the response) is the expected cost of the fixed schedule less that of the policy that reads the price. *COP arbitrage* is charging where the lift is small. A *round trip* is one kilowatt-hour out of a store and what had to enter it earlier. A *reserve as a chance constraint* is a quantile of the need, held on a store. The *guard* is the power held below the chiller's capacity to keep the supply water on its set point. The *ride-through load* is the load the battery carries through an outage. The *wear flip* is the cost per kilowatt-hour out at which arbitrage stops.

Key idea

The tank's move enters the hall's rows only where the chiller's capacity binds, so the shift separates into a stance and a store. The store's schedule on a linear cost depends on the price's mean, the band moves no schedule, and a derate has no row to act on. A policy that reads the price earns a value of recourse that grows faster than the band, and that needs a store free to pick its hours. The value of what is left at the end decides the morning. The band on the load, not on the price, is what a reserve holds, and it matters only near capacity.

21.1.1 The restriction this chapter works under

This chapter lies on the row "a re-solved plan with a store" of Table 7.6, a restriction of the problem (7.1). Of the couplings of Section 7.1.3, coupling 1 is kept, since the moves are the future copies' entries, and coupling 2 is kept, since a price read moves the next move. Coupling 3 is given up: nothing is learned from a tank move. Coupling 4 is split: the load's band leaves the price recursion and enters as a guard and a reserve, and the stance does not respond to the store. The belief, the filter and the stance's multiplier are Chapter 18's, and are cited, not derived again. That chapter owed this one the value of what its hour's plan leaves in the tank.

21.1.2 The two-hall plant and the store (step 0)

The plant is the two-hall plant of Chapter 18 over its whole shift, 144 windows of $\Delta t = 5$ minutes from 19:00 to 07:00. The tank is a state with one row,

$$\begin{gathered} E_k = f E_{k-1} + \Delta t\,(\eta\, q_k^{+} - q_k^{-}), \quad k = 1, \ldots, 144, \\ f = 0.999583, \quad \eta = 0.95, \quad 0 \le E_k \le 600\,\text{kWh}, \quad |q_k| \le 200\,\text{kW}, \end{gathered} \tag{21.1}$$

with E_k the level at the end of window k, $E_0 = 300\,\text{kWh}$, $q_k = q_k^{+} - q_k^{-}$ the move, positive charging, and f the standing loss over one window. Its cost and its one tie to the hall are

$$\begin{aligned} \text{cost} &= \sum_k p_k\, \kappa_k\, \Delta t\,(\text{heat}_k + q_k), \qquad \text{heat}_k + q_k \le 800\,\text{kW}, \\ \kappa_k &= \frac{1}{\text{COP}(T_{s,k}, T_{\text{amb},k})} = \frac{T_{\text{amb},k} + T_{\text{app}} - T_{s,k}}{\eta_{\text{cop}}\,(T_{s,k} + 273.15)}, \end{aligned} \tag{21.2}$$

with $T_{\text{app}} = 15\,\text{K}$, $\eta_{\text{cop}} = 0.45$ and p_k the price. The tariff is 0.22 from 19:00, 0.14 from 21:00, 0.07 from 23:00, 0.10 from 05:00 and 0.18 \$/kWh from 06:00 to 07:00. The price's log factor a_k follows an AR(1) with $\varphi = 0.90$ and stationary spread 0.10.

The scenario node. Every what-if here is closed at the hall: return set points pinned, valves free and determined by the PI's row, and at a stop the actuator pinned again. The chiller's set point T_{sp} is pinned with its lag. The tank's move q_t is pinned or free as each row names. Two more nodes are named where they are used: the capacity node reads the capacity short at 670 kW, and the one-unit node pins T_{sp} at 14°C and r at 30.5°C.

The battery. The UPS battery is electrical, with $E_{\max} = 400\,\text{kWh}$, $P_{\max} = 300\,\text{kW}$ and round-trip efficiency $\eta_{\text{rt}} = 0.90$. It has no COP, its loss is all on the way in, and no self-discharge or wear model is stated. It is in this chapter because it shares the price with the tank, and for no other reason.

Figure 21.1 draws the plant. The hall's block on the left runs from the belief to the stance and the heat; the store's block on the right holds the tank's level and its move. Green is free, gray is pinned or supplied, dashed arrows are readings. The red box in the middle is the capacity row, the only row with an entry on both sides. The price above feeds both stores, and the battery sits beside the tank with no row to it.

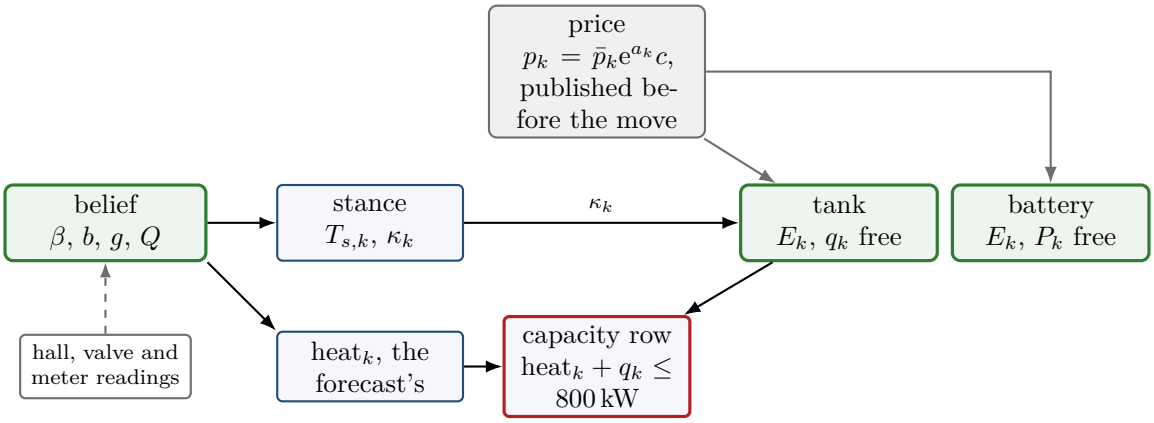


Figure 21.1: Stance and store. The belief fixes the stance alone, and the stance gives the store κ_k. The red capacity row is the only row that holds entries of both, and it binds only when charging or when capacity is short. The battery shares the price and no other row with the tank.

21.1.3 The control engineer's questions (step 1)

Table 21.1 lists the nine questions. The first column is the question in the engineer's words, the second the general form, the third its mathematical form on this plant, the fourth its category, the fifth the computation and the last the subsection.

21.1.4 The evidence as rows (step 2)

The record's rows are those of Chapter 18, read through 20:00 on the book's record, and are not written again. What is new is how they relate to the store's rows and which choices the store rests on.

Which row acts on which entry. With $*$ for an entry that can be nonzero, and columns for the hall's entries (the belief's β, b, g, Q), the level E_k and the move q_k, the pattern is

$$\mathbf{J} = \begin{bmatrix} * & 0 & 0 \\ 0 & * & * \\ 0 & * & 0 \\ * & 0 & * \end{bmatrix} \begin{matrix} \text{hall's readings and physics} \\ \text{tank's row (21.1)} \\ \text{level reading} \\ \text{capacity row, when it binds} \end{matrix} . \tag{21.3}$$

The first two lines have disjoint columns, so the hall's fit needs nothing of the store. Only the last line has both. Where the capacity row is slack the problem separates: the stance is solved from the belief alone, then the store is solved given κ_k. The price is a further row, on a_k alone, and it has no column in either block.

The rows of the price. With $\bar{p}_k$ the tariff and $c = \mathrm{e}^{-s^2/2}$ the factor that preserves the mean,

$$p_k = \bar{p}_k \, \mathrm{e}^{a_k} \, c, \qquad a_{k+1} = \varphi a_k + \varepsilon_k, \qquad \varepsilon_k \sim \mathcal{N}(0,\ s^2(1-\varphi^2)). \tag{21.4}$$

The modeling choices. Each enters here and is tested later.

- *The stance is separated from the store.* This buys two small problems for one large one, and it fails when the capacity row binds or when a charging plan warms the water; Section 21.3.8 measures the second.

Table 21.1: The questions asked of the two stores. The computation "scenario nodes across time" is the expected consequence read at pinned moves of time-indexed copies.

Control engineer's question	General form	Mathematical form	Category	Computation	Answered in
What is a kWh left at 07:00 worth?	expected consequence, (2.23)	$V(E)$, cyclic recursion, against stated values	design	expected consequence	21.3.1
Does the price band change the schedule?	certainty equivalence, (2.23)	$\mathbb{E}$[cost] linear in p at fixed q	design	expected consequence	21.3.2
What is the band worth with a response?	stochastic recursion, (2.23)	$J_k(E, a)$, policy against schedule and perfect knowledge	design	expected consequence	21.3.3
Price or COP, and what does the loss take?	what-ifs and a break-even, (2.23)	scenario nodes; $p_{\text{out}}\kappa_{\text{out}} > p_{\text{in}}\kappa_{\text{in}}/(\eta f^{g+1})$	design	scenario nodes across time	21.3.4
When does the tank serve the limit?	chance constraint on a store, (2.12)	$\Pr(\text{need}_k \le E_k \mid \mathbf{y}) \ge 0.95$	prediction under a threshold	posterior	21.3.5
What reserve does the UPS need, and is arbitrage worth it?	chance constraint, then expected consequence, (2.13), (2.23)	$\Pr(D \le E/P) \ge \alpha$; wear flip	design	expected consequence	21.3.6
What ties the two stores?	the coupling, stated, (2.23)	a sum of costs, no grid row here	design	scenario nodes across time	21.3.7
Does it hold on fresh records?	fresh records, (2.23)	paired net through the plant	validation	expected consequence	21.3.8
What next?	preposterior, means held, (4.73)	saving in \$/shift	sensor placement	preposterior	21.6

- *The window's price is published before the move.* Without it there is no response at all, and Section 21.3.3 prices its age.
- *One AR(1) on the price.* This buys a two-entry recursion. The forecast's log spread at leads 1, 12, 72 and 144 windows is 0.073, 0.130, 0.246 and 0.327, against the AR's stationary 0.10. The two disagree at long lead, so both are swept and 0.10 is the plant's scenario. No correlation with the load is stated.
- *An idle day between shifts, and tomorrow the same as today.* This buys a terminal value from the plant's own rows, and it fails if the day has its own tariff, which none is stated here. Section 21.3.1 prints the regret of the alternatives.
- *The guard.* 30 kW is held below capacity when charging, a stated choice; 20 kW is printed beside it in Section 21.3.8, and nothing here shows that 30 beats 20.
- *The stance's multiplier.* $z = 2.172$ at all windows, Chapter 18's, not recalibrated over 144 windows.
- *The outage.* Lognormal, median 5 minutes, log spread 1.0, the median's log uncertain by 0.5, ride-through load 300 kW. These are invented to be declared, no record stands behind them, and every battery reserve follows from them.

21.2 The posterior of the state (step 3)

Three objects serve the answers: the belief on the hall and the stance it gives, the price's conditional law, and the outage's one entry.

21.2.1 The belief and the stance

The belief is the pair $(\hat{\mathbf{z}}_k, \mathbf{\Sigma}_k)$ of the filter of Chapter 18, 23 entries, on the book's record at window 12. Table 21.2 gives four of them. Each row is an entry; the columns give its mean and spread at windows 12 and 72. The information matrix of the normalized rows is $\mathbf{J}^\mathsf{T}\mathbf{J}$, and with unnormalized rows it is $\mathbf{J}^\mathsf{T}\mathbf{W}\mathbf{J}$; the rows' weights are that chapter's. The unread racks' entries stay at their prior spread of 0.0800 in every window. Their block of $\mathbf{\Sigma}$ at window 12 with the hall's offset, in the order (β_{B3}, b_B), has variances 0.006400 and 0.009964 and a cross entry of 0.000000: no row reads both.

Table 21.2: Four entries of the belief. Means with spreads in parentheses.

Entry [unit]	Window 12	Window 72
β_{B3} [–]	0.5000 (0.0800)	0.5000 (0.0800)
β_{A5} [–]	0.4500 (0.0800)	0.4500 (0.0800)
b_A [K]	0.0296 (0.0999)	0.0334 (0.0428)
b_B [K]	−0.2688 (0.0998)	−0.1694 (0.0428)

The stance turns this pair into the store's price. An inlet i in hall h is a derived quantity, with return set point r_h, sensor offset b_h, airflow share β_i, total load $Q_{t,h}$ and the hall's summed airflow conductance C_h, and its gradient on (β_i, b_h, Q_h):

$$T_i(\mathbf{z}) = r_h - b_h - (1-\beta_i)\frac{Q_{t,h}}{C_h}, \qquad \mathbf{g}_i = \Big(\frac{Q_{t,h}}{C_h},\ -1,\ -\frac{1-\beta_i}{C_h}\Big)^\mathsf{T}, \tag{21.5}$$

the variance being $\mathbf{g}_i^\mathsf{T}\mathbf{\Sigma}\mathbf{g}_i$ on those entries, with the Q_h cross terms set aside and the load's own step variance added. The warmest supply the band allows, with the valve held at $u \le 0.95$

on the hall's conductance G_h, is

$$T_{s,k} = \min_{h} \Big[\min_{i \in h} \Big(27 - z \sqrt{\mathbf{g}_i^{\mathsf{T}} \boldsymbol{\Sigma} \mathbf{g}_i} + (1 - \hat{\beta}_i) \frac{Q_{t,h}}{C_h} \Big) - \frac{Q_{t,h}}{0.95\, G_h} \Big], \tag{21.6}$$

within the set-point range. This is the stance, Chapter 18's hold. Nothing of the store appears in it.

Table 21.3 gives the design's table from the belief at window 12 with the forecast's mean load, at $z = 2.172$. Each row is a window; the columns give the tariff, the ambient, the heat, the stance, its COP, and the product $p\kappa$, the store's price per kilowatt-hour of cooling. At window 13 the stance at $z = 0$, 1.645, 2.172 and 3.341 reads 17.078, 15.761, 15.339 and 14.317°C. The margin to the capacity row on the forecast's mean is 153 kW at 19:00 and 103 kW at 07:00.

Table 21.3: The stance over the shift, belief at window 12, forecast's mean load, $z = 2.172$.

Clock	Tariff [$/kWh]	Ambient [°C]	Heat [kW]	T_s [°C]	COP [–]	$p\kappa$ [$/kWh]
19:00	0.22	22.00	647.0	15.421	6.018	0.03656
20:55	0.22	20.09	655.0	15.263	6.547	0.03360
22:55	0.14	18.08	663.4	15.099	7.215	0.01940
00:55	0.07	16.58	671.8	14.935	7.788	0.00899
02:55	0.07	16.00	680.2	14.771	7.983	0.00877
05:25	0.10	16.77	690.7	14.565	7.524	0.01329
06:55	0.18	17.92	697.0	14.442	7.002	0.02571

21.2.2 The price's conditional law

Rows on the past are conditioned and rows on the future are averaged. Given the published factor a_k, the rows (21.4) give the law at lead ℓ,

$$a_{k+\ell} \mid a_k \sim \mathcal{N}\big(\varphi^{\ell} a_k,\ s^2(1 - \varphi^{2\ell})\big), \qquad \mathbb{E}[p_{k+\ell} \mid a_k] = \bar{p}_{k+\ell}\, c\, \exp\big(\varphi^{\ell} a_k + \tfrac{1}{2} s^2 (1 - \varphi^{2\ell})\big). \tag{21.7}$$

At $\ell \to \infty$ the expected price is the tariff, whatever a_k is. That is what c is for.

21.2.3 The outage's one entry

The battery's unknown is the log of the outage's median, μ, a one-entry state. It has a prior row and no reading, since no outage log exists: the term is $\mathcal{L}(\mu) = (\mu - \ln 5)^2 / (2 \cdot 0.5^2)$, with slope zero at $\hat{\mu} = \ln 5$ and curvature $\Lambda = 1/0.5^2 = 4$, so that $\sigma_\mu^2 = 0.25$. This is a declared prior and not a result. The predictive of the log duration adds the outage's own spread,

$$\ln D \sim \mathcal{N}\big(\hat{\mu},\ 1.0^2 + \sigma_\mu^2\big), \qquad s_D = \sqrt{1 + 0.25} = 1.1180. \tag{21.8}$$

21.3 The answers (step 4)

Each answer reads the stance's κ_k of equation (21.6), the price's law (21.7), the belief's draws or the outage's predictive (21.8). Dollars and percentages are this example's.

21.3.1 The value of a kilowatt-hour left at the end of the shift

Control engineer's question. What is a kWh left at 07:00 worth?

General form. The expected consequence (2.23) of a schedule, with what the shift leaves behind as part of its loss.

Specialized form. Equation (2.23) with the action the schedule $q_{1:144}$ and the loss the cost (21.2) less $V(E_{144})$, at the price's mean. The heat term is fixed by the stance, so the store's own recursion over the level is

$$J_k(E) = \min_q \Big[\bar{p}_k\,\kappa_k\,\Delta t\,q + J_{k+1}\big(fE + \Delta t(\eta q^+ - q^-)\big)\Big], \qquad J_{145}(E) = -V(E), \tag{21.9}$$

with q limited by the capacity row. The stated assumption is that an idle day follows and tomorrow's shift equals today's, so that the value is its own fixed point,

$$V(E) = -J_1(f^{144}E) + \text{constant}, \qquad f^{144} = 0.941753. \tag{21.10}$$

This buys a terminal value that is not chosen, and it fails if the day has a tariff of its own. The score at the cyclic value is circular as an argument for the cyclic rule. What it does give is the regret of the others if tomorrow is today.

Building blocks. κ_k and $p\kappa$ of Table 21.3; the tank's row (21.1); the capacity row; the references $p\kappa$ at 19:00 times the day factor, $p\kappa$ at 06:25 and the recharge price $0.07\kappa/\eta$ at κ at 07:00, 0.1428.

The plant's numbers. The fixed point is reached in 2 passes. Overnight charging is free of its start, since the night has hours enough, so among rules that end above empty the night forgets the terminal value, and the start's value depends only on the evening. The slope $\mathrm{d}V/\mathrm{d}E$ at $E = 50$, 150, 250, 350, 450 and 550 kWh is 0.03412, 0.03339, 0.03264, 0.03186, 0.01979 and 0.01929 \$/kWh, with the kink at about 430 kWh. That is 405 kWh the next evening, what two hours at 200 kW can use. The references are 0.03443 for 19:00 times the day factor, 0.02505 for 06:25 and 0.01052 \$/kWh for the recharge. **Chapter 18's recharge value is three times too low below the kink.** That recharge value uses κ at 07:00, 0.1428, and charging happens near 0.1253. There the same price is 0.00923 \$/kWh and the slope stands 3.70 times above it, against 3.24. Figure 21.2 draws V. The horizontal axis is the level at 07:00, the vertical axis the value above an empty tank; the solid curve is the cyclic value, the dotted lines are flat values at the three references, and the vertical line marks the kink.

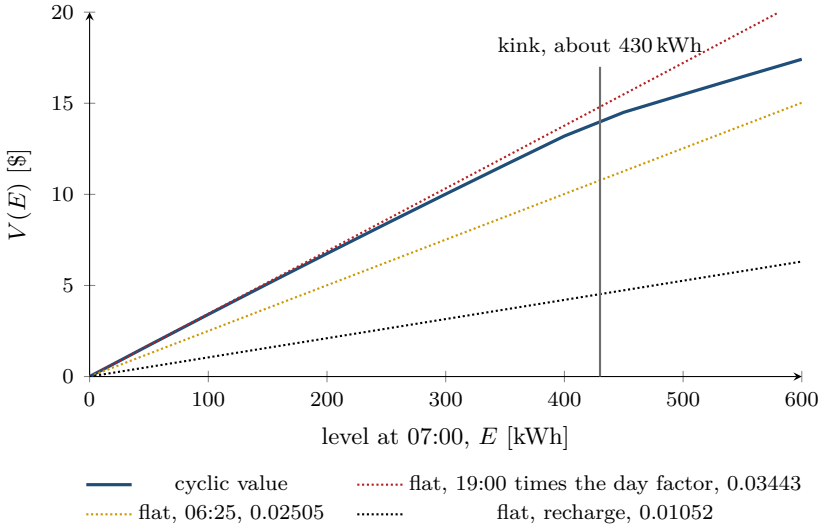


Figure 21.2: The value of what is left at 07:00, equation (21.10), zero at empty. The flat lines are vE at the stated v. Above the kink a further kilowatt-hour cannot be used in next evening's two dearest hours.

Table 21.4 sets the rules side by side, with no tank use at 135.737 $ as a labelled reference. Each row is a rule; the columns give the electricity, the net scored at the cyclic value, the level at 07:00 and the kilowatt-hours out between 06:00 and 07:00.

Table 21.4: Five terminal rules on the shift's schedule. No tank use, a labelled reference, is 135.737 $.

Terminal rule	Electricity [$]	Net at cyclic value [$]	E end [kWh]	Out 06:00–07:00 [kWh]
Nothing left has value	121.220	121.220	0.0	200.0
Recharge at 0.07, Chapter 18's	123.160	116.543	195.9	200.0
Pinned at the start, 300	124.544	114.529	300.0	200.0
Pinned full, 600	130.968	113.553	600.0	0.0
Cyclic recursion	126.670	112.577	430.0	167.3

Figure 21.3 draws what each rule does. In the upper panel the horizontal axis is the clock and the vertical axis the level; the shaded bands are the tariff's. The lower panel is the move under the cyclic rule. The schedule under the cyclic value is out from 19:00 to 20:30, 299 kWh; in from 23:45 to 06:00, 643 kWh; and out from 06:05 to 07:00, 167 kWh. Every rule uses the first evening the same way, 299 kWh out by 20:30. **Among the rules that end above empty, on the first shift, the terminal rule decides the morning only.** "Nothing left has value" also changes the night: it charges only 427 kWh, from 01:15 to 04:50, against 640 to 646 kWh from 23:45. And the evening depends on the start: from the periodic state it is 403 kWh out to 21:05, against 299. The rules with the smaller values empty the tank into the 06:00 price, and pinned full forgoes the morning altogether.

The periodic state under the cyclic value starts the shift at 405.0 and ends it at 430.0 kWh. Its electricity is 123.175 $, 12.563 $ below no tank use (9.26 %). The first shift from 300 kWh, scored at the same value, saves 13.711 $. A flat value of 0 to 0.0088 empties the tank; 0.016 to 0.0243 ends at 397.5 kWh, net 112.746 $; 0.03 and up ends full. So the regret, if tomorrow is today, is 3.966 $ for Chapter 18's rule and 1.952 $ for the sheet's 300. The first is sensitive to which κ prices the recharge: at the night's 0.00923 $/kWh the tank empties, net 121.220 $, and the regret is 8.643 $.

The hour's plan again. Chapter 18's plan is now run with only its terminal changed, each row scored at the cyclic value and the multiplier held. Table 21.5 gives it. Each row is a plan's start and terminal; the columns give the first and last moves, the level at the hour's end, the electricity, the net and the joint with its standard error.

Table 21.5: Chapter 18's hour's plan under the store's own value, multiplier held: 2.172 at windows 12 and 126, 1.906 at window 72 under Chapter 18's stated tie-break (2.055 without it).

Now [–]	Terminal	q_t first, last [kW]	E end [kWh]	Electricity [$]	Net [$]	Joint (se) [–]
12	pinned	0, 0	297.4	21.846	16.030	0.9506 (0.0031)
12	Chapter 18's	−200, −200	97.8	14.953	13.173	0.9506 (0.0031)
12	this recursion	−200, −200	97.8	14.953	13.173	0.9506 (0.0031)
72	Chapter 18's	0, 0	290.6	6.104	0.445	0.9534 (0.0030)
72	this recursion	0, +123.1	350.0	6.657	0.424	0.9534 (0.0030)
126	Chapter 18's	−8, −200	100.0	10.305	6.801	0.9582 (0.0028)
126	this recursion	+49.2, +103.3	371.9	16.442	3.961	0.9732 (0.0023)

The evening's move stands and its worth shrinks: 6.893 $ of electricity against the pinned tank is 2.857 $ net of what the tank gives up. At 05:30 the move reverses, since that chapter's

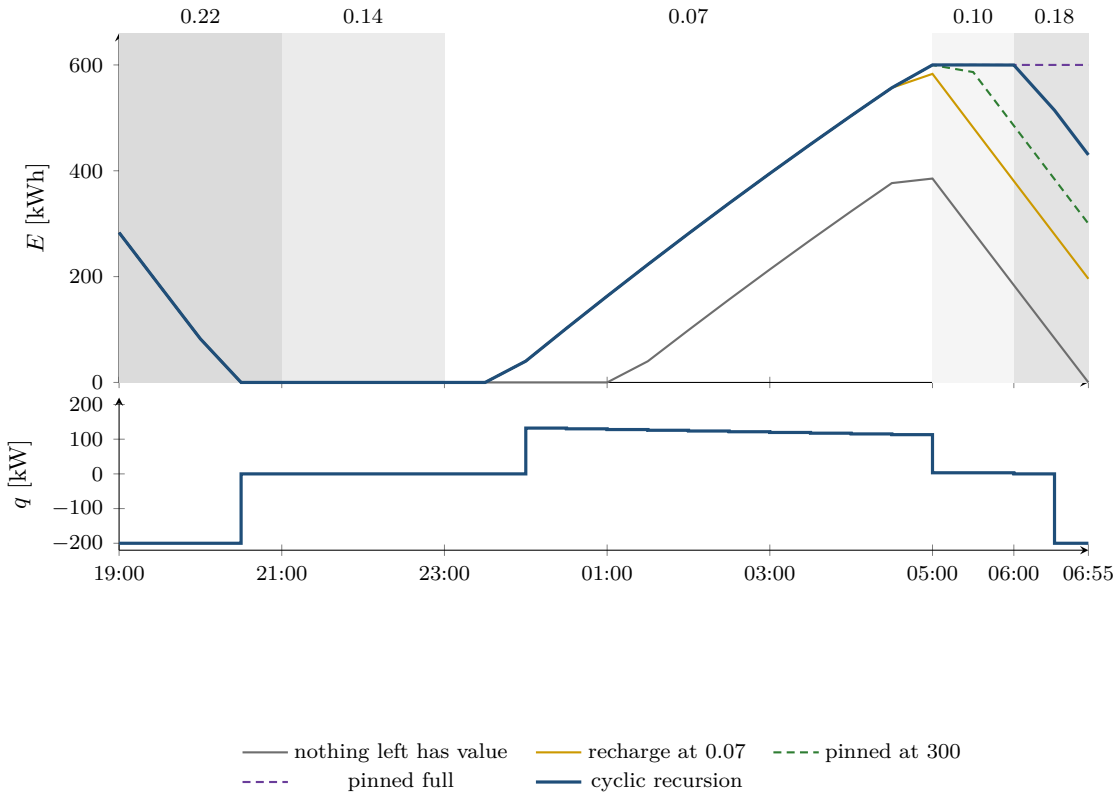


Figure 21.3: The level under five terminal rules, at the end of every sixth window, and the cyclic rule's move below, a held value at each sixth window. Shaded bands and the numbers above them: the tariff in \$/kWh, unshaded at 0.07.

plan discharges and this one charges. Chapter 18's 21.85 and 14.97 reappear as 21.846 and 14.953 on these draws.

21.3.2 Certainty equivalence in price

Control engineer's question. Does the price band change the schedule?

General form. The expected consequence (2.23) over the price's law, with the action fixed before any price is read.

Specialized form. Equation (2.23) with the action a fixed $q_{1:144}$ and the loss the cost (21.2). It is linear in p_k, and q_k, κ_k and the heat are not functions of p_k, so

$$\mathbb{E}[\text{cost}] = \sum_k \mathbb{E}[p_k]\,\kappa_k\,\Delta t\,(\text{heat}_k + q_k). \tag{21.11}$$

The band is not an argument of the minimization, and the schedule from the means is the schedule under the band. Since κ in equation (21.2) is linear in the ambient, the ambient's band enters by its mean as well. Three things would break this: a limit with the load's band in it (Section 21.3.5), a price correlated with the load (none in this plant), and a move that reads the price (Section 21.3.3).

Building blocks. κ_k from the stance; the price's law (21.7) at its mean; the schedule of the last answer.

The plant's numbers. Table 21.6 tests it with sample-average schedules: the mean of the training draws is what such a schedule sees, and it is scored on 20 000 other draws with the mean-preserving band. Each row is a band; the columns give the L1 gap to the mean schedule and the cost over it, at 100 and 2000 training draws. The quadrature for the ambient at 1 K leaves a gap of zero, since the function is linear.

Table 21.6: Sample-average schedules against the mean schedule, evaluated on 20 000 draws, $\varphi = 0.90$.

Band [–]	L1 gap, 100 [kWh]	L1 gap, 2000 [kWh]	Cost over, 100 (se) [$]	Cost over, 2000 (se) [$]
0.05	0.0	0.0	0.000 (0.000)	0.000 (0.000)
0.10	20.3	0.0	0.001 (0.000)	0.000 (0.000)
0.33	253.0	1.1	0.107 (0.003)	0.000 (0.000)

The band moves no schedule; a sample of it does, and that costs. So the sheet's "derate by 10 %" has no row to act on: the mean schedule's expected net stands at 112.610 $ at the plant's band, against 112.593 and 112.692 $ at 0.05 and 0.33, and those three are a spread of draws around one number, each with its standard error of 0.016, 0.033 and 0.110 $.

21.3.3 The value of the mean and of the response

Control engineer's question. What is the band worth with a response?

General form. The expected consequence (2.23) with the action a policy, a move for each level and each price read.

Specialized form. Equation (2.23) with the action the policy on (E, a) and the price's law (21.7) at lead one, by the recursion

$$J_k(E,a) = \min_q \Big[\bar{p}_k \, \mathrm{e}^{a} c \, \kappa_k \, \Delta t \, q + \mathbb{E}\big[J_{k+1}\big(fE + \Delta t(\eta q^{+} - q^{-}),\ a'\big) \mid a\big]\Big], \qquad J_{145}(E,a) = -V(E), \tag{21.12}$$

on grids of levels, moves and factors, whose error is stated in Section 21.4. With the band at zero it is equation (21.9). Three costs are compared, all paired on 4000 price paths: the fixed schedule, the policy, and perfect knowledge of the same draw on 400 of them, the last a physical limit that no policy reaches. The value of recourse is the schedule's cost less the policy's. The fixed schedule and the mean plan re-solved are this problem's restrictions.

Building blocks. κ_k; V of equation (21.10); the rows (21.4); the capacity row as the upper limit on q.

The plant's numbers. Table 21.7 gives the sweep. Each row is a band and a persistence; the columns give the value of the mean against a labelled no-tank reference, the value of the response, the recursion over perfect knowledge and the mean schedule over perfect knowledge. Figure 21.4 draws the rows at $\varphi = 0.90$, with the battery's at the plant's band: the horizontal axis is the band, the vertical axis in dollars on a log scale, curves for the tank and open marks for the battery. The fair reading is the two right-hand columns.

Table 21.7: The value of the mean and of the response, paired, 4000 paths; perfect knowledge on 400.

Band [–]	φ [–]	Value of the mean [$]	Of the response (se) [$]	Recursion over perfect (se) [$]	Mean schedule over perfect (se) [$]
0.05	0.90	13.704	0.011 (0.001)	0.099 (0.003)	0.110 (0.004)
0.10	0.90	13.697	0.135 (0.003)	0.301 (0.007)	0.433 (0.011)
0.20	0.90	13.683	0.579 (0.007)	0.729 (0.017)	1.296 (0.027)
0.33	0.90	13.664	1.465 (0.016)	1.325 (0.037)	2.772 (0.063)
0.10	0.99	13.690	−0.005 (0.000)	0.069 (0.003)	0.063 (0.003)
0.33	0.99	13.649	0.255 (0.005)	0.598 (0.019)	0.840 (0.026)

The mean is worth about 100 **times the response at the plant's band.** The response grows faster than the band, from 0.135 to 1.465 $ for a band 3.3 times as wide, and it needs a price that moves within the window a store can use: $\varphi = 0.99$ at the same marginal band is worth about nothing, inside the grid's error of 0.011 $, and rows below that are not separated from zero. A mean plan re-solved every three windows stands, on 100 paths, −0.090 (se 0.020) against the schedule and +0.058 $ (se 0.009) against the recursion, and on 200 paths −0.109 (se 0.014) and +0.029 $ (se 0.007): 79.2 % **of the response without a recursion** (se 5.2) on the 200, about 60 % on the first 100 alone, at that re-solve rate.

Table 21.8 shows the policy as thresholds. Each row is a window and a level; the columns give the move at each published factor. At 19:25 and 100 kWh it charges 151 kW at $a = -0.20$, holds at −0.10 and −0.05, and discharges from 0 up. At 00:55 and 500 kWh it charges to −0.10, holds at −0.05 and 0, and **discharges at** +0.05 **at the tariff's floor**. At 04:15 it charges whatever the price: the capacity row leaves the night almost no choice of when, and that is one reason the tank's response is small. With the row lifted on charging the response is 0.289 $ (se 0.004), 2.14 times as much, so the row is a factor of about two; the larger factor is the COP of about seven, which Section 21.3.6 shows on the battery. The recursion moves 493.0 kWh out, against the schedule's 466.1, but **the response is mostly re-timing, not extra cycling**. A wear of 0.001 $/kWh in the recursion alone holds the throughput at 465.9 kWh, and the response

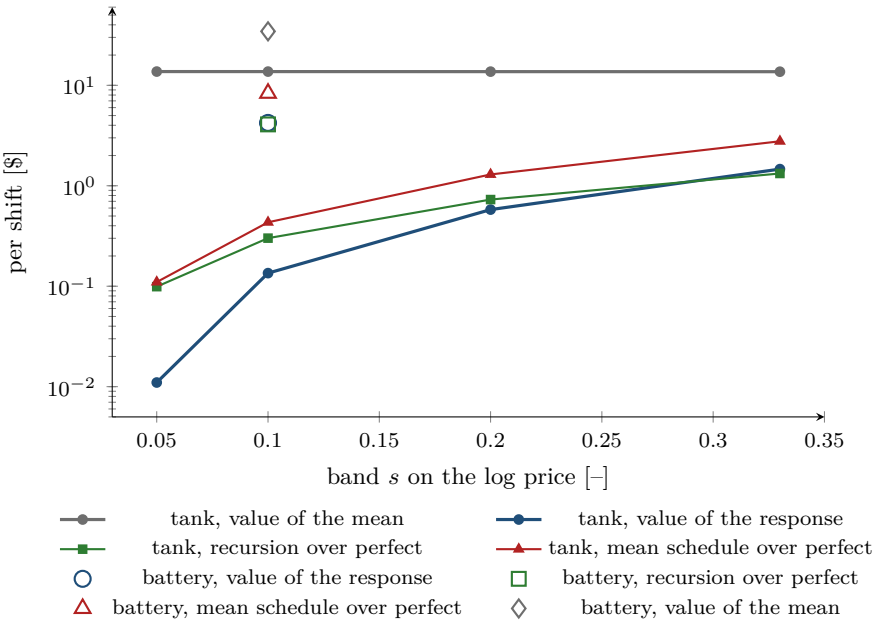


Figure 21.4: What the mean and the response are worth, against the band, $\varphi = 0.90$. Curves: the tank. Open marks: the battery at the plant's band and its R_{95} reserve of 157.3 kWh.

read without the wear is 0.126 $ (se 0.003), 92.7 % of the 0.135. It ends at 430.7 kWh, spread 12.3, against the schedule's 430.0.

Table 21.8: The policy of equation (21.12) at the plant's band: the move in kW by the published log factor a.

Clock	E [kWh]	−0.20	−0.10	−0.05	0	+0.05	+0.10	+0.20
19:25	100	151	0	0	−200	−200	−200	−200
19:25	300	−30	−200	−200	−200	−200	−200	−200
00:55	300	128	128	128	128	0	−30	−200
00:55	500	128	128	0	0	−200	−200	−200
04:15	500	114	114	114	114	114	114	114
06:25	500	100	0	0	0	−200	−200	−200

The price's age matters as well, and how an old price is used matters too. Read now, the net is 112.434 $. Feeding the old factor's conditional mean to the policy built for a price read now, one window old is 0.034 more (se 0.001) and one hour old 0.200 more (se 0.003), above the mean schedule's 0.135 $. That excess belongs to the policy and not to the age: the same policy fed $a = 0$ always, reading nothing, stands 0.095 $ (se 0.002) above the schedule. A recursion built for the age, with the old factor as its state and the price at its conditional mean, is 0.032 (se 0.001) over read-now at one window and 0.145 (se 0.003) at one hour, the latter 0.009 $ (se 0.001) above the schedule and inside the grid's error. **A reading one hour old, properly used, is worth about nothing, not less than nothing.**

21.3.4 Price, COP and the round trip

Control engineer's question. Price or COP, and what does the loss take?

General form. The expected consequence (2.23), read at scenario nodes across time, and at one pair of windows as a break-even.

Specialized form. Equation (2.23) with the action the schedule at a node that flattens the price, κ, both or neither, the hall's stance pinned, the tank free, and the end pinned at 300 kWh so that nothing is stored across the comparison. For one pair, a kilowatt-hour out at window k needs $1/(\eta f^{g+1})$ in at window j, with g windows held between them. The tank's row (21.1) and the cost (21.2) then give

$$\text{pays} \iff p_{\text{out}}\kappa_{\text{out}} > \frac{p_{\text{in}}\kappa_{\text{in}}}{\eta f^{g+1}} \iff \frac{p_{\text{out}}}{p_{\text{in}}} > \frac{\kappa_{\text{in}}/\kappa_{\text{out}}}{\eta f^{g+1}}. \tag{21.13}$$

Building blocks. κ at 19:00, 03:00 and 06:25 from the stance: 0.16618, 0.12532 and 0.13917. Time means: tariff 0.11833 \$/kWh, κ 0.13789 (COP 7.252).

The plant's numbers. Table 21.9 gives the four nodes, and each row names one. Their order-free split of 11.487 \$ gives **price** 9.846 \$ **(85.7 %) and COP** 1.641 \$ **(14.3 %)**, with an interaction of 0.774 \$. Overnight has a second reason, a small one, and it pays alone. Table 21.10 tests the break-even. Each row is a case; the columns give the two values of κ, the windows held, the ratio from equation (21.13) and the ratio found by bisecting the rows themselves.

Table 21.9: Scenario nodes, tank free, end pinned at 300 kWh; saving over the same node with no tank use.

Node	Saving [\$]
Price flat, COP flat	−0.293
Tariff, COP flat	9.166
Price flat, COP's ambient	0.960
Both	11.193

Table 21.10: The round trip's break-even price ratio, by equation (21.13) and by the rows.

Case	κ_{in} [–]	κ_{out} [–]	Windows held [–]	Derived [–]	By bisection [–]
Same COP, next window	0.14000	0.14000	0	1.05307	1.05307
Same COP, 4 h	0.14000	0.14000	47	1.07390	1.07390
Same COP, 12 h	0.14000	0.14000	143	1.11774	1.11774
03:00 in, 06:25 out	0.12532	0.13917	40	0.96421	0.96421
03:00 in, 19:00 next evening	0.12532	0.16618	191	0.85992	0.85992

With COP the break-even price ratio falls below one. At 03:00 charging pays at a price 3.7 % above 06:25's and 16 % above next evening's, losses included. Every tariff pair tested pays, the thinnest being 21:00–23:00 into 06:00–07:00, ratio 1.286 against 1.143, a margin of 0.00278 \$/kWh. The hall's band sits in these rows too: κ at 19:00 over κ at 03:00 is 1.375 at $z = 0$ and 1.326 at $z = 2.172$, and the no-tank electricity is 121.875 and 135.737 \$.

21.3.5 The band on the load and the reserve

Control engineer's question. When does the tank serve the limit?

General form. A threshold on a derived quantity under the posterior, equation (2.12), here held on a store.

Specialized form. Equation (2.12) with f the energy the tank must still hold at the end of window k to serve the deficit to 07:00, for capacity C. It is a function of the load's path, backward from the end,

$$\text{need}_k = \frac{\text{need}_{k+1} + \Delta t\,(\text{heat}_{k+1} - C)^+}{f}, \qquad \text{need}_{144} = 0, \tag{21.14}$$

and the chance constraint on the store, with $R_{95,k}$ the quantile, is

$$\Pr(\text{need}_k \le E_k \mid \mathbf{y}) \ge 0.95 \iff E_k \ge R_{95,k}. \tag{21.15}$$

The load's paths are draws of the belief at window 12 with the forecast's band. The quantile is read directly, since a deficit is a tail on a kink and a normal on the need would miss it. The capacity node is a stated what-if, closed at the hall, with T_{sp} and r set by the stance and q free. Once the floor (21.15) is in the store's rows, the band has re-entered the schedule, and certainty equivalence no longer applies.

Building blocks. $\hat{\mathbf{z}}$, $\boldsymbol{\Sigma}$ at window 12 and 72; the forecast's band; heat_k; the recursion (21.14); the tank's level at 20:00, 298.5 kWh, with no earlier move.

The plant's numbers. From the fit at window 12, the heat at 07:00 has mean 696.8 and spread 15.3 kW, and the largest heat over the shift has a 95th percentile of 724.3 kW. The true ramp's mean heat at 07:00 is 707.0 kW. Table 21.11 sweeps the capacity. Each row is a capacity; the columns give the probability of a deficit at some window, the need at 20:00 by its mean and spread, its 95th and 99th percentiles, the need at the true ramp with no noise, and what the tank has free at R_{95}.

Table 21.11: The band on load. The tank's level at 20:00 is 298.5 kWh. The true ramp is the stated truth's, noiseless.

Capacity [kW]	Pr(deficit) [–]	Need, mean [kWh]	sd [kWh]	R_{95}, R_{99} [kWh]	True ramp [kWh]	Free at R_{95} [kWh]
800	0.0000	0.0	0.0	0.0, 0.0	0.0	298.5
720	0.0906	0.4	2.7	1.0, 11.3	0.0	297.5
700	0.5130	7.2	15.5	40.1, 72.6	5.4	258.4
680	0.9268	49.9	47.1	142.7, 190.2	77.0	155.8
670	0.9870	99.8	65.3	218.2, 271.6	143.5	80.3
660	0.9990	174.7	79.0	309.9, 365.7	230.3	−11.4
640	1.0000	383.7	85.1	524.5, 581.8	449.8	−226.0

The answer to "when": at this plant's 800 kW**, never; the limit takes the tank below about** 720 kW, and R_{95} reaches the tank's level at 661.17 kW. Mean need is 46 % of R_{95} at the node, because the deficit is a tail. Figure 21.5 draws both views. The left panel has the capacity on its horizontal axis and the need at 20:00 on its vertical axis, with the tank's level as a horizontal line. The right panel has the clock on its horizontal axis at 670 kW, with the floor R_{95}, the level under each plan on the rows, and the true ramp's need as a dotted marker. A later fit narrows it: at 01:00 the spread of the need is 59.0 from 20:00's belief and 26.0 kWh from 01:00's, and R_{95} is 198.9 against 135.5 kWh.

On the rows at 670 kW, the mean need alone has net 126.637 \$ and 232 kWh out in the evening. The R_{50} floor has 126.703 \$, and the R_{95} floor has 127.145 \$ with 81 kWh out. **Holding** R_{95} **costs** 5.368 \$ **of electricity and** 151.3 kWh **of evening, and** 0.508 \$ **net** once the 143.1 kWh left at 07:00 is credited at the 800 kW plant's value, which assumes that capacity returns tomorrow. That assumption is needed, since this node cannot be repeated: on the forecast's

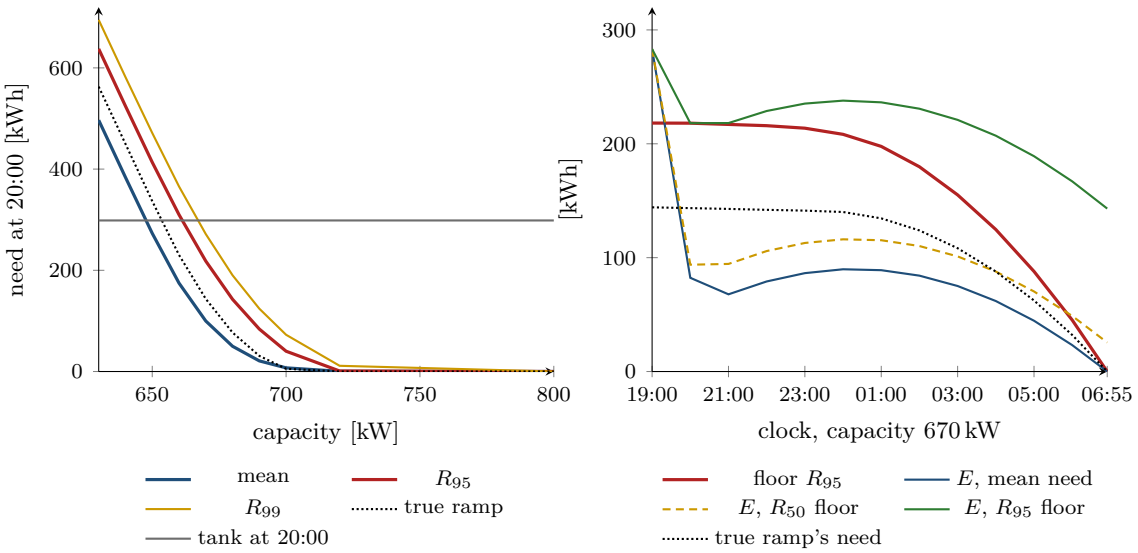


Figure 21.5: Left: the need at 20:00 against capacity, belief at window 12. Right: at 670 kW, the floor and the level under three plans on the rows, at every twelfth window. Dotted: the stated truth's ramp, noiseless.

mean the deficit is 88.0 kWh against 60.8 kWh of recharge room, no plan remains by the tenth shift from full, and at the true ramp the figures are 137.5 and 50.8 kWh and the fifth shift. The nets of Table 21.12 rest on the same value. Table 21.12 runs fresh records at the stated truth, 200 records, with the record's reading of its own load setting the deficit and a guard of 5 kW. Each row is a plan; the columns give the net, the probability that all inlets stay at or below 27°C, the mean largest inlet, the probability that the tank empties before 07:00, and the 95th percentile of the largest T_s.

Table 21.12: Fresh records at the capacity node, 670 kW, 200 records, stated truth.

Plan	Net (se) [$]	Pr(all ≤ 27) [–]	Mean largest inlet [°C]	Pr(tank empty) [–]	T_s P95 [°C]
Mean plan, no floor	122.327 (0.351)	0.000	31.676	1.000	20.246
R_{50} floor	125.804 (0.359)	0.000	30.453	1.000	19.118
R_{95} floor	135.525 (0.391)	1.000	26.540	0.170	15.548
Deficit only, no arbitrage	136.102 (0.402)	1.000	26.539	0.000	15.548

At this truth a mean plan spends what the limit needs, and every record loses its inlets. Its net is lower because it failed: warm water is cheap, and that row is not comparable. The true ramp sits above the forecast's mean by 0.625 spread per hall, but the failure does not rest on that. With the floors and plans unchanged and the true ramp at the forecast's mean of 25 kW per hall, the mean plan keeps no record (Pr = 0.000, mean largest inlet 28.950°C) and R_{50} keeps 0.075. At 18 kW the mean plan keeps 0.995 (26.542°C) and R_{50} all of them. R_{95} keeps all of them at all three, with Pr(tank empty) 0.170, 0.000 and 0.000. The floor is built at $C = 670$ kW while the policy draws the deficit at $C - 5$ kW: on the same draws, Pr(need at 665 ≤ R_{95}) = 0.868, not 0.95, and R_{95} at 665 kW is 262.2 kWh. So the assurance read at the plant's actual draw is about 0.87. Against no arbitrage, R_{95} is −0.577 $ (se 0.023); against perfect knowledge on the rows, 133.689 $, it is +1.837 $ (se 0.032). It empties in 17 % of the records at this truth, late enough that no inlet is lost. That is not coverage, since it is one truth, and no value is placed on a breach here; the 0.95 is stated, and a breach's price belongs to Chapter 17.

The filter at an empty tank. Chapter 20 left this fix here. The tank's row in the filter subtracts the move asked, and an empty tank does not deliver it. On the mean plan's records at the node, which all empty the tank, the belief's largest error averages 107.1 kWh with a stated spread of 0.417 kWh, the two orders that chapter found. Limiting the asked move by the level believed takes it to 8.3 kWh and changes neither that row's net nor its inlets. A loss the row does not know about, as in that chapter, needs a walk on the row as well: 2 kWh per window raises the reading's share from 0.007 to 0.328 and the stated spread to 2.863 kWh, at 0.016 \$ (se 0.001) on the recursion at the plant's capacity. The records in this chapter keep Chapter 18's filter, and on the R_{95} row the limit leaves the net at 135.525 \$ and the walk moves it to 135.532 \$.

The reserve Chapter 20 asked about. Two values have been placed on what the tank holds: V, at arbitrage prices, and the floor, against the load's band. Neither prices a lost unit. What can be stated is the cost of holding for one. At 800 kW the schedule stands at 82.3 kWh at 20:00, 24.7 minutes at 200 kW. A floor all shift of 50, 100 or 194 kWh costs 1.257, 2.544 or 5.042 \$ net per shift, the last being what the 58.2 minutes of the next paragraph deliver. Whether that is worth paying needs a breach's price, which belongs to Chapter 17.

One unit. At 400 kW against 647 kW of heat and 200 kW from the tank, no tank move holds the hall. With T_{sp} pinned at 14°C, r at 30.5°C, and 50 records, the tank idle has its first inlet over 27°C at 9.4 minutes (spread 1.6) and a valve at its stop at 10.0. With the tank out at 200 kW, the first inlet is at 58.2 minutes (spread 3.5) and the stop at 26.6. **The tank buys** 48.8 **minutes, not a hold.** At 200 kW the 300 kWh would last 90 minutes.

21.3.6 The UPS battery

Control engineer's question. What reserve does the UPS need, and is arbitrage worth it?

General form. A predictive tail, equation (2.13), for the reserve, and the expected consequence (2.23) for what it costs.

Specialized form. Equation (2.13) with $y_{\text{new}} = D$ under the predictive (21.8), read at the ride-through load $P = 300$ kW. The stated assumptions are the lognormal form and the median; both are declared, and a site's log would replace them. Assurance α per outage needs

$$\Pr(D \le E/P \mid \mathbf{y}) \ge \alpha \iff R(\alpha) = \frac{P}{60}\exp\!\big(\hat{\mu} + s_D\,\Phi^{-1}(\alpha)\big), \tag{21.16}$$

in kWh with D in minutes. Then equation (2.23) with the action the reserve and the loss the cash of one cycle at 0.22 out and 0.07 in, ending full. The loss being all on the way in, and $U = 400 - R$ the usable energy,

$$\text{saving} = U\Big(0.22 - \frac{0.07}{\eta_{\text{rt}}}\Big) = U \times 0.14222 \ \$/\text{kWh}. \tag{21.17}$$

A wear cost w per kilowatt-hour out enters the same line, and the verdict flips at

$$w^{*} = p_{\text{out}} - \frac{p_{\text{in}}}{\eta_{\text{rt}}}. \tag{21.18}$$

The recursion (21.12) serves here too, with $\kappa = 1$, no heat and a penalty of 1 \$/kWh for not ending full.

Building blocks. $\hat{\mu}$ and $s_D = 1.1180$ of equation (21.8); Φ^{-1}; the price's rows; the battery's constants.

The plant's numbers. Table 21.13 gives the reserve by assurance, and each row is an assurance. **The stated** 120 kWh **rides** 24.0 **minutes, an assurance of** 0.9197 **on this distribution**, and full it is 0.9934. Both depend wholly on the declared numbers. Table 21.14 gives the saving by reserve. Each row is a reserve; the columns give the assurance, the saving from equation (21.17), which the rows match to the printed digit, and the time-averaged probability over the 12 h shift that an outage outlasts the store, the idle day at full not counted, under two charging times. The flat 0.07 band and no self-discharge make every charging hour from 23:00 to 05:00 cost the same, so the program's choice, the band's end, is a tie. **The tie is broken here by name: charge from 23:00, which halves the exposure at no cost.** A full battery with no arbitrage, a labelled reference, reads 0.0066. Figure 21.6 draws the same trade: on the left the assurance on the horizontal axis and the saving on the vertical, with 120 kWh marked; on the right the wear cost on the horizontal axis with the saving on the means and the recursion's on the vertical, the recursion's on 1000 paths.

Table 21.13: The reserve by assurance per outage, on the declared distribution, at 300 kW.

Assurance [–]	Duration [min]	Reserve [kWh]
0.90	21.0	104.8
0.95	31.5	157.3
0.99	67.4	336.9
0.999	158.3	791.4, above the 400

Table 21.14: The saving by reserve, one cycle, end full, tariff means.

Reserve [kWh]	Assurance [–]	Saving [$/shift]	Pr(outlasts), program's tie, band's end [–]	Pr(outlasts), from 23:00 [–]
0	0.0000	56.889	0.6265	0.2528
60	0.7832	48.356	0.1474	0.0643
104.8	0.9000	41.989	0.0721	0.0339
120	0.9197	39.822	0.0590	0.0285
157.3	0.9500	34.524	0.0384	0.0199
250	0.9803	21.333	0.0168	0.0109
336.9	0.9900	8.973	0.0094	0.0077

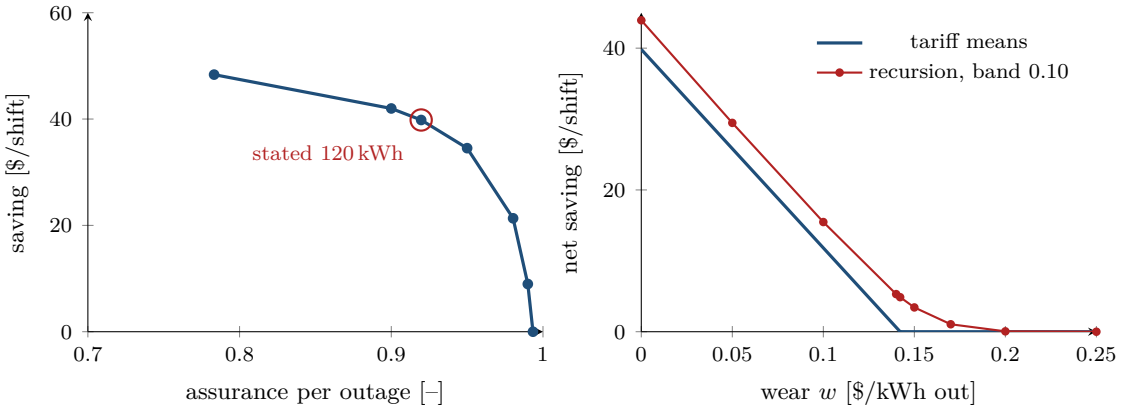


Figure 21.6: The battery. Left: saving against assurance, on the declared outage distribution, one cycle per shift; the last mark is a full battery with no arbitrage. Right: the wear flip at the stated reserve, with the response read.

Arbitrage clears the loss by a wide margin here, a ratio of 3.143 **against** $1/\eta_{\mathrm{rt}} = 1.111$**, before the wear figure and the reserve, which decide it.** Each kilowatt-hour of reserve costs 0.14222 \$/shift, so moving from 120 to 157.3 kWh costs 5.298 \$. The terminal rule returns as well. A rule that values nothing left, or a recharge, gives 90.222 \$ with a morning cycle and ends at the reserve; the cyclic value ends full and gives 39.822 \$, net at the cyclic value 101.422 against 90.222 \$. **The morning cycle is the same mistake as the tank's**, 11.200 \$ here.

With the response, at the 157.3 kWh reserve and ending full, the mean schedule earns 34.461 \$, the recursion 38.684 \$ and perfect knowledge 42.586 \$. **The response is** 4.223 \$ **(se** 0.056**),** 31 **times the tank's**, since the battery's price is not scaled by a COP and it can pick one hour of six. Its power counts too: the response is 3.137 \$ (se 0.039) at 200 kW and 1.416 \$ (se 0.017) at 120 kW. The recursion over perfect is 4.083 (se 0.099) and the schedule over perfect 8.337 \$ (se 0.178).

Equation (21.18) gives w^* of 0.14222 for 0.22 against 0.07, 0.10222 for 0.18, 0.06222 for 0.14, and 0.06444 \$/kWh for 0.22 against 0.14. Bisecting the rows gives 0.14222. **On tariff means the verdict flips at** 0.14222 \$/kWh **out; with the response read it survives there** (+4.871 \$), and it is gone by 0.25 (+0.052, se 0.012, at 0.20; 0.000 at 0.25). No wear figure is stated. This is the number to compare one against.

21.3.7 What ties the two stores

Control engineer's question. What ties the two stores?

General form. The expected consequence (2.23) over both moves, as a sum.

Specialized form. Equation (2.23) with the action the two schedules and the loss the sum of their costs. The pattern (21.3) has no row with both columns: this plant has no grid row, so the schedules separate by construction, on the means and draw by draw. No joint solve is shown, since none would show anything. The sum's net is 78.053 \$, the tank's electricity less value and the battery's cash, and it is a sum.

Building blocks. The two schedules of the earlier answers; κ at 19:00 and at night.

The plant's numbers. What the separation hides is that both stores charge at the tariff's floor. The chiller alone reads 85.0 to 107.5 kW electric with no tank and no battery. Chiller plus battery reads 404.3 kW at 04:55 and −227.2 kW at 19:40; breaking the battery's tie toward 23:00 moves the peak to 23:45 and leaves it at 404.2 kW. Chapters 29 and 30 hold the grid limit, and nothing here says that it is absent there. Per kilowatt-hour stored, at 0.22 against 0.07, the battery earns 0.14222 \$ and the tank 0.02733 \$, COP 6.018 in the evening and 7.980 at night, a ratio of 5.205; with the 192 windows' standing loss the tank's is 0.02656 \$, a ratio of 5.356. **A thermal kilowatt-hour is a fifth of an electrical one here, because it is priced through COP.**

21.3.8 Fresh records

Control engineer's question. Does it hold on fresh records?

General form. The expected consequence (2.23), estimated by paired records through the plant.

Specialized form. Equation (2.23) with the action each plan and the expectation replaced by the mean over 200 records, seed 2103, with the same exogenous paths and noise per plan. Each record's own filter sets its stance. The same charge limit table, from the forecast's mean heat with the guard of 30 kW, serves every plan, so that like is compared with like. Net is electricity less the cyclic value of what is left. These records' price is the tariff times e^{a} with no c, a mean factor of 1.00547 on them against $e^{s^2/2} = 1.00501$; the recursion on them uses $c = 1$, so the response is consistent, while electricity and net carry that factor and V does not.

Building blocks. The stance (21.6) from each record's filter; the policy of equation (21.12); V.

The plant's numbers. Table 21.15 gives them. Each row is a plan; the columns give the electricity, the net, the excess over perfect knowledge, the mean level at the end, and $T_s - T_{sp}$ from 23:00 to 06:00. Paired, the response is 0.080 $ **(se** 0.014**)** against 0.088 $ (se 0.003) on the rows at the same guard, and the mean is worth 13.842 $ (se 0.071). A price one window old loses 0.022 $ (se 0.003). The battery on these prices, on its own row, has cash 34.292 $ under the schedule, 38.628 $ under the recursion and 42.700 $ under perfect knowledge, a response of 4.336 $ (se 0.270). The store's rows against the plant, on these records, differ by 0.099 $ (0.07 %) of electricity, spread 0.132 $.

Table 21.15: Fresh records, 200, paired, at the plant's capacity and band.

Plan	Electricity (se) [$]	Net (se) [$]	Over perfect (se) [$]	E end [kWh]	$T_s - T_{sp}$, night [K]
No tank use, labelled reference	143.092 (0.406)	133.642 (0.406)	14.173 (0.069)	282.5	0.002
Mean schedule	133.892 (0.370)	119.800 (0.370)	0.331 (0.015)	430.0	0.068
Recursion, price read now	133.762 (0.363)	119.719 (0.366)	0.251 (0.010)	428.6	0.068
Recursion, one window old	133.781 (0.364)	119.741 (0.366)	0.272 (0.010)	428.3	0.068
Perfect knowledge of the draw	133.463 (0.363)	119.469 (0.367)	0	426.4	0.068

A guard costs on two counts, charging room and warmer water, and the first is the larger. With a guard of 20 kW on the same records, the schedule's net is 119.335 $ and the recursion's 119.253 $, with $T_s - T_{sp}$ at 0.117 and 0.116 K against 0.068 K, and the response is unchanged at 0.081 $ (se 0.013). Of the 0.465 $ gap, the rows alone, with no water in them, give 0.256 $ (55.1 %): legitimate charging room. That leaves at most 0.209 $ (44.9 %) for a plan that charges into the capacity row being warmer than its stance. On the rows, guard 30 charges from 22:55 to 06:00, with 1.2 kWh at the 0.14 tariff, and costs 0.393 $ against no guard. Guard 20 starts at 23:00. Hence one limit table for all: letting only the recursion read its load for the limit would compare guards and not plans. Pr(all ≤ 27) = 1.000 on every row at this truth is one draw of the unread entries and not coverage, and the joint over 144 windows is at Chapter 18's multiplier, not recalibrated over the shift. The mean largest inlet is 26.520°C on all of them.

21.4 The checks (step 5)

The checks test the recursion against the program, the band's limit, the ordering of the three plans, and that nothing reads the truth.

21.4.1 The recursion against the program and the band's limit

Table 21.16 gives every check. Each row names one; the columns give the two values and the verdict. The first line is the resolution of every recursion number: the gap of 0.0110 $ between the recursion with no band and the program is the grid's error, and no difference below that is read as a result. The rows of the tank are exact in the program; the corresponding limit here is the band's, $\sigma \to 0$, taken to 10^{-4} in the table.

Table 21.16: The checks.

Check	Result	Verdict
Recursion with no band against the program	112.5880, 112.5770 $	gap 0.0110 $, the grid's error
The same, refined, 481 levels and 161 moves	112.5838 $	gap 0.0069 $
Band 0.10, refined against the book's grid	net 112.4276, 112.4336 $; response 0.1415, 0.1355 $	paired gap −0.0060 (se 0.0002)
Band $\sigma \to 0$, taken to 10^{-4}	112.5878 against 112.5770 $	within the grid's error
Ordering on the same 400 paths	perfect 112.1789 ≤ recursion 112.4803 ≤ schedule 112.6117 $	smallest per-path gaps 0.0658, 0.0628, positive
Path model	mean factor 1.00018; spread of a at window 72 0.1007	stated 1, 0.10
Stated truth changed, belief and readings fixed	stance, schedule and table	gaps 0.0
Tie-break removed	electricity less the change in V	0.00000 $
Round trip by run against derivation	Table 21.10	equal to the printed digit
Battery's rows against equation (21.17)	Table 21.14	equal to the printed digit
Chapter 18's multiplier and hour	2.172, 1.906; 21.846, 14.953 $	reproduced
Store's rows against the plant, fresh records	0.099 $ (0.07 %), spread 0.132 $; level, mean absolute error 3.97 kWh	close, stated

21.4.2 What the checks leave open

Monte Carlo standard errors are printed with every row they apply to, and the response of 0.135 $ stands with a grid error of about a tenth of itself: the refined grid reads 0.1415 $. The two response numbers on the plant's band, 0.135 on the rows and 0.080 $ on the fresh records, are not like with like: the rows' has no guard and charges 128 kW at 00:55 in Table 21.8, and the records have the guard of 30 kW. On the rows with the guard, everything else at its mean, the response is 0.095 at 20 kW and 0.088 $ (se 0.003) at 30 kW, against the records' 0.081 and 0.080 $. So the guard accounts for most of the gap, and what remains is inside the records' standard error; the load, ambient and stance held at their means are not shown to matter here. No check here tests the outage's distribution, the wear, or the terminal value's premise of an idle day.

21.5 The report (step 6)

The band on the price changes no schedule. A response to it is worth 0.080 to 0.088 $ per shift on the tank at the logged guard, 0.135 $ with no guard, mostly re-timing, and about 4.3 $ on the battery. The terminal value decides the morning, with a regret of 3.966 $ for the recharge rule at its stated κ. The limit takes the tank only below about 720 kW, and then, at this truth and at the forecast's mean ramp, a mean plan loses every record, though not at a ramp of 18 kW.

21.5.1 The answers and their bands

Table 21.17 gives the answers in the order of Table 21.1. Each row is a question; the columns give the answer, its band or basis, and the stated truth, which here is the true ramp only.

Table 21.17: The report for the two stores. Dollars are per shift and this example's.

Control engineer's question	Answer	Band or basis	Stated truth
What is a kWh left at 07:00 worth?	0.03412 falling to 0.03186, kink at about 430 kWh, then about 0.0195 $/kWh; end at 430.0 kWh; saves 12.563 $ at the periodic state	an idle day and a repeat; regret 3.966, 1.952 $	–
Does the price band change the schedule?	No: gap 0.0 at 2000 draws; a sample of 100 at band 0.33 costs 0.107 $	se 0.003; linear cost, fixed q	–
What is the band worth with a response?	0.135 $ at band 0.10, 0.088 $ with the guard, 1.465 $ at 0.33; recursion over perfect 0.301 $; mean 13.697 $	se 0.003, 0.016; grid's error 0.011 $	–
Price or COP, and what does the loss take?	Price 85.7 %, COP 14.3 % of 11.487 $; break-even 0.96421 at 03:00 into 06:25	equal by run and derivation	–
When does the tank serve the limit?	Never at 800 kW; below about 720 kW; at 670 kW $R_{95} = 218.2$ kWh, at 5.368 $ of electricity, 0.508 $ net if capacity returns	need's spread 65.3 kWh; fresh, Pr(tank empty) 0.170, one truth	ramp 707.0 kW; need 143.5 kWh
What reserve does the UPS need, and is arbitrage worth it?	157.3 kWh at 0.95; 120 kWh reads 0.9197; 39.822 $ at 120, with 0.14222 $/kWh the wear flip on means; charge from 23:00; response 4.223 $	se 0.056; distribution declared, no record; no wear figure stated	–
What ties the two stores?	The price and, at Chapters 29 and 30, the grid; charging together reads 404.3 kW; 0.14222 against 0.02656 $/kWh stored	no grid row here	–
Does it hold on fresh records?	Response 0.080 $; over perfect 0.331, 0.251 $; battery response 4.336 $	se 0.014, 0.270; 200 records, one truth	–
What next?	Section 21.6	means held	–

21.5.2 The entry in the commissioning log

The entry sets the capacity guard at 30 kW, a stated choice. It states the tank's schedule at that guard: out from 19:00 to 20:30, in from 22:55 to 06:00, out to 430 kWh by 07:00, with the terminal value cyclic and stated. It replaces the sheet's "end at 300". It records that the band on price is not a derate. At capacity short it holds the reserve table at R_{95}, with the 0.95 over the forecast, about 0.87 at the plant's draw, and a true ramp above its mean noted. It enters the battery's reserve as a declared assurance, with 120 kWh reading 0.9197 on a distribution no record supports, and it charges the battery from 23:00. It gives the wear flip of 0.14222 $/kWh to compare against a figure when one exists. And it notes for Chapters 29 and 30 that both stores charge together.

21.5.3 Common mistakes

Table 21.18 lists six mistakes these stores make easy. Each row names the mistake, its form here and the correction.

Table 21.18: Six mistakes with a store.

Mistake	On this plant	Correction
Derating a schedule by the price's band	±10 % moves no schedule; gap 0.0 at 2000 draws	the expected cost at fixed q reads the mean, equation (21.11); price the response instead
Saying the band never matters	1.465 $ at band 0.33; 4.223 $ on the battery	it matters past about 0.2, or on a store free to pick its hours
Pinning the tank's end, or valuing it at the recharge price	regret 1.952 and 3.966 $; slope three times the recharge's below the kink, with κ at 07:00	the terminal value is a choice; state it and print the regret
Reading a mean plan's lower net at a short capacity as a saving	122.327 $, every record over its inlets	compare plans that hold the limit; warm water is cheap
Reading Pr = 1.000 on fresh records as coverage	all rows, one truth	one draw of the unread entries
Treating the outage distribution as data, or arbitrage as settled	0.9197 on declared numbers; no wear figure	state the assurance as declared; compare a wear figure with 0.14222 $/kWh

21.6 The next measurement (step 7)

Each candidate adds its rows to the belief by the preposterior (4.73), or its reading to a store's rows, with the means held, so each number flatters, as Chapter 18 found.

21.6.1 The candidates

Table 21.19 orders them by dollars per shift against none, the load forecast's being an upper bound at a what-if and so set after the figure it could exceed. Each row is a candidate; the columns give the figure and the note. For the outage log the rows add precisions: 20 logged durations with a log spread of 1.0 on a prior of 0.5 give a curvature of $4+20 = 24$, so $\sigma_\mu = 0.2041$ and $s_D = 1.0206$.

Table 21.19: The candidates, means held, multiplier kept, against none.

Candidate	$/shift (se)	Note
Permanent inlet readings, B3 and A5, 0.3 K, 12 windows	11.208	T_s rises 1.482 to 1.597 K; spread of β 0.0132, 0.0111 against 0.0800; multiplier kept, not recalibrated
Outage log, 20 durations	3.311	R_{95} from 157.3 to 134.0 kWh, median unchanged; a log may raise it
Load forecast, ramp spread 8 to 2 kW per hall	at most 3.346	only at the capacity node: R_{95} from 218.2 to 126.7 kWh
Published price against none, plant's band	0.135 (0.003)	1.465 (0.016) at band 0.33
Tank level read perfectly, against spread 5 kWh	0.035 (0.001)	against the filter's 1.387 kWh, 0.004 (0.000)

21.6.2 The recommendation

The instrument that is not on the store reads 83 **times the price reading**, 11.208 against 0.135 $, **and more than the tank's own value on the same schedule and first shift**, 9.068 $ (the periodic figure is 12.563 $). It works through the stance, on the chiller's whole heat. It also lowers the tank's own value, from 9.068 to 8.443 $ of electricity saved, since

a warmer supply makes κ smaller. The two racks' rows are therefore the first candidate to price. No instrument cost is stated, and the means are held, so this ranks rows and gives no verdict to buy. A better price forecast, the sheet's suggestion, stands near the bottom of the table. The outage log costs no instrument and would replace a declared number.

21.6.3 What the analysis showed

The sheet offered a ratio, a derate, a column and a reserve. The derate has no row to act on, since a schedule on a linear cost reads the mean. The column decides nothing in the first evening and a good deal in the morning, and a store's value is a choice to be stated. The outline's lead question has a small answer on the tank: charge by the mean, and the band on the price is a rounding error until it is about 0.2 or the store can choose its hours. The band on the load is the band that matters, and only near capacity, where a plan built on the mean spends what the limit needs. The electrical store, worth five times the tank per kilowatt-hour cycled and free to pick its hours and its charging time, is where a response pays, subject to a wear figure and a reserve that this record does not supply. And the largest figure in Table 21.19 is not on either store: it is the two racks no reading touches, with the means held and the multiplier not recalibrated, so it ranks rows and gives no verdict.

21.6.4 Limits of the model

The separation. The stance is separated from the store and there is no joint solve. The rows' κ reads one belief at 20:00, T_s from 15.4 to 14.4°C, against a mean of 14.117 to 14.152°C on the fresh records from their own filters.

The recursion. Ambient and load are at their means in it, and the hall and water stores are ignored at the shift's scale. Its grid's error is 0.011 \$. One AR(1) carries the price, with no correlation with the load, and the window's price published before the move is a choice.

The terminal value. The cyclic value assumes an idle day and a repeat, and no day tariff is stated.

The battery. The outage distribution and the ride-through load are declared. An outage is at most one per evaluation, with its frequency unmodeled; there is no recharge race after one, no wear and no calendar loss.

The limit. No value is placed on a breach, the 0.95 being stated. Fresh records are at one truth and number 200. The joint over 144 windows is not calibrated, and the belief is a Laplace normal, as in Chapter 18.

The numbers. Percentages and dollars are this example's.

The constructed record. The record, the tariff and the outage distribution are constructed from a stated truth and stated choices, so the agreement tests the method on its own model and not the model against a real plant.

Nothing here learns from a move, and the next Part takes up the moves that do, beginning in Chapter 22.

21.7 What to retain

Definition 21.1 (Stance and store)**.** Where a store's move enters the plant's rows only through a capacity row, the shift separates: the stance from the belief alone, equation (21.6), then the store over its level, equation (21.9) or (21.12), given the price of cooling κ_k, with a stated terminal value.

Three conclusions carry beyond this plant as statements of form. A fixed schedule on a linear cost reads the price's mean, equation (21.11), so a band on the price is priced by a policy that reads it and not by a derate; what the response earned here is this plant's, and it needs a store free to choose its hours. A round trip's break-even, equation (21.13), carries the loss and the ratio of the costs of cooling together, and that ratio can bring the price ratio below one. A reserve is a quantile of a need, equations (21.15) and (21.16), and a band that a schedule ignores returns through the floor it has to hold. The chapter's questions join the others in Appendix A.

Table 21.20 lists the chapter's terms with their expressions, units and the subsection that defines each.

Table 21.20: The terms of the chapter and where each is defined. Units for the two-hall plant.

Term	Expression	Units	Meaning	Defined in
Tank's row	eq. (21.1)	kWh	level with loss and efficiency	21.1.2
Price of cooling	κ_k, eq. (21.2)	–	1/COP at the stance	21.1.2
Stance	eq. (21.6)	°C	the hold from the belief alone	21.2.1
Price's law	eq. (21.7)	\$/kWh	conditional at lead ℓ	21.2.2
Terminal value	$V(E)$, eq. (21.10)	\$	cyclic, over an idle day	21.3.1
Certainty equivalence in price	eq. (21.11)	\$	a fixed q reads the mean	21.3.2
Value of recourse	schedule less policy, eq. (21.12)	\$	the value of the response	21.3.3
COP arbitrage	eq. (21.13)	–	charging where the lift is small	21.3.4
Need	eq. (21.14)	kWh	energy to serve the deficit to 07:00	21.3.5
Reserve as a chance constraint	eq. (21.15), (21.16)	kWh	a quantile, on a store	21.3.5
Wear flip	w^*, eq. (21.18)	\$/kWh	where arbitrage stops	21.3.6
Guard	30 kW below capacity	kW	keeps T_s on its set point	21.1.4

21.8 Exercises

Exercise 21.1. *Formulation.* A single hall with a hot-water tank is held by a heater of limited power. Write the stance and the store, and name the only row that ties them. (The stance is the warmest or coolest set point the band on the protected temperature allows, a function of the belief alone, and it fixes the heater's price of heat. The store is $E_k = f E_{k-1} + \Delta t(\eta q_k^+ - q_k^-)$ with cost $\sum p_k \kappa_k \Delta t(\text{load}_k + q_k)$ and a terminal value. The only row with entries of both is the heater's power limit, $\text{load}_k + q_k \le C$.)

Exercise 21.2. *Derivation.* Derive certainty equivalence in price and say why the ambient follows it. (At a fixed q the cost is $\sum_k p_k \kappa_k \Delta t(\text{heat}_k + q_k)$, linear in p_k, and κ_k, the heat and q_k are not functions of p_k, so the expectation passes through and reads $\mathbb{E}[p_k]$ alone. $\kappa = (T_{\text{amb}} + T_{\text{app}} - T_s)/(\eta_{\text{cop}}(T_s + 273.15))$ is linear in T_{amb}, so its expectation is its value at the mean, and the quadrature's gap is zero.)

Exercise 21.3. *Derivation.* Derive the round trip's break-even, and give it for 03:00 in and 06:25 out. (One kilowatt-hour out at k needs $1/(\eta f^{g+1})$ in at j, by iterating the tank's row over $g+1$ windows. It pays if $p_{\text{out}}\kappa_{\text{out}} > p_{\text{in}}\kappa_{\text{in}}/(\eta f^{g+1})$. With $\kappa_{\text{in}} = 0.12532$, $\kappa_{\text{out}} = 0.13917$, $\eta = 0.95$, $f = 0.999583$ and $g = 40$, the ratio is 0.96421.)

Exercise 21.4. *Application.* Give the battery's saving at the stated reserve. (Usable $U = 400 - 120 = 280\,\text{kWh}$; margin $0.22 - 0.07/0.9 = 0.14222\,\$/\text{kWh}$; the product is $39.822\,\$$.)

Exercise 21.5. *Application.* Give the wear cost at which arbitrage stops on tariff means, at 0.22 against 0.07. (Out at 0.22 with a cost w against in at 0.07 with $1/\eta_{\text{rt}}$: $0.22 - w = 0.07/0.9$, so $w^* = 0.14222\,\$/\text{kWh}$.)

Appendix A

Atlas of control questions

This appendix can be read independently of the chapters. Each row is a question a control engineer asks. The first column gives the question in the engineer's words, as the chapter's questions table gives it; the second the general form of its category, written out or cited by equation number; the third its mathematical form for the plant of that chapter, once the plant graph and its state $\mathbf{z}$ are fixed, in that chapter's symbols; the fourth its category and the fifth its computation, the five categories and the five computations of Section 2.3.6; and the last the section in which the question is answered, or the chapter when a whole chapter answers it. The rows of Chapter 2 are the forms of its classification, Table 2.11. Chapter 3 develops the probability rules the rows rest on and holds no questions table, so it has no rows.

Control engineer's question	General form	Mathematical form	Category	Computation	Answered in
Chapter 1. Controlled Plants as Graphs					
What temperature is the hall running at now?	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$, $\mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y})$	$p(T \mid \mathbf{y})$, $T = z_1$	estimation	posterior	1.3.1
What is the process gain to the load?	$\mathbb{E}[f(\mathbf{z}) \mid \mathbf{y}]$, $\mathrm{Var}(f(\mathbf{z}) \mid \mathbf{y})$	$p(K_Q \mid \mathbf{y})$, $K_Q = \mathrm{e}^{-z_2}/G_{\mathrm{ref}}$	estimation of a derived quantity	posterior	1.3.1
Is the hall within its 24°C limit now?	$\Pr(f(\mathbf{z}) < f^* \mid \mathbf{y})$	$\Pr(T < 24^\circ\mathrm{C} \mid \mathbf{y})$	prediction under a threshold	posterior	1.3.2
Can the hall carry 60 kW within 27.5°C at this set point?	$\Pr(f(\mathbf{z}) < f^* \mid \mathbf{y})$	$\Pr(T_2 < 27.5^\circ\mathrm{C} \mid \mathbf{y})$, $T_2 = T_c + (Q_2/G_{\mathrm{ref}})\,\mathrm{e}^{-z_2}$	prediction under a threshold	posterior	1.3.2
How far must the set point move to hold it?	$\Pr(f(\mathbf{z}; r) < f^* \mid \mathbf{y})$ at a scenario node	$\Pr(T_2(r) < 27.5^\circ\mathrm{C} \mid \mathbf{y})$, $T_2(r) = r + (Q_2/G_{\mathrm{ref}})\,\mathrm{e}^{-z_2}$	prediction under a threshold	posterior	1.3.3
Which added sensor would narrow the gain?	preposterior covariance $(\mathbf{\Sigma}^{-1} + \mathbf{J}_a^{\mathsf{T}}\mathbf{J}_a)^{-1}$	$\mathbf{\Sigma}_a$ for each candidate row	design	preposterior	1.6.1
Chapter 2. The Language of Control Questions					
What is unknown, and how well is it known?	(2.9), (2.10)	$p(z_i \mid \mathbf{y})$, $\mathbb{E}[z_i \mid \mathbf{y}]$	estimation	posterior	2.4.1
Does a requirement hold, with the stated assurance?	(2.12)	$\Pr(f(\mathbf{z}) \geq f^* \mid \mathbf{y}) \geq \kappa$, f^* the requirement	prediction under a threshold	posterior	2.4.2
What happens in the coming run?	(2.13)	$\int \Pr(A_{\mathrm{future}} \mid \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	prediction under a threshold	posterior	2.4.4

Control engineer's question	General form	Mathematical form	Category	Computation	Answered in
Is it rising, and when does it cross its limit?	(2.21)	$\Pr(\gamma > 0 \mid \mathbf{y})$; $\Pr(t^* \le s \mid \mathbf{y}) = \Pr(\gamma > 0,\ f_s(\mathbf{z}) \ge f^* \mid \mathbf{y})$	prediction under a threshold for the sign and a crossing by a date; estimation for a forecast and the distribution of a crossing time	trend	2.3.4, Chapter 19
Is it the path or the sensor?	(2.14), (2.15)	$\Pr(M_j \mid \mathbf{y}) \propto p(\mathbf{y} \mid M_j)\ \Pr(M_j)$	hypothesis comparison; diagnosis over a family of faults	evidence	2.5.1, 2.5.2, 2.5.3
What is the expected consequence?	(2.23)	$\mathcal{R}(a) = \mathbb{E}[L(a, \mathbf{z}) \mid \mathbf{y}] = \int L(a, \mathbf{z})\, p(\mathbf{z} \mid \mathbf{y})\, \mathrm{d}\mathbf{z}$	estimation of a derived quantity	expected consequence	2.5.4
Retune or leave the tuning, and is a step test worth its slot?	(2.23), (2.16)	$\arg\min_a \mathcal{R}(a)$; EVSI	design	expected consequence; preposterior	2.5.5, 2.6.1, 2.6.2
Chapter 4. The Anatomy of Probabilistic Inference					
What is the thermal resistance of the air path?	moments, (2.10)	$p(\rho \mid \mathbf{y})$: $\hat{z}_1$, $\sqrt{\Sigma_{11}}$	estimation	posterior	4.2.10, 4.3.6
What was the hall temperature during the test?	moments, (2.10)	$p(T_1 \mid \mathbf{y})$: $\hat{z}_2$, $\sqrt{\Sigma_{22}}$	estimation	posterior	4.3.6
What will the hall temperature be at full load?	moments, quantiles, (2.10), (2.11)	$p(T_2 \mid \mathbf{y})$: $\hat{z}_3$, $\sqrt{\Sigma_{33}}$, credible interval	estimation	posterior	4.3.1, 4.3.4, 4.3.6
Can the hall be released for 60 kW within 28°C?	threshold, (2.12)	$\Pr(T_2 < T_{\max} \mid \mathbf{y})$, with $f = -T_2$	prediction under a threshold	posterior	4.3.5, 4.3.7
How far must the set point move to hold it at 0.95?	threshold at a scenario node, (2.12)	r_{95} from $\Pr(T_2^*(r) < T_{\max} \mid \mathbf{y}) = 0.95$	prediction under a threshold	posterior	4.3.8
Chapter 5. Actuators and Controllers as Graph Rows					
What does closing the loop change in the graph?	moments, (2.10)	the free and pinned entries of each graph; the columns of $\mathbf{J}$	estimation	posterior	5.3.1
What is the plant now?	moments of derived quantities, (2.10)	$p(G, K_Q, W \mid \mathbf{y})$	estimation of a derived quantity	posterior	5.3.2
Lower the set point 1 K: what moves, closed loop and open?	moments at a scenario node, (2.10)	$p(T^*, T_c^*, u^* \mid \mathbf{y})$ at $r^* = 17$°C, against u pinned	prediction	posterior	5.3.3
A load step to 45 kW: what moves, closed loop and open?	moments at a scenario node, (2.10)	the same at $Q^* = 45$ kW	prediction	posterior	5.3.4

Control engineer's question	General form	Mathematical form	Category	Computation	Answered in
Will the valve reach its stop at 17°C and 45 kW?	threshold at a scenario node, (2.12)	$\Pr(u^* \geq 1 \mid \mathbf{y})$	prediction under a threshold	posterior	5.3.5
Chapter 6. Protected Quantities and Chance Constraints					
What does an offset the loop cannot see do to the plant?	moments, (2.10)	$p(b, g, g_w \mid \mathbf{y})$, $p(T^*, u^* \mid \mathbf{y})$	estimation	posterior	6.3.1
What must stay inside its limit, with what probability?	chance constraint, (6.1)	$\Pr(T^*(r) \leq 27.5°\text{C} \mid \mathbf{y}) \geq 1 - \alpha$	prediction under a threshold	posterior	6.3.2
What is the highest set point that holds it?	the same, inverted, (6.17)	r_{95}, with back-off $\Phi^{-1}(0.95)\sqrt{\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}}$	prediction under a threshold	posterior	6.3.3
Is the valve a protected quantity too?	(6.1), gradient depending on r	$\Pr(u^*(r) \leq 1 \mid \mathbf{y}) \geq 1 - \alpha$	prediction under a threshold	posterior	6.3.4
Which set points hold both?	joint chance constraint, (6.20)	$\Pr(T^* \leq 27.5°\text{C},\ u^* \leq 1 \mid \mathbf{y}) \geq 1 - \alpha$	prediction under a threshold	posterior	6.3.5
At what load does no set point hold?	the same, over the load	the window's width against Q^*	prediction under a threshold	posterior	6.3.5
Chapter 7. The Control Problem on the Graph					
What is the problem every controller is solving part of?	expected consequence over policies, (2.23)	the problem (7.1)	design	expected consequence	7.1.1
With no limits, does the band change the move?	the same, restricted, (7.3)	$u_k = K_k \hat{x}_k$, against s_0	design	expected consequence	7.3.1
Can the future be fit like the past?	MAP against expectation, (7.2)	u_0 with the future noise left free	estimation	posterior	7.3.2
What does a response to the next reading save, under a limit?	chance constraint over policies, (6.1)	fixed, linear, free, held after every reading	design	posterior	7.3.3
When the plant is unknown, what does the band do to the move?	the recursion, (7.3)	certainty equivalent, cautious, exact	design	preposterior	7.3.4
Which later chapter is which restriction?	–	the map	–	all five	7.3.5
Chapter 8. Stability Under Parameter Uncertainty					
What does the band do to the critical gain?	moments of a derived quantity, (2.10)	$p(K_c(\mathbf{z}) \mid \mathbf{y})$, with $\rho(\mathbf{M}) = 1$ at K_c	estimation of a derived quantity	posterior	8.3.1
How likely is this tuning to be stable?	threshold at a scenario node, (2.12), (6.1)	$\Pr(\rho(\mathbf{M}(K_p, \mathbf{z})) < 1 \mid \mathbf{y})$	prediction under a threshold	posterior	8.3.2
What is the largest gain that holds 0.95?	the same, inverted, (6.17)	K_{95}: $\Pr(K_c > K_{95} \mid \mathbf{y}) = 0.95$	prediction under a threshold	posterior	8.3.3
And if it must also decay?	the same, on the pair, (6.1)	$\Pr(\lvert\lambda_{\text{pair}}\rvert \leq 0.8609 \mid \mathbf{y})$	prediction under a threshold	posterior	8.3.4
Where does the band come from?	budget of a variance, (4.54)	$\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$ by entry	estimation of a derived quantity	posterior	8.3.5

Control engineer's question	General form	Mathematical form	Category	Computation	Answered in
What does instability look like with stops?	simulation at the truth	$\mathbf{x}_{k+1}$ of the nonlinear loop	prediction	– (simulation at the truth)	8.3.6
Which test would narrow it most?	next measurement, (2.16)	sd K_c after each candidate	design of measurement	preposterior	8.6
Chapter 9. The Kalman Filter as a Time-Stacked Fit					
Is the filter a separate device?	moments, stacked, (9.19)	$p(T_k \mid \mathbf{y}_{60:k})$ against the recursion	estimation	posterior	9.3.1
What is the hall at the last step?	moments, (2.10)	$\hat{T}_{120\mid120}$, $P_{120\mid120}$	estimation	posterior	9.3.2
What do later readings say about earlier steps?	moments, (9.19)	$\hat{T}_{k\mid120}$ against $\hat{T}_{k\mid k}$	estimation	posterior	9.3.3
How fast does the band fall, and to what gain?	moments of a derived quantity, (2.10)	υ_k, $P_{k\mid k}$, the steady limit (9.21)	estimation of a derived quantity	posterior	9.3.4
What did the hall do through the load step?	moments, (9.19)	T_k as a virtual reading, $k = 81, \ldots, 120$	estimation	posterior	9.3.5
With g free, does the recursion still give the marginal?	moments of a derived quantity, (9.19)	$p(g, T_{120} \mid \mathbf{y}_{60:120})$, three routes	estimation of a derived quantity	posterior	9.3.6
Chapter 10. Feedback as Conditioning					
What does one reading do to the belief?	moments, (2.10), (10.5)	$(\hat{\mathbf{z}}, \boldsymbol{\Sigma}) \to (\hat{\mathbf{z}} + \mathbf{K}\iota,\ \boldsymbol{\Sigma} - S\mathbf{K}\mathbf{K}^{\mathsf{T}})$	estimation	posterior	10.3.1
What is known before the reading comes?	moments of a preposterior, (10.8)	$\boldsymbol{\Sigma}_n$ against n	estimation	preposterior	10.3.2
What can more hall readings not remove?	the same, in the limit, (10.9)	$\lim_n \mathrm{sd}\, T^*$	estimation	preposterior	10.3.3
What does one reading do to the window?	threshold at a scenario node, (2.12), (6.17)	r_{95} and the window after y	prediction under a threshold	posterior	10.3.4
Should the PI row act on the reading or on a filtered estimate?	expected consequence, one fixed law, (2.23)	travel, T_c, margin against L	design	expected consequence	10.3.5
What do a step response's rows add?	moments, (2.10)	$p(g_w, b \mid$ stacked rows$)$	estimation	posterior	10.3.6
Chapter 11. Observability and Sensor Loss					
What can these readings resolve?	moments, (2.10)	rank of the reading rows; sd against prior, per entry	estimation	posterior	11.3.1
What does the controller need of them?	moments of a derived quantity, (2.10)	$\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ for T_3, the hold node and T_s^*	estimation of a derived quantity	posterior	11.3.2
The loop's reading drops: what is the loop?	moments at a scenario node, (2.10)	T_3^*, T_r^*, T_s frozen, open loop	prediction	posterior	11.3.3
Can a virtual reading replace it?	moments at a scenario node, (2.10)	closed on the air balance, T_3^*, closed loop	prediction	posterior	11.3.4
Hold or virtual loop, as the load moves?	chance constraint, (6.1)	$\Pr(T_3^* \le 27^\circ\mathrm{C} \mid \mathbf{y}) \ge 0.95$; r_{95}, $T_{s,95}$, crossing loads	prediction under a threshold	posterior	11.3.5

Control engineer's question	General form	Mathematical form	Category	Computation	Answered in
Which other reading matters most?	moments, row removed, (2.10)	the ranking by added band on T_3	estimation	posterior	11.3.6
What does a parent level keep?	marginal of a derived quantity, (2.10)	$p(Q_{\text{tot}}, G_{\text{tot}} \mid \mathbf{y})$ by Schur complement, (11.18)	estimation of a derived quantity	posterior	11.3.7
Which sensor next?	preposterior variance, (2.16), (4.73)	ranking by $\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g}$ against by entry spread	sensor placement	preposterior	11.6.1
Chapter 12. Certainty Equivalence and the Linear-Quadratic Law					
What is the law with no band and no limit?	expected consequence over moves, (2.23)	$\mathbf{K}_k$ from the stacked rows, (12.9)	design	expected consequence	12.3.1
Is the recursion Chapter 9's in disguise?	–	$\mathbf{A} \to \mathbf{A}^{\mathsf{T}}$, $\mathbf{B} \to \mathbf{C}^{\mathsf{T}}$, $\mathbf{Q} \to \mathbf{W}$, $\mathbf{R} \to \mathbf{V}$	–	none: algebra	12.3.2
When does the band change the move?	expected consequence at a belief, (2.23)	v_0 and $\mathbb{E}J$ against $\boldsymbol{\Sigma}$, (12.16)	design	expected consequence	12.3.3, 12.3.4
What does the band cost, if not the move?	moments of a derived quantity, (2.10)	$\sum_k \operatorname{tr}(\mathbf{K}_k^{\mathsf{T}} H_k \mathbf{K}_k \boldsymbol{\Sigma}_{k\mid k})$	estimation of a derived quantity	posterior	12.3.5
How does the law do against the PI?	prediction at a scenario node, closed loop, (2.10)	load and set-point steps, noise	prediction	posterior	12.3.6
What does fit error cost in closed loop?	moments of a derived quantity, (2.10)	regret over $p(\mathbf{z} \mid \mathbf{y})$, (12.19)	estimation of a derived quantity	posterior	12.3.7
Which measurement most reduces the expected cost?	next measurement, (2.23), (4.73)	the law's cost term after each candidate, (12.20)	design of measurement	preposterior	12.6.1
Chapter 13. Integral Action					
Where is the integrator?	expected consequence, the law, (2.23)	a pole of $\mathbf{A}_c$ at $z = 1$	design	expected consequence	13.3.1
What does the cost on the move do?	the same, by cost, (2.23)	steady y against the cost row	design	expected consequence	13.3.2
Which disturbance does the loop assume?	moments, then the law, (2.10)	$\mathbf{E}$, F of the entry; load-step response	design	posterior	13.3.3
How many entries can one reading carry?	moments, (2.10)	rank of the steady rows	estimation	posterior	13.3.4
Does the meter remove the need?	expected consequence, (2.23)	feedforward against the entry, two events	design	expected consequence	13.3.5
What can it not remove?	threshold at a scenario node's entry, (2.12)	$T_c \to r - b$; $\boldsymbol{\Sigma}$ of b	estimation	posterior	13.3.6
Is the derived controller better than the PI?	expected consequence, fresh records, (2.23)	travel, scatter, mean square, margins	design	expected consequence	13.3.7
Does the band of d enter the law?	expected consequence, (2.23)	$\mathbf{K}, K_d$ against $\boldsymbol{\Sigma}$	design	expected consequence	13.3.8
Chapter 14. Tuning as an Objective					
Is this PI an LQ law on the hall's state?	inverse of the expected consequence, (14.8)	$\gamma \geq 1$, $\gamma\lvert\chi_c\rvert^2 - \lvert\chi\rvert^2 \geq 0$	design	– (algebra at the fit), then posterior	14.3.1
And on the model the PI sees?	the same on a reduction, (14.8)	on (e, z), $\mathbf{Q}$ diagonal	design	– (algebra at the fit)	14.3.2

Control engineer's question	General form	Mathematical form	Category	Computation	Answered in
What is it the best PI for?	restricted expected consequence, inverted, (14.8)	$\nabla_K[qV_e + \rho V_m] = 0$ for (q, λ)	design	expected consequence	14.3.3
What does the ratio buy?	the forward map, (2.23)	$s_T/s_m \mapsto K_p$, decay, margin	design	expected consequence	14.3.4
Does the band move the gains?	expected consequence over the posterior, (2.23)	$\min_K \mathbb{E}[J(K, \mathbf{z}) \mid \mathbf{y}]$	design	expected consequence	14.3.5
Does a schedule do better than the band?	the same, by scenario node, (2.23)	$\min_{K(Q^*)}$ against $\min_K$	design	expected consequence	14.3.6
Which test most improves the tuned cost?	next measurement, (2.16)	$\mathbb{E}_{\mathbf{y}}[\min_K \mathbb{E}(J \mid \mathbf{y})]$	design of measurement	preposterior	14.6
Chapter 15. Saturation and Anti-Windup					
Where is the stop, and with what probability is the loop at it?	threshold, (2.12)	$\Pr(m < 0 \mid \mathbf{y})$	prediction under a threshold	posterior	15.3.1
What is the hall air across the kink?	moments of a mixture, (2.10)	$p(T^* \mid \mathbf{y})$, two regimes	prediction under a threshold	posterior	15.3.2
Does Chapter 6's window survive?	joint chance constraint, (6.20)	$\Pr(T^* \le L,\ u^* < 1 \mid \mathbf{y})$	prediction under a threshold	posterior	15.3.3
What is the law at the stop?	expected consequence, the stop kept, (15.13)	bounded least squares against the clipped law	design	expected consequence	15.3.4
Why does the integrator wind, and what stops it?	the Fit line with the wrong input, (15.16)	the estimator's row on c or on $\mathrm{sat}(c)$	estimation	posterior	15.3.5
What does each treatment cost on recovery?	expected consequence of fixed laws, (2.23)	recovery, fall below $r - b$	design	expected consequence	15.3.6
What does the stop reveal?	moments, (2.10)	$p(b, g_w \mid$ rows with y_c live)	estimation	posterior	15.3.7
Chapter 16. Chance-Constrained Set-Point Selection					
What set point holds each inlet?	chance constraint inverted, (2.12)	$x_{95,k} = 27 - m_k - \kappa s_k$; loop, hold, two routes	prediction under a threshold	posterior	16.3.1
What does holding all three cost?	joint chance constraint, (2.12)	$\Pr(\max_k T_k^* \le 27^\circ\mathrm{C} \mid \mathbf{y}) \ge 0.95$	prediction under a threshold	posterior	16.3.2
Does the normal hold?	the same, by draws, (2.12)	weighted quantile of the exact posterior	prediction under a threshold	posterior	16.3.3
Does a tighter posterior raise the set point?	preposterior of a threshold, (2.16)	$\mathbb{E}[x_{95}(y)]$, its sd, $\Pr(x_{95}(y) < x_{95})$	sensor placement	preposterior	16.3.4
What is a degree worth?	expected consequence through a price, (2.23)	$\mathbb{E}[P(T_s^*(x))]$ against the largest r at the truth	decision value	expected consequence	16.3.5
Does the 95 % hold on fresh records?	coverage, (16.18)	over the prior; at one truth; with an offset	validation	posterior	16.3.6
When does no set point hold?	window of set points, (2.12)	ceiling against T_s stops, over Q^*	prediction under a threshold	posterior	16.3.7

Control engineer's question	General form	Mathematical form	Category	Computation	Answered in
What next?	preposterior value, (2.16)	saving in \$/h, paired on the draws	sensor placement	preposterior	16.6.1
Chapter 17. Expected-Cost Set Points					
What does a degree of set point cost?	expected consequence, (2.23)	$E(r)$, $c_e(r) = -E'(r)$	decision value	expected consequence	17.3.1
Where does a price on the excess put it?	the same, minimized, (2.23)	$\Pr(T_{\max} > 27) = c_e(r^*)/c_v$	design	expected consequence	17.3.2
Where does a price per violation put it?	the same, (2.23)	$c_e(r) = C_f\, p(r)$, or a stop	design	expected consequence	17.3.3
Which price does $\alpha = 0.05$ imply?	the inverse, (2.23)	$c_v = c_e/\alpha$, $C_f = c_e/p$	decision value	expected consequence	17.3.4
What does the band do to the optimum?	the same, across $\mathbf{\Sigma}$, (2.23)	location, curvature, interval within a tolerance	design	expected consequence	17.3.5
What is the chance constraint's error if the price is known?	paired, same draws, (2.23)	$J(r_{\mathrm{cc}}) - J(r_{\mathrm{ec}})$; perfect knowledge	decision value	expected consequence	17.3.6
Sum over racks or joint limit?	expected consequence, (2.23)	$\sum_k c_k \Pr(T_k > 27) = c_e$	design	expected consequence	17.3.7
Do the readings rank the same?	preposterior value, (2.16)	$\mathbb{E}[J(r_0) - J(r(y))]$ against the back-off's	sensor placement	preposterior	17.3.8
Does it hold on fresh records?	coverage and realized cost, (2.23)	violations against the α asked	validation	expected consequence	17.3.9
Chapter 18. Model Predictive Control as a Restriction					
Is the plan one fit?	two problems, expected consequence, (2.23)	rank of past rows on the future moves, across copies	design	expected consequence	18.3.1
What does the plan run from?	moments, (2.10)	filter of the fit, $(\hat{\mathbf{z}}_k, \mathbf{\Sigma}_k)$	estimation	posterior	18.3.2
How good are its rows?	moments of a prediction, (2.10)	rows and prediction model against the plant, noiseless	design	posterior	18.3.3
What is the plan over the next hour?	expected consequence with limits, (2.23)	linear rows across copies, min $\mathbb{E}$ cost	design	expected consequence	18.3.4
What do 95 % limits per copy give?	joint chance constraint, (2.12)	$\Pr_J$ at $z = 1.645$	prediction under a threshold	posterior	18.3.5
What does holding the joint cost?	the same, by z, (2.12)	Bonferroni against calibrated, across copies	design	expected consequence	18.3.6
What does a fixed plan give up?	expected consequence against a bound, (2.23)	perfect knowledge of the draw; re-planning's first number	design	expected consequence	18.3.7
Do the stop, capacity, mass, tank matter, and when?	what-ifs, (2.23)	pinned moves at scenario nodes, scan by window	design	expected consequence	18.3.8
Does it hold on fresh records?	threshold on fresh records, (2.12)	coverage	prediction under a threshold	posterior	18.3.9
Chapter 19. Forecast Uncertainty and Constraint Tightening					
What moves an inlet under the loop?	moments at scenario nodes, (2.10)	$\partial T_k/\partial(T_{sp}, r, Q, \beta, g)$ per node	design	posterior	19.3.1
Which entry makes the band, by lead?	moments, (2.10)	$\mathrm{Var}(\max_k T_{k,j} \mid \mathbf{y})$, groups on one at a time	estimation	posterior	19.3.2
Does the band level off under feedback?	moments at nodes, (2.10)	the band by lead, open, closed, affine	prediction under a threshold	posterior	19.3.3

Control engineer's question	General form	Mathematical form	Category	Computation	Answered in
What does each schedule cost?	expected consequence with limits, (2.23)	$z\,\mathrm{sd}_j$ per copy, joint 0.95; against perfect knowledge	design	expected consequence	19.3.4
What does doing nothing risk before the next decision?	threshold and trend, (2.12), (2.21)	$\Pr(\exists k, j \le s : T_{k,j} > 27 \mid \mathbf{y})$, with $\Pr(\gamma > 0,\ f_s \ge 27 \mid \mathbf{y})$ apart	prediction under a threshold	trend	19.3.5
What does waiting one window cost?	expected consequence, (2.23)	plan now against plan after one held window	design	expected consequence	19.3.6
Trend or published forecast, and when?	trend, (2.21)	$\hat{R}$ from readings against 25 kW; rmse, coverage; the valve's $\Pr(t^* \le s \mid \mathbf{y})$	prediction under a threshold	trend	19.3.7
Where does weather enter?	moments at nodes, (2.10)	ambient band at the rows; one unit; a capacity row	design	posterior	19.3.8
Chapter 20. Recourse and the Receding Horizon					
Why take only the first move?	posterior against the planned posterior, (2.10)	innovation against its band; moved entries; plan's second move against the re-solved first	estimation	posterior	20.3.1
What is a response worth?	expected consequence, paired, (2.23)	net, fixed less re-solved, by truth	design	expected consequence	20.3.2
How often must it re-solve?	the same, by interval, (2.23)	5, 15, 60 min, with and without a surprise	design	scenario nodes across time	20.3.3
What does a restated plan hold?	joint chance constraint, restated, (2.12)	conditional at each now against frequency over the shift	prediction under a threshold	posterior	20.3.4
Can the re-solve have no move?	recursive feasibility, (20.13)	least slack at a prohibitive price; the rows' contradiction	design	scenario nodes across time	20.3.5
Where is a response worth more?	what-ifs at scenario nodes, (2.23)	load step, one unit, tank, by node	design	scenario nodes across time	20.3.6
What should be read next?	fresh records with a real reading	saving and value of recourse by instrument	validation	posterior	20.6
Chapter 21. Thermal and Electrical Storage					
What is a kWh left at 07:00 worth?	expected consequence, (2.23)	$V(E)$, cyclic recursion, against stated values	design	expected consequence	21.3.1
Does the price band change the schedule?	certainty equivalence, (2.23)	$\mathbb{E}[\text{cost}]$ linear in p at fixed q	design	expected consequence	21.3.2
What is the band worth with a response?	stochastic recursion, (2.23)	$J_k(E, a)$, policy against schedule and perfect knowledge	design	expected consequence	21.3.3
Price or COP, and what does the loss take?	what-ifs and a break-even, (2.23)	scenario nodes; $p_{\text{out}}\kappa_{\text{out}} > p_{\text{in}}\kappa_{\text{in}}/(\eta f^{g+1})$	design	scenario nodes across time	21.3.4
When does the tank serve the limit?	chance constraint on a store, (2.12)	$\Pr(\text{need}_k \le E_k \mid \mathbf{y}) \ge 0.95$	prediction under a threshold	posterior	21.3.5
What reserve does the UPS need, and is arbitrage worth it?	chance constraint, then expected consequence, (2.13), (2.23)	$\Pr(D \le E/P) \ge \alpha$; wear flip	design	expected consequence	21.3.6

Control engineer's question	General form	Mathematical form	Category	Computation	Answered in
What ties the two stores?	the coupling, stated, (2.23)	a sum of costs, no grid row here	design	scenario nodes across time	21.3.7
Does it hold on fresh records?	fresh records, (2.23)	paired net through the plant	validation	expected consequence	21.3.8
What next?	preposterior, means held, (4.73)	saving in $/shift	sensor placement	preposterior	21.6