

Control Engineering by Derivation

Learning the plant while controlling it, and holding it up when something fails

Draft · 203 pages

Ricardo Marquez, Ph.D.

ricardo-marquez.com

Control Engineering by Derivation

Volume III

Learning, faults and supervision

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. Volume II holds Parts IV and V, Chapters 12 to 21: the control law derived from that problem, then limits and a horizon. This third volume opens with Part VI, in which the move may change what is learned, and continues with Part VII, in which the model itself is in question. Chapter numbers run continuously across the volumes; a reference to a chapter numbered 21 or below points to Volume I or Volume II, 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 the earlier volumes closed, and the atlas of questions in the appendix covers all the chapters written so far.

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 **v**

VI Learning While Controlling **1**

22 Dual control **2**
22.1 A self-tuning option on the sheet 2
22.2 The posterior of the state (step 3) 6
22.3 The answers (step 4) 7
22.4 The checks (step 5) 17
22.5 The report (step 6) 19
22.6 The next measurement (step 7) 22
22.7 What to retain 24
22.8 Exercises 25

23 Excitation design **26**
23.1 The test on the sheet 26
23.2 The posterior of the state (step 3) 29
23.3 The answers (step 4) 30
23.4 The checks (step 5) 46
23.5 The report (step 6) 47
23.6 The next measurement (step 7) 47
23.7 What to retain 50
23.8 Exercises 51

24 Closed-loop identification **52**
24.1 What a loop's record reads 52
24.2 The posterior of the state (step 3) 55
24.3 The answers (step 4) 56
24.4 The checks (step 5) 71
24.5 The report (step 6) 72
24.6 The next measurement (step 7) 74
24.7 What to retain 76
24.8 Exercises 77

25 The adaptive set point **78**
25.1 A filter that has run for three months 78
25.2 The posterior of the state (step 3) 82
25.3 The answers (step 4) 83

25.4 The checks (step 5) . . . 96
25.5 The report (step 6) . . . 96
25.6 The next measurement (step 7) . . . 99
25.7 What to retain . . . 100
25.8 Exercises . . . 101

VII Faults and Reconfiguration 102

26 Fault detection inside the loop 103
26.1 A tight trend at two in the morning . . . 103
26.2 The posterior of the state (step 3) . . . 107
26.3 The answers (step 4) . . . 108
26.4 The checks (step 5) . . . 122
26.5 The report (step 6) . . . 124
26.6 The next measurement (step 7) . . . 126
26.7 What to retain . . . 128
26.8 Exercises . . . 128

27 Isolation and reconfiguration 130
27.1 A flag at half past midnight . . . 130
27.2 The posterior of the state (step 3) . . . 134
27.3 The answers (step 4) . . . 135
27.4 The checks (step 5) . . . 152
27.5 The report (step 6) . . . 153
27.6 The next measurement (step 7) . . . 155
27.7 What to retain . . . 157
27.8 Exercises . . . 158

28 Fail-safe moves 160
28.1 A log line that asks for every fault . . . 160
28.2 The posterior of the state (step 3) . . . 164
28.3 The answers (step 4) . . . 165
28.4 The checks (step 5) . . . 176
28.5 The report (step 6) . . . 177
28.6 The next measurement (step 7) . . . 178
28.7 What to retain . . . 181
28.8 Exercises . . . 182

A Atlas of control questions 183

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.
Learning while controlling (Volume Three)		
m, s	varies	Mean and spread of the belief about the input's gain on the one-node plant (Chapter 22).
$V_j(\cdot)$	varies	Cost to go with j moves left, the value of the recursion on a belief (Chapter 22).
δ_w	K	Deviation of the chilled water from its stated T_w, an entry of the state when T_w is released (Chapter 22).
K_c	1/K	Critical gain of the loop, the gain at which it reverses into oscillation (Chapter 23).
β_k, b_r	–, K	Fraction of the return recirculated to rack k's inlet; offset of the return sensor (the two-unit hall of Chapter 11).
$d_\bullet$	K	Offset of one named sensor, an entry present only where a row names it (Chapter 25).
λ	–	Forgetting factor of the filter, which divides the whole covariance by λ at each hour (Chapter 25).
Faults and reconfiguration (Volume Three)		
T_{sp}, r_h, u_h	°C, °C, –	Chiller set point of the two-hall plant; the return set point and the valve of hall h (Chapter 18).
$T_{r,h}$, $T_{\mathrm{in},k}$	°C	True return air of hall h; inlet of rack k (Chapter 18).
$Q_{t,h}$, G_h, $\sum C_h$	kW, kW/K, kW/K	Heat of hall h; its open-valve conductance; its airflow capacity (Chapter 18).
g_u, b_h	–, K	Log-derate of coil unit u; offset of the return sensor of hall h (Chapter 18).
$\mathbf{p}$, $\mathbf{m}$, $\mathbf{\Lambda}$	varies	What persists across windows, (g_u, β_k, b_h); the mean and information of the frozen reference, the first hour's posterior (Chapter 26).
$e_{j,k}$, $h_{j,k}$	–	Normalized residual of row j at window k; that row's leverage (Chapter 26).
$C^{\pm}_{j,k}$, S_k, ℓ	–	CUSUM of a row, whose allowance is written κ in Chapter 26 and is not the confidence above; a detector's statistic; its level.
FAR, POD, k_0, k^*	–, –, –, –	False-alarm rate per shift; probability of detection; the onset; the first alarm (Chapter 26).

Symbol	Unit	Meaning
π	–	Base rate of a fault per shift (Chapter 26).
M_0, J, δ_j, w_j	–, –, varies, varies	No change; the number of drawings in the family; the added entry of drawing j; its prior width, not a row's weight (Chapter 27).
$\tilde{\mathbf{c}}_{jk}$, $\ln B_j$	varies, –	Projected column of drawing j at window k; its log evidence against no change (Chapter 27).
$\tau(\ell)$	–	The structural twin of the leader ℓ (Chapter 27).
u^*, K_v, T_v, I_v	–, K, s, °C	Target, gain, integral time and integrator of the valve-position loop (Chapter 27).
π_j, $P_j(m)$	–	The ranking's share of drawing j; the probability of the event under it at move m, here not the belief's mean (Chapter 28).
H, H_{in}, $\mathcal{S}$	–	The joint event, inlets and valves; the inlets alone; the set of drawings a worst case reads (Chapter 28).
ℓ_j, s_j	varies	Level and slope of the change of drawing j at the decision window (Chapter 28).
c_v, K h	\$/K h, K h	A stated price on an excursion; kelvin-hours over the inlet limit (Chapter 28).

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 $\mathbf{g}^\top \boldsymbol{\Sigma} \mathbf{g}$.

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

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

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

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

Part VI

Learning While Controlling

Chapter 22

Dual Control

22.1 A self-tuning option on the sheet

The commissioning sheet offers a self-tuning option, priced per year, with one line in its favor: "it probes the plant while it runs, so it keeps learning". The trend beside it shows the valve within 0.01 of steady for a quiet hour. Chapter 7's example said that the lean of a move toward learning was worth under 0.1 %, over two moves.

The engineer would conclude one of two things: that two moves were too few and the option must pay over a season, or that it never pays. The sheet supports neither, and two moves cannot separate them. The question of this chapter is

which move holds the limit and narrows $\mathbf{\Sigma}$*, and when that is worth anything.*

Jargon: dual effect, dual control, active learning, probing, passive learning, cautious control, turn-off, lean, stroke test, two-point belief, perfect reference

The *dual effect* is a move changing the preposterior: what the move does to the plant and what it does to the future $\mathbf{\Sigma}$ are one choice. *Dual control* is the policy that keeps both, the second term of the recursion (7.3). *Active learning* or *probing* is a move chosen partly for what it will teach. *Passive learning*, using the reading, is updating from readings that the move did not choose to produce. *Cautious control* is the move from the recursion with the belief held, re-solved at each step on the updated belief; its *turn-off* is its shrinking to zero when the band swamps the mean gain, so that nothing is learned. The *lean* is the excess of the cautious route's cost over the exact route's, measured by paired draws. A *stroke test* here is a restriction of the recursion: one stated move, then ordinary control. A *two-point belief* is a prior on a gain as two hypotheses, with an exact recursion on (x, p). The *perfect reference* is the cost with the gain told, the physical limit each route is read against.

Key idea

On this plant's own posterior, at its own band, no case was found in which a move made in order to learn pays, and *pays* has two readings here. In dollars there is nothing beyond a payback measured in seasons, at a worth under 0.004 \$/h, on rows whose offsets are not trusted. In capacity there is one move, a set-point test held near the stop, worth 0.42 kW of the 55 kW window over the same steps held still with T_w pinned and 0.20 kW released, with no price on either side; one reference reading adds more. A move teaches only an entry whose column depends on the move, and the entries that hold most of the band have no row at any move. Where a column does depend on a move, ordinary operation already excites it, as long as the rows carry no persistent offset. Learning carries the problem where the cautious route shuts the loop off, and the one such case found is constructed: a belief that the valve's action might be either way, which no fit of this plant carries, and which is a commissioning question with a stroke test for an answer. None found is not none exists; an unstable open loop was run only at $a = 1.05$ on stated beliefs.

22.1.1 The restriction this chapter works under

This chapter lies on the row "The second term kept" of Table 7.6. No coupling of Section 7.1.3 is restricted, and the future $\mathbf{\Sigma}$ is chosen. Horizon and plant class stay small, and this chapter says how small: the recursion is exact only on a one-node reading and a two-point belief. Coupling 3 is what the chapter measures. Section 7.3.4 closed with a debt, a horizon or a plant in which the second term pays, or a plain statement that it does not, and this chapter settles it.

22.1.2 Four plants, smallest question first (step 0)

The two-unit hall gives zero at once, by an identity. An honest case needs smaller plants, so the chapter uses four, in the order of the questions. None of them adds a new plant to the book.

The two-unit hall. The node of Chapter 16 with the state of equation (11.4), on the reduced state of Section 4.2.8. Closed loop: r pinned, T_s free, the controller's row (11.3) present; a stop pins the actuator again.

The same hall with a valve row. The row is the two-hall plant's form on this hall, with a valve reading of 0.01:

$$u = \frac{Q + Q_o}{\sum G\,(r - b_r - T_s)}. \tag{22.1}$$

The test pins T_s below the stop and pins r, with the valve free: closed at the hall. The inlets of equation (11.6) do not contain T_s, so the move holds the limit by construction, and it costs energy.

The closed-loop hall. The hall of Chapters 5 to 15, with the offset b of Chapter 6 and a deviation δ_w of the chilled water from its stated T_w as a fifth entry,

$$\mathbf{z} = (Q,\ g,\ g_w,\ b,\ \delta_w)^{\mathsf{T}}, \qquad T = r - b + \frac{Q}{G}, \qquad u = \frac{Q}{W\,(r - b - T_w - \delta_w)}, \tag{22.2}$$

with G and W the two conductances of equation (5.6). Closed loop: r pinned at the test's value, valve free. This hall has no chiller and no price: T_w is fixed and the coil's duty is fixed by the load.

One node. Chapter 7's reading of the loop, equation (7.4), with x read exactly and costs (7.5), $q = 4.00$ and $\rho = 25.00$. Two speeds carry the same $K = 7.0996\,\mathrm{K}$ per unit: $a = 0.36788$, $\hat{\theta} = 4.48779$; and $a = 0.95$, $\hat{\theta} = 0.35498$, the chapter's slower plant. The noise sd w is 0.10 K, Chapter 7's, or 0.02 K, Chapter 9's, the latter on two-point beliefs only. This node is a law on a stated plant, with no controller's row: closed in the sense that the second move reads the reading, where a fixed sequence would be open.

22.1.3 The control engineer's questions (step 1)

Table 22.1 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 these plants, the fourth its category, the fifth the computation and the last the subsection.

Table 22.1: The questions asked of a move made in order to learn. Each what-if states its node in Section 22.1.2.

Control engineer's question	General form	Mathematical form	Category	Computation	Answered in
Can the set point teach anything?	columns against the move, (2.16)	$\partial\mathbf{J}/\partial r$, T_3^*'s variance by entry	estimation	preposterior	22.3.1
Which move has a line, and what does it add to the readings?	preposterior, no move against move, (2.16)	$\mathbf{\Sigma}$ after n windows, δT_s, joint x	sensor placement	preposterior	22.3.2
What does a set-point move buy at the hall's limit?	window of set points, (2.12)	closure at 0.95, test r, T_w pinned or free	prediction under a threshold	preposterior	22.3.3
Over a long horizon, what is the second term worth?	the recursion, (7.3)	$V_j(x, m, s)$; exact against cautious	design	preposterior	22.3.4
When does a probe start at $x = 0$?	the same	smallest N with $u_0 \neq 0$	design	preposterior	22.3.5
With positive gains or an unknown sign?	the recursion on a two-point belief	$V_j(x, p)$	design	preposterior	22.3.6
Does a bump do it?	restriction of the recursion, (2.23)	one stated move, then cautious	design	expected consequence	22.3.7
What next?	the instrument beside the move, (2.16)	\$/h, kW	sensor placement	preposterior	22.6.1

22.1.4 The evidence and the move as rows (step 2)

The rows of a test window on the two-unit hall. A window i at a pinned T_s adds its own load Q_i and supply $T_{s,i}$ as entries. With $\mathbf{v}_i = (\mathbf{z}^\mathsf{T},\ Q_i,\ T_{s,i})^\mathsf{T}$, its rows, each divided by its standard deviation, are

$$\mathbf{r}_i(\mathbf{v}_i) = \begin{bmatrix} (u(\mathbf{v}_i) - y_{u,i})/0.01 \\ (T_{s,i} - y_{T_s,i})/0.05 \\ (T_{d1}(\mathbf{v}_i) - y_{T_{d1},i})/\sigma_{T_d} \\ (T_{d2}(\mathbf{v}_i) - y_{T_{d2},i})/\sigma_{T_d} \\ (Q_i - y_{Q,i})/3 \\ (T_1(\mathbf{v}_i) - y_{T_1,i})/\sigma_{T_1} \end{bmatrix}, \qquad \mathcal{L} = \tfrac{1}{2}\sum_i \mathbf{r}_i^\mathsf{T}\mathbf{r}_i, \tag{22.3}$$

with $T_{d,u}$ and T_1 of equations (11.5) and (11.6), in which $T_r = r - b_r$, at chapter 11's spreads. At the stop $u = 1$ is known and not read; the first line then has a standard deviation of 10^{-5}. 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 column depends on the move. A move teaches through the line of a row, not through its residual. Equations (11.5) say that on the two-unit hall's closed loop every reading shifts with r and no line changes. Equation (22.1) gives, with $\Delta T = r - b_r - T_s$,

$$\frac{\partial u}{\partial b_r} = \frac{u}{\Delta T}, \qquad \frac{\partial u}{\partial g_v} = -u\,\frac{G_v}{\sum G}, \qquad v = 1, 2, \tag{22.4}$$

so a move of T_s changes the ratio of the b_r column to the g columns. Equation (22.2) does the same with r on the closed-loop hall, because T_w stands pinned across the coil there, where the two-unit hall's supply is free and carries r with it. The one node's row $(x' - ax + \theta u)/\sigma_w$ has the slope u/σ_w on θ, Section 7.1.7. The first case needs the supply free: the two-unit hall's r has no line because T_s moves with it and no row reads a valve. Pin T_s at its stop and the valve row sees $r - b_r - T_s$ again, the same ΔT a T_s move would change. Figure 22.1 draws the five cases. Each row of the figure runs from a move, through the row it reaches, to the column that depends on it, and the last box names the entry that the column can teach.

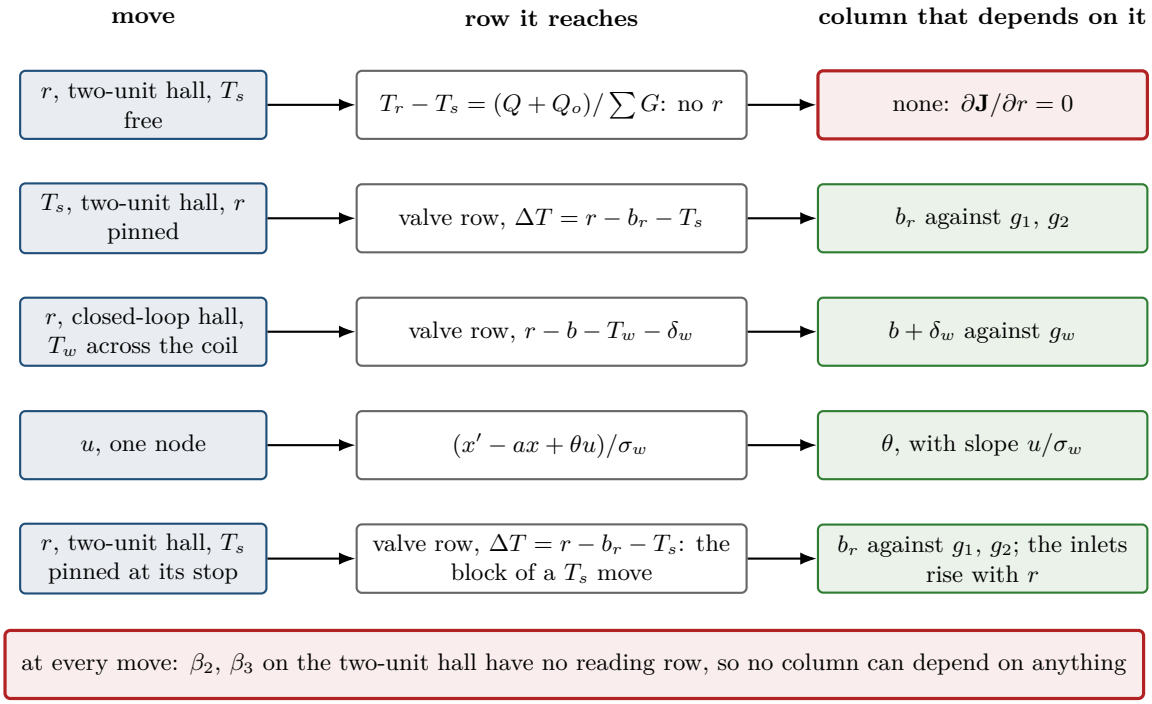


Figure 22.1: Which column depends on which move. Blue: the move. Green: the entries whose column depends on it. Red: no such column. Only a line that changes with the move can make the future $\mathbf{\Sigma}$ a choice.

The modeling choices. Each buys something and can fail.

- *Constant g, β, b_r over a test.* What is learned stays, which helps a test. It fails on a plant whose coils move over a test, and the effect of drift on the comparison was not run: a drifting gain would keep the band open, and it would also leave less time to use what a move taught.

- *Means held.* Rows (22.3) are close to linear in the entries, so $\mathbf{\Sigma}$ does not depend on the readings' values, and on the halls the recursion collapses to a schedule of covariance, a deterministic choice of $(n, \delta T_s)$. It buys closed comparisons. It fails where a reading would move the mean, as Chapter 16's realized figures do.

- *T_w pinned, or a prior of* 0.20 K *on* δ_w. Pinned is the favorable case. The prior is Chapter 15's, a choice with no source.
- *x read exactly and no stop on the one node.* This buys a recursion on three numbers. Its cost is that a real bump is limited by a range the node does not have.
- *A normal prior on θ.* Chapter 7's, with a closed update. It gives the wrong sign real weight at the wide end, and the table says how much.
- *A two-point prior on θ.* Chosen because it is exact and because it removes the wrong-sign weight when both gains are positive. Its factor of four and its sign are stated scenarios with no source, and they are a device.
- *Prices.* Chapter 16's 23:00 window at 0.07 \$/kWh on the two-unit hall. None on the closed-loop hall: its worth is read in kW, and no price is borrowed onto it.

22.2 The posterior of the state (step 3)

Each plant has one belief, and every answer below reads it. Two are fits already derived, cited here with what the moves will act on. The other two are one-entry beliefs and have an exact update.

22.2.1 The two-unit hall's fit

The fit is Chapter 16's, with $\hat{\mathbf{z}}$ of equation (16.7) and $\boldsymbol{\Sigma}$ of equation (16.8). Rack 3's inlet at the node is equation (16.2), whose gradient has the entry $c = 10.4000\,\mathrm{K}$ on β_3 and -1 on b_r on the persistent route. Its variance is therefore read off from three entries of $\boldsymbol{\Sigma}$, with $\Sigma_{\beta_3\beta_3} = 0.6922/c^2 = 0.0800^2$, the survey prior untouched:

$$\mathbf{g}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{g} = c^2\,\Sigma_{\beta_3\beta_3} + \Sigma_{b_r b_r} - 2c\,\Sigma_{\beta_3 b_r} = 0.6922 + 0.0374 - 0.0000\ \mathrm{K}^2. \tag{22.5}$$

A set of test windows adds its information to this fit. Each window's own Q_i and $T_{s,i}$ are eliminated by the Schur complement, leaving

$$\mathbf{F}(\delta) = \mathbf{F}_{zz} - \mathbf{F}_{zb}\mathbf{F}_{bb}^{-1}\mathbf{F}_{bz}, \qquad \mathbf{F} = \mathbf{J}_i^{\mathsf{T}}\mathbf{J}_i \text{ at } T_{s,i} = T_s^{\mathrm{stop}} - \delta, \tag{22.6}$$

with the fit's $T_s^{\mathrm{stop}} = 14.0014^{\circ}\mathrm{C}$ at the window's load. The two arms of the comparison are then

$$\boldsymbol{\Sigma}_{\mathrm{hold}} = \left(\boldsymbol{\Sigma}^{-1} + 2n\,\mathbf{F}(0)\right)^{-1}, \qquad \boldsymbol{\Sigma}_{\mathrm{move}} = \left(\boldsymbol{\Sigma}^{-1} + n\,\mathbf{F}(0) + n\,\mathbf{F}(\delta)\right)^{-1}, \tag{22.7}$$

preposterior covariances in the sense of equation (4.73). The arm with no move holds $2n$ windows at the stop, against n there and n at the move. A window is 5 min.

22.2.2 The closed-loop hall's fit

The fit is Chapter 6's: $\hat{Q} = 30.0581\,\mathrm{kW}$, $\hat{g} = 0.0183$, $\hat{g}_w = -0.1676$ and $\hat{b} = 0.0571\,\mathrm{K}$, with spreads 0.1580 kW, 0.0421, 0.0485 and 0.1899 K. The fifth entry adds the prior row δ_w/σ_w, at the limit of $\sigma_w \to 0$ when T_w is pinned, and otherwise at 0.20 K; released, $\mathrm{sd}\,b$ is 0.1902 K. A test of n steps at $(r_t, 30\,\mathrm{kW})$ adds means of the hall, valve and load readings, with the test's own load eliminated as in equation (22.6), and so

$$\boldsymbol{\Sigma}_a = \left(\boldsymbol{\Sigma}^{-1} + \mathbf{F}(r_t, n)\right)^{-1}. \tag{22.8}$$

22.2.3 The one node's belief on a normal prior

With x read exactly, the state is the one number θ. Its belief is (m, s), and the first step's row and the prior add to the update (7.10), exact because the row is linear in θ. Written for one step, with $\varepsilon = x' - (ax - mu)$ the innovation and $v = u^2 s^2 + \sigma_w^2$ its variance,

$$\frac{1}{s'^2} = \frac{1}{s^2} + \frac{u^2}{\sigma_w^2}, \qquad m' = m - \frac{s^2 u\, \varepsilon}{v}. \tag{22.9}$$

The hall's own band comes from the closed-loop hall's fit. The process gain at 30 kW is a derived quantity of the state, with $\hat{K} = W(r - b - T_w)^2/Q = 6.9982\,\mathrm{K}$ per unit and, by $\mathbf{g}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{g}$, a spread of 0.3377. So $s/\hat{\theta} = 0.0483$ on this plant, and it stays flat toward the stop, 0.0486 at both 45 and 55 kW. The wider values of the tables below, 0.50 to 1.50, are stated scenarios, and they are not this plant's belief.

22.2.4 The two-point belief

Let $\theta \in \{\theta_1, \theta_2\}$ with $p = \Pr(\theta = \theta_2 \mid x_{0:k})$. Now the belief with x is (x, p), carried as the log-odds $\lambda = \ln(p/(1-p))$, and the row of the step gives the exact update

$$\lambda' = \lambda + \tfrac{1}{2}(d_1^2 - d_2^2), \qquad d_i = \frac{x' - (ax - \theta_i u)}{\sigma_w}. \tag{22.10}$$

Its mean and variance are $m(p) = (1-p)\theta_1 + p\theta_2$ and $s^2(p) = p(1-p)(\theta_2 - \theta_1)^2$. At $u = 0$ the two d_i are equal and the belief does not move. Every two-point belief below starts at $p_0 = 0.5$ unless a row says otherwise. The scenarios are θ_i of $\{\frac{1}{2}, 2\}\hat{\theta}$ and $\{\frac{1}{4}, 4\}\hat{\theta}$, positive gains, and $\{-1, +1\}\hat{\theta}$, the unknown sign. That last case is constructed: no chapter's fit of this plant carries such a belief, since Section 22.2.3's 0.0483 places the sign at about twenty spreads from zero.

22.3 The answers (step 4)

The first three answers read the fits of the two halls through preposterior covariances. The next four read the one node's beliefs through the recursion.

22.3.1 The set point has no line

Control engineer's question. Can the set point teach anything?

General form. The preposterior (2.16), with u the band left on a protected quantity.

Specialized form. Equation (2.16) with a the set point r at the closed node. By equation (4.73) the band depends on the candidate only through its lines. On equations (11.5), $T_r - T_s$, the drops $T_r - T_{d,u}$ and the meters contain no r, so

$$\frac{\partial \mathbf{J}}{\partial r} = \mathbf{0}, \qquad \boldsymbol{\Sigma}_a(r) = \boldsymbol{\Sigma}_a \text{ for every } r, \tag{22.11}$$

and the second term of the recursion (7.3) is the same at every set point. It adds nothing to the choice. What remains is how much of the band any move could reach at all, read from equation (22.5).

Building blocks. $\boldsymbol{\Sigma}$ and the split (22.5); on the fresh route, the error of equation (16.4).

Table 22.2: The variance of T_3^* by entry, on the two routes of equation (16.4).

Entry	Persistent [K^2]	Fresh [K^2]	Could a move reach it
β_3, unread	0.6922 (94.9 %)	0.6922 (88.5 %)	no: no row at any move
b_r	0.0374 (5.1 %)	–	only with a line, Section 22.3.2
Fresh error	–	0.0900 (11.5 %)	no: drawn at the node
Cross term	−0.0000	–	–
Total	0.7297	0.7822	at most 5.1 %, persistent; 0, fresh

The plant's numbers. The largest change in any line of $\mathbf{J}$ for r moved by +2 K is 3.55×10^{-9}, which is numerical. Table 22.2 splits T_3^*'s variance. Each row is an entry; the columns give its variance on each route, with its share, and whether any move could reach it.

With the supply free, a set point teaches nothing at this node, by an identity, and the second term is zero. This is the limit Chapter 16 stated, with its reason, and no horizon changes it. The qualifier carries weight. With T_s pinned at its stop and a valve read, an r move has the block of a T_s move, to within 3×10^{-7}, and the next subsection takes up that row.

22.3.2 A move with a line, against the readings

Control engineer's question. Which move has a line, and what does it add to the readings?

General form. The preposterior (2.16), for a move against no move.

Specialized form. Equation (2.16) with the arms (22.7) and u the joint set point (16.12) on the persistent route, means held, priced at the slope of equation (16.16). The move's worth is

$$\Delta x = x_{95}(\mathbf{\Sigma}_{\text{move}}) - x_{95}(\mathbf{\Sigma}_{\text{hold}}), \qquad \text{payback} = \frac{\text{energy of the test}}{\Delta x \cdot \text{slope}}, \tag{22.12}$$

and its reference is b_r told, $\mathbf{\Sigma}$ with the entry's row at infinite precision. Taking the means as held enters here, and it flatters every row of the comparison alike. The arm with the T_d and T_1 rows out asks what is left if those offsets are not trusted, which Chapter 25's season of averaging gives reason to ask. It is the honest upper case.

Building blocks. $\mathbf{\Sigma}$ of Section 22.2.1; $\mathbf{F}(\delta)$ of equation (22.6); x_{95} before any test, 31.0732°C; the slope 0.1199 \$/h per K at 14.3°C; a test's energy from the chiller row at 0.07 \$/kWh.

The plant's numbers. Table 22.3 gives the arms. Each row is a set of rows, a count n and a move δT_s. The next two columns give sd b_r without and with the move. The rest give the joint gain of holding over no test, the move's gain over holding, its \$/h, the test's energy and the hours to pay that back. The valve in the test stands at 0.8892 and 0.8005 for moves of 2 and 4 K. The chiller's floor of 8°C makes 6.0 K the largest legal move, so the table stops there and the 8 K row is kept only as a marked reading below the floor.

The set point moves 0.0369 K **if** b_r **is told, and the rows already there take** 0.0342 K **of it in** 24 **windows, two hours, with no move.** A held load makes the T_d rows read b_r: their line on it is −1 at every T_s. After those two hours sd b_r is 0.0511 K, 0.0699 K after one, and the move adds +0.000050 K at most on the rows run inside the chiller's range. That reading assumes the rows' errors are independent from window to window.

Give each of the T_d and T_1 rows a persistent offset, a stated scenario with no source, and windows no longer average it out. At 288 windows and 4 K an offset of spread 0.05 K leaves the

Table 22.3: A move of T_s with r pinned, against holding, on the persistent route with means held. The hold is $2n$ windows at the stop; the move is n there and n at δT_s. The row marked † breaches the chiller's floor: its energy is unclipped, and it is not a test the chiller can run.

Rows	n [–]	δT_s [K]	sd b_r, hold [K]	sd b_r, move [K]	Hold over no test [K]	Move over hold [K]	Worth [\$/h]	Energy [\$]	Payback [h]
All	12	6	0.0511	0.0506	+0.03425	+0.000050	0.000006	0.7362	1.2×10^5
All	48	6	0.0263	0.0260	+0.03617	+0.000014	0.000002	2.9447	1.8×10^6
All	2016	4	0.0041	0.0041	+0.03684	+0.000000	0.000000	81.87	3.5×10^9
T_d, T_1 out	12	2	0.1715	0.1720	+0.00784	−0.000173	−0.000021	0.2420	never
T_d, T_1 out	48	4	0.1706	0.1498	+0.00812	+0.006526	0.000782	1.9493	2490
T_d, T_1 out	48	6	0.1706	0.1369	+0.00812	+0.010174	0.001220	2.9447	2410
T_d, T_1 out	48	8†	0.1706	0.1270	+0.00812	+0.012740	0.001527	3.9544	2590
T_d, T_1 out	288	6	0.1704	0.0788	+0.00820	+0.022469	0.002694	17.67	6560
T_d, T_1 out	2016	4	0.1703	0.0429	+0.00822	+0.026802	0.003213	81.87	25 500

move +0.00015 K; 0.10 K leaves it +0.0013 K, 3.6 % of what b_r told would give; 0.30 K leaves it +0.0111 K, 30.2 %. So ordinary operation excites b_r on condition that the rows carry no large persistent offset.

Take the rows out altogether and the move has something to add, at most +0.0268 K and 0.0032 \$/h. Inside the chiller's range its shortest payback is 2410 h, at $n = 48$ and 6 K, right on the floor, and 2490 h at 4 K: 0.84 and 0.87 of Chapter 25's 2880 h season, which leave 0.57 and 0.30 \$ at the end of that season. The 1930 h of an 8 K move read a line at 6°C and priced energy clipped at 8°C; unclipped it is 2590 h.

The set point again, with the supply pinned. Pin T_s at its stop and raise r by δ, valve free. The information is that of a T_s move of δ, and no chiller energy is spent, so this is the arm a cheap test would reach for. With the rows out, 48 windows give +0.0021 K at +2 K and +0.0065 K at +4 K; 288 windows at +2 K give +0.0112 K, 0.0013 \$/h. With all rows the worth is under 0.00001 K. What it costs is the inlet. Rack 3's mean goes from 24.98°C to 26.98°C at +2 K, against the 27°C limit, with $P(T_3 \le 27°\mathrm{C}) = 0.508$, and to 28.98°C at +4 K, with 0.008. A test whose hall temperature sits on its limit one time in two does not hold the limit, so this row teaches and is not free.

A move can lose. In the row of $n = 12$ and 2 K with the rows out, sd b_r widens from 0.1715 to 0.1720 K. A window at the move gives up a window at the stop, whose valve is known and not read, and a 2 K lever does not repay that. Told the valve in the test, the same arm stands at 0.1718 K, so the valve reading's 0.01 is part of the loss and not all of it.

Figure 22.2 sets these beside the instrument of Section 22.6.1. Its horizontal axis is the worth in \$/h on a logarithmic scale, and each row is one candidate.

22.3.3 A set-point move at the hall's limit

Control engineer's question. What does a set-point move buy at the hall's limit?

General form. The threshold (2.12) with a joint event, as a window of set points, read over a preposterior covariance.

Specialized form. Equation (2.12) as the joint (6.20) on $T \le 27.5°\mathrm{C}$ and $u \le 1$, from equation (22.2), with $\mathbf{\Sigma}_a$ of equation (22.8). The answer is the load at which that window closes,

$$Q_{\mathrm{cl}}(\mathbf{\Sigma}_a) = \max\Big\{Q : \max_r \ \Pr\big(T^*(r,Q) \le 27.5°\mathrm{C},\ u^*(r,Q) \le 1 \mid \hat{\mathbf{z}}, \mathbf{\Sigma}_a\big) \ge 0.95\Big\}. \qquad (22.13)$$

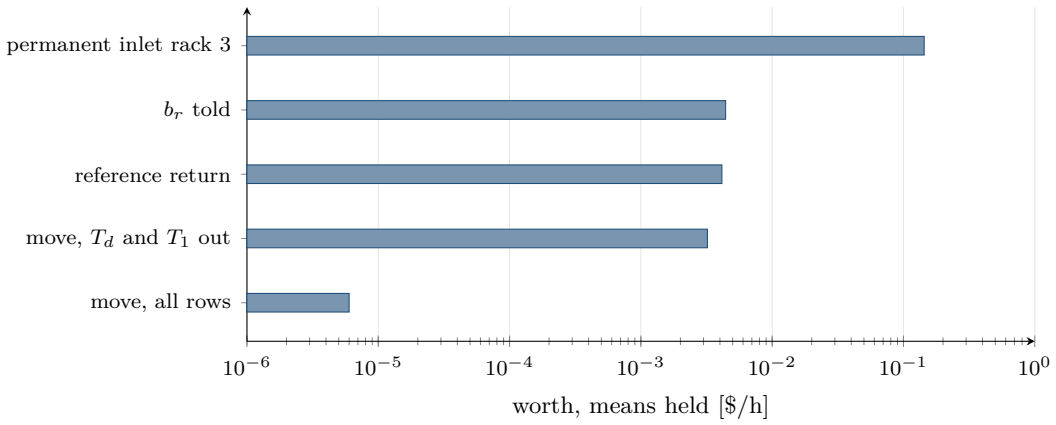


Figure 22.2: What a move could reach against the instrument, on the persistent route with means held. The moves are their largest worth in Table 22.3. Told b_r is the most any move could give on this route, and it is not a route a controller can take.

This is not a price. The physical reference is the same limit at the truth with no band. Holding the means is a choice entering here, so a closure above that reference ranks rows and claims no load, as Chapter 15 said of its own. Here r does have a line, equation (22.2), where it had none in Section 22.3.1. Its valve row reads $b + \delta_w$, and pinning T_w gives all of that sum to b.

Building blocks. $\hat{\mathbf{z}}$ and $\mathbf{\Sigma}$ of Section 22.2.2; $\mathbf{\Sigma}_a$ of equation (22.8); the closure (22.13) by bisection over Q.

The plant's numbers. The test $r = 20°\mathrm{C}$ for 40 steps at 30 kW reproduces Chapter 10's row: $\mathrm{sd}\, b = 0.0902$ and $\mathrm{sd}\, T^* = 0.0989\,\mathrm{K}$ against 0.0901 and 0.0987 K. Table 22.4 gives the closures. Each row is a treatment of T_w and a test; the columns give the two spreads, the closure, the probability that the hall stays under 27.5°C during the test, and the gain over the hold of the same steps, the arm with no move.

The first thing the table shows is that holding is not nothing. The same 40 steps with r left alone already open the window by 0.26 kW, because they narrow the load, $\mathrm{sd}\, Q$ from 0.158 to 0.123 kW; only the move narrows b. So the move's own gain is 0.42 kW with T_w pinned and 0.20 kW released, where $\mathrm{sd}\, b$ stays near 0.15 K because the sum $b + \delta_w$ is what is read. Set against no test the same rows would read 0.68 and 0.45 kW, which credits the move with what holding gives.

Two more readings qualify those figures. The move takes 131 steps, 65.5 min, with its 91 steps of settling; a hold of the same time reads 131 steps and closes at 55.410 kW, so at equal time the move is ahead by 0.26 kW pinned and 0.05 kW released. And the means are held. With a truth drawn from the fit, the readings drawn for each arm, the mean moved by the linear update and the closure read at the moved mean, 300 paired draws give the move over the 40-step hold +0.41 kW (se 0.04) pinned and +0.13 kW (se 0.03) released; over the 131-step hold, +0.21 kW (se 0.04) pinned and −0.05 kW (se 0.03) released. One draw's closure spreads 0.6 to 0.7 kW. So pinned, the move's gain survives each reading at 0.2 to 0.4 kW; released and at equal time, it is not told from zero.

The reference is read on the same basis. On its own it closes at 55.442 kW pinned and 55.453 kW released, below the move's 55.667 and 55.463 kW, the second by a margin finer than the bisection's 0.01 kW. But a reference does not stop the hall, and added to the same 40-step hold it closes at 55.807 kW on both treatments: 0.56 and 0.55 kW over the hold, against the move's 0.42 and 0.20 kW; means free, +0.50 kW (se 0.04) and +0.44 kW (se 0.04). Told b

Table 22.4: The load at which the joint window closes at 0.95, means held. The hold is the same 40 steps with r left at 18°C. The physical limit at the truth is 55.904 kW; a closure above it ranks a row and claims no load.

T_w	Test at 30 kW	sd b [K]	sd T^*(45 kW) [K]	Closure [kW]	P in test [–]	Over hold [kW]
pinned	none	0.1899	0.1559	54.991	–	–
pinned	hold, 40 steps	0.1899	–	55.249	1.0000	–
pinned	hold, 131 steps	–	–	55.410	1.0000	–
pinned	r 20°C, 40 steps	0.0902	0.0989	55.667	1.0000	**0.419**
pinned	r 20°C, 120 steps	0.0744	0.0725	55.925	1.0000	–
pinned	r 21.5°C, 480 steps	0.0496	0.0425	56.205	1.0000	–
pinned	one reference reading	0.0484	–	55.442	–	–
pinned	hold, 40 steps, and reference	–	–	55.807	–	0.559
pinned	b told	–	–	55.474	–	–
prior 0.20 K	none	0.1902	0.1560	55.012	–	–
prior 0.20 K	hold, 40 steps	0.1902	–	55.259	1.0000	–
prior 0.20 K	hold, 131 steps	–	–	55.410	1.0000	–
prior 0.20 K	r 20°C, 40 steps	0.1483	0.1148	55.463	1.0000	**0.204**
prior 0.20 K	r 21.5°C, 480 steps	0.1430	0.0794	55.721	1.0000	–
prior 0.20 K	one reference reading	0.0484	–	55.453	–	–
prior 0.20 K	hold, 40 steps, and reference	–	–	55.807	–	0.548

outright the hall closes at 55.474 kW with no hold; the rest is W and G, which any held steps also read.

This is the one place on the actual plant where a move shows a visible worth, in kW near the stop, at a load far above the window's 30 kW, with no price on either side. It is why the statement of this chapter has two readings of *pays*: in capacity this move has a gain, and the statement that no move pays holds in dollars, and holds here only in the sense that one reference reading adds more and no price exists to set against either. The move's cost is warmer air, and its 40 steps are readings on a settled hall. Chapter 23 owns the design of such a test.

22.3.4 The second term over a long horizon

Control engineer's question. Over a long horizon, what is the second term worth?

General form. The preposterior (2.16) inside the recursion (7.3), with both of its terms, on the belief.

Specialized form. Equation (2.16) with u the negative of the cost to go, and equation (7.3) on the belief of Section 22.2.3, j moves to go,

$$V_j(x,m,s)=\min_u\Big\{\rho u^2+q\big((ax-mu)^2+v\big)+\mathbb{E}_{x'}\big[V_{j-1}(x',m',s')\big]\Big\},\tag{22.14}$$

with $x'\sim\mathcal{N}(ax-mu,\ v)$ and (m',s') of equation (22.9). The belief held, $m'=m$ and $s'=s$, gives a closed form, equation (7.25) carried back,

$$u=\frac{S\,a\,m}{\rho+S(m^2+s^2)}\,x,\qquad S=q+P,\qquad P^+=a^2S\Big(1-\frac{Sm^2}{\rho+S(m^2+s^2)}\Big),\qquad c^+=c+S\sigma_w^2,\tag{22.15}$$

starting from $P=c=0$, so that $\Phi_j=P_j(m,s)x^2+c_j(m,s)$. The recursion writes $V_j=\Phi_j+D_j$, with Φ_j tabulated finely and D_j, what learning changes, on a grid linear in x^2. Five routes are

each a policy: *exact*, the second term kept; *cautious*, law (22.15) re-solved on the updated belief, so that the reading is used when it arrives and its arrival is not planned for; *certainty equivalent*, the same with s dropped; *no update*, the prior's cautious gains throughout; and *perfect*, θ told, Riccati at each draw of the prior, the physical limit and the reference of every row. The lean, on paired draws of θ and w, is

$$L = \frac{1}{n} \sum_{m=1}^{n} (J_m^{\text{cautious}} - J_m^{\text{exact}}). \tag{22.16}$$

The recursion's error is larger than most leans, so the leans are quoted by these draws. A policy read from an imperfect value is still a policy, whatever that value's error, so its cost by draws is valid. At $a = 0.95$, $s/\hat{\theta} = 1.50$ and $x_0 = 0$ the recursion gives 8.7233 for the exact route, where draws of that very policy give 9.9183 (se 0.2975). A negative lean on the recursion's table is that error, never a finding.

Building blocks. The update (22.9) and the band 0.0483 of Section 22.2.3; the law (22.15); the recursion (22.14), whose expectation is a nine-point Gauss–Hermite quadrature over one reading, the quadrature of Section 7.3.4; the lean (22.16), on 1000 paths.

The plant's numbers. At the hall's own band the recursion's routes agree to the digit shown on nearly every line of its table, and the exact route stands at most 0.144 % over the perfect reference, on two moves at $a = 0.36788$, falling to 0.038 % at 20 and to 0.011 % at $a = 0.95$, 40 moves. By draws the lean there lies within 0.00010 of zero at two standard errors, on costs of 0.96098 at $a = 0.36788$, $x_0 = 1$, 20 moves; 25.33184 at $a = 0.95$, $x_0 = 1$, 40 moves; and 8.00191 at $a = 0.95$, $x_0 = 0$, 40 moves. **At this plant's band there is no lean to measure.** Using the reading is worth no more: the no-update route stands within the table's last digits of the cautious one.

A bound says the same without any recursion. No policy costs less on average than the gain told, so cautious less perfect is a ceiling on the lean of any policy, whatever its horizon or its plan. On a two-point belief of the plant's own width, paired on 200 000 paths, that ceiling is 0.022 % of the cost at $a = 0.36788$, 40 moves; 0.007 % at $a = 0.95$, for 40 and for 80 moves; 0.007 % at $a = 1.02$; and 0.028 % at $a = 1.05$. The plant's own band holds this answer shut for every plan, and for an unstable open loop of those rates as well.

Table 22.5 gives the wider scenarios by draws. Each row is a plant, a start, a width and a horizon. The columns give the perfect reference, the three routes' costs with their standard errors, the paired lean, the lean as a share of the cautious cost the exact route over the better of the two non-dual routes, and the prior's weight on the wrong sign. Figure 22.3 draws the leans of the slow plant. Its horizontal axis is $s/\hat{\theta}$, the vertical axis the lean, the bars one standard error and the series the start and horizon.

Table 22.5: The routes by draws on the normal prior, 1000 paired paths, seed 2208. Costs dimensionless. Perfect is a closed form over the prior.

a [–]	x_0 [K]	$s/\hat{\theta}$ [–]	N [–]	Perfect [–]	Exact (se) [–]	Cautious (se) [–]	Certainty equivalent (se) [–]	**Lean (se)** [–]	Share [%]	Exact over better non-dual (se) [–]	$P(\theta < 0)$ [–]
0.368	1	1.50	20	1.0066	1.2660 (0.0179)	1.2694 (0.0175)	1.6598 (0.0335)	+0.0034 (0.0016)	0.27	+0.0034 (0.0016)	0.2525
0.95	1	0.50	40	27.3088	27.8688 (0.5153)	27.8843 (0.5192)	27.8918 (0.5191)	+0.0155 (0.0210)	0.06	+0.0155 (0.0210)	0.0228
0.95	1	1.00	40	27.0616	31.0125 (0.7031)	31.6410 (0.7546)	31.5392 (0.7505)	+0.6284 (0.1980)	1.99	+0.5267 (0.1941)	0.1587
0.95	1	1.50	40	24.7112	29.9090 (0.7341)	31.0656 (0.7793)	30.7300 (0.7750)	+1.1566 (0.2526)	3.72	+0.8210 (0.2341)	0.2525
0.95	0	0.50	40	8.1964	8.6190 (0.2002)	8.6159 (0.1981)	8.6181 (0.1976)	−0.0031 (0.0059)	−0.04	−0.0031 (0.0059)	0.0228
0.95	0	1.00	40	8.0587	9.9062 (0.2924)	9.9284 (0.3001)	9.9214 (0.2986)	+0.0223 (0.0691)	0.22	+0.0153 (0.0701)	0.1587
0.95	0	1.50	40	7.4863	9.9183 (0.2975)	10.5867 (0.3368)	10.3380 (0.3235)	+0.6684 (0.1326)	6.31	+0.4198 (0.1142)	0.2525
0.95	0	1.00	80	17.2811	20.9799 (0.5513)	21.0212 (0.5651)	20.9829 (0.5623)	+0.0413 (0.1479)	0.20	+0.0030 (0.1485)	0.1587
0.95	0	1.50	80	15.9951	21.1254 (0.5539)	22.2595 (0.6320)	21.6514 (0.6043)	+1.1341 (0.3002)	5.09	+0.5260 (0.2604)	0.2525

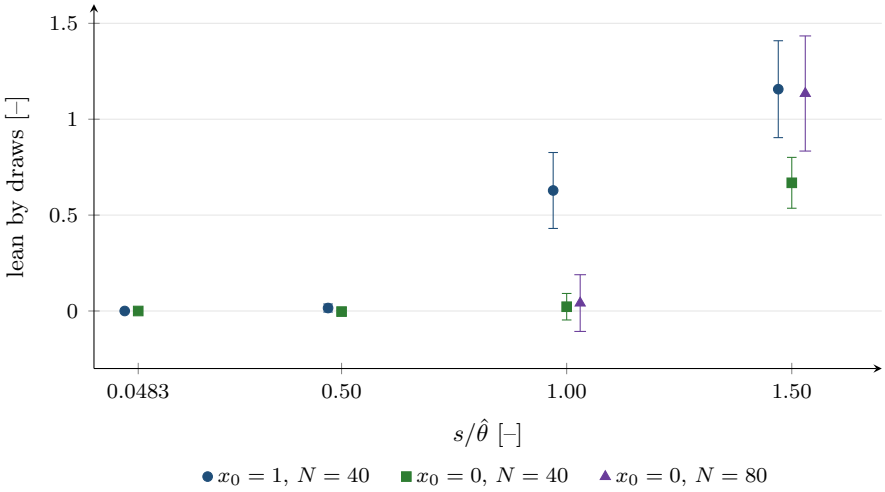


Figure 22.3: The lean by draws on the normal prior, $a = 0.95$, $\mathrm{sd}\, w = 0.10\,\mathrm{K}$, from Table 22.5 and, at 0.0483, the hall's own band. Bars: one standard error. Series are offset slightly along the axis to be read apart. The wide values are stated scenarios, with the prior's wrong-sign weights 0.0228, 0.1587 and 0.2525.

The lean is measured against the cautious route, and at the wide end the certainty equivalent beats that route with no plan at all. So 16 to 54 % of the leans of the wide rows that stand clear of zero need no planned learning: at $s/\hat{\theta} = 1.50$ and $x_0 = 1$, 40 moves, the lean of 1.16 becomes 0.82 (se 0.23) over the better non-dual route; at $x_0 = 0$, 0.67 becomes 0.42 (se 0.11); at 80 moves 1.13 becomes 0.53 (se 0.26). The rest is real. **It appears only where the prior gives the wrong sign real weight, the plant is slow and the horizon is long.** At $s/\hat{\theta} = 0.50$ it is not told from zero, and on the fast plant at the widest prior it is 0.27 %. It does not keep growing as a share: 6.31 % at 40 moves and 5.09 % at 80, within their standard errors. The rise from Chapter 7's 0.06 % to about 6 % needs the slow plant as well as the long horizon, and it leans on a prior that gives the wrong sign 0.2525. The no-update route blows up on such a prior, 5.9×10^{16} at 80 moves and $s/\hat{\theta} = 1$, because on wrong-sign draws $|a - \theta K_j| > 1$. That is the prior's sign weight, as Chapter 7 warned, and it is not a finding about control.

22.3.5 The first probe

Control engineer's question. When does a probe start at $x = 0$?

General form. The preposterior (2.16) inside the recursion, read at its first move.

Specialized form. Equation (2.16) through equation (22.14) at $x = 0$. Replacing (u, x') by $(-u, -x')$ leaves the row's cost, s' and m' of equation (22.9) unchanged, and V_{j-1} is even in x, so the bracket of equation (22.14) is even in u. Then $u = 0$ is stationary, and a probe is a bifurcation to $\pm u^*$ when $u = 0$ stops being the minimum.

Building blocks. The recursion (22.14), scanned finely in u at $(0, \hat{\theta}, s)$.

The plant's numbers. At $a = 0.95$ and $\mathrm{sd}\, w = 0.10\,\mathrm{K}$, for $s/\hat{\theta} = 1.50$, $|u_0|$ is 0 up to $N = 15$, then 0.0102 at 20, 0.0123 at 30 and 0.0129 at 40, with gains of 0.00015, 0.00269 and 0.00532 over $u_0 = 0$. For 0.50 and 1.00 it is 0 at every N to 40. **A probe here is** 0.013 **of valve and its gain is about** 0.05 % **of the cost.** These gains are below the recursion's error, so "starts

near 20" is the recursion's reading and it is stated as such. Chapter 7's "no probing at $x_0 = 0$" is the two-move end of this.

22.3.6 Positive gains and an unknown sign

Control engineer's question. With positive gains or an unknown sign?

General form. The preposterior (2.16) inside the recursion, on a two-point belief.

Specialized form. Equation (2.16) through equation (22.14) with the belief of Section 22.2.4,

$$V_j(x,p) = \min_u \Big\{\rho u^2 + \sum_i \Pr(\theta_i)\Big[q((ax-\theta_i u)^2+\sigma_w^2) + \mathbb{E}_{x'|\theta_i}\big[V_{j-1}(x',p')\big]\Big]\Big\}, \tag{22.17}$$

with p' from equation (22.10). Law (22.15) holds with $m(p)$ and $s^2(p)$, and it gives the cautious route on this belief. This form leaves the wrong-sign objection out when both gains are positive, and it leaves nothing but the sign when they are $\pm\hat{\theta}$. There $m(\frac{1}{2}) = 0$, so the law gives $u = 0$ at every step; with $u = 0$, equation (22.10) leaves $p = \frac{1}{2}$. The certainty-equivalent move and the no-update move are zero for the same reason. All three routes then cost what the uncontrolled plant costs,

$$J_0 = q\sum_{k=1}^{N}\Big(a^{2k}x_0^2 + \sigma_w^2\,\frac{1-a^{2k}}{1-a^2}\Big), \tag{22.18}$$

which is the turn-off of cautious control written out.

Building blocks. The update (22.10); the recursion (22.17), on a fine grid in x^2 and λ; the law (22.15); the closed form (22.18); the lean (22.16) on 4000 paths where it is quoted.

The plant's numbers. Table 22.6 gives the routes on the grid. Each row is a plant, a pair of gains, a start and a horizon. The columns give the costs of perfect, exact, cautious and no update on the grid; then the use of the reading, no update over cautious, the exact route over perfect, and the lean by draws with its standard error where draws were run. The grid's own lean is not a column: Section 22.4 shows it doubled on one row.

Table 22.6: Two-point beliefs, $p_0 = 0.5$, costs on the grid. Lean by draws: 4000 paired paths; – where none were run. The unknown-sign rows are constructed.

a, sd w [–, K]	Gains [$\times\hat{\theta}$]	x_0 [K]	N [–]	Perfect [–]	Exact [–]	Cautious [–]	No update [–]	Using the reading [%]	Exact over perfect [%]	Lean by draws (se) [–]
0.368, 0.10	$\frac{1}{2}$, 2	1	40	1.84412	1.87673	1.87684	1.88872	0.633	1.768	–
0.368, 0.10	$\frac{1}{4}$, 4	1	40	1.95174	1.97769	1.97785	1.99015	0.622	1.330	–
0.95, 0.10	$\frac{1}{2}$, 2	1	40	25.91687	26.06567	26.08203	26.44892	1.407	0.574	−0.0012 (0.0036)
0.95, 0.10	$\frac{1}{4}$, 4	1	40	26.34445	26.56186	26.62036	28.74296	7.974	0.825	0.0312 (0.0045)
0.95, 0.02	$\frac{1}{4}$, 4	1	10	13.85777	14.13503	14.22951	15.92753	11.933	2.001	0.0997 (0.0046)
0.95, 0.10	$\frac{1}{4}$, 4	0	40	7.64213	7.74086	7.76297	8.33574	7.378	1.292	−0.0116 (0.0037)
0.95, 0.10	±1	1	40	25.36862	31.98158	49.08962	49.08962	0	26.067	**17.0782** (0.2973)
0.95, 0.02	±1	1	40	17.83847	20.19098	36.92117	36.92117	0	13.188	**16.6581** (0.0581)
0.95, 0.10	±1	0	40	7.84391	10.75122	12.67547	12.67547	0	37.064	1.8190 (0.1080)
0.368, 0.10	±1	1	40	1.77953	2.32226	2.46926	2.46926	0	30.499	0.1308 (0.0049)

With a factor of 16 between the gains, using the reading is worth up to 12 %, and planning for it little. By draws the part that planning adds at $\{\frac{1}{4}, 4\}$, $a = 0.95$, $x_0 = 1$, 40 moves is 0.0312 (se 0.0045) of 26.6, the grid's cautious cost, which 200 000 draws of the law put at 26.587 (se 0.054). That is about 0.1 to 0.2 %, the grid's own figure being 0.22 %, and

it is not resolved finer, against 7.97 % for using the reading at all. The largest by draws is at $\mathrm{sd}\,w = 0.02\,\mathrm{K}$ and 10 moves, 0.0997 (se 0.0046), 0.69 % of the cautious cost, beside 11.933 % for using the reading. Draws score the stored move read by interpolation, so a lean by draws is what that policy achieved, a lower bound on what the second term could achieve. The two negative values in the table are of this kind: −0.0116 (se 0.0037) says that the stored policy read this way lost to the cautious route at that row, not that learning loses. The grid's leans there are 0.0221 and 0.0164, so these two rows are not resolved between the two figures, and either way they are under 0.3 %. Chapter 7's split survives a long horizon and a positive gain.

The one case where the lean carries the problem is constructed, and it is contrived. At $\pm\hat{\theta}$, $a = 0.95$, $\mathrm{sd}\,w = 0.10\,\mathrm{K}$, 40 moves and $x_0 = 1$, the lean by draws is 17.08 (se 0.30) of 49.2, 34.7 % of the cautious cost, and at $x_0 = 0$ it is 1.82 (se 0.11) of 12.7, 14.3 %. The column for using the reading is 0 there because the cautious move is zero and teaches nothing. Equation (22.18) gives 49.08962 and 12.67547, as the grid does. No fit of this plant carries such a belief: Chapter 5's fit has the sign. The belief asks a commissioning question, whether the valve is wired the way the sheet says.

The turn-off is exact only at p_0 of exactly $\frac{1}{2}$, and that point is an unstable fixed point of the belief. At any other p_0 the mean gain is not zero, so the cautious move is not zero, so it teaches. Table 22.7 gives the sign's lean by paired draws of the stored move against the law as p_0 leaves $\frac{1}{2}$. Each row is a noise and a start; the columns give the lean as a share of the cautious cost at each p_0. At $\mathrm{sd}\,w = 0.02\,\mathrm{K}$ and $x_0 = 1$ the bold 16.66 needs exactly $\frac{1}{2}$: the lean is 4.94 at 0.501, 2.65 at 0.510 and 0.14 at 0.600. At $\mathrm{sd}\,w = 0.10\,\mathrm{K}$ it declines slowly, to 24.8 % at 0.600, so that is the fair case to quote.

Table 22.7: The constructed sign away from $p_0 = \frac{1}{2}$, $a = 0.95$, 40 moves: lean by 4000 paired paths, as a share of the cautious cost.

$\mathrm{sd}\,w$ [K]	x_0 [K]	0.500 [%]	0.501 [%]	0.510 [%]	0.550 [%]	0.600 [%]	0.750 [%]	0.900 [%]
0.10	1	34.5	34.2	33.2	28.9	24.8	15.3	7.3
0.10	0	13.1	13.2	12.5	13.0	11.8	9.8	5.2
0.02	1	45.2	19.6	11.6	4.3	0.7	0.1	−0.0

22.3.7 A bump

Control engineer's question. Does a bump do it?

General form. The expected consequence (2.23) of a restricted policy.

Specialized form. Equation (2.23) with a a first move u_b, thereafter the cautious route, on the belief of Section 22.2.4,

$$J_b(u_b) = \text{bracket of (22.17) at } u = u_b,\ V_{j-1} = V_{N-1}^{\text{cautious}}, \qquad \text{share} = \frac{J^{\text{cautious}} - \min_{u_b} J_b}{J^{\text{cautious}} - J^{\text{exact}}}. \tag{22.19}$$

A restriction cannot beat what it restricts, so the share is at most one. The size of the bump is searched from 0.02 to 0.60, with no stop, which is the node's limit. The bump's use is that it separates the two gains: after it, two possible x' stand $2\hat{\theta}u_b/\sigma_w$ noise spreads apart, by equation (22.10).

Building blocks. The cautious value on the grid of Section 22.3.6; the exact value of the same; the bracket of equation (22.17), in closed quadrature for each size. On the normal prior, the lean (22.16) with a stated move held for 1, 3 or 6 steps.

The plant's numbers. Table 22.8 gives the best bump. Each row is a noise, a pair of gains and a start; the columns give the cautious and exact costs on the grid, then the grid's best size and share, then the best size by draws, how far it separates the gains, and its share by draws. The share by draws takes the paired gain of the bump over the cautious law, 100 000 paths on the sign and 50 000 on the positive gains, over the grid's lean. The draws row is the one to quote: minimizing over a rippled curve picks the quadrature's lows, so the grid's shares run high. Figure 22.4 draws the cost against the size at the unknown sign. Its horizontal axis is the first move on a logarithmic scale and its vertical axis the cost; each solid curve is J_b on the grid at one noise, the circles of its color are the same cost by draws, and the dashed and dotted lines are the cautious and exact costs.

Table 22.8: One stated move, then the cautious route, $a = 0.95$, 40 moves. The unknown-sign rows are constructed. Separation is $2\hat{\theta}u_b/\sigma_w$ at the size by draws. The row marked ‡ is not resolved.

sd w [K]	Gains [$\times\hat{\theta}$]	x_0 [K]	Cautious [–]	Exact [–]	Grid's size [–]	Grid's share [%]	Size by draws [–]	Separation [–]	Share by draws [%]
0.10	±1	1	49.08962	31.98158	0.427	85.1	0.381	2.71	**83.6**
0.02	±1	1	36.92117	20.19098	0.098	99.9	0.110	3.89	**99.7**
0.10	±1	0	12.67547	10.75122	0.193	35.9	0.216	1.54	31.9
0.02	±1	0	0.50702	0.39854	0.039	44.3	0.044	1.57	22.6‡
0.10	$\frac{1}{4}$, 4	1	26.62036	26.56186	0.216	74.0	0.216	–	55.2
0.10	$\frac{1}{4}$, 4	0	7.76297	7.74086	0.020	−52.0	0.020	–	−56.5

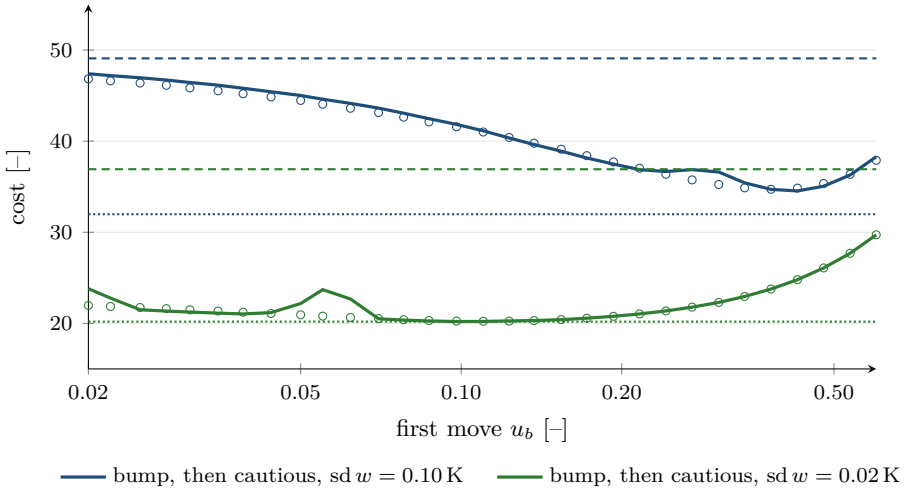


Figure 22.4: The cost of one first move, then the cautious route, at the constructed unknown sign, $a = 0.95$, $x_0 = 1$ K, 40 moves. Solid: on the grid. Circles: by draws. Dashed: the cautious cost, a zero move throughout. Dotted: the exact route. Each color is one noise.

At the sign and at $x_0 = 1$, **one bump recovers** 83.6 % **and** 99.7 % **of the lean, a stroke test's value: one move, then ordinary control.** The figure shows why the size depends on the noise and on the start. From $x_0 = 1$ the bump stops paying once the gains stand about three to four noise spreads apart, 0.381 of valve at sd $w = 0.10$ K and 0.110 at 0.02 K, and beyond

that it only costs ρu^2. The ripple on the grid's curves, largest near 0.055 at $\mathrm{sd}\, w = 0.02\,\mathrm{K}$, is numerical: the nine-point quadrature is taken over a value sharply peaked at $p = \frac{1}{2}$, and the circles show no such rise.

The size that pays at rest is smaller, and a larger one loses money. From $x_0 = 0$ the best bump separates the gains by about 1.5 spreads, 0.216 at $\mathrm{sd}\, w = 0.10\,\mathrm{K}$ and 0.044 at 0.02 K. It breaks even near 0.34 and 0.07, about 2.5 spreads, and past that it costs more than doing nothing: 0.098 at 0.02 K costs 0.603 against a cautious 0.507, and 0.427 at 0.10 K costs 13.83 against 12.68.

The statement is narrower than it sounds. At $x_0 = 0$ one bump recovers 31.9 % at $\mathrm{sd}\, w = 0.10\,\mathrm{K}$, and there the exact route's value is not one move's. The row at 0.02 K and $x_0 = 0$ is not resolved: the grid's first node in x is at 0.125 K while x stands near 0.064 K, so the grid's exact 0.399 is too low. The stored move scores 0.4737 (se 0.0056) by draws, near the bump's 0.4813, and the bump's gain is 22.6 % of the grid's lean and 56 % of the stored move's. No single share is quoted for that row. At $\{\frac{1}{4}, 4\}$ and $x_0 = 1$ the 0.216 is not a test at all: it is the ordinary first move, near the exact route's 0.22132 and below the cautious route's 0.25804. At $x_0 = 0$ the best size on the search is its lower edge, 0.020, and a bump there costs more than it teaches.

On the normal prior the answer depends on the start. At $x_0 = 0$, 40 moves, the 12 bumps of the design gave no bump better than the cautious route past its standard error: the best at $s/\hat{\theta} = 1.00$ is +0.0450 (se 0.0121), and at 1.50 it is −0.0361 (se 0.0569), against a lean by draws of 0.67 (se 0.13); 40 000 paired paths give +0.008 (se 0.002) and −0.003 (se 0.018). Valve bumps of 0.4 held for 6 steps add +46 and +56. At $x_0 = 1$ one larger first move does recover part of the lean: on 40 000 paired paths a first move of 0.35 changes the cost by −0.473 (se 0.036) at $s/\hat{\theta} = 1.50$ and by −0.245 (se 0.024) at 1.00, about two fifths of the leans of 1.16 and 0.63. That move is near the exact route's own first move, +0.267 and +0.219, against the cautious route's +0.226 and +0.245, so it is largely a bolder ordinary move and not a separate test. So no move in this design recovers the normal prior's lean from $x_0 = 0$; from $x_0 = 1$ a bolder first move recovers some of it.

The unstable open loop. Chapter 7's plants are stable, so doing nothing is a fallback. At $a = 1.05$, the same $\hat{\theta}$, it is not. On the two-point grid, $x_0 = 1$, 40 moves, 4000 paired paths, the lean is 0.40 (se 0.03), 0.49 %, at $\{\frac{1}{2}, 2\}$, and 2.63 (se 0.12), 1.25 %, at $\{\frac{1}{4}, 4\}$, larger than the stable plant's figures and still small. At the constructed sign the zero move's cost is 2278 against an exact route's 70.0 on the grid: the turn-off runs away, the case for the stroke test at commissioning is far stronger, and it is still not this plant's belief. At the plant's own band the ceiling of Section 22.3.4 remains 0.028 %. One rate and stated beliefs narrow the gap in the search and do not close it.

22.4 The checks (step 5)

The checks test the two recursions against closed forms, against finer grids and against draws. Their first result decides which table a lean may be read from.

22.4.1 The recursions against closed forms and draws

Table 22.9 lists them. Each row is one check; the columns give what is compared and the two values. The fit's limit $\sigma_p \to 0$ is Chapter 7's stacked row, and it is not repeated.

Table 22.9: The checks on the recursions. Lean figures in the last rows are draws against the grid.

Check	Values	Verdict
Known gain, $s/\hat{\theta} = 0.02$, recursion against Riccati, $N = 2$, 10, 20	0.210576, 0.210438; 0.540927, 0.540778; 0.953857, 0.953694	agree to 2×10^{-4}
Chapter 7's two moves at $s/\hat{\theta}$ 0.10 to 1.50	0.21387, 0.23138, 0.28705, 0.41629, 0.50476 against 0.21387, 0.23138, 0.28704, 0.41565, 0.50437	wide end off by 6×10^{-4}
Coarse grid against finer, $N = 10$, $s/\hat{\theta}$ 0.50: exact; cautious; lean	0.62392, 0.62391; 0.62340, 0.62374; −0.084, −0.028 %	**a negative lean on the table is error**
The same over 0.25 to 1.50	exact to 8×10^{-5}; cautious to 8×10^{-4}	cautious is the weaker
Recursion against draws, normal, $a = 0.95$, $x_0 = 0$, $s/\hat{\theta}$ 1.50, $N = 40$: exact; cautious	8.7233 against 9.9183 (se 0.2975); 9.7368 against 10.5867 (se 0.3368)	not close: leans by draws
Two-point against draws at the sign, $x_0 = 1$: exact; cautious; lean	32.1501 (se 0.1961) against 31.9816; 49.2284 (se 0.4044) against 49.0896; 17.0782 (se 0.2973) against 17.1080	agree
The same, $x_0 = 0$; and $\mathrm{sd}\, w = 0.02\,\mathrm{K}$, $x_0 = 1$: lean	1.8190 (se 0.1080) against 1.9243; 16.6581 (se 0.0581) against 16.7302	agree within 1.3 se
Two-point at $\{\frac{1}{4}, 4\}$, $x_0 = 1$: lean; exact	0.0312 (se 0.0045) against 0.0585; 27.4879 (se 0.3886) against 26.5619	**the grid doubles it: draws quoted**
Order on the whole two-point grid, exact below cautious, $N \leq 40$	largest violation 4.88×10^{-7}	holds
Equal gains against Riccati, $N = 10$, $x_0 = 1$	17.831841, 17.831846	agree
Refined grid at the sign, $N = 20$, $x_0 = 1$: exact; cautious	27.55916, 27.55703; 37.16305, 37.16305	agree
Zero move, equation (22.18), against the grid's cautious route	49.08962, 49.08962; 12.67547, 12.67547; 36.92117, 36.92117; 0.50702, 0.50702	agree
Lean by draws against the grid, positive gains: $\{\frac{1}{4}, 4\}$, $\mathrm{sd}\, w = 0.02\,\mathrm{K}$, $N = 10$; $\{\frac{1}{4}, 4\}$, $x_0 = 0$; $\{\frac{1}{2}, 2\}$	0.0997 (se 0.0046), 0.0945; −0.0116 (se 0.0037), 0.0221; −0.0012 (se 0.0036), 0.0164	first agrees; the others unresolved, small
Lean by draws against the grid, the sign, $a = 0.36788$	0.1308 (se 0.0049), 0.1470	3.3 se apart; draws quoted
Closed-loop hall against Chapter 10's row: $\mathrm{sd}\, b$; $\mathrm{sd}\, T^*$	0.0902, 0.0901; 0.0989, 0.0987	agree
Vectorized closure for the draws against the bisection, pinned: no test; hold; move	54.994, 54.991; 55.246, 55.249; 55.668, 55.667	agree to the bisection's 0.01 kW
An r move with T_s pinned against a T_s move of the same size: largest block difference	2.73×10^{-7}	the same
Cautious cost at $\{\frac{1}{4}, 4\}$, $x_0 = 1$: 200 000 draws of the law; the grid; 4000 paired paths	26.587 (se 0.054); 26.620; 27.519	the 4000 paths drew high: of 26.6, not 27.5

22.5 The report (step 6)

No case was found on this plant's posterior, at its own band, in which a move made in order to learn pays: in dollars nothing beyond a payback measured in seasons at under 0.004 $/h, in capacity one move near the stop, worth 0.42 kW pinned and 0.20 kW released over the same steps held still. With the supply free, a set point teaches nothing, by an identity. With no large persistent offset on them, the readings already teach what a valve move could. A self-tuning option has at most the ceiling of 0.03 % to earn at $s/\hat{\theta} = 0.0483$. The one case that pays much is constructed, an unknown sign, and on it a bump gets most of the value when the plant starts away from zero. None found is not none exists, and an unstable open loop was run only at $a = 1.05$ on stated beliefs.

22.5.1 The answers and their bands

Table 22.10 settles the debt of Section 7.3.4 row by row, the plant first. Each row is a place the search looked; the columns give what a move can teach there, what that is worth and the reference beside it. Table 22.11 gives the answers in the order of Table 22.1, with the truth as its last rows, and Figure 22.5 draws the closures against the physical limit. In that figure the horizontal axis is the closure, each row is a test, the two marks are the two treatments of T_w, and the dotted line is the truth.

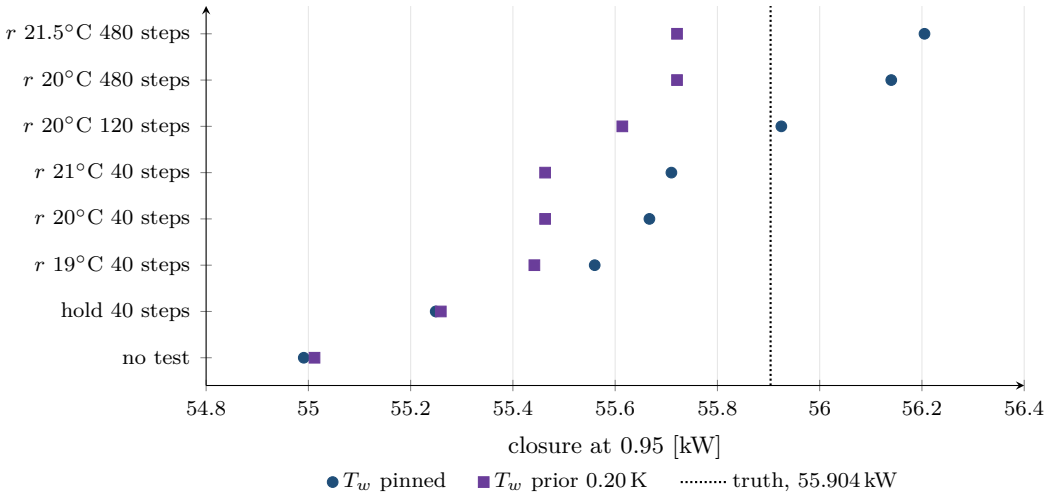


Figure 22.5: The closure against the test, closed-loop hall, means held. Dotted: the physical limit at the truth, no band. Marks to its right rank rows and claim no load.

22.5.2 The entry in the commissioning log

The entry says that this evidence gives no case for paying for probing on this plant, at a band of $s/\hat{\theta} = 0.0483$ with g and β constant; no price for the option was set against it, so it is not a verdict on any offer. It says to run a stroke test when a valve's action or sizing is new, sized by start and noise. At rest the bump separates the gains by about 1.5 noise spreads, $2\hat{\theta}u/\sigma_w$, which is 0.2 of valve at $\mathrm{sd}\, w = 0.10$ K and 0.04 at 0.02 K, and past about 2.5 spreads it costs more than doing nothing; away from rest it separates them by three to four, 0.4 and 0.1.

It says to run a set-point move of 40 settled steps at 30 kW if the load is going near 55 kW, with 65.5 min allowed for it, and it records why that does not contradict the first line. This entry recommends the move for capacity near the stop, where it has a gain of 0.42 kW pinned and 0.20 kW released and no price is set against it; it does not recommend an option that keeps

END OF THE PREVIEW

Control Engineering by Derivation

These are the opening pages. The complete book runs 203 pages.

The full PDF opens with a token. Write to **rmarquez123@gmail.com**, name the book, and I will send you one.

One token opens every book on the site.

Ricardo Marquez, Ph.D.

ricardo-marquez.com