

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.

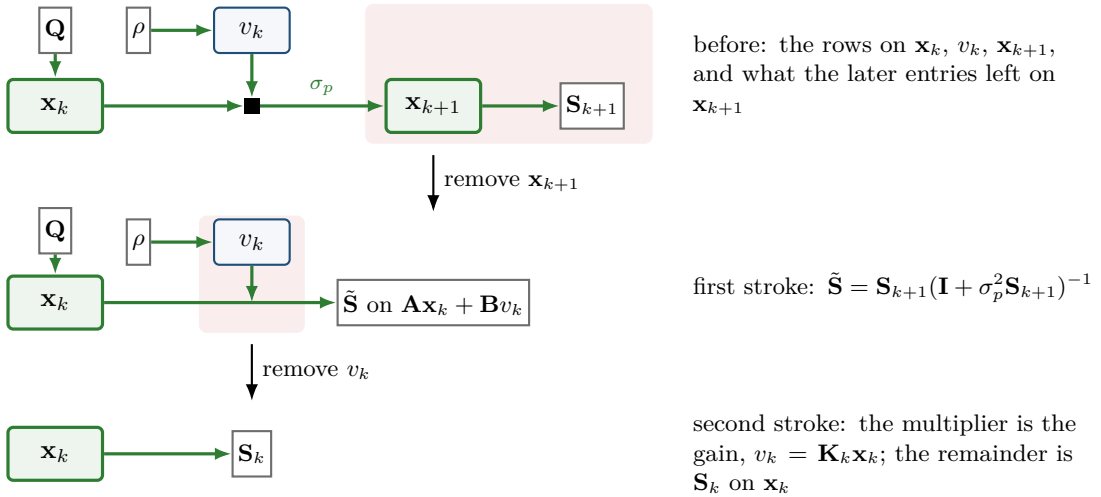


Figure 12.1: The elimination of the stack from the last entry back, equations (12.8) and (12.9). Green: states. Blue: the move. Black square: the physics row at tolerance σ_p. Open boxes: the weights the cost rows and the later entries leave. Shaded: what the next stroke removes. The bottom row has the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

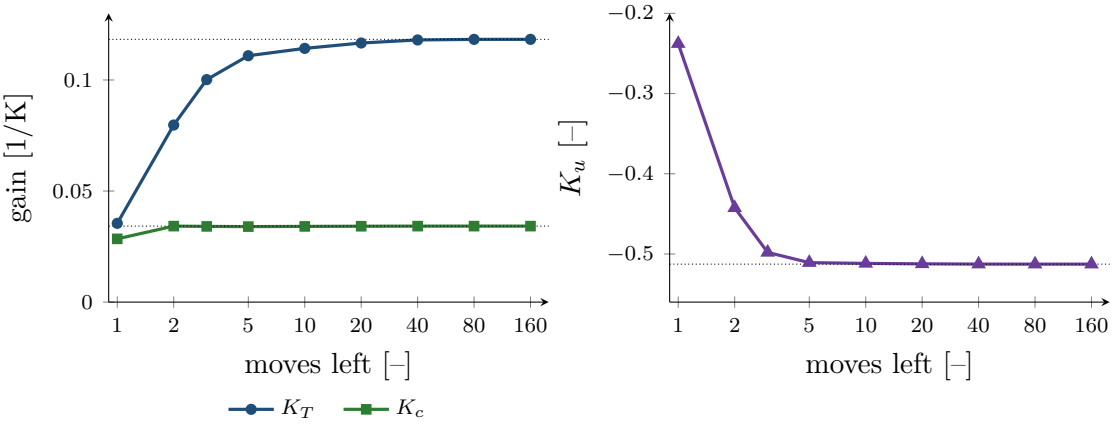


Figure 12.2: The gains of Table 12.2, with 20 and 160 moves left added. Left: the gains on the hall air and the discharge air. Right: the gain on the valve's position. Dotted: the steady values.

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 $s = 1$.

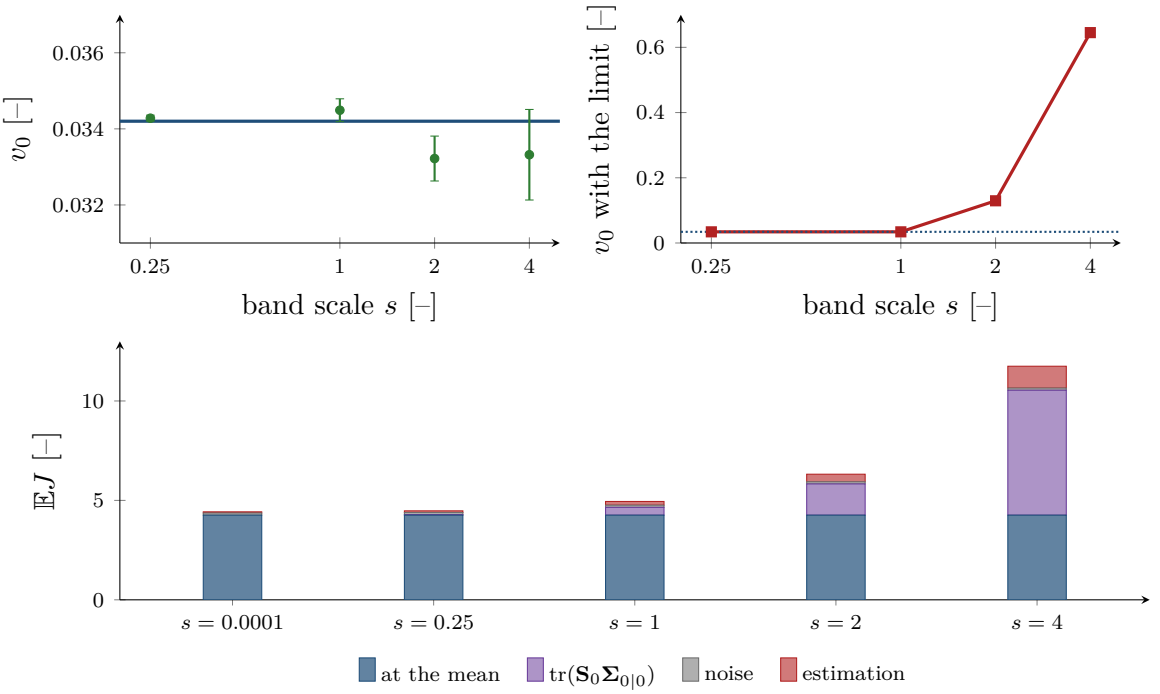


Figure 12.3: Certainty equivalence on the hall. Upper left: the first move against the band's scale, closed form (line) and records (points, one se); the axis is narrow, 0.031 to 0.037. Lower: the terms of equation (12.16) at each s, stacked. Upper right: the first move with the limit of Section 12.3.4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

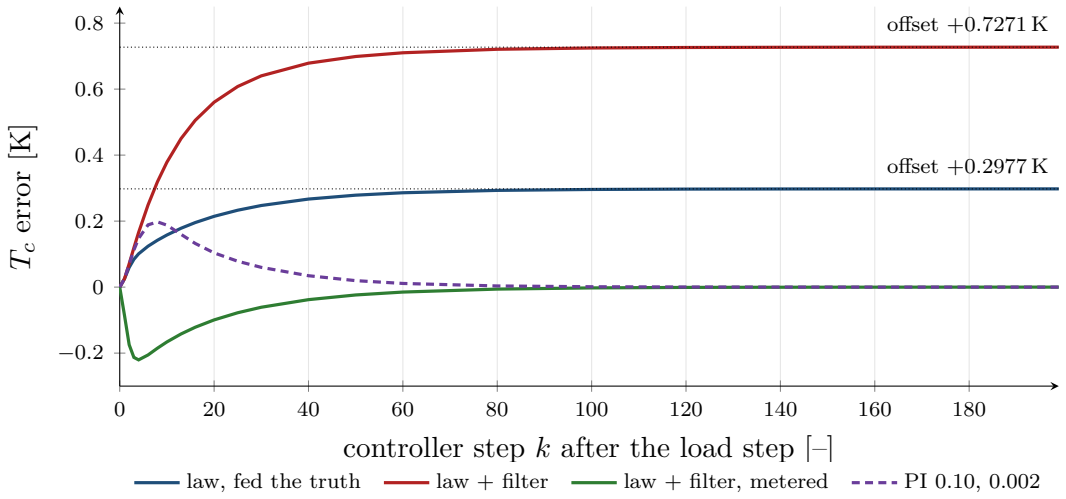


Figure 12.4: The load step from 30 to 45 kW on the plant at the fit, noiseless. Each curve is T_c's error against $r - \hat{b}$. Dotted: the two offsets the law leaves.

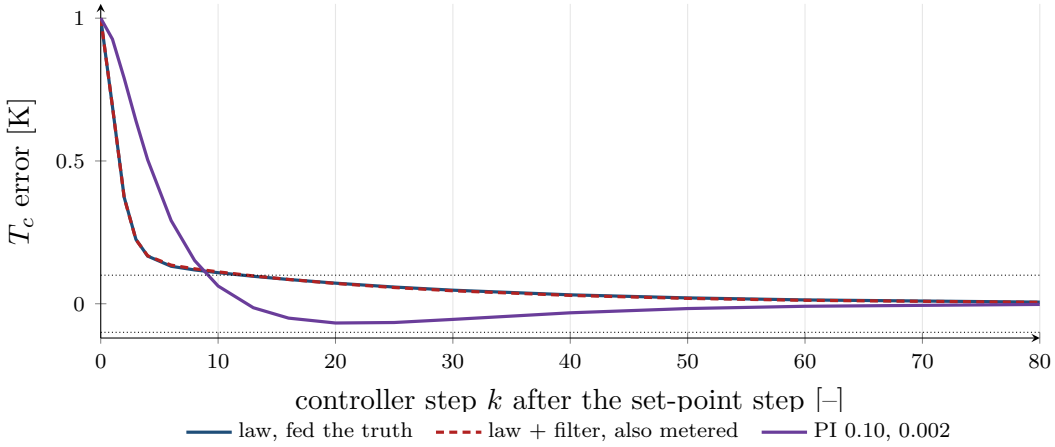


Figure 12.5: The set-point step from 18 to 17°C at 30 kW, same plant, the first 80 of 200 steps. Each curve is T_c's error against the new target. Dotted: the settling band of ±0.1 K. The metered row coincides with the law and filter, because the load does not change.

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.

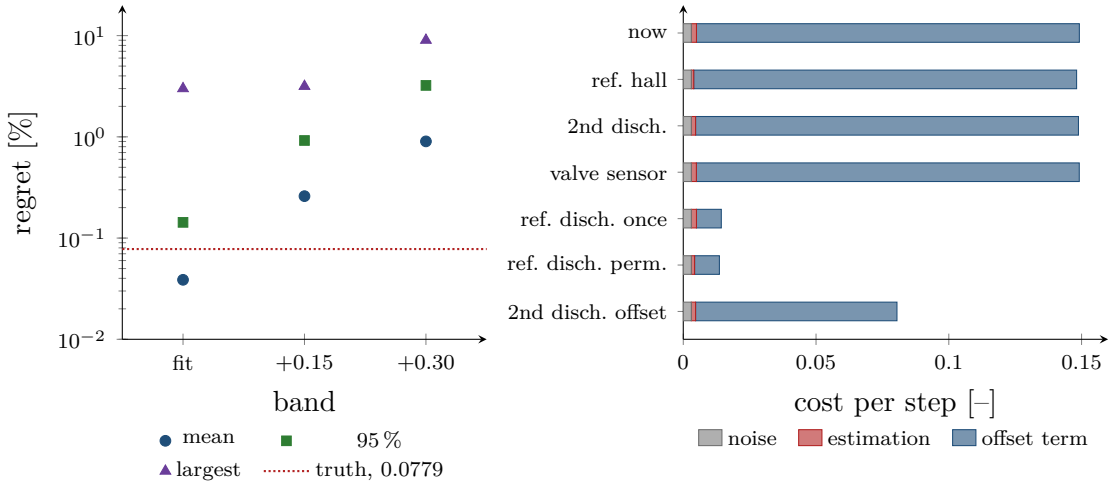


Figure 12.6: Left: the regret of Table 12.11 by band, on a logarithmic axis; the dotted marker is the stated truth. Right: the cost per step of Table 12.15 for each candidate, stacked: the noise term, the estimation term and the offset term.

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,

END OF THE PREVIEW

Control Engineering by Derivation

These are the opening pages. The complete book runs 266 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