INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 19 Control of Biological Systems and Human Performance

CHAPTER 3

Closed-Loop Anesthesia Simulator for Model Inversion with Feedback Correction

Pages 54–85

Edited by John D. Hedengren

How to cite this chapter. "Closed-Loop Anesthesia Simulator for Model Inversion with Feedback Correction," in *Control of Biological Systems and Human Performance*, John D. Hedengren, Ed., vol. 19 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 3, pp. 54–85.

The page numbers in this file are those of the printed volume.

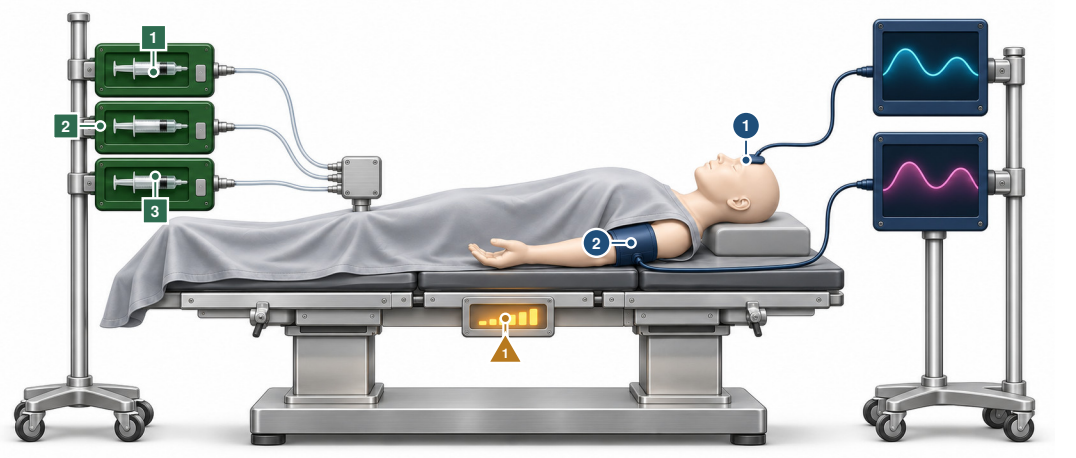
The app for this chapter
https://apmonitor.com/apps/anesthesia/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v19-biological-human/#anesthesia

Copyright © 2026 Advanced Process Solutions, LLC. All rights reserved. Produced by Advanced Process Solutions, LLC.

CHAPTER 3

Closed-Loop Anesthesia Simulator for Model Inversion with Feedback Correction



The controller moves 1 propofol pump, the drug that brings sleep 2 remifentanil pump, an opioid against pain 3 vasopressor pump for blood pressure

The controller reads 1 depth monitor on the forehead 2 blood pressure on the arm

What disturbs it 1 surgical stimulation, up at the incision and down at the closing

Notice that the two drug pumps work from their own estimates of the drug in the brain, and only the forehead monitor shows how deep the sleep really is. Illustration, not to scale.

IN PLAIN TERMS **A Pump That Trusts Its Own Arithmetic**

The main idea. During an operation, two drugs given into a vein keep the patient asleep and free of pain. A pump works out how much drug has reached the brain, and a brain-wave monitor shows how deep the sleep is.

The challenge. The pump's arithmetic fits an average patient, but one patient can need three times the drug of another. The monitor can be fooled into showing a light sleep, and a pump can add drug but never take it back.

Why it matters. Too light a sleep risks memory of the operation, while too deep a sleep or too low a blood pressure does harm.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the history boxes of Section 3.2, the fact boxes of Sections 3.1 and 3.6, the try-it boxes of Section 3.6, and the key points
Study	60 min	The whole chapter with its pause points and the twelve exercises
Practice	25 min	The first page, Section 3.6, Section 3.7, Section 3.8 and the open challenge

The control question An infusion pump holds the drug concentration its population model predicts, which is exact on the patient the model was built from and silently wrong on anybody else. What does an outer loop on a depth monitor add to that model inversion, and which structural choices matter more than its tuning?

Learning objectives

- Compute the rate constants of a three-compartment model from its covariates, and explain why the maintenance infusion falls for hours.
- Evaluate the response surface for a pair of effect-site concentrations, and find the process gain of the depth loop at an operating point.
- Explain why an inner loop closed on the pump's own model cannot detect a patient who differs from the population, and how an outer integral identifies that patient.
- Show by runs why signal validation, a pressure loop and the drug split are structural choices that no tuning replaces.
- Distinguish what a faster loop can and cannot change in the excursion at the end of a case.

Prerequisites Sections 1.4.2 (compartments and model inversion), 1.4.4 (variability between individuals), 1.4.6 (one-sided actuators) and 1.6.2 (velocity form). Section 1.6.3 is needed only for the predictive controller.

The app https://apmonitor.com/apps/anesthesia/

See also Volume 5, Chapter 10 and Volume 5, Chapter 4 for product quality inferred from a model and corrected by a late measurement. Volume 3, Chapter 2 for structure chosen before tuning, and Volume 3, Chapter 3 for how fast an inner loop must be.

3.1 Introduction

Every loop in an introductory course pairs a controlled variable with an instrument that measures it. Intravenous anesthesia has no such instrument. What decides whether a patient is unconscious is the concentration of the drug at its site of action in the brain, and there is no probe for it, no sample to draw and no prospect of either during an operation. The loop is closed regardless. A target-controlled infusion pump carries a published population pharmacokinetic model, feeds that model the infusion it is delivering, and controls the concentration the model predicts [1, 2]. The inner loop that results is exact, immediate and completely blind, because its plant is an equation rather than a patient.

DID YOU KNOW? Sixty Thousand Pumps

By the count of a review published in 2023, about 60 000 target-controlled infusion pumps had been sold in more than 90 countries, and they give intravenous anesthesia to millions of patients each year [3].

The misconception this plant corrects is that an exact loop is a correct loop. With the pump commanded to hold 1.7 µg/mL of propofol and 2 ng/mL of remifentanil, the modeled concentration follows the same trajectory on every patient and sits within 0.01 µg/mL of its target at forty minutes, while the depth index those concentrations actually produce settles anywhere between 52.1 and 70.4 as the patient's propofol requirement moves across the population range. The requirement itself is what varies, by a factor of three between a sensitive and a resistant adult, and no measurement available before induction reveals which one is on the table [4, 5]. Feedback from a real measurement is the only mechanism that finds the difference, and the simulator shows an integral finding it without anybody having told the controller that a parameter was missing.

The structure is general. A fast inner loop closed on a model-derived estimate, corrected by a slow outer loop closed on a real measurement, is the standard answer whenever the variable that matters cannot be measured quickly and a usable model of it exists, and it is the arrangement inferential control was introduced to formalize [6]. Distillation composition inferred from tray temperatures, reactor conversion inferred from a heat balance, a furnace outlet inferred from a firing model and any

modern soft sensor all have the same shape, and all of them fail in the same way when the model is fitted to a population and applied to an individual [7]. The milk powder dryer of Volume 5, Chapter 10 steers its product moisture by such a soft sensor and corrects its bias from a late laboratory sample. Anesthesia is the sharpest instance available because the inner measurement is not merely indirect, it is a calculation, so the inner loop cannot report an error even in principle.

A transfer function or a block diagram of a cascade draws two summing junctions and hides exactly that fact. However, a simulation can put the model concentration and the patient concentration on the same axis and let them separate, and it can move the plant rather than the output, corrupt the measurement in the direction that provokes an overdose, and shut the actuator off at the moment more actuator is not what is needed. This chapter works with a simulation of anesthetic depth in which the model-based inner loop and the feedback outer loop can be switched into service one at a time, so that what each half contributes is measured rather than asserted. It applies the compartment models, the open-loop error of model inversion and the one-sided actuator of Sections 1.4.2, 1.4.4 and 1.4.6 to a full-scale patient.

Section 3.2 traces how infusion pumps and depth monitors came to be combined. Section 3.3 gives the thirteen-state patient, with two independent copies of the pharmacokinetics, one for the patient and one inside the pump, down to every constant needed to rebuild it. Section 3.4 shows that the process gain of the depth loop rises from 8.8 to 18.8 index units per µg/mL as the opioid target moves from zero to 8 ng/mL, and why the maintenance rate never settles. Section 3.5 describes the four loops, the pump's model as a measurement and the optional predictive controller. Section 3.6 quantifies the open-loop error of model inversion against patient sensitivity, works through a thirteen-scenario sequence that separates model error, measurement error and actuator authority, ranks four structural choices by what each is worth on a two-hour case and compares the predictive controller with the closed loop. Section 3.7 describes the real monitor, pumps and controllers, and Sections 3.8 and 3.9 state the limitations and conclude. The pharmacokinetic models and the control algorithm are published ones.

3.2 Automation of Intravenous Anesthesia

Automatic control of anesthesia is older than most of the drugs it is now practiced with. A servo that read the amplitude of the electroencephalogram and moved the delivery of the anesthetic to hold it constant was built and used on patients in 1950, and it worked in the sense that the trace stayed where it was put [8]. What the arrangement lacked was a drug that could be titrated quickly enough to make the loop worth closing, and a measurement whose relationship to consciousness was understood well enough to choose a setpoint.

The problem that stayed was dosing, and it was solved by computing the infusion from a compartment model fitted to a population [1, 2, 9–11], as the first history box below tells. Section 1.4.2 derives the same inversion for two compartments.

> ⧖ HISTORY **From a Bolus to a Target (1992 to 1997)**
>
> An intravenous drug spreads into the tissues while it is being cleared, so the infusion that holds a steady concentration cannot be a steady rate. The first model for automated dosing gave a bolus sized from the target and the volume of distribution, followed by an infusion. Three-compartment models then described how the drug moves between the plasma, well-perfused tissue and fat, and an effect site joined to the plasma by a single rate constant let the target be set where the drug acts. A microprocessor in the pump now computes the rate from the model and the patient's covariates. The name target-controlled infusion appeared in 1992 and was agreed as the generic term in 1997. The first commercial systems entered practice in 1996 with a propofol model whose only covariate was weight [3].

Feedback returned when the measurement improved [12, 13], as the second history box tells, and the combination that this application implements is the one that box ends with, a bispectral index

controller sitting above a target-controlled infusion pump rather than driving the syringe directly. Later work replaced the tuning of that outer loop with a robust design based on the variability of the population rather than on a nominal patient, and evaluated it clinically [5]. Several hundred trials later the method is not standard of care anywhere, and the reasons are the ones this simulator reproduces. The measurement can be corrupted by electrocautery, by muscle activity and by warming devices, and it is corrupted upward, in the direction that makes a light patient of an artifact [14]. The actuator has authority in one direction only. The costs of the two error directions are not comparable, because awareness with recall occurs in roughly one general anesthetic in a thousand and is not undone when the trace recovers [15], while sustained electroencephalographic suppression on the other side is associated with delirium after surgery [16]. A regulator that scores the two directions with one symmetric error integral is answering the wrong question [17]. Equation (1.10) of Section 1.4.6 writes the asymmetric objective that such a loop needs.

> ⧖ HISTORY **Feedback Above the Pump (1987 to 2002)**
>
> In 1987 a team held methohexital anesthesia in 13 volunteers by feedback. The measured variable was the median frequency of the electroencephalogram, with a setpoint of 2 to 3 Hz, and the loop succeeded in 11 of them. The controller rested on a pharmacokinetic and pharmacodynamic model, and its authors saw feedback as a way to adapt population mean data to the individual [12]. A depth index computed from the electroencephalogram came into use in 1996 [14]. In 2002 a closed loop took that index as its controlled variable, a PID algorithm as its controller and a propofol target-controlled infusion system as its actuator. It gave adequate anesthesia in 9 of 10 patients, and in three of them the index oscillated around the setpoint [13].

3.3 Patient Model

The simulated patient is a 40 y old man of 70 kg and 170 cm, and every rate constant below is computed from those covariates rather than typed in. Two drugs are infused. Propofol is the hypnotic and is responsible for most of the depth of anesthesia and for essentially all of the loss of arterial pressure. Remifentanil is a short-acting opioid that contributes to depth mainly through its interaction with propofol and costs very little pressure. A third infusion, a vasopressor, exists only to defend the pressure. Figure 3.1 shows the compartments, the three infusions, the two effect sites, the response surface that produces the depth index, the arterial pressure state and the point at which the monitor is read.

3.3.1 What the Loops See

The controlled variable of the outer loop is the depth index as the monitor reports it, and the controlled variable of the pressure loop is the measured mean arterial pressure. The manipulated variables are the three infusion rates. Between them sits the propofol effect-site target, which is the output of the depth controller and the setpoint of the concentration controller, and which is the only place in the loop where a concentration is ever commanded. Table 3.1 lists the classification. The distinction between the true index and the reported index is load bearing throughout, because the controller acts on the reported one while every performance number in Section 3.6 is computed from the true one.

3.3.2 Pharmacokinetics

Propofol is described by the Schnider three-compartment model with an effect site hung behind the central compartment, as shown in Equation (3.1). The same four equations with the Minto coefficients describe remifentanil. They extend Equation (1.2) by a second peripheral compartment and are written in rate constants rather than clearances.

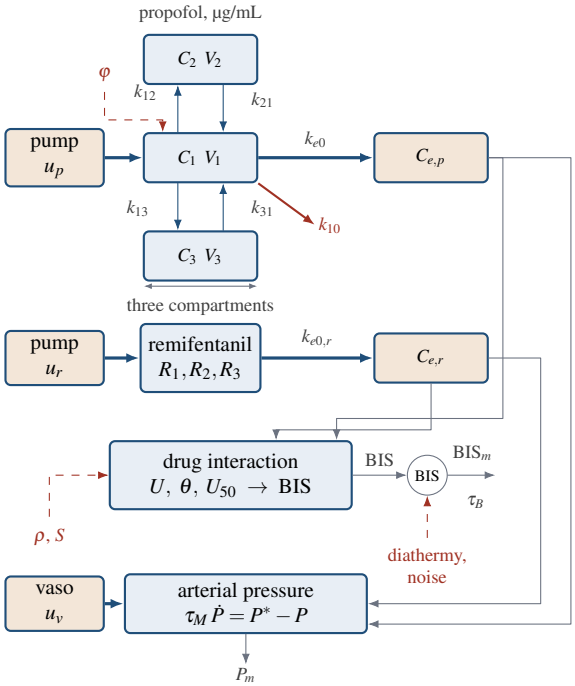


Figure 3.1 Plant of the anesthesia simulator. Propofol enters central compartment C_1 of volume V_1 from pump u_p, exchanges with peripheral compartments C_2 and C_3, is cleared at k_{10} and reaches effect site $C_{e,p}$ through k_{e0}. Remifentanil has the same structure and much faster constants. The two effect-site concentrations enter a response surface that returns the depth index BIS, which the monitor smooths with time constant τ_B and reports as BIS_m. Both concentrations, the surgical stimulation and the vasopressor effect site drive the arterial pressure state. Patient sensitivity ρ and stimulation S scale the concentrations at which the drugs are half maximally effective, blood loss φ shrinks V_1, and diathermy adds a positive offset to the monitor.

$$\frac{\mathrm{d}C_1}{\mathrm{d}t} = \frac{u_p W}{60\,\varphi V_1} - (k_{10} + k_{12} + k_{13})C_1 + k_{21}C_2 + k_{31}C_3 \tag{3.1a}$$

$$\frac{\mathrm{d}C_2}{\mathrm{d}t} = k_{12}C_1 - k_{21}C_2 \tag{3.1b}$$

$$\frac{\mathrm{d}C_3}{\mathrm{d}t} = k_{13}C_1 - k_{31}C_3 \tag{3.1c}$$

$$\frac{\mathrm{d}C_{e,p}}{\mathrm{d}t} = k_{e0}\,(C_1 - C_{e,p}) \tag{3.1d}$$

where C_1, C_2 and C_3 are the central, fast peripheral and slow peripheral concentrations in µg/mL, $C_{e,p}$ is the effect-site concentration, u_p is the commanded infusion in mg/(kg h), W is body weight in kg, V_1 is the central volume in L, φ is the central volume multiplier that represents blood loss, and the rate constants are in reciprocal minutes. The unit arithmetic of the infusion term is worth following once, because a factor of a thousand hides in it. Milligrams per kilogram per hour multiplied by kilograms is milligrams per hour, divided by sixty is milligrams per minute, and divided by a volume in liters is milligrams per liter, which is micrograms per milliliter. The remifentanil infusion u_r is in µg/(kg min), so its own term is $u_r W/(\varphi V_{1,r})$ with no factor of sixty, and micrograms per liter is

Table 3.1 Classification of Signals. Nominal Values Are Those of the Maintenance Start Used by Six of the Thirteen Scenarios.

Symbol	Role	Nominal	Range
BIS_m	CV, measured	50	0–100
P_m	CV, measured	79 mmHg	20–200 mmHg
$C_{e,\mathrm{SP}}$	outer MV, inner SP	1.70	0–8 µg/mL
u_p	MV, propofol	3.43	0–20 mg/(kg h)
u_r	MV, remifentanil	0.149	0–1.0 µg/(kg min)
u_v	MV, vasopressor	0	0–1.0 µg/(kg min)
ρ	DV, unmeasurable	1.0	0.5–2.0
ρ_r	DV, unmeasurable	1.0	0.6–1.6
S	DV, observed	0	0–1
φ	DV, unmeasured	1.0	0.6–1.1
σ	DV, monitor noise	1.0	0–6 index
ΔP_0	DV, baseline offset	0	−20 to +20 mmHg
C_1, C_2, C_3	states, propofol	1.70, 9.55, 13.6	≥ 0 µg/mL
$C_{e,p}$	state, effect site	1.70	≥ 0 µg/mL
R_1, R_2, R_3	states, opioid	4.00, 7.69, 4.20	≥ 0 ng/mL
$C_{e,r}$	state, effect site	4.00	≥ 0 ng/mL
P	state, pressure	79 mmHg	20–200 mmHg
$C_{e,v}$	state, vasopressor	0	0–1.2
x_p, x_r, x_v	states, pumps	3.43, 0.149, 0	at command

nanograms per milliliter. Blood loss enters only through φ, which concentrates the same delivered milligram into a smaller volume.

The Schnider volumes and clearances are functions of the covariates, as shown in Equation (3.2), with lean body mass from the James formula [10].

$$V_1 = 4.27\,\mathrm{L}, \quad V_2 = 18.9 - 0.391\,(A-53), \quad V_3 = 238\,\mathrm{L} \tag{3.2a}$$

$$\mathrm{Cl}_1 = 1.89 + 0.0456\,(W-77) - 0.0681\,(L-59) + 0.0264\,(H-177) \tag{3.2b}$$

$$\mathrm{Cl}_2 = 1.29 - 0.024\,(A-53), \quad \mathrm{Cl}_3 = 0.836 \tag{3.2c}$$

$$L = 1.1\,W - 128\,(W/H)^2 \quad \text{(male)} \tag{3.2d}$$

where A is age in years, H is height in cm, L lean body mass in kg, volumes are in L and clearances in L/min. The rate constants follow as $k_{10} = \mathrm{Cl}_1/V_1$, $k_{12} = \mathrm{Cl}_2/V_1$, $k_{13} = \mathrm{Cl}_3/V_1$, $k_{21} = \mathrm{Cl}_2/V_2$ and $k_{31} = \mathrm{Cl}_3/V_3$, with $k_{e0} = 0.456/\mathrm{min}$ fixed. The female lean body mass expression $1.07W - 148(W/H)^2$ is implemented and is reached through a hidden parameter. The remifentanil parameters are $V_1 = 5.1 - 0.0201(A-40) + 0.072(L-55)$, $V_2 = 9.82 - 0.0811(A-40) + 0.108(L-55)$, $V_3 = 5.42\,\mathrm{L}$, $\mathrm{Cl}_1 = 2.6 - 0.0162(A-40) + 0.0191(L-55)$, $\mathrm{Cl}_2 = 2.05 - 0.0301(A-40)$, $\mathrm{Cl}_3 = 0.076 - 0.00113(A-40)$ and $k_{e0,r} = 0.595 - 0.007(A-40)$ [11]. The two remifentanil volumes, and every volume that appears in a rate-constant denominator, are floored at 1 L so that an extreme covariate set cannot divide by zero. The resulting constants for the default patient are listed in Table 3.2.

The ratio $k_{13}/k_{31} = 55.7$ is the reason this plant has no steady state on the time scale of an operation. At equilibrium the slow compartment holds 55.7 times the central concentration, and it fills with a time constant of $1/k_{31} = 285\,\mathrm{min}$, so an anesthetic of ordinary length is a transient from beginning to end.

Table 3.2 Rate Constants Computed for the Default Patient, a 40 y Old Male of 70 kg and 170 cm With Lean Body Mass 55.3 kg. These Are Derived Quantities and Are Not Displayed by the Application.

Constant	Propofol	Remifentanil	Unit
V_1	4.27	5.12	L
V_2	23.98	9.85	L
V_3	238	5.42	L
k_{10}	0.3836	0.5088	1/min
k_{12}	0.3752	0.4003	1/min
k_{13}	0.1958	0.0148	1/min
k_{21}	0.0668	0.2081	1/min
k_{31}	0.00351	0.0140	1/min
k_{e0}	0.456	0.595	1/min

3.3.3 Response Surface and the Depth Index

The two drugs act on one endpoint and they are synergistic, so the index is computed from a response surface rather than from two independent contributions [18]. Each concentration is normalized by its own half-maximal concentration, the two normalized quantities are added, and the sum is divided by a normalizing term that depends on how the mixture is made up, as shown in Equation (3.3).

$$U_p = \frac{C_{e,p}}{C_{50,p}\,\rho\,(1+k_s S)}, \qquad U_r = \frac{C_{e,r}}{C_{50,r}\,\rho_r\,(1+k_s S)} \tag{3.3a}$$

$$\theta = \frac{U_p}{U_p + U_r}, \qquad U_{50} = 1 - \beta\,\theta\,(1-\theta) \tag{3.3b}$$

$$U = \frac{U_p + U_r}{\max(0.05,\ U_{50})} \tag{3.3c}$$

$$\mathrm{BIS} = E_0 - E_{\max}\,\frac{U^\gamma}{1+U^\gamma} \tag{3.3d}$$

where $C_{50,p}$ and $C_{50,r}$ are the population half-maximal concentrations, ρ and ρ_r are the patient's sensitivities relative to the population, S is surgical stimulation between zero and one, k_s is the stimulation gain, β is the synergy coefficient, γ is the Hill slope, E_0 is the index of an awake patient and $E_{\max}$ is the largest achievable depression. When either drug is absent θ is zero or one and U_{50} is unity, which is the additive case. At an even split $U_{50} = 0.65$ with the default $\beta = 1.40$, so the pair reaches an index of 50 on 35 % less total normalized drug than addition would predict. The index is returned as E_0 when both concentrations are zero.

Equation (3.3a) is where surgical stimulation acts, and the way it acts is the reason this disturbance behaves unlike a load. Stimulation does not add to the index. It multiplies the concentration the patient needs, so the operating point and the process gain move at the same instant. The incision event sets $S = 0.85$, the level the scenarios treat as full stimulation. At that level the effective propofol C_{50} of the default patient rises from 4.16 to 7.17 µg/mL, a factor of 1.72, and on the resistant patient of scenarios 3 and 4 it rises from 6.66 to 11.46 µg/mL.

3.3.4 Arterial Pressure

Arterial pressure is one first-order state driven by the two drug concentrations, the stimulation and the vasopressor, as shown in Equation (3.4).

$$P^{*} = P_0 - a\,C_{e,p} - b\,C_{e,r} + \Delta P_0 + c\,S\left[1 - 0.6\min\left(1,\ \frac{C_{e,r}}{\max(1, C_{50,r})}\right)\right] + 100\,d\,C_{e,v} \tag{3.4a}$$

$$\tau_M \frac{\mathrm{d}P}{\mathrm{d}t} = P^{*} - P, \qquad \tau_v \frac{\mathrm{d}C_{e,v}}{\mathrm{d}t} = u_v - C_{e,v} \tag{3.4b}$$

where P_0 is the awake pressure, a and b are the pressure lost per unit of propofol and of opioid, c is the pressure rise at $S = 1$, d is the vasopressor potency, $C_{e,v}$ is the vasopressor effect site and ΔP_0 is a baseline offset that represents a patient who starts low. The bracket in (3.4a) is the opioid blunting the pressor response to surgery, which is most of what the opioid is given for in practice. The coefficients $a = 6.4\,\mathrm{mmHg}$ per µg/mL and $b = 0.55\,\mathrm{mmHg}$ per ng/mL carry the whole argument about how to split the work between the two drugs, because holding an index of 50 on propofol alone requires 3.74 µg/mL against 1.69 µg/mL with the opioid at 4 ng/mL, and the difference is 11 mmHg taken off the patient.

3.3.5 Actuators and Instruments

All three pumps are first-order lags on the commanded rate with a common time constant, as shown in Equation (3.5), and the state is clipped after every step to the mechanical range of the device.

$$\tau_p \frac{\mathrm{d}x_i}{\mathrm{d}t} = u_{i,\mathrm{cmd}} - x_i, \qquad i \in \{p,\ r,\ v\} \tag{3.5}$$

The depth monitor is a first-order filter on the true index with additive noise and an artifact term, as shown in Equation (3.6).

$$y_k = \mathrm{BIS}_{k-1} + \sigma\,n_k + \delta_k\,(34 + 6\,n_k) \tag{3.6a}$$

$$\hat{y}_k = \alpha_k\,\hat{y}_{k-1} + (1 - \alpha_k)\,\mathrm{clamp}(y_k, 0, 100) \tag{3.6b}$$

$$\mathrm{BIS}_{m,k} = \begin{cases} \mathrm{BIS}_{m,k-1}, & v\,\delta_k = 1 \\ \hat{y}_k, & \text{else} \end{cases} \tag{3.6c}$$

where n_k is a first-order-filtered Gaussian sequence rescaled to a stationary standard deviation of one and advanced once per controller execution while the noise or the artifact is on, so the artifact's fluctuation is the same sequence scaled by six, δ_k is one while diathermy is in use, v is one when signal-quality validation is selected, σ is the noise setting in index units, Δt_c is the controller period and $\alpha_k = \exp(-\Delta t_c / \tau_B)$ with τ_B the monitor smoothing, or one while $v\,\delta_k = 1$, so the hold freezes the smoothing as well. The index the monitor filters is the one computed on the previous step, because everything downstream of the electrodes is at least one sample behind the patient. The artifact is a positive offset of 34 index units, which is the direction a real electromyographic or radio-frequency artifact takes the index [14], and the simulated signal quality indicator falls from 96 to 24 while it is present. The measured pressure is the pressure state through a first-order filter of time constant $\tau_{M,m}$ with no noise.

3.3.6 Parameters, Starting States and Numerics

Table 3.3 collects the constants. Four of them are hidden and never appear on screen. Every other entry is both displayed and adjustable over the range shown, and the derived rate constants of Table 3.2 are neither, because they follow from the covariates.

Two starting conditions are offered. An awake start puts every concentration at zero and the pressure at $P_0 + \Delta P_0$. A maintenance start puts the central compartment and the effect site at 1.7 µg/mL,

Table 3.3 Model Parameters. Adjustable Constants Appear in the Model Panel With the Ranges Shown. Hidden Constants Are Never Displayed.

Symbol	Meaning	Value	Unit	Range
A	age	40	y	18–90
W	weight	70	kg	40–140
H	height	170	cm	145–200
$C_{50,p}$	propofol C_{50}	4.16	µg/mL	1.5–9
$C_{50,r}$	opioid C_{50}	19.0	ng/mL	6–40
γ	Hill slope	1.43	—	0.8–4
β	synergy	1.40	—	0–3.6
E_0	awake index	93	—	85–100
k_s	stimulation gain	0.85	—	0–2
P_0	awake pressure	92	mmHg	60–130
a	propofol vasodilation	6.4	mmHg per µg/mL	0–15
b	opioid vasodilation	0.55	mmHg per ng/mL	0–2
c	pressor response	26	mmHg	0–60
d	vasopressor potency	0.42	mmHg per %	0.05–1.5
τ_M	circulatory lag	0.8	min	0.1–5
τ_v	vasopressor lag	1.6	min	0.2–8
τ_B	monitor smoothing	0.42	min	0.05–2
$\tau_{M,m}$	pressure line lag	0.15	min	0.02–1
τ_p	pump lag	0.12	min	0.02–1
BIS_{lo}	band bottom	40	—	25–55
BIS_{hi}	band top	60	—	45–75
BIS_s	suppression limit	30	—	15–40
P_{lo}	perfusion limit	60	mmHg	45–80
—	sex flag	1	—	hidden
$E_{\max}$	largest depression	93	—	hidden
—	propofol price	0.012	$/mg	hidden
—	remifentanil price	0.34	$/mg	hidden

the fast compartment at $C_1 k_{12}/k_{21} = 9.55\,\mu\mathrm{g/mL}$, the slow compartment at $8C_1 = 13.6\,\mu\mathrm{g/mL}$ and the pressure at the value of P^* in Equation (3.4) with no vasopressor on board, which is $P_0 - aC_1 - bR_1$ when the case starts unstimulated and without offset, with the opioid states set the same way and its slow compartment at $1.05R_1$. That state is not an equilibrium and cannot be, since the slow compartment starts at 8 times the central concentration against 55.7 times at equilibrium, so it is about a seventh full. What it is instead is about an hour into a case, which is what an anesthetist means by steady. One scenario overrides the start to $C_{e,p} = 4.3\,\mu\mathrm{g/mL}$ under full stimulation, which is the concentration a stimulated patient needs, with the concentration target at the same value, so that case starts on the setpoint at a pressure of 81.6 mmHg.

The engine of Section 1.6.1 integrates the model by the fourth-order Runge-Kutta method at a fixed step of 0.01 min (0.6 s). The controllers execute every 0.05 min (3 s), which is a typical update period for a published closed-loop system, and the historian logs twenty samples per simulated minute. The noise generator is seeded at 20260915. After every step the eight concentrations and the vasopressor effect site are floored at zero, the pressure is clipped to 20–200 mmHg, and the three pump states are clipped to their mechanical ranges of 40 mg/(kg h), 2.5 µg/(kg min) and 1.2 µg/(kg min), all of which are wider than the limits the controllers are allowed to command.

3.3.7 Assumptions

The model carries detail where the control problem is and very little elsewhere. Each simplification below breaks down at a stated place.

The pharmacokinetics are population means. Schnider and Minto were fitted to healthy adult volunteers and are known to describe the extremes of age and body habitus poorly [10, 11]. The application inherits that and then makes the residual variability a slider, because the variability rather than the mean is the control problem.

One hypnotic endpoint. The response surface returns a single index. Real anesthesia has at least three endpoints, hypnosis, analgesia and immobility, and the two drugs trade differently against each of them [18]. A controller tuned on one endpoint in this model would have to be re-examined against the others.

The index is algebraic in the concentrations. Equation (3.3d) is instantaneous and the only dynamics are the monitor filter. A real processed-EEG index is a proprietary computation over an epoch with dynamics of its own, it varies between patients at the same concentration, and its relationship to awareness is statistical rather than a threshold [14].

Stimulation is one number. A single quantity between zero and one scales both half-maximal concentrations. Real stimulation is a sequence of discrete events of very different intensities, it differs between the two drugs, and it depends on what is being done at that second.

The circulation is one state. There is no cardiac output, no heart rate, no fluid status, no baroreflex and no arrhythmia, and blood loss appears only as a change in one compartment volume. The application therefore cannot show the cases where a pressure problem is not a vasodilation problem.

The artifact is a step. Diathermy adds a fixed offset while it is on and the signal quality indicator switches between two values. A real artifact varies with electrode position and cautery mode, and a real quality index degrades gradually, which makes validation logic a threshold problem rather than the switch it is here.

Recovery is outside the model. Awakening, respiratory depression and postoperative pain are absent, which is precisely where the cost of a heavy opioid technique appears. The model can therefore show what moving depth onto the opioid buys and cannot show what it spends.

3.4 Dynamic Analysis

3.4.1 Process Gain of the Depth Loop

The gain the outer loop sees is the slope of the response surface at the operating point, and it moves for two independent reasons. Differentiating Equation (3.3d) numerically at the concentration pair that holds an index of 50 gives -8.8 index units per µg/mL with no opioid, -10.0 at 4 ng/mL and -18.8 at 8 ng/mL. The opioid target is an operator setting rather than a disturbance, so a decision taken for a reason that has nothing to do with tuning changes the loop gain by a factor of 2.1. Patient sensitivity moves the same gain again. The propofol concentration that holds an index of 50 at 4 ng/mL of opioid is 0.85 µg/mL at $\rho = 0.5$ and 3.38 µg/mL at $\rho = 2.0$, and over the narrower range $0.6 \leq \rho \leq 1.8$ the requirement still varies by a factor of three. Equation (1.6) gives the open-loop error that such a spread leaves under model inversion alone. A fixed-gain outer loop has to be stable across all of it, which is the argument for tuning the loop against the population rather than against a nominal patient [5].

3.4.2 Synergy and the Choice of Split

The set of concentration pairs that produce a given index is the isobologram, and its curvature is what makes the split between the two drugs a decision rather than a habit. Figure 3.2 draws the curve for an index of 50 on the population patient, on a resistant patient and under full surgical stimulation, against the straight line that additive drugs would give.

Two features of Figure 3.2 matter for control. The curve bows well inside the additive line, so at 8 ng/mL the propofol requirement is 0.84 µg/mL against the 1.99 µg/mL that addition predicts, a factor of 2.4. The two ways the curve moves are not the same shape. Patient sensitivity scales $C_{50,p}$ alone in (3.3a), so the whole curve stretches along the propofol axis by the factor ρ and by nothing else. Stimulation scales both half-maximal concentrations, so the curve moves outward along both

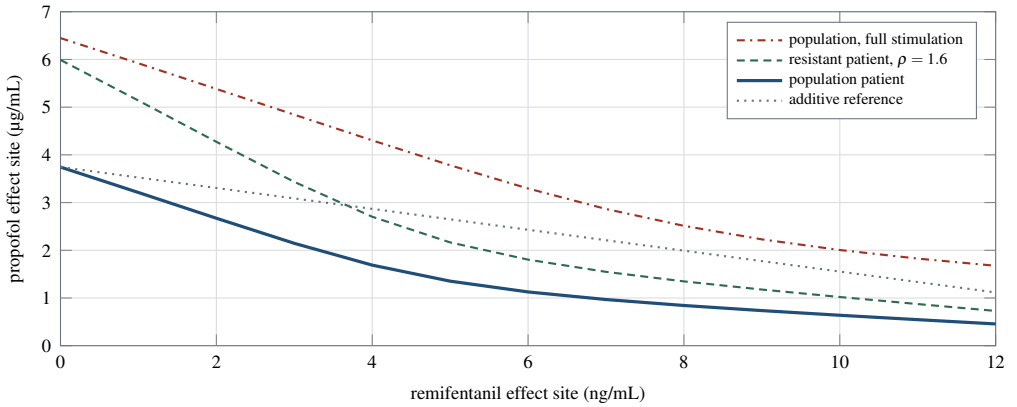


Figure 3.2 Isobologram of concentration pairs holding a depth index of 50, computed from the implemented response surface. The population patient is the thick solid curve, the resistant patient densely dashed and full stimulation dash-dotted. The dotted line joins the population curve at zero opioid to the additive prediction, and the bow of the solid curve below it is the synergy. Moving the opioid from zero to 4 ng/mL cuts the propofol requirement from 3.74 to 1.69 µg/mL, and every unit of that reduction is worth 6.4 mmHg of arterial pressure.

axes by $1 + k_s S$ together and the same opioid concentration buys proportionally less depth than it did before the incision.

🕮 **Pause.** Try Exercise 3.2 before reading on.

3.4.3 Why the Maintenance Rate Never Settles

Holding the effect site at 1.7 µg/mL from the maintenance start requires 3.43 mg/(kg h) at the beginning of the case. With the target-controlled pump in service and nothing disturbed at all, the pump delivers that rate at one minute, and the rate falls to 3.23 mg/(kg h) at one hour and to 2.84 mg/(kg h) at four hours, heading for the 2.39 mg/(kg h) that clearance alone would require once the slow compartment is full. The decay is the third compartment filling, the three-compartment form of the falling rate of Equation (1.3) and Figure 1.3, and it is the clearest illustration in the application of what a model buys. A controller with no model has to discover a requirement that falls by 17 % over four hours through its integral, continuously, for the whole case. A controller with a model gets it for free and spends its integral on the patient instead. With the depth loop closed over the same hour, the rate wanders between 2.7 and 3.8 mg/(kg h), because the outer loop passes the monitor noise through to the target.

3.4.4 Lags and the Speed of the Outer Loop

The propofol effect site is behind the plasma by $1/k_{e0} = 2.19$ min, the monitor adds $\tau_B = 0.42$ min of smoothing, and the pump adds $\tau_p = 0.12$ min. The outer loop therefore faces about two and a half minutes of lag that no tuning removes, and it also faces an inner loop that is fast enough to be treated as instantaneous because that loop works on an equation. The fermenter of Volume 3, Chapter 3 measures what an inner loop that is not fast enough costs a cascade. The default depth integral time of 3.5 min is a little longer than the total lag, which is why identification of an unusual patient takes half an hour rather than a few minutes. Anything that feels responsive on this loop is too fast.

3.4.5 One-Sided Authority

The propofol pump can add drug and cannot remove it. Shutting it off leaves whatever is already in the effect site to leave by redistribution and metabolism, and the effect site can only follow the

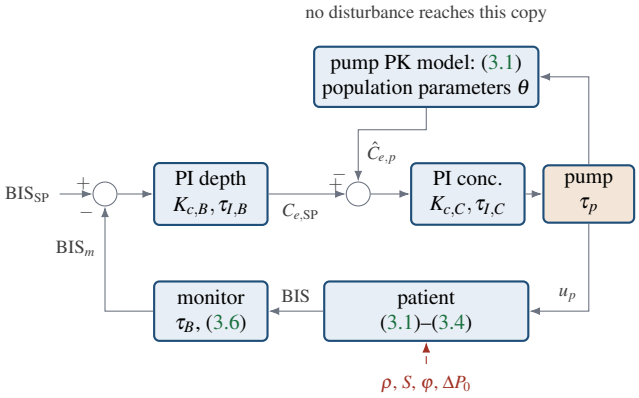


Figure 3.3 Two nested loops with one real measurement. The inner loop is closed on the pump's own integration of Equation (3.1) with population parameters. No disturbance reaches that copy, so it holds its target whatever the patient does, and everything the model gets wrong appears in the outer loop's error where an integral can absorb it.

central compartment down at k_{e0}. The deep excursions in Section 3.6 are of two kinds. In scenario 7 the propofol pump has just been commanded to zero when the index reaches its lowest value, so a different choice of gains trims the excursion only at its margins. At the end of a case the index reaches its lowest value within half a minute of the stimulation going, while the pump is still running, and in that time no change of infusion can reach an effect site that lags the plasma by 2.19 min. Stopping the propofol pump outright at the instant of closing changes the lowest index by less than 0.1 in both scenarios that end a case. A faster response after the closing, from a higher depth gain, a stopped pump or a lower opioid target, ends the excursion sooner while leaving its depth nearly unchanged. What makes it shallower is the moment at which the operator starts lightening, which is decided before the excursion rather than in response to it. Section 1.4.6 gives the time such a plant takes to fall to a limit, and the insulin pump of Chapter 2 meets the same asymmetry.

3.5 Control Structure as Implemented

3.5.1 Four Loops and Four Structures

Figure 3.3 shows the arrangement, a cascade whose inner loop is closed on a model. The depth controller reads the monitor and writes the propofol effect-site target. The concentration controller reads the pump's own copy of the pharmacokinetics and writes the propofol infusion rate. A second concentration controller does the same for remifentanil against an operator-set target, and a pressure controller writes the vasopressor rate. The control structure selector chooses how many of them are in service. In MANUAL the operator sets all three infusion rates directly. In target concentration the two inner loops run and the operator sets the concentration targets. In closed loop the depth controller sets the propofol target as well. A fourth setting replaces the depth and propofol concentration loops with the predictive controller of Section 3.5.4. The vasopressor loop and the signal-quality logic are independent switches, so the four structural choices can be made one at a time.

3.5.2 The Discrete Algorithm

All four controllers are instances of the velocity-form proportional-integral-derivative algorithm of Section 1.6.2, Equations (1.12) to (1.15), executed every $\Delta t_c = 0.05$ min. This chapter writes the action as $\mathscr{A}$, the a of Chapter 1, with $+1$ for a reverse-acting loop whose process gain is positive and -1 for a direct-acting one, since a here is the pressure coefficient of Equation (3.4). Anti-windup is back-calculation with the tracking constant $T_t = \tau_I$ for all four loops, so each scan removes $\Delta t_c / \tau_I$ of

Table 3.4 The Four Loops as Configured. Every Loop Is Proportional-Integral With Derivative Action Disabled, a 3 s Execution Period and No Rate Limit.

Loop	PV	MV (range)	K_c	τ_I
Depth ($\mathscr{A} = -1$)	BIS_m	$C_{e,\mathrm{SP}}$ (0–8)	0.045	3.5
Propofol ($\mathscr{A} = +1$)	$\hat{C}_{e,p}$	u_p (0–20)	5.0	2.5
Opioid ($\mathscr{A} = +1$)	$\hat{C}_{e,r}$	u_r (0–1.0)	0.16	4.0
Pressure ($\mathscr{A} = +1$)	P_m	u_v (0–1.0)	0.012	4.0

the clipped excess, and it runs only when there is integral action to unwind. The clamp acts on the emitted output alone.

Three properties of this implementation matter for the results of this chapter. Derivative action is taken on the measurement rather than on the error, and τ_D is zero for all four loops and is not exposed as a setting, so every loop in the application is a PI controller. No output carries a rate limit, and the only actuator bounds are the absolute ones in Table 3.4. MANUAL mode sets the integral so that the output reproduces the operator's setting exactly, so transfer in either direction is bumpless with no logic of its own. The inner loops are held in that tracking state whenever target control is out of service, and the depth controller is held in it whenever the loop is open, so switching structures mid-case does not bump the pumps.

The depth loop is the only direct-acting one, because more propofol produces a lower index, so the gain from this manipulated variable to this controlled variable is negative. Setting it reverse acting instead sends the target to zero on a light patient. The pressure loop is reverse acting because more vasopressor produces more pressure, and the opposite choice drives the syringe to its stop whenever surgical stimulation pushes the pressure above setpoint, which is most of an operation.

The opioid loop's integral time of 4 min is a little longer than the 1.7 min remifentanil effect-site lag plus the pump. After the awake start of scenario 2 the modeled concentration rises past its 4 ng/mL target to 4.45 at 3.2 min, dips to 3.87 at 5.8 min and stays within 5 % of the target from 4.1 min on, and the true index stays between 49.5 and 49.8 from twenty to forty minutes. From a maintenance start the same loop holds its target without moving. Setting the integral time to 0.7 min on its slider shows what too short an integral time does on this loop: the concentration overshoots to 7.3 ng/mL at about three minutes and then rings with a period of about 5.3 min.

3.5.3 The Measurement That Is a Calculation

The pump's model is a separate integration of Equation (3.1) advanced by explicit Euler at the controller period, driven by the infusion the pump is actually delivering, and built from the population parameters of Equation (3.2). It never sees ρ, S or φ. Explicit Euler at 3 s is adequate for rate constants of a few per minute and is what a device does. The inner loop is a PI controller on that model rather than the effect-site targeting algorithm of commercial pumps [1], so after an awake start the modeled effect site overshoots a 1.7 µg/mL target by 0.19 µg/mL at five minutes, stays within 0.043 µg/mL of it from ten minutes on and is 0.008 µg/mL above it at forty minutes. The consequence of a model that sees no disturbance is stated in the application's own verification suite and is worth restating here. With the central volume reduced by 20 %, the modeled effect-site concentration holds its target to better than 0.05 µg/mL after twenty-five minutes while the patient is more than 0.25 µg/mL past it (0.002 and 0.33 µg/mL in the current build). The inner loop cannot be wrong about its target and cannot be right about the patient.

🕮 **Pause.** Try Exercise 3.1 before reading on.

Signal-quality handling is four lines of configuration and it is the single most consequential switch in the application. With the controller set to trust the monitor, Equation (3.6c) passes the filtered value through whatever the quality indicator says. With validation selected, the controller holds its

last good measurement while the indicator is poor. No tuning substitutes for it, because a corrupted trace and a light patient are the same number.

3.5.4 Model Predictive Control

The fourth value of the structure selector gives the propofol infusion to a receding-horizon optimization of the kind proposed for hypnosis under total intravenous anesthesia [19, 20]. It has the move suppression, hard input box and soft output bands of Section 1.6.3, and it departs from the linear model of that section in two ways. The index enters through the nonlinear response surface, and one parameter of that surface is learned online. It writes u_p itself, while the depth and propofol concentration loops track the applied infusion at every scan with no preload, so a change of structure in either direction is free of a bump. The opioid loop, the pressure loop, the signal-quality logic and the pump's range act as they do in closed loop. Where the cascade corrects the response surface with an integral, the predictive controller corrects it with a parameter identified online [21].

Controller model. The pharmacokinetic half is the pump's own model, the delivered rate and the four concentrations of Equation (3.1) with $\varphi = 1$, built from the covariate entries the pump uses and the declared lag $\tau_p = 0.12\,\text{min}$ of Equation (3.5), discretized with an exact zero-order hold and started at each decision from the pump model's state. Ten minutes of 6 mg/(kg h) from empty give an effect site of 1.9089 µg/mL through the hold and through a fourth-order Runge-Kutta integration of the plant at $\varphi = 1$, and 1.9097 µg/mL through the pump's Euler copy, a residual of 0.044 %. The response half is Equation (3.3) at the declared defaults of Table 3.3 with $\rho_r = 1$, $S = 0$ and an estimate $\hat{\rho}$ in place of ρ, so the monitor is predicted as $\hat{y} = h(\hat{C}_{e,p}, \hat{C}_{e,r}; \hat{\rho}) + \hat{b}$, with h the index of Equation (3.3d) and $\hat{b}$ an output bias. The opioid effect site is assumed to approach its target with $\tau_R = 1.5\,\text{min}$. The surface is not linearized, because a linearization at an awake start is flat. Since h is static and monotone in $\hat{C}_{e,p}$, the setpoint is mapped by bisection into an effect-site reference C_k^* at every step k of the horizon,

$$h(C_k^*, \hat{C}_{e,r,k}; \hat{\rho}) = \text{BIS}_{\text{SP}} - \hat{b} \tag{3.7}$$

capped at 8 µg/mL, and the band edges 40 and 60 are mapped the same way into C_k^+ and C_k^-. At an index of 50 with 4 ng/mL of opioid and $\hat{\rho} = 1$ the reference is 1.6893 µg/mL and the slope of h is $K = 10.04$ index units per µg/mL, against 10.31 identified from target steps in closed loop.

Learned potency. The estimate is a one-parameter recursive least-squares fit of the monitor against the model's own prediction of it, in the potency $g = 1/\hat{\rho}$ bounded so that $0.5 \le \hat{\rho} \le 2$, the declared range of sensitivity. The prediction h_k and the regressor $\partial h_k / \partial g$, both evaluated at the pump model's effect sites, pass through the declared monitor smoothing as F_k and ϕ_k, with $\alpha_B = 1 - \exp(-\Delta t_c / \tau_B)$, as shown in Equation (3.8).

$$z_k = \text{BIS}_{m,k} - F_k + \phi_k\, \hat{g}_{k-1}, \quad L_k = \frac{P_{k-1}\,\phi_k}{\lambda_k + \phi_k^2\, P_{k-1}} \tag{3.8a}$$

$$\hat{g}_k = \text{clamp}(\hat{g}_{k-1} + L_k\,(z_k - \phi_k\,\hat{g}_{k-1}),\ 0.5,\ 2) \tag{3.8b}$$

$$P_k = \min[(P_{k-1} - L_k\,\phi_k\,P_{k-1})/\lambda_k,\ 0.5] \tag{3.8c}$$

Here z_k is the regression target, L_k the gain and P_k the covariance, which starts at 0.5. Three guards keep a quiet signal from walking the estimate away. The forgetting factor λ_k is 0.98 per scan only while the pump model's effect site has moved by more than 0.1 µg/mL in 3 min, and one otherwise. The covariance returns to 0.5 when the depth setpoint changes. Samples are skipped while the signal quality indicator is below 50 and for 1.5 min afterward. Over a quiet maintenance hour of scenario 5 the estimate moves by 0.14, 0.83 and 0.17 % at ρ of 0.5, 1 and 2 and P falls. The bias is $\hat{b}_k = \hat{b}_{k-1} + \alpha_D(\text{BIS}_{m,k} - F_k - \hat{b}_{k-1})$ with $\alpha_D = 1 - \text{e}^{-\Delta t_c/\tau_D}$ and $\tau_D = 1\,\text{min}$, an output disturbance model [22] frozen while validation holds the measurement, which leaves a mean $|\text{BIS} - 50|$ of 0.006 and 0.079 over the last twenty minutes of scenario 5 at ρ of 0.5 and 2 with the learning and the noise off.

Decision and objective. A plan is made every $T_s = 0.25$ min, five scans, over $N = 40$ steps (10 min) in ten move blocks of $1, 1, 1, 2, 2, 3, 4, 6, 8, 12$ steps, and the pump applies the first block's infusion for the 15 s until the next plan. Each plan minimizes, over the ten block moves,

$$J = \sum_{k=1}^{N} K^2 \left(\hat{C}_{e,p,k} - C_k^*\right)^2 + \mu K^2 g_u^2 \sum_{j=1}^{10} \Delta u_{p,j}^2 + \sum_{i=1}^{4} \left(w_i s_i^2 + v_i s_i\right) \quad (3.9)$$

subject to $C_k^- - s_1 \le \hat{C}_{e,p,k} \le C_k^+ + s_2$, $\hat{C}_{e,p,k} \le \hat{C}_{e,p,0} + (P_m - P_{lo})/a + s_3$, $\hat{C}_{e,p,k} \le 8 + s_4$, $s_i \ge 0$ and the hard box $0 \le u_p \le 20$, the clamp of the concentration loop. Here $g_u = 0.318$ µg/mL per mg/(kg h) is the model's effect site after ten minutes of a unit infusion, so a move is priced at μ times the square of the index shift it would cause if held, and each slack is one number for the whole horizon. The move weight is $\mu = 0.1$. The light side of the band carries $w_1 = 10^3 K^2$ and $v_1 = 100K$, and the deep side ten times as much, because the pump can always add drug and can never take any back (Section 3.4). The pressure floor, written through the declared $a = 6.4$ mmHg per µg/mL without the circulatory lag, the opioid or the vasopressor, carries $w_3 = 10^3 a^2$ and $v_3 = 100a$, and the ceiling $w_4 = 10^7$ and $v_4 = 10^6$. The limits are read from the run's parameters as the existing logic reads them. The weight μ was chosen once on scenario 12 from 0.03 to 3 as the smallest value whose infusion travel stays within about 15 % of the closed loop's (76 against 68 mg/(kg h)), and the horizon was not tuned.

Acceptance and fallback. A plan takes effect only with an optimal status, every planned infusion inside the box, the first planned infusion equal to the applied one and a finite cost, slacks and predicted index. Otherwise the historian is marked and the closed loop continues from its tracked state until a valid plan returns. Seven kinds of forced failure on scenario 12 handed over with a step of at most 0.018 mg/(kg h), except 0.657 one minute into the induction. The first plan is solved at $t = 0$. In Chrome the 360 plans of scenario 12 cost a mean of 601 and 610 µs in two runs, which is 0.24 % of one core at the 1 min/s playback speed.

Three models kept apart. The plant is the thirteen-state model at the run's patient, the instrument is Equation (3.6) with its noise and artifact, and the controller's model is the pump's pharmacokinetics, the population surface with $\hat{\rho}$ and $\hat{b}$, and the slope a. It never reads ρ, ρ_r, S, φ or the true index, so $\hat{\rho}$ is the potency of the drug pair as the population surface reads the monitor, which Section 3.6.9 shows is not the patient's sensitivity.

3.6 Scenarios and Measured Results

The thirteen scenarios in Table 3.5 are arranged so that each one adds a single mechanism to the one before. Every scenario fixes the starting state, the control structure, the four tunings, the patient and the timed events, and carries the briefing and completion rule of Section 1.6.4. Table 3.6 reports what the application measures over each run, using its own scorecard quantities. Peak values, minutes outside the band and delivered dose are quoted rather than a bare integrated error, because integrated absolute error keeps accruing after a transient ends and is comparable only within a fixed window.

3.6.1 A Rate Is Not a Trajectory

Scenario 1 makes the point before any controller is involved. Setting the propofol pump to 3.43 mg/(kg h) and the remifentanil pump to 0.149 µg/(kg min), the rates that hold the maintenance state still, and waiting, gives an index of 84.2 at two minutes, 69.5 at five, 61.5 at ten, 56.2 at twenty and 52.4 at forty, still falling. The effect site fills through a plasma that is itself filling a 4.27 L central compartment while two peripheral compartments draw off it, and the sum of the two transfer rate constants $k_{12} + k_{13} = 0.571/\text{min}$ is larger than the clearance $k_{10} = 0.384/\text{min}$. Arriving in two minutes requires overshooting the plasma deliberately and coming off again before the effect

Table 3.5 The Scenario Ladder and What Each Step Introduces

#	Configuration	What it shows
1	MANUAL, awake start	a rate is not a trajectory
2	Target control, $\rho = 1$, $C_{e,\mathrm{SP}} = 1.7$, $C_{e,r} = 4$	model inversion on the patient the model fits
3	Target control, $\rho = 1.6$, $C_{e,r} = 2$	exact inner loop, wrong depth
4	Closed loop, $\rho = 1.6$, incision at 40 min	feedback as identification, then a target at its bound
5	Closed loop, maintenance start, nothing happens	a plant with no steady state
6	Closed loop, incision at 8, closing at 40 min	a disturbance that moves the plant
7	Closed loop, opioid stepped 0.5, 8, 4 ng/mL	synergy and the price of the split
8	Closed loop, diathermy 10–16 and 30–38 min	an artifact in the dangerous direction
9	Target control, 20 % volume loss at 10, loop closed at 45 min	a disturbance the model cannot see
10	Closed loop, $\rho = 0.7$, $\Delta P_0 = -12$ mmHg, vasopressor loop in service	the constraint is on another loop
11	Closed loop, $S = 0.85$ from the start, closing at 20 min	one-sided authority at the end of a case
12	Predictive control, awake start, $\rho = 0.5$, incision at 30, closing at 60 min	a plan that learns the gain it lacks
13	Challenge: 120 min, $\rho = 1.45$, $\rho_r = 0.85$, five events	structure before tuning

Table 3.6 Measured Outcomes Over the Window Shown, Read From the Application's Own Scorecard. The Index Extremes Are of the True Index and Not of the Monitor. In the Six Scenarios That Start Awake the Maximum Is the Awake Value of 93 and the Minutes Above the Band Include the Induction Transit.

#	Configuration	Window	True index		In band	Above 60	Below 30	Lowest P	Propofol
		(min)	lowest	highest	(%)	(min)	(min)	(mmHg)	(mg)
1	MANUAL, maintenance rates	40	52.4	93.0	70	12.0	0.0	80	160
2	Target control, population patient	40	48.0	93.0	95	2.1	0.0	78	215
3	Target control, resistant patient	40	65.9	93.0	0	40.0	0.0	79	215
4	Closed loop, resistant patient	40	51.0	93.0	94	2.5	0.0	65	481
4	Closed loop, through the incision	90	51.0	93.0	87	11.5	0.0	61	1415
5	Closed loop, undisturbed	90	49.5	50.7	100	0.0	0.0	79	339
6	Closed loop, incision and closing	70	37.1	66.8	90	3.7	0.0	69	483
7	Closed loop, opioid swept	110	29.0	57.1	94	0.0	1.3	68	456
8	Diathermy, controller trusts monitor	45	31.8	50.2	63	0.0	0.0	62	333
8	Diathermy, last good value held	45	48.9	50.2	100	0.0	0.0	78	178
9	Blood loss, loop closed at 45	90	43.2	50.7	100	0.0	0.0	76	275
10	Depth against pressure	90	31.1	49.0	29	0.0	0.0	67	340
11	End of case, opioid at 4	90	35.5	50.6	95	0.0	0.0	67	427
11	End of case, opioid at 8	90	33.0	50.8	95	0.0	0.0	73	185
12	Predictive control, sensitive patient	90	35.8	93.0	92.5	3.6	0.0	81	322
13	Challenge as shipped	120	60.0	93.0	17	99.7	0.0	72	482
13	Challenge, structure configured	120	33.5	93.0	91	4.4	0.0	59	795

site catches up. With the target-controlled pump in service that trajectory is computed rather than remembered, and scenario 2 delivers 10.2 mg/(kg h) at one minute, 7.3 at three, 4.8 at ten and 3.8 at thirty, reaching an index of 50 at 3.1 min and dipping to 48.0 at 4.2 min as both inner loops carry their effect sites about a tenth past target before settling.

Table 3.7 Settled Index at 40 min Under Model Inversion and Under Closed-Loop Control, Against Patient Sensitivity. The Modeled Concentration Is the Same in Every Row of the Left Half, 0.008 µg/mL Above Its Target. The Closed-Loop Values Are Read Before the Incision of Scenario 4.

	Target control		Closed loop		
ρ	index	P (mmHg)	index	$C_{e,\mathrm{SP}}$	P (mmHg)
0.70	52.1	80	50.6	1.90	79
0.85	56.3	80	50.6	2.29	77
1.00	59.6	80	50.6	2.68	74
1.30	64.2	80	50.8	3.43	69
1.60	67.4	80	51.1	4.15	65
2.00	70.4	80	51.7	5.03	59

3.6.2 What Model Inversion Cannot Do

Table 3.7 runs the scenario 3 configuration, a commanded 1.7 µg/mL of propofol with the opioid at 2 ng/mL, across the range of patient sensitivity, and runs the scenario 4 configuration, which is the same case with the depth loop closed, beside it.

The left half of Table 3.7 spans 18.3 index units while the pump reports the same modeled concentration in every row, and only the three lowest sensitivities land inside the surgical band, the population patient at 59.6 with 0.4 to spare. The right half spans 1.1 index units, and the mechanism is entirely in the fourth column, where the commanded concentration moves from 1.90 to 5.03 µg/mL. That column is a parameter estimate produced by an integral, on a plant where nobody told the controller which parameter was missing. The last column is the price. Because the depth is bought with a vasodilator, the settled pressure falls from 79 mmHg at $\rho = 0.70$ to 65 mmHg at $\rho = 1.60$ and to 59 mmHg at $\rho = 2.0$, which is one millimeter of mercury below the perfusion limit the application marks.

Identification is not fast, and it cannot be. In scenario 4 the commanded concentration is 3.13 µg/mL at five minutes, 3.83 µg/mL at twenty and 4.07 µg/mL at thirty-nine, against the 4.27 µg/mL this patient requires at 2 ng/mL of opioid. The index stays within five units of setpoint from 9.4 min onward. Half an hour to identify a patient 60 % away from the population is consistent with what published closed-loop systems achieve [5, 13], and it follows directly from the effect-site lag of Section 3.4.

> **TRY IT IN THE APP 4 · Closing the Loop**
>
> Load scenario 4 and press Run. Watch the Propofol target trace under the Depth of hypnosis strip rather than the index itself. The target climbs as the outer integral finds this resistant patient: 3.13 µg/mL at five minutes, 3.83 at twenty and 4.07 at thirty-nine, and the index stays within five units of 50 from 9.4 min onward.

3.6.3 A Disturbance That Moves the Plant

The incision in scenario 6 takes the index from 50.2 to 66.8 within one integration step, because the index is algebraic in the stimulation, and the controller answers by raising the target, which is ordinary. What is not ordinary is that the effective C_{50} moved by a factor of 1.72 at the same instant, so the gain and the operating point changed together and the loop is at its slowest exactly when it is being asked to work hardest. Closing at forty minutes is the same disturbance reversed, and it is the one that costs. The index falls to 37.1 half a minute after the closing and spends 10.3 min below 45. The propofol pump is still delivering 7.2 mg/(kg h) at the lowest index and 4.0 mg/(kg h) two minutes after the closing, because the depth loop lowers the target by only 0.7 µg/mL in the first minute, but stopping the pump outright at the closing would leave the lowest index within 0.1 of the same value.

Scenario 4 shows the other end of the same mechanism. On the resistant patient at full stimulation the depth controller runs its output up over the forty-eight minutes after the incision and reaches

8 µg/mL, the upper bound on the concentration target, at 88.5 min. At 90 min the index is 52.0 and the pressure is 60.6 mmHg. The controlled variable is close to setpoint and the loop has no authority left. The honest reading of that trend is not a tuning problem, because a resistant patient at full stimulation needs more than the hypnotic alone can safely give, and the remedy is a higher opioid target.

3.6.4 The Cost of Believing the Monitor

Scenario 8 runs two episodes of electrocautery totaling fourteen minutes. With the controller trusting the monitor, the displayed index rises while the true index falls to 31.8, the pressure reaches 62 mmHg and the patient receives 333 mg of propofol in forty-five minutes. With signal-quality validation selected and nothing else changed, the lowest true index is 48.9, the lowest pressure is 78 mmHg and the dose is 178 mg. Over the same forty-five minute window the integrated absolute index error falls from 367 to 14 index-minutes, a factor of 26, and the time inside the surgical band rises from 63 % to 100 %. The rejected samples never enter the monitor's smoothing, so after each episode the reading resumes within 1.8 index units of the true index. Every second of the trusting run is invisible while it happens, because a corrupted trace and a light patient produce the same number on the same screen.

> **TRY IT IN THE APP 8 · The Measurement Lies, Upward**
>
> Load scenario 8 and press Run with Signal quality at Trust the monitor. During the cautery from 10 to 16 and 30 to 38 min, watch the monitor line rise on the Depth of hypnosis strip while the true index falls, to 31.8. Reset, set Signal quality to Hold on low quality and run again: the lowest true index is now 48.9.

3.6.5 A Disturbance the Model Cannot See

Scenario 9 removes a fifth of the central volume from a patient on target control. The pump's model reports 1.702 µg/mL against a target of 1.700 µg/mL because it is holding a number it computed itself, while the patient reaches 2.05 µg/mL and the index falls to 43.2. Nothing anywhere in the pump indicates that anything has happened. Closing the loop at forty-five minutes brings the index within one unit of 50 from 63.8 min, about nineteen minutes later, by lowering the commanded concentration. At 89 min the target is 1.18 µg/mL, the model holds 1.09 µg/mL and the patient holds 1.32 µg/mL. The model is still wrong. It is being asked for the wrong number in exactly the way that makes the patient right, which is a more honest description of what an outer loop does than the phrase error correction suggests.

3.6.6 The Constraint Is on Another Loop

Scenario 10 gives a sensitive patient a baseline pressure 12 mmHg low and a surgeon who wants twelve more index units of depth. Taking the setpoint from 50 to 38 moves the propofol effect site from 1.27 to 2.22 µg/mL, and the pressure holds 70 mmHg only because the vasopressor loop takes itself to 0.147 µg/(kg min) by fifty minutes. Moving the same depth onto the opioid brings the vasopressor back to 0.067 µg/(kg min) by ninety minutes, less than half, at the same depth and the same pressure. The loop with the setpoint everyone is discussing is not the loop that limits the plant, and the only way to see that is to have the other one on the screen.

3.6.7 The End of the Case

Scenario 11 ends a case run at full stimulation. It starts with the propofol effect site and the concentration target both at 4.3 µg/mL, the requirement under stimulation, and the index stays between 49.8 and 50.2 until the closing. Closing at twenty minutes takes the index from 50.1 to 35.5 a quarter of a minute later, and it spends 13.0 min below 45 while the propofol rate falls from 8.5 to 3.8 mg/(kg h) over two minutes, because the drug is in the effect site and the only way out is redistribution and metabolism. The dressings at 42 min take the remaining stimulation away and add a second, shallower dip to 44.4. Running the same case with the opioid target at 8 ng/mL instead of 4 ng/mL deepens the dip to 33.0, although the propofol pump is commanded to zero 1.8 min after the

closing. The time below 45 after the closing is shorter, 11.5 min instead of 13.0, because the opioid has taken over part of the depth and less propofol has to wash out, but the higher target costs another 8.9 min below 45 at the start, where the opioid climbs to its target faster than the depth loop takes propofol off and the index reaches 40.8. A deeper dip is the opposite of what the pharmacology of a short-acting opioid suggests. The reason is that nothing in the arrangement asks for less of it, because the target-controlled pump goes on holding 8 ng/mL after the stimulation has gone. A short-acting drug is only short-acting if something turns it down.

> ✦ **DID YOU KNOW? More Opioid, Less Propofol, a Deeper Dip**
>
> From the simulator: raising the opioid target of scenario 11 from 4 to 8 ng/mL cuts the propofol given over 90 min from 427 to 185 mg, yet the lowest true index after the closing falls further, to 33.0 against 35.5 (Table 3.6).

Turning it down works. Cutting the opioid target to 1 ng/mL at the moment closing begins leaves the depth of the dip unchanged at 35.5, because that part is the stimulation falling in one step, and takes the time below 45 from 13.0 to 3.8 min. Lightening six minutes before closing instead, to an index setpoint of 55 with the opioid target cut to 1 ng/mL at the same time, stops the dip at 43.3 with 1.8 min below 45 and no time above 60. A faster depth loop would shorten the time below 45 as the opioid cut does and leave the lowest index where it is. Only anticipation raises the bottom of the dip, and it is feedforward from an event that nobody has instrumented.

🕮 **Pause.** Try Exercise 3.3 before reading on.

3.6.8 Structure Before Tuning

The challenge is a case run to 120 min on a patient who needs 45 % more propofol than the model expects and is slightly more sensitive to the opioid than the population ($\rho_r = 0.85$), with a baseline pressure 8 mmHg low, monitor noise of 1.5 index units, an incision at eight minutes, cautery from twenty-two to twenty-nine, blood loss at forty-two and closing at sixty-two. It loads on target control with weak tuning ($K_c = 0.008$ and $\tau_I = 20$ min on the depth loop), a trusting monitor, the vasopressor by hand and an opioid target of 2 ng/mL. The application judges the case as given on the structure the challenge names and on three quantities of the true index. For the whole run the case must stay in Auto with the depth loop closed, the signal-quality hold selected and the vasopressor loop in service, and the opioid target must be above the shipped 2 ng/mL when the run starts. The largest deviation of the index from 50 after 79.2 min, which the application calls the last third of the run, must not exceed 6, and over the whole run the index must spend less than 2 min below 30 and less than 10 min above 60. Hypotension is displayed beside the verdict and is not judged. Table 3.8 adds the four structural choices one at a time to the shipped configuration and reports the judged quantities beside the pressure and the dose.

Reading Table 3.8 downward separates four different kinds of help. Closing the loop with the default tuning raises the time in the surgical band from 17 to 82 %, cuts the time above 60 from 99.7 to 12.55 min and the late deviation from 10.1 to 5.9, and it is the only change that addresses the patient. On its own it creates a hypotension problem that did not exist before, because it buys the depth with a vasodilator and spends 74.0 min below 60 mmHg with a lowest pressure of 38 mmHg. Signal validation is worth 17.1 min of that hypotension and 86 mg of propofol, and it addresses the measurement rather than the patient. The vasopressor loop is worth another 54.5 min, addresses the constraint and leaves the index untouched. Raising the opioid target addresses the split, and it is the change that takes the time above 60 under its limit and reduces the propofol dose by 43 %, from 1399 to 795 mg. None of the four is a tuning constant, and the tuning is identical in the last five rows. The level loop of Volume 3, Chapter 2 makes the same point on a tank, where each structure rejects the disturbance that enters its own measurement path whatever the tuning. Figure 3.4 draws three columns of the table as bars.

Table 3.8 Structural Choices Added One at a Time on the 120 min Challenge. Each Row Above the Rule Keeps Everything From the Row Above It. $|e|_{\max}$ Is the Largest $|\text{BIS} - 50|$ From 79.2 min, Light Counts Minutes Above an Index of 60, Low P Counts Minutes Below 60 mmHg, $P_{\min}$ Is the Lowest Pressure, and No Row Spends Time Below an Index of 30. The Last Row Adds Only the Opioid Target to the Closed, Tuned Loop. Pass Is the Application's Verdict, Which Also Requires the Three Switches for the Whole Run and an Opioid Target Above 2 ng/mL at the Start.

Added	$\lvert e\rvert_{\max}$	Light (min)	In band (%)	Low P (min)	$P_{\min}$ (mmHg)	Propofol (mg)	Pass
As shipped	10.1	99.7	17	0.0	72	482	no
Closed loop, tuned	5.9	12.55	82	74.0	38	1486	no
Signal validation	5.2	12.55	84	56.9	40	1399	no
Vasopressor loop	5.2	12.55	84	2.35	56	1399	no
Opioid target 5	3.3	4.4	91	1.18	59	795	yes
Opioid 5 only	4.4	4.4	88.5	11.3	53	928	no

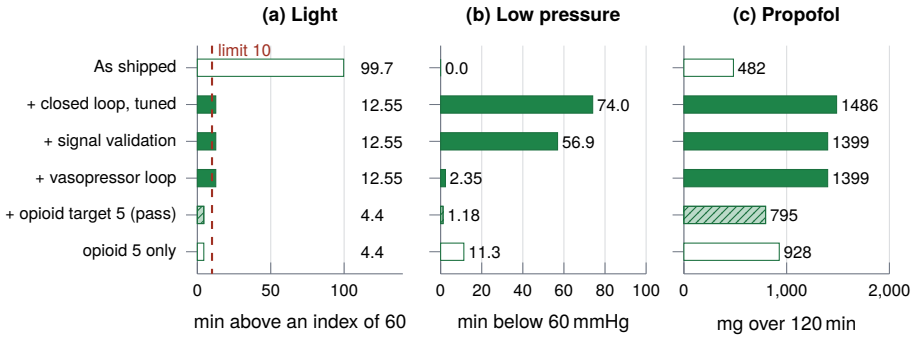


Figure 3.4 The structural choices of Table 3.8 on the challenge: (a) minutes above an index of 60 against the limit of 10 (dashed), with the values in a column at the right, (b) minutes below 60 mmHg and (c) propofol given. White: the shipped configuration and the opioid target raised alone. Solid: the cumulative rows that fail. Hatched: the configured row that passes. Closing the loop cuts the light time and creates hypotension, the vasopressor loop removes most of the hypotension, and only the opioid target brings the light time under its limit and the dose down.

The configured run passes with a late deviation of 3.3, 4.4 min above 60 and none below 30, margins of 2.7 index units, 5.6 min and 2 min, and it gives the same result whether its settings are made before Run or carried through Reset. The shipped configuration lacks all three switches, keeps the opioid at 2 ng/mL and misses the index limits by 4.1 index units and by 89.7 min. It is not a clean case. Of the time above 60, 1.55 min is the induction transit and 2.85 min follows the incision, and the index also spends 6.7 min below the surgical band, 1.3 min of it in an undershoot to 38.8 during induction and 5.4 min after the closing, where it reaches 33.5. That deep excursion is the end-of-case mechanism of scenario 11 arriving on a case that also carries blood loss.

The verdict judges what the challenge text states. The challenge asks for a control structure, a signal-quality gate, a pressure loop and a drug split to be chosen before any tuning, its briefing calls the gate and the closed outer loop structural requirements, and its explanation names the vasopressor loop in service and an opioid target above the shipped one. The last row of Table 3.8 closes the loop and raises the opioid target with the monitor trusted and the vasopressor by hand. Its index meets the three index limits, with a late deviation of 4.4 and 4.4 min above 60, but the pressure spends 11.3 min below 60 mmHg and falls to 53 mmHg, and the verdict fails it on the two missing switches. The index limits alone reject the weaker configurations as well. Target control fails them in each

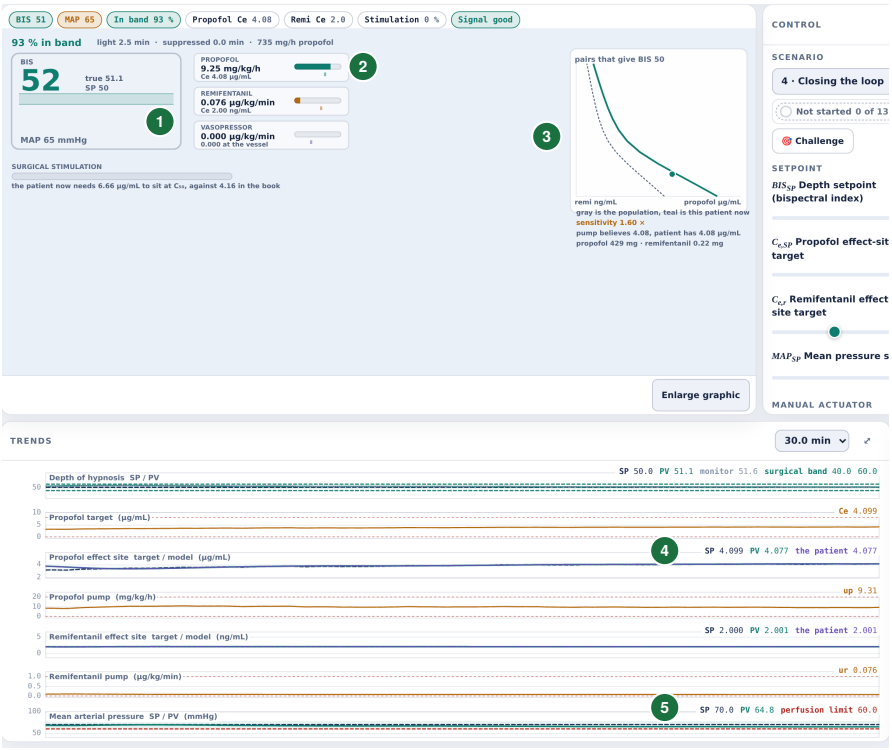


Figure 3.5 Depth held on a patient who needs 60 % more propofol than the population model expects, scenario 4 at 35 min, with the trends covering the last thirty minutes. The true index is 51.1 against a setpoint of 50 while the propofol target below it has been driven to 4.10 µg/mL, where the population model asks for 2.67 at the same opioid concentration of 2.00 ng/mL, and the separation between those two numbers is the patient. The isobologram inset shows the operating point sitting on the patient's own curve rather than on the population curve, the arterial pressure strip shows what the extra propofol costs (64.8 mmHg), and the remifentanil traces show the opioid loop holding its target. Callouts: (1) the depth readout with the true index and the setpoint, (2) the pump tiles, (3) the isobologram inset, (4) the propofol effect-site target and model trend, (5) the mean arterial pressure strip. The crop removes the header, the scorecard and most of the control panel.

of six combinations of targets tried, the full structure with the opioid at 5 ng/mL and the shipped weak tuning spends 46.25 min above 60, and the full structure with the opioid left at 2 ng/mL spends 12.55 min above 60. The ±6 band applies only from 79.2 min, seventeen minutes after the last event, so tracking through the incision, the cautery and the blood loss is judged through the two time limits and the structure, and time between an index of 30 and the bottom of the surgical band at 40 is not counted. The pressure loop is required as a switch rather than through a pressure limit, so a pass says that the structure was in place and the depth held, and the hypotension column of Table 3.8 still has to be read beside it. The application is shown in Figure 3.5 running the scenario that carries the argument, the resistant patient of scenario 4 with the depth loop closed.

3.6.9 Predictive Control Against the Closed Loop

Table 3.9 runs scenarios 4, 5, 6, 8 and 9 for 90 min from their own settings with only the structure changed and the patient sensitivity at 0.5, 1 and 2, and the challenge for 120 min on its suggested settings, with the same seed and instruments in every row. Scenario 8 keeps its trusting monitor, scenario 9 closes the depth loop from the start, and the predictive controller uses scenario 12's settings everywhere. Five rows are compared: the closed loop with its shipped tuning, the predictive controller,

Table 3.9 Integrated Absolute Index Error in BIS·min, With the Depth Below an Index of 40 in BIS·min in Parentheses, Over 90 min for Scenarios 4 to 9 and 120 min for the Challenge (13, Suggested Settings), at Patient Sensitivity ρ. The Predictive Controller Uses the Settings of Scenario 12 Throughout. Scenario 8 Trusts the Monitor, and Scenario 9 Closes the Depth Loop From the Start.

Scenario	ρ	Closed loop	Predictive	Predictive, no learning	Retuned loop	Learned-gain loop
4	0.5	254 (18.8)	155 (0.0)	157 (6.5)	288 (41.8)	333 (15.8)
4	1	297 (0.0)	161 (0.0)	146 (0.0)	216 (4.7)	237 (0.0)
4	2	695 (0.0)	622 (0.0)	632 (0.0)	537 (0.0)	561 (0.0)
5	0.5	82 (2.9)	66 (2.9)	67 (2.9)	120 (2.9)	132 (3.8)
5	1	19 (0.0)	32 (0.0)	29 (0.0)	31 (0.0)	19 (0.0)
5	2	134 (0.0)	51 (0.0)	45 (0.0)	59 (0.0)	74 (0.0)
6	0.5	214 (9.6)	182 (10.1)	162 (10.3)	212 (9.4)	221 (10.2)
6	1	315 (5.8)	138 (6.9)	126 (7.2)	141 (8.4)	218 (6.4)
6	2	548 (0.0)	216 (2.4)	215 (2.3)	297 (8.4)	370 (5.3)
8	0.5	523 (176.8)	666 (324.0)	673 (328.2)	682 (302.4)	576 (110.2)
8	1	454 (88.6)	350 (124.9)	348 (124.9)	484 (203.1)	452 (87.7)
8	2	445 (13.9)	179 (0.0)	174 (0.3)	403 (50.3)	431 (37.7)
9	0.5	103 (2.9)	75 (2.9)	78 (2.9)	204 (2.9)	171 (3.8)
9	1	53 (0.0)	41 (0.0)	41 (0.0)	50 (0.0)	57 (0.0)
9	2	103 (0.0)	59 (0.0)	58 (0.0)	82 (0.0)	86 (0.0)
13	1.45	479 (20.8)	251 (10.7)	242 (10.8)	358 (22.8)	354 (13.9)

Table 3.10 Three Decisive Cases Under the Five Rows. Scenario 12 ($\rho = 0.5$) and Scenario 4 ($\rho = 2$) Over 90 min, the Challenge on Its Suggested Settings Over 120 min. t_b Is the Time to the Band, Below 40 the Minutes and the Depth Under It in BIS·min, IAE in BIS·min, P_{lo} the Minutes Under 60 mmHg, Plans the Optimizations Solved and Rejected.

Row	t_b	Below 40		IAE	Propofol	P_{lo}	Plans
	(min)	(min)	depth		(mg)	(min)	
Scenario 12, sensitive patient							
Closed loop	1.20	6.39	38.2	341	330	0.00	—
Predictive	1.20	3.11	7.5	211	322	0.00	360/0
No learning	1.19	4.86	17.2	207	323	0.00	360/0
Retuned loop	1.20	7.60	60.7	352	336	0.00	—
Learned gain	1.22	7.72	40.2	423	339	0.00	—
Scenario 4, resistant patient							
Closed loop	3.04	0.00	0.0	695	1562	15.65	—
Predictive	2.99	0.00	0.0	622	1595	15.98	360/0
No learning	3.12	0.00	0.0	632	1587	20.40	360/0
Retuned loop	3.00	0.00	0.0	537	1667	44.47	—
Learned gain	3.00	0.00	0.0	561	1638	37.87	—
Challenge, suggested settings							
Closed loop	1.55	6.74	20.8	479	795	1.18	—
Predictive	1.55	3.18	10.7	251	815	0.00	480/0
No learning	1.56	3.21	10.8	242	809	0.00	480/0
Retuned loop	1.56	7.11	22.8	358	828	1.03	—
Learned gain	1.57	4.88	13.9	354	822	0.51	—

the same with the learning off ($\hat{\rho} = 1$, bias only), the closed loop retuned once by the simple internal model control rule with the closed-loop time constant equal to the dead time, and the closed loop with its depth gain scheduled on the learned sensitivity, $K_c = 0.045\,\hat{\rho}$. The retuned loop rests on a first-order-plus-dead-time fit of $\pm 0.2\,\mu g/mL$ target steps from the maintenance start with the noise off, a gain of -10.31 index units per µg/mL, a lag of 1.13 min and a dead time of 1.33 min with a residual of 3.6 % of the response, and gives $K_c = 0.0413$ and $\tau_I = 1.13$ min. The time to the band differs by at most 0.63 min among the rows of any case, and the 19,440 plans include no fallback.

Scenario 12 induces a patient who needs half the propofol the model expects ($\rho = 0.5$) from an awake start, with the opioid target at 4 ng/mL, an incision at 30 min and closing at 60 min. The plan opens the pump to 20 mg/(kg h) and cuts the infusion to 8.1 at 1.25 min and to zero by 1.55 min, because its model already holds enough drug in the plasma to carry the effect site to the reference, while the closed loop stays at or above 18 until 1.45 min and still delivers 4.6 at 2.75 min. The index falls to 27.7 under the closed loop and turns at 39.5 under the plan, with 31.5 against 0.2 BIS·min below the band in the first half hour, as shown in Figure 3.6a. In minutes below 40 over that half hour the closed loop spends 4.03, the plan 0.68 and the plan with the learning off, which turns at 33.9, spends 2.38. Over the case the index error falls from 341 to 211 BIS·min and the closed loop's 1.6 min below an index of 30 disappear (Table 3.10). The closing costs the same under every structure, 2.43 min below 40 against 2.36, because that drug is in the effect site before the stimulation goes.

Switching the corrections off one at a time apportions the induction. With the learning off and the bias frozen at its first value, the plan holds the sensitive patient at the population's concentration and spends 28.2 min of the first thirty below 40 (82.2 BIS·min), worse than the closed loop. Either correction then suffices and the two substitute. Adding the learning first removes 81.9 of the 82.2 BIS·min, while adding the bias first removes 72.6 and the learning then 9.3. Over the whole case only the bias handles the incision (164 against 65 BIS·min from 30 to 50 min), and with the learning off the whole-case error is 207 against 211, a tie, so the learning pays only below the band, 17.2 against 7.5 BIS·min. Given to the closed loop as a gain schedule, the learned number raises its error from 341 to 423, because a lower gain on a sensitive patient slows the recovery without changing the first minutes on the pump's ceiling.

The learned sensitivity reads 0.535 at 2 min and 0.502 at 5 min against the patient's 0.50, and after the incision it climbs to 1.26 at 50 min although the patient has not changed, as shown in Figure 3.6b. The equivalent value, at which the population surface returns the patient's true index at the true effect sites, is 1.28 there. Stimulation multiplies both half-maximal concentrations by 1.72 in Equation (3.3a), and one potency can express that only as a less sensitive patient (on propofol alone the equivalent would be 0.86). After the closing the estimate returns to 0.62 at 80 min against an equivalent 0.60.

Across Table 3.9 the predictive controller lowers the index error on scenarios 4, 6 and 9 at every sensitivity and on scenario 5 at 0.5 and 2. With the hold selected it also lowers the error on scenario 8, 58, 23 and 43 against 81, 24 and 114 BIS·min at ρ of 0.5, 1 and 2. It loses scenario 5 at $\rho = 1$, 32 against 19 BIS·min, because the 1 min bias passes the monitor noise to the plan (2.2 against 0.9 with the noise off, 25.4 with a 3 min bias). It also loses scenario 8 with the monitor trusted, 324.0 against 176.8 BIS·min below 40 at $\rho = 0.5$ with a lowest index of 15.6 against 24.5, and 124.9 against 88.6 at $\rho = 1$ although its index error is lower there. The bias follows the 34-unit artifact within a minute, and on a sensitive patient the inverted surface turns that apparent lightness into a large reference. Freezing the bias on a poor signal would credit the plan with the validation this row has switched off, so it is not done.

The retuned loop beats the plan on scenario 4 at $\rho = 2$, 537 against 622 BIS·min, at a cost of 44.5 min below the 60 mmHg floor against 16.0, all of the plan's minutes falling in the induction and 10.5 min of the retuned loop's after the incision. Its faster integral goes deeper on sensitive patients, 41.8 against 18.8 BIS·min below the band on scenario 4 at $\rho = 0.5$. On the challenge's suggested settings the plan lowers the index error from 479 to 251 BIS·min with a largest late deviation of 1.7 against 3.3 and no time under the pressure floor, yet the verdict fails it because the challenge requires the closed depth loop for the whole run, and the rules were left as they are. From the shipped challenge settings with only the structure changed, the pressure spends 30.6 min below the floor with the vasopressor by hand, so the soft pressure row does not replace the pressure loop.

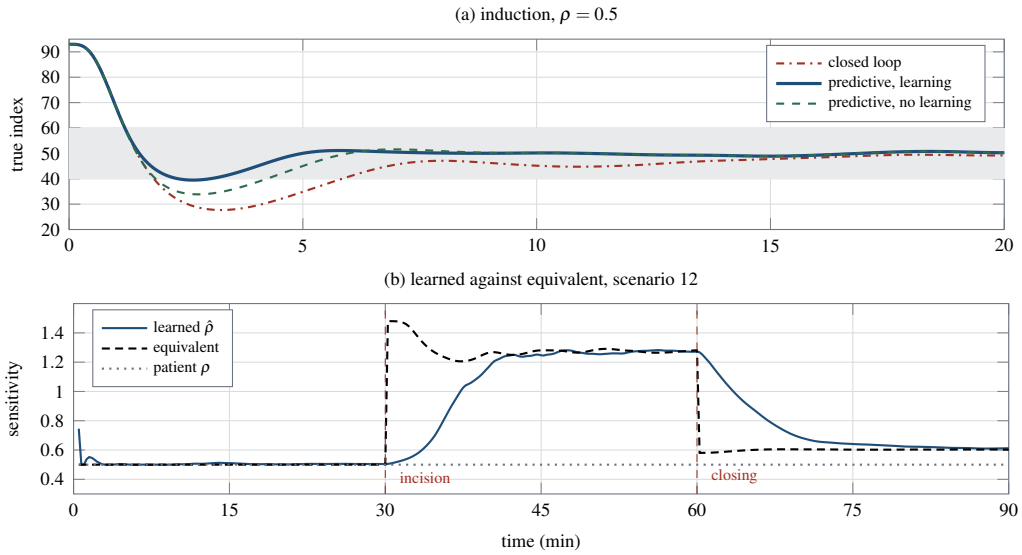


Figure 3.6 Scenario 12 under the closed loop and the predictive controller, from engine data. (a) The true index in the first 20 min against the shaded surgical band, under the closed loop (dash-dotted), the plan (thick solid) and the plan without learning (dashed). The plan stops the pump at 1.55 min and the index turns at 39.5, while the closed loop and the plan without learning fall to 27.7 and 33.9. (b) The learned sensitivity (solid) against the equivalent value (densely dashed) at which the population surface reads the patient's true index at the patient's true effect sites. The patient's sensitivity (dotted) stays at 0.50, and the estimate follows the equivalent to 1.26 under the incision.

3.7 Technology in Practice

The monitor, the pumps and the pressure measurement of Figure 3.1 stand, in an operating room, for a processed electroencephalogram monitor, syringe pumps that carry the pharmacokinetic models, and a blood pressure monitor. Figure 3.7 draws the four loops with that hardware in place of the schematic's symbols.

What Measures It

The depth loop reads a processed electroencephalogram. Its index is computed by the monitor's own algorithm from a stretch of signal, which is why Section 3.3.7 credits it with dynamics of its own. The chapter's monitor of Equation (3.6) stands for it with a first-order filter of $\tau_B = 0.42$ min, noise, and a positive offset of 34 index units during diathermy.

> **TECHNOLOGY** **Sensor: The Processed EEG Depth Monitor**
>
> A depth monitor records the electroencephalogram through electrodes on the head, the function of a Class II electroencephalograph in United States regulation [23], and reduces it to one number. The most common index analyzes the signal in the frequency and phase domains and reports a value from 0 to 100, with 40 to 60 generally taken as right for surgical anesthesia. Electrode placement alone can shift the reading by up to 9 %, and drugs such as ketamine and nitrous oxide give a wake-like signal in a sedated patient [24]. High muscle activity and interference from electrical devices can pass for brain activity and raise the index without being flagged as artifact [14]. On a trend that failure is an index climbing while nothing has changed in the patient.

The pressure loop reads mean arterial pressure. A basic monitoring standard requires that every patient under anesthesia have arterial blood pressure and heart rate determined and evaluated at least

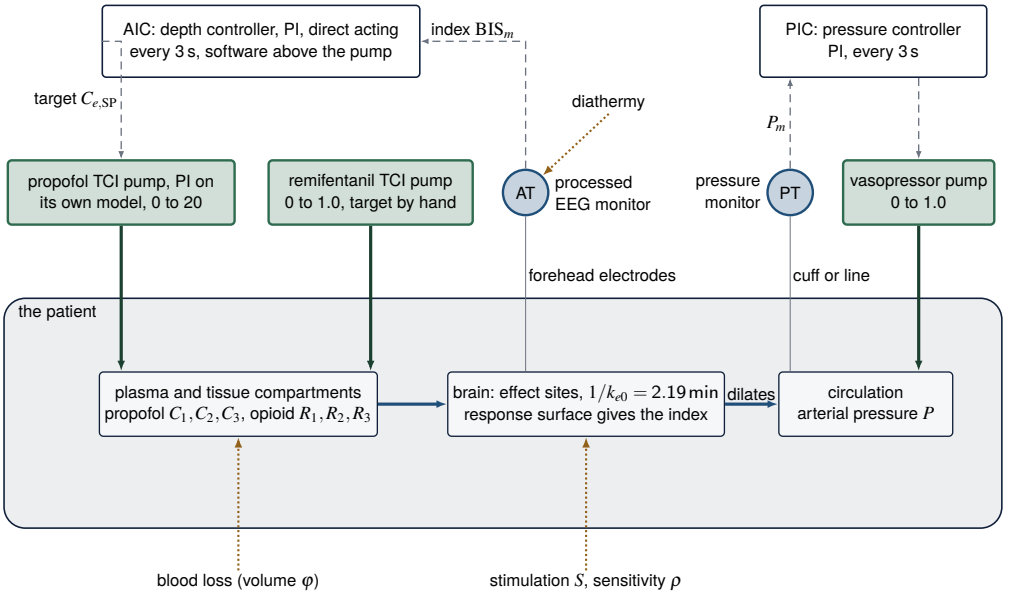


Figure 3.7 The four loops of Figure 3.3 with the hardware of an operating room. The processed EEG monitor (AT) reads the forehead and gives the index to the depth controller AIC, which writes the propofol effect-site target to a target-controlled infusion (TCI) pump. That pump closes its own loop on its internal model, which no disturbance reaches. The pressure monitor (PT) feeds the pressure controller PIC and the vasopressor pump. Pump ranges are those of Table 3.4, in mg/(kg h) for propofol and µg/(kg min) for the others. Thick solid arrows carry drug and its effect, thin dashed arrows signals, dotted arrows disturbances. Every pump fails stopped and none can remove drug.

every five minutes, and it lists a tracing of intra-arterial pressure among the ways to evaluate the circulation continually [25]. The chapter's pressure measurement, a first-order filter with no noise, is closer to a continuous arterial tracing than to a cuff read every few minutes.

What Moves It

Three syringe pumps deliver propofol, remifentanil and the vasopressor. The two anesthetic pumps run target-controlled infusion, in which the processor in the pump computes the rate from its model and the patient's covariates [3]. The inner concentration loops of Figure 3.3 are the chapter's version of that calculation. The fail position of every pump is stopped, and a stopped pump cannot remove drug, which is the one-sided authority of Section 3.4. Equation (3.5) models each pump as a first-order lag of $\tau_p = 0.12$ min with no bolus function and no dead volume.

TECHNOLOGY Final Element: The Target-Controlled Syringe Pump

In United States regulation an infusion pump moves fluid into a patient in a controlled manner by a piston, roller or peristaltic mechanism, may detect air in or blockage of the line and sound an alarm, and is a Class II device [26]. A target-controlled pump adds a microprocessor that holds a pharmacokinetic model. The clinician enters a target for the plasma or the effect site and the patient's covariates, and the processor computes the infusion, because neither concentration can be measured during the case. Such pumps alarm on high line pressure, a stopped infusion, an empty syringe, a power failure and a low battery, and some detect a disconnection from a drop in line pressure [3].

Where the Controller Runs

The inner loop runs inside the pump, on the processor that carries the model [3]. Research systems have placed the depth controller above such a pump: in 2002 a PID algorithm read the depth index and wrote the target of a propofol target-controlled infusion system [13]. Closed-loop target-controlled infusion on a pharmacodynamic endpoint such as the depth index or blood pressure is still experimental, and in routine practice the clinician watches the clinical response and the depth of anesthesia while the pump runs [3], which closes the outer loop by hand. The chapter's four controllers run every 3 s, as Section 3.5 describes.

IN PRACTICE Setting Up an Intravenous Anesthetic

A national audit named a poor grasp of the pharmacology and failed drug delivery as the two most common causes of awareness during total intravenous anesthesia. The setup therefore matters as much as the model. The drug goes into a large peripheral vein that stays visible, through a line free of dead space, kinks, air and leaks, with locking connectors, an anti-siphon valve against overdose from a damaged syringe, and an anti-reflux valve where several infusions share a line. The pressure cuff goes on another limb, a keep-vein-open mode is not used for propofol or remifentanil, and several pumps sit on the pole in a fixed order [3].

What Is Modernizing It

Three developments bear on this loop. Commercial pharmacokinetic models were built from selected populations that often left out children, obese and elderly patients. General-purpose models fitted to pooled data now cover far wider ranges: one propofol model rests on 15 433 concentration measurements from 1033 people, from 27 weeks postmenstrual age to 88 years and from 680 g to 160 kg [27]. Such a model still predicts a population mean, and no measurement before induction reveals where the patient on the table sits in the spread of Table 3.7. Closed loops are reaching beyond hypnosis: a closed-loop vasopressor controller studied in 40 patients after cardiac surgery reduced postoperative hypotension [3]. Research on artificial intelligence in anesthesia now centers on multimodal depth monitoring and automated drug infusion [24].

3.8 Limitations

The most important limitation is what the application is for. It is a teaching model of a control structure, and closed-loop control of anesthetic depth is not standard of care in any jurisdiction [17]. Nothing in the application has been validated against patient data, no result in Section 3.6 describes a person, and no configuration found here should be used to plan or guide the care of a patient. The pharmacokinetic models and the response surface are taken from published work, which makes the trajectories plausible and does not make them clinical. A reader who takes anything from this chapter into an operating room has misread it.

The technical limitations follow the assumptions of Section 3.3.7. The response surface has one endpoint, so the application cannot show a patient who is unconscious and moving, which is the situation the opioid and the neuromuscular blocking agents exist to handle. The circulation is one first-order state, so a pressure problem is always a vasodilation problem and never a volume or

a rhythm problem. The depth index is a deterministic function of two concentrations plus noise, whereas real indices differ between patients at the same concentration and occasionally move in the wrong direction. Recovery and respiratory depression are outside the model, which is exactly where a heavy opioid technique presents its bill, so scenario 7 can price the split in millimeters of mercury and cannot price it in minutes of apnea. The pump is a first-order lag with no bolus function and no mechanical dead volume. Repeatability is itself unrealistic. A fixed integration step and a fixed controller period give the same trajectory on every machine, and neither a real pump nor a real monitor behaves that way.

Two elements are deliberately stylized. Surgical stimulation arrives as a step and leaves as a step, which no operation does, although a step is the one disturbance shape that can be read off a trend without fitting anything. A patient sensitivity that the operator can set on a slider does not exist in an operating room either, although it is the only way to show the quantity the outer loop is estimating while the estimate is being made.

Everything that survived the simplification survived for one reason, which is that a control engineer has to answer for it. What survived is a manipulated variable that acts through two lags in series, a plant gain uncertain by a factor of three before the case starts, a multiplicative disturbance that moves the gain and the operating point together, an inner measurement that is a computation, a measurement corrupted in the direction that provokes an overdose, an actuator with authority in one direction, and a constraint that lives on a different loop from the setpoint under discussion.

The predictive controller's surface is the population one, so its single learned parameter absorbs everything the surface leaves out. Under the incision it reads 1.26 on a patient at 0.50, and on a resistant patient under stimulation the equivalent value of 4.29 lies beyond the bound of 2, so the estimate cannot be used to infer the patient. Each soft constraint has one slack for the whole horizon, the pressure row is a static slope with no circulatory lag or vasopressor, and the opioid is not a decision variable. The comparison is between control structures on a teaching simulator and makes no clinical claim.

> ⚠ **OPEN CHALLENGE** **One Number for Several Endpoints**
>
> A processed index reflects only statistical features of cortical electrical activity, while consciousness arises from networks that reach below the cortex. Indices respond differently to different drugs, proprietary algorithms make monitors disagree, and neuromuscular blockade and the balance of analgesia also move the reading. An index is therefore not a measure of overall adequacy when hypnosis, analgesia and paralysis change together [24]. The single response surface of Section 3.3.7 makes the same simplification. Progress would be monitoring that joins the electroencephalogram with heart rate variability, the pulse wave, muscle activity and near-infrared spectroscopy [24], and a controller that can tell which endpoint has moved.

3.9 Conclusion

This chapter used depth of anesthesia under propofol and remifentanil as a laboratory for a control structure whose inner measurement is a calculation. The simulator integrates thirteen states, two three-compartment pharmacokinetic models with an effect site behind each, an arterial pressure state, a vasopressor effect site and three pump lags, and turns the two effect-site concentrations into a depth index through a response surface with one synergy coefficient. The pump carries its own copy of the pharmacokinetics, built from population parameters and shielded from every disturbance, and that copy is within 0.01 µg/mL of its target at forty minutes on a patient whose requirement is anywhere between 0.85 and 3.38 µg/mL.

Under model inversion alone the settled index moves across 18.3 units as patient sensitivity spans its range, and closing an outer loop on the depth monitor holds the same range of patients within 1.1 units by moving the commanded concentration from 1.90 to 5.03 µg/mL. Signal-quality validation changes the lowest true index during electrocautery from 31.8 to 48.9 and the delivered dose from 333 to 178 mg, and no tuning substitutes for it. At the end of a case the index falls to 35.5 a quarter

of a minute after the stimulation goes, before any change of infusion can reach the effect site, so a higher gain or a lower opioid target can shorten the dip, while lightening before the closing raises its bottom.

The graded challenge judges the structure its text names as well as the index, so a configuration that trusts the monitor and leaves the pressure to the operator, 11.3 min below 60 mmHg, fails it although its index passes. A predictive controller on the pump's own model with one learned potency holds the induction of a patient twice as sensitive as the model at a lowest index of 39.5 against 27.7, although without the learning it ties on index error (207 against 211 BIS·min). A fuller model would add a second endpoint to the response surface, a subcutaneous or intramuscular route with real transport delay, a signal quality index that degrades gradually, and a population of virtual patients rather than one covariate set with a slider. What the application already shows is the division of labor between a model and a measurement.

KEY POINTS

- The pump holds the drug level its average-patient model predicts, so it cannot be wrong about its target and cannot be right about the patient.
- Under model inversion alone the settled index spans 18.3 units across patient sensitivity, and the closed depth loop holds the same patients within 1.1 units by moving the commanded concentration.
- Signal-quality validation raises the lowest true index during electrocautery from 31.8 to 48.9, and no tuning substitutes for it.
- At the end of a case a higher gain or a lower opioid target shortens the dip, while lightening before the closing raises its bottom.

Where else this applies. A model-based estimate corrected by a late measurement returns in Volume 5, Chapter 10, and the one-way pump in Chapter 2.

References

[1] S. L. Shafer and K. M. Gregg, "Algorithms to rapidly achieve and maintain stable drug concentrations at the site of drug effect with a computer-controlled infusion pump," *Journal of Pharmacokinetics and Biopharmaceutics*, vol. 20, no. 2, pp. 147–169, 1992.

[2] M. M. R. F. Struys, T. De Smet, J. I. B. Glen, H. E. M. Vereecke, A. R. Absalom, and T. W. Schnider, "The history of target-controlled infusion," *Anesthesia & Analgesia*, vol. 122, no. 1, pp. 56–69, 2016.

[3] P. U. Bidkar, A. Dey, P. Chatterjee, R. Ramadurai, and J. J. Joy, "Target-controlled infusion – past, present, and future," *Journal of Anaesthesiology Clinical Pharmacology*, vol. 40, no. 3, pp. 371–380, 2024.

[4] T. W. Schnider, C. F. Minto, S. L. Shafer, P. L. Gambus, C. Andresen, D. B. Goodale, and E. J. Youngs, "The influence of age on propofol pharmacodynamics," *Anesthesiology*, vol. 90, no. 6, pp. 1502–1516, 1999.

[5] K. van Heusden, G. A. Dumont, K. Soltesz, C. L. Petersen, A. Umedaly, N. West, and J. M. Ansermino, "Design and clinical evaluation of robust PID control of propofol anesthesia in children," *IEEE Transactions on Control Systems Technology*, vol. 22, no. 2, pp. 491–501, 2014.

[6] B. Joseph and C. B. Brosilow, "Inferential control of processes: Part I. steady state analysis and design," *AIChE Journal*, vol. 24, no. 3, pp. 485–492, 1978.

[7] D. E. Seborg, T. F. Edgar, D. A. Mellichamp, and F. J. Doyle, III, *Process Dynamics and Control*, 4th ed. Hoboken, NJ: John Wiley & Sons, 2016.

[8] R. G. Bickford, "Automatic electroencephalographic control of general anesthesia," *Electroencephalography and Clinical Neurophysiology*, vol. 2, no. 1–4, pp. 93–96, 1950.

[9] E. Krüger-Thiemer, "Continuous intravenous infusion and multicompartment accumulation," *European Journal of Pharmacology*, vol. 4, no. 3, pp. 317–324, 1968.

[10] T. W. Schnider, C. F. Minto, P. L. Gambus, C. Andresen, D. B. Goodale, S. L. Shafer, and E. J. Youngs, "The influence of method of administration and covariates on the pharmacokinetics of propofol in adult volunteers," *Anesthesiology*, vol. 88, no. 5, pp. 1170–1182, 1998.

[11] C. F. Minto, T. W. Schnider, T. D. Egan, E. Youngs, H. J. M. Lemmens, P. L. Gambus, V. Billard, J. F. Hoke, K. H. P. Moore, D. J. Hermann, K. T. Muir, J. W. Mandema, and S. L. Shafer, "Influence of age and gender on the pharmacokinetics and pharmacodynamics of remifentanil," *Anesthesiology*, vol. 86, no. 1, pp. 10–23, 1997.

[12] H. Schwilden, J. Schüttler, and H. Stoeckel, "Closed-loop feedback control of methohexital anesthesia by quantitative EEG analysis in humans," *Anesthesiology*, vol. 67, no. 3, pp. 341–347, 1987.

[13] A. R. Absalom, N. Sutcliffe, and G. N. Kenny, "Closed-loop control of anesthesia using bispectral index," *Anesthesiology*, vol. 96, no. 1, pp. 67–73, 2002.

[14] A. A. Dahaba, "Different conditions that could result in the bispectral index indicating an incorrect hypnotic state," *Anesthesia & Analgesia*, vol. 101, no. 3, pp. 765–773, 2005.

[15] P. S. Sebel, T. A. Bowdle, M. M. Ghoneim, I. J. Rampil, R. E. Padilla, T. J. Gan, and K. B. Domino, "The incidence of awareness during anesthesia: A multicenter United States study," *Anesthesia & Analgesia*, vol. 99, no. 3, pp. 833–839, 2004.

[16] B. A. Fritz, P. L. Kalarickal, H. R. Maybrier, M. R. Muench, D. Dearth, Y. Chen, K. E. Escallier, A. Ben Abdallah, N. Lin, and M. S. Avidan, "Intraoperative electroencephalogram suppression predicts postoperative delirium," *Anesthesia & Analgesia*, vol. 122, no. 1, pp. 234–242, 2016.

[17] G. A. Dumont and J. M. Ansermino, "Closed-loop control of anesthesia," *Anesthesia & Analgesia*, vol. 117, no. 5, pp. 1130–1138, 2013.

[18] C. F. Minto, T. W. Schnider, T. G. Short, K. M. Gregg, A. Gentilini, and S. L. Shafer, "Response surface model for anesthetic drug interactions," *Anesthesiology*, vol. 92, no. 6, pp. 1603–1616, 2000.

[19] Y. Sawaguchi, E. Furutani, G. Shirakami, M. Araki, and K. Fukuda, "A model-predictive hypnosis control system under total intravenous anesthesia," *IEEE Transactions on Biomedical Engineering*, vol. 55, no. 3, pp. 874–887, 2008.

[20] C. M. Ionescu, R. De Keyser, B. C. Torrico, T. De Smet, M. M. Struys, and J. E. Normey-Rico, "Robust predictive control strategy applied for propofol dosing using BIS as a controlled variable during anesthesia," *IEEE Transactions on Biomedical Engineering*, vol. 55, no. 9, pp. 2161–2170, 2008.

[21] E. Furutani, Y. Sawaguchi, G. Shirakami, M. Araki, and K. Fukuda, "A hypnosis control system using a model predictive controller with online identification of individual parameters," in *Proceedings of 2005 IEEE Conference on Control Applications (CCA 2005)*. Toronto, Canada: IEEE, 2005, pp. 154–159.

[22] G. Pannocchia and J. B. Rawlings, "Disturbance models for offset-free model-predictive control," *AIChE Journal*, vol. 49, no. 2, pp. 426–437, 2003.

[23] U.S. Food and Drug Administration, "Electroencephalograph," Code of Federal Regulations, Title 21, Section 882.1400, revised as of April 1, 2024, U.S. Government Publishing Office, govinfo.gov, 2024.

[24] H. Yu, B. Gao, and Y. Li, "Depth of anesthesia monitoring: Current evidence, clinical impact, and future directions," *Journal of International Medical Research*, vol. 54, no. 6, p. 03000605261457781, 2026.

[25] American Society of Anesthesiologists, "Standards for basic anesthetic monitoring," Standard of the American Society of Anesthesiologists, approved 1986, last amended October 15, 2025, asahq.org, 2025.

[26] U.S. Food and Drug Administration, "Infusion pump," Code of Federal Regulations, Title 21, Section 880.5725, revised as of April 1, 2024, U.S. Government Publishing Office, govinfo.gov, 2024.

[27] O. Vandemoortele, L. N. Hannivoort, F. Vanhoorebeeck, M. M. R. F. Struys, and H. E. M. Vereecke, "General purpose pharmacokinetic-pharmacodynamic models for target-controlled infusion of anaesthetic drugs: A narrative review," *Journal of Clinical Medicine*, vol. 11, no. 9, p. 2487, 2022.

3.10 Exercises

3.1 ○ Concept check · 2 min

In scenario 3 the target-controlled pump holds its model's propofol effect site within 0.01 µg/mL of the 1.7 µg/mL target, while the true index settles near 67, above the surgical band. Why does the propofol concentration loop not act on this error?

(a) Its gain is too low to remove an error that builds up over forty minutes

(b) Its measurement is the pump's own model, which no property of the patient reaches

(c) The monitor noise hides the error from the concentration controller

(d) Its integral is held at the pump's upper limit by the anti-windup

3.2 Concept check · 2 min

The incision of scenario 6 sets the surgical stimulation to $S = 0.85$. What does it do to the plant that the depth loop sees?

(a) It adds a fixed number of index units, so it acts like a load at the output

(b) It speeds the effect-site link k_{e0}, so the index answers the pump sooner

(c) It multiplies both half-maximal concentrations, so the operating point and the gain move together

(d) It shrinks the central volume, so the same infusion gives a higher plasma level

3.3 Concept check · 2 min

At the closing of scenario 11 the true index falls from 50.1 to 35.5 within a quarter of a minute. Which statement is right?

(a) A doubled depth loop gain would keep the index above 40, since the pump would be cut sooner

(b) Stopping the propofol pump at the instant of closing would keep the index above 40

(c) An opioid target of 8 instead of 4 ng/mL makes the dip shallower, since the opioid is short-acting

(d) The lowest index comes before an infusion change can reach the effect site, so only earlier lightening raises it

3.4 Theory and calculation · 5 min

The default patient is a man of 40 y, 70 kg and 170 cm. Use Equation (3.2) with $V_1 = 4.27\,\mathrm{L}$, $V_3 = 238\,\mathrm{L}$ and $\mathrm{Cl}_3 = 0.836\,\mathrm{L/min}$.

(a) Find the lean body mass L, the clearance Cl_1 and the rate constant k_{10}. *Checkpoint:* 55.3 kg, 1.638 L/min and 0.384/min.

(b) Find the ratio k_{13}/k_{31} and the time constant $1/k_{31}$, and say what the ratio means. *Checkpoint:* 55.7 and 285 min.

(c) Find the infusion that clearance alone needs to hold $C_1 = 1.7\,\mu\mathrm{g/mL}$ once every compartment is full. *Checkpoint:* 2.39 mg/(kg h).

3.5 Theory and calculation · 6 min

Use the response surface of Equation (3.3) with the defaults of Table 3.3: $C_{50,p} = 4.16\,\mu\mathrm{g/mL}$, $C_{50,r} = 19.0\,\mathrm{ng/mL}$, $\gamma = 1.43$, $\beta = 1.40$ and $E_0 = E_{\max} = 93$, with $\rho = \rho_r = 1$.

(a) Find the value of U that gives an index of 50, and the propofol effect site that gives it with no opioid. *Checkpoint:* 0.900 and 3.74 µg/mL.

(b) Find U_{50} at an even split, $\theta = 0.5$, and say how much less total normalized drug the pair needs than addition would predict. *Checkpoint:* 0.65, or 35 % less.

(c) The pair $C_{e,p} = 1.69\,\mu\mathrm{g/mL}$ and $C_{e,r} = 4\,\mathrm{ng/mL}$ holds an index of 50. Find the index at the same concentrations at $S = 0.85$, within 0.1. *Checkpoint:* 66.7.

3.6 Theory and calculation · 6 min

Scenario 3 runs target control on a patient with $\rho = 1.6$ and $\rho_r = 1$, with no stimulation. At 40 min the model holds 1.708 µg/mL of propofol and 2.00 ng/mL of remifentanil at the effect sites, and the patient's pharmacokinetics are those of the model.

(a) Find the true index from Equation (3.3), within 0.1. *Checkpoint:* 67.4.

(b) The population patient needs 2.67 µg/mL of propofol for an index of 50 at 2.00 ng/mL of opioid. Find what this patient needs, and its ratio to the target the pump holds. *Checkpoint:* 4.27 µg/mL and 2.5.

(c) The depth loop has $K_{c,B} = 0.045\,\mu g/mL$ per index unit. The process gain is -10.0 index units per µg/mL at 4 ng/mL of opioid and -18.8 at 8 ng/mL. Find the magnitude of the loop gain at each. *Checkpoint:* 0.45 and 0.85.

3.7 App challenge · 8 min

Load scenario 3, The same pump, a different patient. Use the default seed, run at speed 5 to 40 min and export the CSV file.

(a) Read the true index, column `bis`, and the model's propofol effect site, column `ce_model`, in the row at 40 min. *Checkpoint:* 67.4 within 0.2, and 1.708 µg/mL within 0.005.

(b) Press the reload button, set Patient sensitivity to propofol to 1.00 before Run, run to 40 min and read `bis` again. *Checkpoint:* 59.6 within 0.2.

(c) Load scenario 4, Closing the loop, run it as shipped to 39 min and read the propofol target, column `ce_target`, at 39 min. *Checkpoint:* 4.07 µg/mL within 0.02.

3.8 App challenge · 8 min

Load scenario 8, The measurement lies, upward. Use the default seed and run each case at speed 5 to 45 min, then export the CSV file.

(a) Run with Signal quality at Trust the monitor, as shipped. Read the lowest true index, column `bis`, the IAE tile and the propofol delivered, column `metric_prop_mg`, in the last row. *Checkpoint:* 31.8 within 0.2, IAE 367 within 3, and 333 mg within 2.

(b) Press Reset, set Signal quality to Hold on low quality before Run and run again. Read the same three values. *Checkpoint:* 48.9, 14 and 178 mg, with the same tolerances.

3.9 App challenge · 8 min

Load the challenge, A whole case. Use the default seed and run each case at speed 20 to 120 min. After each run read the challenge verdict (pass or fail and the worst index error from 79 min), the Light tile and the Hypotension tile.

(a) Run the case as shipped. *Checkpoint:* Fail, worst error 10.1, Light 99.7 min, Hypotension 0.0 min.

(b) Press the reload button, press Reveal suggested tuning and run. *Checkpoint:* Pass, worst error 3.3, Light 4.4 min, Hypotension 1.2 min.

(c) Press the reload button, press Reveal suggested tuning, then set Signal quality to Trust the monitor and Vasopressor to By hand before Run. *Checkpoint:* Fail, worst error 4.4, Light 4.4 min, Hypotension 11.3 min.

Accept the worst error within 0.2 and each tile within 0.3 min. Say what each run shows about structure against tuning.

3.10 Application area · 3 min

An index of 50 needs 3.74 µg/mL of propofol alone, or 1.69 µg/mL with 4 ng/mL of remifentanil. Take $a = 6.4$ mmHg per µg/mL and $b = 0.55$ mmHg per ng/mL from Equation (3.4).

(a) Find the steady difference in arterial pressure between the two splits with no stimulation. *Checkpoint:* About 10.9 mmHg higher with the opioid.

(b) Name two costs of an opioid-heavy split that the model does not price.

3.11 Application area · 4 min

A distillation column holds its top composition with a fast inner loop on a composition inferred from tray temperatures through a model fitted to past operation. A slow outer loop corrects the inferred value from an online analyzer.

(a) Name the counterpart in the anesthesia loop of the inference model, of the analyzer and of a change of feed that the model was not fitted to.

(b) The analyzer reads high for ten minutes during its own calibration. Say what the outer loop does with it and what the anesthesia chapter recommends.

3.12 Review the agent · 5 min

An AI assistant comments on scenario 11 as follows. “The index drops to 35.5 at the closing because the depth loop is too slow: it lowers the propofol target by only a fraction of a microgram per milliliter in the first minute. Double the depth loop gain to 0.090 µg/mL per index unit and the loop will cut the pump at once and keep the index inside the 40 to 60 band through the closing.” Find the flaw in the proposal and say how it could be checked.