INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 8 Batch, Run-to-Run, and Hybrid Control

CHAPTER 9

Aseptic Vial Filling Simulator With Drifting Heads and Late Checkweigher Results

Pages 221–242

Edited by John D. Hedengren

How to cite this chapter. "Aseptic Vial Filling Simulator With Drifting Heads and Late Checkweigher Results," in *Batch, Run-to-Run, and Hybrid Control*, John D. Hedengren, Ed., vol. 8 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 9, pp. 221–242.

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

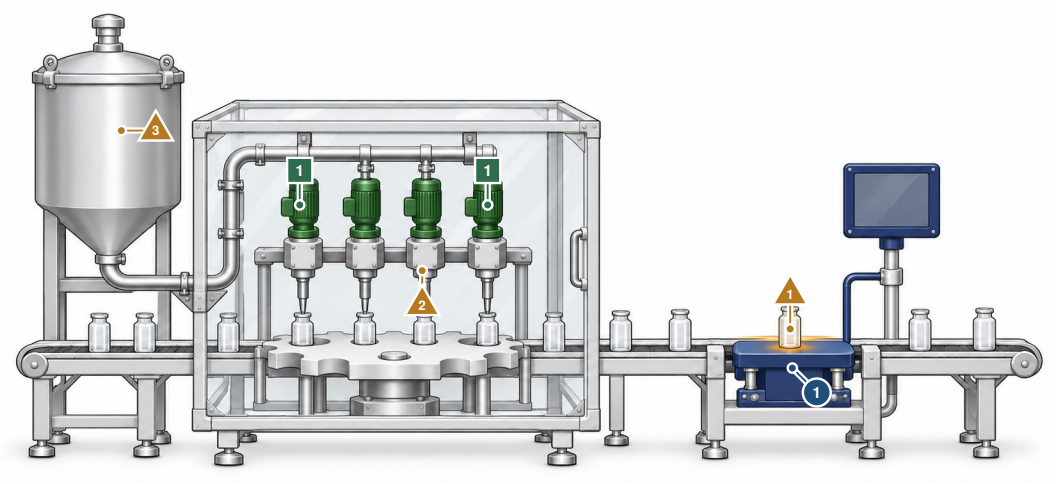
The app for this chapter
https://apmonitor.com/apps/aseptic-vial-filling/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v08-batch-hybrid/#aseptic-vial-filling

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

CHAPTER 9

Aseptic Vial Filling Simulator With Drifting Heads and Late Checkweigher Results



The controller moves 1 the pump behind each of the four needles

The controller reads 1 the scale that weighs a few of the filled vials

What disturbs it 1 weights that arrive after more vials are filled 2 a pump tube that wears 3 a thicker batch of liquid

Notice the distance between the needles and the scale: by the time a vial is weighed, its needle has already filled the next ones. Illustration, not to scale.

IN PLAIN TERMS **Four Needles and a Scale That Reports Late**

The main idea. A machine fills small glass bottles of medicine through four needles, each fed by its own pump. A scale weighs some of the bottles, and a controller uses those weights to nudge each pump up or down.

The challenge. Each weight arrives after the same needle has already filled several more bottles. A controller that treats every late weight as fresh news counts the same error again and again, and the pumps start to swing.

Why it matters. A short bottle that is never weighed can still be shipped, so how often the scale weighs matters as much as the controller.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	the first page, the first fact in the Introduction, the history in Section 9.2, the second try-it and the second fact in Section 9.3, the first try-it in Section 9.6, and the key points
Study	40 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	15 min	the first page, the results (Section 9.6), Technology in Practice (Section 9.7), Limitations (Section 9.8) and the open challenge

The control question A checkweigher weighs one vial per head every few cycles and reports it after later vials are filled. Which correction keeps every head on target without counting a late result twice, and what does the weighing interval decide that no correction can?

Learning objectives

- Explain why a pooled correction keeps the average fill right while one head fills under its target.
- Compute the cycle timing, pump speed ratio and spread factors from the line rate and the viscosity.
- Derive the stability limit of the integral correction with results in flight, and explain why the record-based estimate escapes it.
- Show from the dose records why canceling each whole error widens the fills.
- Choose the weighing interval, correction, guard band and target that release no vial outside its limits with the least giveaway.

Prerequisites Sections 1.4.5 (run-to-run control, the EWMA and integral forms and the metrology delay), 1.6.1 (engine and seed) and 1.6.4 (scenarios, scorecard and data). Chapter 7 applies the integral form to one machine with no result in flight.

The app https://apmonitor.com/apps/aseptic-vial-filling/

See also Volume 5, Chapter 6 for tablet weight held on a press, Volume 5, Chapter 9 for a correction from a valuation a year late, and Volume 5, Chapter 3 for control through a slow analyzer.

9.1 Introduction

Four peristaltic heads on the simulated line fill 72 vials a minute with 1.250 g of an antibody solution whose label minimum is 1.236 g. The whole margin between product given away and an underfilled vial is 14 mg, about 2.5 times the 5.7 mg spread of one replayed dose. In the first case of the application a correction that keeps the average of every weighed group on target lets one wearing head fill 10.4 mg under the target over its last 100 vials while it lifts the other three heads by 0.9 to 6.7 mg, and 32 underfilled vials leave the line accepted because three groups in four are never weighed. The average fill is right in that batch. One needle is not.

Aseptic filling of sterile drug and biological products follows current good manufacturing practice, for which a 2004 guidance replaced the 1987 guideline [1], and the dose is set by the filling pump. Peristaltic pumps are an industrial standard for filling high-concentration antibody formulations into vials and syringes, a precision target of 3 % at 0.3 mL has proved challenging for viscous liquids, and fill-weight variability tracked the height of the liquid drawn back at the end of each fill (the suck-back) rather than any pump setting [2]. In syringe filling the speed of that suck-back has also decided whether a filling needle clogged [3]. An open data descriptor reports 149 847 peristaltic doses from 0.1 to 2.0 mL taken under controlled conditions with calibrated weighing equipment [4, 5], and its authors have compared statistical and ensemble-learning compensation of the doses [6]. Balance readings have served as the feedback signal for the flow of a peristaltic pump [7], and adaptive filtering has been applied to the dynamic weighing of a checkweigher [8].

DID YOU KNOW? Product Worth Weighing

A 2025 dose study notes that biopharmaceutical products can cost up to 10 000 dollars a gram, so even small gains in dosing accuracy save product [4]. The suggested settings of this chapter's challenge give away 13.25 mg per accepted vial (Table 9.5).

Two beliefs are the simulator's target. The first is that a filling line has one fill weight, its average, when each head is its own pump, tube and needle. The second is that a correction which removes every measured error makes the fills tighter. A result is one dose with that dose's own noise, and it describes a vial filled tens of seconds earlier with a command that later corrections have already

changed. Adjusting a process after each deviation can increase the variance of its output, the lesson usually drawn from Deming's funnel, although active control to a target improves the output of the many processes whose disturbances are not of that kind [9], and which case applies rests on assumptions that often go unstated [10]. A measurement that arrives after later units have been processed is the metrology delay of run-to-run control, where delay combined with model error can destabilize a closed loop [11]. The same shape appears wherever discrete units are dosed from late, sampled measurements, from a silicon epitaxy reactor corrected run by run [12, 13] to a feed pump controlled from a balance [7]. Volume 5, Chapter 6 holds tablet weight on a rotary press, and Volume 5, Chapter 9 resets a contribution from a valuation a year late. A dead-time transfer function captures neither the sampling nor the identity of the unit a result describes.

However, open records of doses cannot show how another controller would have changed the doses they contain. In the compensated series of the dataset the command is a function of the compensation's own prediction, so command and disturbance are confounded, and only the uncompensated series recorded at a fixed command can be replayed through a different correction, under independence assumptions that have to be stated. This chapter sets fixed calibration, a pooled correction, a filtered per-head correction and a record-based per-head estimate side by side on a four-head filling simulator whose dose noise is replayed from those uncompensated records, with the tails of the fill distribution, the giveaway, the rejects and the recovery after a restart reported apart. Geometric moving averages [14], adjustment from past data [15], run-to-run control [13, 16] and its delay problem [11] are established, and the comparison rests on a declared action-response surrogate, not on prospective experiments. The simulator has a causal dispense model per head, a chronological split of the records into training and validation thirds, one write-once record per vial whose tare, gross weight, head and command stay aligned through every delay, and a bulk balance that closes to round-off, and the chapter gives every equation and parameter needed to build it again.

Section 9.2 follows the problem's history. Section 9.3 sets out the line: its cycle, the dispense, the drop, weighing, disposition and the parameters with their evidence. Section 9.4 splits the dose records and reads the noise-chasing lesson from them, and Section 9.5 gives the corrections as implemented. Section 9.6 measures them on eight cases and a challenge, the pipeline sweep and a policy comparison over seeds. Section 9.7 sets the line on real instruments and pumps, Section 9.8 states the limitations, Section 9.9 concludes, and the exercises of Section 9.10 run the cases.

9.2 Giveaway, Sampling and Feedback Adjustment

Setting a fill target is an older problem than correcting a filler. A graphical procedure published in 1977 chose the target of a process whose items are each checked against a lower specification limit, balancing the proportion sold at a reduced price against the give-away cost of setting the target further from the limit [17]. The simulated challenge poses that trade with one complication, since most vials are released on the evidence of the weighed vials around them rather than checked themselves.

The filter in every correction of the simulator came from control charts. A geometric moving average, which gives the newest observation the greatest weight and the older ones weights that fall geometrically, was proposed as a chart test in 1959 [14]. Three years later, keeping a variable such as the concentration of consecutive batches close to a target was posed as the estimation of the best adjustment from past data [15]. Whether to adjust at all was argued for decades afterward. Deming's funnel experiment illustrates that active control can increase the variance of a process output, and a 1990 note analyzed when that holds and how control to a target improves the output in the continuous process industries [9]. A 1992 discussion set statistical process monitoring beside feedback adjustment, illustrated the importance of assumptions that often go unstated and derived minimum-cost feedback schemes for simple models [10].

Semiconductor manufacturing turned these ideas into run-to-run control. A condition relating the EWMA weight to the estimated process gain was shown to keep an EWMA controller stable for a process that drifts or wanders [12]. A run-by-run controller that combined statistical process control

with feedback recovered a silicon epitaxy process within 3 runs after a disturbance and reduced its variation by a factor of 2.7 against historical data [13], and a 1997 review recast the methods from statistical and control viewpoints with self-tuning controllers that also handle trends and autocorrelation [16]. As production outran inspection, quality measurements arrived after later operations had started, and a sufficient condition for the stability of a multivariable EWMA controller under any length of metrology delay was derived from the eigenvalues of a model-mismatch matrix [11].

Filling followed a path of its own, from the 1987 guideline to the dose records of 2025, told in the history box below, and the simulator takes no number from the 2004 guidance [1]. The line in this chapter brings the two paths together, a run-to-run correction per head fed by a checkweigher whose results arrive late.

> ⧖ HISTORY **1987 to 2025: Filling Learns From Its Weights**
>
> A 2004 guidance on sterile drug products made by aseptic processing replaced the guideline of 1987. It counts any intervention or stoppage as a risk of contamination, and it names an on-line weight check device as one way to take a repeated manual task out of the critical area [1]. Bench studies of viscous antibody fills by peristaltic pump traced the variation of the suck-back height, which followed the fill weight, to two causes: a liquid stream that separated at an inconsistent point at the end of fill, and pressure variations on suck-back [2]. A dataset published in 2025 weighed one container tare and gross around every dose and calibrated the pump once at the start of every subset of 300 doses [4]. Compensation learned from such records raised dosing accuracy by 53.93 percent at 0.3 mL in a test on the real pump [6].

9.3 The Filling Line as Modeled

The simulated line is a teaching design drawn from published sources, and it copies no commercial filler. Figure 9.1 shows it. Qualified bulk, already sterile-filtered and transferred, waits in one vessel that feeds four peristaltic heads. Empty vials enter four at a time, one per lane, and an intermittent-motion track indexes the group one pitch per cycle. Every vial is tared on its lane's cell at the first pitch, filled under the needle of its lane's head at the second, weighed gross on the sampled cycles $D = 2$ pitches later while still open, given a stopper from a track fed by a vibratory bowl, checked by a presence sensor, and closed out as accepted, rejected or held. The tare and the gross weight belong to the same vial in the same closure state, so no stopper mass has to be compensated.

Each variable in Table 9.1 carries its role. The physical state is a set of event-driven quantities rather than integrated ones: the bulk mass, the liquid in vials, the drip tray, the purge cup, the hold-up in tubes and on needles and the accepted mass make six ledger states, and each head carries its gain G_h, its drop r_h, its relaxation a_h and its stroke counter n_h. Every vial has one sealed record with its head, group, empty mass, tare and its time, command, stroke, fill start and end, true liquid mass, bubble and drip flags, sampling flag, gross weight and its time, report time, net, stopper, sensor reading, aseptic flag, bracketing samples and disposition, each field written once.

9.3.1 The Cycle and Its Events

A cycle at line rate R (vials per minute, four per cycle) lasts $T_c = 60 \cdot 4/R$, of which t_{idx} is the index move and t_{set} a settle, so the fill window and the pump-speed ratio follow as shown in Equation (9.1).

$$T_f = T_c - t_{idx} - t_{set}, \qquad \frac{\omega}{\omega_{ref}} = \frac{T_f(R_{ref})}{T_f(R)} \tag{9.1}$$

where $R_{ref} = 60\,\text{vials/min}$ is the rate the dose statistics are referred to. At 72 vials/min, $T_c = 3.333\,\text{s}$, $T_f = 2.133\,\text{s}$ and the pump runs 1.3125 times its reference speed. A weighing cycle dwells at least $t_w = 3.0\,\text{s}$, which slows a line that weighs every group to 61.5 vials/min before stops. Each event of a cycle happens once, in order, at the process time it falls on: the index, the tare, four dispenses that read their heads' commands at their start, the gross weighing of a sampled group, the stoppers, the close-out, and each result's delivery a latency later. The time step of the engine (Section 1.6.1) is 0.1 s, a thirtieth of a cycle, and nothing is integrated between events.

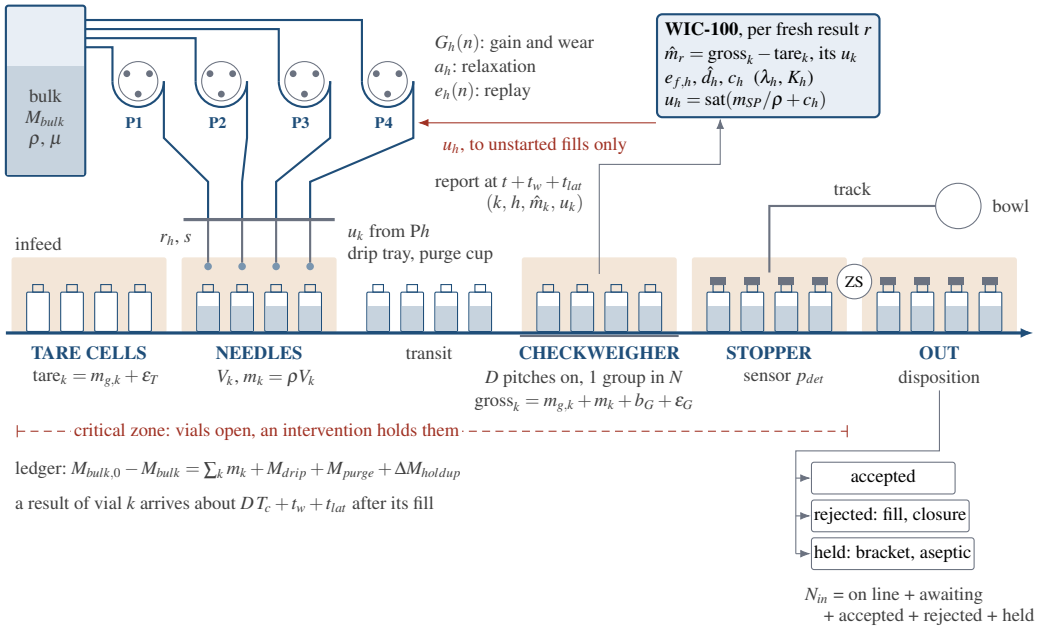


Figure 9.1 Four-head filling line as modeled, after the Line view of the simulator. The bulk (M_{bulk}, density ρ, viscosity μ) feeds heads P1 to P4 with calibration and wear gains G_h, relaxation a_h after a stop and replayed deviations $e_h(n)$ indexed by each head's stroke n. Under the needles vial k receives V_k for the command u_k its head held when the dispense started, with the hanging drop r_h set by the suck-back s. The checkweigher weighs one group in N with zero drift b_G, and WIC-100 reads only the reported records, net of each vial's own tare, and acts on dispenses that have not started. The critical zone runs from the tare cells to the stopper. Process model only, not a construction drawing. Not to scale.

Table 9.1 Variables, Their Roles and Their Reference Values

Variable	Role	Reference	Range
$\hat{m}_r$	CV, reported net of result r	—	1 mg steps
m_{SP}	setpoint, fill target (net)	1.250 g	1.230–1.290 g
u_h	MV, command of head h	1.2136 mL	1.10–1.32 mL
N	decision, weigh 1 group in N	4	1–20 cycles
R	decision, line rate	72	30–100 vials/min
s, n_p	decision, suck-back, purge shots	1, 0	0–3, 0–5
g_b	decision, reject guard band	0	0–5 mg
w_h	DV, tubing wear	0	0–5 %/1000 doses
μ	DV, bulk viscosity	1.0	0.8–12 mPa s
β_G	DV, gross-scale zero drift	0	±40 mg/h
x_e	DV, dose-noise multiplier	1	0–3
—	DV, events	—	lot, jam, glove, infeed
G_h, a_h, r_h	states, head (truth)	1, 0, 8 µL	—
b_G	state, scale zero (truth)	0	—
$e_{f,h}$, $\hat{d}_h$, c_h	states, controller	0	—
M_{bulk}	state, ledger	12 kg	—
m_k	output, true fill (truth)	—	—

9.3.2 The Dispense

The dispense is a causal surrogate, a calibrated volume command times the head's gain plus a recorded deviation, rather than a pump mechanism, as shown in Equation (9.2).

$$\begin{aligned} V_{p,k} &= u_k\, G_h(n)\, \frac{\gamma}{\gamma_0}\left(1 + a_h e^{-j/n_{rel}}\right) + x_e f_\sigma\, e_h(n) \\ V_k &= V_{p,k}\,(1-\beta_{air,k}) + r_h^- - r_h^+, \qquad m_k = \rho V_k \\ G_h(0) &= 1 + \sigma_{cal} z_h, \qquad G_h(n+1) = G_h(n) - 10^{-5} w_h \end{aligned} \tag{9.2}$$

where u_k is the command the head held when the dispense of vial k started, written into the record and never changed, n the head's own stroke count (a purge shot counts), j the doses since the last restart, r_h^- the drop on the needle before the fill and r_h^+ the drop it leaves, $\beta_{air,k}$ the share an air bubble displaces (zero without one), a_h the relaxation left by the last stop as a fraction of a dose, w_h the wear in percent per 1000 doses, x_e a dose-noise multiplier (1 by default) and z_h a unit normal drawn once per head and run. The response of a dose to its command, one gram of fill per gram asked after calibration, is a teaching assumption, because the records cannot identify it (Section 9.4). The pump speed and the viscosity change the delivery and widen the deviation as shown in Equation (9.3).

$$\begin{aligned} \gamma &= 1 + k_\omega\left(\frac{\omega}{\omega_{ref}} - 1\right) + k_\mu \ln\frac{\mu}{\mu_{ref}} \\ f_\sigma &= \left(\frac{\omega}{\omega_{ref}}\right)^{\beta_\omega}\left(\frac{\mu}{\mu_{ref}}\right)^{\beta_\mu} \end{aligned} \tag{9.3}$$

where γ_0 is γ at the rate and viscosity the heads were calibrated at, the start of the batch. At 72 vials/min on water $f_\sigma = 1.146$, and a lot at 6 mPa s takes 2.15 % off every head's delivery and widens its deviation 1.565 times.

9.3.3 Recorded Deviations and the Fitted Alternative

The default deviation $e_h(n)$ is replayed from the validation third of the dataset's uncompensated 1.2 mL series: 18 calibration segments of 300 doses, each less its own mean so that the calibration offset stays the separate state $G_h(0)$, converted from milligrams of water to volume at 998.2 kg/m^3. Head h starts at segment $[b+4(h-1)]$ mod 18 for a seed-chosen b and runs on through the following segments, indexed by its own stroke count so that a stop neither skips nor repeats a recorded dose. The fitted alternative is the mechanism behind the sign of the records' autocorrelation, a white pump term and a drop that leaves with the next fill, as shown in Equation (9.4).

$$e_k = p_k + r_{k-1} - r_k, \qquad \rho_1 = -\frac{\sigma_r^2}{\sigma_p^2 + 2\sigma_r^2} \tag{9.4}$$

where $p_k \sim N(0, \sigma_p^2)$, the drop is $\bar{r} + \sigma_r z$ and ρ_1 is the lag-1 autocorrelation. Fitting σ_p and σ_r to the training statistics gives 2.85 µL and 3.17 µL. Under the replay the drop is held at its mean $\bar{r} = 8$ µL, since the replayed deviation already carries the carry-over. Each stroke draws the same five random numbers whatever the settings, so a pair of runs that differ in a single setting meets identical disturbances.

9.3.4 Drop, Drips, Bubbles and Restarts

The suck-back setting s scales the drop by $\exp[-0.35(s-1)]$ and the drop grows with $(\mu/\mu_{ref})^{0.3}$. A drip, with probability $5\times10^{-4}\exp[-3(s-1)]$ per dose, drops the drop into the drip tray so the next vial lacks it, and a bubble, with probability $5\times10^{-4}\exp[2(s-1)]$, displaces $\beta_{air} = 4$ % of the dose (standard deviation 35 % of that). At the dataset's setting of 1 the two events are rare by construction. A pause longer than 10 s relaxes the tube by $a_h = a_{max}(1 - e^{-T_s/\tau_{rel}})$, with $a_{max} = -1.2$ % of a dose

and $\tau_{rel} = 300\,\mathrm{s}$, which the next doses recover over $n_{rel} = 2$ doses through the factor of Equation (9.2), and a stop longer than $t_{drop} = 30\,\mathrm{s}$ lets each hanging drop fall to the drip tray. A restart purge sends n_p doses per head to the purge cup before the first vial. All of these laws are teaching assumptions shaped by the qualitative findings on suck-back [2, 3].

> **TRY IT IN THE APP** **7 · Stop and Restart**
>
> Load the scenario, open the Restart view and press Run. The line stops twice for 4 min, and the first fills after a restart come out 10.2 to 23.7 mg under the target, with 8 vials rejected and 36 held. Press Reset, set Restart purge shots to 2 and run again: the first fills lie between −9.0 and +2.2 mg, with 2 rejected and 12 held.

9.3.5 Weighing and the Record

The tare and gross readings carry their own noise and quantization, and the gross scale's zero is a hidden state apart from any pump gain, as shown in Equation (9.5).

$$\begin{aligned} \mathrm{tare}_k &= q\,\mathrm{round}\big[(m_{g,k} + \sigma_s z_{T,k})/q\big] \\ \mathrm{gross}_k &= q\,\mathrm{round}\big[(m_{g,k} + m_k + b_G + \sigma_s z_{G,k})/q\big] \\ \hat{m}_k &= \mathrm{gross}_k - \mathrm{tare}_k, \qquad \frac{\mathrm{d}b_G}{\mathrm{d}t} = \beta_G \end{aligned} \tag{9.5}$$

where $m_{g,k}$ is the empty vial's mass (6.5 g, standard deviation 80 mg), $\sigma_s = 0.5\,\mathrm{mg}$ is the dataset scale's stated repeatability bound used as a standard deviation, and $q = 1\,\mathrm{mg}$. A weighing at time t is reported at $t + t_w + t_{lat}$, with $t_{lat} = 0.5\,\mathrm{s}$ by default, 40 s in case 3 and 20 s in the challenge. The first group after a restart and the last group of the batch are always weighed. Without that restart check case 7 accepts 7 underfilled vials instead of 3, with fewer rejects in view. A 2 g reference weight on the gross scale reads b_G plus noise, the only instrument that tells a drifting zero from a drifting head.

9.3.6 Disposition and the Ledger

Each vial is closed out once at OUT by teaching rules. A stopper the sensor finds missing (probability of detection 0.995) rejects it for closure, and a vial open in the critical zone during a glove intervention is held, whatever it weighs. A weighed vial is accepted when its net lies within the fill limits narrowed by the guard band g_b and rejected otherwise. An unweighed vial is accepted only when the weighed vials of its own head before and after it, in fill order, both read inside the narrowed limits, and held otherwise, which leaves it awaiting until its later bracket is reported. The limits are teaching values, 1.236 g (the 1.20 mL label at 1.030 g/mL) and 104 % of it, 1.2854 g. The bulk balance is the ledger of Equation (9.6).

$$M_{bulk,0} - M_{bulk} = \sum_k m_k + M_{drip} + M_{purge} + \Delta M_{holdup} \tag{9.6}$$

where the hold-up is the liquid in four tubes of 2.26 mL and on four needles. The giveaway is the mean of m_k less the label minimum over the accepted vials, a quantity of teaching truth that no instrument on the line reads.

> **DID YOU KNOW?** **Stopping the Line Ships More Vials**
>
> From the simulator: when the stopper bowl jams for 3 min and the line keeps running, 148 vials leave unstoppered, all are found and rejected, and 2201 are accepted. Stopping on a low stopper track leaves none unstoppered and accepts 2354, although the line then makes 62.0 against 66.2 vials/min (Table 9.5).

9.3.7 Parameters and Evidence

Table 9.2 lists every parameter with its default and its label, measured (M) when read from the descriptor or its data, fitted (F) when estimated here from the training segments, and teaching (T)

Table 9.2 Parameters, Defaults and Evidence (M Measured, F Fitted, T Teaching)

Symbol	Meaning	Value	Unit	Ev.
$e_h(n)$	replayed deviation	18 seg.	—	M
σ_{cal}	offset one calibration leaves	0.363	% of dose	F
σ_p, σ_r	fitted pump, drop terms	2.85, 3.17	µL	F
s	suck-back of the records	1	device	M
σ_s, q	scale noise, resolution	0.5, 1	mg	T, M
ρ	bulk density	1.030	g/mL	T
V_{label}	label volume	1.20	mL	T
—	upper limit	104	% of label	T
g	fill per command, calibrated	ρ	g/mL	T
μ_{ref}	viscosity of the records	1.0	mPa s	T
k_ω, β_ω	speed: delivery, spread	−1.0 %, 0.5	—	T
k_μ, β_μ	viscosity: delivery, spread	−1.2 %, 0.25	—	T
$\bar{r}$	hanging drop	8	µL	T
a_{max}, τ_{rel}, n_{rel}	tube relaxation	−1.2, 300, 2	%, s, doses	T
t_{drop}	stop before the drop falls	30	s	T
t_{idx}, t_{set}	index move, settle	0.9, 0.3	s	T
R_{ref}	rate of the records	60	vials/min	T
t_w, D	weighing dwell, position	3.0, 2	s, pitches	T
t_{lat}	result latency	0.5	s	T
m_g	empty vial, mean and SD	6.5, 0.08	g	T
—	stopper track, bowl feed	60, 2.5	–, /s	T
—	track stop, restart levels	12, 40	stoppers	T
p_{det}	stopper detection	99.5	%	T
v	validity window	2.5	% of m_{SP}	T
Δu_{max}	largest correction step	1	% of u_{nom}	T
K_h, λ_h	gain, filter	0.50, 0.30	mL/g, —	T

when no source fixes it. The batch starts from a checkpoint, a primed and calibrated line with 12 kg of bulk and no vial filled.

9.3.8 Assumptions

The dispense is an empirical surrogate, so pressure, compliance, tube geometry and the roller pulsation are absent, and the model cannot say how a different tube or pump would behave. Its deviations are those of water on one pump and one container, and replaying them through another correction is valid only under three independence assumptions: a deviation does not depend on its command within the range the corrections move it, the four heads' deviations are independent of one another, and a deviation depends on the head's stroke count rather than on the line's timing or its stops, with a constant factor f_σ carrying the pump speed. The commands stay within 2.7 % of nominal outside cases 3 and 5, reach 3.4 % in case 5 and 6.7 % in case 3's oscillation. The drop, drips, bubbles, relaxation, viscosity and speed laws are teaching closures. The density does not change with the lot. Stoppering is a presence check, and crimping, container-closure integrity, visual inspection and contamination are not modeled. An aseptic hold is a discrete operational state, never inferred from a weight, so a run with no rejects is not evidence of sterility.

9.4 The Dose Records

The dataset's core 1.2 mL file holds 16 195 doses separated by 54 calibration markers into 54 segments, weighed at 0.001 g, and the replayed segments are used under the dataset's CC BY license [5]. The split is chronological, to avoid leakage between neighboring subsets: the first 36 segments (10 795 doses) train and the last 18 (5400 doses) validate, so that no validation segment lies between two training segments. The statistics are given in Table 9.3. The within-segment spread and the lag-1 autocorrelation differ between the thirds, which is why the application replays the validation third by default and uses the model fitted on the training third only as an option. The fitted model reproduces its training statistics and overstates the validation spread by 10 %.

Table 9.3 The 1.2 mL Records Split by Calibration Segment, and the Next-Dose Prediction Error of an EWMA Restarted at Each Calibration

Statistic, 1.2 mL	Training	Validation
segments, doses	36, 10 795	18, 5400
within-segment SD (mg)	5.31	4.81
SD of the segment means (mg)	4.35	3.90
lag-1, lag-2 autocorrelation	−0.356, +0.084	−0.228, −0.129
drift (mg per 100 doses)	−0.05 ± 0.32	−0.22 ± 0.19
first 3 doses less the rest (mg)	−0.02	+0.10
RMSE, target only, $\lambda = 0$ (mg)	6.82	6.12
RMSE, EWMA $\lambda = 0.05$ (mg)	5.47	4.95
RMSE, last dose, $\lambda = 1$ (mg)	9.33	7.98
Validation RMSE, other volumes	fixed / best / last	λ
0.2 mL (mg)	4.33 / 3.39 / 4.99	0.10
0.3 mL (mg)	5.46 / 3.70 / 5.96	0.05
2.0 mL (mg)	10.31 / 6.71 / 10.76	0.05

Successive doses are negatively correlated, a dose that runs high tends to be followed by one that runs low, and the spread of the segment means (4.35 mg on training) is what one recalibration leaves on a head, which fixes $\sigma_{cal} = 0.363\,\%$. A drop that hangs on the needle and leaves with the next fill gives that sign, consistent with the suck-back mechanism [2], although the records cannot tell it from a reading artifact. A gross reading reused as the next tare would give a lag-1 of about −0.01 at the stated scale repeatability of 0.5 mg, far from −0.36. Within a segment the slope spread matches the 0.35 mg per 100 doses that white noise gives in 300 doses, so no drift is identified, and the first doses after a calibration do not differ from the rest beyond –0.02(48) mg (standard error over 36 segments), which bounds the effect of a short pause but not of a stop of minutes.

The prediction rows of Table 9.3 are the data's own statement of the noise-chasing lesson. An EWMA of the dose error restarted at the target after each calibration, with its weight chosen on training, predicts the next validation dose to 4.95 mg, against 6.12 mg for the target alone and 7.98 mg for the last dose. The order holds at every volume in the lower rows. Under the independence assumption a correction that cancels the predicted error leaves exactly the prediction error, so these are also the corrected-dose spreads of a perfect actuator. They are prediction results, not closed-loop measurements. The records cannot identify the response of a dose to its command, because in the compensated files the command is $V_c = V_{tg}^2/V_p$, a deterministic function of the compensation's prediction V_p of the dose at target V_{tg}, so that $(V_c/V_{tg})V_p \equiv V_{tg}$ and command and disturbance cannot be separated.

9.5 The Corrections as Implemented

The fill-weight controller WIC-100 executes every 0.1 s but changes a command only at a fresh result. Its inbox holds, for each reported vial, the vial and head identifiers, the net $\hat{m}_r$ of that vial's own gross and tare, the command u_r that vial was filled with and the report time. It never reads a true fill, a pump gain or the scale's zero. A result is first checked for validity against the median of its head's last five results, valid or not, as shown in Equation (9.7).

$$\left|\hat{m}_r - \mathrm{med}_5(\hat{m}_h)\right| \le v\, m_{SP} \tag{9.7}$$

where the median is replaced by m_{SP} before three results and $v = 2.5\,\%$ (31.25 mg). An invalid result is counted and changes nothing. The median follows a real shift after three results, so the window cannot lock a head out, while a bubble's 50 mg short dose falls outside it. A window centered on the target instead froze the correction of case 3 for the rest of the batch, with 1020 results ignored and 282 vials accepted against 648. The command of head h is the recipe's $u_{nom} = m_{SP}/\rho$ plus a

correction c_h, updated by one of three laws, as shown in Equation (9.8).

$$\begin{aligned}
\text{per head: } & e_{f,h} \leftarrow \lambda_h(m_{SP} - \hat{m}_r) + (1-\lambda_h)e_{f,h} \\
& c_h \leftarrow c_h + \lim(K_h e_{f,h}) \\
\text{record: } & \hat{d}_h \leftarrow \hat{d}_h + \lambda_h(\hat{m}_r - \rho u_r - \hat{d}_h) \\
& c_h \leftarrow c_h + \lim\big(K_h\rho(-\hat{d}_h/\rho - c_h)\big) \\
\text{pooled: } & \bar{e}_f \leftarrow \bar{\lambda}\,(m_{SP} - \bar{m}_g) + (1-\bar{\lambda}\,)\bar{e}_f \\
& c_h \leftarrow c_h + \lim(\bar{K}\bar{e}_f) \quad \text{for every } h \\
& u_h = \mathrm{sat}(u_{nom} + c_h), \qquad c_h = u_h - u_{nom}
\end{aligned} \tag{9.8}$$

where lim caps a step at $\Delta u_{max} = 1\,\%$ of u_{nom}, sat clamps to 1.10–1.32 mL with the correction held at the clamp so nothing winds up, $\bar{m}_g$ is the mean of a weighed group's valid results and $\bar{\lambda}$, $\bar{K}$ are the means of the four heads' constants. Fixed calibration applies no update. The defaults are $K_h = 0.50\,\mathrm{mL/g}$ and $\lambda_h = 0.30$ on every head, and $u_{nom} = 1.213\,59\,\mathrm{mL}$ at the 1.250 g target. An update changes only dispenses that start after it, because a dispense reads its head's command at its start. In Manual each head's station sets its command, the correction follows the station and the filters keep filtering, so a return to Auto continues from the station's value.

The plain per-head law integrates a filtered error, and each result adds its error to a command that the corrections from earlier results have already moved. When several results of a head are in flight, the same error is counted once for each of them. With the filter at 1.00 it is the integral form of Section 1.4.5, whose characteristic polynomial with a delay of D results is Equation (1.9), and Chapter 7 applies the same form to pack pressure from each weighed part. The record-based law estimates the head's offset $\hat{d}_h$ from each result less the fill its own command would have given, so a late result is compared with the command that produced it and is never counted twice. That is the intercept form of the EWMA run-to-run controller, Equation (1.7) [12, 13], with a partial move of the command toward the estimate's target, and at rest it gives $\rho u_h + \hat{d}_h = m_{SP}$. A sufficient condition for such a controller to stay stable at any length of metrology delay places the model-mismatch eigenvalues within the unit circle centered at one [11]. Here the mismatch is the scalar ratio of a head's true delivery to the declared ρ per milliliter, which stays between 0.943 and 1.006 over every case as shipped, and the record-based immunity measured below is consistent with that condition rather than proved by it. The immunity rests on the declared response gain. A plant whose dose did not follow its command one for one would test it.

Pause. Try Exercise 9.2 before reading on.

The application keeps three models apart. The plant is everything in Section 9.3. The identified model is the data's dose statistics and the declared response $m = \rho u + d$. The controller's model is that declared response and nothing else, so no correction knows the wear, the viscosity, the drop or the scale's zero. An ideal-information policy, which sets each command before each dispense from the head's true gain and the delivery factor, is a teaching bound that reads the simulator's truth and is reported apart.

9.6 Measured Results

Each number in this section is a measurement on the application's own engine, driven headless at seed 20261021 (the default) through the step sequence the page uses, unless a range over seeds is given. Settings named in a row are applied before the batch begins. The giveaway and every count of accepted vials outside the limits are teaching truth. The exported CSV carries each head's reported results, commands and requested commands beside the true fills, gains and scale zero, whose columns carry a "true" suffix in their names, so a replay of the file through a correction has to leave them out by name.

Table 9.4 Code Verification: Identities Checked and Their Residuals

Identity	Result
Bulk ledger, Eq. (9.6), 13 runs with drips, purges, viscosity and scale drift	worst 3.2×10^{-14} relative
Sum of record fills against the in-vial state, hold-up against $4(V_{tube} + r_h)\rho$	0, 1.9×10^{-16}
Each dispense's command equals the correction standing at its start, 14 runs	36 000 dispenses, 0 differ
Updates citing an unreported result, results used twice, nets not of one record	0 of 10 219, 0 of 12 052, 0
Results held back (latency 10^6 s)	0 command updates
Counts reconcile: in = accepted + rejected + held + awaiting + on line	8667 instants, 0 fail
Replayed stroke order through two stops and purges	0 mismatches
Scale zero drifting 30 mg/h: pump gains	bit-identical
Step halved to 0.05 s, 12 runs	0 dispositions or readings changed, fills within 0.002 mg
Playback at 120× against stepping	identical records

9.6.1 Code Verification

The checks that the model does what Section 9.3 says are listed in Table 9.4. They verify bookkeeping and causality, not physics. The bulk ledger in particular closes by construction of the dispense step, so its residual shows that no event loses or doubles liquid. All 46 checks of the application's own test suite pass, its ten completion checks pass, and every case loads identically when chosen first, when chosen after another case and after Reset.

9.6.2 Eight Cases and a Challenge

The nine scenarios appear in Table 9.5 with the measured contrast each one teaches. All share the reference configuration of Table 9.1 and differ from it only where a row says so.

> **TRY IT IN THE APP** **1 · One Needle Drifting, a Pooled Correction**
>
> Load the scenario, open the Heads view and press Run. Watch head 3: over its last 100 vials it fills 10.4 mg under the target while the other three heads are lifted by 0.9 to 6.7 mg, and 32 underfilled vials are accepted. Press Reset, set Correction to Per-head correction and run again: head 3 ends within 0.7 mg.

The counts of accepted underfills in Table 9.5 are the default seed's and they depend on which stretch of the records each head replays. Over seeds 1 to 3 the head means, the spreads and the accepted and held counts keep the direction of every contrast, but the accepted underfills under the per-head correction rise from 2–4 at the default seed to 27–81 in cases 4, 6 and 8, and in case 1 at seeds 2 and 3 the per-head correction lets 36 and 51 underfills through against 30 and 46 under the pooled one. Its accepted count still rises in case 1 at every seed, from 1324–1530 to 1553–1579. The remedy of case 7 lifts the mean first fill after a restart by 8.7 to 13.9 mg at all four seeds, with the mean first fill moving from between −17.4 and −13.2 mg to between −6.6 and +0.1 mg. The scorecard leads with the accepted vials and the giveaway, and its per-head tiles (the integral of the held reported error, command travel and time at a limit) are diagnostics only. In case 4 the command travel of the four heads is 0.69 to 1.19 mL as shipped against 0.10 to 0.14 mL with the filtered correction, the same order as their true spreads.

Pause. Try Exercise 9.1 before reading on.

9.6.3 Pipeline Length and the Plain Law

Case 3 is where the two per-head laws part. With one group in two weighed and 40 s of latency, a result arrives 50.2 s after its fill, its head has filled 12.8 more vials meanwhile, and 5.9 of them were weighed, so about six results per head are always on their way. The plain law at its default gain swings head 1's command from 81 µL under its nominal 1.2136 mL to 78 µL over it. The sweep of Figure 9.2 varies the latency at four gains. At 0.50 mL/g the plain law holds a head spread of 6.05 mg at 0.5 s and 7.86 mg at 5 s, and oscillates from 20 s on (36.7 mg). At 0.25 mL/g it holds to 20 s (6.95 mg) and

Table 9.5 The Nine Scenarios and What Each One Measures at the Default Seed

#	Configuration and events	Measured contrast
1	1600 vials, head 3 wears 5 % per 1000 doses, pooled correction	head 3 −10.4 mg and heads 1, 2, 4 +6.7, +3.4, +0.9 mg over their last 100 vials, 32 underfills accepted. Per head: head 3 within 0.7 mg, 3 accepted, held 59 to 18
2	4800 vials, all heads wearing 1.5–4.5 %, fixed calibration	head 4 from −16.1 mg (10–15 min) to −62.1 mg (last 20 min), 2499 held, 713 rejected, 1588 accepted. Per head from 15 min: 4519 accepted
3	results 40 s late, 1 group in 2, per-head law	50.2 s from fill to result, 5.9 weighed vials of the head meanwhile. Head SD 47–48 mg, 648 accepted. Record-based 4.6–6.8 mg, 2394 accepted
4	filter 1.00, gain 0.97 mL/g	head SD 6.1–9.2 mg, 2141 accepted, 44 underfills accepted. At 0.30 and 0.50 mL/g: 5.0–7.1 mg, 2365, 2
5	a 6 mPa s lot at 300 s	head 1 −8.8 mg for two minutes, within 0.4 mg from 15 min, SD 5.1 to 9.7 mg, 66 underfills. At 56 vials/min 35, at 1.256 g 14 with giveaway 14.8 to 20.2 mg
6	suck-back 2.8	42 bubbles, head SD 8.2–10.0 mg, 33 underfills. At 1.0: 1 bubble, 2. At 0: 27 drips, 6
7	two 4 min infeed stops, no purge	first fills 10.2–23.7 mg under, 8 rejected, 36 held. With 2 purge shots: −9.0 to +2.2 mg, 2 rejected, 12 held, 19.8 g purged
8	stopper bowl jam 3 min, glove intervention, line keeps running	148 unstoppered, all found and rejected, 2201 accepted. Stopping on a low track: 0 unstoppered, 2354 accepted at 62.0 against 66.2 vials/min. 16 aseptic holds both ways
9	challenge: head 3 wearing, results 20 s late, a 4 min stop. Pooled, filter 1.00, gain 0.90 mL/g, 1 in 4	shipped 1598 accepted, 20.9 mg giveaway, 92 under and 18 over accepted. Suggested: 2329, 13.25 mg, 55.8 vials/min, none outside

fails at 30 s (17.5 mg), and at 0.10 mL/g it stays between 5.81 and 5.99 mg to 60 s. The record-based estimate at 0.50 mL/g stays between 5.81 and 5.98 mg at every latency. The same pattern holds when every group or one in four is weighed, and weighing every group lengthens the pipeline in results per head, so that the plain law at 0.10 mL/g reaches 61.3 mg at 60 s with every group weighed while the record-based estimate holds 5.7 mg. Gain at 0.97 mL/g, one result's error divided by the density, chases the noise even at 0.5 s (13.9 mg).

9.6.4 Four Corrections on the Challenge Line

The policy comparison of Table 9.6 uses the challenge line with the suggested settings (every group weighed, a 3 mg guard band, two purge shots and a 1.249 g target) and changes only the correction, then repeats the comparison with one group in four weighed. Each row runs 12 paired seeds on the validation replay. The seeds are not 12 independent noise realizations, since they start the replay on only 9 distinct segments of 18, while every seed draws its own calibration offsets and scale readings. The fill-distribution tail is reported as the 0.1st percentile of all true fills and as the share of true fills under the label minimum, before any disposition, and medians of 12 are rounded to whole vials. The training replay and the fitted model give the same pass count in every row and the same order, apart from the two near-equal rows.

Three results follow. The plain law at the default gain and the pooled correction oscillate on this line, because every vial weighed with 20 s of latency is a long pipeline, and their head spread of 53 to 61 mg sends 1835 to 1881 vials to the reject chute. Fixed calibration chases no noise, but it leaves head 3's wear and each head's calibration offset in place, so its head spread over the batch is 7.1 mg and 19 % of its fills fall under the label minimum. The record-based estimate and the plain law at 0.10 mL/g are equal within the seed spread, 2325 against 2324 accepted and 13.20 against 13.22 mg of giveaway, and both come within 18 vials and 0.06 mg of the ideal-information bound. On this line the record-based estimate's advantage disappears once the plain law is detuned for the pipeline it meets. Figure 9.3 draws the medians of the table as bars against the challenge bar.

A higher gain buys little here. With a 6 mPa s lot arriving at 600 s on the same line, the mean true fill error in the next two minutes is −11.0 mg under the record-based estimate at 0.50 mL/g

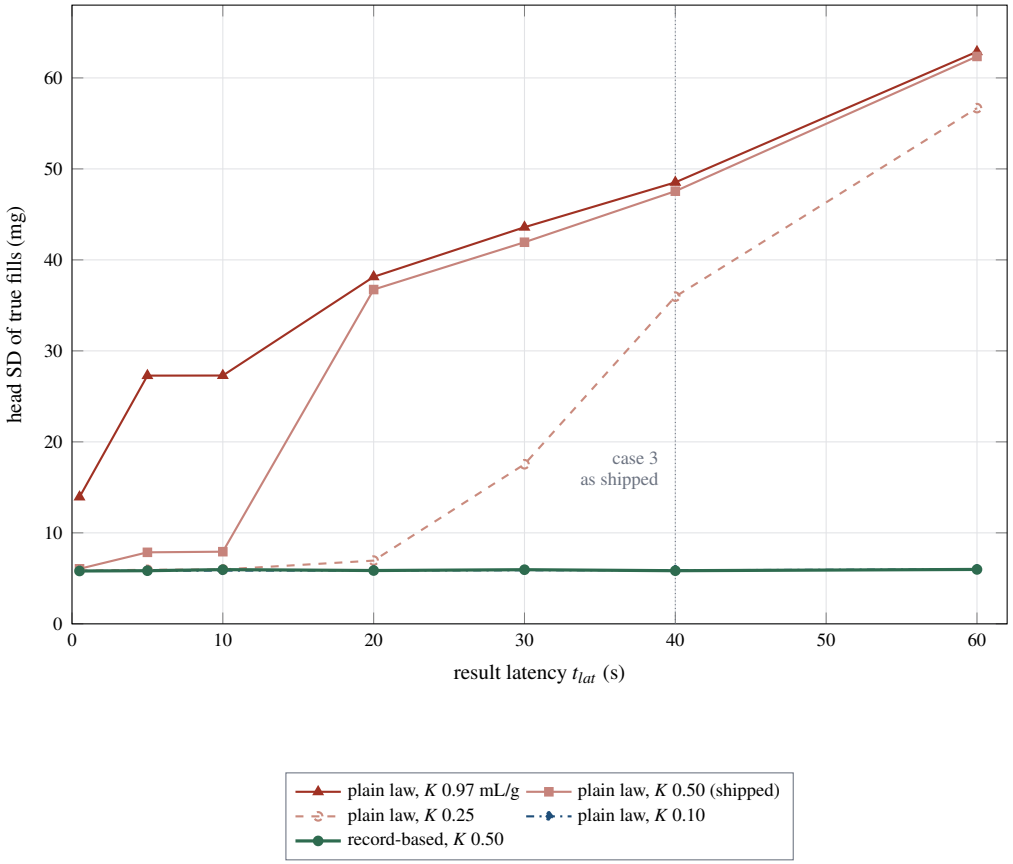


Figure 9.2 Mean head spread of the true fills against the checkweigher's result latency, with the settings of case 3 (one group in two weighed, filter 0.30, default seed). The plain per-head law loses stability at a latency that falls as its gain rises, while the record-based estimate is unaffected. The plain law's gains are marked by triangles and squares on solid lines, open circles on a dashed line and diamonds on a dash-dotted line, in falling order of gain, and the record-based estimate by dots on a thick solid line.

and −12.2 mg under the plain law at 0.10 mL/g, and over the two minutes after that −0.3 mg against +1.6 mg, with 2093 and 2094 vials accepted, while the plain law at 0.25 mL/g already oscillates (672 accepted, medians over 12 seeds). The record-based estimate's value is that its gain can be chosen for the disturbance without knowing the length of the pipeline. Fixed calibration stays 29.5 mg under. The restart recovery is the purge's, not the correction's. With the record-based estimate weighing every group, the first fill of each head after the 4 min stop is −14.6 mg with no purge, −5.1 mg with one shot and −1.2 mg with two (medians over the 12 seeds).

🕮 **Pause.** Try Exercise 9.3 before reading on.

9.6.5 What No Correction Sees

A gross scale whose zero drifts by 30 mg/h in case 1, with head 3's wear removed and the per-head correction on, leaves every pump gain bit-identical, and the correction lowers the commands by 11.1 µL, close to the bias over the density, 11.4 µL. The reported mean of the last 100 vials stays +0.6 mg from the target while the true mean is −9.7 mg, and the accepted underfills rise from 2 to 158. A correction that trusts its scale turns a scale bias into an underfill, and only the reference-weight

Table 9.6 Corrections on the Challenge Line Over 12 Paired Seeds, Validation Replay: Median (Range)

Correction	Pass	Accepted	Giveaway (mg)	Fill rejects	Held	Head SD (mg)	0.1st pct. (mg)	Under LSL (‰)
Every group weighed								
fixed calibration	0	1861 (1717–2124)	14.09	539	0	7.1	−39.9	193 (78–230)
pooled, K 0.50	0	566 (533–588)	22.87	1835	0	61.1	−127.6	427
per head, K 0.50	0	519 (504–551)	23.70	1881	0	53.2	−101.3	437
per head, K 0.10	12	2324 (2313–2352)	13.22	77	0	6.0	−16.1	3.1 (1.7–5.8)
record-based, K 0.50	12	2325 (2307–2348)	13.20	75	0	6.1	−16.2	4.2 (2.9–6.7)
ideal information	12	2342 (2325–2356)	13.26	59	0	5.9	−17.1	3.1 (2.1–4.2)
1 group in 4 weighed								
fixed calibration	0	1793 (1577–2021)	13.45	141	475	7.1	−39.9	193
per head, K 0.10	0	2225 (2137–2301)	12.71	31	146	6.1	−17.7	10.4 (4.6–24.2)
record-based, K 0.50	0	2226 (2146–2301)	12.97	31	144	6.2	−18.7	13.3 (4.2–26.7)
ideal information	0	2292 (2267–2340)	12.98	16	93	5.9	−17.1	3.1 (2.1–4.2)

Every row weighing every group accepts no vial truly outside the limits, so the tail appears as fill rejects. With one group in four, the accepted underfills (teaching truth) are 9 (5–23) under the record-based estimate, 8.5 (5–16) under the plain law at 0.10 mL/g and 6 (3–10) under ideal information, which is why no row passes. The pooled correction and the plain law at 0.50 mL/g accept 1752 and 1861 vials with one group in four and are not listed.

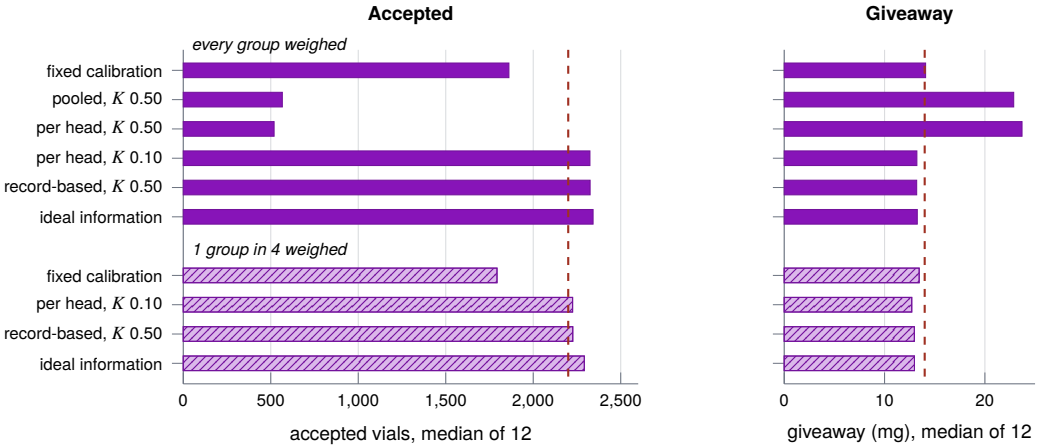


Figure 9.3 The corrections of Table 9.6 on the challenge line, solid with every group weighed and hatched with one group in four. The dashed lines are the challenge bar of 2200 accepted vials and 14 mg of giveaway. The pooled correction and the plain law at 0.50 mL/g lose most vials to the reject chute, while the detuned plain law and the record-based estimate sit 1 vial apart and within 18 vials of the ideal-information bound.

check separates the two. Viscosity and suck-back are the other limits, since they widen the spread that no command can narrow (cases 5 and 6), and they are paid for in line rate, in fill target or in the suck-back setting.

9.6.6 The Challenge

The challenge fills 2400 vials at 72 vials/min with head 3 wearing 4 % per 1000 doses, results 20 s late and a 4 min infeed stop at 15 min. Its target is at least 2200 vials accepted, at most 14 mg of giveaway, at least 55 vials/min with stops included, and no accepted vial truly outside its limits or without a stopper. The bar of 2200 was set so that the verdict holds on the validation replay, on the training replay and on the fitted model at 12 seeds each. As shipped (pooled, filter 1.00, gain

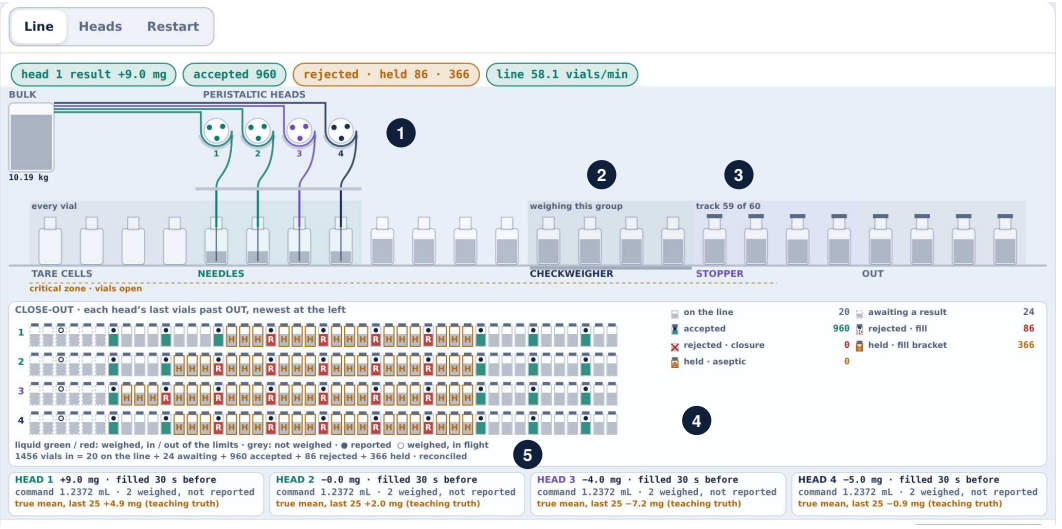


Figure 9.4 The challenge as shipped at $t = 1500$ s in Auto with the pooled correction, cropped to the Line view. The line is filling again after the infeed stop from 903.5 to 1140 s. (1) The bulk vessel feeds the four peristaltic heads, whose needles fill the group after the tare cells. (2) The checkweigher weighs a sampled group gross. (3) The stopper track leads to OUT, where each vial is closed out. (4) The close-out rows show each head's last vials past OUT, many held on their brackets, and reconcile 1456 vials as 20 on the line, 24 awaiting, 960 accepted, 86 rejected and 366 held. (5) Each head's chip shows a last result filled 30 s before its report with two weighed vials in flight and the command of 1.2372 mL at which all four heads stand, and head 3's true mean over its last 25 vials is −7.2 mg (teaching truth).

0.90 mL/g, one group in four, 1.250 g, no purge) the batch fails at all 12 seeds on every noise source, with 1598 accepted, 20.88 mg of giveaway and 92 under and 18 over accepted at the default seed. The suggested settings pass at all 12 on every source, accepting 2307 to 2348 on the validation replay, 2209 to 2350 on the training replay and 2248 to 2279 on the fitted model, at 55.8 vials/min. The settings of Figure 9.4 are the shipped ones at 1500 s.

Each piece of the suggested settings was ablated over six seeds in the application's benchmark file. Without the guard band 2 to 6 underfills pass at every seed, and over 12 seeds a band of 0, 1, 2, 3 and 4 mg passes 1, 5, 11, 12 and 12 of them. Without the purge the batch still passes, with 79 fill rejects against 75 (medians over 12 seeds), so the purge buys rejects rather than the verdict. Weighing one group in four instead fails at every seed even with ideal information, because unweighed underfills pass on their brackets, and filling fuller instead, one in four at 1.256 g, gives 19.40 to 19.95 mg of giveaway and still releases 1 to 4 underfills. Over 252 partial-weighing settings of the record-based and ideal policies, with one group in 2, 3 or 4, targets from 1.245 to 1.265 g and a guard band of 0 or 3 mg, the best passes one of six seeds. On this line the weighing interval, not the correction, is the decision that releases or rejects the tail.

9.6.7 Runtime

A whole challenge batch of 2353 s of line time runs headless in 49 ms, and the suggested batch of 2612 s in 56 ms. In the browser the longest case, 4800 vials in 4187 s, plays in 34.9 s at the top speed of 120 times real time, and every batch takes 12 to 35 s at that speed.

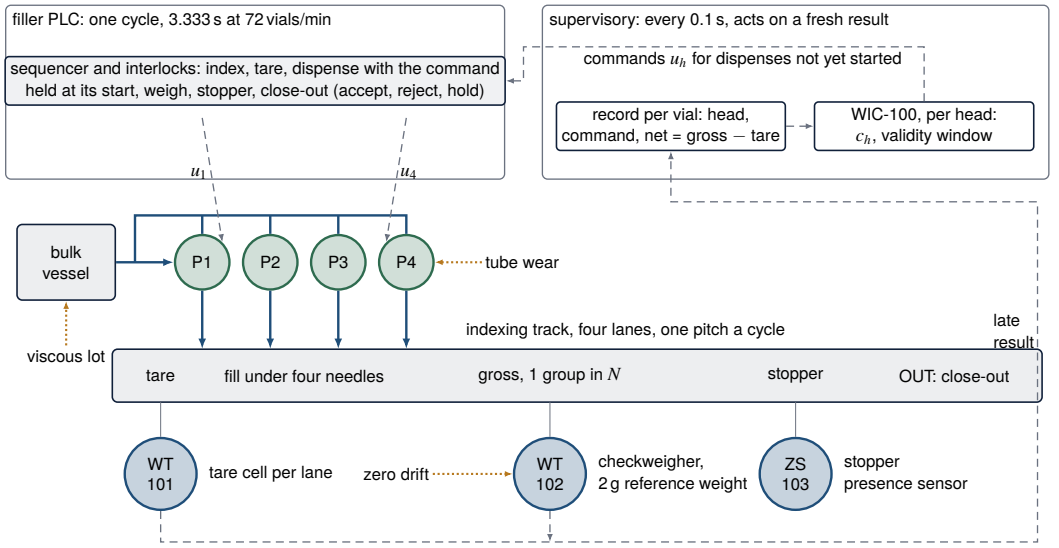


Figure 9.5 The chapter's loop on plant hardware. The filler's sequencer indexes the track, tares each vial, starts each dispense with the command its head holds, weighs one group in N, checks the stopper and closes each vial out. WIC-100 reads only the reported records, each net of its vial's own tare, and moves the commands of dispenses not yet started. Dotted arrows mark disturbances. Tags other than WIC-100 are the book's.

9.7 Technology in Practice

What Measures It

The line of this chapter weighs every vial empty on its lane's tare cell and weighs a sampled group gross on the checkweigher two pitches later, and a presence sensor looks for each stopper. The net of a vial's own gross and tare is the only fill result the controller receives, and it never reads a true fill. Figure 9.5 draws the loop on that hardware. The model's scales add 0.5 mg of noise, round to 1 mg and report each result a latency after the weighing dwell of 3.0 s.

> **TECHNOLOGY** **Sensor: Tare Cells and a Gross Checkweigher**
>
> A weighing cell turns force into an electrical signal. In a load cell, one or more strain gauges are bonded to a metal structure of precisely known elastic properties, and the strain changes their resistance. Vibration from machinery around a weighed mass puts oscillating forces on the cell [18]. The balance behind the dose records resolved 0.002 g, repeated within ± 0.0005 g and settled in under 120 ms [6], and it cancelled vibration in real time [4]. The model's scale noise of 0.5 mg is that repeatability bound used as a standard deviation. Its zero b_G may drift, and a correction that trusts the scale then turns the bias into an underfill while the reported fills stay on target, which only the reference weight exposes.

What Moves It

Each lane has its own head, tube and needle, so each head needs its own command. In the dose study the pump was calibrated by the pump controller's own function, which computes pump parameters from weight measurements to minimize the dosing error [4]. In the model that calibration leaves each head with the offset $G_h(0)$, and the dispense multiplies the command by the head's gain and adds a replayed deviation, so the rollers, the occlusion and the tube geometry are absent (Section 9.3).

TECHNOLOGY Final Element: Four Peristaltic Heads

A peristaltic pump compresses sterile, disposable tubing with rotating rollers, so the liquid touches only the inside of the tube. The same tubing brings variability through tube wear and temperature [6]. In the dose study, new tubing and a newly assembled circuit introduced calibration drift, so the pump was calibrated five times at the start of each experiment, after a 30 s purge that removed air bubbles [4]. At the needle, a fast suck-back drove product down into a plug at the tip that dried quickly, and a slow one prevented clogging [3]. In the model, wear w_h lowers each head's gain G_h, and purge shots prime the heads after a stop.

Where the Controller Runs

The dose study ran its rig from one controller that served both as a programmable logic controller (PLC) and as a PC-based platform [4], and its compensation worked through the pump controller's standard interface rather than through the mechanics of one pump [6]. A PLC can be the primary controller of a discrete process, and control loops in industry run with cycle times from milliseconds to minutes [19]. WIC-100 executes every 0.1 s but changes a command only at a fresh result, which in case 3 arrives 50.2 s after its fill. Read as plant hardware, the sequence of index, tare, dispense, weighing and close-out, with its interlocks, belongs to the filler's PLC, and the per-head correction from the checkweigher is a supervisory function above the pump controllers.

IN PRACTICE Interventions Go in the Batch Record

The European rules for sterile products ask for an authorized list of qualified interventions, both inherent ones such as aseptic assembly and corrective ones such as clearing a component jam. Each line stoppage or intervention goes into the batch record with its time, its duration and the operators involved, and an intervention that was not qualified is weighed in batch disposition. The filling zone and the stopper bowl lie in grade A, the critical zone [20]. In case 8 of this chapter a glove intervention puts 16 vials on aseptic hold, whatever they weigh (Table 9.5).

What Is Modernizing It

Three developments touch this line. The first takes people out of the critical zone. The European rules ask manufacturers to consider barrier systems, isolators, robotic systems, rapid methods and continuous monitoring [20], and the 2004 guidance notes that robotics can further reduce the risk to the product [1]. The second is the pump and its needle. A mechanical design with eight rollers in place of the industry-standard six makes smaller, more frequent compression points and less pulsation [6], and a nozzle tip with a hydrophobic face and a pressure-resistant barrier gave consistent suck-back heights and about 90 percent of fills within 3 percent of target [2]. The third is software. Compensation learned from the pump's own dose records matched or exceeded that mechanical design, and an ensemble of autoregressive models ran in under 1 s per cycle against about 11 s for a recurrent network [6]. The cycle of this chapter's line lasts 3.333 s at 72 vials/min.

9.8 Limitations

The dose deviations are those of purified water on one pump with one container, so the data cannot validate protein viscosity, temperature, sterility, a multihead line, stoppages or scale bias, and every mechanical law beyond the data is a teaching assumption. The density, the limits, the throughput and the giveaway target are teaching values. Actual limits are set by the product's specification, the pharmacopeial requirements for the content of the container and the site's control strategy. The controller comparison rests on the declared action-response surrogate and on the three independence assumptions of Section 9.3, which the records cannot test, and it is a simulation result, not a prospective experiment. The seeds share replay segments, so their ranges understate the variation of a new noise record, and the accepted-underfill counts in particular are default-seed values with the spread reported in Section 9.6. The dispositions are teaching rules, the aseptic holds are discrete states and nothing contaminates a vial, so a simulated accepted vial is not a released vial and a batch with

no rejects is not evidence of sterility. The protection the guard band gives depends on the declared scale noise, a repeatability bound used as a standard deviation. The validity window, the guard band and the restart check are additions to the plan's correction law, each kept because a measurement on this line justified it, and a line with different noise would need them set again.

OPEN CHALLENGE **Filling Real Formulations, Not Water**

Compensation learned from dose records has been validated on purified water only. Whether it adapts to formulations of other viscosities and rheology, and whether it carries over to other pump architectures, control systems and environments, are open questions that the study itself calls critical for future research [6]. The data are scarce, because the technical details and critical parameters of peristaltic filling are rarely published, and those of viscous, high-concentration formulations least of all [2]. Progress would be open dose records of real formulations on several pumps, against which a record-based correction could be tested beyond its water surrogate.

9.9 Conclusion

The simulated filler fills four lanes from one bulk vessel, replays recorded water doses stroke by stroke per head, and keeps one record per vial that ties its weights to its head and its command, so that a correction sees only what a checkweigher reports and only when it reports it. A pooled correction keeps the average right while one wearing head fills 10.4 mg under the target. A correction that cancels each whole error widens the heads' spread from 5.0–7.1 to 6.1–9.2 mg, as the records' own prediction errors of 4.95, 6.12 and 7.98 mg anticipate. With about six results per head in flight the plain per-head law oscillates to a 48 mg spread, while the record-based estimate holds 4.6 to 6.8 mg there and a mean head spread of 5.8 to 6.0 mg at every latency tested. A plain law detuned to 0.10 mL/g matches it on the challenge line, 2324 against 2325 accepted, and the weighing interval, not the correction, decides whether underfills are released. The comparison is a simulation under declared surrogate and independence assumptions. A fuller model would add a pump map with a measured command response, a viscous formulation's own dose records, weighing after closure with stopper-mass compensation, and a prospective test of the corrections on a filler. The simulator shows that a correction has to know which vial a result describes before it can act on it.

KEY POINTS

- Each weight describes a vial filled tens of seconds earlier, by a pump whose command later corrections have already changed.
- A pooled correction keeps the average right while one wearing head fills 10.4 mg under the target.
- With about six results per head in flight the plain per-head law oscillates to a 48 mg spread, while the record-based estimate holds 4.6 to 6.8 mg there.
- A plain law detuned to 0.10 mL/g matches it on the challenge line, 2324 against 2325 accepted.
- The weighing interval, not the correction, decides whether underfills are released.

Where else this applies. Volume 5, Chapter 6 holds tablet weight on a rotary press, and Chapter 7 applies the same integral form to pack pressure from each weighed part (Section 1.4.5).

References

[1] U.S. Food and Drug Administration, "Guidance for industry: Sterile drug products produced by aseptic processing — current good manufacturing practice," Docket FDA-2003-D-0145, Oct. 2004. [Online]. Available: https://www.fda.gov/regulatory-information/search-fda-guidance-documents/sterile-drug-products-produced-aseptic-processing-current-good-manufacturing-practice

[2] W. Shieu, D. Lamar, O. B. Stauch, and Y.-F. Maa, "Filling of high-concentration monoclonal antibody formulations: Investigating underlying mechanisms that affect precision of low-volume fill by peristaltic pump," *PDA Journal of Pharmaceutical Science and Technology*, vol. 70, no. 2, pp. 143–156, 2016.

[3] S. Hanslip, K. G. Desai, M. Palmer, S. Bell, P. Schofield, P. Varma, F. Roche, J. D. Colandene, T. M. Smith, and D. P. Nesta, "Syringe filling of high-concentration mAb formulation: Slow suck-back pump speed prevented filling needle clogging," *Journal of Pharmaceutical Sciences*, vol. 106, no. 12, pp. 3651–3653, 2017.

[4] D. Privitera, A. Mecocci, and S. Bartolini, "Extensive dataset for peristaltic pump accuracy enhancement in pharmaceutical environments," *Scientific Data*, vol. 12, no. 1, p. 1618, 2025.

[5] ——, "DosePI: A comprehensive dataset for peristaltic pump accuracy enhancement in pharmaceutical environments," Dataset, figshare, licensed CC BY 4.0, 2025.

[6] ——, "Ensemble methods for peristaltic pump accuracy enhancement," *IEEE Access*, vol. 13, pp. 125 157–125 179, 2025.

[7] L. de Mare, S. Velut, P. Hagander, and M. Akesson, "Feedback control of flow rate from a peristaltic pump using balance measurements," in *2001 European Control Conference (ECC)*, 2001, pp. 2164–2168.

[8] M. Meller, M. Niedźwiecki, and P. Pietrzak, "Adaptive filtering approach to dynamic weighing: A checkweigher case study," *IFAC Proceedings Volumes*, vol. 47, no. 3, pp. 5927–5932, 2014.

[9] J. F. MacGregor, "A different view of the funnel experiment," *Journal of Quality Technology*, vol. 22, no. 4, pp. 255–259, 1990.

[10] G. Box and T. Kramer, "Statistical process monitoring and feedback adjustment—a discussion," *Technometrics*, vol. 34, no. 3, pp. 251–267, 1992.

[11] R. P. Good and S. J. Qin, "On the stability of MIMO EWMA run-to-run controllers with metrology delay," *IEEE Transactions on Semiconductor Manufacturing*, vol. 19, no. 1, pp. 78–86, 2006.

[12] A. Ingolfsson and E. Sachs, "Stability and sensitivity of an EWMA controller," *Journal of Quality Technology*, vol. 25, no. 4, pp. 271–287, 1993.

[13] E. Sachs, A. Hu, and A. Ingolfsson, "Run by run process control: Combining SPC and feedback control," *IEEE Transactions on Semiconductor Manufacturing*, vol. 8, no. 1, pp. 26–43, 1995.

[14] S. W. Roberts, "Control chart tests based on geometric moving averages," *Technometrics*, vol. 1, no. 3, pp. 239–250, 1959.

[15] G. E. P. Box and G. M. Jenkins, "Some statistical aspects of adaptive optimization and control," *Journal of the Royal Statistical Society Series B: Statistical Methodology*, vol. 24, no. 2, pp. 297–331, 1962.

[16] E. Del Castillo and A. M. Hurwitz, "Run-to-run process control: Literature review and extensions," *Journal of Quality Technology*, vol. 29, no. 2, pp. 184–196, 1997.

[17] W. G. Hunter and C. P. Kartha, "Determining the most profitable target value for a production process," *Journal of Quality Technology*, vol. 9, no. 4, pp. 176–181, 1977.

[18] T. R. Kuphaldt, *Lessons in Industrial Instrumentation*, version 2.33 ed. Open textbook, Creative Commons Attribution 4.0 license, ibiblio.org/kuphaldt/socratic/sinst, 2024.

[19] K. Stouffer, M. Pease, C. Tang, T. Zimmerman, V. Pillitteri, S. Lightman, A. Hahn, S. Saravia, A. Sherule, and M. Thompson, "Guide to operational technology (OT) security," National Institute of Standards and Technology, Gaithersburg, MD, Special Publication 800-82r3, 2023.

[20] European Commission, "EudraLex, volume 4, EU guidelines for good manufacturing practice for medicinal products for human and veterinary use, annex 1: Manufacture of sterile medicinal products," C(2022) 5938 final, Brussels, Aug. 2022. [Online]. Available: https://health.ec.europa.eu/system/files/2022-08/20220825_gmp-an1_en_0.pdf

9.10 Exercises

9.1 ⌕ Concept check · 2 min

In Scenario 1 the pooled correction moves all four heads by the filtered error of each weighed group's average, while head 3's tube wears. What do the true fills of the four heads show over each head's last 100 vials?

(a) All four heads fill within about a milligram of the target, since the group average is held there

(b) Head 3 fills under the target, and the other three heads are pushed over it

(c) Head 3 is brought back to the target, and the other three heads stay where they were calibrated

(d) All four heads fall under the target, because only one group in four is weighed

9.2 Concept check · 2 min

In Scenario 3 each result arrives 40 s after its vial is weighed, and about six results per head are always in flight. The plain per-head law oscillates, and the heads' true fills spread 47 to 48 mg. What is the cause?

(a) The late results carry more scale noise, so the correction chases noise it would otherwise filter out

(b) The validity window rejects most late results, so the heads run without correction between them

(c) Each result's error is added to a command that earlier results have already corrected for the same error

(d) The step limit of one percent of the nominal command is too small for the correction to keep up

9.3 Concept check · 1 min

In Scenario 7 the line stops twice for 4 min and restarts with the commands it had. The first fill of each head after a restart is 10.2 to 23.7 mg under the target. What restores those first fills?

(a) The per-head correction, as soon as the results of the restart group are reported

(b) A higher correction gain, so that the first results move the commands further

(c) A reject guard band, which keeps the short vials inside the fill limits

(d) Purge shots before the first vial, which replace the fallen drop and work the tube

9.4 Theory and calculation · 6 min

The line fills four vials per cycle. Each cycle at rate R (vials per minute) lasts $T_c = 60 \cdot 4/R$, of which $t_{idx} = 0.9\,\mathrm{s}$ is the index move and $t_{set} = 0.3\,\mathrm{s}$ a settle, and the dose statistics are referred to $R_{ref} = 60\,\mathrm{vials/min}$ (Equations (9.1) and (9.3), Table 9.2).

(a) At 72 vials/min, find T_c, the fill window T_f, the pump-speed ratio ω/ω_{ref} and the spread factor f_σ on water. *Checkpoint:* 3.333 s, 2.133 s, 1.3125 and 1.146.

(b) In a weighing cycle the dwell after the index lasts at least $t_w = 3.0\,\mathrm{s}$. Find the line rate before stops when every group is weighed. *Checkpoint:* 61.5 vials/min.

(c) A lot at 6 mPa s replaces water at 1.0 mPa s. Find the change in every head's delivery and the factor on the spread of its doses. *Checkpoint:* −2.15 % and 1.565.

9.5 Theory and calculation · 6 min

Equation (9.4) models one head's dose deviation as $e_k = p_k + r_{k-1} - r_k$, with independent pump terms p_k of standard deviation $\sigma_p = 2.85\,\mu\mathrm{L}$ and drops of standard deviation $\sigma_r = 3.17\,\mu\mathrm{L}$.

(a) Derive the variance of e_k and its lag-1 autocorrelation, and evaluate the standard deviation and ρ_1. *Checkpoint:* 5.31 µL and −0.356.

(b) Convert the standard deviation to milligrams of water at $998.2\,\mathrm{kg/m^3}$ and compare it with Table 9.3. *Checkpoint:* 5.30 mg against 5.31 mg on training.

(c) A predictor that always predicts the target meets both the spread within a segment and the spread of the segment means. Treat the two as independent and estimate its RMSE on the training third. *Checkpoint:* 6.86 mg against 6.82 mg in the table.

9.6 Theory and calculation · 8 min

With the report filter at 1.00 the plain per-head law adds $K_h(m_{SP} - \hat{m}_r)$ to the head's command after each result r. A head delivers $\xi\rho$ grams per milliliter of command, with $\rho = 1.030\,\mathrm{g/mL}$, and each result reaches the controller after D later results of the same head have already moved the command.

(a) Show that the loop gain per result, $\lambda\xi$ in Equation (1.9), is $K_h\rho\xi$, and evaluate $K_h\rho$ for K_h of 0.97, 0.50, 0.25 and 0.10 mL/g. *Checkpoint:* 0.999, 0.515, 0.258 and 0.103.

(b) Put $z = e^{i\omega}$ into Equation (1.9) and show that the loop is stable for a loop gain below $g^* = 2\sin\big(\pi/(4D+2)\big)$. Check $D = 0$ to 3 against Section 1.4.5. *Checkpoint:* 2, 1, 0.618 and 0.445.

(c) With $D = 6$ and $\xi = 1$, which of the four gains are stable? *Checkpoint:* $g^* = 0.241$: only 0.10 mL/g.

9.7 App challenge · 8 min

Open Scenario 1, "1 · One needle drifting, a pooled correction", at the default seed, and run it to the end at 120× (120 s of line time per wall second, about 12 s) on the Heads view.

(a) Read each head's true mean over its last 100 vials on the "last 100" line of its panel, and the tile Accepted outside the limits (true). Tolerance 0.1 mg. *Checkpoint:* Head 3 −10.4 mg, heads 1, 2 and 4 +6.7, +3.4 and +0.9 mg, 32 under.

(b) Press Reset, set Correction on the WIC-100 card to Per-head correction, and run to the end. Read head 3's true mean, the underfills accepted and the fill-bracket holds on the tile Rejected · held. *Checkpoint:* −0.7 mg, 3 under, 18 held against 59.

(c) Read the four commands on the head panels at the end of each run and explain the difference. *Checkpoint:* Pooled: 1.2186 mL on every head. Per head: 1.2084 to 1.2333 mL.

9.8 App challenge · 8 min

Open Scenario 3, "3 · A late checkweigher result", at the default seed, on the Heads view at 120× (about 19 s per batch).

(a) Run it to the end as shipped. Read the SD on each head's panel and the lead tile Accepted vials · giveaway. Tolerance 0.5 mg. *Checkpoint:* SD 46.7 to 48.1 mg, 648 of 2400 accepted.

(b) Press Reset, set Correction on the WIC-100 card to Per-head, record-based, and run to the end. Read the same values. *Checkpoint:* SD 4.6 to 6.8 mg, 2394 accepted.

(c) Press the reload button beside the scenario list, set the correction gain on all four Head cards to 0.10 mL/g, and run to the end. Compare with part (b), and say what the lower gain costs. *Checkpoint:* 2394 accepted, SD 4.7 to 6.8 mg.

9.9 App challenge · 9 min

Open the challenge, "Challenge · Least giveaway, nothing out of limits", at the default seed, and run at 120× (about 20 s per batch).

(a) Run it to the end as shipped. Read the lead tile, Accepted outside the limits (true) and Throughput. *Checkpoint:* 1598 of 2400 at 20.9 mg, 92 under and 18 over, 62.0 vials/min.

(b) Press Reveal suggested tuning, then Reset, and run to the end. Read the same tiles. Tolerance 0.1 mg on the giveaway. *Checkpoint:* 2329 of 2400 at 13.3 mg, none outside, 55.8 vials/min, target met.

(c) Open Advanced on the WIC-100 card, set Reject guard band to 0 mg, press Reset and run again. Read the same tiles and explain the result. *Checkpoint:* 2393 accepted at 12.9 mg, 6 under, target not met.

9.10 Application area · 4 min

The label states 1.20 mL of a product at 1.030 g/mL, and the upper fill limit is 104 % of the label minimum.

(a) Find the label minimum, the upper limit, and the margin from the 1.250 g target to the minimum in standard deviations of a replayed dose (5.7 mg). *Checkpoint:* 1.236 g, 1.2854 g, 14 mg or 2.5 standard deviations.

(b) The shipped challenge gives away 20.88 mg per accepted vial and the suggested settings 13.25 mg. Express the difference for a batch of 2400 vials, and name what the line pays for it. *Checkpoint:* 7.63 mg per vial, about 18.3 g per batch, for 55.8 instead of 62.0 vials/min and about 10 g of purge.

9.11 Application area · 3 min

A quality engineer proposes dropping the reference-weight check, because the per-head correction holds the reported mean of the last 100 vials within 1 mg of the target. In the chapter's run of Scenario 1, with head 3's wear removed, the gross scale's zero drifts by 30 mg/h over the 1407 s batch.

(a) Find the zero offset at the end of the batch and the command change a correction that trusts the scale would make. *Checkpoint:* 11.7 mg and 11.4 µL lower.

(b) Using the chapter's result, say what the proposal would hide. *Checkpoint:* Reported +0.6 mg against a true −9.7 mg; accepted underfills from 2 to 158.

9.12 Review the agent · 5 min

An AI assistant reviews the challenge line and writes: "Weighing every group costs 6 vials per minute. Keep one group in four, the record-based estimate, the 3 mg guard band and the two purge shots, and raise the target from 1.249 to 1.256 g. The giveaway rises to about 20 mg, which the plant accepts for the throughput. A mean 7 mg higher is more than twice the guard band, so the tail clears the label minimum and no underfilled vial is released." Find the flaw, and say how you would check it in the app.