INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 19 Control of Biological Systems and Human Performance

CHAPTER 9

Refolding Insulin Precursor When a Better Fold Can Cost the Campaign

Pages 234–255

Edited by John D. Hedengren

How to cite this chapter. "Refolding Insulin Precursor When a Better Fold Can Cost the Campaign," 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. 9, pp. 234–255.

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

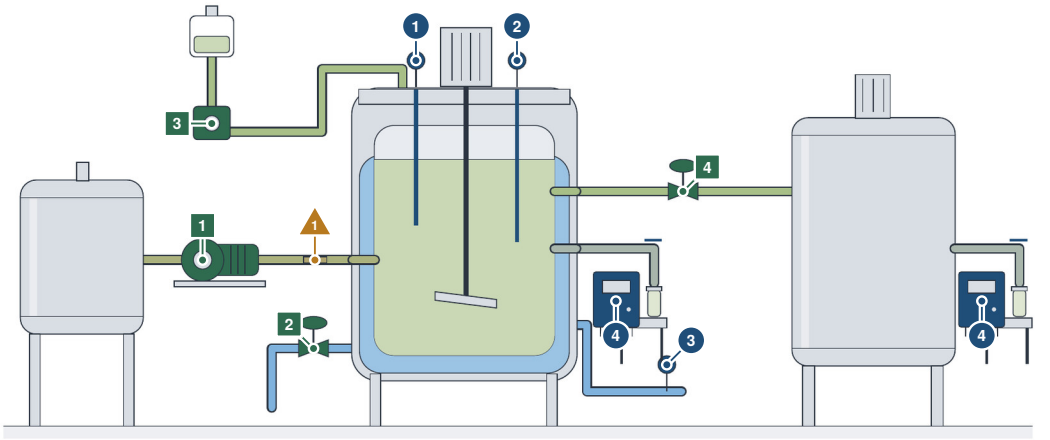
The app for this chapter
https://apmonitor.com/apps/insulin-manufacturing/

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

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

CHAPTER 9

Refolding Insulin Precursor When a Better Fold Can Cost the Campaign



The controller moves 1 the charge pump, which meters the dissolved protein in 2 the valve on the jacket utility 3 the pump that doses base to hold the pH 4 the valve that moves the folded lot on

The controller reads 1 the pH probe in the liquid 2 the temperature probe in the liquid 3 the temperature element on the jacket return 4 the two sample points, each read at a bench beside its vessel

What disturbs it 1 dissolved protein arriving weaker than the recipe assumes

Notice that the two benches, which alone say how much protein has folded, stand off the line, so between their reports every folded, aggregated and unfolded mass the controller uses is an estimate. Illustration, not to scale.

IN PLAIN TERMS **Folding a Protein Without Letting It Clump**

The main idea. Bacteria can be made to build the protein that turns into insulin. They pack it away as a useless lump, so the plant melts the lump and lets the protein take its proper shape in a big stirred tank.

The challenge. Each molecule either folds the right way or sticks to a neighbor and is lost. Sticking wins when the liquid is crowded, so the protein is added slowly, and then the tank is busy so long that the next batch spoils while it waits.

Why it matters. A plant is paid for the medicine it finishes in a week, not for the share of one tank that came out right.

THREE WAYS TO READ THIS CHAPTER

⊘	**Explore**	13 min	the first page, the history of the route (Section 9.2), the two fact segments, both try-it segments and the key points
🕮	**Study**	35 min	the whole chapter with its three pause points and the twelve exercises of Section 9.10
⛰	**Practice**	10 min	the first page, the ten measured cases (Section 9.6), Technology in Practice (Section 9.7), Limitations (Section 9.8) and the open challenge

The control question Unfolded insulin precursor either folds, at a rate first order in its concentration, or aggregates, at a rate second order in it, so a lot charged slowly keeps more protein than a lot charged at once. The refolding vessel is also the bottleneck of the recovery train, so the question is when a charge that folds more protein pays for the vessel hours it costs.

Learning objectives

- Compute the folded share of a bolus charge from the folding and aggregation constants, and predict how temperature and pH move it.
- Explain why the hours a lot occupies the refolding vessel, and not the yield of that lot, decide what a 72-hour campaign delivers.
- Correct a model of a vessel from an assay reported late, and stop a reaction whose product is consumed by the enzymes that made it.
- Choose a wash, a charge strategy, a refolding temperature and the two supervisory rules for a stated campaign, and defend the choice by the accepted insulin per campaign hour.
- State the conditions under which a larger harvest delivers less product, and the measurement that would settle the matter in a plant.

Prerequisites Sections 1.4.3 (competing first- and second-order kinetics and the fold they produce), 1.4.5 (a model corrected at the hour of the draw), 1.6.2 (velocity form) and 1.3 (loop action). Chapter 8 is the nearest neighbor: the same correction from a laboratory sample, on a plant faster than its samples.

The app https://apmonitor.com/apps/insulin-manufacturing/

See also Volume 8, Chapter 11 for the capture step this chapter declares as a recovery map, Volume 8, Chapter 2 for a feed metered to keep an unreacted store small, and Volume 8, Chapter 9 for a lot judged by a sampled check. Volume 5, Chapter 4 for steering on a grade that arrives hours late, and Volume 5, Chapter 6 for an upstream decision that reaches the product minutes later.

9.1 Introduction

An insulin fermentation report quotes grams of precursor per liter of broth, while the plant is paid for grams of insulin that pass their lot limits. In the simulated campaign of this chapter the cells make 3013 g of precursor-equivalent protein and a recipe modeled on a published pilot process isolates 217 g of acceptable insulin, one part in fourteen. Between the two stand a wash, a refolding vessel, an enzyme reaction that removes 53 % of every converted molecule's mass as fragments, and two purification steps. The largest loss the operator can move is in the refolding vessel. Unfolded precursor either folds, at a rate first order in its own concentration, or aggregates, at a rate second order in it [1], so a lot charged all at once loses more protein to aggregate than the same lot fed in slowly.

Two beliefs are what the simulator is designed to correct. The first is that more precursor at the harvest means more insulin at the end. A harvest at 22 h carries 3238 g, 11 % more than the 2905 g of a harvest at 16 h, and the campaign isolates 114.6 g instead of 217.3 g because the recovery train cannot take the lots before their clocks run out. The second is that the fold that keeps the most protein is the best fold. Charged as a bolus to 2.0 g/L, the first lot of 815 g folds 67.7 % and frees the vessel after 18.8 h. Fed by the scheduler under a ceiling of 0.35 g/L on the unfolded concentration it folds 84.7 %, but it holds the vessel for 39.0 h, 665 g of the harvest behind it expires at 52 h, and by the 72 h horizon the campaign has isolated nothing, against 112.7 g after the bolus.

The arrangement is general. Two reactions of different order that compete for one reactant favor the lower-order product when the reactant is kept dilute, which is why preparative refolding dilutes the denatured protein or adds it in pulses [2, 3], and aggregation of non-native protein is a multi-stage process whose kinetics set processing conditions across biotechnology products [4]. A fed-batch polymerization that meters its monomer, a crystallizer that doses its antisolvent and a fermenter fed below the rate that makes byproduct have the same shape. Volume 8, Chapter 2 meters a reagent into a batch for a different reason, to keep the heat release no faster than the feed, and the unreacted store

it builds plays the part the unfolded protein plays here. The scheduling half of the problem is also common. Production scheduling models and solution methods have found industrial use across the process industries [5], and biopharmaceutical campaigns are planned against fermentation titers that are themselves uncertain [6].

The published studies treat the two halves apart. Refolding studies optimize one vessel's yield over a fixed incubation [7], and campaign planning works at the level of batches and titers [6]. This chapter puts them on one screen. Its simulator carries the folding competition, the vessel occupancy and the hold clocks through a single 72-hour campaign, and the question it answers is when a concentration-aware charge pays for the vessel hours it costs, and how far that answer survives errors in the folding constants and the assay delay. The reduced folding law, the velocity-form proportional-integral (PI) loops and the estimator are established methods, and every comparison below is a set of measurements on one declared model: a 34-state model whose protein ledger closes to 1.8×10^{-13}, behind a measurement boundary that keeps every controller and rule away from the simulator's truth.

Section 9.2 follows insulin manufacture from extraction to refolding. Section 9.3 states the model, vessel by vessel, with its parameters and the evidence behind each one, and Section 9.4 checks the code against the identities it should satisfy and the folding law against data withheld from its fit. Section 9.5 describes the four loops, the feed scheduler, the estimator and the transfer and endpoint rules, and Section 9.6 measures ten cases, from the protein ledger to an infeasible harvest and a graded campaign. Section 9.7 describes the instruments and the platform a plant would use, Section 9.8 states what the model cannot be asked, and the exercises of Section 9.10 run the cases that carry the argument.

9.2 From Pancreatic Extract to a Folded Precursor

Insulin reached patients first as an extract of animal pancreas, and that route supplied the drug until recombinant human insulin arrived in the early 1980s [8].

> ⧖ HISTORY **From Pancreatic Extract to a Recombinant Route (1921 to 1983)**
>
> Insulin entered medicine as an extract. Pancreatic extracts that lowered blood sugar in animals had been made for two decades before 1921, but impurities and toxic reactions kept them from patients, and the Toronto work of 1921 and 1922 produced an extract from whole beef pancreas that treated people with diabetes once its purification had been worked out [9]. Extraction from animal pancreas then supplied insulin until recombinant human insulin arrived in the early 1980s [8]. Human insulin from recombinant deoxyribonucleic acid technology became the first commercial health-care product of that technology, and its development began before federal guidelines for large-scale recombinant work existed [8].

The molecule's own biosynthesis pointed to the route used now. In 1967 islet tissue was shown to make a larger single-chain protein that radioactivity passed into insulin, and the name proinsulin was proposed for it [10]. In 1971 proinsulin was converted to insulin in vitro with trypsin and carboxypeptidase B [11], the two enzymes the pilot process simulated here still uses [7]. Recombinant expression followed within a decade. Synthetic genes for the A and B chains were expressed in *E. coli* as fusions to β-galactosidase, the chains were cleaved, purified, mixed, reduced and reoxidized, and insulin was detected [12]. Production later moved from two separately made chains to a proinsulin route on an *E. coli* fermentation [13]. Insulin products commanded annual global sales above $4.5 billion by 2005 [14], and expression hosts and capacity remain a subject of review because demand is expected to outgrow current manufacturing [15].

The cost of the bacterial route is that the protein arrives as an inclusion body. Inclusion bodies are dense deposits of inactive protein that can be isolated by solid-liquid separation, solubilized and refolded in vitro, and folding protocols control the protein concentration, disulfide formation, temperature, pH and solvent to suppress unproductive side reactions [16]. Refolding of inclusion-body protein is cumbersome, often recovers poorly and accounts for the major cost of producing a

Table 9.1 Variables by Role, With Defaults and Ranges

Symbol	Role	Default	Range
$\hat{U}$	CV, estimated unfolded (FIC-304)	$U_{tgt} = 0.35$	0.05–1.5 g/L
T	CV, vessel temperature (TIC-301)	15 °C	8–30 °C
T_j	CV, jacket (TIC-302)	from TIC-301	2–40 °C
pH	CV, vessel pH (AIC-303)	10.6	9–11.5
F	MV, charge pump	—	0–400 L/h
$T_{j,SP}$, T_{sup}	MV, jacket setpoint, utility	—	2–40, −5–45 °C
u_b	MV, base pump (0.25 mol/h at 100 %)	—	0–100 %
t_{ind}, t_{harv}	decision, induction, harvest	6, 16 h	2–14, 6–30 h
V_{lot}, w	decision, lot, wash volumes	250 L, 3	100–350, 0–6
C_{tgt}	decision, charge concentration	2.0 g/L	0.5–4 g/L
c_E, T_{conv}	decision, enzyme dose, temperature	1, 15 °C	0.25–3, 8–30
rules	decision, charge, transfer, endpoint	bolus, time, time	three each
x_a	DV, acid load multiple	1	0–6
κ	DV, chiller capacity	100 %	20–100
s	DV, solubilization strength	100 %	40–120
δ_{pH}, δ_X	DV, pH and biomass probe drift	0, 0	±0.1/h, −3–6 %/h
V_f, X, S, P	states, fermenter	at 16 h: 744 L, 27.1, —, 3.90	kg/m^3
m_H, I_H, V_H	states, harvest tank	2905 g	—
P_A, I_A, V_A, w	states, recovery vessel	—	—
U_S, I_S, V_S	states, solubilized tank	—	—
VU, VN, VA, V	states, refolding (lot 1)	815 g, 408 L	—
I_R, b, T, T_j	states, refolding impurity, base	—	—
n_N, n_I, n_G, n_D	states, conversion (moles)	—	—
I_C, V_C	states, conversion impurity, volume	—	—

recombinant protein in *E. coli* [3]. The quantitative basis for controlling it came in 1991, when the yield of native protein upon refolding was explained by the rates of competing first-order folding and second-order aggregation, so that aggregation dominates at high protein concentration [1]. Dilution, pulsed addition and column-based methods followed as preparative practice [2, 3]. The pilot process that calibrates this simulator refolds a fused precursor directly, without sulfonation, at 1.5 g/L and pH 10.6 for 16 h at 15 °C, and from 200 L of broth it isolates 138.2 g of insulin on average over nine batches [7].

9.3 The Route as Modeled

The simulated route is a proposed teaching route assembled from public evidence, not a reproduction of any manufacturer's process, as shown in Figure 9.1. A fed-batch fermenter FERM-101 makes precursor in inclusion bodies. At the harvest the broth goes to a hold tank with a stability clock. Recovery vessel REC-201 takes one lot at a time through lysis, an inclusion-body wash and solubilization, and passes it to a solubilized hold tank, which frees REC-201 for the next lot. The refolding vessel is charged from that tank through the charge pump of FIC-304, its temperature is held by the cascade TIC-301 to TIC-302 and its pH by AIC-303 on a base pump. When folding is judged complete the lot is acid-precipitated into the conversion vessel, where enzymes cut the precursor to insulin and then slowly over-process it. A declared capture and polish map and a gate on purity and host-cell protein turn each lot into accepted or rejected bulk. Three vessels hold one lot each, so the schedule is a decision rather than a queue.

The variables of Table 9.1 are classified by role. Thirty-four states are integrated, 28 physical and 6 ledgers that accumulate production and losses so that conservation can be checked exactly. The physics is in SI units (seconds, cubic meters, kilograms, moles), and the interface displays hours, liters, grams and g/L, which equals kg/m^3.

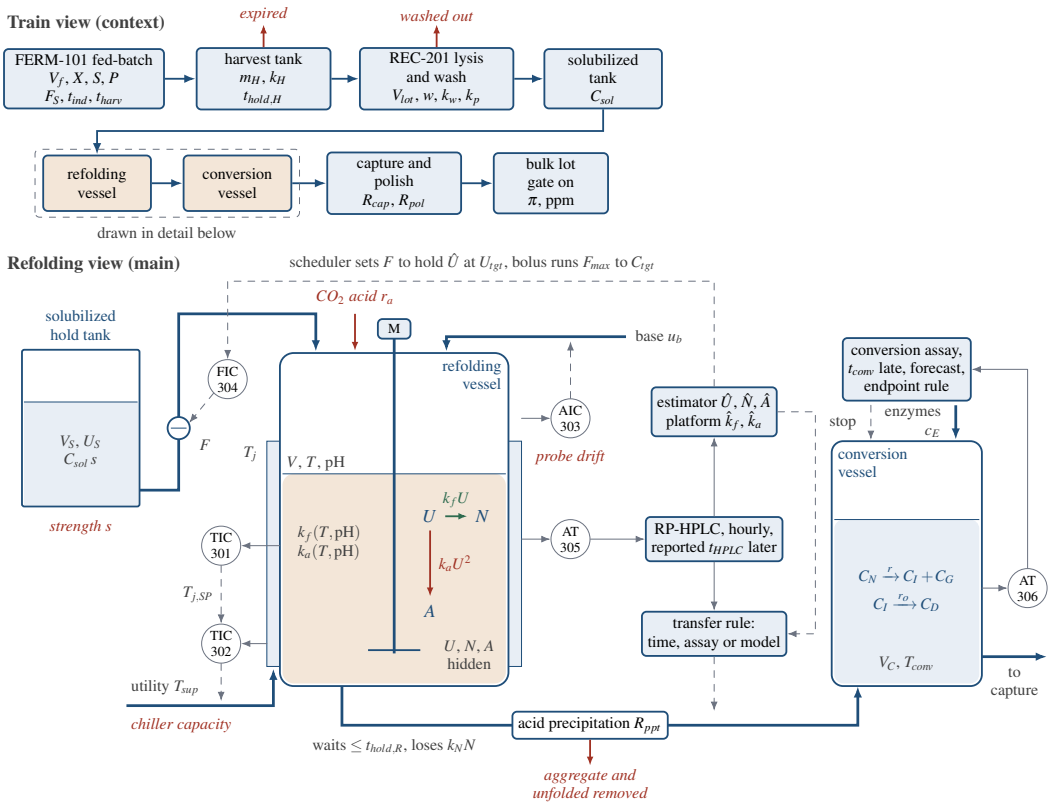


Figure 9.1 Insulin recovery and refolding as modeled. The upper strip is the Train view, from fermenter to bulk lot, with the two vessels drawn in detail below. The lower panel follows the Refolding view: the solubilized hold tank at V_S and concentration $C_{sol}s$ feeds the charge F, the refolding vessel of volume V holds the hidden states U, N and A with folding $k_f U$ and aggregation $k_a U^2$, TIC-301 and TIC-302 cascade from T to the jacket T_j and the utility T_{sup}, and AIC-303 doses base u_b against the acid load r_a. The RP-HPLC assay at AT-305 reports t_{HPLC} late to the estimator of $\hat{U}$, $\hat{N}$ and $\hat{A}$, which sets the scheduler's charge and the transfer rule. The conversion vessel holds C_N, C_I, the fragments C_G and the over-processed C_D, with its own delayed assay and endpoint rule. Disturbances are in red italics. Process model only, not to scale.

9.3.1 The Fermenter That Supplies the Load

The fermenter runs on one substrate with a declared oxygen and cooling capacity, and induction switches on a production rate that decays while most of the growth stops, as shown in Equation (9.1).

$$
\begin{aligned}
\frac{\mathrm{d}V_f}{\mathrm{d}t} &= F_S, \qquad \frac{\mathrm{d}X}{\mathrm{d}t} = \left(\mu - k_d - \frac{F_S}{V_f}\right)X \\
\frac{\mathrm{d}S}{\mathrm{d}t} &= \frac{F_S}{V_f}(S_f - S) - \left(\frac{\mu}{Y_{X/S}} + m_S + \frac{q_P}{Y_{P/S}}\right)X \\
\frac{\mathrm{d}P}{\mathrm{d}t} &= q_P X - \left(k_{P,loss} + \frac{F_S}{V_f}\right)P \\
\mu &= \mu_{max}\frac{S}{K_S + S} f_O f_b^{\iota}, \qquad f_O = \min\left(1, \frac{X_{cap}}{X}\right) \\
q_P &= \iota\, q_{P,max}\, \mathrm{e}^{-(t - t_{ind})/\tau_q}\frac{S}{K_S + S} f_O
\end{aligned}
\tag{9.1}
$$

where ι is 0 before induction and 1 after it, and the feed stops at the working volume $V_{f,max}$. The production $q_P X V_f$ is integrated as the ledger of everything made. Every constant in Equation (9.1) is a teaching assumption. The one anchored number is the load, 3.87 g of precursor-equivalent protein per liter of broth, derived from the pilot's Table 3 as 539.8 g refolded at a 69.7 % step yield from 200 L [7]. That is protein charged to refolding per liter of broth, and the model places it at the harvest, 3.90 g/L in 744 L at 16 h, so after its own hold and wash losses a 250 L lot charges 3.26 g per liter of broth into refolding. The burden $f_b = 0.05$ leaves 5 % of the growth rate after induction, which gives case 2 an interior best hour. The pilot's own culture kept growing after its temperature-shift induction (its Fig. 1), so the value is a strong-burden teaching choice. With $f_b = 0.1$ the best induction hour tried is still near 6 h, and with 0.2 or more the earliest hour tried delivers the most.

9.3.2 Harvest Hold, Wash and Solubilization

At t_{harv} the broth moves to the hold tank with $m_H = PV_f$ of precursor and $I_H = f_{HCP} X V_f$ of host-cell protein. The tank loses precursor at $\dot{m}_H = -k_H m_H$, and whatever remains $t_{hold,H}$ after the harvest is discarded. A lot of $\min(V_{lot}, V_{A,max}, V_H)$ is charged to REC-201 when it and the solubilized tank are free, lysed for t_{lys}, and washed at constant volume, with w the wash volumes passed, as shown in Equation (9.2).

$$\frac{\mathrm{d}P_A}{\mathrm{d}w} = -k_p P_A, \qquad \frac{\mathrm{d}I_A}{\mathrm{d}w} = -k_w I_A, \qquad \frac{\mathrm{d}w}{\mathrm{d}t} = \frac{1}{t_w} \tag{9.2}$$

Three wash volumes leave 2.7 % of the host-cell protein and 84 % of the precursor. Solubilization then sets the lot's volume to $V_A = P_A/(C_{sol} s)$, with s the strength the solubilization reaches, holds it for t_{sol}, and transfers it to the solubilized tank.

9.3.3 The Refolding Competition

The refolding vessel is first charged with buffer to $V_0 = \max\left(0.02, \min(V_{R,max}, m_S/C_{tgt}) - V_S\right)$, so that the lot of mass m_S in volume V_S reaches C_{tgt} when it is all in. The charge then enters at F with concentration $c_U = U_S/V_S$, and the precursor-equivalent masses of unfolded (U), correctly folded (N) and aggregated (A) protein follow Equation (9.3).

$$\begin{aligned} \frac{\mathrm{d}(VU)}{\mathrm{d}t} &= F c_U - V\left(k_f U + k_a U^2\right) \\ \frac{\mathrm{d}(VN)}{\mathrm{d}t} &= V k_f U - h k_N V N, \qquad \frac{\mathrm{d}(VA)}{\mathrm{d}t} = V k_a U^2 \\ \frac{\mathrm{d}V}{\mathrm{d}t} &= F, \qquad \frac{\mathrm{d}V_S}{\mathrm{d}t} = -F, \qquad \frac{\mathrm{d}U_S}{\mathrm{d}t} = -F c_U \end{aligned} \tag{9.3}$$

where $k_a U^2$ is the mass disappearance rate of unfolded protein, so no factor of two appears, and $h = 1$ only while a folded lot waits for the conversion vessel. With the feed off, $V(U + N + A)$ is constant, and a bolus at constant volume folds the fraction

$$\frac{N_\infty}{U_0} = \frac{\ln(1+\alpha)}{\alpha}, \qquad \alpha = \frac{k_a U_0}{k_f} \tag{9.4}$$

which depends on the charge only through α. Feeding the same lot slowly keeps U below U_0 and beats this bound, at the cost of vessel hours. Temperature and pH act on the two constants as shown in Equation (9.5).

$$\begin{aligned} k_f &= k_{f,15}\, \mathrm{e}^{\theta_f (T-15)} \exp\left[-\left(\frac{\mathrm{pH} - \mathrm{pH}^*}{\sigma_{pH}} \right)^2 \right] \\ k_a &= k_{a,15}\, \mathrm{e}^{\theta_a (T-15)}\, \mathrm{e}^{\lambda_{pH} \max(0,\ \mathrm{pH}^* - \mathrm{pH})} \end{aligned} \tag{9.5}$$

with T in °C. Folding doubles per 10 K ($Q_{10} = 2.0$) and aggregation nearly quadruples ($Q_{10} = 3.86$), so a warmer vessel finishes sooner and keeps less. At 1.5 g/L the bolus has $\alpha = 0.72$ and folds 75.3 % at 15 °C, and $\alpha = 1.39$ with 62.7 % at 25 °C.

Pause. Try Exercise 9.1 before reading on.

DID YOU KNOW? A Fed Charge in a Commercial Tank

A commercial refolding step for a two-chain immunotoxin is run in jacketed stainless steel tanks of 950 L nominal volume, and at the 5 L bench scale, with every other condition held, feeding the solubilized protein in place of adding it at once raised the volumetric productivity of the refold by 34 % [17].

9.3.4 Vessel Temperature and pH

The vessel's energy balance carries the jacket, the agitator, the room and the incoming charge, and the jacket relaxes toward the applied utility temperature, as shown in Equation (9.6).

$$\begin{aligned} \frac{\mathrm{d}T}{\mathrm{d}t} &= \frac{UA(T_j - T) + P_{mix} - UA_{amb}(T - T_{amb})}{C_{ves} + \rho c_p V} \\ &\quad + \frac{\rho c_p F (T_{feed} - T)}{C_{ves} + \rho c_p V} \\ C_j \frac{\mathrm{d}T_j}{\mathrm{d}t} &= \frac{C_j}{\tau_j}\left(T_{sup}^{app} - T_j\right) - UA(T_j - T) \\ T_{sup}^{app} &= \max\left(T_{sup},\ 45 - 50\kappa\right) \end{aligned} \tag{9.6}$$

where P_{mix} acts only while there is liquid and the empty vessel keeps its own heat capacity C_{ves}. The chiller event of the challenge sets $\kappa = 0.6$, so the supply can go no colder than 15 °C. The pH is a linear titration of the net base b about the buffer capacity, as shown in Equation (9.7).

$$\mathrm{pH} = \mathrm{pH}^{*} + \frac{b}{\beta V}, \qquad \frac{\mathrm{d}b}{\mathrm{d}t} = q_{b,max}\frac{u_b}{100} - x_a r_a V \tag{9.7}$$

clipped to 6–13, where r_a is the acid load of carbon dioxide absorbed by the alkaline buffer and $q_{b,max} = 0.25\,\mathrm{mol/h}$ of 1 M base.

9.3.5 Precipitation, Conversion and Isolation

At transfer, acid precipitation keeps R_{ppt} of the folded protein, the fraction c_A of the aggregate and $c_{I,ppt}$ of the host-cell protein, and removes the unfolded protein. The conversion vessel counts molecules, as shown in Equation (9.8).

$$\begin{aligned} r &= \frac{k_{cat} c_E f_T C_N}{K_M + C_N}, \qquad r_o = \frac{\rho\, k_{cat} c_E f_T C_I}{K_{M,o} + C_I} \\ \frac{\mathrm{d}n_N}{\mathrm{d}t} &= -rV_C, \qquad \frac{\mathrm{d}n_I}{\mathrm{d}t} = (r - r_o)V_C \\ \frac{\mathrm{d}n_G}{\mathrm{d}t} &= rV_C, \qquad \frac{\mathrm{d}n_D}{\mathrm{d}t} = r_o V_C \end{aligned} \tag{9.8}$$

where $C = n/V_C$, $f_T = \mathrm{e}^{\theta_E (T_{conv} - 15)}$, c_E is the dose as a multiple of nominal, and $k_{cat} c_{E,0}$ is a maximum rate in $\mathrm{mol\,m^{-3}\,s^{-1}}$. One precursor molecule gives one insulin molecule and its cleavage fragments n_G, which carry $MW_P - MW_I = 6.59\,\mathrm{kg/mol}$ of every 12.4 kg/mol, and the over-processed species n_D keeps insulin's mass so that the mass ledger closes exactly. Reaction water is not counted. When the endpoint rule stops the reaction, the declared map gives the lot of Equation (9.9).

$$\begin{aligned} m_{out} &= MW_I\, n_I\, R_{cap} R_{pol} \\ m_{rel} &= MW_P\, n_N + MW_I\, n_D + A_C \\ \pi &= \frac{m_{out}}{m_{out} + c_{rel} m_{rel} + c_{HCP} I_C} \\ \mathrm{ppm} &= 10^6 \frac{c_{HCP} I_C}{m_{out}} \end{aligned} \tag{9.9}$$

A lot is accepted when $\pi \geq 95.0\,\%$ and the host-cell protein is at most 1000 ppm. The carried-over related substances and host-cell protein are booked as purification losses, and accepted mass is insulin only. These limits are teaching values. Real acceptance criteria come from the product's specification, the pharmacopeia and regulatory guidance such as ICH Q6B [18], and a simulated pass is not batch release. The scorecard's lead tile is $\dot{m} = m_{acc}/t_{camp}$, accepted insulin divided by the whole 72 h horizon, so stopping early cannot flatter it, and unprocessed, expired and in-train material are reported beside it.

9.3.6 Instruments

Only the instruments reach a controller. The vessel and jacket temperatures pass first-order lags of 60 and 30 s with 0.05 K of noise, the pH probe a 120 s lag, 0.008 of noise and a drift δ_{pH}, the fermenter's optical probe 4 % noise and a gain drift, and the vessel weights 0.2 %. Two sampled analyzers of the shared framework (Section 1.6) carry the delayed assays. The RP-HPLC assay of the refolding vessel is drawn every hour and reported 0.75 h later with the total protein ($\sigma = 0.03\,\mathrm{g/L}$) and the folded, aggregated and unfolded shares (σ 0.012, 0.010, 0.012). The conversion assay is drawn every 30 min and reported 30 min later with insulin, precursor and related substances (σ 0.012, 0.010, 0.012 g/L). Each reading holds until the next report, and both analyzers are reset when their vessel is charged.

9.3.7 Parameters, Evidence and Numerics

Every parameter is in Table 9.2 with its source label. A value marked S is read in a source, F is fitted to a source trend, D is derived by arithmetic on source numbers, and T is a teaching assumption. Every parameter is adjustable in the application's parameter panel. Inside the engine of Section 1.6.1, states advance by fourth-order Runge-Kutta steps of $\Delta t = 5\,\mathrm{s}$, the controller scan is 30 s, and one row a minute is logged. The fastest dynamics are the jacket turnover (300 s) and the probes, while folding and conversion run over hours. Every irreversible event (induction, harvest, expiry, each lot transfer, the endpoint and isolation) happens once at a step boundary after the integration, and the right-hand side is pure. Each case starts with the fermenter at $V_{f,0} = 0.60\,\mathrm{m^3}$, $X_0 = 6\,\mathrm{kg/m^3}$ and $S_0 = 1.5\,\mathrm{kg/m^3}$, every downstream vessel empty, and the refolding vessel and jacket at the temperature setpoint.

9.3.8 Assumptions

The refolding vessel is one well-mixed liquid. A real charge enters at a point where the local unfolded concentration exceeds the vessel average, which is where a second-order term matters most, so the model can show only the feed rate among the three levers a plant uses (charge point, agitator and rate). The competition has one unfolded species, one correctly folded species and an irreversible aggregate, with no disulfide intermediates or recoverable misfolded monomer. The concentration dependence is the declared second-order form, not a fit (Section 9.4). Conversion is one effective enzyme reaction with one over-processing path, while the pilot uses two enzymes in sequence that make an intermediate this model does not carry. The pH model is a linear titration valid for a few tenths of a unit about the optimum. Everything upstream of the harvest is a teaching model of no particular strain, and purification is a declared recovery and selectivity map, because preparative chromatography is the subject of Volume 8, Chapter 11, where the capacity that decides when to take a column off is itself unmeasured.

9.4 Code Verification and Agreement With the Pilot Study

Code verification and agreement with the literature are kept apart here. The first asks whether the program solves the stated equations, and the second whether the stated equations reproduce a published trend. No claim is made about what a reader learns from either. The identities of the reduced model were checked independently of the application's own tests on six disturbed runs and are listed in Table 9.3. The ledger relates everything the cells made to everything in a vessel, lost, expired or isolated, and it closes to 1.8×10^{-13} relative, or 1.8×10^{-10} in the overloaded case where both hold limits expire.

Table 9.2 Parameters and Their Evidence. S Source, F Fitted, D Derived, T Teaching Assumption

Symbol	Meaning	Value	Unit	Ev.
$k_{f,15}$	folding constant, 15 °C, pH 10.6	4.61×10^{-5}	s^{-1}	F
$k_{a,15}$	aggregation constant at 1.5 g/L	2.21×10^{-5}	$m^3/(kg\,s)$	F
θ_f	folding temperature coefficient	0.0693	K^{-1}	T
θ_a	aggregation temperature coefficient	0.1351	K^{-1}	F
pH^*, σ_{pH}	folding optimum, width	10.6, 1.0	—	S, T
λ_{pH}	aggregation rise below pH^*	1.2	per pH	T
C_{sol}	solubilized concentration	10	kg/m^3	T
$V_{R,max}$, $V_{C,max}$	refolding, conversion volumes	1.0, 1.2	m^3	T
$k_{cat}c_{E,0}$	maximum conversion rate	4.61×10^{-5}	$mol/(m^3\,s)$	F
ρ	over-processing relative rate	0.15	—	T
K_M, $K_{M,o}$	Michaelis constants	0.5, 0.5	mol/m^3	T
θ_E	enzyme temperature coefficient	0.0693	K^{-1}	T
MW_P, MW_I	precursor (112 residues), insulin	12.4, 5.808	kg/mol	D, S
R_{ppt}	folded protein kept at precipitation	0.993	—	S
c_A, $c_{I,ppt}$	aggregate, host protein carried	0.05, 0.40	—	T
R_{cap}, R_{pol}	capture, polish recovery	0.90, 0.85	—	T
c_{rel}, c_{HCP}	carried into the bulk	0.040, 0.005	—	T
k_w, k_p	wash removal, precursor loss	1.2, 0.06	per volume	T
t_{lys}, t_w, t_{sol}, t_{turn}	lysis, per wash, solubilize, turnaround	0.5, 0.5, 1, 0.5	h	T
$V_{A,max}$	recovery vessel volume	0.35	m^3	T
k_H, k_N	losses in harvest hold, while waiting	0.010, 0.012	h^{-1}	T
$t_{hold,H}$, $t_{hold,R}$	harvest, folded-lot hold limits	36, 10	h	T
β, r_a	buffer capacity, acid load	2.0, 0.09	mol/m^3 per pH, per h	T
UA, UA_{amb}	jacket, room conductance	2000, 120	W/K	T
C_j, C_{ves}, τ_j	jacket, shell capacity, turnover	0.209, 0.30, 300	MJ/K, s	T
P_{mix}, T_{amb}, T_{feed}	agitator, room, feed	300, 21, 15	W, °C	T
μ_{max}, K_S	growth, half-saturation	0.52, 0.05	h^{-1}, kg/m^3	T
$Y_{X/S}$, $Y_{P/S}$, m_S	yields, maintenance	0.42, 0.12, 0.035	–, –, h^{-1}	T
$q_{P,max}$, τ_q, f_b	production, decay, burden	0.0444, 5, 0.05	h^{-1}, h, –	T
X_{cap}, k_d, $k_{P,loss}$	capacity, death, loss	38, 0.004, 0.006	kg/m^3, h^{-1}	T
S_f, f_{HCP}, $V_{f,max}$	feed, host protein, volume	600, 0.085, 1.0	kg/m^3, –, m^3	T
P_H	precursor at harvest, per broth	3.87	kg/m^3	D
π_{min}, HCP_{max}	lot limits	95.0, 1000	%, ppm	T
t_{HPLC}, t_{conv}	assay delays	0.75, 0.5	h	T

Table 9.3 Code Verification and Agreement With the Pilot Study

Check	Result
Code verification	
protein ledger, six disturbed campaigns	$\leq 1.8 \times 10^{-13}$ (1.8×10^{-10} case 9)
$V(U+N+A)$ with the feed off, 20 h	$< 10^{-12}$ relative
$\dot{V} = F$, vessel mass equals fed less misfolded	$< 10^{-12}$ relative
$k_a = 0$: $U = U_0 e^{-k_f t}$	$< 10^{-9}$ relative
bolus against Equation (9.4), 15, 20, 25 °C	$< 3 \times 10^{-4}$ in the fraction
moles and fragment mass in conversion	$< 10^{-12}$ relative
wash against e^{-kw}, every lot	3.8×10^{-13}
Δt halved, challenge (suggested)	5.565 and 5.565 g/h
playback in 60 and 3600 s chunks	no difference
Fitted, so code verification of the fit	
bolus, 1.5 g/L, 15 and 25 °C	75.31, 62.67 % (source 75.3, 62.8)
conversion, fixed 16 h, 15 °C	56.6 % molar (source 56.6)
Withheld from the fit	
Fig. 4A, eight digitized points	within 0.11 of the 20 h value
Table 2 at 1.5–2.0 g/L	71.2, 66.4 against 68.0–74.2, 57.7–66.0
Table 2 at 0.5–1.0 g/L	83.9, 76.9 against 53.6–66.0 (opposite)

The pilot study gives three numbers that fix the fitted constants. Refolding at 15 and 25 °C yielded 75.3 and 62.8 %, with the concentration of that pair not stated [7]. The fit assumes the article's optimum of 1.5 g/L, so $k_{a,15}$ scales with that assumption and would be 0.75 times as large at 2.0 g/L.

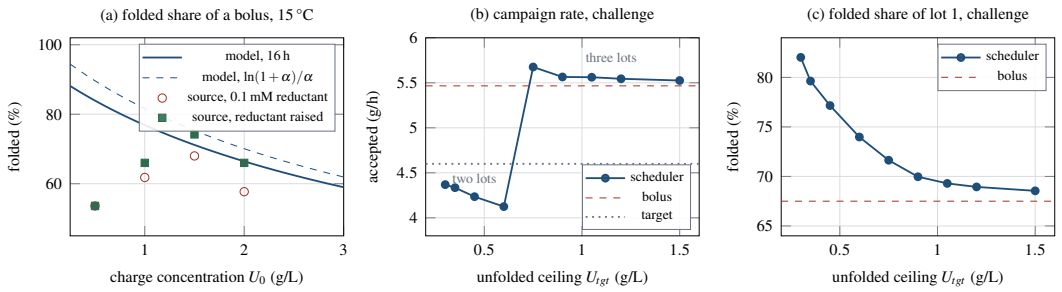


Figure 9.2 The folding law against its source and the campaign it drives. (a) The model's 16 h bolus yield at 15 °C (solid) and its closed form (dashed) against the two concentration series of the pilot study's Table 2, which were withheld from the fit (open circles at constant reductant, filled squares with the reductant raised). (b) Accepted insulin per campaign hour in the challenge with the suggested settings and the scheduler's ceiling U_{tgt} varied (solid, filled circles), against the bolus (dashed) and the target (dotted). (c) The first lot's folded share in the same runs, with the bolus again dashed. Below 0.75 g/L each lot folds better but only two lots fit. At 0.5 g/L the two source series share one point.

The constant $k_{f,15}$ is fitted to the statement that the 15 °C reaction was complete after 15 h [7], and the model reaches 95 % of its final folded mass in 16.5 h. The pilot's enzyme step converted 56.6 % of the refolded precursor molecules (25.8 % by mass, scaled by its residue ratio of 51 to 112), and $k_{cat}c_{E,0}$ is fitted so that a 16 h reaction converts the same, given $\rho = 0.15$. One end-point yield cannot fix both constants, so ρ is a teaching value, and with it the insulin peaks at 72.4 % of the molecules after 7.49 h.

Two comparisons were withheld from the fit. The article's Fig. 4A, sampled every 1.5 h for 20 h and digitized to about 10 g, reaches 95 % of its 20 h value at about 12 h at 15 °C, against 14.0 h in the model, and the eight digitized fractions agree with the model within 0.11. The source curve is nearly linear for 9 h where the first-order model is front-loaded. The second check is the concentration series of the article's Table 2, as shown in Figure 9.2a. In both of its series, one at constant reductant and one with the reductant raised with the protein, the yield rises from 0.5 g/L to 1.5 g/L and falls at 2.0 g/L. The model's 16 h yield falls all the way, from 83.9 to 66.4 %. It agrees with the raised-reductant series at 1.5 and 2.0 g/L (71.2 against 74.2, 66.4 against 66.0) and runs opposite to both series below 1.5 g/L. The second-order law is the plan's declared reduced model and the mechanism of the cases, so it is kept, and every number the simulator gives outside 1.5–2 g/L and 10–30 °C is an extrapolation of a teaching model. These comparisons are agreement with one study's data, not validation of the model for a plant.

9.5 The Controllers and What They Read

The four loops are the velocity-form PID of Section 1.6.2 with the derivative off, executed every 30 s on the measurement packet. With $e_k = \mathrm{SP}_k - \mathrm{PV}_k$, the law is Equations (1.12) to (1.15) written for this chapter, the running sum S_k of Chapter 1 appearing here as s_k and its integral I_k carrying the anti-windup term.

$$\begin{aligned} s_k &= s_{k-1} + K_c(e_k - e_{k-1}), \qquad \hat{u}_k = u_0 + s_k + I_{k-1} \\ u_k &= \min\left(u_{max},\ \max(u_{min},\ \hat{u}_k)\right) \\ I_k &= I_{k-1} + \left[\frac{K_c}{\tau_I}e_k + \frac{u_k - \hat{u}_k}{\tau_I}\right]\Delta t_c \end{aligned} \tag{9.10}$$

where u_0 is the output at which the loop is initialized and $\Delta t_c = 30\,\mathrm{s}$, with back-calculation anti-windup at a tracking time equal to τ_I and the clamp on the emitted output only. In Manual each

loop tracks the operator's output, so transfers are bumpless. TIC-301 sets the jacket setpoint within 2–40 °C at $K_c = 4\,\mathrm{K/K}$ and $\tau_I = 30\,\mathrm{min}$. TIC-302 sets the requested utility supply within −5–45 °C at $K_c = 3\,\mathrm{K/K}$ and $\tau_I = 5\,\mathrm{min}$, and its card shows the requested supply beside the applied one of Equation (9.6) when the chiller is short. AIC-303 drives the base pump within 0–100 % at $K_c = 150\,\%$ per pH and $\tau_I = 30\,\mathrm{min}$. All three are reverse acting, with action $a = +1$ in the convention of Section 1.3, because a warmer jacket warms the vessel, a warmer utility supply warms the jacket and more base raises the pH. TIC-301 and TIC-302 are the outer and inner halves of one cascade, so the inner loop is tuned six times faster than the outer one.

9.5.1 Three Ways to Charge

FIC-304 charges the solubilized lot by one of three strategies. A bolus runs the pump at $F_{max} = 400\,\mathrm{L/h}$ until the tank is empty, and the first lot of case 4 is in after 0.23 h. A fixed rate runs at F_{set}, 60 L/h by default. The scheduler holds the estimated unfolded concentration at its ceiling with a feedforward and a PI trim, as shown in Equation (9.11).

$$\begin{aligned} F_{FF} &= \frac{V\left(\hat{k}_f U_{tgt} + \hat{k}_a U_{tgt}^2\right)}{C_{sol}} \\ F^* &= \max\left(0,\; F_{FF} + \mathrm{PI}(U_{tgt} - \hat{U})\right) \\ F &= \min\left(F^*,\; F_{max},\; \frac{V_{fin} - V}{\Delta t},\; \frac{V_S}{\Delta t}\right) \end{aligned} \tag{9.11}$$

where $\hat{k}_f$ and $\hat{k}_a$ are Equation (9.5) evaluated with the platform's declared constants at the measured temperature and pH, the trim has $K_c = 100\,\mathrm{L/h}$ per g/L and $\tau_I = 30\,\mathrm{min}$, and its output is bounded to $[-F_{FF},\; F_{max} - F_{FF}]$. The volume V_{fin} is the one the charge will reach. The scheduler is a feedforward on a model with feedback on an estimate, and none of the receding-horizon machinery of Section 1.6.3 is used: no objective is minimized and nothing is predicted beyond the present rate. An advanced option stretches the charge so that the lot is ready when the conversion vessel is expected to be free. It never engages in a shipped case or in the challenge, because the conversion vessel is always free first.

> **TRY IT IN THE APP** **4 · A Bolus Charge Against a Fed One**
>
> Load the case and press run. On the Refolding view, watch the competition panel while the first lot goes in as a bolus: folding and aggregation both start at once, and the aggregation rate falls away only as the unfolded protein is used up. The chapter reports 67.7 % of this lot folded.

9.5.2 The Estimator

The estimator integrates Equation (9.3) on the commanded feed with the platform's constants, and corrects it with each RP-HPLC report against the value it had when that sample was drawn, which is the sampled correction of Section 1.4.5 with a fixed weight in place of the logarithmic bias of Equation (1.8), as shown in Equation (9.12).

$$\begin{aligned} \frac{\mathrm{d}(\hat{V}\hat{U})}{\mathrm{d}t} &= F C_{sol} - \hat{V}\left(\hat{k}_f \hat{U} + \hat{k}_a \hat{U}^2\right) \\ \hat{U} &\leftarrow \hat{U} + g\left(U_{assay}(t_s) - \hat{U}(t_s)\right), \qquad g = 0.6 \\ \hat{N}_{plat} &= \hat{V}\hat{U}\,\frac{\ln(1+\hat{\alpha})}{\hat{\alpha}}, \qquad \hat{\alpha} = \frac{\hat{k}_a \hat{U}}{\hat{k}_f} \end{aligned} \tag{9.12}$$

where t_s is the sample time, C_{sol} is the design value (so a weak solubilization is unknown to the estimator), $\hat{N}$ and $\hat{A}$ are corrected the same way, and the assay's masses use the vessel volume measured at t_s. The estimate history is stored every 300 s. A conversion estimate on Equation (9.8) is corrected the same way by each conversion report, and after each report it is integrated forward in

Table 9.4 The Ten Cases

#	Case and departure	Teaching point
1	recipe, Train view, 72 h	precursor mass against insulin
2	induction time moved	titer against usable product
3	one wash volume	host-cell protein against loss
4	transfer on the estimate, Refolding view	bolus against fed charge
5	both model rules, temperature moved	faster lots, less folded
6	acid load ×3, base pump in Manual at 0	pH, folding and aggregation
7	endpoint rule moved	a peak behind a delayed assay
8	harvest at 22 h	a harvest the train cannot take
9	22 h harvest of a faster feed, $t_{hold,H} = 10$ h, 350 L lots	infeasible, logged sweep
10	challenge: 1.5 wash volumes, bolus to 2.5 g/L, acid ×4 at 30 h, chiller 60 % at 46 h	the whole campaign

60 s steps for up to 30 h to forecast the time and height of the insulin peak. The estimator keeps the declared constants when a case perturbs the plant's. Chapter 8 corrects a culture's glucose the same way, from one sample a day, and Volume 5, Chapter 4 steers a cook whose grade arrives hours after the move that set it.

9.5.3 Transfer and Endpoint Rules

The transfer rule empties the refolding vessel after a fixed hold $t_{fold} = 16$ h from the start of the charge, when two successive RP-HPLC folded shares differ by less than 0.01 after 2 h, or when the estimate says the charge can make less than $\varepsilon = 3\,\%$ of $\hat{N}$ more ($\hat{N}_{plat} \le \varepsilon\hat{N}$, after one report and 1 h). The endpoint rule stops the conversion after a fixed $t_{end} = 16$ h, after two successive reports more than 2 % below the lot's best insulin report, or when the forecast puts the peak within 0.5 h. The pilot's methods describe refolding by dilution held for 16 h and an enzyme reaction at 15 °C for 16 h [7], which the bolus and the two time rules follow. The application keeps three models apart: the plant of Section 9.3, the identified behavior a learner reads from the trends, and the controller's model, which is the platform's declared constants inside Equations (9.11) and (9.12) and knows nothing of a perturbed plant.

> **TRY IT IN THE APP** **7 · The Conversion Endpoint Behind a Delayed Assay**
>
> Run this case once with each of the three endpoint rules. The fixed time stops at 16 h with 57.3 % of the molecules converted, the assay rule at 9.01 h with 71.7 and the forecast at 7.51 h with 72.5, against a true peak of 72.5 % at 7.64 h.

9.6 Ten Cases, Measured

The cases are the application's scenarios in the sense of Section 1.6.4, and they are listed in Table 9.4. All start from a recipe modeled on the pilot's, a 6 h induction, a 16 h harvest, 250 L lots washed three times, a bolus to 2.0 g/L at 15 °C and pH 10.6, and fixed 16 h holds, and differ from it only where their rows say so. Every number in this section comes from the implemented model run without a display, advanced by the same step sequence the browser uses, with seed 20261020 and every setting typed in before the clock starts. Loading order does not matter, since each case produced identical logs when opened first, when selected after another case and when rerun from Reset.

9.6.1 The Competition in One Vessel

The opening experiment charges the first lot too quickly, as shown in Figure 9.4. At 21.5 h, 2.5 h after a bolus put 815 g into 408 L, the vessel holds 416 g unfolded, 232 g folded and 167 g aggregated, and aggregation runs at 35 g/h against folding at 69 g/h. The lot ends at 67.7 % folded in case 4, within 2.4 points of the closed form for its $\alpha = 0.96$ (70.1 %). More buffer is the other lever. The same lot charged as a bolus into 815 L for 1.0 g/L folds 78.3 %. The scheduler at 0.35 g/L gains 17.0 points

Table 9.5 Measured Outcomes of Cases 1 to 9

#	Measured
1	3013 g made, 217 g accepted in two lots, 3.02 g/h. Of the made: precipitation 24.3 %, fragments 18.5, wash 14.8, still in the train 14.6, cells and holds 10.4, purification 10.1.
2	precursor at harvest, induction at 3, 6, 7, 10 h: 2166, 2905, 2933, 2399 g. Accepted at 4, 6, 10 h: 188, 217, 179 g.
3	lot 1 at 0, 1, 2, 3, 5 wash volumes: 131.6, 124.3, 117.1, 110.3, 97.8 g at 8349, 2789, 892, 285, 29 ppm. Washed out at 1, 3, 5: 158, 445, 696 g.
4	lot 1, bolus to 2.0 g/L: 67.7 % folded, 30.0 aggregated, 18.8 h. Scheduler at 0.35, 0.6, 0.9 g/L: 84.7, 79.0, 72.1 % in 39.0, 27.8, 18.8 h. At 0.35, 665 g expire and nothing is accepted (bolus 112.7 g).
5	lot 1 at 15, 20, 25, 30 °C: 67.7, 62.1, 55.1, 48.7 % in 18.8, 14.8, 9.5, 6.8 h. Campaign 3.91, 5.04, 4.53, 4.04 g/h.
6	pH falls 0.13 h^{-1} to 8.44, lot folds 46.2 % with 47.9 aggregated, against 65.4 and 30.8 in Auto.
7	lot 1, fixed, assay, forecast: 57.3, 71.7, 72.5 % of molecules at 16.00, 9.01, 7.51 h, 110.3, 137.9, 139.4 g. True peak 72.5 % at 7.64 h. Campaign 3.02, 3.69, 3.82 g/h.
8	harvest at 16, 19, 22 h: 2905, 3131, 3238 g precursor, 217.3, 226.0, 114.6 g accepted, 0, 59, 136 g expired.
9	5961 g harvested, 1618 g expired, 308.1 g in one lot (4.28 g/h). Best of 48 runs 542.5 g with 539 g lost, none clears the tank.

Table 9.6 The Challenge With the Plant Perturbed and the Controller's Constants Kept (Accepted Insulin per Campaign Hour, g/h)

Plant	Scheduler 0.9	Bolus	Verdict
declared	5.56	5.47	both pass
$k_{a,15}$ −20 / +20 %	5.89 / 5.34	5.79 / 5.19	both pass
$k_{f,15}$ −20 %	3.68	5.01	scheduler fails
$k_{f,15}$ +20 %	5.94	5.74	both pass
θ_a −20 %	5.77	5.66	both pass
$k_{cat}c_{E,0}$ −20 %	5.55	5.45	both pass
ρ +50 %	4.98	4.91	both pass
c_{rel} 0.050 / 0.070	5.56 / 0.00	5.47 / 0.00	fail at 0.070
RP-HPLC delay 1.5 h	5.59	5.47	both pass
conversion assay delay 1 h	5.61	5.36	both pass

over the bolus, and its cost appears in the campaign rather than in the vessel. In case 5 temperature trades the same way. The first lot keeps less at every step warmer, but at 20 °C a third lot fits and the campaign makes 5.04 g/h against 3.91 at 15 °C, while at 30 °C the lots are too poor and the rate falls back to 4.04 g/h. With the pH loop in Manual in case 6, the ratio k_aU/k_f at 1 g/L rises from 0.57 at pH 10.47 to 5.6 at pH 9.52 and about 670 at 8.44, and the lot loses 19.2 points of folded protein.

🕮 **Pause.** Try Exercise 9.2 before reading on.

9.6.2 The Wash and the Endpoint

Host-cell protein decides the wash. One wash volume leaves 2789 ppm and rejects a lot of 124.3 g, two pass at 892 ppm, and each further volume costs 6 to 7 g of insulin per lot. The endpoint is the peak of a reaction that consumes its product. The fixed 16 h reaction lands 15.2 points below the peak and leaves the lot at 95.76 % purity, 0.76 points above the limit. The assay rule stops 1.4 h after the true peak, and at half the dose the assay's own scatter stops the reaction 5.8 points short of the forecast's 72.5 %.

> **+ DID YOU KNOW? Twice the Reaction Time, a Fifth Less Insulin**
>
> From the simulator: in case 7 the conversion stopped by the forecast rule ends after 7.51 h with 139.4 g of insulin, while the fixed 16 h reaction ends with 110.3 g (Table 9.5). More than twice the reaction time leaves a fifth less product, because the enzymes go on to cut what they have made.

9.6.3 Titer, Recipe and Concentration-Aware Scheduling

The plan's comparison is run inside the challenge, with its acid upset at 30 h and chiller loss at 46 h, and the stage ledgers are listed in Table 9.7. The fixed recipe, with only the wash corrected to three

Table 9.7 Campaign Ledgers by Policy in the Challenge (Acid Upset at 30 h, Chiller Loss at 46 h, 72 h Horizon, Default Seed). Losses as Percent of the Precursor Made

Policy	Made (g)	Wash (%)	Precip. (%)	Fragm. (%)	Purif. (%)	In train (%)	Expired (g)	Lot 1 folded (%, h)	Buffer (L)	Lots	Accepted (g)	Rate, verdict (g/h)
shipped (1.5 wash volumes, bolus to 2.5 g/L)	3013	7.7	29.9	19.2	10.6	15.1	0	62.3, 16.0	744	0 of 2	0	0.00, fails
fixed recipe (shipped, 3 wash volumes)	3013	14.8	27.2	17.5	9.6	13.7	0	62.3, 16.0	677	2	208	2.88, fails
titer: recipe, harvest at 22 h	3459	13.5	17.4	8.0	4.4	35.3	136	62.2, 16.0	708	1	109	1.51, fails
titer: suggested, harvest at 22 h	3459	14.1	20.0	15.6	9.5	17.7	0	68.5, 12.8	1400	2	296	4.11, fails
titer: suggested, induction at 7 h	3027	14.7	22.0	24.4	14.9	0.0	0	70.0, 14.8	1280	3	405	5.62, passes
suggested settings, bolus	3013	14.8	23.5	23.9	14.7	0.0	0	67.5, 14.8	1283	3	394	5.47, passes
suggested settings, scheduler at 0.9 g/L	3013	14.6	21.8	24.3	14.8	0.0	0	70.0, 14.8	1267	3	401	5.56, passes
suggested settings, scheduler at 0.35 g/L	3013	9.9	9.8	18.8	11.2	0.0	665	79.6, 20.4	857	2	312	4.34, fails

volumes, folds each lot to about 62 % in a 16 h hold, converts 58 % of the molecules, and finishes two lots for 2.88 g/h. Maximizing titer by a later harvest is the wrong lever. Harvesting at 22 h raises the precursor made from 3013 to 3459 g, but the recipe then finishes one lot (1.51 g/h) with 1222 g still in the train, and even the suggested downstream settings finish only two (4.11 g/h). Maximizing titer by the induction hour is harmless here. Induction at 7 h instead of 6 adds 28 g at the harvest and, with the suggested settings, raises the rate from 5.56 to 5.62 g/h, because the 16 h harvest leaves the train time for three lots. The difference between the two is time. Extra precursor helps only when it arrives without delaying the first lot.

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

Concentration-aware scheduling wins narrowly. The suggested settings (three wash volumes, 1.5 g/L charge, 20 °C for both vessels, both model rules and the scheduler at 0.9 g/L) isolate 401 g in three lots, 5.56 g/h, against 5.47 g/h for a bolus with the same other settings. The scheduler's first lot folds 70.0 % against 67.5 in the same 14.8 h, so at a loose ceiling the better fold costs no vessel time. Over six paired seeds the scheduler makes 5.55–5.61 g/h and the bolus 5.43–5.47 g/h, a margin of 0.10–0.16 g/h. The ceiling decides which side of a cliff the campaign is on, as shown in Figure 9.2b and c. Every ceiling from 0.75 to 1.5 g/L fits three lots and passes, and the lowest of them is best, 5.68 g/h with a first lot folded 71.6 %, and 5.64–5.68 g/h over the six seeds. At 0.6 g/L the first lot folds 74.0 % but only two lots fit and the rate drops to 4.12 g/h, and at 0.35 g/L the campaign also loses 665 g to the harvest clock. The suggested set is a passing configuration, not the optimum. The eight policies are drawn side by side in Figure 9.3.

9.6.4 Robustness of the Comparison

The scheduler's advantage rests on the controller's folding constants, as shown in Table 9.6. The plant was perturbed while the estimator kept the declared constants. Slower aggregation, faster folding, a weaker temperature dependence of aggregation and slower assays all leave the ranking intact. A plant that folds 20 % slower than the controller assumes reverses it. The scheduler's model rule then waits longer, its third lot is ready too late, and it fails at 3.68 g/h, while the bolus passes at 5.01 g/h. The 0.75 g/L ceiling fails the same way (3.77 g/h). The purity gate is the most sensitive declared number. Recipe lots sit at 95.76 % and flip to rejected at $c_{rel} = 0.048$, 19 % above the declared 0.040, so cases 1, 3 and 8 depend on it, while the suggested settings pass from 0.030 to 0.060 and fail at 0.070, as does the bolus variant at 0.070.

9.6.5 An Infeasible Harvest and the Challenge

Case 9 delivers 5961 g of precursor into a tank with a 10 h hold limit. As shipped, two 350 L lots leave the tank at 22.0 and 25.5 h, the clock runs out at 32 h and 1618 g (27 %) is discarded. A logged sweep of 48 runs over lot size, charge strategy, refolding temperature and conversion temperature finds no setting that clears the tank. The least loss is 539 g, with three lots and 542.5 g isolated at 7.54 g/h, and the run that isolates the most, 601.1 g, still discards 1618 g. The case is scored for recognizing the limit and getting more through it, not for an impossible target. The challenge's target is 4.6 g/h

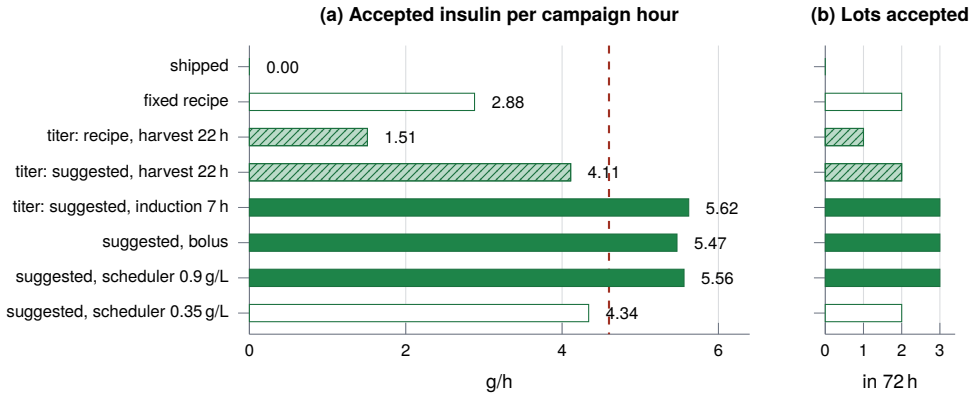


Figure 9.3 The eight policies of Table 9.7 drawn as bars: accepted insulin per campaign hour against the challenge target of 4.6 g/h (dashed line), and the lots the 72 h horizon accepted. Solid bars pass, hatched bars chase titer and white bars fail. Notice that the three passing policies are the three that finish three lots, that the richest harvest of all (two hatched bars) buys at most two, and that the lowest ceiling tried, 0.35 g/L, folds its lots best and still ends below the target.

with every lot accepted, nothing expired, and the vessel within 1.5 °C and 0.3 of its setpoints. As shipped both lots are rejected on host-cell protein at 1654 and 1685 ppm and the chiller loss leaves a 15 °C vessel 1.6 °C high. The suggested settings pass at all six seeds and, at the default seed, hold the vessel within 0.04 °C and 0.12 of pH.

The model runs 850 000 to 875 000 steps per second headless on the build host, so a 72 h campaign of 51 900 steps takes 0.06 s, and in a browser it plays in 36 s at the top speed of 7200 times real time. Halving the step to 2.5 s leaves the challenge's rates unchanged to 0.001 g/h and the first lot's folded share to 0.01 point.

9.7 Technology in Practice

The simulated train carries four loops, two sampled analyzers and two rule sets. Figure 9.5 draws them with the devices a plant would install.

What Measures It

Temperature, pH and vessel weight are read continuously, and the composition is not. A glass electrode gives millivolts per pH unit and a resistance thermometer or a thermocouple in a thermowell gives the vessel and jacket temperatures, which is what the 60 s, 30 s and 120 s lags of Section 9.3 stand for. The share of the protein that has folded has no probe at all. It is drawn as a sample and separated, and the chapter's RP-HPLC analyzer is that instrument, reporting one result per cycle and holding it until the next.

> **TECHNOLOGY Sensor: At-Line Liquid Chromatography**
>
> The folded, aggregated and unfolded shares of a refolding vessel are read by reversed-phase liquid chromatography on a drawn sample. The separation is denaturing, so it tolerates the chaotropic salts of this stream, which interfere with the measuring principle of near-infrared and Raman probes, and at-line means a measurement made close to the process stream. A method developed for inclusion-body work runs in 8.1 min at 78 °C and 1.2 mL/min, short enough to steer a step while it proceeds, where the gel assay it replaces is too slow for in-process control [19]. The chapter's analyzer stands for this device: one sample an hour, the report 0.75 h later, and nothing in between.



Figure 9.4 A bolus charge aggregating, case 4 as shipped at $t = 21.5\,\text{h}$ on the Refolding view, cropped to the plant view. The first lot of 815 g went in as a bolus at 19.0 h to 2.0 g/L in 408 L, and 2.5 h later it is 28.5 % folded, 20.5 % aggregated and 51.0 % unfolded (232, 167 and 416 g, teaching truth). Callout 1 is the vessel with its temperature, pH and assay badges. Callout 2 is the competition panel, which gives $k_f = 0.17\,\text{h}^{-1}$ and $k_a = 0.08\,\text{L}/(\text{g}\,\text{h})$ at 15.0 °C and pH 10.58, the true unfolded 1.021 g/L against the estimate 1.037 g/L, and folding at 69 g/h against aggregation at 35 g/h. Callout 3 is the charge drawn as folded, unfolded and aggregated mass, and callout 4 the latest RP-HPLC report, from the sample drawn at 20.0 h, which reads 12.3 % folded and 10.1 % aggregated with the next report due in 0.3 h. Callout 5 is the status line: 902 g of harvest waits behind the lot.

What Moves It

Four final elements act. The charge pump meters the solubilized lot, a base pump doses the titrant for the pH loop, a utility valve sets the jacket supply through the cascade, and a transfer valve empties the vessel into the conversion stage once a rule fires. Only the first two run continuously, and the chapter treats all four as ideal: a commanded flow becomes that flow. A plant pays for the difference in calibration work.

> **TECHNOLOGY** **Final Element: The Charge Pump**
>
> The solubilized lot is metered into the refolding vessel by a positive-displacement pump. A peristaltic head is common in sterile service, because the liquid touches only a length of replaceable tubing and the delivered volume follows the rollers, so the rate is set by speed alone. The dose is referred to weight. In a study of 149 847 pharmaceutical dosing measurements the pump was calibrated against a balance before every run, fitting a fresh length of tubing shifted that calibration, and one fifth of the recorded sequences drifted within a single calibration period [20]. A worn tube shows on a trend as a vessel weight that climbs more slowly than the commanded charge, which is the one failure the chapter's ideal F cannot produce.

Where the Controller Runs

The loops of a recovery train are hosted on a distributed control system, an architecture that puts a supervisory level over local controllers and is used for pharmaceutical processing among other industries [21]. Supervisory and lower loops run continuously with cycle times from milliseconds to minutes, and loops are nested so that the setpoint of one is the process variable of another [21], which is the cascade TIC-301 to TIC-302 of Section 9.5. The chapter's controller scan is 30 s, slow for a jacket of 300 s turnover and fast for a reaction of hours. The scheduler, the estimator and the transfer

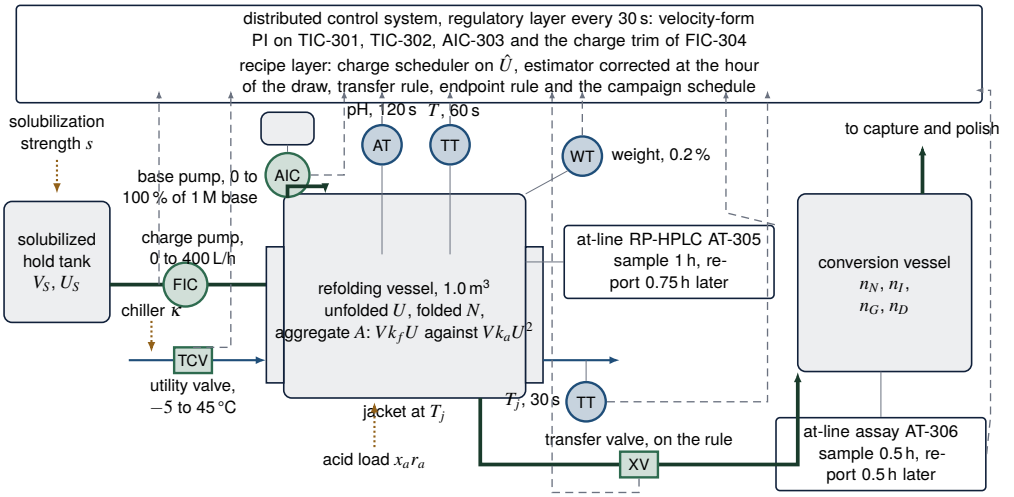


Figure 9.5 The chapter's loops with the devices a plant would install. Solid green lines carry protein, thin blue lines the jacket utility, dashed gray lines the signals and dotted amber arrows the disturbances. Notice that the two composition measurements are the only instruments that are not continuous: each at-line chromatograph reports one result per cycle, after the delay printed beside it.

and endpoint rules belong one layer above the loops, in the recipe or phase logic that sequences a batch, where the delayed assay arrives and the decision to charge, transfer or stop is taken.

> **IN PRACTICE** **Dilute Early, and Quench at the Peak**
>
> Two rules of thumb govern the two reactions in service. Proinsulin is diluted to under 1 mg/mL in the refolding buffer to prevent cross-folding between molecules, and the enzyme reaction that follows is quenched at the maximal insulin point, because the des-threonine species grows with reaction time [22]. Both calls are made on chromatograms of drawn aliquots, which is why each step is run to a measured end and not to a clock. United States drug manufacturing rules require written in-process controls and the testing of in-process material at the completion of significant phases or after long storage [23].

What Is Modernizing It

Three lines of work would change the measurement boundary this chapter draws. The first is in-process chromatography: an at-line method short enough to follow solubilization as it runs turns an end-of-step check into a control signal, and the same monitoring showed that solubilization time alone moves the refolding yield that follows [19]. The second is in-situ optical monitoring. Intrinsic tryptophan and tyrosine fluorescence tracks refolding in real time down to 0.01 g/L, with the shift of the emission peak following the folding intermediates and the intensity following the extent of aggregation, and both correlations have been written into a mechanistic model as an observation function [24]. The third is the control layer itself: state estimation that fills in what no probe reads, and model-based control, including receding-horizon control, of the feed rate of a fed charge [25]. Each one shortens the 0.75 h the simulated estimator must bridge.

9.8 Limitations

The upstream fermenter is a teaching model whose kinetics and induction burden have no source, and only the load it delivers is anchored. The folding law agrees with its source's concentration series only at 1.5–2.0 g/L, the over-processing ratio is a teaching value, purification is a declared map whose carried-over fraction decides every recipe lot within 19 %, and the well-mixed vessel

cannot represent the charge point, where aggregation is worst. The controllers are four PI loops, a feedforward scheduler with a trim, and rule-based decisions. Nothing optimizes the schedule, the ceiling or the harvest, the best ceiling of 0.75 g/L is a sweep result, and the scheduler's margin of 0.10–0.16 g/h disappears if the plant folds 20 % slower than its constants. Quality is scored on simulated attributes, and every limit, hold time and horizon is a teaching value.

OPEN CHALLENGE **Refolding Still Has No Platform Recipe**

Refolding remains the bottleneck of inclusion-body processing. A review of monitoring and control lists the open drawbacks as the large buffer volumes and tank capacities, uneven mixing in those volumes, analytics that struggle at the low protein concentrations dilution forces, kinetics that run into unusable aggregate and generally low yields, with no generic platform procedure and a lack of robust control strategies [25]. Progress would look like the route that review proposes: measure the folding kinetics in real time, estimate from them the states that no probe reads, and close the loop on the feed rate of a fed charge.

9.9 Conclusion

This chapter has worked an insulin recovery train in which the folding competition, the vessel schedule and the hold clocks share one 72-hour campaign, over a protein ledger that closes to 1.8×10^{-13}. A fed charge folds 17 points more of a lot than a bolus and can still leave the campaign with nothing accepted, because the vessel it occupies is the bottleneck. A harvest 11 % richer in precursor delivers half the insulin when it delays the train, while a richer harvest at the same hour helps slightly. In the challenge, concentration-aware charging beats a bolus by 0.10–0.16 g/h over six seeds and the best ceiling tried reaches 5.68 g/h, but both scheduler settings fail when the plant folds 20 % slower than the controller assumes. The refolding constants reproduce the pilot's temperature pair by construction and its concentration series only at 1.5–2.0 g/L, and the upstream kinetics are teaching values. A fuller treatment would add an optimized schedule, a two-enzyme conversion and refolding data at fixed reductant. What the cases show is that the yield of one vessel and the productivity of a campaign are different objectives, and that the second one is what a plant is paid for.

KEY POINTS

- A vessel that keeps more of its protein can still lose the campaign, because everything behind it waits.
- Unfolded precursor folds at a rate first order in its concentration and aggregates at a rate second order in it, so a lot charged slowly keeps more protein.
- The refolding vessel is the bottleneck, and the campaign is scored by accepted insulin per hour, not by one lot's yield.
- A harvest at 22 h carries 11 % more precursor than one at 16 h and the campaign isolates about half as much insulin.
- The scheduler's margin over a bolus disappears when the plant folds 20 % slower than the controller assumes.

Where else this applies. Chapter 8 corrects a culture from a sample reported hours late, Volume 5, Chapter 4 steers a cook whose grade arrives after the move that set it, and Volume 8, Chapter 2 meters a feed to keep an unreacted store small.

References

[1] T. Kiefhaber, R. Rudolph, H.-H. Kohler, and J. Buchner, "Protein aggregation in vitro and in vivo: A quantitative model of the kinetic competition between folding and aggregation," *Nature Biotechnology*, vol. 9, no. 9, pp. 825–829, 1991.

[2] A. P. J. Middelberg, "Preparative protein refolding," *Trends in Biotechnology*, vol. 20, no. 10, pp. 437–443, 2002.

[3] S. M. Singh and A. K. Panda, "Solubilization and refolding of bacterial inclusion body proteins," *Journal of Bioscience and Bioengineering*, vol. 99, no. 4, pp. 303–310, 2005.

[4] C. J. Roberts, "Non-native protein aggregation kinetics," *Biotechnology and Bioengineering*, vol. 98, no. 5, pp. 927–938, 2007.

[5] I. Harjunkoski, C. T. Maravelias, P. Bongers, P. M. Castro, S. Engell, I. E. Grossmann, J. Hooker, C. Méndez, G. Sand, and J. Wassick, "Scope for industrial applications of production scheduling models and solution methods," *Computers & Chemical Engineering*, vol. 62, pp. 161–193, 2014.

[6] K. Lakhdar, S. S. Farid, N. J. Titchener-Hooker, and L. G. Papageorgiou, "Medium term planning of biopharmaceutical manufacture with uncertain fermentation titers," *Biotechnology Progress*, vol. 22, no. 6, pp. 1630–1636, 2006.

[7] C.-K. Kim, S.-B. Lee, and Y.-J. Son, "Large-scale refolding and enzyme reaction of human preproinsulin for production of human insulin," *Journal of Microbiology and Biotechnology*, vol. 25, no. 10, pp. 1742–1750, 2015.

[8] I. S. Johnson, "Human insulin from recombinant DNA technology," *Science*, vol. 219, no. 4585, pp. 632–637, 1983.

[9] L. Rosenfeld, "Insulin: Discovery and controversy," *Clinical Chemistry*, vol. 48, no. 12, pp. 2270–2288, 2002.

[10] D. F. Steiner, D. Cunningham, L. Spigelman, and B. Aten, "Insulin biosynthesis: Evidence for a precursor," *Science*, vol. 157, no. 3789, pp. 697–700, 1967.

[11] W. Kemmler, J. D. Peterson, and D. F. Steiner, "Studies on the conversion of proinsulin to insulin. I. conversion in vitro with trypsin and carboxypeptidase B," *Journal of Biological Chemistry*, vol. 246, no. 22, pp. 6786–6791, 1971.

[12] D. V. Goeddel, D. G. Kleid, F. Bolivar, H. L. Heyneker, D. G. Yansura, R. Crea, T. Hirose, A. Kraszewski, K. Itakura, and A. D. Riggs, "Expression in Escherichia coli of chemically synthesized genes for human insulin," *Proceedings of the National Academy of Sciences*, vol. 76, no. 1, pp. 106–110, 1979.

[13] R. E. Chance and B. H. Frank, "Research, development, production, and safety of biosynthetic human insulin," *Diabetes Care*, vol. 16, no. Supplement 3, pp. 133–142, 1993.

[14] G. Walsh, "Therapeutic insulins and their large-scale manufacture," *Applied Microbiology and Biotechnology*, vol. 67, no. 2, pp. 151–159, 2005.

[15] N. A. Baeshen, M. N. Baeshen, A. Sheikh, R. S. Bora, M. M. M. Ahmed, H. A. I. Ramadan, K. S. Saini, and E. M. Redwan, "Cell factories for insulin production," *Microbial Cell Factories*, vol. 13, p. 141, 2014.

[16] R. Rudolph and H. Lilie, "In vitro folding of inclusion body proteins," *The FASEB Journal*, vol. 10, no. 1, pp. 49–56, 1996.

[17] T. Linke, M. T. Aspelund, C. Thompson, G. Xi, A. Fulton, M. Wendeler, T. M. Pabst, X. Wang, W. K. Wang, K. Ram, and A. K. Hunter, "Development and scale-up of a commercial fed batch refolding process for an anti-CD22 two chain immunotoxin," *Biotechnology Progress*, vol. 30, no. 6, pp. 1380–1389, 2014.

[18] European Medicines Agency, "ICH Q6B specifications: Test procedures and acceptance criteria for biotechnological/biological products," Scientific guideline CPMP/ICH/365/96, effective 1 September 1999, 1999. [Online]. Available: https://www.ema.europa.eu/en/scientific-guidelines/ich-q6b-specifications-test-procedures-acceptance-criteria-biotechnological-biological-products

[19] J. Ebner, D. Humer, R. Klausser, V. Rubus, R. Pell, O. Spadiut, and J. Kopp, "At-line reversed phase liquid chromatography for in-process monitoring of inclusion body solubilization," *Bioengineering*, vol. 8, no. 6, p. 78, 2021.

[20] 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.

[21] 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, Tech. Rep. NIST Special Publication 800-82r3, 2023.

[22] Y. Y. Siew and W. Zhang, "Downstream processing of recombinant human insulin and its analogues production from *e. coli* inclusion bodies," *Bioresources and Bioprocessing*, vol. 8, no. 1, p. 65, 2021.

[23] U.S. Food and Drug Administration, "21 CFR 211.110: Sampling and testing of in-process materials and drug products," Code of Federal Regulations, title 21, revised as of April 1, 2024, 2024. [Online]. Available: https://www.govinfo.gov/content/pkg/CFR-2024-title21-vol4/xml/CFR-2024-title21-vol4-sec211-110.xml

[24] C. L. Igwe, D. F. Müller, F. Gisperg, J. N. Pauk, M. Kierein, M. Elshazly, R. Klausser, J. Kopp, O. Spadiut, and E. Přáda Brichtová, "Online monitoring of protein refolding in inclusion body processing using intrinsic fluorescence," *Analytical and Bioanalytical Chemistry*, vol. 416, no. 12, pp. 3019–3032, 2024.

[25] J. N. Pauk, J. Raju Palanisamy, J. Kager, K. Koczka, G. Berghammer, C. Herwig, and L. Veiter, "Advances in monitoring and control of refolding kinetics combining PAT and modeling," *Applied Microbiology and Biotechnology*, vol. 105, no. 6, pp. 2243–2260, 2021.

9.10 Exercises

9.1 Concept check · 2 min

A lot of solubilized precursor is charged into the refolding vessel as a bolus at 15 °C and pH 10.6. The same lot is then charged into twice as much buffer, so that the unfolded concentration once the charge is in is halved. What happens to the share of the lot that ends up correctly folded?

(a) It is unchanged, because folding is first order and the lot's mass is the same

(b) It rises, because $\alpha = k_a U_0 / k_f$ is halved and $\ln(1+\alpha)/\alpha$ falls as α rises

(c) It falls, because the folding rate $k_f U$ is halved while the hold is the same length

(d) It doubles, because aggregation is second order in the unfolded concentration

9.2 Concept check · 2 min

Folding doubles for each 10 K ($Q_{10} = 2.0$) and aggregation nearly quadruples ($Q_{10} = 3.86$). In case 5 the whole campaign delivers the most insulin with the refolding vessel at 20 °C, not at 15 and not at 30 °C. Why is the best temperature an interior one?

(a) Because the folded share of a single lot is highest at 20 °C

(b) Because the purity gate rejects lots refolded above 20 °C

(c) Because the hours a lot occupies the vessel fall faster than its folded share until a third lot fits, after which every further kelvin is yield given away

(d) Because the chiller cannot hold 15 °C once its capacity falls

9.3 Concept check · 1 min

In case 8 the 22 h harvest carries 11 % more precursor than the 16 h harvest and the campaign isolates about half as much insulin. What is the direct cause?

(a) The extra precursor ages in the hold tank, and all of it expires

(b) The larger lots fold worse, because a bigger charge means a higher unfolded concentration

(c) The extra host-cell protein pushes both lots past the 1000 ppm limit

(d) The six extra fermentation hours are six fewer recovery hours, so the train finishes one lot instead of two and the rest is still inside it at the horizon

9.4 Theory and calculation · 6 min

A lot is charged as a bolus to 2.0 g/L at 15 °C and pH 10.6, where $k_{f,15} = 4.61 \times 10^{-5}\,\mathrm{s^{-1}}$ and $k_{a,15} = 2.21 \times 10^{-5}\,\mathrm{m^3/(kg\,s)}$ (Table 9.2).

(a) Compute $\alpha = k_a U_0 / k_f$ for this charge, remembering that 1 g/L is 1 kg/m^3. *Checkpoint:* $\alpha = 0.96$

(b) Compute the folded share the charge would reach if it were left until every unfolded molecule had gone one way or the other. *Checkpoint:* 70.1 %, to within 0.1 point

(c) Case 4 measures 67.7 % folded for exactly this charge. Explain in two sentences why the measured share is below the value of part (b), and say which way the gap moves if the transfer rule waits longer.

9.5 Theory and calculation · 5 min

The temperature dependence of the two constants is Equation (9.5), with $\theta_f = 0.0693\,\mathrm{K^{-1}}$ and $\theta_a = 0.1351\,\mathrm{K^{-1}}$.

(a) Compute k_f and k_a at 25 °C from their values at 15 °C given in the previous exercise. *Checkpoint:* $k_f = 9.22 \times 10^{-5}\,\mathrm{s^{-1}}$, $k_a = 8.54 \times 10^{-5}\,\mathrm{m^3/(kg\,s)}$

(b) Compute α and the folded share of a bolus to 1.5 g/L at that temperature, and compare with the 75.3 % the same charge reaches at 15 °C. *Checkpoint:* 62.7 %, to within 0.2 point

(c) A colleague proposes running every refolding vessel 10 K warmer because the reaction then takes half as long. State in one sentence what that argument leaves out.

9.6 Theory and calculation · 6 min

The scheduler's feedforward is the first line of Equation (9.11), with $C_{sol} = 10\,\mathrm{g/L}$ the solubilized concentration the controller assumes.

(a) Compute $\hat{k}_f$ and $\hat{k}_a$ at 20 °C, where the suggested settings hold the vessel. *Checkpoint:* $0.235\,\mathrm{h^{-1}}$ and 0.157 L/(g h)

(b) The vessel holds 500 L part way through a charge and the ceiling is $U_{tgt} = 0.9\,\mathrm{g/L}$. Compute the feedforward charge rate F_{FF} in L/h. *Checkpoint:* 16.9 L/h, to within 0.2 L/h

(c) The PI trim is bounded to $[-F_{FF},\ F_{max} - F_{FF}]$. Name two errors in the feedforward that the trim has to absorb, and say which of them it cannot see.

9.7 App challenge · 8 min

Open scenario **4** · **A bolus charge against a fed one** at the default seed, on the Refolding view, with the speed at 7200× (the whole 72 h campaign takes 36 s).

(a) Run the shipped case, in which *Charge strategy* is Bolus and *Charge concentration* is 2.0 g/L, until the first lot has folded and left the vessel. Read the folded and aggregated shares and the hours in the vessel from the LAST CHARGE line of the refolding vessel panel. *Checkpoint:* 67.7 and 30.0 % in 18.8 h, each to within 0.3

(b) Press Reset, set *Charge strategy* to Scheduler on $\hat{U}$ and the *Unfolded ceiling* $\hat{U}$ to 0.35 g/L, and run the same lot again. *Checkpoint:* 84.7 and 13.1 % in 39.0 h

(c) Let each of the two runs reach the 72 h horizon and read the *Insulin accepted* and the *expired* figures on the scorecard. Say in one sentence which charge the plant should buy. *Checkpoint:* bolus 112.7 g accepted and nothing expired; scheduler nothing accepted and 665 g expired, both to within 1 g

9.8 App challenge · 9 min

Open scenario **7** · **The conversion endpoint behind a delayed assay** at the default seed, with the speed at 7200×. Everything is at its shipped value except the *Stop the conversion* rule on the Supervisor card. Each run needs about half a minute to isolate the first lot.

(a) Run with the shipped rule, After a fixed time (16 h). Read the molar conversion from the scorecard tile *Last lot: folded · converted*, and the insulin and purity of the lot from the LOTS ISOLATED list on the Recovery view. *Checkpoint:* 57.3 % of molecules, 110.3 g, 95.76 %, each to within 0.2

(b) Reset and repeat with When the assay falls, then with At the predicted peak. *Checkpoint:* 71.7 and 72.5 %; 137.9 and 139.4 g

(c) The true insulin peak of this lot is 72.5 % of the molecules at 7.64 h of reaction. The assay rule stops at 9.01 h and the forecast at 7.51 h. Explain in three sentences why the assay rule is late by construction and why being early costs the forecast almost nothing.

9.9 App challenge · 8 min

Open the **Challenge** · **The most insulin a campaign can deliver** at the default seed, with the speed at 7200×, and run each configuration for the full 72 h (36 s of wall clock). Enter every setting before pressing Run.

(a) Run the campaign as shipped and read the *Accepted insulin rate* tile and the *worst lot purity · host-cell protein* tile. Say in one sentence why the rate is what it is. *Checkpoint:* 0.00 g/h; both lots rejected near 1670 ppm

(b) Reset and enter the suggested settings: *Wash intensity* 3, *Charge strategy* Scheduler on $\hat{U}$, *Unfolded ceiling* 0.90 g/L, *Charge concentration* 1.5 g/L, *Refolding temperature* 20 °C, *Conversion temperature* 20 °C, and both supervisory rules on their model-based settings (When the estimate says it is done, At the predicted peak). Read the rate and the accepted mass. *Checkpoint:* 5.56 g/h and 401 g in three lots, to within 0.02 g/h

(c) Change only *Charge strategy* to Bolus, Reset and run again. Read the rate, and state what the difference between (b) and (c) is worth against the 0.3 pH and 1.5 °C tolerances the challenge also imposes. *Checkpoint:* 5.47 g/h

9.10 Application area · 3 min

Section 9.3 assumes the refolding vessel is one well-mixed liquid.

(a) Name the three levers a plant uses to keep the unfolded concentration low where the charge enters, and say which one this model can represent. *Checkpoint:* charge point, agitation, feed rate; only the feed rate

(b) Say whether that assumption flatters the bolus or the fed charge, and why.

9.11 Application area · 4 min

A lot is accepted when its purity is at least 95.0 % and its host-cell protein at most 1000 ppm. The recipe's lots sit at 95.76 %, and they flip to rejected when the carried-over related-substances fraction c_{rel} rises from the declared 0.040 to 0.048.

(a) Say what that margin means for the conclusions drawn from cases 1, 3 and 8, and name the measurement a plant would make before trusting them.

(b) Say where real acceptance criteria come from, and what a simulated pass is not. *Checkpoint:* the product specification, the pharmacopeia and regulatory guidance; not a batch release

9.12 Review the agent · 5 min

An assistant is asked whether the plant should adopt the scheduler and writes:

> *"In the challenge the scheduler isolates 401 g against the bolus's 394 g, 5.56 g/h against 5.47 g/h, and the advantage survives every perturbation the chapter reports: aggregation 20 % faster or slower, a weaker temperature dependence of aggregation, a 20 % slower enzyme, a heavier over-processing path, an RP-HPLC assay reported 1.5 h late and a conversion assay reported 1 h late. The scheduler is therefore the robust choice, and the plant should run it at the 0.9 g/L ceiling."*

Find the flaw in this recommendation and say how you would test it on the application.