Biochemical readouts that preserve native function

Reporter storage, finite-time inference and certified recovery in a shared cofactor pool

17 September 2026

Abstract

A reporter that consumes a regenerated cofactor also stores it: while the cofactor sits in a reporter complex it is available neither to the native reaction it is meant to report on nor to the regenerating reaction that would replenish it. We show that this storage makes a reporter's effective turnover activity insufficient to determine its effect on native function. For an explicit mass-action source in which the reporter complex is retained, the complete-recovery native-output loss per reporter product equals $k(p+c)/\big(c(p+k)\big)$, with p the regeneration coefficient, k the native-use coefficient and c the reporter's catalytic release coefficient. The instantaneous-turnover idealisation predicts $k/(p+k)$; the discrepancy is the factor $1+p/c$, which is a factor of six at $p=1$, $c=0.2$. The identity is an invariant of the source rather than of a schedule: it holds for every nonnegative, time-varying association profile of finite total exposure, so no dosing or gating strategy removes it at fixed final reporter product. We then give general design inequalities that make a finite readout simultaneously preserving and discriminating, and a rational witness: over a declared uncertainty region for nine parameters, an eight-unit acquisition suppresses native flux by less than 4.73% throughout acquisition and recovery, while separating two promised functional classes with a residual detector margin of $1214759671/218587500000 > 0.0055$. Ideal shut-off is not required: if the post-acquisition association flux decays at rate γ, an explicit convolution bound converts γ into a certified recovery deadline, five time units at $\gamma \geq 21$ and 11.2 at $\gamma = 1$, while the additional cumulative loss is bounded by $\mathcal{R}_{\mathrm{extra}} \leq (k/a)(1+p/c)uC/\gamma$ and is not zero. Dropping the two-class promise, set inversion of the same envelopes returns an outer interval for the regeneration capacity and hence for the classified native flux, with no distributional assumption added. Replacing the linear rates by increasing differentiable ones yields a divided-slope sandwich for the cost per product that collapses to the linear coefficient in the low-saturation limit. A negative result limits all of this: an unobserved background consumer makes native and total cofactor use observationally indistinguishable, so reporter data alone cannot identify native function without an independent native-specific calibration. Source algebra, the finite integral account, the nonlinear sandwich, the exponential recovery certificates and the exact decision arithmetic are verified in Lean 4; existence, invariance, comparison and convergence are conventional proofs. This is a conditional methods theorem with a prospective biochemical implementation, not an empirically validated assay.

Contents

1 Introduction

1.1 The design question

An assay that reports on a reaction usually competes with it. When a reporter draws on the same cofactor pool as the process being measured, the act of observing changes the quantity observed. The practical question is not whether this happens but how much of it a given readout causes, and

whether a readout can be chosen so that the disturbance is small enough to be tolerable while the record is still informative enough to answer the biological question.

The usual summary of a reporter's demand is its effective turnover activity: how fast it converts cofactor into signal. We show that this summary is insufficient, and we identify exactly what it misses.

Consider a reduced cofactor X regenerated from its oxidised form Z, consumed by a native reaction, and also consumed by a reporter. If the reporter is idealised as an instantaneous consumer, every unit of cofactor it takes is a unit the native reaction did not get, and the accounting is transparent. A real reporter is not instantaneous: it binds the cofactor into a complex B, holds it for a while, and only then releases the products. During that residence the cofactor is denied to the native reaction *and* withheld from regeneration. Two reporters calibrated to identical effective activity, but with different catalytic release rates, therefore impose different native-output costs.

1.2 Contribution

Our contributions are the following.

1. **An exact, schedule-invariant cost identity** (Theorem 3.3). For the explicit-complex source of Section 2, and for *every* nonnegative time-varying association profile of finite total exposure, the complete-recovery native-output loss per unit of reporter product equals

$$\frac{D_\infty}{q_\infty} = \frac{k}{p+k}\Big(1+\frac{p}{c}\Big).$$

   The instantaneous-turnover idealisation predicts the same expression without the factor $1+p/c$. Because the identity is invariant under the schedule, the extra factor cannot be removed by delaying, pulsing or ramping the reporter at fixed final product; it is removed only by changing the chemistry.

2. **General design inequalities, then a rational witness** (Theorem 4.4, Theorem 5.1). We give uniform envelopes for the free-pool deficit and for the accumulated product that hold over a declared uncertainty region, and a criterion under which a finite acquisition is simultaneously preserving and discriminating. Instantiating the criterion gives a nonempty operating region with exact rational constants.

3. **Certified recovery without an ideal switch** (Theorem 6.4, Table 3). The published version of this analysis assumed that new binding stops instantaneously. We replace that assumption by a decaying association flux and convert the decay rate into a certified recovery deadline. We also bound, and show to be nonzero, the *additional* cumulative loss that residual binding causes.

4. **Continuous bounded-error inference** (Proposition 7.1). Removing the two-class promise, set inversion of the same envelopes returns an outer interval for the regeneration capacity, and hence for the classified native flux, with no distributional assumption introduced.

5. **A controlled nonlinear extension** (Theorem 8.1). With increasing differentiable regeneration and native kinetics, the cost per product is sandwiched by divided-slope bounds that collapse to the linear coefficient in the low-saturation limit.

6. **A matching negative result** (Proposition 9.1). An unobserved background consumer of the same cofactor is observationally indistinguishable from native use. No amount of reporter data removes the ambiguity; a native-specific calibration is necessary.

1.3 Relation to prior work

That a downstream load perturbs an upstream module is the content of retroactivity and insulation [3], and shared-resource competition is a recognised design constraint in synthetic circuits [11], including the observation that isolated attenuation does not by itself justify arbitrary interconnection [10]. Fast catalytic turnover with low sequestration is an established design objective [5], and weak coupling can reduce fuel consumption and retroactivity simultaneously [4]. Coupled and auxiliary enzyme assays have a classical transient theory [6, 7], and the conditions under which a Michaelis–Menten reduction is legitimate at low substrate are explicit [8, 12]. Bounded-error parameter estimation by set inversion is a standard alternative to distributional inference [9].

We claim none of these as new. What is new here is narrower and quantitative: an exact, schedule-invariant native-output account for a regenerated pool with an explicitly retained reporter complex, together with a finite-time decision region that respects a prescribed disturbance tolerance under bounded parameter, preparation and detector uncertainty, and a recovery guarantee that survives a non-instantaneous intervention. Our literature comparison is bounded and does not establish exhaustive priority.

1.4 What is proved formally, and what is not

The manuscript separates two levels of evidence throughout, and Table 6 states the split declaration by declaration. Source algebra, the finite integral account, the nonlinear divided-slope sandwich, the exponential recovery certificates and the exact classifier arithmetic are machine-checked in Lean 4 against a pinned Mathlib [2, 13]. Existence, forward invariance, scalar comparison, convergence of the recovery integrals, and every biochemical interpretation are conventional. A compiled inequality does not validate a biological reading, and a conventional proof is not thereby weaker; conflating the two is the failure mode we are trying to avoid.

Nothing here is an experimental result. All parameters are synthetic design requirements chosen to exhibit a nonempty region; their exact rational form supports transparent checking and does not make them measured kinetic constants.

2 The source, the matched reference and the measurement contract

2.1 Reactions and state

Let X and Z denote the reduced and oxidised forms of a shared cofactor, R free reporter and B the cofactor-bound reporter complex. The driven reactions are

$$Z \xrightarrow{p} X, \qquad X \xrightarrow{k} Z + P_N, \qquad X + R \xrightarrow{\alpha} B, \qquad B \xrightarrow{\beta} X + R, \qquad B \xrightarrow{c} Z + P_R + R. \tag{1}$$

Here P_N is the native product whose production rate is the function to be classified, and P_R the reporter product that is actually recorded. External substrates supply the regeneration and the consuming chemistry and are held at fixed activity; the internal cofactor cycle is not an energetically closed system. Figure 1 shows the topology.

With concentrations x, z, b, q for X, Z, B, P_R, conserved totals $x + z + b = C$ and $r + b = R_t$, and the abbreviations

$$a := p + k, \qquad h := \beta + c, \qquad u := \frac{\alpha R_t c}{h},$$

mass action gives

$$\begin{aligned} x' &= p\,(C - x - b) - kx - \alpha x\,(R_t - b) + \beta b, \\ b' &= \alpha x\,(R_t - b) - hb, \qquad q' = cb, \qquad P_N' = kx. \end{aligned} \tag{2}$$

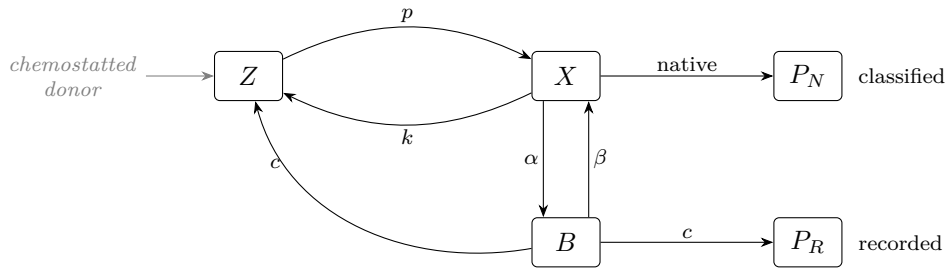

Figure 1: The source (1). Association (α) consumes free reporter R; dissociation (β) and catalytic release (c) return it. The reporter both consumes the reduced cofactor and, while the complex B persists, withholds it from the regeneration channel: the release arrow returns that cofactor to the *oxidised* pool Z, not to X. The recorded observable is the reporter product P_R; the quantity being classified is the rate of native product P_N.

All rates and totals are positive. The regeneration term acts on the *oxidised* pool $z = C - x - b$, not on x; this is what makes the complex a store rather than a mere sink, and it is the point at which an instantaneous-turnover idealisation loses information. Table 1 collects the notation.

Symbol	Meaning	Role
x, z, b	free reduced, free oxidised, bound cofactor	state
$C = x + z + b$	conserved total cofactor	invariant
$R_t = r + b$	total active reporter	invariant
p	effective regeneration coefficient, acting on z	class parameter
k	effective native-use coefficient, acting on x	calibrated
α, β, c	association, dissociation, catalytic release	reporter
$a = p + k$, $h = \beta + c$	relaxation rate, complex clearance rate	derived
$u = \alpha R_t c/h$	effective reporter activity	derived, not primitive
q	accumulated reporter product	recorded
$F = kpC/(p + k)$	stationary reference native flux	classified
η	relative preparation uncertainty	contract
g, E	detector gain, additive error bound	contract

Table 1: Notation. The effective activity u is a derived combination of α, R_t, β, c, not an independent primitive: a design that fixes the four primitives independently must check that the resulting u lies in the admitted interval.

2.2 Matched preparation and reference

Write $\bar{x} := pC/a$ for the reporter-free stationary free-pool level. The preparation satisfies

$$b(0) = q(0) = 0, \qquad x(0) = \bar{x}(1 + s), \quad |s| \le \eta < 1, \qquad x(0) \le C, \tag{3}$$

and the matched reporter-free reference starts from the same free cofactor in the same environment, so that

$$x_0(t) = \bar{x}\big(1 + s\, e^{-at}\big), \qquad F = \frac{kpC}{p + k}. \tag{4}$$

The classified function is the stationary reference flux F. Disturbance, however, is always measured against the *actual* matched trajectory $kx_0(t)$, never against an equilibrium substituted after the

experiment has begun. Define

$$e := x_0 - x, \qquad w := e - b, \qquad I(t) := \int_0^t b, \qquad D(t) := \int_0^t ke. \tag{5}$$

D is the cumulative native-output deficit; w is the deficit net of what is currently in store. Outside the regime $c \geq k$ studied below, D should be read as a signed account until its sign has been established.

2.3 Observation and intervention

The endpoint record is

$$Y = g\, q(T) + \xi, \qquad |\xi| \leq E, \tag{6}$$

with a calibrated positive gain g. The reporter product is assumed stable over the recording interval and the detector linear on the admitted range. Reporter turnovers are not identified with photon counts, and (6) is a deterministic envelope, not an error probability.

Acquisition uses a constant association coefficient on $[0, T]$. Afterwards the experiment intervenes to stop *new* binding, leaving existing complex to dissociate or turn over. Two versions of this intervention appear below: the ideal switch $\alpha \equiv 0$ for $t > T$, and the physically weaker hypothesis

$$0 \leq \alpha(T + s) \leq \alpha_0\, e^{-\gamma s}, \qquad \gamma > 0, \tag{7}$$

where α_0 is the acquisition value. In both cases the intervention must not remove bound cofactor from the system: an operation that withdraws the reporter together with its cargo changes the conservation law and invalidates every account below. This is a physical contract, and Section 11 returns to what would be required to meet it.

The binary decision of Section 5 additionally assumes that the specimen belongs to one of two separated parameter classes. Without that promise a threshold alone certifies nothing, and Section 7 supplies the continuous alternative.

3 Exact accounting of the native-output cost

3.1 Existence and the physical domain

Proposition 3.1 (Invariance and global solutions)**.** *Let $\alpha : [0, \infty) \to [0, \infty)$ be measurable and locally bounded. The set*

$$\Delta := \{(x, b) : x \geq 0,\ b \geq 0,\ x + b \leq C,\ b \leq R_t\}$$

is forward invariant for (2)*, and every solution starting in Δ exists and remains in Δ for all $t \geq 0$.*

Proof. On each face of Δ the vector field points inward. At $x = 0$, $x' = p(C - b) + \beta b \geq 0$ because $C - b \geq 0$ on Δ. At $b = 0$, $b' = \alpha x R_t \geq 0$. At $b = R_t$, $b' = -hR_t < 0$. At $z = C - x - b = 0$, $(x + b)' = pz - kx - cb = -(kx + cb) \leq 0$. The field is polynomial in the state with locally bounded measurable coefficients, hence locally Lipschitz in the state uniformly on compacts, so Carathéodory solutions exist, are unique, and cannot leave the compact set Δ; boundedness rules out finite-time blow-up. The piecewise intervention of Section 2 joins two such solutions continuously. □

3.2 The finite-time account

Proposition 3.2 (Finite-time account)**.** *For matched preparation $e(0) = b(0) = 0$ and any $t \geq 0$,*

$$e' = -ae + b' + (p + c)\, b, \qquad w' = -aw + (c - k)\, b, \tag{8}$$

and consequently

$$a\, D(t) = k\big[(p + c)I(t) + b(t) - e(t)\big]. \tag{9}$$

Proof. Subtract the reference equation $x_0' = p(C - x_0) - kx_0$ from the first line of (2). The association terms cancel against b':

$$e' = x_0' - x' = -p(x_0 - x - b) - k(x_0 - x) + \alpha x(R_t - b) - \beta b = -ae + pb + \alpha x(R_t - b) - \beta b.$$

Substituting $\alpha x(R_t - b) = b' + hb = b' + \beta b + cb$ gives the first identity in (8); subtracting b' and using $w = e - b$ gives the second. Integrating the first identity from 0 to t with $e(0) = b(0) = 0$ yields (9). □

Identity (9) is exact and finite: it contains the terminal complex $b(t)$ and the terminal free-pool deficit $e(t)$ explicitly. Omitting them understates the eventual cost of a short acquisition, because a short window leaves much of its cost in store.

3.3 The complete-recovery cost is a schedule invariant

The next statement is the central result. It requires no shut-off contract at all: only that the reporter is eventually switched off in the weak sense that the total association mass is finite.

Theorem 3.3 (Schedule-invariant cost per reporter product)**.** *Let $\alpha : [0, \infty) \to [0, \infty)$ be measurable and locally bounded, let the preparation be matched as in* (3)*, and suppose the total association mass is finite:*

$$V_0 := \int_0^\infty \alpha(t)\, x(t)\big(R_t - b(t)\big)\, \mathrm{d}t < \infty. \tag{10}$$

Then $b(t) \to 0$ and $e(t) \to 0$, the limits $q_\infty = \lim_{t\to\infty} q(t)$ and $D_\infty = \lim_{t\to\infty} D(t)$ exist and are finite, and

$$\boxed{q_\infty = c \int_0^\infty b = \frac{c\, V_0}{h}, \qquad D_\infty = \frac{k(p + c)}{ac}\, q_\infty = \frac{k}{p + k}\Big(1 + \frac{p}{c}\Big) q_\infty.} \tag{11}$$

In particular the complete-recovery native-output loss per unit of reporter product depends only on p, k and c; it is independent of β, of R_t, of C, of the preparation offset s, and of the entire association schedule.

Proof. Write $v(t) := \alpha(t)x(t)(R_t - b(t)) \geq 0$, so that $b' = v - hb$ and $b(t) = \int_0^t e^{-h(t-\sigma)} v(\sigma)\, \mathrm{d}\sigma$ since $b(0) = 0$. Integrating $b' = v - hb$ on $[0, t]$ gives $h \int_0^t b = \int_0^t v - b(t) \leq V_0$, so $\int_0^\infty b$ converges and $h \int_0^\infty b = V_0$ once we know $b(t) \to 0$. For that, fix $\varepsilon > 0$ and choose T_ε with $\int_{T_\varepsilon}^\infty v < \varepsilon$; for $t > T_\varepsilon$,

$$b(t) \leq e^{-h(t - T_\varepsilon)} \int_0^{T_\varepsilon} v + \int_{T_\varepsilon}^t v < \varepsilon + o(1),$$

so $\limsup b \leq \varepsilon$ for every ε, i.e. $b(t) \to 0$. Since $q' = cb \geq 0$ and $\int_0^\infty b < \infty$, the limit $q_\infty = c \int_0^\infty b = cV_0/h$ exists.

Next, $w' = -aw + (c - k)b$ with $w(0) = 0$ gives $w(t) = \int_0^t e^{-a(t-\sigma)}(c - k)b(\sigma)\, \mathrm{d}\sigma$, so $|w|$ is bounded by $|c - k|$ times the convolution of an integrable function with an exponential kernel; hence $w \in L^1(0, \infty)$ with $a \int_0^\infty w = (c - k) \int_0^\infty b$, and $w(t) \to 0$ by the same argument used for b. Therefore $e = w + b \to 0$ and $e \in L^1(0, \infty)$, so $D_\infty = k \int_0^\infty e$ is finite. Letting $t \to \infty$ in (9) and using $e(t), b(t) \to 0$ gives $aD_\infty = k(p + c)I_\infty$ with $I_\infty = q_\infty/c$, which is (11). □

Remark 3.4. D_∞ is a convergent integral of a *difference* of fluxes. It is never obtained by subtracting two divergent cumulative production totals, and the proof does not require the reference and the measured trajectory to be compared at any particular finite time.

3.4 What the factor means, and what evades it

In the instantaneous-turnover idealisation, in which the complex is eliminated and the reporter acts as a first-order sink of strength u, the same computation gives $D_\infty = kq_\infty/a$. The explicit complex multiplies this by

$$1 + \frac{p}{c}. \tag{12}$$

Table 2 evaluates it at $p = 1$.

catalytic release c	0.2	1	5	20
storage factor $1 + p/c$ at $p = 1$	6	2	1.2	1.05

Table 2: The complete-recovery cost per reporter product, relative to the instantaneous-turnover prediction, for four catalytic release coefficients at equal effective activity u and equal final reporter product.

Two quantities are easy to conflate here, and the distinction is the whole content of (12).

- The *mean residence of one binding episode* is $1/h = 1/(\beta + c)$. Dissociation shortens it.
- The *accumulated bound occupancy per successful reporter product* is $1/c$, because $q_\infty = c\, I_\infty$. Dissociation does not shorten it: a dissociated complex returns its cofactor but produces nothing, so the reporter must bind again.

This is why β disappears from (11) while remaining important for transients, and why increasing β at fixed u and fixed final product is not a substitute for increasing c.

Theorem 3.3 also delimits the claim. The factor is a property of *this* source at fixed final reporter product. It is evaded by changing the chemistry, not the protocol: a reporter that does not consume the cofactor, one that returns it in the reduced form, an externally replenished pool, or simply a larger c. There is no universal lower bound here for every measurement technology, and none is claimed.

4 Uniform envelopes and a design criterion

Throughout this section the reporter is calibrated in the regime

$$c \geq k, \tag{13}$$

which says the reporter releases its cargo at least as fast as the native reaction consumes free cofactor. This is the regime in which the deficit account is signed, and it is the regime the design targets.

4.1 Preservation

Proposition 4.1 (Uniform preservation)**.** *Assume* (3), (13) *and* $0 \leq \alpha(t) \leq \alpha_0$ *for all* t*, with* $u_0 := \alpha_0 R_t c/h$*. Then* $w \geq 0$*, hence* $e \geq b \geq 0$ *and* $x \leq x_0$*, and for all* $t \geq 0$

$$0 \leq b(t) \leq \frac{u_0(1+\eta)\bar{x}}{c}, \qquad 0 \leq e(t) \leq (1+\eta)\,\bar{x}\,\frac{u_0}{a}\Big(1 + \frac{p}{c}\Big), \qquad \frac{e(t)}{x_0(t)} \leq \frac{1+\eta}{1-\eta} \cdot \frac{u_0}{a}\Big(1 + \frac{p}{c}\Big). \tag{14}$$

Proof. By (8), $w' = -aw + (c-k)b$ with $w(0) = 0$, so variation of constants gives $w(t) = \int_0^t e^{-a(t-\sigma)}(c-k)b(\sigma)\,\mathrm{d}\sigma \geq 0$, using $b \geq 0$ from Proposition 3.1 and $c \geq k$. Hence $e = w + b \geq b \geq 0$ and $x = x_0 - e \leq x_0 \leq (1+\eta)\bar{x}$.

Because $R_t - b \leq R_t$ and $x \leq (1+\eta)\bar{x}$,

$$b' \leq \alpha_0 R_t (1+\eta)\bar{x} - hb,$$

and comparison with the scalar linear equation from $b(0) = 0$ gives $b \leq \alpha_0 R_t(1+\eta)\bar{x}/h = u_0(1+\eta)\bar{x}/c$, which is the first bound. Feeding it into the representation of w gives $w \leq (c-k)u_0(1+\eta)\bar{x}/(ca)$, so

$$e = w + b \leq \frac{u_0(1+\eta)\bar{x}}{c}\Big(1 + \frac{c-k}{a}\Big) = (1+\eta)\bar{x}\,\frac{u_0}{a}\Big(1 + \frac{p}{c}\Big),$$

where the last step uses $a + c - k = p + c$. Dividing by $x_0 \geq (1-\eta)\bar{x} > 0$ gives the third bound. Every comparison above is an integration of a differential inequality against its positive integrating factor; no extremum of a nonlinear trajectory at a corner of a parameter box is assumed anywhere. □

Remark 4.2. The hypothesis is $\alpha(t) \leq \alpha_0$, not $\alpha \equiv \alpha_0$ on $[0,T]$ followed by zero. Proposition 4.1 therefore covers acquisition, an ideal switch, and any fading switch (7) whose amplitude does not exceed the acquisition value. Preservation is never the difficult half of the recovery analysis.

4.2 Signal

Proposition 4.3 (Lower and upper signal envelopes)**.** *Let B_0 be any uniform upper bound for $u(1+p/c)/a$ over the admitted parameters and set $m := 1 - \eta - (1+\eta)B_0$. If $m > 0$ then $x \geq m\bar{x}$ throughout, and for constant association on $[0,T]$ with $T \geq 1/h$,*

$$\frac{u\,m\,\bar{x}}{1 + u\,m\,\bar{x}/(cR_t)}\Big(T - \frac{1}{h}\Big) \;\leq\; q(T) \;\leq\; u(1+\eta)\,\bar{x}\,T. \tag{15}$$

Proof. From (14), $x = x_0 - e \geq (1-\eta)\bar{x} - (1+\eta)\bar{x}B_0 = m\bar{x}$. Using $x \geq m\bar{x}$ and $R_t - b \geq 0$,

$$b' \geq \alpha m\bar{x}\,(R_t - b) - hb = \alpha R_t m\bar{x} - \lambda b, \qquad \lambda := h + \alpha m\bar{x} \;\geq h.$$

Comparison from $b(0) = 0$ gives $b(t) \geq (\alpha R_t m\bar{x}/\lambda)(1 - e^{-\lambda t})$, hence

$$q(T) = c\int_0^T b \;\geq\; \frac{c\,\alpha R_t m\bar{x}}{\lambda}\Big(T - \frac{1 - e^{-\lambda T}}{\lambda}\Big) \;\geq\; \frac{c\,\alpha R_t m\bar{x}}{\lambda}\Big(T - \frac{1}{h}\Big).$$

Finally $c\alpha R_t/\lambda = u/(1 + \alpha m\bar{x}/h)$ and $\alpha/h = u/(cR_t)$ give the stated left-hand side. The right-hand side is c times the uniform bound on b from (14) integrated over $[0,T]$. □

4.3 The design criterion

Proposition 4.1 and Proposition 4.3 combine into a criterion stated entirely in terms of interval endpoints. Write $\square_-$ and $\square_+$ for the smallest and largest admitted value of a quantity, and let $\bar{x}_+^{\mathrm{lo}}$ and $\bar{x}_-^{\mathrm{hi}}$ denote the largest reference level attainable in the low class and the smallest attainable in the high class; since $\bar{x} = pC/(p+k)$ increases in p and C and decreases in k, these are attained at interval endpoints.

Theorem 4.4 (Design criterion)**.** *Assume* (3), (13), *$c_- \geq k_+$, and put*

$$B_0 := \frac{u_+}{a_-}\Big(1 + \frac{p_+}{c_-}\Big), \qquad m := 1 - \eta - (1+\eta)B_0, \qquad \sigma := \frac{u_+ C_+}{c_- R_-}.$$

Let δ be the prescribed relative disturbance tolerance.

(i) (Preservation.) *If* $\frac{1+\eta}{1-\eta}B_0 \le \delta$ *and* $m > 0$*, then for every admitted source and preparation, and at every time during acquisition and any subsequent intervention with* $\alpha \le \alpha_0$*,*

$$0 \le \frac{kx_0(t) - kx(t)}{kx_0(t)} \le \delta.$$

(ii) (Discrimination.) *If moreover* $T \ge 1/h_-$ *and*

$$\underbrace{g_- u_- \, m \, \bar{x}_-^{\mathrm{hi}} \, \frac{T - 1/h_-}{1+\sigma}}_{=:L(T)} \; - \; \underbrace{g_+ u_+ (1+\eta) \, \bar{x}_+^{\mathrm{lo}} \, T}_{=:U(T)} \; > \; 2E, \tag{16}$$

then every record compatible with the low class satisfies $Y \le U(T) + E$*, every record compatible with the high class satisfies* $Y \ge L(T) - E$*, and the threshold* $\theta = \big(U(T) + L(T)\big)/2$ *classifies every admitted specimen correctly.*

Proof. Part (i) is the third bound of (14), evaluated at the interval endpoints that maximise it: $u/a \le u_+/a_-$ and $p/c \le p_+/c_-$. Part (ii) applies (15), multiplies by the gain bounds and adds the detector envelope (6). The monotonicity of $\bar{x}$ in p, C, k has the stated derivatives $Ck/(p+k)^2 > 0$, $p/(p+k) > 0$ and $-pC/(p+k)^2 < 0$, so its extrema over the box are attained at endpoints, and the saturation denominator is bounded by $1+\sigma$ using $m \le 1$ and $\bar{x} \le C$. Under (16) the two closed record intervals $[\,\cdot\,, U(T)+E]$ and $[L(T)-E,\,\cdot\,]$ are disjoint and θ lies strictly between them. □

Theorem 4.4 is the useful form for a designer: it says which quantities must be calibrated, and it exposes the trade-off. Larger T helps discrimination linearly and does not affect preservation; larger u helps discrimination but enters preservation directly; smaller c hurts both, through B_0 and through σ.

5 A robust finite-time protocol

Theorem 5.1 (Robust joint protocol)**.** *Consider* (2)–(4) *with the dimensionless intervals*

$$p \in [0.95, 1.05] \ \textit{(low class)}, \qquad p \in [1.95, 2.05] \ \textit{(high class)},$$
$$k \in [0.98, 1.02], \quad C \in [0.99, 1.01], \quad R_t \in [0.99, 1.01], \quad \eta = 0.01,$$
$$u = \frac{\alpha R_t c}{\beta + c} \in [0.079, 0.081], \quad c \in [20, 25], \quad \beta \in [1, 2], \quad g \in [0.98, 1.02], \quad E = 0.01,$$

acquisition $T = 8$ *and recovery* $\tau = 5$*. Then, for every admitted source and preparation, at every time during acquisition and recovery,*

$$0 \le \frac{kx_0(t) - kx(t)}{kx_0(t)} \le \frac{400869}{8492000} = 0.047205\ldots < 0.0473 < 0.05, \tag{17}$$

and in absolute terms $e(t) < 0.032$ *and* $b(t) < 0.0028$ *concentration units. The low class has stationary native flux below* 0.523 *and the high class above* 0.645*. With*

$$U(T) = \frac{102}{100} \cdot \frac{81}{1000} \cdot \frac{101}{100} \cdot \frac{(105/100)(101/100)}{203/100}\, T,$$
$$L(T) = \frac{98}{100} \cdot \frac{79}{1000} \cdot \frac{94}{100} \cdot \frac{(195/100)(99/100)}{297/100} \cdot \frac{T - 1/21}{1005/1000},$$

every compatible low record is at most $U(8)+0.01$ *and every compatible high record is at least* $L(8)-0.01$*, where*

$$U(8)=\frac{126420993}{362500000}=0.34874756\ldots,\qquad L(8)=\frac{56426461}{150750000}=0.37430488\ldots,$$

$$L(8)-U(8)-0.02=\frac{1214759671}{218587500000}=0.00555731\ldots>0.005.$$

The threshold $\theta=(U(8)+L(8))/2=0.36152622\ldots$ *therefore separates the two classes strictly, with per-class slack exceeding* 0.0027*. After the recovery period the outstanding native-output loss is below* $3\cdot10^{-6}$ *concentration units (Theorem 6.2).*

Proof. On this box $a\ge1.93$, $a\le3.07$, $h\ge21$, $h\le27$, $p\le2.05$ and $c\ge20\ge1.02\ge k$, so (13) holds. Take

$$B_0=\frac{0.081\,(1+2.05/20)}{1.93}=\frac{35721}{772000},$$

which dominates $u(1+p/c)/a$ over the box. Exact rational arithmetic gives

$$\frac{1.01}{0.99}B_0=\frac{400869}{8492000}<0.0473,\qquad m=1-0.01-1.01\,B_0=\frac{72820179}{77200000}>0.94,$$

which is (17) and the free-pool floor. For the absolute bounds, $\bar{x}=pC/(p+k)\le(2.05)(1.01)/(2.05+0.98)=41/60$, so by (14) $e\le1.01\cdot(41/60)\cdot B_0<0.032$ and $b\le0.081\cdot1.01\cdot(41/60)/20<0.0028$.

For discrimination, $\sigma=0.081\cdot1.01/(20\cdot0.99)=909/220000<0.005$, so the saturation denominator is at most 1.005, and $T=8\ge1/21$. The extreme reference levels are $\bar{x}_+^{\mathrm{lo}}=(1.05)(1.01)/2.03$ and $\bar{x}_-^{\mathrm{hi}}=(1.95)(0.99)/2.97$. Substituting these, together with $m\ge0.94$, into Theorem 4.4(ii) gives exactly the displayed U and L; the displayed rational values of $U(8)$, $L(8)$, their separation and the threshold are exact. The class flux bounds follow because $F=kpC/(p+k)$ increases in each of p,k,C, so $F\le(1.02)(1.05)(1.01)/2.07<0.523$ in the low class and $F\ge(0.98)(1.95)(0.99)/2.93>0.645$ in the high class. □

Remark 5.2 (What the theorem is not). The region is a *sufficient inner* region. The inequalities are strict, so the admitted design has genuine slack; but failure of the same rounded certificate at a shorter acquisition is not an impossibility proof. With the rounded constants above, the margin $L(T)-U(T)-0.02$ equals -0.00139 at $T=6$, $+0.00208$ at $T=7$ and $+0.00556$ at $T=8$: the printed certificate first closes at $T=7$. Nothing here classifies all protocols, and no optimality is claimed.

Remark 5.3 (Two contracts, not one). The rounded envelopes of Theorem 5.1 and the unrounded envelopes of Theorem 4.4 are both valid, but they define different thresholds. An observation threshold from one contract must never be combined with signal envelopes from the other.

6 Finite recovery, with and without an ideal switch

6.1 The outstanding-loss functional

Proposition 6.1 (Remainder identity)**.** *Suppose the hypotheses of Theorem 3.3 hold, and let* $t\ge0$*. Write* $V(t):=\int_t^\infty v$ *for the remaining association mass. Then the outstanding native-output loss* $\mathcal{R}(t):=D_\infty-D(t)=k\int_t^\infty e$ *satisfies*

$$\mathcal{R}(t)=\frac{k}{a}\Big[w(t)+\frac{p+c}{h}\big(b(t)+V(t)\big)\Big]=\frac{k}{a}\Big[e(t)+\frac{p-\beta}{h}\,b(t)+\frac{p+c}{h}V(t)\Big].\tag{18}$$

In the regime $c\ge k$ *every term of the first expression is nonnegative.*

Proof. Integrating $e' = -ae + b' + (p+c)b$ on $[t,\infty)$ and using $e, b \to 0$ gives $a\int_t^\infty e = e(t) - b(t) + (p+c)\int_t^\infty b = w(t) + (p+c)\int_t^\infty b$. Integrating $b' = v - hb$ on $[t,\infty)$ gives $h\int_t^\infty b = b(t) + V(t)$. Multiplying by k gives the first form; the second follows from $w = e - b$ and $h = \beta + c$, since $-b + (p+c)b/h = (p-\beta)b/h$. □

The first form of (18) is the one to use in bounds: it is a sum of nonnegative terms, so upper bounds may be substituted term by term without worrying about the sign of $p - \beta$.

6.2 Ideal switch

Theorem 6.2 (Ideal-switch recovery)**.** *Under the hypotheses of Theorem 5.1, with $\alpha \equiv 0$ for $t > T$,*

$$\mathcal{R}(T+\tau) \;\le\; \frac{k_+}{a_-}\Big[\Big(e_* + \frac{d_*\, b_*}{h_- - a_-}\Big)e^{-a_-\tau} + \frac{(p+c)_+}{h_-}\, b_*\, e^{-h_-\tau}\Big],$$

where $e_* = 0.032$, $b_* = 0.0028$, $d_* = c_+ - k_- = 24.02$, $a_- = 1.93$, $h_- = 21$ and $(p+c)_+ = 27.05$. At $\tau = 5$ this is below $1.5 \cdot 10^{-6}$, hence below $3 \cdot 10^{-6}$, concentration units.

Proof. After the switch $V \equiv 0$ and $b(T+s) = b(T)e^{-hs} \le b_* e^{-h_- s}$. For w, $w' = -aw + (c-k)b$ gives

$$w(T+s) = w(T)e^{-as} + \int_0^s e^{-a(s-\sigma)}(c-k)\, b(T+\sigma)\,\mathrm{d}\sigma \le e_* e^{-a_- s} + d_* b_* \int_0^s e^{-a_-(s-\sigma)} e^{-h_-\sigma}\,\mathrm{d}\sigma,$$

where we used $w(T) \le e(T) \le e_*$, $a \ge a_-$ for the kernel and $h \ge h_-$ for the forcing. The remaining integral equals $(e^{-a_- s} - e^{-h_- s})/(h_- - a_-) \le e^{-a_- s}/(h_- - a_-)$. Substituting into the first form of (18) and using $k/a \le k_+/a_-$ and $(p+c)/h \le (p+c)_+/h_-$ gives the display. At $\tau = 5$, the degree-16 Taylor partial sum of $e^{9.65}$ exceeds 15208, so $e^{-9.65} < 1/14000$, and the resulting rational expression evaluates to less than $1.21 \cdot 10^{-6}$; using the single crude bound $e^{-105} \le e^{-9.65} < 1/14000$ for both exponentials still gives less than $1.48 \cdot 10^{-6}$. □

6.3 Fading switch

The ideal switch is the mathematically convenient intervention, not the physically available one. We now assume only (7). Write

$$A_* := \sup \alpha_0 C R_t = \frac{u_+ h_+ C_+}{c_-} = 0.1104435,$$

which bounds the association flux amplitude immediately after acquisition, because $v = \alpha x(R_t - b) \le \alpha_0 e^{-\gamma s} C R_t$ and $\alpha_0 = uh/(R_t c)$.

Lemma 6.3 (Two-rate convolution bound)**.** *Let $r, \rho > 0$ and let $\nu > 0$ satisfy $\nu \le r$ and $\nu < \rho$. Then for all $s \ge 0$,*

$$\int_0^s e^{-r(s-\sigma)} e^{-\rho\sigma}\,\mathrm{d}\sigma \;\le\; \frac{e^{-\nu s}}{\rho - \nu}, \qquad \int_0^s e^{-r(s-\sigma)}\sigma\, e^{-\rho\sigma}\,\mathrm{d}\sigma \;\le\; \frac{e^{-\nu s}}{(\rho-\nu)^2}.$$

The same bounds hold with the roles of r and ρ interchanged.

Proof. Since $\nu \le r$ we may dominate the kernel, $e^{-r(s-\sigma)} \le e^{-\nu(s-\sigma)}$ for $\sigma \le s$. Then $\int_0^s e^{-\nu(s-\sigma)} e^{-\rho\sigma}\,\mathrm{d}\sigma = e^{-\nu s}\int_0^s e^{-(\rho-\nu)\sigma}\,\mathrm{d}\sigma \le e^{-\nu s}/(\rho-\nu)$, and similarly $\int_0^s \sigma e^{-(\rho-\nu)\sigma}\,\mathrm{d}\sigma \le (\rho-\nu)^{-2}$. The integrand is symmetric under $(r,\sigma) \leftrightarrow (\rho, s-\sigma)$, which gives the interchanged form. □

Theorem 6.4 (Fading-switch recovery)**.** *Assume the box of Theorem 5.1 and* (7)*. Choose rates $\nu_1, \nu_0 > 0$ admissible for Lemma 6.3 in the sense that ν_1 is admissible for the pair (h_-, γ) with denominator d_1, and ν_0 is admissible for the pair (a_-, ν_1) with denominator d_0. Put*

$$\beta_* := b_* + \frac{A_*}{d_1}, \qquad B_* := e_* + \frac{d_*\,\beta_*}{d_0}.$$

Then for all $s \geq 0$,

$$b(T+s) \leq \beta_* e^{-\nu_1 s}, \qquad w(T+s) \leq B_* e^{-\nu_0 s}, \qquad V(T+s) \leq \frac{A_*}{\gamma} e^{-\gamma s},$$

and consequently

$$0 \;\leq\; \mathcal{R}(T+s) \;\leq\; \frac{k_+}{a_-}\Big[B_* e^{-\nu_0 s} + \frac{(p+c)_+}{h_-}\Big(\beta_* e^{-\nu_1 s} + \frac{A_*}{\gamma} e^{-\gamma s}\Big)\Big]. \tag{19}$$

Moreover the classifier of Theorem 5.1 and the preservation bound (17) *are unaffected: the record is taken at T, and Proposition 4.1 applies verbatim because $\alpha(T+s) \leq \alpha_0$.*

Proof. From $b' = v - hb$ and $b(T+s) \leq b(T)e^{-hs} + \int_0^s e^{-h(s-\sigma)} v(T+\sigma)\,\mathrm{d}\sigma$, with $b(T) \leq b_*$, $h \geq h_- \geq \nu_1$ and $v(T+\sigma) \leq A_* e^{-\gamma\sigma}$, Lemma 6.3 gives $b(T+s) \leq b_* e^{-\nu_1 s} + A_* e^{-\nu_1 s}/d_1 = \beta_* e^{-\nu_1 s}$. Feeding this into $w(T+s) \leq w(T)e^{-as} + d_* \int_0^s e^{-a(s-\sigma)} b(T+\sigma)\,\mathrm{d}\sigma$ with $w(T) \leq e_*$ and $a \geq a_- \geq \nu_0$, and applying Lemma 6.3 again with the pair (a_-, ν_1), gives $w(T+s) \leq B_* e^{-\nu_0 s}$. The tail mass obeys $V(T+s) \leq \int_s^\infty A_* e^{-\gamma\sigma}\,\mathrm{d}\sigma = A_* e^{-\gamma s}/\gamma$, which is finite, so (10) holds and Proposition 6.1 applies; substituting the three bounds into the first, termwise nonnegative, form of (18) gives (19). □

Corollary 6.5 (Five recovery units suffice for a fast fading switch)**.** *If $\gamma \geq 21$, then taking $\nu_1 = 12$ and $\nu_0 = 1.93$ in Theorem 6.4 gives $d_1 = 9$, $d_0 = 10.07$, $\beta_* = 0.0150667$ and $B_* = 0.0678749$, and the outstanding loss at $\tau = 5$ is below $2.4 \cdot 10^{-6}$, and below $4 \cdot 10^{-6}$ if all three exponentials are crudely replaced by $e^{-9.65} < 1/14000$. The original ideal-switch deadline of five recovery units is therefore retained.*

Table 3 and Figure 2 instantiate Theorem 6.4 across decay rates, giving in each case the smallest recovery time, to one decimal place, for which (19) certifies an outstanding loss below $3 \cdot 10^{-6}$. The deadline degrades gracefully: a switch ten times slower than the complex-clearance rate costs roughly a doubling of the recovery wait, not a loss of the guarantee.

decay rate γ	ν_0	ν_1	certified τ	bound on $\mathcal{R}(T+\tau)$
0.5	0.5	0.5	22.7	$2.89 \cdot 10^{-6}$
1	1	1	11.2	$2.89 \cdot 10^{-6}$
2	1.85	2	6.8	$2.62 \cdot 10^{-6}$
5	1.93	5	5.2	$2.50 \cdot 10^{-6}$
10	1.93	10	4.9	$2.90 \cdot 10^{-6}$
21	1.93	12	4.9	$2.80 \cdot 10^{-6}$

Table 3: Certified recovery deadlines from Theorem 6.4, for the parameter box of Theorem 5.1 and target $3 \cdot 10^{-6}$. Each row is verified in exact rational arithmetic, with every exponential replaced by a rational upper bound obtained from a truncated exponential series; each listed τ also fails the same test at $\tau - 0.1$, so the tabulated values are tight for this bound to one decimal place.

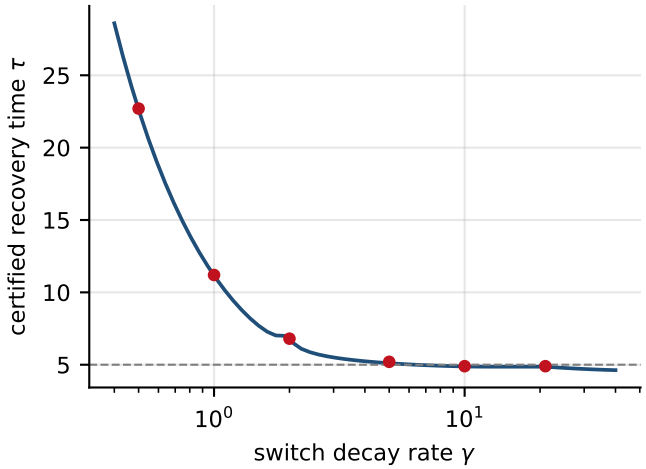

Figure 2: The certified recovery time of Theorem 6.4 against the switch decay rate γ, optimised over the admissible certified rates ν_0, ν_1. Points mark the tabulated rows; the dashed line is the five-unit deadline of Theorem 5.1. The curve flattens once γ exceeds the relaxation rate a, because the remainder is then limited by the cofactor pool's own relaxation rather than by the switch.

6.4 A small remainder is not no additional cost

Theorem 6.4 bounds what is still *outstanding* after the recovery wait. It does not say that a fading switch is as cheap as an ideal one.

Proposition 6.6 (Additional complete loss)**.** *Under* (7)*, the additional complete-recovery reporter product and native-output loss caused by residual binding, relative to an ideal switch from the same acquisition-end state, satisfy*

$$0 \le \Delta q_\infty = \frac{c}{h}\int_T^\infty v \le \frac{cA_*}{h\gamma}, \qquad 0 \le \Delta D_\infty = \frac{k}{a}\Big(1+\frac{p}{c}\Big)\Delta q_\infty \le \frac{k}{a}\Big(1+\frac{p}{c}\Big)\frac{uC}{\gamma}.$$

On the box of Theorem 5.1 the right-hand bound is 0.002270 *at* $\gamma = 21$ *and* 0.04767 *at* $\gamma = 1$*.*

Proof. Both statements are Theorem 3.3 applied to the two schedules that agree on $[0, T]$, together with $\int_T^\infty v \le A_*/\gamma$ and $\alpha_0 C R_t = uhC/c$, so that $cA_*/(h\gamma)$ simplifies to uC/γ before the cost factor is applied. □

At $\gamma = 21$ the extra complete loss bound, 0.0023, is three orders of magnitude larger than the outstanding remainder, $2.4 \cdot 10^{-6}$. The two numbers answer different questions: the remainder says the experiment may stop waiting, the extra loss says what the imperfect switch cost. Numerically, at the nominal parameters of Section 10, the realised extra complete loss is 0.00096 at $\gamma = 21$, 0.00405 at $\gamma = 5$ and 0.02041 at $\gamma = 1$, all within the proposition's bounds.

7 Continuous bounded-error inference

Theorem 5.1 answers a binary question and needs a binary promise. The same envelopes answer a continuous question with no extra assumption, by set inversion [9].

Admit p anywhere in $[0.95, 2.05]$, keeping every other interval as before. The loading and free-pool bounds are unchanged, because their proof already used the enclosing range. Writing the

p-dependence of $\bar{x}$ explicitly, the envelopes of Theorem 4.4 become

$$L_p = A_L \frac{p}{p + k_+}, \qquad U_p = A_U \frac{p}{p + k_-}, \tag{20}$$

with, at $T = 8$,

$$A_L = \frac{0.98 \cdot 0.079 \cdot 0.94 \cdot 0.99}{1.005}\Big(8 - \frac{1}{21}\Big) = \frac{47745467}{83750000} = 0.5700951\ldots,$$
$$A_U = 1.02 \cdot 0.081 \cdot 1.01^2 \cdot 8 = \frac{42140331}{62500000} = 0.6742453.$$

Proposition 7.1 (Outer set and function interval)**.** *For a record y satisfying the contract* (6)*, define*

$$\mathcal{P}(y) := \{p \in [0.95, 2.05] :\ L_p \le y + E \text{ and } U_p \ge y - E\}. \tag{21}$$

Then the true p lies in $\mathcal{P}(y)$. Both envelopes in (20) *are strictly increasing in p, so whenever $0 < y - E < A_U$ and $0 \le y + E < A_L$ the set is the interval $[p_-, p_+] \cap [0.95, 2.05]$ with*

$$p_- = \frac{k_-\,(y - E)}{A_U - y + E}, \qquad p_+ = \frac{k_+\,(y + E)}{A_L - y - E}. \tag{22}$$

Because $F = kpC/(p + k)$ is increasing in each of p, k, C, the classified function obeys

$$\frac{k_- C_-\, p_-}{p_- + k_-} \ \le\ F\ \le\ \frac{k_+ C_+\, p_+}{p_+ + k_+}. \tag{23}$$

Proof. Membership is immediate: $g\,q(T) \in [L_p, U_p]$ and $|Y - gq(T)| \le E$. For the inversion, $L_p \le y + E$ reads $A_L p \le (y + E)(p + k_+)$, i.e. $p\,(A_L - y - E) \le k_+(y + E)$, which gives the upper endpoint when $A_L > y + E$; the lower endpoint is the same computation applied to $U_p \ge y - E$. Outside those ranges the defining inequalities in (21) are used directly: a nonpositive lower requirement imposes no lower constraint, and a record at or beyond the upper asymptote makes the set empty. □

Remark 7.2 (Outer, not identified). $\mathcal{P}(y)$ is an *outer* enclosure. Because L_p and U_p are sufficient envelopes rather than the exact attainable range of $gq(T)$, the set may contain values of p that no admitted source actually realises. An empty $\mathcal{P}(y)$ therefore certifies incompatibility with the declared contract; a nonempty $\mathcal{P}(y)$ does not certify that a compatible source exists.

Example 7.3. Two records, and the intervals they support:

record y	outer set for p	outer set for F
0.30549	[0.9500, 1.2639]	[0.47756, 0.57011]
0.41151	[1.4427, 2.0500]	[0.57775, 0.68792]

The first record excludes the high class, the second excludes the low class, and neither needs the binary promise to do so. The intervals are wider than the binary conclusions of Theorem 5.1, which is the appropriate price for the weaker premise. A record at the binary threshold $\theta = 0.36153$ yields the outer set $p \in [1.0675, 1.9084]$, which meets *neither* promised class: this is exactly what strict separation means, expressed in the continuous contract.

A decision rule built on (23) needs no point estimate: report the functional verdict when the entire interval lies on one side of the biological threshold, and report *unresolved* otherwise. To upgrade such an interval to a frequentist confidence statement one would have to supply a calibration event with a stated coverage probability and prove the deterministic implication on that event; to claim posterior contraction one would need a stochastic observation model, an identifiable parameter class and an asymptotic regime. Section 9 shows why the identifiability requirement is not a formality.

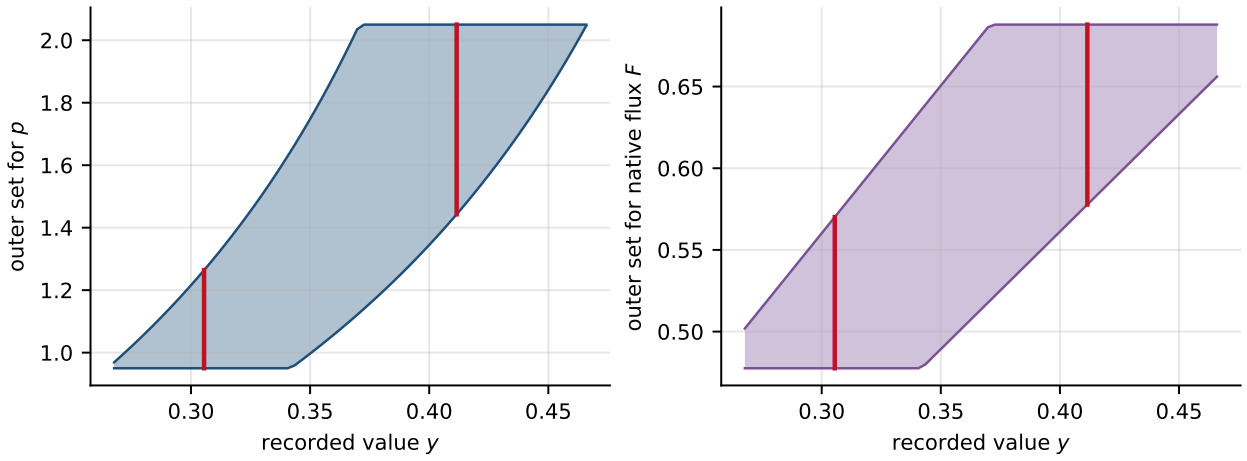

Figure 3: Outer sets from Proposition 7.1 as a function of the record y: left, for the regeneration capacity p; right, for the classified stationary native flux F. Red segments mark the two records of Example 7.3. The sets flatten against the admitted range $[0.95, 2.05]$ at both ends, where the prior physical interval, not the record, is binding. Records outside $[0.2650, 0.4661]$ are incompatible with the declared contract at every admitted p, and the plotted band stops there. A functional threshold may be decided whenever the whole interval lies on one side of it; otherwise the correct report is *unresolved*.

8 Beyond linear native and regeneration kinetics

The effective first-order forms of the native and regeneration channels are a calibration premise, not a theorem. This section replaces them by increasing differentiable functions and asks what survives.

Let regeneration be $g(z)$ and native consumption $n(x)$, both increasing and differentiable on the relevant compact range, with the reporter mechanism and the conservation law unchanged. Using $-b' - cb = -\alpha x(R_t - b) + \beta b$, the source and reference read

$$x' = g(C - x - b) - n(x) - b' - cb, \qquad x_0' = g(C - x_0) - n(x_0). \tag{24}$$

The native deficit is now $D_N(t) = \int_0^t [n(x_0) - n(x)]$, which is not $k \int e$ unless n is linear. Put $z_0 = C - x_0$ and $z = C - x - b = z_0 + w$, and define the divided slopes

$$g(z) - g(z_0) = G(t)\, w, \qquad n(x_0) - n(x) = K(t)\, e, \tag{25}$$

taking the derivative at coincident arguments. Subtracting (24) gives the exact nonlinear replacement for (8):

$$e' = -(G + K)\, e + b' + (G + c)\, b, \qquad w' = -(G + K)\, w + (c - K)\, b. \tag{26}$$

Theorem 8.1 (Divided-slope sandwich)**.** *Suppose that along the matched trajectories*

$$0 \le G_- \le G(t) \le G_+, \qquad 0 < K_- \le K(t) \le K_+, \qquad c \ge K_+,$$

that the total association mass (10) *is finite, and that* $w(0) = 0$*. Then* $w \ge 0$*,* $e = w + b \ge 0$*, and with* $I = \int_0^\infty b$*,* $W = \int_0^\infty w$ *and* $q_\infty = cI > 0$*,*

$$\frac{K_-(G_+ + c)}{c\,(G_+ + K_+)} \;\le\; \frac{D_{N,\infty}}{q_\infty} \;\le\; \frac{K_+(G_- + c)}{c\,(G_- + K_-)}. \tag{27}$$

With constant slopes $G \equiv p$ *and* $K \equiv k$ *both sides equal* $\frac{k}{p+k}(1 + \frac{p}{c})$*, recovering Theorem 3.3.*

Proof. Positivity of w follows from the second equation of (26) by variation of constants, exactly as in Proposition 4.1, using $c \geq K_+ \geq K$. Integrating that equation over $[0,\infty)$ with $w(0) = w(\infty) = 0$ gives the balance $\int_0^\infty (G+K)\, w = \int_0^\infty (c-K)\, b$, whence

$$(G_- + K_-)\, W \;\leq\; \int_0^\infty (G+K) w \;=\; \int_0^\infty (c-K) b \;\leq\; (c-K_-)\, I,$$

and symmetrically $(G_+ + K_+)W \geq (c-K_+)I$. Since $D_{N,\infty} = \int_0^\infty K\,(w+b)$ lies between $K_-(W+I)$ and $K_+(W+I)$,

$$D_{N,\infty} \;\geq\; K_-\Big(\frac{c-K_+}{G_+ + K_+} + 1\Big) I = \frac{K_-(G_+ + c)}{G_+ + K_+}\, I, \qquad D_{N,\infty} \;\leq\; \frac{K_+(G_- + c)}{G_- + K_-}\, I.$$

Dividing by $q_\infty = cI > 0$ gives (27). Setting $G_\pm = p$, $K_\pm = k$ collapses both sides to $k(p+c)/(c(p+k))$. □

Example 8.2 (Reduced Michaelis–Menten rates)**.** For $g(z) = V_g z/(K_g + z)$ the divided slope is exactly $G = V_g K_g/((K_g + z)(K_g + z_0))$, which on $0 \leq z, z_0 \leq C$ lies between $V_g K_g/(K_g + C)^2$ and V_g/K_g; the same formula applies to K for $n(x) = V_n x/(K_n + x)$. Writing $p = V_g/K_g$ and $k = V_n/K_n$, these slope ranges contract to $\{p\}$ and $\{k\}$ as C/K_g and C/K_n tend to zero at fixed p, k, so by continuity (27) converges to $\frac{k}{p+k}(1+\frac{p}{c})$. Concretely, at $V_g = V_n = K_g = K_n = 10$, $C = R_t = 1$, $c = 20$, $\beta = 1$, the sandwich is $[0.4339,\ 0.6300]$ around the linear coefficient 0.525; multiplying K_g and K_n by 10, 100 and 1000 narrows it to $[0.5147, 0.5351]$, $[0.5240, 0.5260]$ and $[0.5249, 0.5251]$. Direct integration of the nonlinear source at the first parameter set gives a complete-loss-per-product ratio of 0.5236, inside the proved sandwich.

Remark 8.3 (Scope). Two cautions. First, low substrate concentration and quasi-steady-state elimination are distinct approximations, and the conditions for the former do not license the latter [8, 12]; a chosen enzyme must justify its actual reduction rather than appeal to a small pool. Second, (27) assumes monotone one-variable rate laws. Product inhibition, allostery, additional cofactor-binding enzymes and multi-substrate kinetics can break the monotonicity, the pool accounting or the slope bounds; in that case the corresponding states must be retained, or the discrepancy proved or calibrated. The numerical operating window of Theorem 5.1 is *not* automatically inherited by the nonlinear model.

9 What the reporter cannot identify

Proposition 9.1 (Native/background alias)**.** *Add an unobserved background channel $X \to Z + P_B$ with rate ℓx, so that the free-pool equation becomes $x' = p(C-x-b) - (k+\ell)x - \alpha x(R_t - b) + \beta b$ while the native output remains $P_N' = kx$. Then the observed trajectories (x, b, q) depend on k and ℓ only through $k + \ell$. In particular $(p, k, \ell) = (1, \frac{1}{2}, \frac{3}{2})$ and $(1, \frac{3}{2}, \frac{1}{2})$ produce identical deterministic records under every common prescribed association schedule and matched preparation, while their stationary native fluxes are $C/6$ and $C/2$. Consequently, for any binary classifier and any observation-noise kernel applied to the common record law, the two class-conditional error probabilities sum to* 1.

Proof. The cofactor and complex equations involve k and ℓ only via the sum, and so does the matched reference x_0 and the stationary level $\bar{x} = pC/(p+k+\ell)$. Uniqueness of solutions (Proposition 3.1) then gives identical (x, b, q) for the two parameter triples. The stationary native fluxes are $kpC/(p+k+\ell)$, namely $C/6$ and $C/2$. Applying a common noise kernel to identical deterministic records gives identical record laws; if a classifier returns "high" with probability r on that law, its error probabilities on the two classes are r and $1-r$. □

Replication does not help: the two hypotheses induce the same law, so no number of repeated measurements of the same record separates them. The positive results of this paper avoid the alias by assumption, not by analysis: Theorem 5.1 requires that there is no unmodelled background consumer and that k has been calibrated independently and narrowly, so that the two promised classes differ in regeneration capacity. What the protocol then classifies is that capacity and the stationary native function it supports. It does not jointly identify unconstrained p and k from reporter data, and a detector calibration alone does not remove the ambiguity. A native-product reference measurement on matched preparations is the relevant additional experiment.

10 Numerical study

All computations in this section are deterministic solutions of the literal source (2). They test the implementation and illustrate the mechanism; the uniform statements are supplied by the proofs, not by the integrations.

Matched reporter comparison. Table 4 compares three protocols at a common cofactor pool, common reporter stock, common effective activity $u = 0.08$, common deadline $T = 8$, common recovery and common detector error. Only the catalytic release coefficient and the activation window differ.

Protocol	low signal	low peak suppression	low complete loss	nominal gap*
Slow catalyst, $c = 1$, $[0, 8]$	0.27040	7.161%	0.28830	0.06639
Fast catalyst, $c = 20$, $[0, 8]$	0.30549	4.023%	0.16134	0.08602
Late fast pulse, $c = 20$, $[4, 8]$	0.15222	4.022%	0.08087	0.03255

Table 4: Numerical solutions at $k = C = R_t = \beta = 1$, $u = 0.08$, stationary preparation, low $p = 1$ and high $p = 2$. *high minus low signal minus 0.02, without gain or parameter uncertainty; the robust statement is Theorem 5.1, not this column. Quantities are dimensionless except where a percentage is shown. No claim of schedule optimality is made.

The slow reporter separates the nominal pair but violates the 5% preservation requirement; the fast reporter meets both, and does so uniformly over the declared uncertainty by Theorem 5.1. Their complete native losses differ by a factor of 1.787 at equal effective activity, and Theorem 3.3 accounts for that number exactly. At equal final reporter product the predicted ratio is $(1 + 1/1)/(1 + 1/20) = 2/1.05 = 1.90476$; the realised complete products are 0.28830 and 0.30731, and $1.90476 \times (0.28830/0.30731) = 1.787$. A reporter selected on effective activity alone would not distinguish the two.

The delayed pulse halves both the recorded product and the complete loss while leaving the peak suppression essentially unchanged (4.022% versus 4.023%): lowering exposure improves the integrated cost but does not defeat the fixed-product identity, and it shrinks the decision margin.

Interior challenges. Sixteen reproducible interior draws from the declared box (eight per class, seed 17092026) were integrated through acquisition and recovery. The maximum residual in the exact account (9) was $1.1 \cdot 10^{-12}$; every draw satisfied the proved suppression bound (17) and the proved recovery remainder. These figures check the implementation, not the continuum coverage, which is the business of Section 4.

Schedule invariance. Four schedules — constant on $[0, 8]$; constant on $[4, 8]$; fading with $\gamma = 5$; delayed and fading with $\gamma = 1$ — give complete-recovery cost ratios D_∞/q_∞ agreeing to 10^{-12} with the predicted value 0.525, as Theorem 3.3 requires.

Fading switches. At the nominal parameters, the complete loss is 0.161338 for an ideal switch and 0.162298, 0.165388 and 0.181518 for fading switches with $\gamma = 21, 5, 1$. The increments, 0.00096, 0.00405 and 0.02041, respect Proposition 6.6.

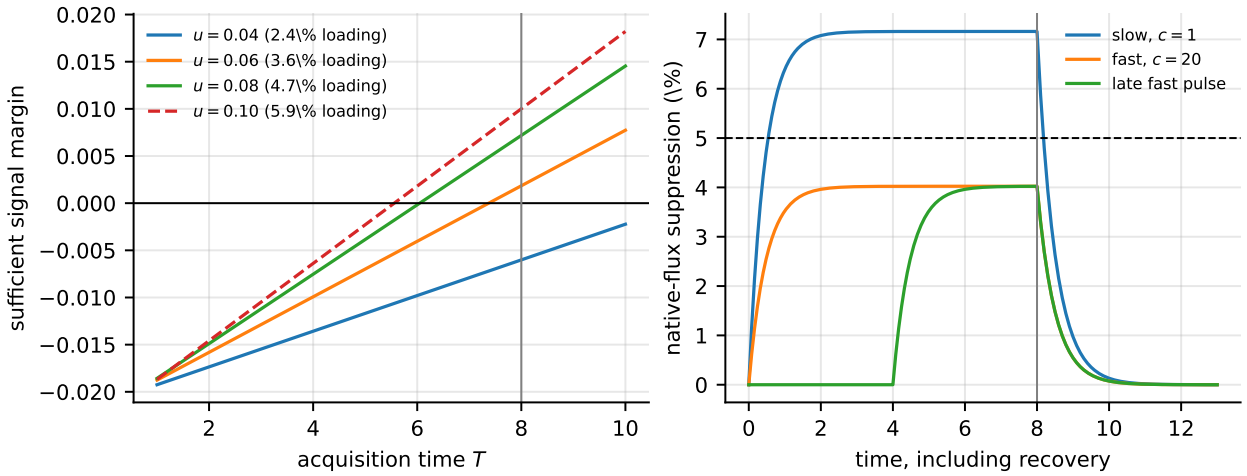

Figure 4: Left: the sufficient signal margin of Theorem 4.4, using unrounded envelopes, as the acquisition time varies at four effective activities; the loading percentage in each label is the corresponding uniform suppression bound, and the dashed curve exceeds the 5% allowance, so preservation and discrimination must be read together. Right: numerical native-flux suppression for the three matched protocols; binding stops at time 8 and the complex is retained through recovery. These are model calculations, not experimental data.

11 Physical translation, calibration and limits

A candidate realisation combines glucose-6-phosphate dehydrogenase regenerating NADPH [14], an NADPH-consuming IDH1 R132H reaction producing 2-hydroxyglutarate as the native product, and a diaphorase/resazurin readout [1]. "Native" here designates the target reaction in the constructed assay, not wild-type physiology. The cited sources support the component directions and the reporter practice; they do not establish this combination, its kinetic intervals, its detector envelope or its switch.

With illustrative scales of 1 μM, 60 s and 100 μL, Table 5 translates the dimensionless requirements.

The nominal low-class complete loss, 0.16134 in dimensionless units, corresponds to 0.16 μM or 16.1 pmol in the illustrative volume. Over the 13-unit horizon the conservative donor-consumption cap is $p_+C_+(T + \tau) = 26.92\,\mu$M, with target and reporter cosubstrate caps of 13.39 and 0.67 μM; these are finite material accounts, not stock recipes, because the supply must additionally keep the calibrated rate laws valid. A longer horizon or a slower switch requires them to be recomputed.

Three implementation frontiers remain, in decreasing order of difficulty.

Quantity	Dimensionless	Dimensional	Measurement required
total cofactor C	0.99–1.01	0.99–1.01 μM	total-pool recovery calibration
regeneration p	0.95–1.05 or 1.95–2.05	0.0158–0.0175 or 0.0325–0.0342 s^{-1}	reporter-free fit on a held-out range
native use k	0.98–1.02	0.0163–0.0170 s^{-1}	native-specific calibration; see Section 9
effective activity u	0.079–0.081	0.001317–0.001350 s^{-1}	insufficient alone: also calibrate c, β, R_t
catalytic release c	20–25	0.3333–0.4167 s^{-1}	turnover/residence experiment
dissociation β	1–2	0.0167–0.0333 s^{-1}	pulse–chase calibration
acquisition, recovery	8, 5	480 s, 300 s	—
switch decay γ	≥ 21 (Cor. 6.5)	time constant ≤ 2.86 s	residual-association measurement
detector error E	0.01	$\pm 0.01\,\mu$M product-equivalent	envelope coverage, not a 95% label

Table 5: Illustrative dimensional translation. These are design requirements to be met and checked, not measured kinetic constants taken from the literature.

1. **The intervention.** Stopping new binding without removing bound cofactor, and without perturbing the native reaction, is the hardest physical requirement. Corollary 6.5 weakens "instantaneous" to "decay constant at most 2.86 s" at the illustrative time scale, and Table 3 prices slower switches in recovery time, which converts an idealisation into a measurable specification: bound the residual association flux, do not assert $\alpha = 0$. Selective trapping of free reporter is a candidate operation, but trapping changes the free-reporter inventory and dissociation replenishes it, and dilution changes more than association alone. Enzyme removal that discards the complex violates the contract outright.

2. **The single-complex reduction.** Multi-step diaphorase chemistry and flavin intermediates storing reducing equivalents are not automatically the same state as cofactor bound in B; the chemical mapping must conserve the correct inventory, or the residual must be proved or calibrated.

3. **The rate laws.** The effective first-order native and regeneration channels must be validated on the actual cofactor range, or replaced using Section 8, which changes the constants.

A minimal prospective validation sequence follows from the theorems rather than from convention: (i) measure active reporter concentration and the binding/dissociation/release behaviour separately, well enough to resolve storage, because u alone does not determine disturbance; (ii) bound native and background consumption independently using a native-product reference on matched preparations, because of Proposition 9.1; (iii) establish the preparation and detector envelopes without tuning the threshold on the same observations used to assess it; (iv) compare theory-selected reporter architectures at common pool, time and detection budgets; (v) record native output through acquisition *and* recovery, including anything released or removed by the switch; (vi) evaluate the

fixed decision rule, or the interval rule of Section 7, on held-out preparations, reporting *unresolved* and *incompatible* outcomes as such. We deliberately do not prescribe a replicate count: the present theory is deterministic, and a sample size requires a statistical model that has not been supplied.

12 Evidence scope

Table 6 states, declaration by declaration, what has been machine-checked and what has not. Compilation used Lean 4.30.0 against the repository's pinned Mathlib with warnings treated as errors; both modules exit zero with empty output and no proof holes, and every exported declaration depends only on `propext`, `Classical.choice` and `Quot.sound`.

Result	Status	Declaration or location
Source subtraction, stored-deficit equation	Lean	`source_deficit_identity`, `stored_deficit_identity`
Finite-time account (9) from the derivative premise	Lean	`finite_account`
Cost coefficient from $q = cI$ and $aD = k(p + c)I$	Lean	`turnover_cost`
Both forms of the remainder (18), incl. residual mass	Lean	`remainder_forms`, `residual_remainder`
Divided-slope sandwich (27)	Lean	`slope_sandwich`
Exponential bound $e^{9.65} > 14000$ from a Taylor partial sum	Lean	`exp_965_gt`, `exp_neg_965_lt`
Ideal-switch and fading-switch recovery certificates	Lean	`ideal_recovery_tail`, `fading_recovery_tail`
Additional complete-loss bound of Prop. 6.6 at $\gamma = 21$	Lean	`fading_extra_loss`
Absolute preservation envelopes, class flux separation	Lean	`absolute_envelopes`, `class_flux_bounds`
Set-inversion endpoints (22), flux monotonicity	Lean	`outer_upper`, `outer_lower`, `flux_monotone`
Loading, free-pool floor and classifier arithmetic	Lean	`parameter_certificate`, `decision_certificate`
Preservation and classification from the envelopes	Lean	`preparation_preservation`, `joint_protocol`
Alias source equality and error-sum arithmetic	Lean	`merged_source`, `alias_fluxes`, `equal_law_error`
Invariance and global existence	conventional	Prop. 3.1
Convergence $b, e \to 0$ and integrability under (10)	conventional	Thm. 3.3
Scalar comparison giving (14) and (15)	conventional	Props. 4.1, 4.3
Convolution bounds and the deadline table	conventional	Lem. 6.3, Thm. 6.4, Tab. 3
Michaelis–Menten slope ranges and the low-saturation limit	conventional	Ex. 8.2
Uniqueness step in the alias argument	conventional	Prop. 9.1
Every biochemical mapping and calibration claim	premise	Sec. 11

Table 6: The formal/conventional boundary. A compiled inequality does not validate a biological interpretation, and a conventional proof is not thereby weaker. The final compiled implication takes the conventionally proved source envelopes as premises; it is not a formalisation of the ODE theory.

A companion script recomputes every constant printed in this paper. Exact values are recomputed in rational arithmetic; every exponential appearing in a bound is replaced by a rational upper bound derived from a truncated exponential series, which has only positive terms and is therefore a valid

lower bound for the exponential itself. The numerical integrations of Section 10 are independent checks of the inequalities and are never used to justify them.

13 Discussion

The defensible claim of this paper is narrow and, we think, useful: a physically explicit reporter imposes a native-output cost that its effective turnover activity does not reveal, that cost is computable and is an invariant of the source rather than of the protocol, and it can be folded into a finite, uncertainty-aware measurement design.

Several negative statements are part of the contribution and should not be dropped when the positive ones are quoted. The storage factor is not a universal law about measurement; it is a property of a regenerated pool with a cofactor-sequestering reporter at fixed final product, and changing the chemistry evades it. The operating region of Theorem 5.1 is sufficient, not necessary, and its failure at a shorter acquisition is not an impossibility theorem. The recovery guarantees are conditional on an intervention that preserves cofactor; Theorem 6.4 makes the requirement measurable but does not supply the reagent. The bounded-error intervals are deterministic enclosures, not confidence sets. And Proposition 9.1 is a genuine obstruction: without an independent native-specific measurement, reporter data cannot identify native function at all, however precise the detector.

What changes for an experimenter is the calibration list. Catalytic residence must be measured, not inferred from activity; the reporter architecture, not merely its dose, is a design variable; the intervention needs a residual association budget rather than an assertion that binding has stopped; and the native channel needs its own reference measurement. Whether a chosen chemistry and a chosen switching operation can meet those premises is an empirical question that this paper states precisely and does not answer.

A Exact constants

Quantity	Exact value	Decimal
loading bound B_0	35721/772000	$0.046270725\ldots$
relative suppression	400869/8492000	$0.047205488\ldots$
free-pool floor m	72820179/77200000	$0.943266567\ldots$
saturation σ	909/220000	$0.004131818\ldots$
absolute deficit bound e_*	49306887/1544000000	$0.031934512\ldots$
absolute complex bound b_*	111807/40000000	0.002795175
association amplitude A_*	1104435/10000000	0.1104435
$U(8)$	126420993/362500000	$0.348747567\ldots$
$L(8)$	56426461/150750000	$0.374304882\ldots$
separation $L(8) - U(8) - 0.02$	1214759671/218587500000	$0.005557315\ldots$
threshold θ	$(U(8) + L(8))/2$	$0.361526225\ldots$
low-class flux ceiling	10403/19906	$0.522565217\ldots$
high-class flux floor	18915/29293	$0.645696246\ldots$
A_L	47745467/83750000	$0.570095128\ldots$
A_U	42140331/62500000	0.674245296
degree-16 Taylor sum for $e^{9.65}$	—	$15208.05\ldots > 14000$

B Reproducing the results

The manuscript source package contains the two Lean modules, their strict verification receipts, and a single script that regenerates every figure, every generated table row and the full list of numerical assertions. The script exits non-zero on any failed assertion. The Lean modules are compiled with the repository's pinned toolchain and manifest; the receipts record the command, the source hashes, the elaborated statements and the axiom list of each exported declaration.

References

[1] Mindy I. Davis, Min Shen, Anton Simeonov, and Matthew D. Hall. Diaphorase coupling protocols for red-shifting dehydrogenase assays. *ASSAY and Drug Development Technologies*, 14(3):207–212, 2016. doi: 10.1089/adt.2016.706.

[2] Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming language. In *Automated Deduction – CADE 28*, Lecture Notes in Computer Science, pages 625–635. Springer, 2021. doi: 10.1007/978-3-030-79876-5_37.

[3] Domitilla Del Vecchio, Alexander J. Ninfa, and Eduardo D. Sontag. Modular cell biology: retroactivity and insulation. *Molecular Systems Biology*, 4(1):161, 2008. doi: 10.1038/msb4100204.

[4] Abhishek Deshpande and Thomas E. Ouldridge. High rates of fuel consumption are not required by insulating motifs to suppress retroactivity in biochemical circuits. *Engineering Biology*, 1(2): 86–99, 2017. doi: 10.1049/enb.2017.0017. Preprint: arXiv:1708.01792.

[5] Abhishek Deshpande and Thomas E. Ouldridge. Optimizing enzymatic catalysts for rapid turnover of substrates with low enzyme sequestration. *Biological Cybernetics*, 114(6):653–668, 2020. doi: 10.1007/s00422-020-00846-6. Preprint: arXiv:1905.00555.

[6] John S. Easterby. Coupled enzyme assays: a general expression for the transient. *Biochimica et Biophysica Acta (BBA) – Enzymology*, 293(2):552–558, 1973. doi: 10.1016/0005-2744(73)90362-8.

[7] Justin Eilertsen and Santiago Schnell. A kinetic analysis of coupled (or auxiliary) enzyme reactions. *Bulletin of Mathematical Biology*, 80(12):3154–3183, 2018. doi: 10.1007/s11538-018-0513-4.

[8] Justin Eilertsen, Santiago Schnell, and Sebastian Walcher. The Michaelis–Menten reaction at low substrate concentrations: pseudo-first-order kinetics and conditions for timescale separation. *Bulletin of Mathematical Biology*, 86(6):68, 2024. doi: 10.1007/s11538-024-01295-z. Preprint: arXiv:2403.05507.

[9] Luc Jaulin and Eric Walter. Set inversion via interval analysis for nonlinear bounded-error estimation. *Automatica*, 29(4):1053–1064, 1993. doi: 10.1016/0005-1098(93)90106-4.

[10] Yili Qian and Domitilla Del Vecchio. Robustness of networked systems to unintended interactions with application to engineered genetic circuits. *IEEE Transactions on Control of Network Systems*, 8(4):1705–1716, 2021. doi: 10.1109/TCNS.2021.3078144. Preprint: arXiv:2009.04098.

[11] Yili Qian, Hsin-Ho Huang, José I. Jiménez, and Domitilla Del Vecchio. Resource competition shapes the response of genetic circuits. *ACS Synthetic Biology*, 6(7):1263–1272, 2017. doi: 10.1021/acssynbio.6b00361.

[12] Lee A. Segel and Marshall Slemrod. The quasi-steady-state assumption: a case study in perturbation. *SIAM Review*, 31(3):446–477, 1989. doi: 10.1137/1031091.

[13] The mathlib Community. The Lean mathematical library. In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2020)*, pages 367–381. ACM, 2020. doi: 10.1145/3372885.3373824.

[14] Aiping Zhu, Roberto Romero, and Howard R. Petty. An enzymatic fluorimetric assay for glucose-6-phosphate: application in an in vitro Warburg-like effect. *Analytical Biochemistry*, 388(1):97–101, 2009. doi: 10.1016/j.ab.2009.02.009.