Certifying fresh microbial conversion from finite challenge records: a sharp two-pool material bound

17 September 2026

Abstract

A finite record of product collected from a microbial preparation is compatible both with fresh conversion and with release of material that was already present. We study this ambiguity in a two-window material model with a mobile product pool, a convertible reserve pool, release, uptake, losses, and an intervening recovery step with pool-specific retention. Given calibrated upper bounds on initial stock B, on the pre-wash reserve J, on credited recovery input H, and retention factors $0 \le e \le s \le 1$, we prove that the fresh credited inventory $F_1 + F_2$ generated during the two windows is at least L, the maximum of five affine expressions in the two observed collections, and that L is *exactly* the minimum of $F_1 + F_2$ over the declared class: an explicit finite material schedule attains it from every admissible initial split of the stock. The lower bound, an equivalent remaining-stock form, the attaining schedule and the resulting exactness statement are verified in Lean 4 against Mathlib. An exact counterexample shows that an initial reserve ceiling cannot be substituted for the pre-wash ceiling, because uptake during the first window moves mobile product into the better retained pool; we also verify a zero-fresh history lying inside the reported output envelopes of our worked record. We then prove that a bound on the initial reserve together with a cumulative uptake cap is a valid alternative premise: the same five expressions, evaluated at $J_0 = \overline{R} + \overline{T}$, remain sound and remain exactly attained, so destructive pre-wash measurement is sufficient but not necessary. Finally we quantify what a wash buys. At a matched material source the mobile retention factor cancels from the certificate, so removing known product is not itself an improvement; the entire benefit is attenuation of first-window measurement uncertainty, worth $(s-e)\big(\epsilon_1 - (J-r)\big)$ in the regime where the reserve-sensitive expression is active. In a synthetic 10 mL example the washed certificate is 1.69 μmol against 1.52 μmol for a matched unwashed protocol, and that 0.17 μmol margin is erased by reserve slack δ or by lost productive activity: strict improvement holds exactly when the retained activity exceeds $\frac{193}{210} + \frac{17}{42}\delta$. All numerical inputs are synthetic. The results are an assay-design theory with a proposed biological application; they are not biologically calibrated, and the certificate does not assert substrate-specific atom attribution, viability, or future function.

Keywords. exact lower certificate; material balance; weak duality; assay design; identifiability; formal verification.

MSC 2020. 90C05, 92C42, 68V20, 62P10.

Contents

1 Introduction

1.1 The decision a challenge record is asked to support

A common experiment presents a microbial preparation with a substrate, collects product after a fixed window, intervenes—washes, dilutes, resuspends—and then collects product again. The question the experiment is meant to answer is whether the preparation *converted* anything: whether the collected product forces the conclusion that new material entered the credited inventory during the observed windows.

The record does not answer this by itself. Product can be exported because it was made, or because it was already present in the vessel, inside cells, on surfaces, or in the medium, and was merely mobilised. Worse, the two explanations are not separated by collecting more product: a preparation that stores mobile product during the first window arrives at the intervention with *more* explanatory material than it started with, precisely in the compartment that an intervention is least able to remove.

This paper isolates that ambiguity and answers it completely, in a deliberately small model. We fix a two-pool material accounting, declare which calibration bounds an experimenter is willing to supply, and compute the exact minimum of the fresh amount over every material history compatible with the record and those bounds. The answer is a maximum of five affine expressions. Because it is a minimum and not merely a bound, it simultaneously certifies what the record forces and identifies the information barrier: no sharper certificate exists under the same premises, and improving the result requires naming an additional measurement or a valid mechanistic restriction that excludes the extremal schedule.

1.2 Contributions

1. **An exact certificate** (§3). For a finite two-window material source with mobile product, convertible reserves, release, uptake, arbitrary losses and pool-specific recovery, the fresh credited inventory satisfies $F_1 + F_2 \geq L$, where L is the maximum of five affine expressions in the observed collections (Theorem 3.2). An equivalent remaining-stock form (Proposition 3.3) exposes the mechanism: L is the first-window deficit plus whatever second-window output cannot be explained by optimally placing the leftover stock in the better retained pool.

2. **Exactness, formally verified** (§4). An explicit finite schedule attains L, starting from *any* admissible split of the initial stock between the two pools (Theorem 4.1), so L is the least element of the set of achievable fresh totals (Corollary 4.2). The attainment argument is machine-checked, not only the inequality.

3. **The reserve premise: necessity and a substitute** (§5). A verified counterexample shows that the initial reserve ceiling cannot be substituted for the pre-wash ceiling: a history with zero fresh conversion receives a false certificate of 1.7 μmol (Theorem 5.1). Nevertheless direct destructive reserve measurement is *not* mathematically mandatory. We prove that a bound on the initial reserve together with a cumulative uptake cap gives a sound certificate with $J_0 = \overline{R} + \overline{T}$ (Theorem 5.3), and that the resulting expression is again exactly attained (Theorem 5.4). Three of the five expressions use no reserve information at all, so positive certification is sometimes possible with none.

4. **What a wash can and cannot buy** (§7). At a matched material source the mobile retention factor cancels from the certificate: removing known product removes it from the explanatory background and from the observed output in equal measure. The residual benefit is attenuation of first-window measurement uncertainty, and it is exactly $(s-e)\big(\epsilon_1-(J-r)\big)$ in the relevant regime, so reserve slack alone can consume it. In the worked synthetic example we compute the full joint budget in exact arithmetic: with reserve slack $\delta \geq 0$ and retained productive activity $a \in [0,1]$, washing strictly improves the certificate if and only if $a > \frac{193}{210} + \frac{17}{42}\delta$ (Theorem 7.4), which is already infeasible at $\delta = 0.2$.

5. **A stated evidence boundary** (§8). We separate machine-checked statements, conventional arguments, exact computations and unvalidated empirical premises, and we say which is which for every claim.

1.3 What is not claimed

The target is *fresh credited inventory*: material that entered a declared, conservatively bounded pool during the two windows. That is a weaker and more defensible object than "de novo production" of a named molecule. In particular a certificate $F_1 + F_2 \geq L > 0$ does not establish gross product-only formation if fresh material may enter convertible intermediates; it does not attribute carbon atoms to a particular substrate; it does not certify viability, growth, abundance or a per-cell rate; it does not certify recovery relative to a demonstrably positive first window; and it says nothing about function after the second window. Each of these is a different target requiring its own observation or biological bridge. A zero certificate is an *unresolved* outcome—a statement about available information—and never a certificate of death or of absent capability.

The numerical example throughout is synthetic. No biological calibration of B, J, e, s, H, ϵ_k or the activity parameter a has been performed, and the percentages we compute are properties of a declared response model, not measured properties of any wash technique or strain.

1.4 Related work

Absolute microbial load carries information that compositional profiling does not [6], but it supplies no reserve-per-cell calibration. Dynamic metabolic flux analysis infers fluxes from transient concentration records together with observability and reconciliation analysis [3], and recent enzyme-constrained variants propagate uncertainty by bootstrapping [5]. We claim neither uncertainty-aware flux inference nor exact linear programming as new. Our contribution is narrower and of a different type: an exact minimum, with its attaining history, for a small material class under explicitly declared calibration bounds. §10.1 states the precise relationship: a *valid* kinetic restriction shrinks the admissible class and can only strengthen a lower certificate, so the comparison is about assumption dependence rather than about method superiority.

For biological motivation we rely on defined-culture cross-feeding studies [1] and on strain-resolved lactate utilisation [2, 4]; §9 explains why a direct lactate challenge tests a downstream conversion rather than a complete community process, and corrects a natural but unsupported choice of reference organism.

1.5 Organisation

§2 fixes the material model and the accounting boundary. §3 proves the certificate and its closed form, §4 its exactness. §5 treats the reserve premise, its necessity and its substitute. §6 is the observation wrapper and reporting rule. §7 is the design analysis. §8 states the evidence boundary, §9 the calibration requirements and a feasibility pilot, and §10 the limitations. Appendix A lists the source modules, receipts and reproduction commands.

2 The material model

2.1 Pools, units and the accounting boundary

Fix one experimental unit: a single cultured aliquot with a preparation identifier. *All* amounts are referred to that original aliquot, never to retained volume and never to surviving cells.

Two pools are tracked.

- P, *mobile extracellular product*: the amount physically available for collection.
- R, *convertible reserves*: everything else in the declared system capable of contributing to credited output. This must include intracellular intermediates, mobilisable biomass carbon, sorbed product and relevant alternative nutrients, each entered at a conservative product-equivalent upper yield.

The credited inventory is measured in product equivalents; in the worked example, μmol of butyrate equivalents, one equivalent per four creditable carbon atoms. Declaring a pool inert because it is inconvenient is not a calibration, and conversely a conservative carbon potential may be much larger than the measured product, which makes B large and the certificate weak. That is an informative limitation of the assay, not a defect of the theorem.

Two conventions must be fixed before any calculation.

1. *The position of unreacted challenge substrate.* If it lies outside the credited pools, then the event counted as fresh conversion is exactly its movement into them. It cannot simultaneously be charged as old explanatory stock and counted again as a new positive source. A complete implementation should list every carbon-bearing input and assign it to initial stock, fresh conversion, recovery input, or another explicitly bounded source.
2. *Collection versus retention.* Q_k is material that physically leaves the boundary in window k, and the two windows are nonoverlapping. A retained endpoint pool is not also counted as exported. For a homogeneous sample of volume v from volume V at concentration c,

$$Vc = (V - v)c + vc,$$

the second term is counted once as collected and the first is carried forward with its state. Wash waste is excluded from Q_1, Q_2 in the protocol declared here and may be reported separately.

2.2 Finite source histories

Definition 2.1 (Window)**.** A *window* is a tuple $w = (p_0, r_0, F, U, T, Q, D_P, D_R) \in \mathbb{R}^8$ of nonnegative reals satisfying the availability constraints

$$U \leq r_0 + F, \qquad U + D_R \leq r_0 + F + T, \qquad T + Q + D_P \leq p_0 + U, \tag{1}$$

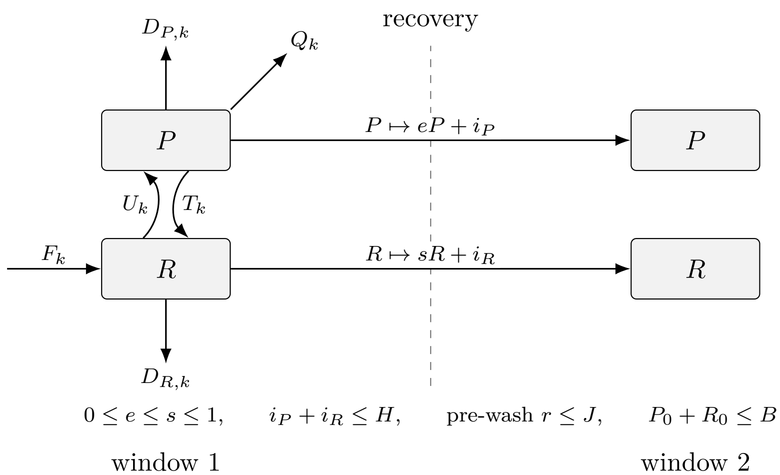

Figure 1: The declared accounts. Arrows are material transfers in a finite schedule, not fitted kinetics. Productive activity is not a material pool; its retention is a separate calibration and does not follow from e or s.

with terminal pools

$$p(w) = p_0 + U - T - Q - D_P, \qquad r(w) = r_0 + F + T - U - D_R. \tag{2}$$

Here F is fresh credited conversion, U release from reserve into mobile product, T uptake or storage from mobile product into reserve, Q collection, and D_P, D_R are unrestricted loss accounts.

Constraints (1) implement a definite event order—fresh addition, release, uptake, collection, loss—with enough material present at each step. A window is therefore an actual nonnegative schedule, not merely an assumed inequality between observables.

Lemma 2.2 (Pool nonnegativity and balance)**.** [pool_nonneg, source_balance] *For every window w, $p(w) \ge 0$ and $r(w) \ge 0$, and*

$$Q + p(w) + r(w) + D_P + D_R = p_0 + r_0 + F. \tag{3}$$

Consequently $Q + p(w) + r(w) \le p_0 + r_0 + F$.

Proof. Nonnegativity is immediate from (1) and (2): $p(w) = (p_0 + U) - (T + Q + D_P) \ge 0$ and $r(w) = (r_0 + F + T) - (U + D_R) \ge 0$. Identity (3) is the sum of the two lines of (2), in which U and T cancel. The inequality follows from $D_P, D_R \ge 0$. □

Equation (3) is why release, uptake and losses need never be identified individually: they cancel or are dropped in one direction. The same summed identity holds for any finer admissible interleaving that respects the same pool availability, though only the displayed finite source is formalised.

Definition 2.3 (History)**.** A *history* is a pair of windows w_1, w_2 together with retention factors $0 \le e \le s \le 1$ and nonnegative recovery inputs i_P, i_R obeying

$$(w_2)_{p_0} \le e\, p(w_1) + i_P, \qquad (w_2)_{r_0} \le s\, r(w_1) + i_R. \tag{4}$$

We write $p = p(w_1)$, $r = r(w_1)$ for the latent pre-recovery state, F_1, F_2 for the two fresh amounts and Q_1, Q_2 for the two collections. The factors e and s are *calibrated upper retention maps* for the two pools, obtained by measuring recovery of the relevant analyte through the actual intervention. A dilution factor is not automatically a process recovery, s does not follow from e, and neither of them bounds retained productive activity, which is a separate parameter and is never inferred from a material retention factor.

Lemma 2.4 (Recovery account)**.** [recovery_account] *For every history, $Q_2 \le e\,p + s\,r + i_P + i_R + F_2$.*

Proof. By Lemma 2.2 applied to w_2, $Q_2 \le (w_2)_{p_0} + (w_2)_{r_0} + F_2$, and (4) bounds the two initial pools. □

2.3 Calibration data and observations

The certificate consumes the following, each of which must be an upper bound valid on the same event:

$$P_0 + R_0 \le B, \qquad r \le J, \qquad i_P + i_R \le H, \qquad 0 \le e \le s \le 1,$$

together with lower bounds $q_1 \le Q_1$, $q_2 \le Q_2$ on the two collections. Section 6 explains how q_k is produced from a reported value and an error allowance. J bounds the reserve *immediately before recovery*; §5 shows that this is not the same as a bound on the initial reserve, and gives an alternative premise that avoids measuring it directly.

2.4 A stoichiometric witness

For the lactate/acetate application the net conversion

$$4\,C_3H_6O_3 + 2\,C_2H_4O_2 \longrightarrow 3\,C_4H_8O_2 + 4\,CO_2 + 2\,H_2 + 2\,H_2O \tag{5}$$

balances carbon, hydrogen and oxygen [balanced_conversion]. This is an accounting witness that a net material conversion is possible. It is not a rate, a thermodynamic feasibility statement for a chosen medium, an organism-specific pathway, or an atom mapping. The mathematical source admits loss to competing products; any further positive source must enter the initial stock or the recovery account.

3 The exact certificate

Throughout, $[z]_+ = \max(0, z)$.

Definition 3.1. For $q_1, q_2, B, e, s, J, H \in \mathbb{R}$ put

$$L(q_1, q_2; B, e, s, J, H) = \max \left\{ \begin{array}{l} 0, \\ q_1 - B, \\ q_1 + q_2 - B - H, \\ eq_1 + q_2 - eB - (s - e)J - H, \\ sq_1 + q_2 - sB - H \end{array} \right\}. \tag{6}$$

We refer to the five entries as the *nonnegativity*, *first-window*, *total*, *reserve-sensitive* and *uniform-retention* expressions. They are a maximum of five affine lower bounds, not five irredundant facets in every parameter slice: at $e = s$ the reserve-sensitive entry is dominated by the uniform one, and other coincidences occur at degenerate parameters.

3.1 Soundness

Theorem 3.2 (Source-bound certificate)**.** [lower_sound, observed_source_bound] *Let a history satisfy* $P_0 + R_0 \le B$*,* $r \le J$*,* $i_P + i_R \le H$*, and let* $q_1 \le Q_1$*,* $q_2 \le Q_2$*. Then*

$$F_1 + F_2 \;\ge\; L(q_1, q_2; B, e, s, J, H).$$

Proof. Lemma 2.2 applied to w_1 and the stock bound give

$$q_1 + p + r \;\le\; Q_1 + p + r \;\le\; P_0 + R_0 + F_1 \;\le\; B + F_1, \tag{7}$$

and Lemma 2.4 with the input bound gives

$$q_2 \;\le\; Q_2 \;\le\; e\,p + s\,r + H + F_2. \tag{8}$$

We produce each entry of (6) as a nonnegative combination of (7), (8) and $r \le J$.

Nonnegativity. $F_1 + F_2 \ge 0$ because fresh amounts are nonnegative.

First-window. From (7) and $p, r \ge 0$, $q_1 - B \le F_1 \le F_1 + F_2$.

Total. Add (7) and (8). Since $e \le 1$ and $s \le 1$ we have $ep + sr \le p + r$, so

$$q_1 + q_2 \le B + H + F_1 + F_2 - (1-e)p - (1-s)r \le B + H + F_1 + F_2.$$

Reserve-sensitive. Multiply (7) by $e \ge 0$ and add (8), writing $ep + sr = e(p+r) + (s-e)r$:

$$eq_1 + q_2 \le eB + eF_1 + (s-e)r + H + F_2 \le eB + (s-e)J + H + F_1 + F_2,$$

using $r \le J$, $s \ge e$ and $eF_1 \le F_1$.

Uniform-retention. Multiply (7) by $s \ge 0$ instead and use $ep \le sp$:

$$sq_1 + q_2 \le sB + sF_1 + H + F_2 - (s-e)p \le sB + H + F_1 + F_2.$$

Taking the maximum of the five conclusions proves the theorem. □

These are explicit weak-duality certificates: fixed nonnegative multipliers applied to material constraints. No linear-programming solver is invoked at run time, and nothing here should be described as a verified general LP-duality engine.

3.2 The remaining-stock form

Proposition 3.3 (Closed form)**.** [lower_closed] *Let* $0 \le e \le s \le 1$ *and* $J \ge 0$*. With*

$$A = [q_1 - B]_+\,, \qquad x = [B - q_1]_+\,, \qquad K(x) = e\,x + (s-e)\min(J, x),$$

one has, for all $q_1, q_2, B, H \in \mathbb{R}$*,*

$$L(q_1, q_2; B, e, s, J, H) = A + [q_2 - H - K(x)]_+\,. \tag{9}$$

Proof. Note $A - x = q_1 - B$ and $A\,x = 0$.

(≤). Write $M = A + [q_2 - H - K(x)]_+$. Since $s \le 1$ and $\min(J, x) \le x$,

$$K(x) \le ex + (s-e)x = sx \le x, \qquad K(x) \le ex + (s-e)J.$$

Both summands of M are nonnegative, so the entry 0 satisfies $0 \le M$, and $q_1 - B = A - x \le A \le M$. For the remaining three entries, using $eA \le A$ and $sA \le A$,

$$
\begin{aligned}
q_1 + q_2 - B - H &= A + \big(q_2 - H - x\big), \\
eq_1 + q_2 - eB - (s-e)J - H &= eA + \big(q_2 - H - ex - (s-e)J\big), \\
sq_1 + q_2 - sB - H &= sA + \big(q_2 - H - sx\big),
\end{aligned}
$$

and in each case the bracket is at most $q_2 - H - K(x) \le [q_2 - H - K(x)]_+$, so each left-hand side is at most M.

($\ge$). If $q_1 \le B$ then $A = 0$ and $x = B - q_1$. When $x \le J$, $K(x) = sx$ and $A + [q_2 - H - K(x)]_+ = \max\{0, sq_1 + q_2 - sB - H\}$, a maximum of two entries of (6). When $x \ge J$, $K(x) = ex + (s-e)J$ and the expression equals $\max\{0, eq_1 + q_2 - eB - (s-e)J - H\}$, again a maximum of two entries. If $q_1 \ge B$ then $x = 0$, $K(0) = 0$ and the expression equals $(q_1 - B) + [q_2 - H]_+ = \max\{q_1 - B,\ q_1 + q_2 - B - H\}$. In every case the right-hand side of (9) is a maximum of entries of (6), hence at most L. □

Form (9) is the mechanism behind the formula. A is first-window formation that cannot be avoided. The credited stock still available after the first collection is x. To explain as much of the second collection as possible an adversary places up to J of that remainder in the better retained reserve pool and the rest in the mobile pool, producing maximal explanatory carryover $K(x)$. Fresh conversion is forced only by the part of q_2 that survives H and $K(x)$.

Remark 3.4 (Scalar specialisation). At $e = s = \rho$ the reserve information is inert, $K(x) = \rho x$, and (9) becomes $[q_1 - B]_+ + \Big[q_2 - H - \rho\,[B - q_1]_+\Big]_+$. This single-pool formula is inherited background for the present work and is not claimed as a discovery here; the two-pool refinement is what requires the pre-wash reserve bound.

Remark 3.5 (Endpoints). No step divides by a retention factor, so $e = 0$, $s = 0$, $e = s = 1$, $B = 0$, $H = 0$ and all equality boundaries are included.

3.3 Monotonicity and sensitivity

The following consequences of (9) are what an interface should report instead of a fixed narrative about which measurement to improve.

Proposition 3.6 (Sensitivity)**.** *Fix $0 \le e \le s \le 1$ and B, H. Then L is nondecreasing in q_1 and q_2, nonincreasing in B, H and J, and:*

- *(i) L does not depend on J once $J \ge x = [B - q_1]_+$; in particular tightening a reserve bound above the remaining stock buys nothing.*
- *(ii) For $0 \le J \le J' \le x$, $0 \le L(J) - L(J') \le (s-e)(J' - J)$, with equality when both second brackets in (9) are positive.*
- *(iii) If $e = s$ then L does not depend on J at all.*

Proof. Monotonicity in q_1, q_2, B, H is entry-wise in (6), all coefficients of q_1 being $0, 1, 1, e, s \ge 0$. For the rest use (9): $J \mapsto K(x) = ex + (s-e)\min(J, x)$ is nondecreasing, is constant for $J \ge x$, and increases by exactly $(s-e)(J' - J)$ on $[0, x]$; $[\cdot]_+$ is nondecreasing and 1-Lipschitz. Part (iii) is $s - e = 0$. □

4 Exactness

Theorem 3.2 would be of limited value if it were not tight: an experimenter needs to know whether an unresolved record reflects a weak argument or a genuine information barrier. It is tight, and the attaining schedule is explicit.

Theorem 4.1 (Attainment). [lower_attained] *Let $q_1, q_2, B, J, H \geq 0$ and $0 \leq e \leq s \leq 1$, and let $R_0 \in [0, B]$ be any initial reserve level. Then there is a history with*

$$P_0 = B - R_0, \quad R_0 = R_0, \quad Q_1 = q_1, \quad Q_2 = q_2, \quad r \leq J, \quad i_P + i_R \leq H,$$

whose first-window uptake satisfies $T_1 \leq [\min(J, B) - R_0]_+$ and whose fresh total is exactly

$$F_1 + F_2 = L(q_1, q_2; B, e, s, J, H).$$

Proof. Write $A = [q_1 - B]_+$, $x = [B - q_1]_+$, $r = \min(J, x)$ and $p = x - r \geq 0$. Set

$$F_1 = A, \qquad U_1 = [R_0 + A - r]_+, \qquad T_1 = [r - R_0 - A]_+, \qquad Q_1 = q_1,$$

with $D_{P,1} = D_{R,1} = 0$. At most one of U_1, T_1 is positive and $U_1 - T_1 = R_0 + A - r$. Consequently

$$r(w_1) = R_0 + A + T_1 - U_1 = r, \qquad p(w_1) = (B - R_0) + U_1 - T_1 - q_1 = B + A - q_1 - r = x - r = p,$$

using $A - x = q_1 - B$. The three availability constraints (1) for w_1 are exactly $U_1 \leq R_0 + A$ (true since $U_1 = [R_0 + A - r]_+$ and $R_0 + A \geq 0$), $r(w_1) \geq 0$ and $p(w_1) \geq 0$, all of which hold. Also $r \leq J$ and $r \leq x \leq B$, so $T_1 = [r - R_0 - A]_+ \leq [\min(J, B) - R_0]_+$.

For the recovery step take $i_P = 0$, $i_R = H$, so $(w_2)_{p_0} = e\,p$ and $(w_2)_{r_0} = s\,r + H$, which satisfy (4) with equality. Since $p + r = x$ and $r = \min(J, x)$,

$$e\,p + s\,r = e\,x + (s - e)\,r = K(x).$$

Now set

$$F_2 = [q_2 - H - K(x)]_+, \qquad U_2 = [q_2 - e\,p]_+, \qquad T_2 = 0, \qquad Q_2 = q_2,$$

$$D_{P,2} = e\,p + U_2 - q_2 \ \geq 0, \qquad D_{R,2} = s\,r + H + F_2 - U_2 \ \geq 0.$$

Nonnegativity of $D_{P,2}$ is $U_2 \geq q_2 - ep$. For $D_{R,2}$, both $0 \leq sr + H + F_2$ and $q_2 - ep \leq sr + H + F_2$ hold—the latter because $F_2 \geq q_2 - H - ep - sr$—so $U_2 = \max(0, q_2 - ep) \leq sr + H + F_2$; this same inequality is the third availability constraint $U_2 \leq (w_2)_{r_0} + F_2$. The remaining two constraints of (1) for w_2 hold with equality by the definition of $D_{P,2}$ and $D_{R,2}$.

Finally $F_1 + F_2 = A + [q_2 - H - K(x)]_+$, which equals L by Proposition 3.3. □

Corollary 4.2 (L is the exact minimum). [lower_is_minimum] *Under the hypotheses of Theorem 4.1, $L(q_1, q_2; B, e, s, J, H)$ is the least element of*

$$\Big\{F_1 + F_2 \ : \ \textit{histories with retention } (e, s),\ P_0 + R_0 \leq B,\ Q_1 \geq q_1,\ Q_2 \geq q_2,\ r \leq J,\ i_P + i_R \leq H\Big\}.$$

Proof. Membership is Theorem 4.1 with $R_0 = 0$; the lower-bound property is Theorem 3.2. □

Remark 4.3 (What sharpness does and does not say). Corollary 4.2 is a statement about the declared class, whose initial state is *bounded* by B and whose uncertainty bounds may vary independently unless an explicit constraint links them. The construction is therefore free to combine unfavourable extremes—maximal stock, maximal recovery input, maximal reserve—that a more detailed model might forbid. It does not show that a particular strain can execute the schedule in a given number of minutes: there are no calibrated transfer rates in the class. If a valid correlation or kinetic restriction excludes the extremiser, the lower bound remains sound and may cease to be attained; that is an opportunity to strengthen the result, and any claim of better performance must name the restriction that does it (§10.1).

5 The reserve premise: necessity, ambiguity, and a substitute

5.1 An initial reserve ceiling is not a pre-wash ceiling

Theorem 5.1 (Storage counterexample)**.** [storage_counterexample, initial_reserve_substitution_unsound] *There is a history with*

$$P_0 = 8, \quad R_0 = 2, \quad Q_1 = 6, \quad T_1 = 2, \quad F_1 = F_2 = 0, \quad e = \tfrac{1}{20},\ s = \tfrac{9}{10},\ i_P = i_R = 0,$$

whose pre-wash state is $(p, r) = (0, 4)$ *and whose second collection is* $Q_2 = \frac{18}{5}$*. Substituting the initial reserve ceiling* 2 *for* J *in* (6) *returns* $L = \frac{17}{10}$*, whereas the true fresh total is* 0*. Evaluated at the correct pre-wash ceiling* $J = 4$ *the formula returns* 0*.*

Proof. Take $w_1 = (8, 2, 0, 0, 2, 6, 0, 0)$; the availability constraints (1) read $0 \le 2$, $0 \le 4$ and $2+6 \le 8$, and (2) gives $p = 0$, $r = 4$. Take $w_2 = (0, \frac{18}{5}, 0, \frac{18}{5}, 0, \frac{18}{5}, 0, 0)$; the recovery inequalities (4) hold with equality since $e \cdot 0 = 0$ and $s \cdot 4 = \frac{18}{5}$. Both fresh amounts are zero. The two evaluations of (6) are arithmetic. □

The defect is storage of initially mobile product, not instrument noise, numerical tolerance or insufficient proof effort. It refutes the unconditional initial-reserve substitution. It does *not* show that destructive pre-wash measurement is the only repair, and it does not make J necessary for every positive certificate: the nonnegativity, first-window, total and uniform-retention entries of (6) contain no reserve information, so

$$L \ \ge\ \max\{0,\ q_1 - B,\ q_1 + q_2 - B - H,\ sq_1 + q_2 - sB - H\}$$

is always available with no reserve calibration whatsoever, and is positive in some regimes.

5.2 The worked record does need it

Necessity is nevertheless real for the example we use in §7.

Theorem 5.2 (The reserve premise is not redundant there)**.** [reserve_premise_not_redundant] *There is a history with* $P_0 + R_0 = 10$*,* $i_P + i_R = \frac{1}{5}$*, retention* $e = \frac{1}{20}$*,* $s = \frac{9}{10}$*, collections* $Q_1 = \frac{29}{5}$ *and* $Q_2 = \frac{381}{100}$*—so* $|Q_1 - 6| \le \frac{1}{5}$ *and* $|Q_2 - 4| \le \frac{1}{5}$*—pre-wash reserve* $r = 4$*, and* $F_1 + F_2 = 0$*.*

Proof. Take $w_1 = (8, 2, 0, 0, 2, \frac{29}{5}, 0, 0)$, giving $p = \frac{1}{5}$, $r = 4$, and $w_2 = (\frac{1}{100}, \frac{19}{5}, 0, \frac{19}{5}, 0, \frac{381}{100}, 0, 0)$ with $i_P = 0$, $i_R = \frac{1}{5}$. Then $e\,p = \frac{1}{100}$ and $s\,r + i_R = \frac{18}{5} + \frac{1}{5} = \frac{19}{5}$, so (4) holds with equality, and the availability constraints are $\frac{19}{5} \le \frac{19}{5}$ and $\frac{381}{100} \le \frac{1}{100} + \frac{19}{5}$. □

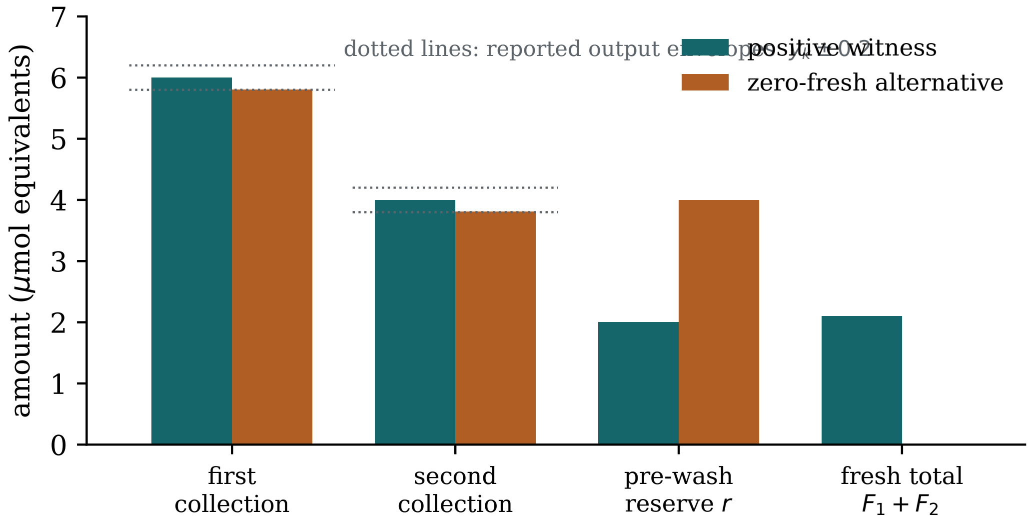

Figure 2: Two exact material histories inside the same output envelopes. They differ in the pre-wash reserve and in the fresh total. An independently justified reserve ceiling separates them; the endpoint record does not.

Both collections of this zero-fresh history lie inside the reported envelopes of the positive witness of §7; the two differ only in the pre-wash reserve, 4 against 2. The premise $r \le 2$ excludes it; the two product measurements alone do not (Figure 2). No claim is made that the alternative also matches substrate, isotope or biomass observations that were not part of the declared record—that is exactly the kind of extra information that would strengthen the certificate.

5.3 Bounded uptake in place of direct reserve measurement

Destructive measurement of the pre-wash reserve is expensive, consumes the preparation, and needs its own transport and heterogeneity calibration. The following shows it is sufficient but not necessary: a bound on the *initial* reserve together with a cap on *cumulative first-window uptake* supplies the same expression.

Theorem 5.3 (Uptake-cap certificate)**.** [lower_sound_uptake, uptake_source_bound] *Let a history satisfy* $P_0 + R_0 \le B$, $i_P + i_R \le H$, $q_1 \le Q_1$, $q_2 \le Q_2$, *and suppose valid bounds*

$$R_0 \le \overline{R}, \qquad T_1 \le \overline{T}, \qquad J_0 := \overline{R} + \overline{T}.$$

Then $F_1 + F_2 \;\ge\; L(q_1, q_2; B, e, s, J_0, H)$.

Proof. From (2) and $U_1, D_{R,1} \ge 0$ [prewash_reserve_from_uptake],

$$r = R_0 + F_1 + T_1 - U_1 - D_{R,1} \;\le\; J_0 + F_1. \tag{10}$$

Note this is *not* the claim $r \le J_0$, and that F_1 is unknown. The four entries of (6) that do not mention J are unaffected, so only the reserve-sensitive entry must be re-derived. Multiply (7) by e and add (8), writing $ep + sr = e(p + r) + (s - e)r$ and using (10):

$$q_2 \le e(B + F_1 - q_1) + (s - e)(J_0 + F_1) + H + F_2 = e(B - q_1) + (s - e)J_0 + H + sF_1 + F_2.$$

Since $s \le 1$ and $F_1 \ge 0$, $sF_1 \le F_1$, giving $eq_1 + q_2 - eB - (s - e)J_0 - H \le F_1 + F_2$. $\square$

The reason the unknown F_1 does no harm is precisely $s \le 1$: material formed during the first window and stored cannot have more than unit retained explanatory value in the second. The bound is therefore not circular.

Theorem 5.4 (Exactness under the uptake premise)**.** [uptake_bound_attained] *Let* $q_1, q_2, B, H, \overline{R}, \overline{T} \ge 0$ *and* $0 \le e \le s \le 1$*. There is a history with* $P_0 + R_0 \le B$*,* $Q_1 = q_1$*,* $Q_2 = q_2$*,* $R_0 \le \overline{R}$*,* $T_1 \le \overline{T}$*,* $i_P + i_R \le H$ *and* $F_1 + F_2 = L(q_1, q_2; B, e, s, \overline{R} + \overline{T}, H)$*.*

Proof. Apply Theorem 4.1 with $J = J_0 = \overline{R} + \overline{T}$ and $R_0 = \min(\overline{R}, B)$. Then $R_0 \le \overline{R}$ and the theorem gives $T_1 \le [\min(J_0, B) - R_0]_+$. If $\overline{R} \le B$ then $R_0 = \overline{R}$ and $\min(J_0, B) - R_0 \le J_0 - \overline{R} = \overline{T}$; if $\overline{R} > B$ then $R_0 = B$ and $\min(J_0, B) - R_0 \le 0$. In both cases $T_1 \le \overline{T}$. □

So the substitution $J := \overline{R} + \overline{T}$ loses nothing: the same expression is again the exact minimum, now over the uptake-bounded class.

Remark 5.5 (How to obtain an uptake cap, and how it can fail). A calibrated upper uptake rate yields a cumulative cap $T_1 \le \int_0^{h_1} v_{\max} X(t)\, dt \le v_{\max} X_{\max} h_1$ in consistent product-equivalent units, which uses a mechanistic constraint without requiring an identified kinetic model. Three cautions. The bound must be a valid *upper* rate for the preparation and conditions, not a best-fit estimate. The cap concerns *gross* uptake into the credited reserve, not net extracellular disappearance: simultaneous release conceals uptake in a net time course. And a cap on uptake of the wrong chemical species does not bound storage of credited product.

Remark 5.6 (Combining the two premises). If both a direct reserve bound J and an uptake pair $(\overline{R}, \overline{T})$ are available and valid, compute both certificates and report the maximum; both are sound. Keep the two premise records distinct. In particular J_0 is not itself a proved bound on the actual pre-wash reserve when $F_1 > 0$, and an interface documenting its input as a measured pre-wash ceiling must not silently accept J_0 without recording the different biological premise. On the history of Theorem 5.1 the uptake route with $\overline{R} = \overline{T} = 2$ gives $J_0 = 4$ and certifies 0, as it must [uptake_route_sound_here].

6 Observation uncertainty and the reporting rule

6.1 A deterministic wrapper

Let y_k be reported output amounts with absolute allowances ϵ_k, and consider the simultaneous event

$$\mathcal{E} = \{\, |Q_1 - y_1| \le \epsilon_1,\ |Q_2 - y_2| \le \epsilon_2,\ \text{and all stated calibration bounds hold} \,\}.$$

On $\mathcal{E}$ put $q_k = [y_k - \epsilon_k]_+$. Since $Q_k \ge 0$ and $Q_k \ge y_k - \epsilon_k$, these are valid lower bounds, and L is nondecreasing in q_1, q_2 by Proposition 3.6.

Proposition 6.1 (Reporting guarantee)**.** [noisy_observation, reporting, resolution] *On* $\mathcal{E}$*, issuing the verdict* at or above *only when* $L(q_1, q_2; B, e, s, J, H) \ge F_\star$ *guarantees* $F_1 + F_2 \ge F_\star$*. If a separately justified upper bound* $U_F \ge F_1 + F_2$ *is supplied, issuing* below *only when* $U_F < F_\star$ *is likewise sound, and the record is* incompatible *when* $L > U_F$ *or when some* $y_k + \epsilon_k < 0$*. All other records are* unresolved*.*

Proof. The first claim is Theorem 3.2 composed with transitivity; the second is immediate. For incompatibility, $L > U_F$ contradicts Theorem 3.2, and $y_k + \epsilon_k < 0$ contradicts $Q_k \ge 0$ on $\mathcal{E}$. Conversely, when the observation intervals meet the nonnegative half-line and $L \le U_F$, Theorem 4.1 produces a history in the class, so the report is not vacuous. □

Two points deserve emphasis. A small observed output does *not* supply U_F: generated material may be lost, consumed or retained, so an upper certificate needs its own carbon or input calibration. And a negative instrument reading is compatible with a nonnegative true amount whenever its interval meets zero; only $y_k + \epsilon_k < 0$ is an inconsistency of the declared observation model, which is an input error rather than a biological diagnosis.

6.2 Amount errors

For a single export measured as volume times concentration, bounded errors $|V - \widehat{V}| \le \epsilon_V$ and $|c - \widehat{c}| \le \epsilon_c$ give

$$|Vc - \widehat{V}\widehat{c}| \le |\widehat{V}|\epsilon_c + |\widehat{c}|\epsilon_V + \epsilon_V \epsilon_c, \tag{11}$$

by writing $Vc - \widehat{V}\widehat{c} = \widehat{V}(c - \widehat{c}) + \widehat{c}(V - \widehat{V}) + (V - \widehat{V})(c - \widehat{c})$ and applying the triangle inequality. Allowances for distinct exports are summed. The calibration feeding ϵ_k must cover dilution, recovery, background subtraction and normalisation as well, and shared bias must not be treated as independent noise that averages away: the wrapper above is a single simultaneous event, not a sequence of independent draws.

6.3 What a probability statement would require

No probability is attached to $\mathcal{E}$ here. If a separately justified simultaneous calibration establishes $\Pr(\mathcal{E}) \ge 1 - \alpha$, then the event that an issued at-or-above claim is false is contained in $\mathcal{E}^{c}$, so

$$\Pr(\text{an issued claim is false}) \le \alpha.$$

This is an unconditional error-event guarantee. It is not a posterior probability, and it is not automatically a bound on the error probability *conditional on issuing* a claim. A nominal "95%" label would require additional statistical work, including a treatment of selection if times or protocols are chosen adaptively. The schedule here is fixed in advance, and repeated timepoints on one preparation are not independent cultures.

6.4 Reporting a rate

Dividing a certified fresh amount by the recorded window duration gives a lower bound on the average rate of credited-inventory generation over that window. Because the target is a completed record, no equilibrium assumption, transient correction or interpolation of an unobserved integral is needed—and none is available: a window-average lower bound is not a steady-state rate.

7 Design: what a wash can and cannot buy

7.1 Cancellation of known mobile product

Write the actual second output of a given material source as $Q_2 = e\,p + s\,r + f - \ell$, where $f = F_2$ and ℓ collects applicable losses. Then

$$Q_2 - e\,p - s\,J - H = f - \ell - s(J - r) - H, \tag{12}$$

[wash_signal_cancellation] and the right-hand side does not contain e. Removing known mobile product removes it from the explanatory background and from the observed output in equal measure. A comparison that holds Q_2 fixed while lowering e therefore manufactures a benefit that the matched physical source does not produce; that comparison is the main way this intervention is oversold.

Quantity (μmol per original aliquot unless noted)	Washed	Unwashed
Original aliquot	10 mL	10 mL
Initial total bound B	10	10
Pre-wash reserve bound J	2	2
Reported first collection y_1	6	6
Reported second collection y_2	4	5.7
Output allowances ϵ_1, ϵ_2	$0.2, 0.2$	$0.2, 0.2$
Mobile retention e	0.05	0.9
Reserve retention s	0.9	0.9
Recovery-input ceiling H	0.2	0.2
Recovery input realised in the witness	0	0
Fresh credited inventory realised in the witness	2.1	2.1
Robust certificate L	**1.69**	**1.52**
Active expression	reserve-sensitive	total / retention (tied)
Verdict at $F_\star = 1.6$	at or above	unresolved

Table 1: Matched-source comparison. Every entry is synthetic. Both arms carry identical reserve information and identical absolute-error contracts; washing adds one preparation operation. No empirical equal-cost claim is implied.

7.2 Precision gain, and the reserve slack that eats it

The conclusion changes when carryover is uncertain rather than known. Assume the matched-source regime: tight initial total, no first-window fresh formation or loss, unclipped lower observations, the reserve-sensitive entry active, and all other quantities equal between arms. Then the value of that entry is $f - \ell - (s-e)(J-r) - H - e\epsilon_1 - \epsilon_2$, and relative to an unwashed arm with $e = s$,

$$G = (s-e)\big(\epsilon_1 - (J-r)\big). \tag{13}$$

The benefit is attenuation of the first-window measurement uncertainty, which enters the washed certificate weighted by e rather than by s [precision_gain_identity, reserve_precision_tradeoff]. It is positive only when the reserve slack $J-r$ is smaller than ϵ_1. Only at a tight reserve ceiling, $J = r$, does (13) reduce to $G = (s-e)\epsilon_1$. Equation (13) is a conditional design identity, not a general statement that washing improves precision.

7.3 A worked synthetic comparison

The source witness is explicit and is a valid history. Window one starts at $(P, R) = (8, 2)$ and collects 6 with no fresh formation, ending at $(2, 2)$. Recovery with $e = 0.05$, $s = 0.9$ gives $(0.1, 1.8)$; fresh formation of 2.1 followed by release of 3.9 permits a second collection of 4 with zero final inventory. In the matched unwashed arm the second window instead starts at $(1.8, 1.8)$ and the same fresh formation permits a collection of 5.7 [witness_useful, witness_satisfies_root].

Remark 7.1 (The first collection is selective). Collecting 6 while leaving the reserve pool at 2 is a *cell-retaining selective product separation*. It must not be pictured as withdrawing 75% of a homogeneous whole culture while retaining all cells and reserves. If the laboratory operation removes whole culture, its actual reserve and biomass losses must be propagated into the next state. One consistent decomposition of the recovery factors is a shared 0.9 retention from sampling followed by an additional mobile-liquid factor 1/18, giving $e = 0.05$ and $s = 0.9$; the activity parameter a below then measures loss *relative to the already sampled unwashed comparator*, and charging the same 0.9 again as an activity loss would double-count a common operation.

With $q_1 = 5.8$ and $q_2 = 3.8$ the five expressions of (6) take the values

$$0, \qquad -4.2, \qquad -0.6, \qquad \mathbf{1.69}, \qquad -0.18,$$

so the reserve-sensitive expression is uniquely active and $L = \frac{169}{100}$ [resolution]. In the matched unwashed arm $q_2 = 5.5$, the total expression is 1.1, the two retention expressions coincide at $\frac{38}{25}$, and $L = 1.52$ [unwashed_value]. The ideal gain is $(s - e)\epsilon_1 = 0.85 \times 0.2 = 0.17$.

Two consistency remarks. First, the realised fresh amount 2.1 exceeds the certified 1.69; the difference is the allowed measurement and recovery-input slack, not a failed conservation check. Second, substituting the reported values without subtracting the allowances gives 1.90, so a claim of "at least 1.8" is *not* protected by the stated error contract even though the nominal arithmetic appears to support it. A shared upward bias of 0.2 in both channels is admitted by $\mathcal{E}$ simultaneously and is not reduced by replication.

Remark 7.2 (Threshold bookkeeping). We use $F_\star = 1.6$ for the principal comparison throughout. At that threshold the washed record is at or above and the matched unwashed record is unresolved. The companion calculator ships a regression default of 1.5, at which *both* arms are at or above; the two numbers serve different purposes and should not be interchanged.

7.4 Exact reserve tolerance

Holding the other washed inputs fixed, the certificate is an explicit piecewise-linear function of the reserve ceiling.

Proposition 7.3 (Reserve tolerance)**.** [reserve_tolerance] *For every $J \in \mathbb{R}$,*

$$L\left(\tfrac{29}{5}, \tfrac{19}{5}; 10, \tfrac{1}{20}, \tfrac{9}{10}, J, \tfrac{1}{5}\right) = \max\Big(0,\ \tfrac{339}{100} - \tfrac{17}{20}J\Big).$$

Consequently the three natural goals have different and separated tolerances:

Goal	Exact condition on the reserve ceiling
Certify strictly positive fresh inventory	$J < \frac{339}{85} \approx 3.98824$
Certify at least 1.6	$J \leq \frac{179}{85} \approx 2.10588$
Strictly beat the unwashed certificate 1.52	$J < \frac{11}{5} = 2.2$

A method can produce a positive certificate while failing the amount target, and can meet the amount target while failing to justify the wash. These should not be merged into a single notion of assay success (Figure 3).

7.5 The joint activity and slack budget

A wash may also cost productive function. Let $a \in [0, 1]$ denote productive activity retained *in addition* to what the sampled unwashed comparator retains (Remark 7.1), and let the reserve ceiling be $J = 2 + \delta$ with slack $\delta \geq 0$. In the declared constant-capacity response family, fresh second-window formation is $2.1a$ and the reported second output is $1.9 + 2.1a$.

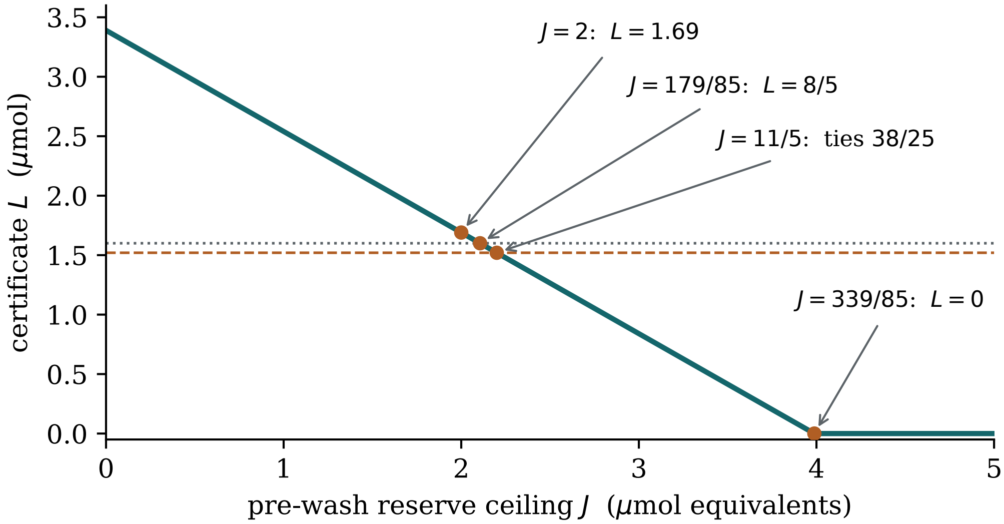

Figure 3: Certificate of the washed worked record as an exact function of the pre-wash reserve ceiling (Proposition 7.3). Unknown reserve, not numerical optimisation, closes the certificate at large J. The dashed line is the matched unwashed certificate 38/25 and the dotted line the target 8/5.

Theorem 7.4 (Exact design budget)**.** [wash_curve_slack, improvement_criterion_slack, target_criterion_slack]
For all $\delta \in \mathbb{R}$ *and all* $a \le 1$*,*

$$L_{\mathrm{wash}}(a,\delta) := L\big(\tfrac{29}{5}, \tfrac{17}{10} + \tfrac{21}{10}a; 10, \tfrac{1}{20}, \tfrac{9}{10}, 2+\delta, \tfrac{1}{5}\big) = \max\Big(0,\ \tfrac{21}{10}a - \tfrac{41}{100} - \tfrac{17}{20}\delta\Big),$$

and consequently

$$L_{\mathrm{wash}}(a,\delta) > \tfrac{38}{25} \iff a > \tfrac{193}{210} + \tfrac{17}{42}\delta, \qquad L_{\mathrm{wash}}(a,\delta) \ge \tfrac{8}{5} \iff a \ge \tfrac{67}{70} + \tfrac{17}{42}\delta.$$

The equality in Theorem 7.4 is a genuine identity for the whole family, not an assumption that the reserve-sensitive expression is active: for $a \le 1$ the total and uniform expressions are at most $2.1 - 2.7 < 0$ and $2.1 - 2.28 < 0$ respectively, so they never win.

Reserve slack δ	Activity to strictly beat 1.52	Activity to certify 1.6
0	$a > 91.9048\%$	$a \ge 95.7143\%$
0.05	$a > 93.9286\%$	$a \ge 97.7381\%$
0.10	$a > 95.9524\%$	$a \ge 99.7619\%$
0.20	would need $a > 100\%$	would need $a \ge 103.8095\%$

The last row is infeasible in the declared range $a \in [0,1]$: at $\delta = 0.2$ the reserve slack alone consumes the whole 0.17 margin, exactly as (13) predicts. At $\delta = 0$ and $a = 0.9$ the washed certificate is 1.48 and *loses* to the unwashed comparator (Figure 4). Note also the boundary conventions: the improvement criterion is strict at $\frac{193}{210}$, where the two certificates tie, and the 1.6 certification boundary is inclusive at $\frac{67}{70}$.

Collecting the terms, if a further washed-only certificate penalty $c \ge 0$ captures additional uncertainty or input costs without double-counting, strict improvement requires

$$0.17 - 0.85\,\delta - 2.1\,(1-a) - c > 0. \tag{14}$$

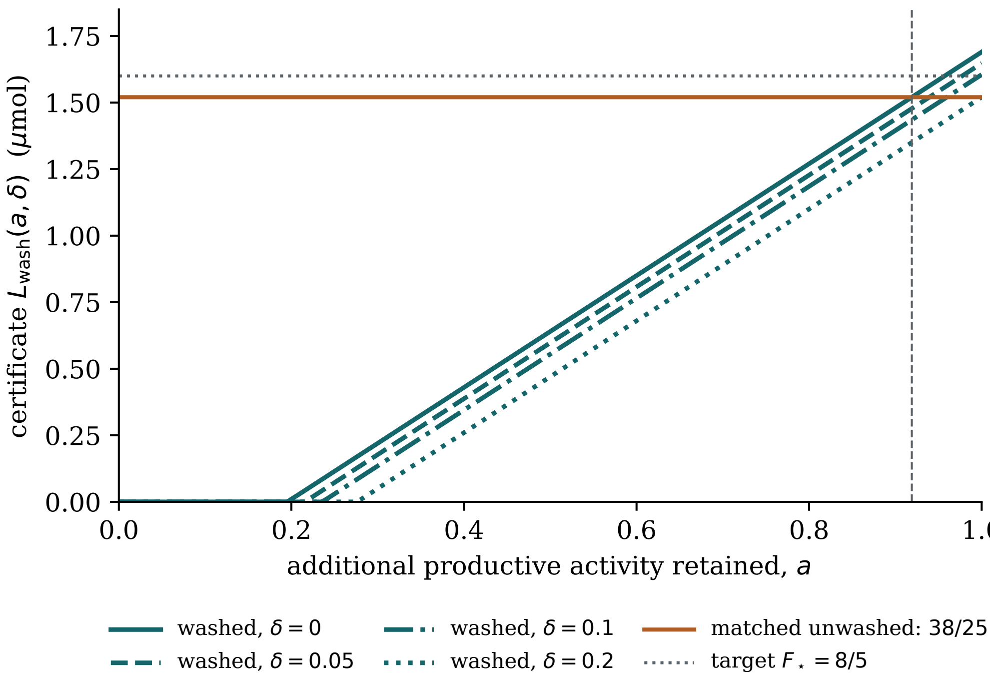

Figure 4: Exact washed certificate against retained activity, for four values of the reserve slack δ (Theorem 7.4). Solid horizontal line: matched unwashed certificate 38/25. This is a family of exact decision bounds, not a fitted dose response.

This is the quantity to estimate *before* investing in a comparative experiment: the question is whether the combined reserve, activity and extra-error costs fit inside 0.17 μmol.

7.6 The uptake route on the same record

Theorem 5.3 lets the same record be certified without destructive reserve measurement, at a price that is equally explicit. With $\overline{R} = 2$ on the worked washed inputs, the substitution $J_0 = 2 + \overline{T}$ and Proposition 7.3 give

$$L \geq \tfrac{8}{5} \iff \overline{T} \leq \tfrac{9}{85} \approx 0.10588, \qquad L > \tfrac{38}{25} \iff \overline{T} < \tfrac{1}{5}.$$

So the uptake route closes the logical gap in the claim that direct reserve measurement is mandatory, but it does not make the required calibration easy: the admissible cumulative uptake is about one tenth of a μmol per original aliquot. Consistently, on the storage counterexample—whose actual uptake is 2—the route correctly certifies nothing.

8 Formal verification and the evidence boundary

8.1 What is machine-checked

The development is a small chain of Lean 4 modules compiled against a pinned Mathlib, with warnings treated as errors. Table 2 maps every claim of this paper to its evidence.

Claim	Declaration	Status
Pool nonnegativity, balance (3), sampling identity	`pool_nonneg`, `source_balance`, `sample_amount`	Lean (`Source`)
Stoichiometric witness (5)	`balanced_conversion`	Lean (`Source`)
Certificate, Theorem 3.2	`lower_sound`, `observed_source_bound`	Lean (`Certificate`)
Observation wrapper, Proposition 6.1	`noisy_observation`, `reporting`	Lean (`Certificate`)
Cancellation identity (12)	`wash_signal_cancellation`	Lean (`Certificate`)
Closed form, Proposition 3.3	`lower_closed`	Lean (`Sharpness`)
Attainment, Theorem 4.1	`lower_attained`	Lean (`Sharpness`)
Exact minimum, Corollary 4.2	`lower_is_minimum`	Lean (`Sharpness`)
Uptake certificate, Theorem 5.3	`lower_sound_uptake`, `uptake_source_bound`	Lean (`Sharpness`)
Uptake exactness, Theorem 5.4	`uptake_bound_attained`	Lean (`Sharpness`)
Root threshold and concrete witness	`resolution`, `witness_useful`, `witness_satisfies_root`	Lean (`Resolution`)
Precision identities (13)	`precision_gain_identity`, `reserve_precision_tradeoff`	Lean (`Design`)
Activity criterion at $\delta = 0$	`wash_curve`, `unwashed_value`, `improvement_criterion`	Lean (`Design`)
Reserve tolerance, Proposition 7.3	`reserve_tolerance`	Lean (`Sensitivity`)
Joint budget, Theorem 7.4	`wash_curve_slack`, `improvement_criterion_slack`, `target_criterion_slack`	Lean (`Sensitivity`)
Storage counterexample, Theorem 5.1	`storage_counterexample`, `initial_reserve_substitution_unsound`, `uptake_route_sound_here`	Lean (`Sensitivity`)
Zero-fresh ambiguity, Theorem 5.2	`reserve_premise_not_redundant`	Lean (`Sensitivity`)
Sensitivity, Proposition 3.6; general interleavings; probability containment	—	Conventional proof here
Printed numerical values; replay of the attaining schedule	`check_paper.py`	Exact rational computation
Biological validity of $B, J, e, s, H, \epsilon_k, a$; lactate-specific attribution	—	Unvalidated empirical premise

Table 2: Claim-to-evidence map. Module names abbreviate `MicrobialFunctionAssay.`⟨`module`⟩.

8.2 What formal verification does not establish

Compiling the algebra substantially reduces the risk of a mistake in the encoded statements and in the binding between the source model and the certificate; the checked declarations depend only on the standard classical axioms of Mathlib, and no axiom asserting the assay conclusion was introduced. It does not make the modelling choices correct. It does not validate the Python interface, the figures, or any biological premise, and it certainly does not “prove the biology”. The honest statement is the one in Table 2: the named declarations passed strict verification, and the conventional, computational and empirical components are listed separately.

9 Calibration requirements and a feasibility pilot

9.1 What must be measured or bounded

Quantity	Principal risk	Useful evidence
Initial total B	Old convertible material omitted	Conservative carbon inventory; media blanks; biomass and sorption terms
Pre-wash J	Initial reserve substituted; aliquots not exchangeable	Direct reserve bound with calibrated transport, or the $(\overline{R}, \overline{T})$ premise of Theorem 5.3
Mobile retention e	Dilution factor mistaken for process recovery	Recovery of the relevant mobile analyte through the actual intervention
Reserve retention s	Assumed equal to e or to a cell count	Conservative retained convertible-material bound
Recovery input H	Wash medium or contamination adds explanatory material	Material-accounted recovery blanks and input ceilings
Activity a	Inferred from material retention or optical density	Matched functional comparison under identical handling
Allowances ϵ_k	Shared bias, volume error or selection omitted	Calibration covering the complete amount measurement, simultaneously valid
Source interpretation	Product-equivalent account read as lactate-specific flux	Declared conversion boundary; isotope bridge if attribution is claimed

Sacrificial subsampling requires more than a mean reserve estimate. For a sample fraction z with measured carbon C and simultaneous carbon error d, a parent-unit bound has the form $J \leq (C+d)/(4z)$ plus a calibrated between-aliquot discrepancy covering heterogeneity, homogenisation, transport, volume scaling and extraction. Removing the subsample changes both retained material and activity, so it must appear in s and a. Optical density may monitor biomass but is not by itself a reserve-carbon measurement or an activity certificate.

9.2 Reference organisms

A direct lactate/acetate challenge of a defined culture is the natural application, and defined-culture cross-feeding between *Bifidobacterium adolescentis* and butyrate producers is documented [1]. Two cautions matter for a pilot.

First, the challenge tests a downstream conversion. It bypasses the upstream carbohydrate-degrading member and therefore cannot certify the complete carbohydrate-to-lactate-to-butyrate community process.

Second, lactate utilisation is strain-specific and should not be inferred from butyrate production generally. Duncan, Louis and Flint isolated lactate-utilising butyrate producers related to *Eubacterium hallii* and *Anaerostipes caccae*, and did not detect significant lactate utilisation in the *Roseburia intestinalis*, *Eubacterium rectale* and *Faecalibacterium prausnitzii* strains they tested under their conditions [2]. A pilot should select a documented strain and substrate phenotype accordingly. Muñoz-Tamayo and colleagues built a lactate/acetate-to-butyrate kinetic model using *E. hallii* L2-7 and the then-described *Anaerostipes coli* SS2/1 [4]; that work is both a directly relevant kinetic comparator and a reminder that substrate behaviour is strain-specific. Historical names should be used when discussing those papers, and current identifiers, taxonomy and material availability verified before procurement.

9.3 The smallest useful study

The immediate empirical question is not a population error rate but whether simultaneous conservative bounds on reserve, recovery and activity leave any usable margin at all—that is, whether (14) can be satisfied.

Use a small number of independently identified preparations, each split into matched washed and unwashed conditions, with destructive reserve measurements on separately designated aliquots where needed. Record initial and pre-wash lactate, acetate and butyrate amounts *with sample volumes*, export volumes, wash waste, recovery inputs, and the reserve-assay transport calibration. Include a deliberate stored-product or uptake perturbation, which is the discriminating intervention for the storage mechanism of Theorem 5.1, and a separate activity-preservation readout. A killed or no-substrate control is a model check, not exact background subtraction, unless its counterfactual mismatch is bounded. Retain incompatible and failed records rather than conditioning them away.

The first deliverable is a table of calibrated or conservatively bounded inputs with the resulting verdicts, including unresolved and incompatible cases, plotted against the feasibility conditions of §7. A failure to obtain $L > 0$ should be decomposed into a stock, reserve, recovery-input, measurement or activity cost before the protocol is changed: if J is too broad, compare direct measurement, the uptake cap and an added state measurement by how much each excludes the current zero-fresh explanation; if activity cost consumes the 0.17 margin, keep the unwashed protocol. The theorem supplies no reason to preserve washing as a design commitment.

The falsifiable prediction is that admitted histories never have total fresh inventory below the issued certificate, and that any wash gain is bounded by the propagated-error gain minus activity, input and loss penalties. A detected missing carbon contribution or an underestimated J invalidates the calibration and must be recorded as such. Positive certification, practical reproducibility and wash superiority are three distinct milestones, and a useful preparation need not achieve the third.

10 Relation to kinetic inference, and limitations

10.1 Kinetic restrictions can only help—if they are valid

The certificate is robust to omitted kinetics because its admissible class allows broad nonnegative transfer histories. The price is that it cannot exploit valid rate or pathway information unless that information is added explicitly. The relationship is elementary and worth stating, because it is often reversed in informal discussion. If $\mathcal{C}_{\mathrm{mat}}$ is the set of histories compatible with the accounting and observations, and a correctly specified kinetic model imposes further constraints, then

$$\mathcal{C}_{\mathrm{kin}} \subseteq \mathcal{C}_{\mathrm{mat}} \quad \Longrightarrow \quad \inf_{h \in \mathcal{C}_{\mathrm{kin}}} F(h) \;\geq\; \inf_{h \in \mathcal{C}_{\mathrm{mat}}} F(h). \tag{15}$$

A valid kinetic restriction therefore *strengthens* a lower certificate. The distinction between the approaches is assumption dependence, not a general superiority of material accounting: if a kinetic model excludes the true history, a narrower interval is misleading; if its assumptions are correct and calibrated, a tighter result is legitimate. Dynamic metabolic flux analysis and its uncertainty-aware variants are serious methods [3, 5] and should not be caricatured as curve fitting. Theorem 5.3 is the smallest useful instance of (15) in this setting: one calibrated mechanistic bound, no identified model.

A fair empirical comparison would give both methods the same observation budget, or report the information difference explicitly, and would let an enhanced material baseline use any extra linear restrictions implied by dense sampling. We do not attempt it here; a hastily implemented straw-man baseline would weaken rather than support the present result.

10.2 Limitations

1. The certificate is conditional on the declared material and observation bounds being simultaneously valid. An omitted carbon-bearing pool invalidates it; no amount of internal consistency detects that omission.

2. Sharpness is sharpness *in the declared class* (Remark 4.3). The class imposes no calibrated transfer rates, so the extremal schedule need not be biologically realisable in a given window.

3. Rectangular treatment of the calibration bounds is sound but not optimal for a genuinely joint uncertainty set; if a validated joint set is available, optimising over it can only improve the result.

4. The wash analysis of §7 is specific to a matched-source comparison with the reserve-sensitive expression active; the 0.17 μmol figure requires the tight reserve ceiling $J = r$ and is not a general property of washing.

5. The activity model behind Theorem 7.4 is a declared constant-capacity family. The percentages it produces are requirements, not observations.

6. Nothing here constrains function after the second window. Two continuations can agree on the entire observed past and then differ between zero and positive fresh conversion, so a future guarantee needs an additional readiness or dynamical premise. Likewise the worked witness has $F_1 = 0$ and therefore cannot support a claim of *recovery* relative to a demonstrably positive first window; that comparison needs sound intervals $[L_1, U_1]$, $[L_2, U_2]$ with $U_1 > 0$ and a criterion such as $L_2 \geq \kappa U_1$.

7. No biological data calibrate any parameter used here, and no claim of priority, global assay optimality or clinical utility is made.

10.3 Outlook

Three continuations are well posed and independent. Empirical calibration of J and of activity retention would decide whether (14) is ever satisfiable in a real preparation. An isotope bridge would move the target from credited inventory to substrate-specific attribution, which is a strictly stronger and currently unsupported claim. And a careful instance of (15)—adding calibrated rate bounds one at a time and measuring how much each excludes the zero-fresh explanation—would show how far the barrier of Corollary 4.2 can be pushed without an identified kinetic model. None of these is a prerequisite for the completed-record statement proved here.

A Reproducibility

Proof modules. The Lean development consists of `Source.lean` (windows, histories, balance, stoichiometric witness), `Certificate.lean` (Theorem 3.2, Proposition 6.1, (12)), `Sharpness.lean` (Proposition 3.3, Theorems 4.1, 5.3, 5.4, Corollary 4.2), `Resolution.lean` (the threshold instance and its concrete witness), `Design.lean` (the precision identities and the $\delta = 0$ activity criterion) and `Sensitivity.lean` (Proposition 7.3, Theorems 7.4, 5.1, 5.2), all in namespace `MicrobialFunctionAssay`. They compile in a pinned `leanprover/lean4:v4.30.0` workspace with `-DwarningAsError=true`. Verification receipts record the toolchain, the manifest hash, the per-file SHA-256 and the axiom report for every requested declaration; across the six modules, 35 requested declarations verify, and all of them depend only on `propext`, `Classical.choice` and `Quot.sound`. No axiom asserting any assay conclusion is introduced.

Numerical audit. `check_paper.py` uses the standard library only. It recomputes every numerical value printed in this paper with exact rational arithmetic, and independently replays the attaining schedule of Theorem 4.1—asserting each nonnegativity and availability constraint—over a grid of degenerate parameter combinations together with randomised rational cases, checking in each instance that the constructed fresh total equals both (6) and (9) and that the uptake bound of Theorem 5.4 holds. At the time of writing it reports 70 printed values reproduced exactly and 73,984 agreeing replays.

Figures. `figures.py` regenerates the three vector figures directly from exact rational evaluations of (6); no plotted quantity is fitted or measured.

Calculator. A companion exact-rational calculator accepts the record keys `y1, y2, eps1, eps2, B, e, s, J, H, threshold`, with optional `fresh_upper`, `provenance` and `protocol`, and returns the verdict of Proposition 6.1 together with the value of every expression in (6) and the active set. Decimal or rational strings preserve exact values. For the principal washed record,

```
{"y1":"6", "y2":"4", "eps1":"1/5", "eps2":"1/5", "B":"10",
"e":"1/20", "s":"9/10", "J":"2", "H":"1/5", "threshold":"8/5"}
```

returns $L = 169/100$, active expression *reserve-sensitive*, verdict *at or above*; changing `e` to `9/10` and `y2` to `57/10` returns $L = 38/25$ and *unresolved*. Following Proposition 3.6, a useful interface should make its "next measurement" advice depend on the active expression: if $e = s$, or if $J \geq [B - q_1]_+$, improving J cannot improve the bound at all, while on a uniquely active reserve-sensitive expression a unit reduction in J improves it by $s - e$ until the active expression changes. An input consuming the uptake premise of Theorem 5.3 must be a separate mode with its own recorded premise, not a relabelling of J.

Availability. The reported assay values are synthetic; no biological dataset was generated for this study. The companion archive contains the proof modules, verification receipts, calculator, example records, figure sources and the numerical audit. A versioned archive identifier, licence and author information are to be inserted before submission.

References

[1] A. Belenguer, S. H. Duncan, A. G. Calder, G. Holtrop, P. Louis, G. E. Lobley and H. J. Flint. Two routes of metabolic cross-feeding between *Bifidobacterium adolescentis* and butyrate-producing anaerobes from the human gut. *Applied and Environmental Microbiology* **72** (2006), 3593–3599. doi:10.1128/AEM.72.5.3593-3599.2006.

[2] S. H. Duncan, P. Louis and H. J. Flint. Lactate-utilizing bacteria, isolated from human feces, that produce butyrate as a major fermentation product. *Applied and Environmental Microbiology* **70** (2004), 5810–5817. doi:10.1128/AEM.70.10.5810-5817.2004.

[3] R. W. Leighty and M. R. Antoniewicz. Dynamic metabolic flux analysis (DMFA): a framework for determining fluxes at metabolic non-steady state. *Metabolic Engineering* **13** (2011), 745–755. doi:10.1016/j.ymben.2011.09.010.

[4] R. Muñoz-Tamayo, B. Laroche, E. Walter, J. Doré, S. H. Duncan, H. J. Flint and M. Leclerc. Kinetic modelling of lactate utilization and butyrate production by key human colonic bacterial species. *FEMS Microbiology Ecology* **76** (2011), 615–624. doi:10.1111/j.1574-6941.2011.01085.x.

[5] O. Pennington et al. A multiscale hybrid modelling methodology for cell cultures enabled by enzyme-constrained dynamic metabolic flux analysis under uncertainty. *Metabolic Engineering* **86** (2024), 274–287. doi:10.1016/j.ymben.2024.10.013.

[6] D. Vandeputte et al. Quantitative microbiome profiling links gut community variation to microbial load. *Nature* **551** (2017), 507–511. doi:10.1038/nature24460.