Chemical Memory under Load: Productive Inheritance and Hidden Operating States

Abstract

Can a chemical state be inherited while the species supporting its reproduction is continuously extracted, and which observations reveal the productive state under feedback? We answer these questions for two operating arrangements of one four-species resident chemistry. Under imposed extraction in $[0.009999, 0.01]$, a literal finite-count compartment protocol completes two growth–division–transfer–recovery cycles with joint probability at least $24482873/24500000$. Both ancestries survive, retained cells recover distinct chemical states, high-state ancestry exceeds an exact bound of approximately 0.69304534, and each batch exports more than five times its initial size. Extraction at least 0.031 separately excludes the designated high stationary state. A reservoir–consumer extension has two fully attracting positive equilibria with identical consumer proportions but uptake differing by a factor greater than 1.8186. An exact resident-field identity connects the arrangements without transferring their inheritance guarantees. We derive a sharper stationary feedback-load cap, composition dissipation, and finite-window readout bounds that recover information lost by proportions. The inheritance, attraction, balance and load-cap results have Lean proofs; additional observation consequences are proved conventionally. The conservative witness uses compartment size 6.5536×10^{22} and 4×10^9 compartments: these are constructive mathematical guarantees, not a calibrated laboratory design.

1 Introduction

Chemical memory has an operational requirement: a distinguishable state must survive the process that uses it. Bistability supplies candidate states, but does not by itself prove that a finite compartment can grow, divide, pass through a neutral bottleneck and recover while exporting the very species that supports its growth. A second difficulty arises when function is inferred from composition. Distinct productive states may look identical to an observer who records only relative consumer abundances.

We study these two difficulties in a common resident chemistry. First we impose a constant extraction coefficient and prove a joint finite-population event covering inheritance, enrichment, recovery and actual output. We then couple the same resident field to a renewable reservoir and self-limited consumers. This produces two locally attracting operating states whose consumer proportions agree exactly although abundance and uptake differ. Their relationship is an equality of resident vector fields: the imposed coefficient becomes the feedback quantity $R\sum_i X_i$. The feedback dynamics, their stability, and their measurements must still be analyzed separately.

Chemical bistability has a long history, including Schlögl's reaction models of nonequilibrium transitions [7]. Structural autocatalytic-set theory asks whether reactions are catalytically closed and can be sustained from a food set [5]; it does not supply the kinetic, stochastic or observation conclusions sought here. Compositional heredity under growth and division has likewise been studied directly. Matsubara et al. [6] analyze conditions for inherited compositional states in compartmentalized autocatalytic systems. Our contribution is the additional extraction load, with

actual complementary molecular division, intact-cell sampling and chemical recovery controlled on one finite probability law. We make no claim to originate compositional inheritance or chemical bistability.

Productive modes and their thermodynamic costs are treated by Despons [2], while Despons, De Decker and Lacoste [3] derive structural constraints on optimal autocatalytic flux. Here, signed production and gross service accounts establish material statements; they do not establish energetic efficiency. The quantitative witness deliberately exposes the cost of the sufficient proof bounds rather than assuming microscopic feasibility.

Loss of absolute scale in compositional observations is also well established. Gloor et al. [4] explain the statistical consequences of compositional microbiome data, and Vandeputte et al. [10] show the importance of absolute microbial load for quantitative profiling. The result here supplies a dynamical example: for each specified positive composition q, one fixed parameter set has two fully attracting equilibria with that same q and separated uptake. Exact transient and finite-window balances identify the additional measurements needed. The conclusions concern the stated mathematical source, not a biological identification of its reaction species.

Three distinctions organize the paper. Chemical readout is separate from the immutable tag used to track ancestry. Collected product under imposed extraction is separate from uptake supporting consumers. Finally, identical stationary proportions do not imply that every time-resolved composition record is uninformative: away from the target composition, derivatives can sometimes reveal total abundance. We state that exception explicitly. Core claims are supported by a frozen Lean development using mathlib [9]; ordinary proofs and exact certificates below explain their mathematical content, with computational provenance confined to the appendices and source package.

2 Model and operating protocol

2.1 Resident chemistry and extraction

A compartment has counts $n = (n_A, n_B, n_z, n_H) \in \mathbb{N}^4$ and integer size m. Concentrations are $x = n/m$. Set

$$e = 10^{-5}, \qquad \eta = 10^{-4}.$$

Table 1 specifies every chemical channel. Activities of the external supply species are maintained at the values implicit in these propensities. Falling factorials are essential: a two-molecule reaction has zero rate if fewer than two reactants are present. Extraction removes one actual intracellular z and increments the collector by one. It remains active during both batch growth and post-transfer recovery.

At fixed size, the zero-growth macroscopic drift is

$$\begin{aligned} \dot{A} &= 6 - 2A + Bz + 2e(B - A^2), \\ \dot{B} &= 27 + A - (1 + z)B - e(B - A^2), \\ \dot{z} &= A - Bz - 16z - 4z^2 + 3H - \rho z, \\ \dot{H} &= 16z + 2z^2 - (2 + \eta)H. \end{aligned} \tag{1}$$

The exact fixed-size count drift differs by

$$m^{-1}(2eA, -eA, 4z, -2z).$$

This finite-volume term and extraction shot noise are included in the stochastic estimates. Equation (1) is used for equilibrium geometry; it is not substituted for the finite-count source.

Channel	Reaction, including external species	Count propensity
1	$A \to B + z$	n_A
2	$B + z \to A$	$n_B n_z/m$
3	$z + F \to H$	$16 n_z$
4	$H \to z + F$	n_H
5	$H \to 2z$	n_H
6	$2z \to H$	$2n_z(n_z - 1)/m$
7	$R_A \to A$	$6m$
8	$A \to R_A$	n_A
9	$R_B \to B$	$27m$
10	$B \to R_B$	n_B
11	$B + G \to 2A$	$e n_B$
12	$2A \to B + G$	$e n_A(n_A - 1)/m$
13	$H \to W_H$	ηn_H
14	$z \to Z_{\text{out}}$	ρn_z

Table 1: The fourteen chemical channels. F, G, R_A, R_B, W_H are external supply or waste species, not additional resident coordinates. Growth is the separate event described below.

2.2 Growth, complementary division and ancestry

Each cell carries an immutable ancestral tag L or H. The tag is inherited by both daughters and does not enter a propensity. Let Q be the available growth precursor and let W_0 be the total size at batch preparation. A growth event has rate

$$\beta n_z, \qquad \beta = \gamma Q/(4W_0). \tag{2}$$

It consumes one Q and one z and increases m by one. At size $2N$, the resulting counts are partitioned: for each species choose a binomial half-sample d_i and place the complementary count $n_i - d_i$ in the other daughter. Both daughters have size N. Molecules are conserved across this division; the daughters are not independent copies of a chemical state.

Write $W = \sum m$ and let C_b, W_b be the cell count and total size of ancestry b. Until the batch endpoint, growth preserves

$$Q + W = 5W_0. \tag{3}$$

The batch starts at $Q = 4W_0$ and ends when $Q = W_0$, so its endpoint has $W = 4W_0$ and exactly $3W_0$ growth events. During the active batch, $\gamma/4 \le \beta \le \gamma$. Ancestral tags permit the statement of inheritance and selection on the same law, without identifying a chemical fluctuation with a change of ancestry.

2.3 Transfer, recovery and finite-law semantics

At the endpoint, choose M distinct live cells uniformly without replacement. Transfer those intact cells into fresh medium; discard the others with their molecules. There is no ancestry-dependent sampling or reset of retained intracellular counts. Disable growth during a recovery interval of length 5376, while keeping resident chemistry and extraction active. The next batch begins from the actual recovered counts, with Q replenished to four times the retained total size.

All probability statements below refer to finite stopped jump models. For a finite model with generator L, choose a clock q at least its total rate in every state. The transition kernel is $K = I + L/q$.

Its time-t law is the Poisson mixture

$$P_t = e^{-qt} \sum_{k \geq 0} \frac{(qt)^k}{k!} K^k. \tag{4}$$

The finite state space contains exit information and event counters. A departure from the admitted chemical region, missed endpoint, exhausted service quota or unsuccessful continuation contributes failure mass. Such mass is retained by the composition and is never renormalized away. Good exits preserve the literal final physical event and its collector increment.

This construction is an exact finite-law definition, not a simulation approximation. Its familiar continuous-time interpretation follows from finite-state uniformization. The formal result is stated for these kernels; it does not require a separately postulated unbounded-population process or a global nonexplosion theorem.

2.4 A second arrangement: reservoir–consumer feedback

Write f_ρ for the four-coordinate field (1). For a fixed positive probability vector $q = (q_1, \ldots, q_n)$, a second arrangement is

$$\dot{x} = f_\lambda(x), \quad \lambda = RS, \quad S = \sum_i X_i, \qquad \dot{R} = d(1 - R) - RzS, \quad \dot{X}_i = X_i(Rz - \mu - X_i/q_i), \tag{5}$$

with $d = 1/20$ and $\mu = 1/2$. Define the uptake $J = RzS = \lambda z$. Projection onto the resident coordinates gives exactly f_{RS}, with no finite-count correction. The feedback term removes the same resident species z as imposed extraction. Here J supports consumers; it is not a harvested-product counter. The coefficients $1/q_i$ encode the chosen composition: the assertion "for every q" constructs a model for that q, within which both equilibria share all parameters.

This identity connects the chemistry, not the protocols or probability laws. No division or transfer is imposed on the feedback communities, and stability with λ frozen would not establish stability of the coupled system. Theorem 8.1 proves the latter directly.

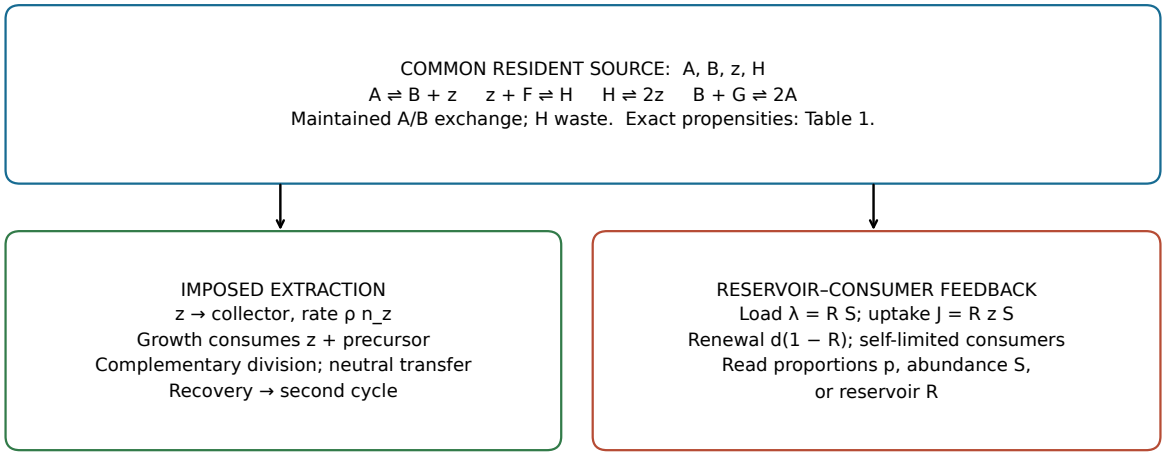

Figure 1: One resident chemistry, two operating arrangements. The left arrangement includes two literal copying cycles and external collection; the right includes reservoir renewal and consumer uptake. The connection is the resident-field identity $\lambda = RS$. It supplies no inheritance theorem for the consumers.

3 Constructive compatibility theorem

Define

$$\begin{gathered} N = 65536000000000000000000, \quad M = 4000000000, \\ \gamma = 10^{-11}, \quad T = 8 \times 10^{11}, \quad \varepsilon = 1/50, \\ b = \tfrac{3}{5}\log 4 - \tfrac{19}{500}, \quad \ell = \log(51/49), \quad g_* = 2(b-\ell) - \log 2. \end{gathered} \tag{6}$$

The deadline $T = 8/\gamma$ applies to each batch. All logarithms are natural. The memory regions are centered at two positive stationary states of (1). They are defined by the rational quadratic forms in Appendix A, with outer energy level $1/32000000$ and ready level $a_* = 1/512000000$.

Theorem 3.1 (Productive chemical memory)**.** *For every $\rho \in [9999/10^6, 1/100]$, there exist positive stationary states s_L, s_H with*

$$(s_L)_z \in [0.98172, 0.98174], \qquad (s_H)_z \in [2.89014, 2.89017].$$

Their outer energy regions are disjoint and attracting for the zero-growth deterministic drift. There is a balanced ready newborn preparation with $M/2$ cells of each ancestry such that the two-cycle law of Section 2 has a joint success event of probability at least

$$P_* = \frac{24482873}{24500000} = 0.999300938775\ldots. \tag{7}$$

On this event:

1. *Both batches reach their precursor endpoints within their deadlines, both transfers retain both ancestries, and the retained cells complete the actual recovery operation in their respective ready regions.*

2. *The final count log-odds gain exceeds g_*; hence the high-ancestry fraction satisfies*

$$f_H^{\text{final}} > \frac{\exp(g_*)}{1+\exp(g_*)} = 0.6930453421\ldots. \tag{8}$$

3. *The first batch exports more than $5NM$ molecules of z. Each batch, more generally, exports more than five times its own initial total size. In particular the separately recorded second batch output exceeds $5NM$.*

4. *All retained low cells have z concentration in $[0.97, 1]$, and all retained high cells have concentration in $[2.88, 3]$.*

Fixed finite service stocks can be selected before either cycle to realize the successful histories, including resident supply and waste, growth precursor and collection during recovery.

The publication Lean declaration directly packages attraction, disjointness, balanced preparation, the joint count/output event and the stationary exclusion below. Ready populations are part of the output type; the successful event has no failed option. Batch timing, actual transfer/recovery and material realizability are supplied by the source-law and companion declarations listed in Appendix D. The fraction in (8) is the elementary monotone conversion of the exact log-odds inequality. Its decimal display is a numerical evaluation, not the formal constant.

Theorem 3.2 (Stationary high-state exclusion)**.** *If $\rho \geq 31/1000$, there is no zero-growth stationary four-coordinate state with $z \geq 2$.*

This theorem excludes a specified stationary support for high memory. It does not assert universal finite-horizon memory failure, exclude all transient behavior, or identify the exact bifurcation. The gap between the constructive interval and 0.031 is not classified by these results.

Quantity	Exact guarantee or witness
Extraction interval	$[9999/10^6, 1/100]$
Initial preparation	$M/2$ cells per ancestry, all size N
Joint success probability	At least 24482873/24500000
Count log-odds gain	Greater than $2(\frac{3}{5}\log 4 - \frac{19}{500} - \log(51/49)) - \log 2$
First and second batch export	Each greater than $5NM$; source bound is $5W_0$ for each batch
Recovery time	5376 per cycle; extraction remains active
Generic ready initial states	Pay $2\log 2$, rather than $\log 2$, in the endpoint phase conversion
Upper extraction obstruction	No zero-growth stationary state with $z \geq 2$ at $\rho \geq 0.031$

Table 2: The final theorem's witnesses and scope. The smaller generic count guarantee is not used to obtain the newborn conclusion.

4 Chemical restoration from the literal source

4.1 Scalar equilibrium reduction

At a stationary point, the first equation plus twice the second gives $60 - (z+2)B = 0$, and the fourth gives H. Define

$$\begin{aligned} B(z) &= \frac{60}{z+2}, & H(z) &= \frac{16z + 2z^2}{2+\eta}, \\ K(z) &= \frac{2(1+2\eta)z^2 - 16(1-\eta)z}{2+\eta}, & A_\rho(z) &= zB(z) + K(z) + \rho z, \\ R_\rho(z) &= 27 - (1+e)B(z) + K(z) + \rho z + eA_\rho(z)^2. \end{aligned} \tag{9}$$

Substitution gives drift $(-2R_\rho, R_\rho, 0, 0)$. Conversely every stationary state with $z > 0$ has these reconstructed coordinates and satisfies $R_\rho(z) = 0$. Thus no unsupported scalar-to-vector correspondence is assumed.

For $\rho \geq 0, z > 0$, reconstructed B, H are positive, and

$$A_\rho(z) = \frac{z[20004(z-3)^2 + 48z + 700056 + 20001\rho(z+2)]}{(z+2)20001} > 0.$$

Moreover, for $s \geq r \geq 0$,

$$R_s(z) - R_r(z) = (s-r)z\{1 + e[A_s(z) + A_r(z)]\} \geq 0.$$

Exact rational endpoint evaluations therefore bracket a root in each stated interval uniformly in ρ. Continuity and the intermediate value theorem provide the two centers. Rational arithmetic, rather than floating-point root finding, establishes these brackets.

4.2 Local energies and deterministic attraction

For each center choose $E_b(y) = y^T P_b y$, with the fixed rational matrices in Appendix A. The source certificates establish

$$\frac{1}{200}\|y\|^2 \leq E_b(y) \leq 60\|y\|^2, \qquad 2y^T P_b F_\rho(s_b + y) \leq -\frac{59}{100}\|y\|^2 \tag{10}$$

whenever $|y_i| \leq 1/400$. The estimates include the nonlinear remainder uniformly over the extraction interval. They are checked by polynomial identities, rational bounds and nonnegative squares.

The outer energy bound implies $|y_i| \leq 1/400$, so (10) applies throughout the region. It gives inward drift and, conservatively, $E_b(y(t)) \leq E_b(y(0))e^{-t/1000}$. Local existence extended within the compact invariant region gives the global forward trajectory and convergence to the center. These are zero-growth deterministic attraction statements. The stochastic growing population requires the additional jump and division estimates below.

4.3 Finite-count exponential control

Put $\alpha = 10^{-12}$ and consider $V(n) = \exp(m\alpha E_b(n/m - s_b))$ at fixed m. The exact quadratic increment separates its linear term and its jump-square term. Extraction contributes the jump $-e_z$ at rate ρn_z, and hence both the drift correction and its positive quadratic noise contribution. Bounding the exponential remainder over the local region yields

$$LV \leq \alpha V\left[-\frac{m}{2}\|n/m - s_b\|^2 + 10^8\right]. \tag{11}$$

For growing cells use the fixed reference exponent $V_N(n, m) = \exp(N\alpha E_b(n/m - s_b))$. On the outer region, consumption of one z, change of denominator from m to $m+1$, and $\beta \leq \gamma$ give

$$LV_N \leq \alpha V_N\left[-\frac{N}{4}\|n/m - s_b\|^2 + 2 \cdot 10^8 + 8 \cdot 10^9 N\gamma^2\right]. \tag{12}$$

These source estimates are combined with a size-dependent exponential weight. Its affine decay coefficient uses denominator 960. At division, concentration fluctuations are controlled for the actual complementary daughter pair; the probability of a bad partition is bounded by

$$d_N = 8\exp[-N/(35 \cdot 10^{12})]. \tag{13}$$

The same partition draw supplies both daughters, so the proof does not multiply two independent success probabilities.

For a finite jump model, $LV \le 0$ implies $KV \le V$ and then $P_t V \le V$. The affine variant follows by the same iteration. An observable bounded below by a barrier on a bad set bounds the probability of that set. Applying these elementary finite-kernel arguments to the cell and population potentials gives the chemical error budget in Appendix B. At most $7M$ divisions occur before the admitted endpoint, and at most $8M$ cells are live. These finite geometric bounds make the union estimates uniform over actual histories.

Recovery uses the counts of the selected cells. Starting with energy at most $8a_*$, the fixed-size extraction kernel returns to energy at most a_* at time 5376 with failure probability at most 10^{-18} per cell. A separately paid finite-service allowance doubles this to $2 \cdot 10^{-18}$. The assertion concerns the terminal census after recovery, not merely a first entrance into the ready region. It is uniform over all permitted selected counts and sizes, which is what allows a second cycle to use the same estimate.

5 Selection, transfer and count enrichment

5.1 Ancestral size gain during a batch

Let $U = \log W_H - \log W_L$. Resident chemistry and extraction preserve cell sizes, and complementary division preserves each ancestry's total size. Only growth changes U. In the admitted regions the size-weighted ancestral z concentrations obey

$$.97 \le a_L \le 1, \qquad 2.88 \le a_H \le 3.$$

The two growth intensities are $\beta a_H W_H$ and $\beta a_L W_L$. The exact changes in U involve $\log(1 + 1/W_b)$, rather than their linear approximation. The source bounds give a lower log-odds drift 1.87β and an upper total-log-size drift 3β. An exponential tilt of the difference between U and $\frac{3}{5} \log W$ supplies a finite-jump deviation allowance $19/500$. Since the endpoint has $W/W_0 = 4$, the resulting batch gain exceeds

$$b = \frac{3}{5} \log 4 - \frac{19}{500},$$

except with probability $\exp(-19N/500000)$, in addition to the chemical and endpoint failures. A separate size-growth barrier gives the missed-deadline term $\exp(-N/2500)$ at $T = 8/\gamma$.

5.2 Neutral sampling at the actual endpoint

Condition on a successful batch endpoint. There are at most $8M$ cells, with sizes between N and $2N$. Uniformly choosing M distinct cells retains a random sum of bounded weights for each ancestry. If the initial ancestry counts are at least pM, the endpoint geometry gives the required lower ancestral membrane totals. The retained weighted sums have the exact without-replacement means and nonpositive pairwise covariances. Chebyshev bounds for the two ancestral sums give failure probability at most

$$D(p, M, \varepsilon) = \frac{32}{\varepsilon^2 pM}. \tag{14}$$

On the successful transfer event, the size-odds loss is at most

$$\log \frac{1 + \varepsilon}{1 - \varepsilon},$$

and each retained ancestry has count fraction at least $(1-\varepsilon)p/16$. Subsequent recovery preserves counts and sizes while restoring chemistry. Thus a complete successful cycle increases U by more than $b-\ell$.

The concentration estimate is conditional on the entire endpoint and uniform over it. No assumption that endpoint sizes or chemical histories are independent of ancestry is made. The proof uses the same finite subset sampling operation that defines the protocol.

5.3 Exact endpoint phase conversion

Define the mean-size phase

$$\phi = \log \frac{W_H/C_H}{W_L/C_L}.$$

For positive ancestral counts,

$$\log(C_H/C_L) = \log(W_H/W_L) - \phi, \qquad -\log 2 \le \phi \le \log 2. \tag{15}$$

The two cycle gains telescope in U; no intermediate phase correction is necessary. Generic initial and final populations can cost $2\log 2$. For the preparation in Theorem 3.1, every cell has size N, so $W_b = NC_b$ and $\phi_{\text{initial}} = 0$. Only the final $\log 2$ penalty remains. Consequently the count gain exceeds g_* in (6).

The preparation is nonvacuous. For each positive center, round $Ns_{b,i}$ down to an integer count. Every concentration error is at most $1/N$, so $E_b \le 240/N^2 \le a_*$. Form $M/2$ copies at each center, all size N. This is the same constructed start used by the probability theorem; the newborn premise is not imposed on an unrelated population.

With initially equal ancestral counts, the initial count log-odds is zero. If the final positive counts are H, L and $\log H - \log L > g_*$, exponentiation gives $H/L > e^{g_*}$. Cross-multiplication by positive denominators proves (8).

5.4 Joint confidence and proof of Theorem 3.1

The complete one-cycle failure bound has the form

$$\mathcal{E}(p) = \mathcal{R} + D(p, M, 1/50), \qquad \mathcal{R} \le 3 \cdot 10^{-6}.$$

The remainder includes chemistry, divisions, deadline, size-odds deviation, output, collection capacity, service and actual selected-cell recovery. It is uniform in p. The initial floor is $p_0 = 1/2$, and a successful first cycle supplies

$$p_1 = \frac{(1-1/50)(1/2)}{16} = \frac{49}{1600}.$$

Therefore

$$D(p_0, M, 1/50) = \frac{1}{25000}, \qquad D(p_1, M, 1/50) = \frac{4}{6125}.$$

Conditioning on each successful first output, the second kernel satisfies the same source theorem with floor p_1. Integration of this uniform bound gives joint success at least

$$1 - \mathcal{E}(p_0) - \mathcal{E}(p_1) \ge 1 - \frac{1}{25000} - \frac{4}{6125} - \frac{6}{10^6} = P_*.$$

The failed branch remains failed in the bind operation. No independence between the failures is required. The chemical lemmas supply return, the batch lemmas supply timing and output, and (15) supplies the final count gain. This proves the joint assertion. Finite resource realization is established in Section 6.

6 Output and material realization

6.1 A common-inventory output barrier

During a batch let $G = 4W_0 - Q$ count growth events and E count extraction events. Let $S_z = \sum n_z$ be the actual live z inventory. Their total intensities are βS_z and ρS_z, respectively. For

$$V = \exp(2G - E),$$

the literal generator is

$$LV = S_z V\{\beta(e^2 - 1) + \rho(e^{-1} - 1)\}. \tag{16}$$

All resident channels preserve G, E. Growth multiplies V by e^2, including when it triggers division, because division does not alter its counter increments. Extraction multiplies V by e^{-1}. Thus the common factor S_z is an exact source identity.

Since $e^2 \le 9$ and $e^{-1} \le 1/2$, the bracket is nonpositive when $16\gamma \le \rho$ and $0 \le \beta \le \gamma$. This condition holds throughout the constructive interval. Initially $V = 1$. At a precursor endpoint with $E \le 5W_0$, we have $G = 3W_0$ and $V \ge e^{W_0}$. The finite-kernel barrier bound therefore yields

$$\mathbb{P}\{Q = W_0,\ E \le 5W_0\} \le e^{-W_0} \le e^{-NM}. \tag{17}$$

This is an unconditional endpoint-and-low-output probability, not a conditional probability given endpoint attainment. The proof includes it in the same failure budget as the deadline and copying events. It requires neither an independent exporter nor a claim that large output alone implies survival.

6.2 Export versus internal production

Let Z be the total intracellular z count. Along an actual event history, let P_z be signed internal z formation by resident channels, G_z the number consumed by growth, D_z the number removed in discarded intact cells, and E physical extraction. Direct summation of event increments gives

$$P_z = Z_{\text{final}} - Z_{\text{initial}} + E + G_z + D_z. \tag{18}$$

The identity includes negative resident production increments, complementary division, intact-cell discard and extraction during recovery. Refilling external precursor does not change intracellular z.

At a ready batch start, the concentration ranges imply $Z_{\text{initial}} \le 3W_0$. Output greater than $5W_0$ already forces positive net internal formation. More sharply, on a completed batch segment without a precursor refill, $G_z = 3W_0$ by (3). Nonnegative final inventory and discard then imply $P_z \ge E$. This segment premise matters: a history containing refills cannot infer its total growth count from its final Q alone.

Equation (18) rules out explaining the batch export entirely by draining prepared intracellular z. It does not identify which atoms originated in fresh food: initially prepared A or H may contribute. Atom-labelled provenance and thermodynamic conversion efficiency are distinct questions.

6.3 Finite service and external maintenance

The formal resident element weights are $(2, 1, 1, 2)$ for (A, B, z, H). The external weights, in order $(F, G, R_A, R_B, W_H, Z_{\text{out}}, Q)$, are $(1, 3, 2, 1, 2, 1, 1)$. Every chemical channel balances these weights. A size material unit has weight two and is formed from one Q and one z. Complementary division and

intact-cell transfer preserve the corresponding material accounts, with discarded material explicitly removed.

Maintained activities are an operating assumption; finite stock is not a claim of autonomous mass-action buffering. To realize the finite event histories, each service counter bounds gross exchanges of every named reservoir. Each event changes each gross account by at most one, so a successful event quota bounds all these accounts without cancelling supply against waste.

For the exact law used in the theorem, choose a collection capacity $J > 0$ satisfying

$$\frac{T\,560NM\rho}{J} \le 10^{-6}.$$

The domain has at most $8M$ cells with $n_z \le 70N$, which gives the numerator. Choose a finite batch clock q_s on each finite domain and a service quota J_s with $Tq_s/J_s \le 10^{-6}$. For each possible retained cell choose a recovery clock q_c and quota J_c with $5376q_c/J_c \le 10^{-18}$. These choices also meet the decay-clock lower requirements.

Let B be one plus the sum of batch quotas over all possible ready starts and let R be one plus the sum of recovery quotas over the finite cell box. Then

$$\mathcal{S} = 2(B + MR) \tag{19}$$

is fixed before the experiment and covers the gross services for two successful cycles. This is the compiled finite resource witness. It is a finite-state sum, not an efficient numerical implementation. Unsaturated hard-budget and ideal event histories agree exactly, so this construction is more than an expectation bound on net resource use.

For newborn initialization, the first precursor preparation is $4NM$ and the second is at most $8NM$, hence total prepared precursor is at most $12NM$. Actual growth consumption is at most $9NM$. These are distinct from the gross resident-service bill. Recovery extraction has its own literal mark and chemical marginal projection; it is included in the stock account even though the main output predicate records only the two batch marks.

7 Stationary obstruction and the extraction gap

To prove Theorem 3.2, multiply $R_{31/1000}(z)$ by the positive denominator $4444888900000000000(z+2)^2$. With $y = z - 5/2$, the resulting polynomial can be written

$$\begin{aligned} P(y) &= q(y) + 268341565475331326 06(y + 1/2)y^2 \\ &\quad + 4448524228728822329y^4 + 136356292472000(y + 1/2)y^4 + 44462224000000y^6, \\ q(y) &= \frac{2646507285317916225}{16} - \frac{373906577279071365}{2}y + \frac{34457435637128464441}{2}y^2. \end{aligned}$$

The quadratic has positive leading coefficient and negative discriminant, verified over the integers. Thus $q(y) > 0$. If $z \ge 2$, then $y + 1/2 \ge 0$, so every remaining term is nonnegative. Hence $R_{31/1000}(z) > 0$. Monotonicity of the residual in ρ proves $R_\rho(z) > 0$ for all larger extraction rates. The stationary reduction then proves the exclusion.

Numerical continuation of the same deterministic source suggests a fold near $\rho = 0.03026905764$. Figure 2 distinguishes those numerical branches from the certified operating band and exclusion region. Neither the exact fold nor a finite-count escape law is needed by the two theorems. A vanishing deterministic restoring rate can guide a future proof, but is not a lower bound on the probability of memory failure in a finite experiment.

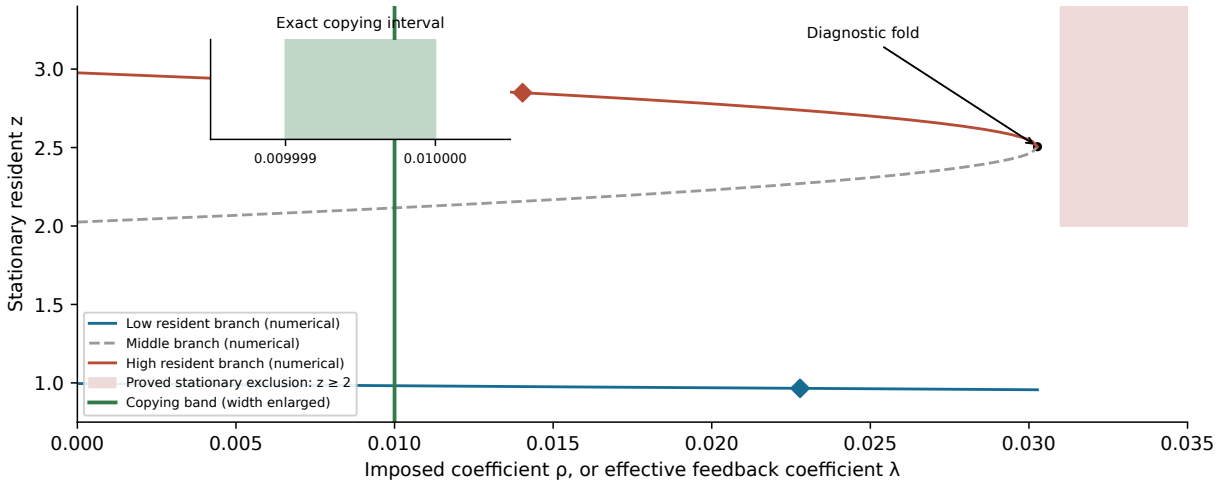

Figure 2: Extraction changes the support of chemical memory. Curves are floating-point zero-growth continuation; dots locate the feedback equilibria at their effective loads. The green operating band is certified; its width is enlarged in the main panel and shown exactly in the inset. The shaded region begins at the proved stationary high-state exclusion $\rho = 0.031$. The intervening region is not classified by the compatibility theorem.

8 Reservoir source and two attracting communities

Let $q_i > 0$, $\sum_{i=1}^n q_i = 1$, and $S = \sum_i X_i$. The full state is $(A, B, z, H, R, X_1, \dots, X_n)$. With $e = 10^{-5}$, the literal field is

$$\begin{aligned} A' &= 6 - 2A + Bz + 2e(B - A^2), & B' &= 27 + A - (z+1)B - e(B - A^2), \\ z' &= A - Bz - 16z + 3H - 4z^2 - RzS, & H' &= 16z + 2z^2 - 2.0001H, \\ R' &= \tfrac{1}{20}(1 - R) - RzS, & X_i' &= X_i(Rz - \tfrac{1}{2} - X_i/q_i). \end{aligned} \tag{20}$$

The functional observable is the actual uptake $J = RzS$.

Theorem 8.1 (Equal composition, different function)**.** *For every positive probability vector q, system (20) has two positive locally attracting equilibria c_-, c_+ with $X_i/S = q_i$ at both equilibria, but $S_- < S_+$ and $J_- < J_+$. Their total abundances lie respectively in*

$$S_- \in [0.04067409, 0.04067412], \qquad S_+ \in [0.07014930, 0.07014933]. \tag{21}$$

Local attraction is in the full community state space, including perturbations of consumer proportions. Nearby initial states admit positive global source trajectories converging to the indicated equilibrium.

Proof. At a positive equilibrium on the prescribed-composition manifold, elimination gives

$$\begin{aligned} J(s) &= s(1/2 + s), & R(s) &= 1 - 20J(s), \\ z(s) &= (1/2 + s)/R(s), & B(s) &= 60/(z(s) + 2), \\ H(s) &= (16z(s) + 2z(s)^2)/2.0001, \\ A(s) &= B(s)z(s) + 16z(s) + 4z(s)^2 - 3H(s) + J(s). \end{aligned} \tag{22}$$

Define $F(s) = A(s) + B(s) + e(A(s)^2 - B(s)) - 33$. This is the same resident stationary relation restricted to feedback balance:

$$F(s) = R_{\lambda(s)}(z(s)), \qquad \lambda(s) = R(s)s.$$

Indeed $\lambda z = J$ and $(z+2)B = 60$ turn the residual in (9) into the displayed F. Thus the two equilibrium constructions are restrictions of one resident relation, while their stability problems remain different.

The aggregate field has residual $(-2F(s), F(s), 0, 0, 0, 0)$. All denominators are positive for $0 \le s \le 3/40$. Exact rational evaluation gives opposite endpoint signs on each interval in (21). The intermediate value theorem supplies a root in each interval. Lifting by $X_i = q_i s$ gives positive full-community equilibria and the claimed composition. Since $s \mapsto s/2 + s^2$ is strictly increasing there, uptake is strictly separated as well.

For attraction let $u = (A, B, z, H, R, S)$ and let u_* be either aggregate equilibrium. The rational matrices P_-, P_+ and enclosing boxes in Appendix C satisfy, throughout the corresponding box,

$$\frac{1}{1000}\|v\|_2^2 \le v^\top P v \le 200\|v\|_2^2, \qquad (v^\top P v)' \le -\frac{1}{400} v^\top P v \tag{23}$$

for the six-dimensional field and $v = u - u_*$. The derivative bound is a secant-field identity on a box, not an eigenvalue computation at an approximate root.

To include all consumer modes, write $y_i = X_i/q_i - S$ and $W = \sum_i q_i y_i^2$. Then $\sum_i q_i y_i = 0$, and summing the consumer equations gives

$$S' = S(Rz - 1/2 - S) - W. \tag{24}$$

The aggregate field differs from the six-dimensional field by $-W$ in its last coordinate. Use the full-community energy

$$E = v^\top P v + 100W. \tag{25}$$

In a sufficiently small neighborhood, the magnitude of the aggregate energy's derivative in the S direction is at most one, $Rz - 1/2 - 2S \le -0.03$, and $y_i \ge -0.01$. Indeed these hold strictly at the center, where $S > 0.03$ and $y = 0$. Direct differentiation gives $W' = 2(Rz - 1/2 - 2S)W - 2\sum_i q_i y_i^3 \le -0.04W$. Together with (23), this yields $E' \le -E/400$. The energy is nonnegative and bounds every full-state coordinate error squared by $2000E$. Choose a small sublevel inside the positive neighborhood and the certified box. A smooth cutoff constructs a global solution; the energy barrier keeps it in the sublevel, where the cutoff equals one. Thus it is an actual source solution, remains positive, and converges to the center. This proves local attraction in the full state space. □

9 What measurements determine function?

For $S > 0$ put $p_i = X_i/S$ and $\chi_q^2(p) = \sum_i (p_i - q_i)^2/q_i$. The weighted dispersion obeys $W = S^2\chi_q^2(p)$.

Theorem 9.1 (Transient and finite-window readout)**.** *Along every actual differentiable source trajectory with $S > 0$,*

$$J = S' + \tfrac{1}{2}S + S^2 + S^2\chi_q^2(p). \tag{26}$$

More generally, the same formula with W in place of $S^2\chi_q^2(p)$ is valid even at $S = 0$. For an abundance measurement M, suppose $0 \le S \le S_{\max}$, $|S'| \le \nu$, $0 \le W \le S^2\xi$ and $|M - S| \le \varepsilon$, with $\xi, \varepsilon \ge 0$. Then

$$\left|J - (M/2 + M^2)\right| \le \nu + S_{\max}^2\xi + (1/2 + 2S_{\max})\varepsilon + \varepsilon^2. \tag{27}$$

On any finite source interval $[a, b]$,

$$\int_a^b J\,dt = S(b) - S(a) + \frac{1}{2}\int_a^b S\,dt + \int_a^b S^2\,dt + \int_a^b W\,dt, \tag{28}$$
$$= \frac{1}{20}\int_a^b (1-R)\,dt - [R(b) - R(a)].$$

Proof. Rearrange (24) and use $RzS = J$. The dispersion identity follows by $X_i/q_i - S = S(p_i/q_i - 1)$. If $e_M = M - S$, then $M/2 + M^2 - (S/2 + S^2) = (1/2 + 2S)e_M + e_M^2$. The triangle inequality proves (27). Source differentiability gives continuity and integrability of every displayed polynomial observable; integrating the total-consumer and reservoir equations proves (28) by the fundamental theorem of calculus. □

At the two equilibria, $S' = 0$ and $p = q$, so abundance alone gives $J = S/2 + S^2$. The rational intervals (21) imply

$$J_- \in [0.0219914265973281,\ 0.0219914440377744], \tag{29}$$
$$J_+ \in [0.0399955742904900,\ 0.0399955934994489]. \tag{30}$$

Their guaranteed gap is $g = 0.0180041302527156$. The threshold $\tau = 0.0309935091641322$ classifies a measurement into the low or high state whenever its absolute uptake error is strictly below $g/2 = 0.0090020651263578$. Conversely, any deterministic predictor based only on the common proportions has the same value at both states; at least one of its two errors is at least half the actual uptake gap, hence at least $g/2$. Equation (27) shows precisely which additional transient, dispersion and abundance-error terms must be paid away from equilibrium.

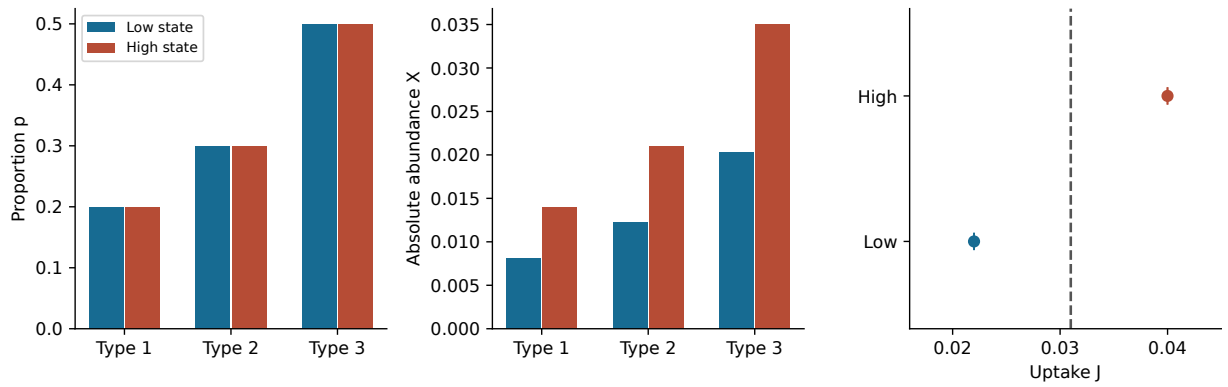

Figure 3: The same proportions $q = (0.2, 0.3, 0.5)$ give distinct absolute abundances and uptake. Abundance bars use interval midpoints for display only. Uptake error bars are the images of the exact rational root intervals; they are narrower than the markers. The dashed threshold has certified separation margin $g/2$.

10 Feedback load, composition relaxation and observation

10.1 A sharper stationary supply bound

The common resident identity permits a precise comparison between imposed extraction and stationary feedback. Retaining consumer self-limitation sharpens the elementary bound $RS \le d/(4\mu) = 0.025$.

Proposition 10.1 (Stationary feedback-load cap)**.** *At every nonnegative equilibrium of* (5)*, including equilibria with absent consumer types,*

$$\lambda = RS \le U := \frac{2239911}{97656250} = 0.02293668864 < 0.023 < 0.031. \tag{31}$$

If all consumers are positive, their proportions equal q*, and the necessary stationary curve is*

$$J = S/2 + S^2, \qquad \lambda = f(S) := S - 10S^2 - 20S^3. \tag{32}$$

Proof. Stationarity gives $J = S/2 + S^2 + W$ and $R = 1 - 20J$, where $W \ge 0$. Since $S \ge 0$,

$$RS = f(S) - 20SW \le f(S).$$

Put $a = 11/250$ and $b = 1 - 20a - 60a^2 = 12/3125$. The exact identity

$$U - f(S) = 10(S - a - b/20)^2 + 20(S + 2a)(S - a)^2 \tag{33}$$

proves $f(S) \le U$ on $S \ge 0$. For a positive equilibrium each consumer equation gives $X_i/q_i = Rz - 1/2$. Summing and using $\sum q_i = 1$ gives $X_i = q_i S$ and $W = 0$, proving (32). □

The scalar maximum of the necessary curve occurs at $S_* = (-\mu + \sqrt{\mu^2 + 3d})/3$, with maximum approximately 0.0229366115755. Thus U is a close rational upper bound, not the exact maximum of the attainable stationary source loads. Not every point on the curve solves the resident equilibrium equations.

Exact interval arithmetic applied to the root enclosures yields the outward enclosures

$$\begin{aligned} \lambda_- &\in [0.0227844647070,\ 0.0227844673247], \\ \lambda_+ &\in [0.0140360482612,\ 0.0140360692085]. \end{aligned} \tag{34}$$

The high operating state has greater uptake but smaller load coefficient, because $J = \lambda z$ and its z is larger. Both load intervals lie outside the copying band. The cap shows that stationary feedback cannot attain the imposed-extraction exclusion coefficient at these supply parameters; it guarantees neither existence nor inheritance. It is not a transient bound: $R = 1, S = 0.1$ gives $RS = 0.1$.

10.2 Why proportions can relax without identifying function

The consumer field yields a replicator-type equation, in the classical sense of frequency dynamics [8]. Its source-specific consequences explain both the loss of information and a transient exception.

Proposition 10.2 (Composition dynamics)**.** *Along a source trajectory with* $S > 0$ *set* $p_i = X_i/S$*,* $r_i = p_i/q_i$*, and* $\kappa = \sum_i p_i^2/q_i = 1 + \chi_q^2(p)$*. Then*

$$p_i' = Sp_i(\kappa - r_i), \tag{35}$$

$$\frac{d}{dt}\chi_q^2(p) = -2S\sum_i p_i(r_i - \kappa)^2 \le 0. \tag{36}$$

If all $p_i > 0$*,* $D(q\|p) = \sum_i q_i \log(q_i/p_i)$ *satisfies*

$$D' = -S\chi_q^2(p). \tag{37}$$

In particular $p = q$ *is invariant, even while* S *and* J *evolve.*

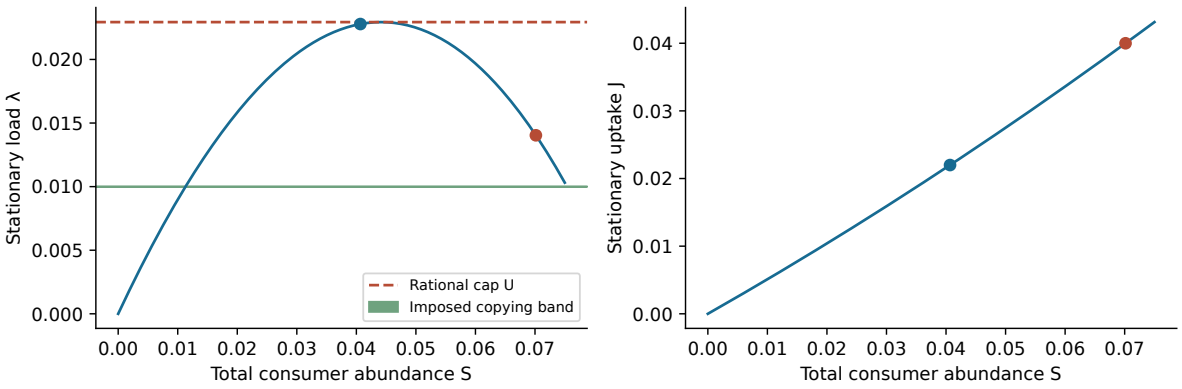

Figure 4: Necessary stationary load and uptake as functions of abundance. Markers use midpoints of the certified abundance enclosures. The shaded copying band belongs to imposed extraction; U is the rational cap proved in (33). The curves are balance relations and do not assert an equilibrium at every point.

Proof. Summing the consumer equations gives $S' = S(Rz - \mu) - S^2\kappa$. The quotient rule cancels the common per-capita term $Rz - \mu$ and proves (35). Differentiating $\kappa - 1$ gives $2S(\kappa^2 - \sum_i p_i r_i^2)$. Since $\sum_i p_i r_i = \kappa$ and $\sum_i p_i = 1$, this equals (36). Finally $D' = -\sum_i q_i p_i'/p_i = -S(\kappa - 1)$. The smooth source equations and uniqueness preserve $p = q$ on every interval with positive S. □

For a quantitative consequence, let $a_0 = \min_i p_i(0)/q_i > 0$. Since $\kappa \geq 1$ and $a_0 \leq 1$, equation (35) points inward on each face $r_i = a_0$. Scalar comparison gives $r_i \geq a_0$. Hence

$$\sum_i p_i(r_i - \kappa)^2 \geq a_0 \sum_i q_i(r_i - \kappa)^2 = a_0\{\chi_q^2(p) + (\kappa - 1)^2\} \geq a_0\chi_q^2(p),$$

and integration yields

$$\chi_q^2(p(t)) \leq \chi_q^2(p(0)) \exp\left[-2a_0 \int_0^t S(s)\, ds\right]. \tag{38}$$

On a global positive trajectory with $\int_0^\infty S = \infty$, composition therefore converges to q. No uniform rate is asserted as abundance vanishes or when a type is initially absent. This argument does not prove convergence of every trajectory to one of the two operating states.

Transient proportions can nevertheless reveal abundance away from q. Write $v_i(p) = p_i(\kappa - p_i/q_i)$. For positive $p \neq q$, $v \neq 0$ (otherwise all p_i/q_i coincide). Thus

$$S = \frac{v(p) \cdot p'}{\|v(p)\|_2^2}. \tag{39}$$

With exact p and derivative error at most $\epsilon_{\dot{p}}$ in Euclidean norm, the inferred abundance error is at most $\epsilon_{\dot{p}}/\|v(p)\|_2$. This follows from Cauchy–Schwarz and becomes poorly conditioned near q. Error in p requires further propagation. Composition-only observation is therefore insufficient uniformly and exactly uninformative on target-composition trajectories; it is not uninformative for every possible movie.

10.3 Finite-precision nonidentification

Proposition 10.3 (Overlapping observations)**.** *Suppose two states have uptake gap $\Delta J > 0$ and proportion vectors satisfying $\|p_+ - p_-\| \leq 2\epsilon_p$. Under observation error at most ϵ_p, every*

deterministic uptake predictor has error at least $\Delta J/2$ on one of the states for an admissible common observation.

Proof. The midpoint of the two proportion vectors lies in both observation balls. If the predictor gives a there, the triangle inequality gives $\Delta J \le |J_+ - a| + |a - J_-|$. At least one error is at least $\Delta J/2$. □

The exact pair already satisfies this conclusion for every $\epsilon_p \ge 0$. For fixed q and any fixed $\epsilon_p > 0$, sufficiently small smooth parameter perturbations also preserve an overlap and a positive uptake gap. Indeed the strict quadratic dissipation implies a Hurwitz, hence invertible, full Jacobian at either equilibrium. The implicit function theorem continues the two equilibria, and local stability and the observed coordinates vary continuously. Shrinking the parameter neighborhood preserves a gap at least $g/2$ and observation distance at most $2\epsilon_p$, giving prediction error at least $g/4$. This is a qualitative conventional consequence, not an explicit parameter-tolerance certificate.

11 Measuring uptake over a recovery window

The identities in Theorem 9.1 avoid numerical differentiation. To use them, measurement and integration errors must be stated in flux units.

Proposition 11.1 (Window error budgets)**.** *Let $h = b - a > 0$ and $\bar{J} = h^{-1}\int_a^b J\,dt$. Assume throughout the window $|M_S - S| \le \epsilon_S$, $0 \le S \le S_{\max}$, and $|M_W - W| \le \epsilon_W$. If Q_h approximates the time average of $\mu M_S + M_S^2 + M_W$ with error at most e_Q, define*

$$\widehat{J}_S = \frac{M_S(b) - M_S(a)}{h} + Q_h[\mu M_S + M_S^2 + M_W],$$
$$E_S(h) = \frac{2\epsilon_S}{h} + (\mu + 2S_{\max})\epsilon_S + \epsilon_S^2 + \epsilon_W + e_Q. \tag{40}$$

Then $|\widehat{J}_S - \bar{J}| \le E_S(h)$. Similarly, if $|M_R - R| \le \epsilon_R$ and the quadrature error in the average reservoir-derived flux is at most $e_{Q,R}$, then

$$\widehat{J}_R = d\{1 - Q_h[M_R]\} - \frac{M_R(b) - M_R(a)}{h},$$
$$|\widehat{J}_R - \bar{J}| \le E_R(h) := d\epsilon_R + 2\epsilon_R/h + e_{Q,R}. \tag{41}$$

Consequently $|\widehat{J}_S - \widehat{J}_R| \le E_S(h) + E_R(h)$.

Proof. Divide the two integrated source identities by h. The endpoint errors contribute at most $2\epsilon_S/h$ or $2\epsilon_R/h$. Write $M_S = S + e$: then $|M_S^2 - S^2| \le 2S_{\max}\epsilon_S + \epsilon_S^2$. Average the pointwise bounds and add the quadrature error. The reservoir integrand is linear with coefficient d. The final assertion is the triangle inequality. □

The rates d, μ are assumed known. If an integration routine bounds the unnormalized integral error, divide that bound by h. If W is omitted, replace ϵ_W with a bound on its average. In particular, recording only mean abundance misses temporal variance:

$$\bar{J} = \frac{S(b) - S(a)}{h} + \mu\bar{S} + \bar{S}^2 + \mathrm{Var}_{[a,b]}(S) + \bar{W}.$$

Longer windows reduce endpoint error but do not remove systematic bias, temporal variability or dispersion. For a Lipschitz integrand with constant L_g on a mesh of maximum width Δt, the left-rectangle average error is at most $L_g\Delta t/2$: sum the integral bounds $L_g\Delta t_i^2/2$ and divide by h. A sampled trace alone supplies no such bound.

11.1 Local recovery and a decision with explicit assumptions

The full-community energy proof gives, inside an admitted invariant sublevel,

$$\mathcal{E}(t) \le \mathcal{E}_0 e^{-t/400}, \qquad |u_j(t) - u_{*,j}| \le \sqrt{2000\mathcal{E}_0}\, e^{-t/800}. \tag{42}$$

Choose a common bound B on the absolute values of R, z, S and their equilibrium values there. Expanding the difference of the three-factor products gives

$$|J(t) - J_*| \le A_J e^{-t/800}, \qquad A_J = 3B^2\sqrt{2000\mathcal{E}_0}. \tag{43}$$

For $A_J, \delta_J > 0$ a sufficient deadline is $\max\{0, 800\log(A_J/\delta_J)\}$. If $A_J = 0$, the error is already zero. These are local source guarantees, not global basin assertions. For a window beginning at t_0, integrating (43) gives

$$|\bar{J} - J_*| \le \delta_{\mathrm{dyn}} := A_J e^{-t_0/800}\frac{800}{h}(1 - e^{-h/800}). \tag{44}$$

The threshold $\tau = 0.0309935091641322$ is a guaranteed branch rule under the two certified alternatives if $E_S + \delta_{\mathrm{dyn}} < g/2$, or with E_R in place of E_S. The recovery sublevel, calibration and two-alternative hypothesis are required inputs. During an unclassified transition, the correct output is unresolved.

Assumed quantity	Value	Role
h, $S_{\max}$	100, 0.08	Window and abundance ceiling
ϵ_S, ϵ_W	0.001, 0.0001	Pointwise measurement ceilings
ϵ_R	0.005	Reservoir measurement ceiling
e_Q, $e_{Q,R}$	10^{-5} each	Average-flux quadrature ceilings
E_S, E_R	0.000791, 0.00036	Resulting uptake-error bounds
δ_{dyn}	0.001	Separately admitted recovery bound
$g/2 - (E_S + \delta_{\mathrm{dyn}})$	0.0072110651263578	Remaining separation margin

Table 3: Illustrative error assumptions and exact resulting arithmetic; these are not instrument specifications.

For example, estimates $\widehat{J}_S = 0.0224$ and $\widehat{J}_R = 0.0222$ differ by $0.0002 < E_S + E_R$. With the assumptions of Table 3, both support the low alternative. Estimates 0.0402 and 0.0401 similarly support the high alternative. Neither conclusion is issued if the local-recovery or two-alternative hypothesis is absent. The supplied calculator also rejects inconsistent balances or observations inconsistent with both stationary intervals. These synthetic examples test the inference rule; they are not simulated or measured chemical data.

12 Quantitative interpretation and remaining scope

12.1 What the constructive witness costs

The count normalization is $x_s = n_s/m$; m is not the total number of resident molecules. If $x_s = C_s/C_*$ and physical time is $\tau = t_* t$, compatible volume is $m/(N_A C_*)$ litres, with $N_A = 6.02214076 \times 10^{23}\,\mathrm{mol}^{-1}$. At the illustrative choice $C_* = 1$ mM, a newborn compartment of size N corresponds to approximately 108.8 litres, and the initial population to 4.35×10^{11} litres. No concentration or time calibration is inferred from the dimensionless theorem.

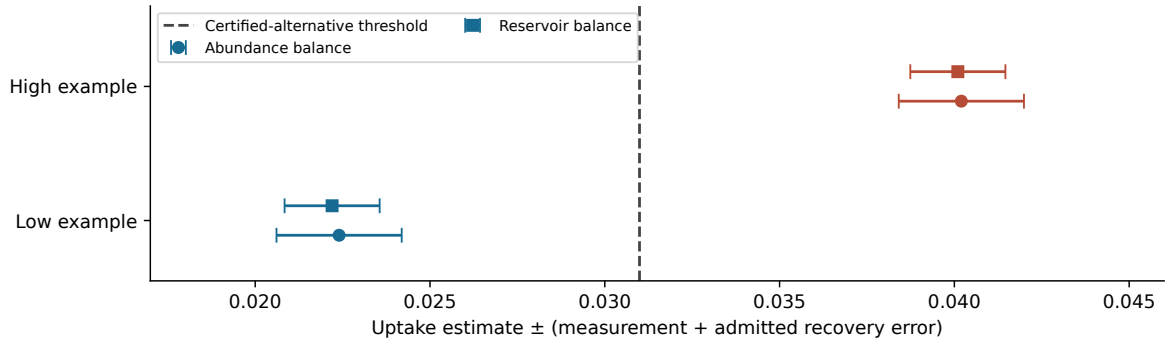

Figure 5: Synthetic low and high readout examples with the stated total errors $E+\delta_{\mathrm{dyn}}$. The vertical line is the separating threshold. Omitting the recovery or two-alternative assumptions produces an unresolved result regardless of apparent separation. No trajectory or instrument performance is inferred from this plot.

Two deadlines plus two recoveries give $2T + 10752 = 1.600000010752 \times 10^{12}$ model time units, about 50701 years if $t_* = 1$ second. Choosing a smaller time unit rescales all rates and leaves $\rho/\gamma \simeq 10^9$ unchanged. A physical bimolecular coefficient would scale as its dimensionless value divided by $t_* C_*$; a first-order coefficient scales by $1/t_*$. Rescaling is therefore not a proof of feasible kinetic constants. The chemical witness has $N\alpha a_* = 128$, with $(\alpha a_*)^{-1} = 5.12 \times 10^{20}$. This exposes

Quantity	Interpretation
$N = 6.5536 \times 10^{22}$	Sufficient compartment-size normalization, not a measured minimum
$M = 4 \times 10^9$	Number retained at each neutral bottleneck
$[0.009999, 0.01]$	Certified extraction width 10^{-6}, about 100 ppm of 0.01
$1/25000 + 4/6125$	Combined upper transfer-failure budget
6×10^{-6}	Combined remaining upper failure budget
Prepared growth precursor	At most $12NM$ across both cycles
Growth consumption	At most $9NM$ across both batches
Gross resident service	Fixed finite witness (19); no efficiency claim

Table 4: Scale and account conventions for imposed extraction. Failure budgets are analytical upper bounds, not measured frequencies.

the conservative restoration and noise estimates that drive the count scale. Transfer dominates the displayed confidence loss, but improving a sampling inequality alone does not remove the chemical scale. Sharper sampling bounds are available in the literature [1]; applying them here requires a new proof for the actual weighted endpoint and a recomputed joint budget.

The successful chemical ranges also supply a simple readout: threshold 1.94 tolerates an absolute error strictly below 0.94 for every retained-cell z concentration. Ancestry and chemical state should still be recorded separately. A small per-cell error probability does not certify correct readout of all four billion retained cells.

12.2 Finite inheritance and the infinite-horizon distinction

The formal publication root covers two actual cycles. Its underlying one-cycle theorem, however, permits every $p > 0$, and the nontransfer remainder is uniform in p. This supports the following

additional conventional consequence for a specified finite repetition of the same kernel.

Corollary 12.1 (Finite repetition of the source kernel)**.** *Fix the parameters and balanced newborn preparation of Theorem 3.1. Compose $K \geq 1$ copies of the same complete cycle kernel, retain failure mass, and record every batch output. Let $\mathcal{A}_K$ mean that all K cycles satisfy their ready-state, ancestry-floor, size-gain and output relations. With $c = 49/800$,*

$$\mathbb{P}(\mathcal{A}_K) \geq \max\left\{0, 1 - \frac{1}{25000}\sum_{j=0}^{K-1} c^{-j} - \frac{3K}{10^6}\right\}. \tag{45}$$

On this event both ancestries survive, each batch exports more than $5W_0$ for its own start, and the final count log-odds gain exceeds $K(b-\ell) - \log 2$. A fixed finite service stock $K(B + MR)$ suffices, with the quota maxima replaced by the finite sums defining (19).

Proof. The cycle output relation supplies readiness and ancestry fractions at least $(1-1/50)p/16 = cp$, as well as a size log-odds increment greater than $b-\ell$ and the actual batch output. Thus a successful history before cycle k has floor $p_k = \frac{1}{2}c^{k-1} > 0$. The source theorem applies to each such history, with failure at most $32/((1/50)^2 p_k M) + 3/10^6 = c^{-(k-1)}/25000 + 3/10^6$. Conditional integration and induction lose at most the sum of these errors; nonnegativity gives the outer maximum. Readiness preserves the admitted state space. Recording finitely many marks adds only a finite history coordinate, and failure remains absorbing. The size gains telescope; equal initial sizes and the final phase bound give the stated count conversion. Each ready start and selected cell is covered by the same finite service sums, so K times the one-cycle stock bound covers every successful history. □

The three- and four-cycle lower bounds are respectively

$$\frac{2373714351}{2401000000} = 0.9886357147\ldots, \qquad \frac{23957912563}{29412250000} = 0.8145555869\ldots.$$

These are conventional corollaries of the verified source kernel, not additional compiled history-root declarations. Their finite iteration is distinct from an indefinite-success claim; for sufficiently large K this sufficient lower bound becomes zero.

A complementary conventional obstruction applies to indefinite retention of two immutable tags. Suppose every successful batch ends with $2M \leq L \leq 8M$ cells, and each uniform bottleneck retains M. Some tag class has at least M cells, so at least one of the $\binom{L}{M}$ equally likely subsets is monochromatic. Conditional tag-loss probability is therefore at least

$$q_M = \binom{8M}{M}^{-1} > 0.$$

With no reintroduction of a lost tag, conditional iteration gives probability of uninterrupted joint success through K cycles at most $(1-q_M)^K \to 0$. No independence is needed. The endpoint range follows here from sizes in $[N, 2N]$, M starting cells and total size $4W_0$: $W_0 \in [NM, 2NM]$ implies $2M \leq L \leq 8M$. Earlier batch failure already violates joint success. This is not a practical extinction-time estimate, and does not preclude long inheritance in an individual lineage or alternative retention rules.

12.3 What is established and what remains separate

The imposed-load result controls a single finite event encompassing output and inheritance. Its output barrier exploits the shared live inventory, while the exact history balance rules out mere

depletion of initial z as the explanation of the exports. It does not identify fresh-food atom provenance or establish thermodynamic efficiency. Recovery is a serviced intervention with growth disabled; maintained activities and finite gross stocks are not an autonomous depleting-reservoir model.

The feedback result controls the full community, including composition perturbations. Its two equilibria demonstrate that the same proportions can conceal different productive operation under the same model parameters. Absolute abundance determines stationary uptake, while transient inference requires accumulation and dispersion or a reservoir measurement. Composition relaxation and the derivative exception clarify why an observation sufficient on one subset may fail uniformly.

The exact resident identity makes the two results comparable without merging their laws. The stationary cap lies below the high-state exclusion coefficient, but neither certified feedback load belongs to the imposed copying interval. Proving inheritance of the feedback communities remains a separate task. Other open targets are substantially smaller finite-count witnesses, quantitatively robust parameter boxes, a certified fold location, and global basins or finite-time loss probabilities. The present paper establishes the stated finite operation and observation guarantees without assuming any of those extensions.

A Rational energy certificates

In coordinate order (A, B, z, H), the energy matrices are exactly

$$P_L = 10^{-6} \begin{pmatrix} 1079122 & -568413 & 2227359 & 3266330 \\ -568413 & 1057785 & -2194388 & -3215160 \\ 2227359 & -2194388 & 6383792 & 9601239 \\ 3266330 & -3215160 & 9601239 & 14651126 \end{pmatrix},$$

$$P_H = 10^{-6} \begin{pmatrix} 976196 & -1500816 & 2954639 & 4411258 \\ -1500816 & 5195401 & -8320868 & -12539182 \\ 2954639 & -8320868 & 14137961 & 21329594 \\ 4411258 & -12539182 & 21329594 & 32242779 \end{pmatrix}.$$

Their lower and upper spectral inequalities in (10) are proved by exact decompositions into positive multiples of squares. The source Jacobian is

$$J_\rho = \begin{pmatrix} -2 - 4eA & z + 2e & B & 0 \\ 1 + 2eA & -1 - z - e & -B & 0 \\ 1 & -z & -B - 16 - 8z - \rho & 3 \\ 0 & 0 & 16 + 4z & -(2 + \eta) \end{pmatrix}.$$

Uniform rational bounds on the reconstructed centers control $P_b J_\rho + J_\rho^T P_b$. The quadratic nonlinear terms are bounded on $|y_i| \le 1/400$ to give the $59/100$ dissipation constant. The exact jump-square expansions then control the exponential generator. All matrix entries above are rational certificates, not rounded eigenvector output.

The outer level is $16a_*$. Batch-active energies are below $8a_*$; the larger outer region supplies the necessary exit barrier. The preparation uses energy at most a_*, and selected-cell recovery returns to that same level. Distinguishing these levels is necessary: a first-entry assertion or an outer-region statement alone does not provide the terminal ready population required by the second cycle.

B Complete sufficient error budget

Let $u = N\alpha a_*$. The chemical and division bound used for one batch is

$$C(N, M, t) = Me^{-7u} + 14Me^{-4u} + \frac{Mtu}{60}e^{-15u/2} + Me^{-u} + \frac{Mtu}{60}e^{-3u/2} + 7Md_N. \tag{46}$$

Its first terms control outer exits and the endpoint spatial condition; d_N is (13). The affine spatial denominator underlying this expression is 960. The complete cycle estimate is

$$\mathcal{E}(p) = C(N, M, T) + e^{-N/2500} + e^{-19N/500000} + e^{-NM} + 2 \cdot 10^{-6} + \frac{32}{\varepsilon^2 pM} + \frac{2M}{10^{18}}. \tag{47}$$

The two 10^{-6} allowances pay collection and batch service. The last term pays selected-cell chemical recovery and its service. The count witness gives $u = 128$. Exact exponential bounds in Lean prove

$$\mathcal{E}(p) - \frac{32}{(1/50)^2 pM} \le 3 \cdot 10^{-6}$$

uniformly in p; no floating-point underflow is treated as a zero probability.

Term	Main scale dependence	Role
$C(N, M, T)$	Exponential in $N\alpha a_*$; linear in M, T	Chemistry and divisions
$e^{-N/2500}$	Exponential in N; deadline $8/\gamma$	Endpoint attainment
$e^{-19N/500000}$	Exponential in N	Size selection fluctuation
e^{-NM}	Exponential in NM	Low batch output
$2 \cdot 10^{-6}$	Chosen quotas	Collection and batch service
$32/(\varepsilon^2 pM)$	Inverse population size and ancestry floor	Neutral transfer
$2M/10^{18}$	Linear in selected count	Recovery and its service

Table 5: Dependence of the sufficient error terms. Extraction enters the chemical centers and certificates and the collection demand; these expressions are not justified outside the certified rate interval.

C Feedback attraction certificates

For the aggregate coordinates (A, B, z, H, R, S) the exact rational matrices are:

$$P_- = \frac{1}{100}\begin{pmatrix} 103 & -51 & 206 & 302 & -153 & -290 \\ -51 & 99 & -200 & -292 & 148 & 281 \\ 206 & -200 & 586 & 881 & -447 & -844 \\ 302 & -292 & 881 & 1346 & -683 & -1293 \\ -153 & 148 & -447 & -683 & 932 & 398 \\ -290 & 281 & -844 & -1293 & 398 & 7151 \end{pmatrix}.$$

$$P_+ = \frac{1}{100}\begin{pmatrix} 109 & -180 & 347 & 519 & -330 & -90 \\ -180 & 597 & -968 & -1458 & 937 & 247 \\ 347 & -968 & 1654 & 2493 & -1592 & -423 \\ 519 & -1458 & 2493 & 3764 & -2401 & -638 \\ -330 & 937 & -1592 & -2401 & 1801 & 409 \\ -90 & 247 & -423 & -638 & 409 & 827 \end{pmatrix}.$$

Coordinate	Low lower	Low upper	High lower	High upper
A	12.76378483	12.76383746	20.62315108	20.62371061
B	20.23475986	20.23476435	12.37243242	12.37244679
z	0.9651939	0.96519458	2.84948539	2.84949104
H	8.65271792	8.65272445	30.91390454	30.91398157
R	0.5601711	0.56017148	0.20008812	0.20008853
S	0.04067408	0.04067413	0.07014929	0.07014934

Table 6: Exact decimal endpoints of the rational aggregate boxes.

If $v = u - u_*$ and $L(u, u_*)$ is the exact secant matrix, so that $f(u) - f(u_*) = L(u, u_*)v$, the certified box inequalities give

$$v^\top[-PL - L^\top P]v \geq \tfrac{1}{2}\|v\|^2, \qquad \tfrac{1}{1000}\|v\|^2 \leq v^\top P v \leq 200\|v\|^2.$$

These imply aggregate decay at rate $1/400$. They follow from exact affine coefficient bounds and positive quadratic decompositions in the supplied Lean modules, rather than an eigenvalue computation at a floating-point center. The strict root enclosures lie in the interiors of the displayed boxes. The full-community argument in Theorem 8.1 adds the dispersion energy and chooses a smaller invariant sublevel inside these boxes; the whole box is not asserted to be a basin.

D Formal scope and reproducibility

The development uses Lean 4.30.0 and the repository's frozen mathlib dependencies. The verifier treats every compiler warning as an error. Authenticated root interfaces report only the standard classical axioms `propext`, `Classical.choice` and `Quot.sound`; numerical root approximations are not imported as axioms. The new stationary-cap proof binds the rational square certificate directly to equilibrium balances, including boundary consumer states.

The source package includes a detailed claim map, strict receipts, source fingerprints, numerical inputs, figure scripts and build instructions. It preserves the exact rational root enclosures, rather than replacing them with solver output. The full sufficient error budget is in Appendix B; no bulk case enumeration is required by the present publication additions.

Propositions 10.2, 10.3 and 11.1, their conditional convergence and recovery-window consequences, Corollary 12.1 and the infinite-horizon bottleneck argument have the conventional proofs given in the text. They are not advertised as new Lean theorems. For Corollary 12.1, the source premise is `productive_cycle_kernel_probability`, whose ancestry-floor hypothesis is $p > 0$; `productive_nontransfer_error` is uniform in p. Exact rational arithmetic checks the formulas and displayed error budgets; finite test instances do not replace their dimension-free proofs. The event-counter interpretation $G_z = 3W_0$ on a completed, unrefilled batch is an induction over the displayed protocol; the formal history inequality is applied with that premise. Physical normalizations and branch-continuation plots are explicitly illustrative.

Mathematical component	Principal declaration or module
Joint newborn copying event	`ProductiveMemory.productive_memory_publication_root`
Finite service witness	`productive_memory_resource_witness`
History and synthesis	`ExtractionHistory`, `ExtractionNetProduction`
Full-community attraction	`RAF1519.Reservoir.R19`
Resident identity	`resident_feedback_identity`
Instantaneous readout	`transient_composition_readout`
Abundance-error bound	`abundance_readout_error`
Window identities	`finite_window_readout`, `reservoir_window_readout`
Local observable recovery	`source_observable_recovery`, `observable_deadline`
Sharper stationary cap	`sharp_stationary_supply_cap`

Table 7: Formal interfaces. Unqualified readout, recovery and cap names are in `RAF1519.Reservoir`; the resource witness is in `ProductiveMemory`.

The stochastic result retains the existing output threshold. A more general exponential e^{aG-bE} would yield stronger output if $\gamma(e^a - 1) \leq \rho(1 - e^{-b})$, but applying it to a stronger literal event additionally needs adequate counter capacity and compatible stopping and collection quotas. No such strengthened joint event is silently substituted here. Likewise, the source bridge does not formalize inheritance of the consumers.

The combined audit interface is `proofs/MemoryUnderLoad/Publication.lean`. Its strict receipt authenticates the current source and transitive local dependencies. The accompanying proof snapshot and manifest identify the exact files used; rebuilding the manuscript itself requires only the supplied LaTeX sources and vector figures.

References

[1] Rémi Bardenet and Odalric-Ambrym Maillard. Concentration inequalities for sampling without replacement. *Bernoulli*, 21(3):1361–1385, 2015. doi: 10.3150/14-BEJ605. URL https://doi.org/10.3150/14-BEJ605.

[2] Armand Despons. Nonequilibrium properties of autocatalytic networks, 2024. URL https://arxiv.org/abs/2404.03347v2. arXiv:2404.03347v2.

[3] Armand Despons, Yannick De Decker, and David Lacoste. Structural constraints limit the regime of optimal flux in autocatalytic reaction networks. *Communications Physics*, 7:224, 2024. doi: 10.1038/s42005-024-01704-8. URL https://doi.org/10.1038/s42005-024-01704-8.

[4] Gregory B. Gloor, Jean M. Macklaim, Vera Pawlowsky-Glahn, and Juan J. Egozcue. Microbiome datasets are compositional: And this is not optional. *Frontiers in Microbiology*, 8:2224, 2017. doi: 10.3389/fmicb.2017.02224. URL https://doi.org/10.3389/fmicb.2017.02224.

[5] Wim Hordijk and Mike Steel. Detecting autocatalytic, self-sustaining sets in chemical reaction systems. *Journal of Theoretical Biology*, 227(4):451–461, 2004. doi: 10.1016/j.jtbi.2003.11.020. URL https://doi.org/10.1016/j.jtbi.2003.11.020.

[6] Yoshiya J. Matsubara, Sandeep Ameta, Shashi Thutupalli, Philippe Nghe, and Sandeep Krishna. Conditions for Darwinian evolution in compartmentalized autocatalytic reaction networks, 2026. URL https://arxiv.org/abs/2211.03155v2. arXiv:2211.03155v2.

[7] Friedrich Schlögl. Chemical reaction models for non-equilibrium phase transitions. *Zeitschrift für Physik*, 253:147–161, 1972. doi: 10.1007/BF01379769. URL https://doi.org/10.1007/BF01379769.

[8] Peter D. Taylor and Leo B. Jonker. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences*, 40(1–2):145–156, 1978. doi: 10.1016/0025-5564(78)90077-9. URL https://doi.org/10.1016/0025-5564(78)90077-9.

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

[10] Doris Vandeputte, Gunter Kathagen, Kevin D'hoe, Sara Vieira-Silva, Mireia Valles-Colomer, João Sabino, Jun Wang, Raul Y. Tito, Lindsey De Commer, Youssef Darzi, Séverine Vermeire, Gwen Falony, and Jeroen Raes. Quantitative microbiome profiling links gut community variation to microbial load. *Nature*, 551:507–511, 2017. doi: 10.1038/nature24460. URL https://doi.org/10.1038/nature24460.