

Reliable repeated harvesting in a finite-molecule autocatalytic reactor:

a source-level finite-horizon operating theorem with a machine-checked proof

15 September 2026

Abstract

Repeated harvesting asks a reactor that has just been thinned to make product and to restore its own catalytic stock before the next withdrawal. At finite molecule number this is a probabilistic question: the withdrawal is a random thinning of individual molecules, the catalyst is regrown by a Markov jump process, and the state from which the next cycle starts is the actual random endpoint of the previous one. We prove a joint finite-horizon operating theorem for one explicitly specified six-species stochastic mass-action reactor with maintained fuel and waste, continuous food supply and washout, and a copy scale V . Each cycle independently retains, withdraws or loses every molecule according to a history-dependent intervention, adds only food, and runs the literal twenty-channel count process for four time units, collecting effluent during the last unit. One event on the full normalized history law requires, in every cycle, return of the actual endpoint to an explicit restart set, two overlapping output thresholds ($\lceil V/56 \rceil$ template equivalents and $\lceil V/1080 \rceil$ free template molecules), two food allowances and a gross driven-service allowance, together with pathwise net-synthesis accounting for every physical realization. For every integer $V \geq 2 \times 10^{11}$, every admitted restart state, every parameter pair in the rectangle $r \in [19, 21]$, $d \in [1/50, 1/25]$ and every controller acting on the complete returned history, the probability that m consecutive cycles all succeed is at least $(1 - e(V))^m \geq 1 - m e(V)$ with $e(V) \leq 101e^{-V/10^{10}}$. Hence a copy scale logarithmic in m/δ achieves failure tolerance δ : at $V = 2 \times 10^{11}$, 48,000 cycles succeed jointly with probability at least 99% and 100 cycles with probability at least 99.9979%. Net internal synthesis is positive after 55 successful cycles from every admitted start. The proof combines exponential tests for a molecular pulse, a backward tilt schedule for seeded catalytic recovery, a positive short-window phase-transport estimate from bound catalyst to free product, marked counters on one reaction law, an identification of the stopped uniformized device with the physical chronological process by a renewal identity and material-height exhaustion, and conditioning on actual returns instead of independence between cycles. The entire chain, including nonexplosion and the identification, is verified in Lean 4 with Mathlib; an exact replay script re-derives every printed constant. The reactor is a schematic benchmark with established catalyst and maintained reservoirs; no molecular implementation is asserted.

1 Introduction

Structural theories of autocatalysis, from Kauffman’s autocatalytic sets [18] through the RAF formalism [16, 29] to the stoichiometric classification of autocatalytic cores [10], say which reaction networks can regenerate their catalysts from a food source. The operational test applied to a replicating system in the laboratory is the serial-transfer protocol [6, 21, 23]: a fraction of the mixture is diluted into fresh food, and the system is called self-sustaining if it regrows to the point where the next transfer can be made, repeatedly. A companion

paper [4] proved a uniform deterministic version of this test for a specified reversible, materially balanced, thermochemically consistent chemistry: for the mass-action concentration equations, one conditioning period carries any admitted preparation into an operating region, and from that region every admissible withdrawal followed by three units of recovery returns the state to the region, with certified template output on the following unit, uniformly over a rectangle of release and cleavage speeds and over state-dependent interventions.

The concentration equations are the large-volume limit of a count process [9, 19], and the deterministic theorem is silent about what happens to a finite number of molecules. The companion paper stated exactly what a finite-copy theorem has to add: a withdrawal kernel specified at the level of molecules, including species-dependent losses and refill rounding; lower tails for the post-pulse catalytic and material coordinates conditional on the actual pre-pulse state; return, output and supplies proved on one continued count law, with free-template export counted as washout events of free molecules; and an iteration of a conditional one-cycle bound that gives m -cycle success at least $(1 - \epsilon)^m$ without an independence assumption between cycles. This paper discharges these four obligations for the same chemistry, and adds a fifth that the deterministic setting cannot express: the state that starts each cycle is the actual random endpoint of the previous one, and the intervention may depend on the complete recorded history.

Main result in words. The reactor is the literal integer-count Markov jump process with twenty labelled channels of Table 1, at copy scale V . An admissible intervention retains each molecule of species i with probability $q\ell_i$, withdraws it with probability $1 - q$, and loses it otherwise, where $q \in [1/4, 3/4]$ and $\ell_i \in [49/50, 1]$; it then adds integer food doses with a bounded rounding error and nothing else. The source runs for four time units; washout is collected on the last unit. A cycle succeeds if its actual endpoint lies in the restart set (material coordinates within $1/160$ of V and weighted catalytic stock at least $V/20$), if the collected template equivalents and free template molecules exceed $\lceil V/56 \rceil$ and $\lceil V/1080 \rceil$, and if the two food expenditures and the gross forward-plus-reverse driven-service count stay below $5V$, $5V$ and $\lfloor V/5 \rfloor$. Theorem 3.1 states that for every integer $V \geq V_0 = 2 \times 10^{11}$ and every controller, the probability that m consecutive cycles all succeed, under the full normalized history law, is at least $(1 - e(V))^m$ with an explicit finite exponential sum $e(V) \leq 101e^{-V/10^{10}}$, and that on this event the net internal template synthesis of every physical realization is at least $m\lceil V/56 \rceil + \lceil V/28 \rceil$ minus the initial inventory. Corollary 3.2 turns this into a design rule: a copy scale $V = \max(V_0, \lceil 10^{10} \log(101m/\delta) \rceil)$ certifies m cycles at failure tolerance δ . At V_0 itself, 48,000 cycles are certified at 99% and 100 cycles at 99.9979%.

Method. Every probability estimate is a Foster–Lyapunov generator inequality [22] used quantitatively over a finite horizon, in the form of exponentially tilted tests on finite stopped kernels obtained by uniformization [13, 17], the technique of the companion finite-copy papers [1, 2]. In the stochastic reaction network literature such inequalities usually serve positive recurrence and moment bounds [8, 14]; here each one yields a number. Seven ingredients are specific to repeated harvesting and are, to our knowledge, new in this combination.

- (i) *The molecular pulse.* The withdrawal is an exact product law over molecules with three categories per molecule. Its retained weighted stock and both retained material coordinates have exponential moment bounds whose exponent is the actual retained mean rather than a worst-case sum of squared weights (Lemma 4.1); every pulse outcome, including loss of all catalyst, keeps its mass in the law.
- (ii) *Seeded recovery by a backward tilt schedule.* Below the target stock $3V/50$ the weighted catalytic stock has count-generator drift at least $\frac{2}{3}Y$, and the exponential-drift coefficient $5/8$ survives the quadratic remainder of the jumps. A tilt that is transported backward

from 1/1000 at the deadline and capped at 101/20000 converts the prepared stock 3V/250 into entry by time 11/4 except with probability $e^{-3V/5 \cdot 10^6}$ plus a Poisson-clock term (Proposition 5.3), with no first-hitting-time construction.

- (iii) *Positive phase transport.* Large weighted stock does not mean large free template at the same instant. A backward weight vector propagated through the actual phase-comparison drifts shows that every catalytic phase funds a free-template observation between 89V and 91V clock steps later, with coefficient at least 1/35 of its stock weight (Lemma 5.6). The resulting low-free-template probability $e^{-V/10^{10}}$ is the dominant term of the certificate.
- (iv) *Occupation by expectation.* A pointwise bound after any burn-in becomes an occupation bound over the whole collection unit through a sum of expectations and Markov’s inequality, without independence between clock steps (Proposition 5.8).
- (v) *Marked counters on one law.* Template washout, free washout, both food feeds and both driven labels are marks on the same sequence of reaction events; their five tails are added by a union bound within one joint law, never multiplied as marginals (Proposition 6.3).
- (vi) *Removing the stopping.* The stopped uniformized kernel is an analytical device. Its estimates transfer to the physical process only on the strict event that the terminal state is still active, and a small exact model shows that the closed restart event alone does not transfer (Remark 7.1). The transfer itself is proved, not assumed: killed chronological laws and stopped clocks satisfy the same first-event renewal equation, whose solution is unique on bounded time intervals, and increasing material-height regions exhaust the countable state space (Section 7).
- (vii) *Conditioning on actual returns.* The m -cycle bound is obtained by integrating the uniform conditional one-cycle bound over successful histories; the restricted history measure is a submeasure of the full law, never renormalized (Proposition 8.1).

Contributions.

- (a) A joint finite-horizon reliability theorem for a maintained finite-molecule autocatalytic reactor under repeated history-dependent withdrawal (Theorem 3.1), with an explicit single-exponential error certificate and a logarithmic sufficient copy scale (Corollary 3.2).
- (b) A pathwise inventory identity on physical realizations of the count process that certifies net covalent synthesis on the same event (Proposition 8.3), with an onset at 55 cycles from every admitted start and at two cycles from a richer witness.
- (c) Explicit same-event operating ratios (Corollary 3.3) and a lower bound showing that the phase-occupation term is the bottleneck of this certificate (Proposition 3.4).
- (d) A complete formalization in Lean 4 [12] with Mathlib [30] of the count source, the pulse, every tail estimate, nonexplosion, the identification of the stopped device with the chronological process, the history law, the repetition argument, the inventory identity and the quantitative corollaries (Appendix D). Unlike the companion stochastic papers, no step of the identification is left to a conventional argument.

Three kinds of evidence. Three kinds of statement appear and are kept apart. Every probability, inventory and sizing statement in Theorem 3.1, Corollaries 3.2 and 3.3 and Proposition 3.4 is a Lean declaration whose statement is described in Appendix D; the prose proofs of Sections 4–8 are expositions of the formal chain and are written so that each constant can be checked by hand. The example scales of Table 2 beyond V_0 and the amount conversions of Section 9 are conventional arithmetic with exact rational Taylor certificates, replayed by the accompanying script. Figures 3 and 2 evaluate proved bounds and closed-form quantities;

nothing in this paper is a simulation, and no simulation of the 2×10^{11} -molecule source is used as evidence.

What is not claimed. The chemistry is schematic: three formal elemental components, maintained fuel and waste activities in the sense of open-network thermodynamics [25, 26], ideal mass action, and one effective trimolecular reverse channel whose elementary implementation is not supplied. No molecule is identified and no rate constant is fitted. The intervention is instantaneous, its thinning is conditionally independent across molecules, and the copy scale is restored externally; handling time, mixing, pumping, separation and purification are outside the model, and the gross driven-service counter measures a specified part of reactor operation, not total dissipation [27]. The theorem starts from established catalytic stock and gives no food-only startup; the necessary initiation scale of [5] is the relevant statement there. Finite missions are certified, not indefinite survival: the count process at fixed V is positive recurrent and catalyst loss eventually occurs. The certificate is sufficient, not sharp; Proposition 3.4 says which term limits it.

Organization. Section 2 specifies the count source, the intervention, the cycle observables and the history process. Section 3 states the results. Sections 4–6 prove the one-cycle bound on the stopped device: pulse and material control (Section 4), recovery, residence and phase transport (Section 5), marked outputs and supplies (Section 6). Section 7 removes the stopping and identifies the physical law. Section 8 proves repetition, inventory and the quantitative corollaries. Section 9 gives a dimensional reading and Section 10 the scope, the verification boundary and open questions. The appendices contain the complete error budget, the quantitative exponential tests, the dimensional coefficients and the formal evidence.

2 The count source and the operating protocol

2.1 Species, channels and observables

The internal species are two foods U, W , the covalent template X , two substrate-binding complexes C_1, C_2 and a duplex Z . A fuel F and a waste P are maintained at unit activity. The count state is $N = (N_U, N_W, N_X, N_{C_1}, N_{C_2}, N_Z) \in \mathbb{N}^6$ and $V \in \mathbb{N}$ is the copy scale. Fix

$$\varepsilon = \frac{1}{500\,000\,000}, \quad \eta = \frac{1}{8\,000\,000\,000}, \quad 19 \leq r \leq 21, \quad \frac{1}{50} \leq d \leq \frac{1}{25}, \quad (1)$$

with (r, d) constant throughout an operation. Table 1 lists the twelve directed chemical channels. In addition, two food feeds $\emptyset \rightarrow U$ and $\emptyset \rightarrow W$ each have propensity V , and each of the six internal species washes out with propensity equal to its count. There are twenty labelled channels; a channel with insufficient reactants has propensity zero, and the falling factorial $N_X(N_X - 1)$ of the duplex reverse channel carries no factorial divisor. This is the count-level realization of [5], whose coefficients are bound formally to the concentration field of [4]; the six pairs balance three formal elemental units and share one vector of standard potentials for every (r, d) in (1).

For label j write $\nu_j \in \mathbb{Z}^6$ for its internal stoichiometric increment and $a_j(N)$ for its propensity. The generator of the count process is

$$\mathcal{L}f(N) = \sum_j a_j(N) \{f(N + \nu_j) - f(N)\}. \quad (2)$$

Four linear observables organize the analysis: the two material totals, a weighted catalytic stock and the template inventory,

$$\begin{aligned} A_N &= N_U + N_X + 2N_{C_1} + 2N_{C_2} + 2N_Z, & B_N &= N_W + N_X + N_{C_1} + 2N_{C_2} + 2N_Z, \\ Y_N &= N_X + \frac{9}{8}N_{C_1} + \frac{7}{5}N_{C_2} + \frac{9}{5}N_Z, & I_N &= N_X + N_{C_1} + N_{C_2} + 2N_Z. \end{aligned} \quad (3)$$

Pair	Full reaction	Forward propensity	Reverse propensity
0	$U + W \rightleftharpoons X$	$\varepsilon N_U N_W / V$	$(\varepsilon/10) N_X$
1	$X + U \rightleftharpoons C_1$	$20 N_X N_U / V$	$20 N_{C_1}$
2	$C_1 + W \rightleftharpoons C_2$	$20 N_{C_1} N_W / V$	$20 N_{C_2}$
3	$C_2 \rightleftharpoons Z$	$20 N_{C_2}$	$2 N_Z$
4	$Z \rightleftharpoons 2X$	$r N_Z$	$r N_X (N_X - 1) / V$
5	$X + F \rightleftharpoons U + W + P$	$d N_X$	$d \eta N_U N_W / V$

Table 1: The twelve chemical channels at unit fuel and waste activities. Together with two food feeds of propensity V and six unit-rate washouts they form the twenty labelled channels of the source.

A and B count the two food-derived elemental units among internal molecules; I counts template equivalents, two for a duplex; Y is the proof device whose weights make the drift favourable at low stock. The exact source identities are

$$\mathcal{L}A_N = V - A_N, \quad \mathcal{L}B_N = V - B_N, \quad \mathcal{L}Y_N = V Y(f_{r,d}(N/V)) + \frac{r N_X}{5V}, \quad (4)$$

where $f_{r,d}$ is the mass-action concentration field of [4]. Internal chemistry conserves both material coordinates, the feeds add one unit each and washout removes the weight of the departing molecule; the last identity says that the only difference between the count drift of Y and the concentration drift is the exact, favourable falling-factorial correction.

Definition 2.1 (Restart set). The admissible restart set at scale V is

$$\mathcal{R}_V = \left\{ N \in \mathbb{N}^6 : \frac{159V}{160} \leq A_N, B_N \leq \frac{161V}{160}, \quad 40N_X + 45N_{C_1} + 56N_{C_2} + 72N_Z \geq 2V \right\}. \quad (5)$$

The integer stock condition is $Y_N \geq V/20$. The set is nonempty at every integer scale, since $(0, 0, V, 0, 0, 0) \in \mathcal{R}_V$, and at scales $1000k$ it contains the lattice witness $k(934, 935, 60, 1, 1, 1)$, whose concentration image is the example state of [4]. Membership in $\mathcal{R}_V$ is an established-stock condition, not a food-only initial condition.

2.2 Interventions and the molecular pulse

Definition 2.2 (Admissible intervention). An intervention is $p = (q, \ell, e_U, e_W)$ with

$$\frac{1}{4} \leq q \leq \frac{3}{4}, \quad \frac{49}{50} \leq \ell_i \leq 1 \quad (i \in \{U, W, X, C_1, C_2, Z\}), \quad |e_U|, |e_W| \leq \frac{1}{200}. \quad (6)$$

Conditional on p and on the current counts N , every molecule independently enters one of three categories with probabilities

$$(\text{retained}, \text{withdrawn}, \text{extra loss}) = (q\ell_i, 1 - q, q(1 - \ell_i)) \quad (7)$$

for a molecule of species i . The retained molecules remain; $\lfloor V(1 - q + e_U) \rfloor$ molecules of U and $\lfloor V(1 - q + e_W) \rfloor$ molecules of W are added, and no catalyst-bearing species. The pulse has zero duration and the copy scale V is restored externally.

Here q is the retained well-mixed fraction, ℓ_i allows up to 2% further species-dependent loss, and e_U, e_W are refill errors. The pulse law is the exact finite product measure over molecules; every outcome, including the loss of every catalytic molecule, keeps its probability mass. Nothing is clipped, rejected or conditioned.




Figure 1: Reaction pathways and one physical cycle. All chemical links are reversible; the four-step templated path from $X + U + W$ to $2X$ is the autocatalytic completion. Food feeds and unit-rate washouts act continuously. Collection marks are switched on at time 3; the actual endpoint starts the next cycle.

2.3 Cycle, counters and success

Definition 2.3 (Cycle and counters). From the post-pulse counts, run the source (2) for four time units. Five integer counters are attached to the labelled events of the same reaction sequence:

- Q_I , template equivalents collected: a washout of X , C_1 , C_2 or Z during $[3, 4]$ adds 1, 1, 1, 2 respectively;
- Q_X , free template collected: a washout of X during $[3, 4]$ adds 1 (so Q_X is included in Q_I);
- C_U, C_W , food expenditure: the pulse doses plus every feed arrival of U , respectively W , during $[0, 4]$;
- G , gross driven service: every forward and every reverse pair-5 event during $[0, 4]$.

The actual endpoint $N(4)$ and the five counters form the cycle record.

Definition 2.4 (Cycle success). A record $(N, Q_I, Q_X, C_U, C_W, G)$ is successful if

$$N \in \mathcal{R}_V, \quad Q_I \geq \lceil V/56 \rceil, \quad Q_X \geq \lceil V/1080 \rceil, \quad C_U, C_W \leq 5V, \quad G \leq \lceil V/5 \rceil. \quad (8)$$

The two output thresholds are overlapping observables of different strength: Q_I counts template material in the effluent whether bound or free, Q_X counts free template molecules only. The three cost coordinates are allowances on the successful event, not caps imposed on the process.

2.4 The history process and admitted controllers

Definition 2.5 (Returned-history process). Fix V , (r, d) , an initial state $N_0 \in \mathcal{R}_V$ and a controller π , an arbitrary function from finite lists of cycle records to admissible interventions. The returned-history process is the Markov chain on finite lists of records whose one-step kernel, from a list h with current counts $N(h)$ (the counts of the newest record, or N_0 for the empty list), applies the intervention $\pi(h)$, samples the pulse and the four-unit marked source cycle from $N(h)$, and prepends the resulting record to h . Its m -step law from the empty list is the full normalized history law $\mathbb{P}$ over m cycles.

Every pulse and source outcome is retained, and every cycle starts from the actual returned counts. The controller sees the complete list of returned populations and counter values; no measurability or continuity is required of it. Section 8 also treats a general measurable history space carrying possibly continuous observations, for which the requirement is that the next record, conditional on the current history and the chosen intervention, have exactly the literal pulse-and-source law, and that the current population equal the newest observed return. That requirement is a specification of the admissible physical model, not an assumption on success probabilities. No within-cycle control of rates is admitted.

3 Main results

For cycle $k \geq 1$ let S_k be the event that its record is successful in the sense of Definition 2.4, and let $E_m = \bigcap_{k=1}^m S_k$. Under $\mathbb{P}$, a finite sequence of labelled reaction events and pulse category outcomes realizing the recorded history exists almost surely; we call such a sequence a physical trace. Let P_{net} denote the signed net internal template synthesis along a trace: the sum over chemical events of the change of I they cause (pair 0 contributes +1 forward and -1 backward, pair 3 likewise, pair 5 contributes -1 forward and +1 backward, and pairs 1, 2, 4 contribute zero). Let $e(V)$ be the explicit finite sum of exponential and polynomial-exponential terms defined in Appendix A.

Theorem 3.1 (Repeated finite-copy operation). *Fix (r, d) in (1), an integer $V \geq V_0 := 200\,000\,000\,000$, a restart state $N_0 \in \mathcal{R}_V$ and a controller π . Then:*

- (a) (One cycle) *For every $V \geq 10^6$, every $N \in \mathcal{R}_V$ and every admissible p , the probability that the cycle started from N with intervention p is successful is at least $1 - e(V)$, where $0 \leq e(V)$ and $e(V) \leq 101e^{-V/10^{10}}$ for $V \geq V_0$.*
- (b) (Repetition) *For every $m \in \mathbb{N}$,*

$$\mathbb{P}(E_m) \geq (1 - e(V))^m \geq 1 - m e(V). \quad (9)$$

- (c) (Accounting) *On E_m every inequality of (8) holds in every cycle, hence the mission totals satisfy $\sum_k Q_{I,k} \geq m[V/56]$, $\sum_k Q_{X,k} \geq m[V/1080]$, $\sum_k C_{U,k}, \sum_k C_{W,k} \leq 5mV$ and $\sum_k G_k \leq m[V/5]$.*
- (d) (Net synthesis) *On E_m a physical trace exists, and every physical trace of the recorded history satisfies*

$$P_{\text{net}} \geq m \left\lceil \frac{V}{56} \right\rceil + \left\lceil \frac{V}{28} \right\rceil - I_{N_0} \geq m \left\lceil \frac{V}{56} \right\rceil + \left\lceil \frac{V}{28} \right\rceil - \left\lfloor \frac{161V}{160} \right\rfloor \geq \left(\frac{m+2}{56} - \frac{161}{160} \right) V. \quad (10)$$

In particular $P_{\text{net}} > 0$ for every $m \geq 55$ and every $N_0 \in \mathcal{R}_V$, and $P_{\text{net}} \geq 13V/1750$ already for $m = 2$ when $I_{N_0} = 8V/125$.

Corollary 3.2 (Logarithmic sizing and fixed-scale missions). *For $m \geq 1$ and $0 < \delta < 1$ put*

$$V_{\log}(m, \delta) = \max \left(V_0, \left\lceil 10^{10} \log \frac{101m}{\delta} \right\rceil \right). \quad (11)$$

At every integer scale $V \geq V_{\log}(m, \delta)$, $\mathbb{P}(E_m) \geq 1 - \delta$. More generally, at any $V \geq V_0$ every mission length $m \leq \lfloor (\delta/101)e^{V/10^{10}} \rfloor$ has $\mathbb{P}(E_m) \geq 1 - \delta$. At $V = V_0$,

$$\mathbb{P}(E_{48000}) \geq 0.99, \quad \mathbb{P}(E_{100}) \geq 0.999979.$$

Corollary 3.3 (Same-event operating ratios). *For $m \geq 1$, on E_m the mission totals Q_I, Q_X, C_U, C_W, G are positive where they appear as denominators and satisfy*

$$\begin{aligned} \frac{Q_I}{4mV} &\geq \frac{1}{224}, & \frac{Q_X}{4mV} &\geq \frac{1}{4320}, & \frac{C_U}{Q_I}, \frac{C_W}{Q_I} &\leq 280, & \frac{C_U + C_W}{Q_I} &\leq 560, \\ \frac{G}{Q_I} &\leq \frac{56}{5}, & \frac{C_U + C_W}{Q_X} &\leq 10800, & \frac{G}{Q_X} &\leq 216. \end{aligned}$$

For output demands D_I template equivalents and D_X free molecules at tolerance δ , the integer scale $\max(V_{\log}(m, \delta), 56D_I/m, 1080D_X/m)$ is sufficient.

Proposition 3.4 (Certificate bottleneck). *For every $V \geq 0$, $e(V) \geq 100e^{-V/10^{10}}$. Consequently, if $m e(V) \leq \delta$ then $V \geq 10^{10} \log(100m/\delta)$. This is a lower limit for the present certificate, not for a physical reactor.*

Remark 3.5 (Quantifiers). The parameter pair is fixed once and arbitrary in its rectangle; the stochastic estimates in fact hold for $0 \leq d \leq 1/25$, and the lower bound $d \geq 1/50$ in (1) is inherited from the thermochemical rectangle of [5]. The intervention class is the full box (6) at every cycle, and the controller is an arbitrary function of the complete returned history. Part (a) holds from $V \geq 10^6$; the larger floor V_0 is where the single-exponential envelope and the quantitative corollaries become available. Nothing in the theorem is conditional on success: $\mathbb{P}$ is the full normalized law, and E_m is an event in it whose probability is bounded below. The costs are allowances on E_m ; spending on unsuccessful runs is not bounded by them.

Remark 3.6 (Two formulations of the history). Part (b) is proved for the concrete list-history process of Definition 2.5. The same product bound is proved for every measurable history model that satisfies the literal conditional-law and actual-return requirements of Section 2.4 (Proposition 8.1); such a model may record continuous within-cycle observations. The theorem does not separately construct a controlled continuous-time filtration for arbitrary raw-path controllers; the conditional-law formulation states exactly which physical models are admitted.

Remark 3.7 (Earlier linear rule). A first compiled version of Corollary 3.2 used the sufficient scale $\max(V_0, \lceil 5 \cdot 10^6 m/\delta \rceil)$, obtained from the bound $V e(V) \leq 5 \cdot 10^6$ on $V \geq V_0$. The logarithmic rule (11) rests on the exact absorption of every polynomial-exponential term into the phase term (Lemma 8.5); at $m = 10^6$ and $\delta = 10^{-2}$ it replaces 5×10^{14} by 2.31×10^{11} (Table 2).

Remark 3.8 (Proof map). Sections 4–6 prove part (a) on a stopped uniformized device: the pulse leaves the material coordinates in a prepared band and the stock at least $3V/250$ (Section 4); the stock enters $Y \geq 3V/50$ by time $11/4$ and stays above $V/20$ through time 4, while the material coordinates return to their restart bands (Section 5); the five counters meet their thresholds (Section 6). Section 7 shows that the strict version of the stopped event transfers to the physical chronological process and that the analytical time split is exactly the physical four-unit law. Section 8 conditions on actual returns for part (b), telescopes the inventory for part (d), and proves the corollaries.

4 Pulse and material control

4.1 The retained stock

Lemma 4.1 (One-molecule retention factor). *Let $0 \leq p \leq 1$, $0 \leq w \leq 9/5$ and $0 \leq s \leq 1/10$. Then*

$$1 - p + p e^{-sw} \leq \exp\left\{-\left(s - \frac{27}{25}s^2\right)pw\right\}.$$

Proof. For $|z| \leq 9/50$ one has $e^z \leq 1 + z + \frac{3}{5}z^2$ (Appendix B). With $z = -sw$, $|z| \leq \frac{9}{50}$, so $1 - p + pe^{-sw} \leq 1 - psw + \frac{3}{5}ps^2w^2 \leq 1 - psw + \frac{3}{5}ps^2 \cdot \frac{9}{5}w = 1 - (s - \frac{27}{25}s^2)pw$, using $w^2 \leq \frac{9}{5}w$. Finally $1 + x \leq e^x$. $\square$

Lemma 4.2 (Pulse stock tail). *Let $N \in \mathcal{R}_V$ and let p be admissible. The retained weighted stock Y' after the pulse satisfies*

$$\mathbb{P}(Y' \leq 3V/250) \leq e^{-1177V/10^9}.$$

Proof. $Y' = \sum_a w_a B_a$ over molecules a , with B_a independent Bernoulli retention indicators of parameters $p_a = q\ell_{i(a)} \in [49/200, 1]$ and weights $w_a \in \{1, 9/8, 7/5, 9/5\}$ for catalytic molecules and 0 otherwise. The exact moment generating function is $\mathbb{E}e^{-sY'} = \prod_a (1 - p_a + p_a e^{-sw_a})$; Lemma 4.1 bounds it by $\exp\{-(s - \frac{27}{25}s^2) \sum_a p_a w_a\}$, and $\sum_a p_a w_a \geq \frac{1}{4} \cdot \frac{49}{50} \cdot Y_N \geq \frac{49}{4000} V$ because $Y_N \geq V/20$ on $\mathcal{R}_V$. The exponential Markov inequality at $s = 1/100$ and level $a = 3V/250$ gives the exponent $sa - (s - \frac{27}{25}s^2) \frac{49V}{4000} = -\frac{1177}{10^9} V$. $\square$

The exponent uses the actual retained mean rather than the sum of squared weights; the retained stock $3V/250$ is one quarter of the deterministic prepared stock $49/4000$, leaving room for the fluctuations of about $V^{1/2}$ molecules.

4.2 The retained material

Lemma 4.3 (Pulse material tails). *Let $N \in \mathcal{R}_V$, $V \geq 10^6$, and let p be admissible. After the pulse, each of A, B lies in the prepared band $[24V/25, 51V/50]$ except with probability at most $4e^{-V/100000}$ in total.*

Proof. Each material coordinate after the pulse is $\sum_a w_a B_a + D$, with material weights $w_a \in \{1, 2\}$ and an integer dose $D = \lfloor V(1 - q + e) \rfloor$. Its mean is $q \sum_a \ell_{i(a)} w_a + D$; over the restart band, the box (6) and the floor, the mean lies in $[\frac{31213}{32000} V - 1, \frac{3231}{3200} V]$, whose distance from each end of the prepared band exceeds $V/100 - 1$. A deviation of magnitude $V/100 - 1$ for a sum of independent bounded terms with weights at most 2 costs at most $e^{-V/100000}$ per sign by the bounded-difference exponential inequality [15]; two coordinates and two signs give $4e^{-V/100000}$. $\square$

Neither lemma conditions on a good pulse: bad pulses keep their mass and are counted in $e(V)$.

4.3 The corridor, uniformization and material exit

Stop the process temporarily when either A or B leaves the corridor $[9V/10, 11V/10]$. On the corridor every count is at most $11V/10$ and the sum of all propensities is at most $3000V$ (the actual sum is below $175V$; the slack is irrelevant). This permits uniformization [13, 17]: a Poisson clock of rate $q_c = 3000V$ selects channel j with probability $a_j(N)/q_c$ and otherwise makes a null step, giving a finite stopped kernel $P = I + \mathcal{L}/q_c$ whose expectations at time t are Poisson mixtures of kernel powers. Counters remain attached to the selected labels. For a linear observable H and tilt s , the exponential test satisfies

$$\frac{\mathcal{L}e^{sH}(N)}{e^{sH(N)}} = \sum_j a_j(N)(e^{s\Delta_j H} - 1) \leq s\mathcal{L}H(N) + \frac{3}{5}s^2 \sum_j a_j(N)(\Delta_j H)^2 \quad (|s\Delta_j H| \leq 9/50), \quad (12)$$

by the quadratic exponential inequality. One kernel step multiplies a test by at most $1 + \mathcal{L}e^{sH}/(q_c e^{sH})$; iteration, the exponential Markov inequality and Poisson summation give tails. Indicators of activity kill a test on exit, and every exit probability appears explicitly in the budget.

Lemma 4.4 (Material corridor exit). *From the prepared band, the probability that A or B leaves the corridor during a run of four time units is at most*

$$M(V) = 4e^{-V/2000} + 48000V e^{-V/2000+1/50}. \quad (13)$$

Proof. For $H \in \{A, B\}$ use the barriers $U_H = e^{(H-11V/10)/100}$ and $L_H = e^{-(H-9V/10)/100}$. At a prepared start each is at most $e^{-V/2000}$; at an exit their sum is at least one. Outside the inner band $|H - V| \leq V/20$ the restoring drift $\mathcal{L}H = V - H$ dominates the quadratic remainder in (12), so the tilted drift is nonpositive; inside it each barrier is at most $e^{-V/2000}$ and a material jump of size at most 2 changes its exponent by at most $1/50$, so the generator of each barrier is at most $q_c e^{-V/2000+1/50}$. Integrating over four time units gives $4 \cdot 3000V \cdot 4 e^{-V/2000+1/50}$ for the four barriers and $4e^{-V/2000}$ for the starts. $\square$

Lemma 4.5 (Terminal material return). *On the corridor, from a start with $|H_0 - V| \leq V/25$, the probability that $|H(4) - V| \geq V/160$ for $H = A$ or $H = B$, on the event of no corridor exit, is at most $4(e^{-V/320000} + e^{-V})$.*

Proof. Let $T_s = \mathbf{1}_{\text{corridor}} e^{s(H-V)}$. The drift $V - H$ contracts the tilt: the compensated test $e^{-2s^2V} T_s$ satisfies $P^n T_s \leq e^{2s^2V} T_{(1-1/q_c)^n s}$ for $|s| \leq 1/100$. After $n \geq 11700V$ steps the contraction factor is at most $e^{-3.9} \leq 1/40$. With $s = \pm 1/1000$ the exponent for the terminal threshold $|H - V| \geq V/160$ is at most $-V/160000 + 2V/10^6 + V/10^6 \leq -V/320000$. The Poisson clock of mean $12000V$ produces fewer than $11700V$ steps with probability at most e^{-V} . Two coordinates and two signs give the factor 4. $\square$

Both the prepared band and the narrow terminal band are kept; the exact cancellation $\mathcal{L}A = V - A$ alone would not imply concentration at the return time.

5 Recovery, residence and phase transport

5.1 Guarded growth and the stock exponential test

Lemma 5.1 (Guarded growth of the count stock). *On the corridor, if $Y_N \leq 3V/50$ then $\mathcal{L}Y_N \geq \frac{2}{3}Y_N$. Moreover, on the corridor with small stock the basal creation channel contributes at least $V/10^9$ to $\mathcal{L}Y_N$, and the rate-weighted squared jumps of Y are at most $6Y_N + V/10^8$.*

Proof. By (4), $\mathcal{L}Y_N = V Y(f_{r,d}(c)) + rN_X/(5V)$ with $c = N/V$, and the last term is nonnegative. The companion paper [4, Lemma 4.6] proves $Y(f_{r,d}(c)) \geq \frac{2}{3}Y(c)$ on the concentration corridor whenever $Y(c) \leq 3/50$, by exhibiting a nonnegative parameterized surplus; the mechanism is that low stock forces both free foods above $3/4$ (the exact floors are $119/150$ and $57/70$), so the templated path carries positive net weighted flux. The basal channel then has propensity $\varepsilon N_U N_W / V \geq \varepsilon (3/4)^2 V > V/10^9$ with Y -increment 1. Every Y -jump has magnitude at most $9/5$; summing squared increments against propensities on the corridor gives the stated linear bound. $\square$

This is a generator inequality, not a claim that random paths grow monotonically.

Lemma 5.2 (Stock exponential test). *On the corridor, for $0 \leq s \leq 1/100$ and $Y_N \leq 3V/50$,*

$$\mathcal{L}e^{-sY_N} \leq -\frac{5}{8}s Y_N e^{-sY_N}.$$

Proof. Apply (12) to $H = -Y$: the right side is $-s\mathcal{L}Y_N + \frac{3}{5}s^2(6Y_N + V/10^8)$. By Lemma 5.1, $-s\mathcal{L}Y_N \leq -\frac{2}{3}sY_N - sV/10^9$. At $s \leq 1/100$ the coefficient of Y_N is at most $-(\frac{2}{3} - \frac{3}{5} \cdot 6 \cdot \frac{1}{100})s = -\frac{473}{750}s \leq -\frac{5}{8}s$, and $\frac{3}{5}s^2V/10^8 \leq sV/10^9$, so the basal term pays for the additive noise source. $\square$

5.2 Seeded recovery

Proposition 5.3 (Seeded recovery deadline). *From a prepared start on the corridor with $Y \geq 3V/250$, the probability that the stopped process is still active and below $Y = 3V/50$ at time $11/4$ is at most*

$$e^{-3V/5\,000\,000} + e^{-V/2}.$$

Proof. Let $T_s = \mathbf{1}_{\{\text{active}, Y < 3V/50\}} e^{-sY}$. Lemma 5.2 and uniformization give $PT_s \leq T_t$ whenever $t \leq (1 + \frac{5/8}{q_c})s$, for $s \leq 1/100$. Start at the deadline with tilt $1/1000$ and transport it backward through $8100V$ steps, increasing the tilt by the factor $1 + \frac{5/8}{3000V}$ per step and capping it at $101/20000$. Grouping the steps into 27 blocks of $300V$ and applying Bernoulli's inequality blockwise, $(1 + \frac{5}{24000V})^{300V} \geq \frac{17}{16}$, and $(17/16)^{27} > 101/20$, so the cap is reached before the start. Hence $P^{8100V} T_{1/1000} \leq T_{101/20000}$ at the start, and by the exponential Markov inequality the probability of being active below $3V/50$ after $8100V$ steps is at most

$$\exp\left\{\frac{1}{1000} \cdot \frac{3V}{50} - \frac{101}{20000} \cdot \frac{3V}{250}\right\} = e^{-3V/5\,000\,000}.$$

The Poisson clock of mean $8250V = \frac{11}{4} \cdot 3000V$ produces fewer than $8100V$ steps with probability at most $e^{-V/2}$. $\square$

The certified exponent is small because the tilt is small; the mechanism is nevertheless comfortable, since the drift coefficient $5/8$ over the deadline $11/4$ gives a growth factor above 5.13 where the ratio of target to prepared stock is 5. The random entry state is continued as it is; no replacement start at the target is used.

Lemma 5.4 (Residence). *From an active state with $Y \geq 3V/50$, the probability that the stopped process is still active and has $Y \leq V/20$ at some fixed later time $t \leq 4$ is at most $e^{-V/10000} + 12000V e^{-V/10000+9/500}$.*

Proof. Use $H_Y = \max(e^{-Y/100}, e^{-3V/5000})$. Below the cap Lemma 5.2 applies and the drift is nonpositive; above it a jump changes Y by at most $9/5$, so $\mathcal{L}H_Y \leq q_c e^{-3V/5000+9/500}$. At the start $H_Y = e^{-3V/5000}$; at a state with $Y \leq V/20$, $H_Y \geq e^{-V/2000}$. Integration over time at most 4 and division by $e^{-V/2000}$ give the bound, since $-3/5000 + 1/2000 = -1/10000$. $\square$

Applying Proposition 5.3 to the actual state at time $11/4$ and Lemma 5.4 through the remaining $5/4$ units yields the recovery-and-residence budget $R(V)$ of Appendix A; material exits are paid separately by Lemma 4.4. The times $11/4$ and $1/4$ partition the analysis only (Section 7).

5.3 Positive phase transport to free template

Large Y does not imply large N_X at the same instant: the stock may sit in the bound phases C_1, C_2, Z . A constant linear corrector fails here (its boundary penalty exceeds the available stock integral), and the successful replacement is a short-window transport estimate that keeps every phase endpoint.

Lemma 5.5 (Phase comparison drifts). *On the corridor,*

$$\begin{aligned} \mathcal{L}N_X &\geq -70N_X + 20N_{C_1} + 38N_Z, & \mathcal{L}N_{C_1} &\geq -70N_{C_1} + 20N_{C_2}, \\ \mathcal{L}N_{C_2} &\geq -70N_{C_2}, & \mathcal{L}N_Z &\geq -70N_Z + 20N_{C_2}. \end{aligned} \tag{14}$$

Proof. Read off (2) with Table 1. For N_X the negative terms are $-(\varepsilon/10)N_X - 20N_X N_U/V - 2rN_X(N_X - 1)/V - dN_X - N_X$; with $N_U, N_X \leq 11V/10$, $r \leq 21$ and $d \leq 1/25$ their coefficient is at most 69.25, and the positive terms include $20N_{C_1}$ and $2rN_Z \geq 38N_Z$. The other three follow in the same way with coefficients 43, 41 and 24; the common value 70 is used for all. $\square$




Figure 2: Backward phase-transport weights (15) in the large- V form $n = \tau V$, divided by the stock weights $(1, 9/8, 7/5, 9/5)$. Free template contributes immediately and decays; the bound phases contribute after a lag through the actual release path. In the certified window $89V \leq n \leq 91V$ every ratio exceeds $1/35$, so $w(n) \cdot N \geq Y_N/35$.

Lemma 5.6 (Backward weights). *Put $q_c = 3000V$, $a = 1 - 70/q_c$, and define the weight vector $w(n) = (w_X, w_{C_1}, w_{C_2}, w_Z)(n)$ by $w(0) = (1, 0, 0, 0)$ and*

$$w(n+1) = \left(aw_X, aw_{C_1} + \frac{20}{q_c}w_X, aw_{C_2} + \frac{20}{q_c}(w_{C_1} + w_Z), aw_Z + \frac{38}{q_c}w_X \right).$$

Then $w(n)$ is nonnegative with $w_X + w_{C_1} + w_{C_2} + w_Z \leq 1$, and

$$w(n) = a^n \left(1, \frac{20n}{q_c a}, \frac{580n(n-1)}{q_c^2 a^2}, \frac{38n}{q_c a} \right). \quad (15)$$

For $89V \leq n \leq 91V$, every coordinate of $w(n)$ is at least $1/35$ times the corresponding weight of Y , so $w(n) \cdot N \geq Y_N/35$.

Proof. The closed form satisfies the recursion, and q_c times the loss of mass in one step is $12w_X + 50w_{C_1} + 70w_{C_2} + 50w_Z \geq 0$. For $n \leq 91V$, $a^n \geq e^{-11/5} \geq 1/12$ by $\log a \geq 1 - 1/a$ and $e^{11/5} \leq 12$. With $n \geq 89V$ and $n-1 \geq 88V$, the four coordinates are at least $\frac{1}{12}$, $\frac{1}{12} \cdot \frac{20 \cdot 89}{3000}$, $\frac{1}{12} \cdot \frac{580 \cdot 89 \cdot 88}{9 \cdot 10^6}$ and $\frac{1}{12} \cdot \frac{38 \cdot 89}{3000}$, which exceed $\frac{1}{35}(1, \frac{9}{8}, \frac{7}{5}, \frac{9}{5})$ coordinatewise; the tightest is C_2 , with 0.0421 against 0.0400. $\square$

The weights are the adjoint iteration of the one-step comparison kernel of (14): $w(n) \cdot N$ is a lower bound for the expected free count n steps ahead, and Figure 2 shows how each phase contributes. Any catalytic phase can fund a future free-template observation; the estimate does not assume that stock is initially free.

Proposition 5.7 (Low free template after a phase window). *Let N be an active corridor state with $Y_N \geq V/20$. For $89V \leq n \leq 91V$, the probability that the stopped process is still active after n clock steps with $N_X \leq V/1000$ is at most $e^{-V/10^{10}}$.*

Proof. Apply (2) to $e^{-s w \cdot N}$ for a nonnegative weight vector w of mass at most one. Every jump changes $w \cdot N$ by at most 2, and the total rate is at most q_c , so for $0 \leq s \leq 1/100$

$$\mathcal{L}e^{-s w \cdot N} \leq e^{-s w \cdot N} \{-s D_w(N) + 7200V s^2\}, \quad (16)$$

where D_w is the weighted right side of (14). One kernel step applied to $e^{-s w(k) \cdot N}$ is bounded by $e^{-s w(k+1) \cdot N} (1 + \frac{7200V}{q_c} s^2) = e^{-s w(k+1) \cdot N} (1 + \frac{12}{5} s^2)$, because $w(k+1) - w(k)$ is exactly $1/q_c$ times the adjoint of the comparison drift (14) applied to $w(k)$, so the drift term of (16) is absorbed by the change of weights. After n steps, $\mathbb{E}[\mathbf{1}_{\text{active}} e^{-s N_X(n)}] \leq e^{-s w(n) \cdot N} (1 + \frac{12}{5} s^2)^n$. With $s = 10^{-6}$, $w(n) \cdot N \geq V/700$ and the threshold $N_X \leq V/1000$, the exponential Markov inequality gives the exponent

$$\left(10^{-9} - \frac{1}{7 \cdot 10^8} + \frac{12}{5} \cdot 91 \cdot 10^{-12}\right)V \leq -\frac{V}{10^{10}}. \quad \square$$

Proposition 5.8 (Occupation and free collection). *From an active state with $Y \geq V/20$ at time 3 that remains active through time 4, the free-template counter fails its threshold, $Q_X < \lceil V/1080 \rceil$, with probability at most*

$$F(V) = e^{-V/(4 \cdot 10^7)} + 100e^{-V/10^{10}} + e^{-V} + e^{-V/200}. \quad (17)$$

Proof. Condition, at every clock step of the collection unit after a burn-in of $90V$ steps, on the state $90V$ steps earlier; Proposition 5.7 bounds the conditional probability of a low-free state by $e^{-V/10^{10}}$ whenever the earlier state had $Y \geq V/20$, which residence guarantees. If B counts low-free steps among the n collection steps, then $\mathbb{E}B \leq ne^{-V/10^{10}}$ and Markov's inequality gives $\mathbb{P}(B \geq n/100) \leq 100e^{-V/10^{10}}$, with no independence between the steps. On the complementary event at least 99% of the collection steps have $N_X > V/1000$, so the free-washout mark has mean at least $1/3000$ per such step. A marked lower-tail tilt of size $1/1000$ applied to the statistic $Q_X + B/(3 \cdot 10^6)$ gives, for $n \geq 2990V$ collection steps, the shortfall bound $\exp\{\frac{1}{1000}(\lceil V/1080 \rceil + \frac{n/100}{3 \cdot 10^6}) - \frac{33n}{10^{11}}\} \leq e^{-V/(4 \cdot 10^7)}$ for $V \geq 10^6$. The precollection quarter unit and the collection unit have clock means $750V$ and $3000V$; their lower tails contribute e^{-V} and $e^{-V/200}$. $\square$

The statistic $Q_X + B/(3 \cdot 10^6)$ is a proof device; the counted output is the actual free washout. The term $100e^{-V/10^{10}}$ is the dominant term of $e(V)$ (Proposition 3.4); the burn-in of $90V$ steps is a lag of about 0.03 time units and is why collection is preceded by three units of uncollected recovery rather than $11/4$.

6 Marked outputs, supplies and the joint counter budget

Lemma 6.1 (Template output). *During residence above $Y = V/20$ the template washout intensity satisfies $I_N \geq V/28$, and the template counter fails its threshold, $Q_I < \lceil V/56 \rceil$, on the active path with probability at most $e^{-V/100000}$ for $V \geq 10^6$.*

Proof. Comparing weights, $Y_N \leq \frac{7}{5}I_N$, so $I_N \geq V/28$. The washout of X, C_1, C_2, Z has total propensity I_N in molecules and mean template mark at least I_N per unit time, hence mean mark at least $1/84000$ per clock step. A lower-tail tilt of size $1/1000$ on a nonnegative mark bounded by 2 multiplies the test by at most $e^{-33/(2.8 \cdot 10^9)}$ per active step. Poisson summation over the collection clock of mean $3000V$ gives the exponent $\frac{V/56+1}{1000} + 3000V(e^{-33/(2.8 \cdot 10^9)} - 1) \leq -V/100000$, where the $+1$ certifies the integer ceiling. $\square$

Lemma 6.2 (Food and gross service). *Over the whole cycle, each food counter exceeds $5V$ with probability at most $e^{-V/300}$, and the gross driven-service counter exceeds $\lfloor V/5 \rfloor$ with probability at most $e^{-V/2000}$, for $V \geq 10^6$.*

Proof. For a counter with intensity at most cV , initial value z and threshold K , an upper-tail tilt of size s over four time units has exponent $-sK + 4cV(e^s - 1) + sz$. Each feed has constant rate V and the pulse dose is at most $151V/200$; with $c = 1$, $s = 1/50$, $K = 5V$ the exponent

is at most $-V/300$. On the corridor the gross pair-5 intensity is $dN_X + d\eta N_U N_W / V \leq \frac{1}{25}(\frac{11}{10} + \frac{121}{100}\eta)V < \frac{9}{200}V$; with $c = 9/200$, $s = 1/20$, $z = 0$ and $K = V/5 - 1$ the exponent is at most $-V/2000$, the -1 certifying the integer floor. $\square$

Proposition 6.3 (Joint counter budget on one law). *On the stopped device, from a prepared start, the probability that the path is active through time 4 and at least one of the five counters fails its threshold in (8) is at most*

$$J(V) = F(V) + e^{-V/100000} + 2e^{-V/300} + e^{-V/2000}. \quad (18)$$

Proof. All five counters are functions of the same sequence of labelled events, retained through the three analytical segments $[0, 11/4]$, $[11/4, 3]$, $[3, 4]$ with the actual intermediate chemical states. The output marginals hold their counters at zero until collection and the supply marginals accumulate from the pulse dose. Propositions 5.8 and Lemmas 6.1–6.2 bound each failure on that joint law; their sum is a union bound within it, not a product of marginal laws. $\square$

Proposition 6.4 (Prepared cycle). *From a start on the prepared band with $Y \geq 3V/250$, the probability that the stopped path exits, fails to return to $\mathcal{R}_V$ at time 4 while active, or fails a counter, is at most $J(V) + T(V)$, where*

$$T(V) = M(V) + R(V) + 4(e^{-V/320000} + e^{-V}) \quad (19)$$

is the restart budget, with $R(V)$ the recovery-and-residence budget of Appendix A.

Proof. Lemma 4.4 pays material exits, Proposition 5.3 and Lemma 5.4 pay stock failures through the actual entry and deadline states, Lemma 4.5 pays the terminal material bands, and Proposition 6.3 pays the counters. The events are combined by a union bound on the joint three-segment law. $\square$

Proposition 6.5 (Stopped one-cycle bound). *For $V \geq 10^6$, $N \in \mathcal{R}_V$ and admissible p , integrating over the pulse law, the stopped device fails the strict event (active terminal state, $N(4) \in \mathcal{R}_V$, all counter thresholds) with probability at most*

$$e(V) = e^{-1177V/10^9} + 4e^{-V/100000} + J(V) + T(V). \quad (20)$$

Proof. Bad pulses (Lemmas 4.2 and 4.3) are charged one; good pulses start on the prepared band with $Y \geq 3V/250$ and are charged $J(V) + T(V)$ by Proposition 6.4. Summing against the pulse masses, which total one, gives (20). $\square$

7 From the stopped device to the physical law

7.1 The strict event

Remark 7.1 (The closed restart event does not transfer). The restart stock inequality is closed ($Y \geq V/20$) and the residence stop is at $Y \leq V/20$, so a stopped path can end in $\mathcal{R}_V$ while frozen at the boundary. A three-state exact model in the accompanying workspace has stopped closed-restart probability 1 and unrestricted probability 1/2 after two steps. This refutes the generic transfer of the closed event; it is not a counterexample to the reactor theorem. The strictly active event in that model has probability 1/4 under both laws.

Lemma 7.2 (Strict terminal activity implies agreement). *Inactive states are absorbing in the stopped kernel. Hence, on the event that the terminal state is active, no earlier exit occurred, and the stopped and literal updates, including every counter increment, agree at every step. The probability of losing the strict event is already contained in $T(V)$.*

7.2 Killed chronology and the stopped clock

Lemma 7.3 (Renewal identity). *Let D be a set of states on which the total rate is bounded. The law of the chronological jump process killed at its first exit from D and the law of the uniformized stopped clock, both with marks and waiting times, satisfy the same first-event renewal equation, whose solution on a bounded time interval is unique. Hence the two killed laws agree, and the stopped-clock expectation of any nonnegative terminal payoff vanishing off D is at most the corresponding expectation under the unrestricted chronological process.*

Proof. Expanding on the first waiting time and first label gives the equation directly for each construction; a null step of the clock carries no mark and no state change, and the survival kernel at the genuine rate solves the resolvent equation of the clock rate, including the case of equal rates. Uniqueness follows by iterating the residual, whose remainder after k iterations on an interval of length t is bounded by the tail of the exponential series of the rate bound. Domination by the unrestricted process is monotonicity of a payoff that vanishes off D . $\square$

The identity preserves waiting times and labels, not merely a state marginal, which is what the counters require.

7.3 Nonexplosion and the chronological semigroup

Lemma 7.4 (Nonexplosion). *For every $V > 0$, $r, d \geq 0$ and every count start, the twenty-channel chronological process has almost surely finitely many events on every bounded time interval.*

Proof. Only the two constant-rate feeds increase the total material $A_N + B_N$; their number on a bounded interval is Poisson and almost surely finite. Every region of bounded material height contains finitely many count states and has bounded total rate, so the process makes finitely many jumps before the next feed. No global rate bound is assumed [cf. 9, 24]. $\square$

Proposition 7.5 (Chronological semigroup and exhaustion). *Exhaust the state space by regions of bounded material height. The killed endpoint laws increase with the region and their supremum is the unrestricted chronological law. The composition identity $P_{t+s} = P_t P_s$ holds for the bounded uniformized clock by Poisson convolution, passes to each killed law, and passes to the limit by monotone convergence, a common larger region handling the two factors. Consequently the analytical split $11/4 + 1/4$ of the uncollected run is exactly the physical three-unit law, the collection unit with marks switched on is the physical marked law, and erasing the counters at time 4 recovers the original count endpoint.*

Every step of Lemmas 7.2–7.4 and Proposition 7.5 is compiled in Lean (Appendix D); in the companion papers [1, 2] the identification was a conventional argument.

7.4 The literal cycle bound

Theorem 7.6 (One cycle on the physical law). *For $V \geq 10^6$, every $N \in \mathcal{R}_V$ and every admissible p , the literal pulse-and-source cycle measure gives the successful records of Definition 2.4 probability at least $1 - e(V)$.*

Proof. By Proposition 6.5 the strict stopped event has probability at least $1 - e(V)$. By Lemma 7.2 and Lemma 7.3 the strict event is dominated by the corresponding event of the killed chronological process, and by Proposition 7.5 the composed killed laws are the literal four-unit marked law. The strict event implies the successful-record event, and the pulse integration is the same on both sides. $\square$

This is part (a) of Theorem 3.1.

8 Repetition, inventory and quantitative consequences

8.1 Conditioning on actual returns

Proposition 8.1 (Product bound). *For the returned-history process of Definition 2.5, and for every measurable history model with the literal conditional law and actual-return property, $\mathbb{P}(E_m) \geq (1 - e(V))^m$ for all m .*

Proof. Let $K(h, \cdot)$ be the one-step history kernel and S the set of histories whose newest record is successful. A successful record places the actual next state in $\mathcal{R}_V$, so Theorem 7.6 gives $K(h, S) \geq 1 - e(V)$ whenever the current state of h is in $\mathcal{R}_V$. Integrating over successful prefixes,

$$\mathbb{P}(E_{m+1}) = \int_{E_m} K(h, S) \mathbb{P}(\mathrm{d}h) \geq (1 - e(V))\mathbb{P}(E_m),$$

and induction gives the product. The measure restricted to successful prefixes is an unnormalized submeasure of the full law whose mass is the success probability; it is never renormalized into a conditional sampling procedure, and no independence between cycles is used. Bernoulli's inequality $(1 - e)^m \geq 1 - me$ gives the linear form. $\square$

This is part (b); part (c) is summation of (8) over the records of a history in E_m .

8.2 Pathwise inventory and net synthesis

Lemma 8.2 (Inventory identity). *Along any physical trace of a recorded history,*

$$P_{\text{net}} = I_{\text{final}} - I_{\text{initial}} + Q_{\text{all washout}} + I_{\text{withdrawn}} + I_{\text{extra loss}}, \quad (21)$$

where $Q_{\text{all washout}}$ counts the template weight of every washout in every cycle, collected or not, and the last two terms count the template weight of withdrawn and extra-lost molecules at every pulse.

Proof. Each chemical event changes I by its synthesis mark, each feed by zero, each washout by minus its template weight, and each pulse by minus the template weight of the molecules that leave. Telescoping over all events and pulse boundaries gives (21). $\square$

Proposition 8.3 (Net synthesis). *On E_m , every physical trace satisfies (10).*

Proof. All removal terms in (21) are nonnegative and $Q_{\text{all washout}} \geq \sum_k Q_{I,k} \geq m[V/56]$. The final state is in $\mathcal{R}_V$, and comparing (40, 45, 56, 72) with (56, 56, 56, 112) gives $40Y_N \leq 56I_N$, so $I_{\text{final}} \geq V/28$, hence $\geq [V/28]$ since inventory is an integer. Initially $I_{N_0} \leq A_{N_0} \leq 161V/160$, and the integer inventory is at most $[161V/160]$. Dropping the ceilings gives the real form. The formal event contains the existence of a trace, so the universal statement over traces is not vacuous. $\square$

At $m = 55$ the real bound equals $13V/1120 > 0$, and for $I_{N_0} = 8V/125$ two cycles give $2[V/56] + [V/28] - 8V/125 \geq 13V/1750$. The statement is net synthesis inside the maintained source; it does not label exported molecules, isolate the contribution of the templated path, or prove pathway necessity.

8.3 The single exponential and the sizing rule

Lemma 8.4 (Two-term envelope). *For every $V \geq 0$,*

$$e(V) \leq 100e^{-V/10^{10}} + 10^6(V + 1)e^{-V/(4 \cdot 10^7)}. \quad (22)$$

Cycles m	Earlier linear rule $\max(V_0, \lceil 5 \cdot 10^6 m / \delta \rceil)$	Logarithmic rule, sufficient scale
100	200,000,000,000	200,000,000,000
10,000	5,000,000,000,000	200,000,000,000
1,000,000	500,000,000,000,000	230,359,000,000

Table 2: Sufficient integer copy scales for $\delta = 0.01$. The last column lists outward-rounded scales certified by exact rational Taylor lower bounds for $e^{V/10^{10}} > 10100m$; it does not assert floating-point ceilings. The necessary scale of this certificate at $m = 100$ is 1.38×10^{11} (Proposition 3.4).




Figure 3: Left: the exact certificate $e(V)$, the two-term envelope (22) and the single exponential $101e^{-V/10^{10}}$, evaluated in logarithmic form for $V \geq V_0$; the first two are visually coincident. Right: the certified mission length $\lfloor (\delta/101)e^{V/10^{10}} \rfloor$ at $\delta = 0.01$. All curves are proved upper bounds, not measured failure frequencies.

Proof. Every term of $e(V)$ other than $100e^{-V/10^{10}}$ decays at least as fast as $e^{-V/(4 \cdot 10^7)}$. The constant coefficients sum to 31, and the two polynomial prefactors, with $e^{1/50}, e^{9/500} \leq 2$, contribute at most $216000V$; both are below $10^6(V+1)$. $\square$

Lemma 8.5 (Residual absorption). *For $V \geq V_0$, $10^6(V+1)e^{-V/(4 \cdot 10^7)} \leq e^{-V/10^{10}}$; hence $e(V) \leq 101e^{-V/10^{10}}$.*

Proof. Put $x = 249V/10^{10}$, so that $e^{-V/(4 \cdot 10^7)} = e^{-V/10^{10}}e^{-x}$ and $x \geq 4980$. The Taylor term gives $e^x \geq x^6/720 \geq 4980^5 x/720$, and $4980^5 \cdot 249V/(720 \cdot 10^{10}) > 10^8 V \geq 10^6(V+1)$. This is an exact polynomial comparison on the whole domain $V \geq V_0$. $\square$

Proof of Corollary 3.2. By Lemma 8.5 and (11), $V/10^{10} \geq \log(101m/\delta)$ gives $me(V) \leq 101me^{-V/10^{10}} \leq \delta$; the mission-length form is the same inequality solved for m . For the fixed-scale statements one exact rational certificate suffices:

$$e^{20} \geq \sum_{k=0}^{36} \frac{20^k}{k!} > 484\,800\,000, \quad (23)$$

so $e(V_0) \leq 101e^{-20} \leq 1/4\,800\,000$. Multiplying by 48 000 gives $1/100$; multiplying by 100 gives $1/48\,000 < 21/10^6$. $\square$

Proof of Proposition 3.4. The term $100e^{-V/10^{10}}$ is a summand of $F(V)$ and all other summands of $e(V)$ are nonnegative. If $me(V) \leq \delta$ then $100me^{-V/10^{10}} \leq \delta$, which is the stated inequality. $\square$

Quantity over 100 cycles at $V = 2 \times 10^{11}$	Count	Approx. pmol
Collected template equivalents	$\geq 357\,142\,857\,200$	0.593050
Collected free X (included above)	$\geq 18\,518\,518\,600$	0.0307507
Each food expenditure	$\leq 100\,000\,000\,000\,000$	166.054
Gross driven events	$\leq 4\,000\,000\,000\,000$	6.64216
Final template inventory	$\geq 7\,142\,857\,143$	0.0118610
Net synthesis, arbitrary restart	$\geq 163\,035\,714\,343$	0.270727
Initial inventory, richer witness	12 800 000 000	0.0212549
Net synthesis, richer witness	$\geq 351\,485\,714\,343$	0.583656

Table 3: Hundred-cycle certified counts and their amount equivalents. Amounts divide the exact count by the SI Avogadro constant $N_A = 6.02214076 \times 10^{23} \text{ mol}^{-1}$; the decimal column is descriptive and is not a sharpened inequality.

8.4 Operating ratios

Proof of Corollary 3.3. On E_m with $m \geq 1$, $Q_I \geq m[V/56] > 0$ and $Q_X \geq m[V/1080] > 0$. Dividing the mission bounds of part (c) by $4mV$, respectively by the output floors, gives each ratio; for instance $G/Q_I \leq m[V/5]/(m[V/56]) \leq 56/5$. The demand rule combines $V_{\log}$ with $m[V/56] \geq D_I$ and $m[V/1080] \geq D_X$. $\square$

Source-running time is $4m$; pulse handling and preparation are excluded. The ratios are operating allowances on one event, not optimized stoichiometric yields, and the collected output is mixed effluent.

9 A dimensional reading

At $V = V_0$ the hundred-cycle event has probability at least 99.9979%. Table 3 lists the certified counts with exact integer rounding. A richer restart witness is

$$N_0 = (186\,800\,000\,000, 187\,000\,000\,000, 12\,000\,000\,000, 200\,000\,000, 200\,000\,000, 200\,000\,000),$$

with $A = B = V_0$, stock $514.6 \times 10^9 \geq 2V_0$ and $I_{N_0} = 12\,800\,000\,000 = 8V_0/125$. Its existence is separate from the reliability guarantee, which holds for every restart state.

Choose concentration and time scales c_*, τ_* and define

$$V = N_A \Omega c_*, \quad t_{\text{phys}} = \tau_* t, \quad k_{\text{dim}} = \frac{k_{\text{norm}}}{\tau_* c_*^{k-1}} \quad (24)$$

for a reaction of full reactant molecularity k ; the reservoir concentrations corresponding to unit activities are c_* . The full driven forward reaction is bimolecular and its reverse trimolecular; substituting the maintained reservoir concentrations produces the effective propensities of Table 1, and Appendix C lists the complete conversion. Table 4 gives two illustrative realizations. A change of units changes the implied physical kinetics; the examples are not calibrated protocols, and the small volume by itself establishes no feasibility claim. What the theorem could serve as is a quantitative benchmark for a maintained cell-free or synthetic reactor: preparation, supply allowances, finite-run output and confidence are specified once a matching source and handling law are justified.

10 Scope, verification boundary and open questions

What the theorem provides. A finite-copy operating certificate for repeated harvesting: a restart set, a withdrawal class, a recovery schedule, output thresholds and supply allowances,

	$c_* = 1 \mu\text{M}, \tau_* = 1 \text{ s}$	$c_* = 1 \text{ mM}, \tau_* = 60 \text{ s}$
Reactor volume Ω	$0.332108 \mu\text{L}$	0.332108 nL
100-cycle source time	400 s	$24\,000 \text{ s (6 h 40 min)}$
Bimolecular coefficient 20	$2 \times 10^7 \text{ M}^{-1}\text{s}^{-1}$	$333.3 \text{ M}^{-1}\text{s}^{-1}$
Washout rate	1 s^{-1}	$(1/60) \text{ s}^{-1}$

Table 4: Two illustrative realizations of the same normalized model at $V = 2 \times 10^{11}$.

and a probability bound for the joint event over any finite number of cycles, uniform over the parameter rectangle and over history-dependent controllers, together with a proof that the exported material is synthesized rather than inherited. In the language of serial transfer [6, 21, 23] it states the transfer fraction, the recovery time, the yield and the confidence for which the finite-molecule system is self-sustaining over a mission. In the language of continuous culture [28] it is a return theorem for a pulsed stochastic chemostat with an impulsive map [20], proved by scalar generator comparison. Relative to the companion deterministic theorem [4, Theorem 3.1], the four obligations it listed are discharged: the pulse is specified at the molecular level (Definition 2.2); the post-pulse tails are conditional on the actual pre-pulse state (Lemmas 4.2–4.3); return, output and supplies are proved on one continued count law with free export counted as free washout (Sections 5–7); and the conditional one-cycle bound is iterated over actual returns (Proposition 8.1). Its entry band [4, Proposition 10.1] is the deterministic shadow of the prepared band and seed used here.

Relation to the companion stochastic certificates. The finite-copy exporter papers [1, 2] certify operation from food alone over a fixed or exponential observation horizon for a reactor that is prepared once and never disturbed. The present theorem starts from established stock and certifies operation across withdrawals; the two are complementary, and the necessary initiation scale of [5, Theorem 6.1] is the statement that governs a food-only start. The finite-count inheritance analysis [3] concerns a different source and is not a prerequisite.

What is Lean-verified and what is conventional. Lean-verified: the twenty-label source and the identities (4); the pulse law, the retention factor and both pulse tails; the material corridor, exit and terminal return; guarded growth, the stock test, seeded recovery, residence; the phase comparison, the backward weights, the window bound, occupation and free collection; the template, food and service tails and their union on one law; the prepared and pulse-integrated cycle bounds; strict safety and the boundary semantics; the renewal identity, nonexplosion, the chronological semigroup and exhaustion, the literal cycle identification; the list-history process, the conditional-law model, the product bound and the full-law event; trace existence and the inventory identity; the envelope, absorption, logarithmic sizing, the fixed-scale certificates, the synthesis corollaries, the ratios and the bottleneck. Appendix D maps each to its declaration. Conventional: the amount conversions and Table 4; the two example scales beyond V_0 in Table 2, certified by exact rational Taylor sums in the replay script; and the prose of this paper, which the Lean kernel does not check.

Limitations. Maintained activities, schematic coefficients and an ideal, conditionally independent molecular pulse are assumptions; the effective reverse driven channel has no supplied elementary implementation, and thermochemical consistency of the rate ratios does not identify molecules. Established stock must be prepared, and the first-seed result of [5] does not by itself provide the restart condition or its preparation cost. Gross driven events measure a specified chemical service; with chemical potentials a signed fuel–waste exchange supports a work calculation for that exchange, but neither is an energy bill for feed, washout, mixing,

pumping, handling, preparation or purification [26, 27]. Output is mixed effluent. The theorem certifies finite missions at a chosen scale, not eternal survival at one volume, food-only startup, autonomous operation with finite baths, a causal knockout comparison or the performance of any named biochemical system.

Open questions. Four seem most natural. First, the phase bottleneck: the certificate’s scale is set entirely by the term $100e^{-V/10^{10}}$ of the occupation estimate, whose tilt 10^{-6} is forced by the quadratic remainder of the phase test over $91V$ clock steps; a sharper transport estimate, for instance with a state-dependent tilt or a two-window argument, would lower V_0 by orders of magnitude, whereas recovery contributes only $e^{-3V/5 \cdot 10^6}$. Second, a large-deviation form of the same question [7]: the present constants are sufficient, not sharp, and no rate function is claimed. Third, feedback: whether a controller acting on the returned history can be chosen to maximize certified yield per unit of food or service, a stochastic control problem on the history chain. Fourth, food-only startup with a subsequent switch to the harvesting protocol, which would join the initiation certificate of [5] to the present restart set through an explicit preparation cycle.

Data and code availability. The repository location will be supplied with the author information. The Lean sources, the strict verification receipts, the figure script and an exact replay script for every printed constant accompany the manuscript source.

A Complete error budget

For $V \geq 0$ define

$$\begin{aligned} M(V) &= 4e^{-V/2000} + 48000V e^{-V/2000+1/50}, \\ R(V) &= e^{-3V/5000000} + e^{-V/2} + M(V) + e^{-V/10000} + 12000V e^{-V/10000+9/500}, \\ F(V) &= e^{-V/40000000} + 100e^{-V/10000000000} + e^{-V} + e^{-V/200}, \\ J(V) &= F(V) + e^{-V/100000} + 2e^{-V/300} + e^{-V/2000}, \\ T(V) &= M(V) + R(V) + 4(e^{-V/320000} + e^{-V}), \\ e(V) &= e^{-1177V/1000000000} + 4e^{-V/100000} + J(V) + T(V). \end{aligned}$$

This is exactly the source’s `oneCycleError`; repeated material terms are retained rather than optimized away. J is the joint five-counter budget and T the restart budget; the two pulse terms are outside the prepared-cycle budget. Figure 4 shows the residual of the two-term envelope in logarithmic form; values far below floating-point range remain finite logarithms and are never interpreted as zero.

B Quantitative exponential tests

This appendix records the inequalities underlying the cycle estimates so that the probabilistic argument can be read without the implementation modules.

Elementary inequalities. For $|z| \leq 9/50$, $e^z \leq 1 + z + \frac{3}{5}z^2$. For $0 \leq x$, $e^x \geq \sum_{k \leq n} x^k/k!$. For $0 < a < 1$, $\log a \geq 1 - 1/a$. For a Poisson variable of mean μ and $k < \mu$, $\mathbb{P}(\text{Poisson}(\mu) \leq k) \leq e^{-(\mu-k)^2/(2\mu)}$ [11].




Figure 4: The polynomial-exponential residual of the envelope (22) against the phase term, for $V \geq V_0$. Lemma 8.5 proves the ordering exactly by a sixth-order Taylor comparison.

Pulse. With independent Bernoulli indicators B_a of parameters p_a and weights $w_a \geq 0$, $\mathbb{E} \exp(s \sum_a w_a B_a) = \prod_a (1 - p_a + p_a e^{s w_a})$ exactly. Stock weights are at most $9/5$ and Lemma 4.1 applies at $s = 1/100$ with $\sum p_a w_a \geq 49V/4000$ and level $3V/250$, giving $-1177V/10^9$. Material weights are at most 2; the unrounded post-pulse material means lie in $[31213V/32000, 3231V/3200]$ and the floor lowers a mean by less than one, leaving the slack required for $[24V/25, 51V/50]$ at $V \geq 10^6$; the bounded-weight exponential estimate gives $e^{-V/100000}$ per sign and coordinate.

Material barriers and terminal return. For $H \in \{A, B\}$ the barriers $U_H = e^{(H-11V/10)/100}$ and $L_H = e^{-(H-9V/10)/100}$ sum to at least one on exit and to at most $4e^{-V/2000}$ over both coordinates at a prepared start; on the corridor each has generator at most $3000V e^{-V/2000+1/50}$, since outside the inner band $|H - V| \leq V/20$ the restoring drift dominates and inside it a jump of size at most 2 shifts the exponent by at most $1/50$. Integration over four units gives $M(V)$. For terminal return the compensated tilt $e^{-2s^2 V} T_s$ contracts its parameter by $1 - 1/(3000V)$ per step; $n \geq 11700V$ gives a factor at most $1/40$; $s = \pm 1/1000$ and $|H_0 - V| \leq V/25$ give exponent at most $-V/320000$ for $|H - V| \geq V/160$; the clock of mean $12000V$ misses $11700V$ with probability at most e^{-V} .

Seeded recovery and residence. Below $3V/50$ the drift is at least $\frac{2}{3}Y + V/10^9$ and the rate-weighted squared stock jumps are at most $6Y + V/10^8$, so $\mathcal{L}e^{-sY} \leq -\frac{5}{8}sY e^{-sY}$ for $s \leq 1/100$. The killed entry test satisfies $PT_s \leq T_t$ for $t \leq (1 + \frac{5/8}{3000V})s$; starting at $1/1000$ and capping at $101/20000$, the block estimate $(1 + \frac{5/8}{3000V})^{8100V} \geq (17/16)^{27} > 101/20$ shows the cap is reached; the exponential Markov inequality gives $\exp\{\frac{1}{1000} \cdot \frac{3V}{50} - \frac{101}{20000} \cdot \frac{3V}{250}\} = e^{-3V/5000000}$; the clock of mean $8250V$ misses $8100V$ with probability at most $e^{-V/2}$. Residence uses $H_Y = \max(e^{-Y/100}, e^{-3V/5000})$ with generator at most $3000V e^{-3V/5000+9/500}$, worth $e^{-V/2000}$ at a lower exit and $e^{-3V/5000}$ at the start; integration over $t \leq 4$ gives $e^{-V/10000} + 12000V e^{-V/10000+9/500}$.

Phase margin and collection tilt. In (15), $a^n \geq 1/12$ for $n \leq 91V$, and $n \geq 89V$, $n - 1 \geq 88V$ give each coordinate at least $1/35$ of its stock weight. The remainder per step is $\frac{12}{5}s^2$; at $s = 10^{-6}$ the total over $91V$ steps is at most $\frac{12}{5} \cdot 91V \cdot 10^{-12}$. The threshold $N_X \leq V/1000$ contributes $V/10^9$ and the initial stock $-V/(700 \cdot 10^6)$; the sum is at most $-V/10^{10}$. After a burn-in of $90V$ steps this bounds low-free occupancy at every later step; with B the number of bad steps among n , $\mathbb{E}B \leq n e^{-V/10^{10}}$ and $\mathbb{P}(B \geq n/100) \leq 100e^{-V/10^{10}}$. The marked tilt of size $1/1000$ on $Q_X + B/(3 \cdot 10^6)$ gives $\exp\{\frac{1}{1000}(K + \frac{B_*}{3 \cdot 10^6}) - \frac{33n}{10^{11}}\} + \mathbb{P}(B \geq B_*)$ with $K = V/1080 + 1$, $B_* = n/100$, $n \geq 2990V$; for $V \geq 10^6$ the first term is at most $e^{-V/40000000}$. The precollection quarter unit and the collection unit have clock means $750V$ and $3000V$, contributing e^{-V} and $e^{-V/200}$.

Pair	Directional molecularities	k_{dim}^+	k_{dim}^-
0	2, 1	$\varepsilon/(\tau_* c_*)$	$\varepsilon/(10\tau_*)$
1	2, 1	$20/(\tau_* c_*)$	$20/\tau_*$
2	2, 1	$20/(\tau_* c_*)$	$20/\tau_*$
3	1, 1	$20/\tau_*$	$2/\tau_*$
4	1, 2	r/τ_*	$r/(\tau_* c_*)$
5	2, 3	$d/(\tau_* c_*)$	$d\eta/(\tau_* c_*^2)$

Table 5: Normalized and dimensional full-reaction coefficients.

Template and supply arithmetic. For a nonnegative mark bounded by 2, the one-step lower-tail test at tilt $1/1000$ is controlled by the mean mark; template washout has mean mark at least $1/84000$ per active step, hence multiplier at most $e^{-33/2800000000}$, and Poisson summation with mean $3000V$ gives $\frac{V/56+1}{1000} + 3000V(e^{-33/2800000000} - 1) \leq -V/100000$ for $V \geq 10^6$. For a supply with rate at most cV , initial count z and threshold K the upper-tail exponent is $-sK + 4cV(e^s - 1) + sz$: $c = 1$, $s = 1/50$, $z \leq 151V/200$, $K = 5V$ give at most $-V/300$; $c = 9/200$, $s = 1/20$, $z = 0$, $K = V/5 - 1$ give at most $-V/2000$ for $V \geq 10^6$.

C Complete dimensional coefficients

Table 5 gives the full chemical rate coefficients under (24). A coefficient for molecularity k has units $\text{M}^{1-k}\text{s}^{-1}$. The effective pair-5 coefficients after substituting $[F] = [P] = c_*$ are d/τ_* (unimolecular forward) and $d\eta/(\tau_* c_*)$ (bimolecular reverse); they must not be confused with the full chemical coefficients. Either food feed becomes a concentration input c_*/τ_* and a molecular arrival rate V/τ_* ; every washout coefficient is $1/\tau_*$; the pulse food counts and retention probabilities remain dimensionless. Preparation and handling durations lie outside the four-unit source interval.

D Formal evidence and reproduction

The frozen environment is Lean 4 (toolchain `leanprover/lean4:v4.30.0`) with the repository’s pinned Mathlib manifest. Verification treats compiler warnings as errors, exports declaration types and reports the axioms used. All declarations live in the namespace `FiniteCopyReactor` under `proofs/FiniteCopyReactor/`; the count source is bound to the realization of [5] through `CommonPhysicalRealization.ExporterSource`. The original root `Resolution.lean` exports `finite_copy_reactor_horizon` (the full-law event with the linear bound), `finite_copy_reactor_result` (the earlier linear sizing with an explicit restart witness), `finite_copy_reactor_hundred` and `finite_copy_reactor_adapted_history`. The publication root `PublicationResolution.lean` adds `publication_logarithmic`, `publication_48000`, `publication_hundred` and `certified_mission_budget`, and its strict receipt authenticates six interfaces: these three together with `full_operational_product`, `HistoryTrace.sharp_uniform_integer` and `operational_ratios`. Each depends only on the standard axioms `propext`, `Classical.choice` and `Quot.sound`; there are no admitted placeholders, no custom axioms and no warning suppression. The publication receipt records exit code 0, empty compiler output, 238 current local dependency sources and a source SHA-256 beginning `bc121e4a`; the original root receipt has 234 dependency sources and root SHA-256 beginning `6ec29251`.

The compiled event `OperationalSuccess N V m` policy is the set of histories of length m whose every record satisfies Definition 2.4, for which a physical trace exists, whose five totals satisfy the bounds of part (c), and whose every trace satisfies the real form of (10); `full_`

`operational_product` bounds its probability under the m -step list-history kernel from the empty list by $(1 - e(V))^m$, and `physical_history_product` does the same for every structure `PhysicalHistory` satisfying the conditional-law and actual-return fields. The intervention structure carries the bounds (6) as proof obligations, and the parameter hypotheses are $19 \leq r \leq 21$ and $0 \leq d \leq 1/25$ in the stochastic chain, $1/50 \leq d$ at the publication interfaces.

Claim	Principal declarations or modules
Literal twenty-label source and identities (4) (Section 2.1)	Source: <code>material_A_generator</code> , <code>material_B_generator</code> , <code>stock_generator_correction</code> , <code>count_guarded_growth</code> ; <code>JointCountSource</code>
Restart set and witnesses (Definition 2.1)	Restart: <code>restart_witness</code> , <code>restart_nonempty</code> ; <code>all_free_restart</code>
Molecular pulse law, retention factor, stock tail (Lemmas 4.1, 4.2)	Pulse: <code>pulseMass_total</code> , <code>pulse_laplace</code> , <code>retention_factor_bound</code> , <code>pulse_chernoff</code> ; <code>PulseStock</code> ; <code>PulseInventory</code>
Pulse material tails and doses (Lemma 4.3)	<code>PulseMaterial</code> , <code>PulseMaterialProbability</code> , <code>PulseDeviation</code> , <code>pulse_food_budget</code>
Corridor, uniformization, exit, terminal return (Lemmas 4.4, 4.5)	<code>MaterialModel</code> , <code>MaterialPulse</code> , <code>MaterialReturn</code> , <code>TerminalMaterial</code> , <code>TerminalMaterialProbability</code> , <code>SwitchedMaterial</code>
Stock test and seeded recovery (Lemma 5.2, Proposition 5.3)	<code>StockExponential</code> , <code>EntrySchedule</code> , <code>EntryProbability</code> : <code>seeded_recovery_deadline</code>
Residence and actual continuation (Lemma 5.4)	<code>ResidencePotential</code> , <code>Residence</code> , <code>ClampedResidence</code> , <code>CollectionResidence</code> , <code>DeadlineStock</code>
Phase drifts, weights, window (Lemmas 5.5, 5.6, Proposition 5.7)	<code>PhaseComparison</code> , <code>PhaseWeights</code> , <code>PhaseWindow</code> , <code>PhaseTime</code> , <code>PhaseProbability</code> : <code>phase_discrete_probability</code>
Occupation and free collection (Proposition 5.8)	Occupation: <code>occupation_after_burnin</code> , <code>bad_occupation_bound</code> ; FreeCollection: <code>free_collection_probability</code> ; <code>FreeThreshold</code>
Template, food, service tails (Lemmas 6.1, 6.2)	<code>Template</code> , <code>TemplateThreshold</code> , <code>SupplyBudget</code> , <code>SupplySource</code> , <code>RoundedService</code> , <code>MarkedUpper</code>
Joint counter budget on one law (Proposition 6.3)	<code>JointMarks</code> , <code>JointCycle</code> , <code>JointTails</code> , <code>JointSupplyTails</code> , <code>JointCounterBudget</code> : <code>joint_counter_budget</code>
Prepared and pulse-integrated cycle (Propositions 6.4, 6.5)	<code>PreparedCycle</code> , <code>StateRestart</code> , <code>PulseCycle</code> : <code>pulse_cycle_failure_bound</code> , <code>oneCycleError</code>
Strict safety and boundary semantics (Lemma 7.2, Remark 7.1)	<code>SafeCycle</code> , <code>StoppingSemantics</code> , <code>CountCycleEvent</code>
Renewal identity and killed comparison (Lemma 7.3)	<code>StoppedAgreement</code> : <code>stopped_clock_eq_killed_chronology</code> , <code>stopped_clock_le_unrestricted</code> ; <code>KilledChronology</code> , <code>SurvivalResolvent</code> , <code>ScalarRenewal</code>
Nonexplosion (Lemma 7.4)	<code>CountNonexplosion</code> , <code>JointTrajectory</code> : <code>joint_trajectory_nonexplosion</code> ; <code>ActualSource</code>

Claim	Principal declarations or modules
Chronological semigroup, exhaustion, literal cycle law (Proposition 7.5)	SourceSemigroup: chronological_endpoint_semigroup; JointTimeComposition: joint_source_cycle_kernel_literal; JointPhysicalEndpoint
One cycle on the physical law (Theorem 7.6)	LiteralCycleBound: literal_cycle_success_bound
History process and product bound (Definition 2.5, Proposition 8.1)	ReturnedHistory, PhysicalHistory, HistorySurvival, HistoryFullLaw, OperationalCorollaries: physical_history_product, full_operational_product
Trace existence and inventory (Lemma 8.2, Proposition 8.3)	PhysicalHistoryInventory: full_history_has_trace, HistoryTrace.inventory_lower; SynthesisCorollaries: restart_final_inventory, HistoryTrace.sharp_uniform_integer, HistoryTrace.positive_after_55, HistoryTrace.rich_two_cycle, HistoryTrace.hundred_uniform_integer, HistoryTrace.hundred_rich_integer
Envelope, absorption, sizing (Lemmas 8.4, 8.5, Corollary 3.2)	ErrorEnvelope: one_cycle_error_envelope; LogarithmicScale: residual_error_absorbed, one_cycle_single_exponential, logarithmic_volume_error, exp_twenty_certificate, mission_48000_budget, hundred_sharp_budget; PublicationResolution
Earlier linear rule (Remark 3.7)	EffectiveScale: volume_times_error, sufficient_volume_error; finite_copy_reactor_result
Ratios (Corollary 3.3)	operational_ratios
Bottleneck (Proposition 3.4)	ErrorBottleneck: one_cycle_error_phase_lower, error_certificate_volume_limit
Theorem 3.1, all parts	finite_copy_reactor_horizon; full_operational_product; publication_budget

Reproduction. From the repository root, with the canonical interpreter and the pinned wrapper (one command per line, continuation lines joined):

```

.\.venv\Scripts\python.exe scripts/check_python.py -quiet
.\.venv\Scripts\python.exe scripts/verify_proof.py
proofs/FiniteCopyReactor/PublicationResolution.lean
-declaration FiniteCopyReactor.publication_logarithmic
-declaration FiniteCopyReactor.publication_48000
-declaration FiniteCopyReactor.publication_hundred
-declaration FiniteCopyReactor.full_operational_product
-declaration FiniteCopyReactor.HistoryTrace.sharp_uniform_integer
-declaration FiniteCopyReactor.operational_ratios

```

`-output [workspace]/PublicationResolution.verify.json`

Here `[workspace]` is the campaign folder that holds the receipts. The wrapper runs Lean serially in the pinned `mathlib4_project` directory with warnings as errors; no dependency update is required. The script `check_paper.py` accompanying this source replays every constant printed above (132 checks: exact generator identities, pulse means and exponents, corridor rates, tilt schedules, phase weights, counter exponents, the envelope, the Taylor certificates, the counts, the ratios and the unit conversions) and `make_figures.py` regenerates the figures from the proved bounds.

References

- [1] Anonymous. Reliable covalent export from an autocatalytic reactor under driven cleavage: explicit finite-copy operating, output and supply certificates over an exponential horizon. Companion manuscript, 2026.
- [2] Anonymous. Finite-copy productive operation in a reversible autocatalytic binding network: an explicit probability guarantee with a machine-checked proof. Companion manuscript, 2026.
- [3] Anonymous. Finite-count chemical inheritance with reversible interactions and bounded food supplies. Companion manuscript, 2026.
- [4] Anonymous. Productive recovery and repeated harvesting in a reversible autocatalytic reactor: a uniform operating theorem with a machine-checked proof. Companion manuscript, 2026.
- [5] Anonymous. Thermochemically consistent realization of a stochastic autocatalytic exporter: chemical completion, inherited operating certificates, cycle affinities and a necessary initiation scale. Companion manuscript, 2026.
- [6] Paul Adamski, Marcel Eleveld, Ankush Sood, Ádám Kun, András Szilágyi, Tamás Czárán, Eörs Szathmáry, and Sijbren Otto. From self-replication to replicator systems en route to de novo life. *Nature Reviews Chemistry*, 4:386–403, 2020. doi: 10.1038/s41570-020-0196-x.
- [7] Andrea Agazzi, Amir Dembo, and Jean-Pierre Eckmann. Large deviations theory for Markov jump models of chemical reaction networks. *The Annals of Applied Probability*, 28(3):1821–1855, 2018. doi: 10.1214/17-AAP1344.
- [8] David F. Anderson and Jinsu Kim. Some network conditions for positive recurrence of stochastically modeled reaction networks. *SIAM Journal on Applied Mathematics*, 78(5):2692–2713, 2018. doi: 10.1137/17M1161427.
- [9] David F. Anderson and Thomas G. Kurtz. *Stochastic Analysis of Biochemical Systems*, volume 1.2 of *Mathematical Biosciences Institute Lecture Series*. Springer, Cham, 2015. doi: 10.1007/978-3-319-16895-1.
- [10] Alex Blokhuis, David Lacoste, and Philippe Nghe. Universal motifs and the diversity of autocatalytic systems. *Proceedings of the National Academy of Sciences*, 117(41):25230–25236, 2020. doi: 10.1073/pnas.2013527117.
- [11] Herman Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *The Annals of Mathematical Statistics*, 23(4):493–507, 1952. doi: 10.1214/aoms/1177729330.
- [12] Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming language. In *Automated Deduction – CADE 28*, volume 12699 of *Lecture Notes in Computer Science*, pages 625–635. Springer, 2021. doi: 10.1007/978-3-030-79876-5_37.
- [13] Winfried K. Grassmann. Transient solutions in Markovian queueing systems. *Computers & Operations Research*, 4(1):47–53, 1977. doi: 10.1016/0305-0548(77)90007-7.

- [14] Ankit Gupta, Corentin Briat, and Mustafa Khammash. A scalable computational framework for establishing long-term behavior of stochastic reaction networks. *PLoS Computational Biology*, 10(6):e1003669, 2014. doi: 10.1371/journal.pcbi.1003669.
- [15] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963. doi: 10.1080/01621459.1963.10500830.
- [16] 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.
- [17] Arne Jensen. Markoff chains as an aid in the study of Markoff processes. *Skandinavisk Aktuarietidskrift*, 36:87–91, 1953. doi: 10.1080/03461238.1953.10419459.
- [18] Stuart A. Kauffman. Autocatalytic sets of proteins. *Journal of Theoretical Biology*, 119(1):1–24, 1986. doi: 10.1016/S0022-5193(86)80047-9.
- [19] Thomas G. Kurtz. The relationship between stochastic and deterministic models for chemical reactions. *The Journal of Chemical Physics*, 57(7):2976–2978, 1972. doi: 10.1063/1.1678692.
- [20] V. Lakshmikantham, D. D. Bainov, and P. S. Simeonov. *Theory of Impulsive Differential Equations*, volume 6 of *Series in Modern Applied Mathematics*. World Scientific, Singapore, 1989. doi: 10.1142/0906.
- [21] Tracey A. Lincoln and Gerald F. Joyce. Self-sustained replication of an RNA enzyme. *Science*, 323(5918):1229–1232, 2009. doi: 10.1126/science.1167856.
- [22] Sean P. Meyn and Richard L. Tweedie. Stability of Markovian processes III: Foster–Lyapunov criteria for continuous-time processes. *Advances in Applied Probability*, 25(3):518–548, 1993. doi: 10.2307/1427522.
- [23] D. R. Mills, R. L. Peterson, and S. Spiegelman. An extracellular Darwinian experiment with a self-duplicating nucleic acid molecule. *Proceedings of the National Academy of Sciences*, 58(1):217–224, 1967. doi: 10.1073/pnas.58.1.217.
- [24] J. R. Norris. *Markov Chains*, volume 2 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge, 1997. doi: 10.1017/CBO9780511810633.
- [25] Matteo Polettini and Massimiliano Esposito. Irreversible thermodynamics of open chemical networks. I. Emergent cycles and broken conservation laws. *The Journal of Chemical Physics*, 141(2):024117, 2014. doi: 10.1063/1.4886396.
- [26] Riccardo Rao and Massimiliano Esposito. Nonequilibrium thermodynamics of chemical reaction networks: wisdom from stochastic thermodynamics. *Physical Review X*, 6(4):041064, 2016. doi: 10.1103/PhysRevX.6.041064.
- [27] Riccardo Rao and Massimiliano Esposito. Conservation laws and work fluctuation relations in chemical reaction networks. *The Journal of Chemical Physics*, 149(24):245101, 2018. doi: 10.1063/1.5042253.
- [28] Hal L. Smith and Paul Waltman. *The Theory of the Chemostat: Dynamics of Microbial Competition*, volume 13 of *Cambridge Studies in Mathematical Biology*. Cambridge University Press, Cambridge, 1995. doi: 10.1017/CBO9780511530043.
- [29] Mike Steel, Wim Hordijk, and Joana C. Xavier. Autocatalytic networks in biology: structural theory and algorithms. *Journal of the Royal Society Interface*, 16(151):20180808, 2019. doi: 10.1098/rsif.2018.0808.
- [30] The mathlib Community. The Lean mathematical library. In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2020)*, pages 367–381. ACM, 2020. doi: 10.1145/3372885.3373824.