

Reliable autocatalytic operation with a finite fuel–waste bath:

joint mission guarantees, reservoir sizing and a sharpened certificate

19 September 2026

Abstract

A reactor that is repeatedly thinned, refed and harvested must recover its catalytic stock, deliver product, and leave a state from which the next operation can succeed. If its fuel is a genuinely finite molecular inventory rather than a chemostatted activity, then every driven event also changes the chemistry of the next operation, and the guarantee has to be proved for that changed stochastic law. We prove a joint finite-horizon operating theorem for an explicitly specified six-species count reactor coupled to a conserved two-species fuel/waste bath of total inventory M and reference capacity R , in which the driven channel $X + F \rightleftharpoons U + W + P$ moves one molecule between the two bath species at every event and its two directed propensities are proportional to the current bath fractions f/R and p/R . One event on the full normalized history law requires, in every one of m externally clocked cycles, return of the actual endpoint to an explicit restart set, two overlapping collected-output thresholds, two food allowances and a gross driven-service allowance, together with an all-prefix bound on the bath composition and a pathwise net-synthesis statement for every physical realization. For every admitted bath state, every parameter in the class $19 \leq r \leq 21$, $d \geq 0$, $dM \leq R/25$ and every controller acting on the returned history, the probability that all m cycles succeed is at least $(1 - e(V))^m \geq 1 - m e(V)$.

We then sharpen the certificate in two independent ways and propagate both through the entire chain. First, the published phase-transport exponent discards a factor: the unrounded rate is $\kappa = 1839/(8.75 \times 10^{12}) = 2.1017 \dots \times 10^{-10}$ rather than 10^{-10} , and a new absorption estimate makes $e(V) \leq 101e^{-\kappa V}$ valid from $V \geq 10^{10}$ instead of 2×10^{11} . Second, on the pure-fuel bath at $d = 1/50$ the two directed service intensities are controlled by *correlated* bath fractions summing to one, which halves the certified per-cycle gross allowance from $\lfloor V/5 \rfloor$ to $\lfloor V/10 \rfloor$ inside the same error budget, and therefore halves the sufficient reservoir at fixed copy scale. Together they give the design rule $V = \max(5 \times 10^{10}, \lceil \kappa^{-1} \log(101m/\delta) \rceil, \lceil 56D_I/m \rceil, \lceil 1080D_X/m \rceil)$, $R = \max(1, \lceil m \lfloor V/10 \rfloor / \rho \rceil)$. The hundred-cycle instance certified at $V = 2 \times 10^{11}$, $R = 4 \times 10^{13}$ with confidence 99.9979% is matched at $V = 9.6 \times 10^{10}$, $R = 9.6 \times 10^{12}$; at the original copy scale the same mission now has failure probability at most 10^{-13} with half the fuel. We also give exact finite-bath endpoint free energies, a finite metered-food corollary and deterministic diagnostics separating productive output from fuel turnover. The whole chain, including nonexplosion, the identification of the stopped device with the physical process, the repetition argument and both refinements, 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 externally implemented feeds and interventions; reliable production does not imply positive net fuel consumption.

1 Introduction

Structural theories of autocatalysis [7, 14, 17, 28] identify which reaction networks can regenerate their catalysts from a food set. The operational test applied in the laboratory is the serial-

transfer protocol [4, 5, 18, 20]: a fraction of the mixture is diluted into fresh food and the system is called self-sustaining if it regrows far enough for the next transfer to succeed, repeatedly. Companion work made that test quantitative for one explicitly specified, reversible and thermochemically consistent chemistry: first for the mass-action concentration equations [2], then at finite molecule number [1], where each cycle is a random thinning of individual molecules followed by four units of a twenty-channel Markov jump process, and m consecutive successes are certified jointly.

All of that assumed maintained fuel and waste: the driven channel ran at activities pinned to one, as in the chemostatted idealization of open-network thermodynamics [21, 22]. A real driven reaction consumes a finite inventory. The present paper removes that idealization for the same chemistry. The fuel F and waste P become two more counted species with conserved total $f + p = M$, held in a bath of reference capacity R that is neither reset nor replenished between cycles, and the two directed propensities of the driven pair become proportional to the current fractions f/R and p/R .

Why this is not a corollary of the maintained-bath theorem. Changing the bath changes the generator, so it changes the law of every observable the earlier proof controls. The reactor’s state space acquires $M + 1$ additional coordinates; the driven forward rate decays as fuel is spent; the reverse rate grows as waste accumulates; and the state that starts each cycle is the actual random endpoint of the previous one, bath included. A trajectory of the finite-bath process is not a trajectory of any fixed-activity process, and no coupling to one is used below. What we do reuse is the *architecture* of the estimates: every local inequality is re-established for the augmented generator, uniformly over a rectangle of admissible instantaneous bath coefficients, and then composed. That uniformity is what makes the finite bath tractable: the source-level estimates never need to know how fast the composition drifts, only that it stays inside the admissible class, which conservation guarantees.

Main result in words. Fix the copy scale V , an admitted bath and an arbitrary controller that sees the complete returned history. Each cycle applies a molecular pulse (every molecule independently retained, withdrawn or additionally lost, then integer food doses and nothing else), runs the literal count process for three time units, and collects effluent for one more. A cycle succeeds if its actual endpoint returns to the restart set $\mathcal{R}_V$, if collected template equivalents and collected 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 count stay inside their allowances. Theorem 3.1 bounds the probability that m consecutive cycles all succeed, under the full normalized history law and with no independence assumption between cycles; Theorem 3.3 adds that on the same event every *chemical prefix* of the mission — not merely the cycle endpoints — has bath composition inside a corridor determined by R ; and Corollary 3.4 inverts this into explicit integer sizes for a declared mission.

The two refinements. The certificate inherited from the maintained-bath analysis is conservative in two places that we can now quantify and repair.

(i) *A discarded factor in the phase exponent.* The dominant error term comes from a backward exponential tilt transported over about $91V$ clock steps, with a linear gain $s(1/700 - 1/1000)$ and a quadratic remainder $(12/5) \cdot 91 s^2$ per unit of V . At the tilt $s = 10^{-6}$ actually used, the exponent coefficient is exactly

$$\kappa = s \left(\frac{1}{700} - \frac{1}{1000} \right) - \frac{12}{5} \cdot 91 s^2 = \frac{1839}{8\,750\,000\,000\,000} = 2.1017142857 \dots \times 10^{-10},$$

more than twice the printed 10^{-10} . Section 5.1 carries κ through the occupation estimate, the joint counter budget and the pulse integration, and proves a new absorption lemma showing

that every remaining term of the budget is dominated by $e^{-\kappa V}$ already from $V \geq 10^{10}$ — a twentyfold reduction of the domain floor. Retuning s inside the same inequality is nearly exhausted (Remark 5.2): the optimal tilt improves κ by under 0.04%, so the remaining conservatism lies in the inequality itself, not its parameter.

(ii) *Correlated bath fractions in the service budget.* The gross driven-service counter was bounded using the worst case of each directed coefficient separately, giving an intensity envelope $\frac{9}{200}V$ on the material corridor. But for a pure bath of total inventory $M = R$ the two fractions satisfy $f/R + p/R = 1$ identically, and the reverse coefficient carries the factor $\eta = 1/(8 \times 10^9)$. Exploiting the correlation gives $\frac{11}{500}V$ instead, and a stopped Chernoff estimate at tilt $1/10$ then certifies the *halved* allowance $\lfloor V/10 \rfloor$ with exponent $-21V/31250 \leq -V/2000$, i.e. inside the error already allocated to this counter. Section 5.2 proves this and — the step that makes it usable — rebuilds the strict stopped-to-physical transfer, the history composition and the prefix sizing corollary around the smaller success event, so that the halved allowance appears in the exported mission conclusion rather than only in an intermediate lemma.

What the refinements buy. Both act on different resources and compose. At δ , m and ρ fixed, (i) reduces the copy scale logarithmically and (ii) halves the reservoir per unit of copy scale. The hundred-cycle instance previously certified at $V = 2 \times 10^{11}$, $R = 4 \times 10^{13}$ with confidence at least 99.9979% is matched at $V = 9.6 \times 10^{10}$, $R = 9.6 \times 10^{12}$: a factor $25/12$ in molecules and $25/6$ in fuel. At the original copy scale, the same hundred cycles now have failure probability at most 10^{-13} using half the fuel. We state the fair comparison explicitly in Section 7: reducing V also reduces the guaranteed output, so a user who fixes the output demand keeps the larger V and banks the improvement as confidence and reservoir instead.

Relation to finite-reservoir stochastic thermodynamics. Finite reservoirs are an established subject and we make no priority claim. Fritz et al. [11] study a Brusselator coupled to finite reservoirs driven by an external pump and quantify the coherence of its oscillations; Remlein et al. [24] derive chemostatted open-network behaviour from hybrid dynamics under rate-scaling and stoichiometric conditions, and Remlein et al. [25] obtain open-network thermodynamics from closed systems in a sequential time-scale-then-abundance limit with an intermediate time window. Those results concern the emergence and thermodynamic consistency of idealized open descriptions. The present theorem answers a different, quantitative question about one fixed source: given a declared mission, how much catalytic preparation, food and finite bath capacity suffice, at finite molecule number and finite horizon, and which guarantee fails to imply which. We neither derive chemostats as an asymptotic limit nor replace those theories. Local detailed balance in the factorial potential form we use is standard [22, 26, 27]; exponential supermartingale and time-uniform concentration methods are standard [8, 15]; Foster–Lyapunov generator inequalities are standard [6, 13, 19]. The model-specific drift, phase-transport and marked-counter estimates, their composition on a single law, and the finite-bath prefix accounting are the technical content.

Three kinds of evidence. Three kinds of statement appear and are kept apart. Every probability, inventory, sizing and detailed-balance statement in Sections 3–5 is a Lean 4 declaration [9, 29] described in Appendix C; the prose proofs are expositions of the formal chain, written so every constant can be checked by hand. The endpoint free-energy simplifications of Section 6.2, the metered-food path coupling of Section 6.3, the tilt-optimality remark and the unit conversions are conventional arguments and are flagged as such at each occurrence. The deterministic trajectories of Section 7.4 integrate the mass-action ODE associated with the same reaction table; they are diagnostics, not evidence about the count process, and no simulation at the certified copy scale is used anywhere.

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	$X + F \rightleftharpoons U + W + P$	$d(f/R) N_X$	$d \eta (p/R) N_U N_W / V$

Table 1: Count propensities per unit of normalized time. Forward means left to right. Only pair 5 changes the bath, by $(f, p) \mapsto (f - 1, p + 1)$ forward and $(f, p) \mapsto (f + 1, p - 1)$ backward. Two food feeds of propensity V and six unit-rate washouts complete the twenty labels.

What is not claimed. The chemistry is schematic: three formal elemental components, one effective trimolecular reverse channel whose elementary implementation is not supplied, ideal mass action, instantaneous interventions, and externally implemented feeds, withdrawals and collection. No molecule is identified and no rate constant is fitted. The theorem starts from established catalytic stock; it is not a food-only startup theorem. Finite missions are certified, not indefinite survival. Most importantly, reliable operation does *not* entail positive net fuel consumption: forward drive destroys a template molecule in this source while reverse drive creates one, the admitted parameter class contains $d = 0$, and Section 8 records the precise additional premises a positive net-service theorem would need.

Organization. Section 2 specifies the source, the bath, the intervention and the cycle. Section 3 states the results. Section 4 proves the one-cycle bound for the changed generator and removes the stopping. Section 5 proves the two refinements and propagates them. Section 6 gives the inventory, design and thermochemical consequences. Section 7 works the missions, the dimensional reading and the diagnostics. Section 8 states scope and open questions. Appendix A records both error budgets, Appendix B the quantitative exponential tests, and Appendix C the claim-to-evidence map and reproduction instructions.

2 The source, the bath and the protocol

2.1 Twenty labels and an evolving bath

The internal species are two foods U, W , the covalent template X , two substrate-binding complexes C_1, C_2 and a duplex Z ; the internal count state is $N = (N_U, N_W, N_X, N_{C_1}, N_{C_2}, N_Z) \in \mathbb{N}^6$ with copy scale $V \in \mathbb{N}$. The bath holds a fuel count f and a waste count p with conserved total $f + p = M$, so the bath state is determined by $f \in \{0, \dots, M\}$; $R > 0$ is its reference capacity. Fix

$$\varepsilon = \frac{1}{500\,000\,000}, \quad \eta = \frac{1}{8\,000\,000\,000}. \quad (1)$$

Table 1 lists the twelve directed chemical channels. In addition the two food feeds $\emptyset \rightarrow U$ and $\emptyset \rightarrow W$ each have propensity V , and each internal species washes out with propensity equal to its count: twenty labelled channels in all. 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. Only pair 5 touches the bath: its forward event changes (f, p) by $(-1, +1)$ and its reverse by $(+1, -1)$. Hence zero fuel disables forward consumption and zero waste disables reverse consumption, with no singular expression at either boundary.

For a full state $s = (N, f, p)$ with labels j , increments ν_j and propensities $a_j(s)$ the generator

is

$$\mathcal{L}h(s) = \sum_{j=1}^{20} a_j(s) \{h(s + \nu_j) - h(s)\}. \quad (2)$$

The admitted parameter class is

$$R > 0, \quad 19 \leq r \leq 21, \quad d \geq 0, \quad dM \leq R/25. \quad (3)$$

Writing $\alpha = df/R$ and $\beta = d\eta p/R$ for the instantaneous driven coefficients, (3) and $f + p = M$ give, *at every reachable bath composition*,

$$0 \leq \alpha \leq \frac{1}{25}, \quad 0 \leq \beta \leq \frac{1}{200\,000\,000\,000}. \quad (4)$$

This rectangle — not any assumption on the speed of bath drift — is the hypothesis used by every local estimate below. Two preparations are named throughout: the *pure bath* $M = R$ with initial state $(R, 0)$, admitting $0 \leq d \leq 1/25$, and the *loaded bath* $M = 2R$ with initial state (R, R) and $d = 1/50$. The bath is never reset or replenished.

2.2 Observables, restart set and intervention

Four linear internal observables organize the analysis:

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

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

Definition 2.1 (Restart set). $\mathcal{R}_V = \left\{ N \in \mathbb{N}^6 : \frac{159V}{160} \leq A_N, B_N \leq \frac{161V}{160}, \quad 40N_X + 45N_{C_1} + 56N_{C_2} + 72N_Z \geq 2V \right\}$.

The integer stock condition is exactly $Y_N \geq V/20$. The set is nonempty at every scale: $(0, 0, V, 0, 0, 0) \in \mathcal{R}_V$, the witness used by the certified instances. Membership is an *established-stock* condition; producing such a seed is not part of the theorem.

Definition 2.2 (Admissible intervention). An intervention is (q, ℓ, e_U, e_W) with $\frac{1}{4} \leq q \leq \frac{3}{4}$, $\frac{49}{50} \leq \ell_i \leq 1$ for each internal species i , and $|e_U|, |e_W| \leq \frac{1}{200}$. Conditional on it, every molecule independently becomes

$$(\text{retained, withdrawn, extra loss}) \quad \text{with probabilities} \quad (q\ell_i, 1 - q, q(1 - \ell_i)), \quad (6)$$

after which $\lfloor V(1 - q + e_U) \rfloor$ molecules of U and $\lfloor V(1 - q + e_W) \rfloor$ of W are added and nothing else. The pulse has zero duration, adds no catalyst, and leaves the bath untouched.

The pulse law is the exact finite product measure over molecules. Every outcome, including loss of all catalyst, keeps its probability mass; nothing is clipped, rejected or conditioned.

2.3 Cycle, counters, success and the history process

Definition 2.3 (Cycle and counters). From the post-pulse full state, run (2) for four time units. Five integer counters are attached to labelled events of the *same* reaction sequence: Q_I , template equivalents collected (a washout of X, C_1, C_2, Z during $[3, 4]$ adds 1, 1, 1, 2); Q_X , free template collected (a washout of X during $[3, 4]$ adds 1, so Q_X is part of Q_I); C_U, C_W , food expenditure (the pulse doses plus every feed arrival during $[0, 4]$); and G , gross driven service (every forward and every reverse pair-5 event during $[0, 4]$). The actual endpoint — internal counts *and* bath — together with the five counters is the cycle record.




Figure 1: The specified process, not a simulated apparatus. Every driven event moves one molecule between the two bath species; the bath is neither reset nor replenished between cycles, so its composition — and hence the two directed propensities of pair 5 — is part of the state carried forward.

Definition 2.4 (Cycle success). A record is successful if

$$N \in \mathcal{R}_V, \quad Q_I \geq \lceil V/56 \rceil, \quad Q_X \geq \lceil V/1080 \rceil, \quad C_U, C_W \leq 5V, \quad G \leq \lfloor V/5 \rfloor, \quad (7)$$

and *sharply* successful if in addition $G \leq \lfloor V/10 \rfloor$.

The two output thresholds overlap and must not be added as though disjoint. The three cost coordinates are allowances *on the successful event*, not caps imposed on the dynamics and not bounds on unsuccessful runs.

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

Every pulse and source outcome is retained and each cycle starts from the actual returned internal *and bath* state; no measurability or continuity is required of π . Section 4.6 also treats an abstract measurable history model, for which the only 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 state equal the newest observed return. No within-cycle control of rates is admitted; the four-unit clocking is external.

3 Main results

For cycle k let S_k be the event that its record is successful, $\hat{S}_k$ that it is sharply successful, and $E_m = \bigcap_{k \leq m} S_k$, $\hat{E}_m = \bigcap_{k \leq m} \hat{S}_k$. Under $\mathbb{P}$ a finite sequence of labelled reaction events and pulse outcomes realizing the recorded history exists almost surely; we call it a physical trace. Let P_{net} denote the signed net internal creation of template equivalents along a trace (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). The two error functions $e(V)$ and $\hat{e}(V)$ are the explicit finite sums of Appendix A.

Theorem 3.1 (Finite-bath repeated operation). *Fix a bath inventory M , parameters satisfying (3), an integer $V \geq V_0 := 2 \times 10^{11}$, an initial full state with $N_0 \in \mathcal{R}_V$, and any controller acting on the returned history. The chronological count process is nonexplosive, and for every $m \in \mathbb{N}$*

$$\mathbb{P}(E_m) \geq (1 - e(V))^m \geq 1 - m e(V), \quad e(V) \leq 101 e^{-V/10^{10}}. \quad (8)$$

On E_m all inequalities of (7) hold in every cycle, so 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 $\sum_k G_k \leq m[V/5]$. Displayed linear lower bounds are understood truncated at zero.

Theorem 3.2 (Sharpened certificate). *Let $\kappa = \frac{1839}{8\,750\,000\,000\,000} = 2.1017142857 \dots \times 10^{-10}$.*

(a) (Every admitted bath.) *There is an explicit error sum $\hat{e}_*(V) \leq \hat{e}(V)$ for which the conclusion of Theorem 3.1 holds verbatim with e replaced by $\hat{e}_*$, and*

$$\hat{e}_*(V) \leq \hat{e}(V) \leq 101 e^{-\kappa V} \quad \text{for all } V \geq 10^{10}. \quad (9)$$

In particular the single-exponential envelope holds from $V \geq 10^{10}$, twenty times below the floor V_0 needed by (8), and $\hat{e}(V) \leq 1/100$ for $V \geq 5 \times 10^{10}$.

(b) (Pure bath, $d = 1/50$.) *If $M = R$ and $d = 1/50$, then for every $V \geq 10^6$, every admitted incoming full state with $N_0 \in \mathcal{R}_V$ and every intervention, one cycle is sharply successful with probability at least $1 - \hat{e}(V)$; consequently, for $V \geq 5 \times 10^{10}$ and every controller,*

$$\mathbb{P}(\hat{E}_m) \geq 1 - m \hat{e}(V), \quad (10)$$

and on $\hat{E}_m$ the mission gross service obeys $\sum_k G_k \leq m[V/10]$ while all conclusions of Theorem 3.1 continue to hold.

Theorem 3.3 (All-prefix inventory and bath tolerance). *Let F_t, P_t count the forward and reverse pair-5 events from the start of the mission to a chemical prefix t , and $J_t = F_t - P_t$. Then at every supported prefix*

$$f_t = f_0 - J_t, \quad p_t = p_0 + J_t, \quad |J_t| \leq F_t + P_t. \quad (11)$$

On E_m every prefix, inside cycles as well as at returns, satisfies $|f_t - f_0| = |p_t - p_0| \leq B_m := m[V/5]$; on $\hat{E}_m$ the same holds with $\hat{B}_m := m[V/10]$. Consequently, for a pure bath with $R \geq B_m/\rho$ (respectively $R \geq \hat{B}_m/\rho$),

$$1 - \rho \leq \frac{f_t}{R} \leq 1, \quad 0 \leq \frac{p_t}{R} \leq \rho \quad \text{at every prefix,}$$

and for the loaded bath both activities lie in $[1 - \rho, 1 + \rho]$, whence for $0 < \rho < 1$

$$\left| \log \frac{f_t/R}{p_t/R} \right| \leq \log \frac{1 + \rho}{1 - \rho}.$$

Corollary 3.4 (Explicit inverse design). *For $m \geq 1$, $0 < \delta < 1$, $0 < \rho < 1$ and nonnegative cumulative output demands D_I, D_X , put*

$$V = \max \left\{ 5 \times 10^{10}, \left\lceil \frac{\log(101m/\delta)}{\kappa} \right\rceil, \left\lceil \frac{56D_I}{m} \right\rceil, \left\lceil \frac{1080D_X}{m} \right\rceil \right\}, \quad R = \max \left\{ 1, \left\lceil \frac{m[V/10]}{\rho} \right\rceil \right\}. \quad (12)$$

Take the pure bath with $r = 20$, $d = 1/50$ and the all-free- X restart witness. Then with probability at least $1 - \delta$, all m cycles are sharply successful, the cumulative collections are at least D_I and D_X , and every chemical prefix has fuel activity in $[1 - \rho, 1]$ and waste activity at most ρ .

The rule is sufficient and deliberately conservative; it is not a global optimum over reactors or protocols, and the maximum is a design rule, not a necessary condition.

Proposition 3.5 (Same-event net synthesis). *On E_m a physical trace exists, and every physical trace of the recorded history satisfies*

$$P_{\text{net}} \geq m \left\lceil \frac{V}{56} \right\rceil + \left\lceil \frac{V}{28} \right\rceil - I_{N_0} \geq m \left\lceil \frac{V}{56} \right\rceil + \left\lceil \frac{V}{28} \right\rceil - \left\lfloor \frac{161V}{160} \right\rfloor \geq \left(\frac{m+2}{56} - \frac{161}{160} \right) V. \quad (13)$$

In particular $P_{\text{net}} > 0$ for every $m \geq 55$ and every admitted seed, at no probability cost.

Proposition 6.2 (exact endpoint bath free energies) and Proposition 6.4 (finite metered food) are stated in Section 6.

Remark 3.6 (Quantifiers). $\mathbb{P}$ is the full normalized law; E_m and $\hat{E}_m$ are events in it whose probability is bounded below, and nothing is conditioned on success. The parameter class (3) is fixed once and is arbitrary within its bounds; in particular $d = 0$ is admitted, at which both driven counters vanish identically. The intervention box is available afresh at every cycle and the controller is an arbitrary function of the complete returned history, which now includes the returned bath. The one-cycle bound holds from $V \geq 10^6$; the larger floors in (8) and (10) are where the single-exponential envelope and the quantitative corollaries become available.

Remark 3.7 (Two allowances, one law). Theorem 3.2(b) does not weaken Theorem 3.1: $\hat{E}_m \subseteq E_m$, and both are events of the same process under the same law. The halved allowance is a *stronger* conclusion bought by a sharper intensity estimate, not a different model or a redefinition of success.

4 The one-cycle bound for the changed generator

This section proves the one-cycle estimate underlying Theorem 3.1. Every statement concerns the full state (N, f, p) and the coefficient rectangle (4); no trajectory of a fixed-activity process is used, and the auxiliary stopped kernels are removed in Section 4.5 before any conditional repetition.

4.1 Material control and the retained stock

Internal chemistry preserves both material coordinates, including the driven pair: $X + F \rightarrow U + W + P$ moves one A -unit and one B -unit from X into U and W . The only positive material jumps are food arrivals, so direct substitution into (2) gives the exact identities

$$\mathcal{L}A_N = V - A_N, \quad \mathcal{L}B_N = V - B_N. \quad (14)$$

These are unchanged by the bath, which is the first reason the finite-bath problem is tractable.

For the pulse, each material coordinate is a sum of independent retained molecule weights plus an integer food dose. The one-molecule estimate

$$1 - p + p e^{-sw} \leq \exp \left\{ - \left(s - \frac{27}{25} s^2 \right) pw \right\} \quad (0 \leq p \leq 1, \ 0 \leq w \leq \frac{9}{5}, \ 0 \leq s \leq \frac{1}{10})$$

applied to the exact product moment generating function of the pulse yields the tails: the admissible bounds put the post-pulse mean of each material coordinate strictly inside $[24V/25, 51V/50]$, with weighted exponential bounds giving failure probability at most $4e^{-V/100000}$, while the retained weighted stock has mean at least $\frac{1}{4} \cdot \frac{49}{50} \cdot \frac{V}{20} = \frac{49V}{4000}$ and lower-tail bound $e^{-1177V/10^9}$ at level $3V/250$ (Appendix B). The exponent uses the actual retained mean rather than a worst-case sum of squared weights.

The analysis is then stopped when either material coordinate leaves $[9V/10, 11V/10]$. On this corridor every species count is at most $11V/10$ and the total propensity — *including* both driven directions, by (4) — is at most $3000V$. This permits uniformization [12, 16]: a Poisson clock of rate $3000V$ selects label j with probability $a_j(s)/(3000V)$ and otherwise makes a null step, with marked counters attached to the selected labels. Exponential tests for the deviations in (14) bound exit during the required time by

$$M(V) = 4e^{-V/2000} + 48000V e^{-V/2000+1/50}, \quad (15)$$

and the same mean-reverting drift with a time-dependent tilt returns A, B to their narrow restart bands at time 4 apart from $4(e^{-V/32000} + e^{-V})$ and the already accounted corridor exits.

4.2 Recovery and residence of the weighted catalyst

The weighted stock Y is chosen so that successive catalytic phases contribute positive drift below a small threshold. Direct source algebra gives the exact identity

$$\mathcal{L}Y_N = VY(f_{r,\alpha,\beta}(N/V)) + \frac{rN_X}{5V}, \quad (16)$$

where $f_{r,\alpha,\beta}$ is the mass-action concentration field of Table 1 with the *instantaneous* coefficients $\alpha = df/R$, $\beta = d\eta p/R$ evaluated at the current bath, and the last term is the exact favourable falling-factorial correction. No derivative of a frozen-bath approximation is taken: the generator acts on an internal observable at the full current state, and the bath enters only through (α, β) , which (4) confines to a rectangle. On the material corridor and below the threshold $Y_N \leq 3V/50$ the parameter bounds give

$$\mathcal{L}Y_N \geq \frac{2}{3}Y_N. \quad (17)$$

This is a generator inequality, not a claim that each random path grows monotonically. Weighted jumps have magnitude at most $9/5$ and their rate-weighted squares are bounded by a linear function of Y_N plus the basal source, so exponentially tilted tests convert the positive drift into a lower-tail estimate for entry into $Y \geq 3V/50$: a backward tilt schedule with exponential-drift coefficient $5/8$, transported from $1/1000$ at the deadline and capped at $101/20000$, gives failure to enter by time $11/4$, before material exit, at most $e^{-3V/5 \cdot 10^6} + e^{-V/2}$. A second tilted test controls descent from $3V/50$ to $V/20$ through the remaining $5/4$ time units, contributing $e^{-V/10000} + 12000Ve^{-V/10000+9/500}$. Together with (15) these form the term $R_{\text{err}}(V)$ of Appendix A. The times $11/4$ and $1/4$ partition the analysis only; the physical uncollected run lasts three units.

4.3 Positive phase transport to free product

Large weighted stock does not mean large free X at the same instant, and this gap supplies the dominant error term. On the material corridor the coordinate drifts satisfy

$$\begin{aligned} \mathcal{L}N_X &\geq -70N_X + 20N_{C_1} + 38N_Z, & \mathcal{L}N_{C_2} &\geq -70N_{C_2}, \\ \mathcal{L}N_{C_1} &\geq -70N_{C_1} + 20N_{C_2}, & \mathcal{L}N_Z &\geq -70N_Z + 20N_{C_2}, \end{aligned} \quad (18)$$

with all off-diagonal coefficients nonnegative. Put $q_c = 3000V$, $a = 1 - 70/q_c$, and initialize a backward weight vector at $(1, 0, 0, 0)$ on (X, C_1, C_2, Z) . One clock step maps w to

$$(aw_X, aw_{C_1} + 20w_X/q_c, aw_{C_2} + 20(w_{C_1} + w_Z)/q_c, aw_Z + 38w_X/q_c),$$

so the weights remain nonnegative with sum at most one, and after n steps equal

$$a^n \left(1, \frac{20n}{q_c a}, \frac{580n(n-1)}{q_c^2 a^2}, \frac{38n}{q_c a} \right). \quad (19)$$

For $89V \leq n \leq 91V$ elementary bounds on a^n give $w(n) \cdot N \geq Y_N/35$: every catalytic phase funds a future free- X observation, so the estimate does not assume that stock is initially free. Figure 3 plots these weights.

To control fluctuations, apply (2) to $e^{-s w \cdot N}$. The jump bound and the total rate give the explicit quadratic remainder

$$\mathcal{L}e^{-s w \cdot N} \leq e^{-s w \cdot N} \{-s D_w(N) + 7200V s^2\}, \quad 0 \leq s \leq \frac{1}{100}, \quad (20)$$

D_w being the weighted right side of (18). Iterating with the backward weights at $s = 10^{-6}$, the transported lower bound $w(n) \cdot N \geq V/700$ (from $Y \geq V/20$) dominates the low- X threshold $N_X \leq V/1000$ and the accumulated quadratic remainder $(12/5) \cdot 91V s^2$. The resulting probability of low free X while the path remains active is $e^{-\kappa V}$ with κ as in Theorem 3.2; Section 5.1 is devoted to keeping this constant rather than rounding it to 10^{-10} .

At every later clock time the same bound applies, and summing expectations with Markov's inequality shows that low- X states occupy at most one percent of collection clock steps except with probability $100e^{-\kappa V}$; this step requires no independence between those states. A marked exponential test for actual X washouts then produces the free-output failure bound on the active path.

4.4 Other marked outputs and expenses

Template washout has intensity I_N , and the coordinatewise comparison $Y_N \leq \frac{7}{5}I_N$ obtained from the four coefficients in (5) gives $I_N \geq V/28$ throughout residence above $Y = V/20$. The marked count test yields template-output failure at most $e^{-V/100000}$, including the one-count margin needed to certify the integer ceiling.

Food additions cost at most $151V/200$ per species from the pulse; each feed has constant rate V for four units, and exponential upper-tail tests give $e^{-V/300}$ per food ceiling $5V$. For the gross driven counter, the material corridor gives

$$\alpha N_X + \beta N_U N_W / V \leq \left(\frac{1}{25} \cdot \frac{11}{10} + \frac{1}{200\,000\,000\,000} \cdot \frac{121}{100} \right) V < \frac{9}{200} V, \quad (21)$$

and the marked upper-tail test with a one-count floor gives $e^{-V/2000}$ for the ceiling $\lfloor V/5 \rfloor$. Section 5.2 replaces (21) by a strictly better estimate.

All five counters are functions of the same sequence of labelled events, so adding their failure probabilities is a union bound within one joint law, never a product of marginals.

4.5 Removing the stopping and identifying the physical law

The finite stopped process is an analytical device, and transferring its estimates requires care. We use the *strict* event on which the terminal state is still active, in addition to the output and restart conditions. Inactive states are absorbing in the stopped kernel, so terminal activity implies that no earlier exit occurred; on this event the stopped and literal updates — including bath updates and counter increments — agree at every step. A small exact model shows this is necessary: for a closed restart event the stopped probability can be 1 while the physical probability is $1/2$, whereas the strict event gives $1/4$ for both.

For a finite region with bounded rates, both the uniformized marked kernel and the chronological jump construction satisfy the same first-event renewal equation; expanding on the first waiting time and label gives that equation directly, and uniqueness follows by iterating

the residual on a bounded time interval, whose remainder is bounded by the corresponding exponential-series tail. Hence the two killed laws agree, preserving waiting times and labels, not merely a state marginal.

To pass to the countable state space, exhaust it by material-height regions. On a bounded interval only the two constant-rate food feeds can increase total material, so their number is almost surely finite; the bath contributes only $M + 1$ values. Every material-height region therefore contains finitely many full states and has bounded total rate, so explosion is impossible, and increasing killed endpoint laws exhaust the unrestricted chronological law. Their deterministic-time composition identities pass to the limit by monotone convergence, so the analytical split $11/4 + 1/4$ is exactly the physical three-unit recovery law followed by the one-unit marked collection.

Integrating over the normalized molecule-pulse distribution — no bad pulse and no failed source path is removed — gives the conditional one-cycle bound: for $V \geq 10^6$, every restart state, every admitted bath and every intervention,

$$\mathbb{P}(S_1 \mid N_0, f_0, \text{intervention}) \geq 1 - e(V). \quad (22)$$

4.6 Conditional histories and the product bound

Write $K(h, \cdot)$ for the next-history kernel, built from the actual full endpoint and the policy choice at h . A successful return has internal state in $\mathcal{R}_V$ while its retained bath still satisfies $f + p = M$, so (22) applies again at that history. With S the event that the newest cycle succeeds,

$$\mathbb{P}(E_{m+1}) = \int_{E_m} K(h, S) d\mathbb{P}(h) \geq (1 - e(V))\mathbb{P}(E_m), \quad (23)$$

and induction with Bernoulli's inequality proves (8). This iterates a *conditional* bound; it assumes neither independence of cycles nor renormalization after rejecting failures, and the restricted history measure is a submeasure of the full law throughout. Almost-sure existence of physical traces is preserved under the iteration.

Finally, at each forward event f decreases by one and p increases by one, with signs reversed at each reverse event; induction on the actual labelled trace gives the first two identities of (11). Pulses leave the bath unchanged. At every chemical prefix the absolute signed change is bounded by the number of driven events seen so far, which is no larger than the final mission gross count because the marks are nonnegative — this is what upgrades an endpoint statement to an all-prefix statement. Summing the per-cycle gross bounds proves Theorem 3.3; division by R gives the activity corridors, and the loaded-bath force bound follows from log monotonicity.

5 Two refinements of the certificate

Sections 5.1 and 5.2 contain the new mathematics of this paper. They are independent: the first reduces the copy scale needed for a given confidence, the second halves the reservoir needed for a given copy scale, and Section 5.3 composes them.

5.1 The unrounded phase exponent and a new absorption estimate

Lemma 5.1 (Exact phase rate). *At tilt s , transported stock lower bound $V/700$, low- X threshold $V/1000$, at most $91V$ clock steps and quadratic remainder $(12/5)s^2$ per step, the exponent of the phase estimate of Section 4.3 is $-\kappa(s)V$ with*

$$\kappa(s) = s \left(\frac{1}{700} - \frac{1}{1000} \right) - \frac{12}{5} \cdot 91 s^2.$$

At the tilt $s = 10^{-6}$ used in the proof, $\kappa := \kappa(10^{-6}) = \frac{1839}{8\,750\,000\,000\,000}$, exactly $\frac{1839}{875} = 2.1017\dots$ times the printed rate 10^{-10} .

Proof. $\kappa(s) = s \cdot \frac{3}{7000} - \frac{1092}{5}s^2$; at $s = 10^{-6}$ this is $\frac{3}{7 \cdot 10^9} - \frac{1092}{5 \cdot 10^{12}} = \frac{1839}{8.75 \times 10^{12}}$. $\square$

Remark 5.2 (Retuning the tilt is nearly exhausted). With $\Delta = 3/7000$ and $C = (12/5) \cdot 91$, the unconstrained maximum of $\kappa(s) = \Delta s - Cs^2$ is at $s_* = \Delta/(2C) \approx 9.8116 \times 10^{-7}$, where $\kappa(s_*) \approx 2.102489 \times 10^{-10}$. This improves κ by about 0.037%. A numerical search for a better tilt inside this same inequality therefore cannot remove orders of magnitude; the remaining conservatism is in the inequality — specifically the worst-case increment used in (20) in place of a channelwise rate-weighted squared increment — and would require a new estimate. Both statements in this remark are conventional arithmetic, replayed by the accompanying script; the improvement carried into the theorems is the recovery of κ itself, which is compiled. (Right panel of Figure 3.)

Carrying κ instead of 10^{-10} through Section 4.3 is mechanical — the window bound, the burn-in composition, the occupation estimate and the marked washout test are all monotone in this constant — but it changes the *envelope*, because the single-exponential form of the certificate is obtained by absorbing every other term of the budget into the phase term. That absorption has to be redone, and it is where the domain floor is determined.

Lemma 5.3 (Residual absorption). *Write $\hat{e}(V)$ for the refined budget of Appendix A. For every $V \geq 0$,*

$$\hat{e}(V) \leq 100 e^{-\kappa V} + 2 \times 10^6 (V + 1) e^{-V/(4 \times 10^7)}, \quad (24)$$

and for every $V \geq 10^{10}$,

$$2 \times 10^6 (V + 1) e^{-V/(4 \times 10^7)} \leq e^{-\kappa V}. \quad (25)$$

Consequently $\hat{e}(V) \leq 101 e^{-\kappa V}$ for $V \geq 10^{10}$.

Proof. For (24): every term of $\hat{e}$ other than the phase term $100e^{-\kappa V}$ is of the form $e^{-V/D}$ with $D \leq 4 \times 10^7$, e^{-cV} with $c \geq 1/(4 \times 10^7)$, or $cV e^{-V/D+a}$ with $D \leq 4 \times 10^7$ and $a \leq 1/50$; since $e^{1/50} \leq 2$, each is bounded by a constant multiple of $e^{-V/(4 \times 10^7)}$ or of $V e^{-V/(4 \times 10^7)}$, and summing the coefficients (Appendix A) gives (24).

For (25), set $c = \frac{1}{4 \times 10^7} - \kappa = \frac{216911}{8\,750\,000\,000\,000}$, so that $e^{-V/(4 \times 10^7)} = e^{-\kappa V} e^{-cV}$ and it suffices that $2 \times 10^6 (V + 1) \leq e^{cV}$. For $V \geq 10^{10}$ we have $x := cV \geq 247$, hence $247^9 x \leq x^{10}$, and $x^{10}/10! \leq e^x$ gives

$$e^{cV} \geq \frac{247^9 c}{3\,628\,800} V > 2.33 \times 10^7 V \geq 2 \times 10^6 (V + 1),$$

the last step because $V \geq 10^{10}$. The margin is a factor above 11, so the floor 10^{10} is comfortable rather than tight. $\square$

Two consequences deserve emphasis. First, the floor for the single-exponential form drops from 2×10^{11} to 10^{10} ; the larger floor 5×10^{10} quoted in Theorem 3.2 is what makes $\hat{e}(V) \leq 1/100$, i.e. what makes the linear bound informative, not what makes the envelope valid. Second, the useful confidence threshold moves: $101 m e^{-\kappa V} \leq \delta$ now solves to $V \geq \kappa^{-1} \log(101m/\delta)$ with $\kappa^{-1} = 4.758 \times 10^9$ instead of 10^{10} .

5.2 Correlated bath fractions halve the service allowance

The estimate (21) bounds each directed coefficient by its own maximum over the admissible class. For the pure bath that is wasteful, because the two coefficients are not free: $f + p = R$ forces $a := f/R$ and $b := p/R$ to satisfy $a + b = 1$ at every reachable composition.

Lemma 5.4 (Correlated gross envelope). *Let $M = R$, $d = 1/50$. On the material corridor $N_U, N_W, N_X \leq \frac{11}{10}V$, the total gross driven intensity satisfies*

$$\alpha N_X + \beta \frac{N_U N_W}{V} = \frac{1}{50} \left(a N_X + \eta b \frac{N_U N_W}{V} \right) \leq \frac{11}{500} V. \quad (26)$$

Proof. With $u = N_U/V$, $w = N_W/V$, $x = N_X/V$ all at most $11/10$, we have $\eta u w \leq \eta(11/10)^2 \leq 11/10$, so $ax + \eta b u w \leq (a + b) \cdot \frac{11}{10} = \frac{11}{10}$, and the factor $1/50$ gives (26). The point is that a and b are weights of a convex combination, whereas (21) maximizes them separately. $\square$

The gain over (21) is the factor $\frac{11}{500} / \frac{9}{200} = \frac{44}{90}$, slightly better than a half.

Lemma 5.5 (Sharp service tail). *Under the hypotheses of Lemma 5.4, for the three-stage stopped marked source over the four time units and every incoming full state,*

$$\mathbb{P}\left(G \geq \frac{V}{10} - 1\right) \leq \exp\left\{-\frac{1}{10}\left(\frac{V}{10} - 1\right) + 4 \cdot \frac{11}{500} V (e^{1/10} - 1)\right\} \leq e^{-V/2000} \quad (V \geq 10^6).$$

Proof. Uniformization turns the intensity bound (26) into a one-step multiplicative bound for the mark-generating function; composing the three stages and applying the exponential Markov inequality at tilt $s = 1/10$ gives the displayed exponent. Using $e^{1/10} - 1 \leq \frac{53}{500}$,

$$-\frac{V}{100} + \frac{11V}{125} \cdot \frac{53}{500} = -\frac{21V}{31250},$$

and $\frac{21}{31250} > \frac{1}{2000}$ absorbs the one-count rounding buffer for $V \geq 10^6$. $\square$

Remark 5.6 (A first attempt that fails, and why it matters). Reusing the independent cap $\frac{9}{200}V$ with the same tilt gives exponent $-\frac{V}{100} + 4 \cdot \frac{9}{200} V \cdot \frac{53}{500} = +\frac{227V}{25000} > 0$: no tail at all. The correlated envelope is therefore not a cosmetic improvement of constants but the step that makes the halved allowance provable, and the failure of the independent version is a weak estimate, not a counterexample.

From a tail to an exported conclusion. Lemma 5.5 alone does not prove Theorem 3.2(b). The success predicate $\lfloor V/5 \rfloor$ is fixed inside the prepared-cycle failure union, the strict stopped-to-physical transfer, the literal-cycle bound, the returned-history iteration and the prefix sizing corollary; supplying a stronger local inequality to a theorem whose stated conclusion still reads $\lfloor V/5 \rfloor$ proves nothing. Accordingly we rebuild that chain around the smaller event. Define the sharp failure set as the original prepared-cycle failure union together with $\{G \geq V/10 - 1\}$, so that its complement gives cycle success and $G \leq \lfloor V/10 \rfloor$; bound its pulse-integrated probability by $\hat{e}_*(V) + e^{-V/2000}$ using the union bound on the one joint law, transfer it through the strict active event to the literal chronological cycle exactly as in Section 4.5, and iterate it over returned histories exactly as in Section 4.6. Summing the per-cycle sharp gross bounds over the mission and composing with the prefix identity (11) yields $|f_t - f_0| \leq m \lfloor V/10 \rfloor$ at every chemical prefix, which is the statement that halves R in Corollary 3.4. Every step of this rebuild is compiled (Appendix C).

Remark 5.7 (Scope of the halving). Lemma 5.4 uses $M = R$ and $d = 1/50$. It does not apply verbatim to the loaded bath $M = 2R$, where $a + b = 2$, nor to $d = 1/25$, where the prefactor doubles; for those the allowance $\lfloor V/5 \rfloor$ of Theorem 3.1 remains the certified value. The phase refinement of Section 5.1, by contrast, applies to every admitted bath.




Figure 2: Left: the hundred-cycle failure bound $101 m e^{-\kappa V}$ against the rounded $101 m e^{-V/10^{10}}$; the dotted line is the confidence of the original certified instance, which the refined bound reaches at $V = 9.6 \times 10^{10}$. Right: sufficient pure fuel R against mission length at $\delta = 10^{-2}$, $\rho = 10^{-1}$ and no additional output demand, under the two per-cycle allowances; markers are the two certified hundred-cycle instances. Both panels evaluate proved bounds; integer rounding is retained in the certified instances but not in the plotted curves.

5.3 Composing the two refinements

The refinements act on different terms of the budget and compose without interaction: the phase refinement changes the constant in the dominant term, the service refinement adds one more exponential of the same magnitude as an already-present term and tightens one conclusion. Writing $\hat{e}(V) = \hat{e}_*(V) + e^{-V/2000}$ for the combined budget, Lemma 5.3 covers both, and Theorem 3.2 follows by the composition of Sections 4.5 and 4.6. The resulting design rule is Corollary 3.4; Figure 2 compares the two certificates and the two reservoir rules.

6 Inventory, design and thermochemical consequences

6.1 Terminal stock strengthens net synthesis

Proof of Proposition 3.5. Telescoping the template inventory along the literal reactions and pulse outcomes gives, for every physical trace,

$$P_{\text{net}} = I_{\text{final}} - I_{N_0} + Q_{\text{all washout}} + I_{\text{withdrawn}} + I_{\text{extra loss}}.$$

All removed inventories are nonnegative and food boluses contain no templates, so the collected output during the designated windows — a subset of all washout — may be substituted as a lower bound. Comparing the four coefficients in (5) gives $40Y_N \leq 56I_N$, so the restart condition $Y \geq V/20$ forces $I_{\text{final}} \geq \lceil V/28 \rceil$. Initially $I_{N_0} \leq A_{N_0} \leq \lfloor 161V/160 \rfloor$ with integer inventory. These give the first two bounds; relaxing the ceilings and the floor gives the third, whose coefficient at $m = 55$ is $13/1120 > 0$ while at $m = 54$ it is negative. $\square$

The statement is that the entire certified output cannot be accounted for by harvesting the initial seed. It does not identify the ancestry of individual molecules, and it is not a statement about net fuel consumption: the two are different observables, related only through the driven pair.




Figure 3: Left: the backward phase-transport weights (19) in their large- V form, with the observation window shaded and the certified coefficient $1/35$ marked; w_{C_2} is scaled by 10^3 to be visible. Right: the phase exponent $\kappa(s)$ of Lemma 5.1; the tilt actually used already attains 2.1017×10^{-10} , and the optimum of this parabola improves it by under 0.04%.

6.2 Exact finite-bath thermochemistry

Use dimensionless free energies in units of $k_B T$. For a species of count n , standard potential g and reference count $C > 0$ set

$$\phi(n; g, C) = ng + \log(n!) - n \log C. \quad (27)$$

The internal standard potentials are those of the chemical realization [3],

$$(g_U, g_W, g_X, g_{C_1}, g_{C_2}, g_Z) = (0, 0, -\log 10, -\log 10, -\log 10, -2 \log 10),$$

and the bath potentials are $g_F = A_0 := \log(8 \times 10^{10})$ and $g_P = 0$, with equal reference counts R .

Proposition 6.1 (Neighbouring-state local detailed balance). *For a supported forward driven step from (u, w, x, f, p) with $x, f \geq 1$ and $d > 0$, evaluating the reverse propensity at the post-forward state $(u + 1, w + 1, x - 1, f - 1, p + 1)$ gives*

$$\frac{a_+(s)}{a_-(s + \nu_+)} = 8 \times 10^9 \cdot \frac{f}{p + 1} \cdot \frac{R_P}{R_F} \cdot \frac{Vx}{(u + 1)(w + 1)} = \exp\{-\Delta G_{\text{total}}\}, \quad (28)$$

with ΔG_{total} computed from (27) over the five changing species.

The identity $\phi(n + 1) - \phi(n) = g + \log(n + 1) - \log C$ applied to each changing species gives (28) directly. The neighbouring convention matters: at $p = 0$ the post-state has $p + 1 = 1$, so no chemical potential is evaluated at $\log 0$, whereas a same-state forward/reverse ratio is undefined there. At $d = 0$ the pair is inactive and no ratio of vanishing propensities is asserted.

Proposition 6.2 (Endpoint bath free energies). *Let $J = F - P$ be the net forward driven service over a mission and W_b the bath free-energy decrease. From a pure bath $(R, 0)$ one has $0 \leq J \leq R$ and*

$$\Delta G_b = -A_0 J - \log \binom{R}{J}, \quad \frac{W_b}{k_B T} = A_0 J + \log \binom{R}{J}. \quad (29)$$

From a loaded bath (R, R) one has $-R \leq J \leq R$ and

$$\Delta G_b = -A_0 J + \log \frac{(R - J)!(R + J)!}{(R!)^2}. \quad (30)$$

Proof (conventional algebra from the compiled potential and inventory identities). The exact inventory identity (11) gives final bath counts $(R - J, J)$ in the pure case. Subtracting $\phi(R; A_0, R) + \phi(0; 0, R)$ from the endpoint potential leaves $-A_0J + \log((R - J)!) + \log(J!) - \log(R!)$, the reference-count terms cancelling because the total bath inventory is conserved; the factorial definition of $\binom{R}{J}$ gives (29). The loaded endpoint is $(R - J, R + J)$ and the same cancellation gives (30). All factorial arguments are nonnegative integers, including at the boundaries, and $0! = 1$. $\square$

No fixed affinity and no Stirling approximation is used. Three distinctions must stay explicit: gross turnover $F + P$ is not net conversion J ; a bath free-energy decrease is not extracted useful work; and detailed balance for the reversible chemistry does not account for pumping, feeding, withdrawal, separation, timing or thermal control. The bath is closed to replenishment, but the reactor-plus-apparatus system is not a closed autonomous chemical system [10, 23].

Corollary 6.3 (Force-budget inversion). *For the loaded bath put $b = \hat{B}_m/R$. Requiring the dimensionless affinity drift to be at most $\zeta > 0$ is certified by $2 \operatorname{artanh}(b) \leq \zeta$, i.e. by*

$$R \geq \hat{B}_m \coth(\zeta/2),$$

with integer rounding applied afterwards. At $\zeta = 0.1$ this is $R/\hat{B}_m \geq 20.0167$; at $\zeta = 0.02$, $R/\hat{B}_m \geq 100.0034$.

This gives a thermodynamic design knob in place of an arbitrary percentage tolerance. It does not apply to “deviation from the initial pure-bath affinity”, since that macroscopic log ratio is singular at zero initial waste. For general initial loadings $f_0, p_0 > 0$ with $|J| \leq B < \min(f_0, p_0)$ the affinity change is exactly $\Delta A(J) = \log(1 - J/f_0) - \log(1 + J/p_0)$, decreasing in J , so its worst absolute value over $[-B, B]$ is $\max\{\log \frac{1+B/p_0}{1-B/f_0}, \log \frac{1+B/f_0}{1-B/p_0}\}$; if a one-sided preparation constrains J , intersect with its support before taking extrema rather than charging impossible directions. (Conventional deterministic consequences of (11).)

6.3 Finite metered food

Proposition 6.4 (Metered no-stockout). *Let P_U, P_W be the food separately charged to the chosen preparation, and allocate initial stocks $H_U = P_U + 5mV + 1$ and $H_W = P_W + 5mV + 1$. Let an external metering mechanism supply each original arrival clock at intensity V while stock remains, pay the literal pulse boluses, and move to an absorbing shutdown if a requested delivery cannot be funded. Then the metered mission succeeds with at least the probability given by Theorem 3.1 (respectively Theorem 3.2), and all finite-bath prefix and activity consequences hold on the same event.*

Proof (compiled inventory and rate agreement; conventional path coupling). That the full propensity vector of the metered process agrees with the original before stockout, and that the cumulative counters are nondecreasing at chemical, cycle and mission prefixes with the boluses included, are compiled. For the transfer, construct both processes with common randomness: while their full states agree and both stocks are positive, use the same exponential waiting draw and the same label draw — their propensity vectors agree there, so time, label, internal update and bath update agree — and the same pulse outcome at each intervention, debiting the specified integer boluses. Nonexplosion (Section 4.5) makes this recursion well defined through the finitely many jumps and the finitely many pulses of a bounded mission. On the original success event, cumulative use of each food is at most $5mV$ at every prefix, so the spare molecule keeps both stocks strictly positive and no shortage occurs; the coupled processes then agree in all states, labels, times and collected outputs, and the original success event is contained in the metered one. $\square$

Quantity	Original certificate	Refined, reduced scale	Refined, original scale
Copy scale V	200,000,000,000	96,000,000,000	200,000,000,000
Per-cycle gross allowance	$\lfloor V/5 \rfloor$	$\lfloor V/10 \rfloor$	$\lfloor V/10 \rfloor$
Pure fuel R at $\rho = 1/10$	40,000,000,000,000	9,600,000,000,000	20,000,000,000,000
Mission confidence at least	0.999979	0.999979	$1 - 10^{-13}$
Template equivalents collected	357,142,857,200	171,428,571,500	357,142,857,200
Free X collected (included above)	18,518,518,600	8,888,888,900	18,518,518,600
Each food, operation only	$\leq 10^{14}$	$\leq 4.8 \times 10^{13}$	$\leq 10^{14}$
Gross driven events	$\leq 4 \times 10^{12}$	$\leq 9.6 \times 10^{11}$	$\leq 2 \times 10^{12}$
Net synthesis, all-free seed	164,285,714,343	78,857,142,929	164,285,714,343

Table 2: One event of the stated probability in each column. The middle column matches the original confidence with 12/25 of the molecules and 6/25 of the fuel; the right column keeps the original copy scale and outputs, halves the fuel, and converts the phase refinement into eight further orders of confidence. Output guarantees scale with V , so the middle column is not uniformly stronger — see Remark 7.1.

A passive food tank whose arrival intensity declines with its remaining concentration is a different model and is not covered. Finite metered food bounds a prescribed finite mission; it does not make the reactor indefinitely autonomous.

7 Worked missions, dimensions and diagnostics

7.1 Two certified hundred-cycle instances

Take $m = 100$, $\rho = 1/10$, $r = 20$, $d = 1/50$, the pure bath and the all-free- X seed $(0, 0, V, 0, 0, 0)$.

Remark 7.1 (The fair comparison). Reducing V reduces the certified output at fixed m . A user who declares the original output demand $D_I = 357,142,857,200$ has $\lceil 56D_I/m \rceil > 2 \times 10^{11}$, so the demand term in (12) pins the copy scale back at the original value and the improvement is banked entirely as confidence and as reservoir — the right-hand column of Table 2. The middle column is the correct comparison only when confidence and mission length, not output, are the binding requirements. Holding all of m , δ , D_I , D_X and ρ fixed, the guaranteed gain of this paper is: the reservoir halves, and the confidence margin improves by the factor $e^{(\kappa-10^{-10})V}$ at whatever V the demands force.

At every chemical prefix of the middle column, $0.9 \leq f/R \leq 1$ and $0 \leq p/R \leq 0.1$. The analogous loaded bath has total inventory $2R$, initial fuel and waste R each, the same $d = 1/50$, and both activities in $[0.9, 1.1]$; by Remark 5.7 it keeps the allowance $\lfloor V/5 \rfloor$, but it does inherit the phase refinement, so the loaded hundred-cycle mission is certified at $V = 9.6 \times 10^{10}$ with the same confidence 0.999979 (Appendix C). None of these bath sizes is asserted to be necessary.

7.2 Which resource controls a mission

For fixed $V \geq 5 \times 10^{10}$, positive integer R , confidence $1 - \delta$, tolerance ρ and finite food stocks H_U, H_W with preparation charges P_U, P_W , the certified horizon is the minimum of

$$\begin{aligned}
m &\leq \left\lfloor \frac{\delta}{101} e^{\kappa V} \right\rfloor, & m &\leq \left\lfloor \frac{H_U - P_U - 1}{5V} \right\rfloor, \\
m &\leq \left\lfloor \frac{\rho R}{\lfloor V/10 \rfloor} \right\rfloor, & m &\leq \left\lfloor \frac{H_W - P_W - 1}{5V} \right\rfloor.
\end{aligned} \tag{31}$$

If $H_i < P_i + 1$ the strict no-stockout certificate admits no positive mission; that is infeasibility, not a negative cycle count. Asymptotically, at fixed output demands, ρ and δ , (12) gives $V = O(1 + \log(m/\delta))$ and $R = O(m(1 + \log(m/\delta))/\rho)$, with food allocations of order mV .

These are sufficient scaling laws; they establish no matching lower bound and do not rule out a better dependence on m obtained by bounding the largest prefix of the *signed* current J directly through its drift and compensated martingale, instead of through $|J| \leq F + P$. For the pure-bath regime, where the reverse coefficient is at most 10^{-11} times the forward one, such cancellation is unlikely to help much; that should be diagnosed before a further theorem is attempted.

7.3 Dimensional reading

Choose reference concentrations c_* for the reactor and $c_{b,*}$ for the bath, and a kinetic time τ . Then $V = N_A c_* \Omega$, $R = N_A c_{b,*} \Omega_b$ and the mission duration is $4m\tau$; R is not a volume until $c_{b,*}$ is specified. At $c_* = c_{b,*} = 1$ mM and $\tau = 60$ s, the middle column of Table 2 has reactor volume 0.159412 nL, bath volume 15.9412 nL — a bath-to-reactor ratio of 100, against 200 for the original instance — and a source-running time of 400 minutes. The corresponding amounts are 0.284664 pmol of collected template equivalents, 0.0147604 pmol of collected free X , at most 79.7059 pmol of each food during operation, at most 1.59412 pmol of gross driven events (an accounting unit, not a species), and 15.9412 pmol of initial pure fuel.

For an order- k internal mass-action term, a dimensionless coefficient a corresponds to $a/(\tau c_*^{k-1})$: at these references, a bimolecular 20 becomes $1000/3 \approx 333.33 \text{ M}^{-1}\text{s}^{-1}$, a unimolecular 20 becomes $1/3 \text{ s}^{-1}$, and washout becomes $1/60 \text{ s}^{-1}$. The driven pair involves contact between two reference volumes and needs its own explicit convention; the same-volume formula must not be applied mechanically to a separated reservoir. Lowering the concentration at fixed counts is not a free improvement — it demands different physical rate constants if the dimensionless source is to be unchanged — and converting counts to nanolitres does not establish a chemical implementation. Mixing, transport across the contact, molecular identities, attainable rate constants, enzyme loadings, osmotic effects and separation efficiency all remain uncalibrated here.

7.4 Deterministic diagnostics

Case	R/V	d	$\min Q_I/V$	$\min Q_X/V$	F/V	P/V
Pure	1	0.02	0.275215	0.128704	0.07239	1.24e-12
Loaded	1	0.02	0.275215	0.128704	0.07239	4.16e-11
Nearly depleted	0.001	0.02	0.277266	0.129788	0.001	3.6e-11
Zero drive	1	0	0.278579	0.130464	0	0

Table 3: The four diagnostic runs of Figure 4. F and P are mission-integrated directed fluxes. Sampled extrema are diagnostics, not rigorous path bounds.

These runs were performed to test one interpretation, namely that sustained output requires abundant fuel, and they reject it for this source: a bath depleted by seven orders of magnitude, and even a bath with the drive switched off entirely, still supports collection at essentially the same level. That is consistent with the structure of the model — forward drive *destroys* a template — and it is the reason the theorem’s gross allowance is an allowance rather than a requirement.

8 Scope, limitations and open questions

The theorem answers whether a specified reactor can repeatedly recover and deliver output while its finite bath stays inside a requested activity corridor, and how large the preparation, food and bath must be for a declared mission. It does not establish strictly positive net




Figure 4: Deterministic mass-action diagnostics for the reaction table, over eight cycles with retained bath, at $r = 20$, $q = 1/4$, $\ell_i = 49/50$, $e_U = e_W = -1/200$. Left: fuel activity. Right: per-cycle integrated collected template equivalents. The nearly depleted bath ends at fuel activity 2.37×10^{-10} and still collects more per cycle than the abundant cases; the zero-drive case produces output with no driven service at all. These are integrated ODE fluxes, not random counts: they are evidence about mechanism, not about probabilities, attractors or the finite-count theorem.

consumption of the named fuel. In this source, forward drive destroys a template and reverse drive creates one; the admitted class includes $d = 0$; and a collected-output lower bound is not a lower bound on the preceding occupation integral. A positive net-service theorem would need a quantitative lower bound on integrated free- X occupancy while fuel is sufficiently abundant, a forward-count lower tail, a reverse-count upper tail, and their joint conditional composition. None of these is assumed here, and the compiled conditional obstruction is simply that a premise $J \geq m j_* V$ together with $f_{\text{end}} \geq 0$ forces $m j_* V \leq f_0$.

Other explicit non-claims: the effective trimolecular reverse channel has no elementary implementation here; the intervention is instantaneous with conditionally independent thinning; the copy scale is restored externally; and the four-unit clocking is external, so this is not an autonomous system. The model's extreme ratios ε and η , the species-dependent loss model and the product-mixture interpretation are at least as consequential for applicability as the copy scale, and a smaller count bound does not address them.

Three directions look worthwhile. (i) The phase estimate still charges every channel a worst-case squared increment; computing the actual rate-weighted quantity $Q_w(N) = \sum_j a_j(N)(w \cdot \nu_j)^2$ channel by channel — food feeds and food washout contribute zero to an internal catalytic observable, $Z \rightleftharpoons 2X$ contributes $(2w_X - w_Z)^2$, the driven pair contributes w_X^2 times its gross intensity — and integrating a certified envelope along the adjoint schedule would sharpen (20) where Remark 5.2 shows retuning cannot. (ii) Targeting the integrated free- X occupation or the collection count directly with a backward equation carrying a reward term may avoid paying for pointwise low- X events and then applying Markov's inequality; an earlier constant-corrector attempt failed on endpoint penalties, so time-dependent weights and all boundary terms must be retained. (iii) A signed-current estimate for $\sup_t |J_t|$, as discussed in Section 7.2, would improve the reservoir rule if directional cancellation is significant. After any such improvement the whole budget must be recomputed: once the phase term is reduced, $e^{-V/(4 \times 10^7)}$ becomes the next obstruction, and a small exponent improvement must never be advertised as a many-orders-of-magnitude copy-number result.

Finally, formal verification supports the theorems under their stated hypotheses. It says nothing about whether those hypotheses describe a particular piece of chemistry, and the calibration questions listed in Section 7.3 are precisely the ones that must be settled before this certificate is applied to a real biochemical process.

A Complete error budgets

For $V \geq 0$ define

$$\begin{aligned} M(V) &= 4e^{-V/2000} + 48000 V e^{-V/2000+1/50}, \\ R_{\text{err}}(V) &= e^{-3V/5000000} + e^{-V/2} + M(V) + e^{-V/10000} + 12000 V e^{-V/10000+9/500}, \\ F(V) &= e^{-V/40000000} + 100 e^{-V/10^{10}} + e^{-V} + e^{-V/200}, \\ J_{\text{err}}(V) &= F(V) + e^{-V/100000} + 2e^{-V/300} + e^{-V/2000}, \\ T(V) &= M(V) + R_{\text{err}}(V) + 4(e^{-V/320000} + e^{-V}), \\ e(V) &= e^{-1177V/10^9} + 4e^{-V/100000} + J_{\text{err}}(V) + T(V). \end{aligned}$$

Here $M, R_{\text{err}}, F, J_{\text{err}}, T$ are error functions, not the bath inventory, capacity, directed counters or net service; J_{err} is the joint five-counter budget and T the restart budget, and the pulse terms sit outside the prepared-cycle budget. This is exactly the one-cycle error of Theorem 3.1, with repeated material terms retained rather than silently optimized away.

The refined budget replaces the single phase term and adds the sharp service term:

$$\hat{F}(V) = e^{-V/40000000} + 100 e^{-\kappa V} + e^{-V} + e^{-V/200},$$

with $\hat{J}_{\text{err}}, \hat{e}_*$ defined from $\hat{F}$ exactly as J_{err}, e are defined from F , and $\hat{e}(V) = \hat{e}_*(V) + e^{-V/2000}$. Every exponential besides the phase term is bounded by a multiple of $e^{-V/(4 \times 10^7)}$, and the two positive offsets have exponent at most $1/50$ with $e^{1/50} \leq 2$; summing coefficients gives (24), and (25) completes Lemma 5.3. For the original budget the analogous statement is $e(V) \leq 100e^{-V/10^{10}} + 10^6(V+1)e^{-V/(4 \times 10^7)}$ with absorption valid from $V \geq 2 \times 10^{11}$, which is the origin of the floor V_0 .

Numerically, at $V_0 = 2 \times 10^{11}$ one has $e^{20} \geq \sum_{k=0}^{36} 20^k/k! > 4.848 \times 10^8$, so $e(V_0) \leq 101/(4.848 \times 10^8)$ and $100 e(V_0) \leq 21/10^6$. For the refined budget at $V = 9.6 \times 10^{10}$ we have $\kappa V \geq 20$, giving the same $21/10^6$; and at $V = 2 \times 10^{11}$ we have $\kappa V \geq 42$ and $e^{42} \geq (e^{20})^2 \geq 2.3503 \times 10^{17}$, giving $100 \hat{e}(V) \leq 10^{-13}$. At $V = 5 \times 10^{10}$, $\kappa V \geq 10$ and $e^{10} \geq 10100$ give $\hat{e}(V) \leq 1/100$. These are rational comparisons against positive Taylor sums, not rounded numerical exponentials.

B Quantitative exponential tests

For a stopped bounded-rate source and a linear observable H , the exponential test obeys

$$\frac{\mathcal{L}e^{sH(N)}}{e^{sH(N)}} = \sum_j a_j(N)(e^{s\Delta_j H} - 1). \quad (32)$$

On the small arguments used here, $e^z \leq 1 + z + \frac{3}{5}z^2$ for $|z| \leq \frac{9}{50}$, so drift plus the rate-weighted sum of squared jumps bounds (32). Uniformization turns a generator bound into a one-step multiplicative bound; iteration, the exponential Markov inequality and Poisson summation supply the stated tails, and indicators of activity kill a test on exit, every resulting exit probability being included explicitly in M, R_{err}, T .

Pulse. For independent Bernoulli retention indicators B_a and nonnegative weights w_a , $\mathbb{E} \exp(s \sum_a w_a B_a) = \prod_a (1 - p_a + p_a e^{s w_a})$. With material weights at most 2 and stock weights at most $9/5$, the parameter bounds give unrounded pulse material means in $[31213V/32000, 3231V/3200]$, and a food floor lowers a mean by less than one, leaving for $V \geq 10^6$ the slack required for $[24V/25, 51V/50]$. For stock, $1 - p + p e^{-s w} \leq \exp\{-(s - \frac{27}{25}s^2)pw\}$ for $0 \leq w \leq 9/5$, $0 \leq s \leq 1/10$; at $s = 1/100$ with $\sum_a p_a w_a \geq 49V/4000$ and level $3V/250$ the exponent is exactly $-1177V/10^9$. Either bad material band requires a centred deviation of magnitude at least $V/100$, giving $e^{-V/100000}$ per sign and coordinate.

Material corridor and terminal return. For each material coordinate use the barriers $U_H = e^{(H-11V/10)/100}$ and $L_H = e^{-(H-9V/10)/100}$. Their sum is at least one on exit and at most $4e^{-V/2000}$ at a prepared start; on the stopped corridor each barrier has generator at most $3000V e^{-V/2000+1/50}$, and integration for time 4 gives $M(V)$. For terminal return, the compensated tilt contracts the test parameter by $1 - 1/(3000V)$ per clock step, so for $n \geq 11700V$ the contraction factor is at most $1/40$; taking $s = \pm 1/1000$ with $|H_0 - V| \leq V/25$ gives exponent at most $-V/320000$ for the threshold $|H - V| \geq V/160$, and the clock of mean $12000V$ misses its cutoff with probability at most e^{-V} .

Recovery and residence. Below $3V/50$ the drift (17) plus its bounded-jump quadratic remainder gives $\mathcal{L}e^{-sY} \leq -\frac{5}{8}sYe^{-sY}$ for $0 \leq s \leq 1/100$; the tilt transported backward over $8100V$ steps from $1/1000$ reaches its cap because $(1 + \frac{5/8}{3000V})^{8100V} \geq (17/16)^{27} > 101/20$, and the exponential Markov inequality at prepared stock $3V/250$ gives $e^{-3V/5 \cdot 10^6}$, while a Poisson clock of mean $8250V$ falls below $8100V$ with probability at most $e^{-V/2}$. For continued residence, $H_Y = \max(e^{-Y/100}, e^{-3V/5000})$ has generator at most $3000V e^{-3V/5000+9/500}$, giving the lower-exit bound $e^{-V/10000} + 12000V e^{-V/10000+9/500}$.

Phase margin and collection. In (19), $a^n \geq 1/12$ for $n \leq 91V$, and $n \geq 89V$, $n - 1 \geq 88V$ give each coefficient at least $1/35$ times its stock weight. The phase exponential remainder per clock step is $(12/5)s^2$, so at $s = 10^{-6}$ its total over $91V$ steps is $(12/5) \cdot 91V \cdot 10^{-12}$; with the low- X threshold contributing $V/10^9$ and the initial stock $-V/(700 \cdot 10^6)$, the combined exponent is at most $-\kappa V$, which is the content of Lemma 5.1. After a burn-in of at least $90V$ clock steps this bounds low-free occupancy at every later step; if B counts bad steps among n collection steps then $\mathbb{E}B \leq ne^{-\kappa V}$ and $\mathbb{P}(B \geq n/100) \leq 100e^{-\kappa V}$. A marked tilt of size $1/1000$ for free washouts then gives the joint shortfall bound

$$\exp\left\{\frac{1}{1000}\left(K + \frac{B_*}{3000000}\right) - \frac{33n}{10^{11}}\right\} + \mathbb{P}(B \geq B_*),$$

with $K = V/1080 + 1$, $B_* = n/100$ and $n \geq 2990V$; for $V \geq 10^6$ the first term is at most $e^{-V/(4 \times 10^7)}$, and the precollection quarter-unit and collection unit clocks of means $750V$ and $3000V$ contribute e^{-V} and $e^{-V/200}$.

Template and supply arithmetic. Template washout has mean mark at least $1/84000$ per active clock step, so the tilt $1/1000$ gives multiplier at most $e^{-33/(2.8 \times 10^9)}$ and Poisson summation with mean $3000V$ gives $(V/56 + 1)/1000 + 3000V(e^{-33/(2.8 \times 10^9)} - 1) \leq -V/100000$ for $V \geq 10^6$. For a supply of rate at most cV with initial count z and threshold K , the upper-tail exponent is $-sK + 4cV(e^s - 1) + sz$: at $c = 1$, $s = 1/50$, $z \leq 151V/200$, $K = 5V$ it is at most $-V/300$; at $c = 9/200$, $s = 1/20$, $z = 0$, $K = V/5 - 1$ it is at most $-V/2000$ for $V \geq 10^6$; and at $c = 11/500$, $s = 1/10$, $K = V/10 - 1$ it is at most $-21V/31250 \leq -V/2000$, which is Lemma 5.5. All follow from the same quadratic exponential inequality, and the $+1$ and -1 threshold margins make the exported integer ceilings and floors immediate.

C Claim-to-evidence map and reproduction

All finite-bath sources reside in `proofs/FiniteReservoir/`, namespace `FiniteReservoir`; the full source, not a declaration count, determines the scope of each checked assertion. The pinned toolchain is Lean 4.30.0 with a checked-in `lake-manifest.json` fixing the Mathlib dependencies. The only logical axioms in the recorded interfaces are the standard Lean foundations `propext`, `Classical.choice` and `Quot.sound`; no `sorry` is accepted, and every receipt was produced with warnings treated as errors.

Claim	Principal source or evidence
Changed source, bath and coefficient rectangle (4)	<code>Reservoir.lean</code> , <code>FiniteModel.lean</code> , <code>CandidateParameters.lean</code>
Material, growth, noise and phase inequalities (14)–(20)	<code>CountSource.lean</code> , <code>StockNoise.lean</code> , <code>StockExponential.lean</code> , <code>PhaseBounds.lean</code> , <code>PhaseExponential.lean</code>
Joint stopped counters, prepared cycle, full pulse	<code>JointCounterBudget.lean</code> , <code>PreparedCycle.lean</code> , <code>SafeCycle.lean</code>
Nonexplosion and the literal chronological law	<code>PhysicalNonexplosion.lean</code> , <code>JointSourceBridge.lean</code> , <code>JointTrajectory.lean</code> , <code>LiteralCycleBound.lean</code>
Full-history joint mission (Theorem 3.1)	<code>finite_bath_horizon</code> in <code>Mission.lean</code> ; <code>full_returned_history_product</code> in <code>FullProduct.lean</code>
Directed inventory and all-prefix tolerance (Theorem 3.3)	<code>BathInventory.lean</code> , <code>BathPrefixes.lean</code> , <code>Sizing.lean</code>
Original inverse design and worked example	<code>pure_inverse_design</code> in <code>InverseDesign.lean</code> ; <code>example_mission</code> in <code>MissionExample.lean</code>
Terminal-stock synthesis (Proposition 3.5)	<code>HistoryTrace.certified_net_lower</code> , <code>ProductionCorollaries.lean</code>
Neighbouring detailed balance (Proposition 6.1)	<code>model_detailed_balance</code> , <code>Thermochemistry.lean</code> , <code>ThermochemicalModel.lean</code>
Correlated gross envelope (Lemma 5.4) and stopped tail (Lemma 5.5)	<code>pure_gross_envelope</code> , <code>pure_supply_cap</code> , <code>pure_sharp_stopped_tail</code> in <code>SharpService.lean</code>
Sharp one-cycle event, strict transfer, literal cycle (Section 5.2)	<code>SharpCycle.lean</code> : <code>literal_sharp_cycle_bound</code>
Unrounded phase rate κ through window, occupation and collection (Lemma 5.1)	<code>SharpPhase.lean</code> : <code>phaseRate</code> , <code>residence_free_window_sharp</code> , <code>joint_free_tail_sharp</code>
Refined counter budget and both literal one-cycle bounds	<code>SharpPhaseCycle.lean</code> : <code>literal_cycle_phase_bound</code> , <code>literal_both_cycle_bound</code>
Phase-refined mission for <i>every</i> admitted bath (Theorem 3.2(a))	<code>SharpGeneral.lean</code> : <code>full_returned_phase_success</code> , <code>finite_bath_horizon_sharp</code> , <code>loaded_hundred_sharp</code>
Envelope, absorption and confidence budgets (Lemma 5.3)	<code>SharpEnvelope.lean</code> : <code>sharp_error_envelope</code> , <code>sharp_residual_absorbed</code> , <code>sharp_single_exponential</code> , <code>sharp_hundred_budget</code>
Sharp history composition, prefix tolerance, reduced-scale instance (Theorem 3.2)	<code>SharpMission.lean</code> : <code>full_returned_sharp_success</code> , <code>pure_prefix_tolerance_sharp</code> , <code>sharp_mission_bound</code> , <code>sharp_example_mission</code>
Refined inverse design (Corollary 3.4) and halved-bath instance	<code>SharpResolution.lean</code> : <code>sharp_inverse_design</code> , <code>sharp_hundred_halved_bath</code> , <code>sharp_finite_reservoir_operation</code>
Endpoint bath energies (Proposition 6.2), force inversion (Corollary 6.3)	Conventional algebra from the compiled potential and inventory identities; replayed in <code>check_paper.py</code>

Claim	Principal source or evidence
Metered food (Proposition 6.4)	<code>MeteredFood.lean</code> for rate agreement and no-stockout; the path coupling is conventional
Tilt optimality (Remark 5.2), amounts and rate units (Section 7.3) Figures 2–4 and Table 3	Conventional arithmetic, replayed in <code>check_paper.py</code> <code>make_figures.py</code> ; the diagnostics integrate the mass-action ODE and are labelled as such

Reproduction. From the repository root, the strict receipts for the seven modules new to this paper are produced by

```
python scripts/verify_proof.py proofs/FiniteReservoir/SharpService.lean
python scripts/verify_proof.py proofs/FiniteReservoir/SharpCycle.lean
python scripts/verify_proof.py proofs/FiniteReservoir/SharpPhase.lean
python scripts/verify_proof.py proofs/FiniteReservoir/SharpPhaseCycle.lean
python scripts/verify_proof.py proofs/FiniteReservoir/SharpEnvelope.lean
python scripts/verify_proof.py proofs/FiniteReservoir/SharpGeneral.lean
python scripts/verify_proof.py proofs/FiniteReservoir/SharpResolution.lean
```

each writing a JSON receipt recording `verified`, the exit code, the source hash, the toolchain and manifest hashes and the compiler diagnostics; all seven report `verified: true` with empty diagnostics. The original root `Resolution.lean` and `ProductionCorollaries.lean` are unchanged and replay the same way. In the manuscript directory, `check_paper.py` rederives every printed constant (92 checks) and `make_figures.py` regenerates the three figures and the diagnostics table. The repository location will be supplied with the author information.

Data and code availability. The repository location will be supplied with the author information.

References

- [1] Anonymous. Reliable repeated harvesting in a finite-molecule autocatalytic reactor: a source-level finite-horizon operating theorem with a machine-checked proof. Companion manuscript, 2026.
- [2] Anonymous. Productive recovery and repeated harvesting in a reversible autocatalytic reactor: a uniform operating theorem with a machine-checked proof. Companion manuscript, 2026.
- [3] Anonymous. Thermochemically consistent realization of a stochastic autocatalytic exporter: chemical completion, inherited operating certificates, cycle affinities and a necessary initiation scale. Companion manuscript, 2026.
- [4] 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.
- [5] Sandeep Ameta, Simon Arsène, Sophie Foulon, Baptiste Saudemont, Ben E. Clifton, Andrew D. Griffiths, and Philippe Nghe. Darwinian properties and their trade-offs in autocatalytic RNA reaction networks. *Nature Communications*, 12:842, 2021. doi: 10.1038/s41467-021-21000-1.

- [6] 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.
- [7] 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.
- [8] 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.
- [9] 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.
- [10] Massimiliano Esposito. Stochastic thermodynamics under coarse graining. *Physical Review E*, 85(4):041125, 2012. doi: 10.1103/PhysRevE.85.041125.
- [11] Jonas H. Fritz, Basile Nguyen, and Udo Seifert. Stochastic thermodynamics of chemical reactions coupled to finite reservoirs: a case study for the Brusselator. *The Journal of Chemical Physics*, 152(23):235101, 2020. doi: 10.1063/5.0006115.
- [12] 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.
- [13] 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.
- [14] 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.
- [15] Steven R. Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform chernoff bounds via nonnegative supermartingales. *Probability Surveys*, 17:257–317, 2020. doi: 10.1214/18-PS321.
- [16] 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.
- [17] Stuart A. Kauffman. Autocatalytic sets of proteins. *Journal of Theoretical Biology*, 119(1):1–24, 1986. doi: 10.1016/S0022-5193(86)80047-9.
- [18] 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.
- [19] 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.
- [20] 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.

- [21] 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.
- [22] 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.
- [23] 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.
- [24] Benedikt Remlein, Massimiliano Esposito, and Francesco Avanzini. What is a chemostat? Insights from hybrid dynamics and stochastic thermodynamics. *The Journal of Chemical Physics*, 162(22):224113, 2025. doi: 10.1063/5.0266086.
- [25] Benedikt Remlein, Massimiliano Esposito, and Francesco Avanzini. Emergence of open chemical reaction network thermodynamics within closed systems. Preprint, [arXiv:2604.14848](https://arxiv.org/abs/2604.14848), 2026.
- [26] Tim Schmiedl and Udo Seifert. Stochastic thermodynamics of chemical reaction networks. *The Journal of Chemical Physics*, 126(4):044101, 2007. doi: 10.1063/1.2428297.
- [27] Udo Seifert. Stochastic thermodynamics, fluctuation theorems and molecular machines. *Reports on Progress in Physics*, 75(12):126001, 2012. doi: 10.1088/0034-4885/75/12/126001.
- [28] 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.
- [29] 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.