Certified one-bit chemical computation:
finite-fuel correction, productive readout, and inheritance

17 September 2026

Abstract

A chemical module computes reliably only if the state that encodes its input survives the act of using it: making the output consumes the same finite material that maintains the state, and dividing the compartment splits that state at random. We construct a seven-species reversible count reaction network, with one fuel species, whose compartment realises a complete one-bit contract. Two preparations with identical molecular support encode the two input values; the network repairs a bounded number of wrong-label molecules, delivers at least four molecules of the corresponding product by a fixed deadline, and returns *both* complementary daughters to the region that encodes the same value. A single probability bound covers all of these events on one trajectory. For an 81-unit preparation we prove a uniform joint guarantee of at least 0.9997 over a parameter box that simultaneously spans a hundredfold range in the correction rate constant, admits product recovery as low as 90%, and admits division bias anywhere in $[0.48, 0.52]$; under ideal operations the same construction gives 0.9998. The proof combines a fixed-time reaction-clock comparison with a downward-rounded integer certificate on a 3 240-state core, so that no floating-point matrix exponential enters the argument. Correction consumes fuel, and its embedded jump chain is exactly solvable: it is the Mabinogion urn, whose wrong-consensus probability $2^{-(r-3)}$ from two minority molecules among r residents is independent of how fast the correcting reaction is. Weighting the exact absorption law by the fuel actually spent yields a 144-unit, two-minority construction with joint probability at least 0.9916, within $5.5 \cdot 10^{-4}$ of the ceiling that the consensus law imposes on its worst admitted state. The same law makes heritable variation quantitative: a specified mixed preparation produces a productive *opposite*-program daughter pair with probability above 0.0074, whereas a pure newborn does so with probability at most $8.1 \cdot 10^{-6}$, because on a pure face the only channel that can create a minority molecule is first-order leakage. No single label-dependent switching rate can therefore describe both classes of newborn. An external product-dependent retention rule converts the output distinction into an explicit finite-population selection bound. The probability theorems are conventional arguments supported by exact finite computation; selected source, partition, rate and arithmetic lemmas are checked in Lean 4. High-order reactions, severe kinetic separation and idealised inventory operations remain assumptions: this is a certified mathematical model of one-bit chemical information processing, not a calibrated biochemical implementation.

1 Introduction

1.1 The contract

Chemical reaction networks compute. Stochastic reaction networks with a finite number of molecules are a well-studied model of computation [5, 8, 30], abstract networks can be compiled into DNA [6, 27, 31], and specific motifs implement specific primitives: approximate majority is computed by a population protocol [1], by a tri-molecular reaction network [7], and by the biological cell-cycle switch [3].

Almost all of this work certifies a *transient output*: the network is started in a state encoding the input, and one asks for the distribution of a designated species at a later time. A physical module that is used repeatedly must satisfy something stronger. Producing the output consumes material. The compartment holding the state divides, and the state is split by a random partition that no controller supervises [15, 16]. If we want to run the module again, the state that remains after paying for the output and after being cut in half must still be an admissible input state — for *both* halves, not for a favourable one.

We therefore fix the following contract, and prove a single probability bound for all of its clauses *on one trajectory and one partition*. Let the logical input be $b \in \{X, Y\}$.

(C1) **Encoding.** b is encoded by the composition of the compartment: an admissible encoding is any state of an explicitly specified finite region $\mathcal{B}_b^{L,k}$. The two regions have identical molecular support, so the encoder cannot be said to hand the network a distinguishing species.

(C2) **Correction.** A preparation may contain up to k molecules of the wrong label. The network must remove them, using its own fuel, before they corrupt the output.

(C3) **Productive readout.** At a fixed deadline T the compartment must contain, and surrender to an imperfect collection step, at least four molecules of the product P_b and at most one molecule of $P_{\bar{b}}$.

(C4) **Inheritance.** After the product is removed and every remaining molecule is assigned at random to one of two daughters, *both* daughters, refilled by a label-blind procedure, must lie in $\mathcal{B}_b^{L,k}$ again.

(C5) **Uniformity.** The guarantee must hold uniformly over all admissible encodings of b, over an interval of rate constants, and over an interval of operation parameters — not at a single tuned point.

The interesting word is *and*. In our network the reaction that copies a resident molecule and the reaction that makes the product have the *same* propensity sf and draw on the *same* finite food pool, so (C3) and (C4) compete for one resource; and the correction reaction of (C2) burns a fuel molecule that is not replenished during the batch. The theorem is a statement about a resource conflict, not about a memory model with an independent reward counter attached.

1.2 Results

The source (Table 1) has seven species X, Y, F, P_X, P_Y, H, W and seven reversible pairs, all with positive rate constants in both directions. Its correction pair, read forward and with the fuel bookkeeping deleted, is $X + X + Y \to 3X$ together with its label swap: exactly the tri-molecular approximate-majority network of Condon et al. [7]. What we add is that each correction event consumes one molecule of fuel H and releases one molecule of waste W, that the reverse reaction is present, and that the correction runs concurrently with the food-consuming reactions that make the output. The results are the following.

- **A joint certificate (Theorem 2.3).** For a compartment holding 81 modelled material units, contract (C1)–(C5) holds with probability at least 9997/10000 in a regime (WO) whose kinetic box spans a hundredfold range in γ and which simultaneously tolerates product recovery $\eta \in [0.9, 1]$ and division bias $\theta \in [0.48, 0.52]$. Restricting either the kinetics or the operations gives 4999/5000.
- **A logical reading (Corollaries 2.4 and 2.5).** Fixed decoders of the two daughters and of the collected product all return the input bit, with total-variation distance at most

$2 \cdot 10^{-4}$ (respectively $3 \cdot 10^{-4}$) between the decoded triple and the constant triple (b, b, b); each decoded output carries at least 0.9972 (respectively 0.9960) bits about a uniformly random input bit.

- **An exact integer certificate (Proposition 3.2).** The terminal payoff of a 3 240-state pure core is bounded below by a downward-rounded integer uniformization recurrence with certified 64-bit accumulator bounds. No floating-point matrix exponential, and no statistical confidence interval, enters any claim.
- **The fuel law (Theorem 5.1, Corollary 5.3).** The embedded correction chain has jump probabilities $(r-j-1)/(r-2)$ and $(j-1)/(r-2)$, independent of the rate constant, of the fuel level and of time. Its unlimited-fuel absorption probability is the Mabinogion urn law of Flajolet and Huillet [10]; from two minority molecules the wrong-consensus probability is exactly $2^{-(r-3)}$. Hence a correction-only route to error δ needs $r \geq 3 + \lceil \log_2(1/\delta) \rceil$ residents — no amount of catalytic speed substitutes for inventory.
- **Spent fuel as a state variable (Lemma 5.4, Theorem 5.5).** Charging each successful trajectory the fuel it actually spent, rather than the fuel it was given, improves the two-minority guarantee from $0.99010\ldots$ to $0.99164\ldots$ at 144 units, and the improved bound is within $5.5 \cdot 10^{-4}$ of the ceiling $127/128$ that Theorem 5.1 imposes on its worst admitted state (Remark 5.6).
- **Variation is not a label-only rate (Theorem 6.2, Corollary 6.3).** A specified mixed preparation yields a productive *opposite*-program daughter pair with probability > 0.0074. A pure newborn yields one with probability at most $8.1 \cdot 10^{-6}$, because on a pure face with no waste the majority channels are both disabled and the only minority-creating channel is first-order leakage. The two classes differ by a factor above 935, so no two-state mutation matrix indexed by the label alone is consistent with the source.
- **Selection (Theorem 7.2).** An external rule that reads only the collected product, never a label, moves the compartment-count odds of a 20-member population from $1:1$ to at least $5:4$ in one round with probability at least $16297339/16384000$, with no independence assumption on the chemical failures.

1.3 What is not claimed

We state the boundary once, and keep to it.

Not universal computation. The module stores and transmits one bit and emits one of two products. "Program" means a choice between two product recipes. We do not compose modules, and the fan-out described here is the copying of one bit into two compartments, not the production of two independent bits usable by arbitrary downstream circuits.

Not a priority claim. Chemical majority computation [1, 3, 7], programmable chemical kinetics [6, 27, 31], experimentally programmable autocatalysis [19], compositional heredity [24, 29] and compartment-level competition and selection [23, 25] all have substantial predecessors. Our contribution is the joint finite-inventory certificate and the exactly solvable finite-fuel correction law, not the individual ingredients. We are also not the first to verify reaction-network statements in a proof assistant [4, 22] or by probabilistic model checking [20, 21]; our formal component is modular and its exact boundary is given in Appendix A.

Not a physical implementation. The correction reaction is fourth order in both directions, and the required kinetic separation is extreme (Section 9). A candidate elementary mechanism that we examined fails for reasons we quantify in Section 8 rather than dismiss. Whether *some* elementary mechanism satisfies the same contract is open, and by Phan et al. [26] the realizability question is itself hard. The species are abstract; no molecular identities are assigned.

Not evolution. We prove one-cycle heritable variation from a specified mixed preparation and a one-round external selection bound. We do not prove a useful recurrent mutation rate from the actual descendants of pure newborns; Theorem 6.2 shows why that is a genuinely different question, and bounds it from the wrong side.

Not an end-to-end formal proof. The probability theorems are conventional arguments supported by exact rational and integer computation. Lean 4 checks the reaction table and its invariants, the repair clock, the partition identity, the uniformization and block-iteration lemmas, the prefix maximum, the exponential estimates, the embedded jump probabilities, the finite-fuel recurrence and the selection tails. It does not replay the 3 240-state recurrence, and an arithmetic theorem about a reported numerator is not a theorem about the probability of the source.

1.4 Organisation

Section 2 gives the source, the protocol and the main theorem, with the logical corollaries. Section 3 proves the exact core certificate, Section 4 the fixed-time coupling and the one-minority regimes. Section 5 develops the finite-fuel correction law and the two-minority regime, Section 6 the asymmetry of heritable variation. Section 7 covers iteration, the material bill and external selection. Section 8 states what the certificate does and does not scale, and Section 9 the dimensional, thermodynamic and mechanistic limits. Section 10 discusses open problems. Appendix A records the evidence classes, the formal-verification scope and the reproduction commands; Appendix B lists the exact rational constants.

2 The source, the protocol and the main theorem

2.1 The reaction source

The state of a compartment is

$$z = (x, y, f, p_X, p_Y, h, w) \in \mathbb{N}^7, \tag{1}$$

the counts of X, Y, F, P_X, P_Y, H, W. Every species carries one *modelled material unit*; this is a formal conservation convention, not an elemental formula or a molecular mass. Table 1 is the entire reaction table. Propensities are count propensities with falling factorials $(n)_k = n(n-1)\cdots(n-k+1)$ and no further combinatorial divisor; a reaction with insufficient reactants has propensity zero. The coefficients are fixed throughout a batch and are identical for the two preparations.

Reversible pair	Forward propensity	Reverse propensity
$S + F \rightleftharpoons 2S, \quad S = X, Y$	sf	$s(s-1)/100$
$S + F \rightleftharpoons S + P_S, \quad S = X, Y$	sf	$sp_S/100$
$X \rightleftharpoons Y$	εx	εy
$2X + Y + H \rightleftharpoons 3X + W$	$\gamma h\, x(x-1)y$	$\beta w\, x(x-1)(x-2)$
$2Y + X + H \rightleftharpoons 3Y + W$	$\gamma h\, y(y-1)x$	$\beta w\, y(y-1)(y-2)$

Table 1: The seven reversible pairs. All rate constants are positive. Replication and production have the *same* propensity sf and consume the same food F: production is not an independent reward process. Each correction event spends one H and creates one W, which enables the reverse correction. Deleting H and W from the last two rows leaves the tri-molecular approximate-majority network of Condon et al. [7].

For a fuel inventory $L \geq 1$ the reachable state space is contained in

$$\Omega_L = \{z : \; x + y + f + p_X + p_Y = 80, \;\; h + w = L, \;\; x + y \geq 1\}. \tag{2}$$

Every transition with positive propensity preserves the two material balances, and reverse replication cannot consume the last resident because its propensity $s(s-1)/100$ vanishes at $s = 1$. Since Ω_L is finite and all rates are finite and nonnegative, Table 1 defines a nonexplosive continuous-time Markov chain (CTMC) on Ω_L; no truncation and no overflow state is introduced. (Both balance statements are machine-checked; see Appendix A.)

2.2 Encodings and the batch protocol

Definition 2.1 (Restart regions)**.** For a minority allowance $k \in \{1, 2\}$ put

$$\mathcal{B}_X^{L,k} = \{z \in \Omega_L : \ x \geq 8, \ 0 \leq y \leq k, \ p_X = p_Y = w = 0, \ h = L\}, \tag{3}$$

and let $\mathcal{B}_Y^{L,k}$ be its image under exchanging $X \leftrightarrow Y$ and $P_X \leftrightarrow P_Y$. The two regions are disjoint. They contain 145 states each for $k = 1$ and 216 each for $k = 2$.

The two regions have the same molecular support: $(8, 1, 71, 0, 0, 1, 0)$ and $(1, 8, 71, 0, 0, 1, 0)$ contain the same species in the same total amount. A later daughter need not retain a minority molecule; the regions are where a compartment must *be*, not a promise that mixed support persists.

Definition 2.2 (Batch operation)**.** Run the source of Table 1, unchanged, from $z \in \mathcal{B}_b^{L,k}$ to the fixed deadline $T = 1 + \tau$. Then:

(O1) **Harvest.** Remove *all* P_X and P_Y. Each removed molecule is credited to the collection independently with probability η; unrecovered product is discarded as waste. Write Q_X, Q_Y for the credited counts.

(O2) **Divide.** Assign every remaining molecule to daughter A with probability θ and otherwise to daughter B, independently conditional on the mother and on θ. The daughters are complementary: there is one draw, not two.

(O3) **Refill.** Restore each daughter's core to 80 using F alone; remove its W and restore H to L; restore the reference operating volume. No resident is added, removed, selected, relabelled or reseeded, and no operation reads b.

Write Z_A^+, Z_B^+ for the refilled daughters. The recovery coins and the partition coins are conditionally independent given the mother, and their parameters may be chosen anywhere in the stated intervals without changing the reaction table.

For $i \in \{X, Y\}$ the *joint event* is

$$\mathcal{J}_i = \{\, Q_i \geq 4, \ \ Q_{\bar{i}} \leq 1, \ \ Z_A^+ \in \mathcal{B}_i^{L,k}, \ \ Z_B^+ \in \mathcal{B}_i^{L,k} \,\}. \tag{4}$$

All four clauses refer to the same trajectory and the same partition. Trajectories outside $\mathcal{J}_i$ are failures of a sufficient proof event; they are not removed from the law, and the process is not conditioned on success at any point.

2.3 Operating regimes and the main theorem

Theorem 2.3 (Productive inheritance)**.** *For each regime of Table 2, each $i \in \{X, Y\}$, each $z \in \mathcal{B}_i^{L,k}$ and every choice of $(\gamma, \beta, \varepsilon, \theta, \eta)$ in the stated intervals,*

$$\mathbb{P}_z(\mathcal{J}_i) \ \geq \ q, \tag{5}$$

with $q = 4999/5000$ in the regimes Original, W and O; $q = 9997/10000$ in regime WO; and $q = 2479/2500$ in regime F.

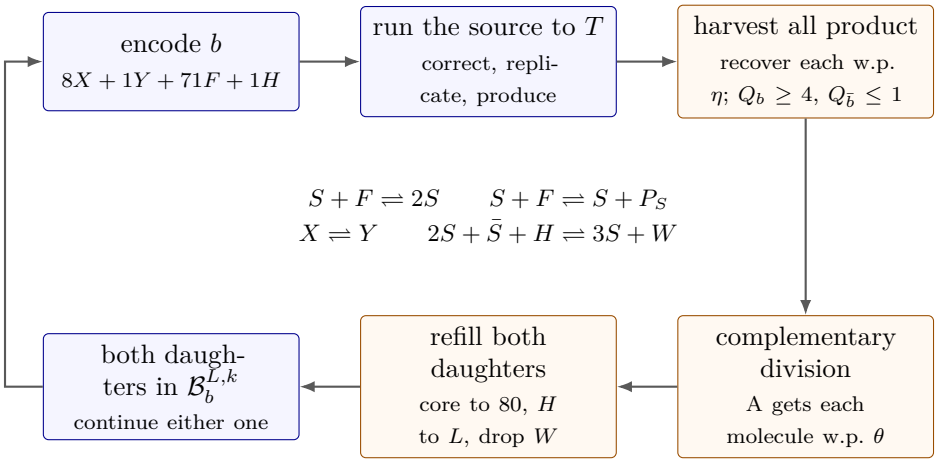

Figure 1: The source and the external batch protocol. Blue: chemistry and the encoded state. Orange: inventory operations, which are label-blind — no step reads b, and no resident molecule is supplied, selected or reseeded. The label is used only to state and score the event (4). This is a mathematical protocol, not a drawing of an implemented vesicle.

	Original	W wider rates	O imperfect ops	**WO** both	**F** two minorities
(L, k)	$(1, 1)$	$(1, 1)$	$(1, 1)$	$(1, 1)$	$(64, 2)$
initial material	81	81	81	81	144
γ	$[10^9, 2 \cdot 10^9]$	$[2 \cdot 10^7, 2 \cdot 10^9]$	$[10^8, 2 \cdot 10^9]$	$[2 \cdot 10^7, 2 \cdot 10^9]$	$[2 \cdot 10^7, 2 \cdot 10^9]$
β	$[10^{-13}, 10^{-12}]$	$[10^{-13}, 6 \cdot 10^{-11}]$	$[10^{-13}, 10^{-11}]$	$[10^{-13}, 6 \cdot 10^{-11}]$	$[10^{-13}, 5 \cdot 10^{-11}]$
ε	$[10^{-10}, 10^{-9}]$	$[10^{-10}, 10^{-7}]$	$[10^{-10}, 10^{-8}]$	$[10^{-10}, 10^{-7}]$	$[10^{-10}, 10^{-7}]$
$\tau = T - 1$	10^{-9}	$1/(8 \cdot 10^7)$	$1/(4 \cdot 10^8)$	$1/(8 \cdot 10^7)$	10^{-7}
θ	$1/2$	$1/2$	$[0.48, 0.52]$	$[0.48, 0.52]$	$[0.48, 0.52]$
η	1	1	$[0.9, 1]$	$[0.9, 1]$	$[0.9, 1]$
certified q	0.9998	0.9998	0.9998	**0.9997**	**0.9916**
exact bound	(19)	(20)	(21)	(22)	(31)

Table 2: Operating regimes; all coefficients not listed are fixed by Table 1. In each coordinate the box WO is the union of the boxes W and O, at the cost of one unit in the fourth decimal; the deadline τ is a protocol choice, and $\tau = 1/(8 \cdot 10^7)$ is the value forced by the slowest admitted γ. Regime F admits two minority molecules and the same imperfect operations. Rates and deadlines are stated in model units; Section 9 discusses their dimensions.

Regimes W and O appear separately because their sharper constant 0.9998 is of independent interest; regime WO is the statement that one does not have to choose. The smallest admissible preparations are $(8, 1, 71, 0, 0, 1, 0)$ and its exchange for $k = 1$, and $(8, 2, 70, 0, 0, 64, 0)$ and its exchange for $k = 2$. These are absolute-count statements: scaling the volume and all populations at fixed concentration destroys the one-minority mechanism, because the falling factorial $y(y - 1)$ is not a concentration.

Theorem 2.3 is proved in Sections 3–5: the terminal payoff of a pure core in Proposition 3.2, the correction prefix and the departure hazard in Lemma 4.1, the one-minority regimes in Theorem 4.2, and regime F in Theorem 5.5.

2.4 The logical reading

Theorem 2.3 is a statement about molecule counts. The following two corollaries say what it means as computation. Fix k and L, and define

- the *compartment decoder* $\mathrm{dec_c}(z) = i$ if $z \in \mathcal{B}_i^{L,k}$, and $\mathrm{dec_c}(z) = X$ otherwise (any fixed

convention outside the two regions will do);

- the *product decoder* $\mathrm{dec_p}(Q_X, Q_Y) = X$ if $Q_X \geq Q_Y$ and Y otherwise.

Corollary 2.4 (Decoded computation)**.** *In each regime of Table 2, for every admissible encoding z of b,*

$$\mathbb{P}_z\Big((\mathrm{dec_c}(Z_A^+),\, \mathrm{dec_c}(Z_B^+),\, \mathrm{dec_p}(Q_X, Q_Y)) = (b, b, b)\Big) \;\geq\; q. \tag{6}$$

Consequently the law of the decoded triple is within total variation $1 - q$ of the point mass at (b, b, b), and each of the three decoded outputs individually equals b with probability at least q.

Proof. On $\mathcal{J}_b$ the two refilled daughters lie in $\mathcal{B}_b^{L,k}$, so both compartment decoders return b; and $Q_b \geq 4 > 1 \geq Q_{\bar{b}}$, so the product decoder returns b. Thus $\mathcal{J}_b$ is contained in the event that the triple equals (b, b, b), and Theorem 2.3 applies. A bound on the probability that a random variable differs from a constant bounds the total variation distance to the point mass at that constant. □

The joint statement is strictly stronger than the three marginal statements, because the three outputs are produced by one trajectory and one partition and are certainly not independent. Note also what the corollary does *not* say: it is a total-variation bound on *decoded* outputs, not a claim that the microscopic terminal distribution is close to a single ideal chemical state.

Corollary 2.5 (Information per batch)**.** *Let B be a uniformly random input bit, encoded by any admissible state of $\mathcal{B}_B^{L,k}$, and let D be any one of the three decoded outputs of Corollary 2.4. Then, with h_2 the binary entropy in bits,*

$$I(B; D) \;\geq\; 1 - h_2(1 - q). \tag{7}$$

Numerically this is at least 0.997254 *bits in regimes W and O,* 0.996056 *bits in regime WO and* 0.930011 *bits in regime F.*

Proof. $\mathbb{P}[D \neq B] \leq 1 - q \leq 1/2$, so Fano's inequality for a binary variable gives $H(B \mid D) \leq h_2(\mathbb{P}[D \neq B]) \leq h_2(1 - q)$, and $I(B; D) = 1 - H(B \mid D)$. □

This is an information-theoretic reading of an existing bound, not a channel-coding theorem and not a measured rate: it counts neither the time nor the material cost of the measurement, both of which are outside the modelled inventory.

3 The finite core and its exact lower certificate

Suppress leakage and correction, and restrict the state to one pure resident face. With resident count n, selected-product stock p and food $f = 80 - n - p$, the four surviving propensities are

$$\underbrace{nf}_{\text{replicate}}, \quad \underbrace{n(n-1)/100}_{\text{reverse}}, \quad \underbrace{nf}_{\text{produce}}, \quad \underbrace{np/100}_{\text{reverse}}. \tag{8}$$

The states $1 \leq n \leq 80$, $0 \leq p \leq 80 - n$ number 3 240. Let G be the row generator of this core chain acting on column payoff vectors. The terminal payoff of the protocol is

$$H_{\theta,\eta}(n, p) = a_p(\eta)\, g_n(\theta), \qquad a_p(\eta) = \mathbb{P}\{\mathrm{Bin}(p, \eta) \geq 4\}, \qquad g_n(\theta) = \mathbb{P}\{8 \leq \mathrm{Bin}(n, \theta) \leq n - 8\}. \tag{9}$$

Two modelling points are built into (9). First, g_n uses *one* binomial draw: if daughter A receives J residents then B receives $n - J$, so requiring $8 \leq J \leq n - 8$ is exactly the requirement that both daughters are admissible. Second, a_p is evaluated at the *terminal stock* p, not at the number of forward production events: the production reaction is reversible and its reverse consumes product. The one-unit joint probability from a core state is $(e^G H_{\theta,\eta})$ evaluated there.

Lemma 3.1 (Endpoint reduction)**.** *For all $\theta \in [0.48, 0.52]$ and $\eta \in [0.9, 1]$ and every state (n, p),*

$$H_{\theta,\eta}(n,p) \;\geq\; H_{0.48,0.9}(n,p).$$

Proof. A binomial count is stochastically nondecreasing in its success parameter (couple all trials to a common uniform vector), so $a_p(\eta) \geq a_p(0.9)$ for $\eta \geq 0.9$; and $a_p \geq 0$, $g_n \geq 0$, so it suffices to treat the two factors separately. For $n < 16$ we have $g_n \equiv 0$. For $n \geq 16$, differentiating the jth binomial mass $\binom{n}{j}\theta^j(1-\theta)^{n-j}$ gives n times the difference of the $(j-1)$st and jth masses for $n-1$ trials, so the sum defining g_n telescopes to

$$g_n'(\theta) = n\binom{n-1}{7}\theta^7(1-\theta)^7\big((1-\theta)^{n-15} - \theta^{n-15}\big). \tag{10}$$

The bracket is nonnegative for $\theta \leq 1/2$ and nonpositive for $\theta \geq 1/2$, so g_n increases up to $1/2$ and decreases after it. Since $g_n(\theta) = g_n(1-\theta)$, we get $g_n(0.52) = g_n(0.48)$ and hence $\min_{[0.48,0.52]} g_n = g_n(0.48)$. □

Lemma 3.1 is what makes a statement about a continuum of operation parameters provable by a finite computation: it is an analytic reduction to one endpoint, not a scan over a grid of parameter values. (The sign of (10) on the left half-interval and the monotonicity of the product are machine-checked; the telescoping identity itself is a conventional step.)

Proposition 3.2 (Exact core certificate)**.** *Uniformly over every pure newborn $(n, 0)$ with $8 \leq n \leq 80$,*

$$(e^G H_{1/2,1})(n,0) \;\geq\; c := \frac{2147232289}{2147483648}, \tag{11}$$

$$(e^G H_{0.48,0.9})(n,0) \;\geq\; c_{\mathrm{O}} := \frac{2147153442}{2147483648}. \tag{12}$$

Both minima are attained at $n = 8$; the range includes $n = 80$, for which the newborn has no food at all.

Certificate and its soundness. All four rates in (8) are integer multiples of $1/100$, and the total exit rate is bounded by $\lambda = 3216$ (Lemma 4.1 proves the sharper 3215.6). Put $P = I + G/\lambda = A/D$ with $D = 100\lambda = 321600$; then A is a nonnegative integer matrix and every row of A sums to D, so P is a stochastic matrix and the uniformization identity $e^G = e^{-\lambda}\sum_{j\geq 0}(\lambda P)^j/j!$ holds [12, 18]. Let $S = 2^{31}$ and

$$e_- := \sum_{m=0}^{25}\frac{(-1)^m}{m!} < e^{-1}, \qquad w_j := \Big\lfloor \frac{S\,e_-}{j!} \Big\rfloor \quad (0 \leq j \leq 18), \tag{13}$$

where the strict inequality is the alternating-series remainder estimate. All w_j are nonnegative and $w_j = 0$ for $j > 12$ at this scale. For an integer vector $0 \leq v \leq S$ define

$$v^{[0]} = v, \qquad v^{[j+1]} = \lfloor Av^{[j]}/D \rfloor, \qquad \mathcal{D}(v) = \Big\lfloor \frac{\sum_{j=0}^{18} w_j v^{[j]}}{S} \Big\rfloor, \tag{14}$$

all roundings coordinatewise and downward. Since A is nonnegative, induction gives $v^{[j]}/S \leq P^j(v/S)$ coordinatewise, and therefore

$$\mathcal{D}(v)/S \;\leq\; \sum_{j\geq 0}\frac{e^{-1}}{j!}P^j(v/S) = e^{G/\lambda}(v/S), \tag{15}$$

where dropping the terms with $j > 18$ is legitimate because every omitted term is nonnegative. Start from $v = \lfloor S\,H_{\theta,\eta} \rfloor \leq S\,H_{\theta,\eta}$ and apply $\mathcal{D}$ exactly $\lambda = 3216$ times. The Markov semigroup

is monotone, so composing the 3216 blocks gives $\mathcal{D}^{3216}(v)/S \le e^G H_{\theta,\eta}$ coordinatewise. Taking the minimum over the 73 newborn coordinates $(n,0)$, $8 \le n \le 80$, yields (11) and (12).

The implementation is exact in signed 64-bit arithmetic: the largest matrix accumulator is bounded by $DS = 690630741196800 < 2^{63}$ and the largest weighted sum by $S\sum_j w_j = 4611686007689969664 < 2^{63}$; every summand is nonnegative and the invariant $0 \le v_i \le S$ is preserved, so the bounds control partial sums as well. No floating-point matrix exponential enters at any stage. The complete state enumeration, the row sums, the outgoing transitions and the minimum are recomputed by the supplied scripts; this execution is exact-computation evidence, and Appendix A states exactly which parts of the argument are additionally kernel-checked. □

Remark 3.3. The rounding is deliberately one-sided in the safe direction, and the resulting slack is small: at $(8,0)$ with ideal operations the certified value (11) is $0.99988295\ldots$ against an unrounded double-precision evaluation of $0.99989142\ldots$, so the certificate costs about $8.5\cdot 10^{-6}$ of the error budget. Increasing the core inventory tightens the payoff quickly — the same computation at $K = 64$ gives only $0.997547\ldots$ and at $K = 96$ gives $0.999983\ldots$ — which is why the core is fixed at 80 resident-plus-product units.

4 A fixed-time coupling for one-minority correction

The comparison chains in this section are proof devices only. The physical source of Table 1 runs unchanged throughout the batch; we never switch rates at a random time and never condition on a stopping time.

Lemma 4.1 (Prefix and departure bounds)**.** *Consider the* repair-only *chain obtained from Table 1 by retaining only the two forward correction reactions. Along it, the total propensity of all suppressed channels is at most*

$$2r(80-r) + \frac{r(r-1)}{100} \;\le\; \frac{16078}{5} = 3215.6 \qquad (8 \le r \le 80), \tag{16}$$

where $r = x + y$, and therefore at most $B = 3216$ in the regimes W, O, WO and F. On a pure face carrying k spent fuel units, leakage and reverse correction have total propensity at most

$$80\varepsilon + 80^3\beta k. \tag{17}$$

Proof. The repair-only chain preserves r and creates no product, so $f = 80 - r$ throughout and $p = 0$. The suppressed channels are forward replication and forward production, each of total propensity $rf = r(80-r)$, their reverses, of total propensity at most $r(r-1)/100$ and 0 respectively, leakage, of propensity $\varepsilon r \le 80\varepsilon$, and reverse correction, of propensity at most $\beta L\,((x)_3 + (y)_3) \le \beta L\, 80^3$ — using $(x)_3 + (y)_3 \le (x+y)^3$. For the polynomial part, multiplying the claimed deficit by 100 gives $16078/5 - 2r(80-r) - r(r-1)/100 \ge 0 \iff (r-40)(199r - 8039) \ge 0$, which holds for every integer r because the two factors change sign between $r = 40$ and $r = 41$; the maximum 3215.6 is attained at $r = 40$. The largest rare-channel allowance used in this paper is that of regime F, namely $80\cdot 10^{-7} + 64\cdot 80^3\cdot 5\cdot 10^{-11} = 1.6464\cdot 10^{-3}$, well inside the remaining 0.4; regime WO needs $80\cdot 10^{-7} + 80^3\cdot 6\cdot 10^{-11} = 3.872\cdot 10^{-5}$. On a pure face both forward correction propensities vanish (each requires a molecule of each label), the product stock contributes only its own reverse, and the spent fuel is $w = k$, which gives (17). □

Theorem 4.2 (One-minority regimes)**.** *Let $k = 1$, $L = 1$, and let c_* denote c under ideal operations and c_{O} for $\theta \in [0.48, 0.52]$, $\eta \in [0.9, 1]$. Then for every $z \in \mathcal{B}_i^{1,1}$,*

$$\mathbb{P}_z(\mathcal{J}_i) \;\ge\; c_* - e^{-56\gamma\tau} - B\tau - 80\varepsilon - 80^3\beta. \tag{18}$$

In particular

$$
\begin{aligned}
q_{\mathrm{Orig}} &\geq c - \tfrac{15}{10^6} = \tfrac{6710000239829}{6710886400000} > \tfrac{4999}{5000}, && (19)\\
q_{\mathrm{W}} &\geq c - \tfrac{999}{12500000} = \tfrac{838695571135489}{838860800000000} > \tfrac{4999}{5000}, && (20)\\
q_{\mathrm{O}} &\geq c_{\mathrm{O}} - \tfrac{187}{12500000} = \tfrac{419359631961841}{419430400000000} > \tfrac{4999}{5000}, && (21)\\
q_{\mathrm{WO}} &\geq c_{\mathrm{O}} - \tfrac{999}{12500000} = \tfrac{419332385763057}{419430400000000} > \tfrac{9997}{10000}. && (22)
\end{aligned}
$$

Proof. Fix $z \in \mathcal{B}_i^{1,1}$ and suppose first $y = 1$. Because $y(y-1) = 0$, the *opposing* forward correction reaction has propensity zero, while the selected one has propensity $\gamma h\, x(x-1)y = \gamma x(x-1) \geq 56\gamma$ for $x \geq 8$ and $h = 1$. In the repair-only chain a single event therefore converts the minority, leaving a pure state with $x' = x + 1 \in [9, 80]$, $p = 0$, $w = 1$, and the chain then stops. Hence the repair-only chain fails to complete by τ with probability exactly $e^{-\gamma x(x-1)\tau} \leq e^{-56\gamma\tau}$. If $y = 0$ no repair is needed and this term is absent.

Couple the literal source to the repair-only chain on $[0, \tau]$ by common reaction clocks: run the same Poisson clocks and let the literal chain follow the comparison chain until the first firing of a suppressed clock. The two agree unless some suppressed clock fires, and the probability of that is at most the expected integrated suppressed propensity, which by (16) is at most $B\tau$. On the complementary event and on completion of the repair, the literal state at the *fixed* time τ is a pure newborn $(n, 0)$ with $8 \leq n \leq 80$ and $w \leq 1$.

Apply the Markov property at the deterministic time τ. On $[\tau, \tau + 1]$ couple the literal source to the pure core (8) by the same device; the only channels present in the literal source and absent from the core are leakage and reverse correction, whose total propensity is bounded by (17) with $k \leq 1$, so the two agree over one time unit except on an event of probability at most $80\varepsilon + 80^3\beta$. On agreement, the terminal payoff of the core is at least c_* by Proposition 3.2 and Lemma 3.1, and by construction the payoff (9) is exactly the probability that $Q_i \geq 4$, $Q_{\bar{i}} = 0$ and both refilled daughters lie in $\mathcal{B}_i^{1,1}$. Subtracting the three error events gives (18). No independence between the error events is used; they are removed by a union bound.

For the numerical instances: in regimes W, O and WO the minimal repair exponent is $56\gamma_{\min}\tau = 14$, and the finite rational inequality $\sum_{j=0}^{18} 14^j/j! > 10^6$ certifies $e^{-14} < 10^{-6}$. The remaining terms are $B\tau$, $80\varepsilon_{\max}$ and $80^3\beta_{\max}$, evaluated at the endpoints of each box, which are the worst cases because (18) is monotone in each of them. The original regime uses exponent 56 and the conservative total loss $15/10^6$. □

Remark 4.3 (The certificate frontier). The same argument proves a family statement. For core inventory K, minimum residents per daughter m, pure-core certificate c, correction rate lower bound $a = m(m-1)\gamma$, suppressed prefix bound B and core horizon t,

$$
\mathbb{P}_z(\mathcal{J}_i) \;\geq\; c - e^{-a\tau} - B\tau - \big(K\varepsilon + K^3\beta\big)t. \tag{23}
$$

The first two terms compete. For fixed $B < a$ the sum $e^{-a\tau} + B\tau$ is minimised at $\tau_* = a^{-1}\log(a/B)$, with value $(B/a)(1 + \log(a/B))$; for $a \leq B$ it is minimised at $\tau = 0$, and this proof architecture then cannot give a small loss at all. At the regime WO values $a = 1.12 \cdot 10^9$ and $B = 3216$ the optimum is $\tau_* = 1.139 \cdot 10^{-8}$ with value $3.95 \cdot 10^{-5}$; our rational choice $\tau = 1/(8 \cdot 10^7)$ gives $4.10 \cdot 10^{-5}$, so nothing material is lost by preferring a transparent deadline. The visible consequence of (23) is that high fidelity here is bought with *kinetic separation*, $a \gg B$, and that a protocol which lowered B — for instance by withholding production food until a fixed repair interval has elapsed — would relax the demand on a by the same factor. Such a staged protocol is a different model and would need its own proof; we do not claim it here.

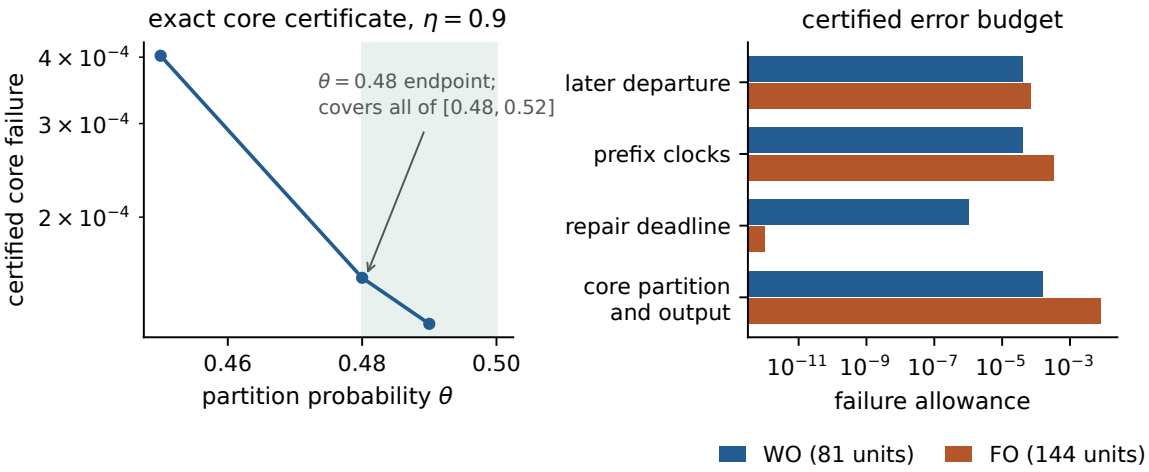

Figure 2: Left: the exact core certificate as a function of the partition probability at recovery 0.9. The theorem uses only the endpoint $\theta = 0.48$ together with Lemma 3.1; the other two points are comparisons, not a proof by scanning. Right: the four terms of the certified error budget in regimes WO and F, on a logarithmic scale. In both regimes the cost of the partition and output requirement dominates; the repair deadline, which is the term that the extreme rate constants buy, is the smallest.

5 Finite fuel, wrong consensus and spent inventory

5.1 The embedded correction chain

Let $r = x + y \geq 3$ be the resident total and let j count Y residents. In the repair-only chain the two correction propensities are

$$\text{down}: \ \gamma h\,(r-j)(r-j-1)j, \qquad \text{up}: \ \gamma h\,j(j-1)(r-j). \tag{24}$$

Their sum is $\gamma h\, j(r-j)(r-2)$, so for $0 < j < r$ and $h > 0$ the embedded jump chain has

$$p_-(j) = \frac{r-j-1}{r-2}, \qquad p_+(j) = \frac{j-1}{r-2}. \tag{25}$$

The common factor γh cancels: *the direction of correction does not depend on the catalytic rate constant, on the amount of fuel remaining, or on time.* Each jump spends one H and creates one W. (The factorisation and the normalised jump probability are machine-checked.)

Theorem 5.1 (Consensus law)**.** *For the embedded chain with unlimited steps, the probability $u_r(j)$ of reaching r before 0 is*

$$u_r(j) = 2^{-(r-3)} \sum_{a=0}^{j-2} \binom{r-3}{a} \quad (1 \leq j \leq r-1), \qquad u_r(0) = 0, \quad u_r(r) = 1, \tag{26}$$

the empty sum being zero. For a fuel inventory L, the probability $v_L(j)$ of reaching 0 within L jumps satisfies

$$v_0(j) = \mathbf{1}_{\{j=0\}}, \quad v_L(0) = 1, \quad v_L(r) = 0, \qquad v_L(j) = p_-(j)v_{L-1}(j-1) + p_+(j)v_{L-1}(j+1). \tag{27}$$

Proof. States 1 and $r-1$ can leave only towards 0 and r respectively, since $p_+(1) = 0$ and $p_-(r-1) = 0$; so $u_r(1) = 0$ and $u_r(r-1) = 1$. Every interior state has a positive finite path to an endpoint, and the state space is finite, so absorption is almost sure and u_r is the

unique bounded harmonic function with these boundary values. Writing $d_j = u_r(j) - u_r(j-1)$, harmonicity gives $d_{j+1} = \frac{r-j-1}{j-1} d_j$ for $2 \le j \le r-2$, whence $d_j = d_2 \binom{r-3}{j-2}$, and summing all differences to 1 gives $d_2 = 2^{-(r-3)}$. This is (26). First-step conditioning with absorbing boundary values gives (27). □

Remark 5.2 (Provenance). Substituting $N = r-2$ and $\kappa = j-1$ turns (25) into "increase with probability κ/N": the embedded chain is the Mabinogion urn, whose absorption probability was computed in closed form by Flajolet and Huillet [10] in an equivalent integral form, and studied further by Huillet [17] and Stenlund [32]. The tri-molecular majority network itself is that of Condon et al. [7], who record the embedded jump probability and then bound it in order to obtain Chernoff-type consensus estimates for large populations with a large initial gap; they do not use an exact absorption probability, which is what a small-count guarantee needs. So (26) should be read as the urn law transported to this network, not as a new identity. What is new here is everything that follows from running the chain with a *finite* fuel inventory: the truncation (27), the joint absorption/ spent-fuel law of §5.2, and its use as a source-level transfer.

Corollary 5.3 (Inventory rule for the correction step)**.** *From exactly two minority molecules,* $u_r(2) = 2^{-(r-3)}$. *Hence the correction step alone achieves wrong-consensus probability at most* δ*, even with unlimited fuel and unlimited speed, if and only if*

$$r \;\ge\; 3 + \lceil \log_2(1/\delta) \rceil. \tag{28}$$

For $\delta = 2 \cdot 10^{-4}$ *this forces* $r \ge 16$*; at* $r = 16$ *the wrong-consensus probability is* $1/8192$*, leaving only* $7.79 \cdot 10^{-5}$ *of the budget for every other failure mode.*

This is the quantitative reason why regime F cannot simply inherit the 0.9998 headline of the one-minority regimes: at its smallest admitted state, $r = 10$, the floor is $1/128$. It also explains why the fix is not "more fuel" and not "larger γ", but *more residents*; and since the restart region (3) demands at least 8 residents in each daughter, raising the residents required at restart would also raise the partition threshold, which is a different trade-off from the one this corollary describes. The corollary is about one subproblem of the contract; it is not an asymptotic inventory bound for the whole module.

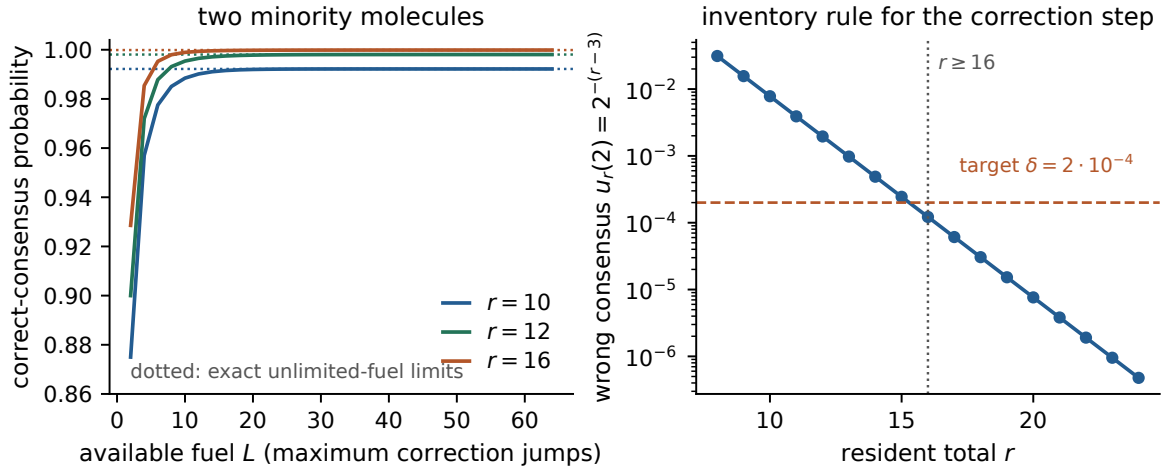

Figure 3: Left: exact correct-consensus probabilities from two minority molecules as a function of the available fuel, by the recursion (27); the dotted lines are the unlimited-fuel limits of Theorem 5.1. More fuel removes the exhaustion loss but cannot remove the wrong-consensus floor at a fixed resident total. Right: the floor $u_r(2) = 2^{-(r-3)}$ against the target $\delta = 2 \cdot 10^{-4}$, i.e. Corollary 5.3.

5.2 Charging the fuel that is actually spent

Let $a_{r,j}(k,b)$ be the probability that the embedded chain started at j is first absorbed at endpoint $b \in \{0,r\}$ on jump $k \le L$. Propagating the transient mass by (25) and removing newly absorbed mass to a table costs $O(rL)$ rational operations. Put

$$v_b = \sum_{k=0}^{L} a_{r,j}(k,b), \qquad s_b = \sum_{k=0}^{L} k\, a_{r,j}(k,b), \tag{29}$$

so that v_b is the absorption mass and s_b is the *unnormalised* first moment of the fuel spent on that event; s_b/v_b is the conditional mean when $v_b > 0$. Transient mass is retained, so absorbed plus transient mass is exactly one at every step.

Lemma 5.4 (Fuel-sensitive source comparison)**.** *Suppose the repair-only chain fails to be absorbed by the deadline τ with probability at most d, the suppressed prefix propensity is at most B, and the pure-core certificate is c_*. Then the probability of productive return to the program represented by endpoint b is at least*

$$(c_* - 80\varepsilon)\, v_b - 80^3 \beta\, s_b - B\tau - d. \tag{30}$$

Proof. On the event that absorption occurs at b after exactly k jumps and before τ, the compartment at time τ is a pure newborn carrying $w = k$ spent fuel units, so by Lemma 4.1 the departure hazard over the subsequent unit interval is at most $80\varepsilon + 80^3\beta k$ and the conditional probability of the joint event is at least $c_* - 80\varepsilon - 80^3\beta k$ by the argument of Theorem 4.2. Summing this conditional bound against the absorption table gives $\sum_k a_{r,j}(k,b)\,(c_* - 80\varepsilon - 80^3\beta k) = (c_* - 80\varepsilon)v_b - 80^3\beta s_b$. Paths not absorbed by τ are discarded at cost at most d, since a terminal probability is at most one; the omitted suppressed clocks cost at most $B\tau$ as before; and the Markov property is applied at the fixed time τ. □

The point of (30) is that it keeps the *dependence* between how the correction ended and how much fuel it burned. The coarse alternative — charging $w = L$ on every successful path, because L units were available — is what the previous version of this construction did, and it costs an order of magnitude in the final bound (Remark 5.6). Spent fuel is the state statistic that controls the downstream hazard, and it must be carried.

Theorem 5.5 (Two minority molecules)**.** *In regime F, uniformly over $\mathcal{B}_i^{64,2}$ and the stated parameter intervals, $\mathbb{P}_z(\mathcal{J}_i) \ge q_{\mathrm{F}}$ where*

$$q_{\mathrm{F}} = \frac{17691542581657548275012690808996041771237242248590086604 7}{17840596158824498513228574618118689204784332800000000000} > \frac{2479}{2500}. \tag{31}$$

Its decimal expansion is $0.9916452580485533\ldots$; with ideal operations ($\theta = 1/2$, $\eta = 1$) the same argument gives $0.9916816872003354\ldots$.

Proof. Every state of the restart region has $r - j \ge 8$ and $j \le 2$, hence $r \ge 8$ and, when $j = 2$, $r \ge 10$. While fuel remains, the total correction propensity from a transient state is $\gamma h\, j(r-j)(r-2) \ge \gamma \cdot 1 \cdot (r-2) \cdot (r-j) \ge 72\gamma$ for $r \ge 10$, $j = 2$. At most $L = 64$ jumps can occur before the fuel is exhausted, so the time to absorption or exhaustion is stochastically dominated by a sum of 64 exponentials of rate $72\gamma \ge 72 \cdot 2 \cdot 10^7$; with $\tau = 10^{-7}$ the exponent is $72 \cdot 2 \cdot 10^7 \cdot 10^{-7} = 144$ and

$$d \;\le\; e^{-144} \sum_{m=0}^{63} \frac{144^m}{m!} \;<\; 10^{-12}, \tag{32}$$

which we verify exactly by dividing the numerator polynomial by the lower Taylor sum $\sum_{m=0}^{300} 144^m/m! \le e^{144}$; the resulting rational bound is below $2.44 \cdot 10^{-14}$. A start with $j = 1$

needs one jump with exponent at least 7168, and a start with $j = 0$ needs none, so the same allowance $d = 10^{-12}$ is valid on the whole restart region.

Now apply Lemma 5.4 with $c_* = c_{\mathrm{O}}$ (legitimate for the whole operation box by Lemma 3.1), $\varepsilon = 10^{-7}$, $\beta = 5 \cdot 10^{-11}$, $B = 3216$, $\tau = 10^{-7}$ and $b = 0$, evaluated by exact rational arithmetic at each of the 216 admissible pairs (r, j) with $j \le 2$ and $r - j \ge 8$. Every value exceeds $2479/2500$, and the minimum is attained at $(r, j) = (10, 2)$, giving (31). At that state $v_0 = 0.992187499369\ldots$ and the conditional mean fuel expenditure is $s_0/v_0 = 2.36775398866\ldots$, far below the 64 units available. Successful daughters are pure and lie in the larger region $\mathcal{B}_i^{64,2}$. □

Remark 5.6 (Sharpness). The bound (31) is close to the best possible for this correction mechanism. Theorem 5.1 gives the ceiling $1 - u_{10}(2) = 127/128 = 0.9921875$ for the worst admitted state, even with unlimited fuel and unlimited catalytic speed; the exact 64-jump truncation reaches $0.9921874993\ldots$, within $6.4 \cdot 10^{-10}$ of it, and the certified source bound (31) is within $5.5 \cdot 10^{-4}$ of it. In other words, of the $8.35 \cdot 10^{-3}$ failure allowance in regime F, $7.81 \cdot 10^{-3}$ is the irreducible wrong-consensus floor of the majority mechanism at $r = 10$ and only $5.5 \cdot 10^{-4}$ is the cost of everything else the theorem covers: the exhaustion of fuel, the terminal partition and output requirement, the clock couplings, the leakage and the reverse correction. For comparison, charging every successful path the full 64 fuel units — because 64 were available — gives only $0.9900657\ldots$, i.e. it loses a further $1.58 \cdot 10^{-3}$, nearly three times the entire remaining budget.

6 Heritable variation and its state dependence

The same absorption table that certifies correct inheritance also certifies its failure mode, and the failure mode is heritable: the opposite endpoint of the correction chain is a compartment running the *other* program, which then makes the other product and transmits the other program to both of its daughters.

Corollary 6.1 (A productive program switch)**.** *In regime F, from the specified mixed preparation* $(8, 2, 70, 0, 0, 64, 0) \in \mathcal{B}_X^{64,2}$*, the probability that* $Q_Y \ge 4$*,* $Q_X \le 1$ *and both refilled daughters lie in* $\mathcal{B}_Y^{64,2}$ *is greater than* $37/5000 = 0.0074$*.*

Proof. Apply Lemma 5.4 with the same constants and $b = r = 10$. Exact recursion gives $v_{10} = 5316911553929185963121618191138170975/6805647338418769269267492148635364229 12$ and $s_{10}/v_{10} = 11.3714239539\ldots$, and (30) evaluates to $0.00748736130\ldots > 37/5000$. The daughter and product clauses come from the pure-core certificate on the *opposite* face, so this is a fully productive opposite-program pair, not a transient leakage event. □

For 128 independent founders of this preparation, the probability that at least one productive opposite pair appears is at least $1 - (1 - 37/5000)^{128} > 0.6135 > 1/2$. Independence of founders is an extra protocol assumption, used only here and not in Theorem 2.3.

Corollary 6.1 is a statement about a mixed preparation. The next theorem shows that it cannot be reinterpreted as a per-generation mutation rate, because a pure newborn behaves entirely differently — and it does so for a structural reason that is visible in the reaction table.

Theorem 6.2 (Pure-start obstruction)**.** *Let* z *be a pure newborn:* $y = 0$*,* $p_X = p_Y = 0$*,* $h = L$*,* $w = 0$ *(any state of* $\mathcal{B}_X^{L,k}$ *with no minority). Then, up to the first time a molecule of* Y *appears:*

(i) both forward correction reactions have propensity zero, since each requires a molecule of each label;

(ii) both reverse correction reactions have propensity zero, since each requires $w \ge 1$ *and* W *can be created only by a forward correction reaction;*

(iii) no other reaction creates Y or W.

Consequently the only channel that can create the first minority molecule is the leakage $X \to Y$, whose propensity is $\varepsilon x \le 80\varepsilon$, and for any deadline T,

$$\mathbb{P}_z(\text{a productive opposite-program daughter pair by } T) \;\le\; 1 - e^{-80\varepsilon T} \;\le\; 80\varepsilon T. \qquad (33)$$

In regime F, with $\varepsilon \le 10^{-7}$ and $T = 1+10^{-7}$, the right-hand side is $8.0000008 \cdot 10^{-6} < 8.1 \cdot 10^{-6}$.

Proof. Items (i)–(iii) are read off Table 1: the forward correction reactions have propensities $\gamma h\, x(x-1)y$ and $\gamma h\, y(y-1)x$, both zero when $y = 0$; the reverse ones carry the factor w, and the only reactions producing W are the forward correction reactions, which are disabled; and replication, production and their reverses change only x, f, p_X. Hence $w = 0$ and $y = 0$ persist until the leakage clock fires. A productive opposite-program pair requires $Q_Y \ge 4$, which requires at least one molecule of P_Y, which requires at least one molecule of Y at some time. The first such molecule must be created by leakage, whose propensity is at most 80ε while it is disabled elsewhere, so the probability that it fires before T is at most $1 - e^{-80\varepsilon T}$, and $1 - e^{-u} \le u$. □

Corollary 6.3 (No label-only switching kernel)**.** *In regime F the mixed preparation of Corollary 6.1 produces a productive opposite-program pair with probability greater than* 0.0074*, while a pure newborn of the same region does so with probability at most* $8.1 \cdot 10^{-6}$ *— a factor above* 935*. No Markov kernel on the two-element label set $\{X, Y\}$ can therefore describe the one-cycle offspring law of this source uniformly over its restart region. A population model must retain at least the minority count, or must use region-dependent bounds.*

Theorem 6.2 improves by a factor of 205.8 the bound we can obtain by charging, in addition to leakage, the reverse correction channel at its maximal rate $64 \cdot 80^3\beta$; that channel is simply not available to a newborn that has not yet made any waste. It is an upper bound, so it does *not* prove that recurrent mutation is zero: a positive lower bound on the switching probability from the actual distribution of pure descendants remains open, and is the missing ingredient for any statement about sustained adaptation. We regard identifying the missing state variable as the useful content of this section.

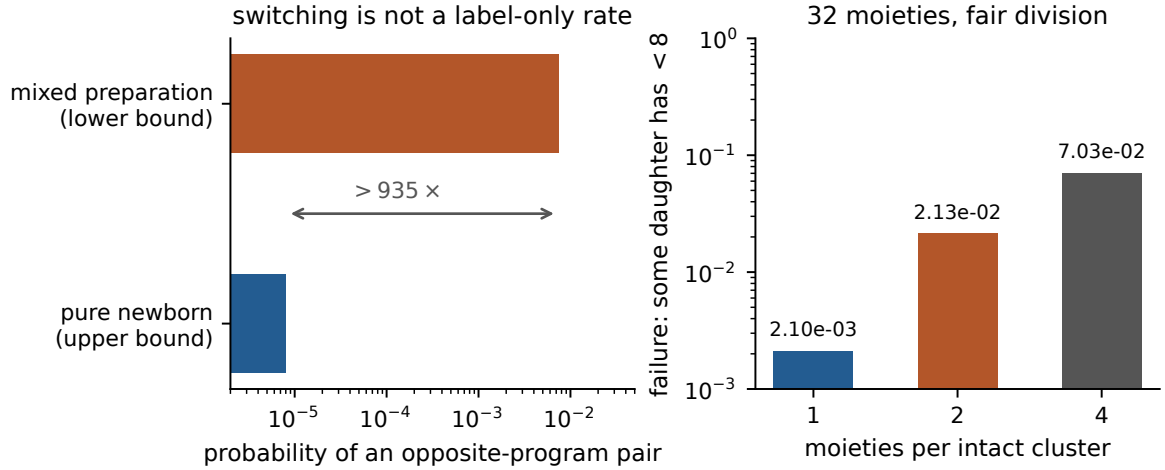

Figure 4: Left: the certified two-sided gap of Corollary 6.3. The two bars are bounds of opposite sense — a lower bound for the mixed preparation, an upper bound for the pure newborn — so the separation is a theorem, not an estimate. Right: intact clusters change the partition law even at equal total moiety count (Proposition 8.1); 32 resident moieties carried by free monomers, dimers or tetramers fail the "at least 8 in each daughter" test with very different probabilities.

7 Iteration, material and external selection

Corollary 7.1 (A preselected lineage)**.** *Fix a regime and designate daughter A, in advance of the partition, as the one that continues. For m inspected divisions the probability that all m joint events occur is at least q^m, and at least $4m$ selected product molecules are collected on that event. In particular ten divisions succeed jointly with probability greater than* 0.998 *in regimes W and O, greater than* 0.997 *in regime WO, and greater than* 0.919 *in regime F.*

Proof. Every successful continued daughter lies in the same restart region, over which Theorem 2.3 is uniform. Hence, conditional on any prior successful history, the next joint event has probability at least q; repeated conditioning gives q^m. No independence of cycles is assumed. Both daughters are checked at each division of this lineage; the statement is not about an exponentially growing family tree. (For a scheduled family tree with at most N divisions, all of them operated by the same protocol, a union bound gives failure at most $N(1-q)$; note that a complete binary tree of depth d has $2^d - 1$ divisions, not d.) □

If the terminal product stock is $p = p_X + p_Y$, the two daughter cores together hold $80 - p$ residents and food after harvesting, so refilling both to 80 supplies exactly $80 + p \le 160$ units of F. Fuel supplied is at most $2L$ and waste removal at most L per division. Table 3 collects the bill, including material that selection later discards.

Quantity (modelled material units)	One-minority regimes	Regime F
Initial material per compartment	81	144
Maximum resident moieties	80	80
Food supplied per division, both daughters	≤ 160	≤ 160
Fuel supplied per division, both daughters	≤ 2	≤ 128
Waste fuel removed per division	≤ 1	≤ 64
Unrecovered product discarded (imperfect recovery)	≤ 80	≤ 80
Initial plus ten divisions of one lineage	≤ 1701	≤ 3024
Initial plus one division for 20 founders	≤ 4860	≤ 8640

Table 3: Supply and disposal bounds. Solvent, containment, compartment construction, separation apparatus and measurement reagents are *outside* this inventory, as is the work needed to reset the volume. "81 molecules in a working device" would therefore misstate the result.

The source is label-symmetric, so selection requires an asymmetric environment. We use an external rule that reads the collected product only — never a program tag, and never the internal state.

Theorem 7.2 (A finite-population selection round)**.** *Start with ten compartments in each region of a one-minority regime. Run one batch, divide, and refill every daughter. Retain both daughters of a mother with $Q_X \ge 4$; otherwise retain each of her daughters independently with probability* 1/2*, discarding a rejected compartment whole. Then the event that every retained compartment is correctly inherited, that exactly* 20 *X-daughters remain and that between* 1 *and* 16 *Y-daughters remain has probability at least*

$$1 - \frac{20}{5000} - \frac{1 + \sum_{j=17}^{20} \binom{20}{j}}{2^{20}} = \frac{16297339}{16384000} = 0.9947106323\ldots \tag{34}$$

On this event the $X : Y$ compartment-count odds increase from 1 : 1 *to at least* 5 : 4*. Requiring at most* 15 *retained Y-daughters gives odds at least* 4 : 3 *with probability at least* 129773087/131072000*. In regime WO the same two bounds become* 16264571/16384000 *and* 129510943/131072000*.*

Proof. A union bound over the 20 mothers bounds the probability that any chemical joint event fails by $20(1-q) = 20/5000$. On the complementary event every X-mother has $Q_X \geq 4$ and so is retained with both daughters, and every Y-mother has $Q_X \leq 1 < 4$ and so has each daughter retained by an independent fair coin; all 40 daughters are correctly inherited. The number of retained Y-daughters is therefore $\mathrm{Bin}(20, 1/2)$, using coins independent of the chemical histories — *no* independence is assumed among the chemical failures themselves. Subtracting the probability of retaining no Y-daughter and the upper tail beyond 16 gives (34); the cutoff 15 corresponds to the tail beginning at 16. Substituting $1-q = 3/10000$ gives the regime WO values. All 40 daughters are replenished before retention, so the round consumes at most the 4860 units of Table 3. □

Remark 7.3 (Readout tolerance). An integer readout error of magnitude at most one can be absorbed by moving the gate to three: a successful X-mother then reads at least three and a successful Y-mother at most two, so the decision is unchanged. If that error guarantee itself fails with probability at most δ_{read} per mother, subtract $20\delta_{\mathrm{read}}$ from (34). This is a specification of the tolerance required of a detector, not a claim that counting four molecules is experimentally easy.

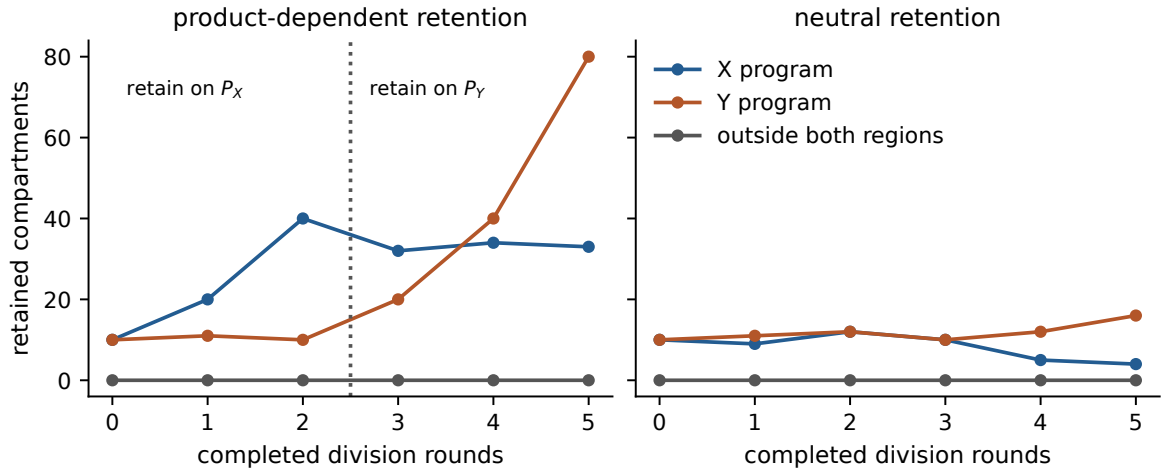

Figure 5: Illustration with the literal regime O chain, actual daughters and external retention coins: two rounds retaining on P_X, then three retaining on P_Y, against a neutral control. Failures and out-of-region states are retained in the population unless removed by the declared rule; none occurred in these particular short runs. These are seeded stochastic trajectories, *not* probability certificates; Theorem 7.2 is the proved one-round statement.

The trajectories of Figure 5 run from $(10, 10)$ to $(40, 10)$ after two P_X rounds and then to $(33, 80)$ after three P_Y rounds, against a neutral control ending at $(4, 16)$, at charged supplies of 27452 and 13158 modelled units. They illustrate response from standing variation. They do not establish a two-type mutation matrix — by Corollary 6.3 no such matrix exists uniformly — and they do not establish sustained adaptation.

8 What the certificate does and does not scale

It is worth separating the three distinct scaling questions that this construction raises, because they have different answers.

Fidelity at fixed inventory. Bounded. Corollary 5.3 shows that with k minority molecules admitted and r residents present, no correction-only route beats $1 - 2^{-(r-3)}$ for $k = 2$, whatever

the rate constants are. A fixed numerical witness therefore does not demonstrate "arbitrarily high fidelity"; that would require a parameterised family in which r, the deadline, the fuel and the partition threshold grow together while (23) stays small, and we do not exhibit one.

Kinetic demand. Explicit. Remark 4.3 isolates $a \gg B$ as the requirement, with a the correction propensity and B the competing propensity, and gives the optimal deadline. This is the quantity that an implementation must supply, and it is the quantity that a staged protocol could reduce on the B side instead of the a side.

Mechanism. Open, with a quantified obstruction. Replacing the fourth-order correction event by elementary associations requires retained intermediates, and those intermediates break the argument in three independent places. We examined one concrete candidate: for each label,

$$2X \rightleftharpoons D_X, \quad D_X + Y \rightleftharpoons C_X, \quad C_X + H \rightleftharpoons E_X, \quad E_X \rightleftharpoons D_X + Z_X, \quad Z_X \rightleftharpoons X + W, \quad (35)$$

with all forward coefficients 1, reverse dimerisation 10^{-4}, the next three reverses 10^{-5} and the last 10^{-4}. The dimer D_X is regenerated by the cycle, so the four ratios of the catalytic turnover, $10^5, 10^5, 10^5, 10^4$, multiply to exactly the thermochemical ratio $\gamma/\beta = 10^{19}$ of the abstract pair. The repair subsystem alone has 12 species; restoring food, both products, replication, production and leakage would give 15 species and 15 reversible pairs.

(a) **Startup.** The initial dimerisation propensity from $8X + 1Y + 1H$ is 56, so at the regime O deadline the probability that even the *first* association occurs is at most $56/(4 \cdot 10^8) = 1.4 \cdot 10^{-7}$. The fast-start guarantee cannot transfer to these coefficients; this is a first-event bound, not a simulation artefact.

(b) **Occupancy.** At horizons long enough for the correction to complete, most residents are bound. In 1000 literal simulations per horizon the number corrected was 0 at $2.5 \cdot 10^{-9}$, 0 at 0.1, 168 at 1 and 998 at 10, with mean bound resident moieties 0, 4.84, 8.599 and 8.003 respectively. Chemical conversion alone does not restore a productive core, because bound residents are not available to the productive reactions.

(c) **Partition.** Intact complexes are partitioned as single objects.

Proposition 8.1 (Cluster partition)**.** *Suppose the residents are carried by intact clusters of weights $w_1, \ldots, w_M$, each assigned to daughter A independently with probability θ. Then the number of resident moieties received by A has generating function $\prod_i[(1-\theta) + \theta z^{w_i}]$, which is not a binomial law in $\sum_i w_i$ unless all $w_i = 1$. In particular, at $\theta = 1/2$ the probability that each daughter receives at least one resident-bearing cluster is $1 - 2^{1-M}$, which depends only on M; and for a threshold of 8 moieties per daughter with 32 moieties in total, the success probability is $2142968775/2147483648 = 0.997898\ldots$ for 32 free monomers but only $32071/32768 = 0.978729\ldots$ for 16 intact dimers.*

Proof. Independent assignment of clusters makes the received weight a sum of independent $\{0, w_i\}$-valued variables, whose generating function is the stated product; the displayed numbers are the exact coefficient sums of that polynomial. For the first claim, each daughter fails to receive a cluster with probability 2^{-M} and the two events are disjoint for $M \geq 1$. □

So equal total moiety count does not justify the payoff (9): the certificate is not transferable merely by matching an eliminated rate expression. This is a quantitative obstruction for the tested candidate, not a universal impossibility theorem about elementary chemical computation, and it points at what a successful mechanism must supply: fast association, fast release, and enough independent resident-bearing units at the moment of division.

Remark 8.2 (Transfer budget). If a proposed implementation admits a proved coupling to the present source that fails with probability at most δ_{lift}, and if its operations differ from Definition 2.2 at cost at most δ_{op}, then a guarantee q transfers only as $q - \delta_{\mathrm{lift}} - \delta_{\mathrm{op}}$. Deriving those two numbers is the substantive task; naming an effective rate constant without a coupling or a semigroup comparison does not discharge it.

9 Dimensions, thermodynamics and scope

For a reaction of total reactant order d with count coefficient c_r at volume V litres, the falling-factorial convention corresponds to

$$k_r = \frac{c_r\,(N_A V)^{d-1}}{t_0}, \qquad t_{\mathrm{physical}} = t_0\, t_{\mathrm{model}}. \tag{36}$$

Replication and production are second order, leakage is first order, and both directions of the correction pair are fourth order, so the raw coefficients γ, β and ε do not share physical dimensions and must not be compared as numbers. Their propensities at a specified count state may be compared, and that is what Lemma 4.1 does.

Quantity	Value	Evidentiary status
Partition and recovery tolerance	$\theta \in [0.48, 0.52]$, $\eta \geq 0.9$	Proved model tolerance (regimes O, WO, F)
Illustrative volume	$V = 1$ fL	Chosen, not measured
Second-order anchor	$10^6\,\mathrm{M^{-1}s^{-1}}$	Illustrative anchor, not a calibration
Implied time unit	$t_0 = 602.214076$ s	Consequence of (36)
Startup interval τ	$1.51\,\mu$s (regime O), $7.53\,\mu$s (WO)	Consequence of the chosen anchor
Correction coefficient at $\gamma = 10^8$	$3.63 \cdot 10^{31}\,\mathrm{M^{-3}s^{-1}}$	Uncalibrated fourth-order demand
Thermochemical ratio	$W/H = \gamma/\beta$	Compatible positive equilibrium activities
Elementary mechanism	not certified	Association, occupancy and partition open (§8)

Table 4: Proved model tolerances, illustrative dimensional choices and unverified physical quantities, kept apart.

The rate table admits positive equilibrium activities $X/F = Y/F = P_X/F = P_Y/F = 100$ and $W/H = \gamma/\beta$, so it is thermodynamically consistent as a formal network. Under an ideal standard-state reading, the fuel pair carries a bias of magnitude $RT\log(\gamma/\beta)$, which is about 108.5 kJ/mol at $\gamma/\beta = 10^{19}$ and 298.15 K; across the regime WO box the ratio ranges over $[3.33 \cdot 10^{17}, 2 \cdot 10^{22}]$ and the bias over $[100.0, 127.3]$ kJ/mol. This is a standard-state ratio, not measured heat, not the dissipation of a cycle, and not a count of ATP equivalents: a physical energy account would also have to charge food conversion, separation, replenishment, volume resetting and discarded compartments [28]. Independent mathematical intervals for γ, β and ε describe a family of formal chemistries rather than three independently tunable knobs of one physical system.

The external protocol is best described as controlled batch operation. Calling it a Maxwell demon would assert a thermodynamic paradox we have not established; calling it autonomous would hide the operations of Definition 2.2, which are substantial. The honest description is that the chemistry is autonomous within a batch and the inventory operations are external, label-blind and charged.

Finally, on the relation to autocatalytic-set theory: the source does contain autocatalytic replication and catalysed production, but we make no formal RAF claim [2, 14]. Such a claim requires an explicit food set, reaction set and catalysis relation and the corresponding reflexive-autocatalytic and food-generation checks, and it matters whether the residents are treated as food, as seeds or as generated catalysts. The initial resident inventory here is supplied externally; it is not a theorem about spontaneous emergence.

10 Discussion

The construction makes one methodological point concretely: *checking that a chemical module produces the right output is weaker than checking that it produces the right output and remains usable.* The two requirements are coupled through the material inventory, and in this network the coupling is exact — copying and production have identical propensity and draw on one food pool, and correction burns a fuel molecule that is then available to drive the reverse reaction. The state statistics that turn out to be necessary are the minority count (Corollary 6.3), the spent fuel (Lemma 5.4) and, for any elementary refinement, the cluster decomposition of the residents (Proposition 8.1). None of these is visible in a concentration description.

Three questions seem to us the natural next steps.

A positive recurrent-variation rate. Theorem 6.2 bounds switching from pure descendants from above. A matching positive lower bound, over the actual reachable distribution of pure newborns across many cycles, is what a population-genetic statement would need; if it turns out to be vanishingly small, the honest conclusion is that a separate, program-blind variation phase must be designed and charged, not that adaptation has been demonstrated.

An elementary mechanism with the same contract. Section 8 gives three necessary conditions — association fast enough at the deadline, release fast enough that residents are free at division, and enough independent resident-bearing clusters — which can be checked for a candidate before any long simulation. Either a certified small elementary model or a parameter-region obstruction would be a substantial result; by Phan et al. [26] the general realizability question is hard.

A tighter formal boundary. The bridge that remains conventional here is the identification of the executed integer recurrence with the CTMC semigroup. A certificate-checking architecture in which the recurrence itself is replayed by the kernel — rather than an arithmetic theorem about its reported output — would close the gap. We do not claim that this is infeasible; we have not benchmarked it. Related efforts in proof assistants [4, 13, 22] and in probabilistic model checking [20, 21] suggest that the obstacle is engineering rather than principle.

Compared with its neighbours, the present result occupies a specific and narrow position. Condon et al. [7] analyse the same correction motif asymptotically, for large populations with a large initial gap; we need a small-count, non-asymptotic statement and therefore an exact absorption law. Kriukov et al. [19] demonstrate programmable autocatalytic functions experimentally; we supply no chemistry. Matsubara et al. [24] and Lu et al. [23] study heredity and compartment competition; we neither originate those ideas nor match their physical scope. What we contribute is a small, fully specified, reversible finite-count source for which the combined requirement — correct output *and* reusable state in both daughters *and* a bounded material bill — is certified uniformly, together with an exactly solvable correction law that says precisely which part of the guarantee is irreducible.

Data and code availability

The repository location and a permanent archive of the sources will be supplied with the author information. The exact certificate, the rational recursions, the figure generation and

an independent replay of every numerical constant printed in this paper are contained in the source package; Appendix A lists the commands.

A Evidence classes, formal scope and reproduction

We use four labels. **C**: conventional proof. **E**: exact finite computation in rational or bounded integer arithmetic. **K**: checked by the Lean 4 kernel [9, 33] with its actual premises. **N**: numerical illustration. A statement may carry several labels, and an arithmetic theorem about a reported numerator (K) is *not* a theorem that the source has that probability.

Statement	Evidence
Reaction table, mass balances, one-event repair update	C, K
Endpoint reduction (10)	C (identity), K (sign on $[0, 1/2]$, product monotonicity)
Core certificate (11)–(12)	C (soundness), E (execution), K (generic uniformization and block-iteration lemmas, accumulator bounds)
Prefix maximum (16) and budgets	C, K
e^{-56}, e^{-14} estimates	K
Erlang deadline bound (32)	C, E
CTMC clock couplings, Markov step at τ	C
Consensus law (26)	C (and classical, Remark 5.2)
Embedded jump probabilities (25), finite-fuel recursion	C, K
Absorption and spent-fuel tables, regime F bounds	C, E
Pure-start obstruction (33)	C
Selection round (34), readout separation	C, K
Cluster partition (35)ff.	C, E, N (the simulated occupancy row)
Five-round population trajectories	N
Full CTMC probability theorem in Lean	*not established*
Physical implementation or calibrated elementary mechanism	*not established*

Table 5: Claims against evidence. The probability theorems of this paper are C/E results with K components, not end-to-end formal proofs.

The Lean development lives in `proofs/TinyProgrammableChemicalFactory/` with entry points `Main.lean` and `Postproof.lean`, verified under a frozen Lean 4.30.0 toolchain with warnings as errors; the decisive declarations print only the standard axioms `Classical.choice`, `Quot.sound` and `propext`. The modules are: `Source` (the 14 channels, rate nonnegativity, vanishing of the opposing correction channel at $y \le 1$, the factor $x(x-1) \ge 56$); `Invariants` (core and fuel mass conservation for every enabled reaction, the one-event repair update and its enabling, the refill food bill); `Partition` (the payoff as a genuine one-draw binomial expectation, complementary restart, impossibility below $n < 2m$); `UniformizationLower` (a finite lower-weight, lower-iterate sum is below the Poissonised kernel; monotone block iteration preserves lower bounds); `IntegerBounds` and `RateBudget` (the e^{-56} and e^{-14} estimates, the integer prefix maximum as $15999r - 199r^2 \le 321560$, the refined budgets for regimes W and O, the accumulator inequalities); `OperationKernel` (the sign of (10) on $[0, 1/2]$); `CorrectionWalk` (the rate factorisation, the embedded down-probability, the recursion (27) and a two-step instance); `SelectionRound` (the exact tails of (34) and the readout separation); and `ConcreteWitness` (the literal repair clock and the corrected inventory for the source itself). What is *not* in Lean is the execution of the 3 240-state recurrence, the CTMC clock couplings, and the identification

of the recurrence output with the semigroup — these are the conventional bridges proved in Sections 3–5.

Everything numerical in this paper is replayed by `check_paper.py`, which recomputes the core certificates, the four regime bounds, all 216 fuel-absorption values, the switch probability, the consensus law against an exact linear solve, the selection tails, the Fano values, the cluster-partition polynomials and every decimal printed in the text, in exact arithmetic where the quantity is rational. From the package directory:

```
python check_paper.py
python make_figures.py
```

and, from the repository root, the source-level computations and the strict Lean verification:

```
python problem_workspaces/RAF_TINY_COMPUTER/experiments/replay.py
python problem_workspaces/RAF_TINY_COMPUTER/experiments/operation_certificate.py
python problem_workspaces/RAF_TINY_COMPUTER/experiments/fuel_absorption.py
python scripts/verify_proof.py (path)/Main.lean
python scripts/verify_proof.py (path)/Postproof.lean
```

where *(path)* abbreviates `proofs/TinyProgrammableChemicalFactory`.

The stochastic illustrations of Figure 5 and of Section 8(b) use a direct Gillespie simulation [11] of the literal 14-channel source with fixed seeds; they are recorded in the package under `data/` and are never used to support a probability claim.

B Exact constants

All denominators below are exact. $S = 2^{31} = 2147483648$.

Core certificates and one-minority bounds.

Symbol	Exact value
c	$2147232289/S = 0.9998829518444836\ldots$
c_{O}	$2147153442/S = 0.9998462358489633\ldots$
q_{Orig}	$6710000239829/6710886400000 = 0.9998679518444836\ldots$
q_{W}	$838695571135489/838860800000000 = 0.9998030318444836\ldots$
q_{O}	$419359631961841/419430400000000 = 0.9998312758489633\ldots$
q_{WO}	$419332385763057/419430400000000 = 0.9997663158489632\ldots$
prefix maximum	$16078/5 = 3215.6$, attained at $r = 40$
selection, cutoff 16	16297339/16384000 (W, O); 16264571/16384000 (WO)
selection, cutoff 15	129773087/131072000 (W, O); 129510943/131072000 (WO)
32 monomers, ≥ 8 each	$2142968775/2147483648 = 0.997897598\ldots$
16 dimers, ≥ 8 each	$32071/32768 = 0.978729248\ldots$

Absorption data at the worst admitted state $(r, j) = (10, 2)$, $L = 64$. The conditional mean fuel expenditures are $s_0/v_0 = 2.36775398866485\ldots$ at the correct endpoint and $s_{10}/v_{10} = 11.371423953928531\ldots$ at the opposite one.

$$v_0 = \frac{675247821429526785890499475507935979615}{680564733841876926926749214863536422912} = 0.9921874993693319\ldots,$$

$$v_{10} = \frac{531691155392918596312161819113817097 5}{680564733841876926926749214863536422912} = 0.0078124993693319\ldots,$$

$$s_0 = \frac{249815737894845679709084563344295708 91}{106338239662793269832304564822427566 08}.$$

Regime F bounds.

$$q_{\mathrm{F}} = \frac{1769154258165754827501269080899604177123724224859008660 47}{1784059615882449851322857461811868920478433280000000000000} = 0.9916452580485533\ldots,$$

$$q_{\mathrm{F}}^{\mathrm{ideal}} = \frac{3538438499888580587371891453332766396796727511986083727 19}{3568119231764899702645714923623737840956866560000000000000} = 0.9916816872003354\ldots,$$

$$q_{\mathrm{switch}} = \frac{1335789892734298336668227655567197469391357788454466047}{1784059615882449851322857461811868920478433280000000000000} = 0.0074873613013967\ldots$$

The wide fractions are simply the exact rational output of the small recursions of Section 5; they are reported so that the thresholds 2479/2500 and 37/5000 can be checked without floating-point arithmetic. The decimal displays are explanatory only — every inequality in the paper is a rational inequality.

References

[1] D. Angluin, J. Aspnes, and D. Eisenstat. A simple population protocol for fast robust approximate majority. *Distributed Computing*, 21(2):87–102, 2008. doi: 10.1007/s00446-008-0059-z.

[2] A. Blokhuis, D. Lacoste, and P. 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.

[3] L. Cardelli and A. Csikász-Nagy. The cell cycle switch computes approximate majority. *Scientific Reports*, 2:656, 2012. doi: 10.1038/srep00656.

[4] H.-L. Chen and X. Huang. Ripple: an open, AI-formalized Lean 4 framework for computing with CRNs, 2026. arXiv:2607.13531.

[5] H.-L. Chen, D. Doty, and D. Soloveichik. Deterministic function computation with chemical reaction networks. *Natural Computing*, 13(4):517–534, 2014. doi: 10.1007/s11047-013-9393-6.

[6] Y.-J. Chen, N. Dalchau, N. Srinivas, A. Phillips, L. Cardelli, D. Soloveichik, and G. Seelig. Programmable chemical controllers made from DNA. *Nature Nanotechnology*, 8(10):755–762, 2013. doi: 10.1038/nnano.2013.189.

[7] A. Condon, M. Hajiaghayi, D. Kirkpatrick, and J. Maňuch. Approximate majority analyses using tri-molecular chemical reaction networks. *Natural Computing*, 19(1):249–270, 2020. doi: 10.1007/s11047-019-09756-4.

[8] R. Cummings, D. Doty, and D. Soloveichik. Probability 1 computation with chemical reaction networks. *Natural Computing*, 15(2):245–261, 2016. doi: 10.1007/s11047-015-9501-x.

[9] L. de Moura and S. 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] P. Flajolet and T. Huillet. Analytic combinatorics of the Mabinogion urn. In *Fifth Colloquium on Mathematics and Computer Science*, volume AI of *DMTCS Proceedings*, pages 549–572, 2008. doi: 10.46298/dmtcs.3591.

[11] D. T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *The Journal of Physical Chemistry*, 81(25):2340–2361, 1977. doi: 10.1021/j100540a008.

[12] W. 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] J. Hölzl. Markov chains and Markov decision processes in Isabelle/HOL. *Journal of Automated Reasoning*, 59(3):345–387, 2017. doi: 10.1007/s10817-016-9401-5.

[14] W. Hordijk and M. 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] D. Huh and J. Paulsson. Non-genetic heterogeneity from stochastic partitioning at cell division. *Nature Genetics*, 43(2):95–100, 2011. doi: 10.1038/ng.729.

[16] D. Huh and J. Paulsson. Random partitioning of molecules at cell division. *Proceedings of the National Academy of Sciences*, 108(36):15004–15009, 2011. doi: 10.1073/pnas.1013171108.

[17] T. Huillet. Reversing the drift of the Ehrenfest urn model and three conditionings. *Journal of Physics A: Mathematical and Theoretical*, 42(34):345005, 2009. doi: 10.1088/1751-8113/42/34/345005.

[18] A. Jensen. Markoff chains as an aid in the study of Markoff processes. *Skandinavisk Aktuarietidskrift*, 1953:87–91, 1953. doi: 10.1080/03461238.1953.10419459.

[19] D. V. Kriukov, J. Huskens, and A. S. Y. Wong. Exploring the programmability of autocatalytic chemical reaction networks. *Nature Communications*, 15:8289, 2024. doi: 10.1038/s41467-024-52649-z.

[20] M. Kwiatkowska, G. Norman, and D. Parker. PRISM 4.0: verification of probabilistic real-time systems. In *Computer Aided Verification (CAV 2011)*, volume 6806 of *Lecture Notes in Computer Science*, pages 585–591. Springer, 2011. doi: 10.1007/978-3-642-22110-1_47.

[21] M. R. Lakin, D. Parker, L. Cardelli, M. Kwiatkowska, and A. Phillips. Design and analysis of DNA strand displacement devices using probabilistic model checking. *Journal of the Royal Society Interface*, 9(72):1470–1485, 2012. doi: 10.1098/rsif.2011.0800.

[22] J. I. Lathrop, J. H. Lutz, R. R. Lutz, H. D. Potter, and M. R. Riley. Population-induced phase transitions and the verification of chemical reaction networks. *Natural Computing*, 23(2):347–363, 2024. doi: 10.1007/s11047-021-09877-9.

[23] H. Lu, A. Blokhuis, R. Turk-MacLeod, J. Karuppusamy, A. Franconi, G. Woronoff, C. Jeancolas, A. Abrishamkar, E. Loire, F. Ferrage, P. Pelupessy, L. Jullien, E. Szathmáry, P. Nghe, and A. D. Griffiths. Small-molecule autocatalysis drives compartment growth, competition and reproduction. *Nature Chemistry*, 16(1):70–78, 2024. doi: 10.1038/s41557-023-01276-0.

[24] Y. J. Matsubara, S. Ameta, S. Thutupalli, P. Nghe, and S. Krishna. Conditions for Darwinian evolution in compartmentalized autocatalytic reaction networks, 2026. arXiv:2211.03155v2.

[25] S. Matsumura, Á. Kun, M. Ryckelynck, F. Coldren, A. Szilágyi, F. Jossinet, C. Rick, P. Nghe, E. Szathmáry, and A. D. Griffiths. Transient compartmentalization of RNA replicators prevents extinction due to parasites. *Science*, 354(6317):1293–1296, 2016. doi: 10.1126/science.aag1582.

[26] T.-M. Phan, D. Merkle, T.-L. Phan, C. Flamm, and P. F. Stadler. On the realizability of abstract reaction networks with real molecules and reactions, 2026. arXiv:2608.04635.

[27] L. Qian and E. Winfree. Scaling up digital circuit computation with DNA strand displacement cascades. *Science*, 332(6034):1196–1201, 2011. doi: 10.1126/science.1200520.

[28] R. Rao and M. Esposito. Nonequilibrium thermodynamics of chemical reaction networks: wisdom from stochastic thermodynamics. *Physical Review X*, 6(4):041064, 2016. doi: 10.1103/PhysRevX.6.041064.

[29] D. Segré, D. Ben-Eli, and D. Lancet. Compositional genomes: prebiotic information transfer in mutually catalytic noncovalent assemblies. *Proceedings of the National Academy of Sciences*, 97(8):4112–4117, 2000. doi: 10.1073/pnas.97.8.4112.

[30] D. Soloveichik, M. Cook, E. Winfree, and J. Bruck. Computation with finite stochastic chemical reaction networks. *Natural Computing*, 7(4):615–633, 2008. doi: 10.1007/s11047-008-9067-y.

[31] D. Soloveichik, G. Seelig, and E. Winfree. DNA as a universal substrate for chemical kinetics. *Proceedings of the National Academy of Sciences*, 107(12):5393–5398, 2010. doi: 10.1073/pnas.0909380107.

[32] D. Stenlund. On the Mabinogion urn model. *Advances in Applied Probability*, 50(2): 327–346, 2018. doi: 10.1017/apr.2018.16.

[33] 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.