The material and kinetic cost of productive chemical memory: exact frontiers for support and proportion encodings

17 September 2026

Abstract

A chemical system that is to be called heritable must do two things on one probability law: make fresh product before a deadline, and hand its own identity to *both* of the compartments it divides into. We fix a single finite protocol — run a batch of K material units to time T, harvest all product, allocate every intact molecule by one complementary coin, refill both daughters — and ask what that joint event costs in material, for two ways of encoding the stored identity.

For *support* memory, where the identity is which of two species is present, we study an explicit reversible binding network with elementary bimolecular steps. The least material budget that guarantees a quota of four fresh product units and two memory-bearing daughters with probability 0.999, uniformly over all 287 admitted preparations and over an entire factor-two band of release rates, is exactly 32. The certificate is a rowwise minimum of endpoint transition kernels evaluated in exact integer arithmetic; it needs no monotonicity in budget, deadline or rate. We then enlarge the uncertainty: the same minimum 32 survives allocation bias $\theta \in [0.39, 0.61]$ and a simultaneous independent multiplicative perturbation of every one of the six rate constants by $\pm 1/300$. We also exhibit a direction in which it genuinely breaks. Untying the two exit channels of the bound complex — so that dissociation and product release no longer share one rate — excludes $K = 32$ by an exact upper certificate and raises the exact minimum to 42. The guarantee is *not* monotone in the product-release rate: doubling release alone also breaks it, because the dissociation channel is what returns food to the pool that replication needs to build carriers.

For *proportion* memory, where both species are always present and the identity is the majority, no elementary network is known. We give a cooperative gate of reaction order 229 using 415 units, and show that the gate's reachable class collapses to two states, which makes the whole architecture class exactly optimisable. Sharpening the partition estimate from a union bound to an exact endpoint minimum raises its guarantee from 0.9995 to 0.9997 per cycle and from 0.995 to 0.997 over ten. More importantly, we compute the class's exact frontier: 253 core units suffice at the same specification and no member of the class does better, so that construction is not optimal. At the support-memory specification the frontier is 230 core units. Because a formal gate may use unbounded reaction order and unbounded rate constants, this is a lower bound on every such realisation: **proportion memory costs at least 231 units where support memory costs 32.** Across five decades the frontier obeys $N \approx 33.3 \ln(1/\delta)$; we prove $N = q + \Theta(\log(1/\delta))$ and identify the balance law behind the constant, in which the corridors of m ordered levels have costs $4k^2z^2$ and the total grows cubically in m.

Concrete minima are conventional exact computer-assisted theorems; the kernel envelope, source conservation, rate bounds, the complementary-binomial event and the exact tail arithmetic are verified in Lean 4 against Mathlib. This is not an end-to-end formalisation of either probability theorem, and neither network is an experimentally calibrated chemistry.

Keywords. chemical heredity; finite-count stochastic reaction networks; random partitioning; robust Markov models; exact certificates; formal verification.

MSC 2020. 60J28, 92C42, 92E20, 68V20, 60J80.

Contents

1 Introduction

1.1 Producing and inheriting are one event

A chemical system is often called heritable when its state persists, and called productive when it converts food into something useful. Either property alone is comparatively cheap. A catalyst that is never consumed persists trivially; a reaction that consumes food makes product trivially. The interesting object is a system that does both *at once, out of the same finite material*, and then divides, and is still both afterwards — in both halves.

That conjunction is what we charge for here. The requirements interfere. Food spent on product is food not spent on carriers, and carriers are what the division stage needs, because a daughter that receives no carrier has lost the program whatever the batch produced. Conversely a preparation rich in carriers is poor in food and may miss its output deadline. The tension is not visible in any calculation that treats production and inheritance separately and multiplies the answers, and it is not visible in a stationary calculation at all, because the deadline is what makes it bite.

1.2 Two encodings of the same bit

Fix the simplest nontrivial memory: one bit, two programs. There are two qualitatively different ways for finite chemistry to hold it.

Support memory. Species X is present and Y is absent, or the reverse. The decoder is a presence test. Absence is dynamically stable here because nothing in the network creates Y from an X-only state, so the bit cannot flip; what can fail is that a daughter receives no molecule at all.

Proportion memory. Both X and Y are always present and the bit is which is in the majority — whether $x/(x+y) > 1/2$. This is the compositional form of heredity studied since [14], and it is strictly harder in a specific finite-count sense: a daughter must receive a count *inside a corridor*, not merely a nonzero count, and the corridor must be re-entered after every division.

The two encodings are usually compared rhetorically. We compare them numerically, on one protocol, at one specification, and the gap turns out to be large and to have an identifiable source.

1.3 Contributions

1. **A common framework** (§2) in which a batch, its deadline, its harvest, one complementary allocation, and the refill that returns both daughters to the restart set are a single probability law with a single material ledger, and repetition along a lineage is an inequality for a substochastic kernel rather than an independence assumption.

2. **Encoding-independent obstructions** (§3). A carrier ceiling $K \geq q + m\lceil 1 + \log_2(1/\delta)\rceil$ for m memory-bearing types; a first-enabling clock bound $p_* \leq 1 - e^{-a(z)T}$ showing that material cannot buy time, with the corollary $k\,K^r T \geq \log(1/\delta)$ that makes reaction order and material explicit substitutes; and, for corridor-based memory, the necessity of $l \leq g \leq u$.

3. **The support-memory frontier** (§4). The exact robust minimum 32, reproduced here from an independent implementation (Theorem 4.2); enlarged robustness to allocation bias and to a simultaneous rate-constant box (§4.5); and a genuine failure direction — untying the complex's two exit channels raises the exact minimum to 42 (Theorem 4.10) and makes the guarantee non-monotone in the product-release rate, for a reason we identify (§4.6).

4. **The proportion-memory frontier** (§5). An endpoint reduction lemma (Lemma 5.5) that replaces a 7 467-state partition computation by four binomial evaluations; sharpened constants for $\mathcal{R}$; and the exact minimum over the whole cooperative-gate class — 253 core units at that construction's own specification (Theorem 5.8), together with its reliability scaling and the cascade law behind it (§5.6).

5. **The comparison** (§6). At one specification, support memory is achieved with 32 units by elementary bimolecular chemistry, while proportion memory needs at least 231 from *every* member of the cooperative-gate class, however large its reaction order and rate constants.

1.4 Evidence classes

We label every assertion.

- **C** a conventional mathematical proof, written out here;
- **E** exact finite arithmetic over $\mathbb{Z}$ or $\mathbb{Q}$, executed by a program whose soundness is proved in **C**;
- **K** a statement compiled in Lean 4 against Mathlib, with a strict receipt (exit zero, warnings as errors, axiom report);
- **N** a floating-point diagnostic or a simulation. No **N** computation is an input to any theorem below.

Both concrete frontiers are **C/E**. The kernel-envelope inequalities, source conservation, the daughter-return and refill identities, the rate bounds, the complementary-binomial event identity and the exact endpoint tail arithmetic are **K**; §7 gives the declaration map and says precisely where the formal boundary lies. We do not claim an end-to-end formal proof of either probability theorem.

1.5 What is not claimed

Neither network is calibrated chemistry: rate constants and the deadline are constructed benchmark values, and $T = 5$ is in model time units, not seconds. Harvest, allocation, refill and volume reset are ideal external operations; the apparatus, the solvent, and the work of separation are excluded from the molecular ledger, and one fuel token is a material count, not a quantity of chemical work. The frontiers are minima *within declared architecture classes*, not universal laws about chemical memory: correlated partitioning, multi-carrier complexes, post-division repair, other encodings and external supplies all change the hypotheses. We make no priority claim for chemical heredity, for programmable autocatalysis, or for robust Markov operators; §1.6 states the relation to the established literature on each.

1.6 Relation to prior work

Random partitioning of molecules at division, including the effect of complexes that partition as units, is established prior art [8]. Resource-limited cellular sensing already separates copies, time and fuel as distinct currencies [6]; our event is productive copying into two actual daughters rather than sensing precision or energy per bit. Transient loading of a module by its downstream load is the subject of retroactivity theory [3]; the load here is four units of food that become product and therefore cannot become carriers. Compositional heredity and selection in compartmentalised autocatalytic networks are studied in [14, 12], programmable autocatalytic responses in [10], and stochastic majority computation by low-order chemical reaction networks in [2] — the last is exactly the elementary mechanism that our proportion-memory gate does *not* provide. Structural bounds in maximal-flux regimes [4] and product-form stationary laws [1] concern hypotheses different from a finite-time chance constraint.

Methodologically, robust sets of finite-state continuous-time Markov models with partially specified parameters, lower transition operators, and numerical schemes with guaranteed error bounds are a developed subject [9, 5]. The envelope of §4.3 is a specialisation of that idea to one explicit chemical source, combined with one-sided integer arithmetic; we claim the source-bound frontier and its transparent certificate, not the operator idea.

2 The common framework

2.1 Batch, deadline, harvest, allocation, refill

Definition 2.1 (Protocol)**.** A *cycle* consists of the following five ideal operations.

(P1) **Batch.** A finite continuous-time Markov chain with generator G on a finite state space determined by a conserved material pool of K units runs from an admitted restart state z for time T.

(P2) **Harvest.** All terminal product is collected and removed.

(P3) **Allocation.** Every remaining intact molecule is assigned to daughter A independently with probability θ (the fair case is $\theta = 1/2$); daughter B receives the complement of that *same* draw. Bound complexes are allocated as single objects.

(P4) **Refill.** Each daughter's food pool is topped up to the core mass and any consumed fuel token is replaced.

(P5) **Restart.** Both daughters are required to lie in the restart set; one of them, designated before any randomness, continues.

No carrier is added, removed, sorted or selected, and there is no post-division repair. The operation is the same for every state and every program, including after failure.

Write $\mathcal{B}$ for the restart set and q for the product quota. The *joint success payoff* at a terminal state is

$$h_q(z) = \mathbf{1}\{\,\text{product} \geq q\,\} \cdot s_\theta(z), \tag{1}$$

where $s_\theta(z)$ is the conditional probability that the allocation (P3) returns both daughters to $\mathcal{B}$. The exact one-cycle probability from z is $(e^{TG}h_q)(z)$, and the quantity to be certified is

$$p_* = \min_{z \in \mathcal{B}} \,(e^{TG}h_q)(z), \tag{2}$$

the worst admitted preparation. Note that (1) is a *single* payoff on a single law: output and inheritance are never computed as separate marginals and multiplied.

2.2 The two payoffs used here

Both encodings instantiate s_θ explicitly.

For support memory, the only way a daughter fails is to receive no carrier-bearing object at all. Conditional on $R \geq 1$ such objects, the two all-to-one-daughter events are disjoint, so

$$s_\theta(R) = 1 - \theta^R - (1-\theta)^R, \qquad s_{1/2}(R) = 1 - 2^{1-R}. \tag{3}$$

For proportion memory with corridors $[l, u]$ per type, a daughter must receive a count *inside* the corridor. With n objects of one type and $I \sim \mathrm{Bin}(n, \theta)$ the daughters receive I and $n - I$, so the return probability of that type is

$$\psi_\theta(n; l, u) = \Pr\{\, l \leq I \leq u \ \text{ and } \ l \leq n - I \leq u \,\}, \tag{4}$$

and s_θ is the product of (4) over the types. The product is across the disjoint allocation objects of different types, never across daughters: (4) retains the exact dependence between the two daughters.

2.3 The material ledger

Material is conserved inside a batch and supplied only at (P4).

Proposition 2.2 (Exact ledger, C)**.** *Let a cycle start from core mass K and harvest p product units. The two daughters together retain $K - p$ core units, so refilling both to K supplies exactly $K + p$ new units. Along a preselected lineage of G inspected cycles in which both daughters are prepared every time,*

$$M_G = K + \sum_{j=1}^{G} (K + p_j). \tag{5}$$

Proof. Allocation is complementary, so the daughters' retained core masses a, b satisfy $a + b = K - p$. Refilling each to K adds $(K - a) + (K - b) = 2K - (K - p) = K + p$. Each daughter's retained mass is at most K, so every refill amount is nonnegative. Summing over cycles and counting the initial preparation once gives (5). [`actual_refill_account`, `refill_identity`] □

The working inventory K is therefore not the lifetime cost. At $K = 32$, $q = 4$ and $G = 10$ the lineage collects at least 40 product units and consumes at least 392 supplied units, the unoperated siblings' prepared inventory included.

2.4 Repetition without independence

Proposition 2.3 (Lineage iteration, C)**.** *Let $\mathcal{K}$ be the substochastic kernel on $\mathcal{B}$ that carries mass to the designated daughter only on full joint success (quota met, both daughters returned). If every row sum of $\mathcal{K}$ is at least p_*, then $(\mathcal{K}^G \mathbf{1})(z) \geq p_*^G$ for every $z \in \mathcal{B}$.*

Proof. Induction on G using positivity: $\mathcal{K}^{G+1}\mathbf{1} = \mathcal{K}(\mathcal{K}^G \mathbf{1}) \geq \mathcal{K}(p_*^G \mathbf{1}) = p_*^G \mathcal{K}\mathbf{1} \geq p_*^{G+1}\mathbf{1}$. [`repeated_conditional_success`] □

No independence between cycles is assumed, and the bound remains valid if the uncertain parameters take different admitted values in different cycles, because p_* is a bound uniform over the whole parameter set. What is bounded is the designated lineage, not an entire binary family tree, whose event and resource bill are different.

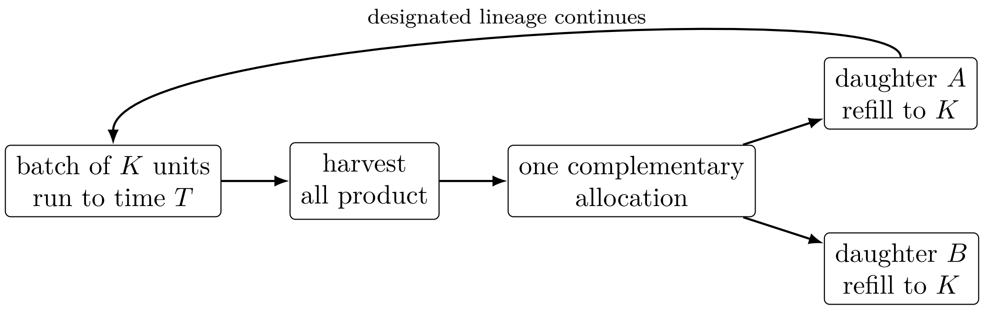

Figure 1: The protocol of Definition 2.1. Success is the single event "quota met *and* both daughters back in the restart set". The same picture serves both encodings; only the restart set and the return payoff s_θ change.

3 Encoding-independent obstructions

Three bounds hold for every source in the framework. They are necessary conditions, never sufficient: a permitted terminal allocation is not a proof that initially product-free chemistry reaches it by the deadline.

3.1 The carrier ceiling

Theorem 3.1 (Carrier ceiling, C)**.** *Suppose the memory is carried by m species, each of which must be present in both daughters, all carriers are unit-weight and allocated independently and fairly, and a quota of q product units is harvested from the same conserved pool of K units. If the joint success probability is at least $1-\delta$ for some $0<\delta<1$, then*

$$K \;\geq\; q + m\,\lceil 1 + \log_2(1/\delta)\rceil\,. \tag{6}$$

Proof. On the quota event the terminal carrier counts $r_1,\dots,r_m$ satisfy $\sum_j r_j \leq K-q$, so $\min_j r_j \leq \lfloor (K-q)/m\rfloor =: r$. Both daughters contain a given species with r copies with probability at most $1-2^{1-r}$ (zero if $r=0$), and the joint event is contained in that one. Averaging over terminal states preserves the inequality, so $1-\delta \leq 1-2^{1-r}$, i.e. $2^{1-r}\leq\delta$, i.e. $r \geq 1+\log_2(1/\delta)$. Since r is an integer, $\lfloor (K-q)/m\rfloor \geq \lceil 1+\log_2(1/\delta)\rceil$, which is (6). [`quota_limits_carriers`, `reliability_requires_carriers`, `material_necessary`] □

At $q=4$ and $\delta=10^{-3}$ this excludes $K\leq 14$ for support memory ($m=1$; the relevant ceiling at $K=14$ is $511/512 < 999/1000$) and $K \leq 25$ for two-type memory. The bound is a ceiling within the stated intact-object partition architecture, not a universal law.

3.2 Material cannot buy time

Theorem 3.2 (First-enabling obstruction, C)**.** *Let z be an admitted restart state at which the total enabled exit rate is $a(z)$, and suppose a positive quota requires at least one reaction to fire. Then for every material budget,*

$$p_* \;\leq\; 1 - e^{-a(z)T}. \tag{7}$$

Consequently a uniform target $1-\delta$ requires $a(z)\,T \geq \log(1/\delta)$ at every admitted restart state.

Proof. The first holding time at z is exponential with rate $a(z)$, so with probability $e^{-a(z)T}$ the chain has not moved by T and the terminal product is zero. $\square$

This is the sharpest structural statement in the paper, because it is completely insensitive to K: no amount of added food removes it. For the support-memory source of §4 the admitted bound seed $(n, c, p, f) = (0, 1, 0, K - 2)$ has exactly two enabled exits, each of rate κ, whence $p_* \leq 1 - e^{-2\kappa T}$; at $T = 5$ and $\delta = 10^{-3}$ every release rate $\kappa < \log(1000)/10 \simeq 0.690776$ is infeasible at every budget. `[bound_seed_admitted, bound_seed_rates]`

Corollary 3.3 (Order and material are substitutes, C)**.** *In stochastic mass action at fixed reference volume, a channel of molecularity r with rate constant k has propensity at most $k\,K^r$ on a pool of K units. Hence a uniform target $1 - \delta$ requires*

$$k\,K^r\,T \;\geq\; \log(1/\delta). \tag{8}$$

Inequality (8) is the quantitative form of the trade-off this paper is about. A network that is denied large k and large r — an elementary network with physically bounded constants — must buy its clock with K. A formal network that is allowed unbounded r and unbounded k gets its clock for nothing, and then pays for something else. Sections 4 and 5 are the two ends of that trade.

3.3 Corridors must contain their own increment

Lemma 3.4 (Corridor necessity, C)**.** *Consider a coordinate whose newborn count lies anywhere in $[l, u]$ and which gains a fixed increment g before division, so that the parent count ranges over $[l + g, u + g]$, and suppose both daughters must return to $[l, u]$ with positive probability uniformly over newborns. Then $l \leq g \leq u$.*

Proof. From the newborn with count l the parent holds $l + g$ objects and the daughters receive I and $l + g - I$; both are at least l only if $l \leq I \leq g$, an empty range when $g < l$. From the newborn with count u the parent holds $u + g$ and both daughters are at most u only if $g \leq I \leq u$, empty when $g > u$. $\square$

So a corridor memory must grow by an amount lying inside its own corridor. Support memory has no analogue of this constraint, and §5 shows it is what makes the difference in cost.

4 Support memory at the elementary frontier

4.1 The source

Fix the label $S \in \{X, Y\}$ and take the three reversible pairs

$$S + F \rightleftharpoons 2S, \qquad S + F \rightleftharpoons C_S, \qquad C_S \rightleftharpoons S + P_S, \tag{9}$$

with forward/reverse constants $(1, 1/100)$, $(2, \kappa)$ and $(\kappa, 1/50)$. Only the selected label is initially present, and no channel creates the other label, so its absence is invariant: this is support-encoded memory, and contamination correction is outside the theorem. Every reaction in (9) is at most bimolecular.

At fixed reference volume write $z = (n, c, p)$ for the counts of free carrier, complex and product, with food reconstructed as $f = K - n - 2c - p$. The state space and restart set are

$$\Omega_K = \{n + 2c + p \le K,\ n + c \ge 1\}, \qquad \mathcal{B}_K = \{(n, c, 0) \in \Omega_K : 1 \le n + c \le K - 1\}, \tag{10}$$

and the six destination increments and propensities are

$$\begin{array}{lll} (+1,0,0)\ \ nf, & (-1,0,0)\ \ n(n-1)/100, & (-1,+1,0)\ \ 2nf, \\ (+1,-1,0)\ \ \kappa c, & (+1,-1,+1)\ \ \kappa c, & (-1,+1,-1)\ \ np/50. \end{array} \tag{11}$$

Zero-rate channels are omitted; in particular there is no extra factor $1/2$ in reverse replication. A complex carries one carrier moiety and one food-derived moiety, so the weighted material invariant is $n + 2c + p + f = K$ while the number of intact carrier-bearing partition objects is $R = n + c$. At $K = 32$, $|\Omega_{32}| = 3\,384$ and $|\mathcal{B}_{32}| = 287$.

Remark 4.1 (No batch extinction). Only reverse replication lowers R, and its propensity $n(n-1)/100$ vanishes at $n \le 1$; every other channel preserves R. The last carrier therefore cannot disappear during a batch, and the relevant failures are exactly the two in (1): missing the quota, or losing a daughter at the allocation. [`jump_conserves`, `final_resident_survives`]

With payoff $h_q(z) = \mathbf{1}\{p \ge q\}(1 - 2^{1-(n+c)})$ define

$$p_*(K, \kappa) = \min_{z \in \mathcal{B}_K} (e^{5G_{K,\kappa}} h_4)(z), \qquad p_{\mathrm{rob}}(K) = \inf_{10 \le \kappa \le 20} p_*(K, \kappa). \tag{12}$$

The generator acts on a column payoff; no transpose is used.

4.2 The robust exact minimum

Theorem 4.2 (Robust exact minimum, C/E)**.** *The least integer $K \ge 2$ with $p_{\mathrm{rob}}(K) \ge 999/1000$ is exactly* 32*. In fact*

$$p_{\mathrm{rob}}(32) \ \ge\ L := \frac{2\,145\,514\,420}{2^{31}} = 0.999083006754\ldots \tag{13}$$

Every $K \le 14$ is excluded by Theorem 3.1 and every $15 \le K \le 31$ by an exact upper certificate at the admitted constant parameter $\kappa = 10$, the largest being $U_{31} = 2\,144\,298\,834/2^{31} = 0.998516955\ldots$ Consequently the same minimum holds for every reliability target ρ with $U_{31} < \rho \le L$.

The proof occupies §§4.3–4.4; the exclusion table is Table 1 and the whole computation is replayed by `scripts/frontier_support.py`. Three features deserve emphasis. The statement is a *robust* minimum: it does not say that every $K < 32$ fails at every release rate, only that no $K < 32$ guarantees the target across the whole band. No monotonicity in K, T or κ is used or proved — the finite list is checked budget by budget. And it is not asserted that every $K > 32$ succeeds.

Remark 4.3 (Where the 17 units go). The ceiling (6) permits $K = 15$; the source requires 32. The 17-unit difference is the price of the kinetics — the deadline, the rate band, the breadth of the restart set, product variability, residual food and complex occupancy together. It is not a measured "sequestration cost" and we do not allocate it to any single mechanism. What we *can* exclude is that it is certificate slack: at $K = 32$ the certified lower value falls about 1.05×10^{-5} below the floating worst-start value, whereas $K = 31$ fails by an exact *upper* bound, so no sharpening of the lower-bound algorithm could rescue it.

4.3 The endpoint-kernel envelope

The band $[10, 20]$ is a continuum, and the terminal probability is not known to be monotone in κ. The certificate avoids both difficulties by applying convexity to the one-step *kernel* and never to a deadline probability.

Choose an integer clock λ at least as large as every exit rate across the band. At $K = 32$ the exit rate $3nf + n(n-1)/100 + 2\kappa c + np/50$ is affine and increasing in κ, so it is largest at $\kappa = 20$, where its maximum over Ω_{32} is 770.4, attained at $(n, c, p) = (16, 0, 0)$ — a state holding no complex, so the value is the same at both endpoints. Hence $\lambda = 771$ is valid throughout the band. Put $P_\kappa = I + G_\kappa/\lambda$. Each P_κ is stochastic, and because the generator is affine in κ,

$$P_\kappa = (1-t)P_{10} + tP_{20}, \qquad t = (\kappa - 10)/10 \in [0, 1]. \tag{14}$$

[`release_affine`]

Lemma 4.4 (Envelope, C; K)**.** *Define the nonlinear operator* $(Mv)(z) = \min\{(P_{10}v)(z), (P_{20}v)(z)\}$*. Then M is monotone, $Mv \le P_\kappa v$ for every κ in the band, and $M^j v \le P_\kappa^j v$ for every $j \ge 0$. For $v \ge 0$ and every $J \in \mathbb{N}$,*

$$e^{-1}\sum_{j=0}^{J}\frac{M^j v}{j!} \;\le\; e^{-1}\sum_{j=0}^{\infty}\frac{P_\kappa^j v}{j!} = e^{G_\kappa/\lambda}v. \tag{15}$$

Proof. Both P_{10} and P_{20} are nonnegative, so each is monotone and hence so is their pointwise minimum. The scalar inequality $\min(a, b) \le (1-t)a + tb$ for $t \in [0, 1]$ together with (14) gives $Mv \le P_\kappa v$ entrywise. The iterate bound follows by induction: $M^{j+1}v = M(M^j v) \le M(P_\kappa^j v) \le P_\kappa(P_\kappa^j v)$, the first step by monotonicity of M and the second by the previous sentence. Finally the omitted terms of the Poisson sum are nonnegative because $v \ge 0$ and both kernels are nonnegative. [`envelope_le_kernel`, `envelope_iterates_lower`, `envelope_nonneg`, `robust_poisson_lower`] ☐

Composing λT such blocks preserves the lower bound for every fixed κ in the band, which is all Theorem 4.2 requires. The rowwise minimum permits a different adverse endpoint at each state and each iteration depth, so the certificate is in general *conservative* relative to any single fixed κ; we make no claim that it equals the fixed-parameter infimum. What must not be done — and is not done here — is to interchange minimisation with matrix exponentiation.

Remark 4.5 (Reusable form). Nothing above is specific to chemistry or to two endpoints. If a family of generators lies in the convex hull of $G_1, \ldots, G_m$ and one common clock is valid for all of them, then $Mv = \min_i P_i v$ certifies every generator in the hull. A parameter box entered affinely lies in the hull of its vertex generators; nonlinear parameter dependence needs a separate enclosure. This is how §4.5 certifies a six-dimensional rate box.

Lemma 4.6 (Complement, C; K)**.** *Because the full stochastic semigroup preserves constants,* $e^{TG}h = 1 - e^{TG}(1-h)$*. Hence a lower certificate ℓ for the payoff $1-h$ yields the upper certificate* $e^{TG}h \le 1 - \ell$*.* [`complement_upper`]

Lemma 4.6 is what makes the exclusions in Theorem 4.2 genuine: a lower bound falling below the target proves nothing, whereas an upper bound below the target at one admitted start excludes the budget outright. One admitted witness suffices to exclude; inclusion requires all 287 starts.

Table 1: Exact two-sided certificates at $\kappa = 10$, $T = 5$, $q = 4$; denominator $2^{31} = 2\,147\,483\,648$ throughout. Every excluded row carries the admitted upper witness $(K-1, 0, 0)$; the included row covers all 287 starts. The largest excluded upper value is at $K = 31$. Appendix B records an independent reimplementation that reproduces every entry.

K	states	lower	upper	K	states	lower	upper
15	428	1015891063	1015896977	24	1522	2072944721	2072959911
16	508	1220976456	1220983211	25	1703	2097931806	2097948254
17	597	1410294481	1410302076	26	1898	2115061624	2115079437
18	696	1576864428	1576872967	27	2107	2126576039	2126595217
19	805	1717234717	1717244202	28	2331	2134178444	2134199094
20	925	1831016096	1831026631	29	2570	2139116888	2139139042
21	1056	1920060288	1920071907	30	2825	2142277641	2142301368
22	1199	1987562186	1987574961	31	3096	2144273535	2144298834
23	1354	2037277175	2037291104	**32**	**3384**	**2145518355**	**2145545336**

4.4 Exact integer implementation

Let $S = 2^{31}$, $D = 100\lambda = 77\,100$, and let $A_i = DP_i$ for the two endpoints; every entry of A_i is a nonnegative integer and every row sums to D. Let

$$e_- = \sum_{j=0}^{25} \frac{(-1)^j}{j!} \in (0, e^{-1}), \qquad w_j = \lfloor S\, e_-/j! \rfloor \;\; (0 \le j \le 18), \tag{16}$$

the alternating partial sum being a valid strict lower bound for e^{-1} because it terminates on a negative term; integer weights vanish beyond $j = 12$. Starting from $v = \lfloor Sh \rfloor$, one block is

$$v_0 = v, \qquad v_{j+1} = \left\lfloor \frac{\min(A_{10} v_j,\, A_{20} v_j)}{D} \right\rfloor, \qquad \mathcal{B}v = \left\lfloor \frac{\sum_{j=0}^{18} w_j v_j}{S} \right\rfloor, \tag{17}$$

all floors and minima componentwise, repeated $\lambda T = 3\,855$ times.

Proposition 4.7 (Soundness and overflow safety, C)**.** *Every intermediate vector of* (17)*, divided by S, is a lower bound for the corresponding quantity in* (15)*; and the computation cannot overflow signed* `int64`*.*

Proof. Rounding is downward at every step and the omitted Poisson terms are nonnegative, so both sources of error have the correct sign; the inductive hypothesis $v_j/S \le M^j(h)$ is preserved by Lemma 4.4 and the monotonicity of $x \mapsto \lfloor x \rfloor$. For the bound, $0 \le v_i \le S$ throughout, so each sparse row product is at most $DS < 2^{63}$ and the weighted accumulator at most $S \sum_j w_j \le S^2 < 2^{63}$. The two payoff columns are carried separately and never added. Threshold comparisons are exact integer cross-products. □

The stationary distribution, floating solvers and trajectory plots supply no input to this certificate. **Trust boundary.** Exact integer arithmetic removes rounding from the recurrence; it does not make the construction of states, destinations, matrices and loops formally verified. That construction, and the identification of the finite generator and its Poissonised semigroup with the physical protocol, are conventional (**C**); see §7.

The robust lower numerators at $K = 32$ are $2\,145\,517\,682$ for the band $[10, 11]$ and $2\,145\,514\,420$ for $[10, 20]$; only the second is needed. The wider band costs just $3\,935$ integer units — about

1.83×10^{-6} — relative to the nominal lower certificate, so certifying a factor-two rate band is nearly free here. The enclosure of the robust value is

$$\frac{2\,145\,514\,420}{2^{31}} \;\le\; p_{\mathrm{rob}}(32) \;\le\; \frac{2\,145\,545\,336}{2^{31}}, \tag{18}$$

of width about 1.44×10^{-5}, the upper bound coming from the admitted value $\kappa = 10$.

4.5 Enlarged robustness: bias and a rate box

Two idealisations in Definition 2.1 invite attack: that the allocation coin is exactly fair, and that the six rate constants are exactly the declared rationals. Both can be removed without recomputing the model from scratch, and the tolerances are far larger than a perturbative estimate suggests.

Allocation bias. Replace the fair coin by independent allocation to daughter A with a common probability θ, the daughters still complementary. The terminal payoff becomes $h_{q,\theta} = \mathbf{1}\{p \ge q\} s_\theta(R)$ with s_θ as in (3). For fixed $R \ge 1$ the function $\theta \mapsto \theta^R + (1-\theta)^R$ is convex and symmetric about $1/2$, so it is maximised over any interval $[\theta_0, 1-\theta_0]$ at the endpoints: certifying at θ_0 certifies the whole interval. Running the envelope certificate with that payoff — rather than bounding the payoff difference — gives the following.

Theorem 4.8 (Bias tolerance, C/E)**.** *For every $\theta \in [39/100,\, 61/100]$ and every $\kappa \in [10, 20]$, the $K = 32$ source meets the target:*

$$\min_{z \in \mathcal{B}_{32}} \Pr_z(\textit{joint success}) \;\ge\; \frac{2\,145\,366\,143}{2^{31}} = 0.999063\ldots > \frac{999}{1000}.$$

Because $\theta = 1/2$ lies in the family, every $K < 32$ remains excluded, so 32 is still the exact robust minimum over this enlarged uncertainty set. At $\theta = 19/50$ the guarantee genuinely fails: the exact upper certificate at the admitted bound seed $(0,1,0)$ is $2\,145\,078\,112/2^{31} < 999/1000$.

An 11 percentage-point tolerance in the allocation bias is a substantive robustness statement: it is 22 times the ± 0.005 that a payoff-perturbation argument charged against the margin, and it is obtained at the same computational cost as the original certificate. Note also that the worst admitted start migrates from the resident-heavy $(31,0,0)$ to the bound seed $(0,1,0)$ as the bias grows, which is why a fixed-worst-start shortcut would be unsound.

A simultaneous rate box. Let each of the six rate constants in (9) be perturbed independently by a factor in $[1-\varepsilon, 1+\varepsilon]$, the two complex-exit constants remaining tied to one another and ranging over $[10(1-\varepsilon), 20(1+\varepsilon)]$. The generator is jointly affine in these five group parameters, so the family lies in the convex hull of its 2^5 vertex generators, and the envelope of Lemma 4.4 applies with $Mv = \min_i P_i v$ over the vertices — equivalently, groupwise, with the rowwise minimum computed one group at a time.

Theorem 4.9 (Rate-box tolerance, C/E)**.** *At $\varepsilon = 1/300$ the $K = 32$ source meets the target uniformly over the whole box and the whole release band, with exact lower certificate $2\,145\,342\,442/2^{31} = 0.999052\ldots$ (integer clock 773). The exact minimum is therefore still 32 under simultaneous $\pm 0.33\%$ uncertainty in every rate constant.*

Since multiplying all six constants by a common factor is a time rescaling, Theorem 4.9 also absorbs a $\pm 0.33\%$ uncertainty in the deadline. At $\varepsilon = 1/250$ our lower certificate falls below the target while the upper certificate stays above it, so that case is undetermined, not excluded; we report it as such.

Table 2: Terminal diagnostics at $K = 32$, $T = 5$, $q = 4$, from the backward equations solved in floating point (evidence **N**; no certificate uses them). The balanced row is the reference decomposition $9.00058131 \times 10^{-4}$ and 4.581103×10^{-6}. The worst admitted start differs between settings, which is exactly why uniformity over the restart set matters.

(κ_d, κ_r)	start	quota fail	partition fail	$\mathbb{E}[R]$	$\mathbb{E}[\text{food}]$	$\mathbb{E}[p]$
$(20, 10)$	$(31, 0, 0)$	2.017×10^{-2}	7.45×10^{-7}	23.37	0.202	8.01
$(20, 10)$	$(0, 1, 0)$	1.469×10^{-3}	1.03×10^{-5}	20.35	0.208	11.03
$(10, 10)$	$(31, 0, 0)$	9.001×10^{-4}	4.58×10^{-6}	20.73	0.154	10.61
$(10, 10)$	$(0, 1, 0)$	4.04×10^{-6}	1.865×10^{-4}	15.99	0.155	15.39
$(10, 20)$	$(31, 0, 0)$	1.74×10^{-5}	1.78×10^{-5}	18.65	0.116	12.94
$(10, 20)$	$(0, 1, 0)$	1.70×10^{-9}	1.325×10^{-3}	12.68	0.108	18.97

4.6 A direction in which it breaks

Theorems 4.8 and 4.9 might suggest that the 32-unit result is robust in every direction. It is not, and the exception is instructive.

In (9) the bound complex has two exits, $C_S \to S + F$ and $C_S \to S + P_S$, and the source gives them the *same* constant κ: the complex is equally likely to dissociate as to release product. Let us untie them, writing κ_d for dissociation and κ_r for release. This changes no molecularity, adds no species and keeps the complex's total lifetime under control; it changes only the branching ratio $\kappa_r/(\kappa_d + \kappa_r)$.

Theorem 4.10 (Branching frontier, C/E). *At the fixed parameters $(\kappa_d, \kappa_r) = (20, 10)$, with $T = 5$, $q = 4$ and target $999/1000$, the budget $K = 32$ is excluded by the exact upper certificate $2\,104\,161\,482/2^{31} = 0.979827\ldots$, and the least admissible budget is exactly 42, with exact lower certificate $2\,145\,488\,787/2^{31} = 0.999071\ldots$ Every $K \le 14$ is excluded by Theorem 3.1 and every $15 \le K \le 41$ by an exact upper certificate at the admitted start $(K-1, 0, 0)$.*

Halving the productive branching fraction therefore costs 10 further material units, a 31% surcharge, without changing a single reaction. The two frontiers are drawn together in Figure 2.

More surprisingly, the guarantee is not monotone in the product-release rate either. At $(\kappa_d, \kappa_r) = (10, 20)$ — the complex now twice as likely to make product as to fall apart — the budget $K = 32$ is again excluded, by the exact upper certificate $2\,144\,643\,824/2^{31} = 0.998678\ldots$ — and this time the excluding witness is the *bound seed* $(0, 1, 0)$, not the resident-heavy start. Raising only the productive channel breaks a specification that the balanced source meets.

Why. The dissociation channel is not futile. It returns the food-derived moiety of the complex to the food pool, and food is the substrate of $S + F \to 2S$, the only channel that increases the carrier count R on which the partition payoff $1 - 2^{1-R}$ depends. The complex therefore arbitrates between the two failure modes of (1), and the diagnostics in Table 2 show both sides of the trade.

Favouring dissociation raises the expected carrier count from 20.7 to 23.4 and drives the expected harvest down from 10.6 to 8.0: the resident-heavy start now misses its quota with probability 2×10^{-2}. Favouring release does the opposite — the quota failure at that start collapses to 1.7×10^{-5} — but the food stock falls, the carrier count at the seed-like starts falls from 16.0 to 12.7, and those starts now lose a daughter with probability 1.3×10^{-3}. The balanced source sits between two distinct failure modes, and the 0.999 specification at $K = 32$ is available only in a window around it.

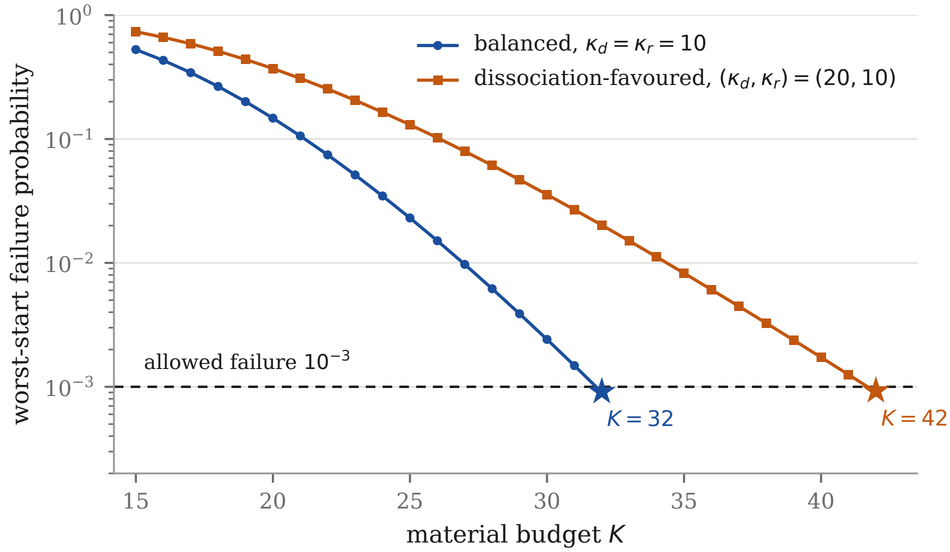

Figure 2: Exact two-sided certificates for the support-memory source. Plotted is the certified upper bound on worst-start failure, i.e. $1 - U_K$; the budgets marked with stars are the exact minima, where the plotted value is the certified lower bound's failure. Discrete certificates, not an interpolated guarantee. Untying the complex's two exits at $(\kappa_d, \kappa_r) = (20, 10)$ shifts the frontier from 32 to 42.

5 Proportion memory with a cooperative gate

5.1 The architecture class

Proportion memory keeps both species at all times, so the failure mode of §4 — a daughter with no carrier — is replaced by the much more demanding requirement that each daughter's counts re-enter a corridor. No elementary network achieving this under a finite-count deadline guarantee is known to us. The route that does work trades reaction order for exact tractability, and we now define the whole class so that it can be optimised rather than merely exhibited.

Definition 5.1 (The class $\mathcal{G}(q)$)**.** Fix integers $1 \le l_Y \le u_Y < l_X \le u_X$, increments $g_X, g_Y \ge 1$ and a quota $q \ge 1$. The species are X, Y, F, H, W, P_X, P_Y, each of unit material weight. The two reversible pairs are

$$\begin{aligned} l_X X + l_Y Y + (g_X + g_Y + q)F + H &\rightleftharpoons (l_X + g_X)X + (l_Y + g_Y)Y + qP_X + W, \\ l_X Y + l_Y X + (g_X + g_Y + q)F + H &\rightleftharpoons (l_X + g_X)Y + (l_Y + g_Y)X + qP_Y + W, \end{aligned} \tag{19}$$

in the fixed-volume falling-factorial convention, with positive rational constants k (forward) and b (reverse). The conserved pools are $x + y + f + p_X + p_Y = N$ and $h + w = 1$. Program 0 has newborn set

$$\mathcal{B}_0 = \{\, l_X \le x \le u_X,\ l_Y \le y \le u_Y,\ f = N - x - y,\ h = 1,\ w = p_X = p_Y = 0 \,\}, \tag{20}$$

program 1 exchanges x and y, and the decoder is the sign of $x - y$. The protocol is Definition 2.1 with $T = 1$: run, harvest all product, allocate every intact molecule by one fair complementary coin, discard the inherited H/W token, refill food to core N and supply one fresh H.

We call the member $(l_X, u_X, g_X; l_Y, u_Y, g_Y; q) = (65, 195, 121; 8, 64, 30; 4)$ the *reference construction* $\mathcal{R}$. Its core is $N = 414$, its initial supply 415 including the fuel token, and both of its reaction directions have molecularity 229.

Lemma 5.2 (Minimal core mass, C)**.** *A member of $\mathcal{G}(q)$ with the stated corridors requires core mass at least*

$$N \;\geq\; c_X + c_Y + q, \qquad c_X := u_X + g_X, \quad c_Y := u_Y + g_Y, \tag{21}$$

and $N = c_X + c_Y + q$ is attainable.

Proof. The gate (19) consumes $g_X + g_Y + q$ food units, and must be enabled at every newborn, in particular the food-poorest one $x = u_X$, $y = u_Y$, where $f = N - u_X - u_Y$. Hence $N - u_X - u_Y \geq g_X + g_Y + q$. □

Lemma 5.3 (Exact two-state closure, C; K)**.** *For every $z \in \mathcal{B}_0$ the reachable class of the full CTMC is exactly $\{z, z^*\}$ with*

$$z^* = (x + g_X,\; y + g_Y,\; f - (g_X + g_Y + q),\; 0,\; 1,\; q,\; 0).$$

Proof. At z the mirror forward channel needs l_X copies of Y and is disabled because $y \leq u_Y < l_X$; both reverse channels need a product that is absent. The selected forward channel is enabled by (20) and Lemma 5.2, and its unique successor is z^*. At z^* both forward channels are disabled because $h = 0$, and the mirror reverse needs P_Y, which is absent; the selected reverse is enabled and returns exactly to z. These are all four channels of the source, so the class is closed. The mirror program is the image under the label permutation. Distinct newborns give distinct classes because subtracting the fixed increment recovers the newborn. [`newborn_only_forward`, `grown_only_reverse`, `forward_result`, `reverse_result`, `exact_closed_pair`] □

This is an exact restriction of the literal chain, not an averaged projection or a quasi-steady-state elimination, and it is what makes the class exactly optimisable: the terminal law is a two-state law, and the entire finite-count difficulty moves to the allocation stage.

5.2 The clock is free; the corridor is not

Lemma 5.4 (Clock, C)**.** *Let a be the forward rate at a newborn and d the reverse rate at z^*. The probability of occupying z^* at time T is exactly*

$$r(a, d) = \frac{a}{a + d}\left(1 - e^{-(a+d)T}\right). \tag{22}$$

If $a \geq A$ and $d \leq D$ then $r(a, d) \geq \frac{A}{A+D}(1 - e^{-AT})$. Moreover for any $A, D > 0$ the choices

$$k = \frac{A}{l_X!\, l_Y!\, (g_X + g_Y + q)!}, \qquad b = \frac{D}{(c_X)_{l_X + g_X}\, (c_Y)_{l_Y + g_Y}\, q!} \tag{23}$$

achieve those bounds uniformly over the newborn set, where $(n)_r = n!/(n - r)!$.

Proof. Equation (22) solves $\dot{\pi} = a - (a + d)\pi$, $\pi(0) = 0$. Both factors are monotone in the stated directions. For (23), the forward propensity is $k\,(x)_{l_X}(y)_{l_Y}(f)_{g_X+g_Y+q}h$ and $(n)_r \geq r!$ whenever $n \geq r$, which holds on the newborn set; the reverse propensity is $b\,(x')_{l_X+g_X}(y')_{l_Y+g_Y}(p_X)_q w$ and each falling factorial is maximised at the largest admitted count, namely c_X and c_Y. [`forward_rate_lower`, `reverse_rate_upper`] □

Because k and b are free positive rationals, $r(a, d)$ can be pushed arbitrarily close to 1 at no material cost. *The clock in $\mathcal{G}(q)$ is free.* By Lemma 5.3 the joint success probability from a newborn z is exactly

$$\Pr_z(\text{success}) = r(a, d) \cdot \psi(x + g_X; l_X, u_X)\, \psi(y + g_Y; l_Y, u_Y), \tag{24}$$

the product over the disjoint X- and Y-allocation objects conditional on the terminal state z^*. Consequently the frontier of $\mathcal{G}(q)$ is determined by the partition stage alone, and any kinetic constraint can only raise it.

5.3 The endpoint reduction

Define, for a coordinate with corridor $[l, u]$ and increment g,

$$\Phi(l, u, g) = \min_{l+g \le n \le u+g} \psi(n; l, u), \tag{25}$$

the uniform two-daughter return probability over all newborns of that coordinate. A naive evaluation of (25) for $\mathcal{R}$ visits 7 467 newborns per program. It is unnecessary.

Lemma 5.5 (Endpoint reduction, C)**.** *Fix $l \le u$. The map $n \mapsto \psi(n; l, u)$ is nondecreasing for $n \le l + u$ and nonincreasing for $n \ge l + u$. Consequently the minimum in* (25) *is attained at an endpoint, and if $l \le g \le u$,*

$$\psi(l+g; l, u) = 1 - 2\Pr\{\mathrm{Bin}(l+g, \tfrac{1}{2}) \le l - 1\}, \qquad \psi(u+g; l, u) = 1 - 2\Pr\{\mathrm{Bin}(u+g, \tfrac{1}{2}) \le g - 1\}. \tag{26}$$

Moreover $\psi(l+g; l, u)$ is nondecreasing in g and nonincreasing in l, and $\psi(u+g; l, u)$ is nonincreasing in g at fixed $u + g$.

Proof. For $n \le l + u$ the constraint $I \ge l$ forces $n - I \le n - l \le u$, so the event is exactly "both daughters receive at least l"; for $n \ge l + u$ the constraint $I \le u$ forces $n - I \ge n - u \ge l$, so it is exactly "both daughters receive at most u". Couple the binomial counts as partial sums of one Bernoulli sequence, so that $\mathrm{Bin}(n+1, \frac{1}{2}) \ge \mathrm{Bin}(n, \frac{1}{2})$ pointwise. The first event is increasing in both daughter counts and the second decreasing, which gives the two monotonicity statements and hence the unimodality. For (26): at $n = l+g$ the failure is $\{I \le l-1\} \cup \{I \ge g+1\}$, two disjoint events of equal probability by the symmetry $I \leftrightarrow n - I$; at $n = u+g$ the failure is $\{I \ge u+1\} \cup \{I \le g-1\}$, likewise. For the last sentence, raising g by one at fixed l raises n by one and leaves the threshold $l - 1$ fixed, and $\Pr\{\mathrm{Bin}(n + 1) \le l - 1\} \le \Pr\{\mathrm{Bin}(n) \le l - 1\}$ by the same coupling; raising l by one raises both n and the threshold by one, and $\{\mathrm{Bin}(n) \le l - 1\} \subseteq \{\mathrm{Bin}(n + 1) \le l\}$. Raising g at fixed $u + g$ raises the threshold $g - 1$ alone. □

So four binomial evaluations replace the whole table, and the search over corridors becomes a two-parameter unimodal problem per coordinate. For $\mathcal{R}$, Lemma 5.5 gives the exact minimum

$$\Phi_* = \Phi(65, 195, 121) \cdot \Phi(8, 64, 30) = 0.999712941541084\ldots, \tag{27}$$

attained at the newborn $(65, 64)$ with productive parent $(186, 94)$; the exact rational value has an 84-digit numerator and is recorded in `data/frontier_proportion.json`. A union bound over the four endpoint tails gives instead $1 - E$ with $E = 4.25767634587 \times 10^{-4}$, hence only the weaker constant 99957/100000. The exact minimum recovers 1.387×10^{-4} of certificate slack, which was conservatism and not a physical loss. `[complementary_binomial_law, endpoint_tail_certificate]`

5.4 Sharpened constants for $\mathcal{R}$

Combining (27) with Lemma 5.4 at $A = 20$, $D = 10^{-5}$, $T = 1$ — the rate bounds of $\mathcal{R}$ — and bounding e^{20} below by its positive 25-term Taylor sum, which exceeds 4×10^8, gives

$$r_* = \frac{20}{20 + 10^{-5}}\Big(1 - \frac{1}{4 \times 10^8}\Big) = 0.999999497500251\ldots \tag{28}$$

and hence the following improvement.

Theorem 5.6 (Sharpened guarantee, C/E). *For $\mathcal{R}$, uniformly over all 7 467 newborns of each program,*

$$\Pr(\textit{joint success}) \;\ge\; r_*\Phi_* = 0.999712439185582\ldots \;>\; \frac{9997}{10000},$$

and ten inspected cycles along a prescribed lineage succeed jointly with probability at least $(9997/10000)^{10} > 0.997$, collecting forty fresh selected products.

The union-bound constants were 0.9995 per cycle and 0.995 over ten; the guaranteed failure budget is thus more than halved, with no change to the source, the protocol, the material account or the output. The improvement is entirely a matter of replacing a union bound by the exact endpoint minimum. We stress what is *not* claimed: $r_*\Phi_*$ is a product of the worst clock factor and the worst allocation factor, which need not occur at the same newborn, so it is a valid lower bound and not the exact minimum of the joint probability. The allocation minimum Φ_* is exact.

Theorem 5.7 (Bias tolerance, C/E). *$\mathcal{R}$ tolerates independent allocation bias $\theta \in [0.49, 0.51]$: the uniform partition probability is at least $0.999614114647\ldots$ and the joint guarantee at least $r_* \cdot 0.999614\ldots = 0.999613612341\ldots > 0.9996$. At $\theta = 0.48$ the uniform partition value drops to $0.999220794519\ldots$ and the 0.9996 guarantee is lost.*

For fixed θ the accepted interval in (4) is $[b, n-b]$ with $b = \max(l, n-u)$, which is symmetric about $n/2$; its binomial mass is symmetric in $\theta \leftrightarrow 1-\theta$ and increases as θ approaches $1/2$, because $\frac{d}{d\theta}\Pr_\theta\{b \le I \le n-b\} = n[\Pr_\theta\{J = b-1\} - \Pr_\theta\{J = n-b\}]$ with $J \sim \mathrm{Bin}(n-1, \theta)$, whose bracket is nonnegative for $\theta \le 1/2$ since the two binomial coefficients agree. The parent-count reduction of Lemma 5.5 holds for every fixed θ, because its coupling argument never used fairness. So interval endpoints in θ and parent-count endpoints in n together reduce the claim to eight exact rational evaluations.

The bias window here, ± 1 percentage point, is an order of magnitude narrower than the ± 11 points of Theorem 4.8. This is the first sign of the structural difference: a corridor is a two-sided constraint, and a biased coin moves the daughter counts off centre in a way that presence memory simply does not notice.

5.5 The exact frontier

We can now answer the question the construction leaves open.

Theorem 5.8 (Frontier of $\mathcal{G}(q)$, C/E). *Let $q = 4$ and let $\rho < 1$ be a target.*

(i) At $\rho = 1999/2000$ — specification of $\mathcal{R}$ — the least core mass admitted by any member of $\mathcal{G}(4)$ is $N = 253$, attained by $(l_X, u_X, g_X; l_Y, u_Y, g_Y) = (37, 126, 75; 1, 36, 12)$ with exact uniform partition probability $0.999529184472\ldots$ and reaction molecularity 130.

(ii) *If in addition every daughter is required to retain at least eight minority molecules ($l_Y \geq 8$, as $\mathcal{R}$ does), the least core mass is $N = 375$, attained by $(62, 172, 110; 8, 61, 28)$, molecularity* 213.

(iii) *At $\rho = 999/1000$ — the support-memory specification of Theorem 4.2 — the least core mass is $N = 230$, attained by $(33, 115, 68; 1, 32, 11)$ with exact uniform partition probability* $0.999033665776\ldots$ *and molecularity* 118.

In each case the value is a minimum over the entire class, for every choice of positive rational rate constants and every deadline.

Proof. By Lemma 5.4 the clock factor can be made to exceed any $\rho' < 1$, so a core mass N is admissible if and only if some corridor choice with $c_X + c_Y + q \leq N$ has $\Phi_X \Phi_Y > \rho$; by Lemma 5.2 the cost is exactly $c_X + c_Y + q$. By Lemma 3.4 we may restrict to $l \leq g \leq u$, so $g \leq c/2$ per coordinate, and by Lemma 5.5 each Φ is the minimum of the two closed forms in (26). The monotonicity statements of Lemma 5.5 show that at fixed (g, c) the value is nonincreasing in l, so the minority takes l_Y at its declared floor and the majority takes $l_X = u_Y + 1$, the smallest value compatible with $u_Y < l_X$.

The search is therefore over (c_Y, g_Y, c_X, g_X) and is exhaustive by the following stopping rules, both of which are sound. First, since $\Phi_X \leq 1$, any admissible pair has $\Phi_Y > \rho$, so minority candidates failing that are discarded. Second, for each minority candidate the majority cost is scanned upward from its minimum $2l_X$ and the first success is taken, which is the true minimum for that candidate; and the enumeration of minority costs stops as soon as $c_Y + q$ reaches the incumbent total, which is valid because $c_X \geq 2$. All arithmetic is exact rational. The witnesses and the exhaustiveness are replayed by `scripts/frontier_proportion.py`. □

Three readings of Theorem 5.8. First, 415 was a construction, not a cost: at its own specification the class does the job with 254 supplied units, a 39% saving, and with reaction order 130 rather than 229. Second, most of the remaining gap to $\mathcal{R}$ is explained by an undeclared constraint — the requirement that each daughter keep at least eight minority molecules, which is a stronger memory than the ratio decoder needs. Making that constraint explicit raises the frontier to 375, much closer to 414. Third, and this is the point of the merged comparison, even the unconstrained frontier is an order of magnitude above the support-memory minimum.

5.6 Reliability scaling and the corridor cascade

Theorem 5.9 (Order of the frontier, C)**.** *For $q \geq 1$ and $\delta \in (0, 1/2)$ let $N(\delta, q)$ denote the frontier of $\mathcal{G}(q)$ at target $1 - \delta$ with $l_Y \geq 1$. Then*

$$q + 2\log_2(2/\delta) \;\leq\; N(\delta, q) \;\leq\; q + 20(z+2)^2 + 6, \qquad z := \sqrt{2\ln(8/\delta)}, \tag{29}$$

so $N(\delta, q) = q + \Theta(\log(1/\delta))$.

Proof. Lower bound. At the minority newborn $y = l_Y$ the parent holds $n_Y = l_Y + g_Y$ minority objects, and the failure event contains the two disjoint all-to-one-daughter events, of total probability $2 \cdot 2^{-n_Y}$. Hence $2^{1-n_Y} \leq \delta$, i.e. $n_Y \geq \log_2(2/\delta)$. By Lemma 3.4, $l_Y \leq g_Y$, so $g_Y \geq n_Y/2 \geq \frac{1}{2}\log_2(2/\delta)$. Again by Lemma 3.4, $u_Y \geq g_Y$, and by Definition 5.1 $l_X > u_Y \geq g_Y$, whence $g_X \geq l_X > g_Y$ and $u_X \geq g_X > g_Y$. Therefore $N = q + u_X + g_X + u_Y + g_Y \geq q + 4g_Y \geq q + 2\log_2(2/\delta)$.

Table 3: The exact frontier of $\mathcal{G}(4)$ with $l_Y \geq 1$, over five decades of target failure. The last column is $N/\ln(1/\delta)$; the penultimate is the cascade prediction $20z_\delta^2$ with z_δ the standard normal upper quantile at level δ.

δ	N	optimal corridors	molecularity	$20z_\delta^2$	$N/\ln(1/\delta)$
10^{-2}	154	$X[21,78]$+45, $Y[1,20]$+7	79	108	33.44
10^{-3}	230	$X[33,115]$+68, $Y[1,32]$+11	118	191	33.30
10^{-4}	305	$X[43,155]$+90, $Y[1,42]$+14	153	277	33.12
10^{-5}	384	$X[54,196]$+114, $Y[1,53]$+17	191	364	33.35
10^{-6}	460	$X[67,232]$+137, $Y[1,66]$+21	231	452	33.30
10^{-7}	537	$X[77,273]$+160, $Y[1,76]$+24	267	541	33.32

Upper bound. By Hoeffding's inequality [7] $\Pr\{|\mathrm{Bin}(n,\frac{1}{2}) - n/2| \geq t\} \leq 2e^{-2t^2/n}$, so the failure of a constraint of half-width $w/2$ at parent count n is at most $\delta/4$ as soon as $w \geq z\sqrt{n}$. Put $s = (z+2)^2$ and take

$$l_Y = 1, \quad g_Y = \lceil s \rceil, \quad u_Y = \lceil 3s \rceil, \quad l_X = u_Y + 1, \quad g_X = \lceil 6s \rceil + 1, \quad u_X = \lceil 10s \rceil + 1.$$

Writing $\sqrt{s} = z+2$ and using $\sqrt{\alpha s + \beta} \leq \sqrt{\alpha}\,\sqrt{s} + \beta/(2\sqrt{\alpha s})$, each of the four conditions $w \geq z\sqrt{n}$ reduces to a polynomial inequality in z that holds for all $z \geq 1$: for the minority lower constraint $s-1 \geq z\sqrt{s+2}$ since $z^2+4z+3 \geq z^2+2z+1$; for the minority upper constraint $2s-1 \geq z\sqrt{4s+2}$ since $2z^2+8z+7 \geq 2z^2+4z+1$; for the majority lower constraint $3s-1 \geq z\sqrt{9s+4}$ since $3z^2+12z+11 \geq 3z^2+7z$; and for the majority upper constraint $4s-1 \geq z\sqrt{16s+4}$ since $4z^2+16z+15 \geq 4z^2+9z$. A union bound over the four gives success at least $1-\delta$, and $z \geq 1$ holds for every $\delta < 1/2$. The core mass is $N = q + u_X + g_X + u_Y + g_Y \leq q + 20s + 6$. □

The two constants in (29) are far apart, and the exact frontier lies between them in a strikingly regular way. Table 3 lists it over five decades.

The ratio $N/\ln(1/\delta)$ is constant to within 1% across the whole range, so empirically $N \approx 33.3\ln(1/\delta)$; we report this as an observation about the computed values, not as a theorem. Its structure, however, can be explained, and the explanation generalises.

The cascade. Suppose all four constraints are tuned to the same Gaussian margin, i.e. each corridor half-width equals z standard deviations of its binomial. Writing $l = \lambda z^2$, $g = \gamma z^2$, $u = \mu z^2$, the two conditions per coordinate become $\gamma - \lambda = \sqrt{\lambda+\gamma}$ and $\mu - \gamma = \sqrt{\gamma+\mu}$, and the level-stacking constraint is $\lambda_{k+1} = \mu_k$. The unique solution starting from $\lambda_1 = 0$ is given by the triangular numbers $T_j = j(j+1)/2$:

$$l_k = T_{2k-2}z^2, \qquad g_k = T_{2k-1}z^2, \qquad u_k = T_{2k}z^2, \tag{30}$$

because $T_{2k-1} - T_{2k-2} = 2k-1 = \sqrt{T_{2k-2}+T_{2k-1}}$ and $T_{2k} - T_{2k-1} = 2k = \sqrt{T_{2k-1}+T_{2k}}$ — both identities exact. The k-th level therefore costs $c_k = (T_{2k}+T_{2k-1})z^2 = 4k^2z^2$, and m ordered levels cost

$$N \;\approx\; q + z^2 \sum_{k=1}^{m} 4k^2 = q + \tfrac{2}{3}\,m(m+1)(2m+1)\,z^2. \tag{31}$$

For $m = 2$ this is $20z^2$, the penultimate column of Table 3; it agrees with the exact frontier to within 1.8% at $\delta = 10^{-6}$ and 0.7% at $\delta = 10^{-7}$, the residual at larger δ coming from the minority's

Table 4: Per-level corridor costs $c_k = u_k + g_k$ from exact computation with an equal per-level budget δ/m, normalised to the second level so that the prediction $c_k \propto k^2$ reads $0.25 : 1 : 2.25 : 4$. The first level is presence-limited rather than corridor-limited and is the outlier the cascade derivation predicts.

δ	c_1	c_2	c_3	c_4	c_k/c_2
10^{-3}	48	201	462	829	0.24 : 1 : 2.30 : 4.12
10^{-5}	78	328	754	1359	0.24 : 1 : 2.30 : 4.14
10^{-7}	108	458	1056	1902	0.24 : 1 : 2.31 : 4.15
prediction	—	—	—	—	0.25 : 1 : 2.25 : 4

lowest constraint, which is a presence bound rather than a corridor bound and is therefore cheaper than the Gaussian balance assumes.

Equation (31) predicts that reading an ordered m-level composition costs $\Theta(m^3 \log(1/\delta))$ material — cubically more than a single bit, because each level must clear the level below it, and clearing a corridor of size c costs $\Theta(\sqrt{c})$ of extra width which then has to be re-grown. Table 4 tests this against exact computation with an equal per-level failure budget (a sufficient design, so each entry is an exact upper bound on the m-level frontier; at $m = 2$ it reproduces the true frontier to within 1.5%). The per-level costs track k^2 closely, and the agreement improves as δ decreases, as the derivation requires.

5.7 The price of the gate

The frontier is cheap in material only because the gate is expensive in mechanism. The member of $\mathcal{G}(q)$ with parameters $(l_X, u_X, g_X; l_Y, u_Y, g_Y; q)$ has reaction molecularity

$$l_X + l_Y + (g_X + g_Y + q) + 1 \tag{32}$$

in each direction — 229 for $\mathcal{R}$ and 130 at the frontier of Theorem 5.8(i). Along the frontier the cascade (30) makes it $(T_2 + T_3 + 1)z^2 = 10z^2$ against a core of $20z^2$, so molecularity is almost exactly $N/2$: it is $79, 118, 153, 191, 231, 267$ down the six rows of Table 3, and is therefore $\Theta(\log(1/\delta))$ as well. Moreover the forward propensity varies enormously over the newborn domain. For $\mathcal{R}$, $\log_{10}$ of the propensity ranges from 63.63 at the vertex $(195, 64)$ to 121.87 at $(118, 14)$, a span of $10^{58.23}$; at the $N = 253$ frontier point the span is $10^{40.01}$. The minimum is at a vertex for a reason: log of a product of falling factorials in x, y and $f = N - x - y$ is a sum of logarithms of positive affine functions, hence concave on the newborn rectangle, and a concave function attains its minimum over a polytope at a vertex. No rescaling removes this span, because multiplying all forward constants by a positive scalar, or changing the global time unit, preserves every dimensionless ratio.

So the material saving of (31) conceals a mechanistic bill, and Corollary 3.3 says where it is paid. This is the honest reading of the proportion-memory result: the class is a mathematical existence witness with an exactly known material frontier, not a proposal for practical chemistry. Replacing a 229-body event by a chain of bimolecular reactions is a separate substantive problem — it changes intermediate stock, sequestration, reaction races, leakage into the wrong program, the deadline distribution and the partition — and a coarse success bound p transfers to a refinement only as $p - \epsilon$, with ϵ a total-variation error over a full cycle including the reset map, and with the intermediates' own inventory charged.

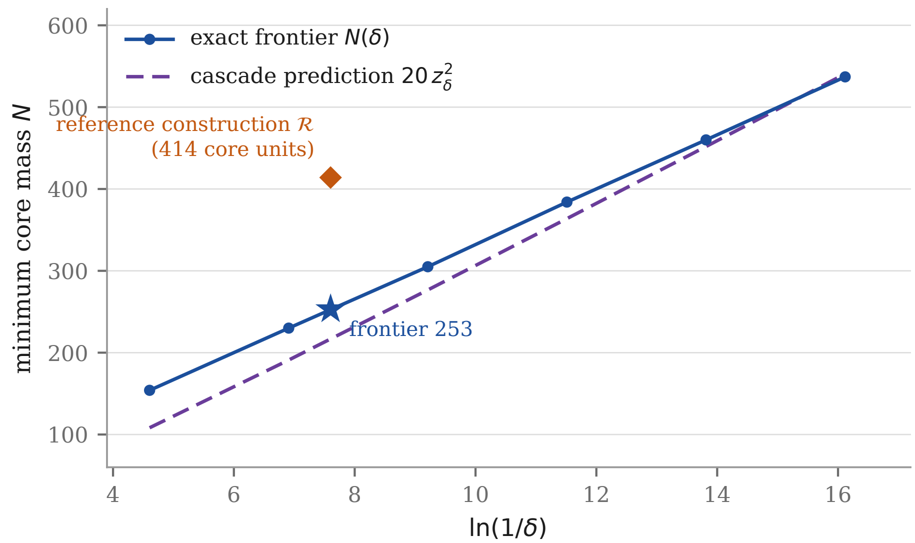

Figure 3: The exact material frontier of the cooperative-gate class $\mathcal{G}(4)$ against target reliability, with the cascade prediction $20z_\delta^2$ of (31). $\mathcal{R}$ sits well above the frontier at its own specification.

6 Reading the two frontiers together

Table 5 is the merged result. Four remarks on how to read it.

The comparison is conservative against proportion memory only in one direction, and it is the direction that matters. The number 32 is the minimum *for one specific elementary network* at a fixed rate band and a fixed deadline; some other elementary support-memory network might do better, so 32 is an upper bound on the true cost of support memory under those constraints. The number 231 is a minimum over an entire class in which reaction order, rate constants and deadline are all free; every member of that class pays at least that much. So the gap of roughly a factor of seven is a lower bound on the gap, not an artefact of comparing a good design with a bad one.

The extra cost is the corridor, not the clock and not presence. The carrier ceiling (6) charges proportion memory only 26 units, against 15 for support memory: keeping both species *present* is nearly free. Lemma 5.4 shows the clock is free in $\mathcal{G}(q)$ outright. What costs 204 further units is Lemma 3.4 and its consequences: a count must land inside a window, the window must be wide enough to absorb binomial fluctuation of order $\sqrt{N}$, and the growth increment that re-centres it must itself lie inside the window. That is the whole of the difference, and (31) is its asymptotic form.

The two encodings fail differently under a biased coin. Support memory tolerates $\theta \in [0.39, 0.61]$ and proportion memory only $[0.49, 0.51]$, a twentyfold difference in tolerance, for the same structural reason: a presence constraint is one-sided and a corridor is two-sided.

Neither frontier is monotone in what it looks like it should be monotone in. Section 4.6 shows that the support-memory frontier moves the wrong way when the productive channel is

Table 5: The two encodings at one specification: quota $q = 4$, target 0.999, uniformity over the whole restart set, one complementary fair allocation, refill to the working inventory, no post-division repair.

	support memory	proportion memory
stored object	which species is present	which species is in the majority
decoder	presence test	$x/(x + y) > 1/2$
daughter must receive	at least one carrier	a count inside a corridor
kinetics	elementary, at most bimolecular	cooperative gate, order 118
rate constants	fixed band $\kappa \in [10, 20]$	free positive rationals
deadline	$T = 5$, fixed	free
carrier ceiling (6)	15	26
exact minimum material	**32**	**231**
status of that number	minimum for this source	minimum over the whole class
allocation-bias tolerance	$\theta \in [0.39, 0.61]$	$\theta \in [0.49, 0.51]$
scaling in $\log(1/\delta)$	—	$\approx 33.3 \ln(1/\delta)$

made faster, because the apparently futile dissociation channel is what recycles the food that builds carriers. Theorem 5.8(ii) shows that demanding *more* of the memory — eight minority molecules rather than one — costs 122 further units, so "retain both species" and "retain eight copies of each" are quantitatively different specifications. In both cases the lesson is the same: these systems have two competing failure modes on one probability law, and intuitions calibrated on either mode alone mislead.

7 Formal verification boundary

The developments compile in Lean 4 [13] (toolchain `leanprover/lean4:v4.30.0`) against a frozen Mathlib [11], under strict receipts: exit code zero, warnings as errors, and an axiom report for each exported declaration. All exported declarations depend only on `propext`, `Classical.choice` and `Quot.sound`; no `sorry` and no numerical-answer axiom is present, and the exact tail certificate uses ordinary `decide` with kernel reduction, not `native_decide`. Receipts are in `companion/verification/`.

What is outside the formal boundary. Five things, stated plainly. The identification of the finite generator and its Poissonised semigroup with the protocol of Definition 2.1 is conventional. The execution of the integer recurrence (17) on 3 384 states is external: the Lean development proves the envelope inequalities and the arithmetic margins, not the recurrence's output, and a theorem conditional on that output does not prove it. The monotone Bernoulli coupling behind Lemma 5.5, the mirror-program symmetry, and the falling-factorial normalisation are conventional. The exhaustive search of Theorem 5.8 is exact rational arithmetic in Python, with its soundness proved here, not in Lean. And a source hash is provenance, not semantics. We regard the resulting status — a conventional proof with exact finite arithmetic and decisive formally verified lemmas — as the accurate description, and neither inflate it to "fully mechanised" nor discount it as unfinished.

Table 6: Declaration map. Namespace `UsefulChemicalMemoryCost` (support memory) and `ProductiveChemicalHeredity` (proportion memory).

statement	declaration
material conservation and survival of the last carrier	`jump_conserves`, `final_resident_survives`
quota charges material; carrier necessity; Theorem 3.1	`quota_limits_carriers`, `reliability_requires_carriers`, `material_necessary`
generator affine in the release parameter	`release_affine`
bound seed admitted, with only its two exits	`bound_seed_admitted`, `bound_seed_rates`
Lemma 4.4 (envelope, iterates, Poisson block)	`envelope_le_kernel`, `envelope_iterates_lower`, `envelope_nonneg`, `robust_poisson_lower`
Lemma 4.6 (complement)	`complement_upper`
complementary daughters return; refill identity	`complementary_daughter_return`, `actual_refill_account`
exact margins around L, U_{31} and 511/512; ten-cycle margin	`exact_frontier_margins`, `ten_cycle_margin`, `publication_arithmetic`
exclusion from one admitted witness; minimum from finite coverage	`robust_exclusion`, `exact_minimum_from_coverage`
Lemma 5.3 (two-state closure)	`newborn_only_forward`, `grown_only_reverse`, `forward_result`, `reverse_result`, `exact_closed_pair`
Lemma 5.4 rate bounds	`forward_rate_lower`, `reverse_rate_upper`
complementary-binomial event; corridor return	`complementary_binomial_law`, `actual_daughters`
exact endpoint tail arithmetic	`fastTail_eq_binomial`, `endpoint_tail_certificate`
$e^{-20} \le 10^{-5}$; rational margins; ledger	`forward_wait_bound`, `joint_margin`, `ten_cycle_margin`, `finite_supply_ledger`, `refill_identity`
Proposition 2.3 (lineage iteration)	`repeated_conditional_success`

8 Reproduction

Every number printed in this paper is regenerated by the scripts in `scripts/`, using only the Python standard library, NumPy and SciPy:

script	produces
`ctmc_certificate.py`	the one-sided integer envelope certifier
`frontier_support.py`	Table 1, Theorems 4.2, 4.8, 4.9, 4.10
`frontier_proportion.py`	Lemma 5.5 check, (27), Theorem 5.8
`proportion_extras.py`	(28), Theorems 5.6, 5.7, Table 3, spans
`cascade.py`	Table 4
`branching_mechanism.py`	Table 2 (**N**)
`figures.py`	the three figures

The certifier reproduces the robust numerator 2 145 514 420 at $K = 32$ with 3 384 states, 287 admitted starts, integer clock 771 and 3 855 blocks, in about six seconds; the whole support-memory frontier computation, including the 28 budgets of Theorem 4.10, takes about six minutes on one core. The proportion-memory frontier of Theorem 5.8 takes under half a second per target. Outputs are stored as JSON under `data/` so the paper can be read without rerunning anything.

9 Limitations and open questions

Scope of the frontiers. Theorem 4.2 is a minimum for one source, and Theorem 5.8 a minimum over one class. Neither is a universal statement about chemical memory. Correlated or clumped partitioning, multi-carrier complexes, post-division repair, daughter selection, changing reference volume, contamination and common-support correction all fall outside the hypotheses, and each would change the arithmetic.

Calibration. The rate constants and deadlines are constructed benchmark values. Converting between concentration-rate and count-rate coefficients for a channel of molecularity r involves a factor $(N_A V)^{1-r}$, so assigning a volume and a physical time scale is a prerequisite to any judgement about diffusion, collision frequency, crowding or enzyme turnover. Increasing K at fixed volume, as the support-memory benchmark does, is not the same as increasing volume at fixed concentration.

The main open problem. Is there an *elementary* network — molecularity at most three, rate constants within a declared physical band — that achieves proportion memory under this protocol at all, and if so at what material cost? Our results bracket the question from one side only: the cooperative-gate class costs at least 231 units even with unbounded order, and Corollary 3.3 says that an elementary network denied large k and large r must buy its clock with K. A refinement of (19) into bimolecular steps would need its intermediates' inventory charged and a coupling error accounted over a full cycle including the reset map; the approximate-majority literature [2] is the natural starting point for the decoder, but it does not by itself preserve a finite material budget or a finite-time joint guarantee.

Secondary questions. A matching uniform sufficient scaling law for the support-memory source remains open — the necessary conditions (6) and (7) are explicit, but we have not proved a sufficient law, and we do not extrapolate one from a finite table. The exact location of the branching window of §4.6, the frontier as a joint function of (q, δ, T), and an end-to-end formalisation of either recurrence are all natural successors. None of them is a hidden condition of the theorems above.

A Exhaustiveness of the support-memory exclusions

For completeness we record the two exclusion mechanisms used in Theorems 4.2 and 4.10, and why together they cover every smaller budget.

Budgets $2 \leq K \leq 14$ are excluded analytically: on the quota event $R \leq K - q \leq 10$, so by (3) the joint success probability is at most $1 - 2^{-9} = 511/512 < 999/1000$, whatever the terminal distribution. Budgets $15 \leq K \leq 31$ (respectively $15 \leq K \leq 41$) are excluded by the exact upper certificates of Table 1, each at the single admitted start $(K-1, 0, 0)$, each obtained from Lemma 4.6 applied to the complementary payoff. No ordering among those upper values is used to infer any uncomputed case: each budget is checked individually. [`robust_exclusion`, `exact_minimum_from_coverage`]

Since the excluding witness is evaluated at $\kappa = 10$, which lies in the band $[10, 20]$, it excludes the robust guarantee as well as the nominal one. The converse does not hold and is not claimed: a budget failing at $\kappa = 10$ may well succeed at some other rate in the band.

B An independent replay

The certificates of §4 were reproduced for this paper from a fresh implementation written directly from (11) and (17), rather than by rerunning the original program. It reconstructs the state space and the destination table from the reaction list, asserts that every row of each integer matrix sums to D and that every entry is nonnegative, asserts the overflow bounds of Proposition 4.7 before iterating, and asserts the enumeration counts $|\Omega_{32}| = 3\,384$ and $|\mathcal{B}_{32}| = 287$ against the closed forms

$$|\Omega_K| = \sum_{c=0}^{\lfloor K/2 \rfloor} \frac{(K-2c+1)(K-2c+2)}{2} - (K+1), \qquad |\mathcal{B}_K| = \sum_{c=0}^{\lfloor K/2 \rfloor} (K-2c+1) - 2.$$

It reproduced all 36 numerators of Table 1, both robust band values, and the diagnostics of Table 2, to the last digit. Agreement between two independent implementations is corroboration of the arithmetic, not a proof of the semantic bridge described in §7.

References

[1] D. F. Anderson, G. Craciun and T. G. Kurtz. *Product-form stationary distributions for deficiency zero chemical reaction networks.* Bulletin of Mathematical Biology **72** (2010), 1947–1970. https://arxiv.org/abs/0803.3042.

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

[3] D. Del Vecchio, A. J. Ninfa and E. D. Sontag. *Modular cell biology: retroactivity and insulation.* Molecular Systems Biology **4** (2008), 161. https://doi.org/10.1038/msb4100204.

[4] A. Despons, Y. De Decker and D. Lacoste. *Structural constraints limit the regime of optimal flux in autocatalytic reaction networks.* Preprint, https://arxiv.org/abs/2306.02366.

[5] A. Erreygers and J. De Bock. *Imprecise continuous-time Markov chains: efficient computational methods with guaranteed error bounds.* Proceedings of ISIPTA '17, PMLR **62** (2017), 145–156. Extended preprint https://arxiv.org/abs/1702.07150.

[6] C. C. Govern and P. R. ten Wolde. *Energy dissipation and noise correlations in biochemical sensing.* Physical Review Letters **113** (2014), 258102. https://doi.org/10.1103/PhysRevLett.113.258102. See also *How biochemical resources determine fundamental limits in cellular sensing*, https://arxiv.org/abs/1308.1449.

[7] W. Hoeffding. *Probability inequalities for sums of bounded random variables.* Journal of the American Statistical Association **58** (1963), 13–30.

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

[9] T. Krak, J. De Bock and A. Siebes. *Imprecise continuous-time Markov chains.* International Journal of Approximate Reasoning **88** (2017), 452–528. https://doi.org/10.1016/j.ijar.2017.06.012.

[10] D. V. Kriukov, J. Huskens and A. S. Y. Wong. *Exploring the programmability of autocatalytic chemical reaction networks.* Nature Communications **15** (2024), 8295. https://doi.org/10.1038/s41467-024-52649-z.

[11] The mathlib Community. *The Lean mathematical library.* Proceedings of CPP 2020, 367–381. https://doi.org/10.1145/3372885.3373824.

[12] Y. J. Matsubara, S. Ameta, S. Thutupalli, P. Nghe and S. Krishna. *Conditions for Darwinian evolution in compartmentalized autocatalytic reaction networks.* Preprint, https://arxiv.org/abs/2211.03155.

[13] L. de Moura and S. Ullrich. *The Lean 4 theorem prover and programming language.* Proceedings of CADE-28, LNCS **12699** (2021), 625–635. https://doi.org/10.1007/978-3-030-79876-5_37.

[14] 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 USA **97** (2000), 4112–4117. https://doi.org/10.1073/pnas.97.8.4112.