

Finite-copy productive operation in a reversible autocatalytic binding network:

an explicit probability guarantee with a machine-checked proof

Author name to be inserted

Affiliation to be inserted

13 September 2026

Abstract

Structural criteria for autocatalysis, such as reflexively autocatalytic food-generated (RAF) sets or stoichiometric autocatalytic cores, certify that a reaction network can in principle amplify its own catalysts. They do not decide whether a reactor that starts with food alone will actually establish catalytic material, keep enough of it in the face of washout and reversible sequestration, and export a specified amount of product within a specified time. We prove such a finite-time operating guarantee for one explicit six-species open stochastic reaction network with exact integer-count propensities: basal and template-catalysed ligation of two food species into a product X , reversible binding of X to its substrates, duplex formation and release, continuous food supply and washout. At copy scale $V = 10^8$ and from the food-only initial state, the event that a weighted catalytic count enters a prescribed band by time 500, stays above a lower threshold thereafter, resources remain in a corridor, and at least 10^7 covalent mass units are exported during $(500, 1000]$ has probability at least $9/10$, for every equilibrium-bias parameter $K \in [8, 12]$ and release parameter $r \in [18, 22]$. Removing only the two catalytic ligation channels leaves a network in which a conservative version of the same export event has probability below $1/5000$. The proof is a chain of generator estimates on finite stopped uniformized kernels: conserved resource units control free food, a weighted catalytic count turns free and bound catalyst into a single drift inequality, a changing exponential test yields entry by a fixed deadline, a return potential pays for residence from the random entry state, and a marked-export tilt yields the lower tail of exported mass. All failures are charged to one union bound under one law; no independence between stages, deterministic approximation, or restart at entry is used. The combined theorem is formally verified in Lean 4 with Mathlib. The result concerns the specified mechanism and parameter box; it is not a claim about generic random reaction networks or experimentally calibrated rates.

1 Introduction

Autocatalysis is the organising idea behind most proposals for how chemistry can bootstrap itself toward biology. Its structural theory is well developed. Kauffman’s reflexively autocatalytic sets [13], the RAF formalism of Hordijk and Steel [10, 11, 20], and the stoichiometric classification of autocatalytic cores by Blokhuis, Lacoste and Nghe [4] all answer a question of the form: does the network contain a subnetwork that can, in principle, make its own catalysts from a food source? On the experimental side, minimal template replicators have been realised in nucleic acids and ribozymes [15, 19, 24], and design principles for minimal autocatalysts based on templated ligation have been analysed in kinetic detail [18].

A structural certificate, or even a positive mean-field growth rate, leaves a different question open. Consider a small open reactor, continuously fed with food and continuously diluted, that contains no catalyst at all at time zero. For it to be productive, four things must happen

together on the same trajectory. Catalytic material must first appear, which at zero catalyst count can only happen through an uncatalysed (basal) channel. Once present, it must be retained against washout. Since the catalyst is a template that binds its substrates and its own product, a large fraction of it is sequestered in complexes at any time, a mechanism known to reduce template replication to sub-exponential (“parabolic”) growth [21]. Finally, a specified amount of covalent product must actually leave the reactor within a specified time. At finite copy numbers these requirements interact: free food is depleted by binding and by consumption, the number of free catalysts is random, and the state of the reactor at the moment catalysis is first established is itself random. The natural notion of success is therefore an *event* on trajectories of the count process (establish, retain, deliver), and the object of interest is its probability.

This paper proves such a guarantee for one explicit mechanism. The reactor has two food species U, W , a covalent product X formed by ligating them, two substrate complexes $C_1 = X:U$ and $C_2 = X:(U, W)$, and a duplex $Z = X:X$. Ligation inside C_2 produces Z , which releases two copies of X ; this is the templated, catalytic route. A basal ligation $U + W \rightarrow X$ with a very small rate provides the only startup route. All five internal reaction pairs are reversible, foods are fed at a constant rate, and every species is washed out at a rate equal to its count. The model is the literal continuous-time Markov chain on integer counts with mass-action propensities [3, 7, 14], including the falling factorial for the homodimerisation $2X \rightarrow Z$. Two parameters are left free: K , the equilibrium bias of the two ligation steps, and r , the rate of duplex release. Our main result (Theorem 3.1) states that at copy scale $V = 10^8$ and for every $(K, r) \in [8, 12] \times [18, 22]$, the operating event described above has probability at least 9/10, with entry deadline 500, observation window $(500, 1000]$ and export threshold $10^7 = 0.1V$ covalent mass units. Deleting exactly the two catalytic ligation channels $C_2 \rightleftharpoons Z$, and changing nothing else, leaves a reactor in which a conservative version of the same export event has probability below 1/5000.

The method is elementary but, we believe, of independent interest. All estimates are generator inequalities of Foster–Lyapunov type [16], which in the stochastic reaction network literature are usually used for long-time properties such as positive recurrence [2, 9]. Here they are used quantitatively over a finite horizon, on finite stopped models, through uniformization [8, 12]: every bound is a statement about a stochastic kernel $P = I + \mathcal{G}/q$ and its Poisson mixture, so that failure probabilities are numbers computed from explicit drift constants rather than asymptotic statements. Four ingredients make the numbers work. Two conserved resource units control free food without clamping it. A weighted catalytic count, in which bound complexes count more than free product, has a uniform positive linear drift at low density; this replaces any assumption about the availability of free catalysts. A *changing* exponential test converts that drift into an entry probability by a fixed deadline without paying the exponential threshold factor that a fixed test would cost. A return potential that is constant before entry and exponential after entry pays for residence from the random entry time and the random landing composition in a single estimate. The lower tail of exported mass follows from a marked-counter tilt. Because every estimate is pointwise in the state, they average against the actual distribution passed from one interval to the next, and the final bound is a union bound on one law.

The combined theorem, including all numerical constants, has been checked in Lean 4 [5] with Mathlib [22]. Appendix A identifies the formal declarations and their relation to the lemmas of this paper. We emphasise what the formal statement covers: the finite stopped law of Section 3, the propensities of Table 1, the fixed initial state, the fixed copy scale and the stated window. Conventional arguments in the text (for instance the identification of the stopped law with the reactor up to the first declared failure, and the mean-field illustration of Section 9) are labelled as such.

Relation to other work. Stochastic models of the emergence of autocatalytic cycles have been studied by simulation [6]; the present result is analytic and is a theorem rather than a sampled frequency. The role of dynamics beyond structure in the evolvability debate [23] and the parameter inventory of Nghe et al. [17] motivate the choice of a fully specified reactor with explicit rates. A companion manuscript [1] proves a probability law for productive operation over an ensemble of *random* polymer networks at a volume scale growing with polymer length; the present paper is its finite-copy, fixed-mechanism complement, and shares with it the insistence on a same-observable disabled comparison.

What is not claimed. The theorem is a finite-horizon statement. The count process is irreducible and positive recurrent on $\mathbb{N}^6$ (every state is reachable from every other through washout, feeds and basal ligation, and the total mass M of (6) has generator $4V - M$, so the Foster–Lyapunov criterion [2, 16] applies), so it visits the catalyst-free state infinitely often and indefinite residence has probability zero; finite-time residence is the correct object. The reversible internal pairs share a common equilibrium vector for every K (Remark 2.2), which is an internal consistency property of the rate table, not a thermodynamic model of the open reactor. No experimental rate has been fitted, the copy scale $V = 10^8$ is not claimed to be optimal, and nothing is asserted about other windows, in particular not about the exploratory window (1, 501] that preceded the present formulation. Chemical-state inheritance, growth and division, and evolution are outside the scope.

Organisation. Section 2 specifies the reactor and its observables. Section 3 defines the operating event and states the theorem together with its error budget. Section 4 collects the two generator inequalities used throughout. Sections 5–8 prove the lemmas and the theorem. Section 9 shows a mean-field illustration, Section 10 discusses scope, and the appendices record the formal verification and the elementary inequalities.

2 The reactor

2.1 Species, reactions and parameters

The reactor contains six molecular species. Two are food species U and W , supplied from outside; think of them as two oligomer fragments. Their covalent ligation gives the product $X = UW$, which is also the template. A template captures its substrates in two steps: X binds U to form the complex $C_1 = X:U$, and C_1 binds W to form $C_2 = X:(U, W)$. Within C_2 the two bound fragments are ligated, giving a duplex $Z = X:X$ of the template with a new copy. The duplex dissociates into two free templates. Every step is reversible: the product can be cleaved, complexes can unbind, the duplex-internal bond can be cleaved, and two free templates can re-associate into a duplex. The state is the count vector

$$n = (u, w, x, c_1, c_2, z) \in \mathbb{N}^6.$$

Fix

$$V = 10^8, \quad \varepsilon = \frac{1}{500\,000\,000}, \quad 8 \leq K \leq 12, \quad 18 \leq r \leq 22, \quad k = K^{-1}. \quad (1)$$

The ten internal channels and their count propensities are listed in Table 1; Figure 1 shows the mechanism. In addition, each food species is fed at rate V (a channel $\emptyset \rightarrow U$ and a channel $\emptyset \rightarrow W$, each with constant propensity V), and every species is washed out at a rate equal to its count (six channels, propensities u, w, x, c_1, c_2, z). There are eighteen channels in total. The count update of a channel subtracts its reactants and adds its products; every positive-rate channel has its reactants available.

Table 1: Internal channels. “Forward” and “reverse” refer to the displayed order. Bimolecular propensities carry the divisor V ; the homodimerisation $2X \rightarrow Z$ uses the falling factorial $x(x-1)$ with no additional factorial divisor. The two channels of the last-but-one row are the catalytic ligation channels removed in the disabled comparison.

Reaction	Meaning	Forward	Reverse
$U + W \rightleftharpoons X$	basal ligation / cleavage	$\varepsilon uw/V$	$\varepsilon k x$
$X + U \rightleftharpoons C_1$	first substrate binding / unbinding	$20 xu/V$	$20 c_1$
$C_1 + W \rightleftharpoons C_2$	second substrate binding / unbinding	$20 c_1 w/V$	$20 c_2$
$C_2 \rightleftharpoons Z$	templated ligation / duplex cleavage	$20 c_2$	$20 k z$
$Z \rightleftharpoons 2X$	duplex release / re-association	$r z$	$r x(x-1)/V$




Figure 1: The mechanism. Solid arrows are the reversible internal channels of Table 1 (labels give forward / reverse coefficients); dashed arrows are washout channels, each with propensity equal to the species count. Washout of X, C_1, C_2, Z carries the covalent export marks 4, 4, 4, 8. The two catalytic ligation channels (red) are the only channels removed in the disabled comparison.

Meaning of the parameters. The number V is a copy-number scale, not a molecule count: it is the divisor in every bimolecular propensity and the feed rate of each food, so that concentrations n/V obey the classical density-dependent scaling [14]. Initially there are $2V$ food molecules and no other species. The parameter K is the common equilibrium constant of the two ligation steps (basal ligation and duplex-internal ligation): both reverse channels carry the factor $k = 1/K$. The parameter r sets the time scale of duplex release relative to the binding and unbinding steps, whose coefficient is 20. The basal coefficient ε is fixed; the time unit is the inverse of the dilution rate. Unbinding and washout rates are of order 1–20 per unit time, so binding is fast compared to dilution and release is faster still. Time and mass are the nondimensional units of the model; Remark 2.3 gives an illustrative conversion.

Remark 2.1 (Startup route). At the initial state $n_0 = (V, V, 0, 0, 0, 0)$, every channel that requires X, C_1, C_2 or Z has propensity zero. The only channel that creates catalytic material is basal ligation, with propensity $\varepsilon uw/V = \varepsilon V = 1/5$ at n_0 . Startup is therefore an immigration problem: roughly one basal product every five time units, each of which is washed out at rate 1 unless the templated cycle amplifies it first. The entry estimate of Section 6 makes this quantitative.

Remark 2.2 (Internal equilibrium vector). In concentration (mass-action) form the vector $(u, w, x, c_1, c_2, z) = (1, 1, K, K, K, K^2)$ balances all five internal pairs: $\varepsilon \cdot 1 \cdot 1 = (\varepsilon/K)K$, $20K \cdot 1 = 20K$, $20K \cdot 1 = 20K$, $20K = (20/K)K^2$, $rK^2 = rK^2$. Thus the rate table is consistent with a single internal equilibrium for every K . The open reactor is fed and diluted and is not assumed to be near this equilibrium; the remark only rules out a hidden internal driving cycle.

2.2 Resource units, catalytic count, mass and export

Internal channels move food atoms between free and bound forms and never create or destroy them. This is expressed by the two resource units

$$A(n) = u + x + 2c_1 + 2c_2 + 2z, \quad B(n) = w + x + c_1 + 2c_2 + 2z, \quad (2)$$

which count U -type and W -type material respectively (X contains one of each; C_1 contains two U and one W ; C_2 and Z contain two of each). Every internal channel preserves A and B ; only feeds and washouts change them. The *resource corridor* is

$$\mathcal{R} = \{n : 0.9V \leq A(n) \leq 1.1V, 0.9V \leq B(n) \leq 1.1V\}. \quad (3)$$

The *weighted catalytic count* is

$$Y(n) = x + \frac{9}{8}c_1 + \frac{7}{5}c_2 + \frac{9}{5}z. \quad (4)$$

Every catalytic form is counted, free or bound. A bound complex counts more than a free template because it has already captured substrate and is closer to producing a new copy; the specific weights are chosen so that the drift of Y has a uniform positive linear margin at low density (Lemma 5.2). Since the largest weight is $9/5$,

$$x + c_1 + c_2 + z \geq \frac{5}{9}Y(n), \quad (5)$$

which is how Y is converted back into a number of molecules subject to washout (Lemma 7.2).

Physical stoichiometric mass and covalent product stock are

$$M(n) = 2u + 2w + 4x + 6c_1 + 8c_2 + 8z = 2(A(n) + B(n)), \quad S(n) = 4x + 4c_1 + 4c_2 + 8z. \quad (6)$$

Here each food fragment has mass 2, so X has mass 4 and a duplex has mass 8; S counts only covalently ligated material, so bound but unligated food in C_1, C_2 is not counted as product. Export is measured at washout: washout of X, C_1, C_2, Z carries the marks 4, 4, 4, 8 respectively (the covalent content of what leaves), and all other channels carry mark 0. Let $L = 10^7$. The recorded export counter C starts at 0, is inactive until time 500, and afterwards updates at every jump by $C \mapsto \min\{L, C + m_j\}$, where m_j is the mark of the channel that fired. The saturation at L keeps the state space finite; $C = L$ at time 1000 means precisely that the uncapped marked sum over $(500, 1000]$ reached at least L .

Remark 2.3 (Illustrative units). If C_{ref} is a reference molar concentration and Ω a volume, $V = N_A C_{\text{ref}} \Omega$ and the nondimensional rate of an m -molecular channel is $k_{\text{phys}} C_{\text{ref}}^{m-1} / d_{\text{ref}}$, where d_{ref} is the dilution rate. At $C_{\text{ref}} = 1 \mu\text{M}$, $V = 10^8$ corresponds to $\Omega \approx 166 \text{ pL}$; if $d_{\text{ref}} = 1/\text{h}$, the certified window is hours 500 to 1000. No dilution rate or microscopic rate has been fitted to an experimental system; the conversion only fixes the meaning of the symbols.

3 Operating event, finite stopped law, and main theorem

3.1 Thresholds and phases

Fix the entry threshold $H = 40\,000$, the residence floor $h = 20\,000$, the entry deadline 500, the observation window $(500, 1000]$, and the export threshold $L = 10^7 = 0.1V$. The operating event, in path language, is: the weighted catalytic count reaches H by time 500; after that first crossing it never falls below h through time 1000; the resource units stay in the corridor $\mathcal{R}$ through time 1000; and at least L covalent units are exported during $(500, 1000]$. Residence is measured from the random entry time, not from time 500, and export is the increment of the counter over a fixed window, not a resident inventory.

To make this an event on a finite state space, the count vector is augmented with a phase $p \in \{0, 1, 2\}$ and the counter $C \in \{0, \dots, L\}$. Initially $p = 0$. After each jump from phase 0, set $p = 1$ if the new state has $Y \geq H$; otherwise keep $p = 0$. After entry, set $p = 2$ at the first jump to a state with $Y < h$; phase 2 is absorbing and retains the failure counts. Resource exit is handled by stopping: the first state outside $\mathcal{R}$ is retained and all its subsequent rates are set to zero. At time 500, any state still in phase 0 is frozen (its rates are set to zero): missing the deadline is a declared failure, and only those paths are stopped. On all other active states the literal count dynamics continue and the export counter is switched on. Reaching $C = L$ does not stop the dynamics: the resource and residence requirements continue to time 1000.

3.2 The finite stopped law

Embed the counts in the finite box $\{0, \dots, 3V + 2\}^6$. Every transition from a state in $\mathcal{R}$ is unchanged by the box: each coordinate of a corridor state is at most $1.1V$ (a coordinate is bounded by $\max\{A, B\}$) and a single jump increases it by at most 2, so it remains below $3V + 3$. The rules of the previous paragraph define two finite generators, $\mathcal{G}_0$ on $(0, 500]$ and $\mathcal{G}_1$ on $(500, 1000]$, acting on functions of (n, p, C) . A common bound for the total rate on active states is

$$\sum_j a_j(n) \leq q := 3000V = 3 \cdot 10^{11}, \quad (7)$$

because each internal channel has propensity at most $27V$ when every count is at most $1.1V$ (for instance $r x(x-1)/V \leq 22 \cdot 1.21V$ and $20 xu/V \leq 24.2V$), the two feeds total $2V$, and the six washouts total at most $6.6V$. Uniformization [8, 12] with this q gives the stochastic kernels $P_i = I + \mathcal{G}_i/q$ and the semigroups

$$T_i(t)f = e^{-qt} \sum_{m=0}^{\infty} \frac{(qt)^m}{m!} P_i^m f, \quad \mathbb{E}_{\text{on}} f = [T_0(500) T_1(500)f](n_0, 0, 0). \quad (8)$$

Equation (8) is the probability law to which the theorem applies. The composition $T_0(500)T_1(500)$ carries the complete distribution of (n, p, C) at time 500 into the second interval; nothing is reset or conditioned at the interface. Up to the first declared failure (resource exit, return below h , or the deadline freeze) the law coincides with the reactor's count process, and on the success event no failure occurs; this identification is a standard property of stopped nonexplosive jump processes and is not part of the formal statement.

For the *disabled* system remove exactly the two channels $C_2 \rightarrow Z$ and $Z \rightarrow C_2$. Keep every other propensity, including basal ligation and cleavage, all binding and unbinding, duplex release and re-association, the feeds, the washouts, the initial state and the output window. Stop only upon resource exit; there is no entry test, residence test or deadline freeze. Write $\mathbb{E}_{\text{off}}$ for the corresponding two-interval expectation.

3.3 Main theorem

Theorem 3.1 (Finite-copy operating guarantee). *Under (1), let*

$$\mathcal{S} = \{n \in \mathcal{R}, p = 1, C = L\} \quad \text{at time 1000}$$

for the enabled law, and let $\mathcal{D} = \{C = L\} \cup \{n \notin \mathcal{R}\}$ at time 1000 for the disabled law. Then, for every $(K, r) \in [8, 12] \times [18, 22]$,

$$\mathbb{P}_{\text{on}}(\mathcal{S}) \geq \frac{9}{10}, \quad \mathbb{P}_{\text{off}}(\mathcal{D}) < \frac{1}{5000}. \quad (9)$$

On $\mathcal{S}$ the recorded covalent export over $(500, 1000]$ is at least $L = 0.1V$ and the physical mass satisfies $M \leq 4.4V$.

The event $\mathcal{S}$ is the operating event of the previous subsection: $p = 1$ at time 1000 records that entry occurred by time 500 (otherwise the state was frozen in phase 0) and that no return below h occurred afterwards (otherwise $p = 2$); $n \in \mathcal{R}$ records that the resource corridor was never left (an exit is absorbing); and $C = L$ records the exported mass. The disabled event $\mathcal{D}$ is deliberately larger than the export event $\{C = L\}$: every resource exit is counted as a possible success, and no entry or residence requirement is imposed. Any bound on $\mathbb{P}_{\text{off}}(\mathcal{D})$ is therefore a bound on the probability that the disabled reactor, whatever else it does, exports L covalent units in the window. Uniformity in (K, r) means that the same numbers hold for each fixed admissible pair; the parameters do not vary along a trajectory.

Figure 2 shows the timeline, and Table 2 the error budget that proves the first inequality. Each row is a pointwise generator estimate under the law (8); the rows are combined by a union bound, not by a product.

4 Two generator inequalities

For a function f of the augmented state write $\mathcal{G}f(n) = \sum_j a_j(n)(f(n + \nu_j) - f(n))$, where the sum runs over the channels, a_j is the propensity and $n + \nu_j$ the updated state (including the phase and counter rules). For a linear observable F of the counts let $\Delta_j F = F(n + \nu_j) - F(n)$ and

$$Q_F(n) = \sum_j a_j(n) (\Delta_j F)^2,$$

the quadratic variation rate of F . Two implications for the finite kernels $P = I + \mathcal{G}/q$ are used throughout.




Figure 2: The operating event. The entry time τ is random; residence is required from τ onward. Export is counted only during $(500, 1000]$, and the resource corridor is required throughout.

Lemma 4.1 (Bounded drift). *If $f \geq 0$ and $\mathcal{G}f \leq b$ pointwise on active states, then $T(t)f \leq f + bt$.*

Proof. $Pf = f + \mathcal{G}f/q \leq f + b/q$, so $P^m f \leq f + mb/q$ by induction, and the Poisson mixture (8) with mean qt gives $T(t)f \leq f + bt$. $\square$

Lemma 4.2 (Exponential decay). *If $f \geq 0$ and $\mathcal{G}f \leq -\lambda f$ pointwise with $0 \leq \lambda \leq q$, then $T(t)f \leq e^{-\lambda t}f$.*

Proof. $Pf \leq (1 - \lambda/q)f$, hence $P^m f \leq (1 - \lambda/q)^m f$, and $\sum_m e^{-qt}(qt)^m(1 - \lambda/q)^m/m! = e^{-\lambda t}$. $\square$

Both conclusions are pointwise in the starting state. Consequently they can be averaged against any distribution, in particular against the distribution at time 500 produced by the first interval. Stopped states have zero rates, so $\mathcal{G}f = 0$ there, and every drift bound of the form $\mathcal{G}f \leq b$ with $b \geq 0$ remains valid after a stop. An event indicator dominated by a nonnegative potential, $\mathbf{1}_E \leq f$, is bounded in probability by the expectation of f . Finally, exponential potentials are handled with the elementary inequality (Appendix B)

$$e^v \leq 1 + v + \frac{3}{5}v^2 \quad (|v| \leq \frac{9}{50}), \quad (10)$$

which for a tilt parameter θ with $|\theta\Delta_j F| \leq 9/50$ gives

$$\sum_j a_j(e^{\theta\Delta_j F} - 1) \leq \theta\mathcal{G}F + \frac{3}{5}\theta^2 Q_F. \quad (11)$$

5 Resource control and local catalytic estimates

Lemma 5.1 (Resource failure). *For either stopped law (enabled or disabled), the probability of resource exit through time 1000 is less than 10^{-4} .*

Proof. Internal channels preserve A and B . Feeds add one unit at rate V ; washout of a species removes its weight in A (or B) at a rate equal to its count. Hence for $D \in \{A, B\}$,

$$\mathcal{G}D = V - D, \quad Q_D \leq V + 2D, \quad |\Delta_j D| \leq 2.$$

For $|\theta| \leq 1/100$, inequality (11) gives

$$\sum_j a_j(e^{\theta\Delta_j D} - 1) \leq \theta(V - D) + \frac{3}{5}\theta^2(V + 2D). \quad (12)$$

Consider the potentials $f_D^+ = e^{(D-1.1V)/100}$ and $f_D^- = e^{-(D-0.9V)/100}$. For f_D^+ ($\theta = 1/100$) the right side of (12) is nonpositive whenever $D \geq 1.05V$, since then $(V - D) + \frac{3}{500}(V + 2D) \leq$

Table 2: Error budget for $\mathbb{P}_{\text{on}}(\mathcal{S}) \geq 9/10$. The final states are partitioned into success and four failure classes; each failure class is bounded by a pointwise estimate under the law (8), and the bounds are added.

Failure class at time 1000	Estimate	Bound	Where
resource exit (absorbed outside $\mathcal{R}$)	exponential potentials in A, B	$< 10^{-4}$	Lemma 5.1
phase 0 (missed deadline), $n \in \mathcal{R}$	changing exponential test	$\leq \frac{1}{12} + \frac{1}{6 \cdot 10^8}$	Lemma 6.1
phase 2 (fell below h), $n \in \mathcal{R}$	return potential	$< 10^{-4}$	Lemma 7.1
phase 1, $n \in \mathcal{R}$, $C < L$	marked-export tilt	$< 10^{-4}$	Lemma 7.2
total failure	union bound	$< \frac{1}{10}$	Section 8

$-0.05V + \frac{3}{500}(3.2V) < 0$; for f_D^- ($\theta = -1/100$) it is nonpositive whenever $D \leq 0.95V$. In the complementary region ($D < 1.05V$ for f_D^+ , $D > 0.95V$ for f_D^-) the potential is at most $e^{-V/2000}$, a jump changes D by at most two, so every destination value is at most $e^{-V/2000+1/50}$; dropping the negative part of the generator and using (7) bounds the generator of each potential there by $qe^{-V/2000+1/50}$.

The sum $F = \sum_{D \in \{A, B\}} (f_D^+ + f_D^-)$ is nonnegative, equals $4e^{-V/1000}$ at n_0 , and is at least 1 at any state outside $\mathcal{R}$. Its drift bound $\mathcal{G}F \leq 4qe^{-V/2000+1/50}$ remains valid on stopped states and on states frozen by any other failure rule. Lemma 4.1 over the two intervals (total time 1000) gives

$$\mathbb{P}(\text{resource exit}) \leq \mathbb{E}F(1000) \leq 4(e^{-100000} + 3 \cdot 10^{14}e^{-50000+1/50}) < 10^{-4}. \quad (13)$$

The last inequality follows from $e^{-a} \leq m!/a^m$ ($a > 0$, Appendix B) with $m = 5$ for both terms. Removing the catalytic ligation channels leaves A, B , their drifts and their quadratic rates unchanged, so the same bound holds for the disabled law. $\square$

Lemma 5.2 (Local catalytic estimates). *On $\mathcal{R} \cap \{Y \leq V/1000\}$,*

$$\mathcal{G}Y \geq \frac{39}{100}Y + \frac{16}{25}\varepsilon V, \quad (14)$$

$$Q_Y \leq 5Y + \frac{25}{16}\varepsilon V, \quad |\Delta_j Y| \leq \frac{9}{5}, \quad (15)$$

$$\sum_j a_j (e^{-s\Delta_j Y} - 1) \leq -\left(\frac{3}{10}s - 3s^2\right)Y - \frac{14}{25}\varepsilon V s, \quad 0 \leq s \leq \frac{2}{25}. \quad (16)$$

Proof. Free food. From (2), $u = A - (x + 2c_1 + 2c_2 + 2z)$ and $w = B - (x + c_1 + 2c_2 + 2z)$, and both subtracted quantities are at most $2Y$. Hence in the corridor $u, w \geq 0.9V - 2Y \geq 0.898V \geq \frac{4}{5}V$ when $Y \leq V/1000$; also $u, w \leq 1.1V$ and $x \leq Y \leq V/1000$.

Drift. Summing the increments of Y over the eighteen channels (each internal pair contributes its forward and reverse propensity times the signed change of Y ; feeds change nothing; washouts subtract the weight of the species) gives the identity

$$\mathcal{G}Y = \varepsilon uw/V - \varepsilon kx + \left(\frac{5u}{2V} - 1\right)x + \left(-\frac{29}{8} + \frac{11w}{2V}\right)c_1 + \frac{11}{10}c_2 + \left(\frac{r}{5} - \frac{9}{5} - 8k\right)z - \frac{r}{5V}(x^2 - x).$$

With $u, w \geq \frac{4}{5}V$, $r \geq 18$ and $k \leq \frac{1}{8}$, the coefficients of x, c_1, c_2, z before the two negative terms $-\varepsilon kx$ and $-rx^2/(5V)$ are at least 1, $\frac{31}{40}$, $\frac{11}{10}$, $\frac{4}{5}$, each at least $\frac{2}{5}$ times the corresponding weight 1, $\frac{9}{8}$, $\frac{7}{5}$, $\frac{9}{5}$ in (4). Since $\varepsilon \leq 10^{-6}$, $k \leq \frac{1}{8}$, $r \leq 22$ and $x/V \leq 10^{-3}$, the two losses total less than $x/100 \leq Y/100$. The term $rx/(5V)$ is nonnegative, and the basal source is $\varepsilon uw/V \geq (0.898)^2 \varepsilon V \geq \frac{16}{25}\varepsilon V$. This proves (14).

Quadratic rate. The internal jumps of Y are ± 1 (basal ligation/cleavage), $\pm \frac{1}{8}$ (first binding), $\pm \frac{11}{40}$ (second binding), $\pm \frac{2}{5}$ (ligation/duplex cleavage) and $\pm \frac{1}{5}$ (release/re-association); feeds and food washouts have zero jump; washouts of X, C_1, C_2, Z have jumps $-1, -\frac{9}{8}, -\frac{7}{5}, -\frac{9}{5}$. Therefore

$$\begin{aligned} Q_Y &= \varepsilon uw/V + \varepsilon kx + \frac{1}{64}(20xu/V + 20c_1) + \left(\frac{11}{40}\right)^2(20c_1w/V + 20c_2) + \left(\frac{2}{5}\right)^2(20c_2 + 20kz) \\ &\quad + \frac{1}{25}(rz + rx(x-1)/V) + x + \left(\frac{9}{8}\right)^2c_1 + \left(\frac{7}{5}\right)^2c_2 + \left(\frac{9}{5}\right)^2z. \end{aligned}$$

Using $uw/V^2 \leq \frac{25}{16}$, $u/V, w/V \leq \frac{5}{2}$, $x/V \leq \frac{5}{4}$ and $x(x-1) \leq x^2$, the coefficients of x, c_1, c_2, z in the resulting upper bound are at most $\frac{481}{160}$, $\frac{343}{64}$, $\frac{267}{40}$, $\frac{113}{25}$, each at most 5 times the corresponding weight; the basal term is at most $\frac{25}{16}\varepsilon V$. This proves (15). (The bounds on $u/V, w/V, x/V$ used here are deliberately loose; they hold in the corridor with room to spare.)

Exponential tilt. With $v = -s\Delta_j Y$ and $0 \leq s \leq \frac{2}{25}$ we have $|v| \leq \frac{18}{125} < \frac{9}{50}$, so (11) gives the bound $-s\mathcal{G}Y + \frac{3}{5}s^2Q_Y$. Substituting (14)–(15), the coefficient of Y is $-\frac{39}{100}s + 3s^2 \leq -\frac{3}{10}s + 3s^2$, and the constant is $-\frac{16}{25}\varepsilon V s + \frac{15}{16}\varepsilon V s^2 \leq -\frac{14}{25}\varepsilon V s$ because $\frac{15}{16}s \leq \frac{15}{16} \cdot \frac{2}{25} < \frac{2}{25}$. This proves (16). $\square$

The margin $\frac{2}{5}$ in the drift is a statement about the whole catalytic population, bound and free. No estimate of the number of *free* templates is needed anywhere in the proof: binding moves material between the terms of Y with a net positive contribution to the weighted drift.

6 Entry by a fixed deadline

A fixed exponential test e^{-sY} with s large enough to give a useful decay rate would cost a factor e^{sH} to convert back into the entry indicator, which at $H = 40\,000$ is useless. The test therefore starts at a very small exponent, where the conversion factor is e , and is sharpened along the uniformized clock at a rate that the drift inequality (16) pays for.

Lemma 6.1 (Entry estimate). *Under the enabled law, the probability of being at time 500 in the corridor and still in phase 0 is at most $\frac{1}{12} + \frac{1}{600\,000\,000}$.*

Proof. Consider the auxiliary chain that follows the enabled count dynamics and is absorbed at the first state outside $\mathcal{R}$ or with $Y \geq H$. It is a device for the no-entry event only. Put

$$F_s(n) = \mathbf{1}_{\{n \in \mathcal{R}, Y(n) < H\}} e^{-sY(n)}, \quad a(s) = \frac{3}{10}s - 3s^2, \quad b = \frac{14}{25}\varepsilon V = \frac{14}{125}.$$

On an active state (in $\mathcal{R}$, $Y < H \leq V/1000$), the destinations at which F_s vanishes only decrease the generator, so (16) gives $\mathcal{G}F_s \leq -(a(s)Y + bs)F_s$. With $P = I + \mathcal{G}/q$ and $1 + v \leq e^v$,

$$PF_s \leq e^{-bs/q} e^{-a(s)Y/q} F_s \leq e^{-bs/q} F_t \quad \text{whenever } 0 \leq s \leq \frac{2}{25}, \quad t \leq s + a(s)/q, \quad (17)$$

because F_t is nonincreasing in t . The no-entry indicator is bounded by $e^{H/40000} F_{1/40000} = e F_{1/40000}$.

Startup schedule. For a cap d and a coefficient c with $c + 3d \leq \frac{3}{10}$, the update $s \mapsto \min\{d, (1 + c/q)s\}$ satisfies $t \leq s + a(s)/q$ for all $s \leq d$, since $cs \leq (\frac{3}{10} - 3s)s = a(s)$. Apply (17) with this update and discard the decay factor. Repeated updates give the cap or the corresponding geometric growth, whichever is smaller. Take $(d, c) = (\frac{1}{20}, \frac{3}{20})$ for $90q$ steps and then $(d, c) = (\frac{2}{25}, \frac{3}{50})$ for $10q$ steps. Bernoulli's inequality in blocks of $10q$ steps gives

$$\left(1 + \frac{3}{20q}\right)^{90q} \geq \left(1 + \frac{3}{2}\right)^9 = \left(\frac{5}{2}\right)^9 > 2000, \quad \left(1 + \frac{3}{50q}\right)^{10q} \geq \frac{8}{5},$$

so the exponent reaches $\frac{1}{40000} \cdot 2000 = \frac{1}{20}$ within the first block of steps and then $\frac{1}{20} \cdot \frac{8}{5} = \frac{2}{25}$ within the second. Thus $P^{100q} F_{1/40000} \leq F_{2/25}$.

Decay at the cap. Keeping $s = \frac{2}{25}$ for a further m steps and retaining the factor $e^{-bs/q}$ in (17),

$$P^{m+100q} \mathbf{1}_{\text{active}} \leq e \cdot e^{-b(2/25)m/q} = \exp\left(1 - \frac{28}{3125} \cdot \frac{m}{q}\right). \quad (18)$$

For $m \geq 399q$ this is at most $e^{1-28 \cdot 399/3125} = e^{-8047/3125} < \frac{1}{12}$; the strict inequality is verified by the Taylor polynomial of degree six of $e^{8047/3125}$, which exceeds 12.9 (Appendix B).

Poisson clock. Let J have the Poisson distribution with mean $500q$. Splitting the Poisson sum at $499q$, the no-entry probability at time 500 is at most $e^{-8047/3125} + \mathbb{P}(J < 499q)$. For $\rho = \frac{999}{1000}$ the Chernoff bound gives

$$\mathbb{P}(J < 499q) \leq \exp[-499q \log \rho + 500q(\rho - 1)] \leq e^{-q/500000} < \frac{1}{600\,000\,000},$$

where the middle inequality uses $\log(1 - u) \geq -u - 2u^2$ at $u = \frac{1}{1000}$ and the last one $e^{-600000} \leq 2/600000^2$ (Appendix B).

Identification. Before either entry or resource exit, the auxiliary chain and the phase-0 part of the enabled law have identical transitions, and both tests vanish after entry or exit. Induction on kernel steps followed by Poisson summation shows that the two no-entry probabilities coincide. No assumption about the state after entry has been introduced. $\square$

7 Residence and exported product

Lemma 7.1 (Return estimate). *Under the enabled law, the probability of a return failure (phase 2) through time 1000 is less than 10^{-4} .*

Proof. On the augmented states define

$$R(n, p) = \begin{cases} e^{-1600}, & p = 0, \\ \exp[-\frac{2}{25}(Y(n) - 20000)], & p \in \{1, 2\}. \end{cases}$$

At the entry jump the destination has $Y \geq H = 40\,000$, so R moves from e^{-1600} to a value at most e^{-1600} : the potential cannot increase at entry. At a return failure ($Y < h$) it is at least 1. For entered states in the corridor with $Y \leq 100\,000 = V/1000$, inequality (16) at $s = \frac{2}{25}$ shows $\mathcal{G}R \leq 0$. For $Y \geq 100\,000$, every positive-rate destination has $Y' \geq 100\,000 - \frac{9}{5}$, so $R(n', p') \leq e^{-799982/125}$; dropping the negative term $-\sum_j a_j R$ and using (7) gives the global bound

$$\mathcal{G}R \leq B_R := 3 \cdot 10^{11} e^{-799982/125}.$$

Absorption, resource stops and the deadline freeze set rates to zero and preserve this bound. Lemma 4.1 applied successively to the two intervals gives

$$\mathbb{P}(\text{return failure}) \leq e^{-1600} + 1000 B_R < 10^{-4}, \quad (19)$$

the last inequality by $e^{-a} \leq m!/a^m$ with $m = 2$ and $m = 6$ respectively. (An exact rational evaluation of the right side gives about $3.9 \cdot 10^{-6}$.) The constant value of R in phase 0 is what makes this a single estimate valid from the initial state: the random entry time and the random binding composition at entry are paid for by the fact that R does not increase at the entry jump. $\square$

Lemma 7.2 (Insufficient export). *Under the enabled law, the probability of being at time 1000 in the corridor and in phase 1 with $C < L$ is less than 10^{-4} .*

Proof. On the second interval use the killed test

$$F(n, p, C) = \mathbf{1}_{\{n \in \mathcal{R}, p=1, C < L\}} e^{-C/8}.$$

States at which F vanishes stay that way: deadline failures are frozen, resource and return failures are absorbing, and a saturated counter remains saturated. On an active state, $Y \geq h = 20\,000$, so by (5) the total washout intensity of catalytic species is

$$x + c_1 + c_2 + z \geq \frac{5}{9}Y \geq \frac{100000}{9}.$$

Each such washout carries a mark of at least 4, and $e^{-4/8} - 1 \leq \frac{2}{3} - 1 = -\frac{1}{3}$ (Appendix B), so these channels contribute at most $-\frac{1}{3} \cdot \frac{100000}{9} F = -\frac{100000}{27} F$ to $\mathcal{G}_1 F$. Every other channel has mark 0 or kills the test, and neither increases $\mathcal{G}_1 F$. Saturation causes no difficulty: if the destination is still active, the counter did not saturate and its increment equals the mark. Hence

$$\mathcal{G}_1 F \leq -\frac{100000}{27} F.$$

Lemma 4.2, $F \leq 1$, and $\mathbf{1}_{\text{active}} \leq e^{L/8} F$ give, from every possible state at time 500,

$$T_1(500) \mathbf{1}_{\text{active}} \leq \exp\left(\frac{10^7}{8} - \frac{500 \cdot 100000}{27}\right) = e^{-16250000/27} < 10^{-4}, \quad (20)$$

using $e^{-a} \leq 1/a$. Averaging this pointwise estimate against the actual time-500 distribution proves the claim. $\square$

The mechanism behind Lemma 7.2 is simple: while the catalytic population sits above the residence floor, at least $\frac{5}{9}h \approx 11\,111$ catalytic molecules are exposed to washout at unit rate, each carrying at least 4 covalent units. Over the window of length 500 the expected export is therefore at least $2.2 \cdot 10^7$, and the tilt converts this into a lower tail at the threshold 10^7 . The estimate uses only the residence floor; it does not use the much larger populations that the mean-field dynamics predict (Section 9).

8 Disabled comparison and proof of the theorem

Lemma 8.1 (Disabled export). *For the disabled law, $\mathbb{P}_{\text{off}}(\mathcal{D}) < \frac{1}{5000}$.*

Proof. Binding, unbinding, duplex release and re-association preserve the covalent stock S of (6): these channels rearrange covalent material without ligating or cleaving. After removal of the two catalytic ligation channels, basal ligation is the only positive source of S and basal cleavage its only internal sink; washout removes S at the export rate $e(n) = 4x + 4c_1 + 4c_2 + 8z$. Channel bookkeeping gives the exact balance

$$\mathcal{G}_{\text{off}}S + e(n) = 4(\varepsilon uw/V - \varepsilon kx). \quad (21)$$

Consider the inventory $I = S + C \geq 0$. On the first interval C is inactive, so $\mathcal{G}_{\text{off}}I = \mathcal{G}_{\text{off}}S \leq 4\varepsilon uw/V$. On the second interval each counter increment is at most the true export mark, so $\mathcal{G}_{\text{off}}I \leq \mathcal{G}_{\text{off}}S + e(n) \leq 4\varepsilon uw/V$ again. On active states $u, w \leq 1.1V$, hence on both intervals

$$\mathcal{G}_{\text{off}}I \leq \frac{121}{25} \varepsilon V = \frac{121}{125}.$$

The bound persists after a resource stop. Since $I(n_0, 0) = 0$, Lemma 4.1 over the two intervals gives $\mathbb{E}_{\text{off}}I(1000) \leq 968$. The pointwise inequality $L \mathbf{1}_{\{C=L\}} \leq I$ then yields

$$\mathbb{P}_{\text{off}}(C = L) \leq \frac{968}{10^7} = \frac{121}{1\,250\,000}.$$

Adding the resource-exit bound of Lemma 5.1,

$$\mathbb{P}_{\text{off}}(\mathcal{D}) < \frac{121}{1\,250\,000} + \frac{1}{10\,000} = \frac{123}{625\,000} < \frac{1}{5000}. \quad (22)$$

No entry or residence requirement was used. $\square$

The comparison is fair in the following sense. The disabled reactor keeps the same initial state, feeds, washouts, basal chemistry, binding and release kinetics, and the same output observable; the only change is that the template no longer ligates its bound substrates. It still produces covalent product through basal ligation, and (21) shows that this source, at most $4\varepsilon uw/V \leq 0.968$ units per unit time, is what bounds its output. The two events $\mathcal{S}$ and $\mathcal{D}$ are not identical, but $\mathcal{D} \supseteq \{C = L\}$, so the second inequality in (9) bounds the probability of the export condition itself, under weaker requirements than the enabled event imposes. No coupling of the two processes is asserted or needed.

Proof of Theorem 3.1. Every final state of the enabled law belongs to exactly one of five classes: success $\mathcal{S}$; a resource failure (retained outside $\mathcal{R}$); a return failure (phase 2, in $\mathcal{R}$); a missed-deadline state (phase 0, in $\mathcal{R}$, frozen at time 500); or an active state with insufficient output (phase 1, in $\mathcal{R}$, $C < L$). Lemma 6.1 bounds the third of these classes because the frozen phase-0 states at time 1000 are exactly the phase-0 states at time 500. Lemmas 5.1, 7.1 and 7.2 bound the others. The union bound gives

$$\mathbb{P}_{\text{on}}(\mathcal{S}) \geq 1 - \frac{1}{12} - \frac{1}{600\,000\,000} - \frac{3}{10\,000} \geq \frac{9}{10}. \quad (23)$$




Figure 3: Deterministic mean-field solution of the six-species system from the food-only state, for three parameter pairs in the certified box and, in grey, for $K = 10$, $r = 20$ with the channels $C_2 \rightleftharpoons Z$ removed. (a) Weighted catalytic count VY during startup; the horizontal lines are the entry threshold H and the residence floor h . (b) Cumulative covalent export over $(500, t]$ in units of V ; the horizontal line is the certified threshold $L/V = 0.1$. These curves illustrate the mechanism; they are not stochastic realisations and play no role in the proof.

All four estimates are pointwise inequalities under the single two-interval law (8); no independence between entry, residence and output is used. Lemma 8.1 gives the second inequality. Finally, $C = L$ means that the uncapped marked sum over $(500, 1000]$ reached at least $L = 0.1V$, and $M = 2(A + B) \leq 4.4V$ on $\mathcal{R}$. $\square$

9 A mean-field illustration

The theorem is analytic and does not rely on simulation. It is nevertheless useful to see what the deterministic mean-field system predicts, because it shows how conservative the certified numbers are. Replacing counts by concentrations n/V and the falling factorial $x(x-1)/V$ by its leading term gives the mass-action system associated with Table 1 (the limit of the count process as $V \rightarrow \infty$ on compact time intervals [14]). Figure 3 shows its solution from the food-only state. The weighted catalytic count VY crosses H at $t \approx 9-10$ for all three parameter choices shown and saturates at about $0.37-0.39V$, far above the residence floor; the cumulative covalent export over $(500, 1000]$ is about $7.3-7.6 \cdot 10^2 V$, some four orders of magnitude above the certified threshold $0.1V$. With the two catalytic ligation channels removed, the same system exports about $4 \cdot 10^{-6} V$ over the window and holds a catalytic count below one molecule. The gap between the mean-field export and the certified threshold is the price of an estimate (Lemma 7.2) that uses only the residence floor h ; the gap between the certified entry deadline 500 and the mean-field entry time ≈ 10 is the price of a low-density immigration estimate that must hold for every trajectory with probability close to one, not on average. The script that produces the figure is part of the supplementary material.

10 Discussion

What has been established. For one explicit open reactor with exact count propensities, the probability that a run started from food alone establishes catalysis by a deadline, retains it, keeps its resources in a corridor, and exports a specified covalent mass in a specified window is at least $9/10$, uniformly over a nontrivial box of thermodynamic bias and release rate; and the export condition alone has probability below $1/5000$ once the two templated-ligation channels

are deleted. The two numbers are separated by more than a factor of 4000 on the same observable. The proof is a sequence of drift estimates on one law; it does not assume clamped free food, independent lineages, a deterministic trajectory, or a restart at entry, and it carries the full distribution across the interface at time 500.

Why the estimates close. Three features of the argument seem to be the ones that matter. First, counting bound forms in the catalytic observable: the weights in (4) make sequestration a redistribution among the terms of Y rather than a loss, so that a uniform margin of $\frac{2}{5}$ survives across the parameter box; without it one would have to control free templates, which fluctuate strongly. Second, conserved resource units in place of a free-food bound: the corridor for A, B is enforced by exponential potentials whose failure probability is astronomically small, and the free food needed by the low-density drift is then a consequence, valid exactly where it is used. Third, count-scale noise: the quadratic rate of Y is of order Y itself (with a moderate constant, 5), so the logarithmic correction Q_Y/Y^2 is already about 10^{-2} at $Y = 10^3$, and $V = 10^8$ suffices. An earlier route through concentration tolerances squared a very small absolute deviation and required astronomically larger volumes; that requirement was a loss in the proof, not a physical obstruction.

The copy scale. $V = 10^8$ enters as the bimolecular divisor and the feed rate. It bounds the number of resident molecules only through the resource corridor, which itself is a probabilistic conclusion (total mass at most $4.4V$ with probability at least $1 - 10^{-4}$); an open reactor has no hard cap. The value is the smallest round scale at which the present constants close comfortably, not an optimum. Whether the same mechanism operates reliably at $V = 10^6$ or 10^7 is open; the entry estimate is the component most sensitive to V , since the basal source εV and the deadline enter it multiplicatively.

Finite time. The result is deliberately a finite-horizon statement. The count process on $\mathbb{N}^6$ is irreducible and positive recurrent, so it returns to the catalyst-free state infinitely often, and no theorem can promise residence forever. The time scales in the theorem (deadline 500, window of length 500) are long compared to the mean-field entry time (≈ 10) and are dictated by the immigration estimate. A shorter certified deadline, or the exploratory window $(1, 501]$ that motivated this work, would require a sharper entry argument; nothing in the present paper rules such windows out.

Limitations and outlook. The theorem concerns the specified mechanism, rate table, initial state, copy scale and window. It is not a theorem about random reaction networks (see the companion manuscript [1] for that setting), and the rates are not calibrated to any experimental replicator, although the mechanism (template binds two substrates, ligates them, releases a duplex) is the one realised in nucleotide-based replicators [19, 24] and analysed in [18]. Spatial structure, saturation of binding sites, side reactions, and energy accounting are outside the model. The internal reversibility and the common equilibrium vector of Remark 2.2 are consistency properties of the rate table, not a thermodynamic analysis of the driven reactor. Natural next steps are an optimisation of the copy scale and of the thresholds H, h ; a deterministic integrated-output companion result; extensions in which the template mechanism sits inside a larger host network with competing reactions; and the connection to inheritance of chemical states, which requires growth and division and is a different question.

On the formal verification. Every inequality in Sections 3–8, including the numerical constants, is part of a Lean 4 development whose main declaration is the combined theorem (Appendix A). The value of the formal check here is not that the individual inequalities are

deep, but that a proof of this kind consists of many bookkeeping identities over eighteen channels, phase rules, saturation rules and two intervals, in which a single sign error or an unsupported reactant would invalidate the number 9/10. The conventional text above is a faithful rendering of the formal argument, written so that it can be checked by a reader without the formal development.

Data and code availability. The Lean sources, verification receipts, the figure script and the constant-checking script accompany the paper as supplementary material and will be released with a public repository tag.

A Formal verification

The formal statement of (9) is the declaration

`RandomViability.Binding.finite_copy_binding_guarantee`

in the file `proofs/RandomViability/BindingOperatingGuarantee.lean` of the accompanying Lean development. Its statement reads as follows (transcribed in ASCII: `R` is the real type, `in Icc a b` is membership in the closed interval, and the elided arguments are the positivity and bound hypotheses on $1/K$ and r that the expectations take as parameters):

```
theorem finite_copy_binding_guarantee (K r : R) (hK : K in Icc 8 12)
  (hr : r in Icc 18 22) :
  (9/10 : R) <= operatingExpectation (1/K) r ...
    (eventIndicator {X | operatingSuccess X}) /\
  disabledExpectation (1/K) r ...
    (eventIndicator {X | X.2.val = 10000000 \
      not resourceGood (boxCounts X.1) 100000000}) < 1/5000
```

where `operatingExpectation` and `disabledExpectation` are the two-interval Poissonized expectations (8) for the enabled and disabled finite laws, `operatingSuccess X` is the conjunction of $n \in \mathcal{R}$, `phase = 1` and $C = 10^7$, and the disabled event is $\mathcal{D}$. The consequence on export and physical mass is the declaration `success_output_and_mass` in the same file. Table 3 lists the principal declarations on the proof path, in the order of the lemmas above; supporting algebra, the finite-kernel library and the Poissonization lemmas are imported by these files.

The recorded strict verification used Lean v4.30.0 with a frozen Mathlib manifest, treated warnings as errors, and accepted no `sorry`. The main file compiled with empty compiler output in 48.75 s together with 43 authenticated local dependency files; the axioms reported for the main declaration are exactly `Classical.choice`, `Quot.sound` and `propext`. The SHA-256 of the main source is

7df4d165a100cecab76cac91cee3a281931f2c9798ad46a507d5c8679a82d786 .

Within the development, verification is reproduced from the repository root by

```
python scripts/verify_proof.py \
  proofs/RandomViability/BindingOperatingGuarantee.lean \
  -timeout 150 \
  -declaration RandomViability.Binding.finite_copy_binding_guarantee
```

which enforces the pinned working directory and toolchain. The proof uses symbolic inequalities over the parameter box; no enumeration of the state space, floating-point evaluation, or simulation enters as evidence. The formal scope is exactly the finite stopped law of Section 3 with the propensities of Table 1, the initial state n_0 , $V = 10^8$, and the window $(500, 1000]$. Two statements in the text are conventional and are not formal exports: the identification of the stopped law with the reactor process up to the first declared failure, and the recurrence remark in Section 10. The mean-field computation of Section 9 is a numerical illustration.

Table 3: Principal formal declarations (namespace `RandomViability.Binding`) and the corresponding steps of this paper.

Step in this paper	Declaration(s)
Rate support and stoichiometry	<code>BindingStoichiometry</code> , <code>BindingRateSupport</code> , <code>BindingBookkeeping</code> (files)
Resource drifts and Lemma 5.1	<code>BindingMoietyDrift</code> (file); <code>evaluated_resource_probability</code>
Lemma 5.2, (14)	<code>actual_count_corridor_growth</code>
Lemma 5.2, (15)	<code>actual_count_variance_bound</code>
Lemma 5.2, (16)	<code>actual_entry_exponential</code>
Lemma 6.1, schedule and decay	<code>entry_hundred_clock_transfer</code> , <code>entry_discrete_after_startup</code>
Lemma 6.1, Poisson clock	<code>poissonized_after_cutoff</code>
Lemma 7.1	<code>return_foster</code> , <code>evaluated_return_budget</code>
Lemma 7.2	<code>output_probability</code>
Union bound (23)	<code>operating_events_cover</code> , <code>operating_success_lower_bound</code>
Lemma 8.1	<code>disabled_success_upper_bound</code>
Theorem 3.1	<code>finite_copy_binding_guarantee</code>

B Elementary inequalities

The following inequalities are used with explicit constants.

- (i) $e^v \leq 1 + v + \frac{3}{5}v^2$ for $|v| \leq \frac{9}{50}$. Indeed $e^v - 1 - v = \sum_{m \geq 2} v^m/m! \leq \frac{v^2}{2} \sum_{j \geq 0} (|v|/3)^j = \frac{v^2}{2(1-|v|/3)} \leq \frac{v^2}{2(1-3/50)} = \frac{25}{47}v^2 < \frac{3}{5}v^2$.
- (ii) $e^{-a} \leq m!/a^m$ for $a > 0$ and every integer $m \geq 0$, since $e^a \geq a^m/m!$. This gives (13) with $m = 5$: $4e^{-100000} \leq 4 \cdot 120/10^{25}$ and $4 \cdot 3 \cdot 10^{14} e^{-50000+1/50} \leq 4 \cdot 3 \cdot 10^{14} \cdot e^{1/50} \cdot 120/(5 \cdot 10^4)^5 < 10^{-8}$; it gives (19) with $m = 2$ and $m = 6$; and (20) with $m = 1$.
- (iii) $\log(1-u) \geq -u - 2u^2$ for $0 \leq u \leq \frac{1}{2}$, since $-\log(1-u) = \sum_{j \geq 1} u^j/j \leq u + u^2 \sum_{j \geq 0} u^j/2 \leq u + u^2$.
- (iv) $e^{-1/2} \leq \frac{2}{3}$, equivalently $e \geq \frac{9}{4}$, and $e^{-600000} \leq 2/600000^2$ is (ii) with $m = 2$.
- (v) $e^{8047/3125} > 12$: with $a = \frac{8047}{3125}$, the partial sum $\sum_{j=0}^6 a^j/j!$ is a rational number exceeding 12.91, and e^a exceeds every partial sum of its Taylor series.
- (vi) Bernoulli: $(1+x)^N \geq 1 + Nx$ for $x \geq 0$ and integer $N \geq 0$.

A short script (`check_constants.py`, supplementary material) re-derives the generator identities of Lemmas 5.2 and 8.1 symbolically and evaluates every numerical bound of the paper in exact rational arithmetic; it is an independent check, not part of the proof.

References

- [1] Author name to be inserted. Productive autocatalysis in random polymer networks: a sharp probability law. Companion manuscript, 2026.
- [2] David F. Anderson and Jinsu Kim. Some network conditions for positive recurrence of stochastically modeled reaction networks. *SIAM Journal on Applied Mathematics*, 78(5): 2692–2713, 2018. doi: 10.1137/17M1161427.

- [3] David F. Anderson and Thomas G. Kurtz. *Stochastic Analysis of Biochemical Systems*, volume 1.2 of *Mathematical Biosciences Institute Lecture Series*. Springer, Cham, 2015. doi: 10.1007/978-3-319-16895-1.
- [4] Alex Blokhuis, David Lacoste, and Philippe Nghe. Universal motifs and the diversity of autocatalytic systems. *Proceedings of the National Academy of Sciences*, 117(41): 25230–25236, 2020. doi: 10.1073/pnas.2013527117.
- [5] Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming language. In *Automated Deduction – CCADE 28*, volume 12699 of *Lecture Notes in Computer Science*, pages 625–635. Springer, 2021. doi: 10.1007/978-3-030-79876-5_37.
- [6] Alessandro Filisetti, Alex Graudenzi, Roberto Serra, Marco Villani, Davide De Lucrezia, Rudolf M. Füchslin, Stuart A. Kauffman, Norman Packard, and Irene Poli. A stochastic model of the emergence of autocatalytic cycles. *Journal of Systems Chemistry*, 2:2, 2011. doi: 10.1186/1759-2208-2-2.
- [7] Daniel T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *The Journal of Physical Chemistry*, 81(25):2340–2361, 1977. doi: 10.1021/j100540a008.
- [8] Winfried K. Grassmann. Transient solutions in Markovian queueing systems. *Computers & Operations Research*, 4(1):47–53, 1977. doi: 10.1016/0305-0548(77)90007-7.
- [9] Ankit Gupta, Corentin Briat, and Mustafa Khammash. A scalable computational framework for establishing long-term behavior of stochastic reaction networks. *PLoS Computational Biology*, 10(6):e1003669, 2014. doi: 10.1371/journal.pcbi.1003669.
- [10] Wim Hordijk and Mike Steel. Detecting autocatalytic, self-sustaining sets in chemical reaction systems. *Journal of Theoretical Biology*, 227(4):451–461, 2004. doi: 10.1016/j.jtbi.2003.11.020.
- [11] Wim Hordijk and Mike Steel. Chasing the tail: the emergence of autocatalytic networks. *Biosystems*, 152:1–10, 2017. doi: 10.1016/j.biosystems.2016.12.002.
- [12] Arne Jensen. Markoff chains as an aid in the study of Markoff processes. *Skandinavisk Aktuarietidskrift*, 36:87–91, 1953. doi: 10.1080/03461238.1953.10419459.
- [13] Stuart A. Kauffman. Autocatalytic sets of proteins. *Journal of Theoretical Biology*, 119(1):1–24, 1986. doi: 10.1016/S0022-5193(86)80047-9.
- [14] Thomas G. Kurtz. The relationship between stochastic and deterministic models for chemical reactions. *The Journal of Chemical Physics*, 57(7):2976–2978, 1972. doi: 10.1063/1.1678692.
- [15] Tracey A. Lincoln and Gerald F. Joyce. Self-sustained replication of an RNA enzyme. *Science*, 323(5918):1229–1232, 2009. doi: 10.1126/science.1167856.
- [16] Sean P. Meyn and Richard L. Tweedie. Stability of Markovian processes III: Foster–Lyapunov criteria for continuous-time processes. *Advances in Applied Probability*, 25(3): 518–548, 1993. doi: 10.2307/1427522.
- [17] Philippe Nghe, Wim Hordijk, Stuart A. Kauffman, Sara I. Walker, Francis J. Schmidt, Harry Kemple, Jessica A. M. Yeates, and Niles Lehman. Prebiotic network evolution: six key parameters. *Molecular BioSystems*, 11(12):3206–3217, 2015. doi: 10.1039/C5MB00593K.

- [18] Yann Sakref and Olivier Rivoire. Design principles, growth laws, and competition of minimal autocatalysts. *Communications Chemistry*, 7:239, 2024. doi: 10.1038/s42004-024-01250-y.
- [19] Dirk Sievers and Günter von Kiedrowski. Self-replication of complementary nucleotide-based oligomers. *Nature*, 369:221–224, 1994. doi: 10.1038/369221a0.
- [20] Mike Steel, Wim Hordijk, and Joana C. Xavier. Autocatalytic networks in biology: structural theory and algorithms. *Journal of the Royal Society Interface*, 16(151):20180808, 2019. doi: 10.1098/rsif.2018.0808.
- [21] Eörs Szathmáry and Irina Gladkih. Sub-exponential growth and coexistence of non-enzymatically replicating templates. *Journal of Theoretical Biology*, 138(1):55–58, 1989. doi: 10.1016/S0022-5193(89)80177-8.
- [22] 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.
- [23] Vera Vasas, Eörs Szathmáry, and Mauro Santos. Lack of evolvability in self-sustaining autocatalytic networks constraints metabolism-first scenarios for the origin of life. *Proceedings of the National Academy of Sciences*, 107(4):1470–1475, 2010. doi: 10.1073/pnas.0912628107.
- [24] Günter von Kiedrowski. A self-replicating hexadeoxynucleotide. *Angewandte Chemie International Edition in English*, 25(10):932–935, 1986. doi: 10.1002/anie.198609322.